

ROMAN STRICKER

CHARACTERIZATION, VERIFICATION & PROTECTION OF
LARGE-SCALE QUANTUM SYSTEMS



CHARACTERIZATION, VERIFICATION & PROTECTION OF LARGE-SCALE QUANTUM SYSTEMS



Doctoral thesis submitted to the faculty for
MATHEMATICS, COMPUTER SCIENCE AND PHYSICS

In partial fulfillment of the requirements for the degree of
DOCTOR OF PHILOSOPHY
(physics)

Carried out at the Institute for Experimental Physics under
supervision of o.Univ.-Prof. Dr. Rainer Blatt in the group for
Quantum Optics and Spectroscopy

ROMAN STRICKER

January 2024

ABSTRACT

The field of experimental quantum information processing has evolved rapidly, from the demonstration of basic building blocks about two decades ago to today's fruitful applications that are driving the development of a functional information processor. While nowadays applications venture into regimes where classical devices are challenged, demonstrating quantum advantage in a meaningful industrial or scientific problem remains an open task. This is partly caused by comparably small system sizes and the quality of gate operations. With ongoing progress in engineering ever larger quantum devices the view has shifted from in-principle demonstrations to being able to deploy quantum machines at scale. Scalability of designs, approaches and devices has become a chief consideration in the continuing developments.

In this thesis, we utilize a medium-scale device based on a string of $^{40}\text{Ca}^+$ ions confined in a linear Paul trap to address scalability challenges for present-day devices. Yet, all methodologies presented herein are hardware-agnostic and may be applied to different platforms just as well. One critical challenge is how to be sure that the output of a quantum computer is correct once it operates in a regime where classically checking is no longer feasible. Crucially, existing methods become resource-devouring when certifying system sizes just beyond a handful of information carriers—so-called quantum bits (qubits). Another critical challenge is the development of concepts that keep devices operable even if some of their quantum bits get lost.

The first experiment reported in this thesis demonstrates a scalable characterization method to obtain the complete tomography information of a multi-qubit system from a single measurement setting. This is achieved by enlarging the underlying Hilbert space, and works independent of the system size. On the postprocessing side, we complement this single-setting tomography with an adapted version of so-called “classical shadow” analysis to efficiently predict arbitrary polynomial functions of the density matrix orders of magnitudes faster than standard methods.

While system characterization is essential to improving setup functionality, large-scale devices have the disadvantage that the computation outcome for certain problems can no longer be confirmed in classical simulations. Based on a novel theory, the second experiment demonstrates the verification of a quantum computation by purely classical means.

Moreover, quantum systems cannot be completely isolated from their environment, and therefore will always be prone to errors. While quantum error correction promises to overcome inherent noise limitations, existing protocols are restricted to correcting errors that change the logical state. Realistic quantum computers, however, suffer not only from such computational errors, but, at a comparable rate, can experience a complete loss of the stored information or the information carriers. We present the first deterministic experiment that corrects qubit losses in real-time. This third work marks an essential step towards error corrected quantum information processors.

Our loss experiments furthermore feature in-sequence measurement and classical feed-forward, which are increasingly prevalent in modern semi-classical algorithms. While this experimental structure becomes more and more accessible, its time evolution might deviate from unitarity and can no longer be described by standard tools. In the fourth work we

develop a novel tomography approach based on quantum instruments to fully capture the dynamics of semi-classical quantum algorithms.

Our newly developed and scalable certification methods are essential to indentifying system limitations and to verify quantum computers, having the potential to significantly advance large-scale device developments. With the inclusion of qubit losses as a standard building block, we have further made a paradigm shift for error correction.

KURZFASSUNG

Die experimentelle Quanteninformationsverarbeitung hat sich in den letzten zwei Jahrzehnten von der Demonstration grundlegender Bausteine zu den heutigen vielversprechenden Anwendungen rasant weiterentwickelt. Diese Anwendungen sind bereits in der Lage die klassischen Möglichkeiten herauszufordern und die Entwicklung eines funktionalen Quantenprozessors stetig voranzutreiben. Der Nachweis des Quantenvorteils eines industriell oder wissenschaftlich wichtigen Problems stellt jedoch eine noch ungelöste Aufgabe dar. Dies ist mitunter auf die vergleichsweise geringen Systemgrößen sowie die Qualität der Gatteroperationen zurückzuführen. Mit der Realisierung immer größerer Systeme hat sich der Blick weg von prinzipiellen Demonstrationen hin zum Einsatz hoch skalierter Informationsprozessoren verschoben. Die Skalierbarkeit von Hardware sowie geeigneter Anwendungen ist somit zu einem Eckpfeiler der laufenden Entwicklungen geworden.

In dieser Arbeit verwenden wir eine lineare Paul-Falle zum Speichern einer Kette von $^{40}\text{Ca}^+$ -Ionen und bilden damit einen mittelgroßen Quantenprozessor. Mit diesem erforschen wir neue Herausforderungen auf dem Weg der Skalierbarkeit von Quantenrechnern. Die hier vorgestellten Methoden sind hardwareunabhängig und können auf anderen Plattformen gleichermaßen eingesetzt werden. Eine der kritischen Herausforderungen ist dabei die Lösung eines Quantenrechners zu überprüfen, sobald dieser in einem Bereich arbeitet, in dem die Nachvollziehbarkeit mit klassischen Mitteln nicht mehr gegeben ist. Bestehende Zertifizierungsmethoden arbeiten hier ressourcenverschlingend sobald die Systemgröße über eine Handvoll von Informationsträgern - den Quantenbits (Qubits) - hinaus wächst. Im Weiteren entwickeln wir Konzepte, die Quantenprozessoren auch dann betriebsfähig halten, wenn ein Teil ihrer Quantenbits verloren geht.

Das erste Experiment über das in dieser Arbeit berichtet wird, demonstriert eine skalierbare Charakterisierungsmethode, welche die vollständige Tomographieinformation eines Multi-Qubit Systems aus einer einzigen Messung heraus extrahiert. Dies geschieht durch Erweiterung des zugrundeliegenden Hilbertraums und funktioniert unabhängig der Systemgröße. Zur effizienten Datenanalyse komplementieren wir diese neue Einzelmessungstomographie mit einer angepassten Version der sogenannten "klassischen Schatten". Dies ermöglicht die Bestimmung beliebiger Polynomfunktionen der Dichtematrix um Größenordnungen schneller als das bisher der Fall war.

Während Systemcharakterisierungen für die Verbesserung der Versuchsaufbauten unerlässlich sind, stoßen wir bei hochskalierten Quantenprozessoren zusätzlich auf Probleme, deren Ergebnisse nicht mehr durch klassische Simulationen überprüfbar sind. Auf der Grundlage einer neuen Theorie erkunden wir mit dem zweiten Experiment die Verifizierung einer Quantenrechnung mit rein klassischen Mitteln.

Darüber hinaus können Quantensysteme nicht vollständig von ihrer Umgebung isoliert werden und bleiben daher fehleranfällig. Bestehende Strategien beschränken die Korrektur jedoch typischerweise auf Fehler, die den logischen Zustand ändern. Realistische Quantenrechner leiden jedoch nicht nur unter derart logischen Fehlern, sondern können in vergleichbarer Häufigkeit einen vollständigen Verlust der gespeicherten Information oder der Informationsträger erleiden. Wir präsentieren hierzu das erste deterministische Experiment, das Qubitverluste in Echtzeit korrigiert. Diese dritte Arbeit liefert einen wesentlichen Beitrag in Richtung fehlerkorrigierter Quanteninformationsverarbeitung.

Bei der Echtzeitkorrektur von Qubitverlusten verwenden wir Insequenzdetektion mit darauf folgender klassischer Sequenzverzweigung. Konkret werden die klassischen Messergebnisse mancher Qubits für den Fortlauf der Quantenrechnung, hier die Korrektur der Qubitverluste, herangezogen. Solche Experimente zählen zur Gattung der semi-klassischen Quantenalgorithmen. Dank der jüngsten technologischen Fortschritte wird diese experimentelle Struktur stetig zugänglicher. Ihre zeitliche Entwicklung wird jedoch von der klassischen Messung unterbrochen und kann daher von einem unitären Verlauf abweichen. In der vierten Arbeit entwickeln wir einen neuartigen Tomographieansatz auf der Grundlage von Quanteninstrumenten, welcher die komplette Dynamik von semi-klassischen Quantenalgorithmen erfasst.

Unsere neuentwickelten und skalierbaren Zertifizierungsmethoden sind unerlässlich, um Systemlimitierungen zu charakterisieren sowie Quantenrechnungen zu verifizieren. Diese Methoden tragen das Potenzial, die Entwicklung skalierbarer Technologien deutlich voranzutreiben. Ferner haben wir mit der Eingliederung von Qubitverlusten als Standardbaustein in Fehlerkorrektur einen Paradigmenwechsel vollzogen.

ACKNOWLEDGMENTS

A considerable part of the work in this thesis was created by the great conditions that Rainer Blatt's working group has created over the many years and that have grown into a unique environment. Unique, not only in terms of scientific exchange, but also in terms of friendly interaction with the many people involved, which never fell short.

With this, I would like to thank above all Professor Rainer Blatt, who has given a greenhorn like me the opportunity to work and develop on one of the most pressing research topics of our time.

Our ambitious experiment was led by a team of the most experienced senior scientists, who always had the sense to work in the right direction. I am talking about Thomas Monz, Philipp Schindler and Martin Ringbauer, who deserve a big thank you for showing a lot of patience in answering the many questions I had. Another thank you goes to Christian Marciniack, who always has something to improve, even if you think you did it right once.

All those long days in the lab seemed short indeed, thanks to colleagues who shared the same interests and who often became friends. In the beginning there were Daniel Nigg, Esteban Martinez and Alexander Erhard, from whom we newbys were able to take over the experiment in such great condition. Those newbys include Mike Meth and Lukas Postler. Mike did a great job in developing our experiment control which opened up many new possibilities. I also appreciate his good humour and his faible for conspiracy. I especially thank Lukas for sharing his technical knowledge in the lab and for going "power play" on the ion-addressing setup when it was already too late to do it carefully. My appreciation also goes Claire Edmunds, Vanya Pogorelov, Georg Jacob and many other colleagues, including those from the university, the IQOQI and the AQT.

I would also like to thank theorists Barbara Kraus, Chris Ferrie, Richard Küng and Markus Müller for their close collaborations and shaping their new ideas to the needs of our experiments, which made much of the scientific success possible.

Outside of the university, I would like to thank many friends who gave me strength to keep a cool head even when things sometimes got stuck at work. Especially worth mentioning is my family with mother Margrith, father Hans and sister Marion with husband Raffaele and their two children Siena and Moreno, who always put me back on ground-potential from all the quantum madness. Last but not least, I owe a lot of this to my lovely girlfriend Carla who was always there by my side and made me feel like I was normal.

CONTENTS

1	Scalability challenges of today’s quantum computers	1
1.1	Quantum computation	1
1.1.1	Quantum states	1
1.1.2	Entanglement	2
1.1.3	Quantum channels	5
1.1.4	Measurement	7
1.1.5	Constructing a quantum computer	9
1.1.6	Higher dimensions	10
1.2	Certification & benchmarks of quantum computers	10
1.2.1	Characterizing quantum systems	12
1.2.2	Verifying quantum computations	14
1.3	Quantum computing with faulty components	16
1.3.1	Quantum error correction	16
1.3.2	Stabilizer codes	18
1.3.3	The surface code	19
1.4	Qubit loss & leakage—beyond computational errors	21
1.5	Semi-classical quantum algorithms	23
2	The trapped-ion quantum information processor	25
2.1	The $^{40}\text{Ca}^+$ -qudit	26
2.2	The linear Paul trap	26
2.3	Laser-ion interaction	28
2.4	The trapped-ion toolbox	31
2.5	Device capabilities	33
3	A scalable approach to characterization	37
3.1	Quantum state tomography—the gold standard	37
3.1.1	Linear inversion reconstruction	38
3.1.2	Maximum likelihood estimation	40
3.2	Quantum process tomography	41
3.3	SIC POVM quantum state tomography	42
3.4	System characterization via classical shadows	45
3.4.1	Classical shadow tomography	46
3.4.2	Efficiently estimating linear system properties	47
3.4.3	Efficiently estimating non-linear system properties	48
3.5	Publication: Experimental single-setting quantum state tomography	51
4	Verifying an untrusted quantum device	87
4.1	Verifiable classical computation	88
4.1.1	Arthur-Merlin protocol	88
4.1.2	Complexity classes	89
4.1.3	Interactive proof system	91
4.1.4	Interactive proofs with trapdoor functions	93
4.1.5	Learning with errors	94
4.2	Verifiable quantum computation by classical means	95
4.2.1	Phrasing a decision problem as an energy measurement	96
4.2.2	Post-quantum secure delegation of energy measurements	98

- 4.2.3 Concept of classical verification 101
- 4.2.4 The step-by-step protocol 102
- 4.3 Publication: Classical verification of quantum computation 106
- 5 Qubit loss protection 129
 - 5.1 Detection of qubit loss 130
 - 5.2 Correction of qubit loss 132
 - 5.2.1 Erasure code—extended four qubit teleportation 132
 - 5.2.2 Loss in the surface code 139
 - 5.3 Publication: Experimental deterministic correction of qubit loss 141
 - 5.4 Non-identity dynamics of the no-loss case 157
 - 5.4.1 The planar color code 159
 - 5.5 Publication: Characterizing quantum instruments 161
- 6 Conclusion & outlook 191
- A List of publications 193

- Bibliography 194

ACRONYMS

AKS	<i>Agrawal–Kayal–Saxena</i>
AM	<i>Arthur–Merlin</i>
bit	<i>binary digit</i>
BPP	<i>bounded-error probabilistic polynomial time</i>
BQP	<i>bounded-error quantum polynomial time</i>
BSB	<i>blue sideband</i>
CNOT	<i>controlled-NOT</i>
CP	<i>completely-positive</i>
CPTP	<i>completely-positive and trace-preserving</i>
DC	<i>Doppler cooling</i>
DLWE	<i>decision learning with errors</i>
GHZ	<i>Greenberger–Horne–Zeilinger</i>
GST	<i>gate-set tomography</i>
H	<i>Hadamard</i>
IC	<i>informationally complete</i>
IP	<i>interactive proof</i>
LI	<i>linear inversion</i>
LWE	<i>learning with errors</i>
MA	<i>Merlin–Arthur</i>
MLE	<i>maximum likelihood estimation</i>
MS	<i>Mølmer–Sørensen</i>
NISQ	<i>noisy intermediate-scale quantum</i>
NP	<i>non-deterministic polynomial time</i>
P	<i>polynomial time</i>
PGC	<i>polarization-gradient cooling</i>
PLS	<i>projected least squares</i>
POVM	<i>positive operator-valued measure</i>
POVMs	<i>positive operator-valued measures</i>
PPT	<i>positive partial transpose</i>
PSPACE	<i>polynomial space</i>
QEC	<i>quantum error correction</i>
QIP	<i>quantum interactive proof</i>
QMA	<i>quantum Merlin–Arthur</i>
QND	<i>quantum non-demolition</i>

QPN	<i>quantum projection noise</i>
QPT	<i>quantum process tomography</i>
QST	<i>quantum state tomography</i>
qubits	<i>quantum bits</i>
qudits	<i>quantum digits</i>
qutrits	<i>quantum trits</i>
RB	<i>randomized benchmarking</i>
RSB	<i>red sideband</i>
SBC	<i>sideband cooling</i>
SIC	<i>symmetric informationally complete</i>
SPAM	<i>state preparation and measurement</i>
TP	<i>trace preserving</i>
VQE	<i>variational quantum eigensolver</i>

PREAMBLE

At the beginning of the 20th century, the development of quantum mechanics fundamentally changed our understanding of the world. Counterintuitive experiments, which could not be explained with existing theories, got the ball rolling and let many research directions flourish over the upcoming decades ranging from nuclear, atomic and particle physics all the way to applications in solid-state physics and quantum chemistry. It did not take much longer until this new understanding enabled technological developments. Most notably, the invention of semiconductors for information processing and the laser over the second half of the century that have had a decisive and lasting impact on our daily lives. The progress continued and starting with the new century, individual quantum particles like atoms, molecules, ions, photons, and artificial atoms could be reliably stored and manipulated, providing a new avenue to exploring the quantum world. Quantum technologies emerged and hold the potential to surpass their classical counterparts, above all, today's supercomputers. Besides, quantum technologies promise a means to efficiently simulate nature, offer secure communication paths and work as accurate microscopic sensors in metrology applications. These applications have recently been designated the four domains of quantum technology in the European Union's quantum flagship program [1] with the aim to drive quantum revolution in Europe. Beyond an increasing amount of governmental funds, quantum technologies have captured a lot of attention as more people than ever join the race for building new hardware and developing new applications. With quantum technology climbing the ladder of technological readiness, more and more industrial actors get involved, bringing more resources into the field with the final goal to accomplish commercial viability.

This thesis contributes primarily to the domains of quantum simulation and computation, which build on the seminal ideas of Richard Feynman from the 1980s: "*Let the computer itself be built of quantum mechanical elements which obey quantum mechanical laws*" [2]. Computers built from quantum systems would thereby be expected to circumvent the inexorable resource demands that prohibit the study of quantum problems by classical means. From then, *Quantum simulations* developed and currently cover promising research paths from nuclear physics [3] to molecular chemistry [4]. Extending Feynman's path, in 1989, David Deutsch proposed a model for a universal quantum computer that generalizes the classical Turing machine [5]. Over time, these ideas began to bear fruit, as the *superposition* of quantum states and fundamental properties without a classical counterpart (*entanglement*) indeed offer a pathway to outperform classical capabilities in a growing number of industrial and scientific applications. Today, quantum algorithms exist ranging from optimization problems [6] such as logistics [7] to programs for quantum machine learning [8] and notably Shor's algorithm for factoring large integer numbers [9], capable to defy current cryptography paradigms—just to name a few. Moreover, quantum effects will limit the further scaling of chip sizes in classical computers, which are slowly but surely entering the quantum realm. The development of quantum computers rapidly evolved from demonstrating basic building blocks [10] about two decades ago to nowadays fruitful applications in quantum sensing, simulations or computations [11]. The current era is marked by so-called *noisy intermediate-scale quantum* (NISQ)-devices [12] featuring tens of quantum information carriers - so-called *quantum bits* (qubits) - and venture into

regimes where classical computers are challenged [13] with the aim to finally demonstrate a quantum speed-up on an impactful problem.

Current quantum devices still fall short of outperforming classical systems in a useful problem, primarily due to persistent challenges in scaling to larger systems and the insufficient quality of the quantum operations. While quantum architectures such as superconducting platforms and trapped-ions achieve register sizes beyond 50 qubits [14, 15], error-rates on the order of 1% per two-qubit operation in such large registers limit their practical applicability.

Here, we employ a NISQ trapped-ion device and explore two branches of scalability bottlenecks, which we identify as timely and critical for pushing the field closer towards real-world applications. The tools we develop to address these challenges are architecture independent making them applicable to a wide range of platforms.

The first branch deals with the certification of quantum computations. This is a pertinent task, not only for pinpointing system limitations, which is vital for progress, but also for enabling trust in a quantum computational outcome that can no longer be classically simulated. *“If a quantum experiment solves a problem which is proven to be intractable for classical computers, how can one verify the outcome of the experiment?”* [16] quoted at a 2004 conference, Daniel Gottesmann aptly highlighted the crux of certification. We will get to the bottom of this somewhat paradox question and demonstrate a framework for certifying computations without requiring any knowledge or trust in the employed device.

Inherent to nature, and something that quantum computers have not been spared from is the influence of environmental noise. Yet, this is expected to persist on future hardware. Our second branch deals with ways of taming noise in quantum devices, marking a cornerstone of current developments, as was distinctly pointed out by David Deutsch *“Without error-correction all information processing, and hence all knowledge-creation, is necessarily bounded. Error-correction is the beginning of infinity”* [17]. In particular, our contributions extend the scope of correctable errors beyond computational errors to the correction of the complete loss of information carriers.

This thesis is organized as follows. Ch. 1 provides broad insights to quantum computation that are presented in an architecture-independent fashion. We then identify scalability issues of existing methods for system certification and computations with noisy hardware. From this four problems emerge, which we address in the remainder of this thesis. In Ch. 2 we will introduce the experimental platform used in this thesis. In Ch. 3 we develop and demonstrate a novel framework for scalable system characterizations that exhibits significant improvements over existing methods on both the quantum measurement and classical post-processing sides. In Ch. 4 we venture into postquantum-secure cryptography and demonstrate the first verification of a quantum computation with purely classical means. This is an important step to gain trust in future large-scale devices whose results cannot be directly verified classically and may not even be trustworthy. In Ch. 5 the first experimental demonstration of qubit loss correction is presented, where each detected loss event triggers a correction step in real-time. The latter experimental structure exemplifies a novel class of advanced tasks, designated as quantum-classical algorithms, which exhibit dynamics that can no longer be captured faithfully by existing characterization tools. In Ch. 5 we introduce a new paradigm for characterizing the most general quantum operations, called quantum instruments. Finally, Ch. 6 puts the results of this thesis into perspective and highlights interesting future questions that drive the ongoing developments.

SCALABILITY CHALLENGES OF TODAY'S QUANTUM COMPUTERS

It is believed that quantum computers outperform their classical predecessors in numerous applications [6–9]. While recently advanced devices are already challenging classical capabilities [13], proving a quantum advantage in a significant industrial or scientific problem is still an open task. There are several reasons for this, including the relatively small system sizes and the quality of gate operations. The scalability of quantum devices has thus become a cornerstone of current hardware developments.

The present chapter provides the basics of quantum computation in Sec. 1.1. Based on this knowledge, we identify critical scalability challenges to be solved in this thesis. This includes two timely branches of research that contribute to the development of larger quantum machines. Namely, the certification of quantum computation in Sec. 1.2 and 1.5 as well as quantum computing with noisy components, discussed over Secs. 1.3-1.4.

1.1 QUANTUM COMPUTATION

The following explanations provide a formal framework for quantum computation with particular emphasis on the concepts used in this work. For a more stringent presentation we refer the interested reader to Ref. [18].

1.1.1 *Quantum states*

The basic unit of information for classical computers is the *binary digit* (bit) that takes the two discrete values 0 and 1. Bits successfully serve for classical information processing with their binary values typically encoded by distinct voltage levels in electrical transistors. In analogy, two levels of a quantum system $|0\rangle$ and $|1\rangle$ manifest a qubit. A qubit differs from its classical counterpart in that its quantum nature allows it to exist not just in the states $|0\rangle$ and $|1\rangle$, but also in any superposition of the two. As such, quantum mechanics formally describes an isolated two-level system by a vector in a *Hilbert space* \mathcal{H}_2 . A general single-qubit pure state is given by

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle \quad \text{with } \alpha, \beta \in \mathbb{C} \quad \text{and} \quad |\alpha|^2 + |\beta|^2 = 1, \quad (1.1)$$

with probability amplitudes $|\alpha|^2$ and $|\beta|^2$ ensuring normalization. Normalization together with complex numbers α and β geometrically constrains qubit state vectors to the surface of a unit sphere. In quantum computation that unit sphere is referred to as *Bloch sphere*, where the computational basis states $\{|0\rangle, |1\rangle\}$ reside at opposite poles, illustrated by Fig. 1.1(a).

In view of real-world experiments a quantum system can only be detected to finite accuracy and might therefore exist in a mixture of multiple states $\{|\psi\rangle_i\}$ with certain

probabilities $\{p_i\}$. A very useful concept to account for this inaccuracy is the *density operator* description

$$\rho = \sum_i p_i |\psi\rangle_i \langle\psi|_i, \quad (1.2)$$

which is a Hermitian, positive-semidefinite operator of trace $\text{tr}(\rho) = 1$. In case a quantum state occupies a distinct state $|\psi\rangle$ with probability 1 it is said to be *pure*. On the other hand, if a quantum system can be in one of several states, each at a certain probability, it is no longer pure and designated *mixed*. We can quantify the purity of normalized states by

$$\mathcal{P} = \text{tr}(\rho^2) \quad \text{with } \rho \in [1/d, 1], \quad (1.3)$$

with its lower bound given by the inverse Hilbert space dimension $1/d$ relating to the complete mixed state $\mathbb{1}/d$. In particular, purity depends non-linearly on the density matrix ρ , which means that it cannot be observed directly on a quantum system.

A general single-qubit density matrix can be written as

$$\rho = \frac{1}{2} \left(\mathbb{1} + r_x X + r_y Y + r_z Z \right) \quad \text{with } \vec{r} \in \mathcal{R} \text{ (real numbers)}, \quad (1.4)$$

with the *Bloch vector* \vec{r} and the *Pauli matrices* $\{X, Y, Z\}$ that geometrically align with the orthogonal sphere axes [18]. Note that this ignores already a global phase. The Paulis (short for Pauli operators) together with the identity matrix $\mathbb{1}$ provide a basis for the space of density matrices, which are the following linear operators on \mathcal{H}_2

$$\mathbb{1} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (1.5)$$

In the literature the Pauli basis is often referred to as standard basis. On the Bloch sphere the pure states are on the surface and the mixed states are inside the sphere. A true quantum state obeys the physical properties of being positive-semidefinite $\rho > 0$ and of unit trace $\text{tr}(\rho) = 1$, ensuring normalization [18].

Let us expand the system size. The states of a system consisting of several qubits also form a Hilbert space due to the superposition principle. The multi-qubit Hilbert space is thus the tensor product of the Hilbert spaces of the individual qubits. As an example, the computational basis states for two qubits read in tensor representation $|0\rangle \otimes |0\rangle = |00\rangle, |01\rangle, |10\rangle$ and $|11\rangle$ or in decimal representation $|0\rangle, |1\rangle, |2\rangle$ and $|3\rangle$. The qubit-pair thus inhabits a four-dimensional Hilbert space $\mathcal{H}_2^{(1)} \otimes \mathcal{H}_2^{(2)}$. This generalizes to the N-qubit case as follows

$$|\psi\rangle = \sum_{i=0}^{2^N-1} \alpha_i |i\rangle \quad \text{with } \alpha_i \in \mathbb{C} \quad \text{and} \quad \sum_{i=0}^{2^N-1} |\alpha_i|^2 = 1, \quad (1.6)$$

with Hilbert space $\mathcal{H}_2^{\otimes N}$ of dimension $\dim(\mathcal{H}_2^{\otimes N}) = 2^N$. While the exponential growth in Hilbert space dimension is at the heart of the potential power of quantum technology, it represents a daunting scalability bottleneck when simulating quantum systems by classical means.

1.1.2 Entanglement

To unleash the full capabilities of quantum computers operating multi-qubit registers becomes imperative. The reason is that multi-qubit systems can exhibit correlations that go




Figure 1.1: **Pictorial representation of single-qubit quantum states.** (a) Semi-positive and trace-preserving pure quantum states can be illustrated as points on the surface of a unit sphere, named Bloch sphere. Mixed states reside within the sphere. The standard qubit basis aligns with the sphere's X, Y and Z axis, denoted by the so-called Pauli basis from Eq. (1.5). Its components uniquely parametrize arbitrary quantum states as exemplified in the single-qubit case of Eq. (1.4). The Pauli operators further constitute a measurement set of six projectors given by the ± 1 eigenvectors of the individual operators, outlined in Sec. 1.1.4. (b) Alternatively, the set of *symmetric informationally complete* (SIC) *positive operator-valued measures* (POVMs) consists of four projectors maximized for intervector spacing. In contrast to the Pauli basis, the non-orthogonal *positive operator-valued measure* (POVM) elements have finite overlap with arbitrary states, where every measurement contributes to the statistical accuracy [19, 20]. SIC POVMs are thus ideal measurements in terms of information gain, thoroughly discussed in Sec. 1.1.4.

beyond classical physics and are called entanglement. Crucially, the properties of entangled states can no longer be fully described by the properties of the individual components, even if spatially far separated. For quantum computers to outperform classical devices, large amounts of entanglement are required [21]. Consequently, entanglement is the subject of many theoretical and experimental studies.

Let us elaborate on the nature of these correlations. A quantum state of a composite system is said to be *separable* if it can be written as a mixture of product states. The two-qubit state $|00\rangle = |0\rangle \otimes |0\rangle$ exemplifies such a product state. On the contrary, if a quantum state ρ_{AB} with bipartite system components A and B cannot be decomposed into a product state of the individual components

$$\rho_{AB} \neq \rho_A \otimes \rho_B, \quad (1.7)$$

the state is said to be *entangled*.

In Ch. 3 we will use different ways to characterize entanglement. Typically, entanglement measures quantify the degree in separability of a bipartite system (A, B). Let us therefore consider the system ρ expressed by its density matrix

$$\rho = \sum_{ijkl} p_{kl}^{ij} |i\rangle\langle j| \otimes |k\rangle\langle l|. \quad (1.8)$$

A way of identifying whether this system is separable or not is by partial transposing either bipartition, e.g. subsystem A, relating to the following transformation

$$\rho^{\Gamma_A} = \sum_{ijkl} p_{kl}^{ij} |i\rangle\langle j| \otimes |l\rangle\langle k|, \quad (1.9)$$

and check whether the transposed outcome is still positive-semidefinite. Peres and Horodecki formulated this property as the so-called *positive partial transpose* (PPT) criterion [22, 23]

in 1996. The criterion states that the partial transpose of either bipartition $\{A, B\}$ is a completely positive map (See Eq. (1.20) below) so it transforms an input quantum state into a valid output quantum state, if and only if the combined system (A, B) is separable.

This can be illustrated by partial transposing qubit A on both a single-qubit density matrix ρ_A and an entangled two-qubit density matrix ρ_{AB} as follows

$$\begin{aligned} \rho_A &= \frac{1}{2} \begin{pmatrix} 1 & -i \\ i & 1 \end{pmatrix} & \rho_A^{\Gamma_A} &= \frac{1}{2} \begin{pmatrix} 1 & i \\ -i & 1 \end{pmatrix} \\ \rho_{AB} &= \begin{pmatrix} 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 \end{pmatrix} & \rho_{AB}^{\Gamma_A} &= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \end{aligned} \quad (1.10)$$

The second line exemplifies the two-qubit case of the famous *Greenberger-Horne-Zeilinger* (GHZ) state [24] that is maximally entangled and will play an important role throughout this thesis. Evidently, the transposed density matrix in the single-qubit case $\rho_A^{\Gamma_A}$ represents another physical quantum state. According to the PPT criterion, there must exist a valid transformation in between. While transpose is a valid quantum operation in the single-qubit case or more generally for any separable state, partial transpose of the entangled state in the second line leads to negative eigenvalues. This means that the transformed density matrix $\rho_{AB}^{\Gamma_A}$ does not anymore represent a physical quantum state. Importantly, negative eigenvalues would not be present if the two subsystems were separable, indicating the presence of entanglement in the original state.

The partial transpose of either bipartition, here A , relates to the so-called *quantum negativity* [25]

$$\mathcal{N}(\rho) = \frac{\|\rho^{\Gamma_A}\|_1 - 1}{2} \quad (1.11)$$

and serves to quantify entanglement for any value greater than zero. Entanglement defined by that means may be correctly referred to as *bipartite entanglement*. Moreover, the negativity is equivalent to the sum of the absolute values of all negative eigenvalues in the transformed outcome, which provides an efficient way to calculate it. Since entanglement is an intrinsic feature of the combined system (A, B) , negativity is symmetric with respect to a system's bipartite components $\|\rho^{\Gamma_A}\|_1 = \|\rho^{\Gamma_B}\|_1$.

Alternatively, one can relate entanglement to entropy. As for example the so-called second-order *Rényi-entropy* [26]

$$S^{(2)}(\rho_A) = -\log_2 \text{tr}(\rho_A^2). \quad (1.12)$$

The order corresponds to the exponent in the density matrix ρ_A , where from the second-order it features a non-linear dependency on it. Eq. (1.12) relates entropy to purity and by that to the length of the Bloch vector. Crucially, the authors of Ref. [27] show that the second order presented here can be efficiently estimated in experiments with randomly selected measurements by relying on the fact that the second order is contained in statistical correlations between the random measurement results. We utilize such a protocol throughout Ch. 3. For the sake of simplicity, we refer to the second order Rényi entropy simply as Rényi entropy from here on.

Let us discuss Eq. (1.12) and consider once more a bipartite system (A, B) . Experimentally, if one observes either bipartition, e.g. A , one receives system part ρ_A , denoted as the *reduced* density matrix. Mathematically, the reduced density matrix of subsystem A is defined by

taking the partial trace $\rho_A = \text{tr}_B(\rho_{AB})$ over the basis of the other subsystem B. As such, Eq. (1.12) relates entropy to the purity of the reduced density matrix. Considering the observations on both subsystems individually, we obtain the reduced density matrices ρ_A and ρ_B . If the total system ρ_{AB} is separable it holds that $\rho_{AB} = \rho_A \otimes \rho_B$ and in our example the entire information about system (A, B) can be recovered from observations on the individual subsystems. In the latter case, the reduced density matrix has purity $\text{tr}(\rho_A^2) = 1$ and yields Rényi-entropy $S^{(2)}(\rho_A) = 0$. However, for a non-separable state it holds that $\rho_{AB} \neq \rho_A \otimes \rho_B$. Therefore, individually observing ρ_A and ρ_B results in a loss of information, contrasted to observations on the combined system ρ_{AB} . A loss of information in turn manifests itself in an increase of entropy. If we turn the argument around, we can observe entropy with Eq. (1.12) and make a statement about entanglement. Conclusively, if we find $S^{(2)}(\rho_A) = S^{(2)}(\rho_B) > S^{(2)}(\rho_{AB})$, the system (A, B) contains entanglement in case the joint state is pure.

Let us illustrate Rényi-entropy on the two-qubit GHZ-state from Eq. (1.10). Since the combined system is pure, it gives $S^{(2)}(\rho_{AB}) = 0$. Upon partial tracing subsystem B, we find ρ_A in a complete mixture

$$\rho_A = \text{tr}_B(\rho_{AB}) = \frac{1}{2}(|0\rangle_A\langle 0| + |1\rangle_A\langle 1|) \quad \text{with} \quad \text{tr}(\rho_A^2) = \frac{1}{2}, \quad (1.13)$$

where the Rényi-entropy takes its maximum $S^{(2)}(\rho_A) = 1$. Analogous to negativity, Rényi-entropy is symmetric with respect of its bipartite system components $S^{(2)}(\rho_A) = S^{(2)}(\rho_B)$.

1.1.3 Quantum channels

As we have seen so far, the properties of qubit systems differ considerably from those of their classical counterparts. A quantum computer can thus be defined as a device that aims to store and manipulate quantum information in such a way that arbitrary computations can be performed, ideally surpassing classical capabilities.

Computations on a quantum device are decomposed into sequences of so-called *quantum gate operations* or *quantum gates*. Conceptually, quantum gates transfer a pure input state $|\psi_{\text{in}}\rangle$ into a pure output state $|\psi_{\text{out}}\rangle$ following a unitary evolution

$$|\psi_{\text{in}}\rangle \xrightarrow{\hat{U}} |\psi_{\text{out}}\rangle = \hat{U} |\psi_{\text{in}}\rangle \quad \text{with} \quad \hat{U}^\dagger \hat{U} = \mathbb{1}, \quad (1.14)$$

that ensures a positive-semidefinite outcome of unit norm $(\|\hat{U} |\psi_{\text{in}}\rangle\|_1)^2 = 1$. In contrast to classical gates, quantum gates are reversible. In view of the state vector representation, we can formulate single-qubit basis states as $|0\rangle = (1, 0)^T$ and $|1\rangle = (0, 1)^T$, where quantum gates are described as unitary 2×2 matrices. For example, the Pauli X operator from Eq. (1.5) performs a bit flip on an arbitrary input state

$$\begin{pmatrix} \alpha \\ \beta \end{pmatrix} \xrightarrow{X} \begin{pmatrix} \beta \\ \alpha \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \end{pmatrix}, \quad (1.15)$$

representing the quantum equivalent to the classical NOT-gate.

As generators of the *special unitary group* $SU(2)$, the Pauli operators can realize arbitrary single-qubit rotations [18]. Herefore angular momentum operators, e.g. $S_x = \frac{\hbar}{2}X$, provide general descriptions of two-level systems similar to the qubit. The vector in the Bloch sphere of Fig. 1.1(a) can further be uniquely defined by the azimuth angle θ and the polar

angle ϕ by means of the representation $[\cos \phi \sin \theta, \sin \phi \sin \theta, \cos \theta]^\top$. For this purpose, arbitrary single-qubit gates refer to rotations around an axis $\vec{n} = (n_x, n_y, n_z)$ given by

$$\hat{U} = e^{i\alpha} \hat{R}_{\vec{n}}(\theta) \quad \text{with} \quad \hat{R}_{\vec{n}}(\theta) = e^{-i\frac{\theta}{2} \vec{n} \cdot \vec{\sigma}} = \cos \frac{\theta}{2} \mathbb{1} - i \sin \frac{\theta}{2} (n_x X + n_y Y + n_z Z). \quad (1.16)$$

Notably, the global phase α is not measurable within the two qubit levels $\{|0\rangle, |1\rangle\}$ and can therefore be neglected. θ is the angle of rotation around the axis \vec{n} , here decomposed into the Pauli basis. For example, $R_x(\pi)$ applied to either basis state $|0\rangle$ or $|1\rangle$ performs a bit flip. A frequently used single-qubit gate is the *Hadamard* (H)

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \quad (1.17)$$

that creates an equally weighted superposition from any input basis state $\{|0\rangle, |1\rangle\}$.

Single-qubit gates act locally on the qubit. A local gate can also be applied to multiple qubits, where the tensor product of the individual single-qubit gates describes the unitarity operation of the composite system. In addition, so-called controlled gates act on two or more qubits, with one or more qubits controlling the action on all others. Such multi-qubit gates are necessary for the creation of correlations.

A well-known and frequently used two-qubit gate in this thesis is the *controlled-NOT* (CNOT)-gate, which flips the second qubit if and only if the first qubit is in state $|1\rangle$

$$U_{\text{CNOT}} = |0\rangle\langle 0|^{(1)} \otimes \mathbb{1}^{(2)} + |1\rangle\langle 1|^{(1)} \otimes X^{(2)} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}. \quad (1.18)$$

This operation generates correlations between the qubits since the first qubit controls the action on the second one. Crucially, successively applying H and CNOT on input state $|00\rangle$ prepares a two-qubit GHZ-state also known as a Bell-state[18]

$$\text{CNOT} \cdot H \otimes \mathbb{1} \cdot |00\rangle = \frac{1}{\sqrt{2}} (|00\rangle + |11\rangle). \quad (1.19)$$

Moreover, the combination of arbitrary single-qubit operations with the two-qubit CNOT establishes a so-called *universal gate-set* that in principle can perform any computation including those of a classical computer [18].

In nature, quantum system are faithfully represented by mixed states described with the density matrix formalism presented in Eq. (1.2). So, in a real-world scenario, quantum operations act between mixed states. A general action from an input density matrix $\rho_{\text{in}} \in \mathcal{H}_{\text{in}}$ to an output density matrix $\rho_{\text{out}} \in \mathcal{H}_{\text{out}}$ is thereby called *quantum channel*

$$\rho_{\text{in}} \xrightarrow{\mathcal{E}} \rho_{\text{out}} = \mathcal{E}(\rho_{\text{in}}), \quad (1.20)$$

described by a linear map \mathcal{E} . To ensure that the channel transforms the system to a valid quantum state, \mathcal{E} must be *completely-positive* (CP) and *trace preserving* (TP). Whereas the latter ensures normalization, the CP condition demands positiv maps across the entire system. The partial transpose in the second line of Eq. (1.10) shows the action of a map that is only positive but not CP with respect to the GHZ-state, leading to unphysicalities in the transformed outcome. The formalism of *completely-positive and trace-preserving* (CPTP) maps

is a powerful tool to describe unitary evolution, measurement processes (see Sec. 1.1.4) and even *open quantum systems*. The latter allow environmental mechanisms to affect the system.

In the case of transformation of pure states into pure states, the channel map \mathcal{E} becomes a unitary operator U related to the quantum gates described in Eq. (1.14). Such a unitary channel can be written as

$$\rho_{\text{in}} \xrightarrow{U} \rho_{\text{out}} = U^\dagger \rho_{\text{in}} U \quad (\text{for pure states}). \quad (1.21)$$

Open quantum systems can incorporate environmental noise mechanisms that inevitably affect any system in a real experiment. Unwanted noise potentially leads to a loss of the definite phase relations between the quantum states, denoted by *decoherence*. The evolution of a quantum system suffering from decoherences along all three axes X , Y and Z simultaneously is described by a so-called *depolarizing channel*

$$\mathcal{E}_{\text{depol}}(\rho) = \left(1 - \frac{3p}{4}\right) \mathbb{1} + \frac{p}{4} (X\rho X + Y\rho Y + Z\rho Z), \quad (1.22)$$

which for $p = 1$ results in a completely mixed state $\mathbb{1}/d$. Depolarizing noise exemplifies a non-unitary process describing the loss of quantum information into the environment with the Bloch sphere shrinking towards its center. Other prominent noise channels restrict decoherence effects to either the X, Y -plane or to the Z -axis referred to as *amplitude damping* or *dephasing* channels, respectively. Amplitude damping can for instance model the physical process of spontaneous emission that shrinks the Bloch sphere towards the $|0\rangle$ -state. Dephasing describes the decay in coherent phase relations and restores classical behaviour where the Bloch sphere shrinks towards the Z -axis [18]. Whatever reflects a particular quantum device best has to be decided case-by-case and might even demand for combinations of the above noise models.

1.1.4 Measurement

In quantum mechanics, measurements are destructive providing classical data, and are crucial for testing and manipulating a physical system.

A useful framework for an experimentalist to quantify system properties are *projective measurements*. Formally, projective measurements on a quantum system relate to *Hermitian* (or self-adjoint) operators $M^\dagger M = \mathbb{1}$, denoted as *observables* M . Every observable has a spectral decomposition $M = \sum_m m P_m$, where P_m is the *projector* onto the eigenspace of M with eigenvalue m . The Pauli operators from Eq. (1.5) exemplify such observables. For instance, Pauli Z , which in computational basis $\{|0\rangle, |1\rangle\}$ has eigenvalues 0 and 1 (classical data) referring to the eigenvectors (or projectors) $P_0 = |0\rangle\langle 0|$ or $P_1 = |1\rangle\langle 1|$. Eigenvalues describe the set of possible measurement outcomes whose eigenvectors form a basis of the underlying Hilbert space. Crucially, the predictions of quantum mechanics are of probabilistic nature with the measurement results relating to expectation values of observables. For instance, a projective measurement on state $|\psi\rangle$ yields eigenvalue m with probability

$$p(m) = \langle \psi | P_m | \psi \rangle. \quad (1.23)$$

In case eigenvalue m was observed, the destructive measurement nature projects the post-measurement state to

$$\frac{P(m) |\psi\rangle}{\sqrt{p(m)}}, \quad (1.24)$$

denoting the wavefunction's motion. In practice, observables represent measurable system properties like position, momentum, angular momentum or energy.

Let us generalize measurement processes in quantum mechanics by a so-called *positive operator-valued measure* (POVM). A POVM represents a probability distribution whose values are positive operators instead of positive numbers. Formally, a POVM is a discrete set of M positive-semidefinite operators E_m acting on a d -dimensional Hilbert space \mathcal{H} summing up to the identity

$$\sum_{m=1}^M E_i = \mathbb{1}. \quad (1.25)$$

The set of $\{E_m\}$ represents the measurement outcomes. Whereas the mixed states introduced above describe the reduced state of a larger system, a POVM serves in an analogous way to express the effect of a projective measurement on a subsystem, performed on a larger system. This is stated by Naimark's dilation theorem [28] from 1940, saying that every POVM emerges as a projective measurement in a higher-dimensional Hilbert space. The theorem is of great use as it allows experimentalists to access any POVM from projective measurements on a higher dimensional system.

The probability distribution of any observable E_m can be computed with the underlying density operator ρ according to Eq. (1.2) and yields expectation values

$$p(m) = \text{tr}(E_m \rho). \quad (1.26)$$

In case of a pure state $\rho = |\psi\rangle\langle\psi|$ this expression simplifies to Eq. (1.23). If a POVM comprises d^2 or more elements it spans the space of Hermitian operators and is said to be an *informationally complete* (IC) POVM. Measurement results of an IC POVM completely sample the underlying Hilbert space and allow the reconstruction of ρ , see Ch. 3.1.

Experimentally, every quantum computation concludes with a measurement that reads the qubit register. Let us imagine a single qubit and a measurement performed with the Pauli observable Z , where in the computational basis $\{|0\rangle, |1\rangle\}$ each outcome projects to either the eigenstate $|0\rangle$ or $|1\rangle$. To experimentally infer the underlying state probabilities, the experiment must be repeated n times to account for all events, e.g., the observation of the ground state $n(P_0)$ yielding the state probability $P_0(n) = n(P_0)/n$. The probability to be in the other basis state $|1\rangle$ is therefore $P_1(n) = 1 - n(P_0)/n$. This probabilistic nature must inherently limit the accuracy of quantum measurements, leading to so-called *quantum projection noise* (QPN) [29], which in the single-qubit case amounts to

$$\Delta P = \sqrt{\frac{P_1(1 - P_1)}{n}}. \quad (1.27)$$

Consequently, many samples of the same experiment must be performed to obtain accurate information about the underlying quantum system.

The combined measurement of the three Pauli observables of a d -dimensional system enables us to reconstruct the underlying density matrix ρ , see Ch. 3.1. The measurement set is therefore IC. An alternative but very useful way of extracting the complete information of a system is given by the non-orthogonal SIC POVM depicted in Fig. 1.1(b) [19]. The set consists of four POVM elements per qubit

$$|S\rangle_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, |S\rangle_2 = \begin{pmatrix} \frac{1}{\sqrt{6}} \\ \frac{1}{\sqrt{3}} \end{pmatrix}, |S\rangle_3 = \begin{pmatrix} \frac{1}{\sqrt{6}} \\ \frac{(-1)^{2/3}}{\sqrt{3}} \end{pmatrix}, |S\rangle_4 = \begin{pmatrix} \frac{1}{\sqrt{6}} \\ -\frac{(-1)^{1/3}}{\sqrt{3}} \end{pmatrix}, \quad (1.28)$$

which are maximized by intervector distance resulting in the given symmetric alignment, boosting the information gain per measurement. This holds, as due to their non-orthogonality, every SIC POVM always overlaps with the underlying system in more than one vector, which is not necessarily the case for orthogonal sets. Imagine the characterization of the $|1\rangle$ -state using orthogonal Pauli observables. Measurements along X and Y would yield random outcomes all the time. Crucially, neglected measurements do not contribute to the final statistics. This becomes substantial in large-scale systems, which is why SIC POVMs provide ideal measurements for reconstructing ρ [19, 20].

Let us think of an informationally complete set of measurements on a quantum system, whose results permit the reconstruction of ρ . The quality of the implemented state can then be assessed by measuring how close the reconstructed outcome is to the target state. The measure to serve this need is called *fidelity* and quantifies the overlap between any two mixed quantum states ρ and σ by calculating [30]

$$F(\rho, \sigma) = \left(\text{tr} \left(\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}} \right) \right)^2. \quad (1.29)$$

It results to 1 for identical states and 0 for orthogonal ones and can take any value in between, provided ρ and σ are physical states. For pure quantum states the fidelity simplifies to $F(\rho, \sigma) = |\langle \psi_\rho, \psi_\sigma \rangle|^2$.

1.1.5 Constructing a quantum computer

David Di Vincenzo formulated a catalogue of the basic requirements for the successful operation of quantum computers, which contains the following five criteria [31]

- (1) A scalable physical system with well-characterized qubits
- (2) The ability to initialize qubit states to a simple fiducial state, typically $|0\dots 0\rangle$
- (3) Long relevant coherence times relative to the gate operation
- (4) A universal set of quantum gates
- (5) Qubit-specific measurement capabilities

Let us briefly discuss these five points. (1) Physical multi-level instances are artificially constrained by two qubit states with well-separated energy levels in terms of extant thermal energy. (2) The successful operation of a quantum computation based on unitary evolution critically depends on the reliable preparation of an initial quantum state. Typically, initialization is performed via cooling the system down to its ground-state $|0\dots 0\rangle$. (3) Relative long coherence times compared to the timescale of the quantum gates. Coherence refers to the duration over which phase relations between quantum states remain intact. It is necessary to store quantum information through effects like superposition and entanglement. Finite and unwanted environmental interactions cause so-called *decoherences*, which limit the coherence time and determine the maximum duration for the quantum task. (4) A universal gate-set is comprised of arbitrary single-qubit rotations (Eq. (1.16)) and, for instance, the two-qubit CNOT (Eq. (1.18)) and enables the implementation of arbitrary operations including those of classical devices. Crucially, any algorithm or computation is decomposable into sequences of individual gates from this set. (5) Results on quantum tasks must be accessible through projective measurements on the qubit register. Due to their probabilistic nature, quantum measurements are repeated multiple times to gain statistical accuracy on the respective outcomes.

1.1.6 Higher dimensions

Guided by classical computer science, quantum computers typically build on their binary processor scheme. Yet, there exist many applications to which a higher dimensional structure more naturally applies, holding the potential to notably simplify the computations. Those prominently feature in quantum chemistry [32] or quantum simulations [3]. Moreover, physical quantum information carriers generally have access to more internal states than just the two qubit levels, see Ch. 2. In this thesis, we utilize higher dimensional quantum systems to study two effects: open quantum systems, where unwanted interactions couple a qubit to an environment, effectively enlarging the underlying Hilbert space. Secondly, we engage quantum multi-level systems to store more than the qubit-information per particle. Based on the above qubit definition, d -dimensional information carriers are referred to as *quantum digits* (qudits). For qudit-based quantum information processing a universal gate-set shall also be defined. Fortunately, the above qubit toolbox readily extends to the higher dimensional qudit case after concatenating any two-level interaction to couple between all the states involved. The set of local qudit gates is represented by the Lie algebra $SU(d)$, which in the $d = 2$ case is generated by the Pauli operators from Eq. (1.5). For example, the three-dimensional case denotes so-called *quantum trits* (qutrits). Accordingly, $SU(3)$ is spanned by the Gell-Mann matrices [33] that naturally generalize from $SU(2)$ and read

$$\begin{aligned} \lambda_1 &= \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \lambda_2 = \begin{pmatrix} 0 & -i & 0 \\ i & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \lambda_3 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \lambda_4 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}, \\ \lambda_5 &= \begin{pmatrix} 0 & 0 & -i \\ 0 & 0 & 0 \\ i & 0 & 0 \end{pmatrix}, \lambda_6 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \lambda_7 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -i \\ 0 & i & 0 \end{pmatrix}, \lambda_8 = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -2 \end{pmatrix}. \end{aligned}$$

The form of the Gell-Mann matrices shows the proposed concatenated two-level interaction structure. Gell-Mann matrices represent CPTP maps since they are traceless and Hermitian ($\lambda_i = \lambda_i^\dagger$) for which holds that $\text{tr}(\lambda_i \lambda_j) = 2\delta_{ij}$ with the Kronecker-Delta δ_{ij} . Notably, this extension generalizes to $SU(d)$ [33].

The so far formal quantum computer framework is given a physical meaning in the upcoming Ch. 2, covering the trapped-ion device used throughout this thesis.

1.2 CERTIFICATION & BENCHMARKS OF QUANTUM COMPUTERS

The construction of a quantum computer remains a fundamental scientific and technological challenge due to the high level of control required over delicate quantum systems, exposed to environmental noise. The transition to scalable hardware is marked by rigorous device testing to pinpoint error sources, to gain trust in the devices based on reproducible results, and finally to confirm the correctness of computational outcomes, when the computations become classically untractable. Hereby, the task to ensure the correct outcome of a computation is referred to as *certification*, while quantifying a device's performance, including tasks such as analyzing gate-sets, detection capabilities and others, is known under the term *benchmarking* [34].

The same features responsible for the quantum computational potential, such as the exponential scaling of Hilbert space size, are also making our life difficult with certification

and benchmarking. For instance, an eight qubit density matrix features $2^8 \times 2^8 - 1 = 65535$ degrees of freedom. While so far all experimental demonstrations could be simulated classically beforehand, operating beyond roughly 50 qubits exceeds the limit of today's classical computation capabilities [35]. Hence, specific tools need to be developed that rely on more than just the classical simulations. The nature of such tools can be very different ranging from hardware specific diagnosis tools that require extensive device access all the way to cryptographically-secure verification techniques working completely device-independent. Importantly, these tools differ in the amount of assumptions made and the information provided out of the system to be analyzed. The usual approach to large system characterization is to keep the procedure efficient by either reducing the information gain or building in assumptions about the system that offer relief from the heavy load of measurements and samples to be performed. In contrast, preknowledge might not always be at hand, particularly in the case of future real-world applications. Moreover, a device or the person controlling it might not be trusted. Following those considerations, Fig. 1.2 illustrates a landscape that categorizes some of the most widely used certification and benchmarking tools.

In this thesis we follow two distinct certification paths on opposite ends of the spectrum of Fig. 1.2, which we assess to be of particular importance with respect to both current and future needs. We leave out benchmarking approaches all together since they were subject to another recently published thesis by our group [36]. For a more elaborate review on certification and benchmarks, we guide the reader to Ref. [34].




Figure 1.2: **Landscape of certification and benchmarking quantum computers adapted from Ref. [34].** These tools differ in the extent of information gain and assumptions about the underlying implementations to analyze, and require more or less device access. We qualitatively review these methods in the text to develop the needs for future methods, while we refer to Ref. [34] for detailed insights.

The first certification path we are pursuing is the development of a scalable characterization tool to analyze computational outcomes without any prior knowledge, given that sufficient device access is provided. In particular, this tool will allow us to predict arbitrary properties of the density matrix of a system. This is helpful from a technical point of view

as to evaluate and improve hardware capabilities. From an application perspective, this can be interesting in quantum simulations, where we might want to characterize e.g. the ground-state of some Hamiltonian [37].

The second certification path follows the development of a cryptographically-secure verification technique to confirm the outcome of a quantum computation by purely classical means, without trust in the devices used, and even when the underlying quantum computation can no longer be classically simulated. This becomes a key requirement on large-scale devices, where classical simulations are no longer viable.

Guided by Fig. 1.2, we elaborate on both certification paths over the subsequent sections discussing previously used methods and their scalability bottlenecks, leaving us with challenges to overcome through the work in this thesis.

1.2.1 Characterizing quantum systems

If full access to a quantum device is provided one can perform rigorous system characterizations to certify the device performance or analyze individual setup functionalities such as gate-sets or detection performances.

The gold-standard characterization tool to certify the device performance is so-called *quantum state tomography* (QST) which, in analogy to medical tomography, can draw a picture of a quantum system from a series of measurements [38]. Crucially, QST enables the reconstruction of the density matrix (Eq. (1.2)) and therefore allows one to evaluate every possible property of the state. For instance, the implementation quality can be quantified by calculating the fidelity with respect to the target state (Eq. (1.29)). Alternatively, one could observe its purity (Eq. (1.3)) to make a statement about the degree of *coherence*. On a correlated system, entanglement can be imparted, which can be quantified by measures like quantum negativity (Eq. (1.11)) or Rényi-entropy (Eq. (1.12)). While offering widespread insights, general QST utilizing the Pauli basis requires three times more measurements for every additional qubit added to the analysis. This becomes even more drastic on qudit systems. It takes our trapped-ion device, considered state-of-the-art [39], about four hours to extract eight qubit QST data and about forty hours for the analogous ten qubit register [40]. Notably these numbers are not even considering the classical data analysis part. However, ten qubits are far from sufficient for demonstrating any quantum speed-up. In addition, performing measurements requires operational overhead incurring uncertainties that cannot be separated from computational errors in the final outcome. So-called *state preparation and measurement* (SPAM) errors additionally mask the quality of implementations characterized by QST.

Fortunately, depending on the property of interest, the full density matrix often provides more information than necessary. Less resource-demanding methods circumvent the density matrix reconstruction and offer more efficient estimates, typically suitable for predicting linear observables. For instance, *direct fidelity estimation* [41, 42] compares the experimental outcome with a target state using significantly fewer measurements than QST. Alternatively, *fidelity witnessing* [43] assesses whether a computational outcome is close to a target state, which is slightly weaker in terms of information gain but requires even less measurements to perform. These techniques offer notable relief on the measurement and sampling requirements at the price of less information gain. Other approaches assume that the experimental data follows a certain Hamiltonian whose parameters, being significantly fewer than those of the density matrix, can be deduced from measurements which is referred to as *Hamiltonian learning* [44]. Further relaxations on measurement and sampling demands are offered by assuming well structured states. This lies at the heart of so-called

compressed sensing tomography that for close to pure quantum states yield more accurate descriptions in comparably less samples [45]. Finally, *tensor network tomography*, which eases the heavy measurement load by limiting the amount of entanglement present to the system under test [46].

We put all this effort into tomography methods because they were believed to be the only way to get a non-linear property of the state. Despite all the efforts, however, they all still scale exponentially or make very strong assumptions. Arbitrary predictions of non-linear system properties like purity and notably entanglement require at least the reduced density matrix, where the properties of interest act on. A novel method called *classical shadows* allows the reconstruction of the reduced density matrix from measurements in random bases [27, 47, 48]. Contrastet to existing tomography methods, classical shadows thereby circumvent the full density matrix reconstruction. However, the measurement pool to randomly source from grows exponentially in the subsystem size, which so far limits the method's applicability. While this is a practical limitation of the underlying experimental methods, classical shadows solve the fundamental one. In addition, many scenarios demand the parallel prediction of multiple system properties, such as *variational quantum eigensolver* (VQE) applications [37], that become very extensive under the scope of many characterization tools, even for linear system properties. Also here, classical shadows solve the fundamental problem and leave only the practical one. While characterization tools are imperative for current and future certification applications, existing tools show only limited scalability.

Apart from certifying quantum states, alternative approaches focus on component benchmarking. So-called *randomized benchmarking* (RB) provides a way to quantify gate fidelities [49] by probing numerous random gate-sequences of different lengths that increase sensitivity by amplifying errors and exclude SPAM-errors. Moreover, RB provides decent device comparability. Even more advanced techniques, such as *gate-set tomography* (GST) [50], offer gate-set characterizations, additionally provide measurement errors [51] and draw detailed attention to system limitations from extensive hardware probing. Such routines serve to improve setup functionalities and further establish trust in a device. To date, benchmarking approaches have been experimentally demonstrated up to around ten qubits [52]. While benchmarking is useful to evaluate setup components, extrapolation from component performance to full system performance is only possible with strong assumptions.

Besides the heavy load in measurements and samples to perform on the quantum side, post-processing and analyzing data with the help of classical computers is cumbersome. Indeed, the matrix dimensions to be processed by these classical devices grow exponentially with the number of qubits. For example, processing six-qubit Pauli QST data can demand a memory load of 3100 MB [40]. Processing ten qubits with this method likely consumes all the memory available on today's desktop hardware. Yet another challenge arises from statistical noise. Recall that quantum measurements are destructive in nature and that state probabilities are derived from many samples of the same experiment that are afflicted with QPN. Since rigorous characterizations with existing methods require at least the reduced density matrix, the number of free parameters ($\sim 4^K$ with subsystem size K) is still subject to exponential growth.

Finally, novel certification approaches are urgently needed as state-of-the-art system sizes of about 50 qubits [14, 15] reach far beyond the capabilities of tomography methods. Improvements are needed not only on the quantum side with respect to the measurements to be performed, but also in the classical hardware for post-processing and analysing the data. Let us summarize these critical scalability issues to overcome as three challenges:

- (1) overcome the exponential growth in required measurement settings to characterize ρ
- (2) decrease resource requirements in classical data analysis
- (3) improve statistics to establish accuracy with fewer experimental samples

As shown in Fig. 1.3, we address all three challenges throughout Ch. 3 to provide a scalable and practical approach to characterizing large systems.

In the upcoming section we move away from extensive hardware testing into the direction of verifying computational outcomes, being particularly free of any preknowledge about the implementation and the device that is used. In that sense, we attempt to act in so-called zero-trust settings.




Figure 1.3: **Today's scalability challenges of large-scale system characterizations.** (1) The measurement load of existing characterization methods is extensive and often subject to exponential growth when adding more information carriers. (2) Data analysis is outsourced to classical devices, where resource requirements likely exceed the capabilities of today's desktop hardware—already for a handful of qubits. (3) Statistical noise due to QPN often requires exponentially many experimental samples to accurately represent the underlying quantum system.

1.2.2 Verifying quantum computations

With all the advantages novel quantum technologies entail, their superior computational power might lead to cases where it is no longer possible to verify a computational outcome through classical simulations. This is where *quantum verification* comes into play, where a powerful quantum computer, called a prover, solves a task that a computationally much weaker verifier wants to confirm, using as few resources as possible.

Let us elaborate on this. While for instance factoring a large number into prime factors is believed to belong to the *non-deterministic polynomial time* (NP) class (see Ch. 4.1.2) only verifiable with polynomial memory on a *deterministic Turing machine* [53], it can be solved efficiently on quantum computers, namely by the famous Shor algorithm [9]. Verifying its outcome on the contrary relates to simply multiplying the resulting prime factors back into the large number—being classically efficient. Other quantum computer applications exist, however, for which not only solving the task but also verifying its outcome is expected to be computationally hard for classical verifiers. Examples include the simulations of quantum many-body systems, whose countless degrees of freedom become classically intractable

already at tens of particles to simulate [54]. For such situations, the successful operation presumes trust in the quantum prover. One might gain trust in a device through in-depth system characterizations, offered by tools like RB [52] or QST. Existing characterization tools are typically resource-devouring and require extensive machine access, which might not always be at hand. Importantly, these methods presume that the device under test or the person controlling it are not malicious. Therefore, the task of quantum verification of computational outcomes that cannot be checked by a classical verifier and for which the quantum prover is not trustworthy is designated as one of the key challenges on the road to scalable quantum computation.

Alongside quantum verification an often coined term is *quantum validation*, which we briefly elaborate on to not confuse it with the verification approaches discussed here. In simple terms, validation is to ask “are you building the right thing?”, whereas verification refers to the question “are you building it right?” [55]. In this sense, verification confirms whether or not a quantum system agrees on certain internal specifications and requirements. One might assess the question “are local gate operations truly local?” Verification can thus be seen as an internal controlling procedure. Validation on the contrary is to assure that a quantum system meets the needs of an external party, most likely a user. In this case, one might ask “can it factorize ten-digit numbers?”

Quantum verification has so far been tackled from many perspectives, under additional assumptions and with a need for additional resources [56]. One early approach is so-called *self-testing* [57] that offers a way to verify the presence of a Bell-state by performing a so-called Bell test [58]. An example would be to confirm the existence of entanglement in the Bell-state from Eq. (1.10). Self-testing rules out the possibility for the quantum computer to cheat. Among others, such methods offer a way to verify quantum advantage and are thus of great interest.

One can go one step further and verify the outcome of a quantum computation or even the whole quantum device. A promising approach is the simultaneous execution of a quantum computation on several devices in order to cross-verify their outcomes [59]. Rather than the output of a computation, devices themselves have also been cross-verified [60]. Cross verification is promising because the computation is performed on each individual machine without extra resources for the verification process itself. The absence of additional computational overhead comes at the price of using multiple independent devices ideally based on different physical platforms. To further eliminate the possibility of cheating, the devices involved must be spacelike separated avoiding any communication between them.

Other techniques offer verification with less additional quantum hardware, relying instead on cryptographically-secure verification protocols between the verifier and the much more powerful quantum prover [61]. Such an approach requires limited quantum resources on the verifier side and provides security through quantum communication or joint entanglement between the prover and the much weaker, only partially quantum verifier. Along those lines, the verifier sends public messages to the prover, but the prover cannot read them, because the initial state of the computation is encoded in single qubits prepared by the verifier. In this way, the verifier can choose to test if the prover is cheating, but the verifier cannot check if the answer given by the prover is correct or not. Computing with encrypted data provides the verifier leverage over the prover to establish trust. So-called *blind verification* has been experimentally demonstrated in Ref. [62]. A downside is the need for additional quantum resources and that the verifier can prepare single qubits and send them in a coherent fashion to the prover.

The most powerful verification approach considers a scenario where both device access and additional quantum resources are absent. In other words, a scenario, where one wants

to verify a computational outcome by purely classical means absent of any trust. Running a quantum computer via cloud access is one such situation, an approach which has become very popular these days since it enables researchers all across the globe to realize their quantum applications. For the longest time, classical verification was believed to be intractable in view of the enormous computational potential of quantum computers. Yet, a recent breakthrough in computer science opens an avenue towards classical verification [63]. Similar to blind verification, the idea is that a weak but now purely classical verifier cleverly outplays a much stronger quantum prover to which it outsources a hard task and gains trust through secure message exchange [63]. This novel idea builds on postquantum-secure cryptography. With regards to Fig. 1.2, we can classify it along blind verification, just without the need for extra quantum resources. We work out an approach for classical verification of a quantum computation and experimentally demonstrate its capabilities throughout Ch. 4.

1.3 QUANTUM COMPUTING WITH FAULTY COMPONENTS

A quantum system can never be completely separated from the environment, so the systems are always prone to errors. In addition, errors may occur due to system instabilities or imperfect calibration routines. Fortunately, quantum computations can be protected from errors through the use of so-called *quantum error correction* (QEC) codes [64], whose study has recently been of major interest in the field. While the following explanations provide a basic understanding to the QEC concepts used throughout this thesis, the interested reader finds further insights in Ref. [65].

1.3.1 *Quantum error correction*

In classical computer science, there is an established method for error correction based on the idea of achieving robustness through redundancy, where a majority vote denotes with high probability the correct computational outcome, while allowing for certain errors to happen. At first glance this appears fruitless on quantum devices since the no-cloning theorem [66] strictly prohibits the duplication of quantum information in the way classical computers do. This is additionally compounded by projective measurements (Sec. 1.1.4), so one cannot simply measure qubits to check for errors. Here, quantum physicists play a trick. Instead of utilizing multiple copies of the same qubit state, a state is encoded across multiple qubits through entanglement. Such a scheme was first proposed by Peter Shor in 1995 [67]. For simplicity, we will restrict the following QEC discussion to qubit systems, while similar ideas exist for the general qudit case as well [68].

An encoded block of physical qubits represents a so-called *logical qubit* establishing redundancy. QEC codes are ideally constructed in such a way that computations with their logical qubits are readily applicable to the quantum computing framework introduced in Sec. 1.1. Moreover, a QEC code ideally shows robustness against errors not only during state preparation, gate operations and readout, but also for the operations necessary for error correction. In this case, the QEC code is said to be *fault-tolerant* [18, 64, 69] which in principle paves the way for executing arbitrary computations.

Let us elaborate on this. A key property for achieving fault-tolerance is that local errors on individual physical qubits do not propagate to other physical qubits from the same logical block. To prevent such error propagation, codes can be designed so that a particular local gate applied individually to all physical qubits of the same logical block yields the

corresponding *logical gate* on the logical qubit. Following these means, certain entangling operations are applied pair-wise between one physical qubit per logical block. The logical entangling operation follows after concatenating all such pairs [70]. This local and bitwise application of gates, each targeting only one qubit per block, is called *transversality* [71] and reduces error propagation. A promising way of performing QEC is by designing codes that support as many logical gates as possible in a transversal fashion. This significantly reduces experimental overhead and the potential for problematic error propagation [72, 73].

Let us take a look at the kind of errors that can occur at the qubit level and what is required to correct them. Similar to classical bits, qubits can suffer from bit flip errors. Besides bit flip errors, the extra phase information stored in qubits provides an additional error source not present to classical bits. Consider hereby the arbitrary single-qubit state $\alpha|0\rangle + \beta|1\rangle$ from Eq. (1.1) that, with α and β complex numbers, can in principle experience an infinite number of errors. Fortunately, as will be pointed out below, the error correction process in modern QEC applications involves destructive measurements that project qubits onto binary outcomes, yielding one error at a time [74]. For individual measurement samples, continuous errors become digitized errors, which result in bit and phase flips, designated as *computational errors* [65]. This digitized error approach indeed suffices to explain all possible noise mechanisms at the qubit level. Those involve *coherent* errors due to systematic drifts, e.g., miscalibrations as well as *incoherent* errors from unwanted environmental interactions [75].

The number of correctable computational errors differs from code to code and strongly relates to its size, i.e., the number of physical qubits encoded per logical block. The correction of at least one error requires a minimum encoding of five physical qubits, which is given by the *quantum Hamming bound theorem* [76]. In the example of the original Shor code [67], nine physical qubits per logical block serve to correct a single computational error. Notably, Kitaev's famous surface code [77, 78], considered state-of-the-art in the field and exemplified in the next paragraph, offers high scaling capabilities because of a very modular structure that can be arbitrarily extended in size. While growing in size, the number of correctable errors or equivalently the tolerable error-rates on physical qubits increases. Following this notion, an important finding of fault-tolerant quantum computation is the *threshold theorem* [64, 69, 79–81]. It states that if the physical error-rate is below the threshold of the QEC code, errors can be suppressed arbitrarily.

Another challenge in QEC is the detection of errors via destructive measurements, see Sec. 1.1.4. While in classical error correction bit registers can be arbitrarily read out to detect errors, doing so in the quantum case destroys the underlying logical information. We will exemplify in the next section how this measurement problem can be circumvented based on so-called *stabilizer codes* that prominently feature all across today's leading QEC applications [65, 67, 72, 79, 82]. Stabilizer codes uniquely allow the detection of errors without affecting the underlying logical information—being stabilized in that sense. Although destructive measurements on qubit states would alter the logical information, quantum mechanics does not forbid all kinds of measurements. Indeed, one can construct a measurement to extract only the error information, which, in contrast, provides no information about the qubit states. Therefore, the error information of code qubits is coupled to ancilla qubits by means of entanglement. The state of the ancilla qubit then contains the error information and can be read out destructively without affecting the logical information of the code qubits. We note that this useful algorithmic structure is often applied in the field of quantum computation and allows one to measure certain information

via ancilla qubits, referred to as *quantum non-demolition* (QND) measurements [83]. We will make use of QND measurements throughout Ch. 5.3.

We have thus far seen that QEC codes conceptually differ from their classical analogues due to various mechanisms intrinsic to the nature of quantum mechanics. Those include the no-cloning theorem, the qubit's extra phase information as well as the destructive measurement character in quantum mechanics. Generally, the basic idea behind QEC is to find logical code words $|0_L\rangle$ and $|1_L\rangle$ entangled across blocks of N physical qubits that form an error robust logical state $|\psi_L\rangle$. To increase the potential for achieving fault-tolerance, the logical code words should allow the realization of as many transversal gates as possible.

QEC codes are classified by the triplet $[N, k, d]$, where N physical qubits form k logical ones with a so-called *code distance* d [65]. The code distance refers to the Hamming distance between code words, which corresponds to the minimum error size that is no longer detectable. In other words, this minimum size error refers to a logical operator that transforms one code word into another—a fatal error which cannot be distinguished from a logical operation. The correction process in QEC is based on majority voting. The code distance therefore relates to $(d - 1)/2$ correctable and $d - 1$ detectable errors. For example, a distance three code can thus correct one and detect two computational errors. The nine-qubit Shor code exemplifies a distance three code, classified by $[9, 1, 3]$ [67].

1.3.2 Stabilizer codes

Let us define a promising class of QEC codes, named stabilizer codes [65, 67, 72, 79, 82], which will play a prominent role in the remainder of this thesis.

To this end, imagine a state $|\psi\rangle$ which is in the $+1$ eigenstate of an operator S_i , where it holds that $S_i |\psi\rangle = +1 |\psi\rangle$. In this case, the operator S_i is said to *stabilize* the state $|\psi\rangle$. Suppose further that a set of states $\mathcal{C} = \{|\psi_i\rangle\}$ is stabilized by the group of operators $\langle S \rangle$. Each element $S_i \in \langle S \rangle$ is then referred to as *stabilizer operator*. We refer to the minimal set of elements that can construct the group $\langle S \rangle$ as the so-called *generators* S . For example, the single-qubit Pauli group can be constructed from the three generators $\{i\mathbb{1}, X, Z\}$, resulting in the $2^3 = 8$ group elements $\langle i\mathbb{1}, X, Z \rangle = \{\pm\mathbb{1}, \pm i\mathbb{1}, \pm X, \pm iX, \pm Y, \pm iY, \pm Z, \pm iZ\}$. Importantly, the stabilizers $\langle S \rangle$ form a subgroup of the N -qubit Pauli group that act trivially on the vector space $\mathcal{C} \subset \mathcal{H}_2^{\otimes N}$ as follows [18]

$$S_i |\psi\rangle = +1 |\psi\rangle \quad \forall S_i \in \langle S \rangle, \quad (1.30)$$

where all stabilizers $\langle S \rangle$ mutually commute.

Let us now discuss how stabilizers can serve in error correction. Assume an $[N, k, d]$ code where N physical qubits represent k logical ones. An N -qubit stabilizer code can therefore be defined by the $+1$ eigenspace $\mathcal{C} \subset \mathcal{H}_2^{\otimes N}$ stabilized by $S = N - k$ generators [65, 79, 82]. To make this clear, let us turn the argument around. An N -qubit logical state then lives in the $+1$ eigenspace of S stabilizer generators and thus has $k = N - S$ degrees of freedom. These k degrees of freedom represent logical qubits. Note that in QEC we denote the $+1$ eigenspace \mathcal{C} as *code-space*.

Moreover, an $[N, k, d]$ code requires $2k$ logical operators to act non-trivially on the code-space. In particular, we need to define $\{Z_L, X_L\}$ operators for each of the k logical qubits that obey the same commutation relations as the Pauli operators for a single qubit. While logical operators commute with all stabilizers, they must anticommute among themselves [65].

Crucially, the stabilizers make it possible to construct the logical code words in a systematic way by applying the projector to the +1 eigenspace of the stabilizer generators to the ground-state [65]

$$|0_L\rangle = \frac{1}{n} \prod_{S_i \in \mathcal{S}} (\mathbb{1} + S_i) |0\rangle^{\otimes N}, \quad (1.31)$$

where n ensures normalization. The other logical basis state $|1_L\rangle$ can be prepared after applying $X_L |0_L\rangle = |1_L\rangle$.

Stabilizers can be measured in a QND fashion after relaying information about their eigenvalue to ancilla qubits. In the absence of errors, every stabilizer S_i applied to a logical code word $|\psi_L\rangle$ projects to the +1 eigenstate by means of Eq. (1.30). So the stabilizers do not induce any action on the logical state. So the logical encoding remains intact, and the quantum task on the logical qubits can continue unchanged. On the other hand, -1 outcomes indicate the presence of errors. Crucially, the measurement outcomes of all stabilizers together allow one to identify which single error has happened, and where in the code it happened. Detected errors can subsequently be corrected by applying the error inverse.

Many experimental efforts all across the leading quantum architectures thus far demonstrate building blocks on the detection and correction of computational errors based on various codes. Out of the many, we highlight the following Refs. [14, 70, 84–93] ranging from proof-of-principle studies all the way to the demonstration of fault-tolerant universal gate-sets.

1.3.3 The surface code

To complement the so far conceptual discussion let us exemplify a state-of-the-art QEC protocol, namely Kitaev's surface code [77, 78]. We hereby follow explanations in Ref. [65].

The surface code generates a set of stabilizers from physical qubits aligned on the edges of a 2D square lattice representing the logical qubit, illustrated by Fig. 1.4(a). Four-body Pauli X- and Z-type stabilizer operators are geometrically aligned with vertices V and plaquettes P of the lattice and defined as follows

$$S_V^X = \prod_{j \in V} X_j \quad \text{Vertex (V)} \quad S_P^Z = \prod_{j \in P} Z_j \quad \text{Plaquette (P)}. \quad (1.32)$$

As neighboring plaquettes share two vertices, and neighboring vertices share two plaquettes, all stabilizers mutually commute. The code features a modular structure with the lattice not being bound to a fixed size. Rather, it can be enlarged in favour of increasing the error robustness, while keeping the stabilizer's next-neighbor interaction structure. The logical code-space is constructed from the stabilizer's common +1 eigenstates

$$S_V^X |\psi_L\rangle = S_P^Z |\psi_L\rangle = +1 |\psi_L\rangle \quad \forall V, P. \quad (1.33)$$

In case of an $[N, k, d]$ code, the logical code words can be constructed by applying the projector on the +1 eigenspace of the stabilizer generators to the ground-state

$$|0_L\rangle = \frac{1}{n} \prod_{S_P^Z \in \mathcal{P}} (\mathbb{1} + S_P^Z) \prod_{S_V^X \in \mathcal{V}} (\mathbb{1} + S_V^X) |0\rangle^{\otimes N}, \quad (1.34)$$

with a normalization factor n .

The surface code defines logical operators by connecting opposite surface boundaries with strings of Pauli operators. Particularly, connecting a string of vertical (horizontal) lattice edges with Pauli Z (X) operators denotes a logical operator Z_L (X_L). To see this, consider that logical operators mutually anti-commute, but commute with the stabilizers, while, importantly, not being stabilizers themselves. Thus, Z_L and X_L act non-trivially within the code-space denoting logical operators. Importantly, their definition is not unique as any path connecting opposite lattice boundaries serves equivalently as logical operator [94]. We will make use of this equivalence to circumvent inoperable qubits in Ch. 5.




Figure 1.4: **The surface code and the correction of computational errors.** (a) Physical qubits align on the edges of a 2D square lattice and form a logical qubit. A complete set of stabilizers is given by local four-body Pauli X- and Z-type operators forming vertices S_V^X and plaquettes S_P^Z of the lattice, respectively. Vertices and Plaquettes are regularly spread across the entire lattice and mutually commute. Connecting two opposite lattice edges with Pauli operators yields logical operators. In particular, vertical (horizontal) strings of Pauli Z (X) yield Z_L (X_L). The stabilizers' common +1 eigenstates define the code-space. Crucially, the surface code's topological structure offers hardware-friendly scaling properties, see text for details. (b) Errors can be read out by continuously measuring all stabilizers, where a shared -1 outcome uniquely pinpoints the error location, which is corrected after applying the inverse of the error.

Based on their periodic structure and next-neighbor interactions, surface codes are classified to the family of so-called *topological* QEC codes [78]. In general, an $L \times L$ lattice supports $L^2 - 1$ stabilizers and indicates an $[[L^2, 1, L]]$ code. Expanding the surface by adding more physical qubits increases the error robustness, while at the same time the mutual commutation of all stabilizers remains as they act locally.

From a hardware perspective the next-neighbor interaction structure is a recognized feature of the surface code as long-range interactions are typically more costly in experimental realizations, including trapped atoms [95], superconducting hardware [96], quantum dots [97] or some trapped-ion approaches [98]. As such, the surface code is a prominent choice for both theoretical and experimental QEC studies, and has thus far been subject to multiple experiments [99–101]. Moreover, the surface code yields tolerable error thresholds at the 10^{-2} level, considered one of the best in that respect [102].

The minimal surface code instance that allows error correction is given by a 3×3 lattice based on nine physical qubits. In that case, the shortest line connecting top to bottom and left to right encounters weight three logical operators resulting in a distance three code. The respective $[[9, 1, 3]]$ -code is robust against a single computational error. Error correction can be achieved upon continuously measuring all stabilizers. Doing so, the eigenvalues

of stabilizers are relayed to ancilla qubits and read out in a QND fashion to not affect the logical information in the absence of errors. Multiple neighboring -1 outcomes uniquely pinpoint the error location, called *syndrome*, illustrated by Fig. 1.4(b). The subsequent correction step applies the error inverse to the faulty qubit and restores the code.

1.4 QUBIT LOSS & LEAKAGE—BEYOND COMPUTATIONAL ERRORS

In quantum computation, the smallest unit of information is formally considered a qubit, representing an idealized quantum mechanical two-level system in analogy to the classical bit. Physical realizations of quantum information carriers, however, rely on systems that exhibit a natural multi-level structure. Ideally, only a subset of those levels, for the qubit it would be two, is manipulated. All explanations given so far rely on this binary approach including the basics of quantum computing (Sec. 1.1), alongside characterization and verification methods (Secs. 1.2.1-1.2.2), as well as schemes for error mitigation (Sec. 1.3).

In a more realistic scenario, a qubit description incorporates several layers of abstraction as illustrated by Fig. 1.5. The multiple levels provide a realistic understanding of a physical qubit implementation giving rise to additional error mechanisms—notably beyond computational ones. In the absence of perfect control, errors may propagate across the layers of Fig. 1.5, potentially leading to a loss of quantum information. For instance, so-called *leakage errors* can occur that couple the qubit to levels outside the computational subspace. Beyond leakage, qubits might get lost altogether. Thus, any erroneous mechanism that transfers a qubit to beyond its two levels or one that makes the qubit inoperable results in what is called *qubit loss*. Note that these arguments also apply to higher dimensional qudits.

On the trapped-ion device considered here, qubit loss manifests itself in a variety of physical incarnations such as the actual loss of particles encoding the qubits, or chemical reactions under which qubits become inoperable. Such hard loss mechanisms almost never occur on experimental timescales as particles can be stably trapped even for days [103]. Chemical reactions due to collisions with background gas are furthermore suppressed by operating the trap at ultra-high vacuum, see Ch. 2. To seek out the limiting loss mechanism we have to take a closer look at the experimental implementation of quantum tasks. Crucially, the experimental performance can often be simplified by addressing higher dimensional states, either to spectroscopically decouple certain constituents (e.g. qubits) from subtasks or to simplify the underlying quantum circuit, see Ch. 2. However, errors in the coherent operations connecting to levels outside the computational subspace result in leakage. Leakage errors can therefore be expected at rates similar to computational errors [52], making their detection and correction inevitable to achieve fault-tolerance.

Apart from trapped-ion devices, all the leading quantum architectures utilize physical multi-level constituents to encode their qubits and thus suffer a realistic chance of leakage. Examples are atomic [104] and molecular systems [105], Rydberg-ions [106], solid-state systems [107] and superconducting hardware [108]. Only for some of them, e.g. photonic systems, harder loss mechanisms, such as the actual loss of the particles encoding the qubits, become relevant on experimental time scales [109].

The influence of loss errors has thus far been subject to multiple conceptual studies. Regarding quantum communication channels, work in Ref. [110] finds a tolerable loss rate of $p_{\text{loss}} < 0.5$ under a simultaneous depolarizing rate of $p_{\text{depol}} < 0.3$ (see Eq. (1.22)). Moreover, a so-called *one way quantum computer* is proposed in Ref. [111] that inhibits unique loss robustness by allowing rates as high as $p_{\text{loss}} < 0.5$.

These tolerable loss rates turn out to be comparatively high contrasted to the thresholds for computational errors given in QEC. For example, the surface code yields computational




Figure 1.5: **Multiple layers of abstraction illustrating the influence of noise at the example of trapped-ion qubits.** A combination of several qubits is used to redundantly encode logical quantum information using QEC, in particular to protect the qubits from unwanted interactions with the environment, see Sec. 1.3 and Sec. 5.4.1. The computational Hilbert space is encoded by the two-levels $4^2S_{1/2}(m = -1/2) = |0\rangle$ and $3^2D_{5/2}(m = -1/2) = |1\rangle$ embedded in the Zeeman structure of $^{40}\text{Ca}^+$, which in turn denotes the physical Hilbert space, see Sec. 2. Absent of perfect control and beyond the scope of many QEC applications, errors can occur at all levels indicated by the figure. Crucially, some of these errors will couple across the respective spaces, referred to as leakage. Depending on the employed QEC codes, these errors may be correctable, or result in a complete loss of information. Moreover, ions can become inoperable or disappear completely from the trap due to collisions with the background gas, which is fatal for the computation, see text.

error thresholds at the 10^{-2} level and is considered one of the best in this respect [102]. This raises the question of how QEC applications behave under the additional and foremost realistic presence of loss errors. The authors of Ref. [112] followed exactly this question and numerically quantified the tolerable parameter regimes for correctable loss and computational errors on topological QEC codes and find both rates to delicately depend on each other. In particular, their schemes combine losses in the surface code [94, 113] with methods from Raussendorf's topological scheme [114–116]. While in the limit of vanishing computational errors, loss errors become tolerable up to $p_{\text{loss}} = 0.25$, a relatively low computational error-rate of $p_{\text{comp}} = 0.005$ already restricts losses to $p_{\text{loss}} = 0.05$ and when computational errors exceed $p_{\text{comp}} > 0.006$, losses are no longer correctable. A more hardware-related study from Ref. [117] quantifies a scalable scenario for fault-tolerant photonic quantum computers that yields correctable loss rates of $p_{\text{loss}} < 0.003$, given the presence of depolarizing noise at a rate of $p_{\text{depol}} < 0.001$.

In QEC applications, loss and computational errors are thereby naturally interwoven, so that addressing only one of them drastically limits the potential for achieving fault-tolerance. Although offering certain stability, most QEC schemes can neither detect nor correct arbitrary losses and instead require additional protocols on top. The occurrence of

losses typically results in fatal errors that deteriorate the quality of the underlying logical information [118, 119].

The detection and correction of qubit loss manifests an inevitable and so far often neglected challenge. We will tackle this challenge throughout Ch. 5 with an in-depth experimental study on account of multiple loss detection and correction instances. These studies are complemented by numerical simulations that estimate realistic parameter regimes for correctable loss and computational errors, with a particular focus on the capabilities of current NISQ-hardware.

1.5 SEMI-CLASSICAL QUANTUM ALGORITHMS

In many modern quantum computational tasks, the unitary evolution gets interrupted multiple times by in-sequence measurements, with the subsequent part of the computation depending on the classical measurement result. Crucially, this operational structure imparts projective measurements that can lead to dynamics featuring non-unitary components, which are not captured by the quantum channel description introduced in Sec. 1.1.3. Ignoring these non-unitary mechanisms bears the risk of missing faulty components, that might severely degrade the quality of the underlying logical information. As a consequence of this non-unitarity, existing characterization tools become invalid and require adaptations.

Let us illustrate such a computational structure using the example of error syndrome detection with the stabilizer formalism in QEC, see Sec. 1.3. Because of the destructive nature of quantum measurements, the stabilizers are read out via coupling the code qubit register with the ancilla qubits in QND fashion [65]. Whereas in the absence of errors the logical encoding remains intact, the presence of an error triggers a correction step, manifesting the concept of classical *feed-forward*. In both cases, the measurement potentially interrupts the unitary evolution as discussed above. It should be noted that such correction of multiple in-sequence detected error events has thus far been subject to multiple experimental studies [120, 121].

Recent advances in quantum devices have led to ever so complex implementations, which increasingly feature the non-unitary elements of in-sequence measurement and classical feed-forward and give rise to a novel class of so-called *semi-classical quantum algorithms*, illustrated by Fig. 1.6. Given the huge capabilities this operational structure offers, in-sequence detection represents a key challenge in current quantum hardware developments. A development that grants access to a broad class of algorithms beyond QEC featuring in quantum network [122], quantum causality [123], measurement uncertainty trade-offs [124, 125], and weak measurements [126–128]. The reliable operation of all these depends on the correct characterization tools, which require conceptual adaptations to those previously presented in Sec. 1.2.

The correct description of semi-classical quantum algorithms is given by so-called *quantum instruments* [129, 130], a formal construct to capture both the classical and the quantum degrees of freedom. Formally, a quantum instrument describes a collection of CP maps \mathcal{E}_j labelled by a classical index j that together are TP, defined as follows

$$\{\mathcal{E}_j\}_{j \in I} \quad \text{with} \quad \text{tr} \left(\sum_{j \in I} \mathcal{E}_j(\rho) \right) = \text{tr}(\rho). \quad (1.35)$$

Each individual linear map \mathcal{E}_j of a quantum instrument describes a *side-channel* associated with a certain classical measurement outcome $j \in I$ where I is the countable set of all possible outcomes. \mathcal{E}_j is considered CP, and importantly, only trace-nonincreasing while not necessarily trace-preserving. As such, the trace to every side-channel j either preserves

or decreases its value throughout the quantum task. Yet, only together as a whole the countable set of side-channels for all positive density operators ρ is CPTP again [131].

The quantum instrument structure allows for non-unitary evolution, particularly in the sense of trace-decreasing side-channels, and thus offers a way to capture the dynamics of arbitrary semi-classical quantum tasks. By that quantum instruments provide a description to the most general quantum operations.




Figure 1.6: **Semi-classical quantum algorithms.** Many of today's leading quantum tasks contain multiple interruptions from in-sequence measurements, where the measurement information is processed classically and the task continues conditional on these outcomes. More precisely, the code register is read out in a QND fashion through coupling to ancilla qubits via entanglement. The relevant part of the ancilla qubit register can subsequently be read out destructively without altering the quantum information in the code register. Depending on the classical measurement outcome, a conditional unitary operation is subsequently applied to the code register denoting the concept of classical feed-forward. As a consequence, measurement outcome related side-channels might follow a non-unitary evolution path that can no longer be captured with existing characterization tools that suffice the quantum channel description introduced in Sec. 1.1.3. Rather, the correct framework to describe the dynamics of these semi-classical algorithms is given by quantum instruments, see text.

We will dive into the characterizations of quantum instruments as part of our qubit loss correction studies in Ch. 5.1. In particular, the correction of losses requires their reliable detection beforehand. The development of a loss detection unit will be described in Ch. 5.1. In analogy to the error syndrome readout in QEC, its working principle relies on mapping a code qubit's information about qubit loss to an ancilla qubit using entanglement, whereafter the ancilla qubit is read out without altering the logical information. Such QND loss detection exemplifies a special case of a semi-classical algorithm. Considering the importance of its correct characterization as part of our loss studies and even beyond, especially in view of the importance of semi-classical quantum algorithms for modern quantum computation, we develop a novel toolbox for quantum instrument tomography at the bottom of Ch. 5.

 THE TRAPPED-ION QUANTUM INFORMATION PROCESSOR

The present chapter gives physical meaning to the so far formal quantum computing framework of Ch. 1 by introducing the experimental setup used in this thesis. A key factor in the realization of a quantum computer is the encoding of the information carriers in physical systems. One of the most promising approaches along those lines are atomic ions, where a certain class of isotopes are always the same and do not suffer from manufacturing flaws. The achievable level of control over ions is thereby only limited by very few fundamental factors.

Ion traps provide spatial confinement of microscopic charged particles [132]. Operated in ultra-high vacuum to prevent collisions with background gases and shielded from unwanted interactions with the environment, such as the Earth's magnetic field, ions enable access to their internal quantum degrees of freedom. In parallel, the development of narrow bandwidth lasers [133] enable their coherent manipulation and detection that together pave the way towards building quantum information processors [134, 135].




Figure 2.1: **Ion trap quantum information processor using $^{40}\text{Ca}^+$.** (a) Each ion along the axial trap direction encodes a qudit with up to seven levels, or a qubit, when restricted by the two levels highlighted in (b). A universal gate-set is realized upon coherent laser-ion interaction using tightly focused laser beams addressing single ions for local gates (bright red) or pairwise ions for entangling gates (dark red), see 2.3. Alternatively, a global beam (not shown) allows for collective operations, mostly used for state preparation and cooling purposes that are imperative for coherent qudit manipulations [136]. Readout is performed via collective fluorescence detection (DET) [87]. (b) Simplified level scheme of relevant energy levels in $^{40}\text{Ca}^+$ -ions suitable for cooling, coherent manipulation and detection of qudits, see text for details [137, 138].

All experiments presented in this thesis are based on qubits or qudits encoded in the internal electronic level structure of trapped $^{40}\text{Ca}^+$ -ions. The setup in Fig. 2.1(a) illustrates an ion trap in which an electric field generates a potential minimum along the axial trap

direction to store and control a crystal of about ten to 15 ions, spatially separated by their mutual Coulomb repulsion and manipulated with coherent and tightly focused laser light.

We give here a brief overview of our setup with special emphasis on the toolbox for quantum computation with trapped ions, which is necessary for understanding the applications presented in the remainder of this thesis. For a more detailed treatment of the experimental setup, we refer to the thesis of A. Erhard in Ref. [36].

2.1 THE $^{40}\text{Ca}^+$ -QUDIT

Calcium comes from the second main column of the periodic table and has a simplified hydrogen-like level structure after ionization. The relevant parts for quantum information processing are illustrated in Fig. 2.1(b). $^{40}\text{Ca}^+$ is a promising candidate as it offers cooling, detection and state manipulation at transition frequencies in the optical or near-infrared regime, where lasers are commercially available.

In particular, the dipole transition $4^2S_{1/2} \leftrightarrow 4^2P_{1/2}$ at 397 nm offers fluorescence detection and *Doppler cooling* (DC) capabilities owed to a short-lived upper energy level with *relaxation time* $T_{[4^2P_{1/2}]}^1 \approx 7.10(2)$ ns [137]. Notably, the $3^2D_{3/2}$ level is metastable with $T_{[3^2D_{3/2}]}^1 \approx 1.168(9)$ s [138] as its only decay channel to the ground-state $4^2S_{1/2}$ is dipole forbidden by selection rules [139]. Spontaneous decay by the transition $4^2P_{1/2} \rightarrow 3^2D_{3/2}$ therefore potentially retains population in $3^2D_{3/2}$, which must be simultaneously re-pumped for keeping DC and detection efficiencies high. Hence, only the two transitions $4^2S_{1/2} \leftrightarrow 4^2P_{1/2}$ and $4^2P_{1/2} \rightarrow 3^2D_{3/2}$ together form a so-called *closed loop cycling transition*.

$^{40}\text{Ca}^+$ features another metastable level $3^2D_{5/2}$ with a relatively long relaxation time $T_{[3^2D_{5/2}]}^1 \approx 1.176(11)$ s [138], whose transition to the ground-state lies in the optical regime at 729 nm. Crucially, this narrow bandwidth *quadrupole transition* $4^2S_{1/2} \leftrightarrow 3^2D_{5/2}$ together with the dipole transition $3^2D_{5/2} \leftrightarrow 4^2P_{3/2}$ at 854 nm with lifetime $T_{[4^2P_{3/2}]}^1 \approx 6.92(2)$ ns [137] forms another closed loop cycling transition. It enables resolved *sideband cooling* (SBC) and state preparation purposes [136].

Moreover, the quadrupole transition $4^2S_{1/2} \leftrightarrow 3^2D_{5/2}$ is particularly well suited for coherently manipulating and storing quantum information. In the presence of a bias magnetic field, the degeneracy of the fine structure lifts and the Zeeman manifolds of the transition allow the encoding of a qudit [140] or a qubit which is typically subject to the magnetic field insensitive components $4^2S_{1/2}(m = -1/2) = |0\rangle$ and $3^2D_{5/2}(m = -1/2) = |1\rangle$ [141], see Fig. 2.1(b). In view of state detection with the dipole transition $4^2S_{1/2} \leftrightarrow 4^2P_{1/2}$ only the lower energy manifold $4^2S_{1/2}$ of the qudit couples resonantly and therefore fluoresces, denoting bright states, while the upper energy manifold $3^2D_{5/2}$ indicates dark states [87]. Through a series of such binary fluorescence detections, a total of seven of the eight states can be identified [140]. An eighth level cannot be resolved as the two ground-states $4^2S_{1/2}(m = \pm 1/2)$ are not separable in the fluorescence detection process.

2.2 THE LINEAR PAUL TRAP

Our setup utilizes a macroscopic linear Paul trap [132] developed by S. Gulde in Ref.[103] that has been successfully operated within ultra-high vacuum for the past two decades. A trap photograph is depicted in Fig. 2.2(a) alongside its construction drawing in Fig. 2.2(b).

This trap has four orthogonally aligned blade electrodes with a length of 30 mm, which together generate an alternating quadrupole electric field for radial confinement. The diagonal center distance of the blade electrodes is 1.6 mm. Axial confinement is provided by two opposing end-cap electrodes maintained at constant voltage and separated by 5 mm. The resulting electric field produces a radial potential minimum along the axial trap direction, in which the ions form a linear string, confined together axially by the constant electric field of the end-cap electrodes, counteracting the mutual Coulomb repulsion between the ions. Additional compensation electrodes along the axial and radial trap directions (not shown) ensure center alignment of the ion string and compensate for construction imperfections.




Figure 2.2: **The Innsbruck trap.** (a) Photograph of the macroscopic linear Paul trap [103] utilized in this thesis. (b) Construction drawing from the trap in (a) with measures given in mm, see text and Ref. [103] for details.

To establish the quadrupole potential for radial confinement, a diagonally opposite pair of blade electrodes is typically powered with radio frequency $\Omega_{\text{trap}}/2\pi = 23.5 \text{ MHz}$ and a voltage amplitude of several 100 V, while the other pair is connected to ground potential. End-cap electrode voltages are around 1 kV. Some freedom applies to these parameters, which are useful to adjust the inter-ion spacing. While moving the string further apart is useful to suppress cross-talk [52] to neighboring ions, increasing trap confinement, on the other hand, improves the potential of laser cooling [136]. The ideal setting has to be chosen on a case-by-case basis and typically depends on the number of ions operated.

The motion of the ions in the resulting trap potential can be approximated by harmonic 3D oscillations around its minimum, which can be decomposed into radial and axial components, the so-called *secular motion*. These secular motions are additionally modulated with the trap drive frequency Ω_{trap} , which is called *micromotion*. Micromotion is typically suppressed by moving the ions to the potential minimum using the compensation electrodes. The above trap parameters yield secular frequencies in the radial direction around $\omega_{\text{radial}}/2\pi \approx 3.4 \text{ MHz}$ and in the axial direction around $\omega_{\text{axial}}/2\pi \approx 1.2 \text{ MHz}$ [141]. A detailed discussion of the equations for the motion of ions in a harmonic trap potential can be found in Ref. [134]. In the following we will restrict our view to a quantum mechanical description of the ions in axial direction, whereby radial directions and micromotion are neglected. In particular, we use the axial center-of-mass mode of the ions to create coupling between them, which is necessary for the generation of entanglement [142], as outlined in the next paragraph.

Finally, the trap is situated inside a vacuum chamber at a gas pressure of around $2 \times 10^{-11} \text{ mbar}$ reliably protecting the ion-crystal from unwanted collisions with background gas—notably reduced to minute timescale [103].

2.3 LASER-ION INTERACTION

Preparing the ground-state, implementing quantum gate operations and finally performing state detection all rely on the interaction between the trapped ions and the laser light [133]. In the following, we derive the necessary interaction terms that form the basis for the universal gate-set, introduced in Sec. 2.4. Explanations are based on Ref. [143] to which we further guide the interested reader.

Let us describe the ion by an effective two-level system based on the narrow bandwidth transition $4^2S_{1/2} \leftrightarrow 3^2D_{5/2}$ with frequency $\nu_0 = (E_D - E_S)/\hbar$. This effective two-level system, together with the ion-motion in the harmonic oscillator trap potential, forms the joint system shown in Fig. 2.3(a). The manipulation of the electronic levels $4^2S_{1/2}$ and $3^2D_{5/2}$ with the laser light additionally enables interaction between different eigenstates of the trap potential, identified by the *motional quantum number* n . The resulting interaction Hamiltonian between laser and joint system can be written as follows

$$\begin{aligned}\hat{H} &= \hat{H}_0 + \hat{H}_1 \\ \hat{H}_0 &= \frac{\hat{p}^2}{2M} + \frac{1}{2}M\omega_{\text{axial}}^2\hat{x}^2 + \frac{1}{2}\hbar\nu_0 Z \\ \hat{H}_1 &= \frac{1}{2}\hbar\Omega(\hat{\sigma}^+ + \hat{\sigma}^-) \left(e^{i(k\hat{x} - \nu_L t + \phi)} + e^{-i(k\hat{x} - \nu_L t + \phi)} \right)\end{aligned}\quad (2.1)$$

with Pauli spin operators $\hat{\sigma}^+ = (X + iY)/2$ and $\hat{\sigma}^- = (X - iY)/2$, the wave number k , position \hat{x} and momentum \hat{p} operator, axial secular frequency ω_{axial} as well as laser phase ϕ and frequency $\nu_L = \nu_0 + \Delta$. The latter detuned from the transition frequency ν_0 by Δ . The term \hat{H}_0 describes the ion by an effective two-level system in the harmonic trap potential, while \hat{H}_1 provides the laser interaction. The coupling constant between laser and ion is given by the so-called *Rabi frequency* Ω , further explained below. We restrict the interaction to a single transition frequency ν_0 , while other levels are considered sufficiently far off-resonant so that interaction with them becomes negligible. The laser is directed along the \vec{x} -axis of the trap and couples only to the axial trap motion.

The Rabi frequency Ω in the example of the quadrupole transition $4^2S_{1/2} \leftrightarrow 3^2D_{5/2}$ is given by [143]

$$\Omega = \left| \frac{eE_0}{2\hbar} \langle S | (\vec{e} \cdot \vec{r}) (\vec{k} \cdot \vec{r}) | D \rangle \right| \approx \frac{kE_0}{2\hbar} e a_0^2 \quad (2.2)$$

with the electric field amplitude E_0 , the position of the valence electron \vec{r} , the polarization \vec{e} , the wave vector \vec{k} as well as the elementary charge e of the single-ionized $^{40}\text{Ca}^+$. Moreover, the Rabi frequency can be estimated by the right term of Eq. (2.2) with the Bohr radius a_0 [143]. The coupling denotes a linear dependency on the electric field amplitude E_0 .

Assuming that the laser frequency is close to resonance with the transition frequency $|\nu_0 - \nu_L| \ll |\nu_0 + \nu_L|$ allows a rotating wave approximation [144] in Eq. (2.1). Let us additionally express the axial motion of the particle in the harmonic trap potential through creation a^\dagger and annihilation a operators, whereafter the laser-ion interaction becomes

$$\begin{aligned}\hat{H}_0 &= \hbar\omega_{\text{axial}} \left(a^\dagger a + \frac{1}{2} \right) + \frac{1}{2}\hbar\nu_0 Z \\ \hat{H}_1 &= \frac{1}{2}\hbar\Omega \left(e^{i\eta(a+a^\dagger)} \hat{\sigma}^+ e^{-i\nu_L t} + e^{-i\eta(a+a^\dagger)} \hat{\sigma}^- e^{i\nu_L t} \right).\end{aligned}\quad (2.3)$$

We next transform this term into the interaction picture through the unitary transformation $\hat{H}_I = \hat{U}^\dagger \hat{H}_1 \hat{U}$ with $\hat{U} = e^{-i\hat{H}_0 t/\hbar}$ that simplifies the above Hamiltonian to

$$\hat{H}_I = \frac{1}{2}\hbar\Omega \left(e^{i\eta(\hat{a}+\hat{a}^\dagger)} \hat{\sigma}^+ e^{-i\Delta t} + e^{-i\eta(\hat{a}+\hat{a}^\dagger)} \hat{\sigma}^- e^{i\Delta t} \right) \quad (2.4)$$

with $\hat{a} = ae^{i\omega_{\text{axial}}t}$ and laser detuning $\Delta = \nu_L - \nu_0$. The parameter η expresses the ratio between the spatial extension of the wave packet Δx and the transition wavelength λ , denoted as *Lamb-Dicke* factor, and given by

$$\eta = k\sqrt{\frac{\hbar}{2M\omega_{\text{axial}}}} = \frac{2\pi}{\lambda}\Delta x. \quad (2.5)$$

If the wavepacket's spatial extension is much smaller than the transition frequency and for small motional quantum numbers n , it holds that $\eta^2(2n+1) \ll 1$, and the ion is said to be in the Lamb-Dicke regime. Under this assumption we can Taylor expand Eq. (2.4) to first order

$$e^{i\eta(\hat{a}^\dagger + \hat{a})} \approx 1 + i\eta(\hat{a}^\dagger + \hat{a}). \quad (2.6)$$

In the Lamb-dicke regime, processes with $\Delta n > 1$ are strongly suppressed, leaving us with three main transitions to consider. Namely, the carrier transition $4^2S_{1/2} \leftrightarrow 3^2D_{5/2}$ with $\Delta n = 0$, the *red sideband* (RSB) with $\Delta n = -1$ at $\nu_0 - \omega_{\text{axial}}$ as well as the *blue sideband* (BSB) with $\Delta n = +1$ at $\nu_0 + \omega_{\text{axial}}$. The latter two named after their frequency detuning with respect to the carrier transition ν_0 . The corresponding interaction terms alongside their coupling strengths Ω read [143]

$$\begin{aligned} \hat{H}_{\text{carrier}} &= \frac{1}{2}\hbar\Omega_{n,n}(\hat{\sigma}^+ + \hat{\sigma}^-) & \text{with } \Omega_{n,n} &= (1 - \eta^2n)\Omega \\ \hat{H}_{\text{RSB}} &= \frac{1}{2}\hbar\Omega_{n-1,n}(\hat{a}\hat{\sigma}^+ - \hat{a}^\dagger\hat{\sigma}^-) & \text{with } \Omega_{n-1,n} &= \eta\sqrt{n}\Omega \\ \hat{H}_{\text{BSB}} &= \frac{1}{2}\hbar\Omega_{n+1,n}(\hat{a}^\dagger\hat{\sigma}^+ - \hat{a}\hat{\sigma}^-) & \text{with } \Omega_{n+1,n} &= \eta\sqrt{n+1}\Omega. \end{aligned} \quad (2.7)$$

Lower coupling strengths for RSB and BSB compared to the carrier transition can be estimated with the Lamb-Dicke parameter from Eq. (2.5). Fig. 2.3 depicts the level scheme of the joint system showcasing coherent carrier and sideband manipulations. Note that the particle's motion serves in both ground-state cooling [136] and for the creation of entanglement—outlined below.




Figure 2.3: **Combined manipulation of the ion-internal electronic and motional trap states via resonant laser-ion interaction.** (a) Illustration of the joint system describing the ion as an effective two-level system S and D coupling to the harmonic oscillator eigenstates n of the trap potential. Energy gaps refer to single motional oscillation quanta given by the axial secular frequency ω_{axial} . We additionally mark the coupling strengths [143] of the individual transitions, see text. (b) Resonant laser-ion interaction of any transition $|S, n\rangle \leftrightarrow |D, n'\rangle$ depicted in (a) results in periodic population transfer between the given pair of states, referred to as Rabi oscillations according to Eq. (2.8).

The time evolution of resonant laser-ion interaction with $\Delta = 0, \pm\omega_{\text{axial}}$ gives rise to so-called *Rabi oscillation* [139]

$$|\psi(t)\rangle = \cos\left(\frac{\Omega_{n',n}t}{2}\right)|S, n\rangle + e^{i\phi} \sin\left(\frac{\Omega_{n',n}t}{2}\right)|D, n'\rangle, \quad (2.8)$$

describing the periodic population transfer between the two target levels, illustrated in Fig. 2.3(b). Depending on the excitation time t any superposition states can be generated. Another, foremost technically interesting case covers off-resonant excitation at a detuning Δ from the transition that yields [139]

$$P_{|S,n\rangle \rightarrow |D,n'\rangle}(t, \Delta) = \frac{\Omega_{n',n}^2}{\Omega_{n',n}^2 + \Delta^2} \sin^2\left(\frac{\sqrt{\Omega_{n',n}^2 + \Delta^2}}{2}t\right), \quad (2.9)$$

where the transition can no longer be excited to high accuracy.

We can further use the above results to derive a multi-ion interaction for the generation of entanglement between the electronic states of ions. In that respect, the axial motion from the harmonic trap potential (see Fig. 2.3) serves as a bus for information transfer between multiple ions. A scheme following this idea was proposed by Mølmer and Sørensen in Ref. [142]. It relies on the bichromatic excitation pattern shown in Fig. 2.4(b) creating a spin-dependent force with the axial motion. As the figure indicates, laser beams with opposite detunings from the carrier $\nu_0 - \nu_L = \pm\delta$ interact with the electronic states S and D and the motional states n . Let us embed this excitation pattern into the interaction Hamiltonian from Eq. (2.4) to derive its action on two ions or more [142]

$$H_I = 2\Omega S_y \cos(\delta t) - \sqrt{2}\eta\Omega S_x \left[\hat{x} \cos(\omega_{\text{axial}} - \delta)t + \hat{x} \cos(\omega_{\text{axial}} + \delta)t + \right. \\ \left. + \hat{p} \sin(\omega_{\text{axial}} - \delta)t + \hat{p} \sin(\omega_{\text{axial}} + \delta)t \right] \quad (2.10)$$

with the Pauli spin operators from Eq. (1.5) acting on N -particles

$$S_x = \frac{1}{2} \sum_{i=1}^N \dots \otimes \mathbb{1}^{(i-1)} \otimes X^{(i)} \otimes \mathbb{1}^{(i+1)} \otimes \dots \\ S_y = \frac{1}{2} \sum_{i=1}^N \dots \otimes \mathbb{1}^{(i-1)} \otimes Y^{(i)} \otimes \mathbb{1}^{(i+1)} \otimes \dots \quad (2.11)$$

as well as position $\hat{x} = \frac{1}{\sqrt{2}}(\hat{a} + \hat{a}^\dagger)$ and momentum operator $\hat{p} = \frac{i}{\sqrt{2}}(\hat{a} - \hat{a}^\dagger)$.

If we choose light intensities sufficiently low $\Omega \ll \delta$ and adjust the laser frequency close to the sidebands with a final detuning $\Delta = \omega_{\text{axial}} - \delta \ll \delta$, terms proportional to S_y and fast oscillating terms of form $\nu_0 + \delta$ disappear. The Hamiltonian then simplifies to [142]

$$H_I = f(t)S_x \hat{x} + g(t)S_x \hat{p} \quad \text{with} \quad f(t) = -\sqrt{2}\eta\Omega \cos(\Delta t) \\ g(t) = -\sqrt{2}\eta\Omega \sin(\Delta t). \quad (2.12)$$

The time evolution of this expression describes a state-dependent force that follows a circular trajectory in phase-space given by the quadrature components f and g . A full circle in phase-space is said to be a *closed loop*. A propagator for this Ansatz is given by [142]

$$\begin{aligned} U_{\text{MS}}(t) = e^{-i\Lambda(t)S_x^2} e^{-iF(t)S_x\hat{x}} e^{-iG(t)S_x\hat{p}} \quad \text{with} \quad & A(t) = -\frac{\eta^2\Omega^2}{\Delta} \left[t - \frac{1}{2\Delta} \sin(2\Delta t) \right] \\ & F(t) = -\frac{\sqrt{2}\eta\Omega}{\Delta} \sin(\Delta t) \\ & G(t) = \frac{\sqrt{2}\eta\Omega}{\Delta} \left[1 - \cos(\Delta t) \right]. \end{aligned} \quad (2.13)$$

For controlled qubit manipulation, we wish to only interact with the electronic states $\{|S\rangle, |D\rangle\}$, while keeping motional modes n unaltered. Thus, we force $F(\tau) = G(\tau) = 0$, which is fulfilled at gate time $\tau = 2\pi K/\delta$. Here K is an integer and defines the number of closed loops in phase-space, where for each closed loop the spin-motion exactly disentangles [142]. In this case, the remaining unitary operation in Eq. (2.13) is proportional to the square spin operator S_x^2 that covers the target qubits, describing a correlated multi-ion process—the so-called *Mølmer-Sørensen* (MS) gate.

2.4 THE TRAPPED-ION TOOLBOX

Laser-ion interaction enables single- and multi-ion operations allowing us to formulate a universal gate-set—which we proceed to explain [18]. We limit the discussion to qubit gates and restrict the interactions by the two levels $|S\rangle$ and $|D\rangle$ in correspondence to the computational basis states $|0\rangle$ and $|1\rangle$. Higher dimensional qudit operations are realized by pairwise concatenation of qubit operations to link any two levels within the Zeeman manifolds, see Ch. 1.1.6. Our setup allows to address all ions along their linear string either individually or pairwise with coherent and tightly focused laser beams, which in the given case establish *full connectivity*, see Fig. 2.1(a).

- **Resonant local gates:** Rabi oscillations from Eq. (2.8) enable addressed local gates on target qubit i , depending on laser pulse duration τ_θ and intensity according to the Rabi frequency Ω from Eq. (2.2). The resulting excitation is described by the so-called *pulse area*

$$\theta = \Omega \cdot \tau_\theta, \quad (2.14)$$


that relates to a rotation around the Bloch sphere's azimuth angle. The corresponding rotation axis ϕ in the equatorial plane (X, Y -plane) can be adjusted by steering the optical phase of the laser pulse. In particular, the phase ϕ is set by the first pulse of an experimental sequence and is accumulated from then on. Shifting this phase by introducing a pause on the control hardware [145] enables gates around arbitrary rotation axes ϕ . With this, we define addressed local gates by

$$R_i^\phi(\theta) = e^{-i\frac{\theta}{2}(\cos\phi X_i + \sin\phi Y_i)} = \cos\frac{\theta}{2}\mathbb{1} - i\sin\frac{\theta}{2}(\cos\phi X_i + \sin\phi Y_i), \quad (2.15)$$

illustrated in Fig. 2.4(a) as rotations of the Bloch vector. The special cases of gates rotating around the X or Y axes are given by

$$X_i(\theta) = R_i^0(\theta) \quad \text{and} \quad Y_i(\theta) = R_i^{\frac{\pi}{2}}(\theta). \quad (2.16)$$

(a) Local gate operation



(b) MS-gate (entangling gate operation)




Figure 2.4: **The trapped-ion universal gate-set.** (a) Arbitrary single-qubit gates are realized by coherent laser-ion interaction defined by the pulse area θ from Eq. (2.14), depending on the laser light intensity and duration, and the rotation axis ϕ , relating to the laser phase. (b) Bichromatic excitation pattern for the multi-qubit MS-gate based on two laser beams with opposite detuning from the carrier transition. Entanglement is created by utilizing the axial trap motion as a bus for information transfer between multiple ions, see Eq. (2.21). The entangling gate can be realized e.g. between ion-pairs with two coherent addressed beams according to Eq. (2.19). Concatenation of two-level interactions realizes higher dimensional multi-qudit gates, see Ref. [140].

- **Off-resonant local gates:** Interactions that are detuned in frequency from the carrier transition induce AC-Stark shifts [139] between the transition levels $|S\rangle$ and $|D\rangle$. The resulting phase shift depends on the laser pulse detuning Δ , intensity Ω and duration τ_θ . Analogous to Eq. (2.14), we find a pulse area for the off-resonant excitation given by [146]

$$\theta_{AC} = \frac{\Omega^2}{4\Delta} \tau_\theta, \quad (2.17)$$

now relating to rotations around the Bloch-sphere's Z-axis. In a controlled way, these Stark shifts can realize local Z-gates on target qubit i as follows

$$Z_i(\theta_{AC}) = e^{-i\frac{\theta_{AC}}{2} Z_i}. \quad (2.18)$$

Alternatively, addressed local Z-gates can be implemented *virtually* on the control hardware by accounting for a phase shift $\Delta\phi$ on all subsequent gate operations. Virtual gates are notably free of calibration errors and cross-talk to neighboring qubits and thereby often the method of choice [70].

- **Two-qubit entangling gates:** A pair of addressed and interferometric stable laser beams grants access to arbitrary two-qubit (i, j) connectivity along the ion string, see Fig. 2.1(a). With the bichromatic excitation scheme shown in Fig. 2.4(b) the coherent pair of beams can realize two-qubit MS-gates [142]

$$MS_{i,j}^X(\theta) = e^{-i\frac{\theta}{4} X_i X_j}. \quad (2.19)$$

It should be noted that the gate becomes full-entangling at pulse area $\theta = \pi/2$. For example, the sequence $|\text{GHZ}\rangle_{01} = Z_0(-\pi/4) \cdot MS_{01}^X(\pi/2) \cdot |00\rangle_{01}$ can prepare the GHZ-state introduced in Eq. (1.19). We typically choose the X-rotation for the MS-gate and absorb potential basis change operations to local phase gates, whereas the gate in principle works for any axis in the equatorial plane (X,Y-plane).

- **Collective gates:** Apart from single or pairwise interactions, gate operations can be applied collectively to the entire ion-register with a global laser beam, manipulating all qubits simultaneously. The N-qubit spin operators S_x and S_y from Eq. (2.11) serve to define collective local gates as follows

$$R^\Phi(\theta) = e^{-i\frac{\theta}{2}(\cos\Phi S_x + \sin\Phi S_y)}. \quad (2.20)$$

Analogously, we obtain a collective N-qubit MS-gate [142]

$$MS^X(\theta) = e^{-i\theta S_x^2}, \quad (2.21)$$

displaying the unique and powerful all-to-all connectivity—a highly recognized feature of trapped ions. In this case, subsets of the qubit-string can be addressed after those not involved in the gate have been temporarily shelved to levels outside their computational subspace. This was our method of choice on the previous ion-addressing setup, described in Ref. [146]. The current addressing setup [140] supports any two-qubit connectivity along the ion-string in favour of more efficient circuit implementations as well as to reduce the error propagation. The latter is crucial for future fault-tolerant QEC applications, see Ch. 1.3. The collective laser beam currently serves for state initialization and laser cooling purposes [136]. Nonetheless, collective operations have the potential to simplify certain implementations that, for example, require simultaneous correlations between multiple qubits, recently demonstrated in Ref. [40]. Crucially, the collective gates in Eq. (2.20) and Eq. (2.21) cannot be combined in a coherent way with addressed single-qubit gates from Eq. (2.15) as no interferometric stability is present between the two laser beams. Rather, collective X, Y-gates from Eqs. (2.20)-(2.21) are accompanied by addressed off-resonant local Z-gates from Eq. (2.18)—or vice versa. This works, because the off-resonant laser light induced AC-Stark shift is not influenced by the specific phase of the light field [87]. We refer to this choice of operations as *refocused* gates.

Finally, Fig. 2.5 depicts optimized gate sequences for the H-gate from Eq. (1.17) and the CNOT-gate from Eq. (1.18), representing well-known and often utilized operations in the field of quantum computation.




Figure 2.5: **Ion trap implementation of useful quantum gates.** (a) Optimized gate sequence for H and CNOT decomposed into the trapped ion toolbox presented here [147]. We remark that a CNOT corresponds to one full-entangling two-qubit $MS^X(-\pi/2)$ alongside four single-qubit gates to correctly set the local phases.

2.5 DEVICE CAPABILITIES

Every experimental run (or sample) begins with a sequence of laser cooling including DC, *polarization-gradient cooling* (PGC) [148] and SBC, accompanied by optical pumping to prepare the ions in the well defined motional ground-state $\bar{n} \approx 0$, a prerequisite for accurate state manipulation, see Ref. [36]. Next, the circuit of interest, decomposed in and optimized with the above described universal gate-set is applied. The gate sequence can

thereby be interrupted multiple times by in-sequence measurements, where the remaining part of the computation depends on the classical output. Our novel control hardware [145] enables such feed-forward in real-time, granting access to a wide class of semi-classical quantum tasks, illustrated by Fig. 2.6(a).

While computations on the qubit level artificially restrict the Zeeman manifolds to $\{4^2S_{1/2}(m = -1/2) = |0\rangle, 3^2D_{5/2}(m = -1/2) = |1\rangle\}$, we hold equivalent control over all eight levels from Fig. 2.1(b). In this thesis we utilize larger dimensional systems with local operations only. Larger dimensions can be useful especially for simplifying gate circuits, for describing open quantum systems, or for quantum tasks to which the higher-dimensional structure is more natural, see Ch. 1.1.6. A general introduction to our qudit toolbox showcasing fruitful examples can be found in Ref. [140]. One or two concatenated qubit gates are necessary to realize a population transfer between any two levels from the Zeeman manifolds. For example, to operate the transition $4^2S_{1/2}(m = -1/2) \leftrightarrow 4^2S_{1/2}(m = +1/2)$ requires the additional state $3^2D_{5/2}(m = -1/2)$ to serve as midlevel.




Figure 2.6: **A single experimental sample.** (a) Each experimental sample (or run) begins with a sequence of DC, PGC and SBC for state initialization, followed by the experimental circuit of interest and a final detection to read the qudit register. The circuit can be interrupted by multiple in-sequence measurements, where the rest of the computation depends on the classical outcomes—denoting the concept of feed-forward. (b) Qudit detection scheme in the example of qutrit and ququart shown on left and right side, respectively. A N -dimensional qudit readout requires for $N - 1$ binary fluorescence detections [140] together with bit flips (target qudit marked as subscribs) after the first detection to transfer population to the readout sensitive $4^2S_{1/2}$ -level. Accounting for $N - 1$ binary outcomes allows the calculation of the N -dimensional qudit state probabilities.

Gate operations can also be applied to a subset of the qudits by temporarily shelving the electronic population of all other qudits in Zeeman levels outside the computational subspace. For example, to apply collective operations to a subset of the qudits, see Sec. 2.4. Moreover, after shelving the information of some qudits in the upper $3^2D_{5/2}$ manifold, the remaining ones can be detected without affecting the electronic states, referred to as *addressed readout*. We remark that scattered photons from in-sequence measurements heat up the ion string. To prevent a degrade in the post-measurement gate operation quality, every in-sequence measurement induced heating is counteracted by a sequence of DC and PGC [36].

Each experimental sample ends with a destructive measurement of the entire qudit register, where a CCD-camera can distinguish between the fluorescence of the individual

ions. Quantum state probabilities are then inferred from multiple samples of the same experiment to gain statistical accuracy. The number of samples relies on the targeted accuracy, and is always chosen in regards of the expected noise level [149]. To reach the single percentage level on statistical noise requires on the order of 100 samples per qudit, see Eq. (1.27).

On final note, let us discuss the time consumption of the individual experimental steps. State preparation takes on the order of milliseconds. Same holds for state detection. Each local bit flip operation requires tens of microseconds. The lower time end for local gates is given by thermal effects in the acousto-optical modulators needed for pulse shaping, that, below a couple of microseconds, start to negatively influence the gate performances [145]. A full-entangling operation requires around hundred microseconds limited by the duration of pulse shaping and the available laser power. Too much power relative to the local gates, however, heats up the light modulators and effectively alters the parameters of subsequent gate operations, thus, lowering their quality. Finally, laser phase noise limits our coherence time T^2 to the hundred millisecond regime setting the time available for storing and manipulating quantum information on our setup [36].

With the present NISQ-device we focus on highly controlled gate operations and the demonstrations of novel applications, while work on scalable hardware is done by other experiments in our group, see thesis by M. Brandl in Ref. [150]. The device capabilities discussed agree well with the Di Vincenzo criteria [31] on the successful demonstration of quantum computations, see Ch. 1.1.5. Moreover, our device compares with the leading quantum computing hardware in terms of the gate operation quality, as recently confirmed by Ref. [39]. Finally, Ref. [151] proposes a hodgepodge of algorithms specifically tailored for the needs of NISQ-hardware in the search for demonstrating quantum advantage, showcasing potential applications that fall under the umbrella of the trapped-ion device presented here.

A SCALABLE APPROACH TO CHARACTERIZATION

The prediction of fidelities, state properties or verifying entanglement characteristics are key ingredients in building and operating quantum hardware. Yet, the development of advanced quantum devices has over the past years significantly raised the number of operable constituents, pushing the limits of existing characterization tools, due to the inexorable exponential Hilbert space growth. Above all, there is QST [38], capable of fully reconstructing a system’s density matrix from a series of measurements, becoming increasingly impractical due to its heavy measurement and sampling requirements as well as requiring classical processing of exponentially large matrices, see Ch. 1.2.1.

In the present chapter, we develop a general characterization framework, realizing the QST capabilities, while alleviating scalability issues at the level of the number of measurement settings to perform, classical data analysis, and statistical uncertainty—as illustrated by Fig. 1.3. We start off by presenting the basic working principles of QST alongside its most established classical state reconstruction techniques in Sec. 3.1. We then extend the QST framework to the reconstruction of quantum channels via so-called *quantum process tomography* (QPT) in Sec. 3.2. QPT will be experimentally demonstrated in Ch. 5 for the characterizations of semi-classical algorithms. While appreciating the rigorous capabilities of tomography methods, we point towards critical scalability bottlenecks limiting the technique’s applicability, before then discussing the development of the scalable methods in Secs. 3.3 and 3.4. The chapter concludes with the experimental demonstrations of these scalable methods, yielding the first publication presented in this thesis in Sec. 3.5.

3.1 QUANTUM STATE TOMOGRAPHY—THE GOLD STANDARD

A d dimensional quantum state is fully described by its density matrix ρ featuring $d^2 - 1$ independent real variables. Having a faithful approximation of ρ at hand allows the prediction of every possible property of the quantum system.

QST aims at reconstructing an unknown ρ from repeated measurements on the system that sample the entire underlying Hilbert space [38, 152]. Experimentally, one performs such measurements on multiple copies of ρ , whereafter ρ is reconstructed from the observed frequencies.

More formally, let the set of measurements $P = \{\Pi_j : j = 1, \dots, M\}$ performed on ρ be projectors that correspond to quantum states $\Pi_j = |\psi_j\rangle\langle\psi_j| \in D(X)$ living in the pure state Hilbert space $X \cong \mathbb{C}^d$ of a d dimensional quantum system. By means of QST, the measurement set P can be categorized by its potential to span $D(X)$ as [153]

- *tomographically incomplete*: $|P| < d^2$ elements
- *tomographically complete*: $|P| = d^2$ linearly independent elements that span $D(X)$
- *tomographically overcomplete*: $|P| > d^2$ elements that span $D(X)$.

To collect enough information about the underlying system for QST, \mathcal{P} must at least be tomographically complete, while it is allowed to be overcomplete as well. Tomographic completeness requires d^2 linearly independent projectors, since they form a basis of all linear operators on \mathcal{H}_N [18]. Specifically, this corresponds to four projectors per qubit, each spanning the Hilbert space \mathcal{H}_2 .

Out of the measurement sets previously discussed in Ch. 1.1, the Pauli basis manifests the archetypal choice for such a task, designated as standard QST. In the single-qubit case, measurements target all three Pauli observables $\{X, Y, Z\}$, yielding an overcomplete set of six projectors $\{-X, +X, -Y, +Y, -Z, +Z\}$, see Ch. 1.1.4. Crucially, our trapped-ion platform provides these six projectors from three measurement settings, each targeting one of the Pauli observables, see Ch. 2. Generalizing to the N -qubit case requires measuring all combinations of Pauli observables, leading to an exponentially growing number of 3^N different settings. Beyond the qubit case, tomography readily extends to higher dimensional qudits [140]. The eight Gell-Mann matrices from Eq. (1.30) represent one possible extension of the standard measurement set in the qutrit case [33], being utilized in Ch. 5.5.

The second measurement set introduced was the SIC POVM from Eq. (1.28), representing a minimal tomographically complete POVM, consisting of four elements per qubit [19].

Let us exemplify the simplest case of standard QST on a single qubit, where the density matrix contains $d^2 = 4$ variables. The *trace preserving* (TP) constraint immediately fixes one variable, so we have $2^2 - 1 = 3$ independent parameters left for the reconstruction procedure. It thus suffices to incorporate state frequencies along one direction of the Pauli matrices $\{\langle X \rangle, \langle Y \rangle, \langle Z \rangle\}$, where ρ can be obtained via [149]

$$\rho = \frac{1}{2} \left(\mathbb{1} + \langle X \rangle X + \langle Y \rangle Y + \langle Z \rangle Z \right), \quad (3.1)$$

giving the parameters introduced in Eq. (1.4) a physical meaning.

For the case of two qubits, ρ features $d^2 - 1 = 15$ independent parameters and therefore requires at least four outcomes per qubit, which remains true in the generally multi-qubit case. Moreover, a tomographically complete or overcomplete measurement set is not an orthonormal basis for $D(X)$, making the state reconstruction of multiple qubits more complicated than in the special case of Eq. (3.1)—outlined below.

3.1.1 Linear inversion reconstruction

Beyond the quantum measurements, a crucial part of QST is the state reconstruction performed in classical data analysis. We open the discussion by introducing the analytic procedure of *linear inversion* (LI) following explanations in Ref. [153].

Born's rule states that the probability of finding the system ρ in the state associated to the projector $\Pi_j \in \mathcal{P}$ is given by the success probability $p_j = \text{tr}(\Pi_j \rho)$, introduced in Eq. (1.26). Following the above description, an experimenter extracts probabilities p_j by preparing N_j samples of ρ and repeatedly measures Π_j . Frequencies from n_j positive outcomes are then obtained by

$$p_j = \frac{n_j}{N_j}. \quad (3.2)$$

Repeating this for all projectors $\Pi_j \in \mathcal{P}$ lies in the nature of LI. If \mathcal{P} was at least tomographically complete, we can attempt to reconstruct ρ from the collection of observed frequencies.

Because a tomographically complete or overcomplete measurement set \mathcal{P} is not an orthonormal basis for $D(X)$, state reconstruction following Eq. (3.1) is not anymore feasible

beyond the single-qubit case. An established way to overcome this problem and reconstruct ρ from the observed frequencies p_j relies on the so-called *dual basis* introduced in Ref. [154] and briefly outlined in the following. To this extent, let us define the so-called *superoperator* spanned by the measurement set P by

$$S_p = \sum_{j=1}^M |\Pi_j\rangle\rangle\langle\langle\Pi_j|, \quad (3.3)$$

with the vectorized projector $|\Pi_j\rangle\rangle$ emerging after stacking all the columns in Π_j . The dual basis $D = \{D_j : j = 1, \dots, M\}$ for P is then defined using the superoperator as $|D_j\rangle\rangle = S_p^{-1}|\Pi_j\rangle\rangle$ with the Moore-Penrose pseudo inverse S_p^{-1} for non-square matrices [155]. Notably, the pseudo inverse relaxes to the square matrix inverse in case of a tomographically complete set [153].

For P being at least tomographically complete, it further holds that $S_p^{-1}S_p = \mathbb{1}$ as otherwise no inverse existed. By using the above relations one can then show that [153]

$$|\rho_{LI}\rangle\rangle = \sum_j \frac{n_j}{N_j} |D_j\rangle\rangle = \sum_j p_j S_p^{-1} |\Pi_j\rangle\rangle = \sum_j S_p^{-1} |\Pi_j\rangle\rangle\langle\langle\Pi_j|\rho_n\rangle\rangle = (S_p^{-1}S_p)|\rho_n\rangle\rangle, \quad (3.4)$$

where we introduce the operator ρ_n on the right hand side, from which the observed frequencies $n_j/N_j = \text{tr}(\Pi_j\rho_n)$ have been generated. Conclusively, we find that our linear inversion estimate is equal to ρ if the superoperator S_p is invertible, being particularly the case for P either tomographically complete or overcomplete.

Finally destacking the columns in $|\rho_{LI}\rangle\rangle$ yields the quested LI density matrix

$$\rho_{LI} = \sum_j \frac{n_j}{N_j} D_j. \quad (3.5)$$

The superoperator from Eq. (3.3) constitutes the largest matrix that needs to be handled in the reconstruction process with LI. Let us therefore discuss its dimensions for different measurement sets to reveal the resource requirements on the desktop hardware utilized for the classical data analysis process. For the exemplified measurement sets of standard Pauli basis and SIC POVM, we receive matrix dimensions of

$$\dim[S_p(\text{Pauli})] = 4^N \times 6^N \quad \text{and} \quad \dim[S_p(\text{SIC POVM})] = 4^N \times 4^N, \quad (3.6)$$

respectively. The products can be interpreted as the total number of entries of the N -qubit density matrix (left) and the number of measurement states of either set $|P|$ (right).

We emphasize, that the inversion of S_p can also be done line-wise in favour of being less memory consuming [153]. Such an approach trades space for time complexity and while allowing the analysis of larger systems, it increases the time cost by introducing significantly more multiplications in the line-by-line inversion process.

A realistic QST scenario can only provide a finite number of experimental samples during the time available for data taking. The bigger the quantum system under study, the more of a problem this becomes. Even absent of computational or measurement errors, finite statistics limit the quality of the reconstructed state due to the influence of QPN, see Eq. (1.27). Consequently, the analytic but completely unconstrained ρ_{LI} is likely to produce state estimates that violate the conditions of positive-semidefiniteness $\rho_{LI} > 0$ and/or unit trace $\text{tr}(\rho_{LI}) = 1$ [152]. While the trace constraint resolves upon normalization $\rho_{LI}/\text{tr}(\rho_{LI})$, the outcome does potentially not anymore represent a positive-semidefinite, i.e., valid quantum state.

It can be shown that LI intrinsically predicts the linear observables most accurately [48], whereas the estimation of non-linear state properties is usually affected by the unphysical nature of the LI estimates under insufficient statistics [40]. For example, the purity from Eq. (1.3), where values greater than 1 indicate the presence of negative eigenvalues in the reconstructed density matrix ρ_{LI} . Moreover, the convergence of such non-linear estimates slows down in the sampling process. We will thoroughly discuss these findings as part of our experimental demonstrations in Sec. [40].

A computationally efficient algorithm that incorporates physicality constraints to LI reconstructed ρ_{LI} was proposed by Smolin et al. as part of Ref. [156]. The proposed algorithm truncates negative eigenvalues in the initial LI-estimated density matrix and seeks out the closest density matrix that obeys physicality constraints under the Frobenius norm. The method is named *projected least squares* (PLS) and while it encouragingly does not add additional time complexity to the reconstruction process of LI, the method comes with very slow convergence behaviour as thoroughly demonstrated in Refs. [40, 157].

3.1.2 Maximum likelihood estimation

The potential failure of ρ_{LI} to be a physical quantum state can be overcome by reformulating QST as a *maximum likelihood estimation* (MLE) problem. The starting point of MLE lies in estimating the likelihood that a certain ρ has generated the observed frequencies. Importantly, the optimization process can incorporate well defined constraints, most typically, positive-semidefiniteness $\rho > 0$ and/or unit trace $\text{tr}(\rho) = 1$. These constraints ensure that the reconstructed density matrices represent true quantum states in any case. The likelihood is then maximized over ρ to find the ρ_{MLE} that is most likely to have produced the observed frequencies of a tomographically complete or overcomplete measurement set under the given constraints. This becomes notably useful in scenarios with insufficient statistics and noisy devices and is often the method of choice, if the underlying system size allows its computation.

More formally, we introduce a slightly relaxed but computationally more efficient version of MLE QST that solves the following optimization problem [153]

$$\begin{aligned} & \text{minimize} && \|W(S|\rho\rangle\rangle - |f\rangle)\|_2 \\ & \text{subject to:} && \rho > 0, \quad \text{tr}(\rho) = 1. \end{aligned} \quad (3.7)$$

Here, S is a vectorization change of basis operator, $|f\rangle$ represents the observed frequencies for a tomographically complete or overcomplete measurement set and W is a diagonal matrix of the statistical weights that incorporates the variances σ_j of measured frequencies

$$S = \sum_{j=1}^M |j\rangle\langle\langle\Pi_j|, \quad |f\rangle = \sum_{j=1}^M \frac{n_j}{N_j} |j\rangle, \quad W = \sum_{j=1}^M \frac{N_j}{\sigma_j} |j\rangle\langle j|. \quad (3.8)$$

Crucially, the cost function in Eq. (3.7) is a convex optimization problem. As the name suggests, the underlying parameter space follows a convex curvature featuring global minima. The convex problem class is fortunately solvable by a wide range of existing algorithms, summarized in Ref. [158]. We additionally remark that the MLE approach from Eq. (3.7) is a relaxation of MLE speeding up the classical computation time and only in the limit of large sample sizes coincides with the original method, further discussed in Ref. [152].

In terms of classical data analysis, MLE is subject to the same matrix dimensions as LI, discussed in Eq. (3.6), which the method further builds into an optimization problem and

therefore scales expectedly worse. The high demand in computer resources makes MLE only useful for qubit numbers in the single digits [149]. While the resulting ρ_{MLE} satisfies physicality constraints in any case, enforcing these constraints can bias the outcome under very few samples [159]. This becomes evident in the estimation of linear observables yielding highly suppressed values for comparably small sample numbers, especially in view of the here much more efficient LI—thoroughly discussed in Ref. [159]. For more technical insights on MLE methods, we refer the reader to Ref. [152].

3.2 QUANTUM PROCESS TOMOGRAPHY

Next, we extend the tomography framework to capture quantum dynamical processes. The essence of so-called QPT is to identify a quantum channel that transforms a density matrix $\rho_{\text{in}} \in \mathcal{D}(X)$ into a density matrix $\rho_{\text{out}} \in \mathcal{D}(Y)$ described by a linear map \mathcal{E}

$$\rho' = \mathcal{E}(\rho) = \sum_{m,n} \chi_{m,n} E_m \rho E_n^\dagger, \quad (3.9)$$

where we additionally defined the $d^2 \times d^2$ -dimensional positive Hermitian χ -matrix on the right hand side. The entries of χ uniquely describe any linear map \mathcal{E} using an orthonormal basis $\{E_n\}$ with $E_n^\dagger E_m = \delta_{nm}$ and $\sum_n E_n^\dagger E_n = \mathbb{1}$ [18]. An example of a quantum channel is the depolarizing channel from Eq. (1.22), which uses the Pauli basis.

In QPT we attempt to extract the χ -matrix in Eq. (3.9) from repeated measurements on many copies of a target quantum process, analogous to QST. Yet capturing the dynamics of an unknown map \mathcal{E} requires not only to probe a complete set of measurements $P = \{\Pi_j\} \in \mathcal{D}(Y)$, as for QST, but a complete set of input states $Q = \{\rho_j\} \in \mathcal{D}(X)$ as well. In that, QPT can be related to QST of various input states, to which a good deal of the previously discussed framework of Sec. 3.1.1 and 3.1.2 applies.

Let us quantify the number of tomography settings consisting of input and measurement states that allow the reconstruction of the process map. The number of parameters in the χ -matrix from Eq. (3.9) is d^4 [18]. Crucially, the TP constraints from a complete set of d^2 input states immediately restricts the free parameters in the χ -matrix down to $d^4 - d^2$. Standard QPT considers once more the Pauli basis. In that case, we can show that four input states together with six Pauli measurement states per qubit provide $4^N \times 6^N = 24^N$ independent parameters for the N -qubit system and thereby suffice to describe all free parameters in the χ -matrix.

Considering measurement outcomes along both directions of each Pauli observable, the N -qubit case requires 12^N tomography settings, i.e., four input states and three measurements per qubit. While in Ch. 1.2.1 we stated four hours of data acquisition for an eight-qubit standard Pauli QST on a state-of-the-art trapped-ion device, QPT on the same register required 33 years of data acquisition [40]. Schemes to relax this discouraging requirement, e.g. *ancilla-assisted QPT*, yield a reduced 4^N settings covering the different input states at the cost of a significant experimental overhead when relaying the measurement information to the ancilla qubits [160]. These hard requirements realistically restrict QPT to only few qubits.

To reconstruct arbitrary quantum channels from the observed frequencies of an, at least, tomographically complete set of input and measurement states, we can in part build on the QST framework presented above. The connection between QST and QPT can be made by utilizing a useful bijection between linear maps and linear operators, referred to as *Choi-Jamiolkowski isomorphism* [161]. Specifically, the so-called *Choi operator* $\Lambda_{\mathcal{E}}$ can fully describe any CP map \mathcal{E} while being isomorphic to a bipartite quantum state $\rho_{\mathcal{E}} \in \mathcal{L}(X \otimes Y)$,

where $\rho_\varepsilon = \Lambda_\varepsilon/d_X$. The isomorphism between Choi operator Λ_ε and channel map \mathcal{E} is given by

$$\mathcal{E}(\rho) = \text{tr}_X[(\rho^T \otimes \mathbb{1}_Y)\Lambda_\varepsilon] \quad \text{with} \quad \Lambda_\varepsilon = \sum_{k,l}^{d-1} |k\rangle\langle l| \otimes \mathcal{E}(|k\rangle\langle l|), \quad (3.10)$$

with an explicit form for the Choi operator Λ_ε on the right hand side, decomposed into the basis $\{|k\rangle\}_{k=0}^{d-1}$ of Hilbert space dimension d . Here, we follow notations in Ref. [162]. The frequencies $\hat{p}_{i,j}$ for observing the outcome state ρ_j from initial state ρ_i under the channel Λ_ε can be expressed by

$$\hat{p}_{i,j} = \text{tr}[\rho_j^\dagger \text{tr}_X[(\rho_i^T \otimes \mathbb{1}_Y)\Lambda_\varepsilon]] = \text{tr}[(\rho_i^T \otimes \rho_j^\dagger)\Lambda_\varepsilon]. \quad (3.11)$$

Let us define the projector $\Pi_{i,j} \equiv \rho_i^* \otimes \rho_j$ considering pure states ρ_i and ρ_j alongside the vectorized operator $|\Lambda_\varepsilon\rangle\rangle = \sum_{i,j}^{d-1} \Lambda_{i,j} |j\rangle \otimes |i\rangle$. We can then identify the trace in Eq. (3.11) with an inner product of the vectorized operators and formulate the frequencies as

$$\hat{p}_{i,j} = \langle\langle \Pi_{i,j} | \Lambda_\varepsilon \rangle\rangle. \quad (3.12)$$

From this, we can express the vector of observed frequencies $|f\rangle$ as well as the superoperator S as follows

$$|f\rangle = \sum_{i,j} f_{i,j} |i,j\rangle \quad S = \sum_{i,j} |i,j\rangle \langle\langle \Pi_{i,j} |. \quad (3.13)$$

Apart from the quantum measurements to be performed on the given set of input states, process reconstruction is a crucial part of QST. LI process reconstruction stays prone to unphysicalities under finite statistics and noisy quantum devices [152]. We therefore mention here only the possibility of LI reconstruction and refer to Ref. [162]. Due to a discouraging 12^N tomography settings, standard QPT mostly targets one or two qubits, a regime, where MLE reconstruction remains efficiently computable.

In analogy to Eq. (3.7), QPT MLE can be defined as follows [162]

$$\begin{aligned} & \text{minimize} && \|WS|\Lambda_\varepsilon\rangle\rangle - W|f\rangle\|_2 \\ & \text{subject to:} && \Lambda_\varepsilon \geq 0, \quad \text{tr}(\Lambda_\varepsilon) = d \end{aligned} \quad (3.14)$$

with W , a weight matrix, that considers the observed frequencies to follow a multinomial distribution. With regards of classical data analysis, the very same algorithms as for QST reliably solve the convex optimization problem from Eq. (3.14) and allow for process reconstruction [158]. For further details about QPT MLE, we guide the reader to Ref. [153].

Experimental QPT will be part of Ch. 5, where we extend the system dimension and experimentally demonstrate combined qubit-qutrit QPT to study leakage dynamics beyond the computational subspace. In particular, this will focus on the various constraints in the reconstruction process of MLE.

3.3 SIC POVM QUANTUM STATE TOMOGRAPHY

Although QST enables the prediction of all possible system properties, the number of measurement settings to be performed grows exponentially with the system size N . This follows the classical data analysis part based on methods from Secs. 3.1.1-3.2 dealing with exponentially large matrices. Playing tricks like parallelization or linewise LI trade space for time complexity which cannot counteract the exponential growth in Hilbert

space towards larger systems. Other characterization methods build in assumptions about the system to characterize or sacrifice the information gain in favour of better scalability, thoroughly discussed in Ch. 1.2.1. However, such relaxed approaches are often unable to predict universal functional properties of the density matrix—especially non-linear ones.

Alternative tomography techniques that try to circumvent the exponential growth in the number of settings, for instance *adaptive tomography* [163], iteratively seek out the ideal measurements to perform on the given system by building on VQE techniques. While ideal measurements promise better statistical accuracy with significantly fewer measurement settings and samples, their implementation often encompasses similar complexity as state preparation itself. The additional time cost required for the iterative search routine limits the methods applicability further.

Given the wide reaching applicability of QST for system characterizations, we now attempt to resolve critical scalability bottlenecks of the method by pursuing the three challenges initially posed in Fig. 1.3.

The first challenge tackles the number of measurement settings for doing tomography. To resolve this issue, we move away from the standard Pauli basis and incorporate the tomographically complete set of SIC POVMs instead. The four non-orthogonal SIC POVM elements from Fig. 1.1(b) are symmetrically aligned and maximized for intervector spacing [164]. Beyond providing tomographically complete information in a single measurement setting (see below), the SIC alignment provides a higher information gain as its non-orthogonal POVM elements almost always have finite overlap with arbitrary states and thereby more likely contribute to the statistical accuracy under realistic sample number constraints.

Accessing non-orthogonal SIC POVMs from projective measurements, however, can be tedious. In a brute force method one would simply rotate the Bloch sphere of every experimental sample such that the target POVM element overlaps with the Z basis, whereafter it can be read out by projective measurement—similar to accessing Pauli X or Y in standard QST. Repeating this for all combinations of SIC POVM elements provides the complete tomographic information. In our ion trap setup this results in four settings per qubit and is notably more complex than extracting the overcomplete six Pauli measurement states from only three projective measurements along X, Y and Z. The situation might differ for photonic platforms that need six projective measurements for the Paulis, where SIC POVMs indeed offer to reduce the number of tomography settings [165]. A smarter but more costly way employs ancilla-assisted measurement strategies [166] to relay measurement relevant information from data to ancilla qubits. This can provide access to the entire measurement information of the SIC POVM elements from a single experimental sample. Getting rid of different measurement settings comes at the price of engaging additional ancilla qubits and costly entangling operations to couple tomography relevant information to them. Because of such practicality issues and outside of photonic systems, SIC POVMs have thus far mostly been avoided in experimental realizations.

We now follow a novel path and attempt to access the SIC POVM elements of a multi-qubit system from a single experimental sample, while avoiding the need of extra ancilla qubits. To this extent, we make use of Naimark’s dilation theorem [28] stating that every POVM emerges as a projective measurement in a higher-dimensional Hilbert space. Let us recall that our trapped $^{40}\text{Ca}^+$ -ions allow encoding of qudits with up to seven levels per ion, out of which already four levels suffice to store the measurement information of the four SIC POVM elements. Together with our recently developed qudit toolbox [140], we then extend the qubit Hilbert space by two additional levels and locally map the non-orthogonal four SIC POVM elements from the qubit to orthogonal ququart states. Each individual ion-ququart thereafter stores the complete tomographic information of

a qubit. This information is experimentally accessible by means of a *four-outcome readout*, notably within the same sample. Such a readout is realized by three sequential in-sequence measurements. Specifically, after the first detection, the population between state $|0\rangle$ and $|1\rangle$ is flipped and likewise before the final and third detection the states $|0\rangle$ and $|2\rangle$ are flipped, illustrated in Fig. 2.6. From the resulting three binary outcomes the ququart state probabilities can be inferred—more details are given in Ch. 2.5. The ion-internal mapping unitary following Naimark’s theorem reads

$$\hat{M} = \begin{pmatrix} \frac{1}{\sqrt{2}} & 0 & 0 & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{6}} & \frac{1}{\sqrt{3}} & \frac{1}{\sqrt{3}} & -\frac{1}{\sqrt{6}} \\ \frac{1}{\sqrt{6}} & \frac{(-1)^{2/3}}{\sqrt{3}} & \frac{(-1)^{4/3}}{\sqrt{3}} & -\frac{1}{\sqrt{6}} \\ \frac{1}{\sqrt{6}} & \frac{(-1)^{4/3}}{\sqrt{3}} & \frac{(-1)^{2/3}}{\sqrt{3}} & -\frac{1}{\sqrt{6}} \end{pmatrix}. \quad (3.15)$$

The corresponding gate sequence realizing this unitary, specifically tailored and optimized for our trapped-ion device results in five local operations

$$M = R(\phi = \frac{\pi}{2}, \theta = \frac{\pi}{2}) \cdot R(-\frac{\pi}{2}, \frac{\pi}{2}) \cdot R(\frac{\pi}{2}, 2 \arctan \sqrt{2}) \cdot R(0, \frac{\pi}{2}) \cdot R(-\frac{\pi}{6}, \pi), \quad (3.16)$$

that must be applied to each qubit. Mapping can be done collectively, i.e., on all qubits at once, or individually using addressed gates, see Ch. 2.4. Subjecting more qubits to the tomography process therefore keeps the resource requirement constant since all the local operations can be parallelized. With this purely local sequence and absent the need of ancilla qubits to store extra information, SIC QST adds only moderate complexity to the underlying experimental implementation, while it provides the complete tomographic information with every experimental sample. Because of this compelling feature, we synonymously refer to SIC QST as *single-setting QST*, resolving the first challenge in Fig. 1.3.

Single-setting QST appears very practical from an experimental point of view as one repeatedly performs the same experiment over and over to gain statistical accuracy. This can be advantageous as changing measurement settings is typically associated with big efforts. Avoiding changes of measurement settings further saves time in classical sequence compiling. For example, the control hardware of our ion trap device requires all sequences to be precompiled before the first experiment begins [145]. As such, in terms of practicality, SIC QST provides benefits over existing tomography approaches. In statistical terms, system properties might converge to a certain accuracy in less samples than necessary for the minimal Pauli QST implementation, i.e, below 3^N samples. Linear observables are good candidates for this, for instance fidelity [167]. Although techniques such as *direct fidelity estimation* [41] enable the prediction of individual linear properties without reconstructing the density matrix, those incorporate preknowledge about the system. Moreover, it is common that numerous observables need to be questioned simultaneously which becomes likewise hard for existing methods that circumvent the density matrix reconstruction. Yet, this is not the case for SIC QST offering all necessary information of a multi-qubit system from a single measurement setting. With a complete set of measurements at hand, one can even pick the properties to analyze after the experiments have been completed. In this case and for linear observables SIC QST already holds the potential to outperform existing tomography methods.

For classical data analysis the same procedures discussed in Secs. 3.1.1-3.1.2 apply to the density matrix reconstruction in SIC QST. While with standard Pauli basis one handles matrices of dimension $4^N \times 6^N$, SIC QST only accounts for $4^N \times 4^N$, being still exponential, but denoting a relevant improvement in view of current NISQ hardware. Beyond state

reconstruction, SIC QST can be applied to the measurement part of QPT, leaving only 4^N settings for the state preparation part. The method has thus the potential to outrun ancilla-assisted QPT, in that it saves the need for extra qubits and costly entangling operations to couple to them [160].

Besides trapped ions many architectures successfully employ higher-dimensional systems to which SIC QST seamlessly applied. These cover atomic and molecular systems [104, 105], photonic platforms [109], Rydberg-ions [106], superconducting devices [108] and solid-states [107].

SIC QST fixes the exponential growth in the number of measurement settings to be performed. Accurate estimates of large and, in particular, non-linear system properties rely on at least the reduced K -qubit density matrix, where the operators of interest act on [47]. Note that the number of independent parameters in the reduced density matrix $d^2 - 1 = 4^K - 1$ scales exponentially with the size of the subsystem K . Fortunately, many quantities of interest are local, which keeps this scaling manageable in a large class of problems. The prediction of non-linear quantities in highly correlated systems, however, remains generally non-scalable. Examples are the entanglement measures from Eq. (1.11) and Eq. (1.12). It is likely that the number of experimental samples needed to reach high accuracy will grow faster than the number of measurement settings. Thus, the specific selection of a tomographic measurement set no longer limits the total time required for the tomography process.

In the upcoming section, we present a novel data analysis method that naturally applies to the SIC QST, revealing interesting scaling properties.

3.4 SYSTEM CHARACTERIZATION VIA CLASSICAL SHADOWS

Accurately reconstructing the (reduced) density matrix requires a number of samples that scales exponentially with the (sub)system size [48]. This holds for the general case of highly correlated quantum states, which are randomly aligned with respect to the chosen tomography measurement set. Exceptions arise from special choices of states, such as eigenstates of the given measurement set. The GHZ-state from Eq. (1.10) characterized by the standard Pauli basis provides such an example, resulting in boosted convergence [40], whose full extent is still not entirely explored. Special cases aside, the number of samples on an accurate description of the (reduced) density matrix remains mostly independent of the QST method and its particular measurement set.

To counteract this daunting bottleneck, Aaronson proposed so-called *shadow tomography* in Ref. [167] to predict multiple system properties in parallel and notably directly from the subset of qubits that the operators act on. His work successfully shows that a polynomial number of state samples can be sufficient to estimate exponentially many properties of the underlying system. Besides being very efficient in terms of the required sample sizes and conceptually interesting, an actual experimental implementation of shadow tomography is very hardware-demanding as it scales to exponentially deep and highly correlated quantum circuits. Even if shadow tomography is feasible, the resulting overhead of gate operations for the characterization procedure likely obscures the quality of the underlying implementation.

Huang, Kueng and Preskill continued this search and consolidated ideas of shadow tomography [167] with rigorous statistical convergence guarantees of QST [157] and an efficient implementation of the stabilizer formalism [82]. Their work establishes a protocol that offers to accurately predict L different function properties of the density matrix ρ , while relying on only $O(\log(L))$ experimental samples. Analogous to the ideas of Aaronson, the

protocol offers the direct reconstruction of the reduced density matrix from measurements on just the involved subset of the qubits. The technique is named classical shadows [48, 168]. Importantly, the authors prove that the stated guarantees are not limited by the system size. So far, the practicality of classical shadows was overshadowed by the requirement of performing random measurement samples drawn from the exponentially large set of tomography measurements [47]. Nonetheless, this is now a technical, rather than a fundamental limitation.

There is a way to overcome this technical limitation by moving from Pauli to SIC QST, where random sampling is built in, since each experimental sample contains the complete tomographic information, see Sec. 3.3. In fact, randomization is naturally integrated into the destructive measurement process. In view of this potential, we attempt to combine SIC QST with classical shadow data analysis in order to overcome the random sampling problem of existing methods and to efficiently access the reduced density matrix.

3.4.1 Classical shadow tomography

As a first step, we formulate an alternative reconstruction procedure for ρ using SIC-based classical shadows, which immediately yields promising scaling properties.

The starting point of QST is the estimation of measurement frequencies for a tomographically complete or overcomplete set of measurements that enable the reconstruction of the density matrix ρ , see Sec. 3.1. So let us consider a projective measurement on an N -qubit system using the tomographically complete SIC QST. On sample-by-sample basis, every qubit is projected onto one of the four SIC POVM elements $\{\frac{1}{2}|\psi_i\rangle\langle\psi_i| : i = 1, 2, 3, 4\}$ living in the two-dimensional Hilbert space \mathbb{H}_2 , illustrated by Fig. 1.1(b). Hence, results on N -qubits can be identified by a string (i_1, \dots, i_N) with $i \in \{1, 2, 3, 4\}$ referring to a specific POVM element. Multiple samples yield the corresponding frequencies $\Pr[i_1, \dots, i_N | \rho]$ from which the density matrix ρ can be reconstructed.

We begin the reconstruction process with a single qubit and weigh each SIC projector $|\psi_i\rangle\langle\psi_i|$ with the corresponding frequency $\Pr[i|\rho]$ of observing this outcome. Hereby, explanations in Ref. [40] show that $\sum_{i=1}^4 \Pr[i|\rho] |\psi_i\rangle\langle\psi_i| = \frac{1}{3}(\text{tr}(\rho)\mathbb{I} + \rho)$ results in a linear map related to a depolarizing channel as in Eq. (1.22) with parameter $p = 1/3$. Fortunately, this linear map has a uniquely defined inverse $\sum_{i=1}^4 \Pr[i|\rho] (3|\psi_i\rangle\langle\psi_i| - \mathbb{I}) = \rho$ that exactly reproduces the requested single-qubit density matrix ρ . Note that this inverse is not a physical operation, but mathematically possible [40].

Each individual measurement outcome, for instance sample m , contributes to the above reconstruction of ρ and therefore denotes a so-called Monte Carlo estimator $\hat{\sigma}_m = (3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I}) \in \mathbb{H}_2$ [40]. Crucially, this sample-based term obeys a tensor structure that readily extends to the multi-qubit case given by

$$(i_1, \dots, i_N) \mapsto \hat{\sigma}_m = \bigotimes_{n=1}^N (3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I}) \in \mathbb{H}_2^{\otimes N}. \quad (3.17)$$

The Monte Carlo estimator $\hat{\sigma}_m$ yields a sample-based approximation of ρ , called a classical shadow—which we proceed to explain. The tensor product of all N single-qubit contributions yields a random $2^N \times 2^N$ matrix. Since each tensor factor in Eq. (3.17) obeys the trace norm, its eigenvalues are -1 or 2 . The eigenvalues of the N -qubit tensor product therefore result in values between $-(2^N - 1)$ and $+2^N$, designating highly unphysical matrices. Crucially, upon averaging over M samples of classical shadow estimators $\hat{\sigma}_m$, the

unphysical properties quickly average out [168], where the true density matrix ρ is found in the asymptotic limit with respect to the sample number

$$\hat{\rho} = \frac{1}{M} \sum_{m=1}^M \hat{\sigma}_m \xrightarrow{M \rightarrow \infty} \rho, \quad (3.18)$$

referred to as *classical shadow tomography*. A more stringent derivation can be found in Ref. [40].

From Eq. (3.18) some useful properties immediately follow. Namely, that the state reconstruction is processed at the smallest possible dimension for full tomography given by the N-qubit density matrix of $\dim(\rho) = 2^N \times 2^N$ and that individual experimental samples are simply accumulated in approximating ρ . The method thereby avoids any costly matrix inversion or convex optimization. This is a huge gain over existing LI, where Pauli and SIC bases require processing much larger matrix dimensions given by $4^N \times 6^N$ and $4^N \times 4^N$, respectively.

Furthermore, the accumulation of individual experimental samples enables so-called *live updates* or *online analysis* of the quantum state. The reconstruction process can start after receiving the first measurement sample, while the experiment continues to run. Once ρ is reconstructed with sufficient accuracy, being the case when the estimators under investigation have converged, the experiment can be terminated. In performing measurements and data analysis in parallel, almost no additional time overhead is created by the state reconstruction procedure. In addition to the smaller matrix dimensions to handle, stated in the previous paragraph, our SIC QST based classical shadow density matrix reconstruction offers more speed-up over existing reconstruction methods the more qubits are involved. In particular, our experimental analysis of a six qubit system in Ref. [40] shows that linear inversion state reconstruction with SIC shadows from Eq. (3.18) is already two times faster than Pauli linear inversion from Eq. (3.5), where the time consumption for SIC shadows grows much more slowly. A comparison with MLE from Eq. (3.7) was only feasible up to five qubits on our desktop hardware, with the SIC reconstruction being already 37 times faster than the Pauli. SIC shadows thereby resolves the second challenge from Fig. 1.3. However, the method still provides a complete reconstruction of the density matrix, with the number of free parameters increasing exponentially as the size of the system is increased. While certain system properties [48] might converge in fewer samples than required for a minimal standard Pauli or SIC based tomography implementation, the specific sample sizes critically depend on the quested estimator, which we thoroughly discuss in the upcoming sections.

3.4.2 Efficiently estimating linear system properties

The tensor structure in Eq. (3.18) offers yet another advantage of classical shadows, namely that the estimation of system observables becomes notably simpler than with existing methods.

Consider an observable O , acting on K out of N qubits, then the expectation value of this observable can be estimated via classical shadows as [40]

$$\begin{aligned}
\hat{o} &= \text{tr}(O\hat{\rho}) = \frac{1}{M} \sum_{m=1}^M \text{tr}(O\hat{\sigma}_m^N) = \\
&= \frac{1}{M} \sum_{m=1}^M \text{tr} \left(O \bigotimes_{n=1}^N (3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I}) \right) = \\
&= \frac{1}{M} \sum_{m=1}^M \text{tr}(O\hat{\sigma}_m^K)
\end{aligned} \tag{3.19}$$

where the shadow estimator allows us to pull the trace into the average, which would not be possible when working with the density matrix. Note that the trace over all $N - K$ qubits in the second line, on which O acts trivially, evaluates to 1. Hence, in the final line of Eq. (3.19) only K qubits remain present to the classical shadow estimator. Moreover, since each experimental sample in SIC QST is tomographically complete, one can estimate many observables in parallel and even choose which properties to analyze after the experiments are complete. Especially for $K \ll N$, the data analysis simplifies significantly in contrast to existing methods.

At this stage, it is important to quantify the number of measurement samples required to predict multiple linear observables in parallel with a given precision. Along those lines we prove in Ref. [40] that estimating $L \gg 1$ subsystem observables O_l each acting on $K \leq N$ qubits requires

$$M \geq \frac{8}{3} 6^K \log(2L/\delta)/\epsilon^2 \tag{3.20}$$

measurement samples to collectively approximate every observable to accuracy ϵ with at least the probability $1 - \delta$. While decreasing ϵ for more precise estimates leads to a quadratic increase in the sample size, estimating more observables L in parallel or increasing the success probability $1 - \delta$ leads only to a logarithmic growth. We note that in a real experiment the lower bound for ϵ is determined by the respective error-rates. The number of measurement samples M scales exponentially only in the subsystem size K . Hence, classical shadows become most efficient in situations where $K \ll N$. We remark that this is a rigorous *a priori* bound utilizing minimal assumptions—more details are given in Ref. [40]. While the presence of such a bound is of conceptual interest, the properties will converge (much) sooner with respect to experiments or numerical simulations, as will be demonstrated in Sec. 3.5.

Predicting multiple linear observables in parallel is necessary in many situations, as for instance in VQE applications [37, 169].

3.4.3 Efficiently estimating non-linear system properties

Estimation of non-linear system properties, such as the Rényi entropy from Eq. (1.12), requires at least the reduced density matrix covering the qubit subsets of interest. Standard QST approaches retrieve the reduced density matrix after partial tracing over qubits not part of the reduced subsystem. The accuracy of the reduced subsystem is therefore naturally bounded by the accuracy of the total system. This represents a daunting bottleneck when estimating non-linear system properties with standard tomography approaches.

For classical shadow analysis, however, we find partial tracing [18] as a linear operation to match the tensor structure of our estimators in Eq. (3.18) well. We remark that each

tensor factor has unit trace $\text{tr}(\rho) = 1$. The shadow estimator therefore allows us to pull the partial trace into the average and to independently extract the subset classical shadow estimator $\hat{\sigma}^K$ for a given subsystem of size $K < N$. This enables the reconstruction of the reduced density matrix ρ_K as follows [40]

$$\rho_K = \text{tr}_{(N-K)}(\rho_N) = \frac{1}{M} \sum_{m=1}^M \hat{\sigma}_{m'}^K, \quad (3.21)$$

imparting only K qubits, while neglecting $N - K$ ones. Classical shadows thus offer the possibility to predict non-linear system properties by direct reconstruction of the reduced density matrix.

Let us illustrate this by the example of the simplest non-linear system property given by the quadratic purity estimator $\text{tr}(\rho^2)$ from Eq. (1.3). Recall that the purity ranges from 1 for a pure state to $1/d$ for a completely mixed one. The trace of ρ^2 can be interpreted as the trace of the product of two copies of ρ . While for classical shadows, the latter converge independently to the true state according to Eq. (3.18), we conclude that the trace of their product converges to the trace of the true state ρ^2 . To derive the corresponding quadratic classical shadow estimator, we use the fact that two distinct classical shadows $\hat{\sigma}_m$ and $\hat{\sigma}_{m'}$, originating from different samples $m \neq m'$, remain statistically independent [48]. Hence, any pair of classical shadow estimators approximates true purity values $\text{tr}(\hat{\sigma}_m \hat{\sigma}_{m'})$. Accounting for all $\binom{M}{2}$ combinations of independent sample-pairs yields the desired estimator [40]

$$\hat{p}_{(M)} = \binom{M}{2}^{-1} \sum_{m < m'} \text{tr}(\hat{\sigma}_m \hat{\sigma}_{m'}). \quad (3.22)$$

Note that there are quadratically more pairs than single estimators, so one naively expects a quadratic improvement in convergence. A proof that this is indeed statistically most efficient can be found in Ref. [48]. While the purity estimator on large-scale systems will remain hard to estimate even for classical shadows, the situation becomes again favorable for comparably small qubit subset purities $K \ll N$ —analogous to the estimation of linear observables above.

Subset purity estimators feature prominently in bipartition Rényi-entropies given by the negative logarithm of the reduced density matrix purities $S^{(2)}(\rho_A) = -\log_2 \text{tr}(\rho_A^2)$, introduced in Ch. 1.1.2. Because the Rényi-entropy is symmetric with respect of its bipartitions (A, B) , only the lower weight subsystem needs to be considered, increasing the resulting statistical accuracy.

We follow up on a measurement budget for the classical shadow purity estimators and quantify the number of measurement samples required to predict multiple estimators in parallel with a given accuracy. A rigorous a priori bound can be derived as [40]

$$M \geq 6L3^K/(\epsilon^2\delta), \quad (3.23)$$

where M measurements suffice to ϵ -approximate L subsystem purities of weight K with at least the probability $1 - \delta$. While, we find a quadratic scaling in accuracy ϵ analogous to Eq. (3.20), observing more purity estimators L in parallel or increasing the success probability $1 - \delta$ leads to a linearly, no longer logarithmically, growing sample size. This is because the individual contributions to the subset purity estimator are not statistically independent. The number of measurements M scales exponentially in the subsystem size K , where an efficient prediction once more requires a comparably small subsystem size $K \ll N$.

The result in Eq. (3.23) reproduces the existing state-of-the-art, which was given by Pauli shadows [48]. However, Pauli shadows have limited applicability because the technique

requires random measurement samples from the exponentially large set of tomography measurements. A technical limitation that the SIC-based classical shadows discussed here overcome, as each outcome samples the full Hilbert space and randomization comes natural from the quantum measurement at the end of each shot. Finally, it should be noted that the properties converge (much) faster in experimental or numerical simulations than given by the scaling of Eq. (3.23). For example, in our experimental analysis of an absolute maximally entangled five qubit state [170, 171] in Ref. [40] the quadratic purity estimator from Eq. (3.22) converges at around 6000 shots, where the purity from SIC linear inversion in Eq. (3.18) is still highly unphysical.

Moreover, for a comparably high number of experimental samples M , accounting for all pairs $\binom{M}{2}$ can become computationally demanding. To soften this computationally heavy requirement, one can average bunches of S samples to reduce the number of combinations $\{m, m'\}$ down to $\binom{M/S}{2}$ in favour of accumulating fewer, but more accurate classical shadow estimators $\hat{\sigma}$. This becomes particularly useful in actual experiments [40], where a suitable batch size S needs to be evaluated on a case basis. In particular, the batch-size depends on the polynomial order of the requested estimators as well as their (sub)system sizes relative to the size of the analyzed quantum state. In the example of the five qubit state mentioned in the previous paragraph a batch-size of 100 shots each shadow estimator kept the data analysis most efficient in regards of both time consumption and convergence.

Beyond the exemplified linear and quadratic properties, arbitrary polynomial functions of the density matrix can in principle be derived with the same guidelines discussed here. See Ref. [40] on further notes. Higher order terms can be interesting for even more in-depth characterization studies.

The classical shadow framework uniquely circumvents the full density matrix reconstruction for predicting non-linear system properties, while it importantly does not rely on any preknowledge or assumptions on the state of the analyzed system. Combined with SIC QST, the randomization, typically limiting classical shadow applications, is no longer required. Measurement budgets from Eq. (3.20) and Eq. (3.23) denote that the capabilities of SIC-based classical shadows lie notably beyond existing characterization methods and suffice to resolve the third challenge from Fig. 1.3 on reducing the sample sizes.

Finally, Fig. 3.1 illustratively combines single-setting QST with classical shadows manifesting our scalable approach for characterizing large-scale systems. Further technical insights as well as experimental demonstrations are covered by the first publication of this thesis, presented in the upcoming Sec. 3.5.




Figure 3.1: **System characterization via classical shadows based single-setting QST.** We introduce a scalable and efficient characterization method for arbitrary quantum systems. Our approach features single-setting QST from Sec. 3.3 based on the SIC POVM which samples the full Hilbert space in every experimental sample through the use of higher-dimensional states. The technique is based on an ion-internal mapping sequence and works independently of the system size. Data analysis is realized via classical shadows, which enables the efficient estimation of non-linear system properties by reconstructing only relevant subsets of qubits and thereby significantly speeds up analysis. Importantly, no pre-knowledge or assumptions about the system to be analyzed are required. Single-setting QST and classical shadows match well as random measurement samples, necessary for classical shadows, are inherent to SIC POVMs. Our approach further processes data at the minimum dimension of the (reduced) density matrix and allows live update analysis in real-time, see text for details.

3.5 PUBLICATION: EXPERIMENTAL SINGLE-SETTING QUANTUM STATE TOMOGRAPHY

PRX Quantum 3, 040310 (2022)

submitted on 30 May 2022, accepted on 22 August 2022 and published on 21 October 2022
<https://doi.org/10.1103/PRXQuantum.3.040310>

Roman Stricker¹, Michael Meth¹, Lukas Postler¹, Claire Edmunds¹, Chris Ferrie², Rainer Blatt^{1,3,4}, Philipp Schindler¹, Thomas Monz^{1,4}, Richard Kueng⁵ and Martin Ringbauer⁵

¹ *Institut für Experimentalphysik, Universität Innsbruck, 6020 Innsbruck, Austria*

² *Centre for Quantum Software and Information, University of Technology Sydney, Ultimo, NSW 2007, Australia*

³ *Institut für Quantenoptik und Quanteninformation, Österreichische Akademie der Wissenschaften, 6020 Innsbruck, Austria*

⁴ *Alpine Quantum Technologies GmbH, 6020 Innsbruck, Austria*

⁵ *Institute for Integrated Circuits, Johannes Kepler University Linz, 4040 Linz, Austria*

The author to the present thesis executed the experiments, analyzed the data and wrote the manuscript.

Experimental Single-Setting Quantum State Tomography

Roman Stricker^{1,*} Michael Meth,¹ Lukas Postler,¹ Claire Edmunds¹,¹ Chris Ferrie,² Rainer Blatt^{1,3,4} Philipp Schindler,¹ Thomas Monz^{1,4} Richard Kueng,⁵ and Martin Ringbauer¹


¹*Institut für Experimentalphysik, Universität Innsbruck, Innsbruck 6020, Austria*

²*Centre for Quantum Software and Information, University of Technology Sydney, Ultimo, NSW 2007, Australia*

³*Institut für Quantenoptik und Quanteninformation, Österreichische Akademie der Wissenschaften, Innsbruck 6020, Austria*

⁴*Alpine Quantum Technologies GmbH, Innsbruck 6020, Austria*

⁵*Institute for Integrated Circuits, Johannes Kepler University Linz, Linz 4040, Austria*

 (Received 30 May 2022; revised 27 July 2022; accepted 22 August 2022; published 21 October 2022)

Quantum computers solve ever more complex tasks using steadily growing system sizes. Characterizing these quantum systems is vital, yet becoming increasingly challenging. The gold-standard method for this task is *quantum state tomography* (QST), capable of fully reconstructing a quantum state without prior knowledge. The measurement and classical computing costs, however, increase exponentially with the number of constituents (e.g., qubits)—a daunting bottleneck given the scale of existing and near-term quantum devices. Here, we demonstrate a scalable and practical QST approach that only uses a single measurement setting, namely symmetric informationally complete (SIC) positive operator-valued measures (POVMs). We implement these nonorthogonal measurements on an ion trap quantum processor by utilizing additional energy levels within each ion—without requiring ancillary ions to assist in measurements. More precisely, we locally map the SIC POVM to orthogonal states embedded in a higher-dimensional system, which we read out using repeated in-sequence detections, thereby providing full tomographic information in every shot. Combining this *SIC tomography* with the recently developed *randomized measurement toolbox* (“classical shadows”) proves to be a powerful combination. SIC tomography alleviates the need for choosing measurement settings at random (“derandomization”), while classical shadows enable the estimation of arbitrary polynomial functions of the density matrix orders of magnitudes faster than standard methods. The latter enables in-depth entanglement characterization, which we experimentally showcase on a five-qubit absolutely maximally entangled state. Moreover, the fact that the full tomography information is available in every shot enables online QST in real time (i.e., while the experiment is running). We demonstrate this on an eight-qubit entangled state (which has $2^8 \cdot 2^8 - 1 = 65\,535$ degrees of freedom), as well as for fast state identification. All in all, these features single out SIC-based classical shadow estimation as a highly scalable and convenient tool for quantum state characterization.

DOI: [10.1103/PRXQuantum.3.040310](https://doi.org/10.1103/PRXQuantum.3.040310)

Quantum systems are prepared in laboratories and in engineered devices such that their state delicately encodes quantum information essential for achieving goals in both science and technology. Any small adjustments, changes in the environment, or active control all change the state. Yet, an accurate mathematical description of the state

is a necessary component for most higher-level tasks. A crucial requirement for ensuring the performance of quantum devices is thus having methods for accurately determining the quantum states that have been prepared. The gold-standard approach for this fundamental task is quantum state tomography (QST) or simply tomography [1]. QST enables the full reconstruction of the system’s quantum state from an exponential number of measurements. Often, however, we are not even interested in the full quantum state, but rather certain features, like entanglement across a particular bipartition. However, it is not clear *a priori* how to access such nonlinear properties without resorting to full QST.

Formally, QST methods use an *informationally complete* set of measurements to reconstruct the complete

*roman.stricker@uibk.ac.at

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.




FIG. 1. Schematic illustration of Pauli and SIC tomography. *Pauli tomography* (a) uses three basis measurements per qubit to obtain tomographic information about an unknown N -qubit system; see (c). Each basis measurement is contingent on one of three possible unitary rotations—red boxes in (a). This produces a total of 3^N different measurement settings that need to be accessed. *SIC tomography* (b), on the other hand, uses the same measurement setting for each qubit; see (d). This nonorthogonal measurement is achieved by isometrically embedding each two-level system (qubit) into a larger four-level system (ququart)—blue boxes in (b)—and subsequently measuring this larger system. The experimental realization of this embedding within each ion is shown in Fig. 2.

description of the quantum state. The optimal measurement for collecting the necessary tomography data has long been known to be the so-called symmetric informationally complete positive operator-valued measures (SIC POVMs) [2,3]. SIC POVMs are constructed from the minimal number of d^2 measurements for a d -dimensional system, which are arranged in a way that maximizes the pairwise distance in Hilbert space. SIC POVMs are known to exist for several low-dimensional systems [4,5], and for qubits, take the form of four nonorthogonal vectors arranged as a tetrahedron in the Bloch sphere; see Fig. 1(d). While SIC POVMs uniquely offer access to complete tomographic information in every single experimental run (shot), implementing these measurements in practice is very challenging, requiring purpose-built setups [6,7], sequential measurement [8,9], or ancilla-assisted schemes [10,11]. Hence, tomography remains almost exclusively performed using the simpler, but overcomplete Pauli basis, requiring 3^N orthogonal measurement *settings*, each with 2^N outcomes for an N -qubit system. The resulting overhead effectively limits full tomography to system sizes of only a few qubits.

From a conceptual point of view, the qubit SIC POVM is favorable, as a single experimental shot already contains complete tomographic information. This distinct advantage has also been recognized in recent theoretical work on adaptive tomography for linear cost functions [11], as well as neural network quantum state tomography [12,13]. Experimentally, it is also much cheaper to repeat the same measurement setting many times than to switch settings an exponential number of times as the system size grows with Pauli tomography. So, this feature can have a significant impact in practice. Moreover, since full tomographic information is contained in every shot, the experimenter is free to stop the tomography at any point, e.g., when certain quantities of interest have converged. In contrast, other QST approaches would require at least one shot for each measurement setting to collect sufficient information in the first place. This discrepancy is particularly relevant when we are not interested in the full density matrix, but only in certain (nonlinear) properties, which often require far fewer shots than the 3^N minimum in Pauli tomography [14,15]. Finally, for randomized measurement schemes [16], where ideally a different measurement setting is required in each shot, the SIC approach obviates this requirement completely (“derandomization”), making these schemes even more practical. Hence, in most situations, SIC tomography has the potential to substantially outperform standard methods for tomography or for the direct estimation of state properties.

Here, we describe our realization of SIC POVMs in a trapped-ion quantum processor and their use for characterizing unknown quantum states. We put an emphasis on demonstrating the speed and robustness obtained from reducing the number of measurement settings in conjunction with new data processing techniques that come with rigorous accuracy guarantees. With our approach, we are able to comfortably reconstruct the full eight-qubit quantum state encoded in the electronic energy levels of calcium ions in *real time* using a standard laboratory computer. Moreover, we demonstrate the simultaneous real-time estimation of Rényi entropies across all bipartitions using a sampling-free classical shadow method [14]. This enables full entanglement characterization of arbitrary (but close to pure) quantum states with orders of magnitude fewer experimental shots than standard QST methods.

I. SIC TOMOGRAPHY

QST aims at reconstructing an unknown quantum state ρ from an informationally complete set of measurements, which spans the entire Hilbert space of the quantum system. The minimal number of measurement outcomes to reconstruct an arbitrary d -dimensional quantum state then is d^2 . An experimenter performs these measurements on many copies of ρ referred to as experimental shots, and attempts to reconstruct ρ from the observed measurement

counts. The standard approach to QST of N -qubit systems combines tomographic measurements of each individual qubit. The overcomplete Pauli basis is a particularly prominent choice; see Fig. 1(c). Three distinct measurement settings are required to evenly cover the Bloch sphere and obtain tomographic single-qubit information. Extending this to N -qubit systems produces 3^N distinct measurement settings that need to be explored; see Fig. 1(a).

In contrast to (single-qubit) Pauli basis measurements, (single-qubit) SIC POVMs provide access to a complete set of tomographic data from a single experimental shot. No change in measurement settings is required. This desirable feature extends to N -qubit measurements: a single measurement setting per qubit suffices to obtain tomographically complete data; see Fig. 1(b). SIC tomography utilizes (tensor products of) the single-qubit SIC POVM depicted in Fig. 1(d) (or a local rotation thereof). Following Naimark's dilation theorem [18], every POVM can be realized as a projective measurement on a higher-dimensional Hilbert space. Using this result, together with a qudit quantum processor [17], we realize the qubit SIC POVM as a projective measurement on a four-level system (ququart) employing two more states within each calcium ion. For this purpose, we map each qubit locally to a ququart using the unitary

$$\hat{M} = \begin{pmatrix} 1/\sqrt{2} & 0 & 0 & 1/\sqrt{2} \\ 1/\sqrt{6} & 1/\sqrt{3} & 1/\sqrt{3} & -1/\sqrt{6} \\ 1/\sqrt{6} & (-1)^{2/3}/\sqrt{3} & (-1)^{4/3}/\sqrt{3} & -1/\sqrt{6} \\ 1/\sqrt{6} & (-1)^{4/3}/\sqrt{3} & (-1)^{2/3}/\sqrt{3} & -1/\sqrt{6} \end{pmatrix}. \quad (1)$$

Here, the four two-dimensional vectors contained in the first two columns represent the measurement vectors of the qubit basis given in Fig. 1(d). The optimized gate sequence for locally mapping the measurement states from qubit to ququart is shown in Fig. 2(b). It consists of five single-qubit rotations, sequentially applied to each qubit, which we optimize such that local phase shifts are absorbed into the rotation angles of the single-qubit operations; see Appendix A 1. Therefore, our single-setting SIC tomography implementation remains the very same, independent of the qubit number. Our approach is particularly well suited to other quantum architectures, since many of today's information carriers manifest multilevel systems. Higher-dimensional systems have been demonstrated, frequently featuring in Rydberg ions [19], atomic and molecular systems [20,21], photonic systems [22], solid states [23], and superconducting platforms [24].

II. RECONSTRUCTION METHODS AND CLASSICAL SHADOWS

Even with a complete set of measurements, reconstructing ρ is computationally demanding, especially if one




FIG. 2. Experimental implementation of SIC tomography from Figs. 1(b) and 1(d). (a) Level scheme of a $^{40}\text{Ca}^+$ ion representing a qubit or ququart with important transitions marked: (blue) dipole transition for cooling and detection, (red) metastable quadrupole transition for encoding qubits and ququarts within the Zeeman submanifold, and (brown) additional transitions for repumping. (b) Gate sequence for locally mapping the SIC POVM from Fig. 1(d) from qubit level to four orthogonal basis states of a ququart denoted in (a). This enables full readout of the SIC POVM in a single experimental shot by means of a four-outcome projective measurement. (c) Experimental realization of SIC tomography composed of cooling (DC, doppler cooling; PGC, polarization gradient cooling; SBC, sideband cooling), preparation of the state to be analyzed, mapping from qubit to ququart according to the SIC POVM, and finally the four-outcome projective measurement. For the latter, three sequential fluorescence detections (DETs) are required [17]; see Appendix A 1 for details.

insists on enforcing physicality constraints. Two standard QST methods in the field are linear inversion (LI) and maximum likelihood estimation (MLE) [25], which can be readily applied to both Pauli and SIC data. LI provides an analytical approach to estimating ρ from a complete set of projectors Π_j that span the entire Hilbert space. Access to (approximations of) the associated probabilities $\hat{p}_j \approx \text{tr}(\Pi_j \rho)$ allows us to reconstruct the underlying state:

$$|\hat{\rho}\rangle\rangle = \sum_j \hat{p}_j \cdot S_p^{-1} |\Pi_j\rangle\rangle \quad (2)$$

with $|\Pi_j\rangle\rangle$ the vectorized projector, obtained by stacking the columns of Π_j . Furthermore, S_p^{-1} denotes the Moore-Penrose pseudoinverse of the measurement superoperator $S_p = \sum_j |\Pi_j\rangle\rangle\langle\langle \Pi_j|$ [26], and $\hat{p}_j = n_j/N_j$ is the observed

frequency of outcome j after averaging over N_j experimental shots. As an unconstrained method, the LI version bears the risk of producing unphysical estimators for ρ featuring negative eigenvalues. This is particularly pronounced when few experimental shots are used and is very problematic for estimating nonlinear observables. Physical constraints are thus typically introduced through MLE, which, following Ref. [26], can be approximated by a convex optimization problem

$$\begin{aligned} & \text{minimize} \quad \|W(S|\hat{\rho}) - |f\rangle\|_2 \\ & \text{subject to} \quad \hat{\rho} \geq 0, \quad \text{tr}(\hat{\rho}) = 1. \end{aligned} \quad (3)$$

Here, $S = \sum_j |j\rangle\langle\langle\Pi_j|$ denotes a change of basis operator, $|f\rangle = \sum_j n_j/N_j |j\rangle$ a column vector of the observed frequencies, and W a diagonal matrix of statistical weights W . Optimization is performed under the constraints that the estimator for ρ is positive semidefinite ($\rho \geq 0$) with unit trace [$\text{tr}(\rho) = 1$], i.e., it must be a valid quantum state. The convex optimization in Eq. (3) is computationally more efficient than full MLE and recovers the latter in the limit of large sample sizes. Nonetheless, the computational complexity remains intractable for anything but very small systems. LI is much more efficient by comparison. Nonetheless, inverting the superoperator S_p also becomes more challenging as the system size increases. Viewed as a matrix, every tomographically complete N -qubit superoperator S_p must have (at least) 4^N rows and (at least) 4^N columns. Performing the inversion row by row can offer some relief in terms of memory load, but the exponential number of multiplications remains challenging. Finally, physical constraints [$\rho \geq 0$ and $\text{tr}(\rho) = 1$] can be incorporated into LI by truncating negative eigenvalues to obtain the closest quantum state under the Frobenius norm [27–29], referred to as *projected least squares* (PLS). It should be noted that more principled, yet ever more computationally challenging, approaches exist [30,31].

So far, we have considered full QST, i.e., experimentally extracting a complete description of ρ , which is traditionally required for predicting certain properties of complex quantum systems, especially nonlinear functions, most prominently purity or entanglement. In large-scale systems, however, predicting such properties becomes very costly independent of the data acquisition (SIC, Pauli) and reconstruction (LI, MLE) method, both in regards to the number of required shots and in regards to the computational power required to analyze the data.

A promising alternative comes in the form of classical shadows [14,15] as a general-purpose method to construct classical descriptions of quantum states using very few experimental shots. Consequently, the classical shadow framework allows for the prediction of L different functions of the state with high accuracy, using order $\log(L)$ experimental shots. Importantly, the number of shots is

independent of the system size and saturates information-theoretic lower bounds. Moreover, target properties can be selected after the measurements are completed. A big drawback of existing classical shadow methods, however, is that they require a different measurement to be sampled randomly for each shot [16], which is demanding and slows down data acquisition. We show in the following that SIC POVMs naturally alleviate this sampling requirement (“derandomization”). SIC POVMs are thus an ideal choice for unlocking the full potential of the classical shadow framework. This has, in parts, been already pointed out in Ref. [11], which explores adaptive SIC tomography for linear cost functions inspired by variational quantum eigensolver (VQE). Instead, we are here interested in a general framework for efficiently predicting general linear and nonlinear properties of the quantum state.

Formally, classical shadows provide an alternative approach for a linear-inversion estimator deduced from SIC measurements on an N -qubit state ρ . Each experimental shot m , containing complete tomographic information, can be assigned to a size- N string $\hat{i}_{m,1}, \dots, \hat{i}_{m,N} \in \{1, 2, 3, 4\}^{\times N}$, where each quartic value keeps track of the SIC POVM outcome observed. For each shot m , an N -qubit estimator for the density matrix $\hat{\sigma}_m = \bigotimes_{n=1}^N (3|\psi_{i_{m,n}}\rangle\langle\psi_{i_{m,n}}| - \mathbb{I})$ is obtained, referred to as a *classical shadow*. A total of M such estimators can be experimentally inferred and accumulated to approximate ρ as

$$\hat{\rho} = \frac{1}{M} \sum_{m=1}^M \bigotimes_{n=1}^N (3|\psi_{i_{m,n}}\rangle\langle\psi_{i_{m,n}}| - \mathbb{I}) \xrightarrow{M \rightarrow \infty} \rho. \quad (4)$$

It is worthwhile to emphasize that each term in this average is highly unphysical, because each contains many negative eigenvalues that also become exponentially large. The crucial insight from Ref. [14] was that these unphysicalities quickly average out. In fact, empirical averages of such unbiased single-shot estimators converge much quicker than physical density matrix estimators that necessarily contain a bias. In the following, we present rigorous convergence guarantees that underscore this claim and refer to Appendix A 8 for additional context and exposition. Another crucial observation is that, compared to standard linear inversion in Eq. (2), the processing of classical shadows is performed in the dimension of the quantum state $2^N \cdot 2^N$, which is the minimal possible dimension for full tomography; see Appendix A 9 b. Moreover, predicting linear observables using classical shadows is even more efficient as it suffices to reconstruct a subset of ρ solely where operators act on. In Appendix A 10, we show how we can formalize these considerations to derive a measurement budget for estimating linear observables. Suppose that we are interested in estimating a total of $L \gg 1$ subsystem observables $\text{tr}(O_l \rho)$, where each O_l only

acts nontrivially on (at most) $K \leq N$ qubits. Then,

$$M \geq \frac{8}{3} 6^K \log(2L/\delta)/\epsilon^2 \quad (5)$$

measurements suffice to jointly ϵ -approximate all observables with probability (at least) $1 - \delta$. We emphasize that this is a novel, rigorous *a priori* bound based on minimal assumptions. In practice, convergence sets in (much) earlier. A full derivation and additional context is provided in the Appendix. For now, we merely point out that improvements of order 2^K are possible for the exponential scaling in case the observables in question have small Hilbert-Schmidt norm, as is the case for fidelities. Apart from linear observables, classical shadows also promise to allow for efficient estimation of nonlinear functions; see Appendix A 11. Whereas the full scope of nonlinear functions is covered in the Appendix, here we focus on a quadratic estimator in form of the (subsystem) purity

$$\hat{p}_{(M)} = \binom{M}{2}^{-1} \sum_{m < m'} \text{tr}(\hat{\sigma}_m \hat{\sigma}_{m'}), \quad (6)$$

as purity generally obeys hard convergence properties and also provides a means to measuring entanglement via Rényi entropies (see below). The latter are given by the negative logarithm of subsystem purities, corresponding to certain bipartitions of the state. Similar to before we can derive a measurement budget, where

$$M \geq 6L3^K/(\epsilon^2\delta) \quad (7)$$

measurements allow for ϵ -approximating L subsystem purities of size (at most) K with probability (at least) $1 - \delta$. We emphasize that this is again a novel, rigorous *a priori* bound based on minimal assumptions; actual convergence sets in much quicker. However, this bound still marks an improvement over the best available results for purity estimation with randomized single-qubit measurements [32,33]. The improvement follows from exploiting the geometric structure of SIC POVMs.

At this point it is also worthwhile to note that the proof techniques behind Eqs. (5) and (7) are novel. The original proof techniques from Ref. [14] do not extend to SIC POVM measurements, because they rely on a powerful structural property—a so-called 3-design [34–37]—that Pauli measurements have, but SIC POVMs do not. We refer the reader to Appendix A 11 for additional context and complete proofs. These results demonstrate that classical shadows in combination with SIC measurements offer a powerful tool set for measuring entanglement in a scalable fashion [32,33]. Whereas our experimental studies primarily focus on quadratic estimators of Rényi entropy [Eq. (9)], classical shadows can be extended to higher-order estimators following the same principles: (i) rewrite a degree- d polynomial as $\text{tr}(O\rho^{\otimes d})$, (ii) replace each ρ with

an independent classical shadow $\hat{\sigma}_m$, and (iii) average over all different subselections of distinct classical shadows. We refer the reader to Refs. [14,32,33] for details. Finally, we remark that classical shadow estimators from Pauli basis measurements can in principle be obtained by randomly sampling over the measurement settings from shot to shot. Although experimentally feasible, this is highly impractical. Remarkably, Pauli basis measurements also appear to lead to slower convergence than SIC measurements.

III. EXPERIMENTAL SETUP AND SIC IMPLEMENTATION

Experimental results in this work are obtained on a trapped-ion quantum processor based on a linear string of $^{40}\text{Ca}^+$ ions, each encoding a single qubit in the (meta)stable electronic states $\{S_{1/2}(m = -1/2) = |0\rangle, D_{5/2}(m = -1/2) = |1\rangle\}$ [38]. A universal set of quantum gate operations is realized upon coherent laser-ion interaction and comprises arbitrary local single-qubit rotations together with two-qubit entangling operations, enabling all-to-all connectivity. A binary qubit measurement is implemented by scattering on the dipole transition, where fluorescence is only observed if the ion is in the $|0\rangle$ state, thereby separating the computational basis states $\{S_{1/2}(m = -1/2) = |0\rangle, D_{5/2}(m = -1/2) = |1\rangle\}$; see Fig. 2(a). Equivalent control over the entire S - and D -state Zeeman manifolds allows for encoding a higher-dimensional quantum decimal digit (qudit) with up to eight levels in each ion, combined with full fluorescence readout of the whole qudit space [17]; see Appendix A 1.

The present work builds on this capability by utilizing up to four levels per ion to implement SIC POVMs. To this extent, two additional levels $D_{5/2}(m = -3/2) = |2\rangle$ and $D_{5/2}(m = +1/2) = |3\rangle$ are taken into account; see Fig. 2(a). Upon applying the mapping sequence depicted in Fig. 2(b), each qubit is locally extended to a ququart where each basis state encodes one SIC vector. A four-outcome projective measurement is implemented by three sequential fluorescence detections, where before the second detection the population between states $|0\rangle$ and $|1\rangle$ is flipped and likewise before the final detection the states $|0\rangle$ and $|2\rangle$ are flipped. This enables us to evaluate the full ququart state probabilities from three binary outcomes. The entire experimental sequence comprising cooling, state preparation, mapping the SIC POVM to the ququart, and four-outcome readout is shown in Fig. 2(c), which we refer to as a single experimental shot. We remark that this SIC tomography procedure works independently on each qubit and that such a single experimental shot delivers the complete tomographic information of the N -qubit system.

IV. RECONSTRUCTION TIME

While SIC POVMs can significantly speed up data acquisition, the classical resources needed for

reconstructing and storing the quantum state ρ is typically an additional bottleneck in QST [see Eqs. (2) and (3)]. In the following, we compare the computational time for reconstructing ρ following various tomography approaches. For the moment, we solely focus on the classical reconstruction time, which is dominated by the size of the involved matrices, and discuss the convergence properties of the various methods later and in Figs. 7, 12, and 14 in the Appendix. For a system of N qubits, we consider tomography data composed of $M = 100 \cdot 3^N$ shots. This corresponds to 100 shots for each measurement setting used in Pauli tomography, which, on the trapped-ion platform, has proven to be a good trade-off accounting for statistics, systematic drifts, and measurement time.

Figure 3 illustrates the classical reconstruction time versus the number of qubits N for experimental data used throughout this manuscript. Whereas absolute time reflects a laboratory desktop computer, relative scaling between methods remains generally valid. Note that the plot is double logarithmic in the number of shots $M = 100 \cdot 3^N$. While MLE methods always obey physical constraints, solving the convex optimization problems is costly and only feasible for small system sizes. We find MLE with SIC measurements to be more efficient, due to handling matrices of maximum size $4^N \cdot 4^N$, in contrast to $6^N \cdot 4^N$ for the overcomplete Pauli basis. However, MLE quickly becomes infeasible as the number of qubits increases. SIC LI suffers an initial offset to Pauli LI due to computing ququart instead of qubit state probabilities, which, for the MLE approaches, is masked in the overhead of convex optimization from Eq. (3). As the number of qubits increases, SIC makes up for this, as the computations are performed in a smaller dimension. Although computationally much cheaper than MLE, even LI becomes increasingly costly due to the memory requirements of processing the inverse superoperator S_p from Eq. (2). Already for six qubits this requires 268.4 and 3100 MB for SIC and Pauli measurements, respectively. Scaled up further, this will rapidly exceed the memory of today's computers. Alternatively, inversion of S_p could be done row by row to reduce memory load, but this would be more time consuming than precalculating the inverse S_p^{-1} as we have done here. While linear observables under LI are proven to quickly converge (see Appendix A 10), nonlinear functions suffer from non-physical properties in the form of negative eigenvalues; see Fig. 4(b). Furthermore, PLS adds negligible computational overhead over LI and is thus neglected in this comparison. PLS does, however, affect the convergence and accuracy of the estimators, as shown and discussed further below.

Finally, we find the best scaling for the SIC-based classical shadows from Eq. (4), where data are processed at the dimension of the density matrix, $2^N \cdot 2^N$, avoiding matrix inversion or optimization altogether. Instead, individual experimental shots are accumulated, offering convenient




FIG. 3. Classical run time comparison for tomographic reconstruction methods. Comparison of the computational run time for state reconstruction using LI [Eq. (2)] or MLE [Eq. (3)] for both Pauli [Fig. 1(c)] and SIC tomography [Fig. 1(d)], as well as SIC-based classical shadows [Eq. (4)] as a function of the system size. For each qubit number N , the reconstruction considers $M = 100 \cdot 3^N$ experimental shots (i.e., 100 shots per Pauli basis; see the text). The analysis is conducted on a standard desktop computer and plotted double logarithmically in the number of experimental shots. We find that both MLE methods require the highest computational resources, with SIC MLE significantly faster as the processed dimension is lower with $4^N \cdot 4^N$ compared to $6^N \cdot 4^N$ for Pauli tomography. LI with SIC measurements shows an initial time offset to Pauli tomography arising from handling ququart (dim = 4^N) instead of qubit data (dim = 2^N), but grows much more slowly with system size. Among the LI methods, we find the classical shadow reconstruction to be the fastest for an increasing number of qubits, as it accumulates each individual shot in dimension $2^N \cdot 2^N$ [see Eq. (4)], avoiding the costly matrix inversion in dimension $4^N \cdot 4^N$ or $6^N \cdot 4^N$ as required for SIC and Pauli tomography, respectively.

updates of ρ for every new set of data. As a consequence of this individual accounting for every shot, the computational complexity of this method grows linearly with the total number of shots. While this linear overhead leads to slightly worse performance for very small systems, it is more than compensated by the improved exponential scaling with qubit number ($2^N \cdot 2^N$ versus at least $4^N \cdot 4^N$) for large systems. Hence, the SIC-based classical shadows clearly outperform all other methods for six or more qubits. Note that, for large-scale systems, the gap between classical shadows and standard LI becomes even larger as row by row LI becomes requisite.

V. ESTIMATING PROPERTIES OF THE STATE

Here we shift our attention towards convergence of the different tomography estimators. In particular, we showcase the classical shadows' unique feature of efficiently predicting nonlinear properties of even large-scale quantum systems. To this end, we experimentally perform




FIG. 4. SIC-based classical shadow tomography of a five-qubit absolutely maximally entangled (AME) state. We compare the convergence of LI [Eq. (2)] and classical shadow tomography using SIC measurements [Eqs. (4)–(6)]. We additionally use PLS following Refs. [27–29], explained in the text. (a) In terms of convergence, we find both SIC LI from Eq. (2) and SIC shadows LI from Eq. (4) to overlap, which is expected due to their similarity. We remark that in terms of reconstruction SIC shadows LI is computationally more efficient; see Fig. 3. PLS incorporates physical constraints [$\text{tr}(\rho) = 1$ and $\rho > 0$], but underestimates the fidelity for a small number of shots and converges very slowly, requiring at least an order of magnitude more shots (indeed slower than MLE, additionally confirmed by experiments in Appendix A 2 and numerical simulations in Appendix A 6). Here, connected data points represent an individual experimental run, while shaded regions represent 1 standard deviation around the mean value of multiple batches of experimental data at identical shot numbers. (b) When estimating purity, LI shows highly unphysical behavior under insufficient statistics and does not converge in the plot. Physical constraints can be corrected by PLS, although at the cost of much slower convergence. On the other hand, classical shadow purity estimators [Eq. (6)] display very quick convergence. The dash-dot line indicates the ideal value. (c) Estimated Rényi entropies [Eq. (9)] from subset purities across all bipartitions converge very quickly for all methods. The lack of difference between the methods is likely due to the small system sizes, and for increasing bipartition sizes, we expect similar behavior as in (b). Note that bipartitions are denoted by tuples (2, 3) qubits and (1, 4) qubits, referring to the number of qubits in each part.

tomography on a five-qubit AME state [39]

$$\begin{aligned}
 2\sqrt{2} |\Omega_{5,2}\rangle = & |00000\rangle + |00011\rangle + |01100\rangle \\
 & - |01111\rangle + |11010\rangle + |11001\rangle \\
 & + |10110\rangle - |10101\rangle. \quad (8)
 \end{aligned}$$

AME states are the most entangled states in the sense that they are maximally entangled in all bipartitions [40]. This makes them interesting for applications in quantum error correction [41], quantum teleportation, quantum secret sharing, and superdense coding [42]. Alas, their general existence remains unknown for all but the smallest systems.

We characterize the five-qubit AME state from Eq. (8) using SIC tomography data. The results in Fig. 4(a) indicate that both LI methods converge very quickly in fidelity, as is likewise expected for all linear observables; see Appendix A 10. In terms of convergence SIC LI from Eq. (2) and SIC shadows LI from Eq. (4) expectantly overlap due to their similarity, whereas the latter is found to be more efficient in terms of reconstruction; see Fig. 3. In contrast, incorporating physical constraints with PLS drastically slows convergence [29], because truncation of negative eigenvalues produces a bias [43]; see also Fig. 12 in Appendix A 6. We repeatedly confirm this bias by experiments in Appendix A 2 as well as numerical simulations in Appendix A 6. We further estimate the state’s

purity, as an example of an archetypal nonlinear function of the full state. Here LI shows highly unphysical results under insufficient statistics that do not converge in the plot, while the classical shadow purity estimators from Eq. (6) converge rapidly after only about 3000 shots; see Fig. 4(b). Finally, most inspired by applications, we probe the state’s entanglement by estimating all second-order Rényi entropies

$$S^{(2)}(\rho_A) = -\log_2 \text{tr}(\rho_A^2) \quad (9)$$

with the reduced density matrix ρ_A for part A of a bipartition (A, \bar{A}) together forming ρ [44]. In Fig. 4(c), we present results on Rényi entropies following Eq. (6) and particularly cover all bipartitions denoted by tuples (1, 4) qubits and (2, 3) qubits, referring to the number of qubits in each part. Note that, for classical shadow prediction, only the subset qubits in the smaller partition need to be taken into account. This leads to a drastic speedup of the analysis for larger scale systems. Moreover, all predictions can be analyzed after the data have been acquired.

For comparison, we also analyze the AME state with Pauli tomography; see Fig. 7 in Appendix A 2. Slowest convergence is consistently found for PLS, which is also notably slower than regular MLE. We additionally confirm this by numerical simulations, considering only shot noise, i.e., statistical noise; see Appendix A 6). Given only a few experimental shots, SIC tomography outperforms

Pauli tomography in the case of MLE reconstruction, likely because the SIC POVM provides the optimal information gain per shot. Curiously, however, for large shot numbers, Pauli MLE starts to outperform SIC MLE in terms of convergence. We suspect this to be a result of the over-completeness of the Pauli basis, where very large shot numbers may lead to improved accuracy for each orthogonal direction. Note that all methods converge to the same point, as verified in numerical simulations; see Fig. 12 in Appendix A 6.

In conclusion, we find SIC tomography to be preferable over Pauli tomography in regards to both classical computation time and convergence speed. At the same time, the underlying classical shadow formalism provides the potential for scaling to large quantum systems. We emphasize that the moderately lower quantitative performance of SIC tomography observed in our data is not inherent to the method, but due to experimental imperfections, i.e., the additional overhead in mapping SIC POVMs to ququarts, and the four-outcome readout; see Appendix A 1. These technical imperfections can be overcome in future devices.

VI. LIVE-UPDATE TOMOGRAPHY

Since the SIC POVM contains full tomographic information within each shot, it provides a unique way to speedup QST. Combined with classical shadows, which work by accumulating estimates in each shot according to Eq. (4), SIC tomography can be performed in real time (or “online”), i.e., a live update is performed for every new set of experimental shots. Apart from reducing time overheads by performing experiment and analysis at the same time, this approach has the big advantage that the experiment can be stopped as soon as all properties of interest are (believed to be) accurately estimated. Based on these ideas, we demonstrate a live reconstruction of a maximally entangled eight-qubit state. Specifically, we use a Greenberger-Horne-Zeilinger (GHZ) state with an additional local $\pi/4$ rotation of all qubits as a proxy for a generic maximally entangled state that is not aligned with any of the tomography bases to provide a fair comparison.

Figure 5(a) presents results on the estimation of fidelity, purity, and Rényi entropies [Eq. (9)], the latter for all possible (2, 6)-qubit bipartitions. Purity of the full eight-qubit state is found to converge the slowest, after around 1000 s. This is still a drastic speedup compared to LI, which, after an order of magnitude longer averaging time still produces unphysical results. We qualitatively conclude that an estimated property has converged by the time the observed values no longer change significantly with additional data, while neglecting slow experimental drifts. Fidelity converges after less than 500 s and Rényi entropies saturate almost immediately on the presented timescale. We remark

that both curves for SIC shadows LI and SIC shadows quadratic overlap as a consequence of the fast converging two-qubit subsets. Figure 8 in Appendix A 3 additionally characterizes data on the eight-qubit state in postprocessing to deliver both uncertainty estimates as well as results on all bipartitions. In Fig. 5(a), live updates are tracked up to 2500 s, which is the limit beyond where the analysis starts to take longer than the data acquisition [due to the quadratic scaling in the number of shots for purity estimation; see Eq. (6)]. We extend the discussion about the time relation between data acquisition and data processing in Fig. 5(b), where we acquire 100 experimental shots in about 2.4 s and show the entire 12 500 s of data taking. Over this time, we observe that the computational time nicely follows the expected quadratic growth with shot number, relating to the number of approximated ρ to compare in Eq. (6). When focusing on entanglement properties, the full-system purity can be excluded from the analysis. We then find that all bipartition Rényi entropies can be estimated in real time throughout the entire time of data acquisition. On top, simulations suggest that even on an 18-qubit state all bipartition Rényi entropies can be estimated live for around 1000 s.

VII. DISCUSSION AND OUTLOOK

We have demonstrated that real-time SIC tomography enables the prediction of Rényi entropies in less experimental shots than required for a minimal Pauli tomography implementation (see Fig. 5). Depending on the state and estimator, SIC tomography has the potential to significantly speed up the prediction of many more properties of quantum states. Beyond Rényi entropies, we challenge SIC tomography in a state identification game partly inspired by Ref. [45]. We show that it excels over other state-of-the-art methods (see Fig. 10 in Appendix A 5). For the challenge in question, SIC tomography required less than 20 shots to correctly identify a randomly chosen four-qubit linear cluster state among 16 orthogonal state possibilities, clearly outperforming any Pauli QST method. Note that such a speedup in favor of SIC tomography will become even more pronounced on up-scaled systems.

In practice, we accumulate 100 shots for each classical shadow approximating ρ following Eq. (6) that we refer to as *batching*, which in this particular case allows us to analyze the data in real time until all properties of interest converge. This postprocessing trick enables a trade-off between computational time and convergence speed, which is studied thoroughly through numerical simulations in Fig. 9 in Appendix A 4. A suitable batch size must be decided on a case-by-case basis. From a practical point of view, the experimental noise might also fall into consideration as it affects the targeted accuracy. Thus, batching experimental shots for the analysis comes as a handy tool for reducing analysis time with a limited effect




FIG. 5. Live-update SIC tomography of a maximally entangled eight-qubit state. We prepare a locally rotated GHZ state and probe the state via live-update SIC tomography. The local $\pi/4$ rotation ensures that the state is not aligned with any tomography basis and can serve as a proxy for an arbitrary entangled state. Data analysis is performed in parallel to data acquisition to get the quickest possible feedback. Shot-by-shot reconstruction is realized via classical shadows [Eqs. (4)–(A20)]. (a) Experimental results on purity (left), fidelity (middle top), and all (2,6)-qubit bipartition Rényi entropies (middle bottom) obtained live in a time regime where the data acquisition time dominates. We find the purity from classical shadows to have converged after less than 1000 s, with the other measures converging significantly faster. Fidelity converges very fast below 500 s and Rényi entropies saturate almost immediately. Note that, for the latter, the curves for SIC shadows LI and SIC shadows quadratic overlap due to the fast convergence for small subsystem sizes. The dash-dot line indicates the ideal value. (b) Comparison of the time for data acquisition (left-pointing triangles) following the expected linear curve along times curves for data analysis utilizing different methods and metrics covering the entire 12 500 s of data taking. When analyzing fidelity, purity, and all (2,6)-qubit bipartition Rényi entropies (circles) as shown in the experimental results of (a), live update is possible for around 2500 s. This is expected, since the shot-by-shot data analysis scales quadratically with the total shot number; see Eq. (6). Estimating only Rényi entropies for all bipartitions (diamonds), on the other hand, remains feasible for the entire time as it overlaps with data acquisition (left-pointing triangles) following the linear curve. Additionally, we simulate the estimation of all Rényi entropies of an 18-qubit (18q) state (crosses) by considering the data acquisition overhead (left-pointing triangles), demonstrating that this in principle also remains feasible to perform live.

on convergence and thereby extends the window for doing real-time analysis.

Though SIC POVMs are optimal for tomography [4], they are not known to exist in every dimension. Indeed, just resolving this issue would have far-reaching consequences for the foundations of mathematics [46]. For a given Hilbert space of dimension $d = 2^N$, a SIC constructed in that space is referred to as a *global* SIC, whereas a measurement composed of N two-dimensional SICs is referred to as a *local* SIC. If they indeed exist, global SIC tomography would be sample optimal in the sense that it saturates fundamental lower bounds from information theory [29,47] (joint measurements across many state copies could still yield further improvements [47,48]). Global SIC measurements would, however, be challenging to realize on quantum hardware. Quadratic circuit sizes (in the number of qubits) may be necessary, because the associated SIC states form a 2-design—a concept closely related to chaos and information scrambling [49]. And recent works provide lower bounds on the minimum circuit depth required to achieve information scrambling [50]. We conclude that, although local SICs may be less efficient than global SICs in terms of measurement complexity [29,51], they are much cheaper to implement. In addition, they are informationally complete and optimal amongst all possible local measurements [51].

To further improve system predictions, one might want to adapt the measurement basis depending on the state to analyze, potentially reducing quantum shot noise and offering a speedup in convergence. The concept of adaptive tomography follows those means, where the measurement setting is adjusted based on the outcomes of prior measurements [11,52–56]. In the most extreme cases, the measurement settings are changed after each shot. Although this has been demonstrated in some settings [57–61], it requires both additional classical computation and physical setting adjustments, rendering it potentially very time consuming.

In contrast, the SIC POVM representation from Fig. 1(d) turns out to be very efficient and practical. Moreover, we studied convergence properties of states with different overlap to either SIC POVM or Pauli basis; see Appendix A 7. Interestingly, purely local states analyzed by SIC significantly increase convergence when a component along one of the nonorthogonal SIC POVMs vanishes, obeying the concept of unambiguous state discrimination [62]. For randomly aligned, or correlated states, however, the effect vanishes, making the local rotation of the SIC POVM irrelevant. However, for estimating Pauli observables, the rotation does matter, with the optimal alignment given such that the overlap with the Pauli basis is symmetric; see Fig. 15(b) in Appendix A 7. This rotated SIC will be notably useful for VQE applications, as those rely on

efficient Pauli observable measurements. Moreover, investigating state dependencies more comprehensively holds the potential for interesting future research.

Finally, we emphasize that the combination of SIC-POVM measurements with the classical shadow formalism is well suited for directly estimating higher-order polynomials of an unknown density matrix ρ . As discussed in Appendix A 12 and following ideas from Refs. [32,33], this opens the door for mixed-state entanglement characterization of large-scale systems in real time, a likely requirement in the development of scalable quantum technology.

DATA AVAILABILITY

The data underlying the findings of this work is available at [63].

ACKNOWLEDGMENTS

R.S., L.P., M.M., C.E., M.R., P.S., T.M., and R.B. gratefully acknowledge funding by the U.S. ARO Grant No. W911NF-21-1-0007. We also acknowledge funding by the Austrian Science Fund (FWF), through the SFB BeyondC (FWF Project No. F7109), by the Austrian Research Promotion Agency (FFG) under Contract No. 872766, by the EU H2020-FETFLAG-2018-03 under Grant Agreement No. 820495, and by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via the U.S. ARO Grant No. W911NF-16-1-0070 and the US Air Force Office of Scientific Research (AFOSR) via IOE Grant No. FA9550-19-1-7044 LASCEM. C.F. is supported by the Australian Department of Industry, Innovation and Science (Grant No. AUSMURI000002). This project has received funding from the European Union's Horizon 2020 research

and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 840450. It reflects only the author's view, the EU Agency is not responsible for any use that may be made of the information it contains. T.M. and R.B. acknowledge support by the IQI GmbH.

M.R. conceived the project. C.F. and R.K. derived the theory results. R.S., L.P., M.M., C.E., P.S., T.M., and M.R. performed the experiments. R.S. analyzed the data. R.K. (theory), M.R. and R.B. (experiment) supervised the project. All authors contributed to writing the manuscript.

The authors declare no competing interests.

Note added.—In the final stages of this project we became aware of independent and complementary research using SIC POVMs on a superconducting quantum processor [64].

APPENDIX: EXPERIMENTAL SINGLE-SETTING QUANTUM STATE TOMOGRAPHY

1. Experimental toolbox

Experimental implementations here and in the main text are performed on a trapped-ion quantum computer, which is schematically shown in Fig. 6(a). The device operates on a string of $^{40}\text{Ca}^+$ ions stored in ultrahigh vacuum using a linear Paul trap. Each ion acts as a qubit encoded in the electronic levels $S_{1/2}(m = -1/2) = |0\rangle$ and $D_{5/2}(m = -1/2) = |1\rangle$ denoting the computational subspace [38].

Quantum state manipulation is realized upon coherent laser-ion interaction. A universal gate set comprises addressed single-qubit rotations with an angle θ around the rotation angle ϕ in the equatorial plane of the form $\mathbb{R}^i(\theta, \phi) = \exp(-i\theta(\cos\phi\sigma_j^x + \sin\phi\sigma_j^y)/2)$ with the Pauli operators $\sigma_j^x = X_j$ or $\sigma_j^y = Y_j$ acting on the j th




FIG. 6. Schematic of the trapped-ion quantum processor using $^{40}\text{Ca}^+$. (a) Each ion within the linear string encodes a qubit ($|0\rangle$, $|1\rangle$) or ququart ($|0\rangle$, $|1\rangle$, $|2\rangle$, $|3\rangle$). A universal gate set is realized upon coherent laser-ion interaction using tightly focused beams, addressing single ions for local gates (bright red) and pairwise ions for entangling gates (dark red). Alternatively, we can globally address all ions simultaneously. Readout is performed via collective fluorescence detection (DET); see the text for details. (b) $^{40}\text{Ca}^+$ level scheme with a dipole transition (397 nm) for cooling and detection, a metastable quadrupole transition (729 nm) for encoding qubits and ququarts within the Zeemann submanifold as well as transitions for repumping at 854 and 866 nm. The labeled transitions (1, 2, 3) provide coherent connection between all ququart states.




FIG. 7. Comparison between SIC and Pauli tomography for a standard set of reconstruction methods. Using the five-qubit AME state from Fig. 4 in the main text, we analyze fidelity, purity, and negativity from Eq. (A1). Because of the increased experimental complexity, we observe that SIC methods generally converge to slightly lower values than Pauli; however, this is not inherent to the method, but rather the implementation. (a) Fidelity converges quickest for LI approaches, as expected for general linear functions; see Appendix A 10. Convergence of PLS, however, is very slow compared to all other methods. SIC MLE at low shot numbers performs better than Pauli MLE as SIC measurements provide the maximum information gain. (b) SIC-based classical shadows (“SIC shadows quadratic”) demonstrate the best convergence for purity. MLE methods converge similarly to (a), but LI methods show very slow convergence with highly unphysical results, making their use problematic in practice. (c) As a commonly used entanglement measure, we evaluate the negativity of the reconstructed states; see the text for details. As expected, this property converges the slowest, since it is a global property of the density matrix (making this non-scalable), with LI again showing highly unphysical results. Bright shaded regions represent 1 standard deviation around the mean value when averaging multiple sets at the given shot number. Black horizontal lines denote ideal values.

qubit, together with two-qubit Mølmer-Sørensen entangling gate operations $MS_{i,j}(\theta) = \exp(-i\theta X_i X_j / 2)$ [65]. Multiple addressed laser beams, coherent among themselves, allow for arbitrary two-qubit connectivity across the entire ion string [17]. Optionally, all ions can be addressed simultaneously using a global beam to enable both collective local operations as well as collective entangling operations $MS(\theta) = \exp(-i\theta \sum_{j < \ell} X_j X_\ell / 2)$. We choose whatever is more efficient for the underlying experiment. Initial state preparation in $|0\rangle$ is reached after a series of Doppler cooling, polarization-gradient cooling, and sideband cooling. Readout is realized by exciting a dipole transition coupled to the lower qubit level $|0\rangle$ and collecting its scattered photons, from which the computational basis states $|0\rangle$ and $|1\rangle$ can be identified. Thereby, a qubit’s state is revealed by accumulating probabilities from multiple experimental runs. The dipole laser collectively covers the entire ion string, which enables a complete readout in a single measurement round. Additional pump lasers support efficient state preparation as well as cooling and prevent the occupation of unwanted metastable states outside the computational subspace $\{|0\rangle, |1\rangle\}$. Beyond the qubit level, we hold equivalent control over the entire S - and D -state Zeeman manifolds with up to eight levels in each ion, allowing us to encode a higher-dimensional quantum decimal digit (qudit); see Fig. 6(b). In this work we make use of up to four levels, denoting a ququart via additionally employing $D_{5/2}(m = -3/2) = |2\rangle$ and $D_{5/2}(m = +1/2) = |3\rangle$ alongside both qubit states. Ququart readout can be performed via three consecutive

fluorescence detections, where before the second detection the population between $|0\rangle$ and $|1\rangle$ is switched and before the third and final detection the population between $|0\rangle$ and $|2\rangle$ is switched; see Fig. 2(c) in the main text. Combining these three binary outcomes enables us to evaluate the ququart state probability within a single experimental run. Note that fluorescence detection in the case of measuring a bright state, even on only a single ion, heats up the entire ion string due to photon scattering as the qubits motion is coupled through the coulomb interaction. This is counteracted by a sequence of Doppler and polarization-gradient cooling after each individual detection, independent of its outcome, to keep the quality of postmeasurement bit-flip operations high and thereby suppress detection errors.

The last paragraph of this experimental setup section is dedicated to technical errors limiting our tomography experiments. Performance on *SIC tomography* is generally found to be moderately lower compared to Pauli tomography; see Fig. 4 in the main text. Evidently, this performance decrease is not inherent to the method, but rather owed to technical errors for two main reasons. (1) The SIC tomography implementation generates an overhead of five local pulses per qubit used for mapping qubit to ququart, depicted in Fig. 2(b), as well as two additional bit flips realizing the four-outcome readout; see Fig. 2(c). In contrast each Pauli *setting* requires just one local pulse per qubit, yet requires three orthogonal measurements per qubit to extract full tomography information. Our trapped-ion setup has a single-qubit gate fidelity of 0.9994(3) estimated from

randomized benchmarking as well as a two-qubit gate infidelity of roughly 0.98(1) estimated from a decay of fully entangling MS gates [17]. The latter two-qubit gate fidelity might slightly fluctuate from pair to pair. Moreover crosstalk to adjacent ions have an influence. (2) After each insequence detection the CCD camera demands for a 3-ms pause to process the data, before the upcoming sequence continues. However, we utilize this time for recoiling the ion string via Doppler and polarization-gradient cooling. During this pause the ions are exposed to amplitude damping due to spontaneous decay from the upper D -state states, having a lifetime of about 1 s. Accounting for all this throughout measurement taking, we observe a loss of fidelity per qubit between Pauli and SIC tomography of less than 1%. Thus, extracting complete tomography information in a single experimental run, i.e., shot, comes at the expense of a more complex experiment, which is however of a technical nature and can be overcome in future devices with reduced single-qubit error rates as well as with faster processing CCD cameras, which nowadays already exist. More importantly, only SIC tomography offers the unique potential to predict nonlinear properties in large-scale systems, as pointed out here further below and in the main text.

2. Comparing tomography methods

We start off by presenting complementary experimental data covering the five-qubit AME state from Fig. 4 in the main text. Whereas in the main text the focus was on scalable approaches, especially SIC-based classical shadows, here we compare these results to Pauli tomography according to Figs. 1(a) and 1(c) using linear inversion and MLE reconstruction following Eq. (3). Generally, linearly reconstructed density matrices exhibit negative eigenvalues in the case of insufficient statistics, which manifest themselves particularly in unphysical values for nonlinear functions of the density matrix, as indicated by purity values above 1. Physical constraints can be imposed on LI, through truncation of negative eigenvalues following Refs. [27–29], which we refer to as PLS. Importantly, PLS adds only a negligible computational overhead to LI. Consequently, the covered set of tomography approaches is representative in the field of quantum computation and quantum information. Figure 7 depicts results on fidelity, purity, and negativity, with the latter being a common measure of entanglement, albeit one that is challenging to access with experiments.

Note that, for five-qubit Pauli tomography, $3^5 = 243$ settings are required, where for this particular case, each setting is repeated 100 times, leading to the stated maximum shot number of 24 300. A hundred shots per setting proves to be a good trade-off in the trapped-ion platform, accounting for both statistics and systematic drifts in the experiment. The moderately lower performance (in terms

of the numerical values) of SIC tomography is not inherent to the method, but comes from technical imperfections due to experimental overhead. In particular, the mapping of the SIC POVM to quaternary states as well as the four-outcome readout, both essential for single-setting tomography, add experimental complexity; see Appendix A 1.

Figure 7(a) shows that the fidelity converges very quickly for LI approaches, as is expected for general linear operators; see Appendix A 10. Interestingly, Pauli MLE performs worse than all other methods at very few shots, while SIC-based methods perform very well, even for low shot numbers, since this measurement extracts the maximal amount of information for a generic state. PLS, on the other hand, despite the computational efficiency, converges much slower than the other methods, even including MLE [29]. For quadratic measures in Fig. 7(b), it becomes clear that LI methods produce highly unphysical results that take a long time to converge, limiting the usefulness of these methods in practice. PLS solves this problem, but again shows very slow convergence. Both problems are solved by the SIC-based classical shadow purity estimator from Eq. (A20), which demonstrates both fast convergence and accurate (physical) estimates. Finally, we study negativity as a commonly used measure of quantum entanglement [66]:

$$\mathcal{N}(\rho) = \frac{\|\rho^{\Gamma_A}\|_1 - 1}{2}. \quad (\text{A1})$$

Here ρ^{Γ_A} represents the partial transpose with respect to subsystem A of a bipartition (A, B) together forming ρ . The 1-norm in Eq. (A1) denotes the absolute sum of all negative eigenvalues given by ρ^{Γ_A} . By construction, the partial transpose of a separable state cannot have negative eigenvalues, such that the negativity vanishes. Quantum negativity is an entanglement monotone: if it is positive then the underlying state must be entangled. The converse, however, need not be true in general [67].

As in the main text, we consider the bipartition (2, 3) for the five-qubit system. We find a significantly slower convergence than for Rényi entropy (see Fig. 4). This is due to the requirement for processing the entire density matrix ρ for quantum negativity as opposed to the classical shadow subsystem purity estimator, which is only evaluated on the smaller partition. The same convergence behavior is confirmed by numerical simulations discussed in Fig. 12 below. Classical shadows, on the other hand, allow for tighter classifications of entanglement [32,33]. The key idea is to probe the presence of negative eigenvalues in the partial transpose by comparing degree- d polynomials in the underlying density matrices. As d increases, these tests become tighter and eventually recover the negativity condition for entanglement [$\mathcal{N}(\rho) > 0$]. Classical shadows allow the estimation of all polynomials involved, but

the classical postprocessing cost becomes less and less favorable as the polynomial degree d increases [33].

We emphasize that, from the given set of tomography schemes, SIC-based classical shadow estimators deliver the best results in terms of both convergence as well as practicability. MLE reconstruction typically fails due to a lack of computational power and LI neglects physical constraints—a shortcoming that becomes very pronounced for nonlinear observables. Incorporating physical constraints by projection (PLS) remains computationally efficient, but leads to considerably poorer convergence behavior. Finally, classical shadows are the only approach for efficiently predicting nonlinear functions, such as mixed-state entanglement [32,33] of large-scale systems.

3. Complementary results on the rotated eight-qubit GHZ state

In the live-update discussion of Fig. 5 in the main text all two-qubit bipartitions were evaluated on top of fidelity and purity. Here we analyze the same data in postprocessing to present results on all bipartitions ((1, 7) qubits, (2, 6) qubits, (3, 5) qubits, and (4, 4) qubits). This evaluation for all possible pairs is not possible in real time on a standard desktop computer. We average until 50 000 shots (roughly 1200 s of data taking), where the SIC-based classical shadow purity estimator [Eq. (A20)] of the eight-qubit state, representing the most demanding property, has converged. Data are taken for a total of 12 500 s. Note that at around 40 000 shots almost no change is visible in the classical shadow purity as well as the respective standard deviation. The latter is due to systematic experimental drifts over the course of the long time measurement. Here we batch 100 shots for each approximate

ρ following Eq. (4), which speeds up the analysis without significant loss in accuracy; see Appendix A 4 for a thorough study of batch sizes considering both analysis time and accuracy. We also observe that the convergence of bipartitions from LI become significantly slower than classical shadows as the subsets get larger. For PLS, individual bipartitions even visually separate [see Fig. 8(c)], indicating very slow convergence behavior, confirming similar observations throughout the main text and Appendix.

Pauli tomography is neglected here as all bipartition Rényi entropies from SIC-based classical shadows converge faster than the time it takes to obtain just a single measurement per Pauli setting. Moreover, there is no efficient way of adding physical constraints to the reconstructed state ρ as MLE for an eight-qubit state is unfeasible on standard desktop computers, and PLS delivers significantly worse convergence. The findings here agree with those of the five-qubit AME state, previously discussed. SIC-based classical shadow estimators demonstrated to be most suitable for predicting nonlinear functions towards larger system sizes. Apart from the exemplified quadratic measures utilized in purity and Rényi entropy, classical shadow estimators support higher polynomial functions following the same principles [32].

4. Classical shadow convergence and practicability

In the live-update studies from Fig. 5(b) in the main text, we demonstrated real-time analysis of ongoing SIC tomography experiments, until all properties of interest were accurately estimated. How long the analysis can be performed in real time, however, depends on the size of the subsystems to be analyzed, as well as the type and number of functions to be estimated in parallel. Whereas time




FIG. 8. Detailed postprocessing analysis of the rotated eight-qubit GHZ state from Fig. 5 in the main text. Results now cover all bipartitions, which is not feasible in real time on a desktop computer. Shaded regions represent 1 standard deviation around the mean value when averaging over multiple sets of a total of 12 500 s of data. (a) We find fidelity from LI to converge before 20 000 shots, i.e., 500 s of data taking. (b) Classical shadow purity as converges well by 40 000 shots, or about 1000 s of data taking, which in the postprocessing analysis here is also confirmed by the saturation of the standard deviation towards more shots. The remaining fluctuations arise from systematic drifts in the experiment over the course of the measurements. (c) SIC-based classical shadows converge even quicker than the time it takes to measure just one shot per Pauli basis. On the other hand, LI and PLS are significantly slower, with PLS even showing a distinct separation between the individual bipartitions [(1, 7) qubits, (2, 6) qubits, (3, 5) qubits, (4, 4) qubits]. Dash-dot horizontal lines denote ideal values.




FIG. 9. Numerical simulations of batching SIC-based classical shadow purities. We plot the purity’s standard deviation from 100 repetitions of each point in the 2D grid using a logarithmic color scheme. A three-qubit linear cluster state is utilized in these numerical simulations, although the particular state and qubit number does not effect the qualitative picture. Regions of similar accuracy are found to have a certain almost vertical extension, indicating that bunching does not significantly degrade accuracy, while greatly reducing computation complexity. The ideal batch size must be evaluated on a case-by-case basis, weighing convergence against analysis time.

consumption for linear observables by means of SIC classical shadows remains constant over the course of data acquisition—individual experimental shots are simply processed and accumulated according to Eq. (A15)—nonlinear functions generally require higher-order products of all combinations from the given set of shots; see Eq. (A20). The resulting scaling is governed by the maximum polynomial order of the function in question minus one. For quadratic functions, in particular, the computation time for every new shot grows linearly with the number of already accumulated shots M . Hence, over the course of the data acquisition, this can eventually become computationally demanding, especially for large problem sizes where a large number of shots must be accumulated. Thus, to perform nonlinear function analysis in real time, one either keeps subsystems (and by that shot requirements) relatively small or uses a so-called batching approach, where multiple shots are bundled to estimate a more accurate ρ and thereby reduce the number of costly higher-order product combinations. In this section, we discuss the potential, as well as limits, of increasing the batch size in contrast to comparing single shots. To develop a quantitative statement, we perform




FIG. 10. Decision game. A quiz master constructs a quantum game where she provides us experimenters, having access to a quantum computer, with a circuit to prepare a state and perform measurements. After the measurement stage, we must report a certain property. Importantly, the property is only revealed after the experiment has been completed, preventing us from simply measuring the property in question. The aim is to win this challenge with as few experimental shots as possible. In the present case a single state from a fixed set is randomly chosen and implemented, which we then have to identify in as few shots as possible. SIC tomography delivers a reward in less than 20 shots, whereas the smallest Pauli tomography instance requires $3^4 = 81$ shots. Distance refers to the minimum fidelity difference between the target state and all others; see Fig. 11(b).

simulations on a three-qubit linear cluster state considering only statistical quantum projection noise. Along those lines, simulated tomography data are sampled from a multinomial distribution considering different sample sizes to mimic experimental shots. Each noisy set of tomography data is then reconstructed by means of SIC-based classical shadows, in particular, focusing on the quadratic purity estimator from Eq. (A20). Since we are only interested in changes in convergence behavior, the choice of state does not affect the qualitative statement of these numerical simulations. We compare 100 different shot numbers for 100 different batch sizes in a practical regime for three-qubit states. On the one extreme we compare all combinations of shadows obtained from a single shot each, which is known to be statistically optimal [see also the discussions around Eq. (A20)]. On the other extreme we compare just two shadows, each linearly more accurate, since they are obtained from averaging half the shots. In between we can trade-off the quality of the individual estimators versus the number of comparisons between estimators. By additionally accounting for analysis time, a sweet spot can be determined on a case-by-case basis.

We analyze the convergence behavior of these different strategies in Fig. 9 by means of the standard deviation of the classical shadow purity estimator, when repeating every point in the two-dimensional (2D) grid 100 times. We plot “shots” against “batch size” using a logarithmic color coding. Darker regions refer to a more accurate purity

estimate for the underlying state. A region at a specific color always shows a certain almost vertical extension, indicating that batching has a very small effect on accuracy. At the same time, however, it offers to speedup the data analysis significantly. We note that the observed pattern qualitatively remains the same for larger system sizes and accordingly more shots. Especially for larger subsystems, where a lot of statistics is required, batching has the potential to significantly speedup analysis. We made use of this method in Fig. 5 in the main text, where 100 experimental shots were used for each classical shadow. This turned out to be sufficient to estimate all relevant target properties in real time.

5. Fast state identification with classical shadows

We now consider a quantum game, where a quiz master targets us experimenters, having access to a quantum computer with a state to prepare and a question about a certain property. Importantly, the question is only revealed after performing the experiment. Such a setup is partly inspired by recent works on quantum-enhanced learning [45]. Here, we follow a game where the goal is to prepare a random target state from a fixed set of 16 states and then pinpoint this target state in as few experimental shots as possible. Figure 10 depicts the results, where SIC tomography allows us to receive a reward in less than 20 shots. In stark contrast, the minimum number of shots for the same task using Pauli tomography is $3^4 = 81$. The figure of merit for this game is the estimated distance between the states as the minimum difference in fidelity between the target state and all others. A reward is obtained as soon as that minimum distance remains positive.

Specifically, in the decision game of Fig. 10, we randomly prepare 1 out of 16 orthogonal four-qubit linear cluster states. These states correspond to all combinations of input states $|\pm\rangle$ on the four qubits. The target state is then identified by comparing the linear-inversion fidelities between the prepared state and all 16 possible targets. Note that fidelity is a good option for state identifications as linear observables generally converge quickest under LI, as we showed experimentally (see Figs. 4 and 5) and also formally derived in Appendix A 10. Figure 11(a) depicts fidelities with respect to all 16 states for SIC and Pauli tomography against the number of experimental shots. As expected, only one of the curves is close to fidelity 1, indicating the target state, whereas all others approach 0 within experimental uncertainties. We remark that, for the minimal Pauli implementation given at 81 shots, the state fidelities bunch at two values distinguishable in the figure due to the limited data. Going to the next higher settings at 162 and 243 shots, errors start to distribute indicated by more traces popping up. The distinguishing performance is even better visualized by plotting the difference between the target state

fidelity and all others in Fig. 11(b). The minimum of these distances (i.e., the worst case) is also shown in the background of Fig. 10 and used as our state distinguishability criterion. Surprisingly, less than 20 experimental shots are required to pinpoint the state, clearly undercutting the minimum for Pauli tomography. This argument can in principle be extended to larger systems, where we have seen that properties such as linear observables or Rényi entropies (see Fig. 5) converge much faster than the 3^N experimental shots required for the smallest instance of Pauli tomography. Fast state identification represents another fruitful example where SIC-based classical shadows not only appear more practical but also impart quicker convergence for certain tasks than Pauli tomography. The difference in performance between the tomography approaches originates from SIC's bigger experimental overhead (Appendix A 1) as well as the fact that the analyzed states are stabilizer states in the Pauli basis. Hence, they favor the analysis using Pauli measurements in terms of convergence; see Appendix. A 7.

6. Numerical simulations comparing tomography approaches

This section aims to support previous experimental findings through numerical simulations of convergence under statistical quantum shot noise. For this purpose, simulated tomography data are generated by sampling from a multinomial distribution, where the sample size is given by the number of shots. For the sake of comparability, we perform these numerical convergence simulations using the five-qubit AME state from Fig. 4 in the main text and Eq. (8). We study both SIC and Pauli tomography and cover all reconstruction methods as experimentally studied in Appendix A 2, i.e., LI, PLS, MLE, and, for quadratic functions, also SIC-based classical shadows. Numerical results are presented in Fig. 12 on a double-logarithmic scale, covering infidelity (to better illustrate convergence), and trace distance as another commonly used property for state distinguishability:

$$T(\rho, \sigma) = \frac{1}{2} \text{tr} \left[\sqrt{[(\rho - \sigma)^\dagger (\rho - \sigma)]} \right]. \quad (\text{A2})$$

LI approaches are again found to converge significantly faster than all other methods in terms of infidelity, while trace distance shows the opposite behavior due to being much more complicated to estimate. This is indicative of the unphysical nature of LI estimates. We further note that linear-inversion infidelities more often produce negative eigenvalues in the case of SIC tomography, since the five-qubit AME state is aligned with the Pauli basis and thus favors this basis for the estimation of linear observables; see Appendix A 7. In contrast, MLE approaches respect those physical boundaries, which result in valid values on all estimators at the cost of a slower




FIG. 11. Fast state identification from SIC tomography. Complementary analysis of the quantum game from Fig. 10. (a) Estimated fidelity with each of the 16 possible states. Via SIC tomography, we manage to identify the state with less than 20 shots, clearly undercutting the smallest Pauli implementation using $3^4 = 81$ shots. (b) For illustration purposes, we plot the fidelity difference between the generated state and all possible states as individual curves, which provides a distance measure. The minimum of these differences is used as a state distinguishability criterion to gain a reward; see the text for details.

convergence. Here we also confirm previous experimental observations (see Fig. 7) that SIC MLE performs better for small shot numbers than Pauli MLE. This is likely due to the fact that SIC POVMs provide the optimal information gain in each shot. Curiously, however, for larger numbers of shots, Pauli MLE eventually converges faster. This might be due to the overcompleteness of the

Pauli basis, but remains to be fully understood. While not specifically presented in this manuscript, we find the very same behavior for other states having different overlaps with the Pauli basis and SIC POVM; hence, this effect does not seem to be due to the choice of state. PLS again produces the slowest converges as we have already seen across our experimental studies.




FIG. 12. Numerical simulations incorporating statistical noise covering all tomography approaches. To this extent, quantum shot noise is incorporated via sampling ideal tomography data from a multinomial distribution, with sample sizes corresponding to the number of shots. We additionally present infidelity (1-fidelity) and complementary trace distance following Eq. (A2) and plot in double-logarithmic scale to more clearly visualize several orders of magnitudes. Note that infidelity estimation with LI sometimes delivers unphysical results, indicated by values out of range of the plot. Note that the smaller number of data points for SIC LI comes from ignoring unphysical results, which are more likely when the target state is misaligned with the measurement basis; see Appendix A 7 for details.




FIG. 13. Experimental convergence investigation of local states with different orientation, i.e., basis overlap. The given states are collectively prepared on eight qubits and partially traced to study multiple subsets. (a) The $|0\rangle$ state perfectly aligned with the first SIC vector. (b) The $|1\rangle$ state orthogonal to the first SIC vector, whose component completely vanishes. Convergence in (b) is significantly faster than in (a) as only three of the nonorthogonal SIC vectors take part. Boosted convergence is related to unambiguous state discrimination [62], limiting nonorthogonal measurements. (c) Superposition state with maximized overlap with one of the SIC vectors. (d) Superposition state with minimum overlap with the SIC vectors. Here, convergence of the state in (d) is better than (a) as the state's information is more regularly distributed over the SIC vectors, not favoring one, resulting in a higher information gain. Also, note that the achievable fidelity generally drops with higher qubit number due to experimental imperfections.

Overall, quantum negativity exhibits the slowest convergence, which confirms what we experimentally observed in Fig. 8. We note that the negativity calculation always requires the full density matrix independent of the bipartition. This is in stark contrast to Rényi-entropy-based measures and general nonlinear functions supported on a subset of the system, which can be estimated efficiently using SIC-based classical shadows. The latter again show the fastest convergence for purity according to Eq. (A20), which will generally be true for classical shadows by construction. To keep simulations efficient, the batch size for the purity estimator is chosen as a constant fraction (0.01) of the total number of shots, which has a negligible effect on convergence; see Appendix A 4. Finally, these numerical simulations could reproduce all features and findings from the experimental studies, thus indicating that there are no principal limitations to our experimental implementation of both SIC and Pauli tomography. Moreover, all reconstruction methods converge to the same values, indicating no principle drawback across the various methods.

7. Overlap with tomography basis in experiment and simulation

Our experimental studies covered by Figs. 4 and 5 focus on maximally entangled states that are differently aligned with respect to Pauli basis and SIC POVM in order to not particularly favor either of them. We resume this discussion in more detail by investigating states of varying orientation, with respect to the measurement

states, to study its effect on both experiments and numerical simulations. To this extent, convergence behavior of first purely local states followed by entangled states is demonstrated.

We start off by presenting experimental results on SIC tomography of local states up to eight qubits having different overlap with the Pauli basis and SIC POVM, depicted in Fig. 13. Analyzed metrics include fidelity and purity using LI and classical shadows as those scale more favorably in the number of qubits than MLE and PLS. While MLE becomes computationally too demanding for moderate systems sizes, PLS shows extremely slow convergence behavior, notably demanding more experimental shots. The given states are collectively prepared and subsequently partial traced to study multiple qubit numbers; see Appendix A 1. Brighter colors denote higher qubit numbers. Among the Pauli basis states $|0\rangle$ and $|1\rangle$, the latter performs significantly better under the SIC POVM. This reflects a general theme, where in using nonorthogonal bases it is preferable if one component vanishes, which is related to the concept of unambiguous state discrimination [62]. In the example of $|1\rangle$ this is indeed true for the first SIC vector aligned with $|0\rangle$ [see Fig. 15(a)]. For superposition states, we find that states that maximize the overlap with one SIC vector [Fig. 13(c)] perform worse than those that feature more even overlap with all SIC vectors [Fig. 13(d)]. Generally, higher qubit number states result in lower fidelities, due to experimental imperfections. We again emphasize the particularly fast convergence of the SIC-based classical shadow purity estimator, being here only moderately slower than fidelity.




FIG. 14. Convergence of SIC and Pauli tomography in numerical simulations for noisy states with different basis overlap. To incorporate noise, a depolarizing channel is applied to the ideal density matrices; see the text for details. Complementary, insets represent numerical results on projection noise only to more distinctively illustrate convergence behavior. (a)–(d) Local states that align with a SIC vector (a),(c) show slower convergence for SIC tomography than those orthogonal to a SIC vector (b),(d), confirming the experimental observations of Fig. 13. Pauli tomography, on the other hand, generally performs best for states aligned with the basis (a),(b). (e)–(h) The differences in convergence performance vanish for SIC tomography when considering entangled states. In the case of the Pauli basis, however, some improved convergence can still be observed for specific states, such as the GHZ state (g), which is aligned with the Pauli basis and thus performs better. Overall, the fast converging situations (more pronounced in the insets) for SIC (b),(d) and Pauli (a),(b),(g) tomography do beyond a certain shot number that depends on the noise, no longer carry unphysical components in the resulting infidelities.

To confirm these results for purely local states and to extend the discussion to entangled states of different orientation, we perform numerical simulations on four qubit states under quantum shot noise as previously explained within Appendix A 6. Figure 14 contains results for SIC tomography as well as now for Pauli tomography for predicting infidelity via LI, which scales more efficiently than MLE, that however would not change the essence

of the statements. On top, we numerically simulate noisy states by applying a depolarizing channel $(1 - p_{\text{depol}})\rho + p_{\text{depol}}\mathbb{1}/d$ with $p_{\text{depol}} = 0.02$ to the ideal density matrix ρ to perform the study under more realistic conditions. This enables us to observe the influence of unphysical properties, especially in regards of LI. Figures 14(a)–14(d) cover local states both along a SIC vector (a),(c) as well as orthogonal to it (b),(d), i.e., in the opposite direction of the




FIG. 15. Various representations of SIC POVM for maximizing information gain. (a) Standard representation of SIC POVM as used throughout this manuscript in both experiments and numerical simulations. (b) Rotated SIC optimized for predicting Pauli observables.

Bloch sphere. These simulations reproduce the effects seen experimentally in Fig. 13, where states aligned with a SIC vector (a),(c) converge more slowly than those orthogonal to it (b),(d). In contrast for Pauli tomography, representing an orthogonal basis, best results are obtained for states that are aligned with the basis, as seen in Figs. 14(a) and 14(b).

In the case of local states, measurements are uncorrelated and the above state-dependent convergence is to be expected. The situation might change when moving to entangled states. In Figs. 14(e) and 14(f) we apply controlled-phase gate states along the SIC vectors from Figs. 14(c) and 14(d) to generate an entangled state that is still (in a sense) aligned with the SIC basis. The resulting convergence is similar for both tomography methods. Curiously, SIC tomography performs equally well as for special states like a GHZ state [Fig. 14(g)] or a more generic rotated GHZ state [Fig. 14(h)]. In contrast, Pauli tomography, which also shows little dependence on the state once entanglement is involved, does outperform for the GHZ state [Fig. 14(g)], which is perfectly aligned with the Pauli basis. Importantly, the rotated GHZ state, which is somewhat randomly aligned with both the SIC POVM and Pauli basis [Fig. 14(h)] shows no difference between the tomography methods. The same rotated state, yet on eight qubits, is utilized for the real-time analysis in Fig. 5 in the main text to make sure that the comparison does not favor any approach. In the case of fast converging situations (see the insets of Fig. 14 without depolarizing noise) negative infidelities, influenced by unphysical properties of LI, disappear beyond a certain shot number depending on the amount of noise. This effect is illustrated by both tomography methods, particularly in Figs. 14(b) and 14(d) for SIC and in Figs. 14(a), 14(b), and 14(g) for Pauli tomography.

Inspired by these findings, it can be beneficial to rotate the SIC POVM used throughout this manuscript from Fig. 15(a) to favor the particular state or application that it is used for. In particular, the alignment depicted in Fig. 15(b), which has equal overlap with every Pauli basis vector, leads to improved prediction of Pauli observables.

Intuitively, this alignment ensures that each shot contains equal information about every Pauli observable. Geometrically, imagine a triangle spanned by the three Pauli basis vectors in the positive direction. Then, the orthogonal state to the first SIC vector ($\{-1/\sqrt{3}, -1/\sqrt{3}, -1/\sqrt{3}\}$) perpendicularly intersects this triangle area through its center. The front area of the SIC tetrahedron orientates parallel with the triangle, thus favoring Pauli eigenstates or stabilizer states, for which we find plenty of applications across the entire field of quantum computation and quantum information. In particular, VQE applications, which rely on the efficient estimation of many Pauli observables, could benefit from this choice of SIC POVM.

Note that experiments on rotated SICs can be straightforwardly realized by changing the local mapping sequence from Fig. 2(b), and do not add further complexity to the implementation.

8. SIC-POVM-based classical shadow framework

Renes *et al.* [2] discovered SIC POVMs, because of their exceptional tomographic capabilities. Ever since, both (tensor products of N) single-qubit SIC POVMs and global 2^N -dimensional SIC POVMs have served as idealized measurements for state reconstruction tasks. Single-qubit SIC POVMs, also known as tetrahedral POVMs, have also been used to acquire training data for neural network quantum state tomography [12,13].

Geometrically speaking, SIC POVMs [2] and the overcomplete Pauli basis [68] both form complex projective 2-designs (the single-qubit Pauli basis is actually a 3-design [35–37]). Roughly speaking, this means that the first two (three) moments exactly reproduce the moments of uniformly (Haar) random states. As detailed below, closed-form expressions for Haar-random moments can then be used to compute measurement operators and estimators analytically. For Pauli basis measurements, this observation culminated in efficient PLS estimators for full state tomography [29], as well as the classical shadow formalism for directly predicting (non)linear properties of the

underlying state [14,32]. Subsequently, some of these ideas have been extended to (single-qubit) SIC POVM measurements. García-Pérez *et al.* [11], in particular, highlighted that SIC POVM measurements can outperform Pauli basis measurements in VQE-type energy measurements and, more generally, for predicting linear state properties.

Here, we build on all these ideas and provide a self-contained derivation of classical shadows from (single-qubit) SIC POVM measurements, as well as rigorous sample complexity bounds for general linear and quadratic property estimation. The actual definition of SIC POVM shadows is virtually identical to existing (Pauli basis) classical shadows, because both form complex projective 2-designs. Sample complexity bounds, on the other hand, require novel proof techniques. They require computing variances that correspond to third-order polynomials in the measurement ensemble. This is comparatively easy for Pauli basis measurements, which form a 3-design. But for SIC POVMs—which only form a 2-design—these existing techniques do not apply. We overcome this drawback with new proof methods that directly use the symmetries within a SIC POVM rather than an abstract 3-design property. To our knowledge, these theoretical arguments are novel and may be of independent interest.

9. Classical shadows from SIC POVM measurements

a. The single-qubit case

Single-qubit density matrices live in the (real-valued) vector space of Hermitian 2×2 matrices that we denote by \mathbb{H}_2 . On this space, single-qubit SIC POVMs are known to form so-called projective 2-designs [2]. Mathematically, this means that discrete averages over (outer products of) SIC vectors reproduce uniform averages over all possible pure states up to second moments [34]. This is captured by the two averaging formulas

$$\frac{1}{4} \sum_{i=1}^4 |\psi_i\rangle\langle\psi_i| = \int |v\rangle\langle v| dv = \frac{1}{2} \mathbb{I} \in \mathbb{H}_2, \quad (\text{A3})$$

$$\frac{1}{4} \sum_{i=1}^4 (|\psi_i\rangle\langle\psi_i|)^{\otimes 2} = \int (|v\rangle\langle v|)^{\otimes 2} dv = \frac{1}{6} (\mathbb{I} \otimes \mathbb{I} + \mathbb{F}). \quad (\text{A4})$$

Here, dv denotes the unique pure state measure (normalized to $\int dv = 1$) that assigns the same infinitesimal weight to each state ($dv = dw$). The two-qubit Slashed Down operator $\mathbb{F}|v\rangle \otimes |w\rangle = |w\rangle \otimes |v\rangle$ acts by permuting tensor factors.

The first averaging formula (A3) confirms that the collection $\{\frac{1}{2}|\psi_i\rangle\langle\psi_i| : i = 1, 2, 3, 4\} \subset \mathbb{H}_2$ forms a valid quantum measurement for single-qubit systems (POVM). Let

$\rho \in \mathbb{H}_2$ be a density matrix, i.e., a Hermitian matrix with unit trace whose eigenvalues are non-negative. Then,

$$\Pr[i|\rho] = \text{tr}\left(\frac{1}{2}|\psi_i\rangle\langle\psi_i|\rho\right) = \frac{1}{2}\langle\psi_i|\rho|\psi_i\rangle \quad \text{for } i = 1, 2, 3, 4$$

is the probability of obtaining POVM outcome i upon measuring state ρ , obeying $\Pr[i|\rho] \geq 0$, because density matrices do not have negative eigenvalues ($\langle\psi_i|\rho|\psi_i\rangle \geq 0$). Moreover, Eq. (A3) ensures proper normalization:

$$\begin{aligned} \sum_{i=1}^4 \Pr[i|\rho] &= \text{tr}\left(2\left(\frac{1}{4} \sum_{i=1}^4 |\psi_i\rangle\langle\psi_i|\right)\rho\right) \\ &= \text{tr}(\mathbb{I}\rho) = \text{tr}(\rho) = 1. \end{aligned}$$

The second averaging property (A4) is more interesting. It ensures that SIC POVM measurements are informationally complete, i.e., we can reconstruct every density matrix ρ based on outcome probabilities. This may not directly be obvious from the display itself, but it follows from the following chain of arguments. If we weigh SIC projectors $|\psi_i\rangle\langle\psi_i|$ with the probability $\Pr[i|\rho]$ of observing this outcome, Eq. (A4) allows us to compute the resulting average. Let $\text{tr}_1(\cdot)$ denote the partial trace over the first of two qubits [$\text{tr}_1(A \otimes B) = \text{tr}(A)B$ and linearly extended to all of $\mathbb{H}_2^{\otimes 2}$]. Then,

$$\begin{aligned} \sum_{i=1}^4 \Pr[i|\rho] |\psi_i\rangle\langle\psi_i| &= \sum_{i=1}^4 \frac{1}{2} \langle\psi_i|\rho|\psi_i\rangle |\psi_i\rangle\langle\psi_i| \\ &= 2\text{tr}_1\left(\left(\frac{1}{4} \sum_{i=1}^4 (|\psi_i\rangle\langle\psi_i|)^{\otimes 2}\right)\rho \otimes \mathbb{I}\right) \\ &= \frac{1}{3} \text{tr}_1((\mathbb{I} \otimes \mathbb{I} + \mathbb{F})\mathbb{I} \otimes \rho) \\ &= \frac{1}{3} (\text{tr}(\rho)\mathbb{I} + \rho), \end{aligned} \quad (\text{A5})$$

where the last equation follows from the interplay between the partial trace and SWAP operator. The final expression is equivalent to applying a depolarizing channel with parameter $p = 1/3$ to the quantum state in question:

$$\mathcal{D}_{1/3}(\rho) = \frac{1}{3}\rho + \left(1 - \frac{1}{3}\right) \frac{\text{tr}(\rho)}{2} \mathbb{I} \in \mathbb{H}_2.$$

Viewed as a linear map on \mathbb{H}_2 , this channel has a uniquely defined inverse:

$$\mathcal{D}_{1/3}^{-1}(A) = 3A - \text{tr}(A)\mathbb{I} \quad \text{for all } A \in \mathbb{H}_2. \quad (\text{A6})$$

Although a linear map, this is not a physical operation. We can, however, use it in the classical postprocessing stage to

counterbalance the effect of averaging over SIC elements. Indeed, linearity and Eq. (A5) ensure that

$$\begin{aligned} \sum_{i=1}^4 \Pr[i|\rho](3|\psi_i\rangle\langle\psi_i| - \mathbb{I}) &= \sum_{i=1}^4 \Pr[i|\rho] \mathcal{D}_{1/3}^{-1}(|\psi_i\rangle\langle\psi_i|) \\ &= \mathcal{D}_{1/3}^{-1} \left(\sum_{i=1}^4 \Pr[i|\rho] |\psi_i\rangle\langle\psi_i| \right) \\ &= \mathcal{D}_{1/3}^{-1}(\mathcal{D}_{1/3}(\rho)) \\ &= \rho. \end{aligned} \quad (\text{A7})$$

The left-hand side of this display features a linear combination involving SIC outcome probabilities $\Pr[i|\rho]$, while the right-hand side exactly reproduces the underlying state ρ . This equips us with a concrete state reconstruction formula—the so-called linear inversion estimator. But, at least at first sight, this formula is only useful if we have precise knowledge of the SIC outcome probabilities $\Pr[i|\rho]$. And, with current quantum technology, these probabilities must be estimated from repeatedly performing SIC POVM measurements on independent copies of ρ and approximating these probabilities by frequencies.

The classical shadow formalism provides an alternative perspective on this estimation process. Suppose that we perform a single SIC POVM measurement of an unknown quantum state ρ (single shot). Then, we obtain a random measurement outcome $\hat{i} \in \{1, 2, 3, 4\}$ with probability $\Pr[\hat{i}|\rho]$ each. Inspired by the left-hand side of Eq. (A7), we can use this outcome \hat{i} to construct a Monte Carlo estimator of ρ :

$$\hat{i} \mapsto \hat{\sigma} = (3|\psi_{\hat{i}}\rangle\langle\psi_{\hat{i}}| - \mathbb{I}) = \mathcal{D}_{1/3}^{-1}(|\psi_{\hat{i}}\rangle\langle\psi_{\hat{i}}|) \in \mathbb{H}_2.$$

This is a *random* 2×2 matrix that can assume four different forms—one for each possible outcome $\hat{i} \in \{1, 2, 3, 4\}$. It does exactly reproduce the underlying quantum state ρ in expectation over the observed single-shot outcome:

$$\mathbb{E}[\hat{\sigma}] = \sum_{i=1}^4 \Pr[i|\rho](3|\psi_i\rangle\langle\psi_i| - \mathbb{I}) = \rho. \quad (\text{A8})$$

Here, we have used the definition of the expectation value (sum of outcomes weighed by their probability), as well as Eq. (A7). It is worthwhile to emphasize that each $\hat{\sigma}$ has the same eigenvalue structure: $\lambda_+ = 2$ and $\lambda_- = -1$. In turn, these random matrices all have unit trace [$\text{tr}(\hat{\sigma}) = 2 - 1 = 1$], but are unphysical in the sense that one eigenvalue is always negative. Equation (A8) represents a physical density matrix ρ as the expectation of four unphysical estimators. This desired expectation value can be approximated by empirically averaging M independently generated Monte Carlo estimators. Suppose that $\hat{\sigma}_1, \dots, \hat{\sigma}_M$ are

M independently and identically distributed (IID) Monte Carlo estimators. Then, their empirical average obeys

$$\hat{\rho} = \frac{1}{M} \sum_{m=1}^M \hat{\sigma}_m \xrightarrow{M \rightarrow \infty} \mathbb{E}[\hat{\sigma}] = \rho,$$

and the rate of convergence can be controlled with arguments from probability theory. This will be the content of the next two subsections.

b. Extension to multiqubit systems

The formalism and ideas presented above readily extend to quantum systems composed of multiple qubits. Let $\rho \in \mathbb{H}_2^{\otimes N} \simeq \mathbb{H}^{2^N}$ be an N -qubit density matrix. We can perform a single-qubit SIC POVM measurement on each of the N qubits. As in the single-qubit case, each such measurement yields one out of four possible outcomes. In total, a single-shot measurement produces a string (i_1, \dots, i_N) of outcomes. There are 4^N such outcomes and the probability of obtaining any one of them is given by

$$\begin{aligned} \Pr[i_1, \dots, i_N | \rho] &= \text{tr} \left(\bigotimes_{n=1}^N \left(\frac{1}{2} |\psi_{i_n}\rangle\langle\psi_{i_n}| \right) \rho \right) \\ &= \frac{1}{2^N} \langle \psi_{i_1}, \dots, \psi_{i_N} | \rho | \psi_{i_1}, \dots, \psi_{i_N} \rangle \\ &\quad \text{for each } i_1, \dots, i_N \in \{1, 2, 3, 4\}. \end{aligned} \quad (\text{A9})$$

Here, we have introduced the shorthand notation $|\psi_{i_1}, \dots, \psi_{i_N}\rangle = \bigotimes_{n=1}^N |\psi_{i_n}\rangle \in (\mathbb{C}^2)^{\otimes N} \simeq \mathbb{C}^{2^N}$. These expressions are non-negative, because the N -qubit density matrix does not have negative eigenvalues. Equation (A3), applied to each qubit separately, moreover ensures proper normalization:

$$\begin{aligned} \sum_{i_1, \dots, i_N=1}^4 \Pr[i_1, \dots, i_N | \rho] &= \text{tr} \left(\bigotimes_{n=1}^N \left(\sum_{i_n=1}^4 \frac{1}{2} |\psi_{i_n}\rangle\langle\psi_{i_n}| \right) \rho \right) \\ &= \text{tr}(\mathbb{I}^{\otimes N} \rho) = \text{tr}(\rho) = 1. \end{aligned}$$

Again, the second averaging property is more interesting: for single-qubit density matrices $\hat{\rho} \in \mathbb{H}_2$, we already know that Eq. (A4) implies that $\sum_{i=1}^4 \frac{1}{2} \langle \psi_i | \rho | \psi_i \rangle = \mathcal{D}_{1/3}(\rho)$, where $\mathcal{D}_{1/3} : \mathbb{H}_2 \rightarrow \mathbb{H}_2$ is the depolarizing channel from Eq. (A6). It is now easy to check that this equation extends to general Hermitian 2×2 matrices:

$$\sum_{i=1}^4 \frac{1}{4} \langle \psi_i | A | \psi_i \rangle |\psi_i\rangle\langle\psi_i| = \mathcal{D}_{1/3}(A) \quad \text{for all } A \in \mathbb{H}_2. \quad (\text{A10})$$

We can use this observation to show that N single-qubit SIC POVM measurements are tomographically complete.

To achieve this, it is helpful to first decompose $\rho \in \mathbb{H}_2^{\otimes N}$ into a sum of elementary tensor products:

$$\rho = \sum_{W_1, \dots, W_N} r(W_1, \dots, W_N) \bigotimes_{n=1}^N W_n, \quad \text{where}$$

$$r(W_1, \dots, W_N) = \text{tr} \left(\bigotimes_{n=1}^N W_n \rho \right) \in [-1, 1]$$

and the summation goes over all four single-qubit Pauli matrices $W_n = \mathbb{I}, X, Y, Z$. Combine this with Eq. (A10) to compute

$$\begin{aligned} & \sum_{i_1, \dots, i_N=1}^4 \text{Pr}[i_1, \dots, i_N | \rho] \bigotimes_{n=1}^N |\psi_{i_n}\rangle\langle\psi_{i_n}| \\ &= \sum_{W_1, \dots, W_N} r(W_1, \dots, W_N) \bigotimes_{n=1}^N \\ & \quad \times \left(\frac{1}{2} \sum_{i_n=1}^4 \langle\psi_{i_n} | W_n | \psi_{i_n}\rangle |\psi_{i_n}\rangle\langle\psi_{i_n}| \right) \\ &= \sum_{W_1, \dots, W_N} r(W_1, \dots, W_N) \bigotimes_{n=1}^N \mathcal{D}_{1/3}(W_n) \\ &= \mathcal{D}_{1/3}^{\otimes N}(\rho), \end{aligned} \quad (\text{A11})$$

where the last equality follows from linearity of depolarizing channels. The final expression is equivalent to applying N -independent, single-qubit depolarizing channels to the N -qubit quantum state ρ . Viewed as a linear map on $\mathbb{H}_2^{\otimes N}$, this tensor product channel has a uniquely defined inverse $\mathcal{D}_{1/3}^{-\otimes N} : \mathbb{H}_2^{\otimes N} \rightarrow \mathbb{H}_2^{\otimes N}$. For elementary tensor products, this tensor product of inverse depolarizing channels factorizes nicely into tensor products. In particular,

$$\begin{aligned} \mathcal{D}_{1/3}^{-\otimes N} \left(\bigotimes_{n=1}^N |\psi_{i_n}\rangle\langle\psi_{i_n}| \right) &= \bigotimes_{n=1}^N \mathcal{D}_{1/3}^{-1}(|\psi_{i_n}\rangle\langle\psi_{i_n}|) \\ &= \bigotimes_{n=1}^N (3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I}). \end{aligned}$$

Again, this is not a physical operation, because it produces matrices with negative eigenvalues. However, we can nonetheless use it in the classical postprocessing stage to counterbalance the N -qubit averaging effect encountered

in Eq. (A11):

$$\begin{aligned} & \sum_{i_1, \dots, i_N=1}^4 \text{Pr}[i_1, \dots, i_N | \rho] \bigotimes_{n=1}^N (3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I}) \\ &= \mathcal{D}_{1/3}^{-\otimes N} \left(\sum_{i_1, \dots, i_N=1}^4 \text{Pr}[i_1, \dots, i_N | \rho] \bigotimes_{n=1}^N |\psi_{i_n}\rangle\langle\psi_{i_n}| \right) \\ &= \mathcal{D}_{1/3}^{-\otimes N} (\mathcal{D}_{1/3}^{\otimes N}(\rho)) \\ &= \rho. \end{aligned} \quad (\text{A12})$$

The left-hand side of this display features a linear combination involving single-qubit SIC outcome probabilities $\text{Pr}[i_1, \dots, i_N | \rho]$, while the right-hand side exactly reproduces the underlying N -qubit density matrix. This provides us with a concrete reconstruction formula for arbitrary N -qubit states. In fact, it is a natural and relatively straightforward extension of the single-qubit linear inversion estimator (A7) to N qubits.

As was the case for single qubits, the classical shadow formalism provides an alternative perspective on such a linear-inversion estimation process. Suppose that we perform N single-qubit SIC POVM measurements of an unknown N -qubit state ρ (single shot). Then, we obtain a random outcome string $(\hat{i}_1, \dots, \hat{i}_N) \in \{1, 2, 3, 4\}^{\times N}$ with probability $\text{Pr}[\hat{i}_1, \dots, \hat{i}_N | \rho]$ each. Inspired by the left-hand side of Eq. (A12), we can use this random outcome string to construct a Monte Carlo estimator of ρ :

$$(\hat{i}_1, \dots, \hat{i}_N) \mapsto \hat{\sigma} = \bigotimes_{n=1}^N (3|\psi_{\hat{i}_n}\rangle\langle\psi_{\hat{i}_n}| - \mathbb{I}) \in \mathbb{H}_2^{\otimes N}. \quad (\text{A13})$$

This is a random $2^N \times 2^N$ matrix that decomposes nicely into tensor products of single-qubit contributions. Each tensor factor contributes a matrix with eigenvalues $\lambda_{n,+} = +2$ and $\lambda_{n,-} = -1$. The eigenvalues of the tensor product $\hat{\sigma}$ then correspond to N -fold products of these two possible numbers. The largest eigenvalue is $+2^N$, the smallest is -2^{N-1} , so $\hat{\sigma}$ is a very unphysical random matrix. The randomness stems from an actual quantum measurement and depends on the underlying quantum state. This ensures that $\hat{\sigma}$ reproduces ρ in expectation, i.e.,

$$\mathbb{E}[\hat{\sigma}] = \sum_{i_1, \dots, i_N=1}^4 \text{Pr}[i_1, \dots, i_N | \rho] \bigotimes_{n=1}^N (3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I}) = \rho, \quad (\text{A14})$$

courtesy of Eq. (A12). This expectation value can now be approximated by empirically averaging M independently generated Monte Carlo estimators, so-called *classical shadows*. Let $\hat{\sigma}_1, \dots, \hat{\sigma}_M \in \mathbb{H}_2^{\otimes N}$ be estimators generated from repeatedly preparing ρ and performing single-qubit

SIC POVM measurements (IID). Then, their empirical average obeys

$$\hat{\rho} := \frac{1}{M} \sum_{m=1}^M \hat{\sigma}_m \xrightarrow{M \rightarrow \infty} \rho, \quad (\text{A15})$$

in full analogy to the single-qubit case. As detailed in the next section, the rate of convergence will depend on the number of qubits N . The larger the space, the longer it takes for convergence to kick in, and this scaling can become unfavorable. It is therefore worthwhile to emphasize another distinct advantage of the tensor product structure of estimators (A13): marginalization to subsystem density operators is straightforward. Let ρ be an N -qubit state, but suppose that we are only interested in a subsystem $K \subset [N] = \{1, \dots, N\}$ composed of only $|K| \leq N$ qubits. Such a subsystem is fully described by the reduced $|K|$ -qubit density matrix $\rho_K = \text{tr}_{-K}(\rho) \in \mathbb{H}_2^{\otimes |K|}$ that results from tracing out all qubits not in K . This partial trace is a linear operation that plays nicely with the tensor product structure in Eq. (A13). Each tensor product factor has unit trace, which ensures that

$$\hat{\sigma}_K = \text{tr}_{-K}(\hat{\sigma}) = \bigotimes_{k \in K} (3|\psi_{i_k}\rangle\langle\psi_{i_k}| - \mathbb{I}) \in \mathbb{H}_2^{\otimes |K|} \quad \text{obeys}$$

$$\mathbb{E}[\hat{\sigma}_K] = \text{tr}_{-K}(\mathbb{E}[\hat{\sigma}]) = \text{tr}_{-K}(\rho) = \rho_K.$$

The object at the very left is a random $2^{|K|} \times 2^{|K|}$ matrix that can be generated from performing a complete N -qubit SIC POVM measurement to obtain outcomes $(\hat{i}_1, \dots, \hat{i}_N) \in \{1, \dots, 4\}^{\times N}$. Subsequently, we only use outcomes that correspond to qubits in K to directly construct an estimator for the subsystem density matrix ρ_K in question. This trick reduces the question of convergence to a problem that only involves $|K|$ qubits, not N . This can be highly advantageous if $|K| \ll N$. What is more, we can use the same N -qubit measurement outcome $(\hat{i}_1, \dots, \hat{i}_N)^{\times N}$ to construct estimators for multiple subsystems $K_1, \dots, K_L \subset [N]$ at once. This allows us to use the same N -qubit measurement statistics to estimate many subsystem properties in parallel.

10. Convergence for predicting linear observables

Suppose that we have access to M -independent Monte Carlo approximations $\hat{\sigma}_1, \dots, \hat{\sigma}_M$ of an unknown N -qubit state ρ . Each of them arises from measuring N single-qubit SIC POVMs on an independent copy of ρ . We can then use these approximations to estimate observable expectation values $\text{tr}(O\rho)$ with $O \in \mathbb{H}_2^{\otimes N}$:

$$\hat{\rho} = \text{tr}(O\hat{\rho}) = \frac{1}{M} \sum_{m=1}^M \text{tr}(O\hat{\sigma}_m).$$

This is an empirical average of M IID random numbers $X_1, \dots, X_M \stackrel{\text{IID}}{\sim} X = \text{tr}(O\hat{\sigma})$ that converges to the true expectation $\mathbb{E}[X] = \text{tr}(O\mathbb{E}[\hat{\sigma}_m]) = \text{tr}(O\rho)$ as M increases. The rate of convergence is controlled by the variance $\text{Var}[X]$.

Lemma 1. Fix an N -qubit observable O and let $\hat{\sigma} \in \mathbb{H}_2^{\otimes N}$ be a (SIC POVM) classical shadow as defined in Eq. (A13). Then,

$$\text{Var}[\text{tr}(O\hat{\sigma})] \leq 3^N \text{tr}(O^2) \quad \text{for any underlying } N\text{-qubit state } \rho.$$

This bound is reminiscent of existing variance bounds for randomized single-qubit Pauli measurements [14], but slightly weaker [random Pauli basis measurements achieve $\text{Var}[\text{tr}(O\hat{\sigma})] \leq 2^N \text{tr}(O^2)$]. The proof exploits the 2-design property of SIC POVM measurements and will be supplied in Appendix A14 below. For now, we use this variance bound to conclude strong convergence guarantees for observable estimation with SIC POVM measurements. The key ingredient is a strong tail bound for sums of IID random variables with known variance and bounded magnitude. The *Bernstein inequality* is a stronger version of the better known Hoeffding inequality; see, e.g., Ref. [69, Corollary 7.31] or [70, Theorem 2.8.4].

Let $X_1, \dots, X_M \stackrel{\text{IID}}{\sim} X \in \mathbb{R}$ be IID random variables with expectation $\mu = \mathbb{E}[X]$ and variance $\sigma^2 = \mathbb{E}[(X - \mu)^2]$ that also obey $|X_m| \leq R$ (with probability 1). Then, for $\epsilon > 0$,

$$\begin{aligned} \Pr \left[\left| \frac{1}{M} \sum_{m=1}^M X_m - \mu \right| \geq \epsilon \right] \\ \leq 2 \exp \left(- \frac{M\epsilon^2/2}{\sigma^2 + R\epsilon/3} \right) \\ \leq \begin{cases} 2 \exp \left(- \frac{3}{8} M\epsilon^2/\sigma^2 \right) & \text{if } \epsilon \leq \sigma^2/R, \\ 2 \exp \left(- \frac{3}{8} M\epsilon/R \right) & \text{if } \epsilon \geq \sigma^2/R. \end{cases} \quad (\text{A16}) \end{aligned}$$

For the task at hand, we write $\hat{\sigma} = (1/M) \sum_{m=1}^M \text{tr}(O\hat{\sigma}_m) = (1/M) \sum_{m=1}^M X_m$, where $X_m = \text{tr}(O\hat{\sigma}_m) \stackrel{\text{IID}}{\sim} X = \text{tr}(O\hat{\sigma})$. For technical reasons, we also assume that $\text{tr}(O^2) \geq (5/9)^N > 2^{-N}$. Note that this is achieved by (i) physical observables [$\text{tr}(O^2) \geq \|O\|_\infty^2 = 1$], as well as quantum states [$O = \rho$ obeys $\text{tr}(\rho^2) \geq \text{tr}[(\mathbb{I}/2^N)^2] = 2^{-N}$]. The random variable X obeys

$$\begin{aligned} \mu = \text{tr}(O\mathbb{E}[\hat{\sigma}]) = \text{tr}(O\rho), \quad \text{Var}[\text{tr}(O\hat{\sigma})] \leq 3^N \text{tr}(O^2) \\ =: \sigma^2, \text{ and } |\text{tr}(O\hat{\sigma})| \leq 5^{N/2} \sqrt{\text{tr}(O^2)} \leq 3^N \text{tr}(O^2) =: R. \end{aligned}$$

The first equality is Eq. (A14), the second bound is Lemma 1, and the last bound is a consequence of the

Cauchy-Schwarz inequality: $|\text{tr}(O\hat{\sigma})| \leq \sqrt{\text{tr}(O^2)}\sqrt{\text{tr}(\hat{\sigma}^2)}$. Since classical shadows are tensor products of single-qubit blocks with eigenvalues $\lambda_+ = 2, \lambda_- = -1$, we can readily conclude that $\text{tr}(\hat{\sigma}^2) = (\lambda_+^2 + \lambda_-^2)^N = 5^N$. The final bound is rather loose and follows from the assumption that $\text{tr}(O^2) \geq (5/9)^N$. Inserting these bounds into Eq. (A16) now ensures that

$$\Pr[|\hat{\sigma} - \text{tr}(O\rho)| \geq \epsilon] = \Pr[|\hat{\sigma} - \text{tr}(O\rho)| \geq \epsilon] \leq 2 \exp \times \left(-\frac{3M\epsilon^2}{8 \times 3^N \text{tr}(O^2)} \right) \quad \text{for all } 0 < \epsilon \leq 1. \quad (\text{A17})$$

This is a bound on the probability of an ϵ -deviation (or more) that diminishes exponentially in the number of Monte Carlo samples (measurements) M . For a fixed confidence $\delta \in (0, 1)$, setting

$$M \geq \frac{8}{3} 3^N \text{tr}(O^2) \log(1/\delta) / \epsilon^2$$

ensures that

$$|\hat{\sigma} - \text{tr}(O\rho)| \leq \epsilon \quad \text{with probability (at least) } 1 - \delta.$$

Note that the required measurement budget M scales exponentially in the number N of involved qubits. At this point it is helpful to remember the marginalization property of classical shadows. Suppose that O is localized in the sense that it only affects a subsystem $\mathbf{K} \subset [N]$ composed of $|\mathbf{K}| \leq N$ qubits. Then, $\text{tr}(O\rho) = \text{tr}(O_{\mathbf{K}}\rho_{\mathbf{K}})$, where $\rho_{\mathbf{K}} = \text{tr}_{-\mathbf{K}}(\rho)$ is the reduced $|\mathbf{K}|$ -qubit density matrix and $O_{\mathbf{K}}$ is the nontrivial part of O . In turn, we can use appropriately marginalized classical shadows $\hat{\sigma}_{m,\mathbf{K}} = \text{tr}_{-\mathbf{K}}(\hat{\sigma}_m)$ to directly approximate this subsystem property:

$$\hat{\sigma} = \frac{1}{M} \sum_{m=1}^M \text{tr}(O\hat{\sigma}_m) = \frac{1}{M} \sum_{m=1}^M \text{tr}(O_{\mathbf{K}}\hat{\sigma}_{m,\mathbf{K}}).$$

The reformulation on the right-hand side now involves only the $|\mathbf{K}| < N$ relevant qubits. We can now redo the argument from above to obtain a measurement budget that only scales exponentially in $|\mathbf{K}|$:

$$\Pr[|\hat{\sigma} - \text{tr}(O\rho)| \geq \epsilon] \leq 2 \exp \left(-\frac{3M\epsilon^2}{8 \times 3^{|\mathbf{K}|} \text{tr}(O_{\mathbf{K}}^2)} \right) \quad \text{for } \epsilon \in (0, 1). \quad (\text{A18})$$

This refinement asserts that the probability of an ϵ -deviation for a single observable estimation diminishes exponentially in the number of measurements. We can use this exponential concentration to bound the probability of a single deviation among many. This allows us to use the same measurement data to predict many observables $\text{tr}(O_1\rho), \dots, \text{tr}(O_L\rho)$ in parallel.

Theorem 1. *Let O_1, \dots, O_L be N -qubit observables that are all localized to (at most) K qubits and fix $\epsilon, \delta \in (0, 1)$. Then,*

$$M \geq \frac{8}{3} 3^K \max_{1 \leq l \leq L} \text{tr}(O_{l,\mathbf{K}_l}^2) \log(2L/\delta) / \epsilon^2 \quad (\text{A19})$$

N -qubit SIC POVM measurements of an unknown state ρ are very likely to ϵ -approximate all observables simultaneously. More precisely, the resulting classical shadows $\hat{\sigma}_1, \dots, \hat{\sigma}_M$ obey

$$\max_{1 \leq l \leq L} \left| \frac{1}{M} \sum_{m=1}^M \text{tr}(O_l \hat{\sigma}_m) - \text{tr}(O_l \rho) \right| \leq \epsilon$$

with probability (at least) $1 - \delta$.

The convergence bound given in Eq. (5) of the main text is a simplified consequence of this result. Note that, by and large, physical observables are normalized in the operator norm: $\|O_l\|_\infty = \|O_{l,\mathbf{K}}\|_\infty = 1$. Equation (A19) features squared Hilbert-Schmidt norms $\|O_{l,\mathbf{K}}\|_2^2 = \text{tr}(O_{l,\mathbf{K}}^2)$ on the $|\mathbf{K}|$ -qubit subsystems in question. These Hilbert-Schmidt norms can be related to the operator norm, which is bounded:

$$\text{tr}(O_{l,\mathbf{K}}^2) = \|O_{l,\mathbf{K}}\|_2^2 \leq 2^{|\mathbf{K}|} \|O_{l,\mathbf{K}}\|_\infty^2 = 2^{|\mathbf{K}|} \quad \text{for all } 1 \leq l \leq L.$$

Here, we have used the fact that each $O_{l,\mathbf{K}}$ is a matrix of size (at most) $2^{|\mathbf{K}|} \cdot 2^{|\mathbf{K}|}$. This implies the bound $\max_{1 \leq l \leq L} \text{tr}(O_{l,\mathbf{K}}^2) \leq 2^{|\mathbf{K}|}$, which can be very pessimistic. Inserting it into Eq. (A19) yields Eq. (5) in the main text.

Proof of Theorem 1. A maximum deviation larger than ϵ occurs if at least one individual prediction $\hat{\sigma}_l$ is further than ϵ off from the actual target $\text{tr}(O_l\rho)$. The union bound, also known as Boole's inequality, tells us that the probability of such a maximum deviation is upper bounded by the sum of individual deviation probabilities. These, in turn, can be controlled via the tail bound from Eq. (A18):

$$\begin{aligned} \Pr \left[\max_{1 \leq l \leq L} \left| \frac{1}{M} \sum_{m=1}^M \text{tr}(O_l \hat{\sigma}_m) - \text{tr}(O_l \rho) \right| \geq \epsilon \right] \\ \leq \sum_{l=1}^L \Pr \left[\left| \frac{1}{M} \sum_{m=1}^M \text{tr}(O_l \hat{\sigma}_m) - \text{tr}(O_l \rho) \right| \geq \epsilon \right] \\ \leq \sum_{l=1}^L 2 \exp \left(-\frac{3M\epsilon^2}{8 \times 3^{|\mathbf{K}_l|} \text{tr}(O_{l,\mathbf{K}_l}^2)} \right). \end{aligned}$$

We see that each of these summands diminishes exponentially in the measurement budget M . The right-hand side of Eq. (A19) ensures that each term contributes at most δ/L to this sum. Since there are L summands in total, we conclude that $\Pr[\max_{1 \leq l \leq L} |(1/M) \sum_{m=1}^M \text{tr}(O_l \hat{\sigma}_m) - \text{tr}(O_l \rho)| \geq \epsilon] \leq \delta$.

$\text{tr}(O\rho) \geq \epsilon] \leq \delta$. This is equivalent to the advertised display. ■

11. Convergence for predicting (subsystem) purities

Classical shadows can also be used to predict nonlinear quantum state properties; see, e.g., Refs. [14,32]. A prototypical example is the purity of an N -qubit density matrix ρ :

$$p(\rho) = \text{tr}(\rho^2) = \text{tr}(\rho\rho) \in (0, 1].$$

The purity equals one if and only if ρ describes a pure quantum state $|\phi\rangle\langle\phi|$. Conversely, it achieves its minimum value for the maximally mixed state: $\rho = (\frac{1}{2}\mathbb{I})^{\otimes N}$ achieves $p(\rho) = 1/2^N \ll 1$. We now describe how to obtain a purity estimator based on classical shadows $\hat{\sigma}_1, \dots, \hat{\sigma}_M$ that arise from measuring N single-qubit SIC POVMs on (independent copies of) ρ . By construction, each $\hat{\sigma}_m$ is a Monte Carlo estimator of ρ . Indeed, Eq. (A14) asserts that $\mathbb{E}[\hat{\sigma}_m] = \rho$ for all $1 \leq m \leq M$. What is more, distinct Monte Carlo estimators $\hat{\sigma}_m$ and $\hat{\sigma}_{m'}$ with $m \neq m'$ are statistically independent. The expectation over statistically independent random matrices factorizes. This ensures that the trace of the product of any two distinct classical shadows reproduces the purity in expectation:

$$\begin{aligned} \text{tr}(\hat{\sigma}_m \hat{\sigma}_{m'}) &\text{ obeys } \mathbb{E}[\text{tr}(\hat{\sigma}_m \hat{\sigma}_{m'})] = \text{tr}(\mathbb{E}[\hat{\sigma}_m] \mathbb{E}[\hat{\sigma}_{m'}]) \\ &= \text{tr}(\rho\rho) = p(\rho) \end{aligned}$$

whenever $m \neq m'$. To boost convergence to this desired expectation, we can form the empirical average over all distinct pairs:

$$\hat{p} = \frac{1}{M(M-1)} \sum_{m \neq m'} \text{tr}(\hat{\sigma}_m \hat{\sigma}_{m'}) = \binom{M}{2}^{-1} \sum_{m < m'} \text{tr}(\hat{\sigma}_m \hat{\sigma}_{m'}). \quad (\text{A20})$$

This formula describes an empirical average of $\binom{M}{2}$ random variables with the correct expectation value $p(\rho)$. This, in turn, ensures that $\mathbb{E}[\hat{p}] = p_2(\rho)$. However, in contrast to before, the individual random variables are not necessarily statistically independent. The first two terms $\text{tr}(\hat{\sigma}_1 \hat{\sigma}_2)$ and $\text{tr}(\hat{\sigma}_1 \hat{\sigma}_3)$, for instance, both depend on $\hat{\sigma}_1$. This prevents us from reusing exponential concentration inequalities, like the Bernstein inequality, to establish rapid convergence to this desired expectation value. More general, albeit weaker, concentration arguments still apply. Chebyshev's inequality, for instance, implies that

$$\begin{aligned} \Pr[|\hat{p} - p(\rho)| \geq \epsilon] &= \Pr[|\hat{p} - \mathbb{E}[\hat{p}]| \geq \epsilon] \leq \frac{\text{Var}[\hat{p}]}{\epsilon^2} \\ &\text{for any } \epsilon > 0. \end{aligned} \quad (\text{A21})$$

In words, the probability of an ϵ -deviation (or larger) is bounded by the variance $\text{Var}[\hat{p}]$ of our estimator divided

by ϵ^2 . This variance can be decomposed into individual contributions:

$$\begin{aligned} \text{Var}[\hat{p}] &= \mathbb{E}[\hat{p}^2] - \mathbb{E}[\hat{p}]^2 \\ &= \mathbb{E}[\hat{p}^2] - \text{tr}(\rho^2)^2 \\ &= \binom{M}{2}^{-2} \sum_{m_1 < m'_1} \sum_{m_2 < m'_2} (\text{tr}(\hat{\sigma}_{m_1} \hat{\sigma}_{m'_1}) \text{tr}(\hat{\sigma}_{m_2} \hat{\sigma}_{m'_2}) \\ &\quad - \text{tr}(\rho^2) \text{tr}(\rho^2)). \end{aligned}$$

Now, note that $\mathbb{E}[\hat{\sigma}_{m_1}] = \mathbb{E}[\hat{\sigma}_{m'_1}] = \mathbb{E}[\hat{\sigma}_{m_2}] = \mathbb{E}[\hat{\sigma}_{m'_2}] = \rho$ implies that these summands vanish unless either two or all four summation indices coincide. A careful case-by-case analysis yields

$$\begin{aligned} \text{Var}[\hat{p}] &= \binom{M}{2}^{-1} 2(M-2) \text{Var}[\text{tr}(\rho\hat{\sigma})] \\ &\quad + \binom{M}{2}^{-1} \text{Var}[\text{tr}(\hat{\sigma}\hat{\sigma}')] \\ &= \frac{4(M-2)}{M(M-1)} \text{Var}[\text{tr}(\rho\hat{\sigma})] \\ &\quad + \frac{2}{M(M-1)} \text{Var}[\text{tr}(\hat{\sigma}\hat{\sigma}')], \end{aligned} \quad (\text{A22})$$

and we refer the reader to [32, Supplemental Material] for details. Here, ρ is the underlying state and $\hat{\sigma}, \hat{\sigma}'$ denote independent instances of a classical shadow approximation. This reformulation contains two variance terms that depend on one (linear contribution) and two independent classical shadows (quadratic contribution), respectively. We can use Lemma 1 to control the first term. Set $O = \rho$ to conclude that

$$\text{Var}[\text{tr}(\rho\hat{\rho})] \leq 3^N \text{tr}(\rho^2) \leq 3^N, \quad (\text{A23})$$

because $\text{tr}(\rho^2) \leq 1$ for any underlying quantum state. Bounding the quadratic variance term requires more work. The following statement is a consequence of the geometric structure of SIC POVM measurements and substitutes existing arguments that rely on 3-design properties that do not apply here.

Lemma 2. *Let $\hat{\sigma}, \hat{\sigma}'$ be independent classical shadows of an underlying N -qubit state ρ . Then,*

$$\text{Var}[\text{tr}(\hat{\sigma}\hat{\sigma}')] = \mathbb{E}[\text{tr}(\hat{\sigma}\hat{\sigma}')^2] - \mathbb{E}[\text{tr}(\hat{\sigma}\hat{\sigma}')]^2 \leq 9^N.$$

We provide a detailed argument in Appendix A15 below. For now, we insert both bounds into Eq. (A22) to

obtain

$$\text{Var}[\hat{p}] \leq \frac{4(M-2)}{M(M-1)} 3^N + \frac{2}{M(M-1)} 9^N \leq \frac{4 \times 3^N}{M-1} + \left(\frac{\sqrt{2} \times 3^N}{M-1} \right)^2.$$

This variance bound diminishes as M increases. For fixed $\epsilon \in (0, 1)$ and $\delta \in (0, 1)$, a measurement budget of $M \geq (5 \times 3^N + 1)/(\delta\epsilon^2)$ ensures that $\text{Var}[\hat{p}] \leq \frac{4}{5}\epsilon^2\delta + \frac{2}{25}\epsilon^4\delta^2 < \epsilon^2\delta$. We can insert this implication into the Chebyshev bound (A21) to obtain a rigorous convergence guarantee for purity estimation:

$$M \geq (5 \times 3^N + 1)/(\delta\epsilon^2) \quad \text{ensures that} \\ \Pr[|\hat{p} - \text{tr}(\rho^2)| \geq \epsilon] \leq \delta.$$

In words, with probability (at least) $1 - \delta$, the purity estimator \hat{p} is ϵ -close to the true purity. Again, the required measurement budget M scales exponentially in the number of qubits involved. For global purities, this exponentially increasing measurement demand quickly becomes prohibitively expensive—a situation that cannot be avoided due to recent fundamental lower bounds [71]. However, once more, the situation changes if we consider *subsystem purities* instead. Let $\mathbf{K} \subseteq [N]$ be a subsystem composed of $|\mathbf{K}|$ qubits. The associated density matrix is $\rho_{\mathbf{K}} = \text{tr}_{-\mathbf{K}}(\rho)$ and we can estimate it by averaging appropriately marginalized classical shadows:

$$\hat{p}_{\mathbf{K}} = \binom{M}{2} \sum_{m \neq m'} \text{tr}(\text{tr}_{-\mathbf{K}}(\hat{\sigma}_m) \text{tr}_{-\mathbf{K}}(\hat{\sigma}'_{m'})) \quad \text{obeys} \\ \mathbb{E}[\hat{p}_{\mathbf{K}}] = \text{tr}(\rho_{\mathbf{K}}^2) = p(\rho_{\mathbf{K}}). \quad (\text{A24})$$

Importantly, this estimation process now depends only on the $|\mathbf{K}| < N$ qubits involved, such that

$$M \geq (5 \times 3^{|\mathbf{K}|} + 1)/(\delta\epsilon^2) \quad \text{ensures that} \\ \Pr[|\hat{p}_{\mathbf{K}} - p(\rho_{\mathbf{K}})| \geq \epsilon] \leq \delta.$$

This scaling is much more favorable, especially for small subsystems ($|\mathbf{K}| \ll N$). Similar to linear observable estimation, we can use this assertion to predict many subsystem purities based on the same classical shadows. A union bound argument, similar to the proof of Theorem 1 above, readily implies the following statement.

Theorem 2. *Suppose that we are interested in predicting L subsystem purities $p(\rho_{\mathbf{K}_i})$ of an unknown N -qubit state ρ . Let $K = \max_{1 \leq i \leq L} |\mathbf{K}_i|$ be the largest subsystem*

size involved and set $\epsilon, \delta \in (0, 1)$. Then,

$$M \geq 6L3^K/(\epsilon^2\delta) \quad (\text{A25})$$

N -qubit SIC POVM measurements on (independent copies of) ρ are likely to ϵ -approximate all subsystem purities simultaneously. More precisely, the resulting subsystem purity estimators $\hat{p}_{\mathbf{K}}$ defined in Eq. (A24) obey

$$\max_{1 \leq i \leq L} |\hat{p}_{\mathbf{K}_i} - \text{tr}(\rho_{\mathbf{K}_i}^2)| \leq \epsilon \quad \text{with probability (at least) } 1 - \delta.$$

The dependence on subsystem size K and accuracy ϵ is virtually identical to convergence guarantees for linear observable prediction; see Theorem 1. However, the dependence on the number of subsystem purities L and the inverse confidence $1/\delta$ now enter linearly, not logarithmically. This is a consequence of the fact that the individual contributions to $\hat{p}_{\mathbf{K}_i}$ are not statistically independent. In turn, we have to resort to Chebyshev's inequality instead of stronger exponential tail bounds like the Bernstein inequality.

It is possible to obtain a scaling proportional to $\log(2L/\delta)$ by using a more sophisticated estimation procedure known as *median of means estimation*; see, e.g., Ref. [14] for details. Practical tests with real data do, however, suggest that median of means estimation actually reduces the approximation quality overall [32]. This is not a contradiction, because statements like Theorem 2 are conservative mathematical statements about the worst-case rate of convergence. In practical applications, convergence can—and usually does—set in much earlier.

12. Convergence for higher-order polynomials and entanglement detection (outlook)

Quadratic estimation with classical shadows readily extends to higher-order polynomials. Such higher-order polynomials can be used, for instance, to probe entanglement in mixed states [32,33]. This is important because quadratic entanglement conditions—like subsystem Rényi entropies (purities)—only apply to global states that are reasonably pure [$\text{tr}(\rho^2) \approx 1$]. To see this, consider the maximally mixed state $\tau = (\frac{1}{2}\mathbb{I})^{\otimes N}$ on N qubits. This state is certainly not entangled, but nonetheless

$$R_2(\tau) = -\log_2[\text{tr}(\rho_{\mathbf{K}}^2)] \\ = -\log_2[\text{tr}((\frac{1}{2}\mathbb{I})^{\otimes |\mathbf{K}|})] \\ = -\log_2(2^{-|\mathbf{K}|}) \\ = |\mathbf{K}| \quad \text{for all subsystems } \mathbf{K} \subseteq N.$$

In words, second Rényi entropy is maximal for all subsystems simultaneously. This, however, is not a consequence of entanglement, but a trivial consequence of the fact that the state is very (maximally) mixed.

Fortunately, there exist entanglement criteria that extend to (very) mixed states. Chief among them is the positive partial transpose (PPT) criterion [72–74]. Let ρ be an N -qubit quantum state, and let $(\mathbf{A}, \bar{\mathbf{A}})$ be a bipartition of the qubits into two disjoint sets. Then, ρ is entangled (across the bipartition) if the partial transpose density matrix is not positive semidefinite (i.e., it has negative eigenvalues):

$$\rho^{T_{\mathbf{A}}} \not\geq 0 \quad \text{implies that } \rho \text{ is entangled} \\ \text{across the bipartition.} \quad (\text{A26})$$

The partial transpose is defined by transposing tensor factors belonging to subsystem \mathbf{A} , i.e., $(\rho_1 \otimes \cdots \otimes \rho_N)^{T_{\mathbf{A}}} = \bigotimes_{a \in \mathbf{A}} \rho_a^T \bigotimes_{\bar{a} \in \bar{\mathbf{A}}} \rho_{\bar{a}}$ and linearly extended to all N -qubit density matrices. A quick sanity check confirms that the maximally mixed state does not pass the PPT condition ($\mathbb{1}^T = \mathbb{1}$):

$$\tau^{T_{\mathbf{A}}} = \bigotimes_{a \in \mathbf{A}} (1/2\mathbb{1})^T \bigotimes_{\bar{a} \in \bar{\mathbf{A}}} (1/2\mathbb{1})^T = \bigotimes_{n \in [N]} (1/2\mathbb{1}) = \tau \geq 0.$$

Very entangled states, like the two-qubit Bell state $|\Omega\rangle = 1/\sqrt{2}(|00\rangle + |11\rangle)$ do, in contrast, have partial transposes with negative eigenvalues:

$$(|\Omega\rangle\langle\Omega|)^{T_1} = (|\Omega\rangle\langle\Omega|)^{T_2} = 1/\sqrt{2}\mathbb{F} \in \mathbb{H}_2^{\otimes 2},$$

where $\mathbb{F}|x\rangle \otimes |y\rangle = |y\rangle \otimes |x\rangle$ denotes the SWAP operator that has one negative eigenvalue [$\lambda_{\min}(\mathbb{F}) = -1$]. We call a state ρ with $\rho^{T_{\mathbf{A}}} \not\geq 0$ a *PPT-entangled state* [with respect to the bipartition $(\mathbf{A}, \bar{\mathbf{A}})$]. The PPT condition is a sufficient, but not necessary, condition for entanglement. It is known that there exist states that are entangled, but nonetheless obey $\rho^{T_{\mathbf{A}}} \geq 0$ [67]. So, it is fruitful to view the PPT criterion as a one-sided test for entanglement: if $\rho^{T_{\mathbf{A}}} \not\geq 0$, we can be sure that the state is entangled. But, $\rho^{T_{\mathbf{A}}} \geq 0$ does not necessarily imply that the state is not entangled (i.e., separable).

The PPT criterion (A26) is conceptually appealing, but it does require full and accurate knowledge of the density matrix ρ . This, in turn, typically requires full state tomography that quickly becomes prohibitively expensive. It is, however, possible to test consequences of $\rho^{T_{\mathbf{A}}} \geq 0$ by comparing moments of the partially transposed density matrix. The simplest consistency check is the so-called p_3 criterion [32]:

$$\rho^{T_{\mathbf{A}}} \geq 0 \quad \implies \quad \text{tr}((\rho^{T_{\mathbf{A}}})^3) \geq \text{tr}((\rho^{T_{\mathbf{A}}})^2)^2 = \text{tr}(\rho^2)^2. \quad (\text{A27})$$

The final simplification follows from the fact that partial transposition preserves the purity. The contrapositive of this implication serves as a (one-sided) test for entangle-

ment: if $\text{tr}((\rho^{T_{\mathbf{A}}})^3) < \text{tr}(\rho^2)^2$ [for some bipartition $(\mathbf{A}, \bar{\mathbf{A}})$] then the underlying state must be PPT entangled (across this bipartition).

Classical shadows can be used to directly estimate the trace moments involved in this test. Indeed, $\text{tr}(\rho^2)$ is just the purity, while $\text{tr}((\rho^{T_{\mathbf{A}}})^3)$ can be rewritten as a linear function on three copies of the underlying state:

$$\text{tr}((\rho^{T_{\mathbf{A}}})^3) = \text{tr}(O\rho \otimes \rho \otimes \rho).$$

We refer the reader to Ref. [32, Eq. (4)] for a precise reformulation. Subsequently, we can approximate this function by averaging over triples of distinct (and therefore independent) classical shadows:

$$\text{tr}(O\rho \otimes \rho \otimes \rho) \approx \frac{1}{6} \binom{M}{3}^{-1} \sum_{m \neq m' \neq m''} \text{tr}(O\hat{\sigma}_m \otimes \hat{\sigma}_{m'} \otimes \hat{\sigma}_{m''}).$$

The convergence analysis from above can, in principle, be extended to this form of cubic approximation. For randomized Pauli basis measurements, this has been done in the supplemental material of Ref. [32]. We leave a parallel treatment of cubic estimation with SIC POVM shadows for future work. The experimental and numerical results from the present work indicate that a SIC-POVM-based approach is expected to be both cheaper and easier than existing approaches based on Pauli basis measurements.

Finally, we point out that Ref. [33] extended the intuition behind the p_3 criterion (A27) to a complete family of polynomial consistency checks that compare polynomials of degree d with polynomials of degree $(d-1)$ and lower. This produces a hierarchy of in total $d_{\max} = 2^N$ consistency checks that is complete in the sense that a state ρ passes all of them if and only if $\rho^{T_{\mathbf{A}}} \geq 0$. Although polynomial estimation with classical shadows becomes more and more challenging as the degree d increases, the lower levels of this hierarchy may still be attainable with (comparatively) modest experimental and postprocessing effort.

13. Technical auxiliary results

14. Variance bounds for observable estimation

Here we prove Lemma 1, which we recall for convenience.

Lemma 3. Fix an N -qubit observable O and let $\hat{\sigma} \in \mathbb{H}_2^{\otimes N}$ be a (SIC POVM) classical shadow as defined in Eq. (A13). Then,

$$\text{Var}[\text{tr}(O\hat{\sigma})] \leq 3^N \text{tr}(O^2)$$

for any underlying N -qubit state ρ .

Proof. The classical shadow $\hat{\sigma}$ is constructed from performing single-qubit SIC POVM measurements on an underlying N -qubit quantum state ρ . Recall from Eq. (A9) that each of the 4^N possible outcome strings $i_1, \dots, i_N \in \{1, 2, 3, 4\}$ occurs with probability

$$\Pr[i_1 \cdots i_N | \rho] = 2^{-N} \langle \psi_{i_1}, \dots, \psi_{i_N} | \rho | \psi_{i_1}, \dots, \psi_{i_N} \rangle \leq 2^{-N}.$$

In words, the probability of any particular outcome occurring is bounded by 2^{-N} . This allows us to bound the dominating part of the variance by

$$\begin{aligned} \mathbb{E}[\text{tr}(O\hat{\sigma})^2] &= \sum_{i_1, \dots, i_N=1}^4 \Pr[i_1 \cdots i_N | \rho] \text{tr}(O(3|\psi_{i_1}\rangle\langle\psi_{i_1}| - \mathbb{I}) \otimes \cdots \otimes (3|\psi_{i_N}\rangle\langle\psi_{i_N}| - \mathbb{I}))^2 \\ &\leq 2^{-N} \sum_{i_1, \dots, i_N=1}^4 \text{tr}(O(3|\psi_{i_1}\rangle\langle\psi_{i_1}| - \mathbb{I}) \otimes \cdots \otimes (3|\psi_{i_N}\rangle\langle\psi_{i_N}| - \mathbb{I}))^2. \end{aligned}$$

Next, we expand the N -qubit observable O in terms of tensor products of single-qubit Pauli matrices $W_1, \dots, W_N \in \{\mathbb{I}, X, Y, Z\}$:

$$O = \sum_{W_1, \dots, W_N} o(W_1, \dots, W_N) W_1 \otimes \cdots \otimes W_N$$

with $o(W_1, \dots, W_N) = 2^{-N} \text{tr}(W_1 \otimes \cdots \otimes W_N O) \in \mathbb{R}$.

Such a decomposition into tensor products allows us to factorize the above bound into a product of single-qubit contributions:

$$\begin{aligned} \mathbb{E}[\text{tr}(O\hat{\sigma})^2] &\leq 2^{-N} \sum_{i_1, \dots, i_N=1}^4 \left(\sum_{W_1, \dots, W_N} o(W_1, \dots, W_N) \text{tr}((3|\psi_{i_1}\rangle\langle\psi_{i_1}| - \mathbb{I})W_1) \times \cdots \times \text{tr}((3|\psi_{i_N}\rangle\langle\psi_{i_N}| - \mathbb{I})W_N) \right)^2 \\ &= \sum_{V_1, \dots, V_N} \sum_{W_1, \dots, W_N} o(V_1, \dots, V_N) o(W_1, \dots, W_N) \\ &\quad \times \prod_{n=1}^N \underbrace{\left(\frac{1}{2} \sum_{i_n=1}^4 \text{tr}((3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I})V_n) \text{tr}((3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I})W_n) \right)}_{f(V_n, W_n)}. \end{aligned}$$

These single-qubit averages can be computed individually. Using the 1-design property (A3) [i.e., $\frac{1}{2} \sum_{i_n=1}^4 \frac{1}{2} \langle \psi_{i_n} | A | \psi_{i_n} \rangle = \text{tr}(A)$], we obtain

$$\begin{aligned} f(V_n, W_n) &= \frac{1}{2} \sum_{i_n=1}^4 \text{tr}((3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I})V_n) \text{tr}((3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I})W_n) \\ &= \frac{9}{2} \sum_{i_n=1}^4 \langle \psi_{i_n} | V_n | \psi_{i_n} \rangle \langle \psi_{i_n} | W_n | \psi_{i_n} \rangle - 3\text{tr}(V_n) \frac{1}{2} \sum_{i_n} \langle \psi_{i_n} | W_n | \psi_{i_n} \rangle \\ &\quad - 3\text{tr}(W_n) \frac{1}{2} \sum_{i_n=1}^4 \langle \psi_{i_n} | V_n | \psi_{i_n} \rangle + \frac{1}{2} \sum_{i_n=1}^4 \text{tr}(V_n) \text{tr}(W_n) \\ &= \frac{9}{2} \sum_{i_n=1}^4 \langle \psi_{i_n} | V_n | \psi_{i_n} \rangle \langle \psi_{i_n} | W_n | \psi_{i_n} \rangle - 4\text{tr}(V_n) \text{tr}(W_n). \end{aligned}$$

Next, we use the 2-design property (A4) of single-qubit SIC POVMs to obtain

$$\begin{aligned} f(V_n, W_n) &= \frac{9}{2} \sum_{i_n=1}^4 \langle \psi_{i_n} | V_n | \psi_{i_n} \rangle \langle \psi_{i_n} | W_n | \psi_{i_n} \rangle - 4 \text{tr}(V_n) \text{tr}(W_n) \\ &= 3(\text{tr}(V_n W_n) + \text{tr}(V_n) \text{tr}(W_n)) - 4 \text{tr}(V_n) \text{tr}(W_n) \\ &= 3 \text{tr}(V_n W_n) - \text{tr}(V_n) \text{tr}(W_n). \end{aligned}$$

For Pauli matrices V_n, W_n , this expression vanishes whenever $V_n \neq W_n$. It equals 2 if $V_n = W_n = \mathbb{I}$ and 6 if $V_n = W_n \neq \mathbb{I}$. That is,

$$f(V_n, W_n) = 2\delta(V_n, W_n)3^{1-\delta(W_n, \mathbb{I})}.$$

Inserting this closed-form expression into the original expression yields

$$\begin{aligned} \mathbb{E}[\text{tr}(O\hat{\sigma})^2] &\leq \sum_{V_1, \dots, V_N} \sum_{W_1, \dots, W_N} o(V_1, \dots, V_N) o(W_1, \dots, W_N) \prod_{n=1}^N 2\delta(V_n, W_n) 3^{1-\delta(V_n, \mathbb{I})} \\ &\leq 3^N 2^N \sum_{W_1, \dots, W_N} o(W_1, \dots, W_N)^2 \\ &= 3^N \sum_{W_1, \dots, W_N} 2^{-N} \text{tr}(W_1 \otimes \dots \otimes W_N O)^2 \\ &= 3^N \text{tr}(O^2), \end{aligned}$$

where the last equation follows from the fact that normalized N -qubit Pauli matrices form an orthonormal basis of $\mathbb{H}_2^{\otimes N}$ with respect to the Hilbert-Schmidt inner product (Parseval's identity). This completes the proof. \blacksquare

15. Variance bounds for purity estimation

Proposition 1 (Purity variance bound). *Let $\hat{\sigma}, \hat{\sigma}' \in \mathbb{H}_2^{\otimes N}$ be two independent classical shadows that arise from performing single-qubit SIC POVM measurements on a K -qubit state ρ . Then,*

$$\text{Var}[\text{tr}(\hat{\sigma} \hat{\sigma}')] = \mathbb{E}[\text{tr}(\hat{\sigma} \hat{\sigma}')^2] - \mathbb{E}[\text{tr}(\hat{\sigma} \hat{\sigma}')]^2 \leq 9^N.$$

The proof strategy behind this statement differs from existing arguments in the literature, most notably Refs. [14,32,33]. These use the fact that Pauli basis measurements form a 3-design, a structural property that does not apply to SIC POVMs. The key idea behind this new proof technique is to note that trace inner products of SIC POVM shadows can only assume very discrete values. Recall that

$$\hat{\sigma} = \bigotimes_{n=1}^N (3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I}) \quad \text{and} \quad \hat{\sigma}' = \bigotimes_{n=1}^N (3|\psi_{j_n}\rangle\langle\psi_{j_n}| - \mathbb{I}),$$

where $i_1, \dots, i_N \in \{1, 2, 3, 4\}$ and $j_1, \dots, j_N \in \{1, 2, 3, 4\}$ record the outcomes of each single-qubit SIC POVM measurement. This tensor product structure then implies that

$$\text{tr}(\hat{\sigma} \hat{\sigma}') = \prod_{n=1}^N \text{tr}((3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I})(3|\psi_{j_n}\rangle\langle\psi_{j_n}| - \mathbb{I})) = \prod_{n=1}^N (9|\langle\psi_{i_n}|\psi_{j_n}\rangle|^2 - 4),$$

and, because $|\psi_{i_n}\rangle, |\psi_{j_n}\rangle \in \mathbb{C}^2$ are SIC vectors, each contribution can only assume one of two discrete values:

$$9|\langle\psi_{i_n}|\psi_{j_n}\rangle|^2 - 4 = \begin{cases} +5 & \text{if } i_n = j_n, \\ -1 & \text{else if } i_n \neq j_n. \end{cases} \quad (\text{A28})$$

So, the magnitude of $\text{tr}(\hat{\sigma} \hat{\sigma}')$ scales exponentially in the number of coincidental measurement outcomes ($i_n = j_n$). This observation can be used to control the variance of this trace inner product. We first illustrate this for $N = 2$ qubits,

which is enough to convey the main gist. The proof of Proposition 1 is then a straightforward, yet somewhat technical, generalization to an arbitrary number of qubits.

In the two-qubit case, $\text{tr}(\hat{\sigma}\hat{\sigma}')^2$ can only assume three values: 25^2 if all single-qubit outcomes coincide, 25 if exactly one single-qubit outcome coincides, and 1 if no outcomes coincide. That is,

$$\begin{aligned} \text{tr}(\hat{\sigma}\hat{\sigma}')^2 &= 25^2\mathbf{1}\{i_1 = j_1 \wedge i_2 = j_2\} + 25\mathbf{1}\{i_1 = j_1 \wedge i_2 \neq j_2\} \\ &\quad + 25\mathbf{1}\{i_1 \neq j_1 \wedge i_2 = j_2\} + \mathbf{1}\{i_1 \neq j_1 \wedge i_2 \neq j_2\}, \end{aligned} \quad (\text{A29})$$

where $\mathbf{1}\{E\}$ denotes the indicator function of the event E . Next, we reexpress these indicator functions in terms of simpler ones:

$$\begin{aligned} \mathbf{1}\{i_1 = j_1 \wedge i_2 \neq j_2\} &= \mathbf{1}\{i_1 = j_1\} - \mathbf{1}\{i_1 = j_1 \wedge i_2 = j_2\}, \\ \mathbf{1}\{i_1 \neq j_1 \wedge i_2 = j_2\} &= \mathbf{1}\{i_2 = j_2\} - \mathbf{1}\{i_1 = j_1 \wedge i_2 = j_2\}, \\ \mathbf{1}\{i_1 \neq j_1 \wedge i_2 \neq j_2\} &= 1 - \mathbf{1}\{i_1 = j_1\} - \mathbf{1}\{i_2 = j_2\} + \mathbf{1}\{i_1 = j_1 \wedge i_2 = j_2\}. \end{aligned}$$

Inserting these reformulations into Eq. (A29) and rearranging terms yields

$$\begin{aligned} \text{tr}(\hat{\sigma}\hat{\sigma}')^2 &= (25^2 - 2 \times 25 + 1)\mathbf{1}\{i_1 = j_1 \wedge i_2 = j_2\} + (25 - 1)\mathbf{1}\{i_1 = j_1\} + (25 - 1)\mathbf{1}\{i_2 = j_2\} + 1 \\ &= (25 - 1)^2\mathbf{1}\{i_1 = j_1 \wedge i_2 = j_2\} + (25 - 1)\mathbf{1}\{i_1 = j_1\} + (25 - 1)\mathbf{1}\{i_2 = j_2\} + 1 \\ &= 8^2 \times 3^2\mathbf{1}\{i_1 = j_1 \wedge i_2 = j_2\} + 8 \times 3\mathbf{1}\{i_1 = j_1\} + 8 \times 3\mathbf{1}\{i_2 = j_2\} + 1, \end{aligned}$$

where we have used the fact that $(25 - 1) = 24 = 8 \times 3$. Now, we are ready to take expectation values. Recall that taking the expectation of an indicator function produces the probability of the associated event:

$$\mathbb{E}[\text{tr}(\hat{\sigma}\hat{\sigma}')^2] = 8^2 \times 3^2\text{Pr}[i_1 = j_1 \wedge i_2 = j_2] + 8 \times 3\text{Pr}[i_1 = j_1] + 8 \times 3\text{Pr}[i_2 = j_2] + 1. \quad (\text{A30})$$

These probabilities for coincidental measurement outcomes can be computed explicitly. This is the content of the following auxiliary result.

Lemma 4. *Suppose that we perform two N -qubit SIC POVM measurements on (distinct copies of) a quantum state ρ and let $\mathbf{K} \subseteq [N] = \{1, \dots, N\}$ be a subset of $K = |\mathbf{K}|$ qubits. Then, the probability that the obtained measurement outcomes are equal ($i_k = j_k$) for all $k \in \mathbf{K}$ obeys*

$$\text{Pr}\left[\bigwedge_{k \in \mathbf{K}} \{i_k = j_k\}\right] = \text{tr}(\rho_{\mathbf{K}} \mathcal{D}_{1/3}^{\otimes K}(\rho_{\mathbf{K}})) \leq 3^{-K},$$

where $\rho_{\mathbf{K}} = \text{tr}_{-K}(\rho)$ is the reduced density matrix supported on the relevant qubit subset and each $\mathcal{D}_{1/3}$ is a single-qubit depolarizing channel.

The proof follows from exploiting the fact that the two SIC POVM measurements are statistically independent, as well as the 2-design property of SIC POVMs. We defer it to the end of this section. For the task at hand, Lemma 3 bounds all remaining probabilities in Eq. (A30). Doing so, conveniently cancels the existing powers of 3 and produces the bound given in Proposition 1 for $K = 2$ qubits:

$$\begin{aligned} \mathbb{E}[\text{tr}(\hat{\sigma}\hat{\sigma}')^2] &= 8^2 \times 3^2\text{Pr}[i_1 = j_1 \wedge i_2 = j_2] + 8 \times 3\text{Pr}[i_1 = j_1] + 8 \times 3\text{Pr}[i_2 = j_2] + 1 \\ &\leq 8^2 + 2 \times 8 + 1 \\ &= (8 + 1)^2 \\ &= 9^2. \end{aligned}$$

This argument can be readily extended to an arbitrary number of qubits.

Proof of Proposition 1. The trace inner product between two N -qubit SIC POVM shadows can only assume discrete values. Indeed, Eq. (A28) states that $\text{tr}(\hat{\sigma}\hat{\sigma}') = \pm 5^c$, where c is the number of coincidental measurement outcomes. We can use indicator functions to single out all possibilities for coincidences and multiplying them with the correct scaling factor provides a closed-form expression of the trace inner product in terms of measurement outcomes alone. In the following, we use $\mathbf{K} \subseteq [N]$ to denote a subset of coincidental indices. The complementary set (where indices must not coincide) will be denoted by $\bar{\mathbf{K}} = [N] \setminus \mathbf{K}$. For the squared trace inner product, we then obtain

$$\begin{aligned} \text{tr}(\hat{\sigma}\hat{\sigma}')^2 &= \sum_{\mathbf{K} \subseteq [N]} 25^{|\mathbf{K}|} \mathbf{1}\left(\bigwedge_{k \in \mathbf{K}} \{i_k = j_k\} \bigwedge_{\bar{k} \in \bar{\mathbf{K}}} \{i_{\bar{k}} \neq j_{\bar{k}}\}\right) \\ &= \sum_{\mathbf{K} \subseteq [N]} 25^{|\mathbf{K}|} \left(\sum_{\mathbf{T} \subseteq \bar{\mathbf{K}}} (-1)^{|\mathbf{T}|} \mathbf{1}\left(\bigwedge_{u \in \mathbf{K} \cup \mathbf{T}} \{i_u = j_u\}\right) \right) \\ &= \sum_{\mathbf{U} \subseteq [N]} \left(\sum_{\mathbf{T} \subseteq \mathbf{U}} (-1)^{|\mathbf{T}|} 25^{|\mathbf{U}| - |\mathbf{T}|} \right) \mathbf{1}\left(\bigwedge_{u \in \mathbf{U}} \{i_u = j_u\}\right). \end{aligned}$$

In the second line, we have reexpressed conditions for noncoincidence ($\{i_{\bar{k}} \neq j_{\bar{k}}\}$) as linear combinations of coincidences on larger subsystems ($\mathbf{K} \cup \mathbf{T}$ with $\mathbf{T} \subseteq \bar{\mathbf{K}}$). The last line follows from introducing the union $\mathbf{U} = \mathbf{K} \cup \mathbf{T}$ and rewriting \mathbf{K} as $\mathbf{K} \setminus \mathbf{T}$. The inner sum over subsets $\mathbf{T} \subseteq \mathbf{U}$ has no effect on the indicator function. The size of such sets ranges from $T = |\mathbf{T}| = 0$ up to $T = |\mathbf{T}| = |\mathbf{U}|$ and, for each T , there are $\binom{|\mathbf{U}|}{T}$ subsets of that size. For a fixed set \mathbf{U} , we therefore obtain

$$\begin{aligned} \sum_{\mathbf{T} \subseteq \mathbf{U}} (-1)^{|\mathbf{T}|} 25^{|\mathbf{U}| - |\mathbf{T}|} &= \sum_{T=0}^{|\mathbf{U}|} \binom{|\mathbf{U}|}{T} (-1)^T 25^{|\mathbf{U}| - T} \\ &= (25 - 1)^{|\mathbf{U}|} = 8^{|\mathbf{U}|} \times 3^{|\mathbf{U}|}, \end{aligned}$$

which considerably simplifies the entire function. Taking the expectation now produces

$$\begin{aligned} \mathbb{E}[\text{tr}(\hat{\sigma}\hat{\sigma}')] &= \sum_{\mathbf{U} \subseteq [N]} 8^{|\mathbf{U}|} \times 3^{|\mathbf{U}|} \mathbb{E}\left[\mathbf{1}\left(\bigwedge_{u \in \mathbf{U}} \{i_u = j_u\}\right)\right] \\ &= \sum_{\mathbf{U} \subseteq [N]} 8^{|\mathbf{U}|} \times 3^{|\mathbf{U}|} \Pr\left[\bigwedge_{u \in \mathbf{U}} \{i_u = j_u\}\right], \end{aligned}$$

and we can use Lemma 3 to complete the argument. Indeed, $\Pr[\bigwedge_{u \in \mathbf{U}} \{i_u = j_u\}] \leq 3^{-|\mathbf{U}|}$ ensures that

$$\mathbb{E}[\text{tr}(\hat{\sigma}\hat{\sigma}')] \leq \sum_{\mathbf{U} \subseteq [N]} 8^{|\mathbf{U}|} = \sum_{U=0}^N \binom{N}{U} 8^U = (8+1)^N = 9^N.$$

This is the advertised bound for the variance of N -qubit purity estimators. ■

Finally, we provide the proof for the bound on coincidental SIC POVM measurement outcomes.

Proof of Lemma 3. For simplicity, we assume that the subset $\mathbf{K} = \{1, \dots, K\} \subseteq [N]$ encompasses the first $K = |\mathbf{K}|$ qubits. The general case works analogously, but notation becomes somewhat cumbersome. We perform two independent single-qubit SIC POVM measurements on (two copies of) an N -qubit quantum state ρ . The probability of getting $K = |\mathbf{K}|$ particular outcomes depends only on the reduced density matrix $\rho_{\mathbf{K}} = \text{tr}_{-\mathbf{K}}(\rho)$ of the relevant qubit subset:

$$\begin{aligned} \Pr[i_1 \dots i_K | \rho_{\mathbf{K}}] &= 2^{-K} \langle \psi_{i_1}, \dots, \psi_{i_K} | \rho_{\mathbf{K}} | \psi_{i_1}, \dots, \psi_{i_K} \rangle \\ &\text{for } i_1, \dots, i_K \in \{1, 2, 3, 4\}. \end{aligned}$$

This observation allows us to rewrite the probability for K coincidental measurement outcomes as

$$\begin{aligned} \Pr\left[\bigwedge_{k=1}^K \{i_k = j_k\}\right] &= \sum_{i_1=1}^4 \dots \sum_{i_K=1}^4 \Pr[i_1 \dots i_K | \rho_{\mathbf{K}}]^2 \\ &= \frac{1}{4^K} \sum_{i_1=1}^4 \dots \sum_{i_K=1}^4 \langle \psi_{i_1}, \dots, \psi_{i_K} | \rho_{\mathbf{K}} | \\ &\quad \times \psi_{i_1}, \dots, \psi_{i_K} \rangle^2. \end{aligned}$$

At this point it is helpful to decompose (one) $\rho_{\mathbf{K}}$ into a linear combination of tensor products, e.g., $\rho_{\mathbf{K}} = \sum_{W_1, \dots, W_K} r(W_1, \dots, W_K) W_1 \otimes \dots \otimes W_K$. Doing so allows us to rewrite

$$\begin{aligned} \Pr\left[\bigwedge_{k=1}^K \{i_k = j_k\}\right] &= 2^{-K} \text{tr}\left(\rho_{\mathbf{K}} \sum_{W_1, \dots, W_K} r(W_1, \dots, W_K) \right. \\ &\quad \times \left(\frac{1}{2} \sum_{i_1=1}^4 |\psi_{i_1}\rangle \langle \psi_{i_1}| \langle \psi_{i_1} | W_1 | \psi_{i_1}\rangle\right) \otimes \dots \\ &\quad \left. \otimes \left(\frac{1}{2} \sum_{i_K=1}^4 |\psi_{i_K}\rangle \langle \psi_{i_K}| \langle \psi_{i_K} | W_K | \psi_{i_K}\rangle\right)\right) \\ &= 2^{-K} \text{tr}\left(\rho_{\mathbf{K}} \sum_{W_1, \dots, W_K} r(W_1, \dots, W_K) \mathcal{D}_{1/3}(W_1) \right. \\ &\quad \left. \otimes \dots \otimes \mathcal{D}_{1/3}(W_K)\right), \\ &= 2^{-K} \text{tr}(\rho_{\mathbf{K}} \mathcal{D}_{1/3}^{\otimes K}(\rho_{\mathbf{K}})), \end{aligned}$$

as advertised. Here, we have used the 2-design property (A4) of SIC POVMs, more precisely, $\sum_{i=1}^4 \frac{1}{2} |\psi_i\rangle \langle \psi_i| \langle \psi_i | A | \psi_i\rangle = \mathcal{D}_{1/3}(A)$. To get the state-independent upper bound, we note that each depolarizing channel is a linear combination between the identity channel

$[\mathcal{I}(A) = A]$ and the projection onto the identity matrix $[\mathcal{T}(A) = \text{tr}(A)\mathbb{I}]$: $\mathcal{D}_{1/3} = \frac{1}{3}(\mathcal{I} + \mathcal{T})$. We also drop the subscript \mathbf{K} in ρ to declutter notation somewhat: $\rho_{\mathbf{K}} \mapsto \rho$. Then,

$$\begin{aligned} 2^{-K} \text{tr}(\rho \mathcal{D}_{1/3}^{\otimes K}(\rho)) &= 2^{-K} \text{tr}(\rho 3^{-K} (\mathcal{I} + \mathcal{T})^{\otimes K}(\rho)) \\ &= \frac{1}{3^K} 2^{-K} \sum_{T \subseteq \{1, \dots, s\}} \text{tr}(\rho \mathcal{T}_T \otimes \mathcal{I}_{\bar{T}}(\rho)) \\ &= 3^{-K} \left(2^{-K} \sum_{T \subseteq \{1, \dots, K\}} \text{tr}(\rho \rho_T \otimes \mathbb{I}_{\bar{T}}) \right) \\ &= 3^{-K} \left(2^{-K} \sum_{T \subseteq \{1, \dots, s\}} \text{tr}(\rho_T^2) \right). \end{aligned}$$

The remaining bracket averages over all subsystem purities $\text{tr}(\rho_T^2)$. Each of them obeys $\text{tr}(\rho_T^2) \leq 1$ and there are a total of 2^K of them (a finite set of size K has 2^K subsets). Upper bounding each of them by 1 produces

$$2^{-K} \text{tr}(\rho \mathcal{D}_{1/3}^{\otimes K}(\rho)) = 3^{-K} \left(2^{-K} \sum_{T \subseteq \{1, \dots, s\}} \text{tr}(\rho_T^2) \right) \leq 3^{-K}.$$

This completes the proof. \blacksquare

-
- [1] M. Paris and J. Řeháček, *Quantum State Estimation*, edited by M. Paris and J. Řeháček, Lecture Notes in Physics, Vol. 649 (Springer Berlin, Heidelberg, 2004).
- [2] J. M. Renes, R. Blume-Kohout, A. J. Scott, and C. M. Caves, Symmetric informationally complete quantum measurements, *J. Math. Phys.* **45**, 2171 (2004).
- [3] A. J. Scott, Tight informationally complete quantum measurements, *J. Phys. A* **39**, 13507 (2006).
- [4] A. J. Scott and M. Grassl, Symmetric informationally complete positive-operator-valued measures: A new computer study, *J. Math. Phys.* **51**, 042203 (2010).
- [5] C. A. Fuchs, M. C. Hoang, and B. C. Stacey, The SIC question: History and state of play, *Axioms* **6**, 00 (2017).
- [6] G. N. M. Tabia, Experimental scheme for qubit and qutrit symmetric informationally complete positive operator-valued measurements using multipoint devices, *Phys. Rev. A* **86**, 062107 (2012).
- [7] Z. Bian, J. Li, H. Qin, X. Zhan, R. Zhang, B. C. Sanders, and P. Xue, Realization of Single-Qubit Positive-Operator-Valued Measurement via a One-Dimensional Photonic Quantum Walk, *Phys. Rev. Lett.* **114**, 203602 (2015).
- [8] M. Proietti, M. Ringbauer, F. Graffitti, P. Barrow, A. Pickston, D. Kundys, D. Cavalcanti, L. Aolita, R. Chaves, and A. Fedrizzi, Enhanced Multiqubit Phase Estimation in Noisy Environments by Local Encoding, *Phys. Rev. Lett.* **123**, 180503 (2019).
- [9] N. Bent, H. Qassim, A. A. Tahir, D. Sych, G. Leuchs, L. L. Sánchez-Soto, E. Karimi, and R. W. Boyd, Experimental Realization of Quantum Tomography of Photonic Qudits via Symmetric Informationally Complete Positive Operator-Valued Measures, *Phys. Rev. X* **5**, 041006 (2015).
- [10] P.-X. Chen, J. A. Bergou, S.-Y. Zhu, and G.-C. Guo, Ancilla dimensions needed to carry out positive-operator-valued measurement, *Phys. Rev. A* **76**, 060303 (2007).
- [11] G. García-Pérez, M. A. Rossi, B. Sokolov, F. Tacchino, P. K. Barkoutsos, G. Mazzola, I. Tavernelli, and S. Maniscalco, Learning to Measure: Adaptive Informationally Complete Generalized Measurements for Quantum Algorithms, *PRX Quantum* **2**, 040342 (2021).
- [12] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, *Nat. Phys.* **14**, 447 (2018).
- [13] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, Reconstructing quantum states with generative models, *Nat. Machine Intell.* **1**, 155 (2019).
- [14] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [15] M. Pains and A. Kalev, An approximate description of quantum states, (2019), Preprint at [ArXiv:1910.10543](https://arxiv.org/abs/1910.10543).
- [16] A. Elben, S. T. Flammia, H.-Y. Huang, R. Kueng, J. Preskill, B. Vermersch, and P. Zoller, The randomized measurement toolbox, arXiv preprint (2022), [ArXiv:2203.11374](https://arxiv.org/abs/2203.11374).
- [17] M. Ringbauer, M. Meth, L. Postler, R. Stricker, R. Blatt, P. Schindler, and T. Monz, A universal qudit quantum processor with trapped ions, (2021), Preprint at [ArXiv:2109.06903](https://arxiv.org/abs/2109.06903).
- [18] M. Naimark, Spectral functions of a symmetric operator, *Bull. Acad. Sci. URSS. Sér. Math. [Izvestia Akad. Nauk SSSR]* **4**, 277 (1940).
- [19] J. Ahn, T. C. Weinacht, and P. H. Bucksbaum, Information storage and retrieval through quantum phase, *Science* **287**, 463 (2000).
- [20] B. E. Anderson, H. Sosa-Martinez, C. A. Riofrío, I. H. Deutsch, and P. S. Jessen, Accurate and Robust Unitary Transformations of a High-Dimensional Quantum System, *Phys. Rev. Lett.* **114**, 240401 (2015).
- [21] C. Godfrin, A. Ferhat, R. Ballou, S. Klyatskaya, M. Ruben, W. Wernsdorfer, and F. Balestro, Operating Quantum States in Single Magnetic Molecules: Implementation of Grover's Quantum Algorithm, *Phys. Rev. Lett.* **119**, 187702 (2017).
- [22] X.-M. Hu, Y. Guo, B.-H. Liu, Y.-F. Huang, C.-F. Li, and G.-C. Guo, Beating the channel capacity limit for superdense coding with entangled ququarts, *Sci. Adv.* **4**, eaat9304 (2018).
- [23] Y. Chi, *et al.*, A programmable qudit-based quantum processor, *Nat. Commun.* **13**, 1166 (2022).
- [24] M. S. Blok, V. V. Ramasesh, T. Schuster, K. O'Brien, J. M. Kreikebaum, D. Dahlen, A. Morvan, B. Yoshida, N. Y. Yao, and I. Siddiqi, Quantum Information Scrambling on a Superconducting Qutrit Processor, *Phys. Rev. X* **11**, 021010 (2021).
- [25] Z. Hradil, Quantum-state estimation, *Phys. Rev. A* **55**, R1561 (1997).
- [26] C. J. Wood, Initialization and characterization of open quantum systems, Ph.D. thesis, University of Waterloo (2015).

- [27] J. A. Smolin, J. M. Gambetta, and G. Smith, Efficient Method for Computing the Maximum-Likelihood Quantum State from Measurements with Additive Gaussian Noise, *Phys. Rev. Lett.* **108**, 070502 (2012).
- [28] T. Sugiyama, P. S. Turner, and M. Muraio, Precision-Guaranteed Quantum Tomography, *Phys. Rev. Lett.* **111**, 160406 (2013).
- [29] M. Guřa, J. Kahn, R. Kueng, and J. A. Tropp, Fast state tomography with optimal error bounds, *J. Phys. A: Math. Theor.* **53**, 204001 (2020).
- [30] R. Blume-Kohout, Optimal, reliable estimation of quantum states, *New J. Phys.* **12**, 043034 (2010).
- [31] J. M. Lukens, K. J. H. Law, A. Jasra, and P. Lougovski, A practical and efficient approach for Bayesian quantum state estimation, *New J. Phys.* **22**, 063038 (2020).
- [32] A. Elben, R. Kueng, H.-Y. R. Huang, R. van Bijnen, C. Kokail, M. Dalmonte, P. Calabrese, B. Kraus, J. Preskill, P. Zoller, and B. Vermersch, Mixed-State Entanglement from Local Randomized Measurements, *Phys. Rev. Lett.* **125**, 200501 (2020).
- [33] A. Neven, J. Carrasco, V. Vitale, C. Kokail, A. Elben, M. Dalmonte, P. Calabrese, P. Zoller, B. Vermersch, R. Kueng, and B. Kraus, Symmetry-resolved entanglement detection using partial transpose moments, *npj Quantum Inf.* **7**, 152 (2021).
- [34] A. Ambainis and J. Emerson, in *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC'07)* (2007), p. 129.
- [35] R. Kueng and D. Gross, Qubit stabilizer states are complex projective 3-designs, arXiv preprint (2015), [ArXiv:1510.02767](https://arxiv.org/abs/1510.02767).
- [36] H. Zhu, Multiqubit Clifford groups are unitary 3-designs, *Phys. Rev. A* **96**, 062336 (2017).
- [37] Z. Webb, The Clifford group forms a unitary 3-design, *Quantum Info. Comput.* **16**, 1379 (2016).
- [38] P. Schindler, D. Nigg, T. Monz, J. T. Barreiro, E. Martinez, S. X. Wang, S. Quint, M. F. Brandl, V. Nebendahl, C. F. Roos, M. Chwalla, M. Hennrich, and R. Blatt, A quantum information processor with trapped ions, *New J. Phys.* **15**, 123012 (2013).
- [39] M. Enřiquez, I. Wintrowicz, and K. Życzkowski, Maximally entangled multipartite states: A brief survey, *J. Phys.: Conf. Ser.* **698**, 012003 (2016).
- [40] W. Helwig and W. Cui, Absolutely maximally entangled states: Existence and applications, (2013), Preprint at [ArXiv:1306.2536](https://arxiv.org/abs/1306.2536).
- [41] M. Grassl, T. Beth, and M. Roetteler, On optimal quantum codes, *Int. J. Quantum Inf.* **02**, 55 (2003).
- [42] S. Muralidharan and P. K. Panigrahi, Perfect teleportation, quantum-state sharing, and superdense coding through a genuinely entangled five-qubit state, *Phys. Rev. A* **77**, 032321 (2008).
- [43] C. Ferrie and R. Blume-Kohout, *Maximum likelihood quantum state tomography is inadmissible*, (2018).
- [44] T. Brydges, A. Elben, P. Jurcevic, B. Vermersch, C. Maier, B. P. Lanyon, P. Zoller, R. Blatt, and C. F. Roos, Probing Ręnyi entanglement entropy via randomized measurements, *Science* **364**, 260 (2019).
- [45] H.-Y. Huang, M. Broughton, J. Cotler, S. Chen, J. Li, M. Mohseni, H. Neven, R. Babbush, R. Kueng, and J. Preskill, *et al.*, Quantum advantage in learning from experiments, (2021), arXiv preprint [ArXiv:2112.00778](https://arxiv.org/abs/2112.00778).
- [46] M. Appleby, S. Flammia, G. McConnell, and J. Yard, SICs and algebraic number theory, *Found. Phys.* **47**, 1042 (2017).
- [47] J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu, Sample-optimal tomography of quantum states, *IEEE Trans. Inform. Theory* **63**, 5628 (2017).
- [48] R. O'Donnell and J. Wright, in *STOC'16—Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing* (ACM, New York, 2016), p. 899.
- [49] D. A. Roberts and B. Yoshida, *Chaos and complexity by design*, *J. High Energy Phys.*, 121, front matter+63 (2017).
- [50] F. G. Brandão, W. Chemissany, N. Hunter-Jones, R. Kueng, and J. Preskill, Models of Quantum Complexity Growth, *PRX Quantum* **2**, 030316 (2021).
- [51] Z. Huangjun, Quantum state estimation and symmetric informationally complete POMs, Ph.D. thesis, National University of Singapore (2012).
- [52] D. H. Mahler, L. A. Rozema, A. Darabi, C. Ferrie, R. Blume-Kohout, and A. Steinberg, Adaptive Quantum State Tomography Improves Accuracy Quadratically, *Phys. Rev. Lett.* **111**, 183601 (2013).
- [53] F. Huszár and N. M. Houlby, Adaptive Bayesian quantum tomography, *Phys. Rev. A* **85**, 052120 (2012).
- [54] C. Ferrie, Self-Guided Quantum Tomography, *Phys. Rev. Lett.* **113**, 190404 (2014).
- [55] Z. Hou, H. Zhu, G.-Y. Xiang, C.-F. Li, and G.-C. Guo, Achieving quantum precision limit in adaptive qubit state tomography, *npj Quantum Inf.* **2**, 1 (2016).
- [56] S. M. Kazim, A. Farooq, J. ur Rehman, M. Sohail, M. Pasha, M. Nadeem, and A. Ali, Adaptive quantum state tomography with iterative particle filtering, *Quantum Inf. Process.* **20**, 348 (2021).
- [57] R. J. Chapman, C. Ferrie, and A. Peruzzo, Experimental Demonstration of Self-Guided Quantum Tomography, *Phys. Rev. Lett.* **117**, 040402 (2016).
- [58] Z. Hou, J.-F. Tang, C. Ferrie, G.-Y. Xiang, C.-F. Li, and G.-C. Guo, Experimental realization of self-guided quantum process tomography, *Phys. Rev. A* **101**, 022317 (2020).
- [59] M. Rambach, M. Qaryan, M. Kewming, C. Ferrie, A. G. White, and J. Romero, Robust and Efficient High-Dimensional Quantum State Tomography, *Phys. Rev. Lett.* **126**, 100402 (2021).
- [60] B. Qi, Z. Hou, Y. Wang, Q. Guo, Z.-Y. Fang, X.-H. Bao, X. Zhu, and R.-B. Liu, Adaptive quantum state tomography via linear regression estimation: Theory and two-qubit experiment, *npj Quantum Inf.* **3**, 19 (2017).
- [61] G. I. Struchalin, I. A. Pogorelov, S. S. Straupe, K. S. Kravtsov, I. V. Radchenko, and S. P. Kulik, Experimental adaptive quantum tomography of two-qubit states, *Phys. Rev. A* **93**, 012103 (2016).
- [62] D. Dieks, Overlap and distinguishability of quantum states, *Phys. Lett. A* **126**, 303 (1988).
- [63] R. Stricker, M. Meth, L. Postler, C. Edmunds, C. Ferrie, R. Blatt, P. Schindler, T. Monz, R. Kueng, and M. Ringbauer, Data for: "Experimental Single-Setting Quantum State Tomography", *PRX Quantum*, (2022), <https://doi.org/10.5281/zenodo.7054827>.
- [64] L. E. Fischer, D. Miller, F. Tacchino, P. K. Barkoutsos, D. J. Egger, and I. Tavernelli, *Ancilla-free implementation of*

- generalized measurements for qubits embedded in a qudit space*, (2022).
- [65] K. Mølmer and A. Sørensen, Multiparticle Entanglement of Hot Trapped Ions, *Phys. Rev. Lett.* **82**, 1835 (1999).
- [66] G. Vidal and R. F. Werner, Computable measure of entanglement, *Phys. Rev. A* **65**, 032314 (2002).
- [67] P. Horodecki, Separability criterion and inseparable mixed states with positive partial transposition, *Phys. Lett. A* **232**, 333 (1997).
- [68] A. Klappenecker and M. Rotteler, in *Proceedings. International Symposium on Information Theory, 2005. ISIT 2005.* (2005), p. 1740.
- [69] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*, Applied and Numerical Harmonic Analysis (Birkhäuser/Springer, New York, 2013), p. xviii+625.
- [70] R. Vershynin, *High-Dimensional Probability*, Cambridge Series in Statistical and Probabilistic Mathematics, Vol. 47 (Cambridge University Press, Cambridge, 2018), p. xiv+284, an introduction with applications in data science, With a foreword by Sara van de Geer.
- [71] S. Chen, J. Cotler, H.-Y. Huang, and J. Li, in *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science—FOCS 2021* (IEEE Computer Soc., Los Alamitos, CA, [2022] ©2022), p. 574.
- [72] A. Peres, Separability Criterion for Density Matrices, *Phys. Rev. Lett.* **77**, 1413 (1996).
- [73] R. Horodecki and M. Horodecki, Information-theoretic aspects of inseparability of mixed states, *Phys. Rev. A* **54**, 1838 (1996).
- [74] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, Quantum entanglement, *Rev. Mod. Phys.* **81**, 865 (2009).

VERIFYING AN UNTRUSTED QUANTUM DEVICE

The developments in quantum devices over recent years have enabled a steady increase in their system sizes [14, 15]. While these advanced systems offer significant new computational capabilities, we are now beginning to enter a regime where the devices can no longer be fully simulated classically. The field of quantum verification fills this gap by developing strategies for a computationally weak verifier (the user) to confirm the correctness of the computational results of a much more powerful quantum prover (the quantum computer), with as few additional resources as possible [56].

Existing quantum verification strategies typically build trust in devices by characterizing a device thoroughly with techniques such as randomized benchmarking [49] or gate-set tomography [50]. Some strategies cross-verify multiple, independent quantum devices to verify a quantum computation [59] or even the quantum devices themselves [60]. Cross-verification stands out because it does not require additional overhead for the verification procedure other than the use of multiple independent quantum devices. Other strategies rely on additional quantum resources to allow the computation with encrypted data [61], e.g., blind verification [62]. Blind verification requires limited quantum resources on the verifier side, but demands a large operational overhead for implementing a cryptographically secure verification protocol. These quantum verification approaches, among others, have been discussed in detail in Ch. 1.2.2. However, additional quantum resources or confidence in the quantum device or the person controlling it, are not always available. For example, when running a quantum computer via cloud access. The most powerful verification path therefore involves a completely untrusted quantum prover and a verifier left with classical resources only. Such purely classical verification has long been considered impossible.

Remarkably, Mahadev in Ref. [63] has recently developed a framework that indeed opens an avenue to classical verification. In simple terms, her protocol is based on an interactive exchange of messages in which a classical verifier first requests a commitment from the quantum computer in the form of a measurement outcome on part of a quantum state. Once received, the verifier randomly asks the prover to provide one of two incompatible measurements on the rest of that state. One of these is used to verify the answer to the computation, and the other is used to verify that the prover is not cheating. The key component of this interactive protocol is a cryptographic technique that allows the verifier to request two types of measurements in such a way that the quantum computer cannot tell them apart.

Here we attempt to classically verify the output of a quantum computation. The relevant building blocks are first discussed at the example of verifying classical computations in Sec. 4.1. In this context, we introduce a verification procedure in Sec. 4.1.1, discuss the problem classes that can be verified in such a way in Sec. 4.1.2, and emphasize their interactive structure in Sec. 4.1.3 to learn that interaction is the basis to many promising

verification approaches. In Secs. 4.1.4-4.1.5 we introduce functions that guarantee encrypted interactions even against quantum computers. With this knowledge, we proceed to the concepts of classical verification of quantum computation in Sec. 4.2. Key components of the protocol are reformulating the quantum computation to be verified as an interactive energy measurement, which is discussed in Secs. 4.2.1-4.2.3, and encoding the energy measurements using the so-called *learning with errors* (LWE) problem which is hard for classical as well as quantum computers, discussed in Sec. 4.2.2. We then present a step-by-step classical verification protocol in Sec. 4.2.4, that is slightly modified from the original proposal in Ref. [63] to make it practical for current NISQ-hardware [54]. The chapter concludes with our publication on the first classical verification of a quantum computation, presented in Sec. 4.3.

4.1 VERIFIABLE CLASSICAL COMPUTATION

Today's classical computing tasks are increasingly outsourced from a comparatively weak computational device, the client, to a much more powerful device, the worker. In this way, the client gains access to the output of a task that exceeds its computational capabilities. Instead of performing the actual task, a dishonest worker could give an answer, which the client cannot check so the client must trust the worker. To fight this lack of security, many strategies have been developed in which a client verifies the output of a stronger worker with only limited computational resources, which is the realm of *verifiable computation* [172].

Interestingly, the concepts of verifiable classical computation [172] largely coincide with those of classical verification of quantum computation [63]. This section aims to provide the overlapping building blocks using classical computation as an example to form a foundation for the upcoming quantum verification by purely classical means, discussed in Sec. 4.2.

4.1.1 *Arthur-Merlin protocol*

Consider the following scenario between King Arthur and the wizard Merlin [173]. Although Arthur rules over a large kingdom, he has limited computing power. Arthur's computer is able to manipulate bit strings according to certain rules and performs one action at a time, which is called a *deterministic Turing machine* [53]. Arthur can solve problems whose time consumption grows polynomially with the size of the problem. The all-powerful Merlin, on the other hand, has unlimited computational power. One day, Arthur is confronted with a task that exceeds his computational abilities and seeks Merlin for help. Crucially, The task is phrased as a "yes" or "no" question, known as a *decision problem*. The wizard agrees and presents the king with the answer. After a while, the initial satisfaction fades and Arthur wonders whether he can trust Merlin's answer, as he himself has no way of checking it directly. Arthur turns to Merlin again and requests evidence to support the wizard's claim. Crucially, the pieces of evidence Arthur asks for must be comprehensible with the king's computational capabilities, i.e. they must be verifiable in polynomial runtime. In a series of such queries, Arthur proves Merlin's answers and accepts or rejects them. A problem is considered verifiable if Arthur accepts a valid proof (e.g. the answer is yes and Merlin claims it is yes) in more than 2/3 of the cases and accepts an invalid proof (e.g. the answer is yes and Merlin claims no) in no more than 1/3 of the cases. In this case, the king is said to be *convinced* with some probability higher than 2/3,

making him a so-called *probabilistic polynomial-time* verifier. This interactive verification model was proposed by Babai in Ref. [174] and named the *Arthur-Merlin* (AM)-protocol.

Let us think of the smallest protocol instance involving a single message from Merlin to Arthur, the so-called *Merlin-Arthur* (MA)-protocol. Examples of verifiable proofs with MA can be found in number theory, e.g. factoring large integers, which is not known to be feasible in polynomial runtime [175]. In a decision version of factoring Arthur might ask Merlin whether the integer x has a factor y besides 1 that satisfies the condition $1 < y \leq k$ for some constant k [176]. Computer scientists and mathematicians formally refer to this decision problem as a *language* $L = \{\langle x, k \rangle : x \text{ has a factor less than } k \text{ besides } 1\}$. Crucially, both the “yes” and “no” outcomes for this language L , formally denoted by $x \in L$ or $x \notin L$, can be verified deterministically in polynomial time. This works in the following way. If Merlin answers “yes”, Arthur requests such a factor y from the wizard and verifies the claim by performing the multiplication $x = y \cdot (x/y)$ in polynomial time proving that $x \in L$. If Merlin’s answer is “no”, Arthur requests a prime factor greater than k and tests its primality using, for example, the *Agrawal–Kayal–Saxena* (AKS) protocol [177], which runs in polynomial time. If AKS succeeds, Arthur can calculate the remaining prime factors one-by-one and finally multiplies them to x and show that $x \notin L$.

While this verification example terminates with a success probability of 1 it would by definition suffice to have a success probability of larger than or equal to $2/3$. The latter condition allows more problems to be verified with MA, which will become clear in the next section.

Let us note that decision problems are very important in the field of verification, since any functional problem can be transformed into a decision problem while preserving the time or space needed for the computation [178]. The factoring example gives an idea of how this is done, so that Arthur gains a leverage over Merlin through the interactive situation. In this respect, the MA-protocol is archetypal for a broad class of verifiable classical computations, which we identify in the upcoming Sec. 4.1.2 based on their time or space required for the computation.

4.1.2 Complexity classes

The resource-based complexity of a computational task, i.e. the way in which its time and space requirements grow with the size of the problem, is denoted by so-called *complexity classes*, an important scientific topic in both mathematics and computer science [178]. While the time complexity indicates the number of steps needed to reach the solution of a given task, the space complexity characterizes the respective memory consumption. The assignment of problems to a complexity class can be proved using abstract computational models, typically based on Turing machines [53]. Some complexity classes overlap with each other, while others are entirely contained in a parent class. In this regard, Fig. 4.1(a) shows an inclusion diagram of the known relationships between the complexity classes relevant to this chapter. The interested reader will find further insights in Ref. [179], on which the following explanations are based on.

We encountered a complexity class without knowing it when we defined a deterministic polynomial-time verifier for the MA-protocol. The respective definition suggests the set of decision problems that can be solved on a deterministic Turing machine, where the time required increases at most polynomially with the size of the input, and denotes the complexity class *polynomial time* (P). Most problems of class P are called efficiently solvable or *tractable*. P represents the starting point and centers the inclusion diagram in Fig. 4.1(a).




Figure 4.1: **Clever delegation of messages allows a computationally weak verifier to verify the output of a much stronger prover and also enables the definition of complexity classes.** (a) Inclusion diagram of the time- or space-based computational resources (*complexity class*) relevant for the verification procedures discussed in this chapter. The diagram shows the relationships between the given classes, which are explained in the text of Sec. 4.1.2. While certain complexity classes are completely contained in a *parent class*, some classes only overlap with each other. The latter is illustrated by the red cross, where a problem in NP is also contained in *bounded-error quantum polynomial time* (BQP), but not in *bounded-error probabilistic polynomial time* (BPP), so that they can be solved with a quantum computer but not with a classical computer. Further explanations of the complexity classes can be found in Ref. [179]. (b) Several rounds of message exchange between a verifier (V) and a computationally stronger prover (P) on the decision task or language L, following notations in Ref. [173]. For a given input x , V attempts to confirm that x agrees with the “yes”-outcome of the decision task L, formally labeled $x \in L$. In this regard, V can make additional queries using its ability to randomly toss coins r , which are kept private, and sends the queries to P. P computes a certificate y_1 that depends on $\{x, r_1\}$ as a proof of its claim and sends it back to V. V evaluates the proof in polynomial runtime, formally denoted by $z_1 = V(x, r_1, y_1)$, whereupon this message exchange continues round by round. V accepts if all tests succeed, where formally at end of the last round $V(x, r_1, z_1, y_1, \dots, y_k) = 1$ holds, and rejects otherwise. Such an *interactive proof* is given by the graph non-isomorphism problem, which is discussed at the bottom of Sec. 4.1.3.

Consider again the set of decision problems solvable in polynomial runtime, now based on a Turing machine that can perform multiple actions in a single computational step, called *non-deterministic Turing machine*. Problems defined in this way range from solvable in polynomial to exponential runtime and denote the class *non-deterministic polynomial time* (NP). Consequently, NP is a parent class of P, as shown in Fig. 4.1(a). It is assumed that factoring a large number into prime factors belongs to NP [175]. As shown in the previous section, factoring can be verified by multiplying prime factors, an efficient task in P. Due to many similar examples, NP is alternatively defined by the set of decision problems that can be verified on a deterministic Turing machine in P.

Let us extend the class P once more by modeling a Turing machine that guarantees the right answer to a decision problem at least $2/3$ of the time, called *probabilistic Turing machine*. This defines the class of *bounded-error probabilistic polynomial time* (BPP) problems. The computational outcomes of quantum computers are also probabilistic in nature. Analogously, if we define the set of decision problems that can be solved on a quantum computer in polynomial runtime with error probability less than $1/3$, we obtain the class

of *bounded-error quantum polynomial time* (BQP). BPP is thus a subclass of BQP. Factoring, for instance, is solvable in BQP using Shor’s algorithm [9].

The computational modeling behind the definition of complexity classes often involves verification procedures, as in the example of NP. Indeed, the MA-protocol in which Arthur requests evidence about the claim Merlin made from a single communication step is no exception. According to the probabilistic nature of the protocol, Arthur belongs to the class BPP, while it is sufficient to consider Merlin in NP [173]. Thus, a problem $L \in \text{MA}$ is given whenever Arthur accepts a valid proof in more than $2/3$ of the cases and accepts an invalid proof in no more than $1/3$ of the cases, denoting the complexity class *Merlin-Arthur* (MA). Finally, if we extend the computational resources by introducing a quantum verifier bounded by polynomial runtime and a quantum prover, we obtain the class *quantum Merlin-Arthur* (QMA).

Switching from time complexity to memory complexity, the set of all decision problems that can be solved by a Turing machine with a polynomial amount of memory is called *polynomial space* (PSPACE). Note that the space complexity remains the same whether we consider a deterministic or non-deterministic Turing machine [180]. PSPACE is parent to all classes discussed so far, see Fig. 4.1(a).

In conclusion, let us note that the exemplary MA-protocol describes an interactive play between Merlin and Arthur, in which the communication forces Merlin to present comprehensible evidence for his claim, giving the king leverage. Analogous instances are used in computational models to define complexity classes [174] and are also the subject to many verification protocols [178, 181]. For example, blind verification of both classical and quantum devices [62] and the approach to classical verification of quantum computation discussed here [63]. In what follows, we generalize such interactive proof models.

4.1.3 Interactive proof system

Suppose a verifier (V) and the untrusted but computationally stronger prover (P) exchange messages about a decision problem. In these messages, P tries to convince V that the answer to the decision problem is indeed “yes” [178, 181]. Even if V is unable to verify the entire statement at once, the conversation forces P to present pieces of comprehensible evidence, and V continues to demand such evidence from P until every attempt to cheat can be detected. This procedure is called *interactive proof system* [178]. The result of interactive proofs is that V can show that P knows the answer, while it cannot solve the problem itself.

The MA-protocol from Sec. 4.1 is an example of an interactive proof system involving a single message [173]. The example formally constitutes a shared input x that V tasks P to prove that it is in the language L. As a proof of its claim, P generates a certificate y that the verifier accepts if $V(x, y) = 1$ holds. Importantly, V must be able to deterministically proof the certificate y by computing an algorithm $V(x, y)$ in polynomial runtime. We remark that in this scenario the specific power of P is irrelevant and it can be assumed all-powerful.

Let us introduce additional rounds of interaction, where we again follow the notation in Ref. [173] and assume all variables to be strings of bits. This ends with several responses y_1, y_2, \dots, y_k from P interleaved with several messages $z_1 = V(x, y_1), z_2 = V(x, y_2), \dots, z_k = V(x, y_k)$ from V to repeatedly gather evidence. V accepts the certificates at the final round k if and only if $V(x, y_1, z_1, y_2, z_2, \dots, y_l) = 1$. Crucially, in such a multi-round proof the all-powerful P can calculate all of V’s responses in advance and thus provide the string y_1, y_2, \dots, y_k in one go. However, this leads again to the definition of the complexity class NP, so we have not yet gained much.

To raise the complexity of problems that can be interactively verified, we additionally provide V with the ability to generate the messages based on random coin tosses. This takes away the possibility for P to calculate everything in advance, which provides V additional leverage, illustrated by Fig. 4.1(b). Let $r_i \in \{0, 1\}$ be the i -th string that Arthur has chosen uniformly at random, and let $(z_1, r_1, y_1, \dots, y_k)$ be the string of communicated sequences between Arthur and Merlin depending on the public input x . We further allow a linear number of rounds $|x|^{O(1)}$ with respect to the length of the input string $|x|$. For the interactive proof to succeed, the resulting probabilities $\Pr[V(x, z_1, r_1, y_1, \dots, y_k) = 1]$ at the end of the k -th round must be greater than or equal to $1 - \epsilon$ if $x \in L$ and less than or equal to ϵ if $x \notin L$ [173]. Let us express this with the equation below

$$\begin{aligned} \Pr_{r \in \{0,1\}^{|x|^{O(1)}}} \left[(P, V(x, z_1, r_1, y_1, \dots, y_k) = 1) \right] &\geq 1 - \epsilon, && \text{if } x \in L \\ \Pr_{r \in \{0,1\}^{|x|^{O(1)}}} \left[(P^*, V(x, z_1, r_1, y_1, \dots, y_k) = 1) \right] &\leq \epsilon, && \text{if } x \notin L \text{ for all provers } P^* \end{aligned} \quad (4.1)$$

with some $\epsilon < 1/2$ and all provers P^* which potentially cheat by altering the protocol. Interactive proof systems satisfy the following two conditions [178] separating the two lines in Eq. (4.1). First, for every correct result, P is able to convince V of its correctness, which is called *completeness*. Second, if the outcome is false, P has no way to convince V of its correctness. This non-existence of a proof when the output is false is called *soundness*.

Due to the probabilistic nature of Eq. (4.1), V 's acceptance is BPP-like, and V is said to be convinced. The complexity class of tasks verifiable in an interactive proof system depends on both the computational boundaries and the capabilities given to V , e.g., additional randomness and security assumptions [178]. Using an interactive proof based on a probabilistic verifier in P and *public coins*, where P could calculate all possible choices in advance, we have previously defined the complexity class MA. Crucially, we can extend the class MA by switching from public to *private coins* and allowing a polynomial number of rounds in the input size $|x|$, yielding the complexity class *interactive proof* (IP). Interestingly, Ref. [181] shows that IP is equal to PSPACE and contains all decision problems that can be solved by a (non-)deterministic Turing machine in polynomial space.

To illustrate an interactive proof, consider two finite graphs \mathcal{G}_0 and \mathcal{G}_1 , each represented by n vertices. If the vertices of either graph can be permuted so that both become equal, we say that the graphs are *isomorphic*. Let us do the opposite and ask whether there is no such permutation $\sigma \in S_n$ for the two graphs [173], yielding the language

$$L = \left\{ \langle \mathcal{G}_0, \mathcal{G}_1 \rangle \mid \mathcal{G}_0 \text{ and } \mathcal{G}_1 \text{ are encodings of graphs and } \forall \sigma \in S_n, \sigma(\mathcal{G}_0) \neq \mathcal{G}_1 \right\}, \quad (4.2)$$

which are then called *non-isomorphic*. Let us construct an interactive proof of this task by a Boolean function f_L that outputs $f_L(\langle \mathcal{G}_0, \mathcal{G}_1 \rangle) = 1$, if and only if the graphs are non-isomorphic, and $f_L(\langle \mathcal{G}_0, \mathcal{G}_1 \rangle) = 0$ otherwise, denoting the so-called *graph non-isomorphism problem* [182]. Since we are allowed to guess which permutation to use, the problem lies in NP and can be verified by definition in P , as described below. The upcoming verification protocol aims to convince V of the output of P with high probability, while allowing both access to \mathcal{G}_0 and \mathcal{G}_1 [173]. In this sense:

- (1) V randomly chooses $c \in \{0, 1\}$
- (2) V randomly chooses a permutation $\sigma \in S_n$ and sends $\sigma(\mathcal{G}_c)$ to P
- (3) P uses its infinite computational power and decides which of the two graphs the given σ is isomorphic to. P responds by sending \mathcal{G}_b for $b \in \{0, 1\}$

(4) V accepts if and only if $b = c$ which can be done effortlessly.

For $\langle \mathcal{G}_0, \mathcal{G}_1 \rangle \in L$ which are non-isomorphic, any permutation of \mathcal{G}_b determines b such that P knows b with certainty and returns it. For $\langle \mathcal{G}_0, \mathcal{G}_1 \rangle \notin L$ isomorphic, on the contrary, the graph sent by V has the same distribution for $c = 0$ and $c = 1$, where P can only guess which one to choose. As such, V 's probability to accept is also bounded by $1/2$. This two-round protocol containing one message from both parties can therefore not convince V . However, we can allow V to send k copies $\mathcal{G}^1, \dots, \mathcal{G}^k$ in parallel to P , where each \mathcal{G}^i is an independent random selection according to the steps 1 and 2 above. Now V only accepts if P answers correctly in all k cases, which results in a verification probability of $1 - 2^{-k}$, whereby the ability of P to cheat converges with 2^{-k} towards 0.

Besides the graph non-isomorphism, factoring large integers lies outside the class P . As shown in Sec. 4.1.1, the verification of prime numbers can be done in a single round of interaction, whereas the interesting thing about factoring is the underlying functionality. While the multiplication of prime factors is computationally efficient, the inverse of the function, i.e., calculating the prime factors, is computationally hard on classical computers [175]. This prominently builds the security foundation of classical cryptography [183]. In cryptography, the prime factors represent the *private key* that is exchanged between some users and gives them access to their shared secret. An eavesdropper's access is restricted to the large integer that is the product of the prime numbers, which is called the *public key*. The eavesdropper's only chance of cracking the code is to laboriously compute the primes, which in a secure protocol, that usually relies on a 2048-bit number, takes longer than the key is actually valid [184]. Yet, knowing a secret, e.g. one of the prime factors, the problem becomes simple. Functions that are defined to be efficient to compute but hard to invert unless someone knows a secret (private key) establish security in many cryptography protocols and are called *trapdoor functions*. In the following, we will explain the significance of trapdoor functions in verification applications.

4.1.4 Interactive proofs with trapdoor functions

The interactive proofs in both AM-protocol and graph non-isomorphism create leverage over a computationally stronger prover by requesting comprehensible pieces of evidence based through interaction or the verifiers's ability to make random choices. Below, we illustrate how trapdoor functions can be used to increase this leverage over the prover.

Consider a computationally weak verifier (V) tasking a much more powerful prover (P) with counting the number of leafs in a tree, as illustrated in Fig. 4.2. Checking the result by simply recounting the number of leafs is computationally far beyond the capabilities of V . Instead, V repeats the query but removes a random number of leafs from the tree beforehand. Importantly, V keeps this number secret from P , who is thus forced to repeat the task entirely. Counting the difference in leafs, however, falls within the computational capabilities of V and provides strong leverage over P , denoting the *trapdoor information*. Many such iterations, where V removes leafs from the tree and P recounts them, will eventually convince V that P is able to correctly solve the task of counting the leafs on the tree.

The example intuitively shows how the difference in the number of leafs provides V comprehensible pieces of evidence and how choosing this difference at random and keeping it secret creates a trapdoor function which cleverly leverages P .

Finally, if we allow the computational entities to be based on quantum computers we end up with a *quantum interactive proof system* [178]. Such a system considers P with unlimited




Figure 4.2: **An interactive proof system based on a trapdoor function, simply illustrated.** Imagine a verifier (V) with limited computational power who wants to solve a task that far exceeds its computational capabilities. Let this task be counting the number of leaves on a tree. A problem where a much more powerful prover (P) steps in and helps. However, P’s trustworthiness is doubted. To assess its honesty, V repeats the query but before removes a random number of leaves from the tree. A number, P is left in the dark about, while it can be counted effortlessly by V. This leaves P no other chance than to completely resolve the problem. By checking the difference in the number of leaves in many iterations, V becomes convinced of the correctness of P’s original answer. The essence of this example is that V gains leverage over P in that it keeps the random number of leaves removed to itself, which is called the *trapdoor information*.

computational power and a quantum V that solves tasks in polynomial runtime based on success probability of at least 2/3 and an error probability of at most 1/3. As such, V is in BQP. This definition denotes the complexity class of *quantum interactive proof* (QIP). Note that IP is similarly defined with a BPP-like verifier. Interestingly, the work of Watrous in Ref. [185] and of Aaronson in Ref. [186], both published in 2010, showed that QIP is equal to PSPACE, which was considered a breakthrough in complexity theory.

4.1.5 Learning with errors

For the above protocols to work, we need to ensure that the encryption in the form of trapdoor functions is secure against attempts by the powerful prover to reverse engineer them. Ref. [187] summarizes a hodgepodge of cryptography and security applications such as verifiable computations that fall under the umbrella of interactive proof systems. However, established ways of distributing keys in classical cryptography are threatened by quantum computers using Shor’s algorithm [9]. At least someday, when scaled-up quantum computers work in a fault-tolerant way this scenario may become a reality. To compensate for this future lack of security, computational tasks have been investigated that are hard for quantum computers to solve and that allow key distribution even in the *post-quantum era* [188].

A promising idea to encrypt messages in cryptography or interactive proofs is based on a set of linear equations with erroneous variables. With only the function outputs publicly available, it would be the task of an eavesdropper to learn from which function domains

the outputs were generated, becoming challenging for high-dimensional equations that are subject to noise. The idea was proposed by Regev in Ref. [189] and is called *learning with errors* (LWE).

Conceptually, LWE starts from \mathbb{Z}_q^m as the m -dimensional ring of integers modulo q [190] and $\mathbb{Z}_q^{m \times n}$ as the $m \times n$ -dimensional ring of integers modulo q . Now imagine a function $f : v = M \cdot t + e$ that transforms an input matrix $M \in \mathbb{Z}_q^{m \times n}$ into an output matrix $v \in \mathbb{Z}_q^m$ depending on a private key $t \in \mathbb{Z}_q^n$ and the influence of noise $e \in \mathbb{Z}_q^m$. Thus, a message of length m can be encrypted with the public key $\{v, M\}$ and the private key $\{t, e\}$.

Let us define the LWE procedure in more detail, and then illustrate how it can be used to encrypt messages:

- (1) choose a fixed vector $t \in \mathbb{Z}_q^n$ uniformly at random
- (2) choose $M \in \mathbb{Z}_q^{m \times n}$ at random from a uniform distribution over $\mathbb{Z}_q^{m \times n}$
- (3) sample an error $e \in \mathbb{Z}_q^m$ from Gaussian distributions in all m dimensions truncated by the norm $\|e\|_1$, which is small compared to the norm of the function output $\|v\|$
- (4) efficiently compute $f : v = M \cdot t + e$ yielding $v \in \mathbb{Z}_q^m$
- (5) distribute $\{v, M\}$ as the public key and keep $\{t, e\}$ secret as the private key

Let us run through this recipe in a one-dimensional example of sending message $\text{msg} = 1$ based on private key $t = 5$ with error $e = 12$. To generate the public key, we choose a random sequence of numbers $M = \{12, 34, 6, 2, 45, 9, 16, 12, 31, 53\}$ longer than the value of t and calculate $f : v = M \cdot t + e$, which gives $v = \{72, 182, 42, 22, 237, 57, 92, 72, 167, 277\}$. Then we encrypt the message msg from a randomly sampled subset $v_{\text{sub}} = \{92, 167, 42, 72, 22\}$ of length t , accumulate its values, and add the message msg on top, giving $\text{msg}_{\text{enc}} = \sum v_{\text{sub}} + \text{msg} = 396$. This encrypted message $\text{msg}_{\text{enc}} = 396$ can be used for secure communication. A receiver knowing the private key can decrypt the message msg by computing the remainder $\text{msg} = \text{msg}_{\text{enc}} \pmod{t} = 1$.

An eavesdropper, on the other hand, can only learn about the message msg by finding out the private key (or trapdoor) t based on the public key $\{M, v\}$. This is only possible by inverting the function f . While the inversion is quite straightforward for the given number strings, the task becomes computationally difficult with the Gaussian distributed error and for high dimensions $m \gg n$, where m and n are the dimensions of the function output and the trapdoor respectively [189].

Geometrically, full-fledged LWE refers to the identification of noisy sites on a high-dimensional lattice where each lattice site is blurred by the error e . Without the error e , the lattice sites become discrete and the task simplifies to a linear equation. However, in the presence of errors e , the only chance to find the trapdoor t is to check all lattice sites, which becomes computationally hard for high dimensions $m \gg n$, more details in Sec. 4.2.2.

Interestingly, Regev proved in Ref. [189] that LWE is hard even for quantum computers [189], making the problem a so-called *post-quantum secure* trapdoor function. LWE is widely used in public key cryptosystems that feature applications in both the current era of classical and the future era of quantum computing, which is discussed in detail in the review work of Ref. [191].

4.2 VERIFIABLE QUANTUM COMPUTATION BY CLASSICAL MEANS

Mahadev's framework for classical verification of quantum computations in Ref. [63] uses an interactive type of measurement protocol that allows a classical verifier to use a

quantum computer that no longer has the ability to cheat. The interactive proof (Sec. 4.1.3) is based on the ability of the classical verifier to make random choices and the inability of the quantum computer to efficiently solve the LWE problem (Sec. 4.1.5). This enables encrypted interactions between verifier and prover. The existence of such a framework represents a milestone in computer science, and testing its experimental capabilities is the next logical step.

Following Mahadev’s protocol, recent work on non-interactive classical verification [192] and an approach to polynomial-time verifiers with zero knowledge [193] have been proposed. Moreover, the authors in Ref. [194] present a protocol on *classical verifiable quantum advantage*.

Here, we attempt to verify the output of a quantum computer classically. Mahadev [63] assumes that the classical verifier (V) is able to solve all decision problems that require a polynomial runtime on a probabilistic Turing machine, i.e. acts in the class of BPP. The quantum prover (P), on the other hand, can solve decision problems from the class BQP, to which it efficiently provides an answer in polynomial runtime. BQP lies within PSPACE, i.e., the class of (quantum) interactive proofs [185, 186]. Crucially, BQP contains problems outside of NP, as illustrated by the red cross in Fig. 4.2(a). An example is simulations of quantum many-body systems, whose countless degrees of freedom can become classically intractable for as few as ten particles [54]. While the latter remain efficiently solvable for P, for V it is no longer clear how they can be verified with the tools presented so far in Sec. 4.1 and therefore represent the relevant problems for the classical verification of quantum computation.

Verifying the output of a quantum computer relates to verifying a quantum state, where the verifier must decide whether or not that state has been correctly prepared by a quantum computer prover, referring to a decision task with “yes”- or “no”-outcome. Since any functional problem can be transformed into a decision task with the same computational complexity [195], this reformulation does not make the verification process more difficult. Mahadev proposes in Ref. [63] to rephrase this decision task as an energy measurement of a specially designed quantum state that indicates the correct computational outcome of the task by a low energy.

4.2.1 Phrasing a decision problem as an energy measurement

Suppose the verifier (V) wants to confirm a decision task $L \in \text{BQP}$ given by the computational outcome of applying an m -qubit quantum circuit \mathcal{C} to the input state $|0\rangle^{\otimes m}$. Analogous to the verifiable classical computations from Sec. 4.1, there is a similar setting for tasks or languages $L \in \text{BQP}$ where the input x is a local Hamiltonian along its ground-state $|\eta\rangle$ and the certificate y of the prover (P) is the energy of this quantum state [63]. Thus, if $x \in L$ there exists a state $|\eta\rangle$ with energy below a certain threshold, while for $x \notin L$ the energy of all states remains above this threshold.

Crucially, there exists a quantum computer approach that is designed so that its problem solution is encoded in the ground-state of a particular Hamiltonian, to which the system then evolves adiabatically, referred to as *adiabatic quantum computers* [196]. In classical verification, where we similarly attempt to encode the solution of a problem $L \in \text{BQP}$ into the ground-state energy of a Hamiltonian, we rely largely on these techniques. Ref. [197] provides the recipe for the efficient construction of such a Hamiltonian $\mathcal{H}(L)$, which consists of $\text{poly}(m)$ terms, each acting on no more than $\log m$ qubits, yielding a so-called $(\log m)$ -local Hamiltonian. The work in Ref. [198] further shows that these Hamiltonians can be constructed from only two-local terms of the Pauli operators X and Z [198]. Hamil-

tonians of this type are considered to be efficiently computable, a prerequisite for our verification studies. An upper energy threshold for the corresponding ground-state can also be estimated using the tools from Ref. [197].

Let us assume, without loss of generality, that P claims the probability \Pr that it has correctly implemented the quantum circuit \mathcal{C} is $\Pr(\mathcal{C}) \geq 1 - \epsilon$, which means that $x \in L$. We refer to this as the “yes”-outcome of the decision problem. One can then show that a Hamiltonian $\mathcal{H}(L)$ constructed based on L and the above means has the following two properties [197]:

- (1) If $\Pr(\mathcal{C}) \geq 1 - \epsilon$, where the prover claims $x \in L$ and this is the case, the smallest eigenvalue λ of $\mathcal{H}(L)$ is below ϵ
- (2) If $\Pr(\mathcal{C}) \leq \epsilon$, where the prover claims $x \in L$ but this is not the case, the smallest eigenvalue of $\mathcal{H}(L)$ is larger than $\frac{3}{4} - \epsilon$.

These properties refer to completeness and soundness, see Sec. 4.1.3. If completeness is satisfied, there exists a state $|\eta\rangle$ with energy less than ϵ , which can be efficiently computed in BQP and is called *clock-state* for reasons that will become clear in a moment [199]. An honest quantum P can prepare the $|\eta\rangle$ -state as a proof of its claim.

We leave further details on the construction of such a problem Hamiltonian $H(L)$ and its lower bound on the ground-state energy ϵ to Refs. [63, 200] and focus instead on the $|\eta\rangle$ -state preparation and the interactive energy measurement on it, which are the key building blocks of Mahadev’s protocol [63].

In this sense, Refs. [63, 200] show that classical V knows how to write down the clock-state $|\eta\rangle$ based on a circuit \mathcal{C} linked to the quantum task L in the following systematic way

- Let U_1, \dots, U_N be the N gates of the m -qubit circuit \mathcal{C} acting on the initial state $|0\rangle^{\otimes m} = |0, \dots, 0\rangle$
- There exists a Hamiltonian $\mathcal{H}(L)$ in Hilbert space of $n_\eta = m + \lceil \log(N + 1) \rceil$ qubits, such that its energy $\langle \eta | \mathcal{H} | \eta \rangle = 1 - p(\mathcal{C})$, with [63, 200]

$$|\eta\rangle = \frac{1}{\sqrt{N+1}} \sum_{n=0}^N U_n \dots U_1 |0 \dots 0\rangle |n\rangle, \quad (4.3)$$

where $U_0 = \mathbb{1}^{\otimes m}$ is the identity. The $|\eta\rangle$ -state is a superposition of all N gates from the circuit \mathcal{C} , accompanied by a second register $|n\rangle$, called *clock-register*. A circuit with N gates requires at most $\lceil \log(N + 1) \rceil$ clock-qubits [197]. For example, consider the $|\eta\rangle$ -state of a three-gate circuit with single-qubit gates U_1, U_2 , and a two-qubit gate U_3 acting on the initial state $|00\rangle$ of output $U_3 U_2 U_1 |00\rangle$. The corresponding clock-state is $|\eta\rangle = 1/2(|00\rangle |0\rangle + U_1 |00\rangle |1\rangle + U_2 U_1 |00\rangle |2\rangle + U_3 U_2 U_1 |00\rangle |3\rangle)$ involving four qubits

- The Hamiltonian $\mathcal{H}(L)$ is a sum of terms of only local operators X and Z [198], and if λ denotes the smallest eigenvalue of $\mathcal{H}(L)$, the following implications hold:
 - (1) $p(\mathcal{C}) > 1 - \epsilon \implies \lambda < \epsilon$
 - (2) $p(\mathcal{C}) < \epsilon \implies \lambda > \frac{3}{4} - \epsilon$.

Suppose for a moment that P sends the qubits of the $|\eta\rangle$ -state one-by-one to V, which is now assumed to have quantum resources. The problem Hamiltonian $\mathcal{H}(L)$ consists of local Z and X terms, so by analogy, local measurements in Z - and X -basis [198, 201]

qualify quantum V to decide whether $p(\mathcal{C}) > 1 - \epsilon$ or $p(\mathcal{C}) < \epsilon$ and thus complete the verification procedure. Since such local energy measurements can be applied to all qubits of the $|\eta\rangle$ -state sequentially, the verification procedure becomes linear in the number of qubits involved [201].

In contrast, if V only has access to classical resources, the quantum measurements must be delegated to quantum P .

The conceptual core of Mahadev's protocol [63] is the delegation of what used to be the quantum aspect of the verifier's tasks to the prover through a series of encrypted measurements that ensure that the prover cannot use the fact that they are measuring the energy themselves to cheat.

4.2.2 Post-quantum secure delegation of energy measurements

Although quantum computers are very powerful, there exist problems that are hard to solve even for them. Analogous to classical cryptography, one such quantum-hard problem is LWE [189], presented in Sec. 4.1.5. In classical verification of quantum computation [63], it is precisely the assumption that LWE is quantum-computationally hard that allows for the construction of post-quantum secure trapdoor functions to delegate the energy measurement from the classical verifier (V) to the quantum prover (P) in an encrypted way.

In particular, to hide whether the measurements on the $|\eta\rangle$ -state from Eq. (4.3) are performed in the basis X or Z , we need to encode two families of functions into a post-quantum secure trapdoor function. These function families are supposed to perform transformations on the quantum state held by P such that they effectively refer to a measurement in either of the two bases. Formally, let \mathcal{F} and \mathcal{G} be the two families of functions along with a pair of elements $\{f_{k,0}, f_{k,1}\} \in \mathcal{F}$ and $\{g_{k,0}, g_{k,1}\} \in \mathcal{G}$, where $k \in K$ is an index based on a finite set K . While the index k secretly communicates the function type and thus serves as a public key, the function elements $\{0, 1\}$ depend on a single qubit state $\{|0\rangle, |1\rangle\}$ from the $|\eta\rangle$ -state, on which the particular measurement is performed. The elements of the two function families are considered as bit string transformations of the form $f_{k,b}, g_{k,b} : \{0, 1\}^{n_x} \rightarrow \{0, 1\}^{n_y}$. The encryption of \mathcal{F} and \mathcal{G} with post-quantum secure trapdoor functions should make it impossible for P to decide whether a certain function $y_k \in \mathcal{F} \cup \mathcal{G}$ originates from \mathcal{F} or \mathcal{G} and thus to know on which basis the measurements are actually carried out. We formally write y_k to emphasize that the function family is hidden. To this extent, we consider the two families to have a private key or trapdoor t_k , i.e. an inverse function y_k^{-1} that can output x_b given an input y such that $f_{k,b}(x_b) = y$ or $g_{k,b}(x_b) = y$ with $b \in \{0, 1\}$. We refer to the image x_b of the inverse function y_k^{-1} as the *preimage*. The trapdoor information t_k is only known to V , while it is kept secret from P . Ideally, P can compute the function outputs of y_k in either case $f_{k,b}(x_b) = y$ or $g_{k,b}(x_b) = y$ if it only knows the public key k without any further information about the function's character.

The linearity of the energy measurement protocol allows us to examine a single qubit $\sum_{b \in \{0,1\}} \alpha_b |b\rangle = \alpha_0 |0\rangle + \alpha_1 |1\rangle$ of the $|\eta\rangle$ -state, see Sec. 4.2.1. To encrypt the measurements, Ref. [63] considers the following state

$$|\phi_k\rangle = \frac{1}{\sqrt{2^{n_x}}} \sum_{b,x} \alpha_b |b\rangle |x\rangle |y_k(b,x)\rangle, \quad (4.4)$$

which V tasks P to prepare. This state involves $1 + n_x + n_y$ qubits with $b \in \{0, 1\}$, $x \in \{0, 1\}^{n_x}$ and $y \in \{0, 1\}^{n_y}$, where we call b the state of the *committed qubit*, x the *preimage register* and

y the *commitment string*. By these means, $n_x + n_y$ additional qubits are used to delegate and perform encrypted measurements on the single committed qubit b from the $|\eta\rangle$ -state. We leave the function sizes n_x and n_y unspecified for now, but discuss examples in Sec. 4.2.4.

We can construct a measurement protocol based on Eq. (4.4) that effectively corresponds to a Z or X measurement on the first qubit b , depending on whether $y_k(b, x) \in \mathcal{G}$ or $y_k(b, x) \in \mathcal{F}$. First, P must measure the last register $|y_k(b, x)\rangle$ in the “physical” Z-basis and pass the resulting bit string to V. By construction, we want the remaining state in the first two registers to project into the product state $\alpha_b |b\rangle |x\rangle$ for an effective Z-measurement (\mathcal{G}) or into the superposition state $\alpha_0 |0\rangle |x\rangle + \alpha_1 |1\rangle |x\rangle$ for an effective X-measurement (\mathcal{F}). P is finally asked to measure the qubits of the first two registers in the “physical” X-basis, which effectively performs either the Z- or the X-basis measurement on the first qubit b , depending on the chosen function family \mathcal{G} or \mathcal{F} .

This projection into either the product or superposition state can be achieved by defining the two function families to be pairs of maps $\{g_{k,0}, g_{k,1}\} \in \mathcal{G}$ and $\{f_{k,0}, f_{k,1}\} \in \mathcal{F}$ being of type one-to-one and two-to-one, respectively. This means that pairs of the first family \mathcal{G} have both disjoint domains and images in any case (injective), whereas pairs from the second family \mathcal{F} can have the same image for disjoint domains (surjective). For domains $x_0 \neq x_1$ it then always holds $g_{k,0}(x_0) \neq g_{k,1}(x_1)$, whereas for domains $x_0 \neq x_1$ with $x_1 \neq x_0 + t_k$ it holds $f_{k,0}(x_0) \neq f_{k,1}(x_1)$. Yet, if someone knows the trapdoor t_k , defined by $x_0 \neq x_1$ with $x_1 = x_0 + t_k$, one can find $f_{k,0}(x_0 + t_k) = f_{k,1}(x_1)$. To separate the two function families, preimages of \mathcal{F} that generate the same image must be identified. By design, this must be computationally efficient with knowledge of t_k and either of $\{x_0, x_1\}$ and quantum computationally hard otherwise. The latter requires calculating the inverse function $f_{k,b}^{-1}$ to find the preimages $x_b = f_{k,b}^{-1}(y)$.

We now qualitatively review how \mathcal{F} and \mathcal{G} can be encoded into a post-quantum secure trapdoor function. Let us therefore relate the two-to-one function family \mathcal{F} to the quantum-hard LWE problem from Sec. 4.1.5, which must be distinguished from the one-to-one function family \mathcal{G} . There are no special requirements for the elements of \mathcal{G} , so they may be efficiently invertible. The distinction of \mathcal{F} and \mathcal{G} leads us to a binary version of LWE called *decision learning with errors* (DLWE) [189]. However, LWE is part of DLWE, so we can seamlessly follow up on our earlier definition presented in Sec. 4.1.5.

Consider $M \in \mathbb{Z}_q^{m \times n}$ and $t_k \in \mathbb{Z}_q^n$, both chosen uniformly at random, together with an error $e \in \mathbb{Z}_q^m$. The error e is drawn at random from a product of truncated Gaussian distributions with a norm $\|e\|_1$ that is considered small compared to the norm of the function output $v \in \mathbb{Z}_q^m$ [189]. In DLWE, P is given the function output v alongside M , denoting the public key, while V, who constructed the problem, holds back both trapdoor information t_k and error e . Based on the public key $\{v, M\}$, the decision task for P is to distinguish, whether the function output v is constructed from $f_k : M \cdot t_k + e$ or a uniformly distributed $g : u \in \mathbb{Z}_q^m$ as follows

$$v = \begin{cases} f_k : M \cdot t_k + e \\ g_k : u \end{cases} . \quad (4.5)$$

Geometrically, DLWE refers to the identification of points on a lattice that are either uniformly distributed $g_k : v = u$ or follow a noisy pattern according to $f_k : v = M \cdot t_k + e$. In the latter case, the function output v lies within a certain radius from a lattice site drawn from the error e , as shown in Fig. 4.3. Crucially, P must learn of the trapdoor information t_k to identify the function type chosen by V, which is only possible by inverting $f : v = M \cdot t_k + e$.

With the given error, this leaves P no choice but to check all lattice sites, which becomes quantum computationally hard at high dimensions $m \gg n$ [189].




Figure 4.3: **Post-quantum secure trapdoor function based on noisy lattice sites.** Imagine two families of functions as pairs of maps $\{g_{k,0}, g_{k,1}\} \in \mathcal{G}$ and $\{f_{k,0}, f_{k,1}\} \in \mathcal{F}$, each of type one-to-one (injective) and two-to-one (surjective), respectively. The index $k \in K$ is based on a finite set K and represents the public key to communicate the chosen function type between verifier V and prover P . The family \mathcal{G} has both disjoint domains and images, i.e. $g_{k,0}(x_0) \neq g_{k,1}(x_1)$ for all $x_0 \neq x_1$, while the family \mathcal{F} can have the same image for two domains $f_{k,0}(x_0 + t_k) = f_{k,1}(x_1)$, depending on a private key or trapdoor t_k . P can compute the function image based on the k received from V and then tries to identify either function family \mathcal{F} or \mathcal{G} by checking whether the given image could have been generated from two preimages. Without knowing the trapdoor t_k , P must invert \mathcal{F} , which by design must be computationally hard. In the classical verification of quantum computations [63], the one-to-one and two-to-one maps are used to encrypt measurements in Z - and X -basis, see Eq. (4.4). V delegates these encrypted measurements to P and by that hides the respective basis. This can be achieved by encrypting \mathcal{F} and \mathcal{G} based on DLWE [189] defined in Eq. (4.5) relating to the identification of noisy lattice sites. Geometrically the shared image $f_{k,0}(x_0 + t_k) = f_{k,1}(x_1)$ lies somewhere in the erroneous area (red shaded). Due to the noise and without knowledge of the trapdoor t_k , P can only separate \mathcal{F} and \mathcal{G} by checking all lattice sites, which becomes quantum computationally hard for high-dimensional lattices [189].

For the measurement delegation in Eq. (4.4), the lattice problem is encoded in the bit string transformation $f_{k,b}, g_{k,b} : \{0, 1\}^{n_x} \rightarrow \{0, 1\}^{n_y}$ and requires $n_x + n_y$ extra ancilla qubits. For realistic security assumptions against a quantum computer, it is assumed that hundreds of additional qubits $n_x + n_y$ are needed to encrypt the energy measurement for each computational qubit of the $|\eta\rangle$ -state [54].

For a summary of the construction of post-quantum secure trapdoor functions, see Ref. [202]. In the upcoming discussion of the classical verification protocol, we relate to the trapdoor functions simply in terms of their bit string lengths. A minimal instance for two-to-one and one-to-one functions that encode X and Z measurements is thus given by a two-bit to two-bit transformation of form $f_{k,b}, g_{k,b} : \{0, 1\}^2 \rightarrow \{0, 1\}^2$, which however

cannot fulfill any security assumptions. For technical explanations on the construction of such relaxed trapdoor functions, see the Appendix of Ref. [200].

4.2.3 Concept of classical verification

The key to classical verification of quantum computation is to turn the verification process into an interactive situation, see Sec. 4.2.1. First, the verifier (V) asks for a commitment from the prover (P), and only after receiving it, requests the measurement results in one of two orthogonal bases, which must be consistent with the original commitment. Importantly, V delegates the measurement bases in an encrypted way by using post-quantum secure trapdoor functions based on DLWE, see Sec. 4.2.2. In this way, P executes the actual verification step, but still cannot cheat because they are kept in the dark about the basis on which the measurements are performed.

Given the quantum task L to be verified, V constructs the corresponding local problem Hamiltonian $\mathcal{H}(L)$ together with a recipe of the clock-state $|\eta\rangle$ in the systematic way of Eq. (4.3), rephrasing the decision problem as an energy measurement. This construction relates the $|\eta\rangle$ -state to the ground-state of $\mathcal{H}(L)$ with energy below a known threshold ϵ [197]. The resulting quantum circuit \mathcal{C} is passed to P, who efficiently prepares it with their quantum power. If done correctly, the state prepared by P has energy below ϵ and is a proof of the correctness of P's claim about the outcome of the decision problem. Crucially, the energy measurement requires quantum resources, so V must delegate the measurements to P. Refs. [198, 201] show that local measurements in X- and Z-basis are sufficient, whereby the measurements can be applied sequentially to each qubit in the $|\eta\rangle$ -state, yielding a linear protocol. While P knows quite a lot about the task to be verified, including the circuit \mathcal{C} to prepare $|\eta\rangle$, the main contribution of Mahadev's classical verification protocol in Ref. [63] is that V hides the measurement bases from P using post-quantum secure trapdoor functions based on DLWE from Eq. (4.5) and thereby prevents any cheating. As such, the given protocol constitutes a quantum interactive proof.

To construct the interactive proof, V tasks P to integrate the trapdoor function by encoding each qubit of the $|\eta\rangle$ -state with Eq. (4.4), resulting in the state $|\phi_k\rangle$. By construction, $|\phi_k\rangle$ requires a one-to-one function (\mathcal{G}) to effectively delegate a Z-measurement and a two-to-one function (\mathcal{F}) for the X-measurement, which are encrypted by means of the first and second case of DLWE in Eq. (4.5). P prepares the state $|\phi_k\rangle$ with the index k secretly communicating the function family (public key) using $n_x + n_y$ additional ancilla qubits that store preimage and image of the trapdoor function. P is then asked to measure the n_y qubits from the commitment register $|y_k(b, x)\rangle$ in the "physical" Z-basis. This projects the first two qubit registers from $|\phi_k\rangle$ into the product state $\alpha_b |b\rangle |x\rangle$ with $b \in \{0, 1\}$ for y_k being the one-to-one function (\mathcal{G}) or into the superposition state $\alpha_0 |0\rangle |x\rangle + \alpha_1 |1\rangle |x\rangle$ for y_k being the two-to-one function (\mathcal{F}). In the latter case two preimages have generated the same trapdoor function image y . P then measures the first two registers in the "physical" X-basis, which effectively performs a Z- or X- measurement depending on the selected function family \mathcal{G} or \mathcal{F} . P finally sends the classical measurement results to V. This part of the protocol must be applied to all qubits in the $|\eta\rangle$ -state and is referred to as "*measurement round*".

To ensure that P cannot cheat, we must additionally verify Eq. (4.4) with respect to the initial commitment, which tells us that P has prepared the correct quantum state. Otherwise, the energy may remain low even though the task was performed incorrectly [63]. To this extent, V interleaves the "*measurement round*" with a round that ensures that P has prepared the state correctly, referred to as "*test round*". This works in the following way.

After P has measured the n_y qubits from the commitment register of $|\phi_k\rangle$ of Eq. (4.4) in the Z-basis, V requests measurements on the remaining $1 + n_x$ qubits of the first two registers in Z-basis and checks whether the function output y obtained by P was actually generated from the selected input, i.e. checks whether $y_k(b, x) = y$ is true. If this test passes with high probability, V can trust the results from the “measurement round”.

Decryption of the measurement bases with the trapdoor information t_k (private key) and linking them to the classical results from the “measurement round” obtained from P finally qualifies V to estimate the energy of the $|\eta\rangle$ -state and to classically verify the quantum task L.

It should be noted that noise in the quantum system of P must, by design [63], increase the measured energy of the $|\eta\rangle$ -state and above a critical level prevents verification [200].

4.2.4 The step-by-step protocol

With current quantum devices, even in the absence of noise, Mahadev’s verification protocol [63] is not feasible for reasonable choices of security parameters, i.e., large enough trapdoor functions based on the DLWE problem from Eq. (4.5). In fact, a realistically secure version of the protocol requires hundreds of extra qubits for each computational qubit to verify [54]. In view of the fast-moving technological developments, we now present a proof-of-principle protocol on classical verification of quantum computation, circling around the original ideas from Ref. [63].

The protocol is presented in Fig. 4.4 and aims at demonstrating all the building blocks for classical verification, while being specifically designed for the capabilities of today’s NISQ-devices [54]. Such a small-scale approach must inherently relax security assumptions in favor of a feasible protocol.

Consider an $|\eta\rangle$ -state of unspecified size constructed by means of Eq. (4.3). The linear structure of the measurement protocol proposed by Ref. [63] allows the energy measurement to be performed separately on each qubit of the $|\eta\rangle$ -state. We therefore assume a single qubit in the arbitrary state $|\psi\rangle = \alpha_0 |0\rangle + \alpha_1 |1\rangle$, whereby the following protocol steps must eventually be applied to all qubits in the $|\eta\rangle$ -state. For the proof-of-principle approach, we use minimum-size trapdoor functions based on a two-bit map $y_k : \{0, 1\}^2 \rightarrow \{0, 1\}^2$, which for $k = 0$ communicates a one-to-one function (\mathcal{G}) to effectively perform a Z-measurement, and for $k = 1$ communicates a two-to-one function (\mathcal{F}) to effectively perform an X-measurement, see Eq. (4.4). A description of this state together with the index k denotes the public key. According to Eq. (4.4), the resulting minimal protocol instance can be partitioned into three qubit-registers, all of which are held by P. The single-qubit in the state $|\psi\rangle$ represents the committed qubit. To delegate the measurement bases $k \in \{0, 1\}$, an additional single qubit is required in the preimage register. To store the two-bit outputs of the trapdoor function y_k , two additional qubits are required in the commitment register. The resulting minimal protocol instance therefore relies on three additional qubits (preimage and commitment register) to secretly delegate and perform the energy measurement on a single computational qubit of the $|\eta\rangle$ -state.

Based on these assumptions, let us go through the protocol from Fig. 4.4 step by step [54]:

- S0 V passes the circuit \mathcal{C} along a description for the $|\eta\rangle$ -state to P, which implements it using its quantum power. The linear protocol structure enables the energy measurement on each qubit $|\psi\rangle$ of the $|\eta\rangle$ -state individually, which is shown in Fig. 4.4.
- S1 V prepares a family of post-quantum secure trapdoor functions y_k , based on DLWE from Eq. (4.5), specified by an index k together with the corresponding trapdoor information




Figure 4.4: **Proof-of-principle protocol on the classical verification of quantum computation [54].**

The classical verifier (V) attempts to confirm the outcome of a decision problem solved by the quantum prover (P). A proof of the correctness of P 's claim is specifically encoded in the energy of the state it prepares, here given by $|\psi\rangle = \alpha_0|0\rangle + \alpha_1|1\rangle$. To verify P 's claim about the outcome of the decision problem, V delegates measurements in bases X and Z to P based on a public key k such that the respective basis is hidden from P . This is achieved by encrypting the measurement bases with the quantum-hard DLWE-problem from Eq. (4.5), which gives V leverage over P . Alternating between “measurement rounds” that target the decision problem under verification, and simpler “test rounds” that ensure that the outputs of P were generated from the correct quantum state prevents any attempt of cheating. Finally, considering the classical statistics from the “measurement rounds” together with the trapdoor t_k qualifies V to evaluate the queried energy, which proves that the decision problem was correctly solved by P .

t_k (private key). There are two types of trapdoor functions, namely one-to-one ($k = 0$) and two-to-one ($k = 1$) maps, which correspond to measurements in Z - and X -basis, respectively. The trapdoor allows V to secretly evaluate the function preimages, which are not computable for P due to the quantum-hard DLWE problem. V then sends an index $k \in \{0, 1\}$ together with a description of the circuit that implements the function depending on k to P (public key).

S2 P implements the quantum state $|\phi_k\rangle = 1/\sqrt{2^{n_x}} \sum_{b,x} \alpha_b |b\rangle |x\rangle |y_k(b,x)\rangle$ according to Eq. (4.4). Here $|b\rangle$ is the single committed qubit from the $|\eta\rangle$ -state which together with the middle qubit $|x\rangle$ constitutes the preimage of the trapdoor function, and the last two qubits in state $|y_k(b,x)\rangle$ represent the trapdoor function image in the commitment register.

S3 V asks P to measure the commitment register, i.e., qubits three and four of state $|\phi_k\rangle$ in the Z -basis, and return the results y . If y_k was chosen to be a one-to-one function ($k = 0$), the state in the committed and preimage register (qubits one and two) is projected into the product state $\alpha_b |b\rangle |x\rangle$, where $y_k(b,x) = y$ holds. If y_k was a two-to-

one function ($k = 1$), the state of qubit one and two is projected into the superposition of the two preimages $\alpha_0 |0\rangle |x_0\rangle + \alpha_1 |1\rangle |x_1\rangle$, where $y_k(0, x_0) = y_k(1, x_1) = y$ holds.

S4 V performs with probability $1/2$ either a “measurement round” or a “test round”.

- a) “*test round*”: P measures the committed and preimage register (qubits one and two) in the Z-basis. V then verifies that P has correctly prepared $|\phi_k\rangle$ by checking $y_k(b, x) = y$. Thus, V is confident that the classical results transmitted by P were indeed generated from the correct quantum state. This ensures that the estimated energy of the $|\eta\rangle$ -state can only be low if P prepared it correctly, which excludes possible coincidences from randomly generated classical numbers.
- b) “*measurement round*”: If the “test round” passes with high probability, P measures the committed and preimage register (qubits one and two) in the X-basis. Depending on the function type y_k , the first qubit is effectively measured in the Z- or the X-basis. V then assigns the measurement bases to the classical results from the “measurement round” using the trapdoor information t_k and determines the $|\eta\rangle$ -state energy, which finally verifies the computational outcome generated by P.

For details on how to estimate the $|\eta\rangle$ -state energy based on the classical results from the “measurement round”, see Ref. [200].

It is worth noting that the protocol has some robustness to noise, a necessary condition to ensure that measurements can be made under experimental conditions without compromising the verification result. This can be shown with the results proposed in Ref. [203], which deal with verification in the presence of limited quantum resources and the absence of perfect experimental control.

In Sec. 4.3 we will follow the above protocol to experimentally demonstrate the classical verification of a single-qubit quantum computation. According to Eq. (4.3), this requires a two-qubit $|\eta\rangle$ -state. The first qubit is subjected to the single-qubit operation to be verified and the second qubit represents the clock-qubit prepared in state $|+\rangle_x = (|0\rangle + |1\rangle)/\sqrt{2}$ to realize a superposition of the identity and the target operation. The above step wise measurement protocol must be applied to both qubits. Using minimal two-bit to two-bit trapdoor functions as in the step-by-step protocol above, requires three additional qubits to delegate the energy measurement on each computational qubit of the $|\eta\rangle$ -state. The resulting eight-qubit protocol can classically verify a single-qubit quantum computation. It should be noted that while this example contains all the necessary protocol steps, the small size of the trapdoor functions cannot satisfy any security requirements [54].

While the specific sizes of trapdoor functions for securely encrypting energy measurements against a quantum computer are generally unknown [54], we can instead discuss the protocol sizes for a given trapdoor function. In general, an m -qubit circuit \mathcal{C} of N gates yields an $|\eta\rangle$ -state with $n_\eta = m + \lceil \log(N + 1) \rceil$ qubits, see Eq. (4.3). To delegate the energy measurement to each qubit in the $|\eta\rangle$ -state, we can choose n_x -bit to n_y -bit maps for the trapdoor functions that require a preimage register of size $n_x = n_y - 1$ and a commitment register of size n_y . Accumulation of all three registers yields a total number of $n_\eta \cdot (n_x + n_y)$ qubits. Crucially, ancilla qubits can, after reinitialization, be reused in all measurement rounds, relaxing the protocol size to $n_\eta + n_x + n_y$. The above proof-of-principle protocol could thus be implemented with only five qubits at the cost of performing n_η insequence measurements.

Finally, let’s talk about the next major protocol implementations. To improve security, we could introduce three-bit to three-bit trapdoor functions, where each committed qubit must be accompanied by two qubits in the preimage register and three qubits in the commitment

register. A minimal two-qubit $|\eta\rangle$ -state for verifying a single-qubit quantum computation therefore yields a twelve-qubit protocol, or a seven-qubit protocol if the ancilla qubits are recycled.

To maintain the security assumptions of the step-by-step protocol with a two-bit to two-bit transformation, but now verifying a correlated two-qubit operation requires a two-qubit clock register to control both qubits in the target operation. This leads to a four-qubit $|\eta\rangle$ -state. Considering three ancilla qubits to measure the energy in each qubit in the $|\eta\rangle$ -state yields a 16-qubit protocol, or an eight-qubit protocol if ancilla qubits can be reused.

Finally, it should be noted that the protocol has a high noise sensitivity with respect to the capabilities of current quantum devices, which makes larger implementations than the verification of a single-qubit quantum computation hardly feasible at the moment [200].

4.3 PUBLICATION: CLASSICAL VERIFICATION OF QUANTUM COMPUTATION

Quantum Sci. Technol. vol. 9, no. 2, 02LT01 (2024)

submitted on 08 February 2023, accepted on 14 February 2024 and published on 26 February 2024

<https://doi.org/10.1088/2058-9565/ad2986>

Roman Stricker¹, Jose Carrasco², Martin Ringbauer¹, Lukas Postler¹, Michael Meth¹, Claire Edmunds¹, Philipp Schindler¹, Rainer Blatt^{1,3,4}, Peter Zoller^{2,3,4}, Barbara Kraus² and Thomas Monz^{1,4}

¹ *Institut für Experimentalphysik, Universität Innsbruck, Innsbruck, Austria*

² *Institute for Theoretical Physics, University of Innsbruck, Innsbruck Austria*

³ *Institute for Quantum Optics and Quantum Information, Austrian Academy of Sciences, Innsbruck, Austria*

⁴ *Alpine Quantum Technologies GmbH, Innsbruck, Austria*

The author to the present thesis executed the experiments, analyzed the data and wrote the manuscript.

LETTER • OPEN ACCESS

Towards experimental classical verification of quantum computation

To cite this article: Roman Stricker *et al* 2024 *Quantum Sci. Technol.* **9** 02LT01

View the [article online](#) for updates and enhancements.

You may also like

- [Classical-Quantum Dual Encoding for Laser Communications in Space](#)
Mathew Winnel, Ziqing Wang, Robert Malaney *et al.*
- [A new calibration method for charm jet identification validated with proton-proton collision events at \$s = 13\$ TeV](#)
The CMS collaboration, Armen Tumasyan, Wolfgang Adam *et al.*
- [Verifiable quantum protocol for dynamic secure multiparty summation based on homomorphic encryption](#)
Mei Luo, Fulin Li, Li Liu *et al.*

The logo for kiutra, featuring a stylized circular icon to the left of the word "kiutra" in a bold, lowercase sans-serif font.

Easy-to-use and Helium-3 free
cryogenics solutions

LEARN MORE

Quantum Science and Technology



LETTER

Towards experimental classical verification of quantum computation

OPEN ACCESS

RECEIVED
8 February 2023

REVISED
19 January 2024

ACCEPTED FOR PUBLICATION
14 February 2024

PUBLISHED
26 February 2024

Roman Stricker^{1,*}, Jose Carrasco², Martin Ringbauer¹, Lukas Postler¹, Michael Meth¹,
Claire Edmunds¹, Philipp Schindler¹, Rainer Blatt^{1,3}, Peter Zoller^{2,3}, Barbara Kraus² and Thomas Monz^{1,4}

¹ Institut für Experimentalphysik, Universität Innsbruck, Innsbruck, Austria

² Institute for Theoretical Physics, University of Innsbruck, Innsbruck Austria

³ Institute for Quantum Optics and Quantum Information, Austrian Academy of Sciences, Innsbruck, Austria

⁴ Alpine Quantum Technologies GmbH, Innsbruck, Austria

* Author to whom any correspondence should be addressed.

E-mail: roman.stricker@student.uibk.ac.at

Keywords: quantum computation, quantum information, quantum verification, quantum physics, quantum algorithms

Supplementary material for this article is available [online](#)

Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Abstract

With today's quantum processors venturing into regimes beyond the capabilities of classical devices, we face the challenge to verify that these devices perform as intended, even when we cannot check their results on classical computers. In a recent breakthrough in computer science, a protocol was developed that allows the verification of the output of a computation performed by an untrusted quantum device based only on classical resources. Here, we follow these ideas, and demonstrate in a first, proof-of-principle experiment the verification of the output of a quantum computation using only classical means on a small trapped-ion quantum processor. We contrast this to verification protocols, which require trust and detailed hardware knowledge, as in gate-level benchmarking, or additional quantum resources in case we do not have access to or trust in the device to be tested. While our experimental demonstration uses a simplified version of Mahadev's protocol we demonstrate the necessary steps for verifying fully untrusted devices. A scaled-up version of our protocol will allow for classical verification, requiring no hardware access or detailed knowledge of the tested device. Its security relies on post-quantum secure trapdoor functions within an interactive proof. The conceptually straightforward, but technologically challenging scaled-up version of the interactive proofs, considered here, can be used for a variety of additional tasks such as verifying quantum advantage, generating and certifying quantum randomness, or composable remote state preparation.

1. Introduction

Quantum computers are now widely believed to be at the brink of solving problems that are classically intractable [1–3]. Yet, for operating these devices in such a quantum advantage regime, one is confronted with the question of how their output can be verified. The answer depends strongly on the task for which the device is used, the level of control, and its quality. In case the user has direct access to the device, for instance, they can perform gate benchmarks [4] and develop a microscopic error model to gain confidence in the device. For such a scenario, several verification and validation schemes have been proposed [5]. However, these techniques are rarely scalable and require detailed hardware knowledge. These requirements can be alleviated somewhat by cross-verifying honest quantum devices [6, 7] to assess their relative performance in a hardware-independent fashion.

The situation becomes significantly more challenging, when the device to be verified cannot be accessed, nor trusted, as might be the case for cloud-access quantum computers. If the user employs a quantum computer to solve a problem within the complexity class NP, such as factoring a large number into its prime factors, the solution to the problem can be efficiently verified with a classical computer. However, it is

believed that quantum computers are capable of efficiently solving problems that can no longer be efficiently verified classically [5, 8]. How can one then rely on the output, given that the quantum computer (or the person operating it) might be malicious and might want to convince the user that the answer to e.g. a decision problem is ‘yes’ when it is actually ‘no’? Harnessing the full power of a quantum device therefore brings with it the necessity to develop protocols for verifying its output. Moreover, in light of recent advancements in remotely accessible quantum computers, there is an increased need for verification protocols for untrusted devices.

Several powerful verification schemes designed for utilizing and testing untrusted quantum devices have been developed, as we briefly summarize below. Some of them provide unconditional security but come with the drawback that the user needs to possess limited quantum capacity [9–14]. In practice though, a user is typically limited to classical resources. In a recent breakthrough, the remaining barrier has been overcome, namely, the requirement for the user to possess quantum resources. This enables, even a purely classical user, under a computational assumption, to verify the output of the significantly more powerful quantum device [15–17]. Understandably, these verification schemes are currently too resource-intensive for practical implementation with existing technology, necessitating further research to enhance their feasibility.

Here, we consider the verification protocol presented in [15], where a purely classical user (verifier) verifies the output of a quantum computer (prover) based on a computational assumption. Realizing the fully secure protocol exceeds the capability of current technology. Therefore, our focus here is on the experimental realization of the necessary subprotocols required to implement such a verification scheme. Specifically, we follow the protocol outlined in [18] for a showcase example of a quantum computation. We relax the security constraints and demonstrate the main ingredients for classical verification tailored to an eight-qubit trapped-ion quantum processor [19]. Our experiment illustrates that this fully classical verification of the output of the computation requires considerably higher fidelity operations of the quantum device than just implementing the underlying quantum computation directly.

1.1. Computing with untrusted quantum devices

Schemes, developed to utilize and verify an untrusted quantum device differ significantly in the resources required by the user and the security level they offer. To highlight certain differences, we briefly summarize some of them here. Schemes, which operate without relying on trust, such as verifiable blind quantum computation [9–14] or schemes utilizing two non-communicating quantum processors [20] have been developed. While these protocols offer unconditional security, they do demand that the user possesses a limited quantum device or has access to two non-communicating quantum processors. For example, the user must be capable of generating and transmitting single qubit states to the quantum computer. A protocol to verify the output of a fault-tolerant quantum computation in a measurement-based model (MBQC) has been presented in [21]. There, verification in a noise setting by a user with a limited quantum device has been demonstrated.

In contrast to the aforementioned schemes, recent achievements in [15–17] have eliminated the necessity for the user to possess quantum resources. There, it has been shown that a purely classical user can verify the output of computation performed by a much more powerful quantum device. This protocol, which we review in appendix E, is no longer unconditionally secure. Instead, it relies on a computational assumption, namely the existence of problems that are hard to solve even by a quantum computer. More precisely, the classical verification protocol presented in [15] is secure against dishonest devices under the computational assumption that there exist trapdoor functions that are post-quantum secure [22–24]. These cryptographic functions have the property that given $f(x)$ it is, even for a quantum computer, not possible to determine the preimage x efficiently. However, having some additional information (trapdoor), the task becomes easy, even for a classical computer. It is widely believed that these functions exist and can be obtained from the Learning with Errors problem [22]. The key idea in the classical verification protocol [15] is to use post-quantum secure trapdoor functions within an interactive proof [25]. There, the user exchanges messages with the quantum device to eventually get convinced of the correctness of the output or to decide that it should be rejected. As shown in [15] the protocol verifies the output of an arbitrary decision problem within the class of BQP (bounded error quantum polynomial time), i.e. a problem that can be solved efficiently by a quantum computer. It fulfills the *completeness* and *soundness* conditions [15]. That is, if the answer provided by the quantum device is true, the honest prover can convince the verifier that it is indeed true (completeness). If the answer is false, no prover, even if dishonest, can convince the verifier that it is true, except with some small probability (soundness).

Here, for the first time, we will showcase the experimental implementation of the essential components needed to execute this protocol. For this purpose, we will consider a simplified version of the protocol with relaxed security constraints. Our results below differ significantly from the recently reported experiments verifying quantum advantage and quantumness [26] (see appendix E). While for both interactive proof

protocols the usage of post-quantum secure functions within the measurement protocol is crucial, the verification of the output of a quantum computation requires additional important steps prior to the measurement. As we explain below, in the first part of the protocol, the answer to the decision problem of interest is encoded in the ground state energy of a Hamiltonian and a corresponding m -qubit state, with m larger than the input size of the decision problem, is prepared by the prover. After this step, these m qubits are measured in a way hidden from the prover. Crucially, this first part of the protocol is not needed to verify quantumness or quantum advantage. This illustrates the difference between the verification protocols and also explains the additional experimental challenges associated with verifying the output of a computation. Due to the resulting experimental difficulties, we realize here a simplified version of the protocol [15], which nevertheless retains all the steps required for a full-fledged verification protocol.

1.2. Classical verification of the output of a quantum computer

Let us now explain the verification protocol for an arbitrary decision problem within the class of BQP [15]. For the experimental realization we will then consider one specific problem. A description of the corresponding quantum n -qubit circuit \mathcal{C} , consisting of T single and two-qubit gates, is sent to the quantum prover Bob (B). He can compute the output of the decision problem ('yes' or 'no') efficiently. Without loss of generality, we assume that B claims the answer is 'yes' (the 'no' case is similar, see appendix D.3). The classical verifier Alice (A) wants to verify this output using only classical means. The correctness of B's answer can be checked with the help of an interactive proof as we will explain in the following.

Following [27], A and B first construct a Hamiltonian H that depends on the prover's answer and the circuit of interest. This Hamiltonian acts on the n system qubits and additionally $\lceil \log(T+1) \rceil$ qubits (the so-called clock register [27, 28], see equation (2)). H can be chosen to consist of only 2-local terms containing the Pauli operators X and Z [27, 29]. Crucially, the ground-state energy of this Hamiltonian $\lambda(H)$ encodes the correct output of the initial computation. More precisely, the energy is below a certain value only if B's answer is correct and larger otherwise [27]. Hence, an honest B can prove that his answer is correct by preparing a state with low energy. To this end, he can prepare efficiently the clock (or history) state $|\eta\rangle$ associated with \mathcal{C} [27]. This state is a superposition, $\sum_{t=0}^T |\Psi_t\rangle|t\rangle$, where $|\Psi_t\rangle$ denotes the state of the system after applying the first t gates of \mathcal{C} to the initial state and the second register denotes the clock register. By construction, this state has low energy, $\langle \eta | H | \eta \rangle$, in case B is honest as outlined below.

If A would be able to measure this energy, the problem is solved [30].

Since A does not possess a quantum device she needs to delegate the energy measurement to B in such a way that B does not learn what is actually measured. This part of the protocol is crucial for the verification and is achieved using post-quantum-secure trapdoor functions [22] within an interactive proof. A constructs these cryptographic functions and keeps the trapdoor information for herself. This information is what allows her to compute the preimages of the function, which B cannot do efficiently. Furthermore, the functions are of two different types (see below), which will determine whether B performs an X or Z measurement. B cannot efficiently differentiate between the two types of functions and thus cannot learn which measurement he implements by following the measurement protocol depicted in figure 1(a). After receiving a description of the function (labeled by k_i in figure 1(a)), B uses it to entangle each of the qubits in the η -state with several auxiliary qubits. Some of them are then measured in the computational basis leading to outcomes \bar{y}_i in figure 1. Importantly, the state of the remaining qubits depends on these outcomes. More precisely, the remaining qubits are in a superposition of computational basis states, which are the preimages of \bar{y}_i . They are known to A. However, B cannot learn them efficiently. At this point the power of interactive proofs comes into play. B is now forced to answer all subsequent questions by A in a way that is consistent with \bar{y}_i . A could now exploit her superiority to verify quantumness and quantum advantage [17]. In a verification protocol this interactive proof is not only used to ensure that B holds a quantum state, which leads (approximately) to the observed measurement outcomes, but also to enable A to determine its energy in a way hidden to B [15].

Since B can prepare a state with low energy only in case his answer was correct this interactive protocol allows A to verify that answer.

A realization of Mahadev's protocol, as presented in [15], requires randomly chosen trapdoor functions with additional properties, especially the hard-core bit property (see appendix E), and with a very large range. Hence, many auxiliary qubits are required in the measurement protocol⁵. More precisely, several hundred auxiliary qubits⁶ are required for the secure delegation of single-qubit measurements, in either the X - or Z -basis, with the basis hidden to the quantum prover. The verification of the output of a quantum computation running on n qubits requires, as explained before [15, 27], $O(n + \log(T+1))$ qubits to prepare

⁵ Note, however, that there are recent proposals like [31] in which some of those requirements can be removed.

⁶ See, for instance, the website www.latticechallenge.org, where functions similar to the ones we need are explicitly constructed.

We experimentally demonstrate the smallest protocol instance, that is, verifying the output of a single-qubit quantum circuit. Let us denote that single-qubit circuit to verify by $\mathcal{C} = U(\alpha)$ parameterized with α as follows

$$U(\alpha) = \cos\alpha Z + \sin\alpha X. \tag{1}$$

For convenience⁷, we choose a promise problem with output ‘yes’ if $p_0(\mathcal{C}) > 3/5$ and ‘no’ if $p_0(\mathcal{C}) < 1/10$, where $p_0(\mathcal{C}) = |\langle 0|\mathcal{C}|0\rangle|^2 = \cos^2\alpha$ (for more details see appendix D.3). A sends a description of the circuit to B, who runs the computation on the quantum computer and obtains an output. Without loss of generality, we assume B claims that the answer is ‘yes’. To verify this output the protocol proceeds with the following steps (see appendix C and figure 1 for details).

Step 1: determination of the Hamiltonian – A determines classically the corresponding Hamiltonian [27] H , as given in equation (C.1). This Hamiltonian acts on two qubits and contains only X and Z operators.

Step 2: preparation of the clock state – B prepares the clock state corresponding to $U(\alpha)$ (figure 1(b)),

$$|\eta\rangle = \frac{1}{\sqrt{2}} [|0\rangle|0\rangle + (U(\alpha)|0\rangle)|1\rangle] \equiv \sum_{b_1, b_2} \alpha_{b_1, b_2} |b_1, b_2\rangle. \tag{2}$$

Its energy is given by $\langle \eta|H|\eta\rangle = 1 - p_0(\mathcal{C}) (= \sin^2\alpha)$. Thus, in case the answer of the problem was indeed ‘yes’ it holds that $\lambda(H) < \langle \eta|H|\eta\rangle < 2/5$. At the same time we have $\lambda(H) \geq \langle \eta|H|\eta\rangle - 2/5$, such that $\lambda(H) > 1/2$ in case the correct answer was ‘no’ (appendix D.2). Consequently, B can only prepare a quantum state with energy below 0.4 in case his output ‘yes’ is indeed the correct output.

Step 3: selection of the trapdoor functions – In order to delegate the measurements of the operators occurring in the Hamiltonian to B, A chooses trapdoor functions y_k labeled by $\mathbf{k} = \mathbf{0}$ or $\mathbf{k} = \mathbf{1}$ to perform Z or X basis measurements respectively, see appendix C for details. We choose the one-to-one function y_0 as the identity and the two-to-one function y_1 as $y_1(z_1, z_2) = (0, 0)$ or $(1, 0)$ for $z_1 = z_2$ or $z_1 \neq z_2$ respectively. Here, $z_i \in \{0, 1\}$ for $i = 1, 2$. For instance, if A wants to measure the term $Z_1 X_2$, she chooses $(k_1, k_2) = (0, 1)$.


Step 4: entangling the qubits to be measured with auxiliary qubits – After receiving k_1, k_2 , the prover B attaches three auxiliary qubits to each qubit of the η -state and implements a unitary operator to generate (depending on the label) one of four 8-qubit entangled states (see figure 1(b))

$$|\phi_{k_1, k_2}\rangle \propto \sum_{b_1, b_2} \alpha_{b_1, b_2} |b_1\rangle_1 |b_2\rangle_2 \times \sum_{w_1=0}^1 |w_1\rangle_3 |y_{k_1}(b_1, w_1)\rangle_{45} \sum_{w_2=0}^1 |w_2\rangle_6 |y_{k_2}(b_2, w_2)\rangle_{78}. \tag{3}$$

Step 5: measurement protocol – We explain here the measurement protocol for a single qubit, namely qubit 1, in the η -state, i.e. consider the state in equation (3) without summing over b_2, w_2 and ignoring qubits 2, 6, 7, 8 (the measurement of qubit 2 is equivalent, see appendix C). Qubit 1 is entangled in a way that depends on \mathbf{k}_1 to the three auxiliary qubits, namely qubits 3, 4, 5 (figure 1(b)). A asks B to measure qubits 4 and 5 in the Z -basis and requests the 2-bit outcome, the commitment, (denoted by \bar{y}_1 in figure 1). After this measurement, qubits 1 and 3 are (i) in a product state $|b_1\rangle|x_1\rangle$ where $y_0(b_1, x_1) = \bar{y}_1$ for $\mathbf{k}_1 = \mathbf{0}$; or (ii) in a superposition of the state $|0\rangle|x_1^{(0)}\rangle$ and $|1\rangle|x_1^{(1)}\rangle$ where $y_1(0, x_1^{(0)}) = y_1(1, x_1^{(1)}) = \bar{y}_1$ for $\mathbf{k}_1 = \mathbf{1}$. Note that A knows which type of function was used and in which state the remaining qubits can be, i.e. she knows the preimage(s) of \bar{y}_1 using the trapdoor information. However, B cannot learn either of them due to the properties of the cryptographic trapdoor functions. This fact implies that A can now ask B for additional measurement outcomes (measuring the remaining qubits) where only she knows the possible outcomes. In order to provide the correct answers, B basically needs to have a quantum state, namely a state of the form from equation (3). Stated differently, the fact that the prover has to commit to an answer (\bar{y}_1) before A asks for additional measurement outcomes, which need to be consistent with the outcome \bar{y}_1 and are unknown to B, enforces B to have a quantum state. These properties can be utilized to verify quantumness and quantum advantage, as recently experimentally reported [26]. In the context considered here, the properties explained above also enable A to verify the output of B. She proceeds by randomly choosing to perform one of the following two steps.

Step 5a: test round – A asks B to measure the remaining qubits in the Z -basis. The outputs, which are known to A (but not to B), need to be consistent with the previous answer. This allows A to test the correct behaviour of B (see figure E.1 for the experimental data and appendix E).

⁷ Note that there is some freedom in choosing the values $(a, b) = (1/10, 3/5)$. In general, for a uniformly generated family of circuits C_n acting on n qubits, the requirement is [27] that the ‘gap’ $b - a$ must be larger than $1/\text{poly}(n)$.



Step 5b: measurement round – A asks B to measure the remaining qubits in the X-basis. Given the output, A can determine the measurement outcome corresponding to her choice of k_1, k_2 (see figure 2 and appendix C). She uses these outcomes to finally determine the energy.

2.2. Experimental demonstrations

Given that B passes the tests of A with sufficiently high probability, a scaled-up version of the protocol (see appendix E) verifies the prover's answer whenever the energy is below 0.4. Figure 2(a) shows the corresponding experimentally measured expectation values following the above protocol. The data covers the whole range $0 \leq \alpha \leq \pi/2$ and, up to experimental imperfections, follows the ideal theory prediction (dashed lines). A then makes use of these expectation values to calculate the energy of the state held by B according to the Hamiltonian H . Results on the total η -state energy are depicted in figure 2(b), where we successfully certify a 'yes' outcome for $\alpha \leq 0.12\pi/2$ (dark gray shaded region), where data-points well undercut the derived energy thresholds 0.4. Notably, our results cover a good fraction of the light gray shaded region, where by construction a proof on the 'yes' outcome exists. Further, all data-points follow the expected $\langle \eta | H | \eta \rangle = \sin^2 \alpha$ behaviour. The stated error-bars represent one standard deviation of quantum projection noise around the mean values from all experimental samples. We use these statistical errors to calculate the significance level as the probability that the verifier, A, incorrectly accepts the result. The significance level thereby states the probability that a measured value below the threshold actually corresponds to a true value above the threshold as a result of statistical uncertainty. For data points in the dark gray shaded region, we receive significance levels $\{0.012, 0.042, 0.015, 0.102, 0.039\}$ for α values $\{0, 0.03\pi/2, 0.06\pi/2, 0.09\pi/2, 0.12\pi/2\}$, respectively.

Due to experimental noise, an energy offset appears with respect to the ideal outcome and remains roughly constant over the course of α . This is confirmed by a simple depolarizing noise model, outlined in appendix I, that accurately describes the measured data points. By construction, any noise will lead to an increase in the measured energy, which in turn prevents the verification protocol from determining the correctness of the computation. Note that for an individual experimental run, the protocol never produces an incorrect result. A computation is either verified as correct, or it cannot be verified.

Before further discussing the experimental realization, let us stress again that the functions in the fully secure version of the protocol need to have much larger domain and range. In addition, they need to satisfy the so-called hard-core bit property (see appendix E for details). To show the main ingredients of classical verification, in this work we simply use instead the random 1-to-1 and 2-to-1 functions $y_0(\cdot)$ and $y_1(\cdot)$, respectively. For this simplified version of Mahadev's protocol, to verify a single-qubit computation we need

in total 8 qubits, 19 single-qubit gates, and 5 entangling two-qubit gates. This compact circuit is a testament to the efficiency of our implementation of the verification protocol, which is made possible, in part, by the all-to-all connectivity of the trapped-ion platform, see appendix A. To emphasize the high control needed over a quantum device to successfully operate the protocol, we quantify the system performance. A simplistic estimate of the fidelity of the experimentally prepared $|\eta\rangle$ -state as measured using the six auxiliary qubits is given by considering error-rates on all individual gates in the respective circuits. Considering single- and two-qubit errors inherent to our setup (see appendix B), we expect a fidelity of 0.873(17). We compare this number to results from figure 2(a). For this we use the fact that the η -state implemented by the circuit in the grey box from figure 1(b) represents a Bell-state for $\alpha = \pi/2$. Hence, averaging expectation values $Z_1 Z_2$ and $X_1 X_2$ at $\alpha = \pi/2$ provides an estimate of the η -state fidelity measured via the six auxiliary qubits, which results in 0.852(8). This is in good agreement with the above mentioned error-model. Moreover, the latter analysis was similarly performed on the direct estimation of the $|\eta\rangle$ -state energy depicted in figure H.1 from the appendix H and leads to an estimated Bell-state fidelity of 0.945(12). The difference between expected and observed outcome is due to experimental errors when extracting the $|\eta\rangle$ -state energy via the auxiliary qubits as required for the classical verification. As a final simpler comparison, note that merely performing, but not verifying the single-qubit computation $U(\alpha)$ to compute $p_0 = |\langle 0|U(\alpha)|0\rangle|^2$ could be achieved with a fidelity of 0.9988(4), see appendix B.

3. Discussion

To conclude, the verification of the output of a quantum computation by purely classical means comes, in general, with stringent requirements on the classical user as well as the quantum prover. On the classical side, finding suitable post-quantum secure trapdoor functions will be crucial for the application of such interactive proof protocols. On the quantum side, a large and powerful quantum computer is required in the sense that secure classical verification is expected to require hundreds of additional qubits per computational qubit and a similar increase in gate performance between merely performing the computation and classically verifying it. As experimental imperfections accumulate in such as verification protocol, its realization is very demanding, as already evident in our demonstration. However, in the continuing effort to relax the constraint to achieve security and in view of wide-ranging applications [17, 26, 32, 33], an inevitable future challenge will be to improve and generalise the protocols both theoretically and technologically, in particular the realisation of interactive proofs using secure post-quantum trapdoor functions.

Data availability statement

The data that support the findings of this study are openly available at [34]. All codes for data analysis are available from the corresponding author upon reasonable request.

Acknowledgments

J C and B K are grateful for the support of the Austrian Science Fund (FWF): stand alone project P32273-N27 and the SFB BeyondC F 7107-N38. R S, L P, M M, C E, M R, P S, T M and R B gratefully acknowledge funding by the U.S. ARO Grant No. W911NF-21-1-0007. We also acknowledge funding by the Austrian Science Fund (FWF), through the SFB BeyondC (FWF Project No. F7109), by the Austrian Research Promotion Agency (FFG) Contracts 872766 and 877616, by the EU H2020-FETFLAG-2018-03 under Grant Agreement No. 820495, and by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via the U.S. ARO Grant No. W911NF-16-1-0070 and the US Air Force Office of Scientific Research (AFOSR) via IOE Grant No. FA9550-19-1-7044 LASCEM. This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 840450. It reflects only the author's view, the EU Agency is not responsible for any use that may be made of the information it contains. P Z acknowledges support by the US Air Force Office of Scientific Research (AFOSR) via IOE Grant No. FA9550-19-1-7044 LASCEM, the European Union's Horizon 2020 research and innovation program under Grant Agreement No. 817482 (PASQuanS), and by the Simons Collaboration on Ultra-Quantum Matter, which is a grant from the Simons Foundation (651440, P Z). P Z and R B acknowledge support by the IQI GmbH.

Author contributions

J C, B K and P Z derived the theory results. R S, M R, L P, M M, C E, P S, and T M performed the experiments. R S analyzed the data. B K and P Z (theory), T M and R B (experiment) supervised the project. All authors contributed to writing the manuscript.

Conflict of interest

The authors declare no competing interests.

Appendix A. Experimental toolbox

Experiments are performed on an ion-trap quantum computer as illustrated by the middle inset of figure 1 from the main text. The setup operates on a linear chain of $^{40}\text{Ca}^+$ ions confined in ultra high vacuum using a linear Paul trap. Each ion acts as a qubit encoded in the electronic levels $S_{1/2}(m = -1/2) = |0\rangle$ and $D_{5/2}(m = -1/2) = |1\rangle$ denoting the computational subspace [19]. Arbitrary qubit manipulation is realized with coherent laser-ion interaction, upon which the setup is capable of implementing a universal set of quantum gate operations. This universal gate-set comprises of addressed single-qubit rotations with an angle θ around the x - and the y -axis of the form $R^{\sigma_j}(\theta) = \exp(-i\theta\sigma_j/2)$ with the Pauli operators $\sigma_j = X_j$ or Y_j acting on the j th qubit, together with two-qubit Mølmer–Sørensen entangling gate operations $MS_{i,j}(\theta) = \exp(-i\theta X_i X_j/2)$ [19]. Multiple addressed laser beams, coherent among themselves, allow for arbitrary two-qubit connectivity across the entire ion string [35]. Initial state preparation in $|0\rangle$ is reached after a series of doppler, polarization-gradient and sideband cooling. Read-out is realized by exciting a dipole transition solely connected to the lower qubit level $|0\rangle$ and collecting its scattered photons, from which the computational basis states $|0\rangle$ and $|1\rangle$ can be identified. Thereby, a qubit's state is revealed by accumulating probabilities from multiple experimental runs. The dipole laser collectively covers the entire ion string, which enables a complete read-out in one measurement round. Additional pump lasers support efficient state preparation as well as cooling and prevent the occupation of unwanted meta-stable states outside the computational subspace $\{|0\rangle, |1\rangle\}$.

Appendix B. Ion-trap implementation & error rates

The particular circuit for the classical verification protocol discussed in the main text is again depicted in figure B.1(a) with a focus on the ion-trap implementation. The circuit demands for local gates, more specifically Hadamards H, as well as two-qubit CNOT-gates—the latter creating pairwise entanglement. Figure B.1(b) follows up on the sub-circuits corresponding to those building blocks suitable and optimized for the ion-trap gate set. Each CNOT gate demands for a full-entangling, two-qubit $MS^{X_i,j}(-\pi/2)$ alongside four single-qubit gates, i.e. single-qubit rotations of type $\theta = \pi/2$ around X, Y or Z. $CU(\alpha)$ from the grey box is realized upon two single-qubit gates acting on the prover-qubit to continuously change basis between CPHASE and CNOT for $\alpha = 0$ and $\pi/2$ respectively. The total number of single-qubit gates is further reduced by compiling the overall circuit. Thus, the final implementation of each circuit Z_1Z_2 , Z_1X_2 , X_1Z_2 and X_1X_2 requires five $MS^{X_i,j}(-\pi/2)$ alongside 19 single-qubit gates ($\theta = \pi/2$ around X, Y or Z).

All results from the main text, covered by figure 2 have been accumulated from 2000 experimental runs in each data point to faithfully estimate the protocol's outcome. The respective number of experimental runs in complementary experiments covered by this appendix have further been stated in the individual figure captions. Generally, for the estimation of quantum projection noise, as stated by error-bars in figures and errors in numbers, the probabilities of measured outcomes were resampled using a multinomial distribution considering the number of experimental runs. If not stated differently the underlying errors were then extracted from the resampled data-set and correspond to 1 standard deviation.

In the following we discuss error-rates inherent to our system. Our average single-qubit fidelity ($\theta = \pi/2$ around X, Y or Z) estimated via randomized benchmarking reads 0.9994(3) [35]. To further improve single-qubit gates on circuit level, we construct each gate out of three gates using various axes following the aim of reducing cross-talk to neighbouring qubits. This results in a slightly lower average fidelity on the composite gate of 0.998(1). However, on the eight-qubit circuits this approach is beneficial, as otherwise cross-talk errors proliferate generally lowering the implementation's quality.

The performance of the two-qubit MS-gates may slightly differ upon the chosen qubit-pair across the ion-string. As such, we characterize the particular five pairs occurring in the X_1X_2 -measurement to build the simple error model below and understand the experimental limitations. Note that the qubit order differs from the circuits depicted in figure B.1 as we optimized for inter-ion spacing to minimize cross-talk. Given




Figure B.1. Ion trap implementation of the classical verification protocol. (a), Eight-qubit circuit implementing our decision problem linked to the outcome of the gate $U(\alpha) = \cos \alpha Z + \sin \alpha X$, according to figure 1 from the main text. The circuit consists of single qubit gates, i.e. Hadamards H, as well as CNOT-gates to create pairwise entanglement. The final read-out in Z-basis is performed via collective fluorescence detection, see text. (b), Efficient ion-trap implementation of the building blocks from (a). We emphasize that one CNOT requires one two-qubit $MS^{X_i, j}(-\pi/2)$ alongside four single-qubit gates. We implement $CU(\alpha)$ using two additional single-qubit gates on the prover-qubit to continuously change between CPHASE and CNOT for $\alpha = 0$ and $\pi/2$ respectively. Compiling the final circuits results in 19 single-qubit gates alongside five $MS^{X_i, j}(-\pi/2)$.

Table B.1. Summary of error-rates on MS-gate pairs according to figure F.1. Results on population, coherence and fidelity are extracted from the exponential decay of a series of MS-gates as depicted in the figure. All errors refer to 1 standard deviation from the exponential fit uncertainty.

MS-gate	Population	Coherence	Fidelity
qubit-pair	$\mathcal{F}_{\text{pop.}}$	$\mathcal{F}_{\text{coh.}}$	$\mathcal{F}_{\text{tot.}}$
(1,8)	0.983(1)	0.984(1)	0.984(1)
(5,8)	0.974(6)	0.981(2)	0.979(3)
(1,7)	0.982(1)	0.986(1)	0.985(1)
(2,5)	0.977(2)	0.982(1)	0.980(1)
(4,7)	0.966(2)	0.975(2)	0.972(1)

similar ion-pairs for Z_1Z_2, Z_1X_2, X_1Z_2 measurements, this is a representative set across the verification circuits. Gate fidelities are estimated in a robust fashion from sequences of MS-gates, see appendix F for the technical details. The resulting populations, coherences and fidelities for each gate, as extracted from figure F.1 in the SI, are summarized in table B.1. Fidelities scatter between 0.972(1) and 0.985(1), with the lower values for qubit-pairs with smaller inter-ion spacing and more neighbours as a consequence of cross-talk and due to position dependent laser beam and focus quality.

Note that calibrations for the verification circuits are independent of the decay series discussed here. The latter only serve to provide fidelity estimates, which we use, along those for the single-qubit gates, to predict outcome qualities of the verification experiments throughout the manuscript. The characterizations also indicate that the overall performance is limited by the entangling gates rather than single-qubit ones. Similar decay rates between population and parity shown in table B.1 led us to conclude that the MS-gate performance, and by that our verification experiments, are dominated by depolarizing noise. We successfully utilize depolarizing noise in an error-model to characterize system limitations—thoroughly discussed at the bottom of appendix I.

We proceed to incorporate these error-rates into a simplistic estimate on the expected fidelity of the η -state as measured using the six auxiliary qubits. Therefore, we accumulate error-rates on all gates from the circuits. The 19 single-qubit gates reduce the fidelity to 0.966(18). Additionally taking MS-gate rates from table B.1 into account, we expect a final fidelity of

$$\mathcal{F} = 0.966(18) \cdot 0.984(1) \cdot 0.979(3) \\ \cdot 0.985(1) \cdot 0.980(1) \cdot 0.972(1) = 0.873(17)$$

for the η -state measurement. We compare this to the results from the experimental implementation depicted in figure 2(a) by averaging Z_1Z_2 and X_1X_2 at $\alpha = \pi/2$ representing an estimate of the η -state's Bell-state fidelity. The result reads:

$$\mathcal{F} \sim \frac{Z_1Z_2(\alpha = \pi/2) + X_1X_2(\alpha = \pi/2)}{2} \\ = \frac{0.891(10) + 0.812(13)}{2} = 0.852(8),$$

and is again in good agreement with the above simplistic error-modelling.

Appendix C. The protocol step-by-step

Here we elaborate on steps 1 – 5 of the verification protocol discussed in the main text, and as implemented in the experiment. In particular, we discuss the details of all the steps of the protocol presented in the main text.

Step 1: determination of the Hamiltonian – The Hamiltonian associated to the circuit $\mathcal{C} = U(\alpha)$ given in equation (1) is defined in the Hilbert space of two qubits and reads $H = H_{\text{out}} + 6H_{\text{in}} + 3H_{\text{prop}}$ (see appendix D.1), where

$$H_{\text{out}} = \frac{1}{2}(1 - Z_1 - Z_2 + Z_1Z_2), \\ H_{\text{in}} = \frac{1}{4}(1 - Z_1 + Z_2 - Z_1Z_2), \\ H_{\text{prop}} = \frac{1}{2}(1 - \cos \alpha Z_1X_2 - \sin \alpha X_1X_2). \quad (\text{C.1})$$

Step 2: preparation of the clock state – As explained in appendix D.2, the ground-energy of H is correlated to the answer of the promise problem and the corresponding clock state (cf (D.5)) is given in equation (2).

Step 3: selection of the trap door function – Our family of functions contains functions transforming two-bits strings to two-bits strings, and consists of two elements labeled by $\mathbf{k} = \mathbf{0}$ and $\mathbf{k} = \mathbf{1}$, respectively. From the main text we recall the definition of the one-to-one function $y_0(z_1, z_2) = (z_1, z_2)$ (identity), and and two-to-one function $y_1(0, 0) = y_1(1, 1) = (0, 0)$ and $y_1(0, 1) = y_1(1, 0) = (1, 0)$, respectively. A chooses a term P_1P_2 in the Hamiltonian that she wants to measure. This determines a pair of labels $(k_1, k_2) \in \{0, 1\}^2$ as

$$k_i = \begin{cases} 0, & P_i \in \{Z_i, 1_i\}, \\ 1, & P_i = X_i. \end{cases} \quad (\text{C.2})$$

She sends (k_1, k_2) to B. Together with k , A generates a trapdoor t_k (see main text) that she keeps private. In the examples considered here, the trapdoor information, t_k , together with an output $y_k(z_1, z_2)$ lead to the preimage(s) of $y_k(z_1, z_2)$.

Step 4: entangling the qubits to be measured with the auxiliary qubits – An honest B prepares the clock state $|\eta\rangle$ given in equation (2) (cf (D.5)). He attaches six auxiliary qubits to it, and performs the unitary transformation $|\eta\rangle|0^{\otimes 6}\rangle \mapsto |\phi_{k_1, k_2}\rangle$, with $|\phi_{k_1, k_2}\rangle$ given by equation (3). The preparation of those states can be done efficiently by a quantum computer.

Step 5: measurement protocol – B is asked to measure some registers of the state he is supposed to hold (the state $|\phi_{k_1, k_2}\rangle$). He is asked to measure the registers (4, 5) (obtaining $\bar{y}_1 \in \{0, 1\}^2$) and the registers (7, 8) (obtaining $\bar{y}_2 \in \{0, 1\}^2$) in the Z basis. He sends (\bar{y}_1, \bar{y}_2) to A. After these measurements, the state of registers (1, 3) and (2, 6) depends on the labels \mathbf{k}_1 and \mathbf{k}_2 , respectively. Let us discuss the cases explicitly:

- In case $(k_1, k_2) = (0, 0)$, the state of the remaining registers will be the product state $|b_1, x_1\rangle_{13}|b_2, x_2\rangle_{26}$, with probability $|\alpha_{b_1, b_2}|^2$, where $y_0(b_i, x_i) = \bar{y}_i$ for $i = 1, 2$.
- In case $(k_1, k_2) = (0, 1)$, the state of the remaining registers will be $\propto \sum_{c=0,1} \alpha_{b_1, c} |b_1, x_1\rangle_{13}|c, x_2(c)\rangle_{26}$, where $y_0(b_1, x_1) = \bar{y}_1$ and $y_1(0, x_2(0)) = y_1(1, x_2(1)) = \bar{y}_2$.

The cases $(k_1, k_2) = (1, 0)$ and $(k_1, k_2) = (1, 1)$ can be computed in a similar way.

After receiving the outcomes \bar{y}_1 and \bar{y}_2 , to which B is now committed to, A randomly chooses (with equal probability) to run either a test or a measurement round. In both cases, B is asked to measure the remaining registers of the state he is supposed to hold.

Step 5a: test round – In a test round, B is asked to measure the registers (1, 3) and (2, 6) in the Z basis and obtains the outcomes $(b_1, x_1) \in \{0, 1\}^2$ and $(b_2, x_2) \in \{0, 1\}^2$, respectively. He sends the results to A. She checks whether $y_{k_i}(b_i, x_i) = \bar{y}_i$ for $i = 1, 2$ as would be the case if B was honest. If this is not the case, she rejects. In figure E.1(a) we show the probability with which A rejects for each possible label.

Step 5b: measurement round – In a measurement round, B is asked to measure the registers (1, 3) and (2, 6), in the X basis and obtains outcomes $(c_1, d_1) \in \{0, 1\}^2$ and $(c_2, d_2) \in \{0, 1\}^2$, respectively. He sends the results to A. For each qubit i , with $i = 1, 2$, there are two options depending on k_i .

- If $k_i = 0$, the i th qubit of the clock state was effectively measured in the Z basis. In this case, A ignores (c_i, d_i) and just computes, with the trapdoor, the preimage of \bar{y}_i under $y_0(\cdot, \cdot)$. In other words, she finds (m_i, x_i) such that $y_0(m_i, x_i) = \bar{y}_i$. Finally, she stores m_i as the result of projecting the i th qubit in the Z basis.
- If $k_i = 1$, the i th qubit of the clock state was effectively measured in the X basis. In this case, A stores

$$m_i = c_i \oplus d_i \cdot (x_{i(0)} \oplus x_{i(1)}), \quad (\text{C.3})$$

as the result of projecting the i th qubit in the X basis. Here, $y_1(0, x_{i(0)}) = y_1(1, x_{i(1)}) = \bar{y}_i$. Note that in order to compute $x_{i(0)}$ and $x_{i(1)}$, A needs to use the trapdoor.

Summarizing, for each (k_1, k_2) , this protocol provides A with a pair of bits (m_1, m_2) . By construction, the random variable (m_1, m_2) has the same statistics⁸ as the measurement outcomes resulting from projecting the qubits of a quantum state (in case of an honest prover the state $|\eta\rangle$) in the basis associated to (k_1, k_2) . Note that the measurement basis is kept secret⁹ from B. This allows A to estimate the expectation values of XZ -type operators corresponding to the state held by B without allowing him to cheat (see figure 2(a)). Moreover, those expectation values can be used to determine the energy with respect to H (see figure 2(b)).

Appendix D. The XZ -type $\log(n)$ -local Hamiltonian and its properties

In this section we present details on the construction of the Hamiltonian corresponding to a general decision problem and the bounds on the ground state energy for the Hamiltonian given in equation (C.1) and discuss the instance where B's answer is 'no'.

More precisely, in appendix D.1, we give the details of the construction of the Hamiltonian associated to a general decision problem given by a circuit \mathcal{C} acting on n qubits. One defines the output of the problem to be 'yes' if $p_0(\mathcal{C}) > b$ and 'no' if $p_0(\mathcal{C}) < a$, where $p_0(\mathcal{C}) = |\langle 0^n | \mathcal{C} | 0^n \rangle|^2$ and $0 \leq a < b \leq 1$. Considering a uniformly generated family of circuits \mathcal{C}_n acting on n qubits, one requires $b - a > 1/\text{poly}(n)$. At the end of this subsection we derive the Hamiltonian for the simple example considered in the main text in case B claims that the answer of the decision problem is 'yes'.

In appendices D.2 and D.3 we focus on the Hamiltonian associated to our example. In appendix D.2 we present the bounds on the ground state energy for the Hamiltonian considered in the main text. Finally, in D.3, we explain how this Hamiltonian must be modified in case B claims that the answer of the problem was 'no'.

D.1. The XZ -type $\log(n)$ -local Hamiltonian

In this section we present the details of the construction of the Hamiltonian H mentioned in the main text associated to an arbitrary circuit $\mathcal{C} = U_T \cdots U_2 U_1$ in case B claims that the answer of the decision problem was 'yes'. This construction is essentially the same as the one presented in [27]. However, here our main concern is not the locality of each term in the Hamiltonian (as shown in the latter reference, it can be made 2-local) but rather to ensure that the total number of required qubits is kept small. As explained below, to achieve this we will use (a) Gray codes [36] and (b) a universal gate set where all the gates are selfadjoint. For a circuit $\mathcal{C} = U_T \cdots U_2 U_1$ acting on n qubits, the Hamiltonian presented in [27] $H = H(\mathcal{C})$ is acting on $n + \lceil \log(T+1) \rceil$ qubits, with $\lceil \cdot \rceil$ the ceiling function. It can be expressed as

$$H = H_{\text{out}} + J_{\text{in}} H_{\text{in}} + J_{\text{prop}} H_{\text{prop}}, \quad (\text{D.1})$$

⁸ In fact, this is only true if one uses post-quantum secure trapdoor claw-free functions [22] and the prover is accepted in a test round with probability 1. In case this probability is only close to 1, one can show that the statistics of the measurement outcomes obtained by A are *close enough* to those of an actual quantum state. In the general case, this is sufficient to prevent B from cheating (see appendix E).

⁹ Again, this is true only when considering the family of functions described in [15]. There it is shown that the labels associated to one-to-one and two-to-one functions are computationally indistinguishable even for a quantum computer (if the problem Learning with Errors is hard for a quantum computer, which is widely believed to be the case [22]).

where J_{in} and J_{prop} are some suitable polynomials of n . The positive semidefinite operators H_{in} , H_{prop} and H_{out} are called input, propagation and output Hamiltonians, respectively, and will be explicitly presented below.

The additional $\lceil \log(T+1) \rceil$ qubit register allows us to encode $T+1$ orthogonal quantum states, representing the time steps. Using Gray's code, we write $|t\rangle$ for each of the $T+1$ orthogonal states such that the representation of two successive values of t differ only in one bit, i.e. they are given by a Gray code like, for instance, $(0, 1, 2, \dots) = (000, 001, 011, \dots)$. The following expressions, which will be useful in describing the Hamiltonians H_{out} , H_{in} and H_{prop} , can be written as products of $\log(n)$ operators that are either X or Z :

$$C(t) = |t\rangle\langle t|, \quad C(t, t-1) = \frac{1}{2}|t\rangle\langle t-1| + \frac{1}{2}|t-1\rangle\langle t|. \quad (\text{D.2})$$

In fact, the $\lceil \log(T+1) \rceil = O(\log(n))$ -local input and output Hamiltonians are given by

$$\begin{aligned} H_{\text{in}} &= \sum_{i=1}^n \frac{1}{2} (\mathbf{1} - Z_i) \otimes C(0), \\ H_{\text{out}} &= (T+1) \frac{1}{2} (\mathbf{1} - Z_1) \otimes C(T), \end{aligned} \quad (\text{D.3})$$

where the first and second factors in the tensor products act in the Hilbert space of the n computational qubits and the $\lceil \log(T+1) \rceil$ qubits encoding the clock states, respectively.

The 2-local Hamiltonian H_{prop} introduced in [27] can be written as a sum of products of only X and Z operators by using a gadget introduced in [29]. Importantly, the resulting Hamiltonian is still 2-local. However, this comes at the price of introducing more ancillary qubits. As we show now, using a Gray code and a universal set of self-adjoint gates leads to a Hamiltonian which is no longer 2-local. However, the number of auxiliary systems is reduced. The result is a XZ -type Hamiltonian that is $\log(n)$ -local (instead of just 2-local).

Without loss of generality [29, 37], we can assume that the circuit $\mathcal{C} = U_T \cdots U_2 U_1$ is written as a sequence of gates U_i that are either (i) 1-local and of the form $U(\alpha) = \cos \alpha Z + \sin \alpha X$ or (ii) CNOT acting on any pair of qubits. In this case, the $O(\log(n))$ -local propagation Hamiltonian is given by

$$\begin{aligned} H_{\text{prop}} &= \sum_{t=1}^T H_{\text{prop}}(t), \\ H_{\text{prop}}(t) &= \frac{1}{2} \mathbf{1} \otimes C(t) + \frac{1}{2} \mathbf{1} \otimes C(t-1) - U_t \otimes C(t, t-1). \end{aligned} \quad (\text{D.4})$$

One can easily verify that the Hamiltonian H constructed in this way can be expressed as a sum of products of only X and Z operators, as desired.

Finally, let us discuss the clock state $|\eta\rangle$ associated to this Hamiltonian and its relation to its lowest eigenvalue $\lambda(H)$. The clock state is given by [27]

$$|\eta\rangle = \frac{1}{\sqrt{T+1}} \sum_{t=0}^T U_t \cdots U_2 U_1 |0\rangle \otimes |t\rangle, \quad (\text{D.5})$$

and it is easy to see that $p_0(\mathcal{C}) = 1 - \langle \eta | H | \eta \rangle$. Hence $\lambda(H) \leq \langle \eta | H | \eta \rangle = 1 - p_0(\mathcal{C})$. It is shown in [27] that $\lambda(H) > \langle \eta | H | \eta \rangle - 1/4$ if one chooses appropriate values $J_{\text{in}} = \text{poly}(n)$ and $J_{\text{prop}} = \text{poly}(n)$ for the coefficients in the Hamiltonian. Thus, if the answer of the problem was 'yes', $\langle \eta | H | \eta \rangle = 1 - p_0(\mathcal{C}) < 1 - b$ (since $p_0(\mathcal{C}) > b$ in this case) and thus $\lambda(H) < \langle \eta | H | \eta \rangle < 1 - b$ as well. If the answer of the problem was 'no', then we have $\lambda(H) \geq \langle \eta | H | \eta \rangle - 1/4 = 1 - p_0(\mathcal{C}) - 1/4 = 3/4 - a$ (since $p_0(\mathcal{C}) < a$ in case the answer was 'no').

As mentioned above, in [27] a two-local Hamiltonian $H^{(\text{two-loc})}$ is presented which has similar properties to H but is acting on more qubits. There, the coefficients in $H^{(\text{two-loc})}$ can be chosen such that $\lambda(H^{(\text{two-loc})}) < 1 - b$ in case the answer was 'yes' and $\lambda(H^{(\text{two-loc})}) > 1/2 - a$ in case the answer was 'no' [27].

In the main text and in what follows, we will choose $J_{\text{in}} = 6$ and $J_{\text{prop}} = 3$ (as we are considering here a Hamiltonian acting on a fixed number of qubits). One can see that for our simple example the values $J_{\text{in}} = 6$ and $J_{\text{prop}} = 3$ are sufficient to establish (D.6) while keeping the trace of the Hamiltonian small enough. This is important because the larger this trace, the larger the impact of errors in the noisy estimation (see the error model given in appendix B). In the next section we explain how, using (D.6), the ground-energy of the Hamiltonian (C.1) can be used to encode the original promise problem.

D.2. Bounds on the ground state energy for the example considered here

As explained above, the Hamiltonian corresponding to the circuit \mathcal{C} encodes the answer to the decision problem [27]. Here we give the details for our particular example, where the circuit $\mathcal{C} = U(\alpha)$ is given in equation (1), the Hamiltonian $H = H_{\text{out}} + 6H_{\text{in}} + 3H_{\text{prop}}$ is given in equation (C.1), and the clock state is given in equation (2).

For our simple Hamiltonian, one can explicitly check (since it is just a matter of numerical diagonalization of a 4×4 matrix) that

$$\lambda(H) > \langle \eta | H | \eta \rangle - 2/5. \quad (\text{D.6})$$

In our particular case, we take the values $(a, b) = (1/10, 3/5)$. So that, if the answer of the problem was ‘yes’, $\langle \eta | H | \eta \rangle = 1 - p_0(\mathcal{C}) < 1 - b = 0.4$. In case the answer of the problem was ‘no’, we have $\langle \eta | H | \eta \rangle = 1 - p_0(\mathcal{C}) > 1 - a = 0.9$ (since $p_0(\mathcal{C}) < 0.1$ in case ‘no’). Now, using the latter inequality and equation (D.6), it follows that $\lambda(H) > \langle \eta | H | \eta \rangle - 2/5 > 1 - a - 2/5 = 3/5 - a = 0.5$ in case the answer of the problem was ‘no’. Summarizing, for our particular promise problem, we have

$$\begin{cases} \lambda(H) < \frac{2}{5} = 0.4, & \text{if “yes”}, \\ \lambda(H) > \frac{1}{2} = 0.5, & \text{if “no”}. \end{cases} \quad (\text{D.7})$$

Recall that we consider here the case where the prover’s output is ‘yes’. We deal in the next subsection with the case in which he outputs ‘no’.

D.3. Details of the case in which B claims that the answer is ‘no’

As mentioned in the main text, one can assume without loss of generality, that the prover B claims that the answer to the problem associated to the circuit \mathcal{C} is ‘yes’. The reason for that is that in case he claims ‘no’ for the circuit \mathcal{C} , this is equivalent to the case where he claims ‘yes’ for the modified circuit $\mathcal{C}' = X\mathcal{C}$. Here we explain in detail the corresponding Hamiltonian for our particular case.

In case B claims that the answer of the promise problem associated to \mathcal{C} is ‘no’, A constructs a Hamiltonian $H^{(\text{no})}$ in the same way (cf D.1) but now associated to the circuit $X\mathcal{C}$. Since the circuit always acts on the fixed initial state $|0\rangle$, one can add an initial Z gate and consider w.l.o.g. in the ‘no’ case the circuit $X\mathcal{C}Z$. Since $\mathcal{C} = U(\alpha_0)$, one obtains $X\mathcal{C}Z = U(\pi/2 - \alpha_0)$.

Hence, the situation in which B claims ‘no’ for $U(\alpha_0)$ is equivalent to the situation in which he claims ‘yes’ for $U(\pi/2 - \alpha_0)$. For this reason, we study the problem in which the promise is that α belongs to $I_1 \cup I_2$, where $I_1 = [0, \arcsin \sqrt{1/10}]$ and $I_2 = [\arcsin \sqrt{9/10}, \pi/2]$. So that both α and $\pi/2 - \alpha$ are possible values of our parameter.

A value $\alpha_0 \in I_1 \cup I_2$ is given and a description of the circuit $\mathcal{C} = U(\alpha_0)$ is sent to B. Then the prover B can run this circuit and measure $p_0(\mathcal{C})$ to decide whether the answer is either ‘yes’ or ‘no’. Once B sends his claim (either ‘yes’ or ‘no’) to A, the Hamiltonian H that A constructs is


$$H = \begin{cases} H^{(\text{yes})} = H(\mathcal{C}) = H(U(\alpha_0)), & \text{if B claims “yes”}, \\ H^{(\text{no})} = H(X\mathcal{C}Z) = H(U(\pi/2 - \alpha_0)), & \text{if B claims “no”}, \end{cases}$$

where $H(U)$ denotes the Hamiltonian of appendix D.1 associated to circuit U . Due to that, the results presented in the main text also show that for any $\alpha = \pi/2 - \alpha_0$, where α_0 is in the interval for which the experimentally estimated energy is sufficiently low (see figure 1), the same conclusion as for α_0 can be drawn.

Appendix E. Sketch of the fully-secure protocol

The aim of this section is to present more details on the main ingredients, which are required for a completely secure protocol, and to discuss main differences to the simplified version presented in this work.

In order to have full security, one uses [15] a family of two-to-one and one-to-one trapdoor functions that are hard to invert even for a quantum computer. Importantly, this family of functions needs to satisfy certain additional technical requirements. First, the two-to-one functions in the family need to have two hardcore bit properties (see [15] for details). Roughly speaking, for a two-to-one function $f: \{0, 1\}^m \rightarrow \{0, 1\}^{m'}$ with the hard-core bit property, the following problem is hard: given $x_0 \in \{0, 1\}^m$ and $f(x_0) = \bar{y} \in \{0, 1\}^{m'}$, find a bit-string $d \in \{0, 1\}^m$ such that $d \cdot (x_0 \oplus x_1) = 0 \pmod{2}$, where $f(x_0) = f(x_1) = \bar{y}$. Note that this property is crucial to ensure security (see [15]). Second, the one-to-one and



the two-to-one functions in the family must *look alike*. That is, it must be computationally hard to decide whether a function in the family is two-to-one or not [15]. As it is unknown how to construct such a family of functions with the required properties, a family of functions that fulfills those requirements not always but with high probability was used in [15].

This family of functions is used in Mahadev's measurement protocol [15] for both the test and the measurement round. The verifier A decides to run each of these rounds with equal probability. In both of these rounds, A can reject B's answer. To this end, A uses the trapdoor information and checks whether the preimages of the measurement outcomes \bar{y}_i exist (see figure E.1 for further explanations and the experimental data).

As explained in the main text, the protocol is used to delegate X- and Z-basis measurements in a, for B indecipherable way. The statistics of X- and Z-basis measurements allows to compute the energy of the state with respect to the Hamiltonian H , corresponding to an arbitrary decision problem, given in equation (C.1). Its ground-energy $\lambda(H)$ encodes the answer to the problem (cf equation (D.7)). In order to determine the energy with respect to H , Mahadev uses the protocol presented in [38], which we recall here. The Hamiltonian is first written as a convex combination (up to some re-scaling) of projectors, as explained in the following. For a 2-local Hamiltonian $H = \sum_{l=1}^L c_l P(l)$ with $P(l) = P_{i(l)}(l) P_{j(l)}(l)$, where $P_i(l) \in \{X_i, Z_i, 1_i\}$, i, j denote the qubits the operator is acting on and $L = \text{poly}(n)$ [27], the Hamiltonian $H' = (1 + H/c)/2$, where $c = \sum_l |c_l|$ is defined. The new Hamiltonian H' is a convex combination of projectors of the form $(1 + s(l)P(l))/2$ with weights $|c_l|/c$, where $s(l) = \text{sign}(c_l)$. In order to determine the energy, A samples, with probability $|c_l|/c$ a term $P(l) = P_{i(l)}(l) P_{j(l)}(l)$, which she then measures. Let $(m_i(l), m_j(l))$ denote the measurement outcomes obtained from projecting qubits $i = i(l)$ and $j = j(l)$ in the eigenbasis of $P_i(l)$ and $P_j(l)$, respectively. We use the notation:

$$m(l) = \begin{cases} (-1)^{m_i(l)+m_j(l)}, & \text{if } P_i(l) \in \{Z_i, X_i\}, \quad P_j(l) \in \{Z_j, X_j\}, \\ (-1)^{m_i(l)}, & \text{if } P_i(l) \in \{Z_i, X_i\}, \quad P_j(l) = 1_j, \\ (-1)^{m_j(l)}, & \text{if } P_j(l) \in \{Z_j, X_j\}, \quad P_i(l) = 1_i, \\ 1, & \text{if } P_i(l) = 1_i, \quad P_j(l) = 1_j. \end{cases} \quad (\text{E.1})$$

Now, A uses the following rule to determine a final bit $r(l) \in \{0, 1\}$:

$$r(l) = \begin{cases} 1, & \text{if } m(l) = s(l), \\ 0, & \text{if } m(l) = -s(l). \end{cases} \tag{E.2}$$

Observe that if B could send a state to A, and A could perform those measurements by herself, the expected value $\langle r \rangle$ of the random variable r coincides with the energy, with respect to H' , of the state sent by B [38]. Hence, after repeating the measurement protocol $N = \text{poly}(n)$ times obtaining $\{r_1, \dots, r_N\}$, A can compute an estimate r^{est} of the expectation value $\langle r \rangle$ of the random variable r . She then ‘accepts’ if $r^{\text{est}} \leq T_0$, with

$$T_0 = (1 + f(H) / c) / 2, \tag{E.3}$$

where $f(H)$ denotes the upper bound on $\lambda(H)$ given in appendix D.1 in case the answer was ‘yes’. For the example considered in the main text we have $T_0 = (1 + 0.4/c)/2$ using equation (D.7). In case the answer to the problem was ‘no’, one has $r^{\text{est}} \geq T_1$, with

$$T_1 = (1 + g(H) / c) / 2. \tag{E.4}$$

where $g(H)$ denotes the lower bound on $\lambda(H)$ given in appendix D.1 in case the answer was ‘no’. For the example considered in the main text we have $T_1 = (1 + 0.5/c)/2$ using equation (D.7).

As the two bounds, T_0, T_1 differ by $1/\text{poly}(n)$, one can run an extended protocol, using polynomially many copies N of the state to differentiate between the two cases (‘yes’ and ‘no’) with an exponentially small error (see [30] and protocol 8.3. in [15]).

Note that for convenience in our experiment we do not estimate $\langle H' \rangle$ by sampling from the probability distribution $|c_l|/c$, but determine $\langle H \rangle$ instead. This amounts to distinguishing energies below 0.4 and above 0.5 (cf figure 2). In the protocol discussed here, one needs to distinguish quantities below $0.4/c$ and above $0.5/c$. However, using several repetitions of the protocol would allow us to distinguish the two cases. In particular, in our example, for $\alpha = 0.12\pi/2$ one has $c = 12.4631$.

A classical verifier A uses Mahadev’s measurement protocol as follows to delegate the previous measurements to B for the extended protocol.

- First, she randomly chooses $N = \text{poly}(n)$ operators P_1, \dots, P_N occurring in the Hamiltonian H' independently with probability $\{|c_l|/c\}_l$. For each choice, l , she defines a vector $h^l \in \{0, 1\}^n$ by setting

$$h_i^l = \begin{cases} 0, & \text{if } P_i \in \{Z_i, 1_i\}, \\ 1, & \text{if } P_i = X_i. \end{cases}$$

Note that $h_i^l = 0$ for any qubit i on which P_l acts trivially. We call the vector $h = (h^1, h^2, \dots, h^N)$ now the basis choice.

- A and B run the measurement protocol explained in the main text for the basis choice h . Let $p_{t,h}$ ($p_{m,h}$) denote the probability that at least one of the tests in the test (measurement) round failed.
- In the measurement round the verifier computes the product of the measurement results for each term $P(l)$ and sets $r(l) = 1$ only if for more than half on the times the product of the measurement results coincides with $s(l)$.

Given $r(l)$, A computes the expectation value of r and thereby the expectation value of H' . As shown in [15] the probability with which A accepts the answer of B is given by

$$P_{\text{accept}} = \frac{1}{2} \sum_h v_h (1 - p_{t,h}) + \frac{1}{2} \sum_h v_h (1 - p_{m,h}) \text{Prob}_h (r^{\text{est}} < T_0). \tag{E.5}$$

Here, v_h denotes the probability with which A samples the basis choice h (depending on the Hamiltonian). Moreover, $\text{Prob}_h (r^{\text{est}} < T_0)$ denotes the probability with which $r^{\text{est}} < T_0$ in case the basis choice h was used for the measurements.

Note that, in the absence of noise, an honest prover B, who would simply prepare N copies of the $|\eta\rangle$ -state, would be accepted with probability exponentially close to 1 [15]. Importantly, in [15], Mahadev showed that A can differentiate between such a honest prover and a dishonest prover. The reason for that is that the probability for accepting a dishonest prover is upper bounded by $3/4$ even for non-vanishing $p_{t,h}$

and $p_{m,h}$ (soundness). Repeating the protocol $\text{poly}(n)$ many times the verifier can distinguish between the cases where the outcome of the problem was ‘yes’ or ‘no’.

As emphasized before, implementing this protocol is highly demanding. This is not only because it requires many auxiliary qubits for each qubit to be measured, but also due to the stringent quality requirements for an honest quantum computer to successfully pass the test rounds. While a detailed study of the influence of realistic noise on the fully-secure version of this protocol remains an open challenge, there has been progress in verifying quantum computations in the presence of noise when users have limited access to quantum devices. In [9], such a protocol is developed within the framework of fault-tolerant MBQC. In this context, Bob sends potentially noisy cluster states to Alice, who then performs the desired computation by measuring some of the qubits. In [21] the authors demonstrated that, under reasonable error models, efficient verification can be achieved.

Let us also finally comment on other applications of interactive proofs using post-quantum secure cryptographic functions. They can be used to verify quantumness and quantum advantage [16, 17], which has been recently demonstrated experimentally [26]. A can use the above ideas to get convinced that B possesses a quantum state and is using it to solve certain computational task efficiently, which would have been impossible for a classical machine. In particular, A can send B the labels k of two-to-one trapdoor claw-free functions $F_k : \{0, 1\}^m \mapsto \{0, 1\}^{m'}$ and ask him to prepare the state

$$\frac{1}{\sqrt{2^m}} \sum_{x=0}^{2^m-1} |x\rangle |F_k(x)\rangle.$$


Then, A asks B for a commitment string $\bar{y} \in \{0, 1\}^{m'}$ in the range of F_k . This can be easily provided by a quantum B. He just needs to measure the last m' registers. Then, B would hold the superposition

$$\frac{1}{\sqrt{2}} |x_0\rangle + \frac{1}{\sqrt{2}} |x_1\rangle,$$

where $F_k(x_0) = F_k(x_1) = \bar{y}$. Recall that B does not know the preimages, x_0, x_1 . However, A can easily compute them knowing the trapdoor information. Now, A asks B to measure the remaining m qubits either in the Z - or X -basis and to send the result to A. We denote by p_A the probability that B sends a preimage of \bar{y} in the first case. An honest B would just measure the first m registers in the Z -basis and obtain a bit-string x_b such that $F_k(x_b) = \bar{y}$, where $b = 0$ or $b = 1$ with probability $1/2$. Let p_B denote the probability that B sends $d \in \{0, 1\}^m$ such that $d \cdot (x_0 \oplus x_1) = 0 \pmod{2}$ with $F_k(x_0) = F_k(x_1) = \bar{y}$, in the second case. It can be easily seen that, in this case, an honest B would just measure the first m registers in the X -basis and obtain a string d that would be accepted. Note that in case of X -basis measurements, the hard-core bit property discussed above comes into play. Recall that the function F_k has the hard-core bit property if, given $x_0 \in \{0, 1\}^m$ and $F_k(x_0) = \bar{y} \in \{0, 1\}^{m'}$, it is hard to find a bit-string $d \in \{0, 1\}^m$ such that $d \cdot (x_0 \oplus x_1) = 0 \pmod{2}$, where $F_k(x_0) = F_k(x_1) = \bar{y}$. Note that, in contrast to a classical device, an honest quantum B would be able to obtain such d without knowing x_0 or x_1 . Using these properties, one can derive interactive proof protocols based on the assumption that the problem Learning with Errors is hard [22] such that $p_A + 2p_B \leq 2$ for the best classical strategy (for m, m' sufficiently large). This means that such a protocol can be used by A to verify the ‘quantumness’ of B and even quantum advantage in case m, m' are large enough [16, 17, 26].

Appendix F. Entangling gate error estimation

In the following we present the technical details on the MS-gate fidelity discussion from appendix B. To characterize the entangling gates, we initialize a GHZ-state for each ion-pair through a series of $(4n + 1)$ with $n \in \{0, 1, 2, 3, 4\}$ full-entangling $\text{MS}^{X_{i,j}}(-\pi/2)$. Here, we use that the GHZ-state’s density matrix ideally consists of only four elements, namely two diagonal terms $|00\rangle$ and $|11\rangle$ referred to as population as well as two off-diagonal ones describing relative coherences. The population can be directly inferred from fluorescence detection of the measured population in p_{00} and p_{11} , whereas the coherence terms are extracted from the contrast in parity oscillations. Averaging population and parity leads to the GHZ-state fidelity. Notably, we implement each series by introducing a pause at the duration of typical local operations after each MS-gate to both mimic realistic experiment conditions, where entangling gates are interleaved with local ones. Results on all gate pairs from the representative $X_1 X_2$ circuit are shown in figure F.1 and thoroughly discussed in appendix B. Note that, qubit orders differ from the circuits in figure B.1 as to maximize inter-ion spacing and by that reduce the influence of cross-talk.



Appendix G. A different decision problem

Here we give the details of our experimental results for a different, but related, decision problem. Specifically we consider the problem where the answer is ‘yes’ for α close to $\pi/2$ in the same circuit $\mathcal{C} = U(\alpha)$ as in the main text. Crucially, from a computer science perspective this case is completely equivalent to the case discussed in the main text. In practice, however, different decision problems, even when they are associated to the same circuit, correspond to different implementations at the hardware level and might thus exhibit different noise sensitivity. For this reason, it might be interesting to verify multiple instances on a given quantum device. To confirm this behaviour, we now consider the circuit $\mathcal{C} = U(\alpha) = \cos \alpha Z + \sin \alpha X$, but defining the answer of the problem to be

$$\begin{cases} \text{“yes”}, & \text{if } p_1(\mathcal{C}) > b, \\ \text{“no”}, & \text{if } p_1(\mathcal{C}) < a, \end{cases} \tag{G.1}$$

under the promise that one of the two cases occurs, where $p_1(\mathcal{C}) = |\langle 1|\mathcal{C}|0\rangle|^2$. The corresponding Hamiltonian $H = H_{\text{out}} + J_{\text{in}}H_{\text{in}} + J_{\text{prop}}H_{\text{prop}}$ (see appendix D.1) must thus be changed such that H_{out} penalizes states with the output-qubit in state $|0\rangle$, such that

$$H_{\text{out}} = (T + 1) \frac{1}{2} (\mathbf{1} + Z_1) \otimes C(T).$$


Assuming now, without loss of generality (see appendix D.3), that B claims that the answer to the decision problem was ‘yes’, the Hamiltonian $H = H_{\text{out}} + 6H_{\text{in}} + 3H_{\text{prop}}$ now reads

$$\begin{aligned} H_{\text{out}} &= \frac{1}{2} (1 + Z_1 - Z_2 - Z_1Z_2), \\ H_{\text{in}} &= \frac{1}{4} (1 - Z_1 + Z_2 - Z_1Z_2), \\ H_{\text{prop}} &= \frac{1}{2} (1 - \cos \alpha Z_1X_2 - \sin \alpha X_1X_2), \end{aligned} \tag{G.2}$$

to be compared with the Hamiltonian of equation (C.1) in the main text. Analogously to the case in the main text, one can show that in this case $\langle \eta|H|\eta\rangle = 1 - p_1(\mathcal{C})$ and $\lambda(H) > \langle \eta|H|\eta\rangle - 2/5 = 3/5 - p_1(\mathcal{C})$ hold (see appendix D.2). This implies that equation (D.7) holds for the Hamiltonian (G.2) and the decision problem (G.1).

Figure G.1 shows the experimental results in this case, again choosing $(a, b) = (1/10, 3/5)$ so that the thresholds of equation (D.7) remain $1 - b = 0.4$ and $3/5 - a = 0.5$. Curiously, despite being formally equivalent to the case in the main text, verification turns out to be slightly more challenging for this problem. This goes to show, that the protocol indeed verifies the output of a device, not the device itself. Hence, just because the protocol successfully verifies one instance, does not mean that all instances can be verified.

A closer inspection of the underlying circuits show that the case $\alpha \rightarrow \pi/2$ generates more entanglement in the system compared to $\alpha \rightarrow 0$. Experimentally, this amplifies the noise in an unfavourable way to prevent



verification for most values of α in this case. One exception is the instance with $\alpha = \pi/2$ which features an energy below the verification threshold for two reasons. First, the term given in the above Hamiltonian proportional to $\cos\alpha Z_1 X_2$ exactly vanishes at $\alpha = \pi/2$. Second, the basis change operation rotating CNOT into CPHASE becomes trivial for $\alpha = \pi/2$.

Error-bars in figure G.1 represent one standard deviation of quantum projection noise, estimated from sampled energy values that we find to be normally distributed around the plotted mean values. These statistical errors are used to calculate the significance level as the probability that the verifier, A, incorrectly accepts the result. Therefore, significance levels are determined by the one-tailed probability that an energy value results above the threshold 0.4 inferred from the normal distribution of sampled energy values. For the instance at $\alpha = \pi/2$ we receive a significance level of 0.425.


We further note that the same depolarizing noise model used in the main text (see section I) also accurately describes the data presented in figure G.1 here. Numerical simulations then suggest that about a 30% reduced depolarizing rate of $\lambda = 0.035$ would be required to verify this case. This demonstrates that there can be large differences in the verifiability of different instances of the same problem on the same hardware.

Appendix H. Direct estimation of η -state energy

For previous attempts, covered in figures 2(b) and G.1(b), the η -state's energy was estimated using the six auxiliary qubits necessary for the trapdoor function implementation to thereby enable classical verification. For comparison, we follow up on the direct estimation of the η -state's energy. To this end, we implement the sub-circuit in the grey box from figure B.1(a) alongside additional operations on prover and clock qubit required to realize basis read-outs according to $Z_1 Z_2$, $Z_1 X_2$, $X_1 Z_2$ and $X_1 X_2$. Figure H.1 contains results on operator values (a) as well as energies covering both decision problems (b) and (c) considering the classical verification of $\alpha \rightarrow 0$ (see main text) and $\alpha \rightarrow \pi/2$ (see appendix G) respectively. In both cases our experimental results undercut the threshold $1 - b = 0.4$ across the relevant region $\alpha \in [0, 0.32] \cup [1.25, 1.57]$, for which our decision problems hold. Comparably good results are obtained due to the less complex experiment using only one full-entangling MS-gate plus on average eight single-qubit gates. Note that, these experiments were likewise performed on an eight-ion string using 400 experimental runs in each data point. Errors represent 1 standard-deviation from quantum projection noise.

We continue to estimate the η -state fidelity by incorporating error-rates inherent to the individual gates as previously done and thoroughly explained in the bottom part of appendix B—there considering the entire protocol. Here, the expected fidelity on the direct estimation reads:

$$\mathcal{F} = 0.998(1)^8 \cdot 0.984(1) = 0.968(8).$$



We compare this number to results from figure H.1(a) by averaging Z_1Z_2 and X_1X_2 at $\alpha = \pi/2$ representing an estimate of the η -state’s Bell-state fidelity. The result reads:

$$\mathcal{F} \sim \frac{Z_1Z_2(\alpha = \pi/2) + X_1X_2(\alpha = \pi/2)}{2} = \frac{0.955(16) + 0.935(17)}{2} = 0.945(12),$$

and is again in good agreement with the above simplistic error-modelling.

Appendix I. Noise model simulations

This section aims to elaborate an error-model, which best describes the experimental data from the measurement rounds depicted in figure 2 from the main text. Choosing a suitable error-model was done upon previous error-rate observations thoroughly discussed at the bottom of appendix B. Those observations distinctly reveal pairwise MS-gates to limit the overall performance of our classical verification implementation. In contrast, single-qubit gates clearly make a smaller contribution, although having an approximately four times higher abundance in the final circuits. The analysis of individual MS-gate pairs, depicted in table B.1, discloses similar decay-rates in population and coherence—the latter characterizing the degree of loss in phase information. Hence, our findings support a simultaneous dephasing along X, Y and Z basis manifesting a so-called depolarizing channel. A fully depolarized state leads to a completely mixed state, which in the single qubit case is illustrated by shrinking the Bloch-sphere towards its center. Based on this, we worked out the following eight-qubit $\Gamma_\lambda = \Delta_\lambda^{\otimes 8}$ depolarizing channel to describe our classical verification results:

$$\Gamma_\lambda(\rho_{\text{ideal}}) = \Delta_\lambda^{\otimes 8}(\rho_{\text{ideal}}) \tag{I.1}$$

$$\Delta_\lambda(\sigma) = \sum_{l=0}^3 K_l \sigma K_l^\dagger$$

where Δ_λ are single-qubit depolarizing channels with Kraus operators


$$K_0 = \sqrt{1 - 3\lambda/4} \mathbf{1},$$

$$K_1 = \sqrt{\lambda/4} X,$$

$$K_2 = \sqrt{\lambda/4} Y,$$

$$K_3 = \sqrt{\lambda/4} Z.$$

Hence λ is the single-qubit depolarizing parameter. In figure I.1 we show the depolarizing channel using the best fitted rate given at $\lambda = 0.05$ represented by dashed-lines on top of the data discussed in the main text. The good agreement between data and noise channel confirms our limitation to be depolarizing noise. According to the channel in equation (I.1) an ideal outcome in the measurement round with respect to the eight-qubit density matrix ρ_{ideal} is expected at a probability of about $(1 - \lambda)^8 \approx 0.66$. The total energy plot from (b) additionally depicts the ground-state energy of the implemented Hamiltonian $H = H_{\text{out}} + 6H_{\text{in}} + 3H_{\text{prop}}$ with each term given by equation (C.1). Note that, for this ground-state energies,



i.e. the smallest numerical eigenvalues $\lambda_{\min}(H^{(\text{yes})})$, the inequality $\lambda_{\min}(H) > \langle \eta | H | \eta \rangle - 2/5$ holds across the entire α range, as discussed in the above section D.2.

Moreover, this identical noise-model ($\lambda = 0.05$) accurately images the experimental outcome on the extra decision problem presented in section G.

ORCID iDs

Roman Stricker  <https://orcid.org/0000-0001-8001-1487>

Martin Ringbauer  <https://orcid.org/0000-0001-5055-6240>

Philipp Schindler  <https://orcid.org/0000-0002-9461-9650>

References

- [1] Arute F *et al* 2019 Quantum supremacy using a programmable superconducting processor *Nature* **574** 505
- [2] Wu Y *et al* 2021 Strong quantum computational advantage using a superconducting quantum processor *Phys. Rev. Lett.* **127** 180501
- [3] Zhong H-S *et al* 2020 Quantum computational advantage using photons *Science* **370** 1460
- [4] Emerson J, Alicki R and Życzkowski K 2005 Scalable noise estimation with random unitary operators *J. Opt. B: Quantum Semiclass. Opt.* **7** 347
- [5] Gheorghiu A, Kapourniotis T and Kashefi E 2019 Verification of quantum computation: an overview of existing approaches *Theory Comput. Syst.* **63** 715
- [6] Elben A *et al* 2020 Cross-platform verification of intermediate scale quantum devices *Phys. Rev. Lett.* **124** 010504
- [7] Greganti C *et al* 2021 Cross-verification of independent quantum devices *Phys. Rev. X* **11** 031049
- [8] Eisert J, Hangleiter D, Walk N, Roth I, Markham D, Parekh R, Chabaud U and Kashefi E 2020 Quantum certification and benchmarking *Nat. Rev. Phys.* **2** 382
- [9] Hayashi M and Morimae T 2015 Verifiable measurement-only blind quantum computing with stabilizer testing *Phys. Rev. Lett.* **115** 220502
- [10] Gheorghiu A, Kashefi E and Wallden P 2015 Robustness and device independence of verifiable blind quantum computing *New J. Phys.* **17** 083040
- [11] Fitzsimons J F and Kashefi E 2017 Unconditionally verifiable blind quantum computation *Phys. Rev. A* **96** 012303
- [12] Aharonov D, Ben-Or M and Eban E and Mahadev U 2017 arXiv:1704.04487
- [13] Takeuchi Y, Morimae T and Hayashi M 2019 Quantum computational universality of hypergraph states with Pauli-X and Z basis measurements *Sci. Rep.* **9** 13585
- [14] Barz S, Kashefi E, Broadbent A, Fitzsimons J F, Zeilinger A and Walther P 2012 Demonstration of blind quantum computing *Science* **335** 303
- [15] Mahadev U 2018 Classical verification of quantum computations *IEEE 59th Annu. Symp. Found. Comput. Sci.* pp 259–67
- [16] Brakerski Z, Koppula V, Vazirani U and Vidick T 2020 Simpler Proofs of Quantumness *15th Conf. Theory of Quantum Computation, Communication and Cryptography (TQC 2020) (Leibniz Int. Proc. in Informatics (LIPIcs))* vol 158, ed S T Flammia pp 8:1–8:14
- [17] Brakerski Z, Christiano P, Mahadev U, Vazirani U and Vidick T 2021 A cryptographic test of quantumness and certifiable randomness from a single quantum device *J. ACM* **68** 1–47
- [18] Carrasco J, Elben A, Kokail C, Kraus B and Zoller P 2021 Theoretical and experimental perspectives of quantum verification *PRX Quantum* **2** 010102
- [19] Schindler P *et al* 2013 A quantum information processor with trapped ions *New J. Phys.* **15** 123012
- [20] Reichardt B W, Unger F and Vazirani U 2012 arXiv:1209.0448
- [21] Fujii K and Hayashi M 2017 Verifiable fault tolerance in measurement-based quantum computation *Phys. Rev. A* **96** 030301
- [22] Regev O and Lattices O 2005 Learning with Errors, Random Linear Codes and Cryptography *Proc. Thirty-Seventh Annual ACM Symp. on Theory of Computing, STOC'05* pp 84–93
- [23] Banerjee A, Peikert C and Rosen A 2012 Pseudorandom functions and lattices *Advances in Cryptology – EUROCRYPT 2012*, ed D Pointcheval and T Johansson (Springer) pp 719–37

- [24] Alwen J, Krenn S, Pietrzak K and Wichs D 2013 Learning with rounding, Revisited *Advances in Cryptology – CRYPTO 2013*, ed R Canetti and J A Garay (Springer) pp 57–74
- [25] Goldwasser S, Micali S and Rackoff C 1989 The knowledge complexity of interactive proof systems *SIAM J. Comput.* **18** 186–208
- [26] Zhu D *et al* 2021 Interactive protocols for classically-verifiable quantum advantage (arXiv:2112.05156 [quant-ph])
- [27] Kempe J, Kitaev A and Regev O 2005 The complexity of the local Hamiltonian problem *SIAM J. Comput.* **35** 1070
- [28] Feynman R P 1986 Quantum mechanical computers *Found. Phys.* **16** 507
- [29] Biamonte J and Love P 2008 Realizable Hamiltonians for universal adiabatic quantum computers *Phys. Rev. A* **78** 012352
- [30] Fitzsimons J F, Hajdušek M and Morimae T 2018 *Post hoc* verification of quantum computation *Phys. Rev. Lett.* **120** 040501
- [31] Kahanamoku-Meyer G D, Choi U V, Vazirani S and Yao N Y 2022 Classically verifiable quantum advantage from a computational Bell test *Nat. Phys.* **18** 918
- [32] Jacak M M, Józwiak P, Niemczuk J and Jacak J E 2021 Quantum generators of random numbers *Sci. Rep.* **11** 16108
- [33] Gheorghiu A and Vidick T 2019 Computationally-secure and composable remote state preparation (arXiv:1904.06320 [quant-ph])
- [34] Data underlying the work: towards experimental classical verification of quantum computation (available at: <https://doi.org/10.5281/zenodo.10091389>) (Accessed 09 November 2023)
- [35] Ringbauer M *et al* 2022 A universal qudit quantum processor with trapped ions *Nat. Phys.* **18** 1053–7 (arXiv:2109.06903 [quant-ph])
- [36] Press W, Teukolsky S, Vetterling W and Flannery B 2007 *Numerical Recipes 3rd edn: the Art of Scientific Computing* 3rd edn (Cambridge University Press)
- [37] Shi Y 2002 Both toffoli and controlled-NOT need little help to do universal quantum computation (arXiv:quant-ph/0205115 [quant-ph])
- [38] Morimae T, Nagaj D and Schuch N 2016 Quantum proofs can be verified using only single-qubit measurements *Phys. Rev. A* **93** 022326

QUBIT LOSS PROTECTION

Most of today’s quantum computer architectures are based on physical multi-level systems artificially constrained to the two levels forming the qubit. The presence of these extra levels leads to additional sources of errors outside the computational subspace, so-called leakage errors (Ch. 1.4), in which these additional levels are populated. In addition to leakage, qubits can become inoperable after chemical reactions or might get lost all together. We denote all these mechanisms as loss errors.

The dominant loss mechanism in trapped-ion devices is leakage, which occurs at rates comparable to computational errors [52]. Besides trapped ions, this is true for most of today’s quantum computer architectures [204–208]. Only in some cases, mechanisms that make the qubit inoperable, such as the physical loss of the particles encoding the qubits, become relevant on experimental timescales [109], thoroughly discussed in Ch. 1.4.

Despite the realistic chance of experiencing leakage, many conceptual ideas of quantum computers ignore the existence of additional levels and thus omit faulty mechanisms beyond the computational ones. A fact with consequences, particularly in view of QEC applications, typically correcting errors that change the logical state. Losses can therefore drastically deteriorate the underlying quantum information. On the contrary, if losses can be detected, many of the leading QEC protocols exhibit immens robustness against them [94, 209]—elaborated on below. However, to successfully exploit this robustness, loss detection and correction routines must be built in, usually by extra protocols.

In the upcoming chapter, we present approaches to implement the building blocks of qubit loss protection in terms of leakage. Particularly, we develop a *quantum non-demolition* (QND) loss detection unit in Sec. 5.1. Sec. 5.2 proceeds with the demonstration of loss correction on account of two different codes. The first code in Sec. 5.2.1 corrects losses based on quantum teleportation, called the erasure code [204]. While the erasure code allows fruitful proof-of-principle demonstrations on a full cycle of qubit loss detection and correction, it cannot correct computational errors. To make up for this, we then build in the possibility to correct computational errors on top by treating loss in the context of the surface code [77, 78]. The surface code stands out because of its comparably high computational error thresholds [102] and a hardware-friendly modular structure, while at the same time, yielding considerable thresholds against losses [94], explained in Sec. 5.2.2. We implement a surface code instance and present the first *deterministic* experiment to resolve in-sequence detected loss-events in real-time. This leads to the third publication of this thesis, presented in Sec.5.3.

The observation of non-unitary dynamics inherent in the semi-classical algorithm structure of our loss detection unit motivates a follow-up work on its correct tomographic reconstruction based on quantum instruments, see Ch. 1.5. In this way, we can capture the full dynamics of loss affected codes, which allows us to estimate the parameter regimes for simultaneously correctable loss and computational errors. Moreover, the tomographic

tools we develop reach beyond existing characterization approaches, which are useful for analyzing a broad class of quantum operations. This work leads to the fourth and final publication in this thesis, presented in Sec.5.5.

5.1 DETECTION OF QUBIT LOSS

To identify a leakage event on a code qubit during an ongoing quantum computation, we aim to transfer the loss information via entanglement to an ancilla qubit. The state of the ancilla qubit can then be destructively detected without affecting the underlying quantum information in QND fashion, see Ch. 1.3. In the absence of loss, such a detection unit must therefore leave both the code and the ancilla qubit unaltered, i.e., it must perform an identity operation on both, whereas we want to signal the presence of a loss by exciting the ancilla qubit.

The two-qubit circuit in Fig. 5.1 follows these ideas. While the bottom code qubit is subject to loss detection, the top ancilla qubit serves for its readout. The protocol is based on two quantum gates. The first, the entangling gate $MS_{a,c}^X(\pi)$ from Eq. (2.19) performs a correlated bit flip on both qubits. More precisely, half of the gate $MS_{a,c}^X(\pi/2)$ transfers an example input state $|00\rangle_{a,c}$ into the maximally entangled Bell-state $1/\sqrt{2}(|00\rangle + |11\rangle)_{a,c}$. The second half $MS_{a,c}^X(\pi/2)$ disentangles the GHZ-state into a product state with both qubits flipped to $|11\rangle_{a,c}$. This correlated bit flip works for arbitrary input states while both qubits have to be present in their computational subspaces. In this no-loss case, the subsequent local bit flip $R_{a,c}^X(\pi) = X_{a,c}$ reverses the first collective bit flip operation on both qubits, completing an overall identity operation. While the code qubit can represent any initial state, the ancilla qubit for loss detection is always prepared in $|0\rangle_a$. In case of a lost code qubit, however, the MS-gate targets only the ancilla qubit on which it must act trivially. This is evident from the exponential argument of the entangling gate, which simplifies to an identity operation $X_a X_a = 1$ when only the ancilla qubit is present in the computational subspace. Thereby, in case of loss on the code qubit, only the second local operation flips the ancilla into $|1\rangle_a$ and signals the loss.

Leakage can be induced in a controlled fashion by partially pumping the lower-energy population of the code qubit out of the computational subspace $\{4^2S_{1/2}(m = -1/2) = |0\rangle, 3^2D_{5/2}(m = -1/2) = |1\rangle\}$ via the carrier transition $R_c^{\text{loss}}(\theta, 4^2S_{1/2}(m = -1/2) \leftrightarrow 3^2D_{5/2}(m = -5/2))$. We proceed to call $R_c^{\text{loss}}(\theta)$ as loss operation with pulse area θ referring to a loss probability $p_{\text{loss}} = 1/2 \sin^2(\theta/2)$. It is noteworthy that we hold equivalent control over the entire Zeeman manifolds and could freely choose a leakage channel sourcing from either of the two qubit basis states, see Ch. 2.5.

To benchmark the loss detection capabilities, we test the unit over the full range of loss probabilities via $R_c^{\text{loss}}(\theta)$ with $\theta \in [0, \pi]$, before a final measurement on both qubits yields their individual population in the upper $3^2D_{5/2}$ -state manifold. Initially, the two-qubit system is prepared in its ground-state $|00\rangle_{a,c}$. Fig. 5.2(a) shows the resulting population correlation between *directly measured loss* on the code qubit and *detected loss* on the ancilla qubit. Error-bars denote one standard deviation of QPN according to Eq. (1.27). Under complete loss $R_c^{\text{loss}}(\pi)$, the ancilla qubit flips to $|1\rangle_a$, while it remains in $|0\rangle_a$ absent of loss $R_c^{\text{loss}}(0)$. A linear fit of the population correlation yields a loss detection efficiency over the readout of the ancilla qubit of 0.96(4).

While the above experiment confirms reliable population transfer between the code and the ancilla qubit, which is essential for loss detection, the measurements do not provide information about the effective unitary operation performed on the code qubit. For loss




Figure 5.1: **Qubit loss detection unit.** We deliberately induce a partial loss by resonantly coupling one of the two qubit states in resonance with a third level outside the computational subspace $\{4^2S_{1/2}(m = -1/2) = |0\rangle, 3^2D_{5/2}(m = -1/2) = |1\rangle\}$. Here, we choose the loss operation $R_c^{\text{loss}}(\theta, 4^2S_{1/2}(m = -1/2) \leftrightarrow 3^2D_{5/2}(m = -5/2))$ corresponding to a loss probability $p_{\text{loss}} = 1/2 \sin^2(\theta/2)$. Loss detection is based on two operations. The first is the entangling $MS_{a,c}^X(\pi)$ from Eq. (2.19) that implements a correlated bit flip, if and only if both qubits are present to the computational subspace, see text. In this no-loss case, the second local bit flip $R_{a,c}^X(\pi) = X_{a,c}$ reverses the first operation leaving both qubits unchanged. Under loss, the MS-gate acts trivially, where only the subsequent local $X_{a,c}$ -gate flips the ancilla qubit to $|1\rangle_{a'}$, signaling the loss. The protocol thus works in a QND fashion and structurally belongs to the class of semi-classical quantum algorithms, see 1.5.

detection, however, it is fundamental to leave the code qubit unchanged so as not to disturb the underlying logical information. For a more quantitative performance analysis in the no-loss case, we verify the target operation on the code qubit with QPT and MLE process reconstruction following Eq. (3.14). Fig. 5.2(b) depicts the resulting χ -matrix according to Eq. (3.9) decomposed into standard Pauli basis from Eq. (1.5) yielding a process fidelity with respect to the targeted identity operation of 0.90(2). For the experiment shown, a loss probability of 0.012 was induced from the state $|0\rangle_c$ before post-selecting the data by the no-loss case. Leakage can be expected at rates similar to those of computational errors [52], which explains the selected, relatively low loss probability. The errors represent one standard deviation of QPN from resampling the QPT data. The χ -matrix's elements in Z further point to coherent phase errors, most likely due to uncompensated AC-Stark shifts [139]. These may be caused by the high light power required for the MS-gate, which potentially phase-shifts all subsequent operations due to heating effects in the pulse-shaping acousto-optic modulators. It should be noted that restricting QPT to the qubit levels is no longer feasible in the loss case, since the process reconstruction method considers all constituents to be in their respective computational subspaces, see 5.5. In terms of loss protection, however, this case is of minor interest, since lost qubits are nonetheless inoperable.

Finally, both results in Fig. 5.2 demonstrate the successful detection of qubit loss making our approach a promising candidate for the upcoming loss correction studies. Note that the protocol is readily applicable to other quantum computer architectures. Let us finally remark that other recent publications present non-destructive loss detection schemes, e.g., for photonic devices [210] or superconducting platforms [211].




Figure 5.2: **Benchmarking the loss detection unit from Fig. 5.1.** (a) The results show the correlations between the induced loss on the code qubit and the detected loss on the ancilla qubit on a full range of loss probabilities. A linear fit yields a detection efficiency via the ancilla qubit readout of $0.96(4)$, see text for details. (b) QPT on the code qubit to reconstruct the χ -matrix according to Eq. (3.9) in the no-loss case that reproduces the expected identity operation with a process fidelity of $0.90(2)$. The exemplified experiment considers a loss probability of 0.012 induced from state $|0\rangle_c$.

5.2 CORRECTION OF QUBIT LOSS

We now examine two protocols for loss correction within the QEC framework, namely erasure [204] and surface code [94]. The erasure code only treats loss correction and has been investigated in several experimental studies so far, mostly using photonic hardware and with the final correction step performed in post-processing [206, 207]. Apart from proof-of-principle demonstrations, however, loss correction in post-processing is pointless, since loss robust implementations require in-sequence corrections in real-time. We introduce the erasure code's framework along its correction capabilities over the upcoming Sec. 5.2.1 and demonstrate experiments in combination with the QND loss detection unit described in the previous section. We then switch to the surface code that besides computational errors holds significant robustness against losses [94], if they can be detected, presented in Sec. 5.2.2.

5.2.1 Erasure code—extended four qubit teleportation

The erasure code, invented by Grassl et. al in Ref. [204] utilizes an extension of quantum state teleportation [212] for transferring a lost state to a still intact ancilla qubit. The authors of Ref. [204] further prove that a logical encoding with four qubits is the minimal instance for the correction of a single lost qubit at a known position. Fig. 5.3 depicts the code's circuit representation, separated into encoding and recoding (or correction) step, which considers the top code qubit has been lost.

Let us review how the code works. An arbitrary state $|\psi\rangle_1 = \alpha|0\rangle_1 + \beta|1\rangle_1$ on the code qubit gets initially encoded across four qubits. After losing one of the qubits involved, the logical information is stored across the three remaining ones. In the following we consider the top code qubit as the lost one. Following the concept of quantum state teleportation, the protocol then requires destructive measurements on the middle qubits two and three along Pauli Z_2 - and X_3 -basis, respectively. The middle qubits thereby serve to store amplitude (Z_2 -measurement) and phase (X_3 -measurement) information of the loss protected state $|\psi\rangle$. Their combined measurement outcome then refers to one of four unitary operations which, applied to the bottom qubit, restores the lost state $|\psi\rangle_4$. The corresponding correction




Figure 5.3: **Erasure code for qubit loss correction.** **Left box:** Four-qubit logical encoding for correcting a single loss at a known position [204]. The given alignment accounts for a loss on the top code qubit and creates redundancy by incorporating the three additional ancilla qubits below. **Right box:** Measurements on qubits two and three along the Z_2 and X_3 bases, respectively, provide an outcome-dependent unitary operation, which applied to the bottom qubit, retrieves the lost state $|\psi\rangle_4$. The correction operation for all four outcomes is depicted in Tab. 5.1. The protocol’s functioning can be seen as an extended four-qubit state teleportation [212].

unitaries are summarized in Tab. 5.1. Crucially, the protocol requires the identification of loss-events prior to correction, e.g., using the QND detection unit from the previous section that requires an additional ancilla qubit.

While the four-qubit code conceptually corrects a single lost qubit at a known position, its small size can never tolerate realistic loss error-rates. To improve robustness against losses, this minimal fragment can be generalized to an n -fold concatenated code operating on the order of 4^n qubits. Let us have a closer look at this. Multiple blocks of the minimal four-qubit code instance from Fig. 5.3 are initially prepared in parallel, where each fragment outputs a single loss-corrected qubit. In the second layer, these loss-corrected qubits are used to form another set of blocks of the minimal four-qubit code instance and again yield one corrected qubit each block. This concatenation continues layer by layer until finally only four qubits are left to form a final block. Losses are thereby *distilled* through all n layers, where finally the code is correctable if no more than one loss arrives at the last layer. If, on the other hand, two or more losses percolate to the final layer, the code is not correctable and requires a deeper distillation process by engaging more layers. This so-called distiller version of the erasure code is thoroughly explained in Ref. [213]. Crucially, in the limit of layers n , the code yields a loss threshold of 16.7%.

We now experimentally demonstrate the minimal erasure code instance. To enable the detection of loss-events, we merge the QND unit from Sec. 5.1 in between encoding and correction step and couple the code qubit to an extra ancilla. This code works in the following way. After logical encoding, partial loss is induced on the top code qubit via $R_1^{\text{loss}}(\theta)$ and subsequently detected by a destructive measurement on the ancilla qubit. Each identified loss-event triggers a correction step, while the logical encoding remains intact in the absence of loss. Fig. 5.4 illustrates the resulting 1+4-qubit protocol.

Let us take a closer look at the circuit implementation of the encoding step. The top code qubit is initially considered in an arbitrary quantum state $|\psi\rangle_1 = \alpha|0\rangle_1 + \beta|1\rangle_1$ along




Figure 5.4: **Erasure code for combined qubit loss detection and correction.** In the present alignment the 1+4-qubit protocol resolves losses on the top code qubit. After encoding and inducing a partial loss via $R_1^{\text{loss}}(\theta)$, the information about the loss is mapped from the code to the ancilla qubit using the detection unit from Fig. 5.1. Each detected loss-event then triggers a correction step on experimental sample (or shot) basis that finally teleports the lost state $|\psi\rangle_1$ onto the bottom qubit. Absent of loss, the logical encoding remains intact.

three ancilla qubits in their ground-state $|000\rangle_{2,3,4}$. The loss-robust logical encoding $|\psi_L\rangle$ is achieved using three CNOT-gates and two local H-gates as depicted in the encoding inset from Fig. 5.4. This circuit prepares logical code words of form

$$|0_L\rangle = 1/\sqrt{4}(|00\rangle + |11\rangle)_{1,2}(|00\rangle + |11\rangle)_{3,4} \quad (5.1)$$

$$|1_L\rangle = 1/\sqrt{4}(|00\rangle - |11\rangle)_{1,2}(|00\rangle - |11\rangle)_{3,4}, \quad (5.2)$$

where the two basis-states $|0_L\rangle$ and $|1_L\rangle$ differ by phase flips $Z_1 Z_3$.

Table 5.1: **Loss correction unitaries for the erasure code from Fig. 5.4.** Measurements on qubits two and three yield outcome-dependent unitary operations, which applied to the bottom qubit, restore the lost state $|\psi\rangle_4$. We abbreviate X-basis eigenstates by $|+\rangle = 1/\sqrt{2}(|0\rangle + |1\rangle)$ and $|-\rangle = 1/\sqrt{2}(|0\rangle - |1\rangle)$.

$M(Z_2, X_3)$	Loss correction unitaries
$ 00\rangle_{23}$	$ \psi\rangle_4 = H_4 X_4 Z_4 \quad (\alpha -\rangle_4 - \beta +\rangle_4)$
$ 01\rangle_{23}$	$ \psi\rangle_4 = H_4 Z_4 \quad (\alpha +\rangle_4 - \beta -\rangle_4)$
$ 10\rangle_{23}$	$ \psi\rangle_4 = H_4 X_4 \quad (\alpha -\rangle_4 + \beta +\rangle_4)$
$ 11\rangle_{23}$	$ \psi\rangle_4 = H_4 \quad (\alpha +\rangle_4 + \beta -\rangle_4)$

Our first experiment targets the encoding of several logical input states. Herefore, Fig. 5.5 provides a specifically tailored gate-sequence for the trapped-ion toolbox based on gate decompositions from Fig. 2.5 that have further been optimized with an in-house developed sequence compiler [214]. For the present implementation, the former addressing setup [146] was utilized, which requires the use of collective MS-gates according to Eq. (2.21). These collective entangling gates are accompanied by refocused local operations comprised of the collective X, Y-gates from Eq. (2.20) and the addressed Z-gates from Eq. (2.18). This is

necessary for maintaining phase coherence between collective and addressed gates, see Ch. 2.4. Implementing the gates in this refocused fashion leads to in total 13 addressed single-qubit gates alongside six collective local gates. Fortunately, collective local gates are much cheaper than addressed single-qubit ones [52]. Note that this last optimization step is no longer displayed in Fig. 5.5.

We experimentally prepare the logical Z_L - and X_L -basis states $\{|0_L\rangle, |1_L\rangle, |0_L\rangle + |1_L\rangle, |0_L\rangle - |1_L\rangle\}$ and use four-qubit standard Pauli QST to evaluate their performance. The comparably low qubit number of the protocol allows MLE state reconstruction according to Eq. (3.7). The reconstruction process thus involves a *completely-positive and trace-preserving* (CPTP) constraint (explained in Sec. 1.1.3), which is particularly useful for experimental data, as they are subject to noise and finite statistics.




Figure 5.5: **Optimized gate sequence for logical encoding of arbitrary input states $|\psi\rangle$ in the erasure code.** The gate sequence according to the left inset of Fig. 5.4 has been specifically tailored for our trapped-ion hardware using gate decompositions from Fig. 2.5 and has further been optimized with our in-house sequence compiler [214], see text for details.

Fig. 5.6 depicts the reconstructed density matrices on all four input states $\{|0_L\rangle, |1_L\rangle, |0_L\rangle + |1_L\rangle, |0_L\rangle - |1_L\rangle\}$ alongside their resulting fidelities $\{0.63(1), 0.67(1), 0.57(2), 0.65(1)\}$ with the ideal outcomes plotted to their left. Errors correspond to one standard deviation of QPN from resampling the QST data.

Next, we consider a simple noise model that takes into account the error-rates of all the individual gates that are present in the experiment. In particular, we consider $p_{MS} = 0.01$ for collective multi-qubit gates and $p_{ADD} = 0.01$ for addressed single-qubit gates according to Ref. [52]. Since the error-rates for the collective local operations are below 0.001, they can be neglected. As such, we expect the fidelities on the logical encoding around

$$\mathcal{F}_{\text{enc}} = \underbrace{(1 - p_{MS})^3 (1 - p_{ADD})^{13}}_{\text{Erasure encoding}} = 0.99^{16} \approx 0.85.$$

The gap between accumulated gate errors from the model and the actual implementation arises mostly from uncompensated AC-Stark shifts, induced by the MS-gate's high light power requirement. Existing phase optimization tools [215] could potentially improve fidelities, at most to the modelled 0.85.

The eventual goal of this chapter is to demonstrate qubit loss detection and correction in a deterministic experiment where detected loss-events trigger a real-time correction step.




Figure 5.6: **Experimentally encoded logical Z- and X-basis states according to Fig. 5.3.** The resulting states are characterized with four-qubit standard Pauli QST using MLE reconstruction. For illustration purposes, ideal and experimentally reconstructed density matrices are printed side-by-side. Color coded bars denote complex phases. Errors on fidelities represent one standard deviation of QPN from resampling the QST data.

Before finally implementing in-sequence detection and correction in such feed-forward fashion, for now we perform the system evaluation with the final correction step (right inset of Fig. 5.4) in post-processing. In particular, we experimentally demonstrate the protocol in Fig. 5.4 for the logical input state $|0_L\rangle$. After logical encoding, loss is induced on the code qubit via $R_c^{\text{loss}}(\theta)$ and subsequently detected with the ancilla qubit measurement. Hereafter, QST is applied to the remaining three qubits which after loss end up in an incoherent mixture of states

$$(|000\rangle + |011\rangle)_{2,3,4} \quad \text{and} \quad (|100\rangle + |111\rangle)_{2,3,4}. \quad (5.3)$$

Fig. 5.7 shows the corresponding experimentally reconstructed three-qubit density matrix. The fidelity with the ideal outcome plotted aside is $0.79(2)$. This fidelity turns out to be notably higher than those of the logical encodings from Fig. 5.6, since here correlations with the code qubit disappear after its loss, increasing the overall quality.




Figure 5.7: **Logical three-qubit encoding after losing the code qubit.** Experimentally reconstructed density matrix with the three remaining qubits, leaving the system in a mixed state according to Eq. (5.3). For illustration purposes, ideal and experimentally from QST derived density matrix are printed side-by-side. Color coded bars denote complex phases. The Fidelity with the ideal density matrix reads $0.79(2)$. Errors represent one standard deviation of QPN from resampling the QST data.

Let us use this logical three-qubit density matrix to perform the final correction step in post-processing. After performing the measurement of the middle qubits two and three and apply the outcome-dependent loss correction unitary from Tab. 5.1 to the bottom qubit, we successfully recover the lost state $|0\rangle_4$ in all four cases. Fig. 5.8 shows the results. Fidelities with respect to the ideal density matrix $|0\rangle\langle 0|_4$ yield $\{0.79(6), 0.77(6), 0.76(6), 0.95(7)\}$, referring to measurement outcomes $\{|00\rangle_{23}, |01\rangle_{23}, |10\rangle_{23}, |11\rangle_{23}\}$, respectively. Errors correspond to one standard deviation of QPN from resampling the QST data. Moreover, results in Fig. 5.8 (a)-(c) show unwanted population in the $|1\rangle_4$ state of around 0.15, which is not present in the $|11\rangle_{23}$ outcome from Fig. 5.8(d). The latter therefore reflects a higher fidelity. Unwanted populations may appear only in some outcomes because phase errors are reflected differently in the recovered state after applying the correction unitary.

This experimental implementation is again limited by the errors in the collective MS-gates and addressed single-qubit gates, analogous to the results on the logical encodings from Fig. 5.6. By incorporating loss detection and correction, we now additionally observe phase errors from AC-Stark shifts [139] on the spectroscopically decoupled ancilla qubits. Such decoupled qubits experience the off-resonant light of the collective beam during loss detection and shift in phase relative to the active qubits. Fortunately, after neglecting the lost qubit and measuring the middle ones, correlations once more vanish and increase the




Figure 5.8: **Density matrices of loss-recovered state $|0\rangle_4$.** The final correction step is performed in post-processing based on the experimentally reconstructed density matrix from Fig. 5.7. Color coded bars denote complex phases. Errors on the fidelity with the ideal density matrix represent one standard deviation of QPN from resampling the QST data. **(a)-(c)** denote unwanted population in $|1\rangle_4$ around 0.15, lowering the fidelity compared to the outcome in **(d)**, see text.

overall fidelity of the loss recovered state $|0\rangle_4$. Modelling the expected overall fidelity by accounting for operational errors on all gates in the sequence results in

$$\begin{aligned} \mathcal{F}_{\text{enc}} \cdot \mathcal{F}_{\text{loss}} \cdot \mathcal{F}_{\text{det}} &= \underbrace{(1 - p_{\text{MS}})^3 (1 - p_{\text{ADD}})^{13}}_{\text{erasure encoding}} \cdot \underbrace{(1 - p_{\text{ADD}})^1}_{\text{induce loss}} \cdot \underbrace{(1 - p_{\text{MS}})^2 (1 - p_{\text{ADD}})^7}_{\text{loss detection}} \\ &= 0.99^{26} \approx 77\%. \end{aligned}$$

The modelled fidelity is in good agreement with the experimental results from Fig. 5.8. For the experiment shown, we have again used an optimized gate sequence consisting of collective X, Y-gates from Eq. (2.20) and addressed Z-gates from Eq. (2.18). Because the correction step was left to post-selection, the respective errors have been excluded from the model.

It should finally be noted that in contrast to the quantum entangled resource presented here, teleportation based on a fully classical resource yields a maximum possible average fidelity of 0.667 [216]. The demonstrated level of performance is thus already better than any classical means of doing teleportation-based loss correction.

The proof-of-principle experiments presented show a full cycle of qubit loss detection and correction. Even though the final loss correction step is done in post-processing, our

experiments draw attention to the practicality issues of the erasure code. These encompass a significant operational overhead due to encoding and recoding steps required at each iteration of loss correction and the need for three additional ancilla qubits to correct a single known loss. The code also does not cover the correction of computational errors, which would require additional protocols for implementing fault-tolerant applications.

5.2.2 Loss in the surface code

The erasure code has the conceptual drawbacks of requiring large operational overheads and not allowing the correction of computational errors. These disadvantages can be overcome by treating loss in the context of the surface code [77, 78], introduced in Ch. 1.3.3. The stabilizer formalism of the surface code bypasses costly decoding and recoding steps, resulting in much lower operational overhead, while its topological structure requires only limited connectivity to neighbors. In addition, the surface code shows significant robustness against losses, if their detection and correction is well treated in the underlying QEC implementation. In all this, the surface code is a promising candidate to resolve errors at all levels of abstraction, as initially motivated by Fig. 1.5.




Figure 5.9: **Qubit loss correction in the surface code as proposed by Stace and Barret in Ref. [94].** (a) Losses can be corrected by deforming logical operators around defect surface sites that define equivalent paths between horizontal and vertical code edges. A code remains correctable as long as such a closed path connecting opposite code edges still exists, see text for details. (b) Left: Minimal surface code excerpt (see Fig. 1.4) for correcting qubit loss. Green plaquettes of the lattice indicate four-body Z-stabilizers, while lattice vertices represent four-body X-stabilizers that define the code with their common +1 eigenspace. For a lost qubit, all affected S_P^Z -stabilizers can be redefined by multiplying neighboring pairs of plaquettes to a superplaquette \tilde{S}_P^Z of weight six, while vertex stabilizers S_V^X can be reused in their shrunken size, i.e. on the remaining three qubits. Shrunken vertex operators \tilde{S}_V^X , however, remain in a state of indeterminate eigenvalues ± 1 after loss, whereon they must be reprojected onto their common +1 eigenspace, which finally revives the logical encoding on the remaining qubits.

We follow the protocol proposed by Stace & Barret in Ref. [94] that provides a loss correction framework for the surface code and to which we refer the reader for further explanation.

To correct for loss, we use that the definition of the logical operators Z_L and X_L , which form straight paths between vertical and horizontal code edges, is not unique. For instance, multiplying Z_L with any plaquette operator S_P^Z yields an equivalent but geometrically deformed logical operator $\tilde{Z}_L = S_P^Z Z_L$ around the plaquette S_P^Z . So there

are many equivalent representations to the two logical operators Z_L and X_L . In Fig. 5.9(a), for example, where only a few qubits are lost, a new logical path on intact qubits can be found to save the code. For too many losses, however, there might occur several code parts where the operator redefinition is no longer feasible, denoted as *percolated regions*. Results in Ref. [217] show that the question of whether a surface code can be saved or not applies geometrically to roping together nails on a square lattice, known as *bond percolation* [217]. For bond percolation in the limit of a large lattice, there is a sharp boundary between correctable and non-correctable losses, yielding $p_{\text{loss}} < 0.5$ [217]. Interestingly, this *bond percolation threshold* readily extends to losses in the surface code, given the absence of computational errors.

On a loss affected surface, the set of stabilizers comprising of plaquette and vertex operators must be redefined too. The code fragment in Fig. 5.9(b) illustrates such stabilizer redefinitions. This works in the following way. If qubit one is lost (right panel), it is helpful to multiply adjacent pairs of plaquettes of the lattice that share the loss to form a so-called superplaquette $S_{p_A}^Z \cdot S_{p_B}^Z = \tilde{S}_p^Z$. This product operator \tilde{S}_p^Z has no more support on the lost qubit, but commutes with all other stabilizers and therefore is a stabilizer itself. Hence, in the absence of computational errors, the product superoperator \tilde{S}_p^Z is preserved in eigenvalue +1. Damaged vertex operators, on the other hand, have no longer support on the lost qubit. Yet, we find the remaining vertex operator that excludes the loss to still commute with all other stabilizers, which can therefore be reused in their shrunken size \tilde{S}_v^X . Due to the loss, however, a measurement of the shrunken vertex operators \tilde{S}_v^X yields random eigenvalues ± 1 as the loss leaves the code in a mixed state. To fix the code, the shrunken vertex operator \tilde{S}_v^X must be reprojected onto their common +1 eigenspace, saving the logical information.

Conclusively, if loss can be detected and corrected before exceeding the code-size dependent threshold, the logical information can be saved by switching to a smaller surface that avoids all losses.

The authors in Ref. [94] further demonstrate that the surface code can remain operable under both losses and computational errors. To this extent, they model local, uncorrelated noise and provide parameter regimes under which both errors are tolerable. It is noteworthy that the correctable computational error and loss rates sensitively depend on each other, whereas in the limit of no computational errors, a loss probability of up to $p_{\text{loss}} = 0.5$ becomes tolerable.

Encouraged by this high potential robustness against losses, we experimentally demonstrate a full qubit loss detection and correction cycle on the surface code, where a detected loss-event triggers a correction step in real-time. Our demonstrations are the first of such deterministic kind and lead to the third publication of this thesis [218], presented in the upcoming Sec. 5.3.

5.3 PUBLICATION: EXPERIMENTAL DETERMINISTIC CORRECTION OF QUBIT LOSS

***Nature* volume 585, pages 207–210 (2020)**

submitted on 06 March 2020, accepted on 26 July 2022 and published on 09 September 2020

<https://doi.org/10.1038/s41586-020-2667-0>

Roman Stricker¹, Davide Vodola^{2,3,4}, Alexander Erhard¹, Lukas Postler¹, Michael Meth¹, Martin Ringbauer¹, Philipp Schindler¹, Thomas Monz^{1,5}, Markus Müller^{2,6,7} and Rainer Blatt^{1,8}

¹ *Institut für Experimentalphysik, Universität Innsbruck, 6020 Innsbruck, Austria*

² *Department of Physics, College of Science, Swansea University, Singleton Park, SA2 8PP Swansea, United Kingdom*

³ *Dipartimento di Fisica e Astronomia dell'Università di Bologna*

⁴ *INFN, Sezione di Bologna, I-40127 Bologna, Italy*

⁵ *Alpine Quantum Technologies GmbH, 6020 Innsbruck, Austria*

⁶ *Institute for Quantum Information, RWTH Aachen University, D-52056 Aachen, Germany*

⁷ *Peter Grünberg Institute, Theoretical Nanoelectronics, Forschungszentrum Jülich, D-52425 Jülich, Germany*

⁸ *Institut für Quantenoptik und Quanteninformation, Österreichische Akademie der Wissenschaften, Otto-Hittmair-Platz 1, A-6020 Innsbruck, Austria*

The author to the present thesis executed the experiments, analyzed the data and wrote the manuscript.

Article

Experimental deterministic correction of qubit loss

<https://doi.org/10.1038/s41586-020-2667-0>

Received: 6 March 2020

Accepted: 26 June 2020

Published online: 9 September 2020

 Check for updates

Roman Stricker¹✉, Davide Vodola^{2,3,4}, Alexander Erhard¹, Lukas Postler¹, Michael Meth¹, Martin Ringbauer¹, Philipp Schindler¹, Thomas Monz^{1,5}, Markus Müller^{4,6,7} & Rainer Blatt^{1,8}

The successful operation of quantum computers relies on protecting qubits from decoherence and noise, which—if uncorrected—will lead to erroneous results. Because these errors accumulate during an algorithm, correcting them is a key requirement for large-scale and fault-tolerant quantum information processors. Besides computational errors, which can be addressed by quantum error correction^{1–9}, the carrier of the information can also be completely lost or the information can leak out of the computational space^{10–14}. It is expected that such loss errors will occur at rates that are comparable to those of computational errors. Here we experimentally implement a full cycle of qubit loss detection and correction on a minimal instance of a topological surface code^{15,16} in a trapped-ion quantum processor. The key technique used for this correction is a quantum non-demolition measurement performed via an ancillary qubit, which acts as a minimally invasive probe that detects absent qubits while imparting the smallest quantum mechanically possible disturbance to the remaining qubits. Upon detecting qubit loss, a recovery procedure is triggered in real time that maps the logical information onto a new encoding on the remaining qubits. Although the current demonstration is performed in a trapped-ion quantum processor¹⁷, the protocol is applicable to other quantum computing architectures and error correcting codes, including leading two- and three-dimensional topological codes. These deterministic methods provide a complete toolbox for the correction of qubit loss that, together with techniques that mitigate computational errors, constitute the building blocks of complete and scalable quantum error correction.

Qubit loss comes in a variety of physical manifestations, such as the loss of particles encoding the qubits in atomic and photonic implementations^{11–14}, but also as leakage out of the two-dimensional (2D) computational qubit subspace in multi-level solid-state¹⁸ and atomic, molecular and optical systems¹¹. Whereas progress has been made in characterizing and suppressing the rate of loss and leakage processes^{19–23}, in many platforms these processes still occur at rates of the same order of magnitude as other errors, such as amplitude damping in trapped-ion qubits encoded in metastable states of optical transitions¹¹. It is known that unnoticed and uncorrected qubit loss and leakage will severely affect the performance of quantum processors^{18,24}; therefore, dedicated protocols to fight this error source have been devised. These protocols include four-qubit quantum erasure codes¹⁰, which have been implemented using photons and post-selective quantum state analysis^{12,13}, as well as protocols proposed to address qubit loss in the surface code^{15,25,26} and 2D colour codes^{27,28}. So far, an experimental implementation of deterministic detection and correction of qubit loss and leakage, both of which will be referred to as ‘loss’ in the following, remains an outstanding challenge.

A general, architecture-independent protocol to protect quantum information against loss errors consists in (i) the initial encoding of logical states into a multi-qubit register, (ii) a quantum non-demolition (QND) measurement scheme that determines the position of potentially lost qubits, (iii) a reconstruction algorithm that, if not too many loss events have occurred, reconstructs the damaged code, and (iv) a final set of measurements that fixes the new code by initializing the new stabilizers.

Here, we encode a single logical qubit in an excerpt of the surface code^{15,16}, which is a topological quantum error-correcting (QEC) code in which physical qubits reside on the edges of a 2D square lattice; see Fig. 1a. The surface code is a Calderbank–Shor–Steane code^{29,30}, for which stabilizer operators S_V^X are associated to each vertex V (blue cross in Fig. 1a) via $S_V^X = \prod_{j \in V} X_j$ and to each plaquette P (green square in Fig. 1a) via $S_P^Z = \prod_{j \in P} Z_j$, where X_j, Y_j, Z_j are Pauli matrices acting on the physical qubit j . All stabilizers mutually commute, and their common +1 eigenspace fixes the code space that hosts the logical quantum states $|\psi_L\rangle$, that is, $S_P^Z |\psi_L\rangle = S_V^X |\psi_L\rangle = |\psi_L\rangle$ for all plaquettes and vertices. The operators that define and induce flips of the logical-basis states

¹Institut für Experimentalphysik, Universität Innsbruck, Innsbruck, Austria. ²Dipartimento di Fisica e Astronomia dell'Università di Bologna, Bologna, Italy. ³INFN, Sezione di Bologna, Bologna, Italy. ⁴Department of Physics, College of Science, Swansea University, Swansea, UK. ⁵Alpine Quantum Technologies GmbH, Innsbruck, Austria. ⁶Institute for Quantum Information, RWTH Aachen University, Aachen, Germany. ⁷Theoretical Nanoelectronics, Peter Grünberg Institute, Forschungszentrum Jülich, Jülich, Germany. ⁸Institut für Quantenoptik und Quanteninformation, Österreichische Akademie der Wissenschaften, Innsbruck, Austria. ✉e-mail: roman.stricker@uibk.ac.at

Article




Fig. 1 | The surface code and correction of qubit loss. **a**, Logical qubits are encoded collectively in many physical qubits (grey circles) that are located on the edges of a 2D square lattice. The code space is defined via four-qubit S^Z and S^X stabilizers acting on groups of qubits that reside around plaquettes (green square) and vertices (blue cross) of the lattice. Logical T^Z and T^X operators are defined along strings of qubits that span the entire lattice along two non-trivial paths, as depicted by the vertical green (horizontal blue) string for T^Z (T^X). Right, logical string operators do not have unique support, but can be deformed by multiplication with stabilizers, as illustrated for T^Z (T^X), which is deformed into \tilde{T}^Z (\tilde{T}^X) by the green plaquette (blue vertex) stabilizer. **b**, Left, excerpt of a qubit lattice suffering the loss (orange arrow) of a physical qubit (white circle). The loss affects two plaquette operators, S^Z_1 and S^Z_2 , and two vertex operators, S^X_1 and S^X_2 . Right, the correction algorithm consists of introducing a new merged Z-stabilizer generator as $\tilde{S}^Z_1 = S^Z_1 S^Z_2$, which does not involve the lost qubit, and two new X stabilizers, \tilde{S}^X_1 and \tilde{S}^X_2 , which have reduced support on three qubits that are unaffected by the loss.

$|0_L\rangle$ and $|1_L\rangle$ are the logical generators T^Z and T^X , respectively. They commute with all stabilizers, and can be chosen as products of X and Z operators along strings that span the entire lattice; see Fig. 1a.

To recover a logical qubit affected by qubit loss, one needs to switch to an equivalent set of stabilizers $\{\tilde{S}^X_p, \tilde{S}^Z_p\}$ and logical operators $\{\tilde{T}^X, \tilde{T}^Z\}$ defined only on qubits that are not affected by losses. For this redefinition we follow the scheme introduced in ref.²⁵ and shown in Fig. 1b. Notably, the logical operators do not have unique support because equivalent operators \tilde{T}^X and \tilde{T}^Z can be obtained by multiplying T^X and T^Z by any subset of stabilizers. For the surface code, this results in the deformation of the string of physical qubits that supports the logical operator; see Fig. 1a. For too many losses, however, finding such an equivalent logical operator might not be possible. Because each loss event results in the deletion of one edge (bond) of the 2D square lattice, the question of whether such a path supporting a logical operator exists corresponds to the classical problem of bond percolation, which for the surface code results in a threshold of tolerable qubit loss rate as high as 50% in the absence of other errors²⁵.

Inspired by the surface code stabilizer structure, we implement a subspace defined by three stabilizers on four qubits that allows us to experimentally explore the reconstruction protocol as described in Fig. 2a. We note that this subspace is neither an error detection nor a correction code for Pauli errors, but the logical information can be made robust to the loss of qubit 1. For the physical realization of this code, we consider a string of $^{40}\text{Ca}^+$ ions confined in a linear Paul trap¹⁷. Each ion represents a physical qubit encoded in the electronic levels $S_{1/2}(m=-1/2) = |0\rangle$ and $D_{5/2}(m=-1/2) = |1\rangle$. Our setup is capable of realizing a universal set of quantum gate operations consisting in (a) single-qubit rotations by an angle θ around the z axis of the form $R^Z_j(\theta) = \exp(-i\theta Z_j/2)$ on the j th ion, (b) collective qubit rotations around the x and y axes of the form $R^\sigma(\theta) = \exp(-i\theta \sum_j (\sigma_j/2))$, with $\sigma = X$ or Y , via a laser beam addressing

the entire register, and (c) multi-qubit Mølmer-Sørensen entangling gate operations³¹ $MS^X(\theta) = \exp(-i\theta \sum_{j<k} (X_j X_k/2))$. This gate set is complemented by single-qubit hiding and unhiding operations in order to apply collective multi-qubit operations to only a subset of qubits¹⁷. Similarly, this technique is used to read out individual qubits within the register without influencing the other qubits; see Supplementary Information for details.

To benchmark the performance of the protocol we introduce qubit loss in a controlled way as leakage to another electronic level outside the computational subspace; see Fig. 3a. Leakage is the dominant form of qubit loss in ion-trap architectures, whereas our protocol is also applicable to other forms of loss and architectures. The qubit that potentially suffers a loss is partially pumped out of its computational subspace $\{S_{1/2}(m=-1/2) = |0\rangle, D_{5/2}(m=-1/2) = |1\rangle\}$ by coherently driving the carrier transition $S_{1/2}(m=-1/2) = |0\rangle \leftrightarrow D_{5/2}(m=-5/2) = |2\rangle$. In the following, this is referred to as the loss operation $R_{\text{loss}}(\phi)$, where the probability of loss from state $|0\rangle$ is given by $\sin^2(\phi/2)$. The loss rate on the logical qubit is the product of the loss probability with the population in $|0\rangle$.

To detect a loss event we implement a QND measurement as shown in Fig. 3a, which signals the loss of a code qubit by a bit-flip on an ancillary qubit prepared in state $|0\rangle$, followed by an addressed readout of the ancillary qubit. The key ingredient of this QND measurement is a two-qubit entangling gate operation $MS^X(\pi)$ that performs a collective bit-flip operation on the code and ancilla qubits if the code qubit is present. If the code qubit has been lost, on the other hand, regardless of whether loss occurs from the $|0\rangle$ or $|1\rangle$ state, this operation acts only on the ancilla, on which it performs an identity operation; see Supplementary Information for details. A subsequent collective bit-flip $R^X(\pi) = X$ will flip the ancilla qubit to $|1\rangle$ before its addressed readout. If no loss occurred, the collective bit-flip induced by $MS^X(\pi)$ will be undone by the $R^X(\pi) = X$ operation, and the ancilla qubit will end in state $|0\rangle$ (ref.¹⁷). The code qubit, on the other hand, will in this case undergo a non-unitary evolution given by (up to normalization) $\rho \rightarrow E\rho E^\dagger$ with $E = |1\rangle\langle 1| + \cos(\phi/2)|0\rangle\langle 0|$, which for small loss rates ($\phi \approx 0$) converges to the identity operation. This is a consequence of the information gain that no loss has occurred in this instance, provided by the ancilla measurement; see Supplementary Information.

We test the loss-detection sub-circuit on the full five-qubit register by driving the loss transition $R_{\text{loss}}(\phi)$ on qubit 1 and measuring the population in the $D_{5/2}$ state on both the code and ancilla qubits. This measurement does not distinguish between the different Zeeman sub-levels of the $D_{5/2}$ state manifold. Figure 3b shows that loss detected by the ancilla qubit matches the loss induced on qubit 1 within statistical uncertainty, indicating that a loss event is reliably detected. The quantified detection efficiency is 96.5(4)%, with a false positive rate of 3(1)% and a false negative rate of 1(1)%.

We note that for very low loss rates, the fidelity of the final state after correcting qubit loss will be limited by imperfections in the QND loss-detection unit; see Supplementary Information for details. To quantify the performance of the QND detection scheme in the absence of loss, we reconstruct the Choi matrix³² of the corresponding non-unitary map using generalized quantum process tomography. The reconstructed Choi matrix shown in Fig. 3c confirms this dynamical behaviour expected in the no-loss case with a process fidelity of 90(2)% with $\sim 20\%$ ($\phi = 0.3\pi$) loss from $|0\rangle$. This demonstrates that information about loss on the code qubit can be reliably mapped onto the ancilla qubit. For general loss-detection purposes, one could use the detection unit to probe all code qubits within the register sequentially.

To investigate the robustness of our minimal-instance logical qubit against loss, we combine the loss-detection unit and the conditional-correction step in a 1+4-qubit algorithm, sketched in Fig. 2a. The experimental sequence for encoding an arbitrary input state of the form $|\psi_L\rangle = \cos(\alpha/2)|0_L\rangle + \sin(\alpha/2)|1_L\rangle$ in our ion-trap quantum computer is given in Supplementary Information. The logical basis states $|0_L\rangle$ and $|1_L\rangle$ encoded by the initial stabilizers read




Fig. 2 | Experimental realization of the 1+4-qubit algorithm aiming at loss detection and correction. **a**, Minimal four-qubit system for the experimental realization of the full loss-correction protocol. The code is defined by three stabilizers, $S_1^Z = Z_1Z_2$, $S_2^Z = Z_1Z_3$ (green squares) and $S_1^X = X_1X_2X_3X_4$ (blue cross) and stores a single logical qubit with logical operators $\tilde{T}^Z = Z_1Z_4$, $\tilde{T}^X = X_4$ and $\tilde{T}^Y = i\tilde{T}^X\tilde{T}^Z$. In the event of the loss (orange arrow) of qubit 1 (white circle), the merged Z stabilizer $\tilde{S}_1^Z = S_1^Z S_2^Z = Z_2Z_3$ and a new X stabilizer $\tilde{S}_1^X = X_2X_3X_4$ with reduced support on the remaining three qubits are introduced for the new encoding. The logical operators equivalent to the initial ones are $\tilde{\tilde{T}}^Z = S_1^Z \tilde{T}^Z = Z_2Z_4$, $\tilde{\tilde{T}}^X = X_4$ and $\tilde{\tilde{T}}^Y = i\tilde{\tilde{T}}^X\tilde{\tilde{T}}^Z$. **b**, Expectation values for logical operators (T), stabilizers (S)

and code space populations (P_{CS}), defined in Supplementary Information, for the logical superposition state $|+i_L\rangle = (|0_L\rangle + i|1_L\rangle)/\sqrt{2}$. A loss rate of 25% was induced on qubit 1. All values are estimated from four-qubit quantum state tomography, with ideal values shaded in the background. Errors correspond to one standard deviation of statistical uncertainty due to quantum projection noise. **c**, In the absence of loss, the logical encoding remains largely intact. **d**, In the case of loss, we reconstruct the code on the three remaining qubits after measuring the shrunk stabilizer of the new encoding, $\tilde{S}_1^X = X_2X_3X_4$, and selecting the appropriate Pauli basis, that is, performing a Pauli frame update in the case of a -1 outcome in the \tilde{S}_1^X measurement.

$|0_L\rangle = (|0000\rangle + |1111\rangle)/\sqrt{2}$ and $|1_L\rangle = (|0001\rangle + |1110\rangle)/\sqrt{2}$. These entangled states are produced with a single fully entangling MS gate, $MS^X(\pi/2)$, acting on all four code qubits, supported by additional local operations. Loss is observed using the QND detection unit, with an ancilla qubit for loss readout. In this smallest excerpt of the surface code, we consider potential qubit loss to happen on qubit 1 only; hence, we probe only qubit 1 using the QND-detection unit as indicated in Fig. 2a. Conditional on the detection of a loss event, our control scheme triggers a real-time

deterministic code restoration via feed-forward. If no loss is detected, the logical states can be verified by measuring the generators of the stabilizer group $\{S_1^Z = Z_1Z_2, S_2^Z = Z_1Z_3, S_1^X = X_1X_2X_3X_4\}$ and the logical operators $\{T^Z = Z_1Z_4, T^X = X_4, T^Y = iT^X T^Z\}$ of the original encoding. If loss occurs, the encoded logical information can be restored by switching to an encoding defined on a smaller subset of three qubits. This is realized by a projective measurement of the shrunk stabilizer of the new encoding $\tilde{S}_1^X = X_2X_3X_4$, which after the loss is in an undetermined state. This




Fig. 3 | Investigating the performance of the QND loss-detection unit. **a**, Circuit representation of the detection unit, which maps potential loss from qubit 1 onto the ancilla qubit. The experimental results in **b** and **c** were extracted from experiments performed on the full five-qubit register, according to Fig. 2. **b**, Population in the $D_{5/2}$ state of qubit 1 (directly measured loss) and ancilla qubit (detected loss) measured after loss detection. Controlled loss of up to 100% from state $|0\rangle$ was introduced. The estimated detection efficiency is 96.5(4)%, which demonstrates that the occurrence of a

loss event can be reliably mapped onto the ancilla qubit and read out in a QND fashion. Errors correspond to one standard deviation of statistical uncertainty due to quantum projection noise. **c**, Reconstructed Choi matrix for a loss of -20° ($\phi = 0.3\pi$) from the $|0\rangle$ state with a process fidelity of 90(2)%, compared to the ideal values denoted by black frames. We find that, as expected, the detection unit performs a non-unitary evolution that deviates from the identity operator owing to measurement back-action; see Supplementary Information.

Article

initializes the three-qubit stabilizer in a +1 (or -1) eigenstate, where the -1 case requires a redefinition of the Pauli basis (Pauli frame update)^{33,34}; see Supplementary Information for details. For this stabilizer readout, a freshly initialized ancilla qubit is needed. In our implementation we recycle the ancilla qubit, previously used for the QND loss detection, because it remains unaffected by the measurement in the loss case. Following this procedure, the initial logical encoding is reconstructed in the smaller subset of three qubits; see Fig. 2a.

We now present the results obtained from the full implementation of the 1+4-qubit algorithm, as shown in Fig. 2. Data were taken for three different input states, namely, the logical basis states $|0_L\rangle$ and $|1_L\rangle$, presented in Supplementary Information, as well as their superposition $|+i_L\rangle = (|0_L\rangle + i|1_L\rangle)/\sqrt{2}$ presented here. To verify the initialization of $|+i_L\rangle$, we reconstruct the experimental density matrix via four-qubit quantum state tomography on the code qubits, yielding a fidelity of 84(1)% with the ideal state. From the reconstructed density matrix we further extract the components of the 'logical' Bloch vector, represented by expectation values of the associated logical operators, the code space population P_{CS} (explained in the Supplementary Information) and the expectation values of the stabilizer generators summarized in Fig. 2b.

After the encoding, partial loss on qubit 1 is induced by coherently exciting the loss transition $R_{\text{loss}}(\phi)$ for different values of ϕ . Here, we present the case of a loss rate of 25%, that is, $\phi = 0.5\pi$, and other values are found in Supplementary Information. Loss is detected by a QND measurement mapping the information of loss onto the ancilla qubit, followed by a projective measurement of the ancilla qubit. The measurement result triggers a real-time deterministic code restoration via feed-forward. If no loss is detected, quantum state tomography on all four code qubits is performed to verify that the initial encoding $|+i_L\rangle$ is still intact, with a fidelity of 66(1)% with respect to the expected state; see Fig. 2c. If loss is detected, the code is switched to the remaining three qubits by a projective measurement of the shrunk stabilizer S_1^A , as illustrated in Fig. 2a, and a Pauli frame update in case of a -1 outcome. Quantum state tomography yields a fidelity of the resulting three-qubit logical state $|+i_L\rangle$ of 78(1)%; see Fig. 2d.

The observed decrease in fidelity after loss detection is mainly due to cross-talk between neighbouring ions resulting in unitary errors on the final state, and dephasing due to laser-frequency and magnetic-field fluctuations. Additionally, in the no-loss case the ancilla qubit has scattered photons during the in-sequence loss detection. This heats up the ion string, decreasing the quality of the subsequent tomography operations.

Our work demonstrates the first deterministic detection and correction of qubit loss. Our building blocks are readily applicable to leading QEC codes, such as the surface and colour codes, and fully compatible with the framework of topological QEC. Although this demonstration is performed on an ion quantum processor, essentially all experimental quantum computing platforms are affected by qubit loss or leakage, and could thus benefit from our methods. Fault-tolerant versions of the presented routines in combination with correction of computational errors constitute required extensions towards the realization of large-scale quantum computers.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information,

acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41586-020-2667-0>.

- Gottesman, D. Theory of fault-tolerant quantum computation. *Phys. Rev. A* **57**, 127–137 (1998).
- Chiaverini, J. et al. Realization of quantum error correction. *Nature* **432**, 602–605 (2004).
- Schindler, P. et al. Experimental repetitive quantum error correction. *Science* **332**, 1059–1061 (2011).
- Nigg, D. et al. Quantum computations on a topologically encoded qubit. *Science* **345**, 302–305 (2014).
- Takita, M., Cross, A. W., Córcoles, A. D., Chow, J. M. & Gambetta, J. M. Experimental demonstration of fault-tolerant state preparation with superconducting qubits. *Phys. Rev. Lett.* **119**, 180501 (2017).
- Linke, N. M. et al. Fault-tolerant quantum error detection. *Sci. Adv.* **3**, e1701074 (2017).
- Córcoles, A. D. et al. Demonstration of a quantum error detection code using a square lattice of four superconducting qubits. *Nat. Commun.* **6**, 6979 (2015).
- Knill, E., Lafamme, R., Martinez, R. & Negrevergne, C. Benchmarking quantum computers: the five-qubit error correcting code. *Phys. Rev. Lett.* **86**, 5811 (2001).
- Yao, X.-C. et al. Experimental demonstration of topological error correction. *Nature* **482**, 489 (2012).
- Grassl, M., Beth, T. & Pellizzari, T. Codes for the quantum erasure channel. *Phys. Rev. A* **56**, 33–38 (1997).
- Brown, N. C. & Brown, K. R. Comparing Zeeman qubits to hyperfine qubits in the context of the surface code: $^{174}\text{Yb}^+$ and $^{171}\text{Yb}^+$. *Phys. Rev. A* **97**, 052301 (2018).
- Lu, C.-Y. et al. Experimental quantum coding against qubit loss error. *Proc. Natl Acad. Sci. USA* **105**, 11050–11054 (2008).
- Bell, B. A. et al. Experimental demonstration of a graph state quantum error-correction code. *Nat. Commun.* **5**, 3658 (2014).
- Morley-Short, S. et al. Physical-depth architectural requirements for generating universal photonic cluster states. *Quant. Sci. Tech.* **3**, 015005 (2018).
- Kitaev, A. Fault-tolerant quantum computation by anyons. *Ann. Phys.* **303**, 2–30 (2003).
- Dennis, E., Kitaev, A., Landahl, A. & Preskill, J. Topological quantum memory. *J. Math. Phys.* **43**, 4452 (2002).
- Schindler, P. et al. A quantum information processor with trapped ions. *New J. Phys.* **15**, 123012 (2013).
- Fowler, A. G. Coping with qubit leakage in topological codes. *Phys. Rev. A* **88**, 042308 (2013).
- Epstein, J. M., Cross, A. W., Magesan, E. & Gambetta, J. M. Investigating the limits of randomized benchmarking protocols. *Phys. Rev. A* **89**, 062321 (2014).
- Xia, T. et al. Randomized benchmarking of single-qubit gates in a 2D array of neutral-atom qubits. *Phys. Rev. Lett.* **114**, 100503 (2015).
- Kwon, M., Ebert, M. F., Walker, T. G. & Saffman, M. Parallel low-loss measurement of multiple atomic qubits. *Phys. Rev. Lett.* **119**, 180504 (2017).
- Brown, N. C. & Brown, K. R. Leakage mitigation for quantum error correction using a mixed qubit scheme. *Phys. Rev. A* **100**, 032325 (2019).
- Hayes, D. et al. Eliminating leakage errors in hyperfine qubits. *Phys. Rev. Lett.* **124**, 170501 (2020).
- Ghosh, J., Fowler, A. G., Martinis, J. M. & Geller, M. R. Understanding the effects of leakage in superconducting quantum-error-detection circuits. *Phys. Rev. A* **88**, 062329 (2013).
- Stace, T. M., Barrett, S. D. & Doherty, A. C. Thresholds for topological codes in the presence of loss. *Phys. Rev. Lett.* **102**, 200501 (2009).
- Varbanov, B. M. et al. Leakage detection for a transmon-based surface code. Preprint at <https://arxiv.org/abs/2002.07119> (2020).
- Bombin, H. & Martin-Delgado, M. A. Topological quantum distillation. *Phys. Rev. Lett.* **97**, 180501 (2006).
- Vodola, D., Amaro, D., Martin-Delgado, M. A. & Müller, M. Twins percolation for qubit losses in topological color codes. *Phys. Rev. Lett.* **121**, 060501 (2018).
- Calderbank, A. R. & Shor, P. W. Good quantum error-correcting codes exist. *Phys. Rev. A* **54**, 1098–1105 (1996).
- Steane, A. M. Error correcting codes in quantum theory. *Phys. Rev. Lett.* **77**, 793–797 (1996).
- Mølmer, K. & Sørensen, A. Multiparticle entanglement of hot trapped ions. *Phys. Rev. Lett.* **82**, 1835–1838 (1999).
- Choi, M.-D. Completely positive linear maps on complex matrices. *Linear Algebra Appl.* **10**, 285–290 (1975).
- Knill, E. Quantum computing with realistically noisy devices. *Nature* **434**, 39–44 (2005).
- Aliferis, P., Gottesman, D. & Preskill, J. Quantum accuracy threshold for concatenated distance-3 codes. *Quantum Inf. Comput.* **6**, 97–165 (2006).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© The Author(s), under exclusive licence to Springer Nature Limited 2020

Data availability

The data underlying the findings of this work are available at <https://doi.org/10.5281/zenodo.3900057>. Source data are provided with this paper.

Code availability

All codes used for data analysis are available from the corresponding author upon reasonable request.

Acknowledgements We gratefully acknowledge funding by the US Army Research Office (ARO) through grant number W911NF-14-1-0103. We also acknowledge funding by the Austrian Science Fund (FWF), through the SFB BeyondC (FWF Project number F71), by the Austrian Research Promotion Agency (FFG) contract 872766, by the EU H2020-FETFLAG-2018-03 under Grant Agreement number 820495, and by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via US ARO Grant number W911NF-16-1-0070. All statements of fact, opinions or conclusions contained herein are those of the authors and should not be construed as representing the official views or policies of

ODNI, the IARPA, or the US Government. We acknowledge support from the Samsung Advanced Institute of Technology Global Research Outreach. This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement number 801110 and the Austrian Federal Ministry of Education, Science and Research (BMBWF). The information provided in this Article reflects only the authors' views; the EU Agency is not responsible for any use that may be made of this information.

Author contributions D.V. and M. Müller derived the theory results. R.S., A.E., L.P., M. Meth, M.R., P.S. and T.M. performed the experiments. R.S. analysed the data. T.M., M. Müller and R.B. supervised the project. All authors contributed to the writing of the manuscript.

Competing interests The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41586-020-2667-0>.

Correspondence and requests for materials should be addressed to R.S.

Peer review information *Nature* thanks Tom Stace and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.

Reprints and permissions information is available at <http://www.nature.com/reprints>.

Supplementary Information: Experimental deterministic correction of qubit loss

Roman Stricker,¹ Davide Vodola,^{2,3} Alexander Erhard,¹ Lukas Postler,¹ Michael Meth,¹ Martin Ringbauer,¹ Philipp Schindler,¹ Thomas Monz,^{1,4} Markus Müller,^{5,6,3} and Rainer Blatt^{1,7}

¹Institut für Experimentalphysik, Universität Innsbruck, Technikerstr. 25, A-6020 Innsbruck, Austria

²Dipartimento di Fisica e Astronomia dell'Università di Bologna, and INFN, Sezione di Bologna, I-40127 Bologna, Italy

³Department of Physics, College of Science, Swansea University, Singleton Park, SA2 8PP Swansea, United Kingdom

⁴Alpine Quantum Technologies GmbH, 6020 Innsbruck, Austria

⁵Institute for Quantum Information, RWTH Aachen University, D-52056 Aachen, Germany

⁶Peter Grünberg Institute, Theoretical Nanoelectronics, Forschungszentrum Jülich, D-52425 Jülich, Germany

⁷Institut für Quantenoptik und Quanteninformation, Österreichische Akademie der Wissenschaften, Otto-Hittmair-Platz 1, A-6020 Innsbruck, Austria

Here we provide further experimental and theoretical results and details on the detection and correction of qubit loss. We start in Sec. I by presenting the quantum circuit specifically tailored for the toolbox given by our ion-trap quantum computer. We continue in Sec. II by explaining our approach to hide certain ions from the dynamics of collective Mølmer-Sørensen entangling gates as well as collective readout operations by shelving their population in Zeeman sublevels outside the computational subspace. In Sec. III we discuss the effective dynamics of the QND qubit loss detection scheme and deliver experimental data characterizing these dynamics. In Sec. IV we provide complementary results on the full 1+4-qubit detection and correction algorithm for a larger number of logical input states and for in total three different qubit loss rates. In Sec. V we present a model, which accounts for dominant experimental imperfections in the QND loss detection circuit and discuss how these limit the performance for current system parameters in the regime of low qubit loss rates.

I. CIRCUIT REPRESENTATION OF THE 1+4 QUBIT LOSS DETECTION AND CORRECTION ALGORITHM

The smallest instance for implementing a correction from qubit losses in the surface code is defined by four physical qubits forming a logical qubit in one plaquette as shown Fig. 2a in the main text. An additional ancilla qubit is required for QND loss detection. This leads to the 1+4-qubit loss detection and correction algorithm under study in the main text. In our detection and correction protocol, loss is considered to happen on qubit 1 only.




Figure S1. Gate sequence of the 1+4-qubit loss detection and correction algorithm. **a**, Encoding sequence implementing the smallest excerpt of Kitaev's surface code employing 4 physical qubits. Logical states of the form $|\psi_L\rangle = \cos(\alpha/2)|0_L\rangle + i\sin(\alpha/2)|1_L\rangle$ with $|0_L\rangle = (|0000\rangle + |1111\rangle)/\sqrt{2}$ and $|1_L\rangle = (|0001\rangle + |1110\rangle)/\sqrt{2}$ are encoded and loss is induced in a controlled fashion on qubit 1. **b**, QND detection unit identifying potential loss events on qubit 1. Conditional on the loss detection, our control scheme either keeps the original code or triggers a real-time deterministic code restoration via feed-forward. **c**, The depicted gate sequence aims at measuring the shrunk stabilizer $\tilde{S}_1^X = X_2X_3X_4$ to reconstruct the code in the smaller subset of the remaining three qubits. For this purpose we reuse the ancilla qubit from the detection-circuit, since it remains unaffected by the measurement in the loss case.

The related 5-qubit gate sequence, optimized for our ion-trap quantum computer, to encode an arbitrary logical input state of the form $|\psi_L\rangle = \cos(\alpha/2)|0_L\rangle + i\sin(\alpha/2)|1_L\rangle$ is depicted in Fig. S1a. In our experimental toolbox the

Mølmer-Sørensen entangling gate operations $\text{MS}^X(\theta) = \exp(-i\theta \sum_{j<\ell} X_j X_\ell / 2)$ ³¹ acts as the entangling gate for multi-qubit operations. A fully-entangling gate $\text{MS}^X(\pi/2)$, acting on all four code qubits, alongside local operations on qubit 4 lead to the GHZ-type logical basis states $|0_L\rangle = (|0000\rangle + |1111\rangle)/\sqrt{2}$ and $|1_L\rangle = (|0001\rangle + |1110\rangle)/\sqrt{2}$.

The subsequent QND loss detection unit in part B of Fig. S1 combines a 2-qubit $\text{MS}^X(\pi)$ followed by a collective bit-flip $\text{R}^X(\pi) = X$. In the absence of loss the $\text{MS}^X(\pi)$ performs a bit-flip on both qubits present in the detection scheme. Whenever qubit 1 is outside the computational subspace spanned by $|0\rangle$ and $|1\rangle$, i.e. loss occurs, the MS-gate couples only to the ancilla qubit performing an identity operation, as can be seen from the argument of the exponential $\sum_{j<\ell} X_j X_\ell = X_j X_j = I$ in the above definition of the MS-gate. The subsequent X operation flips the state of the ancilla qubit to $|1\rangle$ followed by its addressed readout signaling the event of loss. If no loss was detected both gates add up to an overall identity operation leaving the logical encoding unaffected. In this way information about loss is mapped onto the ancilla qubit, which can be read out without influencing the logical encoding. Consequently, the described unit works in a quantum non demolition (QND) way. By probing all code qubits sequentially, one could extend this protocol to check the entire register for loss.

In the absence of loss, the logical encoding remains intact and can be verified by measuring the generators of the stabilizer group $\{S_1^Z = Z_1 Z_2, S_2^Z = Z_1 Z_3, S_1^X = X_1 X_2 X_3 X_4\}$ as well as the logical operators $\{T^Z = Z_1 Z_4, T^X = X_4, T^Y = iT^X T^Z\}$ of the original encoding. If loss is detected on qubit 1, the encoded logical information can be restored by switching to an encoding defined on a smaller subset of three qubits, see Fig. 2a in the main text. The merged Z stabilizer $\tilde{S}_1^Z = S_1^Z S_2^Z = Z_2 Z_3$ and a new X stabilizer $\tilde{S}_1^X = X_2 X_3 X_4$ are introduced. This newly defined shrunk stabilizer $\tilde{S}_1^X = X_2 X_3 X_4$ is, after the loss of qubit 1, in an undetermined state and needs to be measured to initialize the stabilizer in a $+1$ (or -1) eigenstate. Here, the -1 case requires a redefinition of the Pauli basis, as so called Pauli frame update^{33, 34}. The respective gate sequence mapping the syndrome onto the ancilla qubit, which is then read out, is shown in Fig. S1c. For this purpose we reuse the ancilla qubit from the QND detection unit, since it remains unaffected by the projective measurement in the loss case. Finally the logical encoding is restored in a new encoding defined by the remaining three qubits.

II. SPECTROSCOPIC DECOUPLING AND RECOUPLING OF IONS

The circuit for the QND loss detection, depicted in Fig. S1b, requires a 2-qubit entangling operation $\text{MS}^X(\pi)$. This entangling gate operation is performed by a collective laser beam illuminating the entire ion string. However, these operations can be applied to a subset of qubits, by temporarily shelving the electronic populations of qubits not taking part in Zeeman sublevels outside the computational subspace. More precisely, population from the lower qubit state $S_{1/2}(m = -1/2) = |0\rangle$ is spectroscopically decoupled to $D_{5/2}(m = +1/2)$ and population from the upper qubit state $D_{5/2}(m = -1/2) = |1\rangle$ is spectroscopically decoupled to $S_{1/2}(m = +1/2)$. In the main text we refer to this as hiding and unhiding operations. A similar technique can be applied to read out individual qubits in the register without influencing the other qubits. Such addressed readout of (ancilla) qubits is essential to allow us to detect a qubit loss event and trigger a subsequent correction step via feed-forward.

III. QND LOSS DETECTION

This section begins with providing theoretical details of the protocol that introduces the loss and on the QND loss detection. We then also present additional experimental data characterizing the QND loss detection.

Loss from one computational basis state. – The controlled loss operation on a code qubit (q) from $|0\rangle$ is realized by coherently transferring qubit population partially from the computational subspace spanned by $\{|0\rangle = S_{1/2}(m = -1/2)$ and $|1\rangle = D_{5/2}(m = -1/2)\}$ into the state $|2\rangle = D_{5/2}(m = -5/2)$ via a coherent rotation $\text{R}_{\text{loss}}(\phi)$

$$\text{R}_{\text{loss}}(\phi) = |1\rangle\langle 1|_q + \cos\frac{\phi}{2} (|0\rangle\langle 0|_q + |2\rangle\langle 2|_q) + \sin\frac{\phi}{2} (|0\rangle\langle 2|_q - |2\rangle\langle 0|_q). \quad (\text{S1})$$

The QND loss detection is realized by the circuit shown in Fig. S1b. It consists of an MS-gate operation $\text{MS}^X(\pi)$ between the code qubit (q) and the ancilla qubit (a) initially prepared in $|0\rangle$, followed by single-qubit bit flips $\text{R}^X(\pi)$ applied to both the ancilla and the code qubit, and a projective measurement of the ancilla qubit in the computational basis. The two-qubit MS-gate applied to the data qubit (q) and the ancilla qubit (a) realizes the unitary

$$\text{MS}^X(\phi) = \exp\left(-i\frac{\phi}{2} X_a X_q\right) = \left[\cos\left(\frac{\phi}{2}\right) (1 - |2\rangle\langle 2|_q) - i\sin\left(\frac{\phi}{2}\right) X_a X_q\right] + |2\rangle\langle 2|_q, \quad (\text{S2})$$

3

which is generated by $X_i = |0\rangle\langle 1|_i + |1\rangle\langle 0|_i$, for $i = q, a$, respectively, and reduces for $\phi = \pi$ to $\text{MS}^X(\pi) = |2\rangle\langle 2|_q - iX_aX_q$. Note that if both the code qubit q and the ancilla qubit a are initially in the computational subspace, this two-qubit operation realizes a collective bit flip (within the computational subspace). In contrast, if the code qubit is in $|2\rangle$, i.e. outside the computational subspace, the state of the code and ancilla qubit remains unchanged under this operation³¹.

The subsequent single-qubit rotations $\text{R}^X(\pi)$ (bit flips) on both the data and the ancilla qubits are realized by

$$\text{R}_a^X(\pi) = -i(|0\rangle\langle 1|_a + |1\rangle\langle 0|_a) \quad (\text{S3})$$

$$\text{R}_q^X(\pi) = |2\rangle\langle 2|_q - i(|0\rangle\langle 1|_q + |1\rangle\langle 0|_q) \quad (\text{S4})$$

and the final unitary evolution will be given by

$$U = \text{R}_a^X(\pi)\text{R}_q^X(\pi)\text{MS}^X(\pi)\text{R}_{\text{loss}}(\phi) = A^{(0)} \otimes \mathbb{1}_a + A^{(1)} \otimes X_a \quad (\text{S5})$$

where

$$A_q^{(0)} = |1\rangle\langle 1|_q + \cos\frac{\phi}{2}|0\rangle\langle 0|_q + \sin\frac{\phi}{2}|0\rangle\langle 2|_q, \quad (\text{S6})$$

$$A_q^{(1)} = \sin\frac{\phi}{2}|2\rangle\langle 0|_q - \cos\frac{\phi}{2}|2\rangle\langle 2|_q. \quad (\text{S7})$$

If we assume that no population is present initially in the $|2\rangle_q$ state the operators $U^{(0)}$ and $U^{(1)}$ will reduce to

$$A_q^{(0)} = |1\rangle\langle 1|_q + \cos\frac{\phi}{2}|0\rangle\langle 0|_q, \quad (\text{S8})$$

$$A_q^{(1)} = \sin\frac{\phi}{2}|2\rangle\langle 0|_q. \quad (\text{S9})$$

The single qubit process arising from the QND measurement and acting on the code qubit q can be then described by two maps \mathcal{E}_0 and \mathcal{E}_1 defined as follows

$$\mathcal{E}_0: \rho \mapsto A_q^{(0)} \rho A_q^{(0)\dagger} \quad (\text{S10})$$

$$\mathcal{E}_1: \rho \mapsto A_q^{(1)} \rho A_q^{(1)\dagger} \quad (\text{S11})$$

and effectively acting on the system of code qubits as

$$\rho \mapsto \mathcal{E}_0(\rho) \otimes |0\rangle\langle 0|_a + \mathcal{E}_1(\rho) \otimes |1\rangle\langle 1|_a. \quad (\text{S12})$$

This single-qubit dynamics can be also described in the Choi representation by the following single-qubit Choi matrices in the elementary basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$

$$\Phi^{(0)} = \frac{1}{2} \begin{pmatrix} \cos^2\frac{\phi}{2} & 0 & 0 & \cos\frac{\phi}{2} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ \cos\frac{\phi}{2} & 0 & 0 & 1 \end{pmatrix}, \quad \Phi^{(1)} = \frac{1}{2} \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & \sin^2\frac{\phi}{2} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \quad (\text{S13})$$

If we now consider the effects of the controlled loss and subsequent QND loss detection on the logical states we will have that after the measurement of the ancilla the logical states $|0_L\rangle$, the $|1_L\rangle$ and the $|+i_L\rangle = (|0_L\rangle + i|1_L\rangle)/\sqrt{2}$ will become

$$|0_L\rangle \otimes |0_a\rangle \mapsto \begin{cases} |2000\rangle \otimes |1_a\rangle & \text{with probability } p_L = \frac{1}{2} \sin^2\frac{\phi}{2} \\ \frac{\cos\frac{\phi}{2}}{(1 + \cos^2(\phi/2))^{1/2}} (|0000\rangle + |1111\rangle) \otimes |0_a\rangle & \text{with probability } 1 - p_L \end{cases} \quad (\text{S14})$$

$$|1_L\rangle \otimes |0_a\rangle \mapsto \begin{cases} |2001\rangle \otimes |1_a\rangle & \text{with probability } p_L = \frac{1}{2} \sin^2\frac{\phi}{2} \\ \frac{\cos\frac{\phi}{2}}{(1 + \cos^2(\phi/2))^{1/2}} (|0001\rangle + |1110\rangle) \otimes |0_a\rangle & \text{with probability } 1 - p_L \end{cases} \quad (\text{S15})$$

$$|+i_L\rangle \otimes |0_a\rangle \mapsto \begin{cases} \frac{|2000\rangle + i|2001\rangle}{\sqrt{2}} \otimes |1_a\rangle & \text{with probability } p_L = \frac{1}{2} \sin^2\frac{\phi}{2} \\ \frac{\cos\frac{\phi}{2}}{(2 + 2\cos^2(\phi/2))^{1/2}} (|0000\rangle + i|0001\rangle) + \frac{|1111\rangle + i|1110\rangle}{(2 + 2\cos^2(\phi/2))^{1/2}} \otimes |0_a\rangle & \text{with probability } 1 - p_L \end{cases} \quad (\text{S16})$$

Note that for example for the four data qubits initially prepared in the $|+i_L\rangle$ state, and if the ancilla qubit is found in the state $|0\rangle_a$ (i.e. no loss detected), this non-unitary time evolution results in the following (ideal) expectation value of the X -type stabilizer

$$\langle S_1^X \rangle = \langle X_1 X_2 X_3 X_4 \rangle = \frac{4 \cos(\phi/2)}{3 + \cos \phi} \approx 1 - \frac{\phi^4}{128}, \quad (\text{S17})$$

where the approximation in the last step holds for small loss rates, i.e. $\phi \ll 1$.

Symmetric loss from both computational basis states. – Complementary to the scenario of loss from state $|0\rangle$ only, we also investigate the effects of loss occurring symmetrically from *both* computational basis states, and the subsequent QND detection for this case. To this end, we realize a loss operation from $|1\rangle$ by transferring qubit population via the coherent rotation

$$\tilde{\mathbf{R}}_{\text{loss}}(\phi) = |0\rangle\langle 0|_q + \cos \frac{\phi}{2} (|1\rangle\langle 1|_q + |2\rangle\langle 2|_q) + \sin \frac{\phi}{2} (|1\rangle\langle 2|_q - |2\rangle\langle 1|_q), \quad (\text{S18})$$

which can be realized by interleaving the Induce loss operation in Fig. S1a between two single qubit bit-flip operations X_q . The latter effectively interchanges the roles of states $|0\rangle$ and $|1\rangle$.

In this case the single qubit process arising from the QND measurement can be still described by two maps $\tilde{\mathcal{E}}_0$ and $\tilde{\mathcal{E}}_1$ similar to the ones in Eq. (S10) and Eq. (S11) where the non-unitary operators read

$$\tilde{A}_q^{(0)} = |0\rangle\langle 0|_q + \cos \frac{\phi}{2} |1\rangle\langle 1|_q, \quad (\text{S19})$$

$$\tilde{A}_q^{(1)} = \sin \frac{\phi}{2} |2\rangle\langle 1|_q. \quad (\text{S20})$$

and the Choi matrices of the single-qubit dynamics take the form

$$\tilde{\Phi}^{(0)} = \frac{1}{2} \begin{pmatrix} 1 & 0 & 0 & \cos \frac{\phi}{2} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ \cos \frac{\phi}{2} & 0 & 0 & \cos^2 \frac{\phi}{2} \end{pmatrix}, \quad \tilde{\Phi}^{(1)} = \frac{1}{2} \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & \sin^2 \frac{\phi}{2} & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \quad (\text{S21})$$

With the two realizations of the loss channel described above, we can implement two different scenarios, namely (i) a symmetric loss channel and (ii) a quantum erasure channel:

(i) For the symmetric loss channel the loss is implemented via Eq. (S1) with a probability of 50% and via Eq. (S18) with a probability of 50%. In this case the data qubit state will be mapped, if no loss is detected in the QND measurement, to

$$\rho \mapsto \frac{1}{2} (\mathcal{E}_0(\rho) + \tilde{\mathcal{E}}_0(\rho)) \quad (\text{S22})$$

and the corresponding Choi matrix will be

$$\Phi_{\text{sym}}^{(0)} = \frac{1}{2} (\Phi^{(0)} + \tilde{\Phi}^{(0)}) = \frac{1}{4} \begin{pmatrix} 1 + \cos^2 \frac{\phi}{2} & 0 & 0 & 2 \cos \frac{\phi}{2} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 2 \cos \frac{\phi}{2} & 0 & 0 & 1 + \cos^2 \frac{\phi}{2} \end{pmatrix}. \quad (\text{S23})$$

Experimental data from the implementation of this symmetric loss channel is discussed below, in Sec. IV.

(ii) The quantum erasure channel can be implemented in the following way: first a loss event is realized via Eq. (S1), i.e. by inducing partial loss from the state $|0\rangle$, then this is followed by the QND loss measurement. Only if this measurement indicates that no loss has been detected, another partial loss via Eq. (S18), i.e. in this case now from the state $|1\rangle$ is induced, followed by a subsequent QND measurement. The overall sequence can be described by the following process

$$\rho \mapsto (1 - p_L)(1 - \tilde{p}_L) \tilde{A}_q^{(0)} A_q^{(0)} \rho A_q^{(0)\dagger} \tilde{A}_q^{(0)\dagger} + (1 - p_L) \tilde{p}_L \tilde{A}_q^{(1)} \rho \tilde{A}_q^{(1)\dagger} + p_L A_q^{(1)} \rho A_q^{(1)\dagger} \quad (\text{S24})$$

where, for a data qubit in an arbitrary initial state $\alpha |0\rangle_q + \beta |1\rangle_q$, the probabilities are

$$p_L = |\alpha|^2 \sin^2 \frac{\phi}{2} \quad \tilde{p}_L = \frac{|\beta|^2 \sin^2(\phi/2)}{(|\alpha|^2 \cos^2(\phi/2) + |\beta|^2)}. \quad (\text{S25})$$

In this case the process (S24) reduces, as desired, to

$$\rho \mapsto \cos^2 \frac{\phi}{2} \rho + \sin^2 \frac{\phi}{2} |2\rangle \langle 2|_q, \quad (\text{S26})$$

where the data qubit is lost (i.e. ending in state $|2\rangle_q$) with probability $p_{\text{loss}} = \sin^2 \frac{\phi}{2}$, and remains unaffected with probability $1 - p_{\text{loss}}$, independent of its initial state.

Characterization of the experimental loss detection operation. – In the following, complementary experimental data characterizing the QND loss detection unit depicted in Fig. 3 in the main text is presented. We start by further analyzing the performance of mapping loss onto the ancilla qubit. Results presented so far in Fig. 3 in the main text were performed on the full 5-qubit string according to the loss detection and correction circuit. To study the effect of the extra three qubits, i.e. the effect of imperfect hiding and unhiding operations, we repeat the experiments isolated on a 2-qubit string. The loss detection sub-circuit is tested by driving the loss transition $R_{\text{loss}}(\phi)$ on qubit 1 and measuring the population in the $D_{5/2}$ -state on both qubit 1 and ancilla qubit. This measurement does not distinguish between the different Zeeman sublevels of the $D_{5/2}$ -state manifold. Fig. S2 shows that in both cases loss detected by the ancilla qubit matches the loss induced on qubit 1 within statistical uncertainties. The quantified detection efficiency for the full 5-qubit string is 96.5(4)%, with a false positive rate of $3(\pm 1)\%$ and a false negative rate of $1(\pm 1)\%$. In the 2-qubit case the detection efficiency is 99.6(3)% with a false positive rate of $0.6(\pm 0.6)\%$ and a false negative rate of $0.2(\pm 0.1)\%$. The difference in detection efficiency is mainly due to imperfect hiding and unhiding operations induced by single-qubit addressing errors. The error bars in Fig. 3b in the main text and Fig. S2 correspond to 1 standard deviation of statistical uncertainty due to quantum projection noise. For the 5-qubit (2-qubit) case 200 (100) experimental cycles were implemented.




Figure S2. Investigating the performance of the 2-qubit QND loss detection unit from Fig. 3b. In addition to the results on the full 5-qubit string depicted in Fig. 3b in the main text we compare to an isolated experiment on 2-qubits only. Population in the $D_{5/2}$ -state of qubit 1 (directly measured loss) and ancilla qubit (detected loss) measured after loss detection. Controlled loss up to 100% with respect to $|0\rangle$ was introduced. The estimated detection efficiencies for the 5-qubit and 2-qubit system are 96.5(4)% and 99.6(3)%, respectively. Error bars correspond to 1 standard deviation of statistical uncertainty due to quantum projection noise. This demonstrates that the occurrence of a loss event can be reliably mapped onto the ancilla qubit and read out in a QND fashion.

Next, we present our experimental findings on the single qubit process describing the QND detection according to Eq. (S13) for the loss from $|0\rangle$ and according to Eq. (S23) for the symmetric loss. We explicitly focus on the non-unitary map $\Phi^{(0)}$ characterizing the no loss case. Therefore generalized single qubit quantum process tomography was applied to qubit 1, whereupon the single qubit Choi matrices were reconstructed in the elementary basis $\{|00\rangle, \dots, |11\rangle\}$. Experiments were implemented on both the full 5-qubit string as well as isolated on a 2-qubit string. The estimated process fidelities with the ideal non-unitary map $\Phi^{(0)}$ are shown in Fig. S3a together with plots of the associated reconstructed single qubit Choi matrices for loss rates $\phi \in \{0.10\pi, 0.53\pi, 0.81\pi\}$ in Fig. S3b. In order to estimate the uncertainty of the values presented here and in Fig. 3c in the main text we re-sample the data given by generalized quantum process tomography via a multinomial distribution and assigned the respective standard deviation, received from 100 iterations, as the statistical uncertainty.

Data for the non-unitary map $\Phi_{\text{sym}}^{(0)}$ together with the reconstructed single-qubit Choi matrices for the symmetric loss channel are shown in Fig. S4. The data clearly show the symmetric behavior, as expected for equal-weighted loss from both computational basis states. They also show the expected deviation from the identity operation, which - similar as for loss from $|0\rangle$ only - becomes more pronounced for higher loss rates.




Figure S3. Loss from $|0\rangle$: Tomographic reconstruction of the non-unitary single qubit map $\Phi^{(0)}$ of Eq. S13, characterizing the QND measurement in the no loss case. a, Process-fidelities of the non-unitary map $\Phi^{(0)}$, when working on both the full 5-qubit string and isolated on 2-qubits only. Error bars correspond to 1 standard deviation of statistical uncertainty due to quantum projection noise. **b**, Single qubit choi matrices $\Phi^{(0)}$ in the elementary basis $\{|00\rangle, \dots, |11\rangle\}$, reconstructed from the full 5-qubit string for loss rates indicated by the circles in Fig. a above. The Ideal Choi-operators, according to the map $\Phi^{(0)}$ of Eq. S13, are denoted by the underlying black frames.

In the final paragraph of this section we investigate the effect of the controlled loss on the logical state after the ancilla measurement. To follow this idea, we initialize $|+i_L\rangle$ and proceed with the loss detection as shown in Fig. S1. Controlled loss between $\phi = 0.1\pi$ and π is introduced on qubit 1. We study the case where we find the ancilla qubit in $|0\rangle$, i.e. in the absence of loss. Fig. S5 shows the results on the expectation values of the stabilizer generators S_1^X , S_1^Z and S_2^Z . The maximum expectation values of the Z-stabilizers remain unaffected by the loss, whereas the expectation value for S_1^X drops for increasing loss rates ϕ according to Eqs. S16 and S17. The underlying modelled curve for S_1^X represents the ideal outcome biased with the experimentally measured S_1^X value, extracted from the lowest loss rate at $\phi = 0.1\pi$. The error bars correspond to 1 standard deviation of statistical uncertainty due to quantum projection noise. In total 200 experimental cycles were implemented. The results show that the experiment and theory predictions of the effect of loss and QND detection are in good agreement.

IV. ADDITIONAL DATA AND EXPERIMENTAL INFORMATION ON THE 1+4-QUBIT LOSS DETECTION AND CORRECTION ALGORITHM

Here, we provide complementary results on the full 1+4-qubit loss detection and correction algorithm for the logical input states $\{|0_L\rangle, |1_L\rangle, |+i_L\rangle = (|0_L\rangle + i|1_L\rangle)/\sqrt{2}\}$ under three different loss rates $\phi \in \{0.1\pi, 0.2\pi, 0.5\pi\}$. Next to fidelities, expectation values for stabilizers and logical operators the remaining population in the code space P_{CS} was estimated according to $\hat{P}_{\text{CS}}|\psi\rangle = P_{\text{CS}}|\psi\rangle$. Here, \hat{P}_{CS} represents the projector onto the code space, defined as the simultaneous




Figure S4. Loss from $|0\rangle$ and $|1\rangle$ (symmetric): Tomographic reconstruction of the non-unitary single qubit map $\Phi_{\text{sym}}^{(0)}$ of Eq. (S23), characterizing the QND measurement in the no loss case. **a, Process-fidelities of the non-unitary map $\Phi_{\text{sym}}^{(0)}$, when working on both the full 5-qubit string and isolated on 2-qubits only. Error bars correspond to 1 standard deviation of statistical uncertainty due to quantum projection noise. **b**, Single qubit Choi matrices $\Phi_{\text{sym}}^{(0)}$ in the elementary basis $\{|00\rangle, \dots, |11\rangle\}$, reconstructed from the full 5-qubit string for loss rates indicated by the circles in Fig. a above. The Ideal Choi-operators, according to the map $\Phi_{\text{sym}}^{(0)}$ of Eq. (S23), are denoted by the underlying black frames.**

+1 eigenspace given by all generators of the stabilizer group $\{S_1^X, S_1^Z, S_2^Z\}$ and $\{\tilde{S}_1^X, \tilde{S}_1^Z\}$ for the 4-qubit and reconstructed reconstructed 3-qubit logical encoding, respectively. The general expression for the code space projector reads:

$$\hat{P}_{\text{CS}} = \prod_i \frac{1}{2}(1 + S_i^X) \prod_j \frac{1}{2}(1 + S_j^Z) \quad (\text{S27})$$

which becomes

$$\hat{P}_{\text{CS}} = \frac{1}{8}(1 + S_1^X + S_1^Z + S_2^Z + S_1^X S_1^Z S_2^Z + S_1^X S_1^Z + S_1^X S_2^Z + S_1^Z S_2^Z) \quad (\text{S28})$$

with $S_1^X = X_1 X_2 X_3 X_4$, $S_1^Z = Z_1 Z_2$, $S_2^Z = Z_1 Z_3$, $S_1^Z = Z_1 Z_2$, $S_1^X S_1^Z S_2^Z = -X_1 Y_2 Y_3 X_4$, $S_1^X S_1^Z = -Y_1 Y_2 X_3 X_4$, $S_1^X S_2^Z = -Y_1 X_2 Y_3 X_4$, and $S_1^Z S_2^Z = Z_2 Z_3$ for the original code, and

$$\hat{\tilde{P}}_{\text{CS}} = \frac{1}{4}(1 + \tilde{S}_1^X + \tilde{S}_1^Z + \tilde{S}_1^X \tilde{S}_1^Z) \quad (\text{S29})$$

with $\tilde{S}_1^X = X_2 X_3 X_4$, $\tilde{S}_1^Z = Z_2 Z_3$ and $\tilde{S}_1^X \tilde{S}_1^Z = -Y_2 Y_3 X_4$ for the reconstructed code.

All results presented in Fig. 2 in the main text and in Tabs. S1, S2 and S3 were extracted from full 4-qubit quantum state tomography using linear state reconstruction technique. In order to estimate the uncertainty of these values we re-sample the data given by quantum state tomography via a multinomial distribution and assigned the respective standard deviation, received from 100 iterations, as the statistical uncertainty. In order to receive enough experimental data under both loss cases for state reconstruction, we adjusted the number of experimental cycles depending on the induced loss rate. The corresponding values on cycle numbers read: 1000 cycles for $\phi = 0.1\pi$, 600 cycles for $\phi = 0.2\pi$ and 200 cycles for




Figure S5. Experimental investigations on the effect of the controlled loss on the logical state $|+i_L\rangle$ after the ancilla qubit measurement. We explicitly study the no loss case. The maximum expectation values for the Z-stabilizers remain unaffected by the loss, whereas the expectation value for S_1^X drops with increasing loss rate according to Eqs. S16 and S17. The simulated curve represents the ideal outcome biased with the experimental measured S_1^X , extracted from the lowest loss rate at $\phi = 0.1\pi$. The error bars correspond to 1 standard deviation of statistical uncertainty due to quantum projection noise.

$\phi = 0.5\pi$. Tabs. S1, S2 and S3 contain the entire data gained on the 1+4-qubit loss detection and correction algorithm. Each table is assigned to one of the logical input states $\{|0_L\rangle, |1_L\rangle, |+i_L\rangle = (|0_L\rangle + i|1_L\rangle)/\sqrt{2}\}$ and includes data on in total three different loss rates $\phi \in \{0.1\pi, 0.2\pi, 0.5\pi\}$.

encoding											
	P_{CS}	S_1^X	S_1^Z	S_2^Z	$S_1^X S_1^Z S_2^Z$	$S_1^X S_1^Z$	$S_1^X S_2^Z$	$S_1^Z S_2^Z$	T^X	T^Y	T^Z
	0.88(1)	0.84(6)	0.95(1)	0.94(1)	0.80(6)	0.80(6)	0.78(6)	0.95(1)	0.01(2)	-0.01(3)	0.94(1)
no-loss											
$\phi (\pi)$	P_{CS}	S_1^X	S_1^Z	S_2^Z	$S_1^X S_1^Z S_2^Z$	$S_1^X S_1^Z$	$S_1^X S_2^Z$	$S_1^Z S_2^Z$	T^X	T^Y	T^Z
0.1	0.74(1)	0.59(2)	0.85(1)	0.86(1)	0.58(3)	0.59(3)	0.51(3)	0.90(1)	-0.02(1)	-0.13(1)	0.84(1)
0.2	0.72(1)	0.57(4)	0.87(1)	0.87(1)	0.48(3)	0.57(3)	0.49(3)	0.91(1)	-0.02(1)	-0.14(1)	0.84(1)
0.5	0.67(2)	0.40(8)	0.80(2)	0.82(2)	0.52(7)	0.46(7)	0.45(7)	0.92(1)	-0.06(2)	-0.10(3)	0.79(2)
loss											
$\phi (\pi)$	P_{CS}	\tilde{S}_1^X	\tilde{S}_1^Z		$\tilde{S}_1^X \tilde{S}_1^Z$			\tilde{T}^X	\tilde{T}^Y	\tilde{T}^Z	
0.1	0.43(4)	0.19(16)	0.60(5)		-0.05(16)			0.00(3)	0.06(5)	0.53(4)	
0.2	0.61(5)	0.53(15)	0.51(5)		0.41(15)			0.00(3)	-0.06(6)	0.63(4)	
0.5	0.66(4)	0.63(11)	0.59(4)		0.40(12)			0.00(2)	-0.05(4)	0.65(3)	

Table S1. logical state $|0_L\rangle$: Complementary experimental data on the 1+4-qubit loss detection and correction algorithm (see Fig. 2 in the main text) including results on three different loss rates ϕ .

V. IMPERFECTIONS IN THE QND LOSS DETECTION

From the data in the previous section in the case of low loss rates, namely $\phi \in \{0.1\pi, 0.2\pi\}$, we find that the probability of success for reconstructing the code after loss is lower than for the higher loss rate $\phi = 0.5\pi$. We relate this to imperfections in the QND loss detection unit. Let's assume the error on the detection unit is of the same order as the loss

encoding											
	P_{CS}	S_1^X	S_1^Z	S_2^Z	$S_1^X S_1^Z S_2^Z$	$S_1^X S_1^Z$	$S_1^X S_2^Z$	$S_1^Z S_2^Z$	T^X	T^Y	T^Z
	0.88(1)	0.74(8)	0.94(1)	0.95(1)	0.74(7)	0.82(6)	0.88(5)	0.95(1)	-0.01(2)	0.04(3)	-0.93(1)
no-loss											
ϕ (π)	P_{CS}	S_1^X	S_1^Z	S_2^Z	$S_1^X S_1^Z S_2^Z$	$S_1^X S_1^Z$	$S_1^X S_2^Z$	$S_1^Z S_2^Z$	T^X	T^Y	T^Z
0.1	0.78(1)	0.68(3)	0.86(1)	0.85(1)	0.70(2)	0.63(3)	0.64(2)	0.89(1)	-0.02(1)	0.10(1)	-0.83(1)
0.2	0.78(1)	0.67(3)	0.86(1)	0.86(1)	0.64(3)	0.65(3)	0.63(3)	0.90(1)	-0.03(1)	0.09(2)	-0.81(1)
0.5	0.70(1)	0.61(6)	0.78(2)	0.80(2)	0.57(7)	0.51(7)	0.46(7)	0.89(1)	0.04(1)	0.07(3)	-0.74(2)
loss											
ϕ (π)	P_{CS}	\tilde{S}_1^X	\tilde{S}_1^Z	$\tilde{S}_1^X \tilde{S}_1^Z$				\tilde{T}^X	\tilde{T}^Y	\tilde{T}^Z	
0.1	0.41(6)	0.23(18)	0.44(5)	-0.01(16)				0.00(3)	-0.04(5)	-0.43(6)	
0.2	0.61(5)	0.49(14)	0.61(4)	0.36(14)				-0.02(3)	0.01(5)	-0.59(5)	
0.5	0.78(3)	0.80(8)	0.80(3)	0.52(13)				-0.09(3)	0.09(5)	-0.72(3)	

Table S2. logical state $|1_L\rangle$: Complementary experimental data on the 1+4-qubit loss detection and correction algorithm (see Fig. 2 in the main text) including results on three different loss rates ϕ .

encoding											
	P_{CS}	S_1^X	S_1^Z	S_2^Z	$S_1^X S_1^Z S_2^Z$	$S_1^X S_1^Z$	$S_1^X S_2^Z$	$S_1^Z S_2^Z$	T^X	T^Y	T^Z
	0.84(2)	0.74(5)	0.94(1)	0.92(1)	0.78(6)	0.66(7)	0.78(6)	0.92(1)	0.01(2)	0.92(1)	-0.04(3)
no-loss											
ϕ (π)	P_{CS}	S_1^X	S_1^Z	S_2^Z	$S_1^X S_1^Z S_2^Z$	$S_1^X S_1^Z$	$S_1^X S_2^Z$	$S_1^Z S_2^Z$	T^X	T^Y	T^Z
0.1	0.76(1)	0.60(2)	0.88(1)	0.87(1)	0.63(2)	0.60(2)	0.61(3)	0.91(1)	-0.02(1)	0.81(1)	0.00(1)
0.2	0.72(1)	0.58(4)	0.88(1)	0.85(1)	0.52(4)	0.52(4)	0.55(4)	0.89(1)	-0.03(1)	0.79(1)	0.01(1)
0.5	0.66(1)	0.48(3)	0.82(1)	0.82(1)	0.48(4)	0.39(4)	0.41(4)	0.90(1)	-0.07(1)	0.75(1)	0.02(1)
loss											
ϕ (π)	P_{CS}	\tilde{S}_1^X	\tilde{S}_1^Z	$\tilde{S}_1^X \tilde{S}_1^Z$				\tilde{T}^X	\tilde{T}^Y	\tilde{T}^Z	
0.1	0.49(6)	0.25(16)	0.57(5)	0.15(15)				0.11(3)	0.51(5)	-0.01(7)	
0.2	0.61(5)	0.55(14)	0.56(4)	0.33(13)				0.13(3)	0.47(4)	0.06(5)	
0.5	0.80(2)	0.81(4)	0.80(2)	0.59(6)				0.13(1)	0.80(2)	0.01(2)	

Table S3. logical state $|+i_L\rangle$: Complementary experimental data on the 1+4-qubit loss detection and correction algorithm (see Fig. 2 in the main text) including results on three different loss rates ϕ .

rate, then many of the measurement cycles detected as loss will be false positives. This limits the performance for current system parameters in the regime of low qubit loss rates.

In order to quantitatively study this effect, we model imperfections in the QND loss detection by a depolarizing noise-channel on each individual qubit. Since loss is induced on the lower qubit state $S_{1/2}(m = -1/2) = |0\rangle$, complete loss ($\phi = \pi$) of this state leads to an overall loss rate of 50% for the encoded GHZ-state, where half of the population occupies the $D_{5/2}(m = -1/2) = |1\rangle$ state. Taking this into account (with a factor 0.5 in front of the \sin^2 -term) our model reads:

$$\rho \longrightarrow \frac{p}{3} \sum_{k=1}^3 M_{(k)}(\rho) + (1-p)\rho \quad \text{with} \quad M_{(k)}(\rho) = \frac{1}{4} \sum_{i \in \{x,y,z,id\}} \sigma_i^{(k)\dagger} \rho \sigma_i^{(k)} \quad (\text{S30})$$

$$\text{and} \quad p = \frac{p_{\text{QND}}}{p_{\text{QND}} + 0.5 \sin^2(\phi/2)}. \quad (\text{S31})$$

The error-rate p_{QND} includes all noise collected throughout the QND loss detection unit, i.e. errors on the specific gates for mapping loss onto the ancilla qubit, hiding and unhiding operations on qubits 1,2,3 and 4 for the ancilla qubit's

addressed readout complemented by errors on the ancilla readout itself. In Fig. S5 we plot the model against the measured data prepared in logical $|1_L\rangle$, previously presented in Tab. S2. Based on this experimental data we estimate an error-rate of $p_{\text{QND}} = 3.3\%$.




Figure S6. Modelling an imperfect QND-detection scheme under the assumption of an independent depolarizing noise channel on each qubit. The theoretical model (lines) shows good agreement with the experimental data (points with error bars). The error bars correspond to 1 standard deviation of statistical uncertainty due quantum projection noise.

The imperfections predominantly originate from addressing errors when hiding and unhiding the qubits preventing them from taking part in the QND loss detection. For faulty experimental shots, where one of those qubits is not hidden in the upper $D_{5/2}(m = +1/2)$ level, it will affect the loss detection measurement. Hence it is likely to happen that the particular experimental cycle assigns to the wrong loss case. By improving the addressing optics such errors could be further suppressed.

5.4 NON-IDENTITY DYNAMICS OF THE NO-LOSS CASE

For trapped ion devices, it can often be useful to engage larger dimensions through coherent operations to outside the computational subspace, so leakage rates at the level of computational errors are to be expected [52]. For instance, when spectroscopically decoupling certain constituents from subtasks or for addressed readouts [88, 92, 218]. Errors in these operations can result in leakage, denoting the main loss mechanism on many of the leading quantum computer architectures, see Ch. 1.4. In our first attempt to characterize the no-loss map of the QND loss detection unit from Fig. 5.2(b) we have thus considered a loss probability at the level of computational errors of 0.012 [52]. Notably, this loss was induced from computational basis state $|0\rangle$, i.e., asymmetrically with respect to the computational subspace. The resulting no-loss map was subsequently analyzed by QPT with MLE process reconstruction under the CPTP constraint according to Eq. (3.14). The resulting no-loss dynamic had a process fidelity of 0.90(2) with respect to the identity operation, which is within the expected range given the experimental noise as well as the additional operational overhead for implementing process tomography.

Further loss detection studies presented in our publication from Sec. 5.3 probed scenarios up to complete asymmetric loss and revealed non-identity dynamics of the no-loss map. These deviations from identity were initially masked by the experimental noise floor, before becoming visually accentuated towards higher loss probabilities up to 0.5, see Fig. S3 in Sec. 5.3. In the latter case, the no-loss maps beyond 0.95 asymmetric loss have less than 0.5 process fidelity with the expected identity map and covered up non-unitary dynamics which we dismissed when first analyzing the detection unit.

Let us take a closer look at the asymmetric loss discussed here. Consider the loss process as a coherent operation between $|0\rangle$ and a third level outside the computational subspace. This leakage level is then disregarded in the process reconstruction, whereby the coherent loss becomes incoherent. Consequently, the evolution restricted to the qubit subspace becomes

$$\rho \rightarrow E\rho E^\dagger \quad \text{with} \quad E = |1\rangle\langle 1| + \cos(\phi/2)|0\rangle\langle 0|, \quad (5.4)$$

which only for small loss rates $\phi \approx 0$ converges to the identity map $|0\rangle\langle 0| + |1\rangle\langle 1|$. Eq. (5.4) describes the loss of population from the computational subspace, leading to a non-unitary process at the qubit level whenever $\phi > 0$. Importantly, the process is no longer normalized to 1. The non-unitary dynamics are therefore owed to the information gain about loss from the projective ancilla qubit measurement.

In particular, the destructive measurement character interrupts the unitary evolution and yields either of the two measurement outcome related side-channels (see Eq. (1.35)) of the loss detection unit (see Fig. 5.1), namely loss or no-loss map. The subsequent MLE reconstruction, which is constrained at the qubit level, incorrectly forces normalized maps by including the *trace preserving* (TP) constraint. Therefore, the reconstruction of the channel map fails if the third leakage level is disregarded, making the process non-unitary in the qubit subspace. On the other hand, it becomes very costly to include all potential leakage levels in the tomography procedure, which is why the non-unitary dynamics discussed here are often dismissed in the system analysis. These dynamics, if undetected, potentially deteriorate the logical information in QEC applications. We discussed these findings carefully in Ref. [131] along the correct relaxed tomography procedure—particularly by lifting the TP constraint in the MLE reconstruction.

The presented loss detection experiments featured in-sequence measurement and classical feed-forward, classifying them within the semi-classical quantum algorithms, introduced in Ch. 1.5. Examples of semi-classical algorithms can be found primarily in




Figure 5.10: **Loss-corrected QEC codes potentially carry non-unitary elements requiring system characterizations with quantum instruments to fully capture their dynamics.** Beyond QEC, today’s quantum operational tasks increasingly feature in-sequence measurements and classical feed-forward. Due to the quantum measurement’s destructive nature, in-sequence measurements result in classical data. The measurement outcome related side-channels (see Eq. (1.35)) become potentially influenced by non-unitary elements and have to be correctly described with the framework of quantum instruments [129], see text for details. Our QND loss detection unit exemplifies such a quantum instrument by which we guide our rigorous characterization demonstrations, presented in Sec 5.5. Their results allow us to numerically estimate the parameter regimes for correctable loss errors under computational ones. These numerics are performed on a minimal color code instance—illustrated on the bottom left of the present figure and outlined in Sec. 5.4.1.

situations where it is important to keep track of the classical measurement outcomes, most prominently in QND measurements [83, 219].

The dynamics of semi-classical algorithms are correctly described by quantum instruments from Eq. (1.35), which is a collection of CP maps that only together become TP, i.e. unitary. Quantum instruments capture both the classical and quantum outputs of a computation, allowing the inclusion of destructive measurements in the analysis process. In particular, quantum instruments account for trace-nonincreasing, i.e., non-unitary side-channels, as observed, for example, in the characterization of our QND loss detection unit, discussed above. The overall unitary map of a quantum instrument can only be reconstructed when accounting for all employed qubit or qudit levels in the analysis procedure. Such reconstruction is referred to as quantum instrument QPT and allows the identification of subtle dynamical properties that go unnoticed by standard methods but which are necessary for high-precision applications such as realistic QEC. Yet, only observing the full quantum instrument reveals information about when and under what conditions regular QPT on a measurement outcome related side-channel is admissible.

Seeing the increasing occurrence of the semi-classical algorithm structure in modern quantum tasks, we decided to do a follow-up work to develop a general toolbox for quantum-instrument QPT. This aims at identifying erroneous mechanisms such as the above dismissed non-identity dynamics of the no-loss map. This work leads to the fourth publication of this thesis, presented in Sec. 5.5.

Along those lines, we experimentally demonstrate our newly developed quantum-instrument QPT using the QND loss detection unit as an example instrument, where we investigate two cases that demonstrate undetected failure under the standard QPT methods from Ch. 3.1.2. Those include a situation where additional levels beyond the computational subspace are ignored, as well as the no-loss cases, where non-TP maps are involved. In addition, a symmetric loss channel, leaking equally from both computational basis states, is investigated. The extracted information leads to the development of an experimentally informed noise model that accurately explains the dominant technical limitations of our loss experiments. Furthermore, we numerically simulate loss affected QEC codes to quantify the parameter regimes for simultaneously correctable loss and computational errors in view of current NISQ-hardware. For these numerics, we slightly switch gears and study loss in the so-called color code [72]—outlined in the next paragraph.

5.4.1 The planar color code

The color code, invented by Bombin and Delgado in Refs. [72, 73] results from a topological concatenation of the Steane code [220]. In its planar version, physical qubits reside on a 2D lattice forming a logical qubit similar to the surface code explained in Ch. 1.3.3. The geometric structure of the planar color code goes back to the Steane code and thus similarly enables the transversal implementation of the entire *Clifford-group* [221]. The latter is a set of operations that opens access to efficient implementations of numerous applications. Examples are quantum distillation, quantum teleportation or superdense coding [18].

A fault-tolerant universal gate-set [221] can therefore be realized with the color code by complementing the transversal Clifford-group with only a single additional non-Clifford-gate, the so-called *T-gate*. The T-gate, for instance defined by $T_L = e^{-i\pi/8Y_L}$, can be implemented in a fault-tolerant way through a technique called *magic state injection*, as discussed and experimentally demonstrated in Ref. [70]. Although very resource-demanding [222], by enabling a universal gate-set, magic state injection in principle paves the way towards arbitrary fault-tolerant applications with planar color codes.

In a 3D representation of the color code [73], the T-gate becomes transversal, whereas the Clifford-group is not anymore. A universal gate-set could thus be alternatively established through so-called *code switching* between 2D and 3D representations of the color code, precluding the necessity of magic state injection, explained in Ref. [92].

The smallest planar color code instance and the one we are going to study in the work of Sec. 5.5 involves the seven qubit logical instance depicted in the bottom left of Fig. 5.10. The code comprises three links that converge at each vertex and by that separates three differently colored plaquettes of the lattice, such that adjacent plaquettes are distinguished in color—explaining the code’s name. Physical qubits reside at each vertex and together form a logical qubit. Each plaquette inhabits two stabilizers of four-body Pauli X- and Z operators. A single logical qubit can be defined by the remaining degree of freedom after projecting all six stabilizers onto their common +1 eigenstate. Logical operators are further given in transversal fashion by $X_L = \prod_{i=1}^7 X_i$ and $Z_L = \prod_{i=1}^7 Z_i$, while the logical code words are $|0_L\rangle \propto \prod_{i=1}^3 (\mathbb{1} + S_i^x) |0\rangle^{\otimes 7}$ and $|1_L\rangle = X_L |0_L\rangle$. This smallest instance denotes a distance three code and therefore offers to correct a single computational error.

As Fig. 5.10 indicates, the transversal implementation of the Clifford-group comes at the cost of dropping the hardware-friendly next-neighbor interaction structure, potentially leading to operational overhead in experimental implementations. Yet, the color code can remain efficient on trapped-ion devices that feature arbitrary connectivity along the ion-string employed. Multiple experimental efforts thus far demonstrate the correction of computational errors with the color code, while some of them feature trapped-ion devices [70, 88]

Apart from computational errors, the color code shows noteworthy robustness against loss errors. A study entirely dedicated to losses in color codes can be found in Ref. [209], where the authors derive a threshold of $p = 0.461 \pm 0.005$ in the limit of an infinitely large lattice. Similar to surface codes, loss robustness in color codes closely relates to percolation theory on lattices [217]. Finally, the minimal code fragment from Fig. 5.10 tolerates losses on two arbitrary qubits or losses on at least some of the three or four qubit subsets, further outlined in the next section.

5.5 PUBLICATION: CHARACTERIZING QUANTUM INSTRUMENTS

***PRX Quantum* 3, 030318 (2022)**

submitted on 22 October 2021, accepted on 13 July 2022 and published on 03 August 2022
<https://doi.org/10.1103/PRXQuantum.3.030318>

Roman Stricker¹, Davide Vodola^{2,3}, Alexander Erhard¹, Lukas Postler¹, Michael Meth¹,
Martin Ringbauer¹, Philipp Schindler¹, Rainer Blatt^{1,4,5}, Markus Müller^{6,7} and Thomas
Monz^{1,5}

¹ *Institut für Experimentalphysik, Universität Innsbruck, Technikerstr. 25, A-6020 Innsbruck, Austria*

² *Dipartimento di Fisica e Astronomia dell'Università di Bologna*

³ *INFN, Sezione di Bologna, I-40127 Bologna, Italy*

⁴ *Institut für Quantenoptik und Quanteninformation, Österreichische Akademie der Wissenschaften, Otto-Hittmair-Platz 1, A-6020 Innsbruck, Austria*

⁵ *Alpine Quantum Technologies GmbH, 6020 Innsbruck, Austria*

⁶ *Institute for Quantum Information, RWTH Aachen University, D-52056 Aachen, Germany*

⁷ *Peter Grünberg Institute, Theoretical Nanoelectronics, Forschungszentrum Jülich, D-52425 Jülich, Germany*

The author to the present thesis executed the experiments, analyzed the data and wrote the manuscript.

Characterizing Quantum Instruments: From Nondemolition Measurements to Quantum Error Correction

Roman Stricker^{1,*}, Davide Vodola^{2,3}, Alexander Erhard¹, Lukas Postler¹, Michael Meth¹, Martin Ringbauer¹, Philipp Schindler¹, Rainer Blatt^{1,4,5}, Markus Müller^{6,7}, and Thomas Monz^{1,5}

¹*Institut für Experimentalphysik, Universität Innsbruck, Technikerstr. 25, Innsbruck A-6020, Austria*

²*Dipartimento di Fisica e Astronomia, Università di Bologna, Bologna I-40129, Italy*

³*INFN, Sezione di Bologna, Bologna I-40127, Italy*

⁴*Institut für Quantenoptik und Quanteninformation, Österreichische Akademie der Wissenschaften, Otto-Hittmair-Platz 1, Innsbruck A-6020, Austria*

⁵*Alpine Quantum Technologies GmbH, Innsbruck 6020, Austria*

⁶*Institute for Quantum Information, RWTH Aachen University, Aachen D-52056, Germany*

⁷*Peter Grünberg Institute, Theoretical Nanoelectronics, Forschungszentrum Jülich, Jülich D-52425, Germany*



(Received 22 October 2021; accepted 13 July 2022; published 3 August 2022)

In advanced quantum processors, quantum operations are increasingly processed along multiple in-sequence measurements that result in classical data and affect the rest of the computation. Because of the information gain of classical measurements, nonunitary dynamical processes can affect the system, which common quantum channel descriptions fail to describe faithfully. Quantum measurements are correctly treated by so-called quantum instruments, capturing both classical outputs and postmeasurement quantum states. Here we present a general recipe for characterizing quantum instruments and demonstrate its experimental implementation and analysis. Thereby the full dynamics of a quantum instrument can be captured, exhibiting details of the quantum dynamics that would be overlooked with standard techniques. For illustration, we apply our characterization technique to a quantum instrument used for the detection of qubit loss and leakage, which was recently implemented as a building block in a quantum error-correction (QEC) experiment [Nature 585, 207 (2020)]. Our analysis reveals unexpected and in-depth information about the failure modes of the implementation of the quantum instrument. We then numerically study the implications of these experimental failure modes on QEC performance, when the instrument is employed as a building block in QEC protocols on a logical qubit. Our results highlight the importance of careful characterization and modeling of failure modes in quantum instruments, as compared to simplistic hardware-agnostic phenomenological noise models, which fail to predict the undesired behavior of faulty quantum instruments. The presented methods and results are directly applicable to generic quantum instruments and will be beneficial to many complex and high-precision applications.

DOI: [10.1103/PRXQuantum.3.030318](https://doi.org/10.1103/PRXQuantum.3.030318)

The field of quantum computation progresses rapidly and experiments demonstrate ever more complex tasks. The majority of experiments have focused on a unitary evolution of quantum systems together with a single final measurement. In modern systems, the time evolution of a computation may get repetitively interrupted by in-sequence measurements and circuit adaptation conditional

on the result. These measurements are required for many classes of semiclassical algorithms, including quantum error correction (QEC) [1–7], resource-efficient Fourier transform [8,9], and measurement-based quantum computing [10]. Owing to its destructive nature, a quantum measurement produces classical data and changes the quantum state in a nonunitary fashion. Operations including such in-sequence measurements therefore deviate from simple linear unitary evolution and can no longer be described with commonly used methods. For prime examples such as QEC codes or quantum nondemolition (QND) measurements [11–17], it is important to keep track of the measurement outcome in each experimental cycle. More subtly, experimental imperfections in realizations of quantum operations are often caused by undesired coupling to

*Corresponding author. roman.stricker@uibk.ac.at

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

other quantum systems [18–24]. As a consequence, the operations that are performed on what is considered a qubit usually feature a small nonunitary component due to coupling to and ignorance of other relevant degrees of freedom. Such small deviations often go unnoticed when enforcing a unitary description onto the system [19].

The correct framework to describe such quantum classical operations is given by so-called *quantum instruments* [25,26]. A quantum instrument includes both quantum and classical inputs as well as outputs and thereby offers a unified description of state preparation, operations, and quantum measurements [27]. Quantum instruments are commonly used to describe scenarios where one needs to keep track of a classical input or output of a quantum operation, e.g., in the description of quantum networks [28], quantum causality [29], measurement uncertainty trade-offs [30,31], and weak measurements [32–34].

So far, device-independent [35] and self-testing [36,37] protocols have been developed to assess the performance of positive operator-valued measures [38,39] and quantum instruments. However, these methods do not give full information on the dynamics that is required in the context of present high-precision quantum computation [40–46] and QEC.

Here, we present a characterization method for quantum instruments that will be particularly useful to characterize building blocks of quantum information processors. We identify quantum instruments where conventional quantum process tomography fails and introduce relaxed tomography procedures beyond the underlying computational subspace (e.g., qubit levels), suitable for completely reconstructing such quantum instruments. We contrast instrument reconstruction to conventional quantum process tomography that typically applies some form of maximum likelihood estimation (MLE) and demonstrate those to bear the risk of unfaithful reconstruction potentially incorporating nonphysical results. This becomes particularly crucial in high-performance applications. Our detailed experimental analysis is guided by a very general example of a QND measurement dedicated to the detection of qubit loss and leakage featuring in our recent work [11] and conveys processes that can drastically deteriorate the performance of QEC codes, if these loss mechanisms go unnoticed [47–49]. These findings apply to generic QND measurements just as well, featuring, for example, in leakage studies beyond trapped ions [50], real-time stabilizer measurements [51,52] for QEC or in metrology applications [53].

Based on an experimental quantum instrument reconstruction using a modified process tomography scheme, we derive a full instrument description for a faulty QND loss detection unit. We numerically study its effect on a QEC cycle on a low-distance near-term logical qubit. This instrument tomography proves to be particularly useful for assessing the QEC performance, since it allows

us to evaluate the effects of different microscopic processes in the loss detection and to estimate the parameter regimes where QEC becomes beneficial. Importantly, those detailed noise dynamics are only accessible upon full instrument reconstruction, while remaining mostly hidden to conventional process tomography. Although the parameters of the precise modeling remain implementation specific, its results reveal general scaling properties, such as the impact of false-positive and false-negative events on error correcting code performances. Those properties together with our tool and its workflow apply to other architectures just as well.

Our results further highlight the importance of developing microscopic, experimentally informed noise models of faulty quantum instruments over widely used generic hardware-agnostic noise models such as dephasing or depolarizing noise channels.

The methods we develop provide the tools and theoretical framework to reconstruct and characterize quantum instruments, such as QND measurements, which have a prominent role in all quantum computing architectures even beyond QEC, as, for instance, in quantum information and quantum metrology [54].

I. INTRODUCTION TO QUANTUM INSTRUMENTS

Formally, a quantum instrument \mathcal{I} is a set I of trace non-increasing, completely positive (CP) maps $\{\mathcal{E}_j\}_{j \in I}$, labeled by an index $j \in I$, with the property that their sum is trace preserving (TP), $\text{Tr}(\sum_j \mathcal{E}_j(\rho)) = \text{Tr}(\rho)$ for every state ρ ; see Fig. 1. For example, when \mathcal{I} describes a quantum measurement, then $j \in I$ labels the measurement outcomes and \mathcal{E}_j transforms the input state ρ to the eigenstate corresponding to outcome j . In this case, each \mathcal{E}_j will generally be trace decreasing, while the sum of all \mathcal{E}_j will be trace preserving for any orthonormal measurement basis. Performing, for example, a computational basis measurement via an ancilla defines two CP maps $\mathcal{E}_0 = |0\rangle\langle 0|$ and $\mathcal{E}_1 = |1\rangle\langle 1|$. Both are trace decreasing (except for computational basis states), since they measure the overlap between ρ and the computational basis states, but their sum must be trace preserving as they form a complete basis and the measured probabilities must add up to 1.

The quantum instrument for a measurement $\mathcal{I}: \mathcal{H}_1 \mapsto \mathcal{H}_2 \otimes \mathbb{C}^{|I|}$ thus maps the input Hilbert space \mathcal{H}_1 to an output Hilbert space \mathcal{H}_2 of potentially different size, and a classical space $\mathbb{C}^{|I|}$. In practice, one might realize such a measurement by coupling the system to an ancilla and subsequently measuring the ancilla with a set of orthogonal projectors $|j\rangle\langle j| \in \mathbb{C}^{|I|}$; see Fig. 1. This final ancilla measurement extracts the classical measurement outcome j , which identifies which operation \mathcal{E}_j was applied to the system. In the following we focus on the simplest case with two possible measurement outcomes ($|I| = 2$), but




FIG. 1. Example of a quantum instrument. The operation U realizes a generic quantum instrument on the system in initial state ρ_q and writes index j of the applied operation into the state of an ancilla initialized in state $|0\rangle_a$. The ancilla is finally projected onto the computational basis to read out the classical index.

all results can be straightforwardly extended to the general case.

II. TOMOGRAPHY OF QUANTUM INSTRUMENTS

For qubit systems, complete information on their quantum evolution can be gained by quantum process tomography [55]. However, when the evolution is described by a quantum instrument, the constituent maps are, in general, not individually trace preserving. For example, if leakage from the qubit level is present, the tomographic measurements do not probe the full Hilbert space. In this case standard reconstruction techniques such as maximum likelihood estimation [56,57] will not be able to describe the quantum dynamics faithfully, because they force the reconstructed map to be trace preserving. To approach this problem, we rely on a relaxed tomography algorithm that does not enforce trace preservation [19,58,59].

In order to reconstruct the quantum channel \mathcal{E} , we make use of the Choi-Jamiolkowsky (CJ) isomorphism [60] to relate \mathcal{E} to an (unnormalized) map Λ , the *Choi operator*. The correspondence between Λ and \mathcal{E} is given by

$$\mathcal{E}(\rho) = \text{Tr}_1[(\rho^T \otimes \mathbb{I})\Lambda].$$

The Choi operator Λ with respect to the basis $\{|k\rangle\}_{k=0}^{d-1}$ can be explicitly constructed as

$$\Lambda = \sum_{k,l} |k\rangle\langle l| \otimes \mathcal{E}(|k\rangle\langle l|),$$

where d is the dimension of the Hilbert space. Following the notation of Ref. [19], the probability $p_{i,j}$ for observing outcome state ρ_j after preparing state ρ_i and subjecting it to the non-trace-preserving channel described by the Choi

operator Λ is given by

$$\begin{aligned} p_{ij} &= \text{Tr}[\rho_j^\dagger \text{Tr}_1[(\rho_i^T \otimes \mathbb{I})\Lambda]] \\ &= \text{Tr}[(\rho_i^T \otimes \rho_j^\dagger)\Lambda]. \end{aligned} \quad (1)$$

Defining the projector $\Pi_{ij} \equiv \rho_i^* \otimes \rho_j$ with ρ_i and ρ_j representing pure states alongside the column vector $|\Lambda\rangle\rangle = \sum_{i,j} \Lambda_{i,j} |j\rangle \otimes |i\rangle$, obtained by stacking the columns of Λ (similarly for $|\Pi_{ij}\rangle\rangle$), we can identify the trace in Eq. (1) with an inner product of the vectorized operators:

$$p_{ij} = \langle\langle \Pi_{ij} | \Lambda \rangle\rangle. \quad (2)$$

We now define the vector of observed frequencies $|f\rangle$, and the quadratic form S , as

$$\begin{aligned} |f\rangle &= \sum_{i,j} f_{ij} |i,j\rangle, \\ S &= \sum_{i,j} |i,j\rangle \langle\langle \Pi_{ij} |. \end{aligned}$$

The most direct way to reconstruct the non-trace-preserving Choi operator Λ is by inverting the above relation, a technique known as *linear inversion*,

$$\hat{\Lambda} = \arg \min_{\Lambda} \|S|\Lambda\rangle\rangle - |f\rangle\|_2, \quad (3)$$

where $\|\cdot\|_2$ denotes the vector 2-norm, and the estimator $\hat{\Lambda}$ is analytically given by

$$\hat{\Lambda} = \sum_{i,j} p_{ij} \left(\sum_{l,m} |\Pi_{lm}\rangle\rangle \langle\langle \Pi_{lm} | \right)^{-1} |\Pi_{ij}\rangle\rangle.$$

Unfortunately, linear inversion can produce nonphysical results, especially in situations where the true (Choi) state is close to pure [56]. To avoid these problems, we can use modified maximum likelihood estimation by constraining the estimator to be positive semidefinite, i.e., a physical state:

$$\begin{aligned} &\text{minimize} \quad \|WS|\Lambda\rangle\rangle - W|f\rangle\|_2 \\ &\text{subject to} \quad \Lambda \geq 0. \end{aligned} \quad (4)$$

Here $W = \sum_{i,j} \sqrt{N_j/p_j(1-p_j)} |i,j\rangle\langle i,j|$ is a weight matrix, taking into account the multinomial distribution of observed frequencies. Note that in contrast to standard MLE quantum process tomography [61], i.e.,

$$\begin{aligned} &\text{minimize} \quad \|WS|\Lambda\rangle\rangle - W|f\rangle\|_2 \\ &\text{subject to} \quad \Lambda \geq 0, \text{Tr}[\Lambda] = d, \end{aligned} \quad (5)$$

we do not enforce the map to be trace preserving in Eq. (4).

III. EXPERIMENTAL SETUP

The experimental demonstrations are realized on a string of $^{40}\text{Ca}^+$ ions confined in a linear Paul trap in ultrahigh vacuum [62]; see Fig. 2(a). Each ion represents a physical qubit encoded in the metastable electronic states $S_{1/2}(m = -1/2) \equiv |0\rangle$ and $D_{5/2}(m = -1/2) \equiv |1\rangle$, denoting the computational subspace. Upon coherent laser-ion interaction, we realize a universal set of quantum gate operations combining single-qubit rotations by an angle θ around the x or y axis of the Bloch sphere, $R^{\sigma_j}(\theta) = \exp(-i\theta\sigma_j/2)$ with the Pauli operators $\sigma_j = X_j$ or Y_j acting on qubit j , together with two-qubit Mølmer-Sørensen entangling gate operations $R_{ij}^{\text{MS}}(\theta) = \exp(-i\theta X_i X_j/2)$ [63]. Multiple addressed laser beams allow for arbitrary two-qubit connectivity across the entire ion string [64]. Readout is performed through continuous excitation of a dipole transition, solely involving the lower S -state and collecting




FIG. 2. Ion-trap quantum processor and qubit loss detection unit as the example quantum instrument. (a) Schematic of our ion-trap quantum processor, where each ion resides along a linear string representing a single qubit. Quantum gate operations are realized upon coherent laser-ion interaction using tightly focused beams addressing single ions for local gates (bright red) and a pair of ions for entangling gates (dark red). Readout is performed via collective fluorescence detection (DET). See the text for details. (b) A QND qubit loss detection unit as our application example for a quantum instrument. The system qubit is encoded in the computational subspace $\{|0\rangle_q, |1\rangle_q\}$ and is affected by loss to a third level $|2\rangle_q$. For details, see the text. (c) A quantum erasure channel implemented by first inducing partial loss from $|0\rangle_q$ followed by its detection using the gadget from (b). Conditional on the qubit not being lost, the same partial loss is induced from $|1\rangle_q$ and subsequently detected.

its scattered photons, which identifies the qubit's $|0\rangle$ and $|1\rangle$ states. This dipole laser collectively covers the entire ion string. However, we are also able to read out only a subset of the qubit register by shelving electronic populations of certain qubits in the upper D -state manifold, referred to as *addressed readout*. This constitutes an essential building block for realizing the in-sequence detections featuring in QND measurements. Beyond the qubit level, we hold equivalent control over the entire S - and D -state Zeeman manifolds, which allows us to encode a higher-dimensional quantum decimal digit (qudit) in each ion. Implementing our example quantum instrument requires us to take the additional level $D_{5/2}(m = +1/2) \equiv |2\rangle$ into account—forming together with the qubit states a qutrit. A qutrit readout demands for two subsequent measurements to separate both D -state levels, namely $D_{5/2}(m = -1/2) = |1\rangle$ and $D_{5/2}(m = +1/2) = |2\rangle$. Because each measurement scatters photons and heats up the ion string, we counteract every in-sequence measurement with polarization gradient cooling, keeping the quality of postmeasurement gate operations high.

IV. EXAMPLE: QUBIT LOSS DETECTION

We experimentally study an example quantum instrument devised for a QND detection of qubit loss or leakage, which represents a key building block towards fault-tolerant quantum computation. Qubit loss occurs in a variety of physical incarnations such as the actual loss of particles encoding the qubits or chemical reactions that make qubits unutilizable. Those mechanisms occur almost never on experimental timescales as particles can be stably trapped for days and working in ultrahigh vacuum prevents chemical reactions. However, the implementation of quantum computational tasks can often be improved by addressing higher-dimensional states, either to spectroscopically decouple certain constituents (e.g., qubits) from subtasks or to improve the quantum circuit. Furthermore, faulty state initialization bears the risk of leakage to levels outside the computational subspace. This applies architecture independent as all qubits are encoded within multilevel systems. Thus, leakage errors are most representative and typically occur at the same rates as computational errors, making their detection and correction an inevitable challenge. Our example quantum instrument recently played a central role in the experimental detection and correction of qubit loss embedded in a state-of-the-art QEC code [11]. There, the successful detection of a qubit loss event triggered a reconstruction routine, to restore the logical information on the remaining qubits. In the absence of loss, however, the reconstructed maps deviate from the aimed identity operation, owing to the in-sequence ancilla readout, resulting in nonunitary components. When forcing a unitary description, those mechanisms remain undetected

and likely diminish QEC performance. Hence, a proper quantum instrument reconstruction becomes essential.

Before we follow up with the characterization, we deliver essential insights into the nature of our loss detection unit. The dominant loss mechanism in a trapped-ion quantum processor is leakage from the qubit subspace $\{|0\rangle, |1\rangle\}$ to other electronic states, which can occur due to radiative decay from metastable electronic qubit states [65], in Raman transitions [66], or due to imperfections in spectroscopic decoupling pulses [67] when additional electronic states outside the computational subspace are used deliberately. Hence, loss can be induced in a controlled fashion by transferring part of the population from either computational basis state to an auxiliary level $D_{5/2}(m = +1/2) \equiv |2\rangle$, referred to as loss transition $R_{\text{loss}}(\phi)$, denoting a full coherent transfer in the case of $\phi = \pi$. We then apply the QND unit to map the information about a loss of the system qubit (q) onto an ancilla qubit (a), which is subsequently read out. In the language of quantum instruments, this means that one of two possible maps (“loss” or “no loss”) has been applied to the system, with the classical index of the applied map stored in the qubit states $|0\rangle_a$ and $|1\rangle_a$ of the ancilla. Similar QND loss detection protocols have been devised using various other physical platforms [68–70].

Notably, for both ancilla outcomes, the system qubit is subject to a map that is CP, but in general not TP. This nonunitarity of the individual maps leads to several counterintuitive effects. For example, in the present case, the evolution of the system qubit differs from the identity map, even in the case where no loss is detected, if loss occurs asymmetrically, i.e., from only one of the computational basis states. More precisely, for loss restricted to occur from $|0\rangle$, the system qubit follows (up to normalization) a nonunitary evolution given by $\rho_q \mapsto \mathcal{E}_0 \rho_q \mathcal{E}_0^\dagger$ with

$$\mathcal{E}_0 = |1\rangle_q \langle 1| + \cos(\phi/2) |0\rangle_q \langle 0|, \quad (6)$$

considering the coherent loss operation $R_{\text{loss}}(\phi)$. This is a consequence of the information gain in the no-loss case, given by the ancilla measurement [11]. In either case, the reconstruction becomes challenging, since standard reconstruction techniques for quantum process tomography enforce the reconstructed processes to be completely positive and trace preserving, thereby suppressing the deviations from this condition characteristic for quantum instruments. This becomes evident in Fig. 3, where we compare the accuracy of quantum process reconstructions of the “no-loss” dynamics obtained via the standard MLE technique, referred to as the trace-constrained approach from Eq. (5) in contrast to the trace-unconstrained approach from Eq. (4). As a figure of merit we use the total variation distance between the measured frequencies and the measurement outcomes that are predicted from the reconstructed Choi operators. This highlights how the




FIG. 3. Comparison of trace-constrained and trace-unconstrained tomography for the nonunitary map \mathcal{E}_0 from Fig. 2(b). We compute the total variation distance between directly measured frequencies and those predicted from the reconstructed Choi operators. Standard MLE from Eq. (5), referred to as the trace-constrained approach, increasingly fails to capture the underlying dynamics for higher loss probabilities, whereas the trace-unconstrained approach from Eq. (4) matches the predicted outcomes. Error bars correspond to one standard deviation of statistical uncertainty due to quantum projection noise. Further notes characterizing erroneous effects owing to a faulty quantum instrument reconstruction can be found in Fig. 14 of Appendix A 1.

trace-constrained approach can fail to capture the dynamics; an error that might go unnoticed for maps that are close to trace preserving. Further notes on how common tomography fails to capture a quantum instrument’s dynamics is subject to Fig. 14 of Appendix A 1.

V. EXPERIMENTAL RESULTS

We now discuss features associated in experiments with QND measurements that can only be captured using a full description as a quantum instrument. We start by characterizing our example instrument acting on a two-level system (qubit), followed by a complete characterization in a higher-dimensional Hilbert space that captures the entire dynamics of the QND measurement.

A. Implementation of the quantum instrument

We implement the circuit in Fig. 2(b) on a two-ion string studying several input states $\{|0\rangle_q, |- \rangle_{X,q}\} = (1/\sqrt{2})(|0\rangle_q - |1\rangle_q)$, $|- \rangle_{Y,q} = (1/\sqrt{2})(|0\rangle_q - i|1\rangle_q)$, $|1\rangle_q$ on the system qubit for a range of loss probabilities. We apply quantum state tomography for the runs that signal no-loss events, effectively applying the “no-loss” map \mathcal{E}_0 as given by Eq. (6). We focus on the no-loss outcome \mathcal{E}_0




FIG. 4. Bloch vectors after undergoing QND detection in the no-loss case for different loss channels. (a) State vector evolution for asymmetric loss from $|0\rangle_q$ is captured by the color gradient, ranging from 0% loss (bright points) to 100% (dark points) for various input states (i) $|0\rangle_q$, (ii) $|1\rangle_q$, (iii) $|-\rangle_{X,q}$, and (iv) $|-\rangle_{Y,q}$. Notably, the Bloch vectors remain close to the surface of the sphere, independent of the loss probability; see the Appendix A 1. The initial superposition states $|-\rangle_{X,q}$ and $|-\rangle_{Y,q}$ are found transitioning to the basis state not affected by the loss. (b) The erasure channel is realized by consecutively inducing the same partial loss from $|0\rangle_q$ followed by $|1\rangle_q$ and postselecting to both no-loss cases, i.e., both ancilla's $|0\rangle_a$ outcome. Our results support the theory derivation of a map $\propto \rho$ leaving the initial states up to noise unaltered; see the Appendix A 1.

given that in a realistic scenario the system qubit would remain intact, as opposed to the loss case. We find that the superposition input states are *distorted* towards the basis state that is not affected by the loss with increasing loss probability; see Fig. 4(a). This is a consequence to the asymmetry of the loss, occurring only from one basis state, as detailed in Eq. (A9) in the Appendix A 1. Importantly, however, the states display no notable reduction in purity, regardless of the loss probability. More details are given in the Appendix A 1.

The archetypal description of a qubit loss channel features symmetric loss, often referred to as a *quantum erasure channel* [71], where loss occurs with a given probability, irrespective of the qubit state, and the position of the lost qubit is known. Experimentally, we realize this quantum erasure channel sequentially in two steps, by first inducing partial loss from $|0\rangle_q$ followed by its detection, and, conditional on detecting no loss in this first step, inducing the same amount of partial loss, but now from $|1\rangle_q$ in this second step. Experiments are conducted on a three-ion string using a single system qubit (q) and two ancilla qubits a_1 and a_2 as depicted in Fig. 2(c). By observing the evolution of the Bloch vectors in Fig. 4(b) we find that the initial state is preserved up to experimental noise, as derived in Eq. (A12) in the Appendix A 1. The purity is again found independent of the loss; see the Appendix A 1.

These findings are further corroborated by quantum process tomography characterizing the map describing the system qubit dynamics by using the unconstrained reconstruction approach of Eq. (4). In the case of the asymmetric loss previously discussed, the single-qubit Choi operators for the map \mathcal{E}_0 are close to the identity only given little loss on the order of a few percent and clearly deviate for higher loss, revealing their nonunitary behavior; see the left plot in Fig. 5(a) for a low-loss probability and the right plot for a high-loss probability. We note that a standard MLE approach would force unitary maps and thereby prevent the correct reconstruction not displaying this nonunitary behavior. In contrast, for the quantum erasure channel, for both the 2% and 61% loss cases, maps are found close to the identity following the theoretical predictions, depicted in Fig. 5(b).

For higher loss rates, however, we observe a deviation of the reconstructed Choi operator from the predicted channel, quantified by the fidelity between the reconstructed and ideal Choi operator shown in Fig. 5(c). For high loss rates, only few experimental cycles remain in the no-loss case. As a result, error terms, such as state-preparation-and-measurement (SPAM) errors, as well as errors in the implementation of the loss process contribute with a higher relative weight. We can model these additional error terms as depolarizing noise at the level of the Choi operator as a




FIG. 5. Tomographic reconstruction of the maps characterizing our quantum instrument. (a) Single-qubit Choi operators in the elementary basis $\{|00\rangle_q, \dots, |11\rangle_q\}$ describing the QND detection under loss from $|0\rangle_q$. Process fidelities compared to the ideal map for 2% and 61% losses read 0.97(1) and 0.98(1), respectively. Black boxes denote the ideal operator in the higher loss case. (b) On the erasure channel we receive the expected identity map for loss from both qubit states up to about 60% before errors start to dominate. (c) Corresponding process fidelities compared with ideal maps together with the decay model (dashed line) from Eq. (7).

function of the loss rate p_{loss} :

$$\begin{aligned} \Lambda_M(p_{\text{loss}}) \propto & (1 - p_{\text{loss}}) \cdot (1 - p_e) \cdot (1 - p_{\text{spam}}) \cdot \Lambda \\ & + p_{\text{loss}} \cdot p_e \cdot (1 - p_{\text{spam}}) \cdot \mathbb{1}/4 \\ & + p_{\text{spam}} \cdot \mathbb{1}. \end{aligned} \quad (7)$$

Here Λ denotes the ideal Choi operator of the no-loss channel, $\mathbb{1}$ is the identity matrix, representing a fully depolarizing channel, p_e is a generic error rate of the erasure channel, and p_{spam} is the error rate due to SPAM errors. The first term of Eq. (7) describes the ideal channel where no loss happened and the QND detection worked, while the second term is a case where a loss happened, but




FIG. 6. Multiqubit entangled state undergoing QND detection in the no-loss case. As an input, we choose the four-qubit GHZ state $(1/\sqrt{2})(|0000\rangle_q + |1111\rangle_q)$. Loss is induced from $|0\rangle_{q,1}$ on system qubit 1. Results for purity (\circ) and population ratio between the GHZ basis states $|0000\rangle$ and $|1111\rangle$ (\diamond) in analogy to the Bloch-vector picture (Fig. 4) are shown. The purity is found constant, while the population ratio increases towards higher loss probabilities, finally causing a *distortion* to state $|1111\rangle_q$ not affected by the loss. Errors correspond to one standard deviation of statistical uncertainty due to quantum projection noise.

the QND unit failed to detect it as such. The final term describes the contribution from SPAM errors. From a fit to the data, we find that $p_e = 0.09$ and $p_{\text{spam}} = 0.03$ captures the observed drop in fidelity well. From Fig. 5(c) we see that these effects become predominant for high loss rates, while for up to about 60%, a faithful reconstruction of the experimental Choi operator is possible.

The results presented so far cover a single system qubit and reveal potential obstacles of our quantum instrument tomography, which are generally transferable to other experiments utilizing QND measurements. We now go one step further by analyzing these effects on a multiqubit entangled state. Experiments are conducted using four system qubits, initialized in the Greenberger-Horne-Zeilinger (GHZ) state $(1/\sqrt{2})(|0000\rangle_q + |1111\rangle_q)$, accompanied by one ancilla. After state preparation, partial asymmetric loss from $|0\rangle_{q,1}$ on system qubit 1 is induced followed by its detection using the QND-detection unit. The “no-loss” evolution \mathcal{E}_0 is analyzed by four-qubit quantum state tomography. In Fig. 6 the states again show no significant reduction in purity (circles) over the range of measured loss probabilities and by that obscuring the nonunitary effect from our instrument. However, an asymmetric effect is displayed by computing the population ratio of the GHZ basis states $|0000\rangle_q$ and $|1111\rangle_q$ in Fig. 6 (diamonds) showing a *distortion* towards the basis state not affected by loss in analogy to the Bloch vectors in the single-qubit case. The underlying theory curve follows $1 - p_{\text{loss}}$, as can be seen from Eq. (6). The fidelity with the initial GHZ state further remains above 50% within one standard deviation of statistical uncertainty, thus certifying multipartite entanglement independent of the loss probability.

B. Qutrit dynamics and identification of failure modes

The full dynamics of our coherent loss process can be reconstructed by explicitly taking the loss level $|2\rangle_q$ into account. The state of the system ion needs then to be represented by a qutrit with basis states $\{|0\rangle_q, |1\rangle_q, |2\rangle_q\}$. We perform quantum process tomography on the combined system of data qutrit and ancilla qubit. This allows us to study both loss cases by distinguishing the maps dependent upon the ancilla state, and provides more fine-grained

information on the microscopic error processes. The reconstructed Choi operators for both ancilla outcomes and various loss probabilities are given in Fig. 7(a). For the sake of clarity, the operators are color coded by peaks occurring in the absence of loss (blue), peaks denoting the partial loss rotation (orange), and erroneous peaks (red). The latter are restricted to the diagonal for simplicity. Note that these experimentally derived maps on the qutrit level are now faithful descriptions of the instrument, obtained




FIG. 7. Full-system dynamics from combined qutrit-ancilla quantum instrument tomography. (a) Choi operator of the system qutrit evolution in the elementary basis $\{|00\rangle_q, \dots, |22\rangle_q\}$ after postselecting on the ancilla, revealing either loss case (rows) examined for different loss probabilities from $|0\rangle_q$ (columns). The tricolor Choi operators show peaks in the absence of loss (blue), peaks occurring due to partial loss (orange), and erroneous peaks (red). The latter are only color coded on the diagonal for visualization purposes. Process fidelities with the ideal map from top left to bottom right read $\{0.97(1), 0.96(1), 0.95(1), 0.83(1), 0.86(1), 0.84(1)\}$. (b) False-positive and false-negative rates extracted from raw data for loss states $\{|0\rangle_q, |1\rangle_q, 1/\sqrt{2}(|0\rangle_q + |1\rangle_q)\}$ versus the loss probability.

from the qubit-qutrit process that is indeed unitary; see Eq. (A5) in the Appendix A 1.

One key piece of information gained from the full tomography is the dominant failure modes of the experimental realization of the QND-detection unit. In the no-loss case, false negatives are retrieved from diagonal elements $\{|02\rangle_q, |12\rangle_q, |22\rangle_q\}$ corresponding to undetected rotations to level $|2\rangle_q$ outside the computational subspace. Likewise false positives in the loss case are retrieved from the elements $\{|00\rangle_q, |01\rangle_q, |10\rangle_q, |11\rangle_q, |20\rangle_q, |21\rangle_q\}$ corresponding to qubit rotations mistakenly assigned as loss. Note that, for standard tomography restricted to qubit levels, such fine-grained analysis would be precluded for two main reasons. First, the true population in the loss state of the system qutrit cannot be estimated independently from the ancilla outcome in the qubit description. Thus, one cannot reliably assign false-positive and false-negative events by postselecting on the ancilla since some erroneous population adds up to the main peaks $\{|00\rangle_q, |11\rangle_q\}$, blurring the information about the error origin. Second, when tracing over the ancilla, loss state $|2\rangle_q$ would be incoherently added to state $|1\rangle_q$, creating a nonphysical bias under which tomography is likely to break; see the Appendix A 1. For a more quantitative analysis, the corresponding false-positive and false-negative rates are depicted in Fig. 7(b). To avoid errors from the quantum instrument reconstruction, these rates are extracted from the raw data for three different loss states: $\{|0\rangle_q, |1\rangle_q, 1/\sqrt{2}(|0\rangle_q + |1\rangle_q)\}$. Notably, there is a significantly higher false-positive rate owing to their sensitivity on the entangling operation, implementing a correlated two-qubit rotation. This operation shows a higher error rate compared to single-qubit operations [44] and only plays a role in the no-loss case: the reason is that, as under loss, the action of the entangling operation, when it only acts on the ancilla qubit alone, is on purpose trivial and no longer induces a correlated qubit-qutrit flip process. Therefore, the loss map is left with the local bit-flip operations, explaining why false negatives are dominated by single-qubit errors, resulting in smaller rates. For loss detection in a QEC setting, we expect this asymmetry to be quite beneficial, as a false-positive event would merely trigger an unnecessary loss correction, while a false-negative event leads to an undetected loss, which can be catastrophic, i.e., leading directly to uncorrectable logical errors, as will be discussed in the next section.

C. Experimentally informed noise model

We now build noise models to characterize the QND-detection unit, which can then be used to study implications on QEC. From the above phenomenological discussion, we assume that the dominant contributions will come from false-positive and false-negative events, where the latter in particular can have a severe impact. However, extracting the respective rates from tomography data as in

Fig. 7(b) in the presence of SPAM errors can be unreliable if these contributions are of the same magnitude. A rough estimate of the SPAM errors from tomography of the identity yields a fidelity of 0.96(2), which indicates that this is indeed the parameter regime we are dealing with here.

Hence, to describe imperfections in the QND loss detection unit, we instead focus on a microscopic noise model $\mathcal{E}_{\text{noise}}$ defined as (see the Appendix A 2)

$$\rho \mapsto \mathcal{E}_{\text{noise}}(\rho) = U_{\text{noise}}\rho U_{\text{noise}}^\dagger, \quad (8)$$

where the unitary $U_{\text{noise}} = R^{\text{MS}}(\alpha)R^X(\beta)$ describes the dominating error source as correlated bit flips with a rate of $p_{\text{corr}} = \sin^2(\alpha/2)$, resulting from systematic miscalibrations in the two-ion R^{MS} gate, and single-qubit flips with a rate of $p_{\text{single}} = \sin^2(\beta/2)$ from errors in the collective local rotations. Fitting channel $\mathcal{E}_{\text{noise}}$ to the experimental data returns values of $p_{\text{corr}} = 0.045$ and $p_{\text{single}} = 2.47 \times 10^{-4}$, respectively; see the Appendix A 2. The fidelity of the experimental data with respect to this model in the no-loss case is 0.94, compared to 0.91 for the noiseless theory prediction.

In order to validate this model against generic hardware-agnostic noise models typically considered in the quantum information literature, we further add depolarizing and dephasing noise channels [72]. As discussed in detail in the Appendix A 2, by fitting a model that includes all four error channels to the experimental data, we again find the correlated bit-flip error to be dominant. The contributions from depolarizing and dephasing noise are consistently of the order of 0.01 and adding these terms does not significantly improve the fit to the data. From this analysis, we conclude that the microscopic model is the most suitable description of our experimental noise and the resulting imperfections in the QND loss detection, and we thus use this model in the following analysis of the impact of a faulty QND loss detection unit on QEC.

VI. IMPLICATION ON QUANTUM ERROR CORRECTION

In the context of QEC and the pursuit for robust and eventually fault-tolerant quantum computers, qubit leakage and loss errors are known to be particularly harmful to the performance of QEC codes, if they go unnoticed [47–49]. Dedicated protocols to fight qubit loss have been devised, including the four-qubit quantum erasure code [71], which has been implemented in the form of post-selective state analysis protocols using photons [73,74]. Moreover, protocols to cope with qubit loss in elementary quantum codes such as the five-qubit code [75] as well as topological QEC codes including the surface code [76] and color codes [67,77,78] have been developed.

Here, our aims are as follows. (i) To estimate the parameter regimes in which active qubit loss error correction and

detection is expected to reach break even, i.e., to become beneficial for low-distance QEC codes as currently pursued in various efforts [2,11,67,69,79–82]. (ii) Whereas most theory studies exclusively focus on the simple (and ideal) quantum erasure channel to describe loss, we are interested in illustrating the effect of various qualitatively different imperfections in the loss detection process on QEC performance, highlighting the importance of microscopically informed noise models of the components used in QEC of qubit loss. (iii) Finally, to predict the performance of QEC protocols by numerical simulations, it is desirable to develop *effective* few-parameter noise models, informed by experimental data, which can be simulated efficiently, e.g., using stabilizer simulations, to predict the performance of large-scale QEC codes built from noisy components. Here, we are particularly interested to which extent our faulty QND loss detection can be reliably substituted by efficiently simulatable noise models. Whereas the phenomenological studies from Fig. 7(b) pointed to false-positive and false-negative events as the dominant noise contributions, accurately extracting the respective error rates from tomography data is prohibited by SPAM errors. Instead, we here utilize the microscopic noise model covered by Sec. V C, incorporating the dominant error sources of correlated and single-qubit errors. This model best fits our noisy QND loss detection unit, especially in contrast to the widely used generic hardware-agnostic models of dephasing or depolarizing noise that lead to no notable contribution.

A. Qubit loss correction with color codes

To be concrete, we focus on the smallest two-dimensional color code [77], a seven-qubit stabilizer code equivalent to the Steane code [77,83], which is at the focus of current experimental efforts to achieve the break-even point of beneficial and fault-tolerant QEC with low-distance QEC codes [84–87]. The code is obtained by projecting the Hilbert space of seven qubits (Fig. 8) into the +1 eigenspace of six commuting stabilizer generators S_i^x and S_i^z ($i = 1, 2, 3$) [see Fig. 8(a)] that define a two-dimensional code space hosting one logical qubit. Logical X and Z operators are defined as $X_L = \prod_{i=1}^7 X_i$ and $Z_L = \prod_{i=1}^7 Z_i$ and the logical basis states are $|0_L\rangle \propto \prod_{i=1}^3 (\mathbb{1} + S_i^x) |0\rangle^{\otimes 7}$ and $|1_L\rangle = X_L |0_L\rangle$ (see the Appendix A 3). The code is a distance $d = 3$ QEC code ($d = 2n + 1$ with n the number of correctable computational errors), so that one arbitrary computational error (bit and/or phase flip error) on any of the physical qubits is correctable. Note that, besides computational errors, this code also allows one to correct the loss of any two of the seven physical qubits, or even the loss of some, though not all subsets of three or even four qubits (see the Appendix A 3 for more details). We note that, for each of the seven qubits forming the code, we incorporate state $|2\rangle_q$, i.e., adopt a qutrit




FIG. 8. Simulations on the faulty QND loss detection embedded in the seven-qubit color code. (a) A single logical qubit encoded on a triangular planar color code lattice formed of three interconnected plaquettes (lower part). The code space is formed by six stabilizer operators $S_x^{(i)}$ and $S_z^{(i)}$, each acting on a plaquette of four physical qubits [67]. Loss is subsequently detected on all code qubits using a faulty QND circuit (top part). We model this taking into account both correlated and single-qubit over-rotations, representing our leading error mechanisms by treating every qubit as a qutrit. (b) Single QEC cycle of qubit loss detection and correction, including initial controlled induction of loss, followed by faulty QND loss detection operations on the qubit subspace of all physical qutrits and stabilizer measurements triggering respective conditional Pauli corrections.

description, and use this additional level to induce loss of a controllable amount via the coherent rotation in the subspace $\{|0\rangle_q, |2\rangle_q\}$ of the quantum instrument depicted in Fig. 2(b).

We then model one round of qubit loss error detection and correction, depicted in Fig. 8, as follows. Starting from an ideal (noise-free) logical state ρ_L of the seven-qubit code, qubit loss is induced with an independent and equal probability p_{loss} on each of the physical qubits of the register. Subsequently, a noisy QND loss detection unit is sequentially applied to each of the seven qubits, in order to detect the possible occurrence of loss. This faulty unit [Fig. 8(a)] is described by the microscopic noise model $\mathcal{E}_{\text{noise}}(\rho)$ in Eq. (8), where the main error sources are given

by correlated and single-qubit over-rotations with error rate p_{corr} and p_{single} , respectively.

Each data qubit, for which the QND measurement indicates the occurrence of a loss, is replaced by a fresh qubit in the computational basis state $|0\rangle_q$. This is followed by one round of possibly faulty measurements of all six stabilizers of the code. For simplicity, since our focus lies on the QND loss detection, here we model imperfections in each stabilizer measurement by a phenomenological noise model, in which the stabilizer measurement outcome is assumed to be faulty with probability q [88,89]. Since the four-qubit stabilizer operators are typically measured with a circuit involving (at least) four two-qubit gates, we work with 4 times the two-qubit error rate as the error rate of the stabilizer measurement, which results in $q = p_{\text{corr}}$, in what follows. Based on the obtained syndrome (± 1 stabilizer eigenvalues) from the measurement of the stabilizers, Pauli corrections are applied if needed (such a Pauli frame update can be done on the software level and is thus modeled as error-free). Finally, to determine the logical error rate, it is checked whether the original logical state ρ_L has been recovered or not, by evaluating the expectation value of the logical operator corresponding to the initially prepared encoded state.

B. Numerical results

Figure 9 shows the predicted logical error rate of the loss QEC cycle applied to all physical qubits as a function of the physical qubit loss rate p_{loss} for various error rates of faulty stabilizer measurements. At the current two-qubit gate infidelities and associated error rates $p_{\text{corr}} = 0.045$ and $p_{\text{single}} = 2.47 \times 10^{-4}$, the regime of beneficial loss correction, when the logical error rate falls below the physical loss rate p_{loss} , is not reachable. However, a moderate reduction of the two-qubit gate error rate by about 50%, from $p_{\text{corr}} = 0.045$ to about $p_{\text{corr}} = 0.023$, suffices to enter the regime where applying a cycle of faulty loss QEC outperforms storing information in a single physical qubit that can suffer loss.

Furthermore, Figs. 10(a) and 10(b) show the calculations of the logical error rate for the no-loss case $p_{\text{loss}} = 0$, which highlights the effects resulting from imperfections in the QND loss detection unit itself in a full QEC cycle. Here, the imperfections in the QND unit are implemented either with the coherent noise channel or an effective incoherent few-parameter Clifford noise model (details on the error models are given in the Appendix A 2). In Fig. 10(a) the logical error rate is shown as a function of the single-qubit over-rotation rate p_{single} for $p_{\text{corr}} = 0$ and it goes to zero as p_{single}^2 (black lines), as expected, representing the rate of weight-two bit-flip errors, which are uncorrectable by the distance-3 color code. In Fig. 10(b) instead the logical error rate is shown as a function of the correlated over-rotation rate p_{corr} for $p_{\text{single}} = 0$. In this case




FIG. 9. Logical error rates simulated for a loss correction cycle of the seven-qubit color code with faulty stabilizer measurements. The logical error rates are shown as a function of the loss probability p_{loss} induced by the QND-detection scheme of Fig. 8 for different error rates q in the stabilizer readout. The black line (with equation $1 - p_{\text{loss}}$) represents the error rate when no encoding is performed. The logical error rates for the ideal case with no over-rotation errors in the QND loss detection unit are shown with green up-pointing triangles. Blue circles show the logical error rates when the QND-detection unit is simulated with over-rotation parameters ($p_{\text{corr}} = 0.045$, resulting in a stabilizer measurement error rate $q = 0.045$ and $p_{\text{single}} = 2.47 \times 10^{-4}$) coming from the experimental data. Data simulated with $q = p_{\text{corr}} = 0.023$ corresponding to an improvement in the R^{MS} -gate fidelity is shown with orange down-pointing triangles. In the region with $0.03 \lesssim p_{\text{loss}} \lesssim 0.33$, error correction is beneficial in protecting the logical states with respect to storing information in an unencoded single physical qubit.

the error rate goes to zero as p_{corr}^3 (black line), representing the rate of three bit-flip errors. The bit-flip errors from the correlated over-rotations result in false-positive events, where a nonlost qubit is substituted by a fresh qubit before the stabilizer measurement. Since two (detected) losses on any two qubits are correctable, some (detected) three-loss events are not; this results in the observed p_{corr}^3 scaling of the logical error rate. This highlights and explains the different sensitivities of the logical error rate to false-positive and false-negative events where the presence of false-negative events, i.e., overlooked losses, occurs for $p_{\text{single}} \neq 0$ and constitutes the more severe source of errors.

Finally, Figs. 10(c) and 10(d) show comparisons of the logical error rate for the two scenarios, where faults in the QND loss detection unit are modeled as coherent versus incoherent errors, respectively. When $p_{\text{corr}} \neq 0$ or $p_{\text{single}} \neq 0$, the logical error rate goes to a finite value when the loss probability $p_{\text{loss}} \rightarrow 0$ as error processes involving data qubit bit flips arise and lead to a finite failure rate of the error-correction cycle. Moreover, we observe that the incoherent approximation of the coherent error channel slightly underestimates the logical error rate, by a maximum relative factor of 0.51. This behavior is not unexpected, and has also been observed in other contexts, e.g., for an incoherent approximation of coherent crosstalk errors [90]. Overall, the results therefore indicate the reliability of the incoherent approximation of the faulty QND




FIG. 10. Comparison between the coherent and incoherent implementations of the faulty QND loss detection unit. (a),(b). Logical error rates for $p_{\text{loss}} = 0$ as a function of (a) the single-qubit over-rotation rate p_{single} for $p_{\text{corr}} = 0$ and (b) the correlated over-rotation rate p_{corr} for $p_{\text{single}} = 0$ after a round of error correction of the seven-qubit color code following the scheme in Fig. 8 where the imperfections in the QND loss detection unit are implemented either with a coherent or an incoherent noise channel. (c) Logical error rate as a function of the loss probability when the faulty QND loss detection unit is modeled as a coherent channel. Panel (d) is the same as (c), but when errors in the QND loss detection are modeled as an incoherent Clifford channel.

loss detection unit in the QEC cycle. This is important as the latter incoherent model is efficiently simulatable and allows the study of faulty loss correction using stabilizer simulations of larger QEC codes.

VII. DISCUSSION AND OUTLOOK

Intermediate measurements with classical feedforward and the use of higher-dimensional quantum systems are rapidly becoming staple techniques in the toolbox of quantum information science. Beyond the obvious example of quantum error correction, the use of classical feedback to stabilize quantum systems [14,91] is an inevitable requirement for many high-precision applications. In the field of quantum computing the whole idea of measurement-based quantum computing is deeply rooted in measurements and feedforward, while quantum metrology often relies on weak or partial measurements, which must be described by quantum instruments. Similarly, in the field of quantum simulation, in-sequence measurements might be a way to use valuable quantum resources more efficiently in a hybrid quantum classical optimization setting [92,93]. What is common to all these tasks is that the measurement is nondestructive and imparts a backaction onto the postmeasurement state, which will depend on the outcome.

Faithfully characterizing the dynamics of such advanced operations will be key for the next generation of quantum devices, yet conventional methods fall short of this goal. The tools we develop here on the example of a QND measurement for qubit loss detection directly generalize to any quantum instrument, including the examples above. We find that the instrument picture captures essential features of the quantum dynamics, which in our case enable a detailed study of the effect of these instruments on quantum error correction. These results will inform progress on the correction of qubit losses and leakage errors, which represent a dominant obstacle on the path to quantum error correction above break even [51,94]. It will thus be interesting to apply these methods not only to fields where the measurement backaction has such a subtle influence, but also to fields where it is a key part of the operation, such as quantum metrology and sensing.

The presented techniques rely on tomographic reconstruction to guide the development of effective models for the studied quantum instruments. An interesting problem for future research would thus be to generalize and validate SPAM-free characterization techniques such as randomized benchmarking and gate set tomography [17] with respect to quantum instruments with low error rates.

DATA AVAILABILITY

The data underlying the findings of this work is available at [95].

ACKNOWLEDGMENTS

We gratefully acknowledge funding by the U.S. ARO Grant No. W911NF-21-1-0007. We also acknowledge funding by the Austrian Science Fund (FWF), through the SFB BeyondC (FWF Project No. F7109), by the Austrian Research Promotion Agency (FFG) Contract No. 872766, by the EU H2020-FETFLAG-2018-03 under Grant Agreement No. 820495, and by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via the U.S. ARO Grant No. W911NF-16-1-0070. This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 840450, as well as by the European Research Council (ERC) via ERC Starting Grant QNets Grant No. 804247. It reflects only the authors' views, the EU Agency is not responsible for any use that may be made of the information it contains. We acknowledge the use of computational resources from the parallel computing cluster of the Open Physics Hub at the Physics and Astronomy Department in Bologna. T.M. and R.B. acknowledge support by the IQI GmbH.

D.V. and M.M. derived the theory results. R.S., A.E., L.P., M. Meth, M.R., P.S., and T.M. performed the experiments. R.S. analyzed the data. P.S., T.M., M.M., and R.B. supervised the project. All authors contributed to writing the manuscript.

APPENDIX: CHARACTERIZING QUANTUM INSTRUMENTS: FROM NONDEMOLITION MEASUREMENTS TO QUANTUM ERROR CORRECTION

The additional information presented here aims at providing further experimental and theoretical results supporting our findings in more detail. We start off by thoroughly deriving all maps underlying our QND loss detection unit in both cases of asymmetric loss and the quantum erasure channel. This will be complemented by further experimental data, all presented in Appendix A 1. We continue in Appendix A 2 by developing a noise model giving a well-founded description to our experimental limitations. Thereafter, those noise models form the basic building blocks to numerical simulations studying the implications of the loss detection in respect of quantum error correcting codes. We conclude with Appendix A 3 by giving more detailed derivations covering the loss treatment in the seven-qubit color code.

1. Quantum instrument: QND loss detection

This section gives a more thorough introduction to QND detection, serving as our quantum instrument working example, by deriving all maps relevant to our studies. Then, additional experiments are presented addressing the demonstration of QND-detection's principal working ability complemented by results on the higher-dimensional process tomography fully characterizing its underlying maps.

As loss on our setup naturally occurs at rates similar to those of single-qubit errors, we introduce it in a controlled fashion. For instance, from the system qutrit's (q) state $|0\rangle_q$, loss can be induced by coherently transferring part of the population outside the computational subspace into $D_{5/2}(m = +1/2) = |2\rangle_q$ via the rotation

$$R_{\text{loss}}(\phi) = |1\rangle \langle 1|_q + \cos(\phi/2)(|0\rangle \langle 0|_q + |2\rangle \langle 2|_q) + \sin(\phi/2)(|0\rangle \langle 2|_q - |2\rangle \langle 0|_q). \quad (\text{A1})$$

The loss rate ϕ relates to the loss probability via $p_{\text{loss}} = \sin^2(\phi/2)$. Note that loss in general can be induced through an arbitrary state $\alpha|0\rangle_q + \beta|1\rangle_q$ with $|\alpha|^2 + |\beta|^2 = 1$ using a single coherent rotation on the system qutrit before and its inverse after the loss rotation $R_{\text{loss}}(\phi)$. To detect loss, two full entangling $R^{\text{MS}}(\phi/2) \cdot R^{\text{MS}}(\phi/2) = R^{\text{MS}}(\phi)$ couple to the ancilla and system qutrit and realize a collective bit flip only if both qubits are present in their computational subspace:

$$R^{\text{MS}}(\phi) = \exp\left(-i\frac{\phi}{2}X_aX_q\right) = (\cos(\phi/2)(\mathbb{1}_a \otimes \mathbb{1}_q - |2\rangle \langle 2|_q) - i\sin(\phi/2)X_aX_q + |2\rangle \langle 2|_q) \quad (\text{A2})$$

with

$$\mathbb{1}_a = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad \mathbb{1}_q = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad X_a = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},$$

$$X_q = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}. \quad (\text{A3})$$

On the other hand, if the system qutrit occupies a state outside the computational subspace, for instance in $|2\rangle_q$, the R^{MS} gate is subject to an identity operation, which can be seen from the argument of its exponential $X_iX_i = \mathbb{I}$ acting merely on the ancilla qubit. This follows a collective bit flip

$$R_a^X(\pi) = -i(|0\rangle \langle 1|_a + |1\rangle \langle 0|_a),$$

$$R_q^X(\pi) = |2\rangle \langle 2|_q - i(|0\rangle \langle 1|_q + |1\rangle \langle 0|_q). \quad (\text{A4})$$

Consequently, in the absence of loss the effect of the R^{MS} gate is undone, whereas under loss, the ancilla qubit gets excited by the final bit flip, signaling the event of loss. The overall unitary combining loss operation and QND detection is given by

$$U = R_a^X(\pi)R_q^X(\pi)R^{\text{MS}}(\pi)R_{\text{loss}}(\phi) \\ = \mathbb{1}_a \otimes U^{(0)} + X_a \otimes U^{(1)} \quad (\text{A5})$$

with

$$U_q^{(0)} = |1\rangle \langle 1|_q + \cos(\phi/2) |0\rangle \langle 0|_q + \sin(\phi/2) |0\rangle \langle 2|_q, \\ U_q^{(1)} = \sin(\phi/2) |2\rangle \langle 0|_q - \cos(\phi/2) |2\rangle \langle 2|_q. \quad (\text{A6})$$

Taking the additional loss state $|2\rangle_q$ on the system qubit into account and, by that, extending the view from qubit to qutrit, one ends up with a unitary process fully describing this quantum instrument. We emphasize that on the qutrit level the entire dynamics of our detection unit can be captured, which is well exploited by the experiments from Fig. 7 in the main text

However, to pick up the discussion on the nonunitary effects potentially leading to unwanted and erroneous mechanisms, we restrict our view again to the qubit level and further assume that no population is initially present in $|2\rangle_q$. Hence, the operators $U_q^{(0)}$ and $U_q^{(1)}$ reduce to

$$A_q^{(0)} = |1\rangle \langle 1|_q + \cos(\phi/2) |0\rangle \langle 0|_q, \\ A_q^{(1)} = \sin(\phi/2) |2\rangle \langle 0|_q, \quad (\text{A7})$$

leading to single-qubit processes describing the QND detection restricted to the system qubit. We can describe both maps $\{A_q^{(0)}, A_q^{(1)}\}$ by two trace nonincreasing CP maps \mathcal{E}_0 and \mathcal{E}_1 ,

$$\mathcal{E}_0: \rho \mapsto A_q^{(0)} \rho A_q^{(0)\dagger}, \\ \mathcal{E}_1: \rho \mapsto A_q^{(1)} \rho A_q^{(1)\dagger}, \quad (\text{A8})$$

acting on the system qubits as

$$\rho \mapsto |0\rangle \langle 0|_a \otimes \mathcal{E}_0(\rho) + |1\rangle \langle 1|_a \otimes \mathcal{E}_1(\rho), \quad (\text{A9})$$

where the two maps are together unitary again. It is noteworthy that the no-loss map \mathcal{E}_0 initially starting from the superposition state $1/\sqrt{2}(|0\rangle_q + |1\rangle_q)$ would be transitioning to $|1\rangle_q$ as the loss probability from $|0\rangle_q$ increases, which is subject to Fig. 4(a) in the main text. Only for very little loss, $\phi \sim 0$, the no-loss map converges to an identity operation.

Next to having loss asymmetrically with respect to either computational basis state $\{|0\rangle_q, |1\rangle_q\}$, we follow a different, often utilized, scenario called the quantum erasure channel [71]. Its circuit is depicted in Fig. 2(c) of the

main text. First, partial loss is induced from $|0\rangle_q$ followed by its detection. The protocol only continues in the absence of loss by inducing the same partial loss from the other qubit state $|1\rangle_q$ together with its detection. The second part of the map can be expressed via $\{\tilde{A}_q^{(0)}, \tilde{A}_q^{(1)}\}$, where we swap the roles of $|0\rangle_q$ and $|1\rangle_q$. Thus, the quantum erasure channel can be described using the map

$$\rho \mapsto (1 - p_L)(1 - \tilde{p}_L) \tilde{A}_q^{(0)} A_q^{(0)} \rho A_q^{(0)\dagger} \tilde{A}_q^{(0)\dagger} \\ + (1 - p_L) \tilde{p}_L \tilde{A}_q^{(1)} \rho \tilde{A}_q^{(1)\dagger} + p_L A_q^{(1)} \rho A_q^{(1)\dagger} \quad (\text{A10})$$

with probabilities p_L and \tilde{p}_L for any arbitrary input state $\alpha |0\rangle_q + \beta |1\rangle_q$ given by

$$p_L = |\alpha|^2 \sin^2(\phi/2), \quad \tilde{p}_L = \frac{|\beta|^2 \sin^2(\phi/2)}{|\alpha|^2 \cos^2(\phi/2) + |\beta|^2}. \quad (\text{A11})$$

In this case the process reduces to

$$\rho \mapsto \cos^2(\phi/2) \rho + \sin^2(\phi/2) |2\rangle \langle 2|_q, \quad (\text{A12})$$

where the effect of the loss is proportional to the arbitrary input state ρ , indicating that after normalization the initial state can be retrieved independently of the loss probability.

The basic idea of this quantum instrument is the detection of qubit loss, i.e., unwanted leakage to levels outside the computational subspace, that in a realistic scenario would be followed by its correction, representing the scope of our foregoing work [11]. To give the rather formal discussion a physical meaning, we demonstrate the unit's working principle. Partial loss induced from $|0\rangle_q$ via the loss transition $R_{\text{loss}}(\phi)$ is continuously increased and subsequently detected. Note that both qubits are read out yielding their populations in the upper $|D\rangle$ -state manifold referring to directly measured loss in the case of the system qubit and detected loss for the ancilla qubit. In Fig. 11 results are presented for individual and repeated loss detection employing up to two ancilla qubits. Slopes extracted from the linear fit in the repeated detection read 0.938(9) and 0.944(12) for ancillae a_1 and a_2 , respectively. On the individual readouts we get 0.977(2) and 0.995(2) with a resonant crosstalk to the ancilla not participating of 0.005(1) and 0.003(1), respectively. When utilizing a_2 , we end up with a higher detection efficiency because of a better performing $R_{i,j}^{\text{MS}}$ gate on the particular ion pair. Two hundred cycles are taken on this measurement. Errors correspond to one standard deviation of statistical uncertainty due to quantum projection noise.

Next, we complement the results from Fig. 4(a) in the main part revealing a pull towards the state not affected from asymmetric loss by further demonstrating that the purity $\text{Tr}(\rho^2)$ of the associated reconstructed states remains constant across the entire loss probability range;




FIG. 11. Investigating the performance of the QND-detection unit according to Fig. 2(b) in the main text. Population in the $D_{5/2}$ state for the system qubit (directly measured loss) versus transferred excitation on the ancilla qubit (detected loss) in the case of detecting loss repeatedly using both ancilla qubits a_1 and a_2 (left), solely with ancilla a_1 (middle), and ancilla a_2 (right). The imprinted detection efficiencies demonstrate reliable loss mapping onto the ancilla qubit and its readout by means of QND. Errors correspond to one standard deviation of statistical uncertainty due to quantum projection noise.

see Fig. 12(a). The purity value is found independent of the loss and therefore underlining at first glance a correct experimental outcome, whereas only in the Bloch sphere picture [Fig. 4(a)] deviations due to the nonunitary map become visible. Likewise, considerations have been done on the erasure channel, previously discussed in Fig. 4(b) and similarly producing purity values independent of loss, as can be seen in Fig. 12(b).

Next, we estimate the detection correlation of a single loss event by two repeated detections. Such system capabilities emphasize the work on the erasure channel and, more generally, become relevant in a realistic scenario demanding several consecutive readouts, especially when embedded in QEC codes. In Fig. 13 positive correlation occurs for a certain shot when both ancilla qubits agree upon a certain loss event. Furthermore, the data on the repeated readout allow us to quantify false-positive and false-negative rates, manifesting important failure modes of our detection unit. Again, false-positive rates dominate

owing to their strong sensitivity on the entangling R^{MS} gate, as was the case in Fig. 7(b) in the main text. One hundred cycles for $|0\rangle_q$ loss and 200 cycles for $1/\sqrt{2}(|0\rangle_q + |1\rangle_q)$ loss are taken on this measurement. Errors correspond to one standard deviation of statistical uncertainty due to quantum projection noise.

We switch our consideration from qubit to the qutrit level and resume the discussion on the process tomography covering both the ancilla qubit and system qutrit from Fig. 7(a) in the main text. Thereby all presented Choi operators are postselected upon the ancilla outcome denoting the qutrit maps separated by both loss cases. This has the advantage of unitary operators describing the full dynamics of the system qutrit in either loss case that moreover gives an estimation on the QND-detection's dominant failure mode, namely false-positive and false-negative rates. As discussed in detail in the main part of the paper, standard tomography restricted to the qubit level prevents us from getting such fine-grained analysis for two main




FIG. 12. Purity of a single system qubit after undergoing QND detection in the no-loss case for different loss channels. (a) After a single QND detection with loss from $|0\rangle_q$, we find purity values unaffected by the amount of loss for all of the given input states. At high loss probabilities, tomography becomes unreliable due to the low count rates. Furthermore, the superposition states show systematic drifts beyond the given statistical errors, which however do not affect the results. (b) The erasure channel is realized by consecutively inducing the same partial from $|0\rangle_q$ followed by $|1\rangle_q$ and postselecting to both ancilla $|0\rangle_a$ outcomes. The purity of the output state is again unaffected by loss for any of the probed input states. Errors correspond to one standard deviation of statistical uncertainty due to quantum projection noise.




FIG. 13. Correlations between two repeated QND detections according to Fig. 2(b) in the main text. Loss on the system qubit is induced from the imprinted states followed by two repeated detections using ancillae a_1 and a_2 . A positive correlation refers to successfully detecting the same loss event twice, whereas faulty assignments can be separated into false-positive and false-negative events; shown in the lower figure part. Errors correspond to one standard deviation of statistical uncertainty due to quantum projection noise.

reasons. First, reliably assigning false-positive and false-negative events is not possible when postselecting by the ancilla qubit's measurement outcome. Second, when tracing over the ancilla, the loss state $|2\rangle_q$ is incoherently added to the qubit state $|1\rangle_q$, creating an nonphysical bias under which tomography is likely to break, as demonstrated in Fig. 14. Here, we distinguish between tracing before and after tomography reconstruction. On the one hand, when directly tracing on the raw data and subsequently reconstructing the map, it includes coherences owing to the reconstruction technique forcing physical properties. On the other hand, when tracing after process reconstruction, coherences on $|01\rangle_q$ vanish. Both cases draw attention to potential risks on how commonly known process tomography fails to describe quantum instruments.

In the context of the numerical simulations covering implications on QEC however, we make use of the full map capturing the combined ancilla-qutrit dynamics together with the noise models; further discussed below. We present experimentally estimated ancilla-qutrit Choi operators for various loss probabilities in Fig. 15 using the elementary basis according to $\{|0000\rangle_{a,q}, \dots, |1212\rangle_{a,q}\}$. The process tomography of every loss probability required $54 \times 12 = 648$ experimental settings. For the sake of clarity, we plot ideal Choi operators (left column) and the experimental ones (right column) for various loss probabilities separated by rows side by side. Color and saturation refer to the argument and absolute value of the complex matrix entries. Process fidelities with the ideal Choi operator from top to bottom read $\{0.91(1), 0.89(1), 0.85(1)\}$, referring to the loss probabilities $\{0\%, 50\%, 85\%\}$. One hundred cycles are taken for each experiment. In the

no-loss case the expected controlled \hat{X}_a operation signaling a loss event whenever the system qubit occupies level $|2\rangle_q$ is clearly reproduced, as expressed by Eq. (A5) derived at the beginning of this section.

Finally, we present additional data on the system qutrit process tomography according to Fig. 16(a) and loss induced from $|0\rangle_q$ and $|1\rangle_q$ presented in Figs. 17 and 18, respectively. We emphasize that here, similar to the qubit level, certain coherences vanish when tracing over the ancilla, which is no longer covered by the process tomography. Still, the dynamics on the system qutrit clearly captures the population transfer from either basis state $\{|0\rangle_q, |1\rangle_q\}$ to the loss level $|2\rangle_q$. Furthermore, a change in the asymmetric behavior between loss from $|0\rangle_q$ and $|1\rangle_q$ becomes distinctly visible in the qubit subspace.

For the sake of completeness, we present similar Choi operators on the repeated loss detection in Fig. 19, consecutively mapping the same loss event to two different ancilla qubits; shown for loss from $|0\rangle_q$. As the reconstructed Choi operators follow the expected behavior previously observed, their fidelities turn out slightly lower compared to Fig. 17, as expected due to the more complex experiment.

2. Noise model on QND loss detection

Here, we study various noise models in order to find the best suitable description of the experimental limitations underlying our QND detection. Although very small contributions will be precluded by SPAM errors, the resulting models give us a rough estimate as a guide for where to look at upon which a microscopic noise model




FIG. 14. Potential risks on system qubit process reconstruction when partial tracing the ancilla qubit. Left column: partial tracing before reconstructing the map directly on the raw data, previously used in Fig. 7 of the main part. In this case, the loss level $|2\rangle_q$ is incoherently added onto state $|1\rangle_q$. Coherences present in $|01\rangle_q$ originate from the reconstruction technique, forcing physical properties. Right column: postselecting from the already reconstructed qubit-qutrit maps presented in Fig. 15. In contrast to before coherences on $|01\rangle_q$ vanish, whereas the nonphysical bias remains.

for the numerical simulations can be developed. Getting more insights on these error mechanisms is essential when observing implications of the quantum instrument in the context of QEC protocols and is further an essential building block towards fault-tolerant quantum computation; see Sec. VI of the main text.

We refer to Eq. (A5) from above and express the action of the ideal QND map U under a given loss rate ϕ acting on the combined ancilla-qutrit system in terms of the Choi operator $\rho^{\text{CJ}} = \mathbb{1} \otimes U \cdot |\Phi_+\rangle \langle \Phi_+| \cdot \mathbb{1} \otimes U^\dagger$, where $|\Phi_+\rangle \langle \Phi_+|$ is the maximally entangled state of two copies of the ancilla-qutrit system. An erroneous channel $\mathcal{E}_{\text{noise}}$ transforms the Choi operator ρ^{CJ} to $\rho_{\text{noise}}^{\text{CJ}} = (\mathbb{1} \otimes \mathcal{E}_{\text{noise}})(\rho^{\text{CJ}})$. Noise rates entering $\mathcal{E}_{\text{noise}}$ for given model parameters are then extracted by minimizing the distance between modeled noisy Choi operators $\rho_{\text{noise}}^{\text{CJ}}$ and the experimentally determined ones $\rho_{\text{exp}}^{\text{CJ}}$ from Fig. 15. As a measure for the distance in the cost function, we minimize the infidelity:

$$\|1 - \mathcal{F}(\rho_{\text{exp}}^{\text{CJ}}, \rho_{\text{noise}}^{\text{CJ}})\|. \quad (\text{A13})$$

Our initial considerations covered the study of the QND-detection's failure modes, i.e., false-positive and

false-negative rates both quantified in the main part of the paper. Measuring process tomography however comes with overhead in the form of preparation and measurement gates followed by two consecutive detections at the end of each experiment required for reading out the qutrit's state. Therefore, SPAM errors are not to be neglected and lead to a significant bias on false-positive and false-negative rates. With this in mind, we put the failure modes aside and focus on experimental limitations instead. In the following, we consider as models for $\mathcal{E}_{\text{noise}}$ a depolarizing channel, a dephasing channel, and coherent two- and single-qubit over-rotations.

Depolarizing and dephasing channels.—We start off by testing the agnostic models, namely depolarizing and dephasing channels as those represent error mechanisms typically considered in the field of quantum computation [72]. The effect of the latter can be understood by losing phase information between the quantum states involved. Coherences get lost and an arbitrary single-qubit state in the Bloch sphere picture would finally shrink onto the Z axis as no phase information is left. Depolarizing noise can be considered as simultaneous dephasing in the X , Y , and Z bases, eventually leading to a complete mixed state that, for a single qubit, can be illustrated by shrinking




FIG. 15. Combined process reconstruction on the ancilla-qutrit system according to Fig. 16(b). The resulting Choi operators (right column) denoted in the elementary basis ($\{|0000\rangle_{a,q}, \dots, |1212\rangle_{a,q}\}$) describe the whole dynamics of the QND-detection unit under loss from $|0\rangle_q$. Hue relates to phase according to the top right color bar and saturation to the absolute value of the complex entries. Process fidelities with the ideal Choi operators (left column) from top to bottom read $\{0.91(1), 0.89(1), 0.85(1)\}$. Errors correspond to one standard deviation of statistical uncertainty due to quantum projection noise.




FIG. 16. Schematics on higher-dimensional process tomography. (a) Qutrit process tomography solely covering the system qubit (q) together with the loss level $\{|0\rangle_q, |1\rangle_q, |2\rangle_q\}$ undergoing the QND-detection unit by using nine preparation settings together with six measurement settings, resulting in 54 experiments each run. (b) Combined process tomography on ancilla (a) and qutrit (1), capturing the entire dynamics of this quantum instrument using 12 settings on the ancilla qubit (four preparation settings and three measurement settings) alongside 54 settings on the system qutrit, resulting in 648 experiments. (c) Qutrit process tomography on the erasure channel, focusing on the no-loss case, i.e., twice postselecting the ancilla qubit's $|0\rangle_a$ outcome.

the Bloch sphere towards its center. Note that we implement those models such that they act both on the ancilla and the qutrit using only a single noise parameter [96]. The upper row of Fig. 20 depicts the fidelities (top part) for the individual models at the optimized parameters (bottom part). Both results indicate similar improvements compared to the fidelity with the ideal QND map from Eq. (A5). Numbers on fidelities and optimized parameters for depolarizing noise $p_{\text{depol.}}$ and dephasing noise $p_{\text{deph.}}$ are further summarized in Table I. The parameters typically lie around 1% or below, yet the small increase in fidelity indicates other error mechanisms to be more dominant.

Correlated two-qubit over-rotations.—The erroneous peaks in the experimentally estimated Choi operators from Fig. 15 imply that additional rotations should be taken into account by the agnostic models. Those dominant error peaks are found originating from correlated rotations between the ancilla and system qubit, as illustratively labeled in Fig. 21(a). Note that the error terms are restricted to the qubit level and partial coherences are still present. Hence, if the system qutrit's state is $|2\rangle_q$, no correlated error is induced on the ancilla qubit. Therefore, correlated

errors are due to faulty entangling R^{MS} gates. A potential noise model covering correlated rotations in such a way reads

$$\rho \mapsto \mathcal{E}_{\text{noise}}(\rho) = U_{\text{corr}} \rho U_{\text{corr}}^\dagger \quad (\text{A14})$$

with

$$U_{\text{corr}} = \cos \frac{\alpha}{2} \mathbb{1}_a \otimes \mathbb{1}_q + i \sin \frac{\alpha}{2} (X_a \otimes X_q + \mathbb{1}_a \otimes |2\rangle \langle 2|_q), \quad (\text{A15})$$

where α describes the correlated under- and over-rotations and relates to the corresponding error probability via $p_{\text{corr.}} = \sin(\alpha/2)^2$. For comparison, a value $p_{\text{corr.}}$ of 0.5 would induce a maximally entangling two-qubit operation on the ancilla and system qubit. We first test the model alone followed by combining it with depolarizing and dephasing noise. The resulting fidelities (top part) at the optimized model parameters (bottom part) are shown in the second row of Fig. 20 and clearly overcome those on the agnostic models denoting correlated rotations to




FIG. 17. Qutrit process tomography characterizing the QND-detection unit for loss from $|0\rangle_q$ according to Fig. 16(a). (a) System qutrit's Choi operator in elementary basis $\{|00\rangle_q, \dots, |22\rangle_q\}$ after tracing over the ancilla qubit and various loss probabilities denoting the effect of the loss transition transferring population from $|0\rangle_q$ to $|2\rangle_q$. (b) The respective fidelities with the ideal Choi operators covering the complete loss range.

be our leading noise mechanism. The effect of the additional depolarizing and dephasing noise (bottom right) leads to slight improvements. The modeled Choi operator on this combined noise model is plotted for the no-loss case in Fig. 21(b), showing strong similarities to the experimental one and underlining good agreement between the model and experiment. Numbers on fidelities and optimized parameters for all models are summarized in Table I.

Correlated and single over-rotations.—Finally, we combine the action of correlated rotations with single-qubit rotations on the ancilla and the qutrit and we consider the coherent error model given by

$$\rho \mapsto \mathcal{E}_{\text{noise}}(\rho) = R U_{\text{corr}} \rho U_{\text{corr}}^\dagger R^\dagger, \quad (\text{A16})$$

where ρ is the state obtained after the application of the loss operation U of Eq. (A5) [see also Fig. 8(a) of the main

text], U_{corr} is a correlated two-qubit over-rotation defined in Eq. (A15), and $R = R_a^X(\beta)R_q^X(\beta)$ with

$$R_a^X(\beta) = \cos(\beta/2)\mathbb{1}_a - i \sin(\beta/2)X_a, \quad (\text{A17})$$

$$R_q^X(\beta) = \cos(\beta/2)(\mathbb{1}_q - |2\rangle\langle 2|_q) - i \sin(\beta/2)X_q + |2\rangle\langle 2|_q \quad (\text{A18})$$

over-rotations with angle β of the ancilla and the qutrit system that corresponds to the single-qubit flip error rate $p_{\text{single}} = \sin(\beta/2)^2$. After measurement of the ancilla, the quantum process arising from the erroneous channel in Eq. (A16) can be written as

$$\rho \mapsto |0\rangle\langle 0|_a \otimes \mathcal{R}_0(\rho_q) + |1\rangle\langle 1|_a \otimes \mathcal{R}_1(\rho_q), \quad (\text{A19})$$

where ρ_q is the state related to the qutrit only and the processes \mathcal{R}_0 and \mathcal{R}_1 describe the maps that transform the




FIG. 18. Qutrit process tomography characterizing the QND-detection unit for loss from $|1\rangle_q$ according to Fig. 16(a). (a) System qutrit's Choi operator in elementary basis $\{|00\rangle_q, \dots, |22\rangle_q\}$ after tracing over the ancilla qubit and several different loss probabilities denoting the effect of the loss transition transferring population from $|1\rangle_q$ to $|2\rangle_q$. (b) The respective fidelities compared to the ideal Choi operators covering the complete loss range.

qutrit state in the case of no-loss detected (ancilla qubit in $|0\rangle_a$) and of loss detected (ancilla qubit in $|1\rangle_a$). The Choi operators Λ_0 and Λ_1 of maps \mathcal{R}_0 and \mathcal{R}_1 can be computed

for all values of the over-rotated angles α and β . In particular, if we consider small deviations for α and β , Λ_0 and Λ_1 read at second order

$$\Lambda_0 = \begin{pmatrix} 1 - \frac{\alpha^2}{4} - \frac{\beta^2}{2} & (-\frac{\alpha}{4} - \frac{i}{2})\beta & 0 & (-\frac{\alpha}{4} - \frac{i}{2})\beta & 1 - \frac{\alpha^2}{4} - \frac{\beta^2}{2} & 0 & 0 & 0 & \frac{i}{2}\beta \\ (-\frac{\alpha}{4} + \frac{i}{2})\beta & \frac{\beta^2}{4} & 0 & \frac{\beta^2}{4} & (-\frac{\alpha}{4} + \frac{i}{2})\beta & 0 & 0 & 0 & -\frac{1}{4}\beta^2 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ (-\frac{\alpha}{4} + \frac{i}{2})\beta & \frac{\beta^2}{4} & 0 & \frac{\beta^2}{4} & (-\frac{\alpha}{4} + \frac{i}{2})\beta & 0 & 0 & 0 & -\frac{1}{4}\beta^2 \\ 1 - \frac{\alpha^2}{4} - \frac{\beta^2}{2} & (-\frac{\alpha}{4} - \frac{i}{2})\beta & 0 & (-\frac{\alpha}{4} - \frac{i}{2})\beta & 1 - \frac{\alpha^2}{4} - \frac{\beta^2}{2} & 0 & 0 & 0 & \frac{i}{2}\beta \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -\frac{i}{2}\beta & -\frac{1}{4}\beta^2 & 0 & -\frac{1}{4}\beta^2 & -\frac{i}{2}\beta & 0 & 0 & 0 & \frac{\beta^2}{4} \end{pmatrix}, \quad (\text{A20})$$

$$\Lambda_1 = \begin{pmatrix} \frac{\beta^2}{4} & \frac{\alpha\beta}{4} & 0 & \frac{\alpha\beta}{4} & \frac{\beta^2}{4} & 0 & 0 & 0 & (\frac{\alpha}{4} - \frac{i}{2})\beta \\ \frac{\alpha\beta}{4} & \frac{\alpha^2}{4} & 0 & \frac{\alpha^2}{4} & \frac{\alpha\beta}{4} & 0 & 0 & 0 & \frac{\beta^2}{4} - \frac{i\alpha}{2} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{\alpha\beta}{4} & \frac{\alpha^2}{4} & 0 & \frac{\alpha^2}{4} & \frac{\alpha\beta}{4} & 0 & 0 & 0 & \frac{\beta^2}{4} - \frac{i\alpha}{2} \\ \frac{\beta^2}{4} & \frac{\alpha\beta}{4} & 0 & \frac{\alpha\beta}{4} & \frac{\beta^2}{4} & 0 & 0 & 0 & (\frac{\alpha}{4} - \frac{i}{2})\beta \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ (\frac{\alpha}{4} + \frac{i}{2})\beta & \frac{\beta^2}{4} + \frac{i\alpha}{2} & 0 & \frac{\beta^2}{4} + \frac{i\alpha}{2} & (\frac{\alpha}{4} + \frac{i}{2})\beta & 0 & 0 & 0 & 1 - \frac{\beta^2}{4} \end{pmatrix}, \quad (\text{A21})$$

where we have labeled the qutrit basis states in the order $|00\rangle, |01\rangle, |02\rangle, \dots, |22\rangle$. In the next section we discuss how to approximate the channel in Eq. (A19) with Clifford gates.




FIG. 19. Qutrit process tomography on two repeated QND detections for loss from $|0\rangle_q$ according to Fig. 16(a). (a) System qutrit Choi operators mapping loss repeatedly onto ancillae a_1 and a_2 under several different loss probabilities. The processes for which we trace over both ancillae prior to reconstruction denote the loss transition transferring population from $|0\rangle_q$ to $|2\rangle_q$. (b) Fidelities compared to the ideal operators remain approximately constant along all measured loss probabilities and show slightly decreased values compared to the results on single QND detection from Fig. 17.




FIG. 20. Noise-model QND detection. Various noise model describing the experimental limitations on the ancilla-qutrit Choi operator depicted in Fig. 15. The limitations are best described when combining correlated coherent rotations together with depolarizing and dephasing noise. Correlated errors clearly dominate as depolarizing and dephasing errors only lead to minor improvements. The error parameters on the bottom of each plot refer to depolarizing error $p_{\text{depol.}}$, dephasing error $p_{\text{deph.}}$, and correlated error $p_{\text{corr.}}$, the latter according to Eq. (A15). Lines connect the points for clarity.

Effective Clifford channel.—Before deriving the analytical expression for the Clifford channel, the forms of Choi operators Λ_0 and Λ_1 allow us to have a qualitative discussion on the events that will form the Clifford channel approximating Eq. (A19). In Λ_0 and Λ_1 we can easily identify the following events happening to the ancilla-qutrit system. If the ancilla is in $|0\rangle_a$, the qutrit state is left unchanged with probability $1 - \alpha^2/4 - \beta^2/2$ or it undergoes an X_q bit-flip error with probability $\beta^2/4$. When the ancilla is instead in $|1\rangle_a$, the qutrit state is left unchanged in the loss state $|2\rangle \langle 2|_q$ with probability $1 - \beta^2/4$.

We can also identify the origin of the false-negative and false-positive events. From Λ_0 we see that the qutrit will

be projected on the loss state $|2\rangle \langle 2|_q$ with probability $\beta^2/4$, while the ancilla will be in the no-loss detected state $|0\rangle_a$. This corresponds to a false-negative event whose origin can be traced back to the single-qubit over-rotation R of Eq. (A16).

From Λ_1 we see that, when the qutrit is generated in the computational space by $|0\rangle_q$ and $|1\rangle_q$, the ancilla will be found in the loss detected state $|1\rangle_a$. In particular, the qutrit will be left unchanged with probability $\beta^2/4$ and it will undergo an X_q bit-flip error with probability $\alpha^2/4$. These events correspond to false-positive events whose origin can be traced back to the single-qubit over-rotation R and to the correlated over-rotation U_{corr} of Eq. (A16).

TABLE I. Summary on noise model parameters and results. The parameters and fidelities refer to the best suitable model values describing the experimental noise from Fig. 20: depolarizing error $p_{\text{depol.}}$, dephasing error $p_{\text{deph.}}$, and correlated error $p_{\text{corr.}}$ according to Eq. (A15).

Loss (%)	Noise Model Parameters and Results										
	F_{ideal}	Depolarizing		Dephasing		Correlated		Correlated, depol., and deph.			
		$\mathcal{F}_{\text{model}}$	$p_{\text{depol.}}$	$\mathcal{F}_{\text{model}}$	$p_{\text{deph.}}$	$\mathcal{F}_{\text{model}}$	$p_{\text{corr.}}$	$\mathcal{F}_{\text{model}}$	$p_{\text{corr.}}$	$p_{\text{depol.}}$	$p_{\text{deph.}}$
0	0.906	0.912	0.004	0.915	0.013	0.939	0.045	0.948	0.042	0.012	0.002
2	0.899	0.912	0.009	0.908	0.014	0.916	0.021	0.923	0.045	0.010	0.010
50	0.894	0.905	0.007	0.905	0.018	0.924	0.032	0.934	0.037	0.020	0.007
85	0.854	0.861	0.005	0.863	0.017	0.897	0.054	0.903	0.045	0.010	0.010




FIG. 21. Comparison between experimental and noisy-modeled Choi operators. (a) Experimentally estimated map according to Fig. 15(a) with additionally marked transitions denoting correlated errors describing our leading error mechanism; see Eq. (A15). (b) Most suitable noisy-modeled Choi operator combining correlated, depolarizing, and dephasing errors.

The previous considerations on the Choi operators Λ_0 and Λ_1 can be made more precise by explicitly computing the process in Eq. (A19) with the help of Eqs. (A15), (A17), and (A18) and by retaining only the terms that can be written in the Kraus form $P\rho P^\dagger$, where P is a Pauli operator. In this way, we can approximate the channel in Eq. (A19) as

$$\begin{aligned} \rho \mapsto & p_a P_{01} \rho P_{01}^\dagger + p_b X_q \rho X_q + p_c X_a X_q \rho X_q X_a \\ & + p_d X_a P_{01} \rho P_{01}^\dagger X_a + q_a X_a |2\rangle \langle 2| \rho |2\rangle \langle 2| X_a \\ & + q_b |2\rangle \langle 2| \rho |2\rangle \langle 2|, \end{aligned} \quad (\text{A22})$$

where ρ is the density matrix of the whole ancilla and qutrit system, $P_{01} = 1 - |2\rangle \langle 2|_q$ is the projector on the computational space $\{|0\rangle_q, |1\rangle_q\}$ of the qutrit, and the probabilities take the form

$$p_a = \sin^2 \alpha \sin^4 \beta + \cos^2 \alpha \cos^4 \beta \sim 1 - \alpha^2/4 - \beta^2/2, \quad (\text{A23})$$

$$p_b = \sin^2 \beta/4 \sim \beta^2/4, \quad (\text{A24})$$

$$p_c = \sin^2 \alpha \cos^4 \beta + \cos^2 \alpha \sin^4 \beta \sim \alpha^2/4, \quad (\text{A25})$$

$$p_d = \sin^2 \beta/4 \sim \beta^2/4, \quad (\text{A26})$$

$$q_a = \cos^2(\beta/2) \sim 1 - \beta^2/4, \quad (\text{A27})$$

$$q_b = \sin^2(\beta/2) \sim \beta^2/4. \quad (\text{A28})$$

The channel in Eq. (A22) can then be implemented in the following way.

1. If the ancilla is in $|0\rangle_a$, we
 - (a) leave the qutrit state in the computational space with probability $1 - \beta^2/2 - \alpha^2/4$;

- (b) apply an X_q bit-flip error to the qutrit with probability $\beta^2/4$;
- (c) apply an X_q bit-flip error to the qutrit and an X_a bit-flip error to the ancilla with probability $\alpha^2/4$ (corresponding to a false-positive event from the correlated over-rotation);
- (d) leave the qutrit state as it is and flip the ancilla with probability $\beta^2/4$ (corresponding to a false-positive event from the single rotations).

2. If the ancilla is in $|1\rangle_a$, we

- (a) leave the qutrit state in the loss state $|2\rangle_q$ with probability $1 - \beta^2/4$;
- (b) flip the ancilla to the no-loss detection state $|0\rangle_a$ with probability $\beta^2/4$ (corresponding to a false-negative from the single rotations).

The comparison between the coherent channel in Eq. (A16) and the effective Clifford channel previously described is shown in Fig. 22 and in Fig. 10 of the main text.

3. Losses in the seven-qubit code

In this section, we discuss the correction from losses for the seven-qubit color code, in the ideal scenario of perfect QND loss detection and stabilizer measurements. We also assume that losses occur on each qubit independently with loss probability p .

A loss event is correctable if the density matrix of the losses is fully mixed or, more generally, it does not contain any information on the encoded logical state. With this criterion, we can then check the loss events that can be corrected. Obviously, the event [happening with probability




FIG. 22. Comparison between the coherent and incoherent implementations of the faulty QND loss detection unit in the case of no losses. (a) Logical error rate as a function of the correlated over-rotation rate p_1 for the parameter $p_2 = 0.045$ obtained from the experimental data. (b) Logical error rate as a function of the correlated over-rotation rate p_2 for the parameter $p_1 = 2.47 \times 10^{-4}$ obtained from the experimental model.

$P_0 = (1 - p)^7$ where no loss occurs is correctable. The events where one loss occurs are also correctable. To show this, let us consider, for instance, the encoded $|0_L\rangle$ state

$$|0_L\rangle \sim (\mathbb{1} + S_x^{(1)})(\mathbb{1} + S_x^{(2)})(\mathbb{1} + S_x^{(3)})|0\rangle^{\otimes 7}, \quad (\text{A29})$$

where the $S_x^{(j)}$ are the stabilizer generators, and let us suppose that the loss affects qubit q_1 [see Fig. 8(a) of the main text]. By introducing the two orthogonal states $|\chi_0\rangle = P_x^{(2)}P_x^{(3)}|0\rangle^{\otimes 6}$ and $|\chi_1\rangle = X_2X_3X_4P_x^{(2)}P_x^{(3)}|0\rangle^{\otimes 6}$ (where $P_x^{(j)} = \mathbb{1} + S_x^{(j)}$ with $j = 2, 3$ are chosen because the loss does not belong to $S_x^{(j)}$), the state $|0_L\rangle$ can be written explicitly as

$$|0_L\rangle \sim |0_1\rangle|\chi_0\rangle + |1_1\rangle|\chi_1\rangle. \quad (\text{A30})$$

As $|\chi_0\rangle$ and $|\chi_1\rangle$ are orthogonal, the reduced density matrix of the loss q_1 obtained by tracing out the six other qubits will be $\rho_1 \sim |0_1\rangle\langle 0_1| + |1_1\rangle\langle 1_1|$, i.e., it will be fully mixed. Therefore, the events with one loss [happening with probability $P_1 = 7p(1 - p)^6$] can be correctable. A similar reasoning applies to all the events where two losses happen [$P_2 = 21p^2(1 - p)^5$] and to the events where three losses that do not form a logical operator happen as well. The events with three losses that form a logical operator are instead not correctable. There are precisely seven such events [corresponding to the logical operators $\mathcal{L} = \{[1, 2, 5], [1, 3, 6], [1, 4, 7], [2, 3, 7], [4, 3, 5], [5, 6, 7], [2, 4, 6]\}$ in Fig. 8(a) of the main text]. The last one ($[2, 4, 6]$) is given by the product of the logical operator acting on all the seven qubits multiplied by all three stabilizer generators. This implies that the probability to successfully recover the logical state is $P_3 = \left[\binom{7}{3} - 7\right]p^3(1 - p)^4 = 28p^3(1 - p)^4$. In the case of four losses, in seven cases out of $\binom{7}{4} = 35$, the reduced density matrix of the losses does not depend on the encoded logical state. These cases correspond to the losses happening on the qubits of the stabilizer

generators and their products and are given by

$$\mathcal{S} = \{[1, 2, 3, 4], [2, 3, 5, 6], [3, 4, 6, 7], [1, 4, 5, 6], [1, 2, 6, 7], [2, 4, 5, 7], [1, 3, 5, 7]\}. \quad (\text{A31})$$

This can be shown by considering, for instance, four losses happening on the stabilizer $[1, 2, 3, 4]$. A bit of algebra shows that the logical states $|0_L\rangle$ and $|1_L\rangle$ can be written as

$$|0_L\rangle = |G\rangle|000\rangle + X_2X_3|G\rangle|110\rangle + X_3X_4|G\rangle|011\rangle + X_2X_4|G\rangle|101\rangle, \quad (\text{A32})$$

$$|1_L\rangle = |G\rangle|111\rangle + X_2X_3|G\rangle|001\rangle + X_3X_4|G\rangle|100\rangle + X_2X_4|G\rangle|010\rangle, \quad (\text{A33})$$

where $|G\rangle = |0000\rangle + |1111\rangle$ is a GHZ state of qubits 1, 2, 3, 4 where the losses happen. Tracing on qubits 5, 6, 7 transforms any logical state $|\psi_L\rangle = c_0|0_L\rangle + c_1|1_L\rangle$ into a mixture with equal probabilities of the four states $\{|G\rangle, X_2X_3|G\rangle, X_3X_4|G\rangle, X_2X_4|G\rangle\}$ that is independent on the coefficients c_0 and c_1 . Finally, no event with five, six, or seven losses can be corrected. The total probability of a successful correction is given by the sum of all probabilities P_j and reads

$$p_{\text{success}} = (1 - p)^7 + 7p(1 - p)^6 + 21p^2(1 - p)^5 + 28p^3(1 - p)^4 + 7p^4(1 - p)^3 = 1 - 7p^3 + 21p^5 - 21p^6 + 6p^7. \quad (\text{A34})$$

-
- [1] D. Gottesman, Theory of fault-tolerant quantum computation, *Phys. Rev. A* **57**, 127 (1998).
 [2] P. Schindler, J. T. Barreiro, T. Monz, V. Nebendahl, D. Nigg, M. Chwalla, M. Hennrich, and R. Blatt, Experimental repetitive quantum error correction, *Science* **332**, 1059 (2011).

- [3] L. Sun, A. Petrenko, Z. Leghtas, B. Vlastakis, G. Kirchmair, K. M. Sliwa, A. Narla, M. Hatridge, S. Shankar, J. Blumoff, L. Frunzio, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, Tracking photon jumps with repeated quantum non-demolition parity measurements, *Nature* **511**, 444 (2014).
- [4] J. Kelly *et al.*, State preservation by repetitive error detection in a superconducting quantum circuit, *Nature* **519**, 66 (2015).
- [5] J. Cramer, N. Kalb, M. A. Rol, B. Hensen, M. S. Blok, M. Markham, D. J. Twitchen, R. Hanson, and T. H. Taminiau, Repeated quantum error correction on a continuously encoded qubit by real-time feedback, *Nat. Commun.* **7**, 11526 (2016).
- [6] T. Uden, P. Balasubramanian, D. Louzon, Y. Vinkler, M. B. Plenio, M. Markham, D. Twitchen, A. Stacey, I. Lovchinsky, A. O. Sushkov, M. D. Lukin, A. Retzker, B. Naydenov, L. P. McGuinness, and F. Jelezko, Quantum Metrology Enhanced by Repetitive Quantum Error Correction, *Phys. Rev. Lett.* **116**, 230502 (2016).
- [7] V. Negnevitsky, M. Marinelli, K. K. Mehta, H.-Y. Lo, C. Flühmann, and J. P. Home, Repeated multi-qubit readout and feedback with a mixed-species trapped-ion register, *Nature* **563**, 527 (2018).
- [8] T. Monz, D. Nigg, E. A. Martinez, M. F. Brandl, P. Schindler, R. Rines, S. X. Wang, I. L. Chuang, and R. Blatt, Realization of a scalable Shor algorithm, *Science* **351**, 1068 (2016).
- [9] A. Y. Kitaev, Quantum measurements and the Abelian stabilizer problem, *Electron. Colloquium Comput. Complex.* **3** (1996).
- [10] R. Jozsa, *Quantum Information Processing* (2005), Vol. 199, p. 137.
- [11] R. Stricker, D. Vodola, A. Erhard, L. Postler, M. Meth, M. Ringbauer, P. Schindler, T. Monz, M. Müller, and R. Blatt, Experimental deterministic correction of qubit loss, *Nature* **585**, 207 (2020).
- [12] C. Unnikrishnan, Quantum non-demolition measurements: Concepts, theory and practice, *Curr. Sci.* **109**, 2052 (2015).
- [13] D. B. Hume, T. Rosenband, and D. J. Wineland, High-Fidelity Adaptive Qubit Detection through Repetitive Quantum Nondemolition Measurements, *Phys. Rev. Lett.* **99**, 120502 (2007).
- [14] C. Sayrin, I. Dotsenko, X. Zhou, B. Peaudecerf, T. Rybarczyk, S. Gleyzes, P. Rouchon, M. Mirrahimi, H. Amini, M. Brune, J.-M. Raimond, and S. Haroche, Real-time quantum feedback prepares and stabilizes photon number states, *Nature* **477**, 73 (2011).
- [15] M. Hatridge, S. Shankar, M. Mirrahimi, F. Schackert, K. Geerlings, T. Brecht, K. M. Sliwa, B. Abdo, L. Frunzio, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret, Quantum back-action of an individual variable-strength measurement, *Science* **339**, 178 (2013).
- [16] M. S. Blok, C. Bonato, M. L. Markham, D. J. Twitchen, V. V. Dobrovitski, and R. Hanson, Manipulating a qubit through the backaction of sequential partial measurements and real-time feedback, *Nat. Phys.* **10**, 189 (2014).
- [17] K. Rudinger, G. J. Ribeill, L. C. G. Govia, M. Ware, E. Nielsen, K. Young, T. A. Ohki, R. Blume-Kohout, and T. Proctor, Characterizing mid-circuit measurements on a superconducting qubit using gate set tomography, *Phys. Rev. Appl.* **17**, 014014 (2022).
- [18] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, 2006).
- [19] M. Ringbauer, C. J. Wood, K. Modi, A. Gilchrist, A. G. White, and A. Fedrizzi, Characterizing Quantum Dynamics with Initial System-Environment Correlations, *Phys. Rev. Lett.* **114**, 090402 (2015).
- [20] I. Rotter and J. P. Bird, A review of progress in the physics of open quantum systems: Theory and experiment, *Rep. Prog. Phys.* **78**, 114001 (2015).
- [21] J. T. Barreiro, M. Müller, P. Schindler, D. Nigg, T. Monz, M. Chwalla, M. Hennrich, C. F. Roos, P. Zoller, and R. Blatt, An open-system quantum simulator with trapped ions, *Nature* **470**, 486 (2011).
- [22] P. Schindler, M. Müller, D. Nigg, J. T. Barreiro, E. A. Martinez, M. Hennrich, T. Monz, S. Diehl, P. Zoller, and R. Blatt, Quantum simulation of dynamical maps with trapped ions, *Nat. Phys.* **9**, 361 (2013).
- [23] M. Müller, K. Hammerer, Y. L. Zhou, C. F. Roos, and P. Zoller, Simulating open quantum systems: From many-body interactions to stabilizer pumping, *New J. Phys.* **13**, 085007 (2011).
- [24] L. M. Sieberer, M. Buchhold, and S. Diehl, Keldysh field theory for driven open quantum systems, *Rep. Prog. Phys.* **79**, 096001 (2016).
- [25] E. B. Davies and J. T. Lewis, An operational approach to quantum probability, *Commun. Math. Phys.* **17**, 239 (1970).
- [26] M. Ozawa, Quantum measuring processes of continuous observables, *J. Math. Phys.* **25**, 79 (1984).
- [27] J. Dressel and A. N. Jordan, Quantum instruments as a foundation for both states and observables, *Phys. Rev. A* **88**, 022107 (2013).
- [28] G. Chiribella, G. M. D'Ariano, and P. Perinotti, Theoretical framework for quantum networks, *Phys. Rev. A* **80**, 022339 (2009).
- [29] O. Oreshkov, F. Costa, and Č. Brukner, Quantum correlations with no causal order, *Nat. Commun.* **3**, 1092 (2012).
- [30] F. Buscemi and M. F. Sacchi, Information-disturbance trade-off in quantum-state discrimination, *Phys. Rev. A* **74**, 052320 (2006).
- [31] L. Knips, J. Dziewior, A.-L. K. Hashagen, J. D. A. Meinecke, H. Weinfurter, and M. M. Wolf, Measurement-disturbance tradeoff outperforming optimal cloning, *arXiv:preprint:1808.07882* (2018).
- [32] S. Lloyd and J.-J. E. Slotine, Quantum feedback with weak measurements, *Phys. Rev. A* **62**, 012307 (2000).
- [33] Q. Sun, M. Al-Amri, L. Davidovich, and M. S. Zubairy, Reversing entanglement change by a weak measurement, *Phys. Rev. A* **82**, 052323 (2010).
- [34] Y.-S. Kim, J.-C. Lee, O. Kwon, and Y.-H. Kim, Protecting entanglement from decoherence using weak measurement and quantum measurement reversal, *Nat. Phys.* **8**, 117 (2012).
- [35] S. Wagner, J.-D. Bancal, N. Sangouard, and P. Sekatski, Device-independent characterization of quantum instruments, *Quantum* **4**, 243 (2020).
- [36] N. Miklin, J. J. Borkala, and M. Pawłowski, Semi-device-independent self-testing of unsharp measurements, *Phys. Rev. Res.* **2**, 033014 (2020).

- [37] K. Mohan, A. Tavakoli, and N. Brunner, Sequential random access codes and self-testing of quantum measurement instruments, *New J. Phys.* **21**, 083034 (2019).
- [38] E. S. Gómez, S. Gómez, P. González, G. Cañas, J. F. Barra, A. Delgado, G. B. Xavier, A. Cabello, M. Kleinmann, T. Vértesi, and G. Lima, Device-Independent Certification of a Nonprojective Qubit Measurement, *Phys. Rev. Lett.* **117**, 260401 (2016).
- [39] M. Smania, P. Mironowicz, M. Nawareg, M. Pawłowski, A. Cabello, and M. Bourennane, Experimental certification of an informationally complete quantum measurement in a device-independent protocol, *Optica* **7**, 123 (2020).
- [40] J. Helsen, I. Roth, E. Onorati, A. H. Werner, and J. Eisert, A general framework for randomized benchmarking, [arXiv:2010.07974](https://arxiv.org/abs/2010.07974) (2020).
- [41] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, Randomized benchmarking of quantum gates, *Phys. Rev. A* **77**, 012307 (2008).
- [42] J. P. Gaebler, T. R. Tan, Y. Lin, Y. Wan, R. Bowler, A. C. Keith, S. Glancy, K. Coakley, E. Knill, D. Leibfried, and D. J. Wineland, High-Fidelity Universal Gate Set for ${}^9\text{Be}^+$ Ion Qubits, *Phys. Rev. Lett.* **117**, 060505 (2016).
- [43] R. Blume-Kohout, J. K. Gamble, E. Nielsen, K. Rudinger, J. Mizrahi, K. Fortier, and P. Maunz, Demonstration of qubit operations below a rigorous fault tolerance threshold with gate set tomography, *Nat. Commun.* **8**, 14485 (2017).
- [44] A. Erhard, J. J. Wallman, L. Postler, M. Meth, R. Stricker, E. A. Martinez, P. Schindler, T. Monz, J. Emerson, and R. Blatt, Characterizing large-scale quantum computers via cycle benchmarking, *Nat. Commun.* **10**, 5347 (2019).
- [45] D. C. McKay, S. Sheldon, J. A. Smolin, J. M. Chow, and J. M. Gambetta, Three-Qubit Randomized Benchmarking, *Phys. Rev. Lett.* **122**, 200502 (2019).
- [46] A. C. Hughes, V. M. Schäfer, K. Thirumalai, D. P. Nadlinger, S. R. Woodrow, D. M. Lucas, and C. J. Ballance, Benchmarking a High-Fidelity Mixed-Species Entangling Gate, *Phys. Rev. Lett.* **125**, 080504 (2020).
- [47] A. G. Fowler, Coping with qubit leakage in topological codes, *Phys. Rev. A* **88**, 042308 (2013).
- [48] J. Ghosh, A. G. Fowler, J. M. Martinis, and M. R. Geller, Understanding the effects of leakage in superconducting quantum-error-detection circuits, *Phys. Rev. A* **88**, 062329 (2013).
- [49] T. M. Stace, S. D. Barrett, and A. C. Doherty, Thresholds for Topological Codes in the Presence of Loss, *Phys. Rev. Lett.* **102**, 200501 (2009).
- [50] Y. Wu, S. Kolkowitz, S. Puri, and J. D. Thompson, Erasure conversion for fault-tolerant quantum computing in alkaline earth Rydberg atom arrays, (2022).
- [51] C. Ryan-Anderson, J. G. Bohnet, K. Lee, D. Gresh, A. Hankin, J. P. Gaebler, D. Francois, A. Chernoguzov, D. Lucchetti, N. C. Brown, T. M. Gatterman, S. K. Halit, K. Gilmore, J. A. Gerber, B. Neyenhuis, D. Hayes, and R. P. Stutz, Realization of real-time fault-tolerant quantum error correction, *Phys. Rev. X* **11**, 041058 (2021).
- [52] J. Hilder, D. Pijn, O. Onishchenko, A. Stahl, M. Orth, B. Lekitsch, A. Rodriguez-Blanco, M. Müller, F. Schmidt-Kaler, and U. G. Poschinger, Fault-tolerant parity readout on a shuttling-based trapped-ion quantum computer, *Phys. Rev. X* **12**, 011032 (2022).
- [53] B. Koczor, S. Endo, T. Jones, Y. Matsuzaki, and S. C. Benjamin, Variational-state quantum metrology, *New J. Phys.* **22**, 083038 (2020).
- [54] C. Roos, M. Chwalla, K. Kim, M. Riebe, and R. Blatt, ‘Designer atoms’ for quantum metrology, *Nature* **443**, 316 (2006).
- [55] I. L. Chuang and M. A. Nielsen, Prescription for experimental determination of the dynamics of a quantum black box, *J. Mod. Opt.* **44**, 2455 (1997).
- [56] Z. Hradil, Quantum-state estimation, *Phys. Rev. A* **55**, R1561 (1997).
- [57] D. F. V. James, P. G. Kwiat, W. J. Munro, and A. G. White, Measurement of qubits, *Phys. Rev. A* **64**, 052312 (2001).
- [58] T. O. Maciel, R. O. Vianna, R. S. Sarthour, and I. S. Oliveira, Quantum process tomography with informational incomplete data of two J-coupled heterogeneous spins relaxation in a time window much greater than T_1 , *New J. Phys.* **17**, 113012 (2015).
- [59] I. Bongioanni, L. Sansoni, F. Sciarrino, G. Vallone, and P. Mataloni, Experimental quantum process tomography of non-trace-preserving maps, *Phys. Rev. A* **82**, 042307 (2010).
- [60] M.-D. Choi, Completely positive linear maps on complex matrices, *Linear Algebra Appl.* **10**, 285 (1975).
- [61] Christopher J. Wood, Ph.D. thesis, UWSpace, 2015.
- [62] P. Schindler, D. Nigg, T. Monz, J. T. Barreiro, E. Martinez, S. X. Wang, S. Quint, M. F. Brandl, V. Nebendahl, C. F. Roos, M. Chwalla, M. Hennrich, and R. Blatt, A quantum information processor with trapped ions, *New J. Phys.* **15**, 123012 (2013).
- [63] K. Mølmer and A. Sørensen, Multiparticle Entanglement of Hot Trapped Ions, *Phys. Rev. Lett.* **82**, 1835 (1999).
- [64] M. Ringbauer, M. Meth, L. Postler, R. Stricker, R. Blatt, P. Schindler, and T. Monz, A universal qudit quantum processor with trapped ions, [arXiv:2109.06903](https://arxiv.org/abs/2109.06903) [quant-ph] (2021).
- [65] A. Kreuter, C. Becher, G. P. T. Lancaster, A. B. Mundt, C. Russo, H. Häffner, C. Roos, J. Eschner, F. Schmidt-Kaler, and R. Blatt, Spontaneous Emission Lifetime of a Single Trapped Ca^+ Ion in a High Finesse Cavity, *Phys. Rev. Lett.* **92**, 203002 (2004).
- [66] D. Hayes, D. Stack, B. Bjork, A. C. Potter, C. H. Baldwin, and R. P. Stutz, Eliminating Leakage Errors in Hyperfine Qubits, *Phys. Rev. Lett.* **124**, 170501 (2020).
- [67] D. Nigg, M. Müller, E. A. Martinez, P. Schindler, M. Hennrich, T. Monz, M. A. Martin-Delgado, and R. Blatt, Quantum computations on a topologically encoded qubit, *Science* **345**, 302 (2014).
- [68] D. Niemietz, P. Farrera, S. Langenfeld, and G. Rempe, Non-destructive detection of photonic qubits, *Nature* **591**, 570 (2021).
- [69] B. M. Varbanov, F. Battistel, B. M. Tarasinski, V. P. Ostroukh, T. E. O’Brien, L. DiCarlo, and B. M. Terhal, Leakage detection for a transmon-based surface code, *Npj Quantum Inf.* **6**, 102 (2020).
- [70] J. Vala, K. B. Whaley, and D. S. Weiss, Quantum error correction of a qubit loss in an addressable atomic system, *Phys. Rev. A* **72**, 052318 (2005).
- [71] M. Grassl, T. Beth, and T. Pellizzari, Codes for the quantum erasure channel, *Phys. Rev. A* **56**, 33 (1997).

- [72] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, New York, USA, 2011), 10th ed.
- [73] T. C. Ralph, A. J. F. Hayes, and A. Gilchrist, Loss-Tolerant Optical Qubits, *Phys. Rev. Lett.* **95**, 100501 (2005).
- [74] C.-Y. Lu, W.-B. Gao, J. Zhang, X.-Q. Zhou, T. Yang, and J.-W. Pan, Experimental quantum coding against qubit loss error, *Proc. Natl. Acad. Sci. U.S.A.* **105**, 11050 (2008).
- [75] R. Laflamme, C. Miquel, J. P. Paz, and W. H. Zurek, Perfect Quantum Error Correcting Code, *Phys. Rev. Lett.* **77**, 198 (1996).
- [76] A. Kitaev, Fault-tolerant quantum computation by anyons, *Ann. Phys.* **303**, 2 (2003).
- [77] H. Bombin and M. A. Martin-Delgado, Topological Quantum Distillation, *Phys. Rev. Lett.* **97**, 180501 (2006).
- [78] H. Bombin and M. A. Martin-Delgado, Topological Computation without Braiding, *Phys. Rev. Lett.* **98**, 160502 (2007).
- [79] E. Knill, R. Laflamme, R. Martinez, and C. Negrevergne, Benchmarking Quantum Computers: The Five-Qubit Error Correcting Code, *Phys. Rev. Lett.* **86**, 5811 (2001).
- [80] J. Chiaverini, D. Leibfried, T. Schaetz, M. Barrett, R. Blakestad, J. Britton, W. Itano, J. Jost, E. Knill, C. Langer, R. Ozeri, and D. Wineland, Realization of quantum error correction, *Nature* **432**, 602 (2004).
- [81] M. D. Reed, L. DiCarlo, S. E. Nigg, L. Sun, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, Realization of three-qubit quantum error correction with superconducting circuits, *Nature* **482**, 382 (2012).
- [82] D. Ristè, S. Poletto, M. Z. Huang, A. Bruno, V. Vesterinen, O. P. Saira, and L. DiCarlo, Detecting bit-flip errors in a logical qubit using stabilizer measurements, *Nat. Commun.* **6**, 6983 (2015).
- [83] A. M. Steane, Error Correcting Codes in Quantum Theory, *Phys. Rev. Lett.* **77**, 793 (1996).
- [84] A. Bermudez, X. Xu, R. Nigmatullin, J. O’Gorman, V. Negnevitsky, P. Schindler, T. Monz, U. G. Poschinger, C. Hempel, J. Home, F. Schmidt-Kaler, M. Biercuk, R. Blatt, S. Benjamin, and M. Müller, Assessing the progress of trapped-ion processors towards fault-tolerant quantum computation, *Phys. Rev. X* **7**, 041061 (2017).
- [85] A. Bermudez, X. Xu, M. Gutiérrez, S. C. Benjamin, and M. Müller, Fault-tolerant protection of near-term trapped-ion topological qubits under realistic noise sources, *Phys. Rev. A* **100**, 062307 (2019).
- [86] M. Gutiérrez, M. Müller, and A. Bermúdez, Transversality and lattice surgery: Exploring realistic routes toward coupled logical qubits with trapped-ion quantum processors, *Phys. Rev. A* **99**, 022330 (2019).
- [87] D. Amaro, J. Bennett, D. Vodola, and M. Müller, Analytical percolation theory for topological color codes under qubit loss, *Phys. Rev. A* **101**, 032317 (2020).
- [88] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, Topological quantum memory, *J. Math. Phys.* **43**, 4452 (2002).
- [89] Such a model could be refined by adopting a circuit-level description and specific compilations of the stabilizer readout into gates, based, e.g., on recently proposed flag-qubit-based stabilizer readout protocols [85,97].
- [90] P. Parrado-Rodríguez, C. Ryan-Anderson, A. Bermudez, and M. Müller, Crosstalk suppression for fault-tolerant quantum error correction with trapped ions, *Quantum* **5**, 487 (2021).
- [91] A. L. Grimsmo and S. Puri, Quantum Error Correction with the Gottesman-Kitaev-Preskill Code, *PRX Quantum* **2**, 020101 (2021).
- [92] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth, and J. Tenynson, The variational quantum eigensolver: A review of methods and best practices, [arXiv:2111.05176](https://arxiv.org/abs/2111.05176) [quant-ph] (2021).
- [93] C. Kokail, C. Maier, R. van Bijnen, T. Brydges, M. K. Joshi, P. Jurcevic, C. A. Muschik, P. Silvi, R. Blatt, C. F. Roos, and P. Zoller, Self-verifying variational quantum simulation of lattice models, *Nature* **569**, 355 (2019).
- [94] S. Krinner, N. Lacroix, A. Remm, A. D. Paolo, E. Genois, C. Leroux, C. Hellings, S. Lazar, F. Swiadek, J. Herrmann, G. J. Norris, C. K. Andersen, M. Müller, A. Blais, C. Eichler, and A. Wallraff, Realizing repeated quantum error correction in a distance-three surface code, [arXiv:2112.03708](https://arxiv.org/abs/2112.03708) [quant-ph] (2021).
- [95] R. Stricker, D. Vodola, A. Erhard, L. Postler, M. Meth, M. Ringbauer, P. Schindler, R. Blatt, M. Müller, and T. Monz, Data for “Characterizing Quantum Instruments: From Non-demolition Measurements to Quantum Error Correction” (2021). <https://doi.org/10.5281/zenodo.6901982>.
- [96] H.-R. Wei, B.-C. Ren, and F.-G. Deng, Geometric measure of quantum discord for a two-parameter class of states in a qubit–qutrit system under various dissipative channels, *Quantum Inf. Process.* **12**, 1109 (2013).
- [97] R. Chao and B. W. Reichardt, Quantum Error Correction with Only Two Extra Qubits, *Phys. Rev. Lett.* **121**, 050502 (2018).

CONCLUSION & OUTLOOK

The field of quantum information processing is on the upswing, with recently advanced devices using more quantum information carriers than ever allowing fruitful new computational capabilities. While these novel capabilities already challenge classical computers [13], new problems arise as we move towards scalable hardware. Some of those problems have been explored in this thesis and contribute the ongoing developments, not only for the trapped-ion device presented here, but for quantum hardware in general.

The topics presented in this thesis follow broadly two different paradigms. The first is certification of large-scale quantum hardware using classically tractable computation outcomes (single-setting QST in Ch. 3, classical verification in Ch. 4 and characterizing quantum instruments in Ch. 5). The second is a means to keep such large quantum devices functional, even if some of their information carriers get lost due to the unavoidable influence of environmental noise (qubit loss correction in Ch. 5).

First reported in Ch. 3 is a novel single-setting QST offering significant relief from the heavy measurement and sampling requirements of existing tomography tools. The method is practical in many ways, e.g., because it does not require costly changes in measurement settings and allows efficient sequence compiling—both reducing the overall time cost of experiments. We combined single-setting QST with classical shadows data analysis [168] and demonstrated state reconstruction at the minimum dimension of the N -qubit density matrix ($2^N \times 2^N$). This single-setting QST based on classical shadows can furthermore predict arbitrary polynomial function properties of the density matrix only from the subset of qubits on which the operators act on. Avoiding the full density matrix reconstruction makes it orders of magnitudes faster than existing methods. The larger the overall system becomes in terms of the analyzed subsets, the more improvements can be observed over existing characterization tools. Further project extensions could focus on more rigorous entanglement studies covering, e.g., higher-order Rényi-entropies [48]. Another promising research direction is state-dependent convergence behaviour in view of differently orientated measurement sets. This holds the potential to speed-up applications like VQE [37, 169] that rely on the simultaneous prediction of multiple observables [163] by orders of magnitude. From a technical point of view, there is a need to speed up the camera readout to improve the detection of qudits, since a long pause (compared to the time for gate operations) between detections currently leaves the system prone to amplitude damping.

Moving away from rigorous device testing towards cryptographically-secure certification methods, Ch. 4 presented the first experimental verification of a quantum computation by purely classical means [63]. We tested this novel verification technique in view of the capabilities of current NISQ-hardware and drew attention to the significant operational overhead required for the fully secure procedure. Yet even the proof-of-principle case presented pushed the limits of current devices. As such, we slightly amended the original

proposal in favor of a feasible protocol incorporating eight qubits to verify one [54]. This protocol, however, can be straightforwardly scaled up towards larger systems. The next bigger implementation targets the verification of a two-qubit operation requiring twice as many qubits. For a fully secure implementation, the protocol required hundreds of extra qubits for each one to verify and demands orders of magnitude lower noise rates than currently available. In order to offer secure options for devices in the near future, the theoretical concepts of classical verification must be refined to either reduce the operational overhead of auxiliary information carriers or increase their tolerable noise levels.

In-depth system characterizations underpin that quantum systems remain prone to errors due to unwanted environmental interactions. At the same time, advanced quantum tasks often use additional levels beyond the computational subspace to store more information than the qubit or to simplify circuit implementations. The use of larger dimensions, however, bears the risk of leakage errors to beyond the computational subspace. Yet, leakage errors are mostly ignored by today's QEC applications, focusing on errors that change the logical state. Along those lines, we demonstrated the first detection and correction of qubit loss in real-time in Ch. 5. Crucially, we built qubit loss correction into the surface code [79] to in principle allow for computational errors on top of losses and to keep the operational overhead for their combined correction as low as possible. While our demonstrations target the smallest surface code fragment, our experiments cover all the necessary steps for loss correction that are readily applicable to larger surface codes [72, 209]. In a follow-up work, a fault-tolerant implementation in the presence of computational errors could be demonstrated—a fundamental next step towards realistic QEC.

Inspired by the semi-classical algorithm structure of the QND loss detection unit featuring in-sequence measurement and classical feed-forward, we presented a suitable characterization toolbox based on quantum instrument tomography [129] at the end of Ch. 5. In doing so, we moved away from existing characterization methods that limit dynamics to unitary evolution. Beyond loss correction, non-unitary evolution complies with the needs of many novel quantum algorithms [122–128]. Our instrument tomography framework enables the identification of erroneous features that go unnoticed by standard analysis tools, but are critical for high-precision and future fault-tolerant QEC applications. The results provide detailed information about deviations from unitary processes and emphasize the need for careful characterization of quantum computing building blocks—independent of the platform. Using this detailed error information, we were able to numerically investigate the effects of experimental failures on QEC performance, using our loss detection instrument as a building block in QEC protocols. These simulations also provided the parameter regimes for simultaneous qubit loss and computational error correction in the context of a state-of-the-art QEC code—information that is paramount for realizing fault-tolerance.

To finally demonstrate quantum advantage, we need new setup designs that can control hundreds of qubits and reduce error-rates by at least an order of magnitude compared to existing devices. These technological improvements require the support of extensive device testing to determine system limitations and ultimately verify their computational outcomes in a cryptographically-secure way, which we have supported with the work presented here. Since the influence of noise remains present on large-scale devices and the types of errors are more general than just computational ones, loss detection and correction must furthermore become a standard building block in QEC protocols.



LIST OF PUBLICATIONS

All results presented throughout this thesis led to the following journal publications, presented in chronological order:

1. [218] R. Stricker *et al.* Experimental deterministic correction of qubit loss. *Nature* 585, 207–210 (2020).
2. [131] R. Stricker *et al.* Characterizing quantum instruments: from non-demolition measurements to quantum error correction. *PRX Quantum* 3, 030318 (2022).
3. [200] R. Stricker *et al.* Towards experimental classical verification of quantum computation. *Quantum Sci. Technol.* Volume 9, Number 2 (2024).
4. [40] R. Stricker *et al.* Experimental single-setting quantum state tomography. *PRX Quantum* 3, 040310 (2022).

During the course of my PhD work, I additionally contributed to the following journal publications:

5. [223] L. Postler *et al.* Experimental quantification of spatial correlations in quantum dynamics. *Quantum* 2, 90 (2018).
6. [52] A. Erhard, J. J. Wallman *et al.* Characterizing large-scale quantum computers via cycle benchmarking. *Nature Comm.* 10, 5347 (2019).
7. [60] C. Greganti *et al.* Cross-Verification of Independent Quantum Devices. *Phys. Rev. X* 11, 031049 (2021).
8. [92] A. Erhard, H. Poulsen-Nautrup *et al.* Entangling logical qubits with lattice surgery. *Nature* 589, 220–224 (2021).
9. [140] M. Ringbauer *et al.* A universal qudit quantum processor with trapped ions. *Nat. Phys.* 18, 1053–1057 (2022).
10. [70] L. Postler *et al.* Demonstration of fault-tolerant universal quantum gate operations. *Nature* 605, 675–680 (2022).
11. [224] M. Meth, V. Kuzmin *et al.* Probing phases of quantum matter with an ion-trap tensor-network quantum eigensolver. *Phys. Rev. X* 12, 041035 (2022).
12. [225] M. Ringbauer *et al.* Verifiable measurement-based quantum random sampling with trapped ions. *arXiv* : 2307.14424 (2023).

BIBLIOGRAPHY

- [1] M. F. Riedel, D. Binosi, R. Thew, and T. Calarco, “The european quantum technologies flagship programme,” *Quantum Science and Technology*, vol. 2, no. 3, p. 030 501, 2017.
- [2] R. P. Feynman, “Simulating physics with computers,” *International journal of theoretical physics*, vol. 21, no. 6, pp. 467–488, 1982.
- [3] E. Rico *et al.*, “So(3) nuclear physics with ultracold gases,” *Ann. Phys.*, pp. 466–483, 2018.
- [4] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, “Simulated quantum computation of molecular energies,” *Science*, vol. 309, no. 5741, pp. 1704–1707, 2005.
- [5] D. Deutsch, “Quantum theory, the church-turing principle and the universal quantum computer,” in *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, The Royal Society, vol. 400, 1985, pp. 97–117.
- [6] M. Schuld, I. Sinayskiy, and F. Petruccione, “An introduction to quantum machine learning,” *Contemporary Physics*, vol. 56, no. 2, pp. 172–185, 2015.
- [7] A. Luckow, J. Klepsch, and J. Pichlmeier, “Quantum computing: Towards industry reference problems,” *Digitale Welt*, vol. 5, no. 2, pp. 38–45, 2021.
- [8] J. Biamonte *et al.*, “Quantum machine learning,” *Nature*, vol. 549, no. 7671, pp. 195–202, 2017.
- [9] P. Shor, “Algorithms for quantum computation: Discrete logarithms and factoring,” in *Proceedings 35th Annual Symposium on Foundations of Computer Science*, 1994, pp. 124–134.
- [10] F. Schmidt-Kaler *et al.*, “How to realize a universal quantum gate with trapped ions,” *Applied Physics B*, vol. 77, no. 8, pp. 789–796, 2003.
- [11] A. Acín *et al.*, “The quantum technologies roadmap: A european community view,” *New Journal of Physics*, vol. 20, no. 8, p. 080 201, 2018.
- [12] J. Preskill, “Quantum computing and the entanglement frontier,” *Bull. Am. Phys. Soc.*, p. 58, 2013.
- [13] F. Arute *et al.*, “Quantum supremacy using a programmable superconducting processor,” *Nature*, vol. 574, no. 7779, pp. 505–510, 2019.
- [14] R. Acharya *et al.*, “Suppressing quantum errors by scaling a surface code logical qubit,” *Nature*, vol. 614, no. 7949, pp. 676–681, 2023.

- [15] F. Kranzl *et al.*, “Controlling long ion strings for quantum simulation and precision measurements,” *Phys. Rev. A*, vol. 105, p. 052 426, 5 2022.
- [16] S. Aaronson. “The Aaronson \$25.00 prize.” (2004), [Online]. Available: www.scottaaronson.com/blog/?p=28 (visited on 01/11/2024).
- [17] D. Deutsch, *The Beginning of Infinity: Explanations that Transform the World* (Always learning). Penguin Books, 2012.
- [18] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010.
- [19] J. M. Renes, R. Blume-Kohout, A. J. Scott, and C. M. Caves, “Symmetric informationally complete quantum measurements,” *J. Math. Phys.*, vol. 45, pp. 2171–2180, 2004.
- [20] A. J. Scott and M. Grassl, “Symmetric informationally complete positive-operator-valued measures: A new computer study,” *J. Math. Phys.*, vol. 51, p. 042 203, 2010.
- [21] G. Vidal, “Efficient classical simulation of slightly entangled quantum computations,” *Phys. Rev. Lett.*, vol. 91, p. 147 902, 14 2003.
- [22] A. Peres, “Separability criterion for density matrices,” *Phys. Rev. Lett.*, vol. 77, pp. 1413–1415, 8 1996.
- [23] M. Horodecki, P. Horodecki, and R. Horodecki, “Separability of mixed states: Necessary and sufficient conditions,” *Physics Letters A*, vol. 223, no. 1, pp. 1–8, 1996.
- [24] D. M. Greenberger, M. A. Horne, A. Shimony, and A. Zeilinger, “Bell’s theorem without inequalities,” *American Journal of Physics*, vol. 58, no. 12, pp. 1131–1143, 1990.
- [25] G. Vidal and R. F. Werner, “Computable measure of entanglement,” *Phys. Rev. A*, vol. 65, p. 032 314, 2002.
- [26] A. Rényi, “On the foundations of information theory,” *Review of the International Statistical Institute*, vol. vol. 33, no. no. 1, p. 1 14, 1965.
- [27] A. Elben, B. Vermersch, C. F. Roos, and P. Zoller, “Statistical correlations between locally randomized measurements: A toolbox for probing entanglement in many-body quantum states,” *Phys. Rev. A*, vol. 99, p. 052 323, 5 2019.
- [28] M. Naimark, “Spectral functions of a symmetric operator,” *Bull. Acad. Sci. URSS. Sér. Math. [Izvestia Akad. Nauk SSSR]*, vol. 4, pp. 277–318, 1940.
- [29] W. M. Itano *et al.*, “Quantum projection noise: Population fluctuations in two-level systems,” *Phys. Rev. A*, vol. 47, pp. 3554–3570, 5 1993.
- [30] R. Jozsa, “Fidelity for mixed quantum states,” *Journal of Modern Optics*, vol. 41, no. 12, pp. 2315–2323, 1994.
- [31] D. DiVincenzo and IBM, “The physical implementation of quantum computation,” *Fortschritte der Physik*, vol. 48, 2000.

- [32] R. J. MacDonell *et al.*, “Analog quantum simulation of chemical dynamics,” *Chem. Sci.*, vol. 12, no. 28, pp. 9794–9805, 2021.
- [33] M. Gell-Mann, “Symmetries of baryons and mesons,” *Phys. Rev.*, vol. 125, pp. 1067–1084, 3 1962.
- [34] J. Eisert *et al.*, “Quantum certification and benchmarking,” *Nat. Rev. Phys.*, vol. 2, no. 7, pp. 382–390, 2020.
- [35] E. Pednault *et al.*, *Pareto-efficient quantum circuit simulation using tensor contraction deferral*, 2020. arXiv: [1710.05867](https://arxiv.org/abs/1710.05867) [quant-ph].
- [36] A. Erhard, “Towards scalable quantum computation with trapped ions,” Ph.D. dissertation, University of Innsbruck, 6020 Innsbruck, Austria, 2022.
- [37] C. Kokail *et al.*, “Self-verifying variational quantum simulation of lattice models,” *Nature*, vol. 569, no. 7756, pp. 355–360, 2019.
- [38] D. F. V. James, P. G. Kwiat, W. J. Munro, and A. G. White, “Measurement of qubits,” *Physical Review A*, vol. 64, no. 5, 2001.
- [39] E. Pelofske, A. Bartschi, and S. Eidenbenz, “Quantum volume in practice: What users can expect from NISQ devices,” *IEEE Transactions on Quantum Engineering*, vol. 3, pp. 1–19, 2022.
- [40] R. Stricker *et al.*, “Experimental single-setting quantum state tomography,” *PRX Quantum*, vol. 3, p. 040 310, 4 2022.
- [41] S. T. Flammia and Y.-K. Liu, “Direct fidelity estimation from few pauli measurements,” *Phys. Rev. Lett.*, vol. 106, p. 230 501, 23 2011.
- [42] M. P. da Silva, O. Landon-Cardinal, and D. Poulin, “Practical characterization of quantum devices without tomography,” *Phys. Rev. Lett.*, vol. 107, p. 210 404, 21 2011.
- [43] S. Pallister, N. Linden, and A. Montanaro, “Optimal verification of entangled states with local measurements,” *Phys. Rev. Lett.*, vol. 120, p. 170 502, 17 2018.
- [44] M. Holzäpfel, T. Baumgratz, M. Cramer, and M. B. Plenio, “Scalable reconstruction of unitary processes and hamiltonians,” *Phys. Rev. A*, vol. 91, p. 042 129, 4 2015.
- [45] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, “Quantum state tomography via compressed sensing,” *Phys. Rev. Lett.*, vol. 105, p. 150 401, 15 2010.
- [46] B. P. Lanyon *et al.*, “Efficient tomography of a quantum many-body system,” *Nature Physics*, vol. 13, no. 12, pp. 1158–1162, 2017.
- [47] T. Brydges *et al.*, “Probing rényi entanglement entropy via randomized measurements,” *Science*, vol. 364, no. 6437, pp. 260–263, 2019.
- [48] A. Elben *et al.*, “Mixed-state entanglement from local randomized measurements,” *Phys. Rev. Lett.*, vol. 125, p. 200 501, 20 2020.
- [49] J. Helsen, I. Roth, E. Onorati, A. Werner, and J. Eisert, “General framework for randomized benchmarking,” *PRX Quantum*, vol. 3, p. 020 357, 2 2022.

- [50] E. Nielsen *et al.*, “Gate Set Tomography,” *Quantum*, vol. 5, p. 557, 2021.
- [51] K. Rudinger *et al.*, “Characterizing midcircuit measurements on a superconducting qubit using gate set tomography,” *Phys. Rev. Applied*, vol. 17, p. 014 014, 1 2022.
- [52] A. Erhard *et al.*, “Characterizing large-scale quantum computers via cycle benchmarking,” *Nature Communications*, vol. 10, no. 1, p. 5347, 2019.
- [53] M. R. Garey and D. S. Johnson, *Computers and Intractability; A Guide to the Theory of NP-Completeness*. USA: W. H. Freeman Co., 1990.
- [54] J. Carrasco, A. Elben, C. Kokail, B. Kraus, and P. Zoller, “Theoretical and experimental perspectives of quantum verification,” *PRX Quantum*, vol. 2, p. 010 102, 2021.
- [55] B. W. Boehm, “Software engineering economics,” in *Pioneers and Their Contributions to Software Engineering: sd&m Conference on Software Pioneers, Bonn, June 28/29, 2001, Original Historic Contributions*, M. Broy and E. Denert, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2001, pp. 99–150.
- [56] A. Gheorghiu, T. Kapourniotis, and E. Kashefi, “Verification of quantum computation: An overview of existing approaches,” *Theory of Computing Systems*, vol. 63, no. 4, pp. 715–808, 2019.
- [57] I. Šupić and J. Bowles, “Self-testing of quantum systems: A review,” *Quantum*, vol. 4, p. 337, 2020.
- [58] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, “Proposed experiment to test local hidden-variable theories,” *Phys. Rev. Lett.*, vol. 23, pp. 880–884, 15 1969.
- [59] A. Elben *et al.*, “Cross-platform verification of intermediate scale quantum devices,” *Phys. Rev. Lett.*, vol. 124, p. 010 504, 1 2020.
- [60] C. Greganti *et al.*, “Cross-verification of independent quantum devices,” *Phys. Rev. X*, vol. 11, p. 031 049, 3 2021.
- [61] T. Morimae, “Verification for measurement-only blind quantum computing,” *Physical Review A*, vol. 89, no. 6, p. 060 302, 2014.
- [62] S. Barz, J. F. Fitzsimons, E. Kashefi, and P. Walther, “Experimental verification of quantum computation,” *Nature Physics*, vol. 9, no. 11, pp. 727–731, 2013.
- [63] U. Mahadev, “Classical verification of quantum computations,” in *IEEE 59th Annu. Symp. Found. Comput. Sci.*, 2018, pp. 259–267.
- [64] D. Gottesman, “Theory of fault-tolerant quantum computation,” *Phys. Rev. A*, vol. 57, pp. 127–137, 1 1998.
- [65] J. Roffe, “Quantum error correction: An introductory guide,” *Contemporary Physics*, vol. 60, no. 3, pp. 226–245, 2019.
- [66] W. K. Wootters and W. H. Zurek, “A single quantum cannot be cloned,” *Nature*, vol. 299, no. 5886, pp. 802–803, 1982.

- [67] P. W. Shor, "Scheme for reducing decoherence in quantum computer memory," *Phys. Rev. A*, vol. 52, R2493–R2496, 4 1995.
- [68] E. T. Campbell, "Enhanced Fault-Tolerant Quantum Computing in d-Level Systems," *Phys. Rev. Lett.*, vol. 113, p. 230 501, 2014.
- [69] J. Preskill, "Fault-tolerant quantum computation, in introduction to quantum computation and information," *Singapore, World scientific*, 1998.
- [70] L. Postler *et al.*, "Demonstration of fault-tolerant universal quantum gate operations," *Nature*, vol. 605, no. 7911, pp. 675–680, 2022.
- [71] W. H. Zurek and R. Laflamme, "Quantum logical operations on encoded qubits," *Phys. Rev. Lett.*, vol. 77, pp. 4683–4686, 22 1996.
- [72] H. Bombin and M. A. Martin-Delgado, "Topological quantum distillation," *Phys. Rev. Lett.*, vol. 97, p. 180 501, 18 2006.
- [73] H. Bombin and M. A. Martin-Delgado, "Topological computation without braiding," *Phys. Rev. Lett.*, vol. 98, p. 160 502, 16 2007.
- [74] D. Gottesman, "The heisenberg representation of quantum computers," *Cambridge, MA, International Press, Group22: Proceedings of the XXII International Colloquium on Group Theoretical Methods in Physics*, eds. S. P. Corney, R. Delbourgo, and P. D. Jarvis, pp. 32–43, 1999.
- [75] E. Knill and R. Laflamme, "Theory of quantum error-correcting codes," *Phys. Rev. A*, vol. 55, pp. 900–911, 2 1997.
- [76] A. Ekert and C. Macchiavello, "Quantum error correction for communication," *Phys. Rev. Lett.*, vol. 77, pp. 2585–2588, 12 1996.
- [77] E. Dennis, A. Kitaev, A. Landahl, and P. J., "Topological quantum memory," *J. Math. Phys.*, vol. 43, no. 9, p. 4452, 2002.
- [78] A. Kitaev, "Fault-tolerant quantum computation by anyons," *Ann. Phys.*, vol. 303, no. 1, pp. 2–30, 2003.
- [79] A. Y. Kitaev, "Quantum computations: Algorithms and error correction," *Russian Mathematical Surveys*, vol. 52, no. 6, pp. 1191–1249, 1997.
- [80] D. Aharonov and M. Ben-Or, "Fault-tolerant quantum computation with constant error rate," *SIAM Journal on Computing*, vol. 38, no. 4, pp. 1207–1282, 2008.
- [81] E. Knill, R. Laflamme, and W. H. Zurek, "Resilient quantum computation," *Science*, vol. 279, no. 5349, pp. 342–345, 1998.
- [82] D. Gottesman, *Stabilizer codes and quantum error correction*, 1997. arXiv: [quant-ph/9705052](https://arxiv.org/abs/quant-ph/9705052) [quant-ph].
- [83] V. B. Braginsky, Y. I. Vorontsov, and K. S. Thorne, "Quantum nondemolition measurements," *Science*, vol. 209, no. 4456, pp. 547–557, 1980.

- [84] E. Knill, R. Laflamme, R. Martinez, and C. Negrevergne, "Benchmarking quantum computers: The five-qubit error correcting code," *Phys. Rev. Lett.*, vol. 86, p. 5811, 2001.
- [85] J. Chiaverini *et al.*, "Realization of quantum error correction," *Nature*, vol. 432, no. 7017, pp. 602–605, 2004.
- [86] X.-C. Yao *et al.*, "Experimental demonstration of topological error correction," *Nature*, vol. 482, p. 489, 2012.
- [87] P. Schindler *et al.*, "A quantum information processor with trapped ions," *New Journal of Physics*, vol. 15, no. 12, p. 123 012, 2013.
- [88] D. Nigg *et al.*, "Quantum computations on a topologically encoded qubit," *Science*, vol. 345, no. 6194, pp. 302–305, 2014.
- [89] A. D. Córcoles *et al.*, "Demonstration of a quantum error detection code using a square lattice of four superconducting qubits," *Nat. Commun.*, vol. 6, p. 6979, 2015.
- [90] M. Takita, A. W. Cross, A. D. Córcoles, J. M. Chow, and J. M. Gambetta, "Experimental demonstration of fault-tolerant state preparation with superconducting qubits," *Phys. Rev. Lett.*, vol. 119, p. 180 501, 2017.
- [91] N. M. Linke *et al.*, "Fault-tolerant quantum error detection," *Sci. Adv.*, vol. 3, no. 10, e1701074, 2017.
- [92] A. Erhard *et al.*, "Entangling logical qubits with lattice surgery," *Nature*, vol. 589, no. 7841, pp. 220–224, 2021.
- [93] C. Ryan-Anderson *et al.*, "Realization of real-time fault-tolerant quantum error correction," *Phys. Rev. X*, vol. 11, p. 041 058, 4 2021.
- [94] T. M. Stace, S. D. Barrett, and A. C. Doherty, "Thresholds for topological codes in the presence of loss," *Phys. Rev. Lett.*, vol. 102, p. 200 501, 2009.
- [95] J. Kruse, C. Gierl, M. Schlosser, and G. Birkel, "Reconfigurable site-selective manipulation of atomic quantum systems in two-dimensional arrays of dipole traps," *Phys. Rev. A*, vol. 81, p. 060 308, 6 2010.
- [96] P. Groszkowski, A. G. Fowler, F. Motzoi, and F. K. Wilhelm, "Tunable coupling between three qubits as a building block for a superconducting quantum computer," *Phys. Rev. B*, vol. 84, p. 144 516, 14 2011.
- [97] D. A. Herrera-Martí, A. G. Fowler, D. Jennings, and T. Rudolph, "Photonic implementation for the topological cluster-state quantum computer," *Phys. Rev. A*, vol. 82, p. 032 332, 3 2010.
- [98] M. Kumph, M. Brownnutt, and R. Blatt, "Two-dimensional arrays of radio-frequency ion traps with addressable interactions," *New Journal of Physics*, vol. 13, no. 7, p. 073 043, 2011.
- [99] J. Kelly *et al.*, "Scalable in situ qubit calibration during repetitive error detection," *Phys. Rev. A*, vol. 94, p. 032 321, 3 2016.

- [100] J. O’Gorman, N. H. Nickerson, P. Ross, J. J. Morton, and S. C. Benjamin, “A silicon-based surface code quantum computer,” *npj Quantum Information*, vol. 2, no. 1, p. 15 019, 2016.
- [101] S. Krinner *et al.*, “Realizing repeated quantum error correction in a distance-three surface code,” *Nature*, vol. 605, no. 7911, pp. 669–674, 2022.
- [102] D. S. Wang, A. G. Fowler, and L. C. L. Hollenberg, “Surface code quantum computing with error rates over 1%,” *Phys. Rev. A*, vol. 83, p. 020 302, 2 2011.
- [103] S. Gulde, “Experimental realization of quantum gates and the Deutsch-Josza algorithm with trapped $^{40}\text{Ca}^+$ -ions,” Ph.D. dissertation, University of Innsbruck, 6020 Innsbruck Austria, 2003.
- [104] B. E. Anderson, H. Sosa-Martinez, C. A. Riofrío, I. H. Deutsch, and P. S. Jessen, “Accurate and robust unitary transformations of a high-dimensional quantum system,” *Phys. Rev. Lett.*, vol. 114, p. 240 401, 24 2015.
- [105] C. Godfrin *et al.*, “Operating quantum states in single magnetic molecules: Implementation of grover’s quantum algorithm,” *Phys. Rev. Lett.*, vol. 119, p. 187 702, 18 2017.
- [106] J. Ahn, T. C. Weinacht, and P. H. Bucksbaum, “Information storage and retrieval through quantum phase,” *Science*, vol. 287, no. 5452, pp. 463–465, 2000.
- [107] Y. Chi *et al.*, “A programmable qudit-based quantum processor,” *Nature Communications*, vol. 13, no. 1, p. 1166, 2022.
- [108] M. S. Blok *et al.*, “Quantum information scrambling on a superconducting qutrit processor,” *Phys. Rev. X*, vol. 11, p. 021 010, 2 2021.
- [109] X.-M. Hu *et al.*, “Beating the channel capacity limit for superdense coding with entangled ququarts,” *Science Advances*, vol. 4, no. 7, eaat9304, 2018.
- [110] A. M. Steane, “Efficient fault-tolerant quantum computing,” *Nature*, vol. 399, no. 6732, pp. 124–126, 1999.
- [111] M. Varnava, D. E. Browne, and T. Rudolph, “Loss tolerance in one-way quantum computation via counterfactual error correction,” *Phys. Rev. Lett.*, vol. 97, p. 120 501, 12 2006.
- [112] S. D. Barrett and T. M. Stace, “Fault tolerant quantum computation with very high threshold for loss errors,” *Phys. Rev. Lett.*, vol. 105, p. 200 502, 20 2010.
- [113] T. M. Stace and S. D. Barrett, “Error correction and degeneracy in surface codes suffering loss,” *Phys. Rev. A*, vol. 81, p. 022 317, 2 2010.
- [114] R. Raussendorf, S. Bravyi, and J. Harrington, “Long-range quantum entanglement in noisy cluster states,” *Phys. Rev. A*, vol. 71, p. 062 313, 6 2005.
- [115] R. Raussendorf, J. Harrington, and K. Goyal, “Topological fault-tolerance in cluster state quantum computation,” *New Journal of Physics*, vol. 9, no. 6, p. 199, 2007.

- [116] R. Raussendorf and J. Harrington, "Fault-tolerant quantum computation with high threshold in two dimensions," *Phys. Rev. Lett.*, vol. 98, p. 190 504, 19 2007.
- [117] C. M. Dawson, H. L. Haselgrove, and M. A. Nielsen, "Noise thresholds for optical quantum computers," *Phys. Rev. Lett.*, vol. 96, p. 020 501, 2 2006.
- [118] A. G. Fowler, "Coping with qubit leakage in topological codes," *Phys. Rev. A*, vol. 88, p. 042 308, 2013.
- [119] J. Ghosh, A. G. Fowler, J. M. Martinis, and M. R. Geller, "Understanding the effects of leakage in superconducting quantum-error-detection circuits," *Phys. Rev. A*, vol. 88, p. 062 329, 2013.
- [120] P. Schindler *et al.*, "Experimental repetitive quantum error correction," *Science*, vol. 332, no. 6033, pp. 1059–1061, 2011.
- [121] V. Negnevitsky *et al.*, "Repeated multi-qubit readout and feedback with a mixed-species trapped-ion register," *Nature*, vol. 563, no. 7732, pp. 527–531, 2018.
- [122] G. Chiribella, G. M. D'Ariano, and P. Perinotti, "Theoretical framework for quantum networks," *Phys. Rev. A*, vol. 80, no. 2, p. 022 339, 2009.
- [123] O. Oreshkov, F. Costa, and Č. Brukner, "Quantum correlations with no causal order," *Nat. Commun.*, vol. 3, p. 1092, 2012.
- [124] F. Buscemi and M. F. Sacchi, "Information-disturbance trade-off in quantum-state discrimination," *Phys. Rev. A*, vol. 74, no. 5, p. 052 320, 2006.
- [125] L. Knips *et al.*, *Measurement-disturbance tradeoff outperforming optimal cloning*, 2018. arXiv: [1808.07882](https://arxiv.org/abs/1808.07882) [quant-ph].
- [126] S. Lloyd and J.-J. E. Slotine, "Quantum feedback with weak measurements," *Phys. Rev. A*, vol. 62, p. 012 307, 2000.
- [127] Q. Sun, M. Al-Amri, L. Davidovich, and M. S. Zubairy, "Reversing entanglement change by a weak measurement," *Phys. Rev. A*, vol. 82, no. 5, p. 052 323, 2010.
- [128] Y.-S. Kim, J.-C. Lee, O. Kwon, and Y.-H. Kim, "Protecting entanglement from decoherence using weak measurement and quantum measurement reversal," *Nat. Phys.*, vol. 8, no. 2, pp. 117–120, 2012.
- [129] E. B. Davies and J. T. Lewis, "An operational approach to quantum probability," *Commun. Math. Phys.*, vol. 17, no. 3, p. 239, 1970.
- [130] M. Ozawa, "Quantum measuring processes of continuous observables," *J. Math. Phys.*, vol. 25, p. 79, 1984.
- [131] R. Stricker *et al.*, "Characterizing quantum instruments: From nondemolition measurements to quantum error correction," *PRX Quantum*, vol. 3, p. 030 318, 3 2022.
- [132] W. Paul, "Electromagnetic traps for charged and neutral particles," *Rev. Mod. Phys.*, vol. 62, pp. 531–540, 3 1990.

- [133] R. Stricker, "Gatteroperationen hoher Güte in einem optischen Quantenbit," M.S. thesis, Universität Innsbruck, 6020 Innsbruck Austria, 2017.
- [134] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland, "Quantum dynamics of single trapped ions," *Rev. Mod. Phys.*, vol. 75, pp. 281–324, 1 2003.
- [135] H. Häffner, C. Roos, and R. Blatt, "Quantum computing with trapped ions," *Physics Reports*, vol. 469, no. 4, pp. 155–203, 2008.
- [136] C. Raab *et al.*, "Motional sidebands and direct measurement of the cooling rate in the resonance fluorescence of a single trapped ion," *Phys. Rev. Lett.*, vol. 85, pp. 538–541, 3 2000.
- [137] J. Jin and D. A. Church, "Precision lifetimes for the Ca^+ $4p\ 2p$ levels: Experiment challenges theory at the 1% level," *Phys. Rev. Lett.*, vol. 70, pp. 3213–3216, 21 1993.
- [138] A. Kreuter *et al.*, "Experimental and theoretical study of the $3d\ 2D$ -level lifetimes of $^{40}\text{Ca}^+$," *Phys. Rev. A*, vol. 71, p. 032 504, 3 2005.
- [139] C. J. Foot, *Atomic physics* (Oxford master series in atomic, optical, and laser physics). Oxford: Oxford University Press, 2007.
- [140] M. Ringbauer *et al.*, "A universal qudit quantum processor with trapped ions," *Nature Physics*, vol. 18, no. 9, pp. 1053–1057, 2022.
- [141] M. Chwalla, "Precision spectroscopy with $^{40}\text{Ca}^+$ ions in a paul trap," Ph.D. dissertation, University of Innsbruck, 6020 Innsbruck, Austria, 2009.
- [142] A. Sørensen and K. Mølmer, "Entanglement and quantum computation with ions in thermal motion," *Phys. Rev. A*, vol. 62, p. 022 311, 2 2000.
- [143] C. Roos, "Controlling the quantum state of trapped ions," Ph.D. dissertation, University of Innsbruck, 6020 Innsbruck Austria, 2000.
- [144] R. Loudon, *The Quantum Theory of Light* (Oxford science publications), third edition. Oxford University Press, 2003.
- [145] M. Meth, "Dynamische Kontrolle von Laserimpulsen zur Quanteninformationsverarbeitung," M.S. thesis, University of Innsbruck, 2017.
- [146] D. Nigg, "Towards fault tolerant quantum computation," Ph.D. thesis, Institute for Experimental Physics, University of Innsbruck, 2016.
- [147] D. Maslov, "Basic circuit compilation techniques for an ion-trap quantum machine," *New Journal of Physics*, vol. 19, no. 2, p. 023 035, 2017.
- [148] M. Joshi *et al.*, "Polarization-gradient cooling of 1d and 2d ion coulomb crystals," *New Journal of Physics*, vol. 22, no. 10, p. 103 013, 2020.
- [149] T. Monz, "Quantum information processing beyond ten ion-qubits," Ph.D. dissertation, University of Innsbruck, Innsbruck 6020 Austria, 2011.
- [150] M. Brandl, "Towards cryogenic scalable quantum computing with trapped ions," Ph.D. dissertation, University of Innsbruck, 6020 Innsbruck, Austria, 2017.

- [151] K. Bharti *et al.*, “Noisy intermediate-scale quantum algorithms,” *Rev. Mod. Phys.*, vol. 94, p. 015 004, 1 2022.
- [152] Z. Hradil, “Quantum-state estimation,” *Phys. Rev. A*, vol. 55, R1561–R1564, 3 1997.
- [153] C. J. Wood, “Initialization and characterization of open quantum systems,” Ph.D. dissertation, University of Waterloo, 2015.
- [154] G. D’Ariano, L. Maccone, and M. Paris, “Orthogonality relations in quantum tomography,” *Physics Letters A*, vol. 276, no. 1, pp. 25–30, 2000.
- [155] R. Penrose, “On best approximate solutions of linear matrix equations,” *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 52, no. 1, p. 17 19, 1956.
- [156] J. A. Smolin, J. M. Gambetta, and G. Smith, “Efficient Method for Computing the Maximum-Likelihood Quantum State from Measurements with Additive Gaussian Noise,” *Physical Review Letters*, vol. 108, no. 7, p. 070 502, 2012.
- [157] M. Guță, J. Kahn, R. Kueng, and J. A. Tropp, “Fast state tomography with optimal error bounds,” *Journal of Physics A: Mathematical and Theoretical*, vol. 53, no. 20, p. 204 001, 2020.
- [158] L. Vandenberghe and S. Boyd, “Semidefinite programming,” *SIAM Review*, vol. 38, no. 1, pp. 49–95, 1996.
- [159] C. Ferrie and R. Blume-Kohout, *Maximum likelihood quantum state tomography is inadmissible*, 2018. arXiv: [1808.01072](https://arxiv.org/abs/1808.01072) [quant-ph].
- [160] J. B. Altepeter *et al.*, “Ancilla-assisted quantum process tomography,” *Phys. Rev. Lett.*, vol. 90, p. 193 601, 19 2003.
- [161] A. Jamiolkowski, “Linear transformations which preserve trace and positive semidefiniteness of operators,” vol. 3, no. 4, pp. 275–278, 1972.
- [162] M. Ringbauer *et al.*, “Characterizing Quantum Dynamics with Initial System-Environment Correlations,” *Phys. Rev. Lett.*, vol. 114, p. 090 402, 2015.
- [163] G. García-Pérez *et al.*, “Learning to measure: Adaptive informationally complete generalized measurements for quantum algorithms,” *PRX Quantum*, vol. 2, p. 040 342, 4 2021.
- [164] F. G. Brandão, W. Chemissany, N. Hunter-Jones, R. Kueng, and J. Preskill, “Models of quantum complexity growth,” *PRX Quantum*, vol. 2, p. 030 316, 3 2021.
- [165] N. Bent *et al.*, “Experimental realization of quantum tomography of photonic qudits via symmetric informationally complete positive operator-valued measures,” *Phys. Rev. X*, vol. 5, p. 041 006, 4 2015.
- [166] P.-X. Chen, J. A. Bergou, S.-Y. Zhu, and G.-C. Guo, “Ancilla dimensions needed to carry out positive-operator-valued measurement,” *Phys. Rev. A*, vol. 76, p. 060 303, 2007.

- [167] S. Aaronson, "Shadow tomography of quantum states," *SIAM Journal on Computing*, vol. 49, no. 5, STOC18-368-STOC18-394, 2020. eprint: <https://doi.org/10.1137/18M120275X>.
- [168] H.-Y. Huang, R. Kueng, and J. Preskill, "Predicting many properties of a quantum system from very few measurements," *Nature Physics*, vol. 16, no. 10, pp. 1050–1057, 2020.
- [169] J. Tilly *et al.*, "The variational quantum eigensolver: A review of methods and best practices," *Physics Reports*, vol. 986, pp. 1–128, 2022.
- [170] W. Helwig and W. Cui, *Absolutely maximally entangled states: Existence and applications*, 2013. arXiv: [1306.2536](https://arxiv.org/abs/1306.2536) [quant-ph].
- [171] M. Enriquez, I. Wintrowicz, and K. Zyczkowski, "Maximally entangled multipartite states: A brief survey," *Journal of Physics: Conference Series*, vol. 698, p. 012 003, 2016.
- [172] X. Yu, Z. Yan, and A. V. Vasilakos, "A survey of verifiable computation," *Mobile Networks and Applications*, vol. 22, no. 3, pp. 438–453, 2017.
- [173] C. Ré and P. Beame. "Lecture 9 Interactive Proofs and Arthur-Merlin Games," School of computer science & engineering, University of Washington. (2004), [Online]. Available: <https://courses.cs.washington.edu/courses/cse532/04sp/lect09.pdf> (visited on 01/11/2024).
- [174] L. Babai, "Trading group theory for randomness," in *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, ser. STOC '85, Providence, Rhode Island, USA: Association for Computing Machinery, 1985, 421429.
- [175] O. Goldreich, *The Princeton Companion To Mathematics - IV20 Computational Complexity*. Princeton University Press, 2008, p. 583.
- [176] J. Watrous. "An introduction to quantum information and quantum circuits," University of Waterloo. (2011), [Online]. Available: <https://cs.uwaterloo.ca/~watrous/Papers/IntroductionQuantumCircuits.pdf> (visited on 01/11/2024).
- [177] M. Agrawal, N. Kayal, and N. Saxena, "Primes is in P," *Annals of Mathematics*, vol. 160, no. 2, pp. 781–793, 2004.
- [178] S. Arora and B. Barak, *Computational complexity. A modern approach*. English. Cambridge: Cambridge University Press, 2009.
- [179] J. Watrous, "Quantum computational complexity," in *Encyclopedia of Complexity and Systems Science*, R. A. Meyers, Ed. New York, NY: Springer New York, 2009, pp. 7174–7201.
- [180] W. J. Savitch, "Relationships between nondeterministic and deterministic tape complexities," *Journal of Computer and System Sciences*, vol. 4, no. 2, pp. 177–192, 1970.
- [181] S. Goldwasser, S. Micali, and C. Rackoff, "The knowledge complexity of interactive proof systems," *SIAM Journal on Computing*, vol. 18, no. 1, pp. 186–208, 1989.

- [182] T. Vidick, "From operator algebras to complexity theory and back," *Notices of the American Mathematical Society*, vol. 66, p. 1, 2019.
- [183] R. L. Rivest, A. Shamir, and L. Adleman, "A method for obtaining digital signatures and public-key cryptosystems," *Commun. ACM*, vol. 21, no. 2, p. 120–126, 1978.
- [184] C. Gidney and M. Ekerå, "How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits," *Quantum*, vol. 5, p. 433, 2021.
- [185] R. Jain, Z. Ji, S. Upadhyay, and J. Watrous, *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010, QIP = PSPACE*. ACM, 2010, p.573–582.
- [186] S. Aaronson, "Qip = pspace breakthrough: Technical perspective," *Commun. ACM*, vol. 53, no. 12, p. 101, 2010.
- [187] B. Schoenmakers, "Interactive proof," in *Encyclopedia of Cryptography and Security*, H. C. A. van Tilborg and S. Jajodia, Eds. Boston, MA: Springer US, 2011, pp. 619–619.
- [188] D. J. Bernstein, "Introduction to post-quantum cryptography," in *Post-Quantum Cryptography*, D. J. Bernstein, J. Buchmann, and E. Dahmen, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 1–14.
- [189] O. Regev, "On lattices, learning with errors, random linear codes, and cryptography," in *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*, ser. STOC '05, Baltimore, MD, USA: Association for Computing Machinery, 2005, p. 84–93.
- [190] C. Gauss and W. Waterhouse, *Disquisitiones Arithmeticae*. Springer-Verlag, 1986.
- [191] Z. Zheng, K. Tian, and F. Liu, "Lwe public key cryptosystem," in *Modern Cryptography Volume 2: A Classical Introduction to Informational and Mathematical Principle*. Singapore: Springer Nature Singapore, 2023, pp. 99–118.
- [192] G. Alagic, A. M. Childs, A. B. Grilo, and S.-H. Hung, "Non-interactive classical verification of quantum computation," in *Theory of Cryptography*, R. Pass and K. Pietrzak, Eds., Cham: Springer International Publishing, 2020, pp. 153–180.
- [193] T. Vidick and T. Zhang, "Classical zero-knowledge arguments for quantum computations," *Quantum*, vol. 4, p. 266, 2020.
- [194] G. D. Kahanamoku-Meyer, S. Choi, U. V. Vazirani, and N. Y. Yao, "Classically verifiable quantum advantage from a computational bell test," *Nature Physics*, vol. 18, no. 8, pp. 918–924, 2022.
- [195] D. Kozen, *Automata and Computability*. New York: Springer-Verlag, 1997.
- [196] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, *Quantum computation by adiabatic evolution*, 2000. arXiv: [quant-ph/0001106](https://arxiv.org/abs/quant-ph/0001106) [quant-ph].
- [197] J. Kempe, A. Kitaev, and O. Regev, "The complexity of the local hamiltonian problem," vol. 35, no. 5, pp. 1070–1097, 2006.

- [198] J. D. Biamonte and P. J. Love, "Realizable hamiltonians for universal adiabatic quantum computers," *Phys. Rev. A*, vol. 78, p. 012 352, 1 2008.
- [199] R. P. Feynman, "Quantum mechanical computers," *Foundations of Physics*, vol. 16, no. 6, pp. 507–531, 1986.
- [200] R. Stricker *et al.*, "Towards experimental classical verification of quantum computation," *Quantum Science and Technology*, vol. 9, no. 2, 02LT01, 2024.
- [201] T. Morimae, D. Nagaj, and N. Schuch, "Quantum proofs can be verified using only single-qubit measurements," *Phys. Rev. A*, vol. 93, p. 022 326, 2 2016.
- [202] R. Lindner, M. Rueckert, P. Baumann, and L. Nobach. "TU Darmstadt lattice challenge," TU Darmstadt. (2010), [Online]. Available: www.latticechallenge.org (visited on 01/11/2024).
- [203] A. Gheorghiu, M. J. Hoban, and E. Kashefi, "A simple protocol for fault tolerant verification of quantum computation," *Quantum Science and Technology*, vol. 4, no. 1, p. 015 009, 2018.
- [204] M. Grassl, T. Beth, and T. Pellizzari, "Codes for the quantum erasure channel," *Phys. Rev. A*, vol. 56, pp. 33–38, 1 1997.
- [205] N. C. Brown and K. R. Brown, "Comparing zeeman qubits to hyperfine qubits in the context of the surface code: $^{174}\text{Yb}^+$ And $^{171}\text{Yb}^+$," *Phys. Rev. A*, vol. 97, p. 052 301, 2018.
- [206] C.-Y. Lu *et al.*, "Experimental quantum coding against qubit loss error," *Proc. Natl. Acad. Sci.*, vol. 105, no. 32, pp. 11 050–11 054, 2008.
- [207] B. A. Bell *et al.*, "Experimental demonstration of a graph state quantum error-correction code," *Nature Communications*, vol. 5, no. 1, p. 3658, 2014.
- [208] S. Morley-Short *et al.*, "Physical-depth architectural requirements for generating universal photonic cluster states," *Quant. Sci. Tech.*, vol. 3, no. 1, p. 015 005, 2018.
- [209] D. Vodola, D. Amaro, M. A. Martin-Delgado, and M. Müller, "Twins percolation for qubit losses in topological color codes," *Phys. Rev. Lett.*, vol. 121, p. 060 501, 6 2018.
- [210] D. Niemietz, P. Farrera, S. Langenfeld, and G. Rempe, "Nondestructive detection of photonic qubits," *Nature*, vol. 591, no. 7851, p. 570, 2021.
- [211] F. Battistel, B. Varbanov, and B. Terhal, "Hardware-efficient leakage-reduction scheme for quantum error correction with superconducting transmon qubits," *PRX Quantum*, vol. 2, p. 030 314, 3 2021.
- [212] C. H. Bennett *et al.*, "Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels," *Phys. Rev. Lett.*, vol. 70, pp. 1895–1899, 13 1993.
- [213] T. C. Ralph, A. J. F. Hayes, and A. Gilchrist, "Loss-tolerant optical qubits," *Phys. Rev. Lett.*, vol. 95, p. 100 501, 10 2005.

- [214] E. A. Martinez, T. Monz, D. Nigg, P. Schindler, and R. Blatt, "Compiling quantum algorithms for architectures with multi-qubit gates," *New Journal of Physics*, vol. 18, no. 6, p. 063 029, 2016.
- [215] M. Müller *et al.*, "Iterative phase optimization of elementary quantum error correcting codes," *Phys. Rev. X*, vol. 6, p. 031 030, 3 2016.
- [216] S. Massar and S. Popescu, "Optimal extraction of information from finite quantum ensembles," *Phys. Rev. Lett.*, vol. 74, pp. 1259–1263, 8 1995.
- [217] A. Aharony and D. Stauffer, *Introduction To Percolation Theory*. Taylor & Francis, 2003.
- [218] R. Stricker *et al.*, "Experimental deterministic correction of qubit loss," *Nature*, vol. 585, no. 7824, pp. 207–210, 2020.
- [219] D. B. Hume, T. Rosenband, and D. J. Wineland, "High-fidelity adaptive qubit detection through repetitive quantum nondemolition measurements," *Phys. Rev. Lett.*, vol. 99, p. 120 502, 12 2007.
- [220] A. Steane, "Multiple-particle interference and quantum error correction," *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, vol. 452, no. 1954, pp. 2551–2577, 1996.
- [221] S. Bravyi and A. Kitaev, "Universal quantum computation with ideal clifford gates and noisy ancillas," *Phys. Rev. A*, vol. 71, p. 022 316, 2 2005.
- [222] M. E. Beverland, A. Kubica, and K. M. Svore, "Cost of universality: A comparative study of the overhead of state distillation and code switching with color codes," *PRX Quantum*, vol. 2, p. 020 341, 2 2021.
- [223] L. Postler *et al.*, "Experimental quantification of spatial correlations in quantum dynamics," *Quantum*, vol. 2, p. 90, 2018.
- [224] M. Meth *et al.*, "Probing phases of quantum matter with an ion-trap tensor-network quantum eigensolver," *Phys. Rev. X*, vol. 12, p. 041 035, 4 2022.
- [225] M. Ringbauer *et al.*, *Verifiable measurement-based quantum random sampling with trapped ions*, 2023. arXiv: [2307.14424](https://arxiv.org/abs/2307.14424) [quant-ph].