

PAPER • OPEN ACCESS

Quantum data centres: a simulation-based comparative noise analysis

To cite this article: K Campbell *et al* 2025 *Quantum Sci. Technol.* **10** 015052

View the [article online](#) for updates and enhancements.

You may also like

- [Improving mechanical squeezing in a dissipative optomechanical system with quadratic dispersive coupling](#)
Guolong Li, Xingmin Li, Jun Xin et al.
- [The influence of post-growth annealing on the optical properties of InAs quantum dot chains grown on pre-patterned GaAs\(100\)](#)
T V Hakkarainen, V Polojärvi, A Schramm et al.
- [A CNN-based four-layer DOI encoding detector using LYSO and BGO scintillators for small animal PET imaging](#)
Wen He, Yangyang Zhao, Xin Zhao et al.



Join Over 450 Quantum Experts In Washington D.C.

April 14 - 16, 2025 | Conrad Hotel, Washington D.C.

VIEW THE AGENDA

QT

Quantum Science and Technology



PAPER

OPEN ACCESS

RECEIVED

19 September 2024

REVISED

7 November 2024

ACCEPTED FOR PUBLICATION

10 December 2024

PUBLISHED

23 December 2024

Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Quantum data centres: a simulation-based comparative noise analysis

K Campbell^{*} , A Lawey and M Razavi

School of Electronic and Electrical Engineering, University of Leeds, Leeds LS29JT, United Kingdom

^{*} Author to whom any correspondence should be addressed.

E-mail: elkmc@leeds.ac.uk and m.razavi@leeds.ac.uk

Keywords: distributed quantum computing, quantum data centre, quantum noise

Abstract

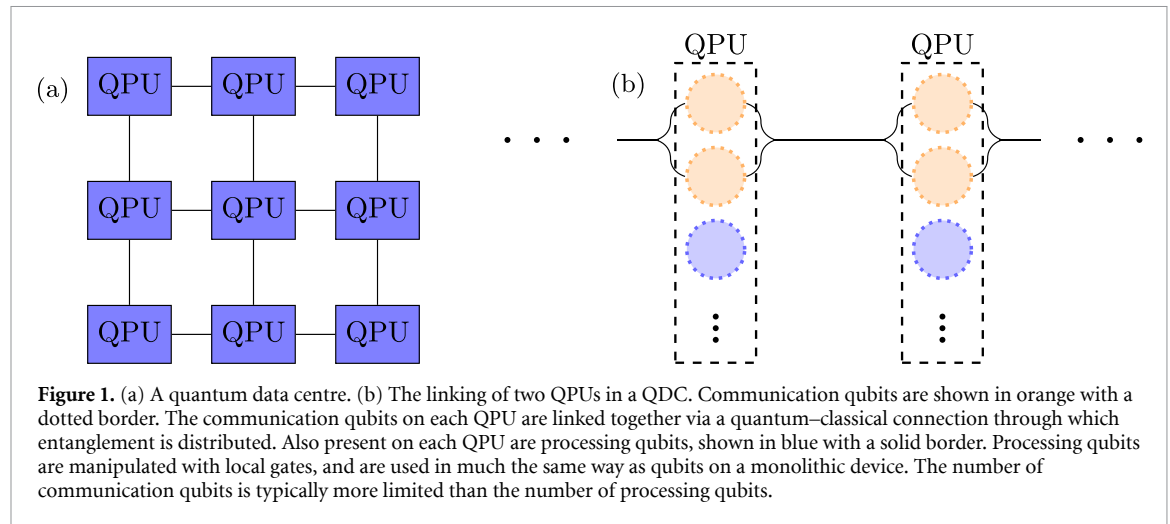
Quantum data centres (QDCs) could overcome the scalability challenges of modern quantum computers. Single-processor monolithic quantum computers are affected by increased cross talk and difficulty of implementing gates when the number of qubits is increased. In a QDC, multiple quantum processing units (QPUs) are linked together over short distances, allowing the total number of computational qubits to be increased without increasing the number of qubits on any one processor. In doing so, the error incurred by operations at each QPU can be kept small, however additional noise will be added to the system due to the latency cost and errors incurred during inter-QPU entanglement distribution. We investigate the relative impact of these different types of noise using a classically simulated QDC with two QPUs and compare the robustness to noise of the two main ways of implementing remote gates, cat-comm and TP-comm. We find that considering the quantity of gates or inter-QPU entangled links is often inadequate to predict the output fidelity from a quantum circuit and infer that an improved understanding of error propagation during distributed quantum circuits may represent a significant optimisation opportunity for compilation.

1. Introduction

Quantum computing offers the chance to revolutionise many fields, from chemistry and economics to cybersecurity and manufacturing [1], through its ability to simulate otherwise inimitable quantum systems [2] and efficiently solve certain problems significantly faster than conventional, classical, computers [3, 4]. However, tapping into the potential of quantum computing requires more qubits than are available in current devices.

A key limitation on the number of qubits to date has been the difficulty in avoiding unwanted interactions between qubits as the number of qubits or parallel gates within the same control system is increased. The correlated errors arising from such interactions are known as cross talk errors [5] and present a significant barrier to quantum error correction, which typically assumes errors to be localised and small [6]. Without quantum error correction, errors propagate rapidly, making it challenging to obtain a meaningful output from circuits acting on data sets large enough to have useful applications [7–9]. Therefore, simply adding more qubits to a quantum computer is not useful if the errors associated with those qubits are not adequately small.

A promising way of increasing the number of high-quality computational qubits is distributed quantum computing. A distributed quantum computer links together multiple distinct quantum computers, which are often referred to as quantum processing units (QPUs) when used in this way. The linking of QPUs creates a larger effective device with more qubits than any of its constituent components. This avoids the need to have many qubits share the same control system and so reduces the number of unwanted interactions. Thus, distributed quantum computing offers a scalable way of increasing the number of qubits while retaining their quality and the quality of local (intra-QPU) gates between qubits on the same QPU.



The price of these improvements relative to the single-processor, or monolithic, quantum computing paradigm is that the need for inter-QPU communication imposes additional sources of error. Whenever a multi-qubit gate acts on qubits that belong to different QPUs, some communication must occur between those QPUs. In most existing proposals for distributed quantum computers, this inter-QPU communication will consist of classical communication and distributing entanglement between qubits on each QPU.

Distributing entanglement between spatially separated entities is technologically challenging. One such challenge is that loss will occur when qubits are transported between QPUs. However, if increasing computational power is the main concern, there is little need to have a large inter-QPU separation, and so channel loss can be kept small. We shall refer to a distributed quantum computer where the distance between its constituent QPUs is sufficiently small that all of the QPUs can be kept in a single warehouse as a quantum data centre (QDC)¹. An example of a QDC is shown in figure 1(a). In the future, QDCs may be used to provide cloud computing services, but here, we focus on scenarios where the entire QDC is used to implement a single algorithm.

The linking of QPUs is shown in more detail in figure 1(b). As illustrated in figure 1(b), it is often helpful to segregate qubits into three different types: processing, communication, and flying qubits. Processing and communication qubits are static and exist within the confines of a QPU. Processing qubits² are exclusively used to carry out the same function that qubits in a monolithic quantum computer would do. They store quantum information and are manipulated using quantum gates. Communication qubits are reserved for being entangled with communication qubits on other QPUs. They are also manipulated with quantum gates as part of inter-QPU gate implementations and may participate in multi-qubit gates with processing qubits. Flying qubits are able to move between QPUs and are used to facilitate entanglement distribution between pairs of communication qubits. Most existing schemes use optical photons as flying qubits, and photons are likely to be uniquely well suited for this role, given their unparalleled speed in low-noise media and extensive use in classical telecommunications. Therefore, we assume photons are the flying qubits in the rest of this work.

Another challenge of distributing entanglement between QPUs is that, due to noise during the entanglement distribution process, the entangled state distributed will not be maximally entangled. Consequently, there is an inherent trade-off in any QDC between increasing the intra-QPU and inter-QPU errors. Intra-QPU errors can be reduced by keeping the individual QPUs small. However, more QPUs would then be required to obtain a given number of computational qubits. Splitting the computation between more QPUs typically means that there will be more inter-QPU gates. As such there will be more inter-QPU communication and the inter-QPU errors will increase. Conversely, when trying to obtain a given number of computational qubits, inter-QPU errors can be reduced by using larger QPUs, leading to higher intra-QPU errors.

¹ Alternative definitions of QDCs exist [10].

² Processing qubits are often referred to as data qubits in the literature [11–15]. However, we feel that the term data qubits is somewhat of a misnomer, implying that the qubits in question are used only for storage. Therefore, we use the term processing qubits throughout this work.

To explore this trade off, it is important to understand the relative impact of different errors in various contexts. To this end, we make the following contributions:

1. We demonstrate the limitations of first-order error analysis relative to full classical simulation in the QDC context.
2. We determine the relative impact of noise in inter-QPU entanglement, intra-QPU CNOT gates, and decohering qubits (as time evolves) on QDC performance.
3. We compare different approaches to inter-QPU (remote) gate operation by calculating the output fidelity for each approach.
4. We provide a pessimistic estimate for the number of inter-QPU gates that could realistically be implemented on a QDC in the upcoming years while retaining a useful output.

The rest of the paper is organised as follows. Firstly, we discuss the relevant literature in section 2. Following that, the problem at hand is elucidated in section 3. Then, the results of our analysis are presented and discussed in sections 4 and 5 for individual inter-QPU gates and larger quantum circuits, respectively. Finally, we summarise our findings and discuss their implications in section 6.

2. Related literature and problem motivation

In a QDC, careful consideration must be given to the implementation of multi-qubit gates. In particular, however a quantum circuit is partitioned between QPUs, it is likely that there will be multi-qubit gates that act on qubits assigned to different QPUs. We will henceforth refer to such inter-QPU gates as remote gates and refer to intra-QPU gates as local gates.

Most methodologies for carrying out remote gates boil down to one of two schemes. In the first scheme, which we call TP-comm (as in [12]), quantum teleportation [16, 17] is used to transfer quantum states from one QPU to another, and then the gate is conducted locally. Alternatively, as advocated for in [18], an entangled state can be maintained between the two QPUs while the local gate is applied. This latter idea is extended in [19], where it is shown that the remote gate methodology proposed in [18] can be decomposed into subprotocols [19], called cat-entanglement and cat-disentanglement, between which one or more local gates are conducted. In doing so, one can implement Shor's algorithm using fewer inter-QPU entangled pairs [20]—which are referred to as entangled bits or ebits³. The term cat-comm [19, 20] is used here to describe implementing a remote gate or gates in this way.

Subsequent work [11–15, 21–27], compiles arbitrary quantum circuits into a form suitable for distributed quantum computing, with a view to minimising the number of ebits used. Many compiler proposals [12, 13, 15] involve a step we call communication assignment [12, 13, 15], in which cat-comm or TP-comm is assigned to remote gates [27] or groups of remote gates [14, 29]. Current work on communication assignment [12, 13, 15] compares cat-comm or TP-comm with respect to the number of ebits they each use or their latency cost. Both metrics are intimately related due to the disproportionate impact of entanglement distribution on latency. Therefore, using either metric implicitly assumes that latency is the most significant optimisation parameter for distributed quantum computing or that entanglement error is the most impactful source of noise⁴. Such assumptions may or may not be true in the QDC setting. Moreover, there are many cases where cat-comm and TP-comm have the same ebit and/or latency cost for implementing a remote gate or group of remote gates [12], and so previous work gives no clear way of distinguishing between cat-comm and TP-comm in such circumstances.

In our work, we instead consider the output fidelity, F_{out} , after a remote gate or quantum circuit. F_{out} indicates the similarity between the ideal output, ρ_{ideal} , and the actual output, ρ_{noisy} , obtained in the presence of noise. Specifically [28],

$$F_{\text{out}} = \left(\text{Tr} \sqrt{(\rho_{\text{ideal}})^{\frac{1}{2}} \rho_{\text{noisy}} (\rho_{\text{ideal}})^{\frac{1}{2}}} \right)^2. \quad (1)$$

Crucially, this allows us to differentiate between cat-comm and TP-comm when the number of ebits is the same, and to understand the robustness of each scheme to different types of error. Using F_{out} , we can also

³ Technically, ebits could refer to intra-QPU entangled pairs too but in this paper, it is the inter-QPU ebits that are relevant. Consequently, ebit is used to refer to inter-QPU entangled pairs specifically in this work.

⁴ Latency considerations are complicated in the quantum computing context by the greater difficulty correcting errors relative to classical computing. In current quantum computers, more noise means more runs of the circuit are required to obtain a given precision in the result, increasing the latency. Moreover, quantum error correction, which is expected to be incorporated into future quantum computers, only works if the errors are small [6]. Therefore, simply arguing that a process takes a long time relative to other processes does not automatically ensure that the overall latency will be lower.

compare the relative impact of different error types. With this far more detailed understanding of cat-comm, TP-comm, and errors, we hope to facilitate more optimal compilation heuristics in the future.

Our work is not the first to consider the errors in the QDC setting. Some of the earliest work to do so is [29]. However, the results of [29] are specific to the quantum phase estimation algorithm [30, 31] and fewer types of noise are considered. The only local error considered is time-dependent dephasing, for which effective suppression techniques exist [32], and comparisons are conducted not between different types of noise, as we do, but between different implementations of the quantum phase algorithm—one in which QPUs can only communicate classically and another in which QPUs can also share quantum entanglement. Consequently, it may be difficult to ascertain from [29] where optimisation efforts for hardware and compilation should focus. Moreover, the figure of merit used by [29] is a cost function which treats the ‘cost’ of running a computation on each QPU and the ‘cost’ of sending measurement outcomes to a central QPU as abstract parameters and relates these ‘costs’ to the number of QPUs. This degree of abstraction makes the *a priori* linking of results to experiment challenging and implicitly assumes that intra-QPU errors are uncorrelated with the number of QPUs. In other words, they assume that the number of qubits and local gates on each QPU are fixed. This neglects the fact that cat-comm and TP-comm require additional local gates, rather than just ebits. By contrast, F_{out} , which we use, can be experimentally measured and has clear physical significance.

Later work [33] uses classical simulation to find the output fidelity of two different distributed implementations of the quantum phase estimation algorithms on a QDC. In one implementation, cat-entanglement and cat-disentanglement are conducted for each remote gate, whereas in the other implementation several remote gates are amalgamated by deferring cat-disentanglement until later in the circuit. The authors consider the impact of adding depolarising noise to inter-nodal entanglement and find that F_{out} declines less quickly with increasing depolarisation probability when the remote gates are amalgamated than when they are not. They also quantify the decrease in output fidelity with the number of devices for different depolarisation probabilities. However, no comparison between cat-comm and TP-comm is made and only entanglement error is considered.

A further limitation of [29, 33] is the specialisation to the quantum phase estimation algorithm. To generalise results, it is necessary to understand how the building blocks of circuits work, as well as how errors propagate through circuits with different structures. For this reason, we consider the impact of errors on individual remote gates, as well as on a wide variety of larger circuits. In this way, we are able to more deeply and generally understand the impacts of different errors and communication schemes than previous work.

3. QDCs: main components

To understand where optimisation of future QDC manufacturing and compilation efforts should focus, it is important to understand the impact of different types of noise on a QDC. In this work, we quantify the extent to which the output of distributed quantum circuits is degraded by the different types of noise present in a QDC. The metric we use to do this is the output error, characterised by $1 - F_{\text{out}}$. We also compare the robustness of cat-comm and TP-comm to error.

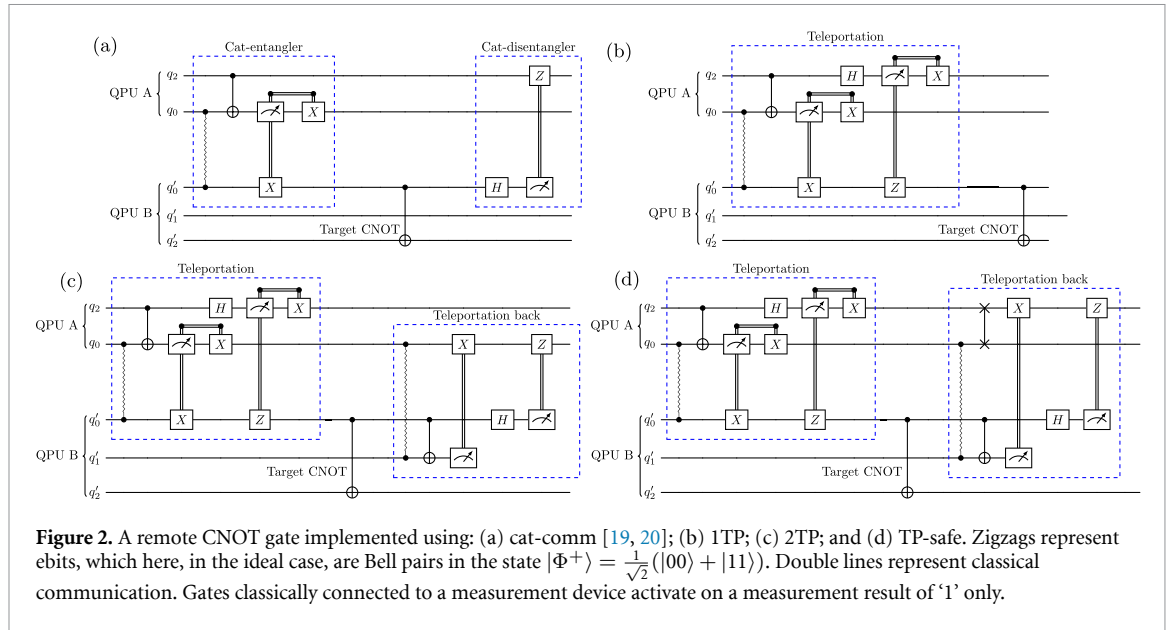
For simplicity, we assume that just two QPUs are present in the QDC we consider. We assume that all remote gates in a given quantum circuit are distributed using the same communication scheme and compare the impact of different communication schemes on the output error. Further details on the schemes considered are provided in section 3.1. We also consider different types of noise, which we model using the methods described in section 3.2. Additional information about how entanglement is distributed between QPUs and how entanglement error could be reduced are discussed in section 3.3.

3.1. Cat-comm and TP-comm circuits

Comparing cat-comm and TP-comm is not as simple as it may appear at first glance. In particular, there is some ambiguity in what it means to implement a remote gate using TP-comm. Therefore, in this work, we consider cat-comm and three different versions of TP-comm, which would be used in different scenarios within larger quantum circuits.

Consider first a remote CNOT gate implemented using cat-comm. The corresponding circuit diagram is shown in figure 2(a). The objective of this circuit is to implement a CNOT gate between QPU A’s processing qubit, q_2 , and QPU B’s processing qubit, q'_2 . q_2 is the control qubit and is in the arbitrary initial state $\alpha|0\rangle + \beta|1\rangle$, where α and β are complex numbers such that $|\alpha|^2 + |\beta|^2 = 1$. q'_2 is the target qubit and is in the arbitrary initial state $|\chi\rangle$. The communication qubits q_0 , from QPU A, and q'_0 from QPU B are also used.

The cat-comm remote CNOT gate consists of three subroutines: cat-entanglement, implementation of the CNOT locally, and then cat-disentanglement. Cat-entanglement entangles q_2 with q'_0 , so that they share the cat-like state, $\alpha|00\rangle + \beta|11\rangle$. In this way, the state of the control qubit, q_2 , on QPU A, is shared with q'_0 ,



on QPU B. The CNOT gate can then be conducted locally between q'_0 and q'_2 before the entanglement between q_2 and q'_0 is destroyed by the cat-disentanglement process shown in figure 2(a). A similar process can be used for any controlled-unitary (CU) gate.

Figure 2(b) shows the circuit diagram for a CNOT gate implemented using our first variant of TP-comm, which we call 1TP. The same qubits that were used in cat-comm are again considered here. In 1TP, the state of q_2 is teleported to q'_0 and then the CNOT gate is applied locally on QPU B. More generally, any local gate could be conducted after the teleportation. Unlike cat-comm, it is not necessary for the gate to be a CU gate.

The next TP-comm scheme, 2TP, is depicted in figure 2(c), again for the specific case of a remote CNOT. 2TP involves first applying 1TP and then teleporting the state of the control qubit back to QPU A. More specifically, the second teleportation is from q'_0 , on QPU B, to the communication qubit, q_0 , on QPU A. An additional communication qubit is also needed on QPU B to facilitate the second teleportation.

Finally, a remote CNOT gate implemented using TP-safe, the final TP-comm scheme, is shown in figure 2(d). TP-safe is almost identical to 2TP except that the states of q_0 and q_2 are swapped during the second teleportation, so that the state of the control qubit is restored to where it first started: the processing qubit q_2 . This is done using a SWAP gate.

It is not immediately obvious which of the TP-comm schemes discussed here is most comparable with cat-comm. On the one hand, cat-comm and 1TP both use exactly the same quantities and types of ebits, measurements, classical communications, and local gates. On the other hand, after cat-comm, q'_0 is free to host any further ebits needed for additional remote gates. By contrast, after 1TP, q'_0 is needed to store the teleported state of the control qubit. To avoid overwriting this crucial quantum information when conducting future remote gates, a second teleportation can be used to free q'_0 . As we consider only two QPUs, the second teleportation must be back to QPU A, as is done in 2TP. More generally, the teleportation could be to any other QPU. Teleportation back to the original QPU is the least favourable scenario, but it can often be necessary when large quantum circuits are compiled. After 2TP, q'_0 on QPU B is left free, but q_0 on QPU A is occupied. To circumvent this issue, we can carry out a SWAP gate between q_0 and q_2 to restore the teleported state back to where it started prior to the remote gate. This is safe to do because we know that q_2 was measured and then re-initialised during the first teleportation and so carries no non-trivial quantum information.

Cat-comm and TP-safe can always be used, regardless of what subsequent remote gates are scheduled. 1TP and 2TP cannot always be used but are less resource-intensive than TP-safe. In circumstances where no further remote gates are scheduled on the QPU that holds a teleported state, there is no need to free the communication qubits, and it will be better to use 1TP or 2TP than TP-safe. Otherwise, TP-safe is the only TP-comm scheme that could be used. TP-safe represents the very worst-case scenario that may be necessary during compilation of a distributed quantum circuit using TP-comm, and 1TP is the very best. By considering 1TP and TP-safe, we aim to provide an upper and lower bound on the performance of TP-comm, while 2TP sits somewhere between those two extremes.

3.2. Error models

Three principle sources of error are considered in this work: imperfections in the distributed ebits, imperfect implementation of local two-qubit gates, and time-dependent memory decoherence.

In cat-comm and TP-comm, we assume that, in the absence of noise, the ebits used have the pure state $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. When modelling imperfections in the ebits, we assume that their state is a Werner state of the form

$$\rho_w = F_w |\Phi^+\rangle\langle\Phi^+| + \frac{1-F_w}{3} (|\Phi^-\rangle\langle\Phi^-| + |\Psi^+\rangle\langle\Psi^+| + |\Psi^-\rangle\langle\Psi^-|), \quad (2)$$

where $0 \leq F_w \leq 1$, and $|\Phi^\pm\rangle = \frac{1}{\sqrt{2}}(|00\rangle \pm |11\rangle)$, $|\Psi^\pm\rangle = \frac{1}{\sqrt{2}}(|01\rangle \pm |10\rangle)$ are the Bell states, which form a basis for all two-qubit entangled states. A pair of qubits which are collectively in a Bell state is referred to as a Bell pair. We refer to $\epsilon_{\text{ebit}} = 1 - F_w$ as the entanglement error.

In reality, the form of the non-ideal entangled states is unlikely to be exactly known and may not be so symmetrical in nature. Nonetheless, any mixture of bipartite entangled states can be transformed into the Werner state via a random bilateral rotation of each of the entangled qubits [34]. Bilateral rotations require only single-qubit gates with relatively low errors and are sometimes done as part of entanglement distillation [34, 35], which is likely to be necessary in the QDC context. Moreover, less symmetrical imperfections are often advantageous for error correction or mitigation [36]. As such, the Werner state is likely to represent a worst-case scenario, which is useful for benchmarking purposes.

Non-ideal application of local gates is modelled by applying a depolarisation channel of the form [37]

$$\rho_{\text{in}} \rightarrow (1 - \epsilon_{\text{cnot}}) U_{i,j} \rho_{\text{in}} U_{i,j}^\dagger + \frac{\epsilon_{\text{cnot}}}{4} \text{Tr}_{i,j}(\rho_{\text{in}}) \otimes \mathbb{1}_{i,j}, \quad (3)$$

to all local two-qubit gates, which in this work are all CNOT gates. Here, ϵ_{cnot} is the error probability; $U_{i,j}$ is an ideal two-qubit gate operation on the i th and j th qubits; ρ_{in} is the input density matrix; $\text{Tr}_{i,j}$ is the partial trace [31] with respect to the subspaces of qubits i and j ; and $\mathbb{1}_{i,j}$ is the identity operation acting on the these subspaces.

Memory decoherence on qubit k is modelled using the depolarisation channel

$$\rho_{\text{in}} \rightarrow e^{-\Delta t r} \rho_{\text{in}} + \frac{(1 - e^{-\Delta t r})}{2} \text{Tr}_k(\rho_{\text{in}}) \otimes \mathbb{1}_k, \quad (4)$$

where Δt is the time since the initialisation of qubit k , r is the hardware dependent depolarisation rate, and ρ_{in} is the input density matrix.

3.3. Entanglement distribution and distillation

Entanglement distribution can be carried out in several ways. For instance, photons can be entangled with communication qubits and then measured out, leaving the communication qubits entangled with one another. The measurement process used is typically a Bell state measurement (BSM) [38]. A BSM interacts and measures a pair of qubits to ascertain whether those qubits are in one of the Bell states. In principle, a BSM can be conducted using a CNOT gate followed by a hadamard gate on the control qubit [38]. This is known as a deterministic BSM. In practice, deterministic BSMs may be difficult to achieve and so partial BSMs using linear optics and photodetectors are frequently used. A partial BSM relies on post-selection of measurement results to probabilistically determine whether a pair of qubits are in one of the Bell states.

Regardless of whether a partial or deterministic BSM is used, most entanglement distribution schemes require the entire entanglement distribution process to be repeated several times. For partial BSMs, this is an intrinsic feature of their probabilistic nature. For deterministic BSMs, repetition is necessitated by the presence of loss in the communication channel. If photons do not reach the detectors in a BSM module, new photons must again be entangled with communication qubits, and sent until they reach the detectors. In this way, loss is heralded by the lack of proper detection events.

The fidelity of the ebits can also be improved by using a more advanced heralding process and imposing more stringent fidelity requirements on the entangled states distributed [38]. This will lower the success rate of entanglement, necessitating a greater number of entanglement distribution attempts. Alternatively, but similarly, entanglement distillation [34, 35, 39, 40] could be used. In entanglement distillation, multiple entangled pairs are distributed between QPUs and local gates are conducted between communication qubits, entangling them. Some of the qubits are then measured out and inferences are made about the fidelity of the remaining qubits. This process is repeated iteratively until the fidelity reaches the desired threshold.

Other options for improving the fidelity of ebits include applying quantum error correction [6] locally on the communication qubits for each QPU and encoding the ebit into a larger many-qubit entangled (cluster) state [38].

Regardless of how loss and noise are mitigated, there is a cost to be paid. All of the aforementioned schemes for improving the fidelity of ebits involve a latency cost due to repetitions of the entanglement process, decoding time, or the addition of local gates. Any additional local gates will also be imperfect, adding further noise into the system. Moreover, in many cases, redundant communication qubits and/or photons are also required. Therefore, there is always a trade-off to be made between improving ebit fidelity and paying the price of doing so.

In this work, we abstract from the details of the entanglement distribution scheme used by assuming that communication qubits are entangled at a fixed rate R . The time taken for a given ebit to be distributed is assumed to always be $\frac{1}{R}$ from when the ebit is requested during the compilation process. Ebit requests are generated when the ebit is first needed, which is unfavourable for latency. That said, it does minimise the time that ebits have to decohere after they have been distributed. With these assumptions, all loss is incorporated into the value of R rather than being explicitly modelled. Noise is accounted for by assuming that the ebit is in the Werner state as discussed in section 3.2.

Despite the abstractions we make, it is important to keep in mind the trade-offs inherent to altering R or reducing the entanglement error. Even with current technology, entanglement error could be reduced from the current state-of-the-art values but, as discussed previously, this would come at the expense of a lower entanglement distribution rate and often requires additional qubits and local gates. These trade-offs are important to keep in mind when contemplating our results.

4. Single remote gates: noise analysis

To understand the behaviour of something as complex as a distributed quantum circuit, it is helpful to first investigate the constituent building blocks. With this in mind, we consider here a single remote CNOT gate, implemented using the schemes in figure 2. The $U3 + \text{CNOT}$ basis [41] is used, meaning that the SWAP gate from TP-safe (see figure 2(d)) is decomposed into three CNOT gates [41]. All other remote gate schemes are conducted exactly as shown in figure 2. Here, we first discuss the error analysis methods that we use in section 4.1. Then, we present and analyse the results of our error analysis in section 4.2.

4.1. Error analysis

4.1.1. Analytical approach

A natural starting point for any error analysis is to derive analytical results where possible. Even analytical results related to a small subset of the broader problem can be used to verify the results obtained by other methods. Here, we present closed form analytical results for arbitrary remote controlled-unitary gates implemented using 1TP and cat-comm in the presence of entanglement error only (EEO). All other errors are assumed to be zero. We defer the derivation to appendix A for brevity.

For 1TP, we obtain

$$F_{\text{out}} = \frac{1 + 2F_w}{3}. \quad (5)$$

Remarkably, this expression is independent of which controlled-unitary operation is applied or the input to that controlled-unitary operation.

For cat-comm, we make the additional limiting assumption that the processing qubits are initially separable, in order to make the calculation tractable. As such, the input state to the remote controlled-unitary gate can be represented as

$$(\alpha|0\rangle + \beta|1\rangle)_{q_2} \otimes |\chi\rangle_{q'_2}, \quad (6)$$

where the subscripts q_2 and q'_2 refer to the qubits shown in figure 2(b), α and β are complex numbers such that $|\alpha|^2 + |\beta|^2 = 1$, and $|\chi\rangle$ is an arbitrary pure state. With this assumption, we find that

$$F_{\text{out}} = F_w + \frac{1 - F_w}{3} \left((|\alpha|^2 - |\beta|^2)^2 + 2|\alpha|^4 |\langle \chi | U_{q'_2} | \chi \rangle_{q'_2}|^2 + 2|\beta|^4 |\langle \chi | U_{q'_2}^\dagger | \chi \rangle_{q'_2}|^2 \right), \quad (7)$$

where $U_{q'_2}$ is the unitary operation applied to q'_2 during the remote CU gate. Unlike equation (5), this does depend on the type of CU gate applied and the input state of the control and target qubits.

Considering the case where a CNOT gate is applied and adding the further limiting assumption that $|\chi\rangle$ is an eigenstate of the computational basis, equation (7) simplifies to

$$F_{\text{out}} = F_w + \frac{1 - F_w}{3} (2|\alpha|^2 - 1)^2. \quad (8)$$

Although, it applies to a very specific situation, equation (8) is a convenient univariate expression for testing other error analysis tools.

4.1.2. First-order approximations

To begin extending beyond the limiting assumptions of the analytical analysis, it is instructive to make first-order approximations to the results. These first-order approximations carry assumptions of their own but allow local multi-qubit gate errors to also be considered. Here, and throughout the remainder of this work, we use the $U_3 + \text{CNOT}$ set of gates and so the only multi-qubit gates considered are CNOTs.

For sufficiently small local gate and entanglement error rates, ϵ_{cnot} and ϵ_{ebit} , respectively, one might expect the output fidelity, F_{out} , to degrade linearly with the number of imperfect CNOT gates, n_{cnot} , or ebits, n_{ebit} , yielding the equation:

$$F_{\text{out}} \approx 1 - n_{\text{cnot}}\epsilon_{\text{cnot}} - n_{\text{ebit}}\epsilon_{\text{ebit}}. \quad (9)$$

Alternatively, retaining the main desired terms in equations (2) and (3), we can approximate F_{out} by

$$F_{\text{out}} \approx (1 - \epsilon_{\text{ebit}})^{n_{\text{ebit}}} (1 - \epsilon_{\text{cnot}})^{n_{\text{cnot}}}. \quad (10)$$

This is consistent with equation (9) in the limit of $n_{\text{ebit}}\epsilon_{\text{ebit}}$ and $n_{\text{cnot}}\epsilon_{\text{cnot}} \ll 1$ and has been used to estimate the fidelity in the presence of gate errors for larger circuits [42].

Both equations (9) and (10) are quite general and can be used to describe local gate errors and entanglement errors for any quantum circuit that adheres to the assumptions used during derivation.

Memory depolarisation is continuous rather than discrete like local gate errors and entanglement errors. For this reason, memory depolarisation is not well described by an equation with the same form as equation (9) or (10), and we do not consider a first-order approximation for memory depolarisation in this work.

4.1.3. Simulation

In order to move past first-order approximations and enable extension to larger circuits, we generate numerical results using classical simulation. The simulator used is an event-based simulator built using Python 3.9 and the associated libraries: NetSquid [43] and nuqasm2 [44]. We modified nuqasm2 for our purposes.

Our simulation package converts arbitrary monolithic quantum circuits, specified using openQASM 2 [41] or as a list of gate tuples, to compiled distributed quantum circuits. The use of communication qubits is handled automatically without explicit user specification and the subroutines associated with each type of remote gate specified in figure 2 can be automatically generated and scheduled. Manual specification of subroutines is also possible. All simulated hardware can have custom or stock noise models specified, and simulated QPUs come with a variety of options, including deciding the size of the communication and processing qubit allocations. Efforts have been made to retain the modularity of the NetSquid package and users can also specify their own compilers, QPUs and connections for use with our simulation package. Further details of our simulation package API and how to access the package will be presented in a separate work.

As the simulator is event-based, anything that happens during the simulation is allocated an amount of time it will take to occur. Time is discretised based on the duration of those events. For example, after a gate or entanglement distribution is carried out, all qubits in the circuit will experience an amount of memory depolarisation correlated to the amount of time that gate or entanglement distribution took. Δt in the memory depolarisation model (equation (4)) is incrementally increased every time such an event occurs in the simulation. The events accounted for are: single-qubit gates, two-qubit gates (CNOTs), measurement, classical communication, and entanglement distribution. For simplicity, the classical communication latency is modelled using NetSquid's fibre delay model, which assumes classical communication occurs in $\frac{d}{2 \times 10^8 \text{ms}^{-1}}$, where d is the inter-QPU distance in metres. Entanglement distribution is carried out using the abstract model described in section 3.3.

4.2. Numerical results

Armed with the tools described in section 4.1, we are able to obtain numerical values for the output error. All simulated results are obtained using either a laptop with 16GB of RAM and an AMD Ryzen 7 4700 processor or a desktop with 32GB of RAM and an i7 processor.

We assume the simulated QDC has two QPUs each possessing just two communication qubits, which is believed to be a realistic limitation for near-term devices [12, 13]. We also assume that the processing qubits

Table 1. The nominal state-of-the-art parameter values used in our simulations.

Parameter	Value	Source
F_w	0.94	With $^{88}\text{Sr}^+$ qubits [45]
ϵ_{cnot}	0.4%	With $^{171}\text{Yb}^+$ qubits [46]
Single-qubit gate time	135 μs	[46]
Two-qubit gate time	600 μs	[46]
Measurement time	6 ms	Inferred from [12, 46, 47] ^a
R	182 Hz	[45]
T_1 ^b	10–100 s	[46]
d	2 m	[45]

^a References [12, 47] put the measurement time at five to ten times the two-qubit gate time for trapped-ion quantum computers. Therefore, we have estimated the measurement time to be ten times the two-qubit gate time from [46].

^b The qubit lifetime. The memory depolarisation rate, r , is given by $r = \frac{1}{T_1}$. When a single state-of-the-art value is required we use the midpoint ($r = 0.055$) of the quoted range.

have an input state of

$$\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)_{q_2} \otimes |0\rangle_{q'_2}, \quad (11)$$

where the subscripts q_2 and q'_2 refer to the corresponding qubits in figure 2. The control state was found to be most susceptible to entanglement errors out of all separable input states that have q'_2 in the state $|0\rangle_{q'_2}$, as discussed in appendix B.

The parameter values used for the different types of error and the durations of different operations are based on state-of-the-art values for trapped-ion systems from the literature, and are displayed in table 1. Where possible, parameters for current commercial hardware (IonQ's Aria quantum computer [46]) are used to inform the state-of-the-art-values, so as to make the simulated set-up more realistically implementable in the near-term.

When we consider a range of parameter values, we refer to the order of magnitude in which the relevant state-of-the-art parameter exists as the state-of-the-art range. To facilitate comparison between entanglement error and local gate error, we also often consider what happens when both error types are varied over the same range. When we do this, we consider a range of entanglement error which is an order of magnitude smaller than the 'state-of-the-art' value. As this entanglement error reduction would require one of the methods from section 3.3, we refer to the resulting range as the distilled range, in reference to entanglement distillation, although entanglement distillation is just one of the possible error reduction methods that could be used.

In the following, we first compare our different error analysis tools. This is done in section 4.2.1 for the building block modules shown in figure 2. After that, we compare cat-comm, 1TP, 2TP, and TP-safe in section 4.2.2. Finally, we compare the impact of different types of error in section 4.2.3.

4.2.1. Comparison of error analysis methods

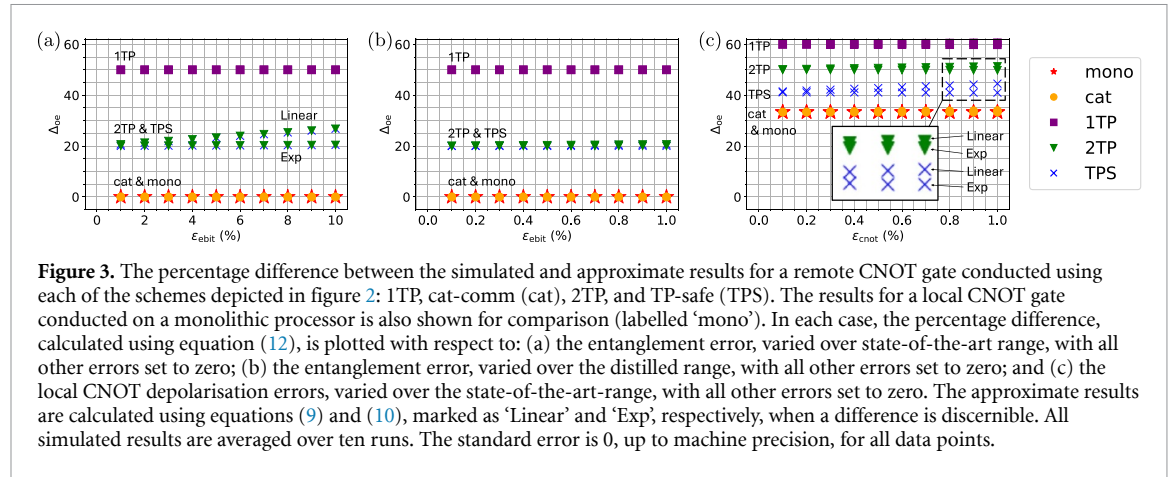
As alluded to in section 4.1.1, analytical results provide a useful robustness check for simulated ones. In our case, we find that the analytical expressions, equations (5) and (7), agree exactly with the simulated results for entanglement error, which suggests that the simulated results are likely to be robust.

The agreement is less strong between the analytical/simulation results and the first-order approximations given by equations (9) and (10). Indeed, we find both first-order expressions to be inadequate, indicating the need for more detailed error analysis, such as the classical simulation of errors conducted in this paper. The failure of the linear approximation is well established for larger monolithic circuits [48], however, here we indicate that clear discrepancies between both first-order approximations and the simulated results can occur for individual remote gates, when error parameters within the state-of-the-art or distilled ranges are used.

To demonstrate the inadequacy of the first-order approximation, we consider the percentage difference between simulated and approximate output errors, Δ_{oe} . This is calculated using the expression:

$$\Delta_{\text{oe}} = \frac{(1 - (F_{\text{out}})_{\text{approx}}) - (1 - (F_{\text{out}})_{\text{sim}})}{1 - (F_{\text{out}})_{\text{sim}}} \times 100. \quad (12)$$

A positive sign to the result indicates that the first-order approximation overestimates the output error, $1 - F_{\text{out}}$, relative to the simulated data, and a negative sign means that the first-order approximation



underestimates the output error. The subscript ‘sim’ refers to simulated data and ‘approx’ to data calculated using equation (9) or (10).

Figure 3 shows Δ_{oe} after a remote CNOT gate is implemented using each of the schemes in figure 2. A local CNOT gate conducted on a single monolithic quantum computer is also considered for comparison. Figures 3(a) and (b) show the Δ_{oe} with respect to the entanglement error, varied over the state-of-the-art and distilled ranges, respectively. All other errors are set to zero. Figure 3(c) shows Δ_{oe} with respect to the local two-qubit gate error. Again all other errors, including entanglement errors, are set to zero.

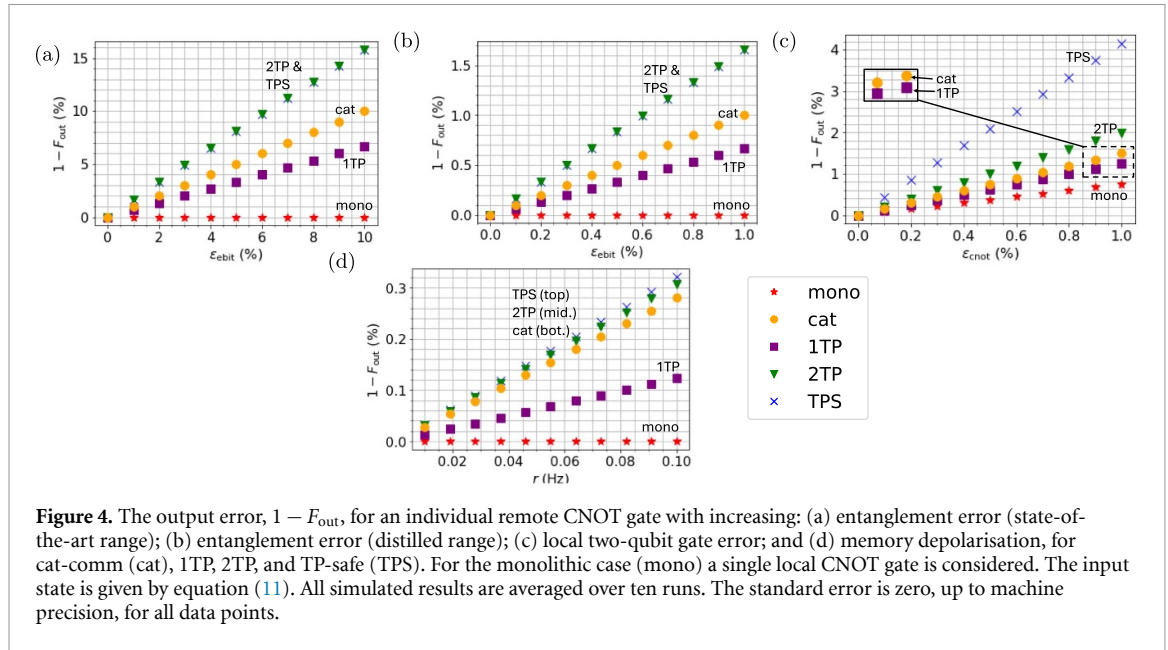
Using figure 3, several observations can be made:

- (i) The first-order approximations provide an upper bound on the output error. All Δ_{oe} values are non-negative in figure 3. This means that the output errors obtained using the simulation never exceed the approximate ones. This observation is resilient to the choice of input state as discussed in appendix C.1.
- (ii) The first-order approximations fail to distinguish between cat-comm and 1TP, for which the number of gates, classical communications, measurements, and ebits is the same. By contrast, the markedly different results for cat-comm and 1TP in figure 3 indicate that such a discrepancy does exist. This is one case where the numerical simulation reveals behaviour that is not apparent from first-order analysis alone.
- (iii) The discrepancy between simulated and approximate results can be large. For the input state chosen, this discrepancy is especially significant for 1TP, 2TP, and TP-safe. Figures 3(a) and (b) indicate that the discrepancy is approximately 20%–50% when only entanglement error is considered and figure 3(c) indicates that it is 40%–60% when local two-qubit gate errors are considered. Therefore, even with just one remote gate, the difference in the output error predicted by the first-order approximations and the simulation can be significant.
- (iv) Good agreement between simulated and first-order results is possible for some schemes. For cat-comm and the monolithic case, the simulated and first-order results agree when only entanglement error is considered. This can be seen from figures 3(a) and (b). This agreement is highly dependent on input state as discussed in appendix C.1.
- (v) The linear and exponential approximations, given by equations (9) and (10), respectively, perform similarly over the investigated error ranges, but, even with just one remote gate, the exponential approximation can be seen to perform slightly better than the linear one for 2TP and TP-safe. This holds for both entanglement error and local two-qubit gate error, within the state-of-the-art range, as can be seen from figures 3(a) and (c).

A key takeaway from these observations is that we cannot rely on first-order approximations alone for a quantitative error analysis of a QDC. Given the complexity of closed-form analytical analysis of arbitrary distributed quantum circuits, numerical simulation is therefore a key tool for understanding the role of errors in the QDC context.

4.2.2. Comparison of remote gate implementations

Using numerical simulation, we are able to compare different ways of implementing remote gates. The remote gate varieties we consider are shown in figure 2. The distributed quantum computing literature [11–14, 18–22, 24–27, 49–51] typically bases such comparisons on the number of EPR pairs required by each



scheme in various contexts, which falls short of the more detailed analysis possible using numerical simulation.

Figure 4 shows the output error, $1 - F_{\text{out}}$, after a remote CNOT gate. We vary one error at a time, respectively varying: entanglement error, over the state-of-the-art range in figure 4(a) and the distilled range in figure 4(b); local two-qubit gate error in figure 4(c); and memory depolarisation rate in figure 4(d). Again, all error parameters not explicitly varied in a given subfigure are set to zero.

Two key observations can be made:

- (i) For the input state given by equation (11), the remote gate schemes are ordered from lowest to highest by their impact on output error, as: 1TP, cat-comm, 2TP, and then TP-safe, where 2TP and TP-safe are equally damaging to the output when only entanglement error is considered. This ordering leads to the next more general observation.
- (ii) Cat-comm and 1TP can be distinguishable. This is consistent with observation (ii) from section 4.2.1, but is easier to see using figure 4. As discussed previously, it is not obvious *a priori* that cat-comm and 1TP could lead to different output errors, as they use exactly the same gates, measurements, ebits and classical communication. The difference between the output error produced by 1TP and cat-comm may offer an optimisation opportunity if either scheme is found to perform better when averaged over all input states. Such averaging is beyond the scope of the current work.

4.2.3. Comparison of error types

Although the relative qualitative performance of different schemes is unchanged by the type of error considered, another important algorithmic consideration for any quantum computer is how best to mitigate or, ideally, correct errors. A greater understanding of which errors have the most impact may help focus such efforts on where they are most needed and assist the efficient deployment of hardware resources to get the best possible output with the least possible effort. To this end, we investigate the impact of the different types of error on the output error.

We gain insight into the relative impact of different error types in two different ways. Firstly, we compare the quantitative values of output error displayed in the different subfigures of figure 4. This allows us to re-use figure 4 to gain insight into the relative impact of errors when entanglement error is varied within the state-of-the-art range. However, this analysis is insufficiently precise to compare errors when entanglement error is varied within the distilled range, where the magnitudes of output error are much smaller. To perform comparison between entanglement error in the distilled range and other types of error, it is useful to directly plot the output error caused by each type of error on the same figure. For this reason, in figure 5, we collect the output fidelity data used in figures 4(b)–(d), and re-express it all on one plot per remote gate scheme. In each plot, the loss of output fidelity due to EEO, two-qubit gate error only (GEO), and memory depolarisation only (MDO) are all shown.

Using figures 4 and 5, we find that:

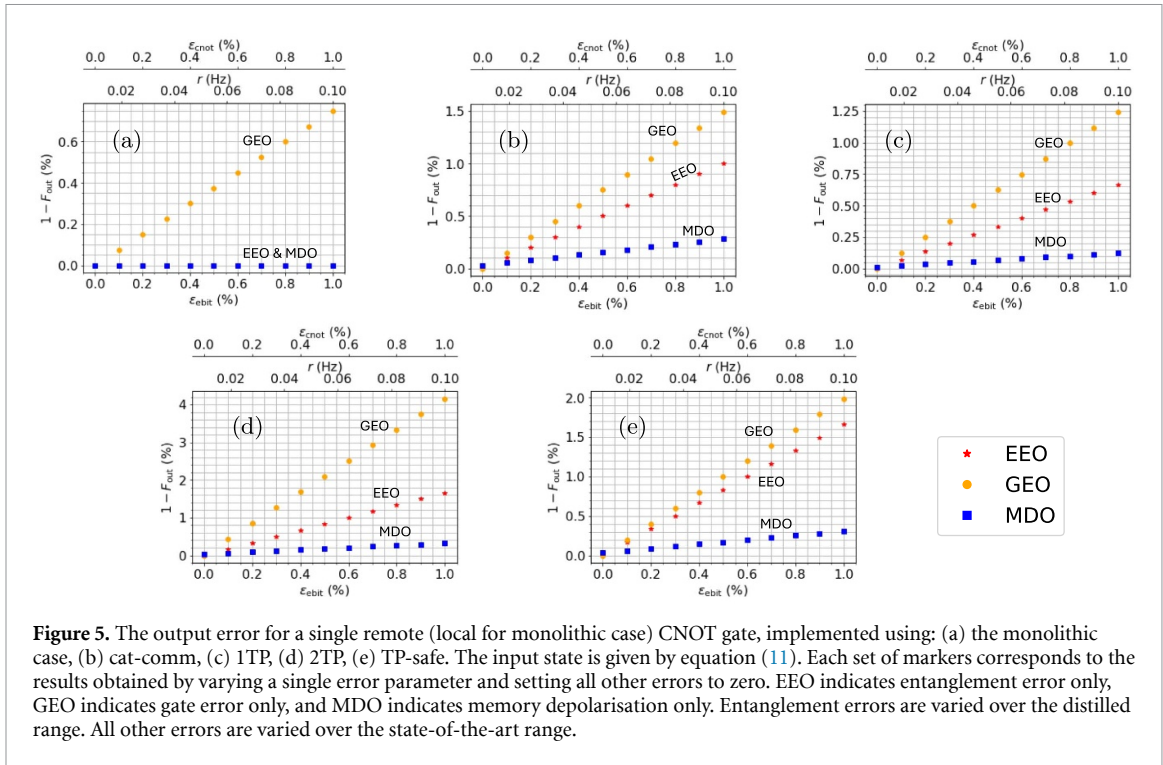


Figure 5. The output error for a single remote (local for monolithic case) CNOT gate, implemented using: (a) the monolithic case, (b) cat-comm, (c) 1TP, (d) 2TP, (e) TP-safe. The input state is given by equation (11). Each set of markers corresponds to the results obtained by varying a single error parameter and setting all other errors to zero. EEO indicates entanglement error only, GEO indicates gate error only, and MDO indicates memory depolarisation only. Entanglement errors are varied over the distilled range. All other errors are varied over the state-of-the-art range.

- (i) When entanglement error is varied over the state-of-the-art range, in our simulation, the error types are ranked from least to most detrimental to output error as: memory depolarisation, local two-qubit gate error, and then entanglement error.

Memory depolarisation is relatively negligible. When only memory depolarisation is non-zero, as in the bottom curve of figures 5(a)–(e), we can see that memory depolarisation leads to significantly less output error than local CNOT error for all considered depolarisation rates.

Local two-qubit gate error is in turn dominated by entanglement error. This can be seen by comparing figure 4(a), in which only entanglement error is non-zero, with figure 4(c), in which only local two-qubit gate error is non-zero. In figure 4(c), only TP-safe ever has an output error exceeding 2% for any local gate error in the considered range and there is no data point where the output error exceeds 4.2%. In figure 4(a), all schemes have an output error exceeding 2% as soon as the entanglement error reaches 4% and at 8% entanglement error, the output error of all schemes is greater than 4.2%.

- (ii) When entanglement error is varied over the distilled range, in our simulation, the ranking of the error types from least to most detrimental to output error changes to: memory depolarisation, entanglement error, and then local two-qubit gate error.

The ordering of entanglement error and local two-qubit gate error flips now that their magnitude is varied over the same range. This is easily seen from figures 5(a)–(e). This finding is not immediately obvious *a priori*, as one might expect the correlations intrinsic to entanglement to propagate errors more rapidly through the system.

Again, memory depolarisation is negligible and its impact is significantly lower than the other types of error shown in figures 5(a)–(e) for the parameter values considered here.

Overall, the findings suggest that entanglement error will be the most urgent optimisation consideration in the near term, however if the entanglement error is reduced, for example by using one of the methods from section 3.3, the local gate errors may well be more significant. The trade-off between improving entanglement errors and introducing additional local errors and/or latency via the methods discussed in section 3.3 would be an interesting and potentially very important point of future study.

5. Larger quantum circuits

The behaviour of individual remote gates can tell us many things, however, when many such gates operate in tandem, any errors in one gate will impact the rest of the system. How this happens may depend on the structure of the circuits in question and so it is not automatically obvious that the trends seen for individual remote gates will also hold for larger circuits. Here, we verify that the observations on the relative impact of

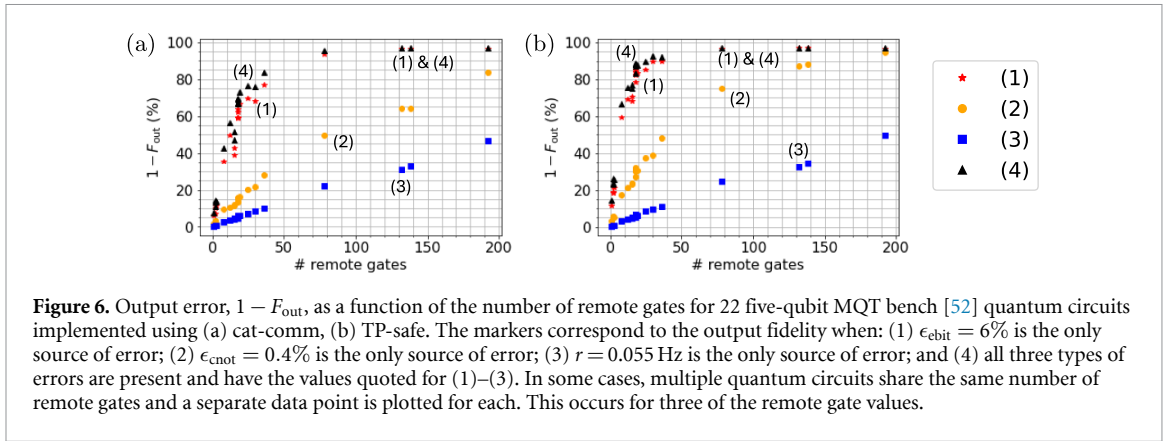


Figure 6. Output error, $1 - F_{\text{out}}$, as a function of the number of remote gates for 22 five-qubit MQT bench [52] quantum circuits implemented using (a) cat-comm, (b) TP-safe. The markers correspond to the output fidelity when: (1) $\epsilon_{\text{ebit}} = 6\%$ is the only source of error; (2) $\epsilon_{\text{cnot}} = 0.4\%$ is the only source of error; (3) $r = 0.055$ Hz is the only source of error; and (4) all three types of errors are present and have the values quoted for (1)–(3). In some cases, multiple quantum circuits share the same number of remote gates and a separate data point is plotted for each. This occurs for three of the remote gate values.

different errors made in section 4.2.3 do in fact hold for larger quantum circuits and discuss some differences between single remote gates and larger circuits. We describe the problem in more detail in section 5.1 and then give the results of our numerical simulations in section 5.2.

5.1. Problem description

We consider 22 different quantum algorithms taken from the MQT bench library [52]. We predominantly consider the five-qubit implementations of these algorithms but do also consider the impact of varying qubit number on each algorithm. These algorithms represent a variety of applications and range from algorithms of historical interest, such as the Deutsch–Jozsa algorithm, to common quantum subroutines, like the quantum Fourier transform, and NISQ applications, such as QAOA and VQE applied to different problems.

Compilation is kept simple. All quantum circuits are distributed by simple bipartitioning between two QPUs. Half of the processing qubits are given to one QPU and half to the other. If there is an odd number of processing qubits, then the additional qubit is arbitrarily allocated to a QPU based on the indices supplied for the original monolithic circuits by MQT bench [52]—the QPU whose qubits have the lower indices are given the extra qubit. Scheduling of operations is done greedily, but, for simplicity, ebits are requested only when needed. No attempt is made to merge remote gates together or otherwise optimise compilation. In this way, we avoid making the results specific to any one distributed quantum computing compiler, which may have imposed unforeseen structural bias on the circuits. The results represent a lower bound on what can be achieved with QDC.

The reasons for considering multiple implementations for TP-comm become apparent for these larger quantum circuits, as it becomes impossible to implement 1TP in many cases and even 2TP is often inapplicable with the simple compilation strategy employed. Therefore, we only consider TP-safe and cat-comm, which are compatible with the simple compilation strategy used. For each circuit considered, results are taken for when all remote gates are conducted with cat-comm and independently when all remote gates are conducted with TP-safe.

Of the error analysis tools discussed in section 4.1, we focus primarily on numerical simulation.

5.2. Numerical results

It is convenient to concisely express the behaviour of all 22 of the circuits investigated in one place. Doing this also allows the very different quantitative values obtained for different circuits to be directly compared. To this end, in figure 6, we consider the output error as a function of the number of remote gates. When multiple circuits share the same number of remote gates, all data points are shown—one for each circuit. Once again, unless otherwise noted, each type of error discussed in section 3.2 is considered separately, with the other errors set to zero. When a given error is being considered, the relevant error parameter is set to the state-of-the-art value stated in table 1. Here, we also include an additional curve corresponding to the output error when all of the errors are set to their state-of-the-art value at once.

Figure 7 shows the same thing as figure 6 except that the entanglement and local gate errors are set to $\epsilon_{\text{ebit}} = 0.5\%$ and $\epsilon_{\text{cnot}} = 0.5\%$, respectively, for comparison. For memory depolarisation, $r = 0.055$ Hz, as before.

Figure 6 shows that observation (i) from section 4.2.3 also holds for larger circuits. Once more, the impact of entanglement error dominates when the state-of-the-art error probability is used, followed by the impact of local gate errors, and then that of memory depolarisation. The one caveat is that, unlike for single remote gates, the impact of memory depolarisation is not negligible for some of the circuits considered.




Figure 7. The output error as a function of the number of remote gates for 22 five-qubit MQT bench [52] quantum circuits implemented using (a) cat-comm, (b) TP-safe. The markers correspond to the output error when: (1) $\epsilon_{\text{ebit}} = 0.5\%$ is the only source of error; (2) $\epsilon_{\text{cnot}} = 0.5\%$ is the only source of error; (3) $r = 0.055$ Hz is the only source of error; and (4) all three types of errors are present and have the values quoted for (1)–(3). In some cases, multiple quantum circuits share the same number of remote gates and a separate data point is plotted for each. This occurs for three of the remote gate values.

Nonetheless, memory depolarisation is still the least impactful type of error. In most cases, the difference in the impact of memory depolarisation from that of the other types of error considered remains quite large.

The magnitudes of the output error observed in figure 6 also offer some insight into what will be achievable with QDCs in the near-term. With as few as 10–20 remote gates in a circuit, the three errors considered can drive the output error to approximately 50% when cat-comm is used. An output error as high as 50% makes quantum advantage very unlikely. Even with entanglement error alone, under 20 remote gates can be tolerated before the output error exceeds 50%. For TP-safe, around five remote gates can be tolerated before the output error exceeds 50%.

When the magnitude of the entanglement error is reduced to fall within the distilled range and we impose that $\epsilon_{\text{ebit}} = \epsilon_{\text{cnot}} = 0.5\%$ (see figure 7), observation (ii) from section 4.2.3 is also found to be true for larger circuits. As with the single remote gates, local errors have the greatest adverse effect on the output error, followed by entanglement error and then memory depolarisation. Unlike for single remote gates, the impact of entanglement error and memory depolarisation is similar when cat-comm is used to implement all remote gates. This can be seen from the proximity of curves (1) and (3) in figure 7(a). For TP-safe, the impacts of entanglement error and memory depolarisation on output error are less similar to each other than for cat-comm, but they are more similar to each other than is observed for single remote gates. This can be seen using curves (1) and (3) of figure 7(b).

Additionally, reducing entanglement error to this extent has a significant quantitative impact on the output error. If cat-comm is used, around 30–40 remote gates could be tolerated before the output error observed in curve (4) of figure 7(a) exceeds 50%. With TP-safe, around 10–20 remote gates could be tolerated. Thus, reducing the entanglement error by an order of magnitude is likely to dramatically improve the number of quantum circuits with which potentially meaningful experiments could be done using a QDC.

Another instructive resource to consider is the number of qubits. The impact of varying the number of qubits is considered for all of the circuits in the MQT bench suite by implementing versions of the same circuits but with a different quantum register size in each case. The effect of this on the output error is shown in figures 8 and 9 for a selection of the MQT Bench circuits with fixed error parameters in the state-of-the-art and distilled ranges, respectively. For each set of data points, one type of error is non-zero and the other types of error are set to zero. The results for the remainder of the circuits considered are shown in the supplementary information [53]. As the standard error is negligible for the five-qubit circuits, which are averaged over ten runs, no such averaging is done for the other circuits, with different numbers of qubits, to avoid incurring the significant increase in runtime associated with repeating the simulation.

Figures 8 and 9 provide further evidence for the previous discussion on the relative impact of each error type. Using figure 8, we again see that, with the state-of-the-art error parameters, entanglement error has the largest adverse impact on output error, followed by local gate errors and then memory depolarisation. When entanglement error is reduced to be within the distilled range (see figure 9), we again find that the impact of local gate errors dominates. However, if all remote gates are implemented with cat-comm, we find that for many of the circuits considered, the impact of memory depolarisation can exceed that of entanglement error when the number of qubits used is high enough. Figure 9(a) shows this happening for the Deutsch–Jozsa algorithm. Further examples can be found in the supplementary information [53]. This is intuitive, as increasing the number of qubits will increase the latency of the circuit, meaning qubits are more likely to undergo memory decoherence.

A less intuitive feature of figures 8 and 9 is the existence of discrete jumps and plateaus in the output error. In some cases, the jumps and plateaus can be easily understood. For example, consider jump j_1 in figure 8(a), which shows $1 - F_{\text{out}}$ as a function of the number of qubits for Deutsch–Jozsa circuits

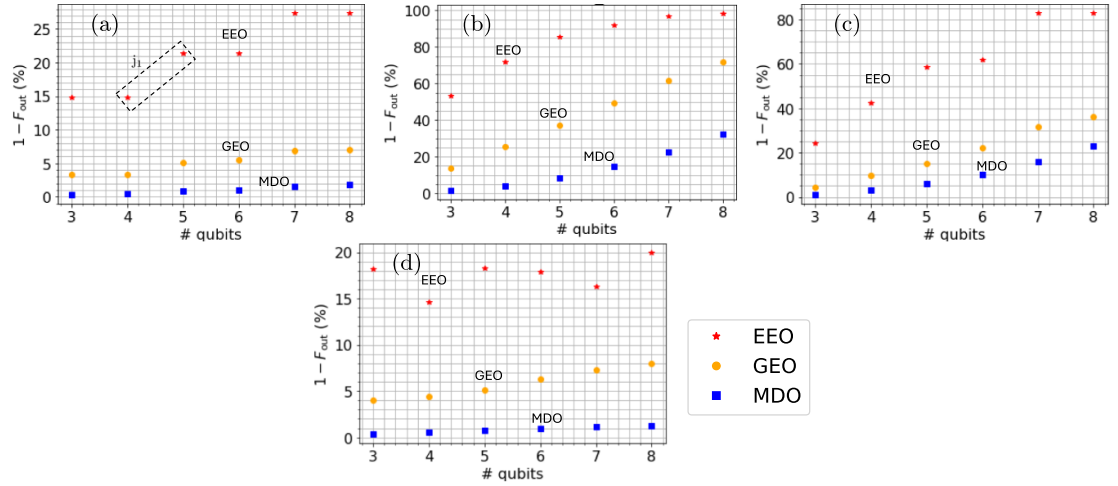


Figure 8. The output error, $1 - F_{\text{out}}$, as a function of the number of qubits used to implement: (a) the Deutsch–Jozsa circuit; (b) a quantum neural network; (c) the variational quantum eigensolver (VQE) applied to portfolio optimisation; and (d) VQE with a TwoLocal ansatz applied to the max-cut problem. All circuits are taken from [52]. In plot (c) cat-comm is used to implement all remote gates and in all other case TP-safe is used. Each set of data points shows the results when one of the error parameters (see the annotations) is non-zero, and all other errors are set to zero. The non-zero error values had the state-of-the-art values given in table 1.




Figure 9. The output error, $1 - F_{\text{out}}$, as a function of the number of qubits used to implement: (a) the Deutsch–Jozsa circuit with cat-comm; (b) the Deutsch–Jozsa algorithm with TP-safe; (c) a quantum walk without ancilla qubits using TP-safe. Each set of data points shows the results when one of the error parameters (see the annotations) is non-zero, and all other errors are set to zero. The non-zero error parameters used are: $\epsilon_{\text{ebit}} = 0.5\%$ for entanglement error only (EEO); $\epsilon_{\text{cnot}} = 0.5\%$ for two-qubit gate errors only (GEO); and $r = 0.055$ Hz for memory depolarisation only (MDO). For ϵ_{ebit} , this means the distilled range is used.

implemented with TP-safe. Jump j_1 occurs when increasing the number of qubits from four to five because the number of ebits increases from two to three when doing so. As such, it is natural for the entanglement error to increase and cause an increase in $1 - F_{\text{out}}$. By contrast, when we consider the five and six qubit circuits, both have the same number of ebits (three) and so we see a corresponding plateau in $1 - F_{\text{out}}$ for the curve in which only entanglement error, ϵ_{ebit} , is non-zero. Similar arguments can be used to explain the other jumps observed in figure 8(a) and 9(a), (b). On the other hand, in some cases, what causes jumps and plateaus remains an open question. Consider, for example, figure 8(d), which shows $1 - F_{\text{out}}$ for circuits implementing the variational quantum eigensolver (VQE) algorithm using TP-safe with increasingly many qubits. The curve for which only the entanglement error is non-zero fluctuates quite a bit, but the number of ebits remains constant at two for all data points. Therefore, unlike for the Deutsch–Jozsa algorithm, variations in $1 - F_{\text{out}}$ cannot be explained by changes in the number of ebits. What causes these fluctuations in $1 - F_{\text{out}}$ for VQE remains an open question. These features of the results allude to the fact that knowing the quantity of resources such as the number of ebits or local CNOT gates is not always enough to predict the error propagation in distributed quantum circuits. We infer that the exact structure of the compiled quantum circuit is also important to the error propagation and suggest caution against overreliance on heuristic optimisation that considers only resource quantities when compiling distributed quantum circuits, although this remains a useful starting point. There is still much to learn about error propagation in QDCs and distributed quantum computers more generally.

6. Conclusion and future outlook

In this work, we have used classical simulation to emulate operations on a QDC. We studied the behaviour of individual remote gates and a variety of larger distributed quantum circuits, in the presence of various types of error.

We found that first-order approximations to error propagation fail to accurately capture the behaviour of even the smallest systems for any finite error.

We also found that for error values obtainable with current trapped-ion hardware, the detrimental impact of imperfect inter-nodal entanglement dominates that of local errors, due to the much higher magnitude of entanglement error. The next most impactful of the errors considered is imperfect implementation of local gates, and time-dependent memory decoherence is negligible in comparison to other forms of error. However, if high-quality entanglement could be generated, so that the magnitudes of the gate and entanglement errors are comparable, then gate errors become more impactful than entanglement errors.

Additionally, we discovered that, despite having the same quantities and types of ebits, gates, classical communications, and measurements, cat-comm and 1TP can have different output errors for the same error types and parameter values. The exact discrepancy in output error is input state dependent.

Finally, we found that if current quantum technology could be successfully integrated to form a single QDC, a circuit containing around ten remote gates could be implemented with cat-comm before the output error exceeded 50%. If the magnitude of the entanglement error could be reduced to the same value as local gate error, a circuit containing at least 30 remote gates could be implemented. These numbers could most likely be significantly improved by using more optimal compilation strategies.

Our results are likely to be most applicable to matter-based QPUs with photonic interconnects, although, it is possible that they could be relevant to other platforms in the future, if a viable method of photon to communication qubit transduction is discovered for those platforms. Viewed collectively, our results indicate that small to medium-sized experimental demonstrations of a QDC are feasible in the near to mid-term. This has been recently corroborated by an experimental implementation of a small two-QPU QDC [54], albeit with a very low entanglement distribution rate (of 9.7 Hz)—to keep the entanglement error low. It may take longer before QDCs can compete with monolithic quantum computers in terms of output error. Our results also indicate which errors should be focused on for QDC hardware and circuit optimisation, and introduce some additional factors that should be considered when developing QDC compilation heuristics.

The further development of QDC compilation heuristics represents a possible avenue for future research, and there remains tremendous scope for future QDC optimisation. It would also be interesting to extend our work to larger QDCs with more QPUs and explore more complicated multipartite entanglement within a QDC network. Further work could include more detailed and specific analysis of certain hardware platforms, such as the one implemented in [54]. For now, comparisons to such work are challenging as they would require more bespoke simulation to mirror the idiosyncrasies of each platform. For example, each platform has its own specific set of native gates, which may differ from ours, and it may be necessary to implement specific error mitigation strategies, such as dynamical decoupling, if we wish to emulate specific devices as they are being implemented in experiments. We defer such investigations to future work.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://doi.org/10.5281/zenodo.13773229> [55].

Acknowledgments

We acknowledge funding from the Leeds Doctoral Scholarship and the UK EPSRC Grants EP/Y037421/1 and EP/X040518/1.

Appendix A. Analytical derivations of the output fidelity for 1TP and cat-comm

In the main text, the working for equations (5) and (8), the analytical expressions of F_{out} for 1TP and cat-comm, respectively, are omitted for brevity. Here, we provide the working for both expressions.

For both calculations, we model entanglement error as the distribution of the Werner state (see equation (2)) rather than the ideal $|\Phi^+\rangle$ state. The Werner state is a classical mixture of the four Bell states. Therefore, it is instructive to consider what happens to 1TP and cat-comm when each of the Bell states is distributed between QPUs.




Figure 10. A simplified version of the ITP quantum circuit shown in figure 2(b). The circuit shown is derived from figure 2(b) by assuming a measurement result of ‘0’ for all measurements in the circuit and generalising to an arbitrary remote gate, rather than a remote CNOT gate. $|\omega\rangle$ is the purified ideal state produced after an ideal teleportation from QPU A to QPU B. To produce $|\omega\rangle$, the system has been combined with a possibly fictitious environment (labelled env. in the figure), so that it can be considered to be pure, as discussed in section 2.5 of [31]. $|\eta\rangle$ is the ideal output from the circuit. l, m , and $n \in \mathcal{N}$ are arbitrary, possibly infinite numbers of qubits, which could be in the QDC itself or an external environment, and may be real or fictitious. R_i is a rotation of the ideal teleported state imposed by the possible non-ideality of the entangled state distributed, and U is a quantum gate local to QPU B’s qubits. The qubit labels q'_0 and q'_2 refer to the same qubits as in the original circuit shown in figure 2(b).

The ITP scheme depicted in figure 2(b) assumes that the $|\Phi^+\rangle$ state, specifically, was distributed between QPUs during the teleportation process. If one of the other Bell states is distributed instead then a Pauli error occurs in the teleported state. We represent this error as a fictitious gate $R_i \in \{\mathcal{I}, \sigma_x, \sigma_z, \sigma_z \sigma_x\}$, where σ_x and σ_z are the Pauli x and y operators, respectively.

Further simplifications can be made using the fact that when any of the Bell states is distributed between QPUs, the output is independent of the BSM results used during the teleportation process, provided that there are no local, intra-QPU, errors of any kind. This is simply verified by calculating the output for each measurement result and can be intuitively understood by noting that any measurement result of qubits on QPU A in figure 2(b) is equally likely to be obtained, regardless of which Bell state is distributed between QPU A and QPU B. Consequently, the circuit shown in figure 2(b) can be greatly simplified by assuming a specific BSM result has occurred, without loss of generality. For simplicity, we assume here that the BSM result is ‘00’, and obtain a simplified circuit diagram for implementing an arbitrary remote gate using ITP, as shown in figure 10.

Combining the results for each of the Bell states, the overall output state, ρ_{out} , can be represented in terms of the ideal input state, ρ_ω , as

$$\rho_{\text{out}} = U \left(\sum_{i=1}^4 p_i R_i \rho_\omega R_i^\dagger \right) U^\dagger, \quad (\text{A1})$$

where p_i is the probability of being in one of the Bell states (corresponding to the coefficients in the Werner state), and U is a unitary gate acting on QPU B from figure 10. For example, if a remote CNOT gate were being considered, as in figure 2(b), U would be the X (or NOT) gate. Here, we consider the more general case in which U is arbitrary and the remote gate need not be a controlled-unitary (CU) gate.

Purifying ρ_ω , by combining it with a, potentially fictitious, environment, E , as discussed in section 2.5 of [31], we obtain:

$$\rho_{\text{out}} = U \left(\sum_{i=1}^4 p_i R_i |\omega\rangle \langle \omega| R_i^\dagger \right) U^\dagger, \quad (\text{A2})$$

where $|\omega\rangle$ is the purified ideal input obtained when $|\Phi^+\rangle$ is distributed between QPU A and QPU B. The environment, E , may include qubits within the quantum data centre and fictitious environmental qubits external to it.

Again using purification, the purified ideal output from the circuit, when the correct Bell state, $|\Phi^+\rangle$, is distributed initially, is given by:

$$|\eta\rangle = U|\omega\rangle. \quad (\text{A3})$$

For an arbitrary pure state, $|\psi\rangle$, the fidelity definition given by equation (1) simplifies to [31]:

$$F(|\psi\rangle\langle\psi|, \rho_{\text{noisy}}) = \langle\psi|\rho_{\text{noisy}}|\psi\rangle. \quad (\text{A4})$$

Inserting equations (A2) and (A3) into (A4) gives an output fidelity F_{out} of

$$\begin{aligned} F_{\text{out}} &= \langle\omega|U^\dagger U \left(\sum_{i=1}^4 p_i R_i |\omega\rangle\langle\omega| R_i^\dagger \right) U^\dagger U |\omega\rangle \\ &= \sum_{i=1}^4 p_i |\langle\omega|U^\dagger U R_i |\omega\rangle|^2, \end{aligned} \quad (\text{A5})$$

which by the unitarity of U simplifies to

$$F_{\text{out}} = \sum_{i=1}^4 p_i |\langle\omega|R_i |\omega\rangle|^2. \quad (\text{A6})$$

Using the Schmidt decomposition [31], $|\omega\rangle = \sum_k \lambda_k |k_{q'_0}\rangle |k_{q'_2, E}\rangle$:

$$F_{\text{out}} = \sum_{i=1}^4 p_i \sum_{j,k} \lambda_j \lambda_k |\langle j_{q'_0} | \langle j_{q'_2, E} | (R_i)_{q'_0} | k_{q'_0} \rangle | k_{q'_2, E} \rangle|^2, \quad (\text{A7})$$

where λ_j and λ_k are non-negative real numbers satisfying $\sum_j \lambda_j^2 = 1$, and $|k_{q'_0}\rangle, |k_{q'_2, E}\rangle$ are orthonormal states for the q'_0 and $\{q'_2, E\}$ systems, respectively (see figure 10).

By the orthonormality of $|j_{q'_2, E}\rangle$ and $|k_{q'_2, E}\rangle$ and the commutativity of the basis states for the q'_0 and $\{q'_2, E\}$ subsystems:

$$\begin{aligned} F_{\text{out}} &= \sum_{j,k} \lambda_j \lambda_k \sum_{i=1}^4 p_i |\langle j_{q'_0} | (R_i)_{q'_0} | k_{q'_0} \rangle \delta_{j,k}|^2 \\ &= \sum_j \lambda_j^2 \sum_{i=1}^4 p_i |\langle j_{q'_0} | (R_i)_{q'_0} | j_{q'_0} \rangle|^2, \end{aligned} \quad (\text{A8})$$

where $\delta_{j,k}$ is the kronecker delta.

From there, we individually consider each term in the sum over index i .

For $i = 1$, $p_i = F_w$, and

$$p_1 |\langle j_{q'_0} | (R_1)_{q'_0} | j_{q'_0} \rangle|^2 = F_w |\langle j_{q'_0} | \mathbb{1} | j_{q'_0} \rangle|^2 = F_w. \quad (\text{A9})$$

For the remaining three terms, $p_i = \frac{1-F_w}{3}$. For $i = 2$, using the Schmidt decomposition, $|\omega\rangle = \sum_m \lambda_m |m_{q'_0}\rangle |m_{q'_2, E}\rangle$:

$$p_2 |\langle j_{q'_0} | (R_2)_{q'_0} | j_{q'_0} \rangle|^2 = \frac{1-F_w}{3} \sum_{m,n} \lambda_m \lambda_n \left(\langle m_{q'_0} | \langle m_{q'_2, E} | \sigma_{x_{q'_0}} | n_{q'_0} \rangle | n_{q'_2, E} \rangle \right)^2. \quad (\text{A10})$$

By the orthonormality of $|m\rangle_{q'_2, E}$ and $|n\rangle_{q'_2, E}$, and the independence of the q'_0 and $\{q'_2, E\}$ subsystems from each other:

$$\begin{aligned} (\text{A10}) &= \frac{1-F_w}{3} \sum_{m,n} \lambda_m \lambda_n \left| \langle m_{q'_0} | (\sigma_x)_{q'_0} | m_{q'_0} \rangle \delta_{m,n} \right|^2 \\ &= \frac{1-F_w}{3} \sum_m \lambda_m^2 \left| \langle m_{q'_0} | (\sigma_x)_{q'_0} | m_{q'_0} \rangle \right|^2. \end{aligned} \quad (\text{A11})$$

Any pure, single-qubit state can be represented in the form $\alpha|0\rangle + \beta|1\rangle$ and so we can say $|m_{q'_0}\rangle = \alpha|0\rangle_{q'_0} + \beta|1\rangle_{q'_0}$ without loss of generality. Using this, equation (A11) becomes

$$\begin{aligned} & \frac{1-F_w}{3} \sum_m \lambda_m^2 |(\alpha^* \langle 0| + \beta^* \langle 1|)(\alpha|1\rangle + \beta|0\rangle)|^2 \\ &= \frac{1-F_w}{3} \sum_m \lambda_m^2 |\alpha^* \beta + \beta^* \alpha|^2 \\ &= \frac{1-F_w}{3} \sum_m \lambda_m^2 (|\alpha^* \beta|^2 + |\beta^* \alpha|^2 + \alpha^* \alpha^* \beta \beta + \alpha \alpha \beta^* \beta^*) \\ &= \frac{1-F_w}{3} (2|\alpha|^2 |\beta|^2 + \alpha^* \alpha^* \beta \beta + \alpha \alpha \beta^* \beta^*), \end{aligned} \quad (\text{A12})$$

where the last line follows from the fact that $\sum_m \lambda_m^2 = 1$, as per the definition of Schmidt decomposition.

The calculations for $i = 3$ and $i = 4$ proceed similarly, except that we replace $\sigma_{x_{q'_0}}$ with $\sigma_{z_{q'_0}}$ and $\sigma_{z_{q'_0}} \sigma_{x_{q'_0}}$, respectively. We obtain:

$$\frac{1-F_w}{3} \left| |\alpha|^2 - |\beta|^2 \right|^2, \quad (\text{A13})$$

for $i = 3$, and

$$\frac{1-F_w}{3} |(\alpha^* \langle 0| + \beta^* \langle 1|)(\alpha|1\rangle - \beta|0\rangle)|^2 \quad (\text{A14})$$

for $i = 4$.

Substituting equations (A9), (A12), (A13) and (A14) into equation (A8) gives

$$F_{\text{out}} = \frac{1 + 2F_w}{3}. \quad (\text{A15})$$

Remarkably, this simple result has no dependence on the local operation, U , applied to QPU B's qubits, provided that U is unitary. The result is also independent of the input state to the processing qubits.

For cat-comm, things are complicated by the measurement on QPU B in the cat-disentanglement subroutine. This means that, although a similar simplified circuit to that depicted in figure 10 can be concocted, the operation U would no longer be unitary and so the cancellations possible when working with the Schmidt decomposition for 1TP no longer apply. Consequently, it is more convenient to use the density matrix formalism throughout the calculation.

The circuit output remains independent of the measurement outcomes, as before, and so there again exist four possible pure output states, one for each of the terms in the Werner state. To make the problem tractable, we make the limiting assumption that the processing qubits are initially separable. This reduces the generality of the results greatly but does allow some initial investigations to be made and means that simulated results can be checked. With these assumptions made, the possible output states, $|\text{out}\rangle$, to the circuit when the ebits used are in the state $|\Phi^\pm\rangle$ or $|\Psi^\pm\rangle$ are:

$$|\Phi^\pm\rangle_{q_0, q'_0} \rightarrow |\text{out}\rangle_{q_2, q'_2} = \alpha|0\rangle_{q_2} |\chi\rangle_{q'_2} \pm \beta|1\rangle_{q_2} U_{q'_2} |\chi\rangle_{q'_2}, \quad (\text{A16})$$

$$|\Psi^\pm\rangle_{q_0, q'_0} \rightarrow |\text{out}\rangle_{q_2, q'_2} = \alpha|0\rangle_{q_2} U_{q'_2} |\chi\rangle_{q'_2} \pm \beta|1\rangle_{q_2} |\chi\rangle_{q'_2}, \quad (\text{A17})$$

where here and hereafter we have highlighted discrepancies between the states in red. $|\chi\rangle_{q'_2}$ is the arbitrary pure state of qubit q'_2 from figure 10.

The output states from equations (A16) and (A17) correspond to the density matrices:

$$\begin{aligned} \rho_{|\Phi^\pm\rangle} &= |\alpha|^2 |0\rangle\langle 0|_{q_2} |\chi\rangle\langle \chi|_{q'_2} \pm \alpha\beta^* |0\rangle\langle 1|_{q_2} |\chi\rangle\langle \chi|_{q'_2} U_{q'_2}^\dagger \\ &\quad \pm \beta\alpha^* |1\rangle\langle 0|_{q_2} U_{q'_2} |\chi\rangle\langle \chi|_{q'_2} \\ &\quad + |\beta|^2 |1\rangle\langle 1|_{q_2} U_{q'_2} |\chi\rangle\langle \chi|_{q'_2} U_{q'_2}^\dagger, \end{aligned} \quad (\text{A18})$$

and

$$\begin{aligned} \rho_{|\Psi^\pm\rangle} = & |\alpha|^2 |0\rangle\langle 0|_{q_2} U_{q_2'} | \chi \rangle \langle \chi |_{q_2'} U_{q_2'}^\dagger \pm \alpha \beta^* |0\rangle\langle 1|_{q_2} U_{q_2'} | \chi \rangle \langle \chi |_{q_2'} \\ & \pm \beta \alpha^* |1\rangle\langle 0|_{q_2} | \chi \rangle \langle \chi |_{q_2'} U_{q_2'}^\dagger + |\beta|^2 |1\rangle\langle 1|_{q_2} | \chi \rangle \langle \chi |_{q_2'}, \end{aligned} \quad (\text{A19})$$

respectively.

Moreover, the ideal input state, $|\omega\rangle$, from before now has the state

$$|\omega\rangle = (\alpha|0\rangle + \beta|1\rangle)_{q_2} (a|0\rangle + b|1\rangle)_{q_2'}, \quad (\text{A20})$$

where a and b are complex numbers such that $|a|^2 + |b|^2 = 1$.

Inserting equation (A20) and equation into (A4) gives

$$\begin{aligned} F_{\text{out}} = & \sum_{i=1}^4 p_i \langle \omega |_{q_2, q_2'} (\rho_i)_{q_2, q_2'} | \omega \rangle_{q_2, q_2'} \\ = & \sum_{i=1}^4 p_i \left(|\alpha|^2 \langle \chi |_{q_2'} \langle 0 |_{q_2} (\rho_i)_{q_2, q_2'} | 0 \rangle_{q_2} | \chi \rangle_{q_2'} \right. \\ & + \alpha^* \beta \langle \chi |_{q_2'} \langle 0 |_{q_2} (\rho_i)_{q_2, q_2'} | 1 \rangle_{q_2} U_{q_2'} | \chi \rangle_{q_2'} \\ & + \beta^* \alpha \langle \chi |_{q_2'} U_{q_2'}^\dagger \langle 1 |_{q_2} (\rho_i)_{q_2, q_2'} | 0 \rangle_{q_2} | \chi \rangle_{q_2'} \\ & \left. + |\beta|^2 \langle \chi |_{q_2'} U_{q_2'}^\dagger \langle 1 |_{q_2} (\rho_i)_{q_2, q_2'} | 1 \rangle_{q_2} | \chi \rangle_{q_2'} \right) \end{aligned} \quad (\text{A21})$$

with ρ_i being the density matrix for one of the states in equations (A18), (A19).

Inserting (A18) and (A19) into (A21) for each i , and using the orthonormality of $\{|0\rangle, |1\rangle\}$ to cancel terms yields

$$\begin{aligned} & p_{1,2} \langle \omega |_{q_2, q_2'} (\sigma_{1,2})_{q_2, q_2'} | \omega \rangle_{q_2, q_2'} \\ = & \left(|\alpha|^4 \left| \langle \chi | \chi \rangle_{q_2'} \right|^2 \pm |\alpha|^2 |\beta|^2 \langle \chi | \chi \rangle_{q_2'} \langle \chi | U_{q_2'}^\dagger U_{q_2'} | \chi \rangle_{q_2'} \right. \\ & \pm |\alpha|^2 |\beta|^2 \langle \chi | U_{q_2'}^\dagger U_{q_2'} | \chi \rangle_{q_2'} \langle \chi | \chi \rangle_{q_2'} \\ & \left. + |\beta|^4 \left| \langle \chi | U_{q_2'}^\dagger U_{q_2'} | \chi \rangle_{q_2'} \right|^2 \right) p_{1,2}, \end{aligned} \quad (\text{A22})$$

for $i = 1$ and $i = 2$, where $p_{1,2} = p_1$ or p_2 .

Using $UU^\dagger = U^\dagger U = \mathbb{1}$ and $\langle \chi | \chi \rangle = 1$, then this simplifies to

$$(|\alpha|^4 \pm 2|\alpha|^2 |\beta|^2 + |\beta|^4) p_{1,2} = (|\alpha|^2 \pm |\beta|^2)^2 p_{1,2}. \quad (\text{A23})$$

For $i = 1$ (+ case), the expression further simplifies to $p_1 = F_w$ because $|\alpha|^2 + |\beta|^2 = 1$, by the normalisation condition of quantum states.

Proceeding in a similar vein, for $i = 3$ and $i = 4$:

$$\begin{aligned} & p_{3,4} \langle \omega |_{q_2, q_2'} (\sigma_{3,4})_{q_2, q_2'} | \omega \rangle_{q_2, q_2'} \\ = & |\alpha|^4 \left| \langle \chi | U_{q_2'}^\dagger | \chi \rangle_{q_2'} \right|^2 \pm |\alpha|^2 |\beta|^2 \langle \chi | U_{q_2'} | \chi \rangle_{q_2'} \\ & \pm |\alpha|^2 |\beta|^2 \left(\langle \chi |_{q_2'} U_{q_2'}^\dagger | \chi \rangle_{q_2'} \right)^2 + |\beta|^4 \left(\langle \chi |_{q_2'} U_{q_2'}^\dagger | \chi \rangle_{q_2'} \right)^2. \end{aligned} \quad (\text{A24})$$

Inserting these terms back into equation (A21), we find that

$$F_{\text{out}} = F_w + \frac{1 - F_w}{3} \left((|\alpha|^2 - |\beta|^2)^2 + 2|\alpha|^4 \left| \langle \chi | U | \chi \rangle_{q_2'} \right|^2 + 2|\beta|^4 \left(\langle \chi | U^\dagger | \chi \rangle_{q_2'} \right)^2 \right), \quad (\text{A25})$$

which is equivalent to equation (8) in the main text. Unlike the expression for 1TP, equation (A25) depends on the exact local operations conducted on QPU B's qubits and the input state of QPU A and QPU B's processing qubits.

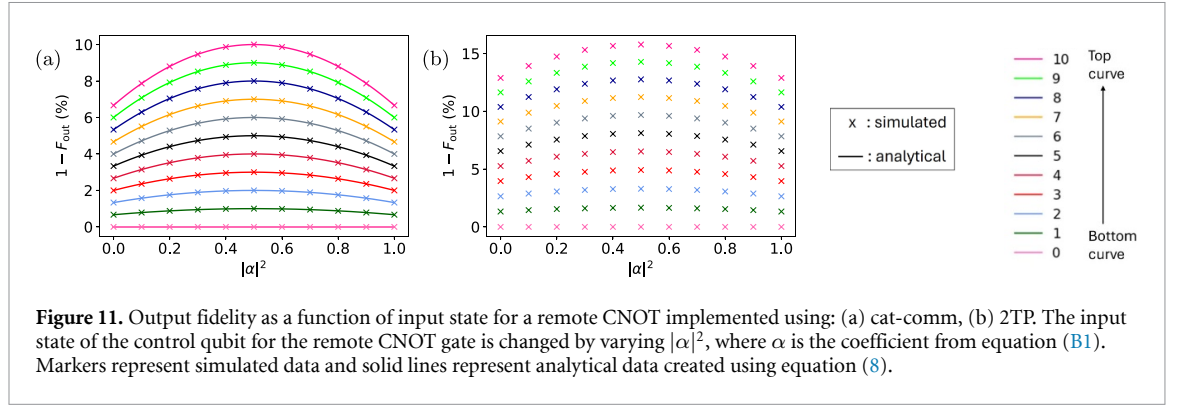


Figure 11. Output fidelity as a function of input state for a remote CNOT implemented using: (a) cat-comm, (b) 2TP. The input state of the control qubit for the remote CNOT gate is changed by varying $|\alpha|^2$, where α is the coefficient from equation (B1). Markers represent simulated data and solid lines represent analytical data created using equation (8).

Appendix B. Variation with respect to input state

Throughout section 4.2 in the main text, we assume that the input state to the remote CNOT gate considered is given by equation (11). Here, we demonstrate that the output error after a remote CNOT gate depends on its input, discuss the implications of our choice of input state and indicate why cat-comm and 1TP differ.

For 1TP, it is clear from equation (5) that the output fidelity, and therefore output error, is unaffected by the input to the remote gate when only entanglement error is considered. However, the same cannot be said for cat-comm, 2TP, and TP-safe. Moreover, for cat-comm, some input state dependence can be seen in the analytical expression for output state fidelity given by equation (7). To fully understand the role of input state would require a non-trivial averaging over all possible quantum states for the control and target qubits, which correspond to q_2 and q'_2 respectively in figure 2. Such an averaging would most likely require a detailed analytical description of the problem and multivariate integration or Monte Carlo analysis. This is beyond the scope of the current work.

Instead, for convenience, we constrain the problem by assuming that the input state has the form:

$$|\text{input}\rangle_{q_2, q'_2} = (\alpha|0\rangle_{q_2} + \beta|1\rangle_{q_2})|0\rangle_{q'_2}, \quad (\text{B1})$$

where q_2 and q'_2 are the processing qubits depicted in figure 2; $|\text{input}\rangle_{q_2, q'_2}$ is the input state of qubits q_2 and q'_2 , prior to the remote CNOT gate taking place; and α and β are complex numbers such that $|\alpha|^2 + |\beta|^2 = 1$.

To identify the input state, with the form of equation (B1), that gives the lowest output fidelity, in figure 11, we show the output error as a function of the input state of q_2 . As noted in the main text, 2TP and TP-safe give identical results when only entanglement error is considered, and so only cat-comm and 2TP are considered in figure 11.

From figure 11, it is apparent that there is a clear maximum in the output error when $|\alpha|^2 = \frac{1}{2}$. This means that the magnitude squared of the coefficients, $\alpha = \beta = \frac{1}{\sqrt{2}}$, in equation (11) correspond to the highest output error of any input state of the form given by equation (B1).

The output error is independent of the relative phase between α and β . For cat-comm, the phase independence of the output error can be seen from equation (7), which has no terms that depend on the relative phases of α and β . For 2TP, we can see the phase independence of the output error by assuming β is real and re-writing equation (B1) as:

$$|\text{input}\rangle_{q_2, q'_2} = (\alpha|0\rangle_{q_2} + e^{i\phi}\sqrt{1-|\alpha|^2}|1\rangle_{q_2})|0\rangle_{q'_2}, \quad (\text{B2})$$

where $0 \leq \phi < 2\pi$ is a real number. We plot the simulated output error as a function of ϕ in figure 12. The flatness of the curves in figure 12, indicates that the results are independent of the phase, ϕ .

The independence of the output error to phase terms in the input, indicates that any input state with $|\alpha|^2 = \frac{1}{2}$ can yield the maximum output error for input states with the form given by equation (B1). We choose $\alpha = \beta = \frac{1}{\sqrt{2}}$ for convenience, recovering equation (11).

To get some insight into the ramifications of our choice of input state, we consider the slightly more general case where only entanglement error is varied and the input state of q_2 and q'_2 is separable. In such a scenario, the input state of q_2 and q'_2 is given by equation (6), which can be re-written, up to a global phase, more explicitly as:

$$|\text{input}\rangle_{q_2, q'_2} = (\alpha|0\rangle_{q_2} + e^{i\phi}\sqrt{1-|\alpha|^2}|1\rangle_{q_2}) \otimes (\gamma|0\rangle_{q'_2} + e^{i\theta}\sqrt{1-|\gamma|^2}|1\rangle_{q'_2}), \quad (\text{B3})$$




Figure 12. Simulated output fidelity as a function of the relative phase, ϕ , in the input state for a remote CNOT implemented using 2TP. ϕ is defined as in equation (B2). $\alpha = \frac{1}{\sqrt{2}}$, where α is the coefficient from equation (B2).

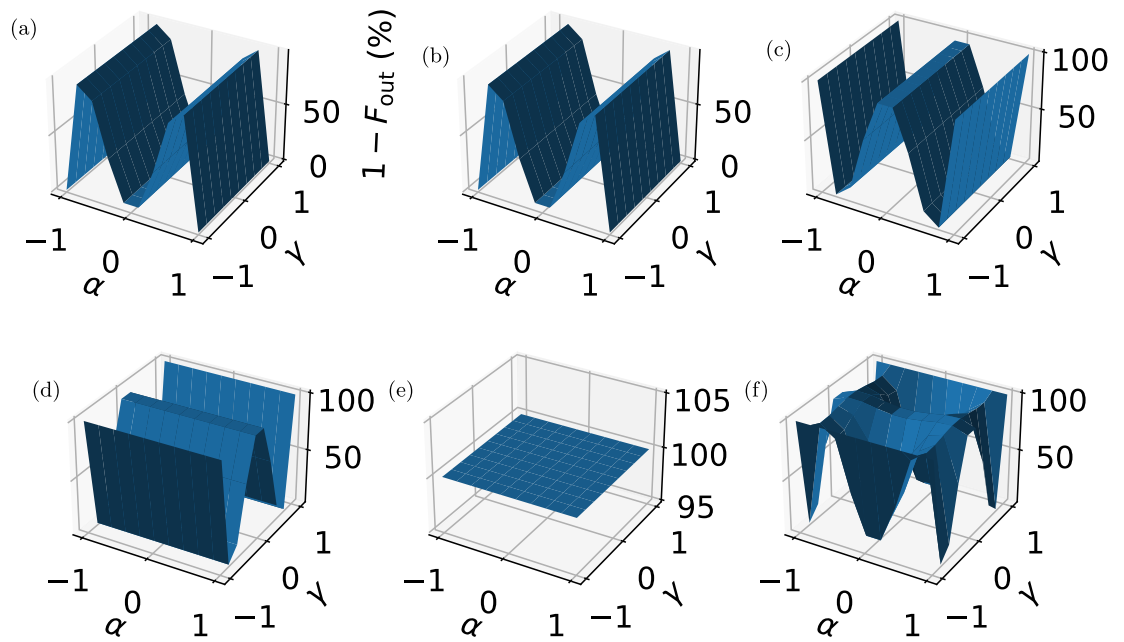


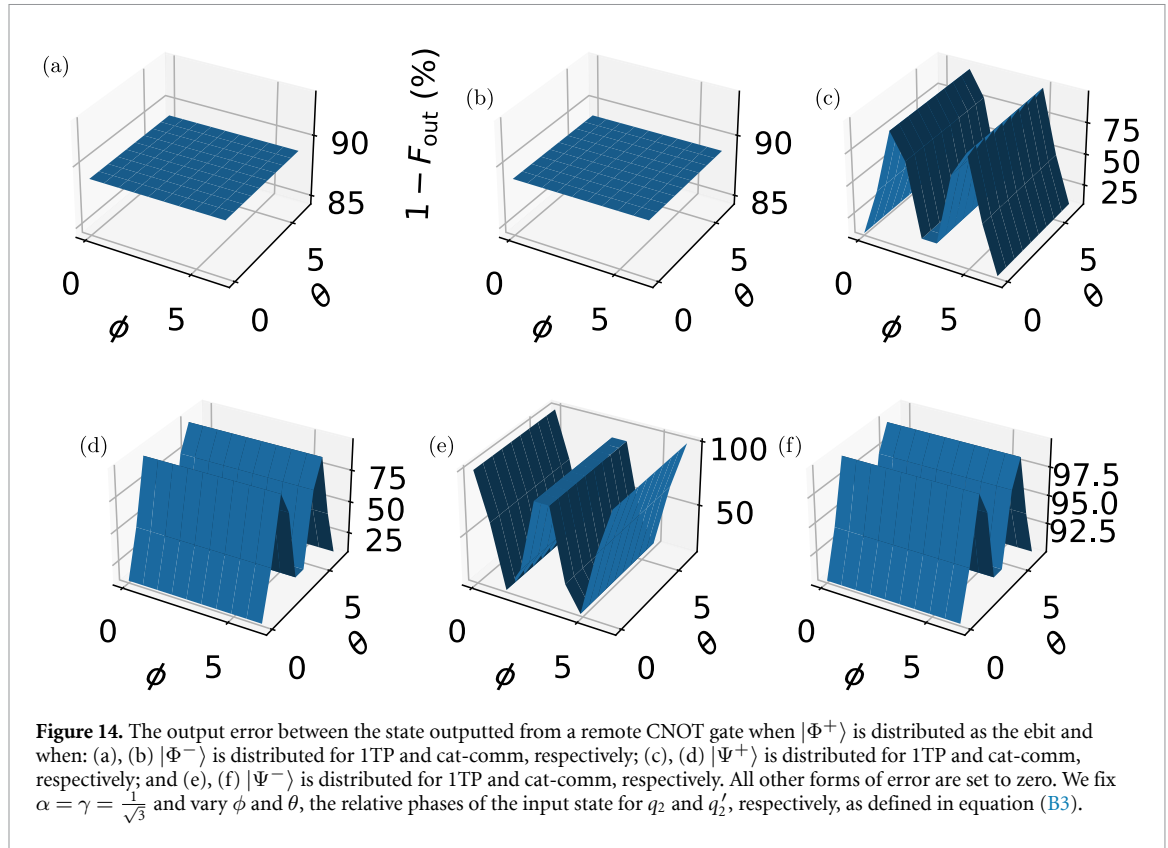
Figure 13. The output error between the state outputted from a remote CNOT gate when: $|\Phi^+\rangle$ is distributed as the ebit and when: (a), (b) $|\Phi^-\rangle$ is distributed as the ebit for 1TP and cat-comm, respectively; (c), (d) $|\Psi^+\rangle$ is distributed as the ebit for 1TP and cat-comm, respectively; and when (e), (f) $|\Psi^-\rangle$ is distributed as the ebit for 1TP and cat-comm, respectively. All other forms of error are set to zero. α and γ are the coefficients from equation (B3).

where α, γ are real numbers with magnitudes $0 \leq |\alpha| \leq 1$ and $0 \leq |\gamma| \leq 1$, and $0 \leq \phi < 2\pi$ and $0 \leq \theta < 2\pi$ are real numbers.

We model the entanglement error by assuming that ebits are in the Werner state, see equation (2), rather than $|\Phi^+\rangle$, as discussed in section 3.2. As the Werner state is just a mixture of the Bell states, we separately consider the output error when each of the Bell states is distributed as an ebit. Figures 13 and 14 show surface plots of the output error between the state outputted from a remote CNOT gate when the ideal state, $|\Phi^+\rangle$, is distributed as an ebit and the state outputted when one of the other Bell states is erroneously distributed as an ebit for different values of α and γ , and ϕ and θ , respectively.

Several observations can be made from figures 13 and 14:

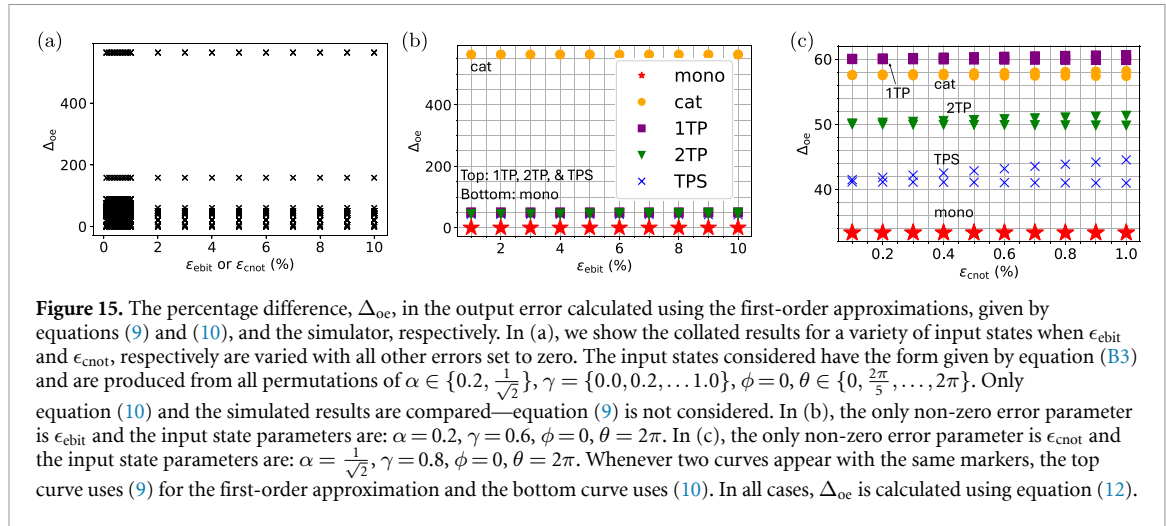
1. Equation (11) is at or near a local maximum in output error for cat-comm. Figures 13(b) and (f), which show the output error with varying α and γ for cat-comm, have maxima at approximately $\alpha = \frac{1}{\sqrt{2}} \approx 0.707$, while figure 13(d), which is also for cat-comm, is constant with respect to α but has maxima at $\gamma = \pm 1$, corresponding to the state $|0\rangle$. Therefore, although the choice of equation (11) is somewhat arbitrary, it does correspond to an especially high entanglement error. That said, equation (11) cannot be said to be a true upper bound on the entanglement error without knowing the global maximum.



2. Cat-comm and 1TP differ due to discrepancies in the circuit output produced by the $|\Psi^\pm\rangle$ terms only. The $|\Phi^\pm\rangle$ terms yield identical outputs for cat-comm and 1TP.
3. Cat-comm is constant with respect to ϕ but varies with respect to θ . This can be seen from figures 14(b), (d) and (f).
4. In the presence of entanglement error only, 1TP gives a constant output error for any input state, but cat-comm does not. This could be seen from equations (5) and (7), but the reason behind it is made clearer by figures 13 and 14. For 1TP, the different output errors resemble out of phase oscillations. When the output error caused by the distribution of one Bell state is low, it is high for a different Bell state. For example, figure 13(a) has maxima at the points where figure 13(c) has minima and slopes up at the α values where figure 13(c) slopes down. The only other term, shown in figure 13(e) is constant with respect to α and γ and so does not disrupt this trend. Similar observations can be made about phase by comparing figures 14(a), (c) and (e). By contrast, for cat-comm the output errors caused by the distribution of each Bell state are often rotated by $\frac{\pi}{2}$ radians relative to each other, as in figures 13(b), (d) and 14(d), (f). This means that the different error terms do not balance out in the same way as they do for 1TP. This accounts for the difference between cat-comm and 1TP.

Appendix C. Verification that results from main text are robust to input state

While observation 1 from appendix B indicates that equation (11) corresponds to an input state of particular interest, it is also clear from appendix B that cat-comm, 2TP, and, by extension, TP-safe will vary with respect to input state. As such, it is important to understand if certain claims made in the main text are robust to input state variation. In particular, it is not clear from the figures shown in the main text alone that observation (i) from section 4.2.1 and any of the observations made in section 4.2.3 hold for arbitrary input states. Although a true proof is not possible without more detailed input state analysis, here we demonstrate that the claims from the main text hold for a variety of input states. In appendix C.1, we demonstrate that observation (i) from section 4.2.1 is robust to variation in input state and make some further comments on the impact of input state variation on the relative difference between first-order and simulated output errors values. In appendix C.2, we verify the observations about the relative impact of different error types made in section 4.2.3 of the main text. Finally, in section C.3, we make some further comments on the observations made about the relative impact of different remote gate schemes in section 4.2.2 of the main text.



C.1. Verification of observations from section 4.2.1

To verify observation (i) from section 4.2.1 and corroborate comments made about the input state dependence of observation (iv), we consider the percentage difference, Δ_{oe} , calculated using equation (12), between the output error computed using the first-order approximation and the simulator, respectively, for various different input states. This is shown in figure 15. Figure 15(a) shows the collated results for a variety of input states when ϵ_{ebit} and ϵ_{cnot} , respectively, are varied with all other errors set to zero. The input states considered all have the form given by equation (B3) and parameter values produced using a permutation of the parameters $\alpha \in \{0.2, \frac{1}{\sqrt{2}}\}$, $\gamma \in \{0.0, 0.2, \dots, 1.0\}$, $\phi = 0$, $\theta \in \{0, \frac{2\pi}{5}, \dots, 2\pi\}$. We average over all of the input states that can be produced in this way. In figure 15(b), the input state $|\text{input}\rangle_{q_2, q'_2} = (0.2|0\rangle_{q_2} + 0.980|1\rangle_{q_2}) \otimes (0.6|0\rangle_{q'_2} + 0.8|1\rangle_{q'_2})$ ⁵ is considered in the presence of a varying non-zero ϵ_{ebit} with all other errors set to zero. Similarly, in figure 15(c), $|\text{input}\rangle_{q_2, q'_2} = \frac{1}{\sqrt{2}}(|0\rangle_{q_2} + |1\rangle_{q_2}) \otimes (0.8|0\rangle_{q'_2} + 0.600|1\rangle_{q'_2})$ ⁵ is considered in the presence of a varying non-zero ϵ_{cnot} with all other errors set to zero.

Figure 15(a) corroborates observation (i) from section 4.2.1, which states that the first order approximations loosely upper bound the simulated output error. All percentage differences, Δ_{oe} , in figure 15(a) are non-negative and so for all of the input states considered, the approximate results are greater than or equal to the exact results. Figures 15(b) and (c) differ both qualitatively and quantitatively from figures 3(a) and (c), with a different ordering of the Δ_{oe} values for the various remote gate schemes and different quantitative values of Δ_{oe} . This indicates that the agreement between the output error predicted by the first-order approximations and the simulation are highly input state dependent, as suggested in observation (iv) from section 4.2.1.

C.2. Verification of all observations from section 4.2.3

The next observations to verify are observations (i) and (ii) from section 4.2.3. Both observations are centred around the relative impact of the different error types introduced in section 3.2 with entanglement error in the state-of-the-art and distilled ranges, respectively. In figure 16, we show the output error when one non-zero error parameter is varied while all other error parameters are fixed at zero. For entanglement error, the state-of-the-art range of values is considered. Much like figure 5 from the main text, in each subfigure, we show the output errors generated by three different error types for a given remote gate scheme, but here, we average each data point over the output errors obtained at a given error parameter value for a variety of input states. Again, the input states considered have the form given by equation (B3) and are produced from all permutations of $\alpha \in \{0.2, \frac{1}{\sqrt{2}}\}$, $\gamma \in \{0.0, 0.2, \dots, 1.0\}$, $\phi = 0$, $\theta \in \{0, \frac{2\pi}{5}, \dots, 2\pi\}$. Averaging is done by taking the mean, in figures 16(a), (c), (e), (g), (i) and the median, in figures 16(b), (d), (f), (h) and (i). When the mean is used, the error bars show the standard deviation and when the median is used, the error bars show the interquartile range. Figure 17 shows the same thing but with entanglement error varied over the distilled range.

Figures 16 and 17 indicate that observations (i) and (ii) from section 4.2.3 do indeed hold for a variety of input states. The output error caused by each error type in figure 16, in which the state-of-the-art parameters are used, is consistently smallest for memory depolarisation. Local gate error yields a larger output error but

⁵ All rounded coefficients are displayed here to three significant figures but floating point precision is used for the generation of data.

⁵ All rounded coefficients are displayed here to three significant figures but floating point precision is used for the generation of data.

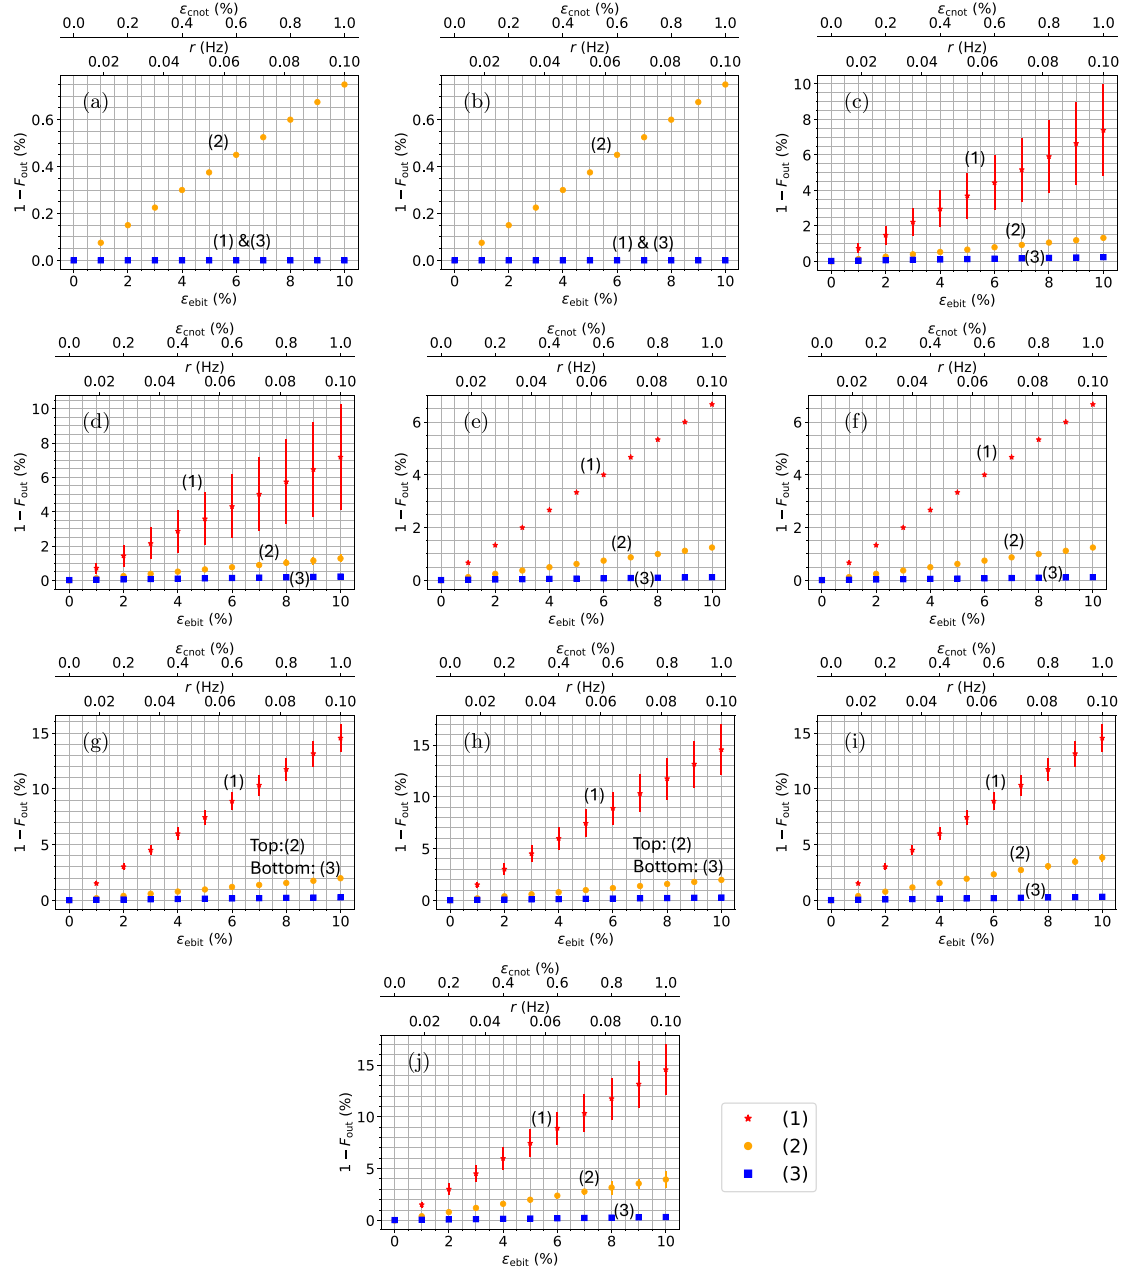


Figure 16. The output error as a function of (1) ϵ_{ebit} within the state-of-the-art range, (2) ϵ_{cnot} , (3) r , for a single CNOT gate implemented using: (a), (b) a monolithic processor; (c), (d) cat-comm; (e), (f) 1TP; (g), (h) 2TP; and (i), (j) TP-safe. The results are averaged over a variety of input states with the form given by equation (B3) and the parameters varied over all permutations of $\alpha \in \{0.2, \frac{1}{\sqrt{2}}\}$, $\gamma = \{0.0, 0.2, \dots, 1.0\}$, $\phi = 0$, $\theta \in \{0, \frac{2\pi}{5}, \dots, 2\pi\}$. The average used is the mean for (a), (c), (e), (g) and (i), with error bars indicating the standard deviation, and the median for (b), (d), (f), (h) and (j) remote gates, with error bars indicating the interquartile range. For curves (1)–(3) on each figure, the non-varied error parameters are set to zero. Due to the low standard error previously observed when averaging over simulation runs, only one simulation run is used for each input state.

its impact is in turn dominated by entanglement error. Again memory depolarisation is relatively negligible. These findings hold true for all remote gate schemes and are true for both the mean and median over input state with neither the standard deviation nor the interquartile range of any data points overlapping—which would have indicated uncertainty in the relative impact of the different error types. All of this is consistent with observation (i)

Similarly, in figure 17, in which the distilled range of entanglement error is used, the output error is smallest when memory depolarisation is non-zero, greater for entanglement error, and greatest for local gate error. This is consistent with observation (ii) from section 4.2.3.

The robustness of observations (i) and (ii) from section 4.2.3 can be further verified by an exhaustive search of plots showing the non-averaged output errors for each input state individually. We perform

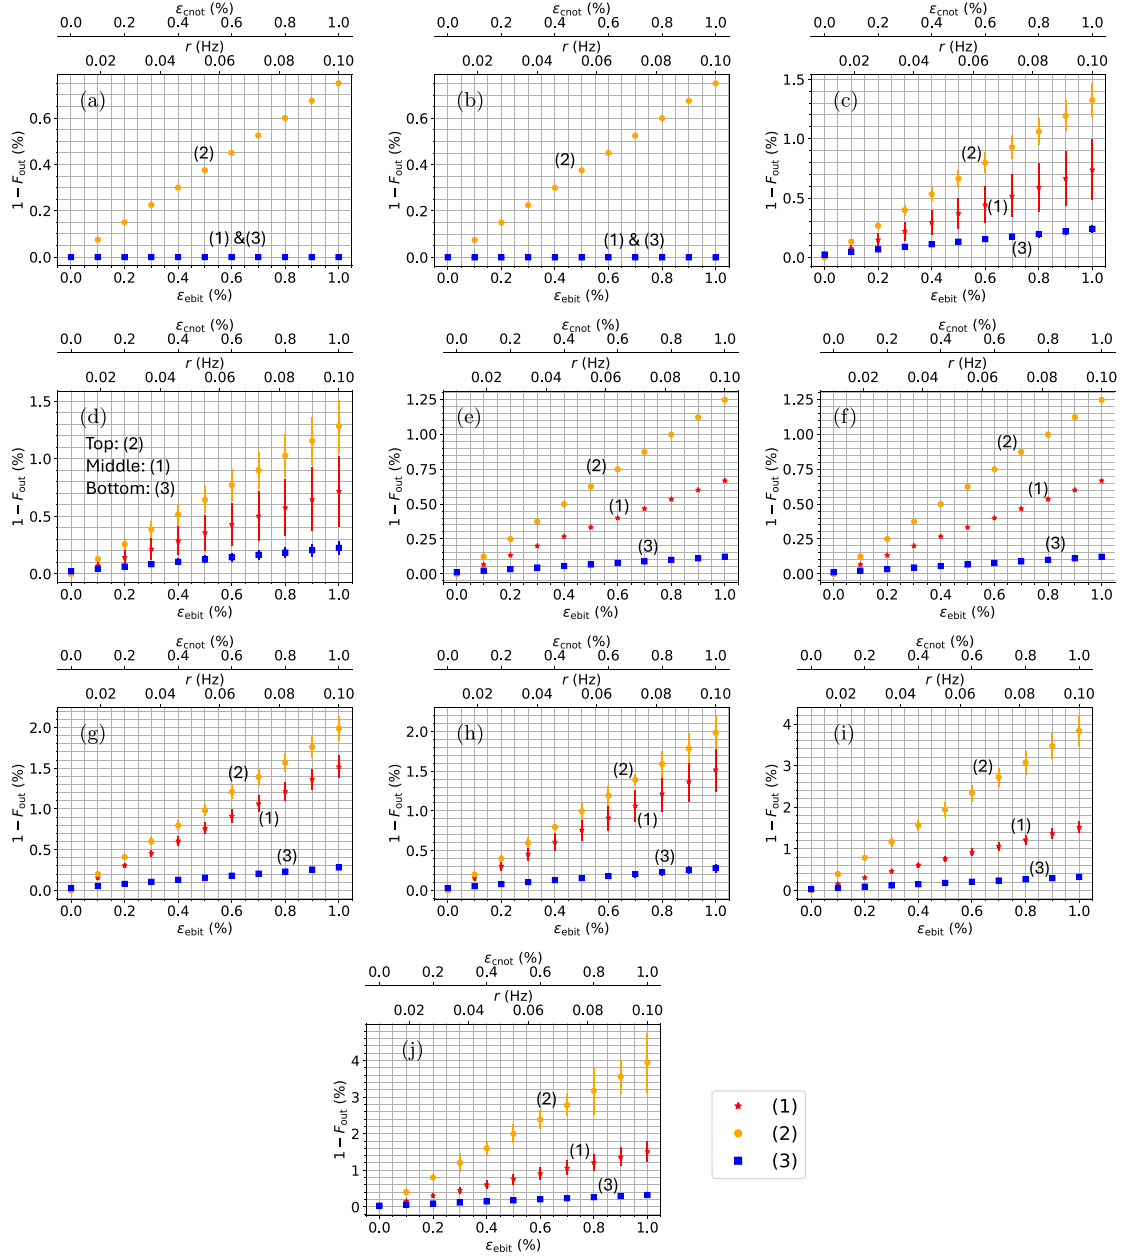


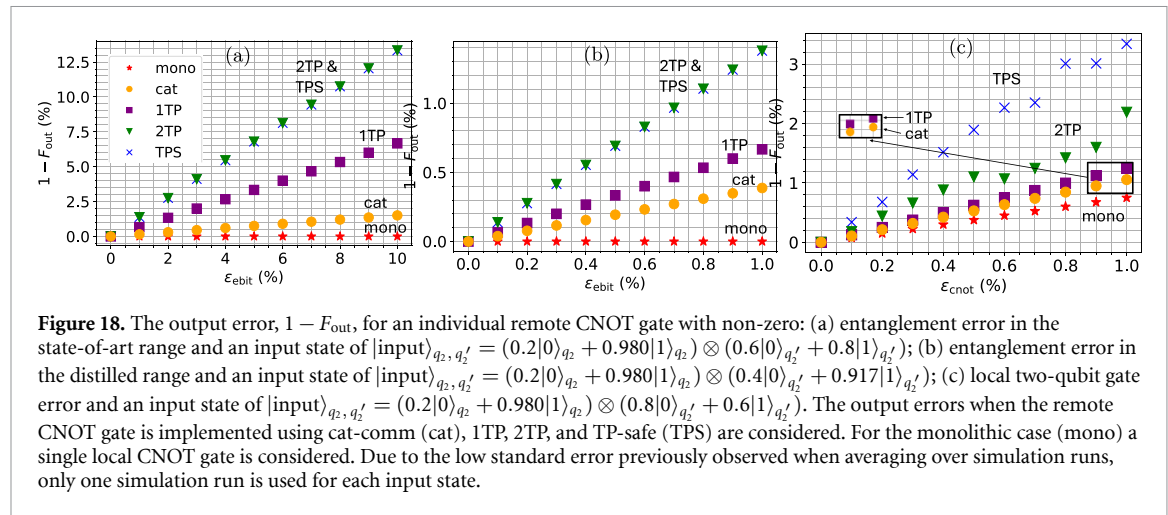
Figure 17. The output error as a function of (1) ϵ_{ebit} within the distilled range, (2) ϵ_{cnot} , (3) r , for a single CNOT gate implemented using: (a), (b) a monolithic processor; (c), (d) cat-comm; (e), (f) 1TP; (g), (h) 2TP; and (i), (j) TP-safe. The results are averaged over a variety of input states with the form given by equation (B3) and the parameters varied over all permutations of $\alpha \in \{0.2, \frac{1}{\sqrt{2}}\}$, $\gamma = \{0.0, 0.2, \dots, 1.0\}$, $\phi = 0$, $\theta \in \{0, \frac{2\pi}{5}, \dots, 2\pi\}$. The average used is the mean for (a), (c), (e), (g) and (i), with error bars indicating the standard deviation, and the median for (b), (d), (f), (h) and (j), with error bars indicating the interquartile range. For curves (1)–(3) on each figure, the non-varied error parameters are set to zero. Due to the low standard error previously observed when averaging over simulation runs, only one simulation run is used for each input state.

such a search and find that the ordering seen in observations (i) and (ii) is seen consistently for all of input states with the form of equation (B3) and parameters $\alpha \in \{0.2, \frac{1}{\sqrt{2}}\}$, $\gamma = \{0.0, 0.2, \dots, 1.0\}$, $\phi = 0$, $\theta \in \{0, \frac{2\pi}{5}, \dots, 2\pi\}$. The relevant plots can be found in the supplementary information [53].

C.3. Comments on section 4.2.2

In section 4.2.2 from the main text, we note that the ordering of remote gate schemes for the input state given by equation (11), is, ordered from lowest to highest by their impact on output error: 1TP, cat-comm, 2TP, and then TP-safe. 2TP and TP-safe are equally damaging to the output when only entanglement error is considered. Here, we confirm that the relative positions of 1TP and cat-comm are specific to input state, as alluded to in the main text and appendix B.

Figure 18 shows a comparison of the output error caused by each of the different remote gate schemes shown in figure 2. Unlike in the main text, the results for variety of different input states are shown. In



figures 18(a) and (b), the entanglement error is varied over the state-of-the-art and distilled ranges, respectively, and all other errors are set to zero. In figure 18(c), the local two-qubit gate error is varied and all other errors are set to zero.

In figures 18(a)–(c), 1TP has a larger output error than cat-comm, in contrast to what is seen and discussed in section 4.2.2 from the main text, in which a different input state, given by equation (11), is used. This is indicative of how input state dependent the relative ordering of the schemes can be. Interestingly, when we consider only a non-zero memory depolarisation, we have not yet found an input state for which 1TP does not outperform cat-comm. The most probable reason for this is that cat-comm has a slightly higher latency than 1TP. 1TP communicates the results of both the measurements that occur in the scheme at the same point in the circuit. This means that both measurements can be combined into a single classical message at the same time. By contrast, for cat-comm, the measurements are split between the cat-entanglement and cat-disentanglement protocols, which must occur sequentially. This means that the measurement results must be communicated between the QPUs in separate messages at different times and so the portion of the latency attributable to classical messaging is doubled. The remaining remote gate schemes retain the order quoted in the main text.

ORCID iDs

K Campbell  <https://orcid.org/0009-0000-5750-2069>

A Lawey  <https://orcid.org/0000-0003-3571-4110>

M Razavi  <https://orcid.org/0000-0003-4172-2125>

References

- [1] Bova F, Goldfarb A and Melko R G 2021 Commercial applications of quantum computing *EPJ Quantum Technol.* **8** 2
- [2] Zalka C 1998 Simulating quantum systems on a quantum computer *Proc. R. Soc. A* **454** 313–22
- [3] Shor P 1997 Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer *SIAM J. Comput.* **26** 1484
- [4] Grover L 1996 A fast quantum-mechanical algorithm for database search *Proc. 28th Annual ACM Symp. on Theory of Computing* pp 212–9
- [5] Sarovar M, Proctor T, Rudinger K, Young K, Nielsen E and Blume-Kohout R 2020 Detecting crosstalk errors in quantum information processors *Quantum* **4** 321
- [6] Roffe J 2019 Quantum error correction: an introductory guide *Contemp. Phys.* **60** 226
- [7] Landauer R 1995 Is quantum mechanics useful? *Phil. Trans. Phys. Sci. Eng.* **353** 367
- [8] Chuang I L, Laflamme R, Shor P W and Zurek W H 1995 Quantum computers, factoring and decoherence *Science* **270** 1633–5
- [9] Shor P W 1995 Scheme for reducing decoherence in quantum computer memory *Phys. Rev. A* **52** R2493
- [10] Liu J, Hann C T and Jiang L 2023 Data centers with quantum random access memory and quantum networks *Phys. Rev. A* **108** 032610
- [11] Baker J, Duckering C, Hoover A and Chong F 2020 Time-sliced quantum circuit partitioning for modular architectures *Proc. 17th ACM Int. Conf. on Computing Frontiers* pp 98–107
- [12] Wu A, Zhang H, Li G, Shabani A, Xie Y and Ding Y 2022 Autocomm: a framework for enabling efficient communication in distributed quantum programs *55th IEEE/ACM Int. Symp. on Microarchitecture (MICRO)* pp 1027–41
- [13] Ferrari D, Cacciapuoti A, Amoretti M and Calet M 2021 Compiler design for distributed quantum computing *IEEE Trans. Quantum Eng.* **2** 1
- [14] Wu A, Ding Y and Li A 2023 Qucomm: optimizing collective communication for distributed quantum computing *Proc. 56th Annual IEEE/ACM Int. Symp. on Microarchitecture (MICRO'23)* (Association for Computing Machinery) pp 479–93

- [15] Ferrari D, Carretta S and Amoretti M 2023 A modular quantum compilation framework for distributed quantum computing *IEEE Trans. Quantum Eng.* **4** 1–13
- [16] Bennett C H, Brassard G, Crépeau C, Jozsa R, Peres A and Wootters W K 1993 Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels *Phys. Rev. Lett.* **70** 1895
- [17] Bouwmeester D, Pan J-W, Mattle K, Eibl M, Weinfurter H and Zeilinger A 1997 Experimental quantum teleportation *Nature* **390** 575
- [18] Eisert J, Jacobs K, Papadopoulos P and Plenio M 2000 Optimal local implementation of nonlocal quantum gates *Phys. Rev. A* **62** 052317
- [19] Yimsiriwattana A and Lomonaco S J Jr 2004 Generalized GHZ states and distributed quantum computing (arXiv:0402148)
- [20] Yimsiriwattana A and Lomonaco S J Jr 2004 Distributed quantum computing: a distributed Shor algorithm (arXiv:0403146v2)
- [21] Dadkhah D, Zomorodi M, Hosseini S E, Plawiak P and Zhou X 2022 Reordering and partitioning of distributed quantum circuits *IEEE Access* **10** 70329
- [22] Daei O, Navi K and Zomorodi-Moghadam M 2020 Optimized quantum circuit partitioning *Int. J. Theor. Phys.* **59** 3804
- [23] Andres-Martinez P and Heunen C 2019 Automated distribution of quantum circuits via hypergraph partitioning *Phys. Rev. A* **100** 032308
- [24] Sundaram R, Gupta H and Ramakrishnan R 2021 Efficient distribution of quantum circuits *35th Int. Symp. on Distributed Computing* vol 209 pp 41:1–41:20
- [25] Ghodsollahee I, Davarzani Z, Zomorodi M, Plawiak P, Houshmand M and Houshmand M 2021 Connectivity matrix model of quantum circuits and its application to distributed quantum circuit optimization *Quantum Inf. Process.* **20** 235
- [26] Zomorodi-Moghadam M, Houshmand M and Houshmand M 2018 Optimizing teleportation cost in distributed quantum circuits *Int. J. Theor. Phys.* **50** 848
- [27] Davis M G, Chung J, Englund D and Kettimuthu R 2023 Towards distributed quantum computing by qubit and gate graph partitioning techniques *IEEE Int. Conf. on Quantum Computing and Engineering (QCE)* vol 1 pp 161–7
- [28] Jozsa R 1994 Fidelity for mixed quantum states *J. Mod. Opt.* **41** 2315
- [29] Cirac J I, Ekert A K, Huelga S F and Macchiavello C 1999 Distributed quantum computation over noisy channels *Phys. Rev. A* **59** 4249
- [30] Kitaev A Y 1995 Quantum measurements and the Abelian stabilizer problem (arXiv:quant-ph/9511026)
- [31] Nielsen M and Chuang I 2010 *Quantum Computation and Quantum Information* 10th edn (Cambridge University Press)
- [32] Ezzell N, Pokharel B, Tewala L, Quiroz G and Lidar D A 2023 Dynamical decoupling for superconducting qubits: a performance survey *Phys. Rev. Appl.* **20** 064027
- [33] Neumann N M P, van Houten R and Attema T 2020 Imperfect distributed quantum phase estimation *Computational Science – ICCS 2020* (Springer) pp 605–15
- [34] Bennett C H, Brassard G, Popescu S, Schumacher B, Smolin J A and Wootters W K 1996 Purification of noisy entanglement and faithful teleportation via noisy channels *Phys. Rev. Lett.* **76** 722–5
- [35] Munro W J, Harrison K A, Stephens A M, Devitt S J and Nemoto K 2010 From quantum multiplexing to high-performance quantum networking *Nat. Photon.* **4** 792
- [36] Li Y and Benjamin S C 2016 Hierarchical surface code for network quantum computing with modules of arbitrary size *Phys. Rev. A* **94** 042303
- [37] Briegel H, Dür W, Cirac J and Zoller P 1998 Quantum repeaters: the role of imperfect local operations in quantum communication *Phys. Rev. Lett.* **81** 5932
- [38] Razavi M 2023 Fiber-based quantum repeaters *Photonic Quantum Technologies* (Wiley) ch 24, pp 675–91
- [39] Deutsch D, Ekert A, Jozsa R, Macchiavello C, Popescu S and Sanpera A 1996 Quantum privacy amplification and the security of quantum cryptography over noisy channels *Phys. Rev. Lett.* **77** 2818
- [40] Pan J-W, Simon C, Brukner C and Zeilinger A 2001 Entanglement purification for quantum communication *Nature* **410** 1067
- [41] Cross A W, Bishop L S, Smolin J A and Gambetta J M 2017 Open quantum assembly language (arXiv:1707.03429 [quant-ph])
- [42] Flannigan S, Pearson N, Low G H, Buyskikh A, Bloch I, Zoller P, Troyer M and Daley A J 2022 Propagation of errors and quantitative quantum simulation with quantum advantage *Quantum Sci. Technol.* **7** 045025
- [43] Coopmans T *et al* 2021 Netsquid, a discrete-event simulation platform for quantum networks *Commun. Phys.* **4** 164
- [44] Woehr J 2022 nuqasm2 (available at: <https://github.com/jwoehr/nuqasm2>)
- [45] Stephenson L J, Nadlinger D P, Nichol B C, An S, Drmota P, Ballance T G, Thirumalai K, Goodwin J F, Lucas D M and Ballance C J 2020 High-rate, high-fidelity entanglement of qubits across an elementary quantum network *Phys. Rev. Lett.* **124** 110501
- [46] IonQ staff 2023 Ionq aria: practical performance (available at: <https://ionq.com/resources/ionq-aria-practical-performance>)
- [47] Metodi T S, Thaker D D, Cross A W, Chong F T and Chuang I L 2005 A quantum logic array microarchitecture: scalable quantum data movement and computation *38th Annual IEEE/ACM Int. Symposium on Microarchitecture (MICRO'05)* (<https://doi.org/10.1109/micro.2005.9>)
- [48] Yu Z and Li Y 2022 Analysis of error propagation in quantum computers (arXiv:2209.01699 [quant-ph])
- [49] Luo M-X and Li H-R 2016 Distributed quantum computation assisted by remote toffoli gate *Cloud Computing and Security*, ed X Sun, A Liu, H-C Chao and E Bertino (Springer) pp 475–85
- [50] Nickerson N 2015 Practical fault-tolerant quantum computing *PhD Thesis* Imperial College London
- [51] Sarvaghad-Moghaddam M and Zomorodi M 2021 A general protocol for distributed quantum gates *Quantum Inf. Process.* **20** 265
- [52] Quetschlich N, Burgholzer L and Wille R 2023 MQT Bench: benchmarking software and design automation tools for quantum computing *Quantum* **7** 1062
- [53] Campbell K 2024 Supplementary information for “Quantum data centres: a simulation-based comparative noise analysis” (Zenodo) [10.5281/zenodo.14382385](https://doi.org/10.5281/zenodo.14382385)
- [54] Main D, Drmota P, Nadlinger D P, Ainley E M, Agrawal A, Nichol B C, Srinivas R, Araneda G and Lucas D M 2024 Distributed quantum computing across an optical network link (arXiv:2407.00835 [quant-ph])
- [55] Campbell K 2024 Data for quantum data centres: a simulation-based comparative noise analysis [Data set] (Zenodo) (<https://doi.org/10.5281/zenodo.13773229>)