

NISQ+: Boosting quantum computing power by approximating quantum error correction

Adam Holmes
Department of Computer Science
University of Chicago
 Chicago, USA
 adholmes@uchicago.edu
Intel Labs
Intel Corporation
 Oregon, USA

Mohammad Reza Jokar
Department of Computer Science
University of Chicago
 Chicago, USA
 jokar@uchicago.edu

Ghasem Pasandi
Department of Electrical and
Computer Engineering
University of Southern California (USC)
 Los Angeles, USA
 pasandi@usc.edu

Yongshan Ding
Department of Computer Science
University of Chicago
 Chicago, USA
 yongshan@uchicago.edu

Massoud Pedram
Department of Electrical and
Computer Engineering
University of Southern California (USC)
 Los Angeles, USA
 pedram@usc.edu

Frederic T. Chong
Department of Computer Science
University of Chicago
 Chicago, USA
 chong@cs.uchicago.edu

Abstract—Quantum computers are growing in size, and design decisions are being made now that attempt to squeeze more computation out of these machines. In this spirit, we design a method to boost the computational power of near-term quantum computers by adapting protocols used in quantum error correction to implement “Approximate Quantum Error Correction (AQEC).” By approximating fully-fledged error correction mechanisms, we can increase the compute volume (qubits \times gates, or “Simple Quantum Volume (SQV)”) of near-term machines. The crux of our design is a fast hardware decoder that can approximately decode detected error syndromes rapidly. Specifically, we demonstrate a proof-of-concept that approximate error decoding can be accomplished online in near-term quantum systems by designing and implementing a novel algorithm in superconducting Single Flux Quantum (SFQ) logic technology. This avoids a critical decoding backlog, hidden in all offline decoding schemes, that leads to idle time exponential in the number of T gates in a program [58].

Our design utilizes one SFQ processing module per physical quantum bit. Employing state-of-the-art SFQ synthesis tools, we show that the circuit area, power, and latency are within the constraints of typical, contemporary quantum system designs. Under a pure dephasing error model, the proposed accelerator and AQEC solution is able to expand SQV by factors between 3,402 and 11,163 on expected near-term machines. The decoder achieves a 5% accuracy threshold as well as pseudo-thresholds of approximately 5%, 4.75%, 4.5%, and 3.5% physical error rates for code distances 3, 5, 7, and 9, respectively. Decoding solutions are achieved in a maximum of ~ 20 nanoseconds on the largest code distances studied. By avoiding the exponential idle time in offline decoders, we achieve a 10x reduction in required code distances to achieve the same logical performance as alternative designs.

Index Terms—Quantum Computing, Quantum Error Correction, Surface Code

I. INTRODUCTION

Quantum computing has the potential to revolutionize computing and have massive effects on major industries including agriculture, energy, and materials science by solving computational problems that are intractable with conventional machines [30], [53]. As we begin to build quantum computing machines of between 50-100 qubits [51] and larger, design decisions are being made to attempt to get the most computation out of a machine, quantified in this work by expanding the “Simple Quantum Volume” (SQV). SQV can be defined as the number of computational qubits of a machine multiplied by the number of gates we expect to be able to perform without error, as in Figure 1. One limiting factor on SQV now is that physical quantum bits (qubits) are extremely error-prone, which means that computation on these machines is bottlenecked by the short lifetimes of qubits. System designers combat this by attempting to build better physical qubits, but this effort is extremely difficult and classical systems can be used to alleviate the burden. Specifically, *quantum error correction* is a classical control technique that decreases the rate of errors in qubits and expands the SQV. Error correction proceeds by encoding a set of *logical* qubits to be used for algorithms into a set of faulty physical qubits. Information about the current state of the device, called *syndromes*, is extracted by a specific quantum circuit that does not disturb the underlying computation. Decoding is the process by which an error correcting protocol maps this information to a set of *corrections* that, if chosen correctly, should return the system to the correct logical state. Fully fault tolerant machines can expand the SQV rapidly by suppressing qubit errors exponentially with




Fig. 1. Boosting the quantum computation power with approximate error correction schemes. A machine with 1024 faulty physical qubits of error rate 10^{-5} has an SQV of $\approx 10^8$. By performing fast, online, approximate decoding, we can trade the number of computational qubits for gate fidelity and boost the SQV by over a factor of 3,402. Moving to a higher code distance raises this increase to a factor of 11,163. NISQ machines are severely limited by gate fidelity, and introducing error mitigation techniques can have dramatic effects on SQV.

the code distance.

While a fully fault-tolerant quantum computer may take many years to construct, it is possible to use the well-developed theory of error correction as inspiration for constructing *error mitigation* protocols that still provide a strong expansion in SQV. In this paper we present an approximate decoding solution specifically targeting execution time and show that we can in fact perform decoding at the speed of syndrome generation for near-term machines. Prior work has suggested and analyzed software solutions for decoding, but relying on hardware-software communication can be slow, especially considering the cryogenic environment of typical quantum computing systems. If decoding occurs slower than error information is generated, the system will generate a backlog of information as it waits for decoding to complete, introducing an exponential time overhead that will kill any quantum advantage (see Section III). A hardware solution proposed here results in the ability to perform logical gates with orders of magnitude better fidelity and at the speed of syndrome generation, resulting in a major expansion in SQV as shown in Figure 1. This relies on an approximate decoding algorithm implemented in superconducting Single Flux Quantum (SFQ) hardware. While the algorithmic design enables the accuracy of the hardware accelerator to be competitive at small scale with existing software implementations, the benefits of implementing the circuitry directly in SFQ hardware are numerous. Specifically, high clock speeds, low power dissipation, and unique gating style allows for our accelerator to be co-located with a quantum chip inside a dilution refrigerator, avoiding otherwise high communication costs.

This work contributes the following:

- 1) We design an approximate decoding algorithm for stabilizer codes based on SFQ hardware, leveraging unique capabilities that the hardware offers,
- 2) We show that using this new error mitigation technique, we can expand the SQV of near-term machines by

factors of between 3,402 and 11,163,

- 3) We use Monte-Carlo simulation based benchmarking of the hardware accelerator, resulting in effective accuracy and pseudo-thresholds,
- 4) We perform system execution time analysis, realistically benchmarking the decoder performance in real time and showing that decoding is likely to be able to proceed at or exceeding the speed of data generation enabling the benefits of fault tolerant quantum computing.
- 5) We show that our online decoder requires 10x smaller code distance than offline decoders when decoding backlog accounted for.

The remainder of the paper is as follows: Section II describes the necessary background of quantum computation and details the specifications of typical quantum computing systems stacks. Section II-B describes quantum error correction and the decoding problem in detail. Section IV describes relevant related work in the area ranging from optimized software implementations of matching algorithms to novel descriptions of neural network based decoders. Section V describes our decoding algorithm, and Section VI describes implementation details of SFQ technology, and the circuit datapaths in detail. Section VII describes our methodology for evaluation, including details of the simulation environment in which our accelerator was benchmarked, details of the metrics used to evaluate performance, and descriptions of novel synthesis tools used to generate efficient layouts of SFQ circuitry. Section VIII presents our accuracy results, a breakdown of the accelerator characterization including area, power, and latency footprints, a timing evaluation, and analysis of the SQV effects. Section IX concludes.

II. BACKGROUND

In this section we discuss the basics of quantum computation, quantum error correction, and a description of the fundamental components of a quantum computing system architecture.

A. Basics of Quantum Computation

Here we provide a brief overview of quantum computation necessary to discuss quantum error correction. For more detailed discussions see [43]. A quantum computing algorithm is a series of operations on two level quantum states called *qubits*, which are quantum analogues to classical bits. A qubit state can be written mathematically as a superposition of two states as $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, where the coefficients $\alpha, \beta \in \mathbb{C}$ and $|\alpha|^2 + |\beta|^2 = 1$. A measured qubit will yield a value of $|0\rangle$ or $|1\rangle$ with probability $|\alpha|^2$ or $|\beta|^2$, respectively, at which point the qubit state will be exactly $|0\rangle$ or $|1\rangle$. Larger quantum systems are represented simply as $|\psi\rangle = \sum_i \alpha_i |i\rangle$ where $|i\rangle$ are computational basis states of the larger quantum system.

Quantum operations (gates) transform qubit states to other qubit states. In this work we will be making use of particular quantum operations known as *Pauli gates*, denoted as $\{I, X, Y, Z\}$. These operations form a basis for all quantum operations that can occur on a single qubit, and therefore

any operation can be written as a linear combination of these gates. Additionally, error correction circuits make use of the Hadamard gate H , an operation that constructs an evenly weighted superposition of basis elements when acting on a basis element. Two-qubit controlled operations will also be used, which can generate entanglement between qubits and are required to perform universal computation.

B. Quantum Error Correction

Qubits are intrinsically fragile quantum systems that require isolation from environmental interactions in order to preserve their values. *Decoherence*, for example the decay of a quantum state from a general state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ to the ground state $|\psi'\rangle = |0\rangle$ happens rapidly in many physical qubit types, often on the order of tens of nanoseconds [57], [59]. This places a major constraint on algorithms: without any modifications to the system, algorithms can only run for a small, finite time frame with high probability of success.

To combat this, *quantum error correction* protocols have been developed. These consist of encoding a small number of *logical* qubits used for computation in algorithms into a larger number of physical qubits, resulting in a higher degree of reliability [12], [20], [39], [58]. In general, developing quantum error correction protocols is difficult as directly measuring the qubits that comprise a system will result in destruction of the data. To avoid this, protocols rely upon indirectly gathering error information via the introduction of extra qubits that interact with the primary set of qubits and are measured. This measurement data is then used to infer the locations of erroneous data qubits.

While many different types of protocols have been developed, this work focuses primarily on the *surface code*, a topological stabilizer code [28] that is widely considered to be the best performing code for the medium-term as it relies purely on geometrically local interactions between physical qubits greatly facilitating its fabrication in hardware, and has been shown to have very high reliability overall [20].

C. The Surface Code

Errors can occur on physical qubits in a continuous fashion, as each physical qubit is represented mathematically by two complex coefficients that can change values in a continuous range. However, a characteristic of the quantum mechanics leveraged by the surface code is that these continuous errors can be *discretized* into a small set of distinct errors. In particular, the action of the surface code maps these continuous errors into Pauli error operators of the form $\{I, X, Y, Z\}$ occurring on the data. This is one of the main features of the code that allows error detection and correction to proceed.

The surface code procedure that accomplishes error discretization, detection, and correction is an error correcting code that operates upon a two-dimensional lattice of physical qubits. The code designates a subset of the qubits as data qubits responsible for forming the logical qubit, and others as ancillary qubits responsible for detecting the presence of errors in the data. This is shown graphically in Figure 2.




Fig. 2. Figure (a) shows a graphical illustration of a surface code mesh. Gray circles indicate data qubits, and nodes labeled X and Z indicate ancillary qubits measuring X and Z stabilizers, respectively. Ancillary qubits are joined by colored edges to the data qubits that they are responsible for measuring. In figure (b) a single data qubit experiences a Pauli X error indicated by red coloring, causing the neighboring Z ancillary qubits to detect an odd parity in their data qubit sets and return $+1$ measurement values indicated by green coloring. In figure (c), the data qubit in red experiences a Pauli Z error, causing the vertically adjacent X ancillary qubits to return $+1$ measurement values. The entire error syndrome strings for either of these two cases would include a string of 12 values, two of which would be $+1$ and the remaining 10 would be 0.

Ancillary qubits interact with all of their neighboring data qubits and are then measured, and the measurement outcomes form the *error syndrome*. This set of operations forms the *stabilizer circuit*, where each ancillary qubit measures a four-qubit operator called a *stabilizer*.

1) *Error Detection*: The ancillary qubits are partitioned into those denoted as X and Z ancilla qubits. These ancilla qubit sets are sufficient for capturing any Pauli error on the data qubits, as Y operators can be treated as a simultaneous X and Z error. The action of the X stabilizer is two-fold: the four neighboring data qubits are forced into a particular state that discretizes any errors that may have occurred on them. Second, the measurement of the X ancilla qubit signals the parity of the number of errors that have occurred on its four neighbors. For example, it yields a $+1$ value if the state of the four neighboring qubits has an even number of Z errors. The same is true of the Z stabilizers – these track the parity of X errors occurring in the neighboring qubits. If an odd number of errors have occurred in either case, the ancilla qubit measurement will yield a $+1$ value, an event known as a *detection event* [23], otherwise these will return values of 0 or -1 depending on convention. We will refer to the ancillary qubits returning $+1$ values as *hot syndromes*. The *error syndrome* of the code is a bit string of length equal to the total number of ancilla qubits, and is composed of all of these measurement values.

Decoding is the process of mapping a particular error syndrome string to a set of corrections to be applied on the device. An example of this process is shown graphically in Figure 2. In this example, the hot syndromes generated by a single data qubit error are marked in red. Each single data qubit error causes the adjacent ancillary qubits to return $+1$ values.

A different situation occurs when strings of data qubit errors cross ancillary qubits, as shown in Figure 3. Here, four consecutive data qubits experience errors which generates hot syndrome measurements on the far left and right of the grid. This is because each ancillary qubit along this chain detects




Fig. 3. Figure (a) shows a data qubit error pattern spanning across ancillary qubits. Each data qubit experiencing error is indicated in red, and the ancillary qubits returning +1 measurement values are indicated in green. Each ancillary qubit that is adjacent to two erroneous data qubits does not signal the presence of any errors, as the parity of the data qubit sets are still even. This creates an *error string* that runs from the ancillary qubit on the left of the grid to the one on the right. Decoding must map these +1 values to the corresponding set of 4 data qubit errors that generated it. Figures (b) and (c) show degeneracy in error syndrome generation by surface code data qubit error patterns. The figures depict two distinct sets of data qubit error patterns that both generate the same error syndromes. Both patterns contain the same number of physical data errors, so these patterns are equally likely assuming independence of errors.

even error parity, so they do not signal the presence of errors. Decoding must be able to pair the two hot syndromes, applying corrections along the chain that connects them.

2) *Error Detection Can Fail*: Notice that in Figure 3 (a), if the data qubits on the left and right endpoints of the chain had also experienced errors, none of the ancillary qubits would have detected the chain. This represents a class of undetectable error chains in the code, and specifically occurs when chains cross from one side of the lattice to the other. The result of these chains are physical errors present in the code that cannot be corrected, and are known as *logical errors*, as they have changed the state of the logical qubit. One important characteristic of the surface code is the minimal number of qubits required to form a logical error. This number is referred to as the *code distance*, d of a particular lattice.

D. Quantum Computing Systems Organization

While qubits are the foundation of a device, a quantum computer must contain many layers of controlling devices in order to interact with qubits. Qubits themselves can be constructed using many different technologies, some of which include superconducting circuits [3], [25], [26], [34], [41], trapped ions [18], [29], [36], [41], [42], and quantum dots [71]. Controlling these devices is often performed by application of electrical signals at microwave frequencies [8], [44], [50], [70].

This work focuses on systems built around qubits that require cryogenic cooling to milliKelvin temperatures [33]. These systems require the use of dilution refrigerators, and typical architectures involve classical controllers located in various temperature stages of the system. Such a system is described schematically in [33], [57], and presents many design constraints. Controllers inside the refrigerator are subject to area and power dissipation constraints [49], [52]. Communication between stages can be costly as well. Many systems are constructed today using control wiring that scales linearly with the number of qubits, which will prohibit the construction of scalable machines [24].

E. Classical Control in Quantum Computing Systems

Error correction classical processing requires high bandwidth communication of the measurement values of many qubits on the quantum substrate repeatedly throughout the operation of the device, encouraging studies of engineering solutions [68], feasibility [56] and controller design [57]. Not

	# qubits	# total gates	# T gates
takahashi_adder	40	740	266
barenco_half_dirty_toffoli	39	1224	504
cnu_half_borrowed	37	1156	476
cnx_log_depth	39	629	259
cuccaro_adder	42	821	280

TABLE I
CHARACTERISTICS OF THE SIMULATED BENCHMARKS.

only are instruction streams primarily dominated by quantum error correction operations [37], [38], but also the classical controller responsible for error correction processing must be tightly coupled to the quantum substrate. If communicating between the quantum substrate and error correcting controller is subject to excessive latencies, the execution of fault tolerant algorithms will be completely prohibited.

III. MOTIVATION: DECODING MUST BE FAST

Decoding must be done quickly for the surface code to perform well. During actual computation on a surface code error corrected device, there exist gates called T -gates that require knowledge of the current state of errors on the device before they can execute.¹ If decoding is slower than the rate at which syndromes are generated, an algorithm will create a *data backlog*. While the machine is waiting for decoder to process the backlog, more syndrome data is accumulating on the device, which must be processed before executing the subsequent T -gate. Over time, this results in latency overhead that is exponentially dependent upon the number of such gates. Specifically, the overhead scales as $(\frac{r_{\text{gen}}}{r_{\text{proc}}})^k = f^k$, where r_{gen} is the rate of data generation, r_{proc} is the rate of decoder processing, each in bauds, f is the decoding ratio, and k is the number of T gates in the quantum algorithm. An exponentially slow quantum computer eliminates all of its usefulness.

Figure 4 shows the exponential latency overhead due to data backlog. The proof of this is summarized as follows (for more details see [58]): suppose $f > 1$. This implies that there will be a time t_0 in the application where we encounter a T gate and must wait for syndrome data to be decoded before continuing. Let Δ_{gen} be the amount of time that the machine must stall for processing this data. During this time an additional $D_1 = r_{\text{gen}} \times \Delta_{\text{gen}}$ bits of syndrome data is generated, which can be processed in time $\Delta_{\text{proc}} = r_{\text{gen}} \Delta_{\text{gen}} / r_{\text{proc}} = f \Delta_{\text{gen}}$. The backlog problem begins to be noticeable at this point, where during processing of the first block D_1 , we generate a *new*

¹Errors commute and can be post-corrected for other gates, but not T -gates.




Fig. 4. Exponential latency overhead when $f = (\frac{r_{\text{gen}}}{r_{\text{proc}}}) > 1$. X-axis shows the compute time if there is no backlog and y-axis shows the actual wall clock time; if there is no backlog we expect wall clock time to be the same as the compute time (line a). Every time we encounter a T-gate we need to decode all the syndromes up until that gate before we can continue the execution [58]. When we encounter the first T-gate at time T_0 , we need to finish the decoding of the data generated during t_0 (not all the data is already decoded as decoding rate is slower than data generation rate) and it takes R_0 to do that. During R_0 where our quantum system is idle, more syndromes are generated and when we encounter the second T-gate at $T_1 + R_0$, we need to finish decoding those syndromes in addition to the syndromes generated during t_1 before continuing the program execution. **The syndrome data generated during the idle periods is the key reason behind data backlog creation which leads to exponential latency overhead.**

block $D_2 = r_{\text{gen}} \times \Delta_{\text{proc}} = fD_1 > D_1$ in size. Then, at the next T gate this process repeats, and we again generate a block of data of size $D_3 = fD_2 = f^2D_1$ bits. Hence, by the k 'th T gate, we generate an overhead of $f^k D_1$ bits to process, exponential in *the decoder's performance ratio*.

As a specific example, consider a multiply-controlled NOT operation on 100 logical qubits from [32]. This algorithm contains ~ 2356 gates, of which 686 are T -gates after decomposition. Assuming that a syndrome generation cycle time is approximately 400 ns [27], and the best prior decoder requires 800 ns to execute [6], the ratio $(r_{\text{gen}}/r_{\text{proc}}) = 2$, and the execution time is intractable.

Figure 5 shows a simulation of real quantum subroutines each composed of a different number of T gates as denoted in Table I. The exponential overhead scaling shows that as decoders become slower than the rate at which data is being generated (which occurs for “syndrome data processing ratios” over 1), the overheads quickly become intractable. Regardless of the effectiveness of the decoder, if it operates at a processing ratio higher than 1 then it will impose exponentially high latency overheads on algorithm execution. The algorithms all draw inspiration from [2]. Barenco-half-dirty-Toffoli is a logarithmic depth multi-control Toffoli gate using $O(n)$ ancilla bits. It performs the same computation as the “cnx-log-depth” gate with a different circuit. The “cnu-half-borrowed” gives an implementation of a multi-control Toffoli using $O(n)$ dirty




Fig. 5. Running times of fault tolerant quantum algorithms with decoders of varying efficiency. The X-axis plots $\frac{r_{\text{gen}}}{r_{\text{proc}}}$. To the left of 1, data is processed as fast as it is generated, whereas rates to the right of 1 indicate that the decoder is slower than syndrome data is generated. The T -gates require synchronization with the decoder in order to execute. Prior work [6] claims that fast neural network inference decoders can perform inference in ~ 800 ns, which places the decoder at approximately the 1.5 - 2 region for a system generating syndromes in the 400-500ns range. Our decoding results show that time to solution never exceeds 20ns, placing it below 1. Clearly computation becomes intractable quickly for slow decoders.

ancilla, meaning the initial states of these bits does not need to be known. The Cuccaro adder is a linear depth implementation of a reversible $A + B$ adder, i.e. two registers of the specified length added together. It has a carry in and a carry out bit as well. The Takahashi adder is an optimized version of the Cuccaro adder [54].

This is the primary motivation for this work – the hardware decoder must be able to execute faster than syndrome data are generated as a prerequisite for tractable fault tolerant computation.

IV. RELATED WORK

Early work focused on the development of and modifications to the minimum-weight perfect matching algorithm (MWPM) [16], [17] to adapt it to surface code decoding [21], [22]. This resulted in a claimed constant time algorithm after parallelization [19].

Other work has constructed maximum likelihood decoders (MLD) based on tensor network contraction [5]. This work is computationally more expensive than minimum-weight perfect matching, but is more accurate.

Neural networks have been explored as possible solutions to the decoding problem as well [1], [6], [60]–[66]. Feed-forward neural networks and recurrent neural networks have been explored in combination with lookup tables to form decoders. The primary distinguishing factor in these systems is that the networks function as *high level decoders* in that they predict both a sequence of error corrections on data qubits along with the existence of a logical error. In this sense, they operate at a higher level than both the MWPM and MLD decoders, seemingly at the cost of execution time with respect to training complexity.

Lastly, more customized algorithms have been developed specifically targeting the surface code decoding problem,

including renormalization group decoders [15], union-find decoding [9], [10], and others [14], [69].

The primary distinguishing factor of our work is that the decoder design is guided by practical system performance. Accuracy has been sacrificed in order to achieve quantum advantage. While the proposed decoder design may not achieve logical error suppression at the same order as some other algorithms, the ability to perform the algorithm in SFQ hardware at or exceeding the speed of syndrome generation is achieved, as is satisfaction of system design constraints.

V. DECODER OVERVIEW AND DESIGN

In this section we describe decoding in terms of a maximum-weight matching problem, followed by details of our approximate decoding algorithm, and demonstrate how we make efficient use of unique features of SFQ gates to implement the algorithm in hardware.

A. Maximum Weight Matching Decoding

The decoding problem requires that the maximally likely set of error chains be reported as a solution, given a particular error syndrome. This can be formulated as a matching problem. Specifically, given an error syndrome string $S \in \{-1, 1\}^n$, we can construct a complete graph on vertices associated with each ancillary qubit that reported an error. The weight of each edge between vertices is proportional to likelihood of a path between these ancillary qubits on the original surface code grid graph. The goal is therefore to find the maximally likely pairing of the syndromes using these weights, one method for doing so is to solve a maximum-weight perfect matching problem.

B. A Greedy Approach

Our decoding algorithm is based upon a greedy approximation to the maximum-weight matching problem. The algorithm calculates all distances $d(v_i, v_j)$ between vertices and sorts them in ascending order $d_1, d_2, \dots, d_{k'}$ where $k' = \binom{k}{2}$. All of the corresponding probability weights are calculated, transforming this ordering to a descending order of likelihood. Then, for each edge e in descending order, add e to the solution M if it forms a matching. This means that it adds another two distinct vertices into M that were not already present. To account for boundary conditions, we introduce a set of external nodes connected to the appropriate sides of the lattice, and connected to one another with weight 0. Under this formulation, the algorithm is a 2-approximation of the optimal solution [13].

C. SFQ-Based Decoder

In this section, we introduce the functional design of our SFQ-based decoder and give some rationale for each aspect of its design. As a reminder, Single Flux Quantum is classical logic implemented in superconducting hardware that does not perform any quantum computation. It is a medium used to express our classical algorithm. The decoder is placed above the quantum chip layer; it receives measurement results from




Fig. 6. Baseline solution to find the two closest hot syndrome modules. Step1: two decoder modules have “1” hot syndrome input. Step2: the hot syndrome modules propagate grow signals. Step3: the grow signals meet at an intermediate module. Step4: the intermediate module sends pair signals in the opposite direction. Step5: pair signals arrive at the hot syndrome modules. Step6: decoding is complete. Note that the decoder modules that receive a pair signal are considered as part of the error chain that has occurred.

ancillary qubits as input, and returns a set of corrections as output. For scalability, our decoder design is built out of a two dimensional array of modules implemented in SFQ logic circuits that we refer to as *decoder modules*. These are connected in a rectilinear mesh topology. Modules are identical and there is one module per each data and ancillary qubit, denoted as *data qubit modules* and *ancilla qubit modules*, respectively. Each decoder module has one input called the *hot syndrome input* that comes from the measurement outcome of the physical quantum bits and determines if the module corresponds to a hot syndrome (note that this input can be “1” only for ancilla qubit modules). Each module contains one output called the *error output* that determines if the module is contained in the error chain (this output can be “1” for all of the decoder modules). In addition, each module has connections to adjacent modules (left, right, up and down).

Our approximate decoder algorithm proceeds as follows. First, the algorithm finds the two modules with “1” hot syndrome input, called *hot syndrome modules*, that are closest together. Next, the algorithm reports the chain of modules connecting them as the correction chain. Finally, it resets the hot syndrome input of the two modules and searches for the next two closest hot syndrome modules. The decoder continues this process until no module with “1” hot syndrome input exists. This is graphically displayed in Figure 6.

Baseline Solution: Our baseline design finds the two closest hot syndrome modules as shown in Figure 6 as follows: 1) every hot syndrome module sends *grow* signals to all the adjacent modules in all four directions; each adjacent module propagates the grow signal in the same direction. Grow signals propagate one step at each cycle. 2) When two grow signals intersect at an *intermediate module*, we generate a set of *pair* signals and back-propagate these to their hot syndrome origins. All of the decoder modules that receive pair signals are




Fig. 7. Scenarios where the SFQ decoder chooses the wrong chain where (a) no reset/boundary/equidistant mechanisms are employed, (b) no boundary/equidistant mechanisms are employed, and (c) no equidistant mechanism is employed.

part of the error chain. Note that more than one intermediate module might exist, however, only one of them is effective and sends the pair signals. For example, in Figure 6, two intermediate modules receive the grow signals, and the decoder is hardwired to be effective (ineffective) when it receives grow signals from up and left directions (down and right directions). Intermediate module refers to the effective one. The baseline solution does not show accuracy or pseudo-threshold behavior and demonstrates poor logical error rate suppression, see the incremental results presented in Section VIII in Figure 10.

Reset Mechanism: One flaw of the baseline system is the lack of a mechanism to reset the decoder modules after two hot syndrome modules are paired. Grow signals of the paired modules continue to propagate, potentially causing these modules to pair incorrectly with other hot syndrome modules, ultimately resulting in an incorrect error chain reported. Figure 7 (a) shows an incorrect matching due to this behavior. To mitigate this, we add a reset mechanism that resets the decoder modules each time hot syndrome modules are paired and the error chain connecting them is determined. Adding the reset mechanism to the baseline system improves the performance somewhat, but does not yet achieve tolerable accuracy.

Boundary Mechanism: Another explanation for the low performance of the baseline solution is that it never pairs hot syndrome modules with boundaries. For example, if two hot syndrome modules are far from each other but are close to boundaries, the error chain with the maximum likelihood is the one that connects the hot syndrome modules to the boundaries. Figure 7 (b) shows this behavior occurring on a machine. We implement a mechanism that enables pairing the hot syndrome modules with boundaries. To do this, we add decoder modules that surround the surface boundaries called *boundary module* (one per each quantum bit located at a boundary). Our solution treats boundary modules as hot syndrome modules but they do not grow and can pair only with non-boundary modules. Note that when two modules are paired, the hot syndrome input of only the non-boundary modules is reset; boundary modules are always treated as hot syndrome modules. Adding the boundary mechanism to the baseline solution augmented with the reset mechanism further increases the accuracy of the decoder.

Equidistant Mechanism: Finally, the last major reason for inefficiency of the baseline is that it does not properly handle the scenarios in which multiple hot syndrome modules are spaced within equal distances of one another, resulting in a

set of pairs that are all equally likely. The baseline solution augmented with reset and boundary mechanisms works properly only if no non-boundary hot syndrome module has an equal distance to more than one other hot syndrome module; otherwise the solution pairs it with all the hot syndrome modules with equal distance. However, this is not the desired output. We need a more intelligent solution to break the tie in the aforementioned scenario, and pair the hot syndrome module to only one other module. This is shown in Figure 7 (c).

To resolve these equidistant degenerate solution sets, we introduce a request – grant policy that allows for the hardware to choose specific subsets of these pairs to proceed. 1) Similar to the baseline solution, the non-boundary hot syndromes first propagate grow signals. 2) An intermediate module receives two grow signals from two different directions, and it sends *pair_request* signals in the opposite directions. *Pair_request* signals continue to propagate until they arrive at a module with “1” hot syndrome input. 3) The modules with “1” hot syndrome input send *pair_grant* signals in the opposite direction of the received *pair_request* signals. Note that multiple *pair_request* signals might arrive at a module with “1” hot syndrome at the same time, but it gives grant to only one of them. 4) An intermediate module receives *pair_grant* signals from two different directions and sends pair signals in the opposite directions. 5) Pair signals continue to propagate until they arrive at a module with “1” hot syndrome input. Boundary modules do not send grow signals but they send *pair_request* signals when they receive grow signals; they also send pair signals when they receive *pair_grant* signals.

VI. IMPLEMENTATION

A. SFQ Implementation of Greedy Decoding

SFQ is a magnetic pulse-based fabric with switching delay of *Ips* and energy consumption of 10^{-19} J per switching. In addition, availability of superconducting microstrip transmission lines in this technology makes it possible to transmit picosecond waves with half of speed of light and without dispersion or attenuation. The combination of these properties together with fast two-terminal Josephson junctions, makes this technology suitable for high speed processing of digital information [31], [35], [40], [55], [67]. SFQ logic families are divided into two groups: ac-biased and dc-biased; Reciprocal Quantum Logic (RQL) [31], and Adiabatic Quantum Flux Parametron (AQFP) [55] are in the first group, and Rapid Single Flux Quantum (RSFQ) [40], Energy-efficient RSFQ (ERSFQ) [35], and energy-efficient SFQ (eSFQ) [67] are examples of the second group. The dc-biased logic family with higher operation speed (as high as 770GHz for a T-Flip Flop (TFF) [7]) and less bias supply issues are more popular than ac-biased logic family.

Our algorithm requires modules to propagate signals one step at each cycle. One approach to implement our algorithm is to use synchronous elements such as flip-flops in decoder modules. However, standard CMOS style flip-flops are very expensive in SFQ logic (e.g., one D-Flip-Flop occupies $72.4\times$

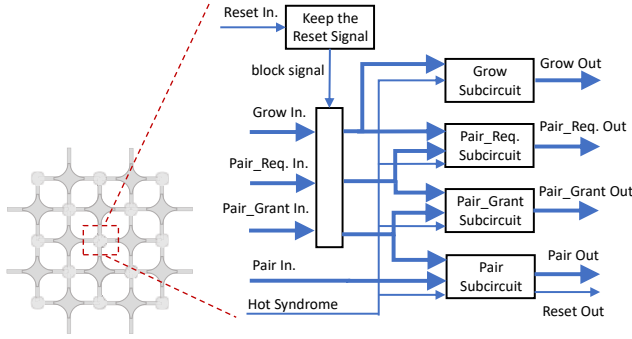


Fig. 8. Overview of decoder module microarchitecture.

more area and consumes $117\times$ more power compared to a 2-input AND gate). On the other hand, SFQ gates have a unique feature that we utilize to implement our algorithm without flip-flops. Unlike CMOS gates, most of the SFQ gates (except for mergers, splitters, TFFs, and I/Os) require a clock signal to operate [48]. Thus, we do not need to have flip-flops and signals can propagate one SFQ gate at each cycle.

As described earlier, our decoder requires resetting the decoder modules each time two hot syndrome modules are paired. We have a global wire that passes through all the modules and is connected to each module using splitter gates. Thus, if we set the value of the global wire, all of the decoder modules receive the reset signal at the same time, as the splitter gates do not require clock signals to operate. If a module receives a reset signal, it blocks the module inputs using 2-input AND gates (one input is *module_input* and the other input is *Reset*). In order to reset a decoder module completely, we need to block the module inputs for as many cycles as the depth of our SFQ-based decoder because the SFQ gates work with clock cycles and one level of gates is reset at each cycle. Thus, we use a simple circuit to keep the reset signal “1” for as many cycles as the circuit depth. In each module, we pass the reset signal that comes from the global wire to a set of m cascaded buffer gates where m is the circuit depth, and the module inputs are blocked if the reset signal that comes from the global wire is “1” or at least one of the buffers has “1” output.

B. Datapath and Subcircuit Design

Figure 8 shows an overview of our decoder module microarchitecture. Our decoder consists of five main subcircuits.

Grow Subcircuit: this subcircuit receives hot syndrome input and 4 grow inputs (from 4 different directions), and produces 4 grow output signals. Grow outputs are “1” if the hot syndrome input is “1” or if the module is passing a grow signal generated by another module.

Pair_Req Subcircuit: this subcircuit is responsible for setting the value of pair_request outputs which are “1” if two grow signals meet at an intermediate module or if the module is passing a pair_request signal that arrived at one of its input ports. The module does not pass the pair_request input signal if the hot syndrome input is “1”; in that case, the module generates a pair_grant signal instead.

Pair_Grant Subcircuit: this module determines the value of pair_grant outputs which are “1” if the module is a hot syndrome module and gives grant to a pair_request signal, or if the module is passing a pair_grant input signal to the adjacent module.

Pair Subcircuit: this subcircuit sets the value of pair outputs which are “1” if two pair_grant signals meet at an intermediate module or if a pair input signal is “1” and the hot syndrome input is not “1”. If both the pair input and hot syndrome input are “1”, the module does not pass the pair signal and instead generates a global reset signal that reset all of the decoder modules and also resets the hot syndrome input. Note that the reset signal resets everything except the subcircuit responsible for passing the pair signals because it is possible that the intermediate module does not have equal distance from the paired hot syndrome modules and we do not want to stop the propagation of all the pair signals in the system when the closer module receives a pair signal (while the farther module has not received a pair signal yet). The SFQ implementation of this subcircuit is shown in Figure 9.

Reset Subcircuit: this subcircuit is responsible to keep the reset signal “1” for as many cycles as the depth of our circuits. The depth is 5 in our circuits, thus reset subcircuit blocks grow, pair_req and pair_grant inputs for 5 cycles in order to reset the module.

VII. METHODOLOGY

Simulation Techniques: In order to effectively benchmark the performance of a stabilizer quantum error correcting code, techniques must be used to simulate the action of the code over many cycles. This is referred to elsewhere in literature as *lifetime simulation* [62], or simply Monte Carlo benchmarking. We constructed a simulation environment that simulates the action of the stabilizer circuits. A cycle refers to one full iteration of the stabilizer circuit. At each step within the cycle, errors are stochastically injected into the qubits and propagated through the circuits. Ancillary qubits are measured, and the outcomes are reported in the error syndrome. This syndrome is then communicated directly to the decoder simulator, which returns the corresponding correction. The correction is applied and the surface is checked for a logical error. The ratio of the number of logical errors to the number of cycles run in simulation is used as the primary performance metric.

Evaluation Performance Metrics: In our evaluations, we use the stabilizer circuits as the primary benchmark. These circuits are replicated for every ancillary qubit present in a surface code lattice. Many different lattices are also analyzed, ranging in size from code distances 3 to 9.

As performance metrics, we focus on *accuracy thresholds* and *pseudo-thresholds*. The former is the physical error rate at which the code begins to suppress errors effectively across multiple code distances. Below this threshold, the logical error rate P_L decreases as the code distance d increases. Above threshold these relationships invert, and P_L grows with d due to decoder performance: the presence of many errors causes the decoding problem to become too complex. In many




Fig. 9. Pair subcircuit after SFQ specific optimizations and mapping. Triangular shapes at the bottom represent the primary inputs of the circuit and those at the top of the circuit show primary outputs. *DFF* is SFQ DRO DFF inserted for path balancing. Splitter (balanced) trees are also shown. Splitter is an asynchronous SFQ gate that receives a pulse at its input and after its intrinsic delay, it produces two almost identical output pulses. We insert splitters at the output of an SFQ gate (or a primary input) with more than one fanout.

Cell	Area (μm^2)	JJ Count	Delay (ps)
AND2	3500	16	8.7
OR2	3500	14	6.0
XOR2	3500	18	6.3
NOT	3500	12	13.0
DRO DFF	3000	11	6.8

TABLE II
THE LIBRARY OF ERSFQ CELLS AND CORRESPONDING CHARACTERISTICS USED FOR SYNTHESIZING THE CIRCUIT INTO SFQ HARDWARE. JOSEPHSON JUNCTION COUNT IS LISTED IN THE SECOND COLUMN.

cases, this leads to corrections that complete what would have otherwise been short error chains, forming logical errors, a process that amplifies as code distances increase.

Pseudo-threshold refers to the performance of a single code distance, and is the physical error rate at which the logical error rate is equal to the physical rate, i.e. $P_L = p$. This can be (and often is) different across different code distances. Better error correcting codes will have higher pseudo-threshold values, as well as higher accuracy thresholds.

Error Models: The Monte Carlo simulation environment requires a model of the errors on the quantum system. We choose to focus on the *depolarizing channel* model [20], [39], [43], [58], parameterized by a single value p : Pauli X , Y , and Z errors occur on qubits with probability $p/3$. During simulation Pauli errors are sampled i.i.d for injection on each data qubit. We present analysis of a variation of the model, the *pure dephasing channel* [9], [10] comprised solely of Pauli Z errors occurring on qubits with probability p . The decoder will be operated symmetrically for both X and Z errors, allowing for simple extrapolation from these results.

Single Flux Quantum Circuit Synthesis: An ERSFQ library of cells is used in this paper to reduce the total power consumption (including the static and dynamic) of the surface code decoder as much as possible. Table II lists characteristics of this library. As seen, this library contains four logic gates including AND2, OR2, XOR2, and NOT, and

it has a Destructive Read-Out D-Flip-Flop (DRO DFF) cell. Area of all logic cells are the same and it is equal to $3500 \mu\text{m}^2$. However, area of the DRO DFF is less than the area of these gates ($3000 \mu\text{m}^2$). DRO DFFs are different from standard CMOS style flip-flops: they are specially designed for SFQ circuits and are usually used for path balancing. In Table II, the total number of Josephson junctions (as a measure of complexity and cost) used in designing each gate together with the intrinsic delay of each cell is reported.

The decoder circuit and its sub-circuits are synthesized by employing ERSFQ specific logic synthesis algorithms and tools [45]–[48]. These algorithms are designed to reduce the complexity of the final synthesized and mapped circuits in terms of total area and Josephson junction count. This is achieved by reducing the required *path balancing* DFF count for realizing these circuits. Please note that for correct operation of dc-biased SFQ (including ERSFQ) circuits, these circuits should be fully path balanced; this means that in a Directed Acyclic Graph (DAG) that represents an SFQ circuit, length of any path from any primary input to any primary output in terms of the gate count should be the same. In most of the SFQ circuits this property does not hold in the beginning. Therefore, some path balancing DFFs should be inserted into shorter paths to maintain the full path balancing property. In the algorithms we employed for mapping these circuits, a dynamic programming approach is used to ensure minimization of the total number of DFFs to maintain the balancing property [46], [47]. In addition, a depth minimization algorithm together with path balancing is employed [48] to reduce the logical depth (length of the longest path from any primary input to any primary output in terms of the gate count) of the final mapped circuit. This helps to reduce the latency of the mapped SFQ circuit. As mentioned before, SFQ logic gates are pulsed-based, meaning

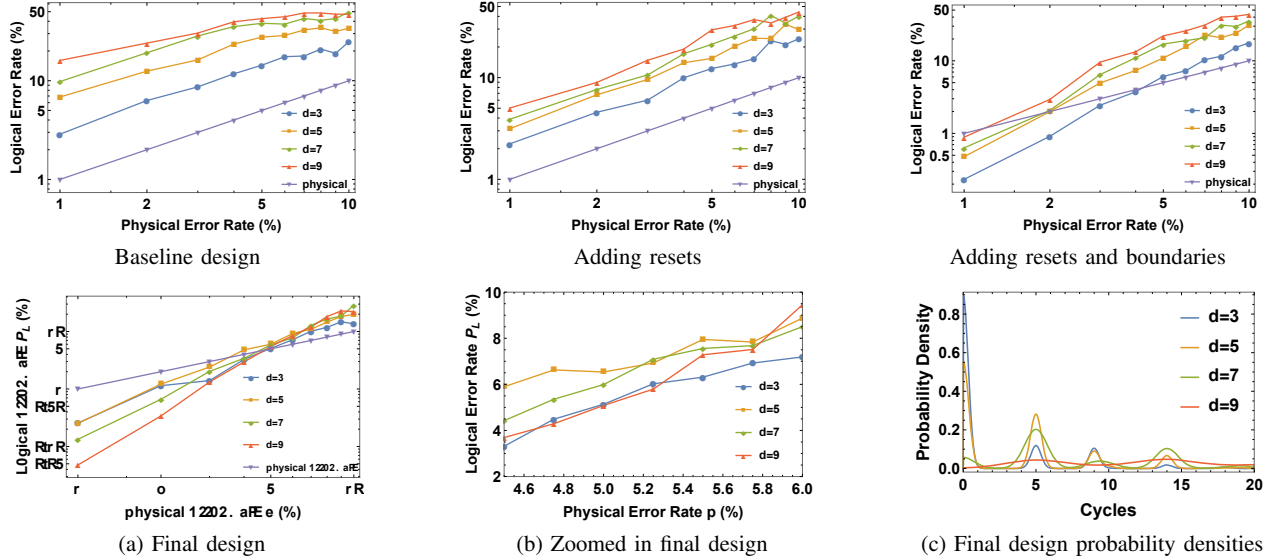


Fig. 10. Top row: Logical error rate performance of each incremental design step. The addition of resets and boundaries each contribute heavily to the realization of pseudo-thresholds, and have a dramatic effect on reducing the minimum achievable logical error rates for each code distance. Bottom row: Results for our final design, including support for reset, boundary, and equidistant mechanisms. (a) Error rate scaling for the proposed decoder. An accuracy threshold is evident at approximately 5% physical error rate, while pseudo-thresholds span the range from $\sim 3.5\% - 5\%$. (b) Logical error rates near the 5% physical error rate value. (c) Truncated unnormalized estimated probability distributions for the execution cycles required by each code distance in simulation. Window shows up to 20 cycles for comparison across code distances. Notice that while distances 3, 5, 7 display peaks centered at 0, 5, 9, and 14 cycles.

that the presence of a pulse represents a logic-“1” and the absence of a pulse represents a logic-“0”. Each gate is clocked, and as an example, the SFQ NOT gate behaves as follows. After the clock pulse arrives, when there are no input pulses, a pulse is generated at the output of the gate representing a “1”. On the other hand, when there is an input pulse, no pulses are generated at the output, meaning a “0”. Each pulse is a single quantum of magnetic flux ($\phi_0 = \frac{h}{2e} = 2.07\text{mV}\times\text{ps}$) [40]. To simulate the SFQ circuits for verifying their correct functionality, we use the Josephson simulator (JSIM) [11].

VIII. EVALUATIONS

In this section we evaluate the performance of our proposed decoder design, both in terms of circuit characteristics including power, area, and latency, as well as error correction performance metrics of accuracy and pseudo-thresholds. We also analyze the execution time of our system, relying upon described operating assumptions and circuit synthesis results.

Threshold Evaluations: To gauge the performance of our design, we use the threshold metrics described in Section VII. Figure 10 (a) shows the central performance result, while the top row of Figure 10 shows the effect of all of the incremental design decisions on the overall performance. This evaluation simulates the performance of the decoder across a range of physical error rates. A pseudo-threshold range of between 3.5% and 5% is observed, and an accuracy threshold appears at approximately the 5% error rate. For code distance 5, the pseudo-threshold is below the accuracy threshold. This highlights the difference between these metrics – an error correcting protocol like the surface code can perform well even

though particular code distances may still be amplifying the physical error rates (i.e. $P_L > p$). It is important to consider both types of thresholds when evaluating decoder performance.

An interesting behavior is observed for code distance $d = 3$. This lattice performs at or surpassing the performance of all other lattices from the 3% physical error rate and above. Below this point, the lattice begins to taper off, and ultimately it converges with the distance 5 lattice. Boundary conditions were highly prioritized in our design, causing this effect. In particular, the decoder is designed such that error chains that terminate at the boundaries are more likely to be correctly identified than other patterns. This choice was made as smaller lattices are more dominated by these edge effects than larger lattices. The smallest lattice in our simulations shows this anomalous behavior, as it contains a disproportionate amount of boundary patterns. In larger lattices, syndromes are less likely to terminate in boundaries, reducing this effect.

Figure 10 (b) highlights the desired threshold behavior. Examining the 6% error rate, code distance 9 is outperformed by code distance 7. Moving to the lowest physical error rate in the window, we find that the lattices perform in the order $d = 9, d = 3, d = 7, \text{ and } d = 5$, ordered from lowest to highest logical error rate. Barring the anomalous $d = 3$ behavior described above, this is accuracy threshold behavior indicative of successful error correction performance.

Performance Analysis: To quantify the approximation factor of our design we compare the performance to that of an ideal decoder by fitting to an exponential analytical model. The achievable error rates by the surface code ideally can be described by $P_L \approx 0.03(p/p_{\text{th}})^d$ [20] when a minimum

Circuit	Logical Depth	Latency (ps)	Total Area (μm^2)	Power Consumption (μW)
AND_GATE	1	8.7	3500	0.026
OR_GATE	1	6.0	3500	0.026
OR_GATE_7_INPUTS	3	18.0	33000	0.338
NOT_GATE	1	13.0	3500	0.026
Pair_Grant Subcircuit	5	115.0	293500	3.38
Pair Subcircuit	5	115.0	303500	3.53
Pair_Req./Grow Subcircuit	5	115.0	406500	4.75
Full_Circuit	6	168.0	1143000	13.44

TABLE III

EXPERIMENTAL SYNTHESIS RESULTS FOR THE SFQ DECODER. SHOWN ARE ALL GATES UTILIZED IN THE SYNTHESIS, AS WELL AS SUBMODULES THAT COMPRISE THE MAIN CIRCUIT. PAIR_REQ. AND GROW SUBCIRCUITS HAVE BEEN COMBINED INTO A SINGLE SUBCIRCUIT.

Code Distance	Max	Average	Standard Deviation
3	3.86	0.29	0.59
5	9.58	0.74	1.13
7	14.7	2.06	2.05
9	19.8	3.93	3.21

TABLE IV

DECODER EXECUTION TIME IN NANoseconds ACROSS EACH CODE DISTANCE STUDIED AND ACROSS ALL SIMULATED ERROR RATES.

Code Distance	3	5	7	9
c_2	0.650	0.429	0.306	0.323

TABLE V

EMPIRICAL PARAMETER ESTIMATION GIVEN A MODEL OF THE FORM $P_L \approx c_1(p/p_{th})^{c_2 \cdot d}$. SHOWN ARE ESTIMATED c_2 PARAMETER VALUES.

weight matching decoder is used in software. This model is valid for the error models we consider, as [20] uses and fits “class-0” Pauli errors in the same fashion. Using a model of the form $P_L \approx c_1(p/p_{th})^{c_2 \cdot d}$, we fit values of c_1, c_2 for each code distance at physical error rates below accuracy threshold, and collect c_2 values in Table V. C_2 coefficients describe the *effective code distance* for our system, and capture the approximation factor we introduce. For code distances 3 and 5, we find that the approximate decoder is roughly 65% and 43% of the optimal distance respectively. This is the trade-off made by our system in order to fit the timing and physical footprint requirements of the system.

Notice that this accuracy tradeoff results a net resource reduction for our design over other proposed designs as shown in Figure 11. The data backlog imposes delays into the system that decrease the logical accuracy of any decoder that incurs this backlog. As the backlog builds up, the number of required syndrome detection cycles builds up as well, resulting in a new effective logical error rate as one logical gate now requires many more syndrome detection cycles to occur. The SFQ decoder pays an accuracy price for speed, but when the backlog is taken into consideration this tradeoff results in a significant performance gain over alternative designs.

Synthesis Results and Circuit Characterization: Table III shows experimental results for the surface code decoder circuit presented in this paper using the aforementioned ERSFQ library of cells described in Section VII. The full circuit demonstrates a cycle latency of 168 ps, and an area and power footprint of 1.143 mm^2 and $13.44 \mu\text{W}$, respectively. The full

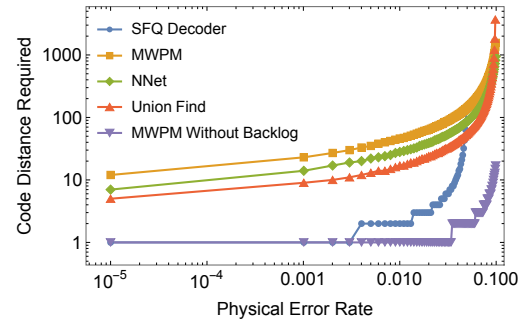


Fig. 11. Comparison of required code distances of different decoders to execute an algorithm consisting of 100 T-gates. Compared are the SFQ Decoder, minimum weight perfect matching decoder (MWPM) [20], neural network decoder [1], union find decoder [9], and a theoretical MWPM decoder with no backlog, across both code distances and physical error rates.

decoder is comprised of a mesh of these circuit modules, requiring a single module per individual qubit. This means that for systems of code distance 9 comprised of 289 qubits, the decoder required will be of size 330.33 mm^2 and will dissipate 3.88 mW of power. Typical dilution refrigerators are capable of cooling up to 1 – 2 Watts of power in the 4-Kelvin temperature region [33], enabling the co-location of a decoder mesh of size 87×87 , which would protect a single qubit of code distance $d = 44$, or 100 qubits of code distance $d = 5$. These values are estimations given modern day SFQ and cryogenic dilution refrigerator technology, much of which is subject to change in the future.

Execution Time Evaluation: The most important characteristic that the SFQ decoder aims to optimize is real-time execution speed. Previous works have described the syndrome generation time to be between 160 – 800 ns for superconducting devices that we are focusing on in this study. [27], [57].

In practice the time to solution is much lower than the upper bound of $O(n)$ on the greedy algorithm. Table IV contains the empirically observed statistics of our decoder operation. The maximum cycles to solution is well approximated by a linear scaling with a leading coefficient of ~ 15.75 . Estimated probability distributions describing the required cycles to solution for each code distance are shown in Figure 10 (c).

Comparison to existing approximation techniques: Trad-

ing the accuracy for decoding speed has been utilized in prior work. Union-find [9] achieves a significant speed-up over the minimum weight perfect matching algorithm, while the accuracy threshold decreases by only 0.4%. Despite this, the union-find decoding time is still longer than the syndrome generation time ($> 2X$ longer) thus exposing it to the exponential latency overhead caused by the data backlog. In contrast to prior approximation techniques, decoding time in our design is faster than syndrome generation time and thus it does not incur exponential latency overhead, enabling a practical implementation of error-correcting quantum machines.

Effect on SQV: The net effect of our design is to expand the SQV achievable by near-term machines. An example of this is a small 100 physical qubit system in which 10^3 gates are performed per qubit, a machine that is conjectured will exist in the near future [4] and shown in red in Figure 1. By utilizing the scaling equation described in Section VIII, we see that a homogeneous system of 78 logical qubits each of code distance 3 is capable of performing $\sim 4.36 \times 10^6$ gates per qubit. This expands SQV from $100 \times 10^3 \rightarrow 78 \times (4.36 \times 10^6) \approx 3.4 \times 10^8$, increasing by a factor of 3402. This can be pushed farther by going to the small qubit count limit, constructing a machine of 40 logical qubits each of code distance 5 with logical error rate 8.96×10^{-10} , yielding SQV of $\approx 1.12 \times 10^9$, an increase of 11, 163. These effects are captured in Figure 1. Not all applications benefit from these expansions in the same fashion, but our techniques allow for machines to be used in ways that are tailored to individual applications, and enable much more computation to be performed on the same machine.

IX. CONCLUSION

In the design of near-term quantum computers, it is vital to enable the machines to perform as much computation as possible. By taking inspiration from quantum error correction, we have designed an “Approximate Quantum Error Correction” error mitigation technique that expands the “Simple Quantum Volume” of near-term machines by factors between 3,402 and 11,163. Our design focuses on the construction of an approximate surface code decoder that can rapidly decode error syndrome, at the cost of accuracy.

Using SFQ synthesis tools, we show that the area and power are within the typical cryogenic cooling system budget. In addition, our accelerator is based on a modular, scalable architecture that uses one decoder module per each qubit. Most importantly, our decoder constructs solutions in real-time, requiring a maximum of ~ 20 ns to compute the solution in simulation. This allows our decoding accelerator to achieve 10x smaller code distance when compared to offline decoders when accounting for decoding backlog. Thus, it is a technique that can effectively boost the Simple Quantum Volume of near-term machines.

ACKNOWLEDGMENT

This work is funded in part by EPiQC, an NSF Expedition in Computing, under grants CCF-1730449; in part by STAQ

under grant NSF Phy-1818914; in part by DOE grants DE-SC0020289 and DE-SC0020331; in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via the U.S. Army Research Office grant W911NF-17-1-0120.

REFERENCES

- [1] P. Baireuther, M. Caio, B. Criger, C. W. Beenakker, and T. E. O’Brien, “Neural network decoder for topological color codes with circuit level noise,” *New Journal of Physics*, vol. 21, no. 1, p. 013003, 2019.
- [2] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, “Elementary gates for quantum computation,” *Physical review A*, vol. 52, no. 5, p. 3457, 1995.
- [3] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell *et al.*, “Superconducting quantum circuits at the surface code threshold for fault tolerance,” *Nature*, vol. 508, no. 7497, p. 500, 2014.
- [4] L. S. Bishop, S. Bravyi, A. Cross, J. M. Gambetta, and J. Smolin, “Quantum volume,” 2017.
- [5] S. Bravyi, M. Suchara, and A. Vargo, “Efficient algorithms for maximum likelihood decoding in the surface code,” *Physical Review A*, vol. 90, no. 3, p. 032326, 2014.
- [6] C. Chamberland and P. Ronagh, “Deep neural decoders for near term fault-tolerant experiments,” *Quantum Science and Technology*, vol. 3, no. 4, p. 044002, 2018.
- [7] W. Chen, A. Rylyakov, V. Patel, J. Lukens, and K. Likharev, “Rapid single flux quantum t-flip flop operating up to 770 ghz,” *IEEE Transactions on Applied Superconductivity*, vol. 9, no. 2, pp. 3212–3215, 1999.
- [8] J. M. Chow, J. M. Gambetta, A. Córcoles, S. T. Merkel, J. A. Smolin, C. Rigetti, S. Poletto, G. A. Keefe, M. B. Rothwell, J. Rozen *et al.*, “Universal quantum gate set approaching fault-tolerant thresholds with superconducting qubits,” *Physical review letters*, vol. 109, no. 6, p. 060501, 2012.
- [9] N. Delfosse and N. H. Nickerson, “Almost-linear time decoding algorithm for topological codes,” *arXiv preprint arXiv:1709.06218*, 2017.
- [10] N. Delfosse and G. Zémor, “Linear-time maximum likelihood decoding of surface codes over the quantum erasure channel,” *arXiv preprint arXiv:1703.01517*, 2017.
- [11] J. A. Delpert, K. Jackman, P. Le Roux, and C. J. Fourie, “Josim - superconductor spice simulator,” *IEEE Transactions on Applied Superconductivity*, vol. 29, no. 5, pp. 1–5, 2019.
- [12] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, “Topological quantum memory,” *Journal of Mathematical Physics*, vol. 43, no. 9, pp. 4452–4505, 2002.
- [13] D. E. Drake and S. Hougardy, “A simple approximation algorithm for the weighted matching problem,” *Information Processing Letters*, vol. 85, no. 4, pp. 211–213, 2003.
- [14] G. Duclos-Cianci and D. Poulin, “Fast decoders for topological quantum codes,” *Physical review letters*, vol. 104, no. 5, p. 050504, 2010.
- [15] G. Duclos-Cianci and D. Poulin, “A renormalization group decoding algorithm for topological quantum codes,” in *2010 IEEE Information Theory Workshop*. IEEE, 2010, pp. 1–5.
- [16] J. Edmonds, “Maximum matching and a polyhedron with 0, 1-vertices,” *Journal of research of the National Bureau of Standards B*, vol. 69, no. 125-130, pp. 55–56, 1965.
- [17] J. Edmonds, “Paths, trees, and flowers,” *Canadian Journal of mathematics*, vol. 17, pp. 449–467, 1965.
- [18] C. Figgatt, A. Ostrander, N. M. Linke, K. A. Landsman, D. Zhu, D. Maslov, and C. Monroe, “Parallel entangling operations on a universal ion-trap quantum computer,” *Nature*, vol. 572, no. 7769, pp. 368–372, 2019.
- [19] A. G. Fowler, “Minimum weight perfect matching of fault-tolerant topological quantum error correction in average $o(1)$ parallel time,” *arXiv preprint arXiv:1307.1740*, 2013.
- [20] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, “Surface codes: Towards practical large-scale quantum computation,” *Physical Review A*, vol. 86, no. 3, p. 032324, 2012.
- [21] A. G. Fowler, A. C. Whiteside, and L. C. Hollenberg, “Towards practical classical processing for the surface code,” *Physical review letters*, vol. 108, no. 18, p. 180501, 2012.

- [22] A. G. Fowler, A. C. Whiteside, and L. C. Hollenberg, "Towards practical classical processing for the surface code: timing analysis," *Physical Review A*, vol. 86, no. 4, p. 042313, 2012.
- [23] A. G. Fowler, A. C. Whiteside, A. L. McInnes, and A. Rabbani, "Topological code autotune," *Physical Review X*, vol. 2, no. 4, p. 041003, 2012.
- [24] D. P. Franke, J. S. Clarke, L. M. Vandersypen, and M. Veldhorst, "Rent's rule and extensibility in quantum computing," *arXiv preprint arXiv:1806.02145*, 2018.
- [25] X. Fu, M. Rol, C. Bultink, J. van Someren, N. Khammassi, I. Ashraf, R. Vermeulen, J. de Sterke, W. Vlothuizen, R. Schouten *et al.*, "A microarchitecture for a superconducting quantum processor," *IEEE Micro*, vol. 38, no. 3, pp. 40–47, 2018.
- [26] X. Fu, M. A. Rol, C. C. Bultink, J. Van Someren, N. Khammassi, I. Ashraf, R. Vermeulen, J. De Sterke, W. Vlothuizen, R. Schouten *et al.*, "An experimental microarchitecture for a superconducting quantum processor," in *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture*. ACM, 2017, pp. 813–825.
- [27] J. Ghosh, A. G. Fowler, and M. R. Geller, "Surface code with decoherence: An analysis of three superconducting architectures," *Physical Review A*, vol. 86, no. 6, p. 062318, 2012.
- [28] D. Gottesman, "Stabilizer codes and quantum error correction," *arXiv preprint quant-ph/9705052*, 1997.
- [29] H. Häffner, W. Hänsel, C. Roos, J. Benhelm, M. Chwalla, T. Körber, U. Rapol, M. Riebe, P. Schmidt, C. Becher *et al.*, "Scalable multiparticle entanglement of trapped ions," *Nature*, vol. 438, no. 7068, p. 643, 2005.
- [30] M. B. Hastings, D. Wecker, B. Bauer, and M. Troyer, "Improving quantum algorithms for quantum chemistry," *arXiv preprint arXiv:1403.1539*, 2014.
- [31] Q. P. Herr, A. Y. Herr, O. T. Oberg, and A. G. Ioannidis, "Ultra-low-power superconductor logic," *Journal of applied physics*, vol. 109, no. 10, p. 103903, 2011.
- [32] A. Holmes, S. Johri, G. G. Guerreschi, J. S. Clarke, and A. Matsuura, "Impact of qubit connectivity on quantum algorithm performance," *arXiv preprint arXiv:1811.02125*, 2018.
- [33] J. Hornibrook, J. Colless, I. C. Lamb, S. Pauka, H. Lu, A. Gossard, J. Watson, G. Gardner, S. Fallahi, M. Manfra *et al.*, "Cryogenic control architecture for large-scale quantum computing," *Physical Review Applied*, vol. 3, no. 2, p. 024010, 2015.
- [34] J. Kelly, R. Barends, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, I.-C. Hoi, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner, A. N. Cleland, and J. M. Martinis, "State preservation by repetitive error detection in a superconducting quantum circuit," *Nature*, vol. 519, no. 7541, pp. 66–69, 2015. [Online]. Available: <http://www.nature.com/doi/10.1038/nature14270>
- [35] D. Kirichenko, S. Sarwana, and A. Kirichenko, "Zero static power dissipation biasing of rsfq circuits," *IEEE Transactions on Applied Superconductivity*, vol. 21, no. 3, pp. 776–779, 2011.
- [36] B. Lekitsch, S. Weidt, A. G. Fowler, K. Mølmer, S. J. Devitt, C. Wunderlich, and W. K. Hensinger, "Blueprint for a microwave trapped ion quantum computer," *Science Advances*, vol. 3, no. 2, p. e1601540, 2017.
- [37] J. E. Levy, M. S. Carroll, A. Ganti, C. A. Phillips, A. J. Landahl, T. M. Gurrieri, R. D. Carr, H. L. Stalford, and E. Nielsen, "Implications of electronics constraints for solid-state quantum error correction and quantum circuit failure probability," *New Journal of Physics*, vol. 13, no. 8, p. 083021, 2011.
- [38] J. E. Levy, A. Ganti, C. A. Phillips, B. R. Hamlet, A. J. Landahl, T. M. Gurrieri, R. D. Carr, and M. S. Carroll, "The impact of classical electronics constraints on the solid-state logical qubit memory," *arXiv preprint arXiv:0904.0003*, 2009.
- [39] D. A. Lidar and T. A. Brun, *Quantum error correction*. Cambridge university press, 2013.
- [40] K. K. Likharev and V. K. Semenov, "Rsfq logic/memory family: A new josephson-junction technology for sub-terahertz-clock-frequency digital systems," *IEEE Transactions on Applied Superconductivity*, vol. 1, no. 1, pp. 3–28, 1991.
- [41] N. M. Linke, D. Maslov, M. Roetteler, S. Debnath, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, "Experimental comparison of two quantum computing architectures," *Proceedings of the National Academy of Sciences*, vol. 114, no. 13, pp. 3305–3310, 2017.
- [42] D. Maslov, "Basic circuit compilation techniques for an ion-trap quantum machine," *New Journal of Physics*, vol. 19, no. 2, p. 023035, 2017.
- [43] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information*. Cambridge university press, 2010.
- [44] G. Paraoanu, "Microwave-induced coupling of superconducting qubits," *Physical Review B*, vol. 74, no. 14, p. 140504, 2006.
- [45] G. Pasandi and M. Pedram, "Balanced factorization and rewriting algorithms for synthesizing single flux quantum logic circuits," in *Proceedings of the 2019 on Great Lakes Symposium on VLSI (GLSVLSI)*, 2019, pp. 183–188.
- [46] G. Pasandi and M. Pedram, "A dynamic programming-based path balancing technology mapping algorithm targeting area minimization," in *Proc. IEEE/ACM Int. Conf. Comput. Aided Des. (ICCAD)*, 2019.
- [47] G. Pasandi and M. Pedram, "PBMap: A path balancing technology mapping algorithm for single flux quantum circuits," *IEEE Transactions on Applied Superconductivity*, vol. 29, no. 4, pp. 1–14, 2019.
- [48] G. Pasandi, A. Shafaei, and M. Pedram, "SFQmap: A technology mapping tool for single flux quantum logic circuits," in *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 2018, pp. 1–5.
- [49] B. Patra, R. M. Incandela, J. P. Van Dijk, H. A. Homulle, L. Song, M. Shahmohammadi, R. B. Staszewski, A. Vladimirescu, M. Babaie, F. Sebastiano *et al.*, "Cryo-cmos circuits and systems for quantum computing applications," *IEEE Journal of Solid-State Circuits*, vol. 53, no. 1, pp. 309–321, 2018.
- [50] J. Plantenberg, P. De Groot, C. Harmans, and J. Mooij, "Demonstration of controlled-not quantum gates on a pair of superconducting quantum bits," *Nature*, vol. 447, no. 7146, p. 836, 2007.
- [51] J. Preskill, "Quantum computing in the nisc era and beyond," *Quantum*, vol. 2, p. 79, 2018.
- [52] F. Sebastiano, H. Homulle, B. Patra, R. Incandela, J. van Dijk, L. Song, M. Babaie, A. Vladimirescu, and E. Charbon, "Cryo-cmos electronic control for scalable quantum computing," in *Proceedings of the 54th Annual Design Automation Conference 2017*. ACM, 2017, p. 13.
- [53] K. M. Svore and M. Troyer, "The quantum future of computation," *Computer*, vol. 49, no. 9, pp. 21–30, 2016.
- [54] Y. Takahashi, S. Tani, and N. Kunihiro, "Quantum addition circuits and unbounded fan-out," *arXiv preprint arXiv:0910.2530*, 2009.
- [55] N. Takeuchi, D. Ozawa, Y. Yamanashi, and N. Yoshikawa, "An adiabatic quantum flux parametron as an ultra-low-power logic device," *Superconductor Science and Technology*, vol. 26, no. 3, p. 035010, 2013.
- [56] S. S. Tannu, D. M. Carmean, and M. K. Qureshi, "Cryogenic-dram based memory system for scalable quantum computers: a feasibility study," in *Proceedings of the International Symposium on Memory Systems*. ACM, 2017, pp. 189–195.
- [57] S. S. Tannu, Z. A. Myers, P. J. Nair, D. M. Carmean, and M. K. Qureshi, "Taming the instruction bandwidth of quantum computers via hardware-managed error correction," in *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture*. ACM, 2017, pp. 679–691.
- [58] B. M. Terhal, "Quantum error correction for quantum memories," *Reviews of Modern Physics*, vol. 87, no. 2, p. 307, 2015.
- [59] Y. Tomita and K. M. Svore, "Low-distance surface codes under realistic quantum noise," *Physical Review A*, vol. 90, no. 6, p. 062320, 2014.
- [60] G. Torlai and R. G. Melko, "Neural decoder for topological codes," *Physical review letters*, vol. 119, no. 3, p. 030501, 2017.
- [61] S. Varsamopoulos, K. Bertels, and C. G. Almudever, "Designing neural network based decoders for surface codes," *arXiv preprint arXiv:1811.12456*, 2018.
- [62] S. Varsamopoulos, K. Bertels, and C. G. Almudever, "Designing neural network based decoders for surface codes," *arXiv preprint arXiv:1811.12456*, 2018.
- [63] S. Varsamopoulos, K. Bertels, and C. G. Almudever, "Decoding surface code with a distributed neural network based decoder," *arXiv preprint arXiv:1901.10847*, 2019.
- [64] S. Varsamopoulos, K. Bertels, and C. G. Almudever, "Comparing neural network based decoders for the surface code," *IEEE Transactions on Computers*, 2019.
- [65] S. Varsamopoulos, B. Criger, and K. Bertels, "Decoding small surface codes with feedforward neural networks," *Quantum Science and Technology*, vol. 3, no. 1, p. 015004, 2017.
- [66] S. Varsamopoulos, B. Criger, and K. Bertels, "Decoding small surface codes with feedforward neural networks," *Quantum Science and Technology*, vol. 3, no. 1, p. 015004, 2017.
- [67] M. H. Volkmann, A. Sahu, C. J. Fourie, and O. A. Mukhanov, "Experimental investigation of energy-efficient digital circuits based on esfq

- logic," *IEEE Transactions on Applied Superconductivity*, vol. 23, no. 3, pp. 1 301 505–1 301 505, 2013.
- [68] F. Ware, L. Gopalakrishnan, E. Linstadt, S. A. McKee, T. Vogelsang, K. L. Wright, C. Hampel, and G. Bronner, "Do superconducting processors really need cryogenic memories?: the case for cold dram," in *Proceedings of the International Symposium on Memory Systems*. ACM, 2017, pp. 183–188.
- [69] J. Wootton, "A simple decoder for topological codes," *Entropy*, vol. 17, no. 4, pp. 1946–1957, 2015.
- [70] C.-P. Yang, S.-I. Chu, and S. Han, "Possible realization of entanglement, logical gates, and quantum-information transfer with superconducting-quantum-interference-device qubits in cavity qed," *Physical Review A*, vol. 67, no. 4, p. 042311, 2003.
- [71] D. Zajac, T. Hazard, X. Mi, E. Nielsen, and J. Petta, "Scalable gate architecture for a one-dimensional array of semiconductor spin qubits," *Physical Review Applied*, vol. 6, no. 5, p. 054013, 2016.