

Quantum Randomness

Conrad Strydom

*Thesis presented in partial fulfilment of the requirements for
the degree of Master of Science (Physics)
in the Faculty of Science at Stellenbosch University.*



Supervised by Prof. M S Tame and Dr. G W Bosman

March 2023

Declaration

By submitting this thesis electronically, I declare that the entirety of the work contained therein is my own, original work, that I am the sole author thereof (save to the extent explicitly otherwise stated), that reproduction and publication thereof by Stellenbosch University will not infringe any third party rights and that I have not previously in its entirety or in part submitted it for obtaining any qualification.

Date: March 2023

Abstract

Randomness is a vital resource with many important applications in information theory. In particular, random numbers play a ubiquitous role in cryptography, simulation and coordination in computer networks. When randomness is generated using classical techniques, the unpredictability relies on incomplete knowledge which can introduce ordered features and compromise the application. This thesis explores the use of quantum techniques to generate true randomness and its application to quantum computing.

The analogue of random numbers in quantum information are random unitary operators sampled from the uniform Haar ensemble, which are used in a number of quantum protocols. Unfortunately, these cannot be generated efficiently and so pseudorandom ensembles called unitary t -designs are frequently used as a substitute. In the first part of this thesis we investigate t -designs realised using a measurement-based approach on IBM quantum computers. In particular, we implement an exact single-qubit 3-design on IBM quantum computers by performing measurements on a 6-qubit graph state. We show that the ensemble of unitaries realised was a 1-design, but not a 2-design or a 3-design under the test conditions set, which we show to be a result of depolarising noise. We obtain improved results for the 2-design test by implementing an approximate 2-design, which uses a smaller 5-qubit graph state, but the test still does not pass for all states due to noise.

To obtain a theoretical understanding of the effect of noise on t -designs, we investigate the effect of various noise channels on the quality of single-qubit t -designs. We show analytically that the 1-design is affected only by amplitude damping, while numeric results obtained for the 2-, 3-, 4- and 5-design suggest that a $2t$ -design is significantly more sensitive to noise than a $(2t - 1)$ -design and that, with the exception of amplitude damping, a $(2t + 1)$ -design is as sensitive to noise as a $2t$ -design.

Next, we test our approximate measurement-based 2-design on an important application in quantum computing, namely noise estimation. For this, we propose an interleaved randomised benchmarking protocol for measurement-based quantum computers that can be used to estimate the fidelity of any single-qubit measurement-based gate. We demonstrate our protocol on IBM quantum computers by estimating the fidelity of a universal single-qubit gate set using graph states of up to 31 qubits. Estimated gate fidelities show good agreement with those calculated from process tomography, which shows that our approximate measurement-based 2-design is of sufficient quality for use in randomised benchmarking, despite not passing our test for all states.

While IBM quantum computers provide a sophisticated platform for randomness generation, they are not specifically designed for this task. We therefore investigate randomness generation on custom-built hardware, by integrating an on-chip nanowire waveguide into an optical time-of-arrival based quantum random number generation setup. Despite loss, we achieve a random number generation rate of 14.4 Mbits/s. The generated bits did not require any post-processing to pass industry standard tests. Our experiment demonstrates an order of magnitude increase in generation rate and decrease in device size compared to previous studies.

Opsomming

Willekeurigheid is 'n noodsaaklike hulpbron met vele belangrike toepassings in inligtingsteorie. Lukrake getalle speel veral 'n alomteenwoordige rol in kriptografie, simulاسie en koördinasie in rekenaarnetwerke. Sonder willekeurigheid deur klassieke tegnieke gegenereer word, steun die onvoorspelbaarheid op onvolledige kennis wat geordende kenmerke kan meebring en die toepassing skaad. Hierdie tesis ondersoek die gebruik van kwantum tegnieke om ware willekeurigheid te skep, asook die toepassing daarvan in kwantum berekeninge.

In kwantum inligting stem lukrake getalle ooreen met lukrake eenheidsoperatore gesteekeproef vanuit die eenvormige Haar-samevatting wat in 'n aantal kwantum protokolle gebruik word. Ongelukkig kan hierdie nie doeltreffend geskep word nie en dus word pseudo-lukrake samestellings genaamd eenheids- t -ontwerpe, gereeld as plaasvervangers gebruik. In die eerste gedeelte van hierdie tesis ondersoek ons t -ontwerpe gerealiseer deur 'n metings-gebaseerde benadering op IBM kwantum rekenaars. In besonder, implementeer ons 'n presiese enkel-kubit 3-ontwerp op IBM kwantum rekenaars deur metings uit te voer op 'n 6-kubit grafiektoestand. Ons toon dat die samestelling van eenheidsoperatore wat gerealiseer is, 'n 1-ontwerp, maar nie 'n 2-ontwerp of 3-ontwerp onder die vasgestelde toetstoestande is nie, vanweë depolariserende geraas. Ons bekom verbeterde resultate vir die 2-ontwerp toets deur implementering van 'n benaderde 2-ontwerp, wat 'n kleiner 5-kubit grafiektoestand benut, maar die toets slaag weens geraas steeds nie vir al die toestande nie.

Om 'n teoretiese begrip van die invloed van geraas op t -ontwerpe te verkry, ondersoek ons die invloed van verskeie geraaskanale op die kwaliteit van enkel-kubit t -ontwerpe. Ons toon analities dat die 1-ontwerp slegs deur amplitude damping geraak word, terwyl numeriese resultate verkry vir die 2-, 3-, 4- en 5-ontwerp aandui dat 'n $2t$ -ontwerp aansienlik meer sensitief is vir geraas as 'n $(2t - 1)$ -ontwerp en dat, met uitsondering van amplitude damping, 'n $(2t + 1)$ -ontwerp ewe sensitief is vir geraas as 'n $2t$ -ontwerp.

Ons toets volgende ons benaderde metings-gebaseerde 2-ontwerp op 'n belangrike toepassing in kwantum berekeninge, naamlik geraasskatting. Hiervoor stel ons 'n tussengelaagde lukrake peilmerking protokol vir metings-gebaseerde kwantum rekenaars voor wat gebruik kan word om die betroubaarheid van enige enkel-kubit metings-gebaseerde poort te skat. Ons demonstreer ons protokol op die IBM kwantum rekenaars deur die betroubaarheid van 'n universele enkel-kubit poort-stel te skat deur gebruik te maak van grafiektoestande van tot 31 kubits. Geskatte poort-betroubaarhede toon goeie ooreenkoms met dié bereken van proesestomografie, wat toon dat ons benaderde metings-gebaseerde 2-ontwerp van voldoende gehalte is vir gebruik in lukrake peilmerking, alhoewel dit nie ons toets vir alle toestande slaag nie.

Alhoewel IBM kwantum rekenaars 'n gesofistikeerde platform vir die skepping van willekeurigheid verskaf, is hulle nie spesifiek vir hierdie doel ontwerp nie. Ons ondersoek dus die skepping van willekeurigheid op doelgerigte hardware, deur 'n op-skyfie nanodraad golfgids te integreer met 'n optiese aankomstyd-gebaseerde kwantum lukrake getal-skeppende stelsel. Ondanks verlies, behaal ons 'n lukrake getalskeppingstempo van 14.4 Mbits/s. Die gegenereerde binêre syfers het geen verdere verwerking vereis om aan nywerheidsstandaarde te voldoen nie. Ons eksperiment demonstreer 'n grootte-orde toename in die skeppingstempo en afname in toestel-grootte vergeleke met vorige studies.

Acknowledgements

I acknowledge the use of IBM Quantum services for the work in Chapters 3 and 5. The views expressed in this thesis are those of the author, and do not reflect the official policy or position of IBM or the IBM Quantum team. I thank Taariq Surtee and Barry Dwolatzky at the University of Witwatersrand and Ismail Akhalwaya at IBM Research Africa for access to the IBM processors through the Q Network and the Africa Research Universities Alliance.

I would like to extend my most sincere gratitude to my supervisor, Prof. Mark Tame. Throughout the past two years he has given invaluable guidance on obtaining theoretical results, designing experiments, analysing and interpreting results and drafting manuscripts for publication in peer-reviewed journals. Above all, I thank Prof. Tame for the many hours spent in the laboratory to assist me in building my first optical experimental setup. I also thank Dr. Gurthwin Bosman for co-supervising me on this thesis.

I humbly thank Prof. Şahin Özdemir and Dr. Sina Soleymani, from Pennsylvania State University in the United States of America, for their willingness to collaborate with Prof. Tame and me on the work presented in Chapter 6. Dr. Soleymani fabricated the on-chip nanowire plasmonic waveguide used in the experiments and Prof. Özdemir provided valuable comments and insights. I also thank Dr. Damian Markham and Dario Trincherro for valuable comments and insights on the work in Chapters 3 and 4 respectively.

Furthermore, I thank Marié Theron, the Science Faculty librarian at Stellenbosch University, who assisted in finding many of the sources cited in this thesis. I also extend a special thank you to Analize Jooste, who assisted in translating the abstract to Afrikaans, Ilse Strydom, who assisted in the search for spelling and grammatical errors, and Emma King, who assisted in designing the template for this thesis.

Last, but certainly not least, I will be eternally grateful for the funding granted through the Stellenbosch University Postgraduate Scholarship Programme, without which none of this work would have been possible.

Publications

Some of the work in this thesis has been published or submitted to various journals for publication. This includes:

- The work on the implementation of single-qubit measurement-based t -designs using IBM processors in Chapter 3 is adapted from an article in the international journal Nature Scientific Reports [*Sci. Rep.* **12**, 5014 (2022)]. Prof. Mark Tame, my supervisor, is a co-author of this article. I hereby acknowledge that the work in Chapter 3 is an extension of a research project which I submitted to Stellenbosch University in 2020 for the purpose of obtaining my Bachelor of Science Honours (Physics). However, all the experimental and numerical results presented in Chapter 3 were newly generated while completing my Master of Science (Physics) and none of these results were previously submitted for the purpose of obtaining a qualification.
- The work investigating the effect of noise channels on the quality of unitary t -designs in Chapter 4 is adapted from a preprint [arXiv:2203.13771] of which Prof. Mark Tame is a co-author.
- The work on measurement-based interleaved randomised benchmarking using IBM processors in Chapter 5 is adapted from an article in the international journal Physica Scripta [*Phys. Scr.* **98**, 025106 (2023)] of which Prof. Mark Tame is a co-author.
- The work on quantum random number generation using an on-chip nanowire plasmonic waveguide in Chapter 6 is adapted from a preprint [arXiv:2306.13490] of which Dr. Sina Soleymani and Prof. Şahin Özdemir, from Pennsylvania State University in the United States of America, as well as Prof. Mark Tame are co-authors. Dr. Soleymani fabricated the on-chip nanowire plasmonic waveguide used in the experiments. Prof. Özdemir and Prof. Tame contributed to the analysis of the results and the discussions and interpretations. Prof. Mark Tame also wrote portions of the Appendix E.

I am the primary author of all published work included in this thesis. I performed all the experiments and calculations for all the experimental and theoretical results in all the publications. I also wrote and ran the code which generated all the numerical results and wrote the first draft of each publication. However, all the authors of the respective publications conceived the idea, designed the experiments, analysed the results and contributed to the discussions and interpretations.

Contents

Declaration	i
Abstract	ii
Opsomming	iii
Acknowledgements	iv
Publications	v
Contents	vi
List of Figures	x
List of Tables	xv
1 Introduction	1
1.1 Quantum information	1
1.2 Quantum computing	1
1.3 Quantum randomness	2
1.3.1 Random numbers	2
1.3.2 Random unitary operators	2
1.4 Randomness generation on quantum hardware	3
1.4.1 IBM's cloud-based superconducting hardware	3
1.4.2 Custom-built on-chip plasmonic hardware	5
1.5 Thesis outline	5
2 Background	7
2.1 Quantum information processing	7
2.1.1 The qubit	7
2.1.2 Unitary evolution	8
2.1.3 Projective measurements	9
2.1.4 Multi-qubit systems	10
2.2 Quantum information processing in the presence of noise	11
2.2.1 The density operator formulation	11
2.2.2 Noise channels	12
2.3 Quantum computing	13

2.3.1	The circuit model	13
2.3.2	Measurement-based quantum computing	13
2.4	IBM's cloud-based superconducting hardware	15
2.5	Quantum randomness	15
3	Implementation of single-qubit measurement-based t-designs using IBM processors	17
3.1	Introduction	17
3.2	Background	17
3.2.1	Measurement-based t-designs	17
3.2.2	Quantum state tomography	18
3.2.3	Quantum process tomography	18
3.2.4	Quantum readout error mitigation	19
3.3	Experiments	20
3.3.1	Implementation	20
3.3.2	Process tomography results	22
3.3.3	Relative frequencies	22
3.3.4	Testing for a t-design	24
3.3.5	Approximate 2-design	27
3.4	Conclusion	29
4	Investigating the effect of noise channels on the quality of unitary t-designs	31
4.1	Introduction	31
4.2	Background	32
4.2.1	Noise channels	32
4.2.1.1	Flip channels	32
4.2.1.2	Phase damping channel	32
4.2.1.3	Amplitude damping channel	33
4.2.1.4	Depolarising noise channel	33
4.3	Noise modelling	33
4.4	Results	34
4.4.1	Implementation	35
4.4.2	Numeric results	35
4.4.2.1	Flip channels	36
4.4.2.2	Phase damping channel	37
4.4.2.3	Amplitude damping channel	38
4.4.2.4	Depolarising noise channel	39
4.5	Conclusion	40
5	Measurement-based interleaved randomised benchmarking using IBM processors	42
5.1	Introduction	42
5.2	Background	44
5.2.1	Measurement-based quantum computing	44
5.2.2	Measurement-based t-designs	44
5.2.3	Interleaved randomised benchmarking	44

5.3	Measurement-based interleaved randomised benchmarking protocol	46
5.4	Adjustments to protocol for implementation on IBM processors	47
5.5	Experiments	48
5.5.1	Implementation of protocol	48
5.5.2	Results for universal gates	49
5.5.3	Noisier gates	51
5.6	Conclusion	53
6	Quantum random number generation using an on-chip nanowire plasmonic waveguide	56
6.1	Introduction	56
6.2	Experimental setup	56
6.3	Results	60
6.4	Conclusion	63
7	Conclusion	64
	Appendices	67
A	Implementation of single-qubit measurement-based t-designs using IBM processors	68
A.1	Qubits used for the 3-design	68
A.2	Depolarising noise	69
A.3	Identity implementation	71
A.4	Qubits used for the 2-design	73
A.5	Qubits used for the identity	75
B	Investigating the effect of noise channels on the quality of unitary t-designs	76
B.1	Proof of well-definedness of noise models	76
B.1.1	Proof for the model where noise is applied before the unitary operations	76
B.1.2	Proof for the model where noise is applied after the unitary operations	77
B.2	Analytic results for the 1-design	77
B.2.1	Analytic results for the model where noise is applied before the unitary operations	77
B.2.2	Analytic results for the model where noise is applied after the unitary operations	78
B.3	Numeric results for the model where noise is applied after the unitary operations	78
B.3.1	Flip channels	78
B.3.2	Phase damping channel	79
B.3.3	Amplitude damping channel	80
B.3.4	Depolarising noise channel	80
B.4	Proof of equivalence of noise models for the depolarising noise channel	80
C	State dependence of the effect of noise channels on the quality of single-qubit t-designs	81
C.1	Visualisation of regions of acceptable quality	81
C.1.1	Implementation	81
C.1.2	Numeric results for the model where noise is applied before the unitary operations	81
C.1.2.1	Bit flip channel	81
C.1.2.2	Phase flip channel	82

C.1.2.3	Bit and phase flip channel	83
C.1.2.4	Phase damping channel	83
C.1.2.5	Amplitude damping channel	83
C.1.2.6	Depolarising noise channel	84
C.1.3	Numeric results for the model where noise is applied after the unitary operations . . .	85
C.1.3.1	Flip channels	85
C.1.3.2	Phase damping channel	85
C.1.3.3	Amplitude damping channel	86
C.1.3.4	Depolarising noise channel	86
C.2	Numeric results for different truncations of the polar angle	87
C.2.1	Numeric results for the model where noise is applied before the unitary operations . .	87
C.2.1.1	Bit flip channel	87
C.2.1.2	Phase flip channel	87
C.2.1.3	Bit and phase flip channel	88
C.2.1.4	Phase damping channel	88
C.2.1.5	Amplitude damping channel	89
C.2.1.6	Depolarising noise channel	89
C.2.2	Numeric results for the model where noise is applied after the unitary operations . . .	89
C.3	Numeric results for different truncations of the azimuthal angle	90
C.3.1	Numeric results for the model where noise is applied before the unitary operations . .	90
C.3.2	Numeric results for the model where noise is applied after the unitary operations . . .	91
D	Measurement-based interleaved randomised benchmarking using IBM processors	92
D.1	Qubits used for fidelity estimation of universal gates	92
D.2	Fitting procedure	93
D.3	Process tomography	94
D.4	Qubits used for fidelity estimation of noisier gates	96
E	Quantum random number generation using an on-chip nanowire plasmonic waveguide	97
E.1	Power transmission factor of nanowire plasmonic waveguide	97
E.2	Polarisation dependence of photon detection rate	97
E.3	Proof of uniformity for time-of-arrival scheme	98
E.4	Higher order photon events and correction factor	99
	Bibliography	101

List of Figures

2.1	The Bloch sphere for visualising the state of a qubit.	8
2.2	Quantum circuit for preparing the 3-qubit linear cluster state.	13
2.3	Measurement-based processing with a n -qubit linear cluster state. Step 1 depicts the initialisation of the qubits. Step 2 depicts the entangled cluster state (after controlled phase gates have been applied to neighbouring qubits) in addition to the measurements performed on each qubit. Step 3 depicts the state resulting from the measurements.	14
3.1	General quantum circuit for implementation of the exact measurement-based 3-design proposed by Turner and Markham [88] on the <i>ibmq_toronto</i> quantum processor. Here ‘in’ represents the set of gates applied to construct the input state and ‘out’ represents the set of gates applied and measurements done to perform tomography on the sixth qubit. The angles for the z -rotation gates are $\phi_2 = \phi_4 = \frac{\pi}{4}$ and $\phi_3 = \arccos \sqrt{1/3}$	20
3.2	Example of process tomography results for the random unitary that agreed least with the ideal case generated for measurement outcome $\mathbf{m} = 00000$ with the exact 3-design on the <i>ibmq_toronto</i> quantum processor. The diagram at the top shows the entangled 6-qubit linear cluster state with the measurements performed on each qubit. The χ matrix obtained without quantum readout error mitigation is shown on the left, the χ matrix obtained with quantum readout error mitigation is shown in the middle and the ideal χ matrix is shown on the right. The real part of each matrix is shown above and the imaginary part of each matrix is shown below.	21
3.3	Example of process tomography results for the random unitary that agreed most with the ideal case generated for measurement outcome $\mathbf{m} = 11001$ with the exact 3-design on the <i>ibmq_toronto</i> quantum processor. The diagram at the top shows the entangled 6-qubit linear cluster state with the measurements performed on each qubit. The χ matrix obtained without quantum readout error mitigation is shown on the left, the χ matrix obtained with quantum readout error mitigation is shown in the middle and the ideal χ matrix is shown on the right. The real part of each matrix is shown above and the imaginary part of each matrix is shown below.	22
3.4	Distribution of channel fidelities for the 32 random unitaries generated with the exact 3-design on the <i>ibmq_toronto</i> quantum processor. (a) Raw shows the distribution without quantum readout error mitigation. (b) Processed shows the distribution with quantum readout error mitigation.	24
3.5	Distribution of relative frequencies with which the 32 random unitaries are generated with the exact 3-design on the <i>ibmq_toronto</i> quantum processor. (a) Raw shows the distribution without quantum readout error mitigation. (b) Processed shows the distribution with quantum readout error mitigation.	24

4.1	Effect of the bit flip channel (see Eq. (4.1)) on the quality of the (a) 2-design and (b) 4-design for the model where noise is applied before the unitary operations, for different truncation radii r_t	36
4.2	ϵ versus t for the bit flip channel with $p = 0.5$ (see Eq. (4.1)) for the model where noise is applied before the unitary operations, for different truncation radii r_t	36
4.3	Effect of the phase damping channel (see Eq. (4.4)) on the quality of the (a) 2-design and (b) 4-design for the model where noise is applied before the unitary operations, for different truncation radii r_t	38
4.4	ϵ versus t for the phase damping channel with $\lambda = 0.5$ (see Eq. (4.4)) for the model where noise is applied before the unitary operations, for different truncation radii r_t	38
4.5	Effect of the amplitude damping channel (see Eq. (4.8)) on the quality of the 2-design for the model where noise is applied before the unitary operations, for different truncation radii r_t . The full set of results is shown on the left and the region in which the anomaly occurs is shown enlarged on the right.	39
4.6	Effect of the amplitude damping channel (see Eq. (4.8)) on the quality of the 4-design for the model where noise is applied before the unitary operations, for different truncation radii r_t . The full set of results is shown on the left and the region in which the anomaly occurs is shown enlarged on the right.	39
4.7	ϵ versus t for the amplitude damping channel, with the parameter λ (see Eq. (4.8)) taken to be the turning point of ϵ versus λ , for the model where noise is applied before the unitary operations, for different truncation radii r_t	40
4.8	Effect of the depolarising noise channel (see Eq. (2.27)) on the quality of the (a) 2-design and (b) 4-design for the model where noise is applied before the unitary operations, for different truncation radii r_t	40
4.9	ϵ versus t for the depolarising noise channel with $p = 0.5$ (see Eq. (2.27)) for the model where noise is applied before the unitary operations, for different truncation radii r_t	41
5.1	Qubit topologies of the processors used in our demonstration, with the qubits used shaded grey.	43
5.2	Quantum circuit for the implementation of the interleaved sequence for the 3-qubit T gate with $m = 1$ on the <i>ibm_hanoi</i> quantum processor. The angle $\phi = \frac{\pi}{4}$ represents the different basis measurements used and ‘out’ is the final qubit on which quantum state tomography is performed.	49
5.3	Sequence fidelities $F(m)$ obtained for $m \in \{1, 2, 3\}$ to estimate the fidelity of the 2-qubit measurement-based implementation of the Hadamard gate and the 3-qubit measurement-based implementation of the T gate, using the 5-qubit approximate measurement-based 2-design proposed in Sec. 3.3.5, on the <i>ibm_hanoi</i> quantum processor. Reference and interleaved sequence fidelities were fit to Eqs. (5.3) and (5.4) respectively (see Appendix D.2). (a) Hadamard gate shows the reference sequence fidelities (in blue) and the interleaved sequence fidelities for the 2-qubit Hadamard gate (in red). (b) T gate shows the same reference sequence fidelities (in blue) and the interleaved sequence fidelities for the 3-qubit T gate (in red).	50

5.4	Sequence fidelities $F(m)$ obtained for $m \in \{1, 2, 3\}$ to estimate the fidelity of the 4-qubit measurement-based implementation of the Hadamard gate and the 5-qubit measurement-based implementation of the T gate, using the 5-qubit approximate measurement-based 2-design proposed in Sec. 3.3.5, on the <i>ibmq_brooklyn</i> quantum processor. Reference and interleaved sequence fidelities were fit to Eqs. (5.3) and (5.4) respectively (see Appendix D.2). (a) Hadamard gate shows the reference sequence fidelities (in blue) and the interleaved sequence fidelities for the 4-qubit Hadamard gate (in red). (b) T gate shows the same reference sequence fidelities (in blue) and the interleaved sequence fidelities for the 5-qubit T gate (in red).	52
5.5	Sequence fidelities $F(m)$ obtained for $m \in \{1, 2, 3\}$ to estimate the fidelity of the 6-qubit measurement-based implementation of the Hadamard gate and the 7-qubit measurement-based implementation of the T gate, using the 5-qubit approximate measurement-based 2-design proposed in Sec. 3.3.5, on the <i>ibmq_brooklyn</i> quantum processor. Reference and interleaved sequence fidelities were fit to Eqs. (5.3) and (5.4) respectively (see Appendix D.2). (a) Hadamard gate shows the reference sequence fidelities from Fig. 5.4 (in blue) and the interleaved sequence fidelities for the 6-qubit Hadamard gate (in red). (b) T gate shows the same reference sequence fidelities from Fig. 5.4 (in blue) and the interleaved sequence fidelities for the 7-qubit T gate (in red).	53
6.1	Quantum random number generation using an on-chip nanowire plasmonic waveguide. (a) Experimental Setup shows the experimental setup used to investigate time-of-arrival based quantum random number generation using an on-chip nanowire plasmonic waveguide. The labels used are: single-mode optical fibre (SM), beam expander (BE), neutral density filter (NDF), half-wave plate (HWP), quarter-wave plate (QWP), polarising beamsplitter (PBS), diffraction-limited microscope (DLM), fibre coupler (FC) and multi-mode optical fibre (MM). (b) Nanowire Plasmonic Waveguide shows a top view of the on-chip nanowire plasmonic waveguide used in the experiments. (c) Time-of-arrival Scheme illustrates the implemented variation of the time-of-arrival scheme, in which random numbers are obtained from the arrival times of photons relative to an external time reference [53]. (d) Atomic Force Microscope Images shows atomic force microscope (AFM) images of the fabricated on-chip nanowire plasmonic waveguide. These include an AFM image of the entire nanowire plasmonic waveguide (left), an AFM image of the top tapering (top centre), an AFM height profile of the top tapering (bottom centre), an AFM image of the top grating (top right) and an AFM height profile of the top grating (bottom right).	57
6.2	Pearson correlation coefficient of the generated sample with 1-bit to 15-bit delays of itself. . .	59
6.3	Pearson correlation coefficient of the (a) raw sample and (b) shuffled sample with 1-bit to 15-bit delays of itself.	61
A.1	Qubit topology of the <i>ibmq_toronto</i> quantum processor. The connecting lines between qubits indicate the qubit pairs for which the CX gate is supported at the hardware level. The qubits used for the exact 3-design implementation are shaded grey.	68
A.2	ϵ versus p for the 2-design test with the middle term in inequality (3.6) replaced by $\mathbb{E}_H^2((\epsilon(\rho))^{\otimes 2})$, for different truncation radii r_t . A plot of the full set of results is shown above and a plot focused on the region of interest is shown below.	70

A.3	ϵ versus p for the 2-design test with the middle term in inequality (3.6) calculated by repeatedly applying depolarising noise to a state, in between the five individual unitary operations that are applied to the input state in the exact 3-design described in Sec. 3.2.1, for different truncation radii r_t . A plot of the full set of results is shown above and a plot focused on the region of interest is shown below.	71
A.4	Process tomography results (average χ matrix) for the implementation of the identity operation with a 3-qubit linear cluster state on the <i>ibmq_toronto</i> quantum processor. The diagram at the top shows the entangled 3-qubit linear cluster state with the measurements performed on each qubit. The χ matrix obtained without quantum readout error mitigation is shown on the left, the χ matrix obtained with quantum readout error mitigation is shown in the middle and the ideal χ matrix is shown on the right. The real part of each matrix is shown above and the imaginary part of each matrix is shown below.	72
A.5	Process tomography results (average χ matrix) for the implementation of the identity operation with a 5-qubit linear cluster state on the <i>ibmq_toronto</i> quantum processor. The diagram at the top shows the entangled 5-qubit linear cluster state with the measurements performed on each qubit. The χ matrix obtained without quantum readout error mitigation is shown on the left, the χ matrix obtained with quantum readout error mitigation is shown in the middle and the ideal χ matrix is shown on the right. The real part of each matrix is shown above and the imaginary part of each matrix is shown below.	73
A.6	Process tomography results (average χ matrix) for the implementation of the identity operation with a 7-qubit linear cluster state on the <i>ibmq_toronto</i> quantum processor. The diagram at the top shows the entangled 7-qubit linear cluster state with the measurements performed on each qubit. The χ matrix obtained without quantum readout error mitigation is shown on the left, the χ matrix obtained with quantum readout error mitigation is shown in the middle and the ideal χ matrix is shown on the right. The real part of each matrix is shown above and the imaginary part of each matrix is shown below.	74
B.1	Effect of the (a) bit flip channel (see Eq. (4.1)) and (b) phase damping channel (see Eq. (4.4)) on the quality of the 2-design for the model where noise is applied after the unitary operations, for different truncation radii r_t	79
B.2	Effect of the amplitude damping channel (see Eq. (4.8)) on the quality of the (a) 2-design and (b) 3-design for the model where noise is applied after the unitary operations, for different truncation radii r_t	79
C.1	Regions of acceptable quality for the 2-design affected by the bit flip channel (see Eq. (4.1)) for the model where noise is applied before the unitary operations, for different p	82
C.2	Regions of acceptable quality for the 4-design affected by the bit flip channel (see Eq. (4.1)) for the model where noise is applied before the unitary operations, for different p	82
C.3	Regions of acceptable quality for the 2-design affected by the phase flip channel (see Eq. (4.2)) for the model where noise is applied before the unitary operations, for different p	83
C.4	Regions of acceptable quality for the 4-design affected by the phase flip channel (see Eq. (4.2)) for the model where noise is applied before the unitary operations, for different p	83
C.5	Regions of acceptable quality for the 2-design affected by the bit and phase flip channel (see Eq. (4.3)) for the model where noise is applied before the unitary operations, for different p	84

C.6	Regions of acceptable quality for the 4-design affected by the bit and phase flip channel (see Eq. (4.3)) for the model where noise is applied before the unitary operations, for different p .	84
C.7	Regions of acceptable quality for the 2-design affected by the phase damping channel (see Eq. (4.4)) for the model where noise is applied before the unitary operations, for different λ .	85
C.8	Regions of acceptable quality for the 4-design affected by the phase damping channel (see Eq. (4.4)) for the model where noise is applied before the unitary operations, for different λ .	85
C.9	Regions of acceptable quality for the 2-design affected by the amplitude damping channel (see Eq. (4.8)) for the model where noise is applied before the unitary operations, for different λ .	86
C.10	Regions of acceptable quality for the 4-design affected by the amplitude damping channel (see Eq. (4.8)) for the model where noise is applied before the unitary operations, for different λ .	86
C.11	Regions of acceptable quality for the 2-design affected by the depolarising noise channel (see Eq. (2.27)) for the model where noise is applied before the unitary operations, for different p .	87
C.12	Regions of acceptable quality for the 4-design affected by the depolarising noise channel (see Eq. (2.27)) for the model where noise is applied before the unitary operations, for different p .	87
C.13	Regions of acceptable quality for the 2-design affected by the bit flip channel (see Eq. (4.1)) for the model where noise is applied after the unitary operations, for different p .	88
C.14	Regions of acceptable quality for the 2-design affected by the phase damping channel (see Eq. (4.4)) for the model where noise is applied after the unitary operations, for different λ .	88
C.15	Regions of acceptable quality for the 2-design affected by the amplitude damping channel (see Eq. (4.8)) for the model where noise is applied after the unitary operations, for different λ .	89
C.16	Regions of acceptable quality for the 3-design affected by the amplitude damping channel (see Eq. (4.8)) for the model where noise is applied after the unitary operations, for different λ .	89
C.17	Effect of the (a) phase flip channel (see Eq. (4.2)) and (b) phase damping channel (see Eq. (4.4)) on the quality of the 2-design for the model where noise is applied before the unitary operations, for different truncations of the polar angle θ_t , for a fixed truncation radius of $r_t = 0.95$.	90
C.18	Effect of the amplitude damping channel (see Eq. (4.8)) on the quality of the 2-design for the model where noise is applied before the unitary operations, for different truncations of the polar angle θ_t , for a fixed truncation radius of $r_t = 0.95$. The full set of results is shown on the left and a zoomed-in region is shown enlarged on the right.	90
D.1	Qubit topology of the <i>ibm_hanoi</i> quantum processor. The connecting lines between qubits indicate the qubit pairs for which the CX gate is supported at the hardware level. The qubits used in the implementations are shaded grey.	92
D.2	Qubit topology of the <i>ibmq_brooklyn</i> quantum processor. The connecting lines between qubits indicate the qubit pairs for which the CX gate is supported at the hardware level. The qubits used in the implementations are shaded grey.	94
E.1	Photon detection rate versus waveplate angle. Data points are an average of five repetitions and the errors are given by the standard deviation.	98

List of Tables

3.1	Channel fidelities for the 32 random unitaries, corresponding to the 32 different measurement outcomes, generated with the exact 3-design on the <i>ibmq_toronto</i> quantum processor. ‘Raw’ shows the channel fidelities without quantum readout error mitigation. ‘Processed’ shows the channel fidelities with quantum readout error mitigation.	23
3.2	Relative frequencies with which the 32 random unitaries, corresponding to the 32 different measurement outcomes, are generated with the exact 3-design on the <i>ibmq_toronto</i> quantum processor. ‘Raw’ shows the relative frequencies without quantum readout error mitigation. ‘Processed’ shows the relative frequencies with quantum readout error mitigation.	25
3.3	Summary of test results for the ensemble of unitaries generated using the <i>ibmq_toronto</i> quantum processor. ‘Raw’ shows the results without quantum readout error mitigation. ‘Processed’ shows the results with quantum readout error mitigation. ‘Radius’ is the truncation radius considered for a test. The column with ‘uniform’ shows the values of ϵ obtained when replacing the experimentally determined relative frequencies with uniform probabilities.	26
3.4	Fraction of states for which the ensemble of unitaries generated using the <i>ibmq_toronto</i> quantum processor passed the different tests. ‘Raw’ shows the fractions without quantum readout error mitigation. ‘Processed’ shows the fractions with quantum readout error mitigation. The column with ‘uniform’ shows the fractions obtained when replacing the experimentally determined relative frequencies with uniform probabilities.	27
3.5	Channel fidelities for the 16 random unitaries, corresponding to the 16 different measurement outcomes, generated with the approximate 2-design on the <i>ibmq_sydney</i> quantum processor. ‘Raw’ shows the channel fidelities without quantum readout error mitigation. ‘Processed’ shows the channel fidelities with quantum readout error mitigation.	28
3.6	Relative frequencies with which the 16 random unitaries, corresponding to the 16 different measurement outcomes, are generated with the approximate 2-design on the <i>ibmq_sydney</i> quantum processor. ‘Raw’ shows the relative frequencies without quantum readout error mitigation. ‘Processed’ shows the relative frequencies with quantum readout error mitigation.	29
3.7	Summary of test results for the ensemble of unitaries generated using the <i>ibmq_sydney</i> quantum processor. ‘Raw’ shows the results without quantum readout error mitigation. ‘Processed’ shows the results with quantum readout error mitigation. ‘Radius’ is the truncation radius considered for a test. The column with ‘uniform’ shows the values of ϵ obtained when replacing the experimentally determined relative frequencies with uniform probabilities. The column with ‘ideal’ shows the expected values of ϵ for the approximate 2-design for the truncation radii considered.	29

3.8	Fraction of states for which the ensemble of unitaries generated using the <i>ibmq_sydney</i> quantum processor passed the different tests. ‘Raw’ shows the fractions without quantum readout error mitigation. ‘Processed’ shows the fractions with quantum readout error mitigation. The column with ‘uniform’ shows the fractions obtained when replacing the experimentally determined relative frequencies with uniform probabilities.	30
5.1	Haar-averaged gate fidelities obtained for the different measurement-based gates on the <i>ibmq_brooklyn</i> quantum processor. ‘Est Fidelity’ shows the fidelity estimated using our measurement-based interleaved randomised benchmarking protocol. ‘Fidelity Range’ shows the range of fidelities calculated from process tomography results obtained for the three different sets of qubits used in the interleaved sequences (see Appendix D.3).	54
6.1	ENT Statistical Test Suite results for the generated sample. ‘Generated’ shows the values obtained using the generated sample. ‘Expected’ shows the expected values for a true random sample.	60
6.2	NIST Statistical Test Suite results for the generated sample. ‘Req’ shows the minimum number of sequences which need to pass a test for the sample to pass the test. ‘Prop’ shows the number of sequences of the generated sample which passed each test. For tests which involve more than five subtests (marked with *) the median of the results is presented.	60
6.3	ENT Statistical Test Suite results for the raw sample and the shuffled sample. ‘Raw’ shows the values obtained using the raw sample. ‘Shuffled’ shows the values obtained using the shuffled sample. ‘Expected’ shows the expected values for a true random sample.	62
6.4	NIST Statistical Test Suite results for the raw sample and the shuffled sample. ‘Req’ shows the minimum number of sequences which need to pass a test for the samples to pass the test. ‘Prop’ shows the number of sequences of the raw sample or the shuffled sample which passed each test. For tests which involve more than five subtests (marked with *) the median of the results is presented.	62
A.1	Calibration information for the <i>ibmq_toronto</i> quantum processor as obtained at the time of running the circuits for the exact 3-design. The single-qubit calibration information for the relevant qubits is shown on the left. T_1 and T_2 are the amplitude and phase damping time constants respectively of the qubits. The CX error rates for relevant qubit pairs are shown on the right.	69
A.2	Values of p inferred for implementations of the identity with different linear cluster states on the <i>ibmq_toronto</i> quantum processor. ‘Raw’ shows the values of p without quantum readout error mitigation. ‘Processed’ shows the values of p with quantum readout error mitigation. . .	74
A.3	Calibration information for the <i>ibmq_sydney</i> quantum processor as obtained at the time of running the circuits for the approximate 2-design. The single-qubit calibration information for the relevant qubits is shown on the left. T_1 and T_2 are the amplitude and phase damping time constants respectively of the qubits. The CX error rates for relevant qubit pairs are shown on the right.	75

A.4	Calibration information for the <i>ibmq_toronto</i> quantum processor as obtained at the time of running the circuits for the identity implementation. The single-qubit calibration information for the relevant qubits is shown on the left. T_1 and T_2 are the amplitude and phase damping time constants respectively of the qubits. The CX error rates for relevant qubit pairs are shown on the right.	75
D.1	Calibration information for the <i>ibmq_hanoi</i> quantum processor averaged over the time period during which circuits were run and data was obtained. The single-qubit calibration information for the relevant qubits is shown on the left. T_1 and T_2 are the amplitude and phase damping time constants respectively of the qubits. The CX error rates for relevant qubit pairs are shown on the right.	93
D.2	Calibration information for the <i>ibmq_brooklyn</i> quantum processor averaged over the time period during which circuits were run and data was obtained. The single-qubit calibration information for the relevant qubits is shown on the left. T_1 and T_2 are the amplitude and phase damping time constants respectively of the qubits. The CX error rates for relevant qubit pairs are shown on the right.	95

Chapter 1

Introduction

1.1 Quantum information

The field of quantum information science is a reconciliation of two revolutionary theories which emerged in the twentieth century, namely information theory, which is the mathematical theory describing the processing, storing and transfer of information, and quantum mechanics, which is the physical theory describing the interaction of matter at the atomic and subatomic scales. Quantum information science comprises quantum computing [1], quantum communication [2] and quantum sensing [3]. In quantum computing, the non-classical features of quantum mechanics, such as superposition and entanglement, are employed to speed up a variety of information processing tasks. In quantum communication, the inherent randomness of quantum mechanics is exploited to enable the provably secure transfer of information, while in quantum sensing, quantum properties such as entanglement are used to improve the precision of measurements. Quantum computing is discussed further in the next section as it is the main theme of quantum information science that this thesis is concerned with.

1.2 Quantum computing

Quantum computers employ the non-classical features of quantum mechanics to substantially speed up certain computational tasks such as prime factorisation [4], unstructured searching [5], simulating many body systems [6], machine learning [7] and combinatorial optimisation [8]. Hence, quantum computers have the potential to revolutionise artificial intelligence [9], computer-aided medicine design [10], financial modelling [11] and climate modelling [12] — to name but a few — and thereby address some of the greatest challenges facing the world today. This has led companies such as IBM, Google, IonQ and PsiQuantum to begin to build small quantum computers on which one can test simple quantum algorithms. Recent demonstrations of quantum supremacy show that quantum computers are indeed capable of performing some specialised computational tasks several orders of magnitude faster than the fastest available classical supercomputers [13–16]. For the most part, however, noise or errors typically prevents the successful realisation of sophisticated quantum algorithms on current quantum computers, which has led to these devices being dubbed noisy intermediate scale quantum (NISQ) computers [17].

In future fault-tolerant quantum computers, noise may be eliminated through the use of quantum error correcting codes [18–21]. Quantum error correcting codes allow the error rates for logical qubits to be made arbitrarily small, by using multiple physical qubits for each logical qubit, provided that the error rates for physical qubits are below a certain threshold (known as the fault-tolerance threshold). Since the error rates for physical qubits on NISQ computers are typically well above the fault-tolerance thresholds of most quantum error cor-

recting codes, quantum error correction is generally not feasible on NISQ computers. Nevertheless, quantum error mitigation techniques can be used to substantially reduce the effect of noise on NISQ hardware [22]. In quantum error mitigation, additional experiments are performed to measure noise levels, and the results from these experiments are then used to adjust the results obtained for the experiments or protocols of interest. A number of quantum algorithms or protocols have been developed specifically for implementation on NISQ computers [23]. A large part of this thesis (Chapters 3–5) focuses on the implementation of protocols which generate and use randomness on NISQ computers.

The most common model for quantum computing is the ‘circuit model’, where quantum computing is performed by explicitly applying unitary operations (or gates) from a universal set [24, 25]. ‘Measurement-based’ quantum computing is a competing model, where the entire computation is carried out by performing adaptive single-qubit measurements on an entangled resource state [26–29]. Provided that the entangled resource state can be generated, quantum computing can be reduced to performing single-qubit measurements and multi-qubit entangling operations do not need to be performed on demand. This is highly advantageous in linear optical systems [30–32], where these entangling operations cannot be performed deterministically. Measurement-based quantum computing also has benefits in physical systems such as superconducting systems [33], cold atoms [34] and quantum dots [35], where the entangled resource state is easy to generate, since the qubits to be entangled are spatially close to each other. The main focus of this thesis is on measurement-based quantum computing, and the practical generation of randomness in this model and its application.

1.3 Quantum randomness

1.3.1 Random numbers

Random numbers are of fundamental importance in information theory. In particular, they are used extensively in cryptography [36], simulation [37] and fundamental physics tests [38], as well as in lotteries, machine learning, cryptocurrencies and coordination in computer networks [39]. When classical techniques are used to generate random numbers, the unpredictability relies on incomplete knowledge which can result in the random numbers being more predictable than anticipated. This can in turn result in their utility being compromised. On the other hand, when random numbers are generated using quantum mechanical techniques, unpredictability is guaranteed by the inherent randomness of quantum mechanics [39–41].

Early quantum random number generators of the twentieth century were based on radioactive decay [42–44]. Over the past two decades, these have largely been superseded by photonic quantum random number generators, since much higher generation rates can be achieved with photonics [40]. A great variety of quantum random number generation schemes have been successfully realised experimentally using photonics, including branching paths [45–47], time-of-arrival [48–56], photon counting [57–59], vacuum fluctuations [60–65] and laser phase fluctuations [66–71]. More recently, several of these schemes have been realised in the form of on-chip quantum random number generators [72–80], which is a focus of part of this thesis (Chapter 6). Random number generation schemes have also been implemented on cloud-based superconducting quantum computers [81–83], trapped ions [84] and magnetic tunnel junctions [85, 86].

1.3.2 Random unitary operators

The analog of random numbers in quantum information is unitary operators chosen randomly with respect to the Haar measure on $U(2^n)$, the group of unitary transformations on a n -qubit system. Unfortunately, the

1.4. Randomness generation on quantum hardware

resources required to sample randomly with respect to the Haar measure on $U(2^n)$ scale exponentially with n [87]. A unitary t -design is a pseudorandom ensemble of unitary operators of which the statistical moments match those of the uniform Haar ensemble either approximately or exactly up to some finite order t [88, 89]. Hence, a t -design is by definition a $(t - 1)$ -design. These t -designs can rarely be distinguished from the true random ensemble, and so they are often used as a substitute.

In particular, 1-designs are used for encrypting quantum data [90, 91]; and 2-designs are used for randomised benchmarking [92–100], characterising correlations within multipartite quantum systems [101] and formulating quantum mechanical models of black holes [102]. Furthermore, 3-designs are used for detecting entanglement [103–106] and solving black-box problems [107]; and 4-designs are used for quantum state distinction [108] and estimating the self-adjointness of quantum noise [109]. Higher order t -designs also find applications in noise estimation for even t [110].

In random circuit constructions, t -designs on n qubits are realised by applying gates selected randomly from a universal set to qubits from a n -qubit system. Approximate n -qubit t -designs can be realised efficiently, since the resources required (the number of gates and random bits) scale polynomially with n and t [111–114]. Very efficient random circuit constructions for exact n -qubit 2-designs, where the resources required scale almost linearly with n , have also been devised [115]. More recently, random circuit constructions for general exact n -qubit t -designs were proposed [109]. However, they are only feasible for small systems, since the number of gates required scales exponentially with n and t for large n . Random circuit constructions have two major disadvantages, namely that they require a source of classical randomness, which can be expensive if it needs to be reliable, and that they require the reconfiguring of physical quantum gates, which is likely to introduce noise.

A measurement-based approach [88, 89], inspired by measurement-based quantum computing, avoids both of these problems at the cost of additional qubits. Measurement-based t -designs are realised by performing a deterministic sequence of single-qubit measurements on a highly entangled graph state. Turner and Markham present a measurement-based protocol for realising an exact single-qubit 3-design, which requires a 6-qubit graph state, and discuss measurement-based protocols for realising higher order approximate single-qubit t -designs, which require larger graph states [88]. Efficient approximate n -qubit measurement-based t -designs, where the number of qubits in the entangled resource state scale polynomially with n and t , have also been found [89]. It is still unknown whether exact single-qubit measurement-based t -designs exist for $t > 3$ or whether exact multi-qubit measurement-based t -designs exist at all. It is within this context that the realisation of measurement-based t -designs using IBM quantum hardware is explored in this thesis.

1.4 Randomness generation on quantum hardware

1.4.1 IBM's cloud-based superconducting hardware

In 2016, IBM released the first cloud-based quantum computing platform. Since then, their processors have improved significantly, and IBM now has superconducting quantum computers ranging from 5 qubits to 127 qubits which are accessible through their website [116]. Random number generation has been explored extensively on the IBM quantum computers [81–83]. In previous work [82], we showed that, with post-processing, IBM quantum computers can be used to generate random numbers of sufficient quality for cryptographic applications, thereby demonstrating a substantial improvement over earlier implementations [81] which did not use post-processing. Source-independent quantum random number generation schemes have also been implemented on the IBM quantum computers [83]. Recently, IBM released their own cloud-based quantum random

1.4. Randomness generation on quantum hardware

number generation platform, with post-processing available via the University of Cambridge.

In this thesis, we go beyond the simple task of generating random numbers and investigate the generation of random unitary operators on the IBM quantum computers (Chapter 3). To this end, we implemented the exact single-qubit measurement-based 3-design of Ref. [88] and our own approximate single-qubit measurement-based 2-design on IBM quantum processors. These were implemented by performing single-qubit measurements on 6-qubit and 5-qubit graph states respectively. Since measurement errors are responsible for a significant amount of noise on IBM quantum processors, and since this noise is predominantly classical, we performed quantum readout error mitigation to improve results [117]. Both the exact 3-design implementation and the approximate 2-design implementation passed our test for a 1-design, but not for a 2-design or a 3-design. Further numerical investigations show that depolarising noise is likely what prevented these implementations from passing the test for a 2-design and a 3-design, and we discuss this in detail.

To investigate the extent to which the results obtained in these numerical investigations generalise, a further study was conducted in which we determined the effect of a variety of common noise channels on the quality of t -designs for $t \in \{1, 2, 3, 4, 5\}$ (Chapter 4). Two noise models were considered, namely one in which noise is applied before the unitary operations of the t -design and one in which noise is applied after the unitary operations. This is in line with the noise models used in randomised benchmarking [92–100], one of the primary applications of t -designs. For the model where noise is applied before the unitary operations, we were able to show analytically that the quality of the single-qubit 1-design is completely unaffected by an arbitrary noise channel, and for the model where noise is applied after the unitaries, we showed that the 1-design is unaffected by noise, unless amplitude damping is applied. We obtained numeric results for the 2-design, 3-design, 4-design and 5-design. These results suggest that a $2t$ -design is significantly more sensitive to noise than a $(2t - 1)$ -design and that, with the exception of the amplitude damping channel, a $(2t + 1)$ -design is as sensitive to noise as a $2t$ -design. We also found large variations in sensitivity to noise throughout the state space, with t -designs generally being most sensitive to noise for pure states and least sensitive to noise for the maximally mixed state. These findings may be helpful for researchers studying and developing applications using t -designs under realistic conditions. While we hope that our work will encourage research into the effect of noise on the quality of multi-qubit t -designs, we also note that there are many protocols which exclusively use single-qubit t -designs [101, 103, 104, 118] for which our results may have direct consequences.

Finally, we test the implementation of our approximate measurement-based 2-design on the IBM processors on an important practical problem in quantum computing, namely noise estimation (Chapter 5). To this end, we developed an interleaved randomised benchmarking [95, 96, 119–121] protocol for measurement-based quantum computers, in which any single-qubit measurement-based 2-design can be used to estimate the fidelity of any single-qubit measurement-based gate. We then demonstrated our protocol by using our approximate measurement-based 2-design to estimate the fidelity of the Hadamard gate and the T gate on IBM quantum computers. Since the Hadamard gate and the T gate form a universal single-qubit set [122], our experiments provide a proof-of-concept demonstration of fidelity estimation of measurement-based gates from a universal single-qubit set. In our demonstration, we use entangled resource states of up to 31 qubits. In all the experiments, estimated gate fidelities show good agreement with gate fidelities determined using alternative methods. Our study highlights the usefulness of IBM quantum processors for single-qubit measurement-based quantum computing and shows how to practically characterise noisy quantum logic gates in this setting. The results of our work contribute to the ongoing goal of building up to advanced multi-qubit quantum computing on NISQ processors using the measurement-based model.

1.4.2 Custom-built on-chip plasmonic hardware

While an IBM quantum computer provides a sophisticated platform on which to generate randomness, it is not specifically built for this task. Custom-built hardware, on the other hand, can be tailored for a specific application such as randomness generation. Recently, the branching paths scheme for quantum random number generation was demonstrated using an on-chip plasmonic beamsplitter [123]. The primary benefit of plasmonics [124,125] is that it enables the confinement of light to subwavelength scales, well below the diffraction limit, which allows quantum information processing protocols to be carried out at a much smaller scale than in the dielectric systems typically used in photonics [126, 127]. Despite being inherently lossy, plasmonic components have been successfully employed in quantum sensing [128–136], and have shown potential for entanglement generation [137] and quantum state engineering [138].

The time-of-arrival scheme for quantum random number generation has two major advantages over the branching paths scheme previously demonstrated using plasmonics, namely it requires one detector instead of two, and substantially higher bit rates are possible as multiple bits of randomness can be extracted from a single photon. So far, the time-of-arrival scheme has not been demonstrated using plasmonics. Moreover, most time-of-arrival generators are bulky and use a highly attenuated coherent source as a single-photon source, where photons simply propagate through free space to a single-photon detector [48–56]. An integrated circuit time-of-arrival generator has been realised recently by integrating a semi-coherent silicon LED source directly on a detector chip [73]. An alternative and more flexible option for on-chip time-of-arrival based quantum random number generation is to embed an on-chip source [139–142] and detector [143–146] within a plasmonic waveguide.

In this thesis (Chapter 6) we report a time-of-arrival based quantum random number generation scheme using an on-chip nanowire plasmonic waveguide. In particular, we integrate an on-chip nanowire plasmonic waveguide into an optical time-of-arrival based quantum random number generation setup and test its performance in the presence of loss. Despite loss, we managed to achieve a random number generation rate of 14.4 Mbits/s. Furthermore, the generated bits did not require any post-processing to pass the ENT [147] and NIST [148] Statistical Test Suites. By increasing the light intensity, we were able to increase the generation rate to 41.4 Mbits/s, however these bits required a shuffle to pass all the tests. Our work demonstrates the successful integration of an on-chip nanoscale plasmonic component into a custom-built time-of-arrival based quantum random number generation setup. Although the current experimental setup employs an off-chip source and detector, future work on the integration of an on-chip source [139–142] and detector [143–146] would yield a fully integrated nanophotonic quantum random number generator chip with a footprint an order of magnitude smaller than its dielectric counterpart. This opens up new opportunities in compact and scalable quantum random number generation.

1.5 Thesis outline

This thesis is structured as follows. In Chapter 2, we review some of the tools and concepts that are used extensively in this thesis. In Chapter 3, we discuss our implementation of the exact single-qubit measurement-based 3-design of Ref. [88] and our own approximate single-qubit measurement-based 2-design on the IBM quantum processors. In Chapter 4, we present our investigation of the effect of noise on the quality of unitary t -designs. In Chapter 5, we present our measurement-based interleaved randomised benchmarking protocol, as well as the results obtained for the demonstration of this protocol on the IBM quantum computers. Our experimental work on the integration of an on-chip nanowire plasmonic waveguide into an optical time-of-

1.5. Thesis outline

arrival based quantum random number generation setup is discussed in Chapter 6. A summary of the thesis, concluding remarks and a description of possible future work are given in Chapter 7. Supplementary appendices follow, in which we present further details about the experiments, complementary investigations and important proofs.

Chapter 2

Background

2.1 Quantum information processing

2.1.1 The qubit

In classical information theory, the most basic unit of information is the binary digit or bit. Bits can take only two values, namely 0 and 1. The analog of the bit in quantum information theory is the quantum bit or qubit [149] — a two-state quantum mechanical system. Examples of qubits include a photon with two polarisation directions and an atom with a ground state and an excited state. The two states of a qubit are denoted by $|0\rangle$ and $|1\rangle$. Unlike bits, qubits are not restricted to the states $|0\rangle$ and $|1\rangle$. Qubits can be in any linear combination or superposition of these states, that is, the most general form for the state of a qubit is

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle, \quad (2.1)$$

where $\alpha, \beta \in \mathbb{C}$ and $|\alpha|^2 + |\beta|^2 = 1$. Hence, the state of a qubit can be represented by a normalised vector in a two-dimensional complex vector space. In particular, we can write

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle = \alpha \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \beta \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix}, \quad (2.2)$$

so that the state of a qubit is represented by a normalised vector in \mathbb{C}^2 , the vector space of two-dimensional complex column vectors. The reason for the restriction to normalised vectors will become clear in Sec. 2.1.3.

The states $|0\rangle$ and $|1\rangle$ form an orthonormal basis for \mathbb{C}^2 . This basis is called the computational basis or the Pauli Z -basis. Other important bases for \mathbb{C}^2 include the Pauli X -basis, which is denoted by $\{|+\rangle, |-\rangle\}$, where

$$|\pm\rangle = \frac{1}{\sqrt{2}} (|0\rangle \pm |1\rangle), \quad (2.3)$$

and the Pauli Y -basis, which is denoted by $\{|+_y\rangle, |-_y\rangle\}$, where

$$|\pm_y\rangle = \frac{1}{\sqrt{2}} (|0\rangle \pm i |1\rangle). \quad (2.4)$$

The constraint $|\alpha|^2 + |\beta|^2 = 1$ allows us to rewrite Eq. (2.1) as

$$|\psi\rangle = e^{i\gamma} \left(\cos\left(\frac{\theta}{2}\right) |0\rangle + e^{i\phi} \sin\left(\frac{\theta}{2}\right) |1\rangle \right), \quad (2.5)$$




Figure 2.1: The Bloch sphere for visualising the state of a qubit.

where $\gamma, \phi, \theta \in \mathbb{R}$. It turns out that the global phase factor $e^{i\gamma}$ has no experimentally observable consequences, and so we may choose $\gamma = 0$, so that $e^{i\gamma} = 1$, for simplicity. This yields

$$|\psi\rangle = \cos\left(\frac{\theta}{2}\right)|0\rangle + e^{i\phi}\sin\left(\frac{\theta}{2}\right)|1\rangle, \quad (2.6)$$

which shows that there is a one-to-one correspondence between the states of a qubit and the points on the surface of a three-dimensional unit sphere, which is commonly referred to as the Bloch sphere (see Fig. 2.1). The Bloch sphere is a very useful tool for visualising the state of a qubit.

2.1.2 Unitary evolution

Quantum information is processed by performing operations on qubits so as to manipulate their states. The operations which can be performed on single qubits are represented by unitary operators or matrices acting on the vector space \mathbb{C}^2 [150]. For a unitary operator U applied to a qubit initially in the state $|\psi\rangle$, we have

$$|\psi'\rangle = U|\psi\rangle, \quad (2.7)$$

where $|\psi'\rangle$ is the state of the qubit after applying U . Restricting to unitary evolution and disallowing more general linear transformations ensures that $|\psi'\rangle$ is also a normalised vector in \mathbb{C}^2 . For the remainder of this section we introduce some common unitary operators that play an important role in quantum computing.

The simplest unitary operator is the identity operator,

$$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad (2.8)$$

which leaves the state of a qubit to which it is applied unchanged. The Pauli X , Y and Z operators are given by

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad (2.9)$$

and they transform the state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ of a qubit to which they are applied as follows:

$$\begin{aligned} X|\psi\rangle &= \beta|0\rangle + \alpha|1\rangle, \\ Y|\psi\rangle &= -i(\beta|0\rangle - \alpha|1\rangle), \\ Z|\psi\rangle &= \alpha|0\rangle - \beta|1\rangle. \end{aligned}$$

Hence, the Pauli X operator performs a bit flip, in the sense that it interchanges the coefficients of the computational basis states $|0\rangle$ and $|1\rangle$, and the Pauli Z operator performs a phase flip, in the sense that it inverts the relative phase by changing the sign of the coefficient of $|1\rangle$. The Pauli Y operator, which can be written as $Y = iXZ$, performs both a bit flip and a phase flip. The Pauli operators together with the identity operator form the Pauli group, which is an exact unitary 1-design.

Other unitary operators which are used extensively in quantum computing include the Hadamard (H) gate, the S gate and the T gate. These are given by

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \quad S = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}, \quad T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}, \quad (2.10)$$

in matrix form. More general unitary operators can be defined in terms of variable parameters, for example the unitary operator which rotates the state of a qubit around the z -axis of the Bloch sphere by an angle ϕ is defined by

$$R_z(\phi) = \begin{pmatrix} e^{-i\phi/2} & 0 \\ 0 & e^{i\phi/2} \end{pmatrix}. \quad (2.11)$$

2.1.3 Projective measurements

Quantum information processing would be of little use if there were no way to extract processed information. Information can be extracted from qubits by performing measurements. In this section, we discuss projective measurements [151]. While more general measurements exist [152], they are outside the scope of this thesis. A projective measurement is a measurement of a physical observable, represented by a Hermitian operator or matrix acting on \mathbb{C}^2 for single qubits. We note that the spectral decomposition of a Hermitian operator M is given by

$$M = \sum_m m P_m, \quad (2.12)$$

where P_m is the projector onto the eigenspace of M with eigenvalue m . Hence, the projective measurement of the observable represented by M can be fully described by the orthogonal projection operators P_m and the eigenvalues m , which are the possible outcomes of the measurement. When a projective measurement is performed on a qubit in the state $|\psi\rangle$, the outcome m occurs with probability

$$p(m) = \langle\psi| P_m |\psi\rangle, \quad (2.13)$$

and the state of the qubit after a measurement in which the outcome m occurred is given by

$$|\psi_m\rangle = \frac{P_m |\psi\rangle}{\sqrt{p(m)}}. \quad (2.14)$$

Furthermore, the average or expectation value of the observable represented by M is given by

$$\langle M \rangle = \langle \psi | M | \psi \rangle. \quad (2.15)$$

The most common projective measurement in quantum computing is a measurement in the computational basis or Pauli Z -basis, for which the projectors are given by $|0\rangle\langle 0|$ and $|1\rangle\langle 1|$. As the latter name suggests, the Hermitian operator associated with this projective measurement is the Pauli Z operator. Applying Eqs. (2.13) and (2.14), we find that a computational basis measurement performed on a qubit in the state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ yields $|0\rangle$ with probability $|\alpha|^2$ and $|1\rangle$ with probability $|\beta|^2$. The normalisation constraint, $|\alpha|^2 + |\beta|^2 = 1$, imposed in Sec. 2.1.1 ensures that the probabilities for the different computational basis measurement outcomes sum to one. We note that even though a qubit can be in any superposition of the computational basis states $|0\rangle$ and $|1\rangle$, a computational basis measurement causes its state to collapse to either $|0\rangle$ or $|1\rangle$.

Other common projective measurements include measurements in the Pauli X -basis, for which the projectors are given by $|+\rangle\langle +|$ and $|-\rangle\langle -|$, and measurements in the Pauli Y -basis, for which the projectors are given by $|+_y\rangle\langle +_y|$ and $|-_y\rangle\langle -_y|$. The Hermitian operators associated with these measurements are the Pauli X and Y operators respectively. More generally, we can consider projective measurements in a basis $\{U|0\rangle, U|1\rangle\}$, for which the projectors are given by $U|0\rangle\langle 0|U^\dagger$ and $U|1\rangle\langle 1|U^\dagger$, where U is any unitary operator acting on \mathbb{C}^2 . Since quantum computers usually only support computational basis measurements at the hardware level, these more general measurements are often realised by applying U^\dagger and measuring in the computational basis.

2.1.4 Multi-qubit systems

Just as classical information processing systems are composed of multiple bits, quantum information processing systems are composed of multiple qubits. The state of a n -qubit system composed of qubits in the states $|\psi_1\rangle, |\psi_2\rangle, \dots, |\psi_n\rangle \in \mathbb{C}^2$ can be represented by the vector

$$|\psi\rangle = |\psi_1\rangle \otimes |\psi_2\rangle \otimes \dots \otimes |\psi_n\rangle = |\psi_1\rangle |\psi_2\rangle \dots |\psi_n\rangle, \quad (2.16)$$

that is, by the tensor product of the vectors representing the states of the individual qubits in the system [153]. Multi-qubit states of this form are called separable states. In general, the state of a n -qubit system can be represented by any normalised vector in $(\mathbb{C}^2)^{\otimes n}$, the n -fold tensor product space of \mathbb{C}^2 . Multi-qubit states which cannot be expressed in the form of Eq. (2.16), that is, multi-qubit states which cannot be separated or factorised into a tensor product of single-qubit states, are called entangled states. Much of the computational power of quantum computers is derived from entangled states.

The operations which can be performed on a n -qubit system are represented by unitary operators or matrices acting on $(\mathbb{C}^2)^{\otimes n}$. Two 2-qubit unitary operators which play a very important role in quantum computing are the controlled not (CX) gate and the controlled phase (CZ) gate, which are given by

$$CX = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \quad CZ = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}, \quad (2.17)$$

in matrix form [154]. The controlled not gate performs a bit flip on the second qubit (the target qubit) if the first qubit (the control qubit) is in the state $|1\rangle$, and leaves the target qubit unchanged if the control qubit is in

the state $|0\rangle$. Similarly, the controlled phase gate performs a phase flip on the target qubit if the control qubit is in the state $|1\rangle$, and leaves the target qubit unchanged if the control qubit is in the state $|0\rangle$. When the control qubit is in a non-trivial superposition of the states $|0\rangle$ and $|1\rangle$, the controlled not and phase gates result in the control and target qubits becoming entangled. Hence the controlled not and phase gates can be used to generate entangled states.

While a great variety of entangled states exist, the only ones considered in this thesis are graph states. Graph states are used in measurement-based quantum computing [26–29], measurement-based t -designs [88, 89] and a wide range of other protocols, including quantum secret sharing [155, 156], quantum sensing [157, 158] and quantum games [159, 160]. A n -qubit graph state is defined in relation to a connected graph with n vertices. Such a graph state is made by preparing each qubit in the state $|+\rangle$ and then applying a controlled phase gate to a pair of qubits whenever their corresponding vertices are connected by an edge in the corresponding graph [161]. In this thesis, we are mostly interested in linear cluster states, which are graph states corresponding to a graph in which the degree of each vertex is less than or equal to 2 (excluding rings).

2.2 Quantum information processing in the presence of noise

The formulation reviewed in Sec. 2.1 is useful for describing quantum information processing in the ideal scenario, where qubits never interact with their surrounding environment, and operations and measurements can be implemented with perfect accuracy. In practice, any experimental realisation of a quantum information processing system is subject to noise, as a result of interactions with its surrounding environment and other imperfections. Noise can cause the state of a qubit to evolve in an uncontrolled manner, so that the resulting state is not completely known. To represent these noisy states, a more general formulation of quantum information processing, namely the density operator formulation, is needed.

2.2.1 The density operator formulation

In the density operator formulation, the state of a qubit is represented by a positive operator or matrix with unit trace acting on \mathbb{C}^2 [162]. Such an operator or matrix is called a density operator or matrix. In particular, when a qubit is in any known state $|\psi\rangle \in \mathbb{C}^2$, its state can be represented by the density operator,

$$\rho = |\psi\rangle\langle\psi|, \quad (2.18)$$

in the density operator formulation. States of this kind are called pure states. When the state of a qubit is not completely known, but it is known that the qubit is in one of a number of pure states $|\psi_i\rangle$ each with probability p_i , then its state can be represented by the density operator

$$\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|. \quad (2.19)$$

States of this kind are called mixed states.

We can still use the Bloch sphere (see Fig. 2.1) to visualise the state of a qubit in the density operator formulation. To this end, note that any single-qubit density matrix can be written in the form

$$\rho = \frac{1}{2} \begin{pmatrix} 1+z & x-iy \\ x+iy & 1-z \end{pmatrix}, \quad (2.20)$$

where $x, y, z \in \mathbb{R}$ and $r = \sqrt{x^2 + y^2 + z^2} \leq 1$. Hence, there is a one-to-one correspondence between single-qubit density matrices and points of the Bloch sphere. In particular, pure states (which are represented by density matrices with $r = 1$) correspond to points on the surface of the Bloch sphere as before, while mixed states (which are represented by density matrices with $r < 1$) correspond to points within the Bloch sphere. Furthermore, we note that the maximally mixed state, which is given by

$$\rho = \frac{1}{2}I, \quad (2.21)$$

corresponds to the point at the centre of the Bloch sphere.

We conclude this section with a brief overview of unitary evolution, projective measurements and multi-qubit systems in the density operator formulation. For a unitary operator U applied to a qubit initially in the state ρ , we have

$$\rho' = U\rho U^\dagger, \quad (2.22)$$

where ρ' is the state of the qubit after applying U . When a projective measurement of an observable represented by a Hermitian operator M , with spectral decomposition given by Eq. (2.12), is performed on a qubit in the state ρ , the outcome m occurs with probability

$$p(m) = \text{Tr}(P_m \rho), \quad (2.23)$$

the state of the qubit after a measurement in which the outcome m occurred is given by

$$\rho_m = \frac{P_m \rho P_m^\dagger}{p(m)}, \quad (2.24)$$

and expectation value of the observable represented by M is given by

$$\langle M \rangle = \text{Tr}(M \rho). \quad (2.25)$$

Furthermore, the state of a n -qubit system is represented by a density operator acting on $(\mathbb{C}^2)^{\otimes n}$ in the density operator formulation. In particular, the state of a n -qubit system composed of qubits in the states $\rho_1, \rho_2, \dots, \rho_n$ can be represented by the density operator

$$\rho = \rho_1 \otimes \rho_2 \otimes \dots \otimes \rho_n. \quad (2.26)$$

2.2.2 Noise channels

The evolution of the state of a qubit as a result of noise can be modelled by a noise channel. In general, the action of a noise channel on a density operator ρ is described by a completely positive, trace-preserving map ε and the resulting density operator is denoted by $\varepsilon(\rho)$. An example of a noise channel which shows up repeatedly in this thesis is the depolarising noise channel [163]. The depolarising channel is described by

$$\varepsilon(\rho) = \frac{p}{2}I + (1-p)\rho, \quad (2.27)$$

that is, a state ρ is replaced by the maximally mixed state with probability p . Other examples of noise channels are given in Sec. 4.2.1.

2.3. Quantum computing




Figure 2.2: Quantum circuit for preparing the 3-qubit linear cluster state.

2.3 Quantum computing

2.3.1 The circuit model

The circuit model is the most common model for quantum computing. In the circuit model, quantum computations are carried out by applying unitary operators, or quantum logic gates, from a universal set [24, 25]. A universal set of gates is a finite set of unitary operators or gates such that any unitary operator can be approximated to arbitrary accuracy by a finite sequence of gates from that set [122]. The most common universal set is $\{H, T, CX\}$. For single-qubit quantum computing, this set can be reduced to $\{H, T\}$.

In the circuit model, quantum algorithms or protocols are represented visually by quantum circuits, which are analogous to the circuit diagrams used in classical computing [164]. However, while the circuits in classical computing are actual electronic circuits, the logic gates in a quantum circuit represent the unitary operators which are applied to qubits in a given protocol. An example of a quantum circuit is shown in Fig. 2.2. In Chapters 3 and 5, it will become clear that this circuit represents a protocol for preparing the 3-qubit linear cluster state. All three qubits are initially in the computational basis state $|0\rangle$. Hadamard gates are then applied to Qubits 1 and 3, after which a controlled not gate is applied to Qubits 1 and 2, with Qubit 1 as the control qubit (indicated by the \cdot symbol) and Qubit 2 as the target qubit (indicated by the \oplus symbol). Thereafter, another controlled not gate is applied to Qubits 2 and 3, with Qubit 3 as the control qubit and Qubit 2 again as the target qubit, after which a Hadamard gate is applied to Qubit 2. Finally, all three qubits are measured in the computational basis (indicated by the meter symbol). Note how the quantum circuit is read from left to right.

2.3.2 Measurement-based quantum computing

Measurement-based quantum computing is a competing model for quantum computing. In the measurement-based model, quantum computing is realised by performing adaptive single-qubit measurements on an entangled resource state, such as a graph state or cluster state [26–29]. In particular, 2D cluster states are entangled resource states for universal quantum computing [28] and fault-tolerant quantum computing can be achieved using 3D cluster states [165]. In this thesis, we concentrate on single-qubit measurement-based quantum computing with linear cluster states [161].

We briefly review measurement-based processing with linear cluster states [161]. The measurement-based protocol for implementing unitary operations with a n -qubit linear cluster state is illustrated in Fig. 2.3. The first qubit is prepared in the input state, $\rho_{\text{in}} = |\psi_{\text{in}}\rangle\langle\psi_{\text{in}}|$, to which the implemented unitary operation is to be applied. The remaining qubits are prepared in the state $|+\rangle$ (Step 1), and qubits are then entangled by applying controlled phase gates to neighbouring qubits (Step 2). This prepares the n -qubit linear cluster state with ρ_{in} encoded within. To implement a unitary operation, each qubit $i \in \{1, 2, \dots, n-1\}$ is measured in the basis $\{|^+\phi_i\rangle, |^-\phi_i\rangle\}$, where

$$|^{\pm}\phi_i\rangle = \frac{1}{\sqrt{2}} \left(|0\rangle \pm e^{-i\phi_i} |1\rangle \right), \quad (2.28)$$




Figure 2.3: Measurement-based processing with a n -qubit linear cluster state. Step 1 depicts the initialisation of the qubits. Step 2 depicts the entangled cluster state (after controlled phase gates have been applied to neighbouring qubits) in addition to the measurements performed on each qubit. Step 3 depicts the state resulting from the measurements.

that is, at an angle ϕ_i in the Pauli XY plane of the Bloch sphere, which we will often refer to simply as a measurement in the ϕ_i -direction. Each measurement is logically equivalent to applying the unitary operation

$$U_{m_i}(\phi_i) = X^{m_i} H R_z(\phi_i), \quad (2.29)$$

to $|\psi_{\text{in}}\rangle$, where $m_i \in \{0, 1\}$ is the measurement outcome (occurring with probability $p_{m_i} = \frac{1}{2}$ for all m_i). Hence, these measurements result in qubit n being prepared in the output state, $\rho_{\text{out}} = U_{\mathbf{m}}(\phi) \rho_{\text{in}} U_{\mathbf{m}}^\dagger(\phi)$ (Step 3), with the implemented unitary operation given by

$$U_{\mathbf{m}}(\phi) = U_{m_{n-1}}(\phi_{n-1}) \cdots U_{m_2}(\phi_2) U_{m_1}(\phi_1), \quad (2.30)$$

where ϕ and \mathbf{m} denote ordered lists of angles and measurement outcomes (occurring with probability $p_{\mathbf{m}} = \frac{1}{2^{n-1}}$ for all \mathbf{m}) respectively.

Even though the implemented unitary operation is probabilistic and depends on the measurement outcomes, any deterministic quantum computation can be realised by employing adaptive measurement feedforward [29, 161]. In particular, unitary operations which are deterministic up to a single known Pauli gate can be implemented by performing the measurements on qubits 1 to $n - 1$ sequentially and allowing future measurement angles to depend on past measurement outcomes. The desired quantum computation is then realised deterministically by applying the required and known Pauli correction to qubit n .

As an example, we give the measurement-based implementations of the Hadamard gate and the T gate. When all measurement outcomes are zero, the 2-qubit linear cluster state with the measurement angle $\phi_1 = 0$ implements the Hadamard gate and the 3-qubit linear cluster state with the measurement angles $\phi_1 = \frac{\pi}{4}$ and $\phi_2 = 0$ implements the T gate. When non-zero measurement outcomes occur, the desired gate can still be implemented by applying the appropriate Pauli correction to the final qubit [161]. In particular, the Pauli X operator must be applied to implement the Hadamard gate when $m_1 = 1$. For the T gate, the Pauli X operator must be applied when $m_1 = 0$ and $m_2 = 1$, the Pauli Y operator must be applied when $m_1 = 1$ and $m_2 = 1$, and the Pauli Z operator must be applied when $m_1 = 1$ and $m_2 = 0$. Hence the Hadamard gate and the T gate can be realised deterministically with fixed measurement angles and a Pauli correction, and in these specific cases do not require adaptive measurement feedforward. However, implementations of more complicated deterministic operations, such as arbitrary single-qubit rotations, require both adaptive measurement feedforward and a Pauli correction [28].

2.4 IBM's cloud-based superconducting hardware

Superconducting quantum computers ranging from 5 qubits to 127 qubits are available via IBM's cloud-based quantum computing platform [116]. In this thesis, the work on the IBM quantum computers forms part of the theory component. In particular, the IBM quantum computers are used to obtain supporting experimental results for the theoretical work in Chapters 3 and 5. As such, a detailed review of the underlying superconducting hardware is outside the scope of this thesis.

In short, the IBM superconducting quantum computers employ transmon qubits [166], fabricated using microelectronics lithography [167, 168], which are a sophisticated variation of conventional charge qubits. The computational basis states $|0\rangle$ and $|1\rangle$ are the ground and first excited states respectively of an anharmonic quantum oscillator. Unitary operations are carried out by applying microwave pulses of a specific duration and frequency to the transmon qubits. Similarly, measurements in the computational basis are realised through the interaction of microwave photons with the transmon qubits. An example of a multi-qubit entangling operation for adjacent transmon qubits is given in Ref. [169].

2.5 Quantum randomness

Randomness plays a ubiquitous role in information processing tasks. In classical computing, randomness typically takes the form of uniformly distributed random bits, that is, binary digits which are assigned the value 0 with probability $\frac{1}{2}$ and the value 1 with probability $\frac{1}{2}$. In quantum computing, on the other hand, randomness typically takes the form of random unitary operators distributed uniformly over $U(2^n)$, the group of unitary operators acting on $(\mathbb{C}^2)^{\otimes n}$. The mathematical definition of uniformly distributed random unitary operators is somewhat more involved than that of uniformly distributed random bits.

For a random unitary operator to be a well defined concept, we need to turn the group $U(2^n)$ into a probability distribution by equipping it with an appropriate measure. Since $U(2^n)$ is a finite-dimensional compact Lie group, a general element of $U(2^n)$ can be written in terms of a finite number of real parameters, each of which lie in some closed interval, that is, the underlying manifold of $U(2^n)$ is a closed and bounded region in some Euclidean space. Hence the manifold of $U(2^n)$ has a finite Euclidean volume, which ensures that the associated distribution is normalisable. In particular, we can obtain a valid normalised probability distribution by introducing a measure such that an integral of unity over the entire $U(2^n)$ with respect to that measure yields 1. At this point, one may be tempted to naively use the normalised Euclidean volume element associated with the Euclidean volume of the underlying manifold of $U(2^n)$. While this measure would give a valid probability distribution, it is not unitarily invariant, in the sense that the normalised Euclidean volume of a subset of $U(2^n)$ may change if a unitary operator is applied to the elements of the subset. As a result, unitary operators chosen randomly with respect to this measure would show a bias towards certain subsets of $U(2^n)$, instead of being distributed uniformly over $U(2^n)$, which is rarely of any use in quantum computing.

The appropriate measure to introduce in order to obtain a uniform probability distribution is the Haar measure [170]. The Haar measure is the unique unitarily invariant measure on $U(2^n)$, that is, it is the unique measure dU on $U(2^n)$ such that for all $V \in U(2^n)$ and all integrable functions f ,

$$\int_{U(2^n)} f(U) dU = \int_{U(2^n)} f(UV) dU = \int_{U(2^n)} f(VU) dU. \quad (2.31)$$

The Haar measure is also the measure with respect to which we must integrate when averaging over $U(2^n)$ or subsets thereof. In particular, the expectation of an operator ρ acting on $((\mathbb{C}^2)^{\otimes n})^{\otimes t}$ with respect to the Haar

2.5. Quantum randomness

measure on $U(2^n)$ is defined by

$$\mathbb{E}_H^t(\rho) = \int_{U(2^n)} U^{\otimes t} \rho (U^{\otimes t})^\dagger dU. \quad (2.32)$$

Finally, we note that the number of real parameters needed to describe a general element of $U(2^n)$ grows exponentially with n . Hence, sampling unitary operators randomly with respect to the Haar measure on $U(2^n)$ is inefficient for large n .

Consequently, a pseudorandom ensemble in the form of a unitary t -design is frequently used as an efficient substitute. A unitary t -design is an ensemble of unitary operators of which the statistical moments match those of the uniform Haar ensemble either approximately or exactly up to some finite order t . Thus formally an ensemble of unitaries $\{p_i, U_i\}$ is an exact n -qubit unitary t -design if for all operators ρ acting on $((\mathbb{C}^2)^{\otimes n})^{\otimes t}$,

$$\mathbb{E}_H^t(\rho) = \sum_i p_i U_i^{\otimes t} \rho (U_i^{\otimes t})^\dagger, \quad (2.33)$$

and $\{p_i, U_i\}$ is an ϵ -approximate t -design if there exists an ϵ such that for all operators ρ acting on $((\mathbb{C}^2)^{\otimes n})^{\otimes t}$,

$$(1 - \epsilon) \mathbb{E}_H^t(\rho) \leq \sum_i p_i U_i^{\otimes t} \rho (U_i^{\otimes t})^\dagger \leq (1 + \epsilon) \mathbb{E}_H^t(\rho), \quad (2.34)$$

where the matrix inequality $A \leq B$ holds if $B - A$ is positive semidefinite [88, 89]. An exact t -design can also be defined as an ϵ -approximate t -design with $\epsilon = 0$.

Since the positive semidefinite property defines a partial order (the Loewner order) on Hermitian matrices, inequality (2.34) is a natural generalisation of an error bound inequality from scalars to Hermitian matrices. However, an interpretation of ϵ in terms of defining an error range for the Haar ensemble expectation determined with an ϵ -approximate t -design is unclear, as the Haar ensemble expectation $\mathbb{E}_H^t(\rho)$ is a matrix comprising many different scalar entries, not a single scalar expectation value. This is further complicated by the fact that inequality (2.34) need not be symmetric, that is, the ϵ required to satisfy the left inequality may differ from the ϵ required to satisfy the right inequality, and only the larger of these, which is the ϵ required to satisfy inequality (2.34), is known. It is also unclear how ϵ can be linked to a distance measure. Nevertheless, it is clear that at a fundamental level, ϵ quantifies an ϵ -approximate t -design's ability to replicate the moments of the uniform Haar ensemble. The smallest possible ϵ is zero, for which we recover an exact t -design, and any larger value quantifies the deviation from an exact t -design, which is unbounded in theory. In practice, an arbitrarily chosen bound, which depends on the application at hand, is typically enforced, as we will see in Chapters 3 and 4.

Chapter 3

Implementation of single-qubit measurement-based t -designs using IBM processors

3.1 Introduction

In previous experiments, multi-qubit pseudorandom ensembles — where the expected distribution of matrix elements of unitary operators sampled from the uniform Haar ensemble is reproduced — have been realised using a nuclear magnetic resonance quantum processor [171], and single-qubit unitary 1-designs and 2-designs have been realised using photons [172]. In this chapter, we implement the single-qubit unitary 3-design of Ref. [88] using the measurement-based model on IBM superconducting quantum computers. To this end, we performed single-qubit measurements on a 6-qubit linear cluster state. Since measurement errors are responsible for a significant amount of noise on IBM quantum processors, and since this noise is predominantly classical, we performed quantum readout error mitigation to improve results [117]. The exact 3-design implementation passed our test for a 1-design, but not for a 2-design or a 3-design. Further investigations, presented in Appendix A, suggest that depolarising noise is likely what prevented the implementation from passing our test for a 2-design and a 3-design. We managed to obtain improved results for the 2-design test by implementing an approximate 2-design, in which measurements were performed on a smaller 5-qubit linear cluster state, but the test still did not pass for all states.

This chapter is structured as follows. In Sec. 3.2, we discuss measurement-based t -designs as well as the state and process tomography techniques used to analyse results and the quantum readout error mitigation technique used to improve results. In Sec. 3.3, we describe the implementations of the exact 3-design and approximate 2-design and present the results obtained. Some concluding comments are given in Sec. 3.4. Further details about the implementations and analysis of the results are presented in Appendix A.

3.2 Background

3.2.1 Measurement-based t -designs

In addition to being entangled resource states for single-qubit measurement-based quantum computing (see Sec. 2.3.2), linear cluster states can be used to implement single-qubit unitary t -designs by entirely foregoing adaptive measurement feedforward and Pauli corrections, that is, by fixing the measurement angles ϕ and

considering the ensemble of unitaries $\{p_{\mathbf{m}}, U_{\mathbf{m}}(\phi)\}$ for all measurement outcomes \mathbf{m} . In particular, Turner and Markham [88] show that the 6-qubit linear cluster state with the measurement angles $\phi_1 = 0, \phi_2 = \frac{\pi}{4}, \phi_3 = \arccos \sqrt{1/3}, \phi_4 = \frac{\pi}{4}$ and $\phi_5 = 0$ implements an exact 3-design (in the sense that the ensemble $\{p_{\mathbf{m}}, U_{\mathbf{m}}(\phi)\}$, which consists of the 32 unitary operators corresponding to the 32 possible measurement outcomes \mathbf{m} , is an exact 3-design). In this chapter, we use little endian encoding to present lists of measurement outcomes, that is, for a n -qubit linear cluster state a list of measurement outcomes is presented as a bit string in which the leftmost bit is the outcome of the measurement on qubit $n - 1$ and the rightmost bit is the outcome of the measurement on qubit 1. This is in correspondence with how measurement outcomes are presented on IBM processors.

3.2.2 Quantum state tomography

Quantum state tomography can be used to infer the output state of an experimental realisation of a protocol, such as a unitary operator implemented with a linear cluster state, for a given input state [173]. The protocol of interest is executed multiple times for a fixed input state ρ_{in} , to obtain multiple copies of the corresponding output state. When the output state is a single-qubit state, these copies are used to infer the expectations values of the Pauli X , Y and Z observables by averaging the outcomes of measurements performed in the respective Pauli bases. The output state corresponding to ρ_{in} is then given by

$$\rho_{\text{out}} = \frac{1}{2} \begin{pmatrix} 1 + \langle Z \rangle & \langle X \rangle - i \langle Y \rangle \\ \langle X \rangle + i \langle Y \rangle & 1 - \langle Z \rangle \end{pmatrix}, \quad (3.1)$$

where $\langle X \rangle$, $\langle Y \rangle$ and $\langle Z \rangle$ are the expectation values of the respective Pauli observables.

3.2.3 Quantum process tomography

Quantum process tomography can be used to determine the extent to which a protocol is realised by an experimental implementation. We consider a method proposed for single-qubit protocols by Nielsen and Chuang [173, 174]. Given any input state ρ_{in} , we write the output state as

$$\varepsilon(\rho_{\text{in}}) = \sum_{mn} E_m \rho_{\text{in}} E_n^\dagger \chi_{mn}, \quad (3.2)$$

where $E_0 = I$, $E_1 = X$, $E_2 = -iY$, $E_3 = Z$ and χ is a 4 by 4 matrix. Since the operators E_i are fixed, the implementation of the protocol is fully characterised by the χ matrix or process matrix, and so process tomography amounts to determining χ . The entries of χ depend on the action of the protocol on the probe input states, $|0\rangle$, $|1\rangle$, $|+\rangle$ and $|+_y\rangle$, which is determined by performing state tomography (see Sec. 3.2.2) to infer the output states for these input states. In particular,

$$\chi = \frac{1}{4} \begin{pmatrix} I & X \\ X & -I \end{pmatrix} \begin{pmatrix} \rho'_1 & \rho'_2 \\ \rho'_3 & \rho'_4 \end{pmatrix} \begin{pmatrix} I & X \\ X & -I \end{pmatrix}, \quad (3.3)$$

where the submatrices of the middle matrix are defined by

$$\begin{aligned}\rho'_1 &= \varepsilon(|0\rangle\langle 0|) \\ \rho'_4 &= \varepsilon(|1\rangle\langle 1|) \\ \rho'_2 &= \varepsilon(|+\rangle\langle +|) + i\varepsilon(|+_y\rangle\langle +_y|) - \frac{1+i}{2}(\rho'_1 + \rho'_4) \\ \rho'_3 &= \varepsilon(|+\rangle\langle +|) - i\varepsilon(|+_y\rangle\langle +_y|) - \frac{1-i}{2}(\rho'_1 + \rho'_4),\end{aligned}$$

where $\varepsilon(|0\rangle\langle 0|)$, $\varepsilon(|1\rangle\langle 1|)$, $\varepsilon(|+\rangle\langle +|)$ and $\varepsilon(|+_y\rangle\langle +_y|)$ denote the output states determined for the respective probe input states. Once constructed, we can use the process matrix to quantify the reliability with which a protocol is realised by an experimental implementation or channel by calculating the channel fidelity, which is given by

$$F(\chi, \tilde{\chi}) = \text{Tr} \left(\sqrt{\sqrt{\chi} \tilde{\chi} \sqrt{\chi}} \right), \quad (3.4)$$

where χ is the process matrix for the ideal protocol and $\tilde{\chi}$ is the process matrix for the experimental implementation of the protocol obtained from process tomography. The channel fidelity ranges from 0 to 1, where 0 indicates that the experimental implementation or channel deviates maximally from the ideal protocol and 1 indicates a perfect implementation of the protocol.

3.2.4 Quantum readout error mitigation

As a result of measurement errors, actual quantum states and protocols are often more similar to expected states and protocols than tomography results would suggest. Since measurement errors on IBM quantum processors are mostly classical [117], quantum readout error mitigation can be used to obtain tomography results which more accurately reflect the prepared states and implemented protocols, as has been successfully done in a number of recent studies which also involved measurements on highly entangled states on IBM quantum processors [175–177]. To mitigate readout errors in a n -qubit experiment (the main experiment), we first use quantum detector tomography [178] to construct a 2^n by 2^n calibration matrix Λ . The entries of Λ are the conditional probabilities of measuring each of the 2^n possible combinations of computational basis states, given that a specific combination of computational basis states was prepared, for all 2^n possible combinations of computational basis states. In particular, each column of Λ contains the 2^n conditional probabilities associated with one of the 2^n prepared combinations of computational basis states. These conditional probabilities are determined in a series of separate experiments, in which each of the 2^n possible combinations of computational basis states is prepared on the n qubits to be used in the main experiment and sufficient computational basis measurements are done to infer the associated conditional probabilities. Once constructed, Λ can be used to correct classical measurement errors in the main experiment by multiplying Λ^{-1} by \mathbf{p}_{exp} , the column vector containing the relative frequencies obtained in the main experiment. As a result of other noise, such as gate errors, the resulting vector, $\Lambda^{-1}\mathbf{p}_{\text{exp}}$, may be non-physical (relative frequencies may be negative or may not sum to one). We therefore use qiskit's built-in method [179], which solves a constrained optimisation problem (least squares method), to find the closest physical relative frequency vector to $\Lambda^{-1}\mathbf{p}_{\text{exp}}$.

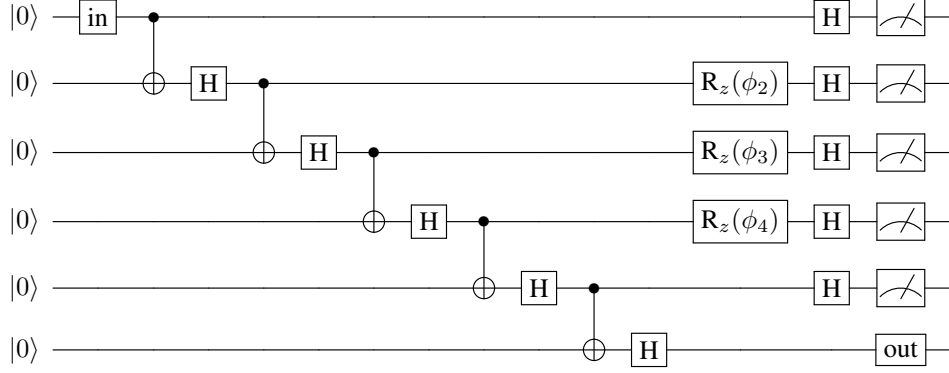


Figure 3.1: General quantum circuit for implementation of the exact measurement-based 3-design proposed by Turner and Markham [88] on the *ibmq_toronto* quantum processor. Here ‘in’ represents the set of gates applied to construct the input state and ‘out’ represents the set of gates applied and measurements done to perform tomography on the sixth qubit. The angles for the z -rotation gates are $\phi_2 = \phi_4 = \frac{\pi}{4}$ and $\phi_3 = \arccos \sqrt{1/3}$.

3.3 Experiments

3.3.1 Implementation

The exact measurement-based 3-design proposed by Turner and Markham [88] and described in Sec. 3.2.1 was implemented on six physical qubits of the *ibmq_toronto* quantum processor. Appendix A.1 provides more details on the *ibmq_toronto* quantum processor and how the logical qubits 1 to 6 were mapped onto the physical qubits of this processor. The four process tomography probe states were considered as input states. For each input state, we prepared the 6-qubit linear cluster state and performed the appropriate single-qubit measurements. Quantum state tomography was then done on the sixth qubit to construct the output state obtained for each of the 32 different measurement outcomes. A general quantum circuit for the implementation is shown in Fig. 3.1.

Qubits are initialised in the state $|0\rangle$ by default on IBM processors, and so the state $|1\rangle$ was prepared by applying the Pauli X operator, the state $|+\rangle$ was prepared by applying a Hadamard gate and the state $|+_y\rangle$ was prepared by applying a Hadamard gate, followed by a S gate. Since IBM processors do not support controlled phase gates at the hardware level, we converted the Hadamards and controlled phase gates, needed to prepare the 6-qubit linear cluster state, into Hadamards and controlled not gates using

$$CZ = (I \otimes H)CX(I \otimes H) \quad (3.5)$$

and then eliminated redundant Hadamards using the fact that $H^2 = I$. Doing so ensured that redundant Hadamards, which would have increased noise in the results due to gate errors, were removed from the circuit. Preparing the 6-qubit linear cluster state with controlled phase gates would have resulted in redundant Hadamards being introduced by the transpiler. Since IBM processors can only perform measurements in the computational basis, measurements in the ϕ -direction were realised by applying $R_z(\phi)$, followed by a Hadamard, and measuring in the computational basis. Quantum state tomography was done using qiskit’s built-in method [180], which uses maximum-likelihood estimation to ensure that the density matrices constructed from the data are physical (i.e. that they have unit trace and are positive). The state tomography results were used to do process tomography for each of the 32 different measurement outcomes to determine the extent to which the 32 corresponding unitary operations performed on the input states in the implementation matched the expected unitary operations. The results are presented in Sec. 3.3.2.

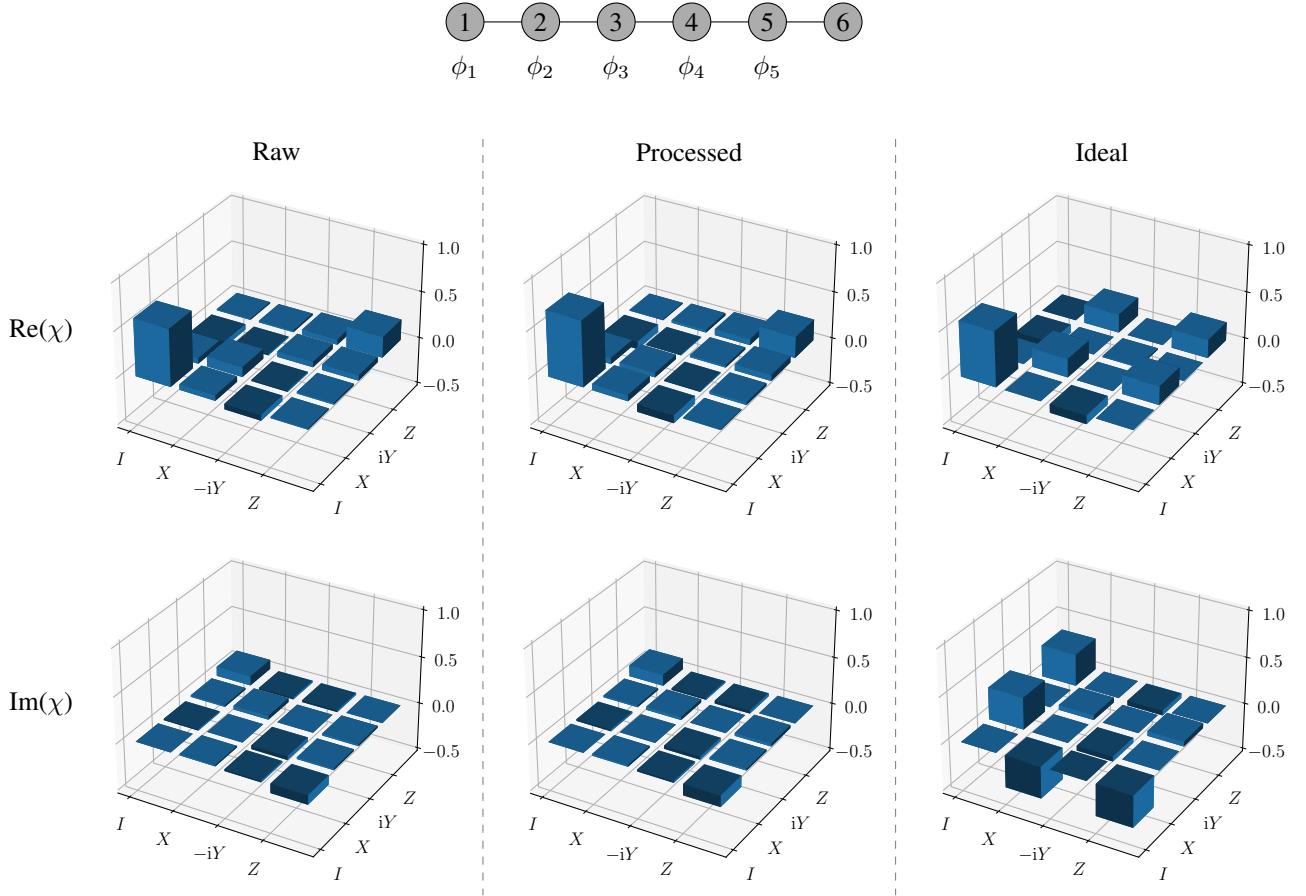


Figure 3.2: Example of process tomography results for the random unitary that agreed least with the ideal case generated for measurement outcome $\mathbf{m} = 00000$ with the exact 3-design on the *ibmq_toronto* quantum processor. The diagram at the top shows the entangled 6-qubit linear cluster state with the measurements performed on each qubit. The χ matrix obtained without quantum readout error mitigation is shown on the left, the χ matrix obtained with quantum readout error mitigation is shown in the middle and the ideal χ matrix is shown on the right. The real part of each matrix is shown above and the imaginary part of each matrix is shown below.

Each of the twelve circuits needed for process tomography (three circuits for state tomography to construct the output state for each of the four probe input states) was run five times with 8000 shots on the *ibmq_toronto* quantum processor. The counts obtained in the five different runs of the same circuit were then combined to obtain an effective run with 40000 shots for each of the twelve circuits. This was done to decrease statistical noise in the tomography results. The procedure was repeated ten times to obtain ten sets of process tomography results (χ matrices) for each of the 32 different unitary operations. Associated results, such as channel fidelities, quoted in the sections which follow, are an average of these ten repetitions and the errors quoted are the standard deviations.

To obtain the conditional probabilities needed to construct the calibration matrix, required to mitigate readout errors in the tomography results, we prepared each of the 64 possible combinations of computational basis states on the same six qubits of the *ibmq_toronto* quantum processor as was used for the 3-design implementation and measured these qubits in the computational basis. Each of the 64 circuits was run with 8000 shots and no combination of counts was done. Readout errors in the raw tomography data (i.e. the counts obtained by running the various circuits) were then mitigated as described in Sec. 3.2.4. Results obtained using both the raw and the processed (error mitigated) data are presented in the sections which follow.

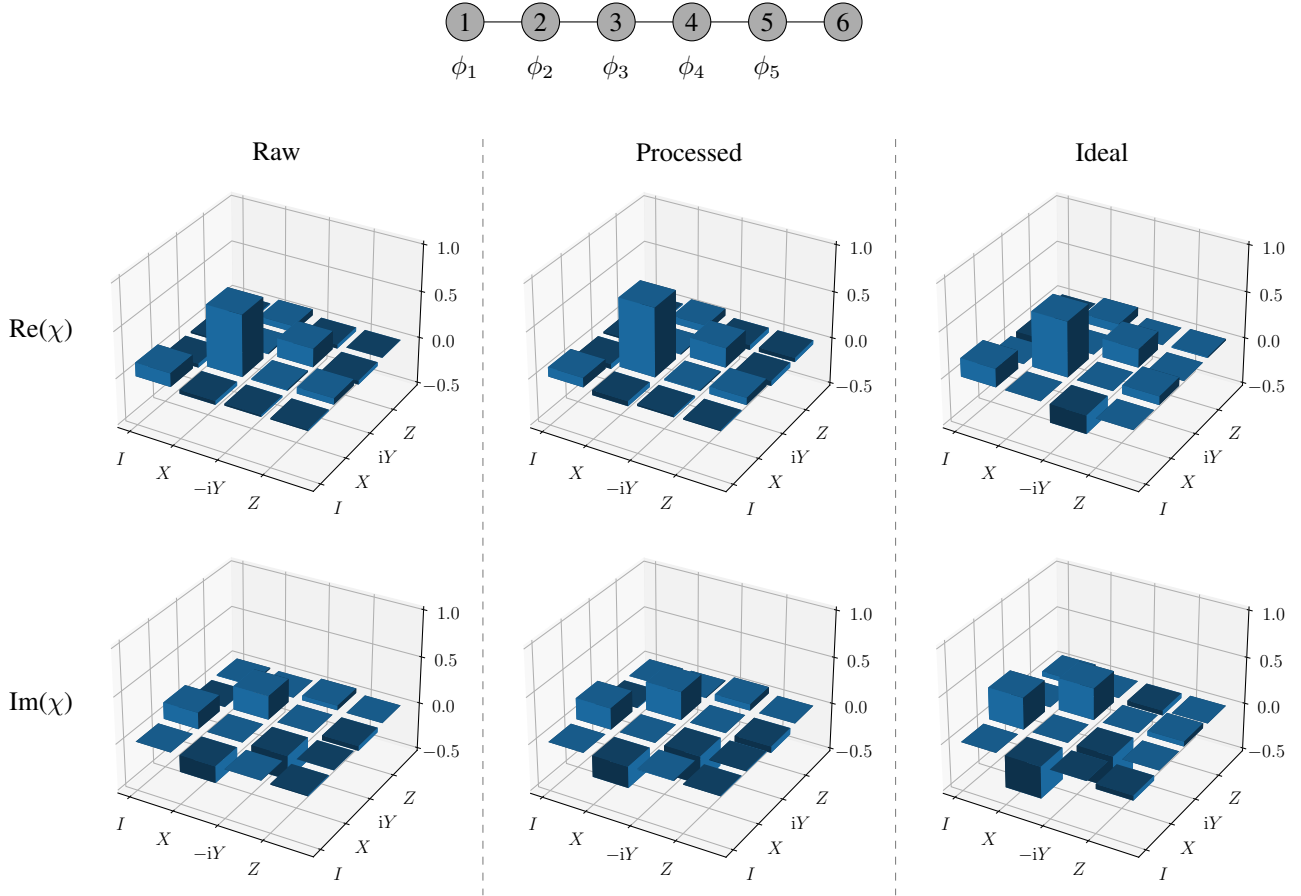


Figure 3.3: Example of process tomography results for the random unitary that agreed most with the ideal case generated for measurement outcome $m = 11001$ with the exact 3-design on the *ibmq_toronto* quantum processor. The diagram at the top shows the entangled 6-qubit linear cluster state with the measurements performed on each qubit. The χ matrix obtained without quantum readout error mitigation is shown on the left, the χ matrix obtained with quantum readout error mitigation is shown in the middle and the ideal χ matrix is shown on the right. The real part of each matrix is shown above and the imaginary part of each matrix is shown below.

3.3.2 Process tomography results

Process tomography results obtained for the two random unitaries, generated with the exact 3-design on the *ibmq_toronto* quantum processor, which showed the least and most agreement with theoretical predictions, are shown in Figs. 3.2 and 3.3 respectively. Channel fidelities for each of the 32 different random unitaries generated with the exact 3-design implementation are given in Table 3.1 and the distribution of these channel fidelities is displayed in Fig. 3.4. The average channel fidelity is (0.8220 ± 0.0325) without quantum readout error mitigation and (0.8754 ± 0.0361) with quantum readout error mitigation. Quantum readout error mitigation improved all the channel fidelities, which confirms that classical measurement errors were responsible for a significant amount of noise in the exact measurement-based 3-design implementation, and would have resulted in channel fidelities which greatly underestimate the reliability with which the expected unitary operations are realised in the implementation, if left uncorrected.

3.3.3 Relative frequencies

Due to errors that occur when gates are applied and measurements are made, the relative frequencies with which the 32 random unitaries are generated, with the exact 3-design on the *ibmq_toronto* quantum processor, do not exactly match the expected uniform probabilities of $\frac{1}{32} = 0.03125$. To determine the relative frequency

Outcome	Fidelity (Raw)	Fidelity (Processed)
00000	0.7246 ± 0.0064	0.7670 ± 0.0073
00001	0.8090 ± 0.0098	0.8688 ± 0.0106
00010	0.7820 ± 0.0051	0.8397 ± 0.0057
00011	0.8013 ± 0.0106	0.8596 ± 0.0119
00100	0.8061 ± 0.0077	0.8657 ± 0.0084
00101	0.8506 ± 0.0084	0.9192 ± 0.0096
00110	0.8498 ± 0.0060	0.9157 ± 0.0073
00111	0.8358 ± 0.0086	0.8949 ± 0.0095
01000	0.7819 ± 0.0106	0.8295 ± 0.0115
01001	0.8356 ± 0.0082	0.8919 ± 0.0090
01010	0.7865 ± 0.0035	0.8295 ± 0.0041
01011	0.8194 ± 0.0066	0.8715 ± 0.0074
01100	0.8215 ± 0.0061	0.8755 ± 0.0071
01101	0.8426 ± 0.0068	0.8920 ± 0.0075
01110	0.8272 ± 0.0056	0.8708 ± 0.0066
01111	0.8538 ± 0.0059	0.9042 ± 0.0065
10000	0.7503 ± 0.0088	0.7953 ± 0.0101
10001	0.8423 ± 0.0084	0.9056 ± 0.0096
10010	0.8135 ± 0.0072	0.8656 ± 0.0080
10011	0.8203 ± 0.0069	0.8830 ± 0.0081
10100	0.7992 ± 0.0096	0.8569 ± 0.0103
10101	0.8597 ± 0.0066	0.9202 ± 0.0069
10110	0.8645 ± 0.0054	0.9240 ± 0.0060
10111	0.8460 ± 0.0067	0.9035 ± 0.0079
11000	0.8104 ± 0.0078	0.8593 ± 0.0088
11001	0.8774 ± 0.0053	0.9353 ± 0.0060
11010	0.8238 ± 0.0084	0.8689 ± 0.0088
11011	0.8419 ± 0.0066	0.8876 ± 0.0076
11100	0.8013 ± 0.0070	0.8447 ± 0.0077
11101	0.8287 ± 0.0097	0.8765 ± 0.0102
11110	0.8513 ± 0.0047	0.8936 ± 0.0052
11111	0.8467 ± 0.0071	0.8959 ± 0.0081

Table 3.1: Channel fidelities for the 32 random unitaries, corresponding to the 32 different measurement outcomes, generated with the exact 3-design on the *ibmq_toronto* quantum processor. ‘Raw’ shows the channel fidelities without quantum readout error mitigation. ‘Processed’ shows the channel fidelities with quantum readout error mitigation.

with which each unitary is generated, a separate investigation was conducted in which the 6-qubit linear cluster state was prepared on the same six physical qubits of the *ibmq_toronto* quantum processor as was used for the 3-design implementation (see Appendix A.1) and the first five qubits were measured in the same way as in the 3-design implementation. The relative frequency of each set of measurement outcomes is the relative frequency with which the random unitary corresponding to that set of measurement outcomes is generated. The required circuit was run five times with 8000 shots each and counts were once again combined to obtain an effective run with 40000 shots. This was repeated ten times, so that relative frequencies quoted are an average of ten repetitions and the errors quoted are the standard deviations. The relative frequencies with which each of the 32 different random unitaries are generated with the exact 3-design on the *ibmq_toronto* quantum processor are presented in Table 3.2 and the distribution of these relative frequencies is displayed in Fig. 3.5. The average relative frequency is (0.03125 ± 0.00355) without quantum readout error mitigation and (0.03125 ± 0.00375) with quantum readout error mitigation. We note that even though the relative frequencies with which the different unitaries are generated deviate from the expected uniform probabilities, the average relative frequency is equal to the expected probability.




Figure 3.4: Distribution of channel fidelities for the 32 random unitaries generated with the exact 3-design on the *ibmq_toronto* quantum processor. (a) Raw shows the distribution without quantum readout error mitigation. (b) Processed shows the distribution with quantum readout error mitigation.




Figure 3.5: Distribution of relative frequencies with which the 32 random unitaries are generated with the exact 3-design on the *ibmq_toronto* quantum processor. (a) Raw shows the distribution without quantum readout error mitigation. (b) Processed shows the distribution with quantum readout error mitigation.

3.3.4 Testing for a t -design

The definition of an approximate t -design as given by inequality (2.34) leads naturally to a simple method for testing whether the ensemble of 32 unitaries generated with the 6-qubit linear cluster state on the *ibmq_toronto* quantum processor is at least an approximate single-qubit t -design. Although this definition applies to any density matrix acting on the tensor product space $(\mathbb{C}^2)^{\otimes t}$, we restrict ourselves to density matrices which are t copies of an arbitrary single-qubit density matrix for the purposes of testing. This is sufficient for quantifying the extent to which the unitaries are able to randomise single-qubit states, which is our primary interest here, and can at least provide a lower bound on the ϵ for which the ensemble of unitaries is an ϵ -approximate t -design with more general states included. The restriction to copies of single-qubit density matrices has two major advantages, namely that the test is computationally feasible for all t , since the number of parameters that need to be varied when creating samples of density matrices would otherwise grow exponentially with t , and that the results can be interpreted geometrically, since single-qubit states correspond to points in the Bloch

Outcome	Frequency (Raw)	Frequency (Processed)
00000	0.03836±0.00095	0.03664±0.00105
00001	0.03130±0.00057	0.02987±0.00059
00010	0.03372±0.00075	0.03196±0.00083
00011	0.02737±0.00064	0.02632±0.00074
00100	0.02861±0.00097	0.02702±0.00108
00101	0.03335±0.00095	0.03183±0.00108
00110	0.03447±0.00091	0.03367±0.00104
00111	0.03305±0.00082	0.03253±0.00091
01000	0.03535±0.00094	0.03567±0.00107
01001	0.02721±0.00076	0.02699±0.00086
01010	0.03530±0.00085	0.03636±0.00094
01011	0.02716±0.00040	0.02750±0.00048
01100	0.03079±0.00081	0.03089±0.00098
01101	0.03349±0.00069	0.03428±0.00080
01110	0.03477±0.00111	0.03628±0.00131
01111	0.02986±0.00087	0.03124±0.00096
10000	0.03788±0.00110	0.03733±0.00123
10001	0.02789±0.00068	0.02666±0.00074
10010	0.03399±0.00069	0.03379±0.00078
10011	0.02408±0.00051	0.02289±0.00057
10100	0.02548±0.00065	0.02432±0.00075
10101	0.03348±0.00089	0.03325±0.00103
10110	0.03032±0.00066	0.03009±0.00075
10111	0.03238±0.00073	0.03287±0.00081
11000	0.03188±0.00075	0.03221±0.00086
11001	0.02757±0.00050	0.02828±0.00057
11010	0.03254±0.00131	0.03376±0.00148
11011	0.02753±0.00081	0.02886±0.00095
11100	0.03008±0.00094	0.03121±0.00112
11101	0.03017±0.00096	0.03116±0.00110
11110	0.03430±0.00104	0.03673±0.00125
11111	0.02631±0.00111	0.02759±0.00127

Table 3.2: Relative frequencies with which the 32 random unitaries, corresponding to the 32 different measurement outcomes, are generated with the exact 3-design on the *ibmq_toronto* quantum processor. ‘Raw’ shows the relative frequencies without quantum readout error mitigation. ‘Processed’ shows the relative frequencies with quantum readout error mitigation.

sphere. For the purposes of testing, we therefore consider the adapted inequality

$$(1 - \epsilon)\mathbb{E}_H^t(\rho^{\otimes t}) \leq \sum_i p_i \rho_i'^{\otimes t} \leq (1 + \epsilon)\mathbb{E}_H^t(\rho^{\otimes t}), \quad (3.6)$$

where p_i are the experimentally determined relative frequencies and

$$\rho_i' = \sum_{mn} E_m \rho E_n^\dagger \chi_{mn}^{(i)}, \quad (3.7)$$

where $\chi^{(i)}$ are the process matrices determined by doing process tomography for the different unitaries on the *ibmq_toronto* quantum processor. The test amounts to generating a sample of single-qubit density matrices and finding, for each density matrix in the sample, the smallest possible ϵ such that inequality (3.6) is satisfied. The largest ϵ found is the one which ensures that inequality (3.6) is satisfied for all density matrices in the sample and is therefore the test result.

Test	Raw			Processed		
	Radius	ϵ	ϵ (uniform)	Radius	ϵ	ϵ (uniform)
1-design	1.00	0.0777 ± 0.0066	0.0760 ± 0.0072	1.00	0.0683 ± 0.0054	0.0677 ± 0.0072
2-design	0.68	0.4543 ± 0.0074	0.4559 ± 0.0063	0.75	0.4538 ± 0.0188	0.4464 ± 0.0179
3-design	0.66	0.4590 ± 0.0061	0.4592 ± 0.0070	0.69	0.4814 ± 0.0062	0.4696 ± 0.0058

Table 3.3: Summary of test results for the ensemble of unitaries generated using the *ibmq_toronto* quantum processor. ‘Raw’ shows the results without quantum readout error mitigation. ‘Processed’ shows the results with quantum readout error mitigation. ‘Radius’ is the truncation radius considered for a test. The column with ‘uniform’ shows the values of ϵ obtained when replacing the experimentally determined relative frequencies with uniform probabilities.

We set the passing criterion for the test to $\epsilon \leq 0.5$, by which we simply mean that, for the purposes of this chapter, we consider the quality of an approximate t -design acceptable if $\epsilon \leq 0.5$. In practice, some applications of approximate t -designs may require a smaller value of ϵ . We also note that, given an ensemble of unitaries which is an exact t -design or an approximate t -design with $\epsilon_t \leq 0.5$, it is generally not possible to say whether this ensemble of unitaries is also an approximate $(t + 1)$ -design with $\epsilon_{t+1} \leq 0.5$, as this depends on the ensemble. Therefore an ensemble of unitaries which passes our test for an approximate t -design, for a given t , may or may not pass our test for an approximate $(t + 1)$ -design.

To generate a sample of single-qubit density matrices, we first generate a representative sample of points in the Bloch sphere using spherical coordinates. We generate 10 evenly spaced values of r in the range $[0, 1]$, 10 evenly spaced values of ϕ in the range $[0, 2\pi)$ and 10 evenly spaced values of θ in the range $[0, \pi]$. Using the standard conversions from spherical to cartesian coordinates,

$$\begin{aligned} x &= r \sin \theta \cos \phi \\ y &= r \sin \theta \sin \phi \\ z &= r \cos \theta, \end{aligned}$$

we compute x , y and z for all combinations of the sampled values of r , ϕ and θ , thereby obtaining 1000 points in the Bloch sphere. Using Eq. (2.20), we obtain a sample of 1000 density matrices.

Using our sample of 1000 density matrices, we applied the test for the 1-design, the 2-design and the 3-design to the ensemble of unitaries generated using the *ibmq_toronto* quantum processor. The expected unitary operations of the exact 3-design given in Sec. 3.2.1 were used to compute $\mathbb{E}_H^t(\rho^{\otimes t})$ for $t = 1, 2, 3$. The test for the 1-design passed. The tests for the 2-design and the 3-design did not pass, as ϵ diverged for states close to the surface of the Bloch sphere. Nevertheless, inequality (3.6) could be satisfied for states close to the centre of the Bloch sphere. This was investigated further by re-applying the tests for the 2-design and the 3-design, this time truncating the values of r considered when generating density matrices so that ϵ did not exceed 0.5. The test results are summarised in Table 3.3. Applying quantum readout error mitigation improved the test results. The changes in the values of ϵ resulting from replacing the experimentally determined relative frequencies with uniform probabilities, are mostly within the error margins. This suggests that non-uniformity did not significantly impair the quality of the ensemble.

The divergence in ϵ observed for states close to the surface of the Bloch sphere is likely a result of pure states becoming inaccessible due to depolarising noise and was investigated further by applying the test for the 1-design, the 2-design and the 3-design to an exact 3-design combined with a depolarising channel. The full study is presented in Appendix A.2 and shows that depolarising noise is a very good noise model for the data.

Test	Raw		Processed	
	Frac	Frac (uniform)	Frac	Frac (uniform)
1-design	1.0000	1.0000	1.0000	1.0000
2-design	0.3834 ± 0.0027	0.3858 ± 0.0034	0.5648 ± 0.0079	0.5768 ± 0.0081
3-design	0.3473 ± 0.0048	0.3534 ± 0.0054	0.5315 ± 0.0061	0.5454 ± 0.0073

Table 3.4: Fraction of states for which the ensemble of unitaries generated using the *ibmq_toronto* quantum processor passed the different tests. ‘Raw’ shows the fractions without quantum readout error mitigation. ‘Processed’ shows the fractions with quantum readout error mitigation. The column with ‘uniform’ shows the fractions obtained when replacing the experimentally determined relative frequencies with uniform probabilities.

Our measurement-based implementation of the identity operation (see Appendix A.3) shows that depolarising noise is the predominant type of noise in measurement-based processes on IBM processors, providing further confirmation that depolarising noise is indeed what prevented the tests for the 2-design and the 3-design from passing. Urbanek *et al.* recently proposed a method for mitigating depolarising noise in quantum computations where the final outcome is an expectation value [181]. Since the final outcome of our implementation is a quantum process or channel, and not an expectation value, this method is unfortunately not applicable here. However, their method may have potential for improving results in applications of measurement-based t -designs where the final outcome is an expectation value.

To determine the fraction of the Bloch sphere for which a test passes, we consider 8000 evenly spaced points in a cube which encloses the Bloch sphere. Points in the Bloch sphere then correspond to valid states. For each valid state, we determine whether inequality (3.6) can be satisfied with $\epsilon \leq 0.5$. The fraction of the Bloch sphere for which a test passes is given by the number of states for which the inequality can be satisfied divided by the number of states considered. The fraction of states for which the ensemble of unitaries generated using the *ibmq_toronto* quantum processor passed the test for the 1-design, the 2-design and the 3-design are given in Table 3.4. The fraction of states for which the test for the 2-design and the 3-design pass is substantially improved by quantum readout error mitigation. This confirms that classical measurement noise is responsible for many states failing to satisfy inequality (3.6) and, if left uncorrected, would result in test results which greatly underestimate the extent to which the various designs are realised by our implementation on the *ibmq_toronto* quantum processor.

3.3.5 Approximate 2-design

Turner and Markham [88] show that there is no set of measurement angles such that the ensemble of unitaries generated by performing single-qubit measurements on the 5-qubit linear cluster state is an exact 2-design. However, applying our test for an approximate 2-design to the expected ensemble of unitaries generated for the 5-qubit linear cluster state with the measurement angles $\phi_1 = 0$, $\phi_2 = \frac{\pi}{4}$, $\phi_3 = \frac{\pi}{4}$ and $\phi_4 = 0$ yields $\epsilon = 0.5$. Hence this ensemble is an approximate 2-design, with our passing criterion, although it must be noted that the ensemble does not resemble an exact 2-design closely and that the passing criterion is satisfied only for the subset of density matrices acting on $\mathbb{C}^2 \otimes \mathbb{C}^2$ which are tensor products of pairs of single-qubit states. We implemented this approximate measurement-based 2-design on five physical qubits of the *ibmq_sydney* quantum processor. Appendix A.4 provides more detail on the *ibmq_sydney* quantum processor, the qubits that were used and why the *ibmq_sydney* quantum processor was used for this experiment instead of the *ibmq_toronto* quantum processor. Generation of process tomography results for the 16 different unitary operations corresponding to the 16 different measurement outcomes, determining of relative frequencies, combining of counts to reduce statistical noise and construction of calibration matrices for quantum readout error mitigation were

Outcome	Fidelity (Raw)	Fidelity (Processed)
0000	0.7931±0.0033	0.8947±0.0038
0001	0.8871±0.0042	0.9851±0.0047
0010	0.8382±0.0051	0.9220±0.0061
0011	0.8506±0.0040	0.9242±0.0047
0100	0.8231±0.0062	0.8900±0.0069
0101	0.8912±0.0053	0.9649±0.0053
0110	0.8399±0.0051	0.8978±0.0058
0111	0.8885±0.0041	0.9378±0.0047
1000	0.7991±0.0039	0.8978±0.0044
1001	0.9039±0.0053	0.9944±0.0063
1010	0.8365±0.0058	0.9061±0.0062
1011	0.8885±0.0059	0.9527±0.0061
1100	0.8487±0.0052	0.9183±0.0054
1101	0.9052±0.0055	0.9639±0.0055
1110	0.8571±0.0030	0.9053±0.0031
1111	0.9027±0.0060	0.9448±0.0066

Table 3.5: Channel fidelities for the 16 random unitaries, corresponding to the 16 different measurement outcomes, generated with the approximate 2-design on the *ibmq_sydney* quantum processor. ‘Raw’ shows the channel fidelities without quantum readout error mitigation. ‘Processed’ shows the channel fidelities with quantum readout error mitigation.

all done in the same way as for the exact 3-design implementation on the *ibmq_toronto* quantum processor.

Channel fidelities for each of the 16 different random unitaries are given in Table 3.5. The average channel fidelity is (0.8596 ± 0.0356) without quantum readout error mitigation and (0.9312 ± 0.0323) with quantum readout error mitigation. The average channel fidelity is larger than the average channel fidelity for the exact 3-design implementation, reflecting reduced noise in the implementation with the smaller cluster state with fewer qubits. The relative frequencies with which each of the 16 different random unitaries are generated are given in Table 3.6. The average relative frequency is (0.06250 ± 0.01127) without quantum readout error mitigation and (0.06250 ± 0.00871) with quantum readout error mitigation. The average relative frequency is once again equal to the expected uniform probability of $\frac{1}{16} = 0.0625$.

We applied our test for the 1-design and the 2-design to the ensemble of unitaries generated using the *ibmq_sydney* quantum processor. The test for the 1-design passed, but the test for the 2-design did not. Test results are summarised in Table 3.7 and the fraction of states for which each test passed is conveyed in Table 3.8. The values of ϵ obtained for the ensemble of unitaries generated using the *ibmq_sydney* quantum processor, for a given truncation radius, are not much larger than the expected values for the approximate 2-design. This suggests that inherent deviations from an exact 2-design, present in the approximate 2-design considered, had a more significant effect on the quality of the ensemble of unitaries than noise on the *ibmq_sydney* quantum processor.

The fraction of states for which the test for the 2-design passed with quantum readout error mitigation is almost double that without quantum readout error mitigation. This confirms that the noise in this implementation was also predominantly classical measurement errors. The effect of readout errors was more pronounced in this implementation, likely because gate errors of the qubits used are much smaller (see Appendix A.4). We note that the fraction of states for which the ensemble of unitaries generated with the 5-qubit linear cluster state on the *ibmq_sydney* quantum processor passed the test for the 2-design, especially with quantum readout error mitigation, is larger than that of the ensemble of unitaries generated with the 6-qubit linear cluster state on the *ibmq_toronto* quantum processor. Hence, even though the approximate 2-design considered does not closely resemble an exact 2-design, the ensemble of unitaries generated with this approximate 2-design implementation

3.4. Conclusion

Outcome	Frequency (Raw)	Frequency (Processed)
0000	0.07835±0.00104	0.06486±0.00113
0001	0.06705±0.00151	0.06268±0.00169
0010	0.07319±0.00100	0.06585±0.00122
0011	0.05817±0.00074	0.05930±0.00091
0100	0.07856±0.00119	0.07513±0.00138
0101	0.05442±0.00073	0.05646±0.00096
0110	0.07869±0.00115	0.08351±0.00144
0111	0.04837±0.00095	0.05567±0.00121
1000	0.06820±0.00120	0.05981±0.00137
1001	0.05429±0.00100	0.05296±0.00120
1010	0.06837±0.00120	0.06680±0.00135
1011	0.04801±0.00088	0.05098±0.00107
1100	0.06347±0.00120	0.06286±0.00150
1101	0.04988±0.00149	0.05518±0.00191
1110	0.06642±0.00137	0.07384±0.00168
1111	0.04458±0.00125	0.05411±0.00164

Table 3.6: Relative frequencies with which the 16 random unitaries, corresponding to the 16 different measurement outcomes, are generated with the approximate 2-design on the *ibmq_sydney* quantum processor. ‘Raw’ shows the relative frequencies without quantum readout error mitigation. ‘Processed’ shows the relative frequencies with quantum readout error mitigation.

Test	Raw				Processed			
	Radius	ϵ	ϵ (uniform)	ϵ (ideal)	Radius	ϵ	ϵ (uniform)	ϵ (ideal)
1-design	1.00	0.1505±0.0056	0.1468±0.0073	0.0000	1.00	0.1505±0.0056	0.1397±0.0065	0.0000
2-design	0.69	0.4488±0.0035	0.4690±0.0052	0.2739	0.81	0.4623±0.0103	0.4865±0.0073	0.3589

Table 3.7: Summary of test results for the ensemble of unitaries generated using the *ibmq_sydney* quantum processor. ‘Raw’ shows the results without quantum readout error mitigation. ‘Processed’ shows the results with quantum readout error mitigation. ‘Radius’ is the truncation radius considered for a test. The column with ‘uniform’ shows the values of ϵ obtained when replacing the experimentally determined relative frequencies with uniform probabilities. The column with ‘ideal’ shows the expected values of ϵ for the approximate 2-design for the truncation radii considered.

more closely resembles a 2-design than the ensemble of unitaries generated with the exact 3-design implementation. This is as a result of significantly reduced noise in the implementation with the smaller 5-qubit cluster state, compared to the implementation with the larger 6-qubit cluster state.

3.4 Conclusion

The exact measurement-based 3-design of Ref. [88] was implemented by performing single-qubit measurements on a 6-qubit linear cluster state, prepared on the *ibmq_toronto* quantum processor. To infer the ensemble of unitaries realised in the implementation, we performed quantum process tomography for all possible measurement outcomes. This ensemble of unitaries passed our test for a 1-design, but not for a 2-design or a 3-design. Further studies, presented in Appendices A.2 and A.3, strongly suggest that depolarising noise prevented the tests for the 2-design and the 3-design from passing. Therefore, for measurement-based t -designs to be effectively realised for $t > 1$ on superconducting systems, such as IBM quantum processors, a significant amount of work will need to be done to reduce or mitigate depolarising noise in these devices.

The noteworthy improvement in results obtained by applying quantum readout error mitigation confirms that classical measurement errors are indeed responsible for a substantial amount of noise on IBM quantum pro-

Test	Raw		Processed	
	Frac	Frac (uniform)	Frac	Frac (uniform)
1-design	1.0000	1.0000	1.0000	1.0000
2-design	0.3984 ± 0.0017	0.4104 ± 0.0020	0.6773 ± 0.0095	0.6925 ± 0.0050

Table 3.8: Fraction of states for which the ensemble of unitaries generated using the *ibmq_sydney* quantum processor passed the different tests. ‘Raw’ shows the fractions without quantum readout error mitigation. ‘Processed’ shows the fractions with quantum readout error mitigation. The column with ‘uniform’ shows the fractions obtained when replacing the experimentally determined relative frequencies with uniform probabilities.

cessors in this instance. It also shows the importance of mitigating these errors, as not doing so would lead to results that give an inaccurate account of the actual implementations realised on these processors. The ensemble of unitaries realised by our approximate measurement-based 2-design implementation, in which single-qubit measurements were performed on a 5-qubit linear cluster state prepared on the *ibmq_sydney* quantum processor, showed improved results for the 2-design test as a result of reduced noise for the smaller 5-qubit cluster state. This clearly demonstrates the advantage of keeping entangled resource states used in measurement-based processes small. It also shows that in experimental realisations (where noise is present), the quality of a noisy approximate t -design may be better than the quality of a noisy exact t -design, if the implementation of the approximate t -design is significantly less sensitive to noise than the implementation of the exact t -design.

Chapter 4

Investigating the effect of noise channels on the quality of unitary t -designs

4.1 Introduction

Irrespective of whether t -designs are implemented using the measurement-based technique discussed in Chapter 3 or with random circuit constructions, experimental realisations of t -designs are subject to noise [172]. In this chapter, we investigate the effect of noise on the quality of unitary t -designs for single qubits, similar to the way the effect of noise on randomised benchmarking [182] and the variational quantum eigensolver [183] have been simulated. Since the extent to which applications of t -designs are affected by noise in the underlying t -design is likely to differ, depending on the application, we studied the effect of noise on the quality of t -designs without reference to any particular application. We determined the effect of the bit flip channel, the phase flip channel, the bit and phase flip channel, the phase damping channel, the amplitude damping channel and the depolarising noise channel on the quality of t -designs for $t \in \{1, 2, 3, 4, 5\}$. We did not investigate t -designs beyond $t = 5$, since there are few known constructions for exact t -designs beyond $t = 5$ and few known applications of t -designs beyond $t = 4$. We considered two noise models — one in which noise is applied before the unitary operations of the t -design and one in which noise is applied after the unitary operations, in line with the noise models used in randomised benchmarking [92–100]. We showed analytically that the quality of the single-qubit 1-design is completely unaffected by an arbitrary noise channel, for the model where noise is applied before the unitary operations. For the model where noise is applied after the unitaries, we showed that the quality of the 1-design is unaffected by noise, unless amplitude damping is applied. Numeric results obtained for the 2-design, 3-design, 4-design and 5-design suggest that a $2t$ -design is significantly more sensitive to noise than a $(2t - 1)$ -design and that, with the exception of the amplitude damping channel, a $(2t + 1)$ -design is as sensitive to noise as a $2t$ -design. Furthermore, these numeric results show large variations in sensitivity to noise throughout the state space, with t -designs generally being most sensitive to noise for pure states and least sensitive to noise for the maximally mixed state.

This chapter is structured as follows. In Sec. 4.2, we discuss the various noise channels that we use to study noisy t -designs. In Sec. 4.3, we describe the two noise models considered. Our main numeric results for the 2-design, 3-design, 4-design and 5-design are presented in Sec. 4.4. A summary of the results and concluding remarks are given in Sec. 4.5. Analytic results for the 1-design and further numeric results for the higher order t -designs, as well as some important proofs, are included in Appendix B, while Appendix C covers complementary numeric results which give a geometric picture of the state dependence of the effect of noise

channels on the quality of single-qubit t -designs.

4.2 Background

4.2.1 Noise channels

An introduction to noise channels is given in Sec. 2.2.2. In this section, we introduce the noise channels that we use to study noisy t -designs. We consider four different types of single-qubit noise channels. These types of noise channels occur in many different physical systems [184].

4.2.1.1 Flip channels

We consider the bit flip channel and the phase flip channel, as well as the bit and phase flip channel. The bit flip channel [163] is described by

$$\varepsilon(\rho) = pX\rho X + (1-p)\rho, \quad (4.1)$$

that is, a bit flip is applied to a state ρ with probability p . The phase flip channel [163] is described by

$$\varepsilon(\rho) = pZ\rho Z + (1-p)\rho, \quad (4.2)$$

that is, a phase flip is applied to a state ρ with probability p . The bit and phase flip channel [163] is described by

$$\varepsilon(\rho) = pY\rho Y + (1-p)\rho, \quad (4.3)$$

that is, a bit and phase flip is applied to a state ρ with probability p .

4.2.1.2 Phase damping channel

Phase damping is information loss from a quantum system without energy loss. The phase damping channel [163] is described by

$$\varepsilon(\rho) = E_0\rho E_0^\dagger + E_1\rho E_1^\dagger, \quad (4.4)$$

where

$$E_0 = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-\lambda} \end{pmatrix}, \quad E_1 = \begin{pmatrix} 0 & 0 \\ 0 & \sqrt{\lambda} \end{pmatrix}, \quad (4.5)$$

with $\lambda \in [0, 1]$. The advantage of this parameterisation is that it leads to a convenient description of maximal phase damping if we set $\lambda = 1$. This parameterisation is related to the conventional parameterisation of the phase damping channel by

$$e^{-\frac{t}{2T_2}} = \sqrt{1-\lambda}, \quad (4.6)$$

where t is the time and T_2 is the phase damping time constant, so that the phase damping rate is given by $\Gamma_{\text{PD}} = \frac{1}{2T_2}$. The parameter λ in the phase damping channel is related to the parameter p in the phase flip channel by

$$p = \frac{1}{2} \left(1 + \sqrt{1-\lambda} \right). \quad (4.7)$$

4.2.1.3 Amplitude damping channel

Amplitude damping is energy loss from a quantum system. Energy loss occurs when the computational basis state $|1\rangle$ (excited state) decays into the computational basis state $|0\rangle$ (ground state). The amplitude damping channel [163] is described by

$$\varepsilon(\rho) = E_0 \rho E_0^\dagger + E_1 \rho E_1^\dagger, \quad (4.8)$$

where

$$E_0 = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-\lambda} \end{pmatrix}, \quad E_1 = \begin{pmatrix} 0 & \sqrt{\lambda} \\ 0 & 0 \end{pmatrix}, \quad (4.9)$$

with $\lambda \in [0, 1]$. This parameterisation once again has the advantage that we can describe maximal amplitude damping by setting $\lambda = 1$. The parameterisation is related to the conventional parameterisation of the amplitude damping channel by

$$e^{-\frac{t}{2T_1}} = \sqrt{1-\lambda}, \quad (4.10)$$

where t is the time and T_1 is the amplitude damping time constant, so that the amplitude damping rate (decay rate) is given by $\Gamma_{\text{AD}} = \frac{1}{2T_1}$.

4.2.1.4 Depolarising noise channel

Depolarising noise is the simplest noise model for incoherent gate errors on NISQ computers such as the IBM quantum processors [185]. The depolarising channel is described by Eq. (2.27). Recall that in the depolarising channel, a state ρ is replaced by the maximally mixed state with probability p . The depolarising channel can also be written as

$$\varepsilon(\rho) = \frac{p}{3} (X\rho X + Y\rho Y + Z\rho Z) + (1-p)\rho, \quad (4.11)$$

that is, a bit flip, a phase flip and a bit and phase flip are each applied with probability $\frac{p}{3}$. Specialised error mitigation techniques are available to reduce the effect of depolarising noise on quantum computers [181]. However, these methods can only be applied in applications of t -designs where the final outcome is an expectation value.

4.3 Noise modelling

Our models for a noisy t -design rely heavily on the definition of an approximate t -design, as given by inequality (2.34). Recall that an ensemble of unitary operators $\{p_i, U_i\}$ is an ϵ -approximate single-qubit t -design if there exists an ϵ such that for all operators ρ acting on $(\mathbb{C}^2)^{\otimes t}$,

$$(1 - \epsilon) \mathbb{E}_H^t(\rho) \leq \sum_i p_i U_i^{\otimes t} \rho (U_i^{\otimes t})^\dagger \leq (1 + \epsilon) \mathbb{E}_H^t(\rho). \quad (4.12)$$

We note that while there are many state-independent quantifiers of the extent to which a given ensemble of unitary operators deviates from an exact unitary t -design, such as the frame potential [186, 187], it is unclear how these can be applied in the context of noise modelling, since noise channels act on states and cannot be applied to the unitary operators directly. Even though the definition of an approximate single-qubit t -design applies to any density matrix acting on $(\mathbb{C}^2)^{\otimes t}$, we restrict our noise models to density matrices which are t copies of an arbitrary single-qubit density matrix, as was done in our test for an approximate t -design in Sec. 3.3.4. This has two major benefits, namely that numeric results can be obtained efficiently for all t , since

the number of parameters that need to be varied when creating samples of density matrices remains constant with increasing t , and that numeric results can be analysed geometrically, since single-qubit states correspond to points in the Bloch sphere. We therefore quantify the effect of a noise channel ε on the quality of an exact single-qubit t -design $\{p_i, U_i\}$ using the smallest possible ϵ such that the inequality,

$$(1 - \epsilon)\mathbb{E}_H^t(\rho^{\otimes t}) \leq \tilde{\mathbb{E}}_H^t(\rho) \leq (1 + \epsilon)\mathbb{E}_H^t(\rho^{\otimes t}), \quad (4.13)$$

holds for all single-qubit density matrices ρ . This ϵ quantifies the noisy t -design's ability to replicate the moments of the uniform Haar ensemble and represents a lower bound in the more general definition of an approximate single-qubit t -design where ρ can be any density matrix acting on $(\mathbb{C}^2)^{\otimes t}$.

The definition of $\tilde{\mathbb{E}}_H^t(\rho)$ depends on the noise model. Inspired by the noise models typically used in randomised benchmarking [92–100], we consider a noise model in which noise is applied before the unitary operations, for which we define

$$\tilde{\mathbb{E}}_H^t(\rho) = \sum_i p_i \left(U_i \varepsilon(\rho) U_i^\dagger \right)^{\otimes t}, \quad (4.14)$$

as well as a noise model in which noise is applied after the unitary operations, for which we define

$$\tilde{\mathbb{E}}_H^t(\rho) = \sum_i p_i \left(\varepsilon \left(U_i \rho U_i^\dagger \right) \right)^{\otimes t}. \quad (4.15)$$

In both models, the same noise channel ε is applied to each single-qubit state in the t -fold tensor product. Both noise models are well-defined in the sense that the value of ϵ obtained is independent of the choice of ensemble and a general property of the t -design for a given t . This is proven in Appendix B.1.

Based on the fact that for $t \geq 2$, a t -design transforms any noise channel into a depolarising channel [92–94, 98–100], one might expect our two noise models to be equivalent. However, since we are studying the effect of a noise channel on the quality of a t -design, not the effect of a t -design on a noise channel, equivalence of the noise models is a question of whether the noisy Haar ensemble expectations given by Eqs. (4.14) and (4.15) are equal, not a question of whether the resulting depolarising channels are equal for the two models. Our noise models are therefore not generally equivalent. Furthermore, the t -fold tensor product makes it difficult to find a relation between the two noisy Haar ensemble expectations by commuting noise through the unitary operations.

We note that a third noise model, in which noise is applied during the unitary operations, could also be considered due to the finite time duration for these operations in an experimental realisation. However, this is dependent on the method used to implement the unitaries in an experiment (e.g. with control pulses) and so we focus on the former two models which are implementation independent.

4.4 Results

Analytic results for the 1-design are presented in Appendix B.2. For the model where noise is applied before the unitary operations, we were able to show that the quality of the 1-design is completely unaffected by an arbitrary noise channel, and for the model where noise is applied after the unitaries, we showed that the quality of the 1-design is unaffected by noise, unless amplitude damping is applied. Furthermore, we showed that $\epsilon = \lambda$ quantifies the effect of the amplitude damping channel on the quality of the 1-design for the model where noise is applied after the unitaries.

When $t > 1$, $\mathbb{E}_H^t(\rho^{\otimes t})$ depends on the state ρ , which makes it very difficult to obtain results analytically, since inequality (4.13) contains variables other than ϵ and the noise parameter (p or λ) and so it is very difficult

to obtain an expression for ϵ in terms of the noise parameter. Numeric results were therefore obtained for the 2-design, 3-design, 4-design and 5-design. Few exact single-qubit t -designs exist for $t > 5$, and so it is hard to obtain results for $t > 5$. It is also of little interest at present, since there are only a few known applications of t -designs for $t > 4$.

With the notable exception of the amplitude damping channel, numeric results obtained for the 3-design are identical to those obtained for the 2-design and numeric results obtained for the 5-design are identical to those obtained for the 4-design. Hence, in the sections which follow and in the appendices referenced, we only present numeric results for the 2-design and the 4-design, unless amplitude damping was applied. We comment on the identical nature of different t -designs in due course.

4.4.1 Implementation

To obtain numeric results we require samples of single-qubit density matrices. In Appendix A.2, where we numerically investigated the effect of depolarising noise on the quality of the 2-design and the 3-design, we found that the ϵ needed to satisfy inequality (4.13) for pure states is very large even for small values of the depolarising noise parameter p (see Eq. (2.27)). For a comprehensive numerical investigation, we therefore do not simply consider a sample of single-qubit density matrices distributed over the entire Bloch sphere, but rather consider various samples of density matrices restricted to different regions of the Bloch sphere. Opening the study of noise in regions of the Bloch sphere may provide useful information that could be used in specific applications of t -designs, for instance those where the full state space is not required.

To generate a sample of density matrices, we first generate 11 evenly spaced values of $r \in [0, r_t]$, 11 evenly spaced values of $\theta \in [0, \theta_t]$ and 11 evenly spaced values of $\phi \in [0, \phi_t]$, where r_t , θ_t and ϕ_t are the points of truncation of the radial coordinate, the polar angle and the azimuthal angle respectively. Together these truncation points define the region of the Bloch sphere being considered. Unless otherwise stated, $r_t = 1$, $\theta_t = \pi$ and $\phi_t = 2\pi$ was used so that the entire Bloch sphere was considered. We then convert all $11^3 = 1331$ possible combinations of the generated values of the spherical coordinates (r, θ, ϕ) into the cartesian coordinates (x, y, z) , and obtain a sample of 1331 density matrices using Eq. (2.20).

Given $t \in \{2, 3, 4, 5\}$, a noise channel, a noise model and a sample of density matrices we obtain ϵ numerically as follows. For each single-qubit density matrix ρ in the sample, we calculate $\mathbb{E}_H^t(\rho^{\otimes t})$ and $\tilde{\mathbb{E}}_H^t(\rho)$ using the icosahedral group [188] (an exact unitary 5-design and therefore also an exact unitary t -design for any $t \leq 5$) and determine the smallest possible ϵ such that inequality (4.13) is satisfied. The largest ϵ found is the smallest possible ϵ such that inequality (4.13) holds for all density matrices in the sample and is therefore the value with which we quantify the effect of the given noise channel on the quality of the t -design.

4.4.2 Numeric results

This section covers numeric results obtained for the model where noise is applied before the unitary operations. Numeric results obtained for the model where noise is applied after the unitary operations follow similar trends for a given noise channel and are presented in Appendix B.3. The primary difference is that the values of ϵ obtained for a given t and noise parameter (p or λ) are generally slightly smaller for the model where noise is applied after the unitary operations. For the depolarising noise channel, we prove in Appendix B.4 that the values of ϵ obtained for the two noise models are equal, and for the other noise channels, numeric results obtained for the model where noise is applied after the unitary operations reflect the transformation of the noise channel into a depolarising noise channel, an effect exploited in randomised benchmarking with 2-designs [92–94, 98–100].

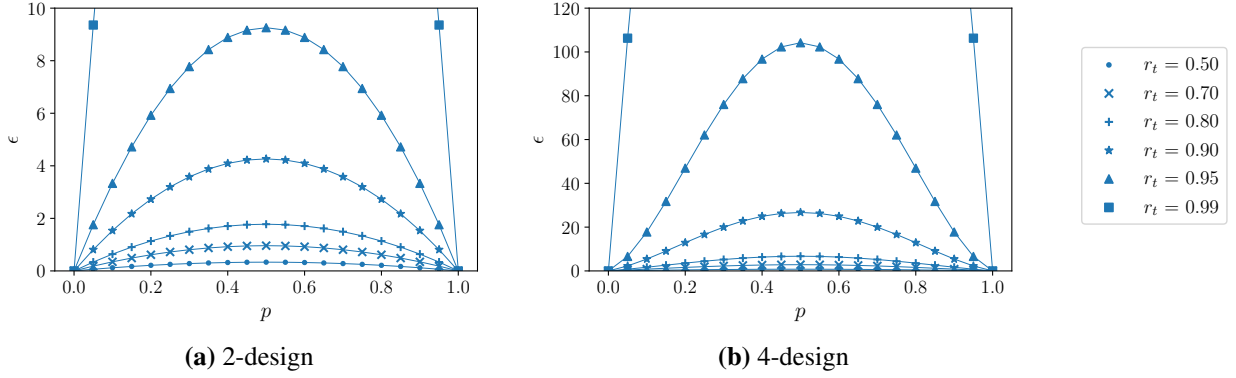


Figure 4.1: Effect of the bit flip channel (see Eq. (4.1)) on the quality of the (a) 2-design and (b) 4-design for the model where noise is applied before the unitary operations, for different truncation radii r_t .




Figure 4.2: ϵ versus t for the bit flip channel with $p = 0.5$ (see Eq. (4.1)) for the model where noise is applied before the unitary operations, for different truncation radii r_t .

4.4.2.1 Flip channels

The effect of the bit flip channel (see Eq. (4.1)) on the quality of the 2-design is shown in Fig. 4.1a. For each truncation radius considered (truncation angles fixed at $\theta_t = \pi$ and $\phi_t = 2\pi$), ϵ versus p is a parabola with maximum at $p = 0.5$. The maxima increase with increasing truncation radius, which shows that as the set of states considered is expanded to include states closer to the pure states at $r_t = 1$, the 2-design becomes more sensitive to bit flips. The symmetry of ϵ versus p around $p = 0.5$ can be explained as follows. Note that states along the x -axis of the Bloch sphere, which are eigenstates of the Pauli X operator, are unaffected by bit flips. However, applying a bit flip to a state off the x -axis and its reflection in the x -axis with probability $p < 0.5$ shifts both states by the same distance towards the x -axis of the Bloch sphere. On the other hand, applying a bit flip to the state and its reflection with probability $p' = (1 - p) > 0.5$ shifts the state across the x -axis, to where its reflection was shifted when applying a bit flip with probability p , and shifts the state's reflection to where the state was shifted when applying a bit flip with probability p . Therefore, applying a bit flip to all states in a sphere of radius r_t with probability p results in the same set of states as applying a bit flip to all states in that sphere with probability $1 - p$. Since ϵ must ensure that inequality (4.13) is satisfied for all states considered, ϵ depends only on the effect of the bit flip channel on the set of states in the sphere considered (not on the effect on individual states). Hence the ϵ computed for a bit flip with probability p is equal to the ϵ computed for a bit flip with probability $1 - p$, so that ϵ versus p is symmetric about $p = 0.5$.

The effect of the bit flip channel on the quality of the 4-design is shown in Fig. 4.1b. For each truncation radius, the maximum of ϵ versus p still occurs at $p = 0.5$, but ϵ versus p now has a more sinusoidal shape. The

values of ϵ obtained for the 4-design are up to an order of magnitude larger than those obtained for the 2-design, for a fixed p and r_t . This shows that the 4-design is significantly more sensitive to bit flips than the 2-design. To visualise the variation in sensitivity, we plot ϵ versus t for $p = 0.5$. As can be seen in Fig. 4.2, ϵ versus t is a step function. There is no increase in sensitivity to bit flips from $t = 2$ to $t = 3$ or from $t = 4$ to $t = 5$, but a significant increase in sensitivity from $t = 3$ to $t = 4$.

Numeric results obtained for the phase flip channel (see Eq. (4.2)) and the bit and phase flip channel (see Eq. (4.3)) are identical to those obtained for the bit flip channel (shown in Figs. 4.1a, 4.1b and 4.2) as they are simply rotated versions of the same channel. To further investigate similarities and differences in the effect of these three channels on the quality of t -designs, we determine and visualise the region of the Bloch sphere (which we will refer to as the region of acceptable quality) for which a noisy t -design is able to replicate the moments of the uniform Haar ensemble, with a predefined accuracy, up to order t . This investigation is presented in Appendix C.1. For each of the three flip channels, we find that the shape of the region of acceptable quality is similar to the shape into which the Bloch sphere is deformed by the relevant channel. For example, the region of acceptable quality is an ellipsoid along the x -axis for the bit flip channel. Since bit flips are performed by applying the Pauli X operator to a state, states along the x -axis, which are closer to the eigenstates of the Pauli X operator, are less affected by bit flips, and so the quality remains acceptable for states along the x -axis even for a large bit flip probability. The regions of acceptable quality for the three flip channels are thus identical up to a rotation, for a fixed p and t , which explains why ϵ versus p is the same for all three flip channels, for a fixed r_t and t , as the full range of θ and ϕ is considered.

To further analyse the dependence of ϵ versus p on the region of the Bloch sphere considered, we vary the truncation of the polar angle θ_t and the truncation of the azimuthal angle ϕ_t . These investigations are included in Appendices C.2 and C.3 respectively. We find that the phase flip channel is the only flip channel for which ϵ versus p has a non-trivial dependence on θ_t . As θ_t is increased from 0 to $\frac{\pi}{2}$, the sample of density matrices is expanded to include states which are further from the eigenstates of the Pauli Z operator, and therefore more sensitive to phase flips, which results in a non-trivial dependence for the phase flip channel. On the other hand, the states along the positive z -axis, which are among the furthest from the eigenstates of the Pauli X and Y operators, and therefore among the most sensitive to bit flips, and bit and phase flips, are included in the sample of density matrices for all θ_t , and so the results for the bit flip channel and the bit and phase flip channel are independent of θ_t . The results are independent of ϕ_t for all three flip channels. For the bit flip channel and the bit and phase flip channel, this can be attributed to the fact that the states along the positive z -axis are included in the sample of density matrices for all ϕ_t . For the phase flip channel, the independence of ϕ_t can be attributed to the fact that the smallest ϵ such that inequality (4.13) holds for a given state remains unchanged when that state is rotated about the z -axis.

4.4.2.2 Phase damping channel

The effect of the phase damping channel (see Eq. (4.4)) on the quality of the 2-design is shown in Fig. 4.3a. For each truncation radius, ϵ increases linearly with λ . The linear relation between ϵ and the parameter λ in the phase damping channel can be attributed to the fact that both ϵ and λ are quadratic functions of the parameter p in the phase flip channel (see Fig. 4.1a and Eq. (4.7) respectively). The gradient of ϵ versus λ increases with increasing truncation radius, similar to the way in which the maximum of ϵ versus p increases with increasing truncation radius for the phase flip channel.

For the 4-design, ϵ versus λ has a more exponential shape (see Fig. 4.3b). For a fixed λ , ϵ versus t is a step function, as shown in Fig. 4.4, just as ϵ versus t is a step function for a fixed phase flip probability p . We also

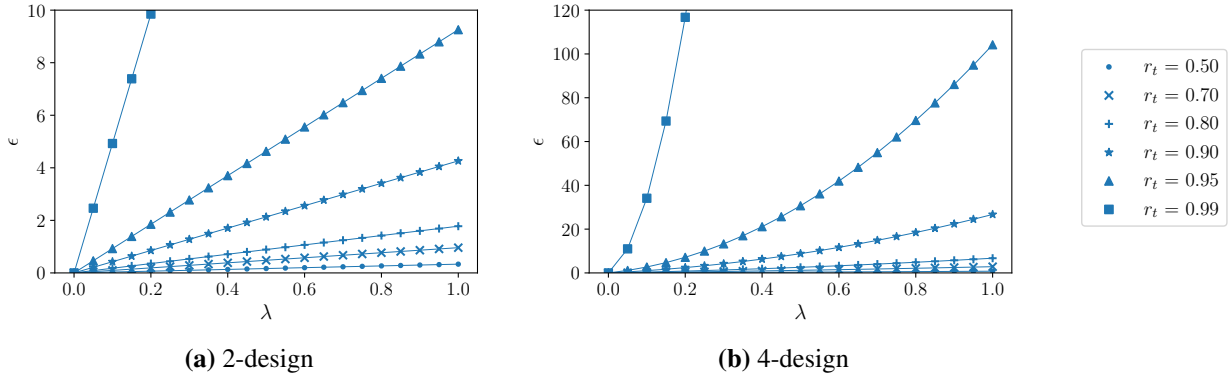


Figure 4.3: Effect of the phase damping channel (see Eq. (4.4)) on the quality of the (a) 2-design and (b) 4-design for the model where noise is applied before the unitary operations, for different truncation radii r_t .




Figure 4.4: ϵ versus t for the phase damping channel with $\lambda = 0.5$ (see Eq. (4.4)) for the model where noise is applied before the unitary operations, for different truncation radii r_t .

investigate variations in sensitivity to phase damping throughout the Bloch sphere. We find that both the shape of the region of acceptable quality and the dependence of ϵ versus λ on θ_t and ϕ_t are similar to that of the phase flip channel (see Appendices C.1, C.2 and C.3 respectively).

4.4.2.3 Amplitude damping channel

The effect of the amplitude damping channel (see Eq. (4.8)) on the quality of the 2-design is shown in Fig. 4.5. For the most part, ϵ versus λ is a parabola with maximum either at or close to $\lambda = 0.5$, but an anomaly occurs for large λ and small r_t , where at a given λ , the trend spontaneously changes to strictly increasing. Just as for the bit flip channel, the maxima of ϵ versus λ increase with increasing truncation radius. The similarities with the bit flip channel are to be expected, considering that the amplitude damping channel actually performs a bit flip on the state $|1\rangle$ with a given probability (the difference being that the state $|0\rangle$ is never flipped by the amplitude damping channel). Bearing in mind that the amplitude damping channel shrinks and shifts any subsphere of states in the Bloch sphere up to the state $|0\rangle$, the anomaly can be interpreted as follows for a given r_t . At the turning point of ϵ versus λ , the south pole of the shifted sphere crosses that sphere's initial equator. The anomaly, where the trend changes to strictly increasing, occurs at the point where the south pole of the shifted sphere crosses its initial north pole. In the limit $\lambda \rightarrow 1$ (maximal amplitude damping), all spheres are reduced to the state $|0\rangle$, which explains why $\epsilon = 1$ for all r_t at $\lambda = 1$.

Just as for the other noise channels, numeric results obtained for the 3-design are identical to those obtained for the 2-design. The effect of the amplitude damping channel on the quality of the 4-design is shown in

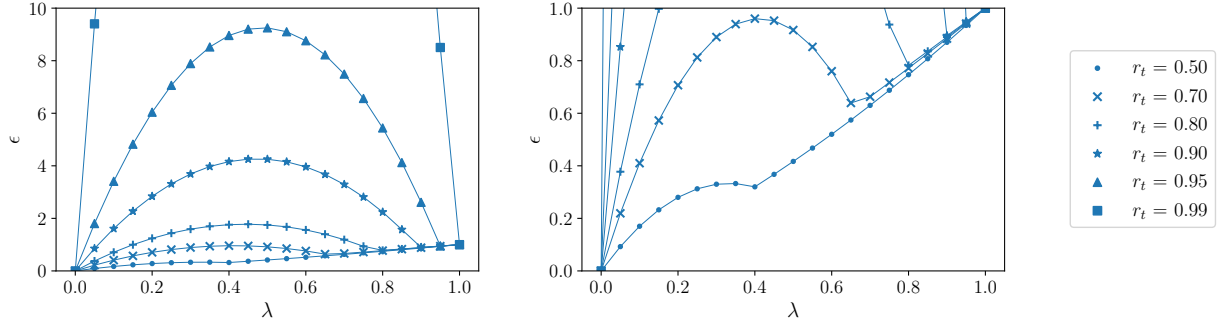


Figure 4.5: Effect of the amplitude damping channel (see Eq. (4.8)) on the quality of the 2-design for the model where noise is applied before the unitary operations, for different truncation radii r_t . The full set of results is shown on the left and the region in which the anomaly occurs is shown enlarged on the right.

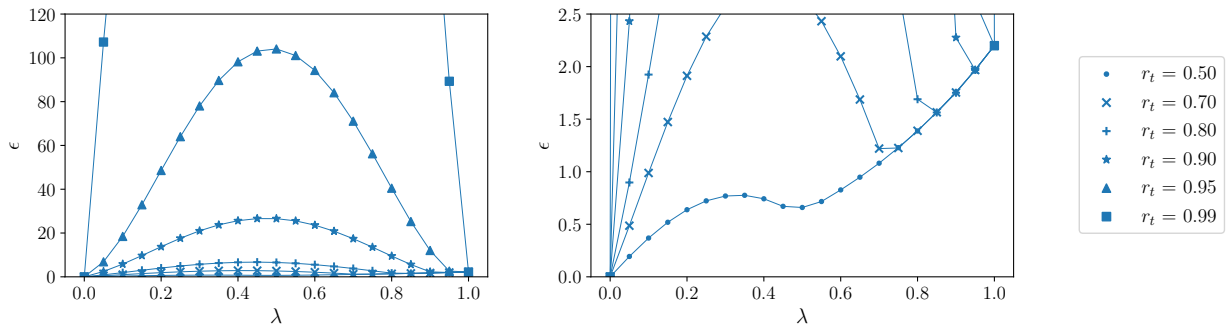


Figure 4.6: Effect of the amplitude damping channel (see Eq. (4.8)) on the quality of the 4-design for the model where noise is applied before the unitary operations, for different truncation radii r_t . The full set of results is shown on the left and the region in which the anomaly occurs is shown enlarged on the right.

Fig. 4.6. The maxima of ϵ versus λ occur at the same values of λ as for the 2-design, and the anomaly still occurs for large λ and small r_t , but ϵ versus λ has a more sinusoidal shape. Numeric results obtained for the 5-design are almost identical to those obtained for the 4-design. The only difference is that for the 4-design ϵ increases to 2.20 as $\lambda \rightarrow 1$, whereas for the 5-design ϵ increases to 4.33 as $\lambda \rightarrow 1$. We compare the values of ϵ obtained for different t -designs, for a fixed λ and r_t , by plotting ϵ versus t for different truncation radii, each time using the turning point of ϵ versus λ as our fixed value of λ (see Fig. 4.7). We find that ϵ versus t is once again a step function and see a significant increase in sensitivity to amplitude damping from $t = 3$ to $t = 4$. However, we note that for larger λ and smaller r_t there is also a slight increase in sensitivity to amplitude damping from $t = 4$ to $t = 5$. The region of acceptable quality for the amplitude damping channel, as well as the dependence of ϵ versus λ on θ_t and ϕ_t is analysed in Appendices C.1, C.2 and C.3 respectively.

4.4.2.4 Depolarising noise channel

Finally, the effect of the depolarising channel (see Eq. (2.27)) on the quality of the 2-design is shown in Fig. 4.8a. Similar results are included in Appendix A.2. For each truncation radius considered, ϵ increases linearly with p , up to about $p = 0.4$, after which the increase becomes more gradual. In Fig. 4.8b, we see that for the 4-design, ϵ first increases rapidly with p , then increases linearly with p for $p \in [0.3, 0.6]$, after which the increase becomes more gradual. Just as for all the other noise channels, the values of ϵ obtained for the 4-design are up to an order of magnitude larger than those obtained for the 2-design, for a fixed p and r_t , once again confirming that the 4-design is significantly more sensitive to noise than the 2-design. As expected, ϵ versus t is a step function for a fixed p and r_t (see Fig. 4.9).




Figure 4.7: ϵ versus t for the amplitude damping channel, with the parameter λ (see Eq. (4.8)) taken to be the turning point of ϵ versus λ , for the model where noise is applied before the unitary operations, for different truncation radii r_t .

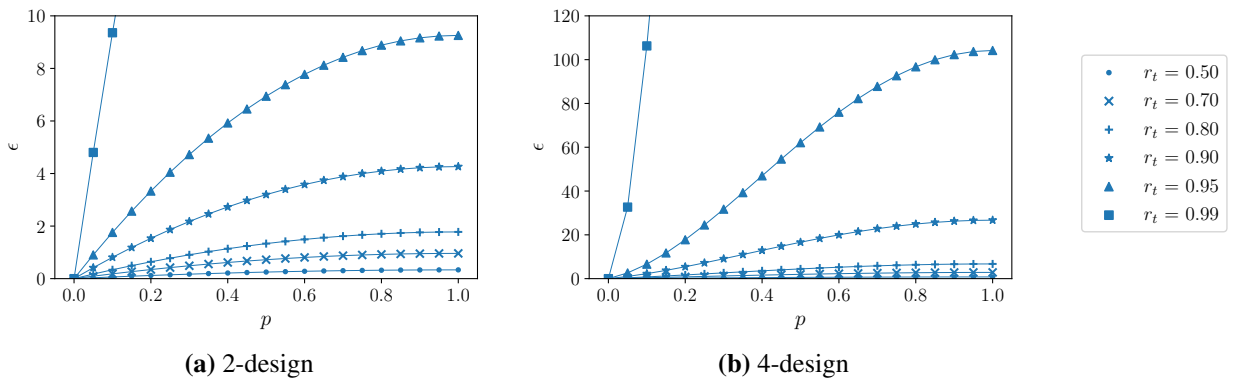


Figure 4.8: Effect of the depolarising noise channel (see Eq. (2.27)) on the quality of the (a) 2-design and (b) 4-design for the model where noise is applied before the unitary operations, for different truncation radii r_t .

As illustrated in Appendix C.1, the region of acceptable quality for the depolarising noise channel has a spherical shape. As such, the numeric results obtained for ϵ versus p are independent of θ_t and ϕ_t (see Appendices C.2 and C.3 respectively).

4.5 Conclusion

We studied the effect of different types of noise on the quality of single-qubit t -designs. The noise channels we investigated were the bit flip channel, the phase flip channel, the bit and phase flip channel, the phase damping channel, the amplitude damping channel and the depolarising noise channel. We quantified the effect of a noise channel on the quality of a t -design using the smallest possible ϵ such that a test inequality, adapted from the defining inequality for an ϵ -approximate t -design, holds for all density matrices in a given sample. Two noise models were considered, namely a noise model in which noise is applied before the unitary operations and a noise model in which noise is applied after the unitary operations, in line with the noise models used in randomised benchmarking [92–100].

We showed analytically that for the model where noise is applied before the unitary operations, the quality of the 1-design is completely unaffected by an arbitrary noise channel, and for the model where noise is applied after the unitaries, the quality of the 1-design is unaffected by noise, unless amplitude damping is applied (see Appendix B.2). For the 2-design, 3-design, 4-design and 5-design, results were obtained numerically using the icosahedral group [188]. With the exception of the amplitude damping channel, ϵ versus t is a step function




Figure 4.9: ϵ versus t for the depolarising noise channel with $p = 0.5$ (see Eq. (2.27)) for the model where noise is applied before the unitary operations, for different truncation radii r_t .

for a fixed p or λ . We see a significant increase in sensitivity to noise from $t = 1$ to $t = 2$ and an even larger increase in sensitivity to noise from $t = 3$ to $t = 4$, but no increase in sensitivity to noise from $t = 2$ to $t = 3$ or from $t = 4$ to $t = 5$, unless amplitude damping is applied. Based on these results, we conjecture that for any t , a $(2t + 1)$ -design is as sensitive to noise as a $2t$ -design, for any noise channel which deforms the Bloch sphere, but does not shift the Bloch sphere. We note that it may be possible to prove this with induction on t using recently discovered random circuit constructions for exact t -designs [109], but such a proof evaded the author.

For all the noise channels considered and for both noise models, ϵ increases with increasing truncation radius, for a fixed t and noise parameter (p or λ). Hence t -designs become increasingly sensitive to noise as the set of states considered is expanded to include states further from the maximally mixed state at $r_t = 0$ (for which the sensitivity to noise is least) and closer to the pure states at $r_t = 1$ (for which the sensitivity to noise is greatest). To further investigate variations in sensitivity to noise throughout the Bloch sphere, we determined the region of acceptable quality (region of the Bloch sphere for which a noisy t -design is able to replicate the moments of the uniform Haar ensemble, with a predefined accuracy, up to order t) for each of the noise channels (see Appendix C.1). For the model where noise is applied before the unitary operations, the shape of the region of acceptable quality for each noise channel is similar to the shape into which the Bloch sphere is deformed by the relevant noise channel.

For the model where noise is applied after the unitary operations, the region of acceptable quality has a spherical shape for all the noise channels considered. Hence our numeric results reflect the transformation of a noise channel into a depolarising channel, an effect exploited in randomised benchmarking with 2-designs [92–94, 98–100], when the noise is applied after the unitary operations. For the depolarising noise channel, our two noise models are equivalent (proven in Appendix B.4). For the other noise channels, t -designs generally show reduced sensitivity to noise for the model where noise is applied after the unitary operations (see Appendix B.3), which seems to suggest that the process by which a noise channel is transformed into a depolarising channel (so the quality of t -designs is affected equally for all states at a given radial distance from the maximally mixed state) mitigates the effect of the noise channel on the quality of t -designs.

Chapter 5

Measurement-based interleaved randomised benchmarking using IBM processors

5.1 Introduction

Noise is not limited to experimental realisations of t -designs, but is present in experimental realisations of any quantum information processing protocol on current hardware. Characterising the noise is important, since noise is often the dominant factor preventing the successful realisation of sophisticated quantum algorithms or protocols on NISQ computers [17, 177, 185]. A complete characterisation of errors on a quantum computer can be obtained by performing quantum process tomography (see Sec. 3.2.3) for all its elementary operations or hardware implemented gates and calculating the associated gate fidelities [173, 174, 189, 190]. However, this is very resource intensive, since the number of experiments that need to be performed grows exponentially with the number of qubits. This method also requires that state preparation and measurement (SPAM) errors are negligible, which is rarely the case in current hardware. Randomised benchmarking, which provides a partial characterisation of errors, avoids both these problems [191].

Standard randomised benchmarking can be used to estimate the average gate fidelity of a set of gates on a quantum computer (usually the Clifford group or a subset thereof) [92–94, 98, 192–199]. Interleaved randomised benchmarking can be used to estimate the fidelity of individual Clifford gates [95, 119, 120]. In interleaved randomised benchmarking, noise parameters are estimated for sequences of randomly chosen Clifford gates and sequences of the Clifford gate of interest interleaved between random Clifford gates. The inferred noise parameters are then used to estimate the fidelity of the Clifford gate of interest. Special interleaved randomised benchmarking protocols for estimating the fidelity of non-Clifford gates such as the T gate have been proposed [96, 121]. Since the Clifford gates together with the T gate form a universal set [122], these protocols enable fidelity estimation of individual gates from a universal set. Note that in this chapter, the term fidelity or gate fidelity refers to the Haar-averaged gate fidelity defined in Sec. 5.2.3 and not to the channel fidelity defined in Sec. 3.2.3.

Both standard and interleaved randomised benchmarking have been implemented in a great variety of physical systems. These include trapped ions [119, 196, 200], superconducting systems [95, 201–204], nuclear magnetic resonance quantum processors [205], cold atoms [206, 207] and quantum dots [208]. Several variations of randomised benchmarking have also recently been demonstrated in trapped ions [209] and superconducting systems [109, 210].

Simple adaptations of standard randomised benchmarking have been suggested for measurement-based

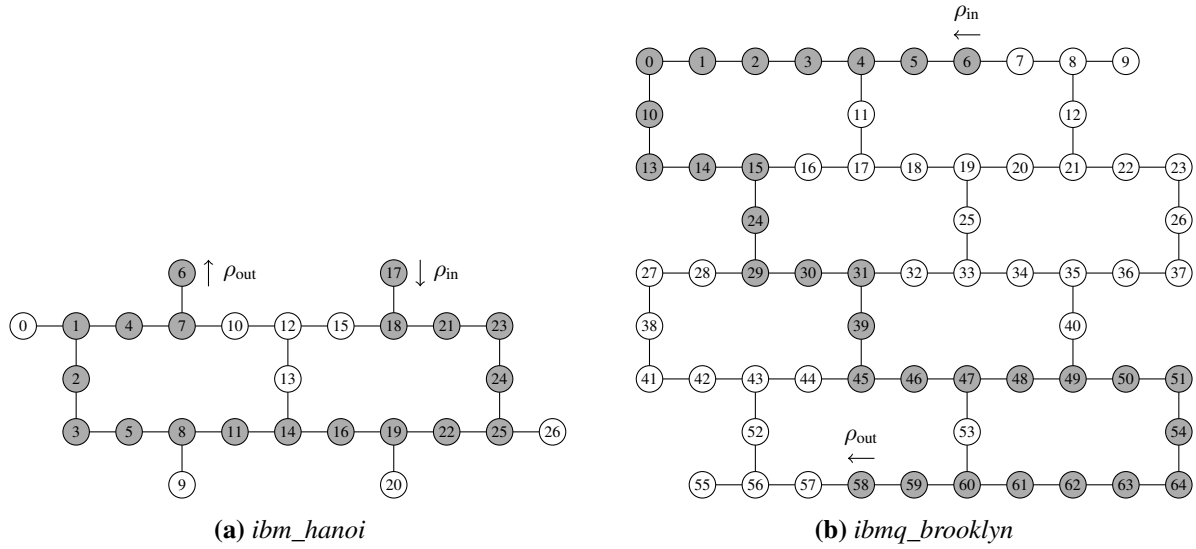


Figure 5.1: Qubit topologies of the processors used in our demonstration, with the qubits used shaded grey.

quantum computers [97]. In this chapter, we expand on the work of Ref. [97] and propose an interleaved randomised benchmarking protocol for measurement-based quantum computers. In our measurement-based interleaved randomised benchmarking protocol, any single-qubit measurement-based 2-design can be used to estimate the fidelity of any single-qubit measurement-based gate. We test our protocol by using the approximate measurement-based 2-design proposed and studied in Sec. 3.3.5 to estimate the fidelity of the Hadamard gate and the T gate on IBM superconducting quantum computers. Even though IBM quantum computers are primarily optimised for quantum computing in the circuit model, a variety of measurement-based protocols have been successfully implemented on these systems (see Chapter 3 and Refs. [177, 211]). The IBM quantum computers were chosen for our experiments, since systems optimised for quantum computing in the measurement-based model, such as PsiQuantum’s photonic systems, are not readily accessible and other cloud-based systems optimised for the circuit model, such as IonQ’s trapped ions, had too few qubits to prepare the large entangled resource states required for our protocol at the time of starting the experiments.

Since the Hadamard gate and the T gate form a universal single-qubit set [122], our experiments provide a proof-of-concept demonstration of fidelity estimation of measurement-based gates from a universal single-qubit set. In our demonstration, we use entangled resource states of up to 19 qubits on the *ibm_hanoi* quantum processor (see Fig. 5.1a) and entangled resource states of up to 31 qubits on the *ibmq_brooklyn* quantum processor (see Fig. 5.1b). In Chapter 3, where we implemented single-qubit measurement-based t -designs on IBM processors, the implementations were not tested on any application and the entangled resource states used did not exceed six qubits. The work in this chapter therefore demonstrates significant progress in practical quantum computing on superconducting systems in the measurement-based model. In all our experiments, estimated gate fidelities show good agreement with gate fidelities calculated from process tomography results.

This chapter is structured as follows. In Sec. 5.2, we summarise single-qubit measurement-based quantum computing, single-qubit measurement-based t -designs and interleaved randomised benchmarking. In Sec. 5.3, we present our measurement-based interleaved randomised benchmarking protocol. In Sec. 5.4, we discuss adjustments to the protocol that are required for implementation on the IBM quantum processors. A description of the experiments performed and the results obtained is presented in Sec. 5.5. A summary of the chapter and concluding remarks are given in Sec. 5.6. Further details about the experiments are given in Appendix D.

5.2 Background

5.2.1 Measurement-based quantum computing

We reviewed single-qubit measurement-based quantum computing with linear cluster states in Sec. 2.3.2. Recall that, when all measurement outcomes are zero, the 2-qubit linear cluster state with the measurement angle $\phi_1 = 0$ implements the Hadamard gate and the 3-qubit linear cluster state with the measurement angles $\phi_1 = \frac{\pi}{4}$ and $\phi_2 = 0$ implements the T gate. When some of the measurement outcomes are non-zero, the desired gate can still be realised by applying the appropriate Pauli correction to the final qubit [161].

5.2.2 Measurement-based t -designs

The definitions of an exact and an ϵ -approximate unitary t -design are given in Sec. 2.5. In this chapter, we are primarily interested in unitary 2-designs, which are sufficient for randomised benchmarking [92]. Single-qubit measurement-based t -designs were introduced in Sec. 3.2.1. Recall that the 6-qubit linear cluster state with the measurement angles $\phi_1 = 0$, $\phi_2 = \frac{\pi}{4}$, $\phi_3 = \arccos \sqrt{1/3}$, $\phi_4 = \frac{\pi}{4}$ and $\phi_5 = 0$ implements an exact single-qubit 2-design [88] and the 5-qubit linear cluster state with the measurement angles $\phi_1 = 0$, $\phi_2 = \frac{\pi}{4}$, $\phi_3 = \frac{\pi}{4}$ and $\phi_4 = 0$ implements an approximate single-qubit 2-design with $\epsilon = 0.5$ (see Sec. 3.3.5). In Chapter 3, we implemented both the exact measurement-based 2-design and the approximate measurement-based 2-design on IBM processors. Neither implementation passed our test for a 2-design under the test conditions set. However, the test results showed that the approximate 2-design implementation more closely resembled a 2-design than the exact 2-design implementation as a result of reduced noise for the smaller 5-qubit cluster state. This is why we use the approximate 2-design, and not the exact 2-design, in our randomised benchmarking experiments in this chapter (presented in Sec. 5.5). We note that the implications of performing randomised benchmarking with an ϵ -approximate 2-design, as opposed to an exact 2-design, are not yet well understood theoretically. However, numerical investigations have shown that the estimated fidelity obtained when using an exact 2-design can differ from the estimated fidelity obtained when an exact 2-design is not used [197]. We find that the results obtained with the approximate 2-design are consistent with those obtained using quantum process tomography.

5.2.3 Interleaved randomised benchmarking

Unitary 2-designs can be used in interleaved randomised benchmarking, which provides an estimate of the Haar-averaged fidelity of a noisy implementation of a unitary operation or gate [95]. The Haar-averaged fidelity of a noisy implementation of an ideal gate G is defined by

$$F_G(\varepsilon, \tilde{\varepsilon}) = \int \left(\text{Tr} \left(\sqrt{\sqrt{\varepsilon}(|\psi\rangle\langle\psi|)\tilde{\varepsilon}(|\psi\rangle\langle\psi|)\sqrt{\varepsilon}(|\psi\rangle\langle\psi|)} \right) \right)^2 d\psi, \quad (5.1)$$

where $\varepsilon(|\psi\rangle\langle\psi|) = G|\psi\rangle\langle\psi|G^\dagger$ is the channel representing the ideal implementation of G and $\tilde{\varepsilon}(|\psi\rangle\langle\psi|)$ is the channel representing the noisy experimental implementation of G [97]. The average gate error is then given by $1 - F_G(\varepsilon, \tilde{\varepsilon})$, which is useful for quantifying the overall reliability of the implementation. However, for some applications, such as determining thresholds for fault-tolerant quantum computing, the worst case error is required [212]. For these applications, the average gate error can be used to obtain bounds on the worst case error [190, 213–215].

We briefly review the interleaved randomised benchmarking protocol proposed by Magesan *et al.* [95], in which the fidelity of individual Clifford gates [216] can be estimated. However, we relax the restriction to

Clifford gates and review a variation of the protocol in which any single-qubit unitary 2-design \mathcal{U} can be used to estimate the fidelity of any single-qubit unitary operation or gate G . Section III A of Ref. [217] explains that the interleaved method of Ref. [95] holds for the variation of the protocol reviewed here. The only benefit of the restriction to Clifford gates is that the inverse of a sequence of Clifford gates can be efficiently computed on a classical computer as a result of the Gottesman-Knill theorem [218]. Restricting the protocol to single-qubit gates has the same benefit, since any single-qubit system can be efficiently simulated on a classical computer [161]. The key assumptions of the protocol are that noise from any gate is time-independent, noise from the 2-design \mathcal{U} is gate-independent and noise from the inverse of any sequence of gates is independent of the sequence.

Randomised benchmarking relies heavily on a unitary 2-design's ability to transform an arbitrary noise channel into a depolarising channel [93–95], a property which was studied extensively in Chapter 4. In interleaved randomised benchmarking, two experiments are performed, namely an experiment to determine the reference depolarising noise parameter p_{ref} and an experiment to determine the interleaved depolarising noise parameter p_{int} . These parameters are then used to estimate the Haar-averaged gate fidelity using

$$F_G \approx 1 - \frac{d-1}{d} \left(1 - \frac{p_{\text{int}}}{p_{\text{ref}}} \right), \quad (5.2)$$

where $d = 2$ for single qubits.

The reference depolarising noise parameter p_{ref} is determined as follows:

1. Prepare the qubit in an arbitrary, but fixed, initial state $\rho = |\psi\rangle\langle\psi|$.
2. Choose m unitary operators uniformly at random from the 2-design \mathcal{U} and apply the resulting sequence of operators, $U_m \cdots U_2 U_1$, to the state ρ .
3. Compute and apply the inverse of the sequence of gates applied in step 2.
4. Measure the qubit in the basis $\{|\psi\rangle\langle\psi|, I - |\psi\rangle\langle\psi|\}$. This measures the probability that the initial state is unchanged by the sequence applied in step 2 followed by its inverse (known as the survival probability).
5. Repeat steps 1 to 4 for different sequences of a fixed length m and average the survival probability over the different sequences to obtain the sequence fidelity $F_{\text{ref}}(m)$ for a given sequence length m . Guidance on choosing the number of repetitions is given in Refs. [213, 219].
6. Repeat steps 1 to 5 for different sequence lengths m and extract the reference depolarising noise parameter p_{ref} by fitting the resulting data for $F_{\text{ref}}(m)$ versus m to the exponential decay model

$$F_{\text{ref}}(m) = A_{\text{ref}} p_{\text{ref}}^m + B_{\text{ref}}. \quad (5.3)$$

The parameters A_{ref} and B_{ref} absorb SPAM errors, as well as the error of the inverse applied in step 3. The relation between the number of repetitions chosen in step 5 and confidence intervals for the extracted parameters is discussed in Ref. [213].

The interleaved depolarising noise parameter p_{int} is determined as follows:

1. Prepare the qubit in the same fixed initial state $\rho = |\psi\rangle\langle\psi|$.
2. Choose m unitary operators uniformly at random from the 2-design \mathcal{U} and apply the interleaved sequence $GU_m \cdots GU_2 GU_1$ to the state ρ .

3. Compute and apply the inverse of the sequence of gates applied in step 2.
4. Measure the qubit in the basis $\{|\psi\rangle\langle\psi|, I - |\psi\rangle\langle\psi|\}$. This once again measures the survival probability.
5. Repeat steps 1 to 4 for a fixed m and average the survival probability over the different sequences to obtain the sequence fidelity $F_{\text{int}}(m)$.
6. Repeat steps 1 to 5 for different m and extract the interleaved depolarising noise parameter p_{int} by fitting the resulting data for $F_{\text{int}}(m)$ versus m to the exponential decay model

$$F_{\text{int}}(m) = A_{\text{int}} p_{\text{int}}^m + B_{\text{int}}. \quad (5.4)$$

Then, using Eq. (5.2) together with p_{ref} and p_{int} we obtain the gate fidelity F_G .

5.3 Measurement-based interleaved randomised benchmarking protocol

We now present an interleaved randomised benchmarking protocol for measurement-based quantum computers (which is an extension of the standard randomised benchmarking protocol proposed for measurement-based quantum computers by Alexander *et al.* [97]). In particular, we explain how the experiments to determine the reference and interleaved depolarising noise parameters (described in Sec. 5.2.3) can be implemented on a single-qubit measurement-based quantum computer. To this end, let \mathcal{U} be a single-qubit measurement-based 2-design implemented by a $(k+1)$ -qubit linear cluster state with the fixed measurement angles $\phi = (\phi_1, \dots, \phi_k)$ and let G be a single-qubit measurement-based gate implemented by a $(\ell+1)$ -qubit linear cluster state with the measurement angles $\theta = (\theta_1, \dots, \theta_\ell)$ (possibly requiring both adaptive measurement feedforward and a Pauli correction).

On a measurement-based quantum computer, the reference depolarising noise parameter p_{ref} can be determined as follows:

1. For a given sequence length m , prepare a $(mk+1)$ -qubit linear cluster state. This chooses the initial state to be the Pauli X -basis state $\rho = |+\rangle\langle+|$, which is the natural choice for measurement-based quantum computing.
2. Repeat the measurement pattern $\phi = (\phi_1, \dots, \phi_k)$ m times along the length of the cluster state. This implements the desired random sequence, $U_m \cdots U_2 U_1$, where the inherent randomness of the measurement-based process ensures that each U_i is chosen uniformly at random from the 2-design \mathcal{U} . Since the measurement angles are fixed, the mk measurements can be performed simultaneously. The measurement outcomes can be used to determine which sequence of gates was implemented.
3. Compute the inverse of the sequence of gates implemented in step 2 and apply this inverse by performing a measurement basis rotation on the final qubit. The inverse is applied in this way, and not as a measurement-based operation, to ensure that noise from the inverse is independent of the sequence.
4. Measure the final qubit in the Pauli X -basis.
5. Repeat steps 1 to 4 for a fixed sequence length m and determine $F_{\text{ref}}(m)$.
6. Repeat steps 1 to 5 for different sequence lengths m and fit the resulting data for $F_{\text{ref}}(m)$ versus m to Eq. (5.3).

On a measurement-based quantum computer, the interleaved depolarising noise parameter p_{int} can be determined as follows:

1. For a given m , prepare a $(m(k + \ell) + 1)$ -qubit linear cluster state.
2. Repeat the measurements $(\phi_1, \dots, \phi_k, \theta_1, \dots, \theta_\ell)$ m times along the length of the cluster state. This implements the desired interleaved sequence, $GU_m \cdots GU_2GU_1$. Since adaptive measurement feedforward may be required for the implementation of G , the $m(k + \ell)$ measurements need to be performed sequentially. If Pauli corrections are required, these can be applied after each set of $k + \ell$ measurements, by performing a measurement basis rotation, before proceeding with the next set of $k + \ell$ measurements.
3. Compute the inverse of the sequence of gates implemented in step 2 and apply this inverse by performing a measurement basis rotation on the final qubit.
4. Measure the final qubit in the Pauli X -basis.
5. Repeat steps 1 to 4 for a fixed m and determine $F_{\text{int}}(m)$.
6. Repeat steps 1 to 5 for different m and fit the resulting data for $F_{\text{int}}(m)$ versus m to Eq. (5.4).

5.4 Adjustments to protocol for implementation on IBM processors

Since IBM quantum processors did not support dynamic circuit execution at the time of performing the experiments, our measurement-based interleaved randomised benchmarking protocol had to be adjusted to enable implementation on the IBM hardware available at the time. Without dynamic circuit execution, it is not possible to include the inverse (step 3 of the experiments to determine the depolarising noise parameters) in the required quantum circuits, since the inverse depends on the sequence being implemented in step 2, which is only known once the measurements have been performed. One solution is to construct a different circuit for each possible inverse, run each circuit a sufficient number of times to obtain the random sequence corresponding to the implemented inverse, and then use post-selection to eliminate runs in which this desired sequence was not obtained. Since there are 2^n possible inverses for the sequences implemented by performing measurements on a $(n + 1)$ -qubit linear cluster state, we would need 2^n different circuits, and since each random sequence occurs with probability $\frac{1}{2^n}$, all but 1 in every 2^n runs would be eliminated, and each of the 2^n circuits would need to be run at least 2^n times to have a reasonable chance of obtaining the sequence corresponding to the implemented inverse. Hence the number of runs required for post-selection scales like 2^{2n} , which is generally prohibitively expensive.

We therefore employ an alternative strategy in our randomised benchmarking experiments in Sec. 5.5. We perform full quantum state tomography (see Sec. 3.2.2) on the final qubit (after performing the measurements in step 2 of the experiments to determine the depolarising noise parameters) for each of the possible measurement outcomes, apply the correct inverse for each implemented sequence to the appropriate constructed density matrix by performing matrix multiplication, and then extract the survival probability from each resulting density matrix. Since each of the 2^n random sequences implemented by performing measurements on a $(n + 1)$ -qubit linear cluster state occurs with probability $\frac{1}{2^n}$, the associated circuit must be run $3(500)(2^n)$ times to obtain 500 data points for tomography (which requires three basis measurements) for each of the 2^n possible measurement outcomes. Hence the circuit must be run $3(500)(2^n)/8000 = 3(2^n)/16$ times if each run has 8000 shots. In the experiments in Sec. 5.5, 8192 shots (the maximum number of allowed shots on IBM quantum processors at the

time of starting the experiments) were used for each run, instead of just 8000 shots, to increase the likelihood of obtaining at least 500 data points for tomography even when the measurement outcomes are not uniform as a result of noise. Although this method has the disadvantage that the inverse is performed through classical post-processing, and not as a quantum mechanical operation, it scales like 2^n , which is much better than post-selection, which scales like 2^{2n} . Since we assume that noise from the inverse is independent of the sequence, noise from the inverse would in any case not affect the depolarising noise parameters which are used to estimate the fidelity, and so performing the inverse through classical post-processing does not affect the results.

Since adaptive measurement feedforward cannot be implemented without dynamic circuit execution either, our protocol could only be used to estimate the fidelity of measurement-based gates which can be implemented with fixed measurement angles and a Pauli correction (such as the 2-qubit Hadamard gate and the 3-qubit T gate given in Sec. 5.2.1) on the IBM quantum hardware available at the time of performing the experiments. As the Pauli corrections required in the interleaved sequences depend on the measurement outcomes, these also cannot be performed without dynamic circuit execution. We therefore average the survival probability, not only over the different random unitaries, but also over the different byproducts that result from omitting the Pauli corrections, when calculating the interleaved sequence fidelity in the experiments in Sec. 5.5. When applying the inverse of an interleaved sequence with matrix multiplication, the inverse of each individual byproduct is applied, as determined by the measurement outcomes.

5.5 Experiments

5.5.1 Implementation of protocol

We implemented our measurement-based interleaved randomised benchmarking protocol, with the adjustments discussed in Sec. 5.4, first on the *ibm_hanoi* quantum processor, a 27-qubit superconducting IBM quantum computer. Details of the qubits used can be found in Appendix D.1. We used the 5-qubit approximate measurement-based 2-design proposed in Sec. 3.3.5 to estimate the fidelity of the 2-qubit measurement-based implementation of the Hadamard gate and the 3-qubit measurement-based implementation of the T gate given in Sec. 5.2.1 on the *ibm_hanoi* quantum processor. To this end, we implemented reference sequences with lengths $m \in \{1, 2, 3\}$ and interleaved sequences with $m \in \{1, 2, 3\}$. For each m (reference or interleaved), we prepared the required linear cluster state, performed the appropriate single-qubit measurements on all but the final qubit, and then performed full quantum state tomography on the final qubit for each of the possible measurement outcomes to infer the output state for each corresponding random sequence. As an example, the quantum circuit for the implementation of the interleaved sequence for the 3-qubit T gate with $m = 1$, which requires a 7-qubit linear cluster state, is shown in Fig. 5.2. The largest number of qubits was used for the implementation of the interleaved sequence for the 3-qubit T gate with $m = 3$, which required $(m(k + \ell) + 1) = (3(4 + 2) + 1) = 19$ qubits, as shown in Fig. 5.1a.

Since controlled phase gates are not supported at the hardware level on IBM quantum processors, linear cluster states were prepared using Hadamards and controlled not gates, instead of Hadamards and controlled phase gates, as recommended by Mooney *et al.* [220]. This prevented redundant Hadamards, which would have unnecessarily increased noise, from being introduced during transpilation. Since qubits can only be measured in the computational basis on IBM processors, single-qubit measurements at an angle ϕ in the Pauli XY plane were carried out by applying $R_z(\phi)$, followed by a Hadamard, and measuring the qubit in the computational basis.

To obtain the data required to determine the sequence fidelity ($F_{\text{ref}}(m)$ or $F_{\text{int}}(m)$) for a given m , the

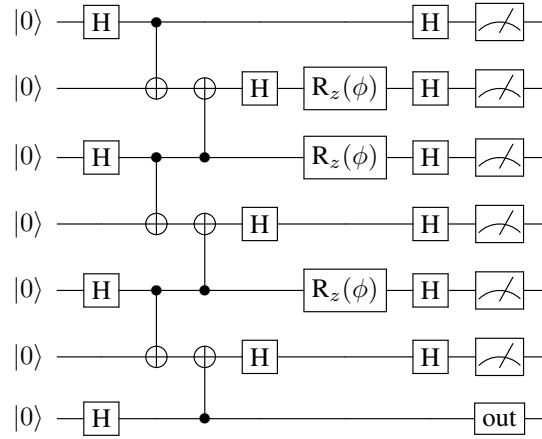


Figure 5.2: Quantum circuit for the implementation of the interleaved sequence for the 3-qubit T gate with $m = 1$ on the *ibm_hanoi* quantum processor. The angle $\phi = \frac{\pi}{4}$ represents the different basis measurements used and ‘out’ is the final qubit on which quantum state tomography is performed.

relevant circuit was run $3(2^n)/16$ times with 8192 shots (see Sec. 5.4) on the *ibm_hanoi* quantum processor (where n is the number of single-qubit measurements performed on the $(n + 1)$ -qubit linear cluster state in the implementation). The data (counts) obtained in repeated runs were combined and grouped according to measurement outcomes. The output density matrix for each random sequence was then constructed from the tomography data obtained for the corresponding set of measurement outcomes. To ensure that the density matrices constructed from the data are physical (i.e. that they have unit trace and are positive) we employed qiskit’s built-in method [180], which uses maximum-likelihood estimation to find the closest physical density matrix to a density matrix constructed from tomography data. The appropriate inverse was then applied to each density matrix and the survival probability was extracted from each resulting density matrix. The sequence fidelity ($F_{\text{ref}}(m)$ or $F_{\text{int}}(m)$) presented for a given m in Sec. 5.5.2 is the average of these survival probabilities, and the error is given by the standard deviation.

In interleaved randomised benchmarking experiments, sequence fidelities are typically determined for sequence lengths up to $m = 80$ or longer [95, 203]. Here, the resources required to implement longer sequences prevented us from considering sequence lengths beyond $m = 3$. For example, the implementation of the interleaved sequence for the 3-qubit T gate with $m = 3$ required a 19-qubit linear cluster state and the relevant circuit had to be run $3(2^{18})/16 = 49152$ times to obtain the data required to determine the sequence fidelity. Since this is already extremely resource intensive, longer sequences would not be feasible, as the number of runs required grows exponentially with the length of the sequence. Hence, even though the adjustments discussed in Sec. 5.4 have enabled us to obtain data for a proof-of-concept demonstration, dynamic circuit execution remains essential for the efficient scaling of our measurement-based interleaved randomised benchmarking protocol in future.

5.5.2 Results for universal gates

Sequence fidelities obtained to estimate the fidelity of the 2-qubit measurement-based implementation of the Hadamard gate and the 3-qubit measurement-based implementation of the T gate (a universal single-qubit set), using the 5-qubit approximate measurement-based 2-design proposed in Sec. 3.3.5, on the *ibm_hanoi* quantum processor, are shown in Fig. 5.3. The large uncertainties in the sequence fidelities reflect the gate-dependence of noise from the approximate measurement-based 2-design (see Sec. 3.3.5). A Monte Carlo

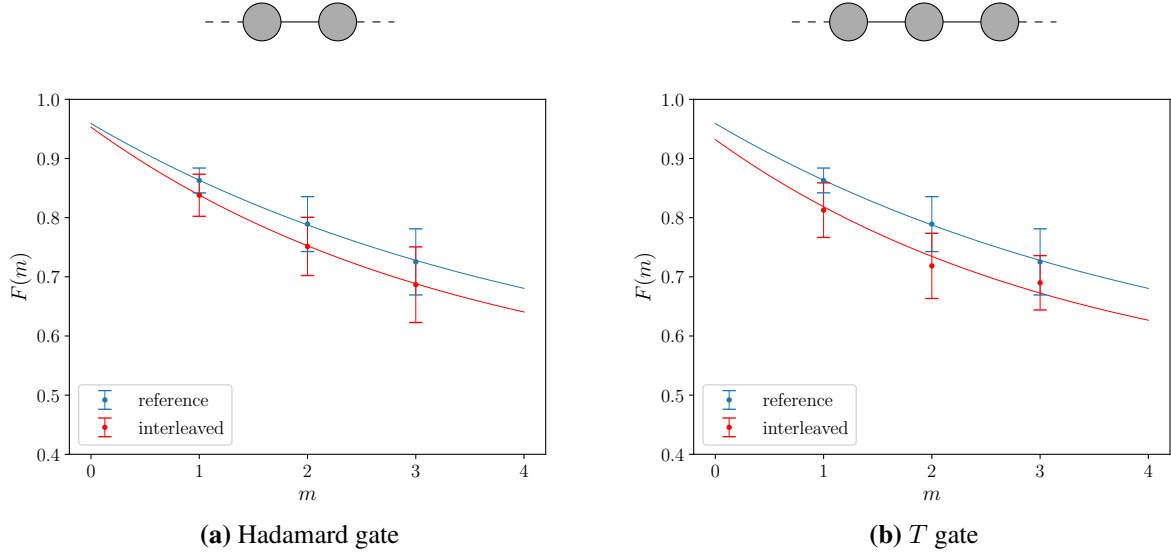


Figure 5.3: Sequence fidelities $F(m)$ obtained for $m \in \{1, 2, 3\}$ to estimate the fidelity of the 2-qubit measurement-based implementation of the Hadamard gate and the 3-qubit measurement-based implementation of the T gate, using the 5-qubit approximate measurement-based 2-design proposed in Sec. 3.3.5, on the *ibm_hanoi* quantum processor. Reference and interleaved sequence fidelities were fit to Eqs. (5.3) and (5.4) respectively (see Appendix D.2). (a) Hadamard gate shows the reference sequence fidelities (in blue) and the interleaved sequence fidelities for the 2-qubit Hadamard gate (in red). (b) T gate shows the same reference sequence fidelities (in blue) and the interleaved sequence fidelities for the 3-qubit T gate (in red).

method which takes these uncertainties into account was used to fit the reference and interleaved sequence fidelities to the exponential decay model given by Eqs. (5.3) and (5.4) respectively. The fitting procedure used, the fitting constraints imposed and the estimation of Haar-averaged gate fidelities from the fitting parameters are discussed in Appendix D.2. The estimated fidelity of the 2-qubit measurement-based implementation of the Hadamard gate is (0.977 ± 0.073) and the estimated fidelity of the 3-qubit measurement-based implementation of the T gate is (0.972 ± 0.070) . The large uncertainties in the estimated fidelities can be attributed both to the large uncertainties in the sequence fidelities and to the small number of sequence fidelities. It is unclear to what extent the use of an ϵ -approximate 2-design, as opposed to an exact 2-design, contributed to the large uncertainties in the estimated fidelities. It is also not clear how ϵ is related to these uncertainties or how the relation is affected by the Monte Carlo method used to estimate the fidelities. These aspects require further theoretical analysis.

We now compare the fidelities estimated using our measurement-based interleaved randomised benchmarking protocol to fidelities calculated from process tomography results. To this end, we performed quantum process tomography on each of the three respective sets of qubits of the *ibm_hanoi* quantum processor on which the 2-qubit Hadamard gate and the 3-qubit T gate were implemented in the interleaved sequences, and used the results to calculate the Haar-averaged gate fidelity of each of the three implementations of the 2-qubit Hadamard gate and the 3-qubit T gate. Further details about the process tomography method used, the implementation of the method and the calculation of Haar-averaged gate fidelities from process tomography results are given in Appendix D.3. Fidelities obtained for the three different implementations of the 2-qubit Hadamard gate, on the three different sets of qubits used in the interleaved sequences, range from 0.948 to 0.972. For the 3-qubit T gate, the fidelities obtained from process tomography results range from 0.928 to 0.939.

For both gates, the uncertainty of the gate fidelity calculated using process tomography results is smaller than the uncertainty of the gate fidelity estimated using our measurement-based interleaved randomised benchmarking protocol. We also note that the estimated fidelities are slightly larger than those obtained using process

tomography results. Nevertheless, the agreement between our estimated gate fidelities and the gate fidelities calculated from process tomography results is somewhat remarkable considering that we used a very weak approximate 2-design (see Sec. 3.3.5), the implementation of this 2-design on the IBM quantum processors did not even pass our test for an approximate 2-design for all states (see Sec. 3.3.5), noise from this 2-design is not entirely gate-independent (see Sec. 3.3.5) and the fitting parameters were inferred from very few sequence fidelities (only three). This clearly demonstrates the robustness of interleaved randomised benchmarking. It also provides motivation for considering weak approximate 2-designs, such as the one proposed in Sec. 3.3.5. It is unclear to what extent the use of a weak ϵ -approximate 2-design, as opposed to an exact 2-design, contributed to the over-estimation of gate fidelities. However, one would expect that in general, an increase in ϵ results in an increase in the positive difference between the estimated gate fidelity and the gate fidelity determined from process tomography results.

5.5.3 Noisier gates

We next investigate the extent to which our measurement-based interleaved randomised benchmarking protocol is able to detect noise variations in different measurement-based implementations of a gate. To this end, we artificially increase noise in the measurement-based implementations of the Hadamard gate and the T gate. One option is to perform the single-qubit measurements in these measurement-based implementations at measurement angles which deviate from the required measurement angles. However, this has the disadvantage that it is not easy to predict whether a measurement-based gate will be more or less noisy than the measurement-based 2-design used to estimate its fidelity. We therefore artificially increase noise in the measurement-based implementations of the Hadamard gate and the T gate by increasing the length of the linear cluster states used in the implementations.

Note that the 3-qubit linear cluster state with the measurement angles $\phi_1 = 0$ and $\phi_2 = 0$ implements the identity operation when all measurement outcomes are zero (see Appendix A.3). By appending this measurement-based implementation of the identity operation to the 2-qubit Hadamard gate and the 3-qubit T gate given in Sec. 5.2.1, we obtain measurement-based implementations of the Hadamard gate and the T gate with longer cluster states. In particular, when all measurement outcomes are zero, the 4-qubit linear cluster state with the measurement angles $\phi_1 = 0$, $\phi_2 = 0$ and $\phi_3 = 0$ implements the Hadamard gate and the 5-qubit linear cluster state with the measurement angles $\phi_1 = \frac{\pi}{4}$, $\phi_2 = 0$, $\phi_3 = 0$ and $\phi_4 = 0$ implements the T gate. Furthermore, provided that all measurement outcomes are zero, the 6-qubit linear cluster state with the measurement angles $\phi_1 = 0$, $\phi_2 = 0$, $\phi_3 = 0$, $\phi_4 = 0$ and $\phi_5 = 0$ also implements the Hadamard gate and the 7-qubit linear cluster state with the measurement angles $\phi_1 = \frac{\pi}{4}$, $\phi_2 = 0$, $\phi_3 = 0$, $\phi_4 = 0$, $\phi_5 = 0$ and $\phi_6 = 0$ also implements the T gate. When non-zero measurement outcomes do occur, the desired gate can be realised by simply applying the appropriate Pauli correction to the final qubit [161], that is, these implementations all use fixed measurement angles and a Pauli correction and do not require adaptive measurement feedforward.

We estimated the fidelity of the 4-qubit and 6-qubit measurement-based implementations of the Hadamard gate, and the 5-qubit and 7-qubit measurement-based implementations of the T gate, using the 5-qubit approximate measurement-based 2-design proposed in Sec. 3.3.5, on the *ibmq_brooklyn* quantum processor, a 65-qubit superconducting IBM quantum computer. Details of the qubits used are provided in Appendix D.4. We once again implemented reference sequences with lengths $m \in \{1, 2, 3\}$ and interleaved sequences with $m \in \{1, 2, 3\}$, for each of the gates considered. The *ibmq_brooklyn* quantum processor was used for these experiments, instead of the *ibm_hanoi* quantum processor, as the *ibm_hanoi* quantum processor has too few qubits to implement some of the required sequences, such as the interleaved sequence for the 7-qubit T gate.

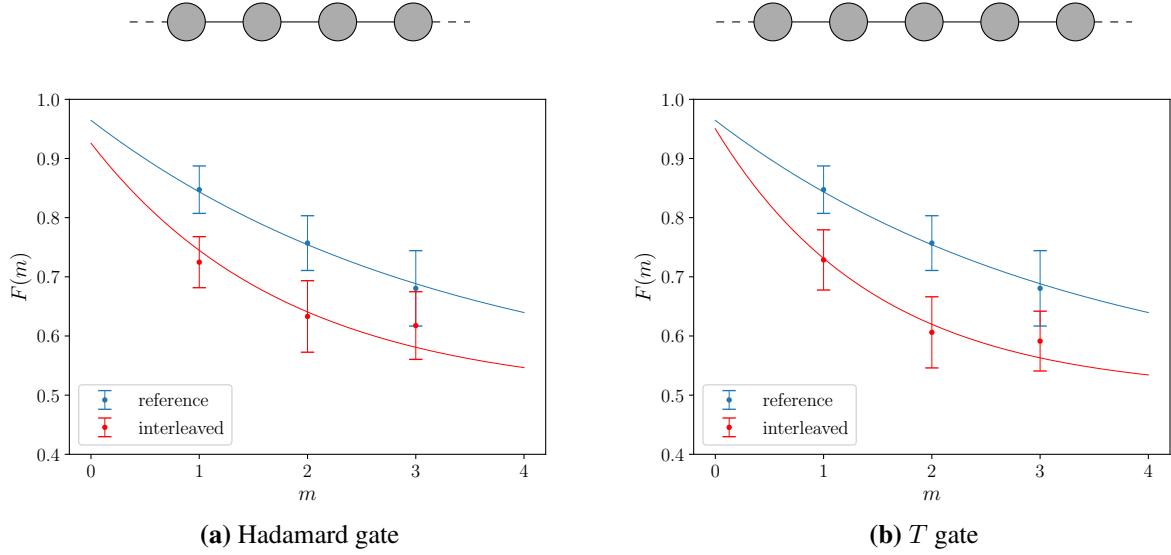


Figure 5.4: Sequence fidelities $F(m)$ obtained for $m \in \{1, 2, 3\}$ to estimate the fidelity of the 4-qubit measurement-based implementation of the Hadamard gate and the 5-qubit measurement-based implementation of the T gate, using the 5-qubit approximate measurement-based 2-design proposed in Sec. 3.3.5, on the *ibmq_brooklyn* quantum processor. Reference and interleaved sequence fidelities were fit to Eqs. (5.3) and (5.4) respectively (see Appendix D.2). (a) Hadamard gate shows the reference sequence fidelities (in blue) and the interleaved sequence fidelities for the 4-qubit Hadamard gate (in red). (b) T gate shows the same reference sequence fidelities (in blue) and the interleaved sequence fidelities for the 5-qubit T gate (in red).

with $m = 3$, which requires a 31-qubit linear cluster state, as shown in Fig. 5.1b. Sequences were implemented, classical post-processing was done to determine the survival probability for each random sequence and sequence fidelities were calculated in the same way as in the experiments on the *ibm_hanoi* quantum processor (see Sec. 5.5.1).

The only difference in the experiments on the *ibmq_brooklyn* quantum processor is that for the interleaved sequences for the measurement-based gates defined here, we need not perform quantum state tomography for each of the possible measurement outcomes to infer the output state for each random interleaved sequence. This can be understood as follows. For the 2-qubit Hadamard gate and the 3-qubit T gate there is a one-to-one correspondence between the outcomes of single-qubit measurements in the measurement-based implementations and the random byproducts that result from omitting the Pauli corrections. This results in a one-to-one correspondence between random measurement outcomes and random interleaved sequences, since the random interleaved sequences consist of random unitary operators interleaved with random byproducts. In contrast, there are 8, 16, 32 and 64 possible outcomes for the single-qubit measurements performed in the respective 4-qubit, 5-qubit, 6-qubit and 7-qubit measurement-based implementations defined here, but there are only four possible byproducts, corresponding to the four possible Pauli corrections, for each implementation. As a result, there is no longer a one-to-one correspondence between random measurement outcomes and random interleaved sequences. Therefore, to infer the output state for each random interleaved sequence, we need not perform quantum state tomography for each of the possible measurement outcomes, since tomography data obtained for different measurement outcomes which correspond to the same random interleaved sequence can be grouped.

We note that, for the interleaved sequences for the measurement-based gates defined here, it is in fact not feasible to perform quantum state tomography for each of the possible measurement outcomes, since the number of possible measurement outcomes grows exponentially with the length of the cluster state. For example, the circuit implementing the interleaved sequence for the 7-qubit T gate with $m = 3$ would need to be run

5.6. Conclusion

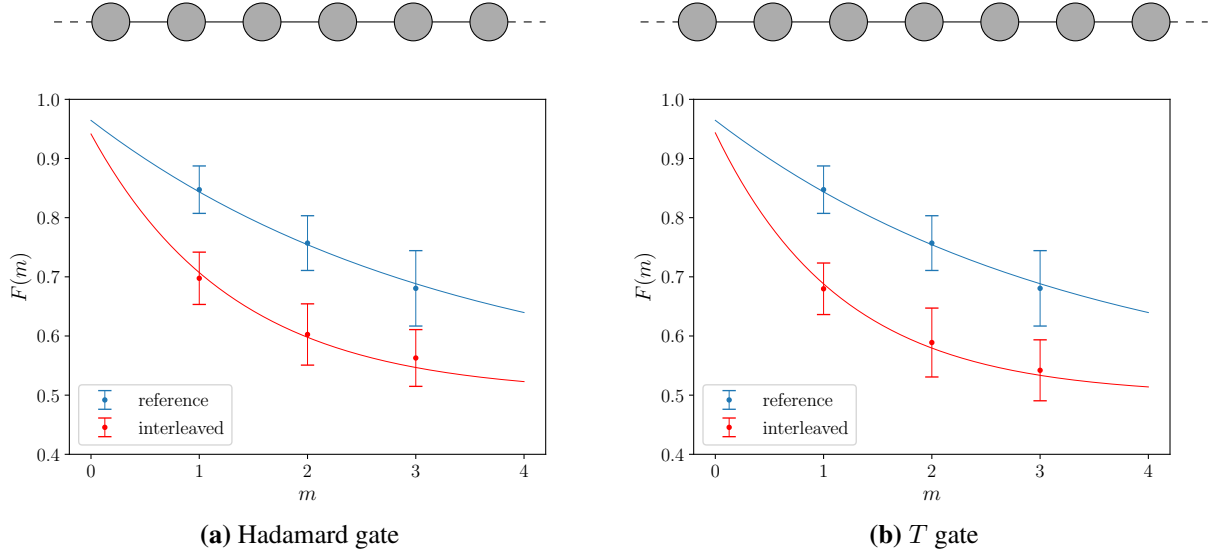


Figure 5.5: Sequence fidelities $F(m)$ obtained for $m \in \{1, 2, 3\}$ to estimate the fidelity of the 6-qubit measurement-based implementation of the Hadamard gate and the 7-qubit measurement-based implementation of the T gate, using the 5-qubit approximate measurement-based 2-design proposed in Sec. 3.3.5, on the *ibmq_brooklyn* quantum processor. Reference and interleaved sequence fidelities were fit to Eqs. (5.3) and (5.4) respectively (see Appendix D.2). (a) Hadamard gate shows the reference sequence fidelities from Fig. 5.4 (in blue) and the interleaved sequence fidelities for the 6-qubit Hadamard gate (in red). (b) T gate shows the same reference sequence fidelities from Fig. 5.4 (in blue) and the interleaved sequence fidelities for the 7-qubit T gate (in red).

$3(2^{30})/16 \approx 201$ million times to obtain the data required to perform tomography for each of the possible measurement outcomes. By performing tomography for each random interleaved sequence, and grouping tomography data obtained for different measurement outcomes corresponding to the same random sequence, we were able to drastically reduce the number of times each circuit had to be run. In particular, to obtain the data required to determine the interleaved sequence fidelity for a given m , for the 4-qubit, 5-qubit, 6-qubit and 7-qubit measurement-based gates defined here, the relevant circuit was run the same number of times as for the 3-qubit T gate (see Sec. 5.5.1), since the number of possible byproducts, and therefore the number of random interleaved sequences, is the same as for the 3-qubit T gate.

Sequence fidelities obtained to estimate the fidelity of the 4-qubit Hadamard gate and the 5-qubit T gate, using the 5-qubit approximate measurement-based 2-design proposed in Sec. 3.3.5, on the *ibmq_brooklyn* quantum processor are shown in Fig. 5.4. Furthermore, sequence fidelities obtained to estimate the fidelity of the 6-qubit Hadamard gate and the 7-qubit T gate are shown in Fig. 5.5. Table 5.1 shows the estimated gate fidelities, as well as the range of Haar-averaged gate fidelities calculated from process tomography results (see Appendix D.3). Both the estimated gate fidelities and the gate fidelities obtained from process tomography results decrease as the length of the cluster state in the implementation increases, which reflects the expected increase in noise resulting from increasing the length of the cluster state. This shows that, even with very little data, our measurement-based interleaved randomised benchmarking protocol is able to detect large noise variations in different measurement-based implementations of a gate.

5.6 Conclusion

We proposed an interleaved randomised benchmarking protocol for measurement-based quantum computers, which is an extension of the standard randomised benchmarking protocol proposed for measurement-based quantum computers by Alexander *et al.* [97]. In our measurement-based interleaved randomised benchmarking

Gate	Est Fidelity	Fidelity Range
4-qubit Hadamard gate	0.894 ± 0.092	0.831–0.895
5-qubit T gate	0.849 ± 0.085	0.821–0.851
6-qubit Hadamard gate	0.820 ± 0.079	0.693–0.835
7-qubit T gate	0.791 ± 0.081	0.702–0.803

Table 5.1: Haar-averaged gate fidelities obtained for the different measurement-based gates on the *ibmq_brooklyn* quantum processor. ‘Est Fidelity’ shows the fidelity estimated using our measurement-based interleaved randomised benchmarking protocol. ‘Fidelity Range’ shows the range of fidelities calculated from process tomography results obtained for the three different sets of qubits used in the interleaved sequences (see Appendix D.3).

protocol, any single-qubit measurement-based 2-design can be used to estimate the fidelity of any single-qubit measurement-based gate. We tested our protocol by using the approximate measurement-based 2-design proposed in Sec. 3.3.5 to estimate the fidelity of measurement-based implementations of the Hadamard gate and the T gate (a universal single-qubit set) on the remotely accessible IBM superconducting quantum computers. Since IBM processors did not support dynamic circuit execution at the time of performing the experiments, our protocol had to be adjusted to enable implementation on the superconducting quantum hardware available at the time. In particular, it was not possible to implement the inverse of a random sequence as a quantum mechanical operation without dynamic circuit execution, since the sequence, and therefore its inverse, is only known after the required single-qubit measurements have been performed. By preparing linear cluster states of up to 31 qubits, performing single-qubit measurements on all but the final qubit, performing quantum state tomography on the final qubit to infer the output state for each random sequence, and by using classical post-processing to apply the inverse of each sequence and extract the survival probability, we were able to determine reference sequence fidelities for sequence lengths $m \in \{1, 2, 3\}$ and interleaved sequence fidelities for $m \in \{1, 2, 3\}$ for each of the gates considered.

In our adjusted protocol, the resources required to obtain the data needed to determine the sequence fidelity scale exponentially with the length of the sequence, which is why the number of sequence fidelities determined here is much smaller than in typical interleaved randomised benchmarking experiments [95, 203]. Hence, even though our adjustments have enabled us to obtain data for a proof-of-concept demonstration, dynamic circuit execution remains essential for the efficient scaling of our measurement-based interleaved randomised benchmarking protocol. We note that our measurement-based interleaved randomised benchmarking protocol could be implemented as presented in Sec. 5.3, without any adjustments, on a measurement-based architecture or a circuit-based architecture which supports dynamic circuit execution. A measurement-based architecture or a circuit-based architecture which supports dynamic circuit execution would therefore eliminate the exponential scaling with sequence length as well as the need for quantum state tomography and classical post-processing. Recent work on circumventing dynamic circuit execution on IBM processors for the measurement-based model via a delayed choice strategy [211] is an interesting direction, although such a strategy comes at the expense of adding further entangling gates, which may introduce additional noise. Dynamic circuit execution is thus an important addition to IBM processors for the measurement-based model, and the recent addition thereof opens up many opportunities for quantum computing in the measurement-based model.

In all the experiments, estimated gate fidelities show good agreement with gate fidelities calculated from process tomography results. Even though some gate fidelities are slightly over-estimated and the uncertainties of estimated gate fidelities are larger than the uncertainties of those calculated from process tomography results, the experimental results clearly demonstrate the robustness of interleaved randomised benchmarking if one takes into consideration that we used a very weak approximate 2-design (see Sec. 3.3.5), the implementation

5.6. Conclusion

of this 2-design on the IBM quantum processors did not even pass our test for an approximate 2-design for all states (see Sec. 3.3.5), noise from this 2-design is not entirely gate-independent (see Sec. 3.3.5) and the fitting parameters were inferred from very few sequence fidelities. Furthermore, by artificially increasing noise in the measurement-based implementations of the Hadamard gate and the T gate, we were able to show that, even with very little data, our measurement-based interleaved randomised benchmarking protocol is able to detect large noise variations in different measurement-based implementations of a gate.

Chapter 6

Quantum random number generation using an on-chip nanowire plasmonic waveguide

6.1 Introduction

In Chapters 3 and 5, we explored randomness generation in the form of matrices and their application on IBM's cloud-based superconducting quantum computers. One can also think about using these systems for other types of randomness generation applications, such as generating random scalars, or in other words random numbers. While these systems provide a sophisticated platform for such a kind of randomness generation, they are not specifically designed for this task and more efficient platforms should be considered. In this chapter, we investigate randomness generation on custom-built on-chip plasmonic hardware. To this end, we integrate an on-chip nanowire plasmonic waveguide into an optical time-of-arrival based quantum random number generation setup. We achieved a random number generation rate of 14.4 Mbits/s, despite loss, and the generated bits did not require any post-processing to pass the industry standard ENT [147] and NIST [148] Statistical Test Suites. Furthermore, we were able to increase the generation rate to 41.4 Mbits/s, by increasing the light intensity, but the resulting bits required a shuffle to pass all the tests.

This chapter is structured as follows. In Sec. 6.2, we describe our custom-built optical experimental setup, our on-chip nanowire plasmonic waveguide and the variation of time-of-arrival scheme used. In Sec. 6.3, we present the test results obtained for various samples of bits generated using our custom-built quantum random number generation setup. A summary of the results and concluding remarks are given in Sec. 6.4. Complementary experimental investigations and theoretical analyses are presented in Appendix E.

6.2 Experimental setup

The experimental setup used to investigate time-of-arrival based quantum random number generation using a nanowire plasmonic waveguide is shown in Fig. 6.1a. The plasmonic waveguide used in the experiments comprises a gold nanowire 70 nm in diameter and just over 3 μm in length with tapering and a grating with a period of 740 nm on either end (see Fig. 6.1b). Each 11-step grating is 2 μm in width and 70 nm in height. In the optical setup, highly attenuated coherent laser light is focused onto the input grating of the nanowire plasmonic waveguide using a diffraction-limited microscope (DLM) objective. At the input grating, photons are converted to surface plasmon polaritons, which propagate through the waveguide to the output grating, where they are converted back to photons [123]. Photons are collected from the output grating of the nanowire

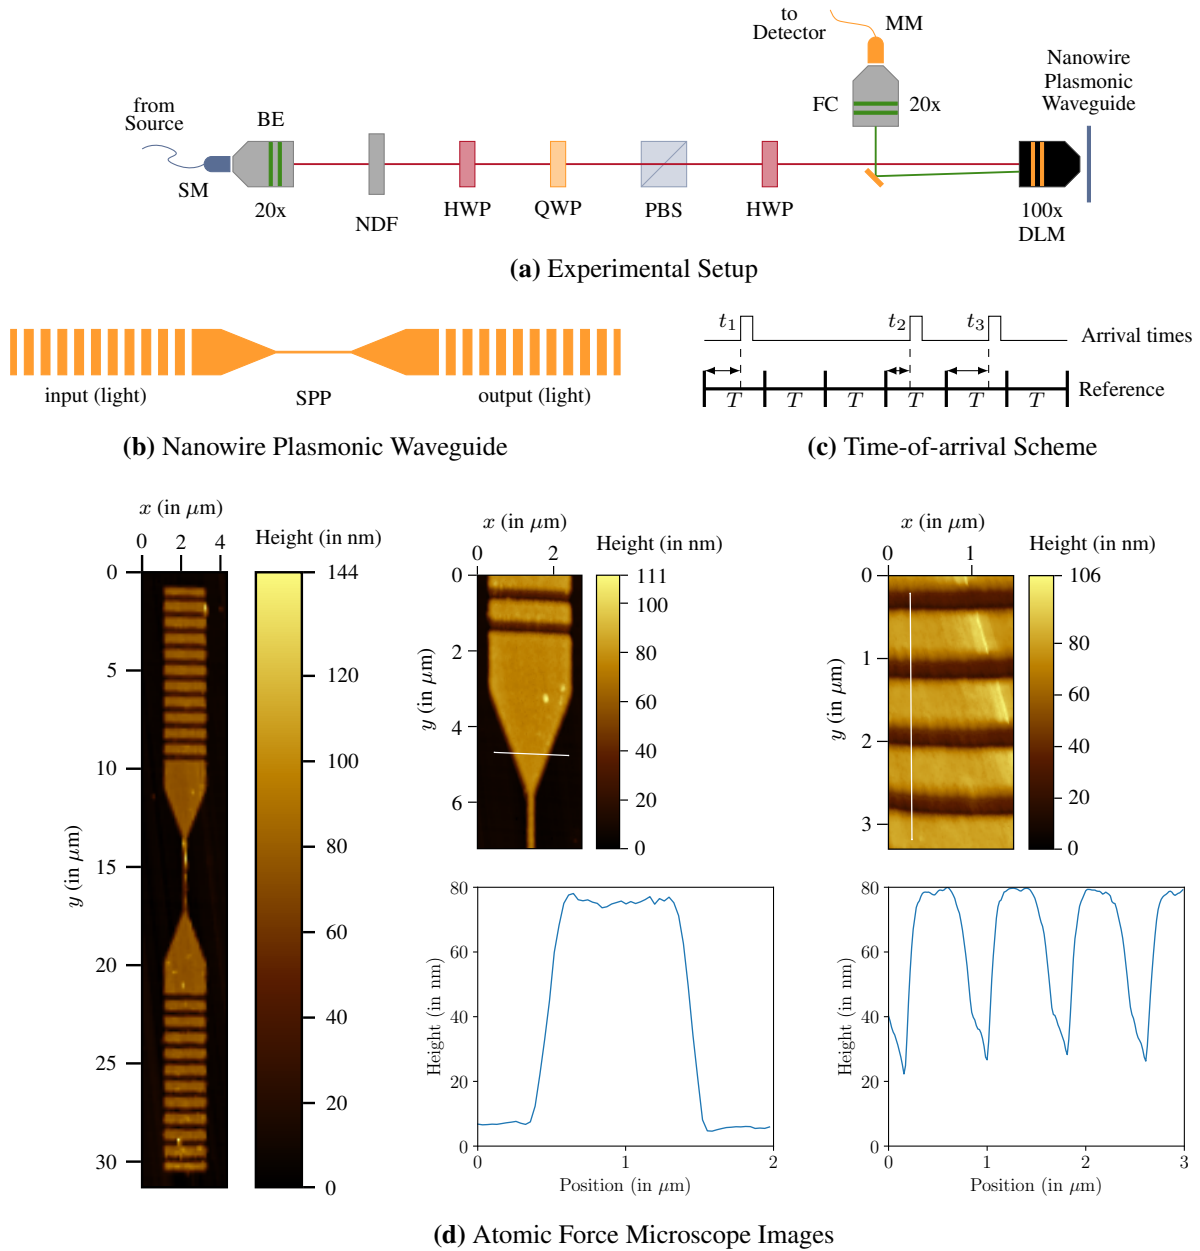


Figure 6.1: Quantum random number generation using an on-chip nanowire plasmonic waveguide. (a) Experimental Setup shows the experimental setup used to investigate time-of-arrival based quantum random number generation using an on-chip nanowire plasmonic waveguide. The labels used are: single-mode optical fibre (SM), beam expander (BE), neutral density filter (NDF), half-wave plate (HWP), quarter-wave plate (QWP), polarising beamsplitter (PBS), diffraction-limited microscope (DLM), fibre coupler (FC) and multi-mode optical fibre (MM). (b) Nanowire Plasmonic Waveguide shows a top view of the on-chip nanowire plasmonic waveguide used in the experiments. (c) Time-of-arrival Scheme illustrates the implemented variation of the time-of-arrival scheme, in which random numbers are obtained from the arrival times of photons relative to an external time reference [53]. (d) Atomic Force Microscope Images shows atomic force microscope (AFM) images of the fabricated on-chip nanowire plasmonic waveguide. These include an AFM image of the entire nanowire plasmonic waveguide (left), an AFM image of the top tapering (top centre), an AFM height profile of the top tapering (bottom centre), an AFM image of the top grating (top right) and an AFM height profile of the top grating (bottom right).

plasmonic waveguide using the same DLM objective and are then sent to a single-photon detector. The arrival times of photons at the detector, relative to an external reference (see Fig. 6.1c), are then used to obtain random numbers [53], as will be explained in more detail later.

The temporal degree of freedom of photons generated during stimulated emission is a true source of randomness [55]. In this work we employ a $\lambda = 785$ nm continuous-wave laser (Thorlabs LPS-785-FC) operating in the stimulated emission regime. A polarisation-preserving single-mode optical fibre (SM) connects the continuous-wave laser to a beam expander (BE), through which polarised coherent laser light enters the optical setup. The collimated beam passes through a neutral density filter (NDF), a half-wave plate (HWP), a quarter-wave plate (QWP), a polarising beamsplitter (PBS) and a second HWP. The NDF, along with loss in the optical setup and the plasmonic waveguide, ensures that light reaching the detector is attenuated down to an appropriate level for single-photon detection. The first HWP, the QWP and the PBS are used to purify the polarisation of light from the laser. In particular, the PBS transmits only horizontally polarised photons and the preceding HWP and QWP adjust the incident polarisation so as to minimise loss in the PBS. The second HWP rotates the resulting horizontally polarised light so as to maximise the conversion of photons to surface plasmon polaritons at the input grating of the nanowire plasmonic waveguide, which can only be achieved when the polarisation vector is perpendicular to the gratings [123]. A 100x DLM objective is used to focus the beam onto the input grating of the nanowire plasmonic waveguide at a spot size of about $2\ \mu\text{m}$.

The plasmonic waveguide is fabricated on a silica glass substrate with a refractive index of 1.5255 and a thickness of 1 mm using a combination of electron beam lithography and electron beam deposition. A resist is first spin coated on the silica glass substrate and 20 nm of gold is deposited so that the surface becomes conductive. Electron beam lithography is used to define the regions for the nanowire, the taperings and the gratings. The gold is then etched and resist developed. Next a lift-off technique is employed, where first a 5 nm thick adhesion layer of titanium is deposited and then the desired 70 nm thick gold layer on the silica glass substrate. The unexposed resist and gold is then lifted off with acetone, IPA and de-ionised water. Atomic force microscope (NT-MDT Smena) images of the fabricated on-chip nanowire plasmonic waveguide are shown in Fig. 6.1d.

The power transmission factor of the nanowire plasmonic waveguide was measured to be $\eta_{\text{wg}} = 6.7 \times 10^{-5}$ (see Appendix E.1). Losses occur in the waveguide as a result of scattering during the conversion between photons and surface plasmon polaritons at the gratings, as well as scattering from the tapering regions, and as a result of absorption during the propagation of surface plasmon polaritons along the nanowire. One can compensate for these losses by increasing the light intensity. Unlike in a previously demonstrated plasmonic quantum random number generator [123], there is no limit on the amount by which one can increase the intensity of the coherent source, since the time-of-arrival scheme does not require the nanowire plasmonic waveguide to operate in the single-excitation regime [53].

Photons are collected from the output grating of the nanowire plasmonic waveguide by the same DLM objective that was used to focus the input beam onto the input grating. A knife-edge mirror is then used to reflect these photons into a fibre coupler (FC), which is connected to a single-photon avalanche diode (SPAD) detector (Excelitas SPCM-AQRH-15) by a multi-mode optical fibre (MM). The polarisation dependence of the photon detection rate [221, 222] confirms that the collection optics is indeed capturing out-coupling photons from the output grating and not scattered photons from the input beam (see Appendix E.2). The SPAD detector used in the experiments has a dead time of 24 ns, a timing resolution of 350 ps and a maximum dark count rate of 50 counts/s. For data collection, the SPAD detector is connected to a channel of a Picoquant TimeHarp 260, which is connected to a PC. The TimeHarp is capable of recording the arrival time of a photon at the detector




Figure 6.2: Pearson correlation coefficient of the generated sample with 1-bit to 15-bit delays of itself.

to a precision of 25 ps.

We implement a variation of the time-of-arrival scheme, first proposed by Nie *et al.* [53], in which random numbers are obtained from the arrival times of photons relative to an external reference (see Fig. 6.1c). The benefit of this variation compared to earlier variations [48, 51, 52], in which random numbers are obtained directly from the difference of consecutive photon arrival times, is a significant reduction in bias. In the variation proposed by Nie *et al.*, the external time reference is divided into time intervals of an arbitrary but fixed length T , each of which are in turn divided into N bins of equal width. Provided that at most one photon detection can occur in a time interval of length T , it follows that a photon detection occurring in a time interval of length T occurs in each of the N bins with probability $\frac{1}{N}$ (shown in Appendix E.3). Hence the numbers of the bins in which photon detections occur are uniformly distributed random $\log_2(N)$ -bit unsigned integers. Device imperfections which can result in deviations from uniformity and degrade the quality of the random numbers generated in an experiment are discussed in Refs. [53, 55].

For our experiment, we set $T = 12.8$ ns. This ensures that T is less than the dead time of the SPAD detector, which in turn ensures that at most one photon detection can occur in a time interval of length T . Furthermore, we set $N = 2^8 = 256$, and so we can extract a random 8-bit unsigned integer from each photon arrival time. We note that the bin width, $\frac{T}{N} = 50$ ps, is greater than the precision of the TimeHarp, but less than the timing resolution of the SPAD detector. However, results from previous experiments [53, 55] suggest that the timing resolution of the detector does not significantly affect the uniformity or the quality of the random numbers.

For data collection, we adjust the light intensity so as to give a photon detection rate of 1.8 Mcounts/s. This corresponds to an average time interval of $0.56 \mu\text{s}$ between photon detections, which is much greater than the dead time of the SPAD detector. This ensures that the majority of photons arriving at the SPAD detector are indeed detected. Random 8-bit unsigned integers extracted from recorded photon arrival times are converted to binary form, resulting in a sample of binary digits or bits. We generated a sample of 866,893,768 bits in 60 s, which corresponds to a random number generation rate of about 14.4 Mbits/s. This is an order of magnitude improvement in speed compared to both a previous plasmonic quantum random number generator [123] and previous on-chip time-of-arrival generators [73]. In the next section, we apply a number of industry standard tests [147, 148] to the first 800 Mbits generated, which we will refer to as the generated sample.

Test	Generated	Expected
Entropy	7.999998	8.000000
χ^2 Distribution	9.08%	10–90%
Arithmetic Mean	127.503	127.500
Monte Carlo value for π	3.14177173	3.14159265
Serial Correlation Coefficient	0.001500	0.000000

Table 6.1: ENT Statistical Test Suite results for the generated sample. ‘Generated’ shows the values obtained using the generated sample. ‘Expected’ shows the expected values for a true random sample.

Test	Req	Prop	p-value
Frequency	783	792	0.634516
Block Frequency	783	792	0.138267
Cumulative Sums 1	783	789	0.928563
Cumulative Sums 2	783	796	0.482223
Runs	783	794	0.739918
Longest Run of Ones	783	795	0.894201
Binary Matrix Rank	783	793	0.757297
Discrete Fourier Transform	783	793	0.021262
Non-overlapping Template*	783	792	0.573621
Overlapping Template	783	788	0.805107
Universal Statistical	783	795	0.487074
Approximate Entropy	783	790	0.562080
Random Excursions*	472	479.5	0.472780
Random Excursions Variant*	472	481	0.540922
Serial 1	783	796	0.934318
Serial 2	783	793	0.219006
Linear Complexity	783	794	0.444226

Table 6.2: NIST Statistical Test Suite results for the generated sample. ‘Req’ shows the minimum number of sequences which need to pass a test for the sample to pass the test. ‘Prop’ shows the number of sequences of the generated sample which passed each test. For tests which involve more than five subtests (marked with *) the median of the results is presented.

6.3 Results

As a first test, we employ the Pearson correlation coefficient [223] to detect short-ranged correlations in the generated sample. The Pearson correlation coefficient is a real number in the interval $[-1, 1]$, where a positive value suggests a positive correlation, a negative value suggests a negative correlation and a value close to zero suggests no correlation. As can be seen in Fig. 6.2, short-ranged correlations in the generated sample are negligible. Furthermore, the relative frequency of zeros and ones in the generated sample is 0.49995 and 0.50005 respectively. Hence the bias in the generated sample is also negligible.

For further quality assessment, we apply the ENT Statistical Test Suite [147] to the generated sample. The ENT Statistical Test Suite comprises five simple tests in which the sample being tested is used to compute five important values. These include the entropy per byte, which quantifies the information density of the sample, the χ^2 distribution, which is known to be extremely sensitive to flaws in random number generators, the arithmetic mean of 8-bit unsigned integers extracted from the sample, which aids in detecting bias, the Monte Carlo value for π , which provides a practical test of the generator’s suitability for use in simulation, and the serial correlation coefficient, which quantifies correlations between adjacent bytes in the sample. The quality of the sample being tested can then be assessed by comparing the values obtained using the sample to the known values for a true random sample. The ENT Statistical Test Suite results for the generated sample are

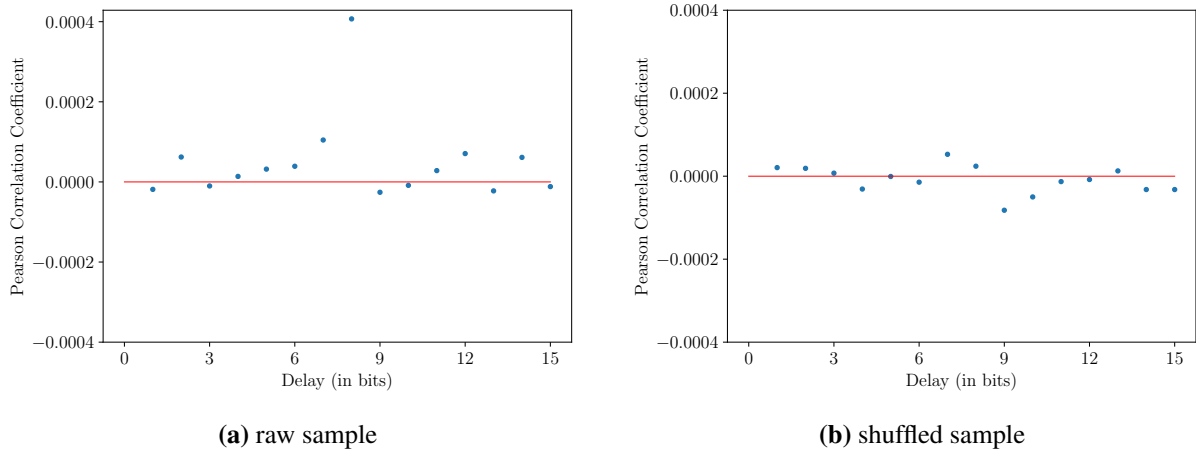


Figure 6.3: Pearson correlation coefficient of the (a) raw sample and (b) shuffled sample with 1-bit to 15-bit delays of itself.

given in Table 6.1. For all five tests, the values obtained using the generated sample show good agreement with the expected values for a true random sample.

As a final assessment of the quality, we apply the NIST Statistical Test Suite [148] to the generated sample. The NIST Statistical Test Suite consists of 15 stringent tests, which are primarily aimed at assessing a random number generator's suitability for use in cryptographic applications. To apply one of these tests to a sample of bits from a generator, the sample is first divided into sequences of a fixed length. The test is then applied to each sequence and a p -value, which can be used to assess the uniformity of the distribution of the test results obtained for the individual sequences, is determined. For the sample to pass the test, a sufficient number of sequences must pass the test and the p -value must be greater than or equal to 0.0001. The NIST Statistical Test Suite results for the generated sample are shown in Table 6.2. For each test, the generated sample was divided into 800 sequences 1 Mbit in length. The default values were used for the block length, except in the Block Frequency test, where the block length was adjusted from $2^7 = 128$ to $2^{14} = 16384$. The generated sample passed all 15 NIST tests. Hence our quantum random number generation setup is capable of generating random numbers of sufficient quality for cryptographic applications — without employing any classical post-processing. This is a significant improvement compared to a previous plasmonic quantum random number generator [123] and previous on-chip time-of-arrival generators [73], which required a randomness extractor to pass the NIST Statistical Test Suite.

In principle, the random number generation rate can be increased by increasing the light intensity, which increases the photon detection rate. However, with an increased photon detection rate, the detector dead time becomes non-negligible and the photon counts obtained in an experiment would typically need to be multiplied by a non-unit correction factor to compensate for the resulting underestimation of photon counts. We now investigate the effect of detection with a non-unit correction factor on the quality of the random numbers generated in an experiment. To this end, we adjust the light intensity so as to give a photon detection rate of 5.2 Mcounts/s. In Appendix E.4, we show that despite the increased photon detection rate, higher order photon number within a time interval of length T is negligible. However, with an increased detection rate, the average time interval between photon detections decreases to $0.19 \mu\text{s}$, which is closer to the dead time of the SPAD detector. The correction factor is about 1.143 (see Appendix E.4), which means that on average, one in every seven photons arriving at the SPAD detector are not detected. We generated 1,242,469,056 bits in 30 s, which corresponds to a random number generation rate of about 41.4 Mbits/s. We then applied the same industry

Test	Raw	Shuffled	Expected
Entropy	7.999997	7.999998	8.000000
χ^2 Distribution	<0.01%	32.17%	10–90%
Arithmetic Mean	127.510	127.524	127.500
Monte Carlo value for π	3.14142925	3.14063437	3.14159265
Serial Correlation Coefficient	0.004634	−0.000002	0.000000

Table 6.3: ENT Statistical Test Suite results for the raw sample and the shuffled sample. ‘Raw’ shows the values obtained using the raw sample. ‘Shuffled’ shows the values obtained using the shuffled sample. ‘Expected’ shows the expected values for a true random sample.

Test	Req	Raw Sample		Shuffled Sample	
		Prop	p-value	Prop	p-value
Frequency	783	794	0.465415	794	0.465415
Block Frequency	783	789	0.028577	794	0.134365
Cumulative Sums 1	783	793	0.598138	794	0.394631
Cumulative Sums 2	783	795	0.629311	795	0.701879
Runs	783	791	0.734904	794	0.346453
Longest Run of Ones	783	790	0.798139	786	0.816537
Binary Matrix Rank	783	793	0.964295	788	0.759756
Discrete Fourier Transform	783	793	0.052778	791	0.455937
Non-overlapping Template*	783	792	0.525357	792	0.550606
Overlapping Template	783	795	0.373203	784	0.729870
Universal Statistical	783	788	0.130557	795	0.053627
Approximate Entropy	783	792	0.379555	792	0.549331
Random Excursions*	483	490	0.538512	489	0.597986
Random Excursions Variant*	483	491	0.209902	491	0.314481
Serial 1	783	786	0.196920	795	0.587791
Serial 2	783	792	0.722284	794	0.467799
Linear Complexity	783	788	0.324821	795	0.737414

Table 6.4: NIST Statistical Test Suite results for the raw sample and the shuffled sample. ‘Req’ shows the minimum number of sequences which need to pass a test for the samples to pass the test. ‘Prop’ shows the number of sequences of the raw sample or the shuffled sample which passed each test. For tests which involve more than five subtests (marked with *) the median of the results is presented.

standard tests to the first 800 Mbits generated, which we will refer to as the raw sample.

We find that short-ranged correlations in the raw sample are mostly negligible, with non-negligible correlations present only between bits at 8-bit intervals (see Fig. 6.3a). To remove these correlations, we deterministically rearrange or shuffle the bits in the raw sample. In what follows, we will refer to the resulting sample of bits as the shuffled sample. As can be seen in Fig. 6.3b, shuffling the bits in the raw sample indeed removed the 8-bit interval correlations. The relative frequency of zeros and ones is 0.49993 and 0.50007 respectively, in both the raw sample and the shuffled sample, and so the bias is negligible in both samples.

The ENT Statistical Test Suite results for the raw sample and the shuffled sample are given in Table 6.3. For the raw sample, the χ^2 distribution and the serial correlation coefficient deviate significantly from the expected values for a true random sample. The large serial correlation coefficient for the raw sample seems to suggest that undetected photons result in correlations between adjacent bytes extracted from consecutive photon arrival times. In contrast, the χ^2 distribution and the serial correlation coefficient obtained using the shuffled sample show excellent agreement with the expected values for a true random sample. The negligible serial correlation coefficient for the shuffled sample shows that rearranging the bits in the raw sample, so that the eight consecutive bits which make up a given byte are extracted from eight different non-consecutive photon

arrival times, removes the correlations between adjacent bytes. Finally, we note that both the raw sample and the deterministically shuffled sample passed the NIST Statistical Test Suite (see Table 6.4). The fact that the raw sample also passed the NIST tests confirms that a deterministic shuffle, which essentially just transforms short-ranged correlations into long-ranged correlations, is sufficient to improve the quality of the raw sample, and that more sophisticated randomness extraction schemes are not needed.

6.4 Conclusion

We demonstrated the successful integration of a nanowire plasmonic waveguide into an optical time-of-arrival based quantum random number generation setup. Despite the presence of loss in the plasmonic waveguide and in the optical setup, we managed to achieve a random number generation rate of 14.4 Mbits/s. This is an order of magnitude improvement in speed compared to both a previous plasmonic quantum random number generator [123] and previous on-chip time-of-arrival generators [73]. Furthermore, unlike these previous devices, our generator did not require any classical post-processing to pass the NIST Statistical Test Suite. While we were able to increase the generation rate to 41.4 Mbits/s, the resulting bits only required a shuffle to pass all the tests.

Chapter 7

Conclusion

In this thesis we explored the use of quantum mechanical methods to generate and use randomness. Within this context, we investigated the practical generation of random matrices and scalars. In particular, we studied the generation of random unitary operators on IBM’s cloud-based superconducting quantum computers, going beyond previous work in which we investigated quantum random number generation on these systems [82]. We also investigated quantum random number generation on custom-built on-chip plasmonic hardware.

We started in Chapter 1 with an introduction to randomness and its importance in quantum information science. Then in Chapter 2, we outlined the basic tools and concepts that would be used in the remainder of the thesis.

In Chapter 3, we implemented the exact single-qubit measurement-based unitary 3-design of Ref. [88] on IBM quantum processors by performing measurements on a 6-qubit linear cluster state. By analysing process tomography results, we were able to show that the ensemble of unitary operators realised was a 1-design, but not a 2-design or a 3-design under the test conditions set, which we showed to be a result of depolarising noise during the measurement-based process. We obtained improved results for the 2-design test by implementing an approximate 2-design, in which measurements were performed on a smaller 5-qubit linear cluster state, but the test still did not pass for all states as a result of noise. This suggests that the practical realisation of measurement-based t -designs on superconducting quantum computers will require further work on the reduction of depolarising noise in these devices.

In Chapter 4, we conducted a theoretical investigation into the effect of noise on the quality of single-qubit unitary t -designs. While we hope that this investigation will encourage research into the effect of noise on the quality of multi-qubit t -designs, we also note that there are many protocols which exclusively use single-qubit t -designs [101, 103, 104, 118] for which our study may have direct consequences. The noise channels we studied were bit flips, phase flips, bit and phase flips, phase damping, amplitude damping and depolarising noise. We considered two noise models, in line with the noise models used in randomised benchmarking [92–100]. The first had noise applied before the t -design unitary operations, while the second had noise applied after the unitary operations. We showed that the single-qubit 1-design is completely unaffected by an arbitrary noise channel, for the model where noise is applied before the unitary operations. For the model where noise is applied after the unitaries, we showed that the 1-design is unaffected by noise, unless amplitude damping is applied.

Based on numeric results obtained for the 2-design, 3-design, 4-design and 5-design in Chapter 4, we conjecture that a $(2t + 1)$ -design is as sensitive to noise as a $2t$ -design, for any noise channel which deforms the Bloch sphere, but does not shift the Bloch sphere. While it may be possible to prove this with induction on t using recently discovered random circuit constructions for exact t -designs [109], such a proof evaded the

author. Further work in this direction is needed. Developing and studying noise models which use the definition of a t -design in terms of a polynomial function [92] may also help to uncover this even/odd behaviour. We note that numeric results also revealed substantial variations in sensitivity to noise throughout the Bloch sphere. In particular, t -designs appear to be most sensitive to noise when acting on pure states and least sensitive to noise for the maximally mixed state. For depolarising noise, we showed that our two noise models are equivalent, and for the other noise channels, numeric results obtained for the model where noise is applied after the unitaries reflect the transformation of the noise channel into a depolarising channel, an effect exploited in randomised benchmarking with 2-designs [92–94, 98–100]. Future work going beyond states that are t -fold tensor products of single-qubit states and their geometric interpretation, as well as investigations into the effect of noise on the quality of multi-qubit t -designs, will help to elucidate further behaviour of t -designs under the effects of noise. This kind of work will be helpful for researchers studying and developing applications using t -designs under realistic conditions.

In Chapter 5, we went on to propose an interleaved randomised benchmarking protocol for measurement-based quantum computers, in which any single-qubit measurement-based 2-design can be used to estimate the fidelity of any single-qubit measurement-based gate. Future work here could involve developing interleaved randomised benchmarking protocols in which multi-qubit measurement-based 2-designs [89] can be used to estimate the fidelity of multi-qubit measurement-based gates. Obstacles that would need to be overcome in this regard include computing the inverse of a random multi-qubit sequence, since multi-qubit systems cannot generally be efficiently simulated on a classical computer, and applying this inverse in such a way that noise from the inverse is independent of the sequence, since the inverse of a multi-qubit sequence cannot be applied by performing a single-qubit measurement basis rotation. Random measurement-based Clifford gates may help in this regard, where they may be used to bound the fidelity of non-Clifford gates that form a universal set [217]. Another option is to combine Clifford group and Pauli group gates within the interleaved sequence of a measurement-based non-Clifford gate [96]. Future work investigating the implications of performing randomised benchmarking with an approximate 2-design, as opposed to an exact 2-design, is also needed, since it is still unknown whether exact multi-qubit measurement-based 2-designs exist.

We also demonstrated our single-qubit measurement-based interleaved randomised benchmarking protocol on IBM superconducting quantum computers by using our approximate measurement-based 2-design to estimate the fidelity of measurement-based implementations of the Hadamard and T gates — a universal single-qubit gate set. To this end, measurements were performed on linear cluster states of up to 31 qubits. In all the experiments, our estimated gate fidelities showed good agreement with those calculated from process tomography results. By artificially increasing noise, we were able to show that our protocol detects large noise variations in different implementations of a gate. This demonstration highlights the usefulness of cloud-based superconducting systems for single-qubit measurement-based quantum computing and shows how to practically characterise noisy quantum logic gates in this setting. In future experiments, our protocol could be implemented on other physical systems, most notably custom-built linear optical systems [30–32, 224–227], one of the most promising physical systems for measurement-based quantum computing.

Finally, in Chapter 6, we investigated quantum random number generation on custom-built on-chip plasmonic hardware. To this end, we integrated an on-chip nanowire plasmonic waveguide into an optical time-of-arrival based quantum random number generation setup. Despite the presence of loss, we achieved a random number generation rate of 14.4 Mbits/s. Furthermore, the generated bits did not require any classical post-processing to pass industry standard tests. By increasing the light intensity, we were able to increase the generation rate to 41.4 Mbits/s, although the resulting bits required a shuffle to pass all tests. This work demon-

strates the successful integration of an on-chip nanoscale plasmonic component into a quantum random number generation setup. We note that although our current setup relies on an off-chip source and the detection is also done off-chip, future work on the integration of an on-chip source [139–142] and detector [143–146] would enable a self-contained quantum random number generator chip with a footprint an order of magnitude smaller than its dielectric counterpart. This would lead to new opportunities in compact and scalable quantum random number generation.

Broadly speaking, the work in this thesis contributes to an improved understanding of the randomness generation capabilities of current quantum hardware. This will be of interest to researchers studying and developing applications of randomness in quantum information science with the aim of realising these applications on NISQ hardware. Future work could involve implementing randomness-based algorithms with important applications in quantum information science on IBM’s cloud-based superconducting quantum computers, as well as developing fully integrated on-chip nanophotonic hardware for quantum random number generation.

Appendices

Appendix A

Implementation of single-qubit measurement-based t-designs using IBM processors

A.1 Qubits used for the 3-design

The exact measurement-based 3-design was implemented on six physical qubits of the *ibmq_toronto* quantum processor. This processor was chosen for its low error rates compared to other processors available at the time of performing the experiments. The qubit topology of the *ibmq_toronto* quantum processor is shown in Fig. A.1. The qubits 1 to 6 of the 6-qubit linear cluster state in the 3-design implementation were mapped onto the physical qubits 16, 19, 22, 25, 24 and 23 of the *ibmq_toronto* quantum processor, in such a way that the input state was prepared on qubit 16 and the output state was retrieved from qubit 23. These qubits were chosen as they form one of the few sets of six connected qubits (see Fig. A.1) in which all the qubits generally have relatively low error rates. The relevant calibration information as obtained at the time of running the circuits for the exact 3-design implementation, is shown in Table A.1.




Figure A.1: Qubit topology of the *ibmq_toronto* quantum processor. The connecting lines between qubits indicate the qubit pairs for which the CX gate is supported at the hardware level. The qubits used for the exact 3-design implementation are shaded grey.

Qubit	T_1 (μ s)	T_2 (μ s)	\sqrt{X} Error	Readout Error
16	123.42	135.48	0.000297	0.0154
19	114.98	123.10	0.000434	0.0116
22	110.50	148.90	0.000330	0.0191
25	125.41	114.64	0.000323	0.0108
24	121.49	155.26	0.000180	0.0093
23	98.92	40.59	0.000441	0.0553

Qubit Pair	CX Error
16–19	0.00758
19–22	0.01024
22–25	0.01053
25–24	0.01099
24–23	0.00892

Table A.1: Calibration information for the *ibmq_toronto* quantum processor as obtained at the time of running the circuits for the exact 3-design. The single-qubit calibration information for the relevant qubits is shown on the left. T_1 and T_2 are the amplitude and phase damping time constants respectively of the qubits. The CX error rates for relevant qubit pairs are shown on the right.

A.2 Depolarising noise

The action of the depolarising channel on a single-qubit state ρ is described by Eq. (2.27) and the resulting state is denoted by $\varepsilon(\rho)$. Recall that in the depolarising channel, the state ρ is replaced by the maximally mixed state with probability p . We investigated the effect of depolarising noise on an exact t -design’s ability to accurately reproduce the moments of the uniform Haar ensemble. In particular, we carried out the test for an approximate t -design with the middle term in inequality (3.6) replaced by $\mathbb{E}_H^t((\varepsilon(\rho))^{\otimes t})$, the expectation of the uniform Haar ensemble computed from a state to which depolarising noise has been applied.

We managed to obtain test results analytically for the 1-design. Using the Pauli 1-design, we showed that $\mathbb{E}_H^1(\rho) = \frac{1}{2}I$ for all states ρ , which we used to show that $\epsilon = 0$ for all $p \in [0, 1]$. Hence depolarising noise has no effect on a 1-design’s ability to reproduce the first moment of the uniform Haar ensemble. Since $\mathbb{E}_H^2(\rho^{\otimes 2})$ and $\mathbb{E}_H^3(\rho^{\otimes 3})$ depend on the state ρ , analytic test results for the 2-design and 3-design are much harder to find. Using a sample of 1000 density matrices, obtained as described in Sec. 3.3.4, and computing $\mathbb{E}_H^2(\rho^{\otimes 2})$ and $\mathbb{E}_H^3(\rho^{\otimes 3})$ for each density matrix ρ using the exact 3-design described in Sec. 3.2.1, we obtained results numerically for the 2-design and the 3-design. Test results obtained for the 2-design for different values of p and different truncation radii r_t are plotted in Fig. A.2. For all truncation radii, ϵ increases linearly with p , up to about $p = 0.4$, after which the increase becomes more gradual. The values of ϵ obtained for states close to the surface of the Bloch sphere are very large, even for small p . This shows that the second moment of the uniform Haar ensemble is very sensitive to depolarising noise. Test results for the 3-design are identical to that of the 2-design, shown in Fig. A.2. This suggests that the third moment of the uniform Haar ensemble is unaffected by depolarising noise.

To determine whether this depolarising noise model is a good noise model for the exact 3-design implementation on the *ibmq_toronto* quantum processor, we attempt to infer a consistent value for the parameter p from the test results in Table 3.3. Given r_t and ϵ obtained in a test for the 2-design or the 3-design, we simply read off the corresponding value of p from the plot in Fig. A.2. Using the results for the 2-design test (without quantum readout error mitigation) we infer $p = 0.31$ and using the results for the 3-design test we infer $p = 0.36$. For the results with quantum readout error mitigation, we infer $p = 0.20$ using the 2-design test results and $p = 0.32$ using the 3-design test results. Since the values inferred from the 2-design test and the 3-design test are different in both cases, we conclude that this depolarising noise model is not the correct noise model for the 3-design implementation.

We now consider a depolarising noise model which more closely resembles the way in which depolarising noise occurs in t -designs realised using a measurement-based approach. In this model, the test for an approximate t -design is performed with the middle term in inequality (3.6) calculated by repeatedly applying

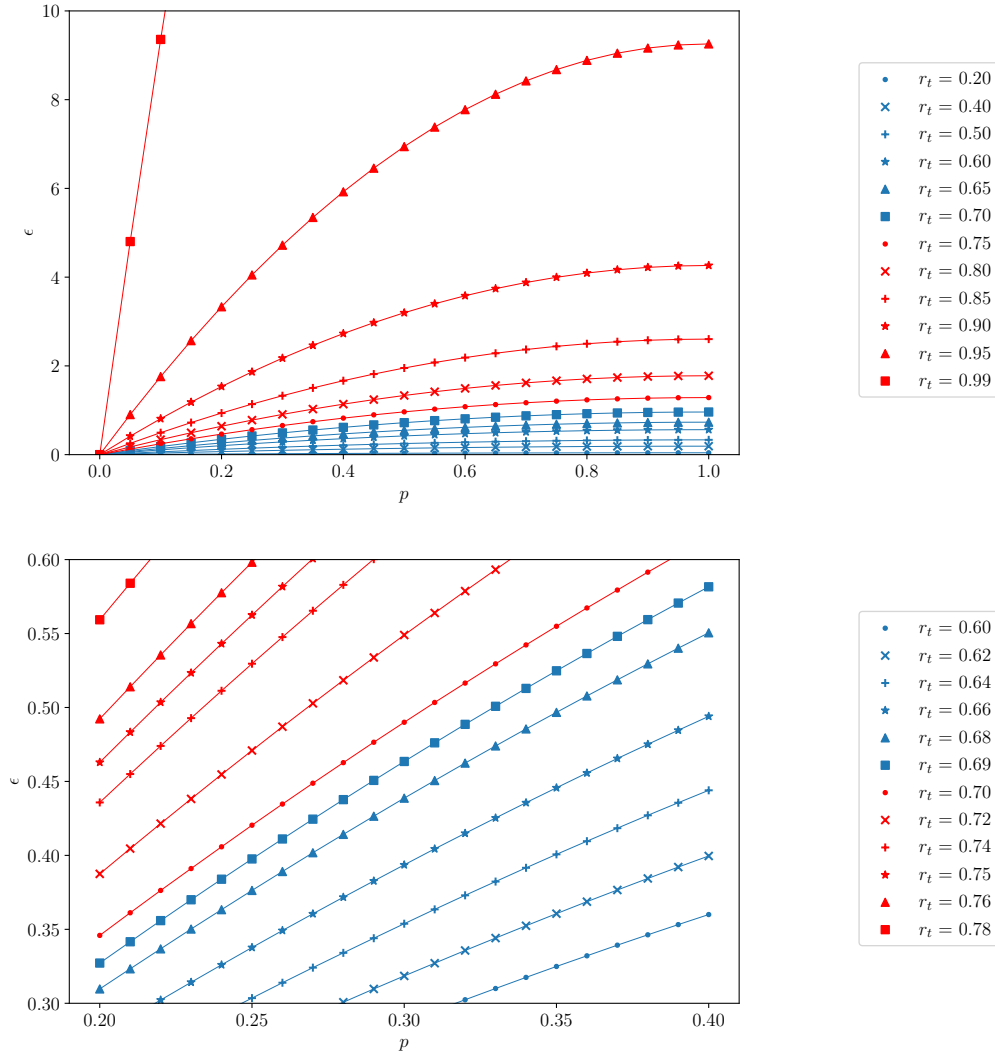


Figure A.2: ϵ versus p for the 2-design test with the middle term in inequality (3.6) replaced by $\mathbb{E}_H^2((\epsilon(\rho))^{\otimes 2})$, for different truncation radii r_t . A plot of the full set of results is shown above and a plot focused on the region of interest is shown below.

depolarising noise to a state, in between the five individual unitary operations that are applied to the input state in the exact 3-design described in Sec. 3.2.1. Our analytic results for the 1-design carry over to this model. Test results obtained numerically for the 2-design are shown in Fig. A.3. For this model, ϵ versus p starts to plateau at much smaller values of p . As a result of repeated applications of depolarising noise, the largest possible ϵ , for a given truncation radius, is reached for much smaller p . Results for the 3-design test are once again identical to that of the 2-design. Considering this model and using Fig. A.3, we infer $p = 0.06$ using the results for the 2-design test (without quantum readout error mitigation) and $p = 0.07$ using the results for the 3-design test. For the results with quantum readout error mitigation, we infer $p = 0.04$ using the 2-design test results and $p = 0.06$ using the 3-design test results. The values of p inferred here from the two test results are close enough to be considered consistent in both cases. We therefore conclude that this is a very good noise model for the exact 3-design implementation on the *ibmq_toronto* quantum processor. From the values of p inferred, it is also clear that quantum readout error mitigation has reduced depolarising noise in the implementation.

A.3. Identity implementation

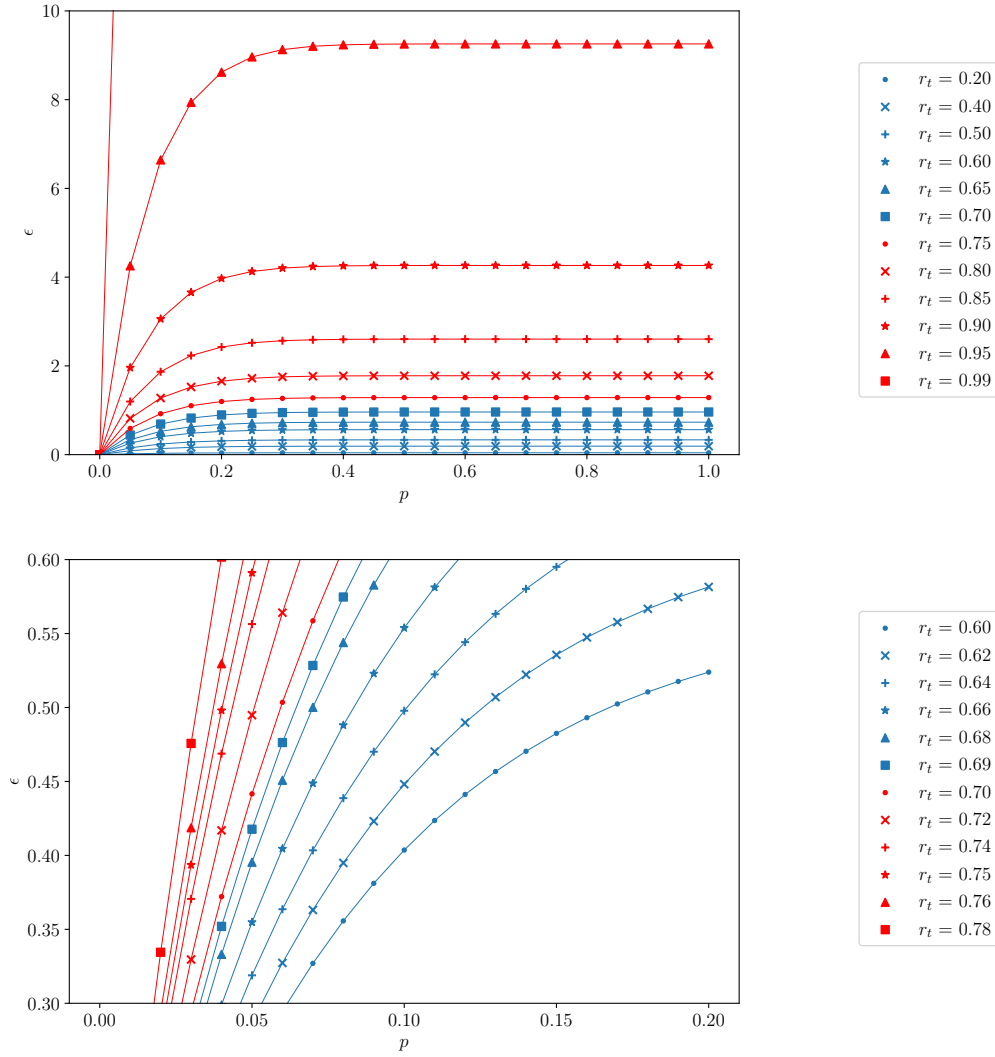


Figure A.3: ϵ versus p for the 2-design test with the middle term in inequality (3.6) calculated by repeatedly applying depolarising noise to a state, in between the five individual unitary operations that are applied to the input state in the exact 3-design described in Sec. 3.2.1, for different truncation radii r_t . A plot of the full set of results is shown above and a plot focused on the region of interest is shown below.

A.3 Identity implementation

In Sec. 2.3.2, we reviewed single-qubit measurement-based quantum computing with linear cluster states and discussed measurement-based implementations of the Hadamard gate and the T gate. A n -qubit linear cluster state can be also used to implement the identity operation for odd n . Note that when each qubit $i \in \{1, 2, \dots, n-1\}$ is measured in the Pauli X -basis, that is in the direction $\phi_i = 0$, Eq. (2.29) reduces to $U_{m_i}(0) = X^{m_i}H$. Hence when the measurement outcomes $m_i = 0$ for all i , Eq. (2.30) reduces to $U_{\mathbf{m}}(\mathbf{0}) = I$. When some of the measurement outcomes are non-zero, the identity can still be implemented by applying the appropriate Pauli correction to the final qubit [161], just as in the measurement-based implementations of the Hadamard gate and the T gate. Measurement-based implementations of the identity operation can be used to determine the type of noise on a set of qubits. For depolarising noise, we expect non-zero real entries along the diagonals of χ matrices obtained by performing process tomography.

We implemented the identity operation by performing single-qubit measurements on 3-qubit, 5-qubit and 7-qubit linear cluster states prepared on the same qubits of the *ibmq_toronto* quantum processor as was used for the exact 3-design implementation. Appendix A.5 provides more detail on the qubits used. Generation

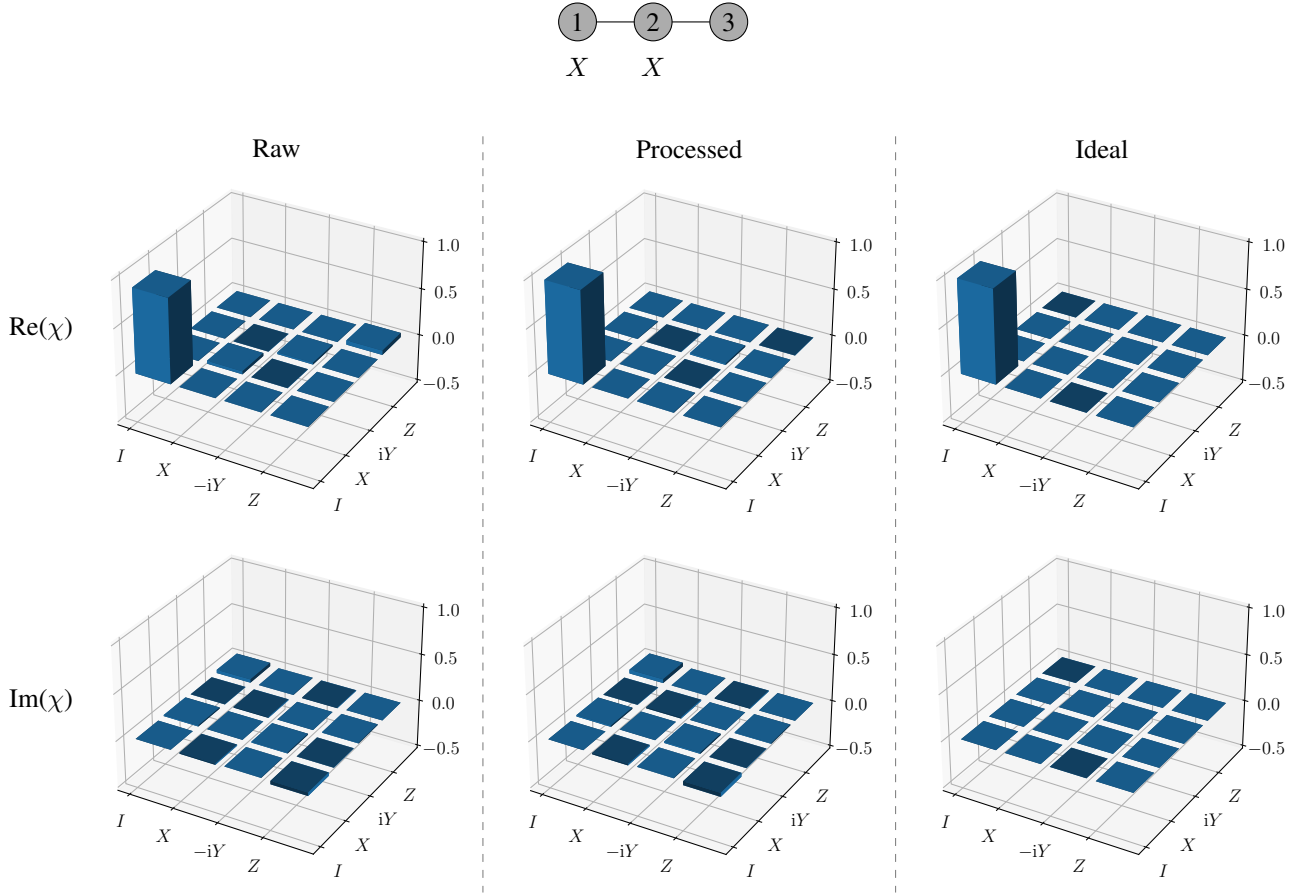


Figure A.4: Process tomography results (average χ matrix) for the implementation of the identity operation with a 3-qubit linear cluster state on the *ibmq_toronto* quantum processor. The diagram at the top shows the entangled 3-qubit linear cluster state with the measurements performed on each qubit. The χ matrix obtained without quantum readout error mitigation is shown on the left, the χ matrix obtained with quantum readout error mitigation is shown in the middle and the ideal χ matrix is shown on the right. The real part of each matrix is shown above and the imaginary part of each matrix is shown below.

of process tomography results, combining of counts to reduce statistical noise and construction of calibration matrices for quantum readout error mitigation were done in much the same way as for the exact 3-design implementation on the *ibmq_toronto* quantum processor. The only significant difference is that we applied the appropriate Pauli corrections to the density matrices of the output states constructed by state tomography before using them to do process tomography. The different χ matrices obtained by doing process tomography for the different measurement outcomes were used to calculate an average χ matrix for each cluster state implementation.

Process tomography results (average χ matrices) obtained for the implementation of the identity operation with the 3-qubit, 5-qubit and 7-qubit linear cluster states are displayed in Figs. A.4, A.5 and A.6 respectively. Non-zero real entries are clearly visible along the diagonals of constructed χ matrices, which confirms that depolarising noise was indeed the predominant type of noise for these qubits. Comparing the process tomography results for the different cluster states, we see that as the length of the linear cluster state used in the implementation increases, so does the depolarising noise in the implementation.

We now infer a value for the parameter p for the depolarising noise present in each implementation of the identity. To this end, we model depolarising noise in one of these measurement-based implementations as a channel in which depolarising noise is applied to the input state n times for an implementation with a n -qubit linear cluster state. For each implementation, we considered 10000 evenly spaced values of p in the range 0

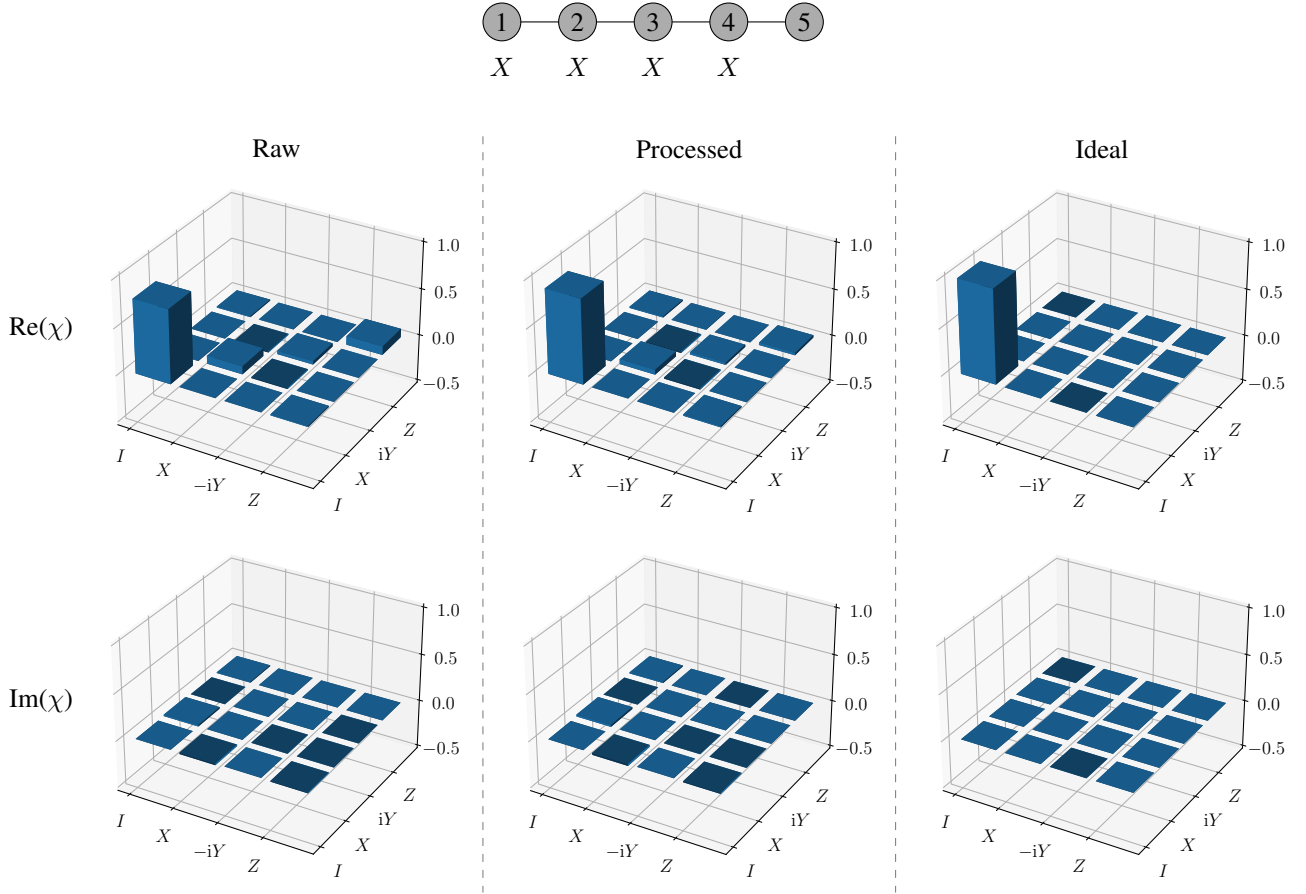


Figure A.5: Process tomography results (average χ matrix) for the implementation of the identity operation with a 5-qubit linear cluster state on the *ibmq_toronto* quantum processor. The diagram at the top shows the entangled 5-qubit linear cluster state with the measurements performed on each qubit. The χ matrix obtained without quantum readout error mitigation is shown on the left, the χ matrix obtained with quantum readout error mitigation is shown in the middle and the ideal χ matrix is shown on the right. The real part of each matrix is shown above and the imaginary part of each matrix is shown below.

to 1. For each p , we determined the χ matrix for the corresponding model channel and calculated the channel fidelity for the implementation of the identity using the average χ matrix obtained from process tomography. Our inferred value of p for a given implementation, is the one which yields the channel fidelity which is closest to 1. The values of p inferred for the different measurement-based implementations of the identity operation are given in Table A.2. These values quantify the increase in depolarising noise resulting from increasing the length of the linear cluster state. The values of p are greatly reduced by applying quantum readout error mitigation, which suggests that classical measurement errors are responsible for a substantial amount of depolarising noise in the implementations. Finally, we note that the values of p inferred for the identity implementation with the 5-qubit linear cluster state agree very well with the values of p inferred for the exact 3-design implementation with a 6-qubit linear cluster state on the same set of qubits. This shows that our methods used to infer the values of p are consistent.

A.4 Qubits used for the 2-design

The approximate measurement-based 2-design was implemented on five physical qubits of the *ibmq_sydney* quantum processor. Its qubit topology is identical to that of the *ibmq_toronto* quantum processor shown in Fig. A.1. The qubits 1 to 5 of the 5-qubit linear cluster state in the 2-design implementation were mapped

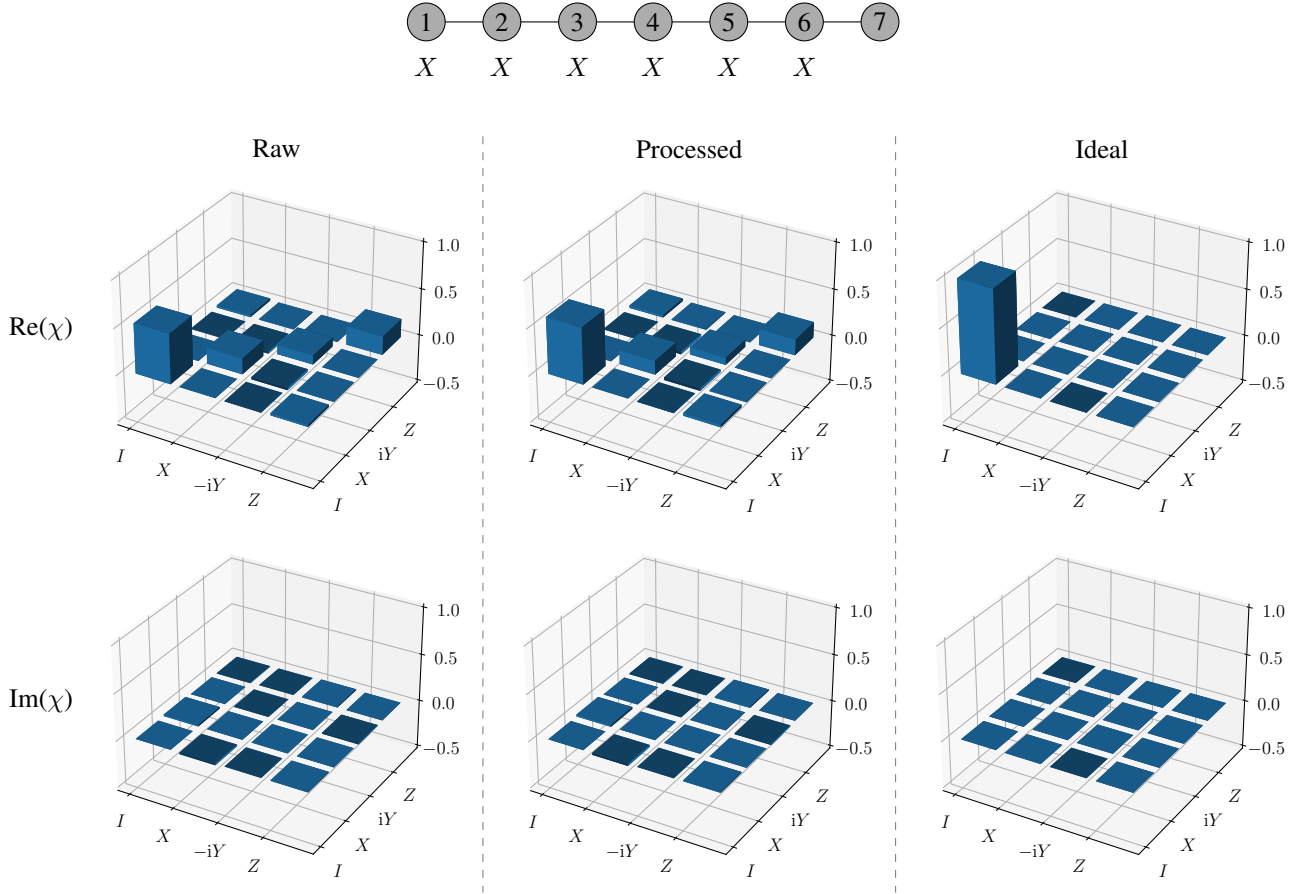


Figure A.6: Process tomography results (average χ matrix) for the implementation of the identity operation with a 7-qubit linear cluster state on the *ibmq_toronto* quantum processor. The diagram at the top shows the entangled 7-qubit linear cluster state with the measurements performed on each qubit. The χ matrix obtained without quantum readout error mitigation is shown on the left, the χ matrix obtained with quantum readout error mitigation is shown in the middle and the ideal χ matrix is shown on the right. The real part of each matrix is shown above and the imaginary part of each matrix is shown below.

Cluster State	p (Raw)	p (Processed)
3-qubit	0.042	0.004
5-qubit	0.062	0.027
7-qubit	0.127	0.099

Table A.2: Values of p inferred for implementations of the identity with different linear cluster states on the *ibmq_toronto* quantum processor. ‘Raw’ shows the values of p without quantum readout error mitigation. ‘Processed’ shows the values of p with quantum readout error mitigation.

onto the physical qubits 13, 14, 16, 19 and 22 of the *ibmq_sydney* quantum processor, in such a way that the input state was prepared on qubit 13 and the output state was retrieved from qubit 22. The relevant calibration information as obtained at the time of running the circuits for the approximate 2-design implementation is shown in Table A.3. These five connected qubits (see Fig. A.1) were chosen as they have unusually low error rates (in particular CX error rates) — much lower than the error rates of any five connected qubits on the *ibmq_toronto* quantum processor. This is why the *ibmq_sydney* quantum processor was used for this investigation instead of the *ibmq_toronto* quantum processor. The *ibmq_sydney* quantum processor was not considered for the exact 3-design implementation, as it does not have any six connected qubits with lower error rates than the six connected qubits of *ibmq_toronto* quantum processor. In particular, the error rates of qubit 25,

A.5. Qubits used for the identity

Qubit	T_1 (μs)	T_2 (μs)	\sqrt{X} Error	Readout Error
13	154.36	162.92	0.000172	0.0097
14	94.78	208.81	0.000206	0.0295
16	97.54	121.99	0.001255	0.0198
19	102.87	89.78	0.000382	0.0271
22	114.71	170.64	0.000233	0.0441

Qubit Pair	CX Error
13–14	0.00576
14–16	0.00621
16–19	0.01155
19–22	0.01168

Table A.3: Calibration information for the *ibmq_sydney* quantum processor as obtained at the time of running the circuits for the approximate 2-design. The single-qubit calibration information for the relevant qubits is shown on the left. T_1 and T_2 are the amplitude and phase damping time constants respectively of the qubits. The CX error rates for relevant qubit pairs are shown on the right.

Qubit	T_1 (μs)	T_2 (μs)	\sqrt{X} Error	Readout Error
16	123.42	135.48	0.000297	0.0154
19	114.98	123.10	0.000434	0.0116
22	110.50	148.90	0.000330	0.0191
25	125.41	114.64	0.000323	0.0108
24	121.49	155.26	0.000180	0.0093
23	98.92	40.59	0.000441	0.0553
21	76.60	56.22	0.000508	0.0177

Qubit Pair	CX Error
16–19	0.00758
19–22	0.01024
22–25	0.01053
25–24	0.01099
24–23	0.00892
23–21	0.01652

Table A.4: Calibration information for the *ibmq_toronto* quantum processor as obtained at the time of running the circuits for the identity implementation. The single-qubit calibration information for the relevant qubits is shown on the left. T_1 and T_2 are the amplitude and phase damping time constants respectively of the qubits. The CX error rates for relevant qubit pairs are shown on the right.

which is connected to the qubits used for the approximate 2-design implementation, are typically large.

A.5 Qubits used for the identity

The identity operation was implemented on the *ibmq_toronto* quantum processor (see Fig. A.1 for the qubit topology) by performing single-qubit measurements on linear cluster states of different lengths. Qubits 1 to 3 of the 3-qubit linear cluster state were mapped onto the physical qubits 16, 19 and 22, qubits 1 to 5 of the 5-qubit linear cluster state were mapped onto the physical qubits 16, 19, 22, 25 and 24 and qubits 1 to 7 of the 7-qubit linear cluster state were mapped onto the physical qubits 16, 19, 22, 25, 24, 23 and 21. Qubits were chosen in this way to ensure maximum possible overlap with the qubits used for the exact 3-design implementation. This allowed us to compare depolarising noise parameters inferred for the two implementations. The relevant calibration information as obtained at the time of running the circuits for the identity implementation is shown in Table A.4.

Appendix B

Investigating the effect of noise channels on the quality of unitary t -designs

B.1 Proof of well-definedness of noise models

Let $\{p_i, U_i\}$ and $\{q_i, V_i\}$ be exact unitary t -designs and let

$$\varepsilon(\rho) = \sum_k E_k \rho E_k^\dagger \quad (\text{B.1})$$

be a noise channel. We note that

$$\mathbb{E}_H^t(\rho^{\otimes t}) = \sum_i p_i \left(U_i \rho U_i^\dagger \right)^{\otimes t} = \sum_i q_i \left(V_i \rho V_i^\dagger \right)^{\otimes t} \quad (\text{B.2})$$

by the definition of an exact unitary t -design. Let $\tilde{\mathbb{E}}_{H,U}^t(\rho)$ and $\tilde{\mathbb{E}}_{H,V}^t(\rho)$ denote $\tilde{\mathbb{E}}_H^t(\rho)$ determined using $\{p_i, U_i\}$ and $\{q_i, V_i\}$ respectively, each with noise applied. For both noise models, we will show that $\tilde{\mathbb{E}}_{H,U}^t(\rho) = \tilde{\mathbb{E}}_{H,V}^t(\rho)$ for all ρ , from which it follows that the smallest ϵ such that inequality (4.13) holds for all density matrices is the same for the two ensembles.

B.1.1 Proof for the model where noise is applied before the unitary operations

Let ρ be a density matrix. Since $\varepsilon(\rho)$ is also a density matrix, it follows from Eq. (4.14) that $\tilde{\mathbb{E}}_{H,U}^t(\rho) = \mathbb{E}_H^t((\varepsilon(\rho))^{\otimes t})$ and $\tilde{\mathbb{E}}_{H,V}^t(\rho) = \mathbb{E}_H^t((\varepsilon(\rho))^{\otimes t})$, so that $\tilde{\mathbb{E}}_{H,U}^t(\rho) = \tilde{\mathbb{E}}_{H,V}^t(\rho)$.

B.1.2 Proof for the model where noise is applied after the unitary operations

Let ρ be a density matrix. Substituting Eq. (B.1) into Eq. (4.15) and using the algebraic properties of the tensor product, we obtain

$$\begin{aligned}\tilde{\mathbb{E}}_{H,U}^t(\rho) &= \sum_i p_i \left(\sum_k E_k U_i \rho U_i^\dagger E_k^\dagger \right)^{\otimes t} \\ &= \sum_i p_i \sum_{k_1} \sum_{k_2} \cdots \sum_{k_t} E_{k_1} U_i \rho U_i^\dagger E_{k_1}^\dagger \otimes E_{k_2} U_i \rho U_i^\dagger E_{k_2}^\dagger \otimes \cdots \otimes E_{k_t} U_i \rho U_i^\dagger E_{k_t}^\dagger \\ &= \sum_{k_1} \sum_{k_2} \cdots \sum_{k_t} \sum_i p_i \left(\bigotimes_{j=1}^t E_{k_j} \right) \left(U_i \rho U_i^\dagger \right)^{\otimes t} \left(\bigotimes_{j=1}^t E_{k_j}^\dagger \right).\end{aligned}\quad (\text{B.3})$$

Multiplying both sides of Eq. (B.2) by $\bigotimes_{j=1}^t E_{k_j}$ from the left and by $\bigotimes_{j=1}^t E_{k_j}^\dagger$ from the right, we get

$$\sum_i p_i \left(\bigotimes_{j=1}^t E_{k_j} \right) \left(U_i \rho U_i^\dagger \right)^{\otimes t} \left(\bigotimes_{j=1}^t E_{k_j}^\dagger \right) = \sum_i q_i \left(\bigotimes_{j=1}^t E_{k_j} \right) \left(V_i \rho V_i^\dagger \right)^{\otimes t} \left(\bigotimes_{j=1}^t E_{k_j}^\dagger \right).\quad (\text{B.4})$$

Performing a term-by-term replacement in the t -fold summation of Eq. (B.3) by substituting in Eq. (B.4) yields

$$\tilde{\mathbb{E}}_{H,U}^t(\rho) = \sum_{k_1} \sum_{k_2} \cdots \sum_{k_t} \sum_i q_i \left(\bigotimes_{j=1}^t E_{k_j} \right) \left(V_i \rho V_i^\dagger \right)^{\otimes t} \left(\bigotimes_{j=1}^t E_{k_j}^\dagger \right),\quad (\text{B.5})$$

from which it follows that $\tilde{\mathbb{E}}_{H,U}^t(\rho) = \tilde{\mathbb{E}}_{H,V}^t(\rho)$ if we then perform our original calculation in Eq. (B.3) in reverse.

B.2 Analytic results for the 1-design

Using the Pauli 1-design, one can show that $\mathbb{E}_H^1(\rho) = \frac{1}{2}I$ for all density matrices ρ . Let

$$\varepsilon(\rho) = \sum_k E_k \rho E_k^\dagger\quad (\text{B.6})$$

be an arbitrary noise channel. For the model where noise is applied before the unitary operations, we will show that the quality of the 1-design is completely unaffected by an arbitrary noise channel, and for the model where noise is applied after the unitary operations, we will show that the quality of the 1-design is unaffected by noise, unless amplitude damping is applied.

B.2.1 Analytic results for the model where noise is applied before the unitary operations

For any density matrix ρ , $\varepsilon(\rho)$ is a density matrix, and so it follows from Eq. (4.14) that $\tilde{\mathbb{E}}_H^1(\rho) = \mathbb{E}_H^1(\varepsilon(\rho)) = \frac{1}{2}I$. It therefore follows that inequality (4.13) can be satisfied with $\epsilon = 0$, and so the quality of the 1-design is completely unaffected by an arbitrary noise channel.

B.2.2 Analytic results for the model where noise is applied after the unitary operations

Let ρ be a density matrix and let $\{p_i, U_i\}$ be an exact 1-design. Substituting Eq. (B.6) into Eq. (4.15), we have

$$\begin{aligned}
 \tilde{\mathbb{E}}_H^1(\rho) &= \sum_i p_i \sum_k E_k U_i \rho U_i^\dagger E_k^\dagger \\
 &= \sum_k E_k \left(\sum_i p_i U_i \rho U_i^\dagger \right) E_k^\dagger \\
 &= \sum_k E_k \left(\frac{1}{2} I \right) E_k^\dagger \\
 &= \frac{1}{2} \sum_k E_k E_k^\dagger \\
 &= \frac{1}{2} I,
 \end{aligned} \tag{B.7}$$

where we have recognised $\mathbb{E}_H^1(\rho) = \frac{1}{2} I$ and assumed that E_k are Hermitian matrices, which is the case for all the noise channels defined in Sec. 4.2.1 except the amplitude damping channel. Hence, inequality (4.13) can again be satisfied with $\epsilon = 0$ for all the noise channels defined in Sec. 4.2.1 except the amplitude damping channel.

For the amplitude damping channel, we have

$$\tilde{\mathbb{E}}_H^1(\rho) = \frac{1}{2} \sum_k E_k E_k^\dagger = \frac{1}{2} \begin{pmatrix} 1 + \lambda & 0 \\ 0 & 1 - \lambda \end{pmatrix}, \tag{B.8}$$

and so,

$$\tilde{\mathbb{E}}_H^1(\rho) - (1 - \epsilon) \mathbb{E}_H^1(\rho) = \frac{1}{2} \begin{pmatrix} \epsilon + \lambda & 0 \\ 0 & \epsilon - \lambda \end{pmatrix}. \tag{B.9}$$

It follows that for $(1 - \epsilon) \mathbb{E}_H^1(\rho) \leq \tilde{\mathbb{E}}_H^1(\rho)$ to hold, we must have $\epsilon \geq \lambda$, so that all the eigenvalues of $\tilde{\mathbb{E}}_H^1(\rho) - (1 - \epsilon) \mathbb{E}_H^1(\rho)$ are non-negative, which ensures that $\tilde{\mathbb{E}}_H^1(\rho) - (1 - \epsilon) \mathbb{E}_H^1(\rho)$ is positive semidefinite. Similarly, $\epsilon \geq \lambda$ ensures that $\tilde{\mathbb{E}}_H^1(\rho) \leq (1 + \epsilon) \mathbb{E}_H^1(\rho)$ holds. Hence, inequality (4.13) can be satisfied with $\epsilon = \lambda$ for the amplitude damping channel, and so $\epsilon = \lambda$ quantifies the effect of the amplitude damping channel on the quality of the 1-design for the model where noise is applied after the unitary operations.

B.3 Numeric results for the model where noise is applied after the unitary operations

B.3.1 Flip channels

The effect of the bit flip channel (see Eq. (4.1)) on the quality of the 2-design is shown in Fig. B.1a. For the model where noise is applied after the unitary operations, ϵ versus p is a parabola with maximum at $p = 0.5$, just as for the model where noise is applied before the unitary operations (see Fig. 4.1a). However, the values of ϵ obtained for a given p and r_t are slightly smaller than those obtained for the model where noise is applied before the unitaries. For the 4-design, ϵ versus p for the model where noise is applied after the unitary operations has the same sinusoidal shape as for the model where noise is applied before the unitaries (see Fig. 4.1b), but the values of ϵ are slightly smaller, just as for the 2-design. Hence t -designs show reduced sensitivity to bit flips

B.3. Numeric results for the model where noise is applied after the unitary operations

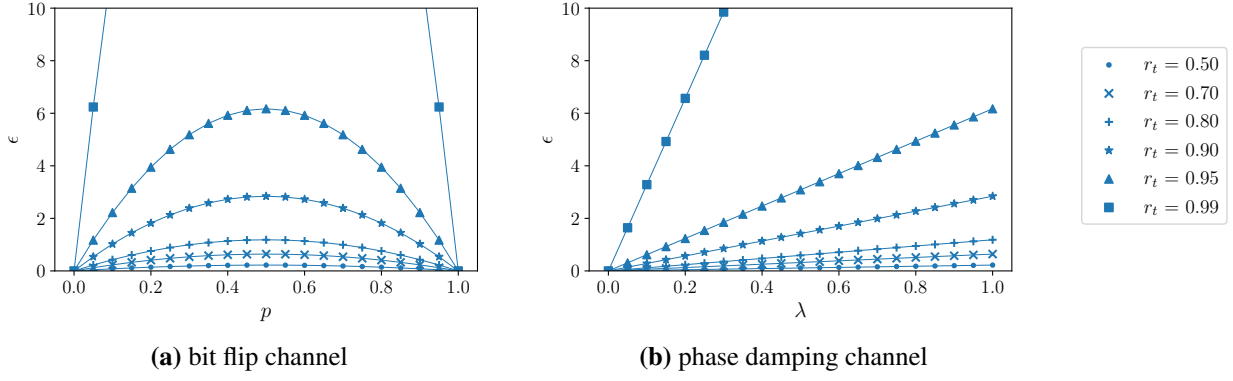


Figure B.1: Effect of the (a) bit flip channel (see Eq. (4.1)) and (b) phase damping channel (see Eq. (4.4)) on the quality of the 2-design for the model where noise is applied after the unitary operations, for different truncation radii r_t .

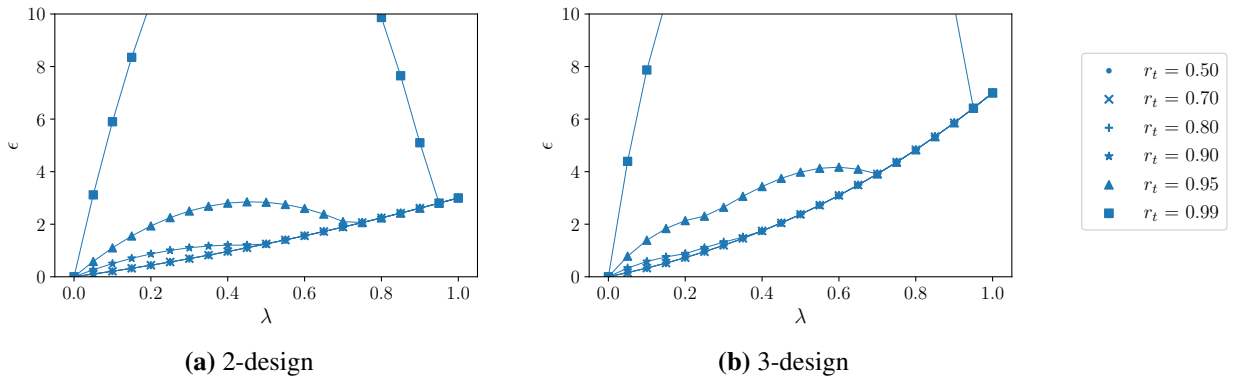


Figure B.2: Effect of the amplitude damping channel (see Eq. (4.8)) on the quality of the (a) 2-design and (b) 3-design for the model where noise is applied after the unitary operations, for different truncation radii r_t .

when applied after the unitary operations.

Numeric results obtained for the phase flip channel (see Eq. (4.2)) and the bit and phase flip channel (see Eq. (4.3)) are identical to those obtained for the bit flip channel (shown in Fig. B.1a). We again investigate similarities and differences by determining the regions of acceptable quality (see Appendix C.1). For all three flip channels, the region of acceptable quality has a spherical shape, that is, it is similar in shape to the region of acceptable quality for the depolarising noise channel for the model where noise is applied before the unitary operations. Hence, we observe the transformation of each of the three flip channels into a depolarising channel when the flip channels are applied to states which have been randomised by the unitary operators, which explains why the results are the same for all three flip channels. A 2-design's ability to transform an arbitrary noise channel into a depolarising noise channel is exploited in randomised benchmarking [92–94, 98–100].

B.3.2 Phase damping channel

For the phase damping channel (see Eq. (4.4)), ϵ versus λ for the model where noise is applied after the unitary operations has the same shape as for the model where noise is applied before the unitary operations, for both the 2-design (shown in Fig. B.1b) and the 4-design, but the values of ϵ are slightly smaller for the model where noise is applied after the unitaries. Just as for the phase flip channel, the region of acceptable quality for the phase damping channel has a spherical shape for the model where noise is applied after the unitaries (see Appendix C.1).

B.3.3 Amplitude damping channel

For the amplitude damping channel (see Eq. (4.8)), ϵ versus λ for the model where noise is applied after the unitary operations is once again similar to ϵ versus λ for the model where noise is applied before the unitary operations, for the 2-design (shown in Fig. B.2a), the 3-design (shown in Fig. B.2b), the 4-design and the 5-design. The most notable differences are that the values of ϵ at the maxima are much smaller, the value of ϵ attained for $\lambda = 1$ is much larger, and the anomaly occurs for much smaller λ and much larger r_t for the model where noise is applied after the unitary operations. It is also worth noting that numeric results obtained for the 3-design differ significantly from those obtained for the 2-design (see Figs. B.2a and B.2b) and that numeric results obtained for the 5-design differ significantly from those obtained for the 4-design. The regions of acceptable quality are once again discussed in Appendix C.1.

B.3.4 Depolarising noise channel

For the depolarising noise channel, we were able to show that the values of ϵ obtained for the two noise models are equal. The proof is given in Appendix B.4.

B.4 Proof of equivalence of noise models for the depolarising noise channel

Let $\{p_i, U_i\}$ be an exact unitary t -design. For the model where noise is applied before the unitary operations, we substitute Eq. (2.27) into Eq. (4.14) to obtain

$$\tilde{\mathbb{E}}_H^t(\rho) = \sum_i p_i \left(U_i \left(\frac{p}{2} I + (1-p)\rho \right) U_i^\dagger \right)^{\otimes t} = \sum_i p_i \left(\frac{p}{2} I + (1-p) U_i \rho U_i^\dagger \right)^{\otimes t} \quad (\text{B.10})$$

and for the model where noise is applied after the unitary operations, we substitute Eq. (2.27) into Eq. (4.15) to obtain

$$\tilde{\mathbb{E}}_H^t(\rho) = \sum_i p_i \left(\frac{p}{2} I + (1-p) U_i \rho U_i^\dagger \right)^{\otimes t}. \quad (\text{B.11})$$

Hence, for the depolarising noise channel, $\tilde{\mathbb{E}}_H^t(\rho)$ is the same for the two noise models, so that the smallest ϵ such that inequality (4.13) holds for all density matrices is the same for the two noise models.

Appendix C

State dependence of the effect of noise channels on the quality of single-qubit t -designs

C.1 Visualisation of regions of acceptable quality

We determine and visualise regions of the Bloch sphere for which noisy t -designs are able to replicate the moments of the uniform Haar ensemble, with a predefined accuracy, up to order t . In the sections which follow, we will refer to these regions as regions of acceptable quality.

C.1.1 Implementation

To determine regions of acceptable quality, we first generate 20 evenly spaced values of $x \in [-1, 1]$, 20 evenly spaced values of $y \in [-1, 1]$ and 20 evenly spaced values of $z \in [-1, 1]$, thereby obtaining a sample of $20^3 = 8000$ coordinates (x, y, z) in a cube which encloses the Bloch sphere. Points which lie in the Bloch sphere, that is points such that $\sqrt{x^2 + y^2 + z^2} \leq 1$, correspond to valid states for which we obtain density matrices using Eq. (2.20). We now visualise the region of acceptable quality for a given $t \in \{2, 3, 4, 5\}$, a given noise channel and a given noise model as follows. For each density matrix ρ , we calculate $\mathbb{E}_H^t(\rho^{\otimes t})$ and $\tilde{\mathbb{E}}_H^t(\rho)$ using the icosahedral group [188] and plot the point (x, y, z) corresponding to ρ if and only if inequality (4.13) can be satisfied with $\epsilon \leq 0.5$.

Note that we have arbitrarily chosen $\epsilon = 0.5$ as a threshold for acceptable quality, merely for the purposes of visualisation. In practice, the value of ϵ required for the quality of a noisy t -design to be acceptable really depends on the application. Furthermore, we note that it is not possible to use the ϵ for which an exact $(t - 1)$ -design is an approximate t -design as a threshold value for a noisy t -design, since this ϵ depends on the ensemble of unitaries that make up the exact $(t - 1)$ -design. In fact, there are exact $(t - 1)$ -designs for which $\epsilon > 0.5$ when used as approximate t -designs.

C.1.2 Numeric results for the model where noise is applied before the unitary operations

C.1.2.1 Bit flip channel

Regions of acceptable quality for the 2-design affected by the bit flip channel (see Eq. (4.1)) are shown in Fig. C.1, for various different p . As expected, the size of the region of acceptable quality decreases, and

C.1. Visualisation of regions of acceptable quality




Figure C.1: Regions of acceptable quality for the 2-design affected by the bit flip channel (see Eq. (4.1)) for the model where noise is applied before the unitary operations, for different p .

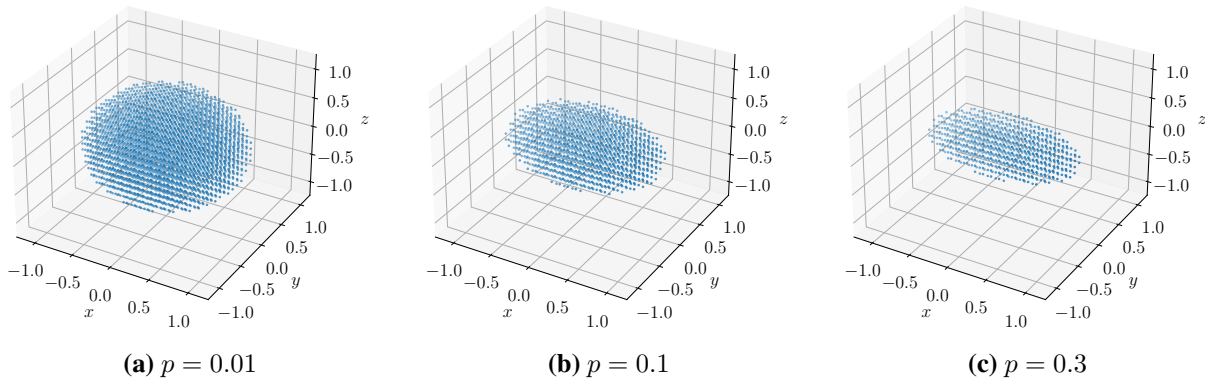


Figure C.2: Regions of acceptable quality for the 4-design affected by the bit flip channel (see Eq. (4.1)) for the model where noise is applied before the unitary operations, for different p .

its shape becomes more pronounced, as the bit flip probability p increases. For all p , the region of acceptable quality is an ellipsoid along the x -axis, that is, the region of acceptable quality is similar in shape and orientation to the region into which the Bloch sphere is deformed by the bit flip channel. This can be understood as follows. Since bit flips are performed by applying the Pauli X operator to a state, states along the x -axis, which are closer to the eigenstates of the Pauli X operator, are less affected by bit flips. Hence the quality remains acceptable for states along the x -axis even for a large bit flip probability. For the 4-design, the region of acceptable quality is similar in shape and orientation to the region of acceptable quality for the 2-design, but smaller in size for a fixed p (see Fig. C.2). This simply confirms that the 4-design is more sensitive to bit flips than the 2-design.

C.1.2.2 Phase flip channel

Regions of acceptable quality for the 2-design affected by the phase flip channel (see Eq. (4.2)) are shown in Fig. C.3. For all p , the region of acceptable quality is an ellipsoid along the z -axis, that is, the region of acceptable quality is similar in shape and orientation to the region into which the Bloch sphere is deformed by the phase flip channel. The shape and orientation of the region of acceptable quality can be attributed to the fact that states along the z -axis, which are closer to the eigenstates of the Pauli Z operator, are less affected by phase flips since phase flips are performed by applying the Pauli Z operator to a state. For the 4-design, the regions of acceptable quality are again similar in shape and orientation, but smaller in size (see Fig. C.4).

C.1. Visualisation of regions of acceptable quality

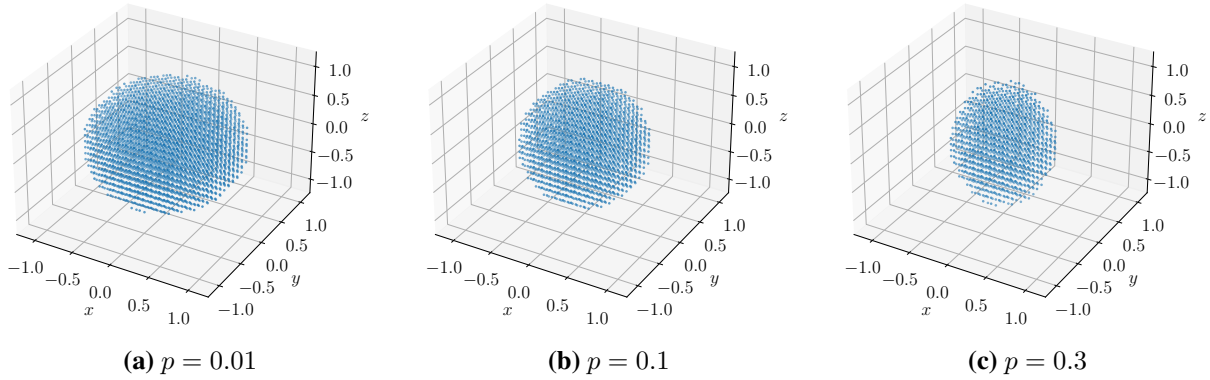


Figure C.3: Regions of acceptable quality for the 2-design affected by the phase flip channel (see Eq. (4.2)) for the model where noise is applied before the unitary operations, for different p .




Figure C.4: Regions of acceptable quality for the 4-design affected by the phase flip channel (see Eq. (4.2)) for the model where noise is applied before the unitary operations, for different p .

C.1.2.3 Bit and phase flip channel

As shown in Figs. C.5 and C.6, the regions of acceptable quality are similar in shape and orientation to the region into which the Bloch sphere is deformed by the bit and phase flip channel, that is, an ellipsoid along the y -axis. Since bit and phase flips are performed by applying the Pauli Y operator to a state, states along the y -axis, which are closer to the eigenstates of the Pauli Y operator, are less affected by bit and phase flips.

C.1.2.4 Phase damping channel

Regions of acceptable quality for the 2-design affected by the phase damping channel (see Eq. (4.4)) are shown in Fig. C.7, for different λ . Just as for the phase flip channel, the regions of acceptable quality are ellipsoids along the z -axis. This is to be expected, since the phase damping channel is equivalent to the phase flip channel, up to a reparameterisation (see Eq. (4.7)). For the 4-design, the regions of acceptable quality are again similar in shape, but smaller in size (see Fig. C.8).

C.1.2.5 Amplitude damping channel

Regions of acceptable quality for the 2-design affected by the amplitude damping channel (see Eq. (4.8)) are shown in Fig. C.9, for various different λ . For small λ , the region of acceptable quality is shifted up slightly towards the state $|0\rangle$, which seems to resemble the way in which the amplitude damping channel shrinks and shifts the Bloch sphere up towards the state $|0\rangle$. However, the quality of the 2-design is not acceptable for states

C.1. Visualisation of regions of acceptable quality

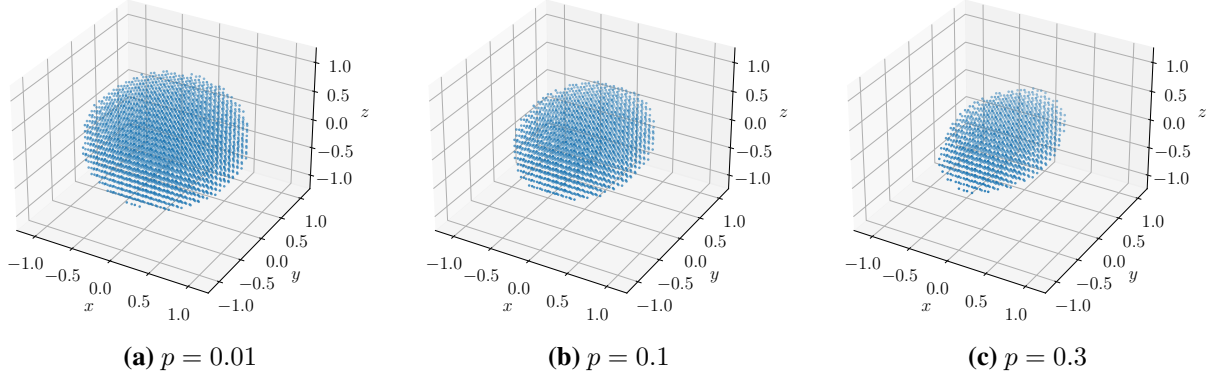


Figure C.5: Regions of acceptable quality for the 2-design affected by the bit and phase flip channel (see Eq. (4.3)) for the model where noise is applied before the unitary operations, for different p .




Figure C.6: Regions of acceptable quality for the 4-design affected by the bit and phase flip channel (see Eq. (4.3)) for the model where noise is applied before the unitary operations, for different p .

close to $|0\rangle$ for large λ . This is a result of setting the threshold for acceptable quality at $\epsilon = 0.5$. As discussed in Sec. 4.4.2.3, $\epsilon \rightarrow 1$ as $\lambda \rightarrow 1$ (see Fig. 4.5), so that states close to $|0\rangle$ remain above the threshold of $\epsilon = 0.5$ for large λ . Just as for the other noise channels, regions of acceptable quality for the 3-design are identical to those obtained for the 2-design. Regions of acceptable quality for the 4-design affected by the amplitude damping channel are shown in Fig. C.10. For small λ , the region of acceptable quality is shifted higher than for the 2-design, and for large λ , the region around $|0\rangle$ for which the quality is not acceptable is much larger than for the 2-design. Regions of acceptable quality for the 5-design affected by the amplitude damping channel are very similar to those obtained for the 4-design.

C.1.2.6 Depolarising noise channel

Regions of acceptable quality for the 2-design affected by the depolarising noise channel (see Eq. (2.27)) are shown in Fig. C.11. The size of the region of acceptable quality decreases with increasing p , as expected. For all p , the region of acceptable quality is a sphere centred at the origin, that is, the region of acceptable quality is similar in shape and orientation to the region into which the Bloch sphere is deformed by the depolarising channel. Since the depolarising channel replaces a state by the maximally mixed state with probability p , states near the centre of the Bloch sphere, which are closer to the maximally mixed state, are least affected by depolarising noise. Just as for the other noise channels considered, the region of acceptable quality for the 4-design is similar in shape to that of the 2-design, but smaller in size (see Fig. C.12), once again confirming




Figure C.7: Regions of acceptable quality for the 2-design affected by the phase damping channel (see Eq. (4.4)) for the model where noise is applied before the unitary operations, for different λ .




Figure C.8: Regions of acceptable quality for the 4-design affected by the phase damping channel (see Eq. (4.4)) for the model where noise is applied before the unitary operations, for different λ .

that the 4-design is more sensitive to noise than the 2-design.

C.1.3 Numeric results for the model where noise is applied after the unitary operations

C.1.3.1 Flip channels

Regions of acceptable quality for the 2-design affected by the bit flip channel (see Eq. (4.1)) are shown in Fig. C.13. For each p , the region of acceptable quality is a sphere centred at the origin, that is, the region of acceptable quality is similar in shape and orientation to the region of acceptable quality for the 2-design affected by the depolarising noise channel for the model where noise is applied before the unitary operations. For the phase flip channel (see Eq. (4.2)) and the bit and phase flip channel (see Eq. (4.3)), the regions of acceptable quality are identical to those for the bit flip channel (shown in Fig. C.13). Thus we observe the transformation of each of the three flip channels into a depolarising channel when the flip channels are applied to states which have been randomised by unitary operators from the 2-design. For all three flip channels, the region of acceptable quality for the 4-design is similar in shape and orientation to that of the 2-design, but smaller in size.

C.1.3.2 Phase damping channel

For the phase damping channel (see Eq. (4.4)), the regions of acceptable quality for both the 2-design (shown in Fig. C.14) and the 4-design are spheres centred at the origin, just as for the phase flip channel.

C.1. Visualisation of regions of acceptable quality

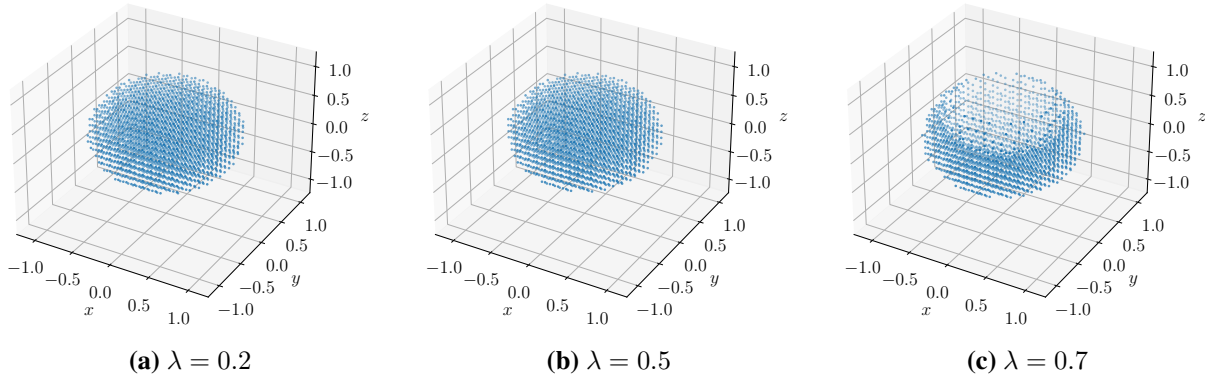


Figure C.9: Regions of acceptable quality for the 2-design affected by the amplitude damping channel (see Eq. (4.8)) for the model where noise is applied before the unitary operations, for different λ .




Figure C.10: Regions of acceptable quality for the 4-design affected by the amplitude damping channel (see Eq. (4.8)) for the model where noise is applied before the unitary operations, for different λ .

C.1.3.3 Amplitude damping channel

Regions of acceptable quality for the 2-design affected by the amplitude damping channel (see Eq. (4.8)) are shown in Fig. C.15, for different λ . For small λ , the region of acceptable quality is once again a sphere centred at the origin, that is, the region of acceptable quality is once again similar to the region of acceptable quality for the 2-design affected by the depolarising noise channel for the model where noise is applied before the unitary operations. For large λ , the region of acceptable quality disappears. As discussed in Appendix B.3, the ϵ attained for $\lambda = 1$ is much larger for the model where noise is applied after the unitary operations (see Fig. B.2a). Hence, all states in the Bloch sphere remain above the threshold of $\epsilon = 0.5$ for large λ . Regions of acceptable quality for the 3-design affected by the amplitude damping channel are shown in Fig. C.16. For small λ , the region of acceptable quality is a hollow spherical shell centred at the origin. Thus when amplitude damping is applied after the unitary operations, it is states close to the maximally mixed state, and not states close to $|0\rangle$, for which the quality is not acceptable. For large λ , the region of acceptable quality once again disappears. For the 4-design and the 5-design, the region of acceptable quality vanishes for $\lambda = 0.2$, $\lambda = 0.5$ and $\lambda = 0.7$.

C.1.3.4 Depolarising noise channel

For the depolarising noise channel, we prove that the two noise models are equivalent (see Appendix B.4).




Figure C.11: Regions of acceptable quality for the 2-design affected by the depolarising noise channel (see Eq. (2.27)) for the model where noise is applied before the unitary operations, for different p .

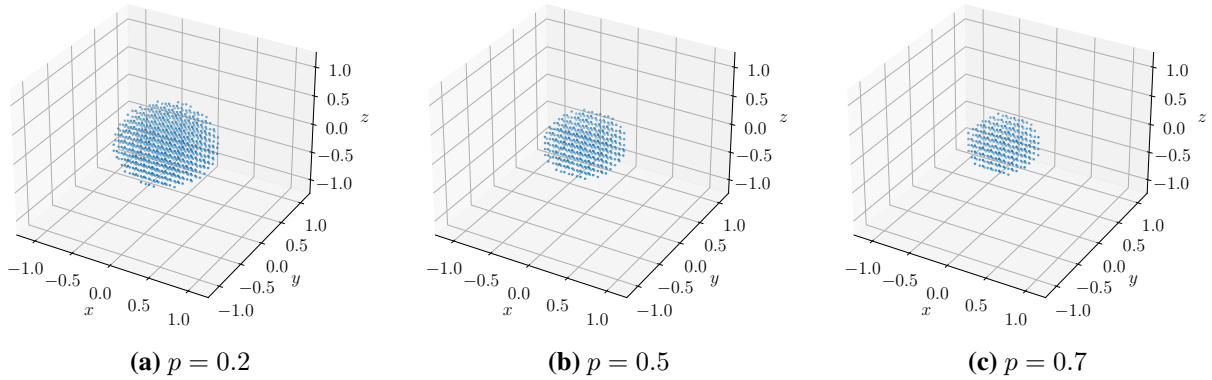


Figure C.12: Regions of acceptable quality for the 4-design affected by the depolarising noise channel (see Eq. (2.27)) for the model where noise is applied before the unitary operations, for different p .

C.2 Numeric results for different truncations of the polar angle

We investigate the dependence of ϵ versus p (or ϵ versus λ) on the truncation of the polar angle θ_t . In particular, we consider $\theta_t \in \{\frac{\pi}{6}, \frac{\pi}{3}, \frac{\pi}{2}, \frac{2\pi}{3}, \frac{5\pi}{6}, \pi\}$ for a fixed truncation radius of $r_t = 0.95$ (with $\phi_t = 2\pi$) and obtain results numerically using the method described in Sec. 4.4.1.

C.2.1 Numeric results for the model where noise is applied before the unitary operations

C.2.1.1 Bit flip channel

For the bit flip channel (see Eq. (4.1)), ϵ versus p is independent of θ_t . This can be explained as follows. For a fixed truncation radius, states along the positive z -axis are among the furthest from the eigenstates of the Pauli X operator. Hence these states are among the most sensitive to bit flips and therefore require the largest ϵ to satisfy inequality (4.13). Since the states along the positive z -axis are included in the sample of density matrices for all θ_t , the value of ϵ obtained for a given p is simply this largest ϵ for all θ_t , so that ϵ versus p is the same for all θ_t .

C.2.1.2 Phase flip channel

For the phase flip channel (see Eq. (4.2)), the maximum of ϵ versus p increases as we increase θ_t , up to $\theta_t = \frac{\pi}{2}$, after which the maximum remains more or less constant, for both the 2-design (shown in Fig. C.17a) and the

C.2. Numeric results for different truncations of the polar angle

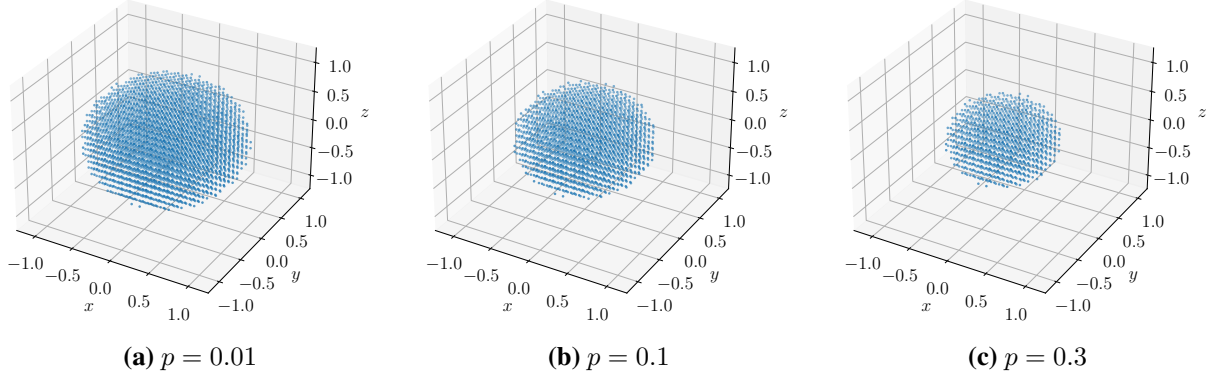


Figure C.13: Regions of acceptable quality for the 2-design affected by the bit flip channel (see Eq. (4.1)) for the model where noise is applied after the unitary operations, for different p .

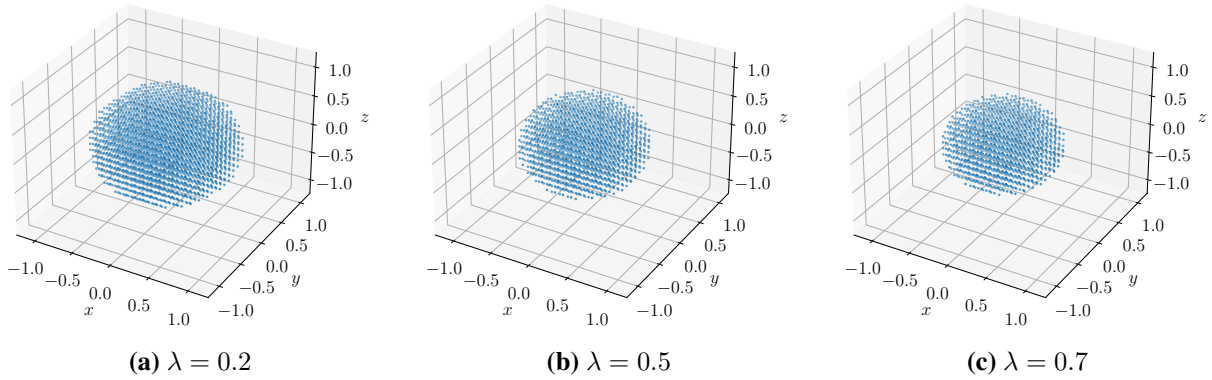


Figure C.14: Regions of acceptable quality for the 2-design affected by the phase damping channel (see Eq. (4.4)) for the model where noise is applied after the unitary operations, for different λ .

4-design. As θ_t increases from 0 to $\frac{\pi}{2}$, the sample of density matrices is expanded to include states which are further from the eigenstates of the Pauli Z operator and therefore more sensitive to phase flips, which results in ϵ increasing. When $\theta_t = \frac{\pi}{2}$, the states along the equator of the Bloch sphere, which are furthest from the eigenstates of the Pauli Z operator and therefore most sensitive to phase flips, are included in the sample of density matrices and so further increasing θ_t does not further increase ϵ .

C.2.1.3 Bit and phase flip channel

For the bit and phase flip channel (see Eq. (4.3)), ϵ versus p is independent of θ_t . Just as for the bit flip channel, the states along the positive z -axis, which are included in the sample of density matrices for all θ_t , require the largest ϵ to satisfy inequality (4.13), and so the value of ϵ obtained for a given p is simply this largest ϵ for all θ_t . The states along the positive z -axis are among the furthest from the eigenstates of the Pauli Y operator and are therefore among the most sensitive to bit and phase flips, which is why they require the largest ϵ to satisfy inequality (4.13).

C.2.1.4 Phase damping channel

For the phase damping channel (see Eq. (4.4)), the gradient of ϵ versus λ increases as we increase θ_t , up to $\theta_t = \frac{\pi}{2}$, after which the gradient remains more or less constant, for both the 2-design (shown in Fig. C.17b) and the 4-design. As expected, the dependence of ϵ versus λ on θ_t for the phase damping channel is similar to the

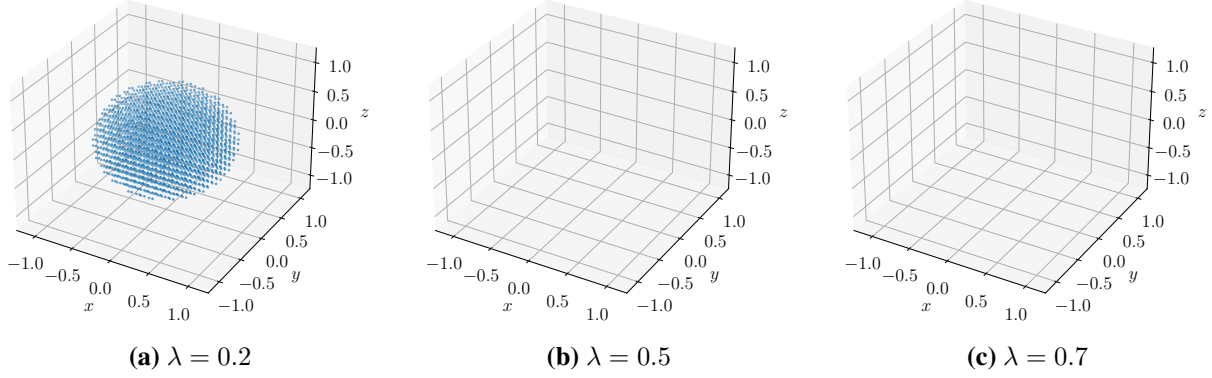


Figure C.15: Regions of acceptable quality for the 2-design affected by the amplitude damping channel (see Eq. (4.8)) for the model where noise is applied after the unitary operations, for different λ .

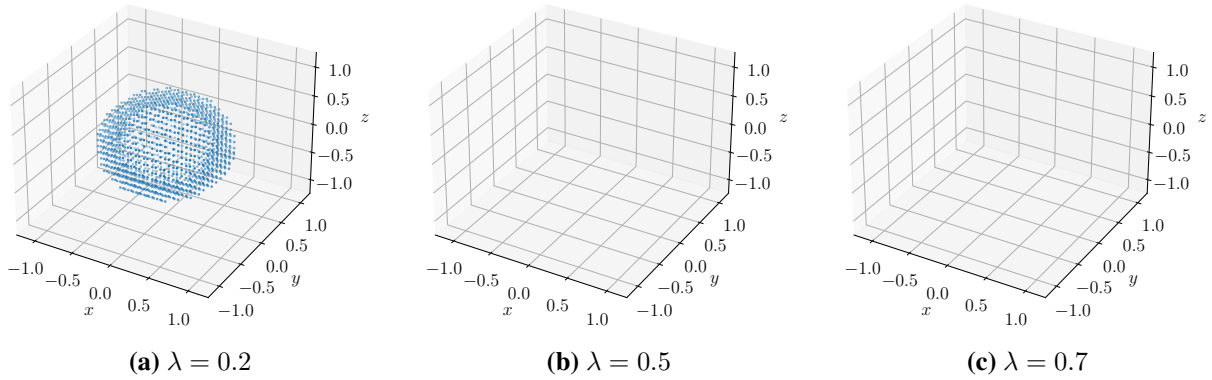


Figure C.16: Regions of acceptable quality for the 3-design affected by the amplitude damping channel (see Eq. (4.8)) for the model where noise is applied after the unitary operations, for different λ .

dependence of ϵ versus p on θ_t for the phase flip channel.

C.2.1.5 Amplitude damping channel

For the amplitude damping channel (see Eq. (4.8)), the maximum of ϵ versus λ increases as we increase θ_t , for the 2-design (shown in Fig. C.18), the 3-design, the 4-design and the 5-design. As θ_t increases, the sample of density matrices is expanded to include states which are further from the state $|0\rangle$ and therefore more sensitive to amplitude damping, which results in ϵ increasing.

C.2.1.6 Depolarising noise channel

For the depolarising noise channel (see Eq. (2.27)), ϵ versus p is independent of θ_t . This is because all states at a given radial distance from the maximally mixed state satisfy inequality (4.13) with the same value of ϵ , irrespective of the polar angles of the states.

C.2.2 Numeric results for the model where noise is applied after the unitary operations

For the model where noise is applied after the unitary operations, numeric results are independent of θ_t for all six noise channels. This reflects our observations from Appendix C.1, that is, when the noise channel is applied after the states have been randomised by the unitary operations, all states at a given radial distance from the

C.3. Numeric results for different truncations of the azimuthal angle




Figure C.17: Effect of the (a) phase flip channel (see Eq. (4.2)) and (b) phase damping channel (see Eq. (4.4)) on the quality of the 2-design for the model where noise is applied before the unitary operations, for different truncations of the polar angle θ_t , for a fixed truncation radius of $r_t = 0.95$.




Figure C.18: Effect of the amplitude damping channel (see Eq. (4.8)) on the quality of the 2-design for the model where noise is applied before the unitary operations, for different truncations of the polar angle θ_t , for a fixed truncation radius of $r_t = 0.95$. The full set of results is shown on the left and a zoomed-in region is shown enlarged on the right.

maximally mixed state satisfy inequality (4.13) with the same value of ϵ , irrespective of the polar angles of the states.

C.3 Numeric results for different truncations of the azimuthal angle

We investigate the dependence of ϵ versus p (or ϵ versus λ) on the truncation of the azimuthal angle ϕ_t . In particular, we consider $\phi_t \in \{\frac{\pi}{6}, \frac{\pi}{3}, \frac{\pi}{2}, \frac{2\pi}{3}, \frac{5\pi}{6}, \pi, \frac{7\pi}{6}, \frac{4\pi}{3}, \frac{3\pi}{2}, \frac{5\pi}{3}, \frac{11\pi}{6}, 2\pi\}$ for a fixed truncation radius of $r_t = 0.95$ (with $\theta_t = \pi$) and obtain results numerically using the method described in Sec. 4.4.1.

C.3.1 Numeric results for the model where noise is applied before the unitary operations

Numeric results are independent of ϕ_t for all six noise channels. For the bit flip channel and the bit and phase flip channel, this can be attributed to the fact that states along the z -axis, which require the largest ϵ to satisfy inequality (4.13), are included in the sample of density matrices for all ϕ_t , so that the value of ϵ obtained is simply this largest ϵ for all ϕ_t . For the phase flip channel, the phase damping channel, the amplitude damping channel and the depolarising noise channel, this can be attributed to the fact that the smallest ϵ such that inequality (4.13) holds for a given state remains unchanged when that state is rotated about the z -axis, so that the value of ϵ obtained is completely independent of ϕ_t .

C.3.2 Numeric results for the model where noise is applied after the unitary operations

For the model where noise is applied after the unitary operations, numeric results are also independent of ϕ_t for all six noise channels.

Appendix D

Measurement-based interleaved randomised benchmarking using IBM processors

D.1 Qubits used for fidelity estimation of universal gates

To perform fidelity estimation for a universal single-qubit set, we implemented reference sequences with lengths $m \in \{1, 2, 3\}$ and interleaved sequences with $m \in \{1, 2, 3\}$, for both the 2-qubit Hadamard gate and the 3-qubit T gate, on the *ibm_hanoi* quantum processor. The qubit topology of the *ibm_hanoi* quantum processor is shown in Fig. D.1, with the qubits used in the implementations shaded grey. These qubits were chosen for their low error rates compared to other qubits. For each m (reference or interleaved), the required linear cluster state was prepared along the length of shaded qubits, starting at qubit 17. For example, to implement the interleaved sequence for the 3-qubit T gate with $m = 1$, the required 7-qubit linear cluster state was prepared on the qubits 17, 18, 21, 23, 24, 25 and 22, single-qubit measurements were performed on the qubits 17, 18, 21, 23, 24 and 25, and quantum state tomography was performed on qubit 22. All nineteen shaded qubits, from qubit 17 through to qubit 6, were used to implement the interleaved sequence for the 3-qubit T gate with $m = 3$, which required a 19-qubit linear cluster state.

The required data was obtained over a period of 16 days on the *ibm_hanoi* quantum processor. On each day on which circuits were run, the relevant calibration information was recorded at 19:00 GMT. Table D.1 shows the average of the calibration information recorded during this period, with the uncertainties given by the standard deviation.




Figure D.1: Qubit topology of the *ibm_hanoi* quantum processor. The connecting lines between qubits indicate the qubit pairs for which the CX gate is supported at the hardware level. The qubits used in the implementations are shaded grey.

D.2. Fitting procedure

Qubit	T_1 (μ s)	T_2 (μ s)	\sqrt{X} Error	Readout Error	Qubit Pair	CX Error
17	135.10 \pm 25.32	75.47 \pm 04.72	0.000316 \pm 0.000045	0.0149 \pm 0.0039	17–18	0.00691 \pm 0.00093
18	174.96 \pm 33.66	173.48 \pm 38.68	0.000360 \pm 0.000055	0.0139 \pm 0.0053	18–21	0.00545 \pm 0.00100
21	137.16 \pm 20.95	24.48 \pm 01.93	0.000268 \pm 0.000023	0.0124 \pm 0.0027	21–23	0.01434 \pm 0.00174
23	152.68 \pm 33.73	45.98 \pm 02.76	0.000363 \pm 0.000052	0.0223 \pm 0.0079	23–24	0.02561 \pm 0.00231
24	153.91 \pm 27.54	31.14 \pm 04.12	0.000355 \pm 0.000048	0.0129 \pm 0.0010	24–25	0.01961 \pm 0.00434
25	178.63 \pm 30.70	70.14 \pm 13.07	0.000148 \pm 0.000040	0.0171 \pm 0.0028	25–22	0.00696 \pm 0.00170
22	162.43 \pm 40.46	102.68 \pm 24.47	0.000212 \pm 0.000128	0.0134 \pm 0.0051	22–19	0.00919 \pm 0.00163
19	153.41 \pm 40.11	149.65 \pm 56.73	0.000133 \pm 0.000016	0.0065 \pm 0.0007	19–16	0.00713 \pm 0.00270
16	114.35 \pm 45.09	143.93 \pm 59.36	0.000238 \pm 0.000117	0.0142 \pm 0.0012	16–14	0.01480 \pm 0.00534
14	148.31 \pm 34.02	27.22 \pm 00.00	0.000456 \pm 0.000522	0.0113 \pm 0.0053	14–11	0.01167 \pm 0.01659
11	158.01 \pm 40.07	197.11 \pm 47.48	0.000129 \pm 0.000027	0.0145 \pm 0.0012	11–8	0.00397 \pm 0.00088
8	170.04 \pm 38.89	260.68 \pm 69.31	0.000145 \pm 0.000065	0.0117 \pm 0.0020	8–5	0.01939 \pm 0.00614
5	136.51 \pm 24.61	148.65 \pm 39.80	0.000550 \pm 0.000221	0.0127 \pm 0.0053	5–3	0.00622 \pm 0.00199
3	164.39 \pm 38.07	247.30 \pm 72.86	0.000120 \pm 0.000016	0.0070 \pm 0.0007	3–2	0.00646 \pm 0.00088
2	149.60 \pm 32.88	236.10 \pm 42.03	0.000120 \pm 0.000025	0.0071 \pm 0.0033	2–1	0.00348 \pm 0.00040
1	165.78 \pm 31.71	149.88 \pm 43.46	0.000219 \pm 0.000071	0.0115 \pm 0.0059	1–4	0.00709 \pm 0.00129
4	133.92 \pm 42.34	15.79 \pm 00.00	0.000167 \pm 0.000018	0.0080 \pm 0.0011	4–7	0.01015 \pm 0.00119
7	175.04 \pm 36.21	216.79 \pm 46.00	0.000129 \pm 0.000014	0.0138 \pm 0.0020	7–6	0.00541 \pm 0.00090
6	148.74 \pm 30.86	184.66 \pm 67.86	0.000336 \pm 0.000223	0.0155 \pm 0.0071		

Table D.1: Calibration information for the *ibm_hanoi* quantum processor averaged over the time period during which circuits were run and data was obtained. The single-qubit calibration information for the relevant qubits is shown on the left. T_1 and T_2 are the amplitude and phase damping time constants respectively of the qubits. The CX error rates for relevant qubit pairs are shown on the right.

D.2 Fitting procedure

A Monte Carlo method, which takes uncertainties into account, was used to fit the reference and interleaved sequence fidelities to the exponential decay model given by Eqs. (5.3) and (5.4), respectively. For each sequence fidelity, one million points were sampled from a Gaussian distribution with mean equal to the sequence fidelity and standard deviation equal to the uncertainty in the sequence fidelity. For a given reference or interleaved sequence, sampled points obtained for the sequence fidelities for $m \in \{1, 2, 3\}$ were fit to Eq. (5.3) or Eq. (5.4), to obtain one million values for each fitting parameter. A given fitting parameter, inferred for a given reference or interleaved sequence, is the average of these one million values, and the uncertainty is given by the standard deviation.

Since each set of fitting parameters (p , A and B) was inferred from only three sequence fidelities, tight fitting constraints were imposed to ensure that the fitting parameters inferred from the data are physical. In particular, the constraint $B \in [0.48, 0.52]$ was imposed to ensure that the asymptote of the exponential decay curve is close to 0.5, as we expect it to be for depolarising noise. Furthermore, the constraint $A \in [0.4, 0.5]$ was imposed to ensure that the y -intercept of the exponential decay curve (which is given by $A + B$) is slightly less than one, as we expect it to be when state preparation and measurement errors are present, but do not dominate.

Finally, a similar Monte Carlo method was used to estimate the Haar-averaged gate fidelity of a given measurement-based gate from the appropriate reference depolarising noise parameter p_{ref} and interleaved depolarising noise parameter p_{int} . One million points were sampled from a Gaussian distribution with mean equal to p_{ref} and standard deviation equal to the uncertainty in p_{ref} , and similarly for p_{int} . One million values were then obtained for the estimated gate fidelity, from the sampled points for p_{ref} and p_{int} , using Eq. (5.2). The estimated gate fidelity for a given measurement-based gate is the average of these one million values, and the uncertainty is given by the standard deviation.




Figure D.2: Qubit topology of the *ibmq_brooklyn* quantum processor. The connecting lines between qubits indicate the qubit pairs for which the CX gate is supported at the hardware level. The qubits used in the implementations are shaded grey.

D.3 Process tomography

For a given measurement-based gate, quantum process tomography was performed on each of the three sets of qubits on which that gate was implemented in the interleaved sequences, that were implemented to obtain the data required to estimate the fidelity of the given measurement-based gate. To this end, we employed the quantum process tomography method proposed for single-qubit protocols by Nielsen and Chuang [173, 174]. This single-qubit process tomography method was reviewed extensively in Sec. 3.2.3. The method amounts to performing quantum state tomography (see Sec. 3.2.2) to infer the output states for four different probe input states, namely $|0\rangle$, $|1\rangle$, $|+\rangle$ and $|+_y\rangle$. The state tomography results are then used to construct a 4 by 4 process matrix χ , which completely characterises a given implementation of a single-qubit protocol.

To perform quantum process tomography for one of the three implementations of a measurement-based gate, on one of the three sets of qubits used in the interleaved sequences, the required linear cluster state was prepared on the set of qubits (with a given probe input state prepared on the first qubit), single-qubit measurements were performed on all but the final qubit, and quantum state tomography was performed on the final qubit to infer the output state, for each byproduct, for the given probe input state. Each of the twelve circuits needed for process tomography (three circuits for state tomography to infer the output state for each of the four probe input states) was run with 8192 shots on the relevant processor. For each byproduct, matrix multiplication was once again used to apply the appropriate Pauli correction to the density matrices, constructed by performing state tomography for each of the four probe input states. The Pauli corrected density matrices were used to construct a process matrix for each byproduct. An average process matrix for a given measurement-based implementation was then calculated from the process matrices obtained for the different byproducts.

The channel fidelity (see Eq. (3.4)) of one of the three implementations of a measurement-based gate G , on one of the three sets of qubits used in the interleaved sequences, was calculated using

$$F_G^C(\chi, \tilde{\chi}) = \text{Tr} \left(\sqrt{\sqrt{\chi} \tilde{\chi} \sqrt{\chi}} \right), \quad (\text{D.1})$$

Qubit	T_1 (μ s)	T_2 (μ s)	$\sqrt{\chi}$ Error	Readout Error
6	71.40±10.65	93.79±14.82	0.000592±0.000205	0.0389±0.0185
5	86.52±18.82	104.75±22.24	0.000278±0.000071	0.0220±0.0029
4	81.47±12.85	83.96±13.72	0.000370±0.000186	0.0219±0.0060
3	78.78±19.86	97.25±25.10	0.000344±0.000108	0.0218±0.0074
2	74.44±16.09	74.86±12.80	0.000561±0.000511	0.0285±0.0164
1	87.94±14.94	101.13±17.78	0.000698±0.000916	0.0384±0.0322
0	100.62±20.16	126.49±23.14	0.000331±0.000157	0.0216±0.0093
10	78.63±11.96	92.33±18.00	0.000437±0.000503	0.0256±0.0227
13	68.27±12.25	12.48±00.00	0.000439±0.000100	0.0186±0.0048
14	70.60±14.50	67.43±13.69	0.000725±0.000933	0.0231±0.0098
15	82.73±15.47	74.37±16.32	0.000526±0.000205	0.0305±0.0178
24	72.06±10.64	83.89±13.06	0.000408±0.000104	0.0173±0.0039
29	73.86±10.95	68.66±08.13	0.000434±0.000071	0.0208±0.0052
30	85.58±15.80	28.83±02.31	0.000344±0.000288	0.0219±0.0036
31	71.68±13.27	81.86±14.98	0.000339±0.000053	0.0164±0.0022
39	61.83±07.62	78.85±10.48	0.000378±0.000092	0.0231±0.0083
45	75.02±11.88	45.20±06.90	0.000352±0.000079	0.0237±0.0045
46	65.20±09.15	79.14±14.64	0.000448±0.000155	0.0211±0.0054
47	72.76±13.58	78.98±15.52	0.000366±0.000065	0.0232±0.0056
48	70.91±11.21	77.97±13.41	0.000503±0.000350	0.0319±0.0214
49	68.51±15.99	79.81±16.21	0.000477±0.000765	0.0274±0.0304
50	77.36±13.30	72.40±11.81	0.000329±0.000045	0.0145±0.0026
51	66.58±10.78	78.75±13.68	0.000447±0.000174	0.0260±0.0123
54	78.62±12.69	81.20±12.27	0.000302±0.000068	0.0168±0.0025
64	44.31±11.38	46.44±17.45	0.000765±0.000585	0.0451±0.0325
63	67.51±15.56	67.97±12.16	0.000419±0.000176	0.0306±0.0114
62	80.82±21.53	88.21±15.81	0.000517±0.000641	0.0358±0.0295
61	74.77±12.04	90.44±19.60	0.000387±0.000231	0.0266±0.0074
60	67.55±14.51	78.66±16.35	0.000471±0.000259	0.0228±0.0049
59	76.66±16.90	90.30±22.01	0.000600±0.000495	0.0280±0.0178
58	71.75±13.95	81.18±10.42	0.000900±0.000398	0.0524±0.0314

Qubit Pair	CX Error
6–5	0.00951±0.00206
5–4	0.00743±0.00243
4–3	0.00893±0.00609
3–2	0.01202±0.00332
2–1	0.01803±0.01495
1–0	0.01559±0.01385
0–10	0.00913±0.00448
10–13	0.01282±0.00480
13–14	0.01490±0.00824
14–15	0.01438±0.00651
15–24	0.01101±0.00295
24–29	0.01207±0.00398
29–30	0.01093±0.00331
30–31	0.00928±0.00338
31–39	0.00952±0.00183
39–45	0.00972±0.00222
45–46	0.01035±0.00234
46–47	0.01126±0.00148
47–48	0.01103±0.00537
48–49	0.01393±0.01103
49–50	0.01242±0.00679
50–51	0.00788±0.00112
51–54	0.01040±0.00134
54–64	0.01586±0.00855
64–63	0.01822±0.00534
63–62	0.01206±0.01430
62–61	0.01245±0.00374
61–60	0.01291±0.00595
60–59	0.01098±0.00438
59–58	0.01780±0.00942

Table D.2: Calibration information for the *ibmq_brooklyn* quantum processor averaged over the time period during which circuits were run and data was obtained. The single-qubit calibration information for the relevant qubits is shown on the left. T_1 and T_2 are the amplitude and phase damping time constants respectively of the qubits. The CX error rates for relevant qubit pairs are shown on the right.

where χ is the process matrix for the ideal implementation of G and $\tilde{\chi}$ is the average process matrix obtained by performing quantum process tomography for the given implementation of G . The Haar-averaged gate fidelity of each of the three implementations of G was then calculated using

$$F_G = \frac{d(F_G^C(\chi, \tilde{\chi}))^2 + 1}{d + 1}, \quad (\text{D.2})$$

where $d = 2$ for single qubits [92]. Based on the channel fidelities calculated from process tomography results in Chapter 3, we expect the uncertainty in the fidelity of a given implementation of G to be small compared to the range of fidelities for the three different implementations of G on the three different sets of qubits. We therefore determined only a single fidelity for each implementation of G , with no estimated uncertainty, to avoid unnecessary use of processor time.

D.4 Qubits used for fidelity estimation of noisier gates

To perform fidelity estimation for the noisier measurement-based implementations, we implemented reference sequences with lengths $m \in \{1, 2, 3\}$ and interleaved sequences with $m \in \{1, 2, 3\}$, for the 4-qubit and 6-qubit Hadamard gate as well as the 5-qubit and 7-qubit T gate, on the *ibmq_brooklyn* quantum processor. The qubit topology of the *ibmq_brooklyn* quantum processor is shown in Fig. D.2, with the qubits used in the implementations shaded grey. These qubits were once again chosen for their low error rates compared to other qubits. For each m (reference or interleaved), the required linear cluster state was prepared along the length of shaded qubits, starting at qubit 6. For example, to implement the interleaved sequence for the 5-qubit T gate with $m = 1$, the required 9-qubit linear cluster state was prepared on the qubits 6, 5, 4, 3, 2, 1, 0, 10 and 13, single-qubit measurements were performed on the qubits 6, 5, 4, 3, 2, 1, 0 and 10, and quantum state tomography was performed on qubit 13. All thirty-one shaded qubits, from qubit 6 through to qubit 58, were used to implement the interleaved sequence for the 7-qubit T gate with $m = 3$, which required a 31-qubit linear cluster state.

The required data was obtained over a period of 86 days on the *ibmq_brooklyn* quantum processor. On each day on which circuits were run, the relevant calibration information was recorded at 19:00 GMT. Table D.2 shows the average of the calibration information recorded during this period, with the uncertainties given by the standard deviation.

Appendix E

Quantum random number generation using an on-chip nanowire plasmonic waveguide

E.1 Power transmission factor of nanowire plasmonic waveguide

To determine the power transmission factor of the nanowire plasmonic waveguide, we note that we can write the net power transmission factor from the rear aperture of the DLM objective to the SPAD detector (see Fig. 6.1a) as $\eta = \eta_{\text{DLM}}\eta_{\text{wgd}}\eta_{\text{col}}$, where η_{DLM} is the input power transmission factor of the DLM objective, η_{wgd} is the power transmission factor of the nanowire plasmonic waveguide, η_{col} is the power transmission factor of the collection optics (which includes the output power transmission factor of the DLM objective, the knife-edge mirror, the fibre coupler (FC) and the multi-mode optical fibre (MM)). By calculating the ratio of the power measured at the focal point and rear aperture of the DLM objective, we found that $\eta_{\text{DLM}} = 0.94$. To determine η_{col} , we used the MM fibre to connect the FC to the continuous-wave laser. By calculating the ratio of the power measured at the focal point of the DLM objective and the output power of the laser, we found that $\eta_{\text{col}} = 0.30$. This is under the assumption that loss in the collection optics is symmetric. Finally, we determined η by calculating the ratio of the power at the SPAD detector (P_{out}) and the rear aperture of the DLM objective (P_{in}). We measured $P_{\text{in}} = 0.24 \mu\text{W}$ and calculated $P_{\text{out}} = \frac{Rhc}{\lambda} = 4.6 \text{ pW}$ for a photon detection rate of $R = 1.8 \text{ Mcounts/s}$. Hence $\eta = 1.9 \times 10^{-5}$. Combining all of the above results, it follows that $\eta_{\text{wgd}} = 6.7 \times 10^{-5}$.

E.2 Polarisation dependence of photon detection rate

We investigate the dependence of the photon detection rate on the polarisation of the input beam. To this end, we use the second HWP in our experimental setup (see Fig. 6.1a) to adjust the polarisation of the input beam. A plot of photon detection rate versus waveplate angle is shown in Fig. E.1. As expected, the photon detection rate has a sinusoidal dependence on the waveplate angle [221, 222]. This confirms that the collection optics is indeed capturing out-coupling photons from the output grating of the nanowire plasmonic waveguide and not scattered photons from the input beam.




Figure E.1: Photon detection rate versus waveplate angle. Data points are an average of five repetitions and the errors are given by the standard deviation.

E.3 Proof of uniformity for time-of-arrival scheme

We note that for laser light, the probability of k photons arriving in a time interval of length t is given by $P_t(k) = e^{-\lambda t} \frac{(\lambda t)^k}{k!}$ [228], where $\bar{k}_t = \lambda t$ is the mean number of photons in a time interval of length t and λ is the time-independent mean photon flux.

Consider a time interval of length T divided into two equal sections of duration $\tau_1 = \tau_2 = \tau = \frac{T}{2}$. Let y represent the case that one photon arrives in the time interval of length T (with probability $p(y)$) and let x represent the case that one photon arrives in the first section $[0, \tau_1]$ and not in the second section (with probability $p(x)$). The probability that one photon arrives in the first section given that one photon arrives in the time interval of length T is

$$p_1 = p(x | y) = \frac{p(x \wedge y)}{p(y)} = \frac{p(y | x)p(x)}{p(y)}. \quad (\text{E.1})$$

We then have that $p(y | x) = 1$, $p(x) = P_{\tau_1}(1)P_{\tau_2}(0)$ and $p(y) = P_T(1)$, which gives

$$p_1 = \frac{P_{\tau_1}(1)P_{\tau_2}(0)}{P_T(1)} = \frac{e^{-\lambda\tau}\lambda\tau e^{-\lambda\tau}}{e^{-\lambda T}\lambda T} = \frac{\tau}{T} = \frac{1}{2}, \quad (\text{E.2})$$

where we have used $T = 2\tau$. Similarly we have the probability that one photon arrives in the second section $[\tau_1, T]$ and not in the first section as

$$p_2 = \frac{P_{\tau_2}(1)P_{\tau_1}(0)}{P_T(1)} = \frac{1}{2}. \quad (\text{E.3})$$

These results extend naturally to the more general case where the time interval of length T is divided into N equal sections of duration $\tau_i = \tau = \frac{T}{N}$ for $i = 1, \dots, N$. In particular,

$$p_i = \frac{P_{\tau_i}(1)P_{(i-1)\tau}(0)P_{(N-i)\tau}(0)}{P_T(1)} = \frac{e^{-\lambda\tau}\lambda\tau e^{-\lambda(N-1)\tau}}{e^{-\lambda T}\lambda T} = \frac{\tau}{T} = \frac{1}{N}. \quad (\text{E.4})$$

The above can be generalised to k photons in a time interval of length T with N sections and we have the

E.4. Higher order photon events and correction factor

probability that at least one photon arrives in section i given that no photons occurred before that section as

$$p_i = \frac{\left(\sum_{m=2}^k \frac{P_\tau(m)}{P_\tau(1)} P_{(N-i)\tau}(k-m) + P_{(N-i)\tau}(k-1) \right) P_\tau(1)}{P_T(k)} P_{(i-1)\tau}(0)$$

$$= \frac{1}{P_T(k)} \sum_{m=1}^k \frac{P_\tau(m)}{P_\tau(1)} P_{(N-i)\tau}(k-m) P_\tau(1) P_{(i-1)\tau}(0).$$

Substituting in the $P_k(j)$ and using $0^0 = 1$ as convention, one finds

$$p_i = \sum_{m=1}^k \frac{c(k, m)(N-i)^{k-m}}{N^k}$$

$$= \left(1 - \frac{i-1}{N}\right)^k - \left(1 - \frac{i}{N}\right)^k,$$

which is the form given in Ref. [53].

E.4 Higher order photon events and correction factor

For laser light, the probability of k photons arriving in a time interval of length t is given by $P_t(k) = e^{-\lambda t} \frac{(\lambda t)^k}{k!}$ [228], where $\bar{k}_t = \lambda t$ is the mean number of photons in a time interval of length t and λ is the time-independent mean photon flux.

For a photon detection rate R , there are R dead times within one second, each of which is τ_d in duration, where τ_d is the detector dead time. The total number of undetected photons in one second is therefore $R\lambda\tau_d$, and the total number of detected and undetected photons is $R(1 + \lambda\tau_d)$. Hence we have that the mean number of photons for one second is

$$\bar{k}_1 = \frac{\bar{k}_T}{T} = R(1 + \lambda\tau_d), \quad (\text{E.5})$$

which gives $\bar{k}_T = RT + R\lambda T\tau_d$. Substituting in $\lambda = \frac{\bar{k}_T}{T}$ on the right hand side of the previous equation and rearranging gives

$$\bar{k}_T = R \frac{T}{(1 - R\tau_d)}. \quad (\text{E.6})$$

Using $R = 5.2 \text{ Mcounts/s}$, $T = 12.8 \text{ ns}$ and $\tau_d = 24 \text{ ns}$ we obtain $\bar{k}_T = 0.076$. We then have that $\lambda = \frac{\bar{k}_T}{T}$, and using this we find that the relative probability of a single photon in a time interval of length T compared to the case of higher order photon number is $p_1 = \frac{P_T(1)}{1 - P_T(0)} = 0.96$ and for two photons $p_2 = \frac{P_T(2)}{1 - P_T(0)} = 0.037$. Thus higher order photon number within a time interval of length T is negligible.

However, undetected photons may introduce some correlations in the random numbers generated. A correction factor takes into account the non-negligible detector dead time and that any counts a detector measures is an underestimate of the true counts. The correction factor is the ratio of total photons to photons detected and is given by

$$c_F = \frac{R(1 + \lambda\tau_d)}{R} = 1 + \lambda\tau_d. \quad (\text{E.7})$$

Substituting in

$$\lambda = \frac{\bar{k}_T}{T} = \frac{RT}{(1 - R\tau_d)} \frac{1}{T} = \frac{R}{(1 - R\tau_d)} \quad (\text{E.8})$$

we get

$$c_F = 1 + \frac{R\tau_d}{(1 - R\tau_d)} = (1 - R\tau_d)^{-1}. \quad (\text{E.9})$$

For $R = 5.2 \text{ Mcounts/s}$ and $\tau_d = 24 \text{ ns}$, we have $c_F = 1.143$. This means that one in every $(1.143 - 1)^{-1} = 7$ photons arriving at the detector are not detected.

Bibliography

- [1] Ladd, T. D., Jelezko, F., Laflamme, R., Nakamura, Y., Monroe, C. & O'Brien, J. L. Quantum computers. *Nature* **464**, 45–53 (2010).
- [2] Gisin, N. & Thew, R. Quantum communication. *Nat. Photonics* **1**, 165–171 (2007).
- [3] Dowling, J. P. & Seshadreesan, K. P. Quantum optical technologies for metrology, sensing, and imaging. *J. Light. Technol.* **33**, 2359–2370 (2015).
- [4] Shor, P. W. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Sci. Comput.* **26**, 1484–1509 (1997).
- [5] Grover, L. K. Quantum mechanics helps in searching for a needle in a haystack. *Phys. Rev. Lett.* **79**, 325–328 (1997).
- [6] Feynman, R. P. Simulating physics with computers. *Int. J. Theor. Phys.* **21**, 467–488 (1982).
- [7] Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N. & Lloyd, S. Quantum machine learning. *Nature* **549**, 195–202 (2017).
- [8] Hauke, P., Katzgraber, H. G., Lechner, W., Nishimori, H. & Oliver, W. D. Perspectives of quantum annealing: methods and implementations. *Rep. Prog. Phys.* **83**, 054401 (2020).
- [9] Chauhan, V., Negi, S., Jain, D., Singh, P., Sagar, A. K. & Sharma, A. K. Quantum computers: a review on how quantum computing can boom AI in 2022 *2nd International Conference on Advance Computing and Innovative Technologies in Engineering (ICACITE)* 559–563 (IEEE, 2022).
- [10] Cao, Y., Romero, J. & Aspuru-Guzik, A. Potential of quantum computing for drug discovery. *IBM J. Res. Dev.* **62**, 6:1–6:20 (2018).
- [11] Egger, D. J., Gambella, C., Marecek, J., McFaddin, S., Mevissen, M., Raymond, R., Simonetto, A., Woerner, S. & Yndurain, E. Quantum computing for finance: state-of-the-art and future prospects. *IEEE Trans. Quantum Eng.* **1**, 3101724 (2020).
- [12] Berger, C., Di Paolo, A., Forrest, T., Hadfield, S., Sawaya, N., Stęchły, M. & Thibault, K. Quantum technologies for climate change: preliminary assessment. *arXiv:2107.05362* (2021).
- [13] Arute, F. *et al.* Quantum supremacy using a programmable superconducting processor. *Nature* **574**, 505–510 (2019).
- [14] Wu, Y. *et al.* Strong quantum computational advantage using a superconducting quantum processor. *Phys. Rev. Lett.* **127**, 180501 (2021).

- [15] Zhong, H. S. *et al.* Quantum computational advantage using photons. *Science* **370**, 1460–1463 (2020).
- [16] Zhong, H. S. *et al.* Phase-programmable Gaussian boson sampling using stimulated squeezed light. *Phys. Rev. Lett.* **127**, 180502 (2021).
- [17] Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018).
- [18] Shor, P. W. Scheme for reducing decoherence in quantum computer memory. *Phys. Rev. A* **52**, R2493–R2496 (1995).
- [19] Steane, A. M. Error correcting codes in quantum theory. *Phys. Rev. Lett.* **77**, 793–797 (1996).
- [20] Laflamme, R., Miquel, C., Paz, J. P. & Zurek, W. H. Perfect quantum error correcting code. *Phys. Rev. Lett.* **77**, 198–201 (1996).
- [21] Devitt, S. J., Munro, W. J. & Nemoto, K. Quantum error correction for beginners. *Rep. Prog. Phys.* **76**, 076001 (2013).
- [22] Cai, Z., Babbush, R., Benjamin, S. C., Endo, S., Huggins, W. J., Li, Y., McClean, J. R. & O’Brien, T. E. Quantum error mitigation. arXiv:2210.00921 (2022).
- [23] Bharti, K. *et al.* Noisy intermediate-scale quantum algorithms. *Rev. Mod. Phys.* **94**, 015004 (2022).
- [24] Deutsch, D. E. Quantum computational networks. *Proc. Math. Phys. Eng. Sci.* **425**, 73–90 (1989).
- [25] Barenco, A., Bennett, C. H., Cleve, R., DiVincenzo, D. P., Margolus, N., Shor, P., Sleator, T., Smolin, J. & Weinfurter, H. Elementary gates for quantum computation. *Phys. Rev. A* **52**, 3457–3467 (1995).
- [26] Raussendorf, R. & Briegel, H. J. A one-way quantum computer. *Phys. Rev. Lett.* **86**, 5188–5191 (2001).
- [27] Raussendorf, R. & Briegel, H. J. Computational model underlying the one-way quantum computer. *Quantum Inf. Comput.* **2**, 443–486 (2002).
- [28] Raussendorf, R., Browne, D. E. & Briegel, H. J. Measurement-based quantum computation on cluster states. *Phys. Rev. A* **68**, 022312 (2003).
- [29] Briegel, H. J., Browne, D. E., Dür, W., Raussendorf, R. & Van der Nest, M. Measurement-based quantum computation. *Nat. Phys.* **5**, 19–26 (2009).
- [30] Nielsen, M. A. Optical quantum computation using cluster states. *Phys. Rev. Lett.* **93**, 040503 (2004).
- [31] Browne, D. E. & Rudolph, T. Resource-efficient linear optical quantum computation. *Phys. Rev. Lett.* **95**, 010501 (2005).
- [32] Walther, P., Resch, K. J., Rudolph, T., Schenck, E., Weinfurter, H., Vedral, V., Aspelmeyer, M. & Zeilinger, A. Experimental one-way quantum computing. *Nature* **434**, 169–176 (2005).
- [33] Tanamoto, T., Liu, Y. X., Fujita, S., Hu, X. & Nori, F. Producing cluster states in charge qubits and flux qubits. *Phys. Rev. Lett.* **97**, 230501 (2006).
- [34] Vaucher, B., Nunnenkamp, A. & Jaksch, D. Creation of resilient entangled states and a resource for measurement-based quantum computation with optical superlattices. *New J. Phys.* **10**, 023005 (2008).

- [35] Weinstein, Y. S., Hellberg, C. S. & Levy, J. Quantum-dot cluster-state computing with encoded qubits. *Phys. Rev. A* **72**, 020304(R) (2005).
- [36] Shannon, C. E. Communication theory of secrecy systems. *Bell Syst. Tech. J.* **28**, 656–715 (1949).
- [37] Metropolis, N. & Ulam, S. The Monte Carlo method. *J. Am. Stat. Assoc.* **44**, 335–341 (1949).
- [38] Bell, J. S. On the Einstein-Podolsky-Rosen paradox. *Physics* **1**, 195–200 (1964).
- [39] Ma, X., Yuan, X., Cao, Z., Qi, B. & Zhang, Z. Quantum random number generation. *npj Quantum Inf.* **2**, 16021 (2016).
- [40] Herrero-Collantes, M. & Garcia-Escartin, J. C. Quantum random number generators. *Rev. Mod. Phys.* **89**, 015004 (2017).
- [41] Mannalath, V., Mishra, S. & Pathak, A. A comprehensive review of quantum random number generators: concepts, classification and the origin of randomness. arXiv:2203.00261 (2022).
- [42] Isida, M. & Ikeda, H. Random number generator. *Ann. Inst. Stat. Math.* **8**, 119–126 (1956).
- [43] Vincent, C. H. The generation of truly random binary numbers. *J. Phys. E: Sci. Instrum.* **3**, 594–598 (1970).
- [44] Schmidt, H. Quantum-mechanical random-number generator. *J. Appl. Phys.* **41**, 462–468 (1970).
- [45] Stefanov, A., Gisin, N., Guinnard, O., Guinnard, L. & Zbinden, H. Optical quantum random number generator. *J. Mod. Opt.* **47**, 595–598 (2000).
- [46] Jennewein, T., Achleitner, U., Weihs, G., Weinfurter, H. & Zeilinger, A. A fast and compact quantum random number generator. *Rev. Sci. Instrum.* **71**, 1675–1680 (2000).
- [47] Shafi, K. M., Chawla, P., Hegde, A. S., Gayatri, R. S., Padhye, A. & Chandrashekar, C. M. Multi-bit quantum random number generator from path-entangled single photons. arXiv:2202.10933 (2022).
- [48] Ma, H. Q., Xie, Y. & Wu, L. A. Random number generation based on the time of arrival of single photons. *Appl. Opt.* **44**, 7760–7763 (2005).
- [49] Stipčević, M. & Rogina, B. M. Quantum random number generator based on photonic emission in semi-conductors. *Rev. Sci. Instrum.* **78**, 045104 (2007).
- [50] Dynes, J. F., Yuan, Z. L., Sharpe, A. W. & Shields, A. J. A high speed, postprocessing free, quantum random number generator. *Appl. Phys. Lett.* **93**, 031109 (2008).
- [51] Wayne, M. A., Jeffrey, E. R., Akselrod, G. M. & Kwiat, P. G. Photon arrival time quantum random number generation. *J. Mod. Opt.* **56**, 516–522 (2009).
- [52] Wahl, M., Leifgen, M., Berlin, M., Röhlicke, T., Rahn, H. J. & Benson, O. An ultrafast quantum random number generator with provably bounded output bias based on photon arrival time measurements. *Appl. Phys. Lett.* **98**, 171105 (2011).
- [53] Nie, Y. Q., Zhang, H. F., Zhang, Z., Wang, J., Ma, X., Zhang, J. & Pan, J. W. Practical and fast quantum random number generation based on photon arrival time relative to external reference. *Appl. Phys. Lett.* **104**, 051110 (2014).

- [54] Yan, Q., Zhao, B., Hua, Z., Liao, Q. & Yang, H. High-speed quantum random number generation by continuous measurement of arrival time of photons. *Rev. Sci. Instrum.* **86**, 073113 (2015).
- [55] Banerjee, A., Aggarwal, D., Sharma, A. & Yadav, G. Unpredictable and uniform RNG based on time of arrival using InGaAs detectors. arXiv:2010.12898 (2020).
- [56] Banerjee, A., Sethia, A., Mogiligidha, V., Krishnan, R. K., AR, M., Rajamani, S. & Shenoy, V. Experimental demonstration of high-entropy time of arrival based optical QRNG qualifying stringent statistical tests. arXiv:2108.06112 (2021).
- [57] Fürst, H., Weier, H., Nauerth, S., Marangon, D. G., Kurtsiefer, C. & Weinfurter, H. High speed optical quantum random number generation. *Opt. Express* **18**, 13029–13037 (2010).
- [58] Ren, M., Wu, E., Liang, Y., Jian, Y., Wu, G. & Zeng, H. Quantum random number generator based on a photon-number-resolving detector. *Phys. Rev. A* **83**, 023820 (2011).
- [59] Applegate, M. J., Thomas, O., Dynes, J. F., Yuan, Z. L., Ritchie, D. A. & Shields, A. J. Efficient and robust quantum random number generation by photon number detection. *Appl. Phys. Lett.* **107**, 071106 (2015).
- [60] Gabriel, C., Wittmann, C., Sych, D., Dong, R., Maurer, W., Andersen, U. L., Marquardt, C. & Leuchs, G. A generator for unique quantum random numbers based on vacuum states. *Nat. Photonics* **4**, 711–715 (2010).
- [61] Shen, Y., Tian, L. & Zou, H. Practical quantum random number generator based on measuring the shot noise of vacuum states. *Phys. Rev. A* **81**, 063814 (2010).
- [62] Symul, T., Assad, S. M. & Lam, P. K. Real time demonstration of high bitrate quantum random number generation with coherent laser light. *Appl. Phys. Lett.* **98**, 231103 (2011).
- [63] Shi, Y., Chng, B. & Kurtsiefer, C. Random numbers from vacuum fluctuations. *Appl. Phys. Lett.* **109**, 041101 (2016).
- [64] Zheng, Z., Zhang, Y., Huang, W., Yu, S. & Guo, H. 6 Gbps real-time optical quantum random number generator based on vacuum fluctuation. *Rev. Sci. Instrum.* **90**, 043105 (2019).
- [65] Haylock, B., Peace, D., Lenzini, F., Weedbrook, C. & Lobino, M. Multiplexed quantum random number generation. *Quantum* **3**, 141 (2019).
- [66] Williams, C. R. S., Salevan, J. C., Li, X., Roy, R. & Murphy, T. E. Fast physical random number generator using amplified spontaneous emission. *Opt. Express* **18**, 23584–23597 (2010).
- [67] Qi, B., Chi, Y. M., Lo, H. K. & Qian, L. High-speed quantum random number generation by measuring phase noise of a single-mode laser. *Opt. Lett.* **35**, 312–314 (2010).
- [68] Xu, F., Qi, B., Ma, X., Xu, H., Zheng, H. & Lo, H. K. Ultrafast quantum random number generation based on quantum phase fluctuations. *Opt. Express* **20**, 12366–12377 (2012).
- [69] Nie, Y. Q., Huang, L., Liu, Y., Payne, F., Zhang, J. & Pan, J. W. The generation of 68 Gbps quantum random number by measuring laser phase fluctuations. *Rev. Sci. Instrum.* **86**, 063105 (2015).

- [70] Zhang, X. G., Nie, Y. Q., Zhou, H., Liang, H., Ma, X., Zhang, J. & Pan, J. W. Note: fully integrated 3.2 Gbps quantum random number generator with real-time extraction. *Rev. Sci. Instrum.* **87**, 076102 (2016).
- [71] Huang, M., Chen, Z., Zhang, Y. & Guo, H. A phase fluctuation based practical quantum random number generator scheme with delay-free structure. *Appl. Sci.* **10**, 2431 (2020).
- [72] Sanguinetti, B., Martin, A., Zbinden, H. & Gisin, N. Quantum random number generation on a mobile phone. *Phys. Rev. X* **4**, 031056 (2014).
- [73] Khanmohammadi, A., Enne, R., Hofbauer, M. & Zimmermann, H. A monolithic silicon quantum random number generator based on measurement of photon detection time. *IEEE Photon. J.* **7**, 1–13 (2015).
- [74] Tisa, S., Villa, F., Giudice, A., Simmerle, G. & Zappa, F. High-speed quantum random number generation using CMOS photon counting detectors. *IEEE J. Sel. Top. Quantum Electron.* **21**, 23–29 (2015).
- [75] Abellan, C., Amaya, W., Domenech, D., Muñoz, P., Capmany, J., Longhi, S., Mitchell, M. W. & Pruneri, V. Quantum entropy source on an InP photonic integrated circuit for random number generation. *Optica* **3**, 989–994 (2016).
- [76] Raffaelli, F., Sibson, P., Kennard, J. E., Mahler, D. H., Thompson, M. G. & Matthews, J. C. F. Generation of random numbers by measuring phase fluctuations from a laser diode with a silicon-on-insulator chip. *Opt. Express* **26**, 19730–19741 (2018).
- [77] Raffaelli, F., Ferranti, G., Mahler, D. H., Sibson, P., Kennard, J. E., Santamato, A., Sinclair, G., Bonneau, D., Thompson, M. G. & Matthews, J. C. F. A homodyne detector integrated onto a photonic chip for measuring quantum states and generating random numbers. *Quantum Sci. Technol.* **3**, 025003 (2018).
- [78] Huang, L. & Zhou, H. Integrated Gbps quantum random number generator with real-time extraction based on homodyne detection. *J. Opt. Soc. Am. B* **36**, B130–B136 (2019).
- [79] Roger, T., Paraiso, T., De Marco, I., Marangon, D. G., Yuan, Z. & Shields, A. J. Real-time interferometric quantum random number generation on chip. *J. Opt. Soc. Am. B* **36**, B137–B142 (2019).
- [80] Bai, B., Huang, J., Qiao, G. R., Nie, Y. Q., Tang, W., Chu, T., Zhang, J. & Pan, J. W. 18.8 Gbps real-time quantum random number generator with a photonic integrated chip. *Appl. Phys. Lett.* **118**, 264001 (2021).
- [81] Tamura, K. & Shikano, Y. Quantum random numbers generated by a cloud superconducting quantum computer in *International Symposium on Mathematics, Quantum Theory, and Cryptography* 17–37 (Springer, 2021).
- [82] Strydom, C. & Tame, M. S. Random number generation using IBM quantum processors in *The Proceedings of SAIP2021, the 65th Annual Conference of the South African Institute of Physics* 630–635 (SAIP, 2021).
- [83] Li, Y., Fei, Y., Wang, W., Meng, X., Wang, H., Duan, Q. & Ma, Z. Quantum random number generator using a cloud superconducting quantum computer based on source-independent protocol. *Sci. Rep.* **11**, 23873 (2021).

- [84] Um, M., Zhang, X., Zhang, J., Wang, Y., Yangchao, S., Deng, D. L., Duan, L. M. & Kim, K. Experimental certification of random numbers via quantum contextuality. *Sci. Rep.* **3**, 1627 (2013).
- [85] Choi, W. H., Lv, Y., Kim, J., Deshpande, A., Kang, G., Wang, J. P. & Kim, C. H. A magnetic tunnel junction based true random number generator with conditional perturb and real-time output probability tracking in *2014 IEEE International Electron Devices Meeting* 12.5.1–12.5.4 (IEEE, 2014).
- [86] Ng, H. J., Yang, S., Yao, Z., Yang, H. & Lim, C. C. W. Provably-secure randomness generation from switching probability of magnetic tunnel junctions. arXiv:2206.06636 (2022).
- [87] Knill, E. Approximation by quantum circuits. arXiv:quant-ph/9508006 (1995).
- [88] Turner, P. S. & Markham, D. Derandomising quantum circuits with measurement-based unitary designs. *Phys. Rev. Lett.* **116**, 200501 (2016).
- [89] Mezher, R., Ghalbouni, J., Dgheim, J. & Markham, D. Efficient quantum pseudorandomness with simple graph states. *Phys. Rev. A* **97**, 022333 (2018).
- [90] Lancien, C. & Majenz, C. Weak approximate unitary designs and applications to quantum encryption. *Quantum* **4**, 313 (2020).
- [91] Hayden, P., Leung, D., Shor, P. W. & Winter, A. Randomising quantum states: constructions and applications. *Commun. Math. Phys.* **250**, 371–391 (2004).
- [92] Dankert, C., Cleve, R., Emerson, J. & Livine, E. Exact and approximate unitary 2-designs and their application to fidelity estimation. *Phys. Rev. A* **80**, 012304 (2009).
- [93] Magesan, E., Gambetta, J. M. & Emerson, J. Scalable and robust randomized benchmarking of quantum processes. *Phys. Rev. Lett.* **106**, 180504 (2011).
- [94] Magesan, E., Gambetta, J. M. & Emerson, J. Characterizing quantum gates via randomized benchmarking. *Phys. Rev. A* **85**, 042311 (2012).
- [95] Magesan, E. *et al.* Efficient measurement of quantum gate error by interleaved randomized benchmarking. *Phys. Rev. Lett.* **109**, 080505 (2012).
- [96] Harper, R. & Flammia, S. T. Estimating the fidelity of T gates using standard interleaved randomized benchmarking. *Quantum Sci. Technol.* **2**, 015008 (2017).
- [97] Alexander, R. N., Turner, P. S. & Bartlett, S. D. Randomized benchmarking in measurement-based quantum computing. *Phys. Rev. A* **94**, 032303 (2016).
- [98] Wallman, J. J. Randomized benchmarking with gate-dependent noise. *Quantum* **2**, 47 (2018).
- [99] Merkel, S. T., Pritchett, E. J. & Fong, B. H. Randomized benchmarking as convolution: Fourier analysis of gate dependent errors. *Quantum* **5**, 581 (2021).
- [100] Helsen, J., Xue, X., Vandersypen, L. M. K. & Wehner, S. A new class of efficient randomized benchmarking protocols. *npj Quantum Inf.* **5**, 71 (2019).
- [101] Liu, Z., Zeng, P., Zhou, Y. & Gu, M. Characterizing correlation within multipartite quantum systems via local randomized measurements. *Phys. Rev. A* **105**, 022407 (2022).

- [102] Hayden, P. & Preskill, J. Black holes as mirrors: quantum information in random subsystems. *J. High Energy Phys.* **9**, 120 (2007).
- [103] Huang, H. Y., Kueng, R. & Preskill, J. Predicting many properties of a quantum system from very few measurements. *Nat. Phys.* **16**, 1050–1057 (2020).
- [104] Elben, A. *et al.* Mixed-state entanglement from local randomized measurements. *Phys. Rev. Lett.* **125**, 200501 (2020).
- [105] Zhou, Y., Zeng, P. & Liu, Z. Single-copies estimation of entanglement negativity. *Phys. Rev. Lett.* **125**, 200502 (2020).
- [106] Liu, Z., Tang, Y., Dai, H., Liu, P., Chen, S. & Ma, X. Detecting entanglement in quantum many-body systems via permutation moments. *Phys. Rev. Lett.* **129**, 260501 (2022).
- [107] Brandão, F. G. S. L. & Horodecki, M. Exponential quantum speed-ups are generic. *Quantum Inf. Comput.* **13**, 0901 (2013).
- [108] Ambainis, A. & Emerson, J. Quantum t -designs: t -wise independence in the quantum world in *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC'07)* 129–140 (IEEE, 2007).
- [109] Nakata, Y. *et al.* Quantum circuits for exact unitary t -designs and applications to higher-order randomized benchmarking. *PRX Quantum* **2**, 030339 (2021).
- [110] Zhang, L., Zhu, C. & Pei, C. Average fidelity estimation of twirled noisy quantum channel using unitary $2t$ -design. *Chin. Phys. B* **28**, 010304 (2019).
- [111] Harrow, A. W. & Low, R. A. Random quantum circuits are approximate 2-designs. *Commun. Math. Phys.* **291**, 257–302 (2009).
- [112] Diniz, I. T. & Jonathan, D. Comment on “Random quantum circuits are approximate 2-designs”. *Commun. Math. Phys.* **304**, 281–293 (2011).
- [113] Brandão, F. G. S. L., Harrow, A. W. & Horodecki, M. Local random quantum circuits are approximate polynomial-designs. *Commun. Math. Phys.* **346**, 397–434 (2016).
- [114] Nakata, Y., Hirche, C., Koashi, M. & Winter, A. Efficient unitary designs with nearly time-independent Hamiltonian dynamics. *Phys. Rev. X* **7**, 021006 (2017).
- [115] Cleve, R., Leung, D., Liu, L. & Wang, C. Near-linear constructions of exact unitary 2-designs. *Quantum Inf. Comput.* **16**, 0721 (2016).
- [116] IBM Quantum Experience, <https://quantum-computing.ibm.com/>. Accessed 22 February 2022.
- [117] Maciejewski, F. B., Zimborás, Z. & Oszmaniec, M. Mitigation of readout noise in near-term quantum devices by classical post-processing based on detector tomography. *Quantum* **4**, 257 (2020).
- [118] Elben, A., Flammia, S. T., Huang, H. Y., Kueng, R., Preskill, J., Vermersch, B. & Zoller, P. The randomized measurement toolbox. *Nat. Rev. Phys.* **5**, 9–24 (2023).
- [119] Gaebler, J. P. *et al.* Randomized benchmarking of multiqubit gates. *Phys. Rev. Lett.* **108**, 260503 (2012).

- [120] Carignan-Dugas, A., Wallman, J. J. & Emerson, J. Bounding the average gate fidelity of composite channels using the unitarity. *New J. Phys.* **21**, 053016 (2019).
- [121] Carignan-Dugas, A., Wallman, J. J. & Emerson, J. Characterizing universal gate sets via dihedral benchmarking. *Phys. Rev. A* **92**, 060302(R) (2015).
- [122] Nielsen, M. A. & Chuang, I. L. Universal quantum gates in *Quantum Computation and Quantum Information: 10th Anniversary Edition* 188–202 (Cambridge University Press, 2010).
- [123] Francis, J. T., Zhang, X., Özdemir, Ş. K. & Tame, M. S. Quantum random number generation using an on-chip plasmonic beamsplitter. *Quantum Sci. Technol.* **2**, 035004 (2017).
- [124] Tame, M. S., McEnery, K. R., Özdemir, Ş. K., Lee, J., Maier, S. A. & Kim, M. S. Quantum plasmonics. *Nat. Phys.* **9**, 329–340 (2013).
- [125] Xu, D., Xiong, X., Wu, L., Ren, X. F., Png, C. E., Guo, G. C., Gong, Q. & Xiao, Y. F. Quantum plasmonics: new opportunity in fundamental and applied photonics. *Adv. Opt. Photonics* **10**, 703–756 (2018).
- [126] Politi, A., Matthews, J. C. F., Thompson, M. G. & O’Brien, J. L. Integrated quantum photonics. *IEEE J. Sel. Top. Quantum Electron.* **15**, 1673–1684 (2009).
- [127] Schröder, T., Mouradian, S. L., Zheng, J., Trusheim, M. E., Walsh, M., Chen, E. H., Li, L., Bayn, I. & Englund, D. Quantum nanophotonics in diamond. *J. Opt. Soc. Am. B* **33**, B65–B83 (2016).
- [128] Lee, C., Lawrie, B., Pooser, R., Lee, K. G., Rockstuhl, C. & Tame, M. S. Quantum plasmonic sensors. *Chem. Rev.* **121**, 4743–4804 (2021).
- [129] Fan, W., Lawrie, B. J. & Pooser, R. C. Quantum plasmonic sensing. *Phys. Rev. A* **92**, 053812 (2015).
- [130] Pooser, R. C. & Lawrie, B. Plasmonic trace sensing below the photon shot noise limit. *ACS Photonics* **3**, 8–13 (2015).
- [131] Lee, C., Dieleman, F., Lee, J., Rockstuhl, C., Maier, S. A. & Tame, M. S. Quantum plasmonic sensing: beyond the shot-noise and diffraction limit. *ACS Photonics* **3**, 992–999 (2016).
- [132] Lee, J. S., Huynh, T., Lee, S. Y., Lee, K. G., Lee, J., Tame, M. S., Rockstuhl, C. & Lee, C. Quantum noise reduction in intensity-sensitive surface-plasmon-resonance sensors. *Phys. Rev. A* **96**, 033833 (2017).
- [133] Lee, J. S., Yoon, S. J., Rah, H., Tame, M. S., Rockstuhl, C., Song, S. H., Lee, C. & Lee, K. G. Quantum plasmonic sensing using single photons. *Opt. Express* **26**, 29272–29282 (2018).
- [134] Dowran, M., Kumar, A., Lawrie, B. J., Pooser, R. C. & Marino, A. M. Quantum-enhanced plasmonic sensing. *Optica* **5**, 628–633 (2018).
- [135] Kongsuwan, N., Xiong, X., Bai, P., You, J. B., Png, C. E., Wu, L. & Hess, O. Quantum plasmonic immunoassay sensing. *Nano Lett.* **19**, 5853–5861 (2019).
- [136] Mpofu, K. T., Lee, C., Maguire, G. E. M., Kruger, H. G. & Tame, M. S. Experimental measurement of kinetic parameters using quantum plasmonic sensing. *J. Appl. Phys.* **131**, 084402 (2022).

- [137] Dieleman, F., Tame, M. S., Sonnefraud, Y., Kim, M. S. & Maier, S. A. Experimental verification of entanglement generated in a plasmonic system. *Nano Lett.* **17**, 7455–7461 (2017).
- [138] Urii, S. A., Tashima, T., Zhang, X., Asano, M., Bechu, M., Güney, D. O., Yamamoto, T., Özdemir, Ş. K., Wegener, M. & Tame, M. S. Active control of a plasmonic metamaterial for quantum state engineering. *Phys. Rev. A* **97**, 053810 (2018).
- [139] Aharonovich, I., Castelletto, S., Simpson, D. A., Su, C. H., Greentree, A. D. & Prawer, S. Diamond-based single-photon emitters. *Rep. Prog. Phys.* **74**, 076501 (2011).
- [140] Akimov, A. V., Mukherjee, A., Yu, C. L., Chang, D. E., Zibrov, A. S., Hemmer, P. R., Park, H. & Lukin, M. D. Generation of single optical plasmons in metallic nanowires coupled to quantum dots. *Nature* **450**, 402–406 (2007).
- [141] Kolesov, R., Grotz, B., Balasubramanian, G., Stöhr, R. J., Nicolet, A. A. L., Hemmer, P. R., Jelezko, F. & Wrachtrup, J. Wave-particle duality of single surface plasmon polaritons. *Nat. Phys.* **5**, 470–474 (2009).
- [142] Huck, A., Kumar, S., Shakoor, A. & Andersen, U. L. Controlled coupling of a single nitrogen-vacancy center to a silver nanowire. *Phys. Rev. Lett.* **106**, 096801 (2011).
- [143] Hadfield, R. H. Single-photon detectors for optical quantum information applications. *Nat. Photonics* **3**, 696–705 (2009).
- [144] Falk, A. L., Koppens, F. H. L., Yu, C. L., Kang, K., De Leon Snapp, N., Akimov, A. V., Jo, M. H., Lukin, M. D. & Park, H. Near-field electrical detection of optical plasmons and single-plasmon sources. *Nat. Phys.* **5**, 475–479 (2009).
- [145] Heeres, R. W., Dorenbos, S. N., Koene, B., Solomon, G. S., Kouwenhoven, L. P. & Zwiller, V. On-chip single plasmon detection. *Nano Lett.* **10**, 661–664 (2010).
- [146] Heeres, R. W., Kouwenhoven, L. P. & Zwiller, V. Quantum interference in plasmonic circuits. *Nat. Nanotechnol.* **8**, 719–722 (2013).
- [147] Walker, J. A pseudorandom number sequence test program. <https://www.fourmilab.ch/random/> (2008).
- [148] Rukhin, A. *et al.* A statistical test suite for random and pseudorandom number generators for cryptographic applications. <https://nvlpubs.nist.gov/nistpubs/legacy/sp/nistspecialpublication800-22r1a.pdf> (2010).
- [149] Nielsen, M. A. & Chuang, I. L. Quantum bits in *Quantum Computation and Quantum Information: 10th Anniversary Edition* 13–16 (Cambridge University Press, 2010).
- [150] Nielsen, M. A. & Chuang, I. L. Single qubit gates in *Quantum Computation and Quantum Information: 10th Anniversary Edition* 17–20 (Cambridge University Press, 2010).
- [151] Nielsen, M. A. & Chuang, I. L. Projective measurements in *Quantum Computation and Quantum Information: 10th Anniversary Edition* 87–90 (Cambridge University Press, 2010).
- [152] Nielsen, M. A. & Chuang, I. L. Quantum measurement in *Quantum Computation and Quantum Information: 10th Anniversary Edition* 84–86 (Cambridge University Press, 2010).

- [153] Nielsen, M. A. & Chuang, I. L. Composite systems in *Quantum Computation and Quantum Information: 10th Anniversary Edition* 93–96 (Cambridge University Press, 2010).
- [154] Nielsen, M. A. & Chuang, I. L. Multiple qubit gates in *Quantum Computation and Quantum Information: 10th Anniversary Edition* 20–22 (Cambridge University Press, 2010).
- [155] Markham, D. & Sanders, B. C. Graph states for quantum secret sharing. *Phys. Rev. A* **78**, 042309 (2008).
- [156] Bell, B. A., Markham, D., Herrera-Martí, D. A., Marin, A., Wadsworth, W. J., Rarity, J. G. & Tame, M. S. Experimental demonstration of graph-state quantum secret sharing. *Nat. Commun.* **5**, 5480 (2014).
- [157] Friis, N., Orsucci, D., Skotiniotis, M., Sekatski, P., Dunjko, V., Briegel, H. J. & Dür, W. Flexible resources for quantum metrology. *New J. Phys.* **19**, 063044 (2017).
- [158] Shettell, N. & Markham, D. Graph states as a resource for quantum metrology. *Phys. Rev. Lett.* **124**, 110502 (2020).
- [159] Paternostro, M., Tame, M. S. & Kim, M. S. Hybrid cluster state proposal for a quantum game. *New J. Phys.* **7**, 226 (2005).
- [160] Prevedel, R., Stefanov, A., Walther, P. & Zeilinger, A. Experimental realization of a quantum game on a one-way quantum computer. *New J. Phys.* **9**, 205 (2007).
- [161] Nielsen, M. A. Cluster-state quantum computation. *Rep. Math. Phys.* **57**, 147–161 (2006).
- [162] Nielsen, M. A. & Chuang, I. L. General properties of the density operator in *Quantum Computation and Quantum Information: 10th Anniversary Edition* 101–105 (Cambridge University Press, 2010).
- [163] Nielsen, M. A. & Chuang, I. L. Examples of quantum noise and quantum operations in *Quantum Computation and Quantum Information: 10th Anniversary Edition* 373–386 (Cambridge University Press, 2010).
- [164] Nielsen, M. A. & Chuang, I. L. Quantum circuits in *Quantum Computation and Quantum Information: 10th Anniversary Edition* 22–24 (Cambridge University Press, 2010).
- [165] Raussendorf, R., Harrington, J. & Goyal, K. A fault-tolerant one-way quantum computer. *Ann. Phys.* **321**, 2242–2270 (2006).
- [166] Koch, J., Yu, T. M., Gambetta, J., Houck, A. A., Schuster, D. I., Majer, J., Blais, A., Devoret, M. H., Girvin, S. M. & Schoelkopf, R. J. Charge-insensitive qubit design derived from the Cooper pair box. *Phys. Rev. A* **76**, 042319 (2007).
- [167] Foroozani, N. *et al.* Development of transmon qubits solely from optical lithography on 300 mm wafers. *Quantum Sci. Technol.* **4**, 025012 (2019).
- [168] Tsioutsios, I., Serniak, K., Diamond, S., Sivak, V. V., Wang, Z., Shankar, S., Frunzio, L., Schoelkopf, R. J. & Devoret, M. H. Free-standing silicon shadow masks for transmon qubit fabrication. *AIP Adv.* **10**, 065120 (2020).
- [169] Chow, J. M. *et al.* Simple all-microwave entangling gate for fixed-frequency superconducting qubits. *Phys. Rev. Lett.* **107**, 080502 (2011).

- [170] Low, R. A. Pseudo-randomness and learning in quantum computation. (University of Bristol, 2009).
- [171] Emerson, J., Weinstein, Y. S., Saraceno, M., Lloyd, S. & Cory, D. G. Pseudorandom unitary operators for quantum information processing. *Science* **302**, 2098–2100 (2003).
- [172] Matthews, J. C. F., Whittaker, R., O’Brien, J. L. & Turner P. S. Testing randomness with photons by direct characterization of optical t -designs. *Phys. Rev. A* **91**, 020301(R) (2015).
- [173] Nielsen, M. A. & Chuang, I. L. Quantum process tomography in *Quantum Computation and Quantum Information: 10th Anniversary Edition* 389–394 (Cambridge University Press, 2010).
- [174] Chuang, I. L. & Nielsen, M. A. Prescription for experimental determination of the dynamics of a quantum black box. *J. Mod. Opt.* **44**, 2455–2467 (1997).
- [175] Mooney, G. J., White, G. A. L., Hill, C. D. & Hollenberg, L. C. L. Whole-device entanglement in a 65-qubit superconducting quantum computer. *Adv. Quantum Technol.* **4**, 2100061 (2021).
- [176] Mooney, G. J., White, G. A. L., Hill, C. D. & Hollenberg, L. C. L. Generation and verification of 27-qubit Greenberger-Horne-Zeilinger states in a superconducting quantum computer. *J. Phys. Commun.* **5**, 095004 (2021).
- [177] Skosana, U. & Tame, M. S. Demonstration of Shor’s factoring algorithm for $N = 21$ on IBM quantum processors. *Sci. Rep.* **11**, 16599 (2021).
- [178] Lundeen, J. S., Feito, A., Coldenstrodt-Ronge, H., Pagnell, K. L., Silberhorn, C., Ralph, T. C., Eisert, J., Plenio, M. B. & Walmsley, I. A. Tomography of quantum detectors. *Nat. Phys.* **5**, 27–30 (2009).
- [179] TensoredFilter, <https://qiskit.org/documentation/stubs/qiskit.ignis.mitigation.TensoredFilter.html>. Accessed 12 April 2021.
- [180] StateTomographyFitter, <https://qiskit.org/documentation/stubs/qiskit.ignis.verification.StateTomographyFitter.html>. Accessed 12 April 2021.
- [181] Urbanek, M., Nachman, B., Pascuzzi, V. R., He, A., Bauer, C. W. & De Jong, W. A. Mitigating depolarizing noise on quantum computers with noise-estimation circuits. *Phys. Rev. Lett.* **127**, 270502 (2021).
- [182] Epstein, J. M., Cross, A. W., Magesan, E. & Gambetta, J. M. Investigating the limits of randomized benchmarking protocols. *Phys. Rev. A* **89**, 062321 (2014).
- [183] Zeng, J., Wu, Z., Cao, C., Zhang, C., Hou, S. Y., Xu, P. & Zeng, B. Simulating noisy variational quantum eigensolver with local noise models. *Quantum Eng.* **3**, 1–14 (2021).
- [184] Caruso, F., Giovannetti, V., Lupo, C. & Mancini, S. Quantum channels and memory effects. *Rev. Mod. Phys.* **86**, 1203 (2014).
- [185] Georgopoulos, K., Emary, C. & Zuliani, P. Modelling and simulating the noisy behaviour of near-term quantum computers. *Phys. Rev. A* **104**, 062432 (2021).
- [186] Renes, J. M., Blume-Kohout, R., Scott, A. J. & Caves, C. M. Symmetric informationally complete quantum measurements. *J. Math. Phys.* **45**, 2171–2180 (2004).

- [187] Gross, D., Audenaert, K. & Eisert, J. Evenly distributed unitaries: on the structure of unitary designs. *J. Math. Phys.* **48**, 052104 (2007).
- [188] Barends, R. *et al.* Rolling quantum dice with a superconducting qubit. *Phys. Rev. A* **90**, 030303(R) (2014).
- [189] Poyatos, J. F., Cirac, J. I. & Zoller, P. Complete characterization of a quantum process: the two-bit quantum gate. *Phys. Rev. Lett.* **78**, 390–393 (1997).
- [190] Sanders, Y. R., Wallman, J. J. & Sanders, B. C. Bounding quantum gate error rate based on reported average fidelity. *New J. Phys.* **18**, 012002 (2016).
- [191] Boone, K. Concepts and methods for benchmarking quantum computers. (University of Waterloo, 2021).
- [192] Emerson, J., Alicki, R. & Życzkowski, K. Scalable noise estimation with random unitary operators. *J. opt., B Quantum semiclass. opt.* **7**, 347–352 (2005).
- [193] Lévi, B., López, C. C., Emerson, J. & Cory, D. G. Efficient error characterization in quantum information processing. *Phys. Rev. A* **75**, 022314 (2007).
- [194] Proctor, T., Rudinger, K., Young, K., Sarovar, M. & Blume-Kohout, R. What randomized benchmarking actually measures. *Phys. Rev. Lett.* **119**, 130502 (2017).
- [195] Carignan-Dugas, A., Boone, K., Wallman, J. J. & Emerson, J. From randomized benchmarking experiments to gate-set circuit fidelity: how to interpret randomized benchmarking decay parameters. *New J. Phys.* **20**, 092001 (2018).
- [196] Knill, E., Leibfried, D., Reichle, R., Britton, J., Blakestad, R. B., Jost, J. D., Langer, C., Ozeri, R., Seidelin, S. & Wineland, D. J. Randomized benchmarking of quantum gates. *Phys. Rev. A* **77**, 012307 (2008).
- [197] Boone, K., Carignan-Dugas, A., Wallman, J. J. & Emerson, J. Randomized benchmarking under different gatesets. *Phys. Rev. A* **99**, 032329 (2019).
- [198] Brown, W. G. & Eastin, B. Randomized benchmarking with restricted gate sets. *Phys. Rev. A* **97**, 062323 (2018).
- [199] Hashagen, A. K., Flammia, S. T., Gross, D. & Wallman, J. J. Real randomized benchmarking. *Quantum* **2**, 85 (2018).
- [200] Ballance, C. J., Harty, T. P., Linke, N. M., Sepiol, M. A. & Lucas, D. M. High-fidelity quantum logic gates using trapped-ion hyperfine qubits. *Phys. Rev. Lett.* **117**, 060504 (2016).
- [201] Chow, J. M., DiCarlo, L., Gambetta, J. M., Motzoi, F., Frunzio, L., Girvin, S. M. & Schoelkopf, R. J. Implementing optimal control pulse shaping for improved single-qubit gates. *Phys. Rev. A* **82**, 040305(R) (2010).
- [202] Córcoles, A. D., Gambetta, J. M., Chow, J. M., Smolin, J. A., Ware, M., Strand, J., Plourde, B. L. T. & Steffen, M. Process verification of two-qubit quantum gates by randomized benchmarking. *Phys. Rev. A* **87**, 030301(R) (2013).

- [203] Barends, R. *et al.* Logic gates at the surface code threshold: superconducting qubits poised for fault-tolerant quantum computing. *Nature* **508**, 500–503 (2014).
- [204] McKay, D. C., Sheldon, S., Smolin, J. A., Chow, J. M. & Gambetta, J. M. Three-qubit randomized benchmarking. *Phys. Rev. Lett.* **122**, 200502 (2019).
- [205] Ryan, C. A., Laforest, M. & Laflamme, R. Randomized benchmarking of single- and multi-qubit control in liquid-state NMR quantum information processing. *New J. Phys.* **11**, 013034 (2009).
- [206] Olmschenk, S., Chicireanu, R., Nelson, K. D. & Porto, J. V. Randomized benchmarking of atomic qubits in an optical lattice. *New J. Phys.* **12**, 113007 (2010).
- [207] Xia, T., Lichtman, M., Maller, K., Carr, A. W., Piotrowicz, M. J., Isenhower, L. & Saffman, M. Randomized benchmarking of single-qubit gates in a 2D array of neutral-atom qubits. *Phys. Rev. Lett.* **114**, 100503 (2015).
- [208] Veldhorst, M. *et al.* An addressable quantum dot qubit with fault-tolerant control-fidelity. *Nat. Nanotechnol.* **9**, 981–985 (2014).
- [209] Mayer, K., Hall, A., Gatterman, T., Halit, S. K., Lee, K., Bohnet, J., Gresh, D., Hankin, A., Gilmore, K. & Gaebler, J. Theory of mirror benchmarking and demonstration on a quantum computer. arXiv:2108.10431 (2021).
- [210] Liu, Y., Otten, M., Bassirianjahromi, R., Jiang, L. & Fefferman, B. Benchmarking near-term quantum computers via random circuit sampling. arXiv:2105.05232 (2021).
- [211] Yang, Z. P., Ku, H. Y., Baishya, A., Zhang, Y. R., Kockum, A. F., Chen, Y. N., Li, F. L., Tsai, J. S. & Nori, F. Deterministic one-way logic gates on a cloud quantum computer. *Phys. Rev. A* **105**, 042610 (2022).
- [212] Aliferis, P., Gottesman, D. & Preskill, J. Quantum accuracy threshold for concatenated distance-3 codes. *Quantum Inf. Comput.* **6**, 97–165 (2006).
- [213] Wallman, J. J. & Flammia, S. T. Randomized benchmarking with confidence. *New J. Phys.* **16**, 103032 (2014).
- [214] Wallman, J. J. Bounding experimental quantum error rates relative to fault-tolerant thresholds. arXiv:1511.00727 (2015).
- [215] Kueng, R., Long, D. M., Doherty, A. C. & Flammia, S. T. Comparing experiments to the fault-tolerance threshold. *Phys. Rev. Lett.* **117**, 170502 (2016).
- [216] Gottesman, D. Theory of fault-tolerant quantum computation. *Phys. Rev. A* **57**, 127–137 (1998).
- [217] Kimmel, S., Da Silva, M. P., Ryan, C. A., Johnson, B. R. & Ohki, T. Robust extraction of tomographic information via randomized benchmarking. *Phys. Rev. X* **4**, 011050 (2014).
- [218] Gottesman, D. The Heisenberg representation of quantum computers in *Group 22: Proceedings of the XXII International Colloquium on Group Theoretical Methods in Physics* 32–43 (International Press, 1999).

- [219] Granade, C., Ferrie, C. & Cory, D. G. Accelerated randomized benchmarking. *New J. Phys.* **17**, 013042 (2015).
- [220] Mooney, G. J., Hill, C. D. & Hollenberg, L. C. L. Entanglement in a 20-qubit superconducting quantum computer. *Sci. Rep.* **9**, 13465 (2019).
- [221] Elson, J. M. & Ritchie, R. H. Photon interactions at a rough metal surface. *Phys. Rev. B* **4**, 4129–4138 (1971).
- [222] Di Martino, G., Sonnefraud, Y., Kéna-Cohen, S., Tame, M. S., Özdemir, Ş. K., Kim, M. S. & Maier, S. A. Quantum statistics of surface plasmon polaritons in metallic stripe waveguides. *Nano Lett.* **12**, 2504–2508 (2012).
- [223] Edwards, A. L. The correlation coefficient in *An Introduction to Linear Regression and Correlation* 33–46 (W. H. Freeman, 1976).
- [224] Adcock, J. C., Vigliar, C., Santagati, R., Silverstone, J. W. & Thompson, M. G. Programmable four-photon graph states on a silicon chip. *Nat. Commun.* **10**, 3528 (2019).
- [225] Wang, J., Sciarrino, F., Laing, A. & Thompson, M. G. Integrated photonic quantum technologies. *Nat. Photonics* **14**, 273–284 (2020).
- [226] Bartolucci, S. *et al.* Fusion-based quantum computation. *Nat. Commun.* **14**, 912 (2023).
- [227] Bartolucci, S., Birchall, P., Bonneau, D., Cable, H., Gimeno-Segovia, M., Kieling, K., Nickerson, N., Rudolph, T. & Sparrow, C. Switch networks for photonic fusion-based quantum computing. *arXiv:2109.13760* (2021).
- [228] Loudon, R. Semiclassical theory of optical detection in *The Quantum Theory of Light Third Edition* 117–123 (Oxford University Press, 2000).