Assessing Teleportation of Logical Qubits in a Distributed Quantum Architecture under Error Correction

John Stack
jstack@ncsu.edu
North Carolina State University
Raleigh, North Carolina, USA

Ming Wang
mwang42@ncsu.edu
North Carolina State University
Raleigh, North Carolina, USA

Frank Mueller
fmuelle@ncsu.edu
North Carolina State University
Raleigh, North Carolina, USA

Abstract

Quantum computing is facing challenges in terms of scaling to thousands of qubits and implementing quantum error correction (QEC). Scaling efforts focus on connecting multiple smaller quantum devices in a distributed manner while error correction, as a means to overcome noisy physical qubits, is being addressed by developing denser codes with protocols for logical qubits and logical quantum gates. Teleportation of quantum states becomes an important operation as it transfers states from one node to another node within a distributed device. For physical qubits, today's high quantum network noise rates prevent the teleportation of states with useful accuracy. By employing QEC, we show that logical qubits can be teleported between nodes under Surface Code and qLDPC encodings with very low logical error rates, even with network noise in near-term regimes. We use circuit-level simulations to assess physical and network noise regimes ranging from 10^{-1} to 10^{-6}. This is a wider range than typically studied in circuit level simulations and understanding the behavior of QEC codes in these regimes is necessary for achieving accurate computation.

1 Introduction

The noise level of physical qubits in quantum computing remains untenable for realizing the most promising quantum algorithms. This problem is addressed by employing QEC codes that encode a logical qubit into many physical ones and then engaging in error correction rounds to reduce noise at the logical level below the threshold of the noise of physical qubits. However, logical qubits are built over a multitude of physical qubits requiring quantum devices to scale to large number of qubits. Useful fault tolerant (FT) quantum computing requires physical qubit counts of at least 10^3 or greater, in part because classical simulation of quantum states is feasible up to 50-60 qubits but not beyond [4].

At this scale, individual devices begin to run into limitations on the number of qubits that they can store effectively [11]. This is due to, e.g., connectivity requirements, cooling requirements, optical address accuracy and limitations on the number of atoms or ions that can be trapped. Consequentially, the idea of networking individual quantum computers together to form a distributed quantum computer (DQC) has been proposed.

What constitutes a DQC is ambiguous, such a device can be realized in many different ways. A DQC may be implemented using classical post-processing and several QCs that are purely classically connected [8]. This is known as circuit knitting. Essentially, a circuit is sliced space-wise either around gates, which is called wire cutting and corresponds to non-local gates, or through gates, which is called gate cutting and corresponds to teleportation. The two (or more) circuits produced by the slicing are executed separately and their results are then combined to infer the results of the original circuit. Unfortunately, the sampling overheads to perform circuit knitting at scale increase exponentially in the number of gates or wires cut [8]. Therefore, this approach is infeasible for circuits with more than a low degree of entanglement.

Other architectures involve quantum interactions between nodes. This allows the state vector realized by the DQC to expand exponentially with respect to the total number of qubits in the DQC. These architectures are either photonic, and therefore naturally distributed, or they involve shuttling between chiplets (see [1, 16]), and/or the generation of "ebits" between entirely separate devices. Our approach in this paper is specific to the ebit case, i.e., we do not consider photonic devices.

In the shuttling case, a circuit is compiled onto a physical device such that if there is a gate between qubits that reside on different modules, one of those qubits is shuttled such that it now is on the same module as the other and can be interacted with it. In recent work [1], it was found that this can be done on the order of $2{\times}10^3$ per second with infidelities below 10^{-5}. Whilst these numbers are very promising, this is not a truly distributed architecture. This approach will eventually be limited by the same constraints as before as mentioned in [1].

A truly distributed architecture will involve some form of quantum communication between nodes. This involves the generation of an entangled resource state of $n \geq 2$ qubits between multiple nodes that allows for operations between computational qubits located on different nodes. This resource state can be generated in many ways.

One of the most well-studied methods is the $n = 2$ case called entanglement heralding [6]. This is the case we study and simulate in this paper. Here, pairs of "matter qubits" are simultaneously pulsed, temporarily putting them into an excited state where they return to their original energy levels by each releasing a photon, which is entangled with themselves. These photon pairs are then each collected and directed to a non-polarized beam splitter and are then detected. A non-polarizing beam splitter then erases "which-path" information and the photon-detectors therefore project the state of the matter qubit pair into a Bell state [6].

These Bell states can then be used as a resource to carry out either non-local gates, e.g., a CNOT between two qubits that are located on a different devices [10], or to teleport a qubit [27]. How qubits are initially allocated to nodes and how qubits on different nodes are interacted with each other is a compilation decision. Compilation that focuses on minimizing non-local operations in the case where qubits are fixed to certain nodes and only non-local gates are used to interact qubits across different nodes has been studied in [2].

Generated ebits are typically much nosier than local operations with a current maximum experimental fidelity of 97% [18]. This is up to 100 times nosier than current local operations. At this error rate it makes them difficult to use, un-corrected, in even shallow circuits [3]. In [18] a 97% fidelity ebit was used to execute several non-local controlled-Z gates with fidelity 86%. This is clearly insufficient for larger circuits. Entanglement distillation can be used to improve their fidelity [5]. But this is typically intensive in the number of ebits required (e.g., 10+ ebits per distilled ebit) and still produces ebits with infidelities. It is therefore necessary to study the effect of ebit infidelities on the execution of logical operations.

Contributions: In this paper, we realize the first full-circuit simulations of two logical operations that utilize ebits. These operations form the basis of a fault tolerant distributed quantum computer. The first operation is a transversal non-local CNOT between error corrected code blocks located on different nodes. This comprises a fault tolerant implementation of "gate teleportation". The second operation is fault tolerant quantum teleportation of an entire code block between nodes. The first operation enables operations between qubits on different nodes of a DQC, while the second enables movement of qubits between nodes of a DQC. We conduct our study both for Bivariate-bicycle codes [7], a type of quantum low-density parity check (qLDPC) code that can be efficiently realized on neutral atom computers, and the Surface Code [12].

We simulate physical and ebit noise down to the 10^{-6} level. This is far lower than other works and allows us to explore whether circuits, QEC codes and a decoder have error floors at this level. We find that they do not.

The paper is structured as follows: In Section 2, we introduce quantum error correction (QEC), decoding and the different approaches to QEC in a distributed architecture. We also we contrast our work to prior research. In Section 3, we discuss the design of the distributed architecture that we are considering. We introduce the non-local fault tolerant CNOT circuit and the fault tolerant teleportation circuit that we simulate. We also describe the noise model behind our simulations and how they are performed. In Section 4 we discuss the results of our simulations, and in Section 5 we summarize our contributions.

2 Background and Related Work

2.1 Quantum Error Correction (QEC)

Let $\mathcal{P} = \langle X, Y, Z\rangle$ denote the single-qubit Pauli group. The $n-$qubit Pauli group is then defined as $\mathcal{P}_n = \{P_1 \otimes P_2 \otimes \cdots \otimes P_n | P_i \in \mathcal{P}\}$. An $[\![n, k, d]\!]$ stabilizer code, defined by the stabilizer group $\mathcal{S} \subset \mathcal{P}_n$, encodes k qubits of logical information into a n-qubit physical qubit block and can correct up to $\lfloor (d-1)/2 \rfloor$ errors. The code space C is the common +1 eigenspace of $\mathcal{S}$, given by

$$C = \{|\psi\rangle \mid s\,|\psi\rangle = +|\psi\rangle\}. \tag{1}$$

Therefore, to construct a valid stabilizer code, $\mathcal{S}$ must be an Abelian subgroup of $\mathcal{P}_n$ and $-1 \notin \mathcal{S}$. For qubit stabilizer codes, the stabilizer group can also be represented by an $(n-k) \times 2n$ matrix,

$$H = \left[\begin{array}{cccc|cccc} s_{X_1} & s_{X_2} & \cdots & s_{X_n} & s_{Z_1} & s_{Z_2} & \cdots & s_{Z_n} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \end{array}\right] \tag{2}$$

where each row $(\boldsymbol{s}_X|\boldsymbol{s}_Z)$ corresponds a generator $s = P_1 \otimes \cdots \otimes P_n$ in $\mathcal{S}$. The entries of H are determined as follows:

$$\begin{cases} s_{X_i} = 0,\ s_{Z_i} = 0 & \text{if} \quad P_i = I \\ s_{X_i} = 1,\ s_{Z_i} = 0 & \text{if} \quad P_i = X \\ s_{X_i} = 1,\ s_{Z_i} = 1 & \text{if} \quad P_i = Y \\ s_{X_i} = 0,\ s_{Z_i} = 1 & \text{if} \quad P_i = Z \end{cases} \tag{3}$$

Calderbank-Steane-Shor (CSS) codes form a special subclass of stabilizer codes in which the stabilizer group $\mathcal{S}$ can be decomposed into two disjoint subsets containing only X-type and Z-type stabilizers. Consequently, the parity-check matrix of a CSS code takes the block-diagonal form:

$$H = \left[\begin{array}{c|c} H_X & \mathbf{0} \\ \mathbf{0} & H_Z \end{array}\right]. \tag{4}$$

Since $\mathcal{S}$ is an Abelian group, the commutativity condition requires that $H_X H_Z^T = \mathbf{0}$ must be satisfied to ensure the validity of the CSS code.

An error that exceeds a codes's correction capability can cause a logical error. That is, after error correction, the residue error acts as a logical operator and changes the logical state of the code. Such errors are undetectable as they commute with all stabilizers.

Decoding: A QEC cycle involves a syndrome extraction round, measurement of stabilizer qubits and then correction

of inferred errors (either physically, or in a Pauli frame). Each stabilizer qubit corresponds to a row of either H_X or H_Z. The result of its measurement indicate the parity of the data qubits of that row.

Using these measurements, we can therefore produce a syndrome vector $\mathbf{s} \in \mathbb{F}_2^{n-k}$. Assuming that CSS codes are used for error correction and errors are of the form $E = \bigotimes_{i=1}^{n} P_i$, represented as a binary vector $\mathbf{e} = (\mathbf{e}_X|\mathbf{e}_Z) \in \mathbb{F}_2^{2n}$, then we can infer that the errors that have occurred on the data qubits related to the syndrome $\mathbf{s}$ by the following equation:

$$\begin{aligned} \mathbf{s}_X &= H_X \odot \mathbf{e}_Z = H_X \mathbf{e}_Z \mod 2 \\ \mathbf{s}_Z &= H_Z \odot \mathbf{e}_X = H_Z \mathbf{e}_X \mod 2. \end{aligned} \tag{5}$$

Given this relation, a decoder for an $[\![n, k, d]\!]$ stabilizer code is a classical algorithm that infers a recovery operation $\hat{E} \in \mathcal{P}_n$ from measured syndromes $\mathbf{s}_X, \mathbf{s}_Z$ that removes the syndrome. The decoder solves the inverse problem, $H_{\{X,Z\}} \odot \hat{\mathbf{e}}_{\{Z,X\}} = \mathbf{s}_{\{Z,X\}}$, to estimate $\hat{\mathbf{e}}$. The same syndrome can correspond to multiple errors $\mathbf{e}$ as H is not full rank. In these cases, a decoder typically chooses the error with the lowest Hamming weight, $\hat{\mathbf{e}}$, i.e., the most likely error (highest probability).

2.2 Surface Codes

The rotated surface code (SC) is a topological QEC code defined on a square lattice of $d \times d$ data qubits (aka. plaquette) with parameters $[\![d^2, 1, d]\!]$. The stabilizer group $\mathcal{S}$ consists of X- and Z-type operators acting on each plaquette, with bulk stabilizers of weight four and edge stabilizers of weight two. For a coordinate (i, j) in the rotated lattice, the X-stabilizers $S_X^{(i,j)}$ and Z-stabilizers $S_Z^{(i,j)}$ are defined as:

$$S_X^{(i,j)} = \bigotimes_{(k,l)\in\partial(i,j)} X_{k,l}, \quad S_Z^{(i,j)} = \bigotimes_{(k,l)\in\partial(i,j)} Z_{k,l}, \tag{6}$$

where $\partial(i, j)$ denotes the qubits adjacent to plaquette (i, j). This code is also a CSS code. Logical operators are pairs of $\overline{X}$ or $\overline{Z}$ strings spanning the lattice's diagonal with minimum weight d. Under independent circuit-level Pauli noise, the code achieves a threshold $p_{\text{th}} \approx 1\%$ [12]. The rotated surface code's constant weight stabilizers and nearest-neighbor connectivity mean that it is well suited to many hardware implementations and has been featured in many physical experiments [4, 28].

2.3 qLDPC Bivariate Bicycle (BB) Codes

Bivariate bicycle (BB) codes form a class of CSS codes defined by two integers, l, m, and two bivariate polynomials: $a(x, y), b(x, y)$, over the quotient ring $\mathbb{F}_2[x, y]/(x^l - 1, y^m - 1)$ [7]. Let I_i denote the identity matrix of size i, and let S_j be the cyclic permutation matrix of size j, defined as $S_j = I_j >> 1$, where $>>$ is a cyclic shift. We can identify the variates x and y as matrices

$$x = S_l \otimes I_m, \quad y = I_l \otimes S_m. \tag{7}$$

This representation allows the polynomials $a(x, y)$ and $b(x, y)$ to be naturally expressed as $lm \times lm$ matrices A and B, respectively. Since $xy = yx$ holds due to the mixed-product property of the Kronecker product, matrix multiplication among these representations remains commutative. The parity-check matrices for BB codes are then given by

$$H_X = [A|B], \quad H_Z = [B^T|A^T]. \tag{8}$$

This construction guarantees a valid CSS code because the commutativity of A and B ensures that$H_X H_Z^T = AB + BA = 2AB = 0$ in $\mathbb{F}_2$, satisfying the CSS code commutativity condition.

BB codes main advantage over the surface code is their high encoding as a function of code distance. Encoding rate $r = \frac{k}{n+c}$ is the number of logical qubits encoded in a code over the number of physical n and stabilizer qubits c used [7]. The surface code has encoding rate $r \approx \frac{1}{2d^2} = \frac{1}{2n}$. But BB codes have far better encoding rates. The codes considered in this paper have encoding rates up to 10 times higher than similar distance surface codes. Additionally, they have low-weight stabilizers that can be fixed to degree 6 [7] (unlike some other qLDPC codes) and can be efficiently implemented on neutral atom computers [26].

2.4 Transversal Gates

A transversal gate in QEC is a logical operation that can be implemented by applying the corresponding physical gate to each data qubit in a code block. Logical operations are those that preserve distance and the code space. For a code with n-qubit logical states, applying a physical operation U to each data qubit acts as: $U^{\otimes n}$ (single-block), or for paired interactions such as a CNOT, $U_{\text{CNOT}}^{\otimes n}$ between blocks. Transversal operations do not propagate errors within a code block and are fault tolerant, as they preserve the code space.

By the Eastin-Knill theorem [9], no QEC code can perform *universal* fault tolerant quantum computation using only transversal gates. CSS codes, however, have sufficient transversal gates and measurement to carry out non-local CNOTs and teleportation. All CSS codes have transversal CNOT, Hadamard, Pauli gates and logical measurement in Z or X basis. Implementation of a transversal Hadamard requires some care as it requires swapping X and Z operators. T-gates, however, are known to be non-transversal for CSS codes, yet are required for universal quantum computing and realized either by distillation factories or code switching. We restrict our work to transversal gates.

2.5 Fault Tolerant (FT) DQC

There have been many simulations and analyses of FT DQC. They typically fall consider three architectures: (1) small network nodes that perform syndrome extraction operations using GHZ states [19, 24], (2) codes extended across multiple nodes where the minority of syndrome extraction operations

require a non-local 2-qubit gate [17, 20, 25], and (3) architectures where syndrome extraction operations use only local gates, but non-local transversal/lattice surgery/teleportation operations are used to carry out computation between code blocks [15, 23].

Type 1 architectures have the benefit of using only small nodes. There are minimal connectivity issues. Optical switches ensure all-to-all connectivity. But given that almost all operations are mediated by a noisy GHZ state, which is often a product of noisy ebit Bell states, it is unlikely that high performance will be achieved. For example, in [24], even with local physical error rates approaching 10^{-5} and ebit noise on the order of 10^{-3}, the logical error rate is either higher than the physical error rate or is not significantly lowered below the threshold (physical error rate).

In type 2 architectures we see better performance. In the extended code architecture of [20], in contrast, we see error suppressive behavior consistently below threshold with ebit noise a fixed ratio of 14 times greater than physical noise. Their architecture consists of multiple surface code patches, where each edge of a patch is interlinked with another patch via non-local CNOTs. Depending on the size of the patches, non-local gates are only used on a minority of stabilizer qubits.

Other type 2 architecture work considers Floquet codes distributed over multiple devices [25]. Floquet codes have weight-2 stabilizer checks, which makes the number of non-local resources needed when distributed over multiple nodes minimal since only *some* stabilizers involve non local operations, and those only need a single ebit. In [25] the authors simulate a distributed version of the $[\![576, 10, 12]\!]$ H64-f3 code that uses 576 data qubits distributed over 32 devices. They ran the code for 64 rounds of extraction and used that to calculate the per-round logical error rate. They found logical noise below the threshold for physical error rate less than 10^{-4} and ebit noise better than 10^{-2}. At a physical error rate of 5×10^{-5} and ebit error rate of 5×10^{-3}, their logical error rates are 100x better than the physical error rate.

Research exploring architecture type 3, where entire code blocks, separated across devices, are used for computation rather than memory experiments, is more limited. In [15], the concept of generating ebit pairs between nodes and a central routing card, enabling operations between nodes, is explored. These ebit pairs are used to carry out distributed lattice surgery between surface code blocks on different nodes. Although the authors only simulate the merging of code blocks via lattice surgery, this is equivalent to $Z \times Z$ measurement. The fidelity with which this process can be completed gives a good indication of how one can realize non-local gates or logical teleportation as lattice surgery merge is the primary step in lattice surgery-based approaches to these. Their work focuses on the surface code and does not explicitly explore circuit level simulations of full teleportation or non-local CNOT circuits while our work does.

In [23], the authors perform circuit level simulations of extended surface code patches, specifically those obtained after a lattice surgery merge operation, where some parts of the patch are connected via operations that are especially noisy either due to the construction of the quantum computer or due to the use of ebits and gate teleportation. They carefully place the "seam", the barrier across which operations are noisier, and construct their circuits such to avoid hook errors. Just as in [20], they find noise across the seam 10 times larger than physical error rate to be easily be tolerated with only a slight reduction in threshold. Whilst the circuit analysis in our work is less sophisticated than [23], we do simulate entire FT operations from end to end, which goes beyond their work. In addition, we simulate a lower noise regime (up to 10^{-6} vs. 10^{-3}) than their work, a larger variety of seam (ebit in our case) noise rates and BB codes as well as SCs.

Some other work explores similar topics to this paper but does not fall into the above categorization. In [22], logical teleportation of the $[\![7, 1, 3]\!]$ Steane code is performed on a trapped ion device in two different ways, (1) via transversal gates and (2) lattice surgery. The former requires three code blocks: the first being the state to teleport, the other two being a logical Bell state. The latter requires only two code blocks: the state to teleport and a logical $|+\rangle$ state. They find the transversal protocol to teleport with lower error rates (fidelities of 99.97% vs 99.5%).

In [5], the authors explores the utility of qLDPC codes to effectively distill ebits. For a $[\![n, k, d]\!]$ code, the process takes in n noisy ebits and produces k distilled ebits. For example, given 1020 ebit pairs of noise 0.04, they can distill these to 136 ebit pairs of noise 10^{-4}. However, their model does not explore the use of these ebits for error corrected computation in a circuit built on logical qubits. In contrast, we perform end-to-end circuit level simulations of non-local CNOTs and teleportation. Nonetheless, their work complements ours as follows. When considering options for implementing quantum hardware, a manufacturer may need to teleport across nodes of a distributed device a $[\![n, k, d]\!]$ code block at a p_L logical error rate using ebits with an error that is too high. In this context and by using [5]'s techniques, a distillation process can be realized that produces n less noisy ebits and therefore uses a $[\![n_2, n, d_2]\!]$ code, where n_2 is the number of ebits to be distilled and d_2 is dependent on the desired distilled fidelity.

Our work makes contributions beyond the discussed prior work in two ways. We fully simulate at the circuit level across two DQC devices for two primitive FT operations, (1) non-local FT CNOTs and (2) FT teleportation. We assess the viability of the two techniques for both the surface code and different qLDPC codes, specifically Bivariate Bicycle codes. Using advanced HPC simulation techniques, we explore a noise floor significantly lower than that in most other papers, namely at 10^{-6}. We find that codes can tolerate ebit noise up

100 times worse than physical local operations for execution of these primitive operations with logical error rate far below physical error rate.

3 Design: FT logical teleportation

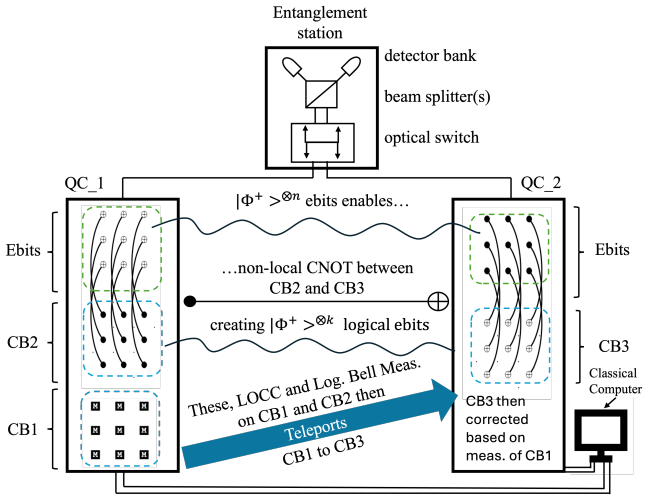

Figure 1. Possible DQC architecture. Nodes are interconnected via an entanglement station that contains an optical switch, beam splitter(s) and detectors which enable the generation of ebits between nodes. This diagram illustrates how CB1 in QC_1 can be teleported to CB3 in QC_2.

3.1 Non-local CNOT

Fault tolerant non-local CNOTs involve transversal gates, measurements and ebits. For a pair of CSS $[\![n, k, d]\!]$ code blocks, CB1 and CB2, n ebit pairs are used. Fig. 3 shows how a fault tolerant non-local CNOT can be implemented between CB1 and CB2 located on different devices. A physical CNOT is performed between every data qubit k in CB1 and ebit k. Then, every ebit on CB2 is measured (in our case in the Z-basis). The parity of each measurement (0 or 1) is used next to conditionally apply an X gate to ebit k located on CB2. A CNOT is subsequently performed between ebit k and data qubit k on CB2. A Hadamard gate is applied to every ebit on CB2 before every ebit on CB2 is then measured (in our case, again, in the Z-basis). Finally, the parity of the measurement determines whether or not a Z gate is conditionally applied on CB2. The above combination of operations results in a FT CNOT between CB1 and CB2, with each logical qubit on CB1 being a control for its corresponding target logical qubit on CB2. The circuit is identical (other than differences due to the number of qubits or in the syndrome extraction rounds, the transversal operations applied are identical) between the surface codes and BB codes.

Using this protocol, we can perform teleportation of CB1 from QC1 to QC2 using CB2 and CB3, where both are initially in the $|0\rangle$ state. The non-local CNOT is used to create logical Bell states. Using these, we can then follow a similar procedure to the teleportation procedure in [22]. The circuit we use is shown in Fig. 2. A possible overall architecture is shown in Fig. 1.

3.2 Noise model and exact simulation

We use STIM [13], a fast stabilizer simulator. For BB codes, our noise model and syndrome extraction cycle is identical to that of [7]. For the Surface Code, our noise model and syndrome extraction circuit is identical to those contained in [13]. For both code families, every time a qubit is initialized (after measurement or in the beginning), it suffers bit flip noise p. Every 1 or 2-qubit gate has 1 or 2-qubit depolarizing noise p applied. Measurement error is simulated by applying bit flip noise p to the qubit (or ebit) being measured immediately prior to its measurement. The BB code syndrome extraction circuit features idle noise and in that case there is idle error in the from of depolarizing noise following the noise model in [7]. In addition, any ebit-data qubit interaction has the same noise as between two data qubits. Ebit creation noise is modeled as 2-qubit depolarizing noise applied to the ebits after they have been perfectly initialized.

The structure of the teleportation circuit is depicted in Figure 3. Code blocks CB1, CB2 and CB3 are noisily initialized before 4 syndrome extraction cycles are applied to each code block. Next, the logical qubits of CB2 and CB3 are entangled into logical Bell states using a transversal Hadamard on CB2 and a non-local CNOT between CB2 and CB3. Three rounds of syndrome extraction are subsequently applied to these two code blocks. These differ from previous syndrome extraction rounds as the transversal Hadamard has changed the basis, so what were X-check qubits are now initialized in $|0\rangle$ state and Z-check qubits are initialized in $|+\rangle$ state. The rest of the syndrome extraction cycle is the same. We then perform logical Bell state measurement on CB1, with corrections made to CB3. Finally, we perform one round of syndrome extraction on CB3 before measuring in the Z-basis. During simulation, noise is applied to each data qubit before being read out. As is common in other works, STIM detectors are only applied to Z stabilizers (and Z stabilizers that are briefly X stabilizers due to the transversal Hadamard) as there should be little difference between X and Z logical error rates. The circuit is, again, identical (other than differences due to the number of qubits or the syndrome extraction rounds, the transversal operations applied are identical) between the surface codes and BB codes. Note that when we simulate just a non-local CNOT, we nosily initialize the two code blocks, apply 4 syndrome extraction cycles to each, perform the non-local CNOT between the two code blocks and then perform 3 syndrome extraction cycles on each before measuring both code block in the Z basis. Detectors are again only applied on Z stabilizers

A detector error model (DEM) is then extracted from the STIM circuit. A new parity check matrix is derived from the

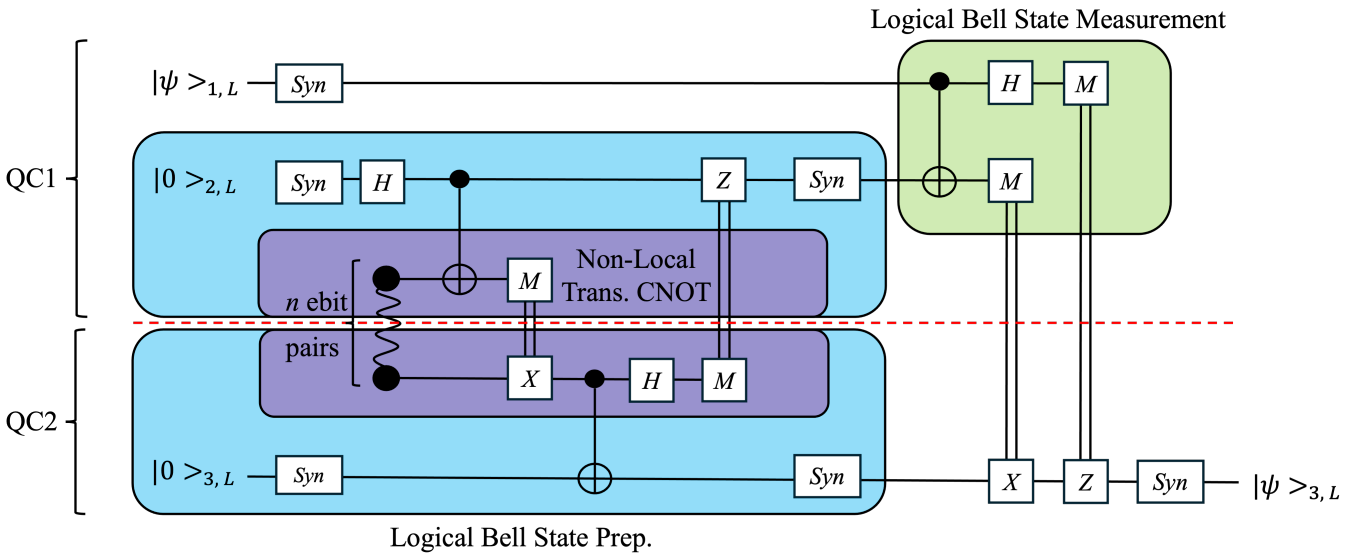

Figure 2. A circuit for fault tolerant teleportation of logical code block $|\psi\rangle$ from quantum computer QC1 to QC2 using two $|0\rangle$ code blocks, n ebit pairs and local operations only. The blue Logical Bell State Preparation subcircuit uses ebits to create k logical Bell states between code block CB2 and CB3. This resource is then used to teleport CB1 from QC1 to CB3 on QC2.

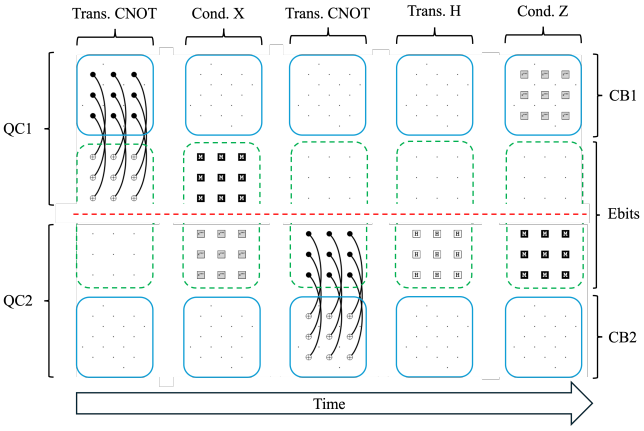

Figure 3. Non-local transversal CNOT carried out between CB1 and CB2 located on QC1 and QC2, respectively. The red line indicates the physical separation between devices. This operation involves a sequence of transversal CNOTS between code blocks and ebits, measurements and (conditional) single-qubit gates. Blue squares indicate logical code blocks. Green dashed lines indicate blocks of ebits.

DEM, where each "stabilizer" (row) corresponds to a detector. Each "data qubit" (column) now represents a distinct error mechanism that flips one or more detectors and possibly an observable. In the case of just the non-local CNOT circuit, there are $2k$ logical observables for a $[\![n, k, d]\!]$ code block. Each logical observable is the Z basis measurement of a logical qubit in one of the two code blocks. In the case of the teleportation circuit there are k logical observables. Each logical observable is the Z basis measurement of the logical qubits of code block 3 as defined in Fig. 2. The syndrome in both cases is a vector, where each entry represents whether a detector is 0 or 1 after the circuit has been noisily executed. This PCM and syndromes are then decoded using the BP-OSD decoder [21] with a maximum of 10^4 BP iterations and `OSD_ORDER` = 7.

The code used to generate STIM circuits and decode them involved heavily modifying and adding to the code in the Github repository corresponding to [14]. Their code can generate arbitrary single code block BB or surface code STIM quantum memory experiments and decode them using [21]. We created new classes and functions that support the generation of STIM circuits featuring logical operations between an arbitrary number of code blocks using local physical gates or ebits.

4 Results

Tab. 1 shows the difference in circuit-level distance for a number of codes while contrasting non-local CNOTs vs. teleportation. The calculated circuit-level distance is a heuristic generated by the STIM function `search_for_undetectable_logical_errors`. We observe that circuits using BB code blocks have a circuit-level distance upper bounded by values *lower* than their code distance whereas circuits using surface codes (SC) have a circuit level distance upper bounded by values *equal* to their distance. Circuit distance refers to the lowest weight undetectable error that flips a logical observable. As calculated, for SC, this quantity is upper bounded by the code distance.

The discrepancy between the SC and BB codes is expected. This is because whilst the syndrome extraction circuit for BB codes utilized here is one of the most highly performant, it is not circuit distance preserving even in simple memory experiments, as shown by [7].

Table 1. Upper bound on circuit-level distance for different SC and BB code blocks. The circuits in question are the non-local CNOT circuit involving 2 code blocks and a full teleportation circuit involving 3 code blocks. A dash indicates a result could not be found in feasible time.

Code	d non-local CNOT	d teleport
$[\![18, 4, 4]\!]$ BB	3	3
$[\![25, 1, 5]\!]$ SC	5	5
$[\![54, 4, 8]\!]$ BB	7	5
$[\![49, 1, 7]\!]$ SC	7	7
$[\![144, 12, 12]\!]$ BB	-	10
$[\![121, 1, 11]\!]$ SC	11	11

Tab. 2 shows the number of gates and measurements used by the non-local CNOT circuit as well as the size of the Parity Check Matrix (PCM) used for decoding the circuits for different codes. Notice that BB codes use up to almost 100% more 2-qubit gates than surface codes with similar distance. This is partially explained by the fact that BB codes are weight-6 whereas SC are weight-4, i.e., each stabilizer measurement involves at least 50% more CNOTs between data qubits. In addition, the number of columns in the PCM for BB codes is typically double the number for comparable SC. This means there are more possible error mechanisms that affect detectors and observables, increasing the complexity of the decoding task.

Fig. 4-6 show logical error rate (LER) against physical error rate (PER) for different ebit noise for a non-local transversal CNOT between pairs of either BB codes or SC. PER refers to p as defined in Sec. 3.2. Given that we are not performing memory experiments, LER is not per round. It instead refers to the probability that a logical error has occurred on any observable during the circuit. Error bars indicate the region where the Bayes' factor is ≤ 1000. Assuming that logical errors are binomially distributed, the regions indicated by the error bars are those where the probability of the LER being in that region is no more than 1000 times less than the best hypothesis, i.e., the LER is the number of logical errors divided by the number of shots. The error bars have the same meaning for all subsequent figures. If a PER lacks data points, the omission is attributed to the wall clock time too small to generate a statistically significant amount of errors at this level.

Fig. 4 highlights the performance of small codes of similar physical qubits n, one BB (solid) and the other of SC type (dashed), with LER (y-axis) over PER (x-axis) and for different ebit error rates (colors). The break-even threshold is indicated by a black dotted line, where below threshold data points indicate that logical encoding qubits gives a better error rate than their physical encoding. Whilst the $[\![25, 1, 5]\!]$ code achieves below threshold performance for ebit noise $\leq 10^{-3}$ at $p \approx 10^{-3}$. The $[\![18, 4, 4]\!]$ does not reach below threshold until significantly lower PER, and even then only for ebit noise $\leq 10^{-4}$.

Table 2. One (1q) and two qubit (2q) gates and measurements (M) used by a non-local logical CNOT circuit involving two instances of code blocks for different codes. The size of the space-time PCM generated by the STIM DEM is also shown, where rows correspond to the syndrome and number of detectors in the circuit and columns to the number of distinct error mechanisms within the circuit.

Code	1q	2q	M.	rows	cols
$[\![18, 4, 4]\!]$ BB	54	1548	288	144	1314
$[\![25, 1, 5]\!]$ SC	75	1170	386	192	847
$[\![54, 4, 8]\!]$ BB	162	4644	864	432	3942
$[\![49, 1, 7]\!]$ SC	147	2450	770	384	1809
$[\![144, 12, 12]\!]$ BB	432	12,384	2304	1152	10,512
$[\![121, 1, 11]\!]$ SC	363	6402	1922	960	4813

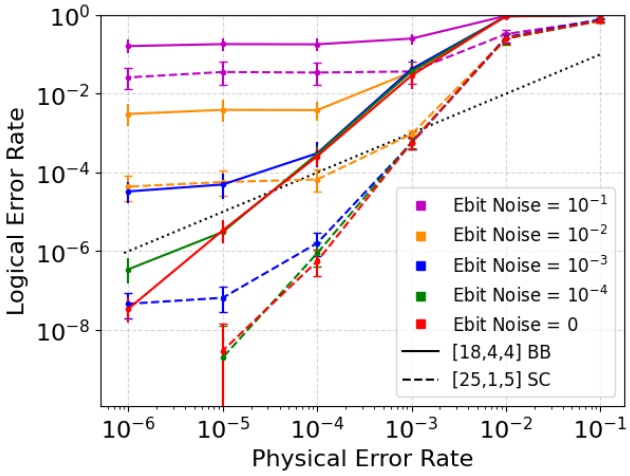

Figure 4. Physical noise rate against LER for different ebit noise rates for non-local CNOT between two identical $[\![18, 4, 4]\!]$ BB code blocks or $[\![25, 1, 5]\!]$ SC blocks.

Fig. 5 depicts medium codes using the same metrics. We observe that the difference between the performance of the BB codes and SC is smaller than in Fig. 4. Both are past the threshold at p on the order of 10^{-3} and perform similarly with a fixed difference in LER of approximately 10x.

Fig. 6 indicates how larger codes perform. We observe that the difference between the performance of the BB codes and SC is slightly smaller than in Fig. 5. Both are far past the threshold at p on the order of 10^{-3}, with a LER at this p on

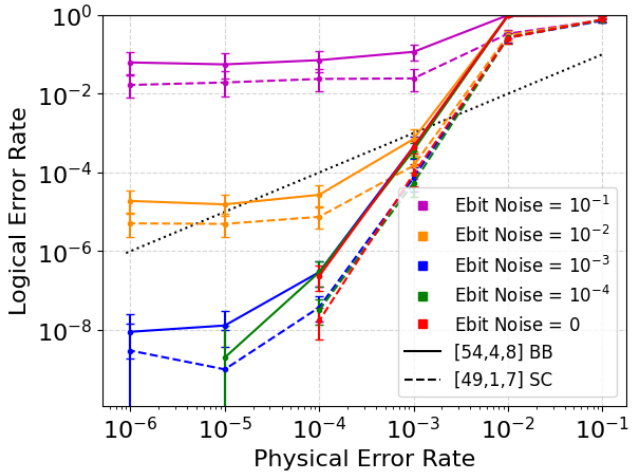

Figure 5. Physical noise rate against LER for different ebit noise rates for non-local CNOT between two identical $[\![54, 4, 8]\!]$ BB code blocks or $[\![49, 1, 7]\!]$ SC code blocks.

the order of 10 times better than in Fig. 5. The BB and SC LER have a similar gradient with, again, a fixed difference in LER of approximately 10x.

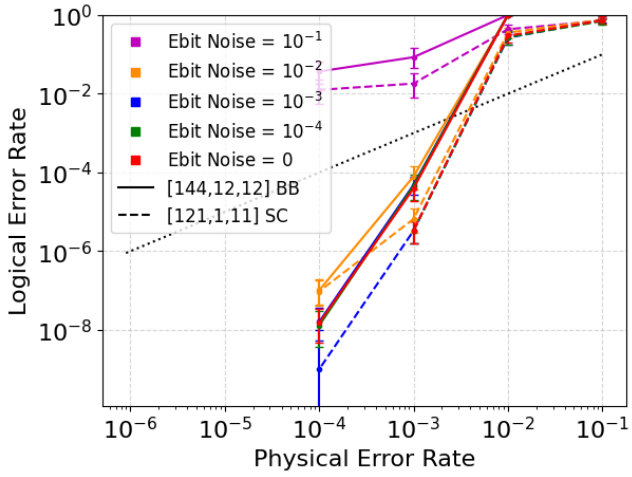

Figure 6. Physical noise rate against LER for different ebit noise rates for non-local CNOT between two identical $[\![144, 12, 12]\!]$ BB code blocks or $[\![121, 1, 11]\!]$ SC code blocks.

Tab. 3 indicates the number of gates and measurements used by the non-local CNOT circuit and the teleportation circuit, as well as the size of the PCM used for decoding the circuits. The trends are resembling Tab. 2 with BB codes requiring up to almost 100% more 2-qubit gates, more measurements, and double the PCM columns compared to their similar distance SC counterparts. Compared to Tab. 2, we also see that the teleportation circuit has approximately 50% more 2-qubit gates than the non-local CNOT circuit. This is due to the circuit requiring 3 rather than 2 code blocks.

Table 3. One (1q) and two qubit (2q) gates and measurements (M) used by a teleportation circuit involving three instances of code blocks for different codes. The size of the space-time PCM generated by the STIM DEM is also shown, where rows correspond to the syndrome and number of detectors in the circuit and columns to the number of distinct error mechanisms within the circuit.

Code	1q	2q	M	rows	cols
$[\![18, 4, 4]\!]$ BB	126	2106	414	180	1512
$[\![25, 1, 5]\!]$ SC	175	1595	556	240	1008
$[\![54, 4, 8]\!]$ BB	378	6318	1242	540	4536
$[\![49, 1, 7]\!]$ SC	343	3339	1108	480	2154
$[\![144, 12, 12]\!]$ BB	1008	16,848	3312	1440	12,096
$[\![121, 1, 11]\!]$ SC	847	8723	2764	1200	5730

Fig. 7-9 show LER against PER for different ebit noise for fault tolerant teleportation of an entire code block using only transversal operations, as described in Fig. 2. Physical error rate again refers to p as defined in Sec. 3.2. LER is defined identically to the non-local gate case. Error bars and any missing data points have the same meaning and reasons as in Fig. 4-6.

Fig. 7 plots the performance of small codes. Whilst the $[\![25, 1, 5]\!]$ code achieves below threshold performance for ebit noise $\leq 10^{-3}$ as p approaches 10^{-4}, the $[\![18, 4, 4]\!]$ does not approach the threshold for non-perfect ebit noise until $p \leq 10^{-6}$.

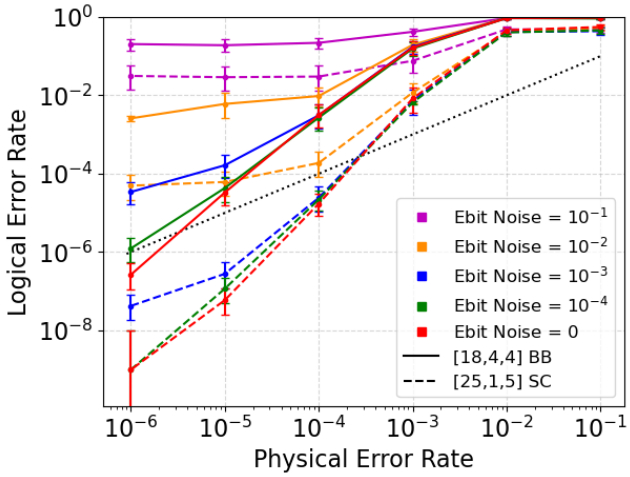

Figure 7. Physical noise rate against LER for different ebit noise rates for teleportation involving three identical $[\![18, 4, 4]\!]$ BB code blocks or $[\![25, 1, 5]\!]$ SC code blocks.

Fig. 8 indicates the performance of medium codes. Whilst the $[\![49, 1, 7]\!]$ code achieves good below threshold performance for ebit noise $\leq 10^{-3}$ as p decreases past 10^{-3}, the

$[\![54, 4, 8]\!]$ code only passes below the threshold as p decreases past 10^{-4}.

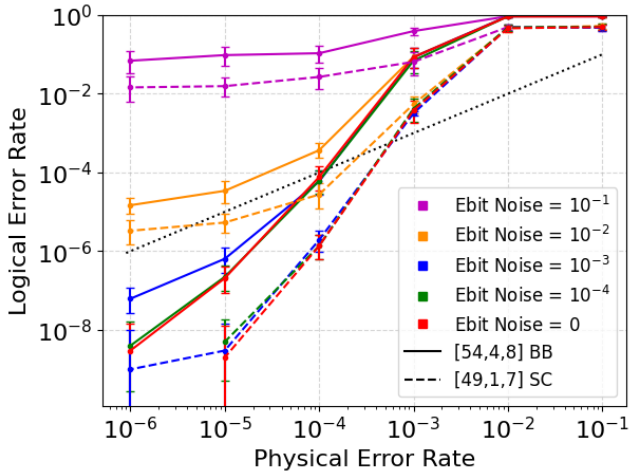

Figure 8. Physical noise rate against LER for different ebit noise rates for teleportation involving three identical $[\![54, 4, 8]\!]$ BB code blocks or $[\![49, 1, 7]\!]$ SC code blocks.

Fig. 9 indicates the performance of larger codes. The performance difference between the two code families is less pronounced here. They both achieve excellent below threshold performance as p approaches 10^{-4} with ebit noise $\leq 10^{-3}$, with the $[\![121, 1, 11]\!]$ code achieving 1000 times better LER than p for this ebit noise.

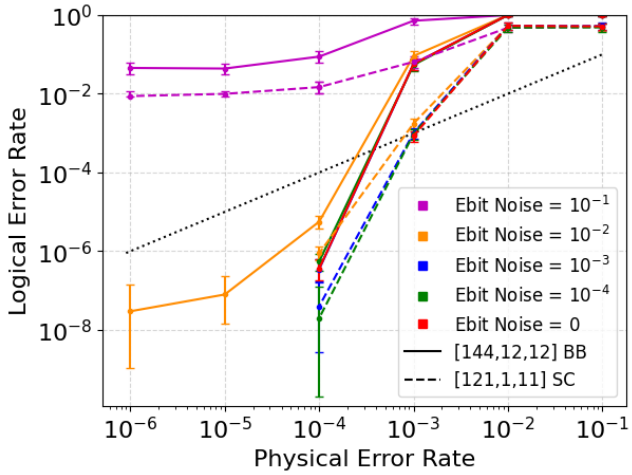

Figure 9. Physical noise rate against LER for different ebit noise rates for teleportation involving three identical $[\![144, 12, 12]\!]$ BB code blocks or $[\![121, 1, 11]\!]$ SC code blocks.

Overall, we observe that across all codes strong error suppressive behavior pushes results below the threshold for PER approaching 10^{-4} and ebit noise $\leq 10^{-3}$, except for the smallest $[\![18, 4, 4]\!]$ BB code. Codes are able to tolerate ebit noise up to 100 times higher than physical noise before approaching an ebit-caused error floor. But usually ebit noise 10 times larger that PER has only a small effect on LER. The extent to which noisy ebits cause this error floor is reduced as the distance of a code increases. For example, compare the relative difference between ebit noise 10^{-2} and 10^{-3} in Fig. 4 vs. Fig. 6, or in Fig. 7 compared to Fig. 9.

In general, similar distance SC performs better than their respective BB codes. This can be explained as follows. (1) The SC teleportation or non-local CNOT circuit has up to half as many CNOTs that the same sized BB code. (2) The syndromes considered in the BB codes are weight-6 compared to weight-4 for the SC, which contributes to a PCM with approximately twice as many error mechanisms (columns) than a comparable SC code so that decoding BB codes becomes more difficult. Further, as discussed above, the syndrome extraction circuit for BB codes is not circuit distance preserving, and this affects their performance. Recall though that the SC has significantly lower encoding capacity in logical qubits, k, i.e., BB encoding is far more desirable for state capacity at the logical level.

However, as the distance of both codes increases, this relative difference to each other reduces. Compare again Fig. 4 vs. Fig. 6, or Fig. 7 compared to Fig. 9.

For the non-local CNOT circuit, we observe that for all codes other than $[\![18, 4, 4]\!]$, $p \leq 10^{-3}$ gives below threshold performance. As distance increases, the performance of both the BB and SC increases as expected, other than a very large jump from $[\![18, 4, 4]\!]$ to $[\![54, 4, 8]\!]$.

It is notable that in a distributed architecture with all the additional gates and noise that entails, codes can produce a LER between 10,000 and 100,000 times lower than the PER, e.g., for the $[\![144, 12, 12]\!]$ BB code and $[\![121, 1, 11]\!]$ SC, respectively, as shown in Fig. 6.

For logical teleportation, codes (other than, again, the $[\![18, 4, 4]\!]$ BB code) need a PER below the range of $\geq 10^{-4}$ to $\geq 10^{-3}$, depending on the code, to achieve below threshold performance. That being said, in the regimes we were able to simulate in feasible wall clock time, we see codes achieving LERs between 1000 and 10,000 times better than PER as shown in Fig. 9. Given that we are simulating 3 code blocks, this is outstanding. Again, like in the non-local CNOT circuit, codes can generally withstand ebit noise 10 times worse than physical noise without a major effect on performance. But ebit noise 100 times worse can still allow for up to 1000 times lower LER than p, e.g., in Fig. 6.

Teleportation performance is significantly worse than non-local CNOT performance. Depending on the code it is on the order of 100 times worse. This difference can be partially explained by the fact the logical teleportation circuit has approximately 25% more CNOTs, 33% more measurements and 10-30% larger PCMs. Additionally, decoding across 3

code blocks rather than 2, makes it more complex than the fault tolerant non-local gate.

This vast difference in LER between non-local logical CNOT and logical teleportation shows the importance of compilers for distributed quantum computing. Comparing the physical, non fault-tolerant, circuits for non-local CNOTs and teleportation [10] gives little reason to think there would be a significant difference in fidelity. Clearly, a circuit will need to be compiled such that logical teleportations are used only when they are more efficient than using approximately 100 non-local CNOTs.

That being said, this trend may change for larger BB codes given their encoding performance. The non-local CNOT operation we have simulated is transversal, i.e., every logical qubit in the control code block operates on its corresponding qubit in the target code block. This may not necessarily always be desired. Teleporting 100 logical qubits via a BB code would likely be much more useful even though it may be 100 times noisier than performing a non-local gate.

Exploring logical operations between larger codes than those considered in this work motivates the development of more efficient computational tools. Simulations of the largest codes in this work took up to 2 days with 250-450 cores, depending on experiments. Almost all computational time was spent on decoding. Decoding performance using BP-OSD is approximately polynomial [21] in the size of the PCM, the size of which increases approximately linearly with qubit size. Clearly, simulating at low-noise regimes codes that are 10 times larger than those in this work will be a very difficult computational task without changes in tooling.

4.1 Observations

We contribute the following novel results:

1. Operations between nodes in a distributed quantum computer can be carried out with LER up to 100,000 lower than PERs in near term regimes (PER $\leq 10^{-4}$) using ebits that are between 10 and 100 times noiser than PER.
2. Codes can be teleported between nodes in a distributed quantum computer with a LER up to between 1000 and 10,000 times lower than PERs in near term regimes (PER $\leq 10^{-4}$).
3. Teleportation carries a LER approximately 100 times worse than a non-local CNOT at similar code distance, ebit noise and PER.
4. For medium and large codes, surface codes carry a LER 10 times lower than BB codes of similar distance. This effect, however, appears to reduce as distance increases.

5 Conclusion

This paper contributes circuit-level simulations of surface and BB codes of non-local CNOT gates and teleportation in a DQC setup. We explore a wide range of ebit noise and physical noise rates to assess which codes result in noise levels under the threshold. We additionally explore lower physical noise rates than typically considered in prior work to provide a path for future manufacturing choices and, interestingly, do not encounter decoder-induced noise floors. Overall, we find that high performance in terms of noise reduction is possible even with relatively noisy ebits.

Acknowledgments

This work was supported in part by NSF awards MPS-2410675, PHY-1818914, PHY-2325080, MPS-2120757, CISE-2217020, and CISE-2316201 as well as DOE DE-SC0025384.

References

[1] M. Akhtar, F. Bonus, F. R. Lebrun-Gallagher, N. I. Johnson, M. Siegele-Brown, S. Hong, S. J. Hile, S. A. Kulmiya, S. Weidt, and W. K. Hensinger. 2023. A high-fidelity quantum matter-link between ion-trap microchip modules. *Nature Communications* 14, 1 (Feb. 2023), 531. doi:10.1038/s41467-022-35285-3

[2] Pablo Andres-Martinez, Tim Forrer, Daniel Mills, Jun-Yi Wu, Luciana Henaut, Kentaro Yamamoto, Mio Murao, and Ross Duncan. 2024. Distributing circuits over heterogeneous, modular quantum computing network architectures. *Quantum Science and Technology* 9, 4 (Oct. 2024), 045021. doi:10.1088/2058-9565/ad6734

[3] James Ang, Gabriella Carini, Yanzhu Chen, Isaac Chuang, Michael Demarco, Sophia Economou, Alec Eickbusch, Andrei Faraon, Kai-Mei Fu, Steven Girvin, Michael Hatridge, Andrew Houck, Paul Hilaire, Kevin Krsulich, Ang Li, Chenxu Liu, Yuan Liu, Margaret Martonosi, David McKay, Jim Misewich, Mark Ritter, Robert Schoelkopf, Samuel Stein, Sara Sussman, Hong Tang, Wei Tang, Teague Tomesh, Norm Tubman, Chen Wang, Nathan Wiebe, Yongxin Yao, Dillon Yost, and Yiyu Zhou. 2024. ARQUIN: Architectures for Multinode Superconducting Quantum Computers. *ACM Transactions on Quantum Computing* 5, 3, Article 19 (Sept. 2024), 59 pages. doi:10.1145/3674151

[4] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. 2019. Quantum supremacy using a programmable superconducting processor. *Nature* 574, 7779 (Oct. 2019), 505–510. doi:10.1038/s41586-019-1666-5

[5] J. Pablo Bonilla Ataides, Hengyun Zhou, Qian Xu, Gefen Baranes, Bikun Li, Mikhail D. Lukin, and Liang Jiang. 2025. Constant-Overhead Fault-Tolerant Bell-Pair Distillation using High-Rate Codes. doi:10.48550/arXiv.2502.09542 arXiv:2502.09542 [quant-ph].

[6] Sean D. Barrett and Pieter Kok. 2005. Efficient high-fidelity quantum computation using matter qubits and linear optics. *Phys. Rev. A* 71 (Jun 2005), 060310. Issue 6. doi:10.1103/PhysRevA.71.060310

[7] Sergey Bravyi, Andrew W. Cross, Jay M. Gambetta, Dmitri Maslov, Patrick Rall, and Theodore J. Yoder. 2024. High-threshold and low-overhead fault-tolerant quantum memory. *Nature* 627, 8005 (March 2024), 778–782. doi:10.1038/s41586-024-07107-7 arXiv:2308.07915 [quant-ph].

[8] Lukas Brenner, Christophe Piveteau, and David Sutter. 2023. Optimal wire cutting with classical communication. arXiv:2302.03366 [quant-ph] https://arxiv.org/abs/2302.03366

[9] Bryan Eastin and Emanuel Knill. 2009. Restrictions on Transversal Encoded Quantum Gate Sets. *Phys. Rev. Lett.* 102 (Mar 2009), 110502. Issue 11. doi:10.1103/PhysRevLett.102.110502

[10] J. Eisert, K. Jacobs, P. Papadopoulos, and M. B. Plenio. 2000. Optimal local implementation of nonlocal quantum gates. *Physical Review A* 62, 5 (Oct. 2000), 052317. doi:10.1103/PhysRevA.62.052317

[11] Marco Fellous-Asiani, Jing Hao Chai, Robert S. Whitney, Alexia Auffèves, and Hui Khoon Ng. 2021. Limitations in Quantum Computing from Resource Constraints. *PRX Quantum* 2, 4 (Nov. 2021), 040335. doi:10.1103/PRXQuantum.2.040335

[12] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. 2012. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A* 86 (Sep 2012), 032324. Issue 3. doi:10.1103/PhysRevA.86.032324

[13] Craig Gidney. 2021. Stim: a fast stabilizer circuit simulator. *Quantum* 5 (July 2021), 497. doi:10.22331/q-2021-07-06-497

[14] Anqi Gong, Sebastian Cammerer, and Joseph M. Renes. 2024. Toward Low-latency Iterative Decoding of QLDPC Codes Under Circuit-Level Noise. arXiv:2403.18901 [quant-ph] https://arxiv.org/abs/2403.18901

[15] Charles Guinn, Samuel Stein, Esin Tureci, Guus Avis, Chenxu Liu, Stefan Krastanov, Andrew A. Houck, and Ang Li. 2023. Co-Designed Superconducting Architecture for Lattice Surgery of Surface Codes with Quantum Interface Routing Card. arXiv:2312.01246 [quant-ph] https://arxiv.org/abs/2312.01246

[16] Peter Kaufmann, Timm F. Gloger, Delia Kaufmann, Michael Johanning, and Christof Wunderlich. 2018. High-Fidelity Preservation of Quantum Information During Trapped-Ion Transport. *Physical Review Letters* 120, 1 (Jan. 2018), 010501. doi:10.1103/PhysRevLett.120.010501

[17] Ying Li and Simon C. Benjamin. 2016. Hierarchical surface code for network quantum computing with modules of arbitrary size. *Phys. Rev. A* 94 (Oct 2016), 042303. Issue 4. doi:10.1103/PhysRevA.94.042303

[18] D. Main, P. Drmota, D. P. Nadlinger, E. M. Ainley, A. Agrawal, B. C. Nichol, R. Srinivas, G. Araneda, and D. M. Lucas. 2025. Distributed quantum computing across an optical network link. *Nature* 638, 8050 (Feb. 2025), 383–388. doi:10.1038/s41586-024-08404-x

[19] Naomi H. Nickerson, Joseph F. Fitzsimons, and Simon C. Benjamin. 2014. Freely Scalable Quantum Technologies Using Cells of 5-to-50 Qubits with Very Lossy and Noisy Photonic Links. *Phys. Rev. X* 4 (Dec 2014), 041041. Issue 4. doi:10.1103/PhysRevX.4.041041

[20] Joshua Ramette, Josiah Sinclair, Nikolas P. Breuckmann, and Vladan Vuletić. 2024. Fault-tolerant connection of error-corrected qubits with noisy links. *npj Quantum Information* 10, 1 (10 Jun 2024), 58. doi:10.1038/s41534-024-00855-4

[21] Joschka Roffe, David R. White, Simon Burton, and Earl Campbell. 2020. Decoding across the quantum low-density parity-check code landscape. *Phys. Rev. Res.* 2 (Dec 2020), 043423. Issue 4. doi:10.1103/PhysRevResearch.2.043423

[22] C. Ryan-Anderson, N. C. Brown, C. H. Baldwin, J. M. Dreiling, C. Foltz, J. P. Gaebler, T. M. Gatterman, N. Hewitt, C. Holliman, C. V. Horst, J. Johansen, D. Lucchetti, T. Mengle, M. Matheny, Y. Matsuoka, K. Mayer, M. Mills, S. A. Moses, B. Neyenhuis, J. Pino, P. Siegfried, R. P. Stutz, J. Walker, and D. Hayes. 2024. High-fidelity teleportation of a logical qubit using transversal gates and lattice surgery. *Science* 385, 6715 (2024), 1327–1331. doi:10.1126/science.adp6016 arXiv:https://www.science.org/doi/pdf/10.1126/science.adp6016

[23] Mohamed A. Shalby, Renyu Wang, Denis Sedov, and Leonid P. Pryadko. 2025. Optimized noise-resilient surface code teleportation interfaces. doi:10.48550/arXiv.2503.04968 arXiv:2503.04968 [quant-ph].

[24] Siddhant Singh, Fenglei Gu, Sébastian de Bone, Eduardo Villaseñor, David Elkouss, and Johannes Borregaard. 2024. Modular Architectures and Entanglement Schemes for Error-Corrected Distributed Quantum Computation. arXiv:2408.02837 [quant-ph] https://arxiv.org/abs/2408.02837

[25] Evan Sutcliffe, Bhargavi Jonnadula, Claire Le Gall, Alexandra E. Moylett, and Coral M. Westoby. 2025. Distributed quantum error correction based on hyperbolic Floquet codes. arXiv:2501.14029 [quant-ph] https://arxiv.org/abs/2501.14029

[26] Joshua Viszlai, Willers Yang, Sophia Fuhui Lin, Junyu Liu, Natalia Nottingham, Jonathan M. Baker, and Frederic T. Chong. 2024. Matching Generalized-Bicycle Codes to Neutral Atoms for Low-Overhead Fault-Tolerance. arXiv:2311.16980 [quant-ph] https://arxiv.org/abs/2311.16980

[27] Anocha Yimsiriwattana and Samuel J. Lomonaco Jr. 2004. Generalized GHZ States and Distributed Quantum Computing. http://arxiv.org/abs/quant-ph/0402148 arXiv:quant-ph/0402148.

[28] Youwei Zhao, Yangsen Ye, He-Liang Huang, Yiming Zhang, Dachao Wu, Huijie Guan, Qingling Zhu, Zuolin Wei, Tan He, Sirui Cao, Fusheng Chen, Tung-Hsun Chung, Hui Deng, Daojin Fan, Ming Gong, Cheng Guo, Shaojun Guo, Lianchen Han, Na Li, Shaowei Li, Yuan Li, Futian Liang, Jin Lin, Haoran Qian, Hao Rong, Hong Su, Lihua Sun, Shiyu Wang, Yulin Wu, Yu Xu, Chong Ying, Jiale Yu, Chen Zha, Kaili Zhang, Yong-Heng Huo, Chao-Yang Lu, Cheng-Zhi Peng, Xiaobo Zhu, and Jian-Wei Pan. 2022. Realization of an Error-Correcting Surface Code with Superconducting Qubits. *Phys. Rev. Lett.* 129 (Jul 2022), 030501. Issue 3. doi:10.1103/PhysRevLett.129.030501