

Delft University of Technology

Neural-Network Decoders for Quantum Error Correction Using Surface Codes A Space Exploration of the Hardware Cost-Performance Tradeoffs

Overwater, Ramon W.J.; Babaie, Masoud; Sebastiano, Fabio

DOI 10.1109/TQE.2022.3174017

Publication date 2022 **Document Version** Final published version

Published in IEEE Transactions on Quantum Engineering

Citation (APA) Overwater, R. W. J., Babaie, M., & Sebastiano, F. (2022). Neural-Network Decoders for Quantum Error Correction Using Surface Codes: A Space Exploration of the Hardware Cost-Performance Tradeoffs. *IEEE Transactions on Quantum Engineering*, *3*, Article 3101719. https://doi.org/10.1109/TQE.2022.3174017

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Received February 25, 2022; revised April 8, 2022; accepted April 27, 2022; date of publication May 10, 2022; date of current version June 3, 2022.

Digital Object Identifier 10.1109/TQE.2022.3174017

Neural-Network Decoders for Quantum Error Correction Using Surface Codes: A Space Exploration of the Hardware Cost-Performance Tradeoffs

RAMON W. J. OVERWATER^{1,2}, MASOUD BABAIE^{1,2} (Member, IEEE), AND FABIO SEBASTIANO^{1,2} (Senior Member, IEEE)

¹QuTech, Delft University of Technology, 2600 Delft, GA, The Netherlands ²Department of Quantum and Computer Engineering, Delft University of Technology, 2600 Delft, GA, The Netherlands

Corresponding authors: Ramon W. J. Overwater; Fabio Sebastiano (e-mail: r.w.j.overwater@tudelft.nl; f.sebastiano@tudelft.nl).

This work was supported by Intel. All data and code are available [1] at doi: 10.4121/16539786.

ABSTRACT Quantum error correction (QEC) is required in quantum computers to mitigate the effect of errors on physical qubits. When adopting a QEC scheme based on surface codes, error decoding is the most computationally expensive task in the classical electronic back-end. Decoders employing neural networks (NN) are well-suited for this task but their hardware implementation has not been presented yet. This work presents a space exploration of fully connected feed-forward NN decoders for small distance surface codes. The goal is to optimize the NN for the high-decoding performance, while keeping a minimalistic hardware implementation. This is needed to meet the tight delay constraints of real-time surface code decoding. We demonstrate that hardware-based NN-decoders can achieve the high-decoding performance comparable to other state-of-the-art decoding algorithms whilst being well below the tight delay requirements (≈ 440 ns) of current solid-state qubit technologies for both application-specific integrated circuit designs (< 30 ns) and field-programmable gate array implementations (< 90 ns). These results indicate that NN-decoders are viable candidates for further exploration of an integrated hardware implementation in future large-scale quantum computers.

INDEX TERMS Application-specific integrated circuit (ASIC), complementary metal-oxide semiconductor (CMOS), CMOS integrated circuits, combinational circuits, cryo-CMOS decoding, cryogenic electronics, digital integrated circuits, error correction codes, feedforward neural networks (NNs), field programmable gate array (FPGA), fixed-point arithmetic, machine learning, NNs, pareto analysis, quantum computing, quantum-error-correction (QEC) codes, supervised learning, surface codes (SCs).

I. INTRODUCTION

For certain problems, quantum-computing algorithms have been demonstrated to run with polynomial time complexity, where classical counterparts would scale with an exponential time complexity [2]–[5]. This speed-up is ascribed to the use of quantum bits (qubits) that, unlike classical bits, can exploit quantum effects, such as superposition, entanglement, and interference [6], [7]. Unfortunately, the information stored in the qubits can be lost via decoherence, due to their sensitivity to their environment. The errors due to decoherence can be mitigated by adopting quantum-error-correction (QEC) schemes that encode multiple imperfect *physical* qubits into a *logical* quantum state, similar to classical error correction. However, while classical bits can be simply copied to introduce redundancy, the quantum no-cloning theorem prevents the copying of qubits [8], [9], thus calling for ad-hoc QEC schemes.

The surface code (SC), a planar form of the toric code [10], is among the most popular QEC schemes thanks to its high error threshold, scalable 2-D structure and the need for only next-neighbor interactions [11]. This makes it suitable for integration in promising solid-state qubit technologies, such as superconducting qubits [12], [13] and quantum-dot-based qubits [14]. Although encoding a logical state in



FIGURE 1. Schematic representation of the four smallest distances of the rotated SC. The white dots represent the data qubits, the blue dots the *X*-ancillas and the red dots the *Z*-ancillas. The connections between the qubits correspond to the local interactions whilst performing the measurement round.

an SC is straightforward, detecting the errors occurring on the physical qubits typically requires a complex decoder [15]-[18], as physical qubits cannot be directly measured without losing quantum information. In addition to the computational complexity of QEC decoding algorithms, the decoder should run orders of magnitude faster than the decoherence process affecting the physical qubits, with a required execution time well below 1 μ s for typical solidstate qubits. This stringent timing requirement has raised the question whether the decoders need to be implemented in hardware instead of running in software for even faster inference [19]. Furthermore, hardware decoders would be preferred to support the scalability of quantum computers. Promising candidates for large-scale quantum computers comprise large arrays of cryogenic solid-state qubits controlled by local electronics also operating at cryogenic temperatures to ensure compactness and reliability by avoiding long interconnects between several temperature stages [20]–[28]. Thus, the QEC decoder must also run at cryogenic temperature and an integrated hardware implementation is favorable to minimize the area occupation (for compactness) and the power dissipation (to comply with the limited cooling budget of cryogenic refrigerators). Recent work has demonstrated hardware-based decoders that run fast enough [29]–[32], but research is lacking on the hardware implementation of a decoding solution that has the potential to outperform them all: neural networks (NNs).

NN decoders have attracted large interest, thanks to their fast and constant inference time and state-of-the-art decoding performance [19], [33]–[40]. The hardware requirements for NN decoders have been estimated before [19], [38], but the tradeoffs between hardware cost and performance have not been explored. This work bridges this gap by focusing on the hardware implementation of an NN decoder for SC QEC [33]. First, the relation between decoder performance and the NN design parameters, such as the number of layers and their size, the neuron transfer function, signal quantization, and symmetries, are explored. Then, the tradeoffs between the decoding performance (error rate, computing delay) and the hardware cost (area, power) are evaluated for an implementation on both an application-specific integrated circuit (ASIC) and a commercial FPGA, with explicit attention to the cryogenic operation of both platforms.

This work demonstrates that hardware NN-based decoders can achieve the high-decoding performance comparable to other state-of-the-art decoding algorithms while satisfying with ample margin the tight delay requirements of current solid-state qubit technologies. The hardware cost in terms of silicon area and power quickly increases with NN size and SC distance. The obtained decoding times are low enough for future work to explore further optimization of the hardware costs.

The rest of this article is organized as follows. First, Section II gives a short background on decoding the SC. This is followed by Section III, which shows the proposed decoder. Next, Section IV outlines the simulation setup. The decoding performance results and design space exploration are shown in Section V. The results are combined with the hardware cost estimations in Section VI. The results are then discussed in Section VII. Finally, Section VIII concludes this article.

II. DECODING THE SC

The SC, as shown in Fig. 1, is a simple 2-D scalable structure of physical qubits (denoted by the dots) that only requires local interactions between qubits (as illustrated by the lines in between the dots). Only a brief overview of the SC operation is given in this section; the interested reader is referred to [11] for a complete treatment. This work focuses on rotated SCs [41], which use the least amount of physical qubits per logical qubit. Each code has a distance d, meaning that a perfect decoder can correctly identify a maximum of (d-1)/2 physical errors. A rotated SC of distance d consists of a $d \times d$ grid of data qubits (white dots in Fig. 1) that encode a single logical qubit. The $d^2 - 1$ colored dots in Fig. 1 represent two types of ancilla qubits. These X- and Z-ancillas can be measured to find the errors on the adjacent data qubits without destroying the quantum state of the encoded logical qubit. This measurement outcome is called the error syndrome and needs to be continuously measured to detect errors in every so-called SC cycle. The task of the decoder is to find the errors on the data qubits from this error syndrome.



FIGURE 2. CNOT gate sequence. The number on each data qubit indicates the order in which the four adjacent data qubits are addressed by both the *X*-ancilla (blue, on the left) and the *Z*-ancilla (red, on the right). This process is performed on each ancilla qubit of the rotated SC shown in Fig. 1.



FIGURE 3. (Left) Quantum circuits for the commonly used SC cycle employing Hadamard and CNOT gates. (Right) Equivalent circuit with CZ and $R_y(\pm \frac{\pi}{2})$ on the right based on [42]. The $R_y(\pm \frac{\pi}{2})$ are denoted as \pm . In both figures, the top circuit shows the circuits for the *X*-ancillas and the bottom circuit for the *Z*-ancillas. The colored qubits are the ancillas. The other gray qubits are the data qubits surrounding this ancilla. Top to bottom this is the same order as the 1 to 4 shown in Fig. 2 in the CNOT dance. Although some $R_y(\pm \frac{\pi}{2})$ seem to cancel, those rotations are needed as the data qubits are interacting with other ancilla qubits in between.

A. SC CYCLE

During an SC cycle, all ancillas are first initialized into the ground state $|0\rangle$. Next, the *X*-ancillas (*Z*-ancillas) are brought onto the *x*-axis (*z*-axis) using a Hadamard gate (identity gate¹).

A sequence of CNOT gates is then performed, as shown in Fig. 2, to entangle each ancilla with its four adjacent data qubits. In case an ancilla is at the edge of the SC, only the two neighboring data qubits are used. Finally, the *X*-ancillas are brought back onto the *z*-axis and all ancillas are measured in the *z*-basis. This whole cycle is shown in Fig. 3 on the left.

The measurement on each ancilla will return either +1 or -1, reflecting the parity of the four (or two) adjacent data qubits. The set of all ancilla measurements in an SC is called the *error syndrome*. After all the ancillas are pushed into a state, there are still some degrees of freedom of the SC. These degrees of freedom define the logical state of the logical qubit. Any operation that does not change the error syndrome and logical state is called a *stabilizer*. The following section shows this in more detail. Measuring the ancillas puts the

¹The identity gate (idling) is shown to keep the two execution flows synchronized.



FIGURE 4. (a) Product of four stabilizers on the distance 5 SC. (b) Single logical *X*-operations. (c) Two logical *Z*-operations. (d) Two illustrations of the product of a logical *X*-operation and an *X*-stabilizer.

total SC into an eigenstate of all stabilizers where the error syndrome represents the eigenvalues.

After the first measurement cycle, the SC is initialized and all ancillas will either be +1 or -1. These do not represent errors, but the random initial quiescent state of the data qubits. Repeated measurement cycles will keep it in the same quiescent state. Any change in the error syndrome after measurement indicates a deviation from the quiescent state and, thus, an error.

B. LOGICAL OPERATIONS AND ERRORS

To understand how a stabilizer does not change the error syndrome [see Fig. 4(a)]. Performing four X-operations (shown in blue) on the data qubits around X-ancilla 3, does not change the parity of the (13, 14, 16) Z-ancilla measurements and, thus, does not change the error syndrome. Similar reasoning applies to the Z-operations around Z-ancilla 15. A product of two stabilizers is performed around X-ancillas 2 and 4. Due to the double X-operation on data qubit 16, an identity operation is performed on that qubit. It can be seen that this product of two stabilizers also does not change any of the adjacent Z-stabilizer measurements and also forms a continuous loop of single qubit operations.

In general, a product of X- or Z-stabilizers will always form a closed chain (or loop) of X- or Z-single-qubit operations. Thus, these loops will never change the quiescent state and the error syndrome. Since an error is modeled here as a random and nonintentional operation on a data qubit, errors that, by chance, form loops will not change the error syndrome and the logical state.

Next, Fig. 4(b) shows a chain of *X*-operations running between the top and bottom edge. Such an edge-to-edge chain performs a logical *X*-operation. One can check that this does not affect any of the *Z*-ancilla parity measurements.



FIGURE 5. Decomposition of the error *E* into a product of stabilizers *S*, a logical error *L*, and a pure error *P*. Of this product, only the pure error leads to a change in the error syndrome outcome.

Similarly, Fig. 4(c) shows two logical Z-operations, which should correspond to a logical *I*-operation. Not altering the logical state and error syndrome, those operations could also be written as a product of stabilizers. In general, any even number of logical operations can be written as a product of stabilizers and every odd number of logical operations can be written as a single logical operation and a product of stabilizers.

Finally, Fig. 4(d) shows a product of a logical X-operation as a chain between data qubits 2 and 22 with a stabilizer around X-ancilla 7. A product of these two results in a chain with the same shape as the one drawn between data qubit 0 and 20 and with the same effect on the logic state. This exemplifies that any odd number of chains, not necessarily straight, between the top and bottom will result in a logical Xoperation. Similarly, chains between the left and right sides of the SC result in logical Z-operations.

C. PURE ERRORS

The errors and operators discussed in the previous section do not change the error syndrome. Note that these operation chains do not end in the center of the SC. If they do end in the center, the error syndrome will change. For instance, the *Z*-error chain shown in Fig. 5 denoted by the error *E* starts at the edge at data qubit 10 and ends in the center at data qubit 17. As this is not a product of just stabilizers and logical operators, the error syndrome will be different, in this case at *X*-ancilla 8.

As Fig. 5 shows, any stabilizer *S* or logical operators *L* can be applied on top of the error *E* without changing the error syndrome. We call any state that gives the same error syndrome as the original error, and thus is only separated by stabilizers and logical operators, a pure error *P* [19], [43]. This pure error can, thus, be logically different from the original error *E*, but it will always give the same error syndrome. To phrase it differently, any error can be decomposed into a product of stabilizers, a logical operator, and a pure error.



FIGURE 6. Sketch of the logical versus physical error rate of an unencoded qubit (black line) and five different distances of the SC. The p_{th} is shown by colored circles and the decoder threshold by the gray circle. For a larger code distance, both the slope and the p_{th} increase.

D. DECODING

The purpose of the decoder is to identify an error configuration on the data qubits that produces the same error syndrome as was measured and is logically equivalent to the actual data error configuration. In other words, the decoder must output any data error configuration that only differs from the actual data error configuration by a product of stabilizers. This can then either be used to immediately correct the errors or can be tracked for later correction using Pauli frames [44]. The error syndrome must be the same to ensure that the SC returns to a logical state. The logical error must also be the same to prevent logical errors during the computation. As stabilizers do not influence either the error syndrome or the logical state, they can be neglected.

Since many data qubit configurations produce the same error syndrome, the error syndrome generation is a noninvertible function, thus making it impossible to unambiguously find the real data qubit configuration. This constitutes the main challenge for the decoder implementation and necessarily requires the decoder to make an arbitrary choice, which can be optimal if it is the one occurring with the highest probability.

E. DECODING PERFORMANCE

In order to quantify the decoding performance, the relation between the error rate of the logical qubit and the error rate of the physical qubits must be analyzed, as shown in Fig. 6. In this figure, both the logical error rate for an increasing SC distance (colored lines) and the error rate for an unencoded physical qubit (black line) are shown. The physical error rate for which the decoder achieves approximately² the same performance independent from the SC distance is defined as the *decoder threshold*.

 $^{^2 \}mathrm{In}$ practice, it is possible that all lines do not cross exactly in a single point



FIGURE 7. Model fit of (1) on the simulation of the MWPM decoder for the distances 3, 5, 7, and 9. The error bars represent the 99.9% confidence interval.

For any physical error rate in the following, the decoder threshold, it pays off to invest in a larger distance. The decoder threshold is often used as a single parameter to quantify the performance of a decoding algorithm. However, as shown in this sketch, operating at this physical error rate will be outperformed by a single unencoded qubit.

The physical error rate at which the logical qubit will outperform the physical qubit is called the *pseudo-threshold* (p_{th}) . The p_{th} is different for every distance and is used to compare decoders at the same SC distance. A higher p_{th} is preferred as it allows obtaining an advantage of using QEC with worse qubits. Even for a fixed physical error rate well below the p_{th} , a higher p_{th} will still give a lower logical error rate, assuming a constant slope of the lines in Fig. 6. Thus, the decoder slope is also an important parameter. Both the slope and the p_{th} increase when going to larger distances, but due to the exponential relation, the slope typically dominates the decoding performance at lower physical error rates.

As this article mainly compares decoders operating at a certain SC distance, we will focus on comparing the $p_{\rm th}$ and the decoder slope. The proposed decoders will be benchmarked against the minimum weight perfect matching (MWPM) algorithm [15], also known as Blossom or Edmonds' algorithm. Although better decoders exist [16], the MWPM algorithm is adopted as benchmark, as commonly done in prior works [16], [17], [19], [29]–[33], [35]–[37], [39], [40].

As an example, Fig. 7 shows the simulated performance of the MWPM decoder for the smallest four SC distances. The data are fitted using (1), adapted from [11, eq. 11] where ϵ_p and ϵ_l are, respectively, the physical and logical error rate, and p_{th} , s and c are fitting parameters representing the p_{th} , the slope for $\epsilon_p \ll p_{\text{th}}$ and the flattening of the curve for increasing physical error rates, respectively. The good fitting

Operation	Transmons	Single-electron spin qubit (Silicon)
Single qubit gate	20 ns [47]	1 μs [47]
Two qubit gate	40 ns [47]	0.1 μs [47]
Measurement	200 ns [48]	1 μs [49]
SC cycle duration	440 ns	5.4 μs

of the model up to the decoder threshold in Fig. 7 indicates that interpolation in the logarithmic domain is necessary to calculate the p_{th}

$$\epsilon_l = p_{\rm th} \left(\frac{\epsilon_p}{p_{\rm th}}\right)^{s \cdot (1 - c \cdot \epsilon_p)}.$$
(1)

F. HARDWARE REQUIREMENTS AND COSTS

In addition to the decoding performance, decoder implementations must also be compared based on their hardware requirements (delay) and hardware costs (area and power).

When using quantum error detection with Pauli frames, the main requirement is the minimum decoder throughput to avoid a data backlog [11], [44]. In principle, the decoder can run in parallel with the main algorithm execution and, to ensure the throughput, the decoding delay should be just lower than the measurement cycle, but not necessarily much smaller. However, tracking of errors is not enough when using non-Clifford gates, and the physical correction of errors is needed [44]. Since such a correction must be performed before the next cycle after the error detection, the decoding can only take a fraction of the cycle time.

The maximum allowed delay in case of transmons and silicon-based single-electron spin qubits are estimated in Table I assuming the SC cycle shown in Fig. 3. For the targeted qubit technologies, the circuit on the left of Fig. 3 can be replaced with the circuit on the right [42], [45], as the *cz* and $R_y(\pm \pi/2)$ gates can be performed faster and more accurately in those physical platforms. The minimum reported duration for each operation is used to obtain the most stringent constraint on the delay. Also, the initialization of the ancillas is neglected [42], [45] to get a lower bound of the estimated cycle time. For quantum error detection, the delay must be below 440 ns. However, when including a correction step for using non-Clifford gates, the delay needs to be as small as possible. This work will, thus, strive to minimize the delay and report the corresponding power and area.

Since the decoder is used once per QEC cycle and the cycle duration is fixed by the qubit technology, the hardware cost in terms of power is accounted for by computing the energy per decoding cycle. The dissipated energy per cycle should be as small as possible to allow for the largest number of logical qubits before running into the cooling power limitations. When fully integrating the decoder with the qubits on the same chip (or in the same package), the area must also be as small as possible to ease the integration requirements [46].

III. PROPOSED DECODER

In this work, we focus on decoding the smallest four rotated SCs, see Fig. 1. The goal is to obtain a decoder that has the high-decoding performance, runs fast enough to avoid a data backlog, and can be efficiently implemented in hardware.

NNs are a promising solution for several reasons. First, they have shown higher p_{th} compared to other decoders, such as the MWPM algorithm [19]. They can also adapt to many error models during training, perhaps even tailored to a specific qubit technology or even an individual quantum computing sample. After training, their execution (inference) time is constant and independent of the input. On the one hand, the inference time of NN decoders in hardware implementations has been estimated before [19], [38], but did not satisfy the throughput requirement. On the other hand, these analyzes do suggest that an optimized design on an ASIC could meet such a requirement. Finally, their regular structure makes them well suited for hardware optimization by parallelization and pipelining.

However, they also have their drawbacks. A large enough training dataset is needed to avoid overfitting. Even though any dataset can be generated using an error model, the size requirement of this dataset can still be a problem [33]. Next, NNs are quite complex and self-trained algorithms that are difficult to thoroughly understand, thus risking unexpectedly failing in untested situations. On top of that, there are a lot of additional parameters that need to be optimized during training [33], [38], making the search space for finding the optimum solution even greater. The main challenge, however, is that NNs are not well suited for direct application to the decoding problem. Fig. 8(a) shows the NN in such a direct, so-called low-level decoder (LLD) application. In this configuration, the NN takes in the error syndrome and guesses the error on every data qubit. The goal is that this data qubit configuration returns the correct logical error and also results in the same error syndrome as was measured. The problem is that the NN has no notion of what such a valid solution entails. This will limit the chance that a valid data error configuration is obtained. Consequently, a rerun of the algorithm is needed until a satisfying solution is found. To circumvent this limitation, we adopt the solution proposed in [19] and use a high-level decoder (HLD). In an HLD, the task of obtaining any correct error syndrome is performed by a pure error decoder (PED). This reduces the task of the NN to finding the type of logical error, allowing the NN to be a classifier, a task that is well suited to NNs, see Fig. 8(b).

This work focuses on fully connected feed-forward NNs. Although they show limits in scalability, those can be solved by opting for more complex topologies, such as convolutional neural networks (CNNs) [39]. However, for near-term small-distance SCs, we deem that the advantages of fully connected NNs still outweigh their disadvantages. The following sections will explain the basic functionality of the NN and the PED chosen in this work.



FIGURE 8. (a) NN in an LLD. This takes the error syndrome as inputs and gives the data qubit errors as output. (b) NN together with a PED in an HLD. Here, the PED gives the data qubit errors. The NN outputs the expected logical error that the PED makes compared to the actual data qubit errors.

A. PURE ERROR DECODER

The only task of the PED is finding a configuration for the data qubit errors that produces the error syndrome measured by the ancillas [19], [43]. As shown in Fig. 5, this pure error will only differ from the actual data qubit errors by a product of stabilizers and logical operators. As the product of stabilizers does not influence the error syndrome or the logical error, it can be neglected. Thus, the only significant difference between the error estimation given by the PED and the effective error is a logical error. Leaving the task of the NN to guess the logical error.

The benefit of this approach is that the guess made by the PED does not need to be the most probable. As a result, the PED can be optimized for other properties, and in this work, we focus on the following three main points.

- 1) *Software simulation speed:* As the PED must run every time the NN is run or trained, the speed of the PED must be maximized.
- 2) *Hardware simplicity:* The area and power of the PED must be minimized.
- 3) Exploiting symmetries: The SC is characterized by several symmetries that can be exploited in the training of the NN. However, as the NN also learns on the basis of the PED output, the PED should also show the same symmetries as the SC for fully optimizing the NN training.

The algorithm for the PED illustrated in Fig. 9(d) complies with the three abovementioned optimization targets.



FIGURE 9. Figure illustrating the steps to obtain the PED used in this work. (a) Boundaries where the error chains end for the ancillas of that given color. (b) Data qubits are grouped to the corresponding ancillas at the boundary. (c) Chains of equal length and equal distance apart for all the Z-ancillas, routing them to the corresponding edge. (d) Chains of equal length and distance apart for all ancillas.

To understand how this PED is obtained, we recall Fig. 5, which shows that pure errors form chains of contiguous ancilla errors from the inner part to the boundaries. This means that our PED must find chains that connect all the ancilla errors to the boundaries corresponding to the appropriate logical error. These boundaries are shown in Fig. 9(a). For this discussion, we will first focus on the top half (*Z*-ancillas). The full decoder can then be obtained by rotating this Algorithm 3 times by 90° .

For any distance, one semiplane has $(d^2 - 1)/4$ (6 for the example in Fig. 9) ancillas to be routed to the edge. As highlighted in Fig. 9(b), at the edge, there are (d + 1)/2 (3) ancillas and d (5) data qubits. As we need only one data qubit per ancilla, we only need (d + 1)/2 (3) data qubits as well. For symmetry, we choose to route each ancilla to the data qubits on the boundary that are equally spaced, as shown in Fig. 9(c). To be invariant to translations, all (d + 1)/2 (3) error chains are kept equidistant when moving toward the boundary. Since we have $(d^2 - 1)/4$ (6) ancillas, each chain will be (d-1)/2 (2) ancillas long. By rotating this scheme 3 times by 90° , each ancilla is routed to the boundary as in Fig. 9(d). Combining this pattern with the numbering as shown, an algorithm to be executed in software or hardware can be derived. This algorithm has minimized the length of the longest chains, making the hardware as fast as possible. As it turns out this means that all chains have equal lengths.

The resulting algorithm is just a series of XOR gates and can be described by the iterative formula in (3) with initial



FIGURE 10. Illustration of a computing node in an NN. The picture shows node *j* in layer *I*.

step 2. Here, *i* indicates the step in the algorithm, starting from the center at i = 0 to the edge at i = (d - 1)/2 - 1. $E(q_i)$ is the error on data qubit with number q_i , and $E(a_i)$ is the error on the ancilla a_i

$$E(q_0) = E(a_0) \tag{2}$$

$$E(q_i) = E(a_i) \oplus E(q_{i-1}). \tag{3}$$

The indices q_i and a_i correspond to the data and ancilla qubits in Fig. 9(d) and can be calculated as

$$q_{i} = \left[\frac{d-1}{2} + r \cdot (i+1) + 1\right] \cdot [t \cdot d + (1-t)] - 1$$
$$+ 2 \cdot c \cdot [d \cdot (1-t) - t]$$
(4)

$$a_i = \left[\frac{d^2 - 1}{4}\right] \cdot \left[1 + 2 \cdot t\right] + \left[\frac{r - 1}{2} + r \cdot i\right] \cdot \left[\frac{d + 1}{2}\right] + c$$
(5)

where *t* is either 0 or 1 for an *X*- or *Z*-chain, respectively, *r* is the rotation of the algorithm, being -1 or +1 for left or right for *X*-chains and up or down for *Z*-chains and *c* is the specific chain. For example, if there is an error on *X*-ancilla 8, we take a look at t = 0, r = +1, and c = 2. If we plug these values into (4) and (5), we obtain

$$q_i = 23 + i \tag{6}$$

$$a_i = 8 + 3i. \tag{7}$$

The initial step i = 0 says that there is an error at $q_0 = 23$ because there is an error at ancilla $a_0 = 8$. Next, there is also an error at $q_1 = 24$, as there is no error at ancilla $a_1 = 11$. This results in the same pure error, as shown in Fig. 5 *P*. The output of the PED is the sum of all data errors after running this iterative process over all chains on both sides of both ancilla types.

B. NEURAL NETWORK

As mentioned earlier, an NN is used to determine the logical error made by the error estimation of the PED. The input to the NN is the error syndrome that consists of all the ancilla measurements and its output is the estimated logical error, i.e., one of the possible logical errors. This work uses a fully connected feed-forward NN, which is a regular multilayered structure consisting of computing nodes. Every node in a layer is connected to all the nodes in the previous and following layer, as illustrated in Figs. 10 and 11. There are two main variations, which are sparsely connected or CNNs,



FIGURE 11. Illustration of a fully connected feed-forward NN. This work uses two hidden layers and two outputs as depicted here.

which only connect a selection of nodes between two layers, and recurrent neural networks (RNNs), which connect the outputs of a layer back to its inputs, thus obtaining memory. Effectively these two variations, respectively, decrease or increase the number of inputs seen by every node, compared to the basic fully connected feed-forward neural network.

Every node sums its weighted inputs, adds a bias to the resulting sum and applies a nonlinear function to the result to generate the output. This is expressed analytically as

$$a_{j}^{(l)} = \sum_{i} W_{j,i}^{(l-1)} \cdot y_{i}^{(l-1)} + b_{j}^{(l-1)}$$
(8)
$$y_{j}^{(l)} = f\left(a_{j}^{(l)}\right)$$
(9)

where y_i^l is the output of node *i* on layer l, $W_{j,i}^{(l-1)}$ is the weight to be applied to the output of node *i* of the previous layer (l-1) when contributing to the node *j* of the layer l, $b_j^{(l-1)}$ is the bias, $a_j^{(l)}$ is the accumulated output, and $f(\cdot)$ is a nonlinear transfer function (or activation function). The nonlinearity of the transfer function is crucial to avoid the whole neural network collapsing into a single-linear layer.

A neural network always contains an output layer. As the name suggests, all the nodes in this layer produce the outputs of the neural network. The vector of inputs is sometimes called the input layer. However, as can be seen in Fig. 11, this layer does not contain any nodes. If more layers are used between the input layer and the output layer, they cannot be directly observed, and are, hence, called hidden layers.

Even though a single hidden layer is enough to map any function [50], having multiple layers reduces the number of nodes needed in each layer. Previous work showed us that two hidden layers perform better than a single layer in terms of decoding accuracy [33]. Since adding another layer did not yield any significant improvement, we will focus on two hidden layers in this work.

The number of inputs of every node depends on the number of nodes of the previous layer, except for the first hidden



FIGURE 12. Overview of the simulation setup used in this work.

layer. In this layer, the number of inputs is equal to the number of ancilla qubits, i.e., $d^2 - 1$.

The number of nodes in the output layer depends on the classification scheme. One can use the classification scheme shown in Fig. 8(b), using four nodes to represent the different errors. This can either be no error, called a logical identity I, or one of the logical X, Y, or Z errors. These logical errors represent the logical difference between the PED output and the actual data errors. However, this can again lead to different output nodes competing and deciding independently. For this reason, we choose only two output nodes, one for signaling a logical X error and the second one for a logical Z-error. This has the added benefit of reducing the number of output nodes, and thus, the weights and size of the neural network, whilst still keeping the four output classes as no or both X and Z give I and Y, respectively.

This work opts for the simplest implementation of a fully connected feed-forward neural network as a first step toward the hardware implementation of NN QEC decoders. More complex architectures include CNNs, which only connect a selection of nodes between two layers, and RNNs, which connect the outputs of a layer back to its inputs to implement memory capabilities. While CNNs and RNNs would be more suited for this application when considering scalability and a more realistic error model, they would only require small modifications in the hardware of the individual nodes, as CNNs and RNNs, respectively, decrease or increase the number of inputs seen by every node compared to the basic fully connected feed-forward neural network. However, including these options is beyond the scope of this initial study, as it would drastically increase the search space in the performance/hardware-cost tradeoffs. At the same time, thanks to the similarity in hardware, the proposed results can form the basis for future extensions to these more complex architectures.

IV. METHODS FOR SIMULATION AND TRAINING

Before delving into the design and optimization of the different parameters of the neural network, the details about the simulation infrastructure are first described. The flow of the simulation setup is shown in Fig. 12. First, to generate a realistic error pattern for both NN training and evaluation, the data qubit errors are sampled using the depolarizing error model as discussed in the following section. Those are then fed to the SC simulator to obtain the corresponding error syndrome. The error syndrome is passed to the PED, which returns the pure error. The pure error is compared with the actual data qubit errors and the logical difference between the two is saved as the target output for the neural network. The error syndrome is also given to the neural network, which produces a logical error estimate. By comparing such an estimate to the target logical difference, the correctness of the estimation can be derived, which can then be used to assess during the training of the NN or to evaluate its performance.

A. SAMPLING

Since this work only focuses on feed-forward neural networks that are unable to deal with measurement errors due to their lack of memory, we adopted the depolarizing error model without measurement errors. The error model is implemented by applying a random physical error on each data qubit in every cycle chosen among a X-, Y-, or Z-error with equal probabilities p/3. This sampling is done on the fly just before the neural network is run, without pregenerating a dedicated training and testing dataset.

Prior work [33] investigated the optimal way of generating such a dataset without overfitting. To generate their dataset, they sampled a large number of data qubit error configurations and recorded the resulting error syndrome and logical error of the PED. For every error syndrome, a distribution of the four resulting logical errors (I_L, X_L, Y_L, Z_L) was saved. The neural network was then trained on this finite error syndrome dataset with the target output being the corresponding logical error distribution until it reached a certain accuracy on such a training dataset. However, due to the huge space of possible error syndrome for SC distances larger than 5, only a small set of all possible error syndromes is represented in the dataset. Furthermore, the logical error distribution for each error syndrome is also undersampled. Especially for the rarer error syndromes, only a single logical error might be sampled. This will inevitably result in an incorrect training data set, with no guarantee of generalization and a large risk of overfitting to incorrect data.

This work proposes a different method. Our simulation setup does not generate a predetermined finite dataset. Instead, this work keeps sampling new data on the fly only providing a single syndrome with the corresponding logical error at each step, which will always be correct. The training procedure itself will then average out all of these points and the neural network will learn the error distribution. Overfitting is avoided as all the data will be new, and the neural network will only benefit from training longer. In addition, since the dataset is uncorrelated, the resulting logical error rate at the end of the training will always represent the performance of the neural network. Thus, instead of finishing training by matching the training dataset, our work stops training when the current decoding performance is saturating or deemed sufficient.

The work in [33] found that training at a certain physical error rate will optimize the performance of the neural network at that physical error rate. Because we want to optimize the performance at the p_{th} , we sample at the physical error

rate corresponding to the p_{th} of the MWPM algorithm for that distance.

B. TRAINING AND TESTING

The training is done using the ADAM optimizer [51] with a batch size of 4992. As we have no finite dataset to optimize for, we trained the neural network for 300 000 batches. This results in a total dataset of $\approx 1.5 \times 10^9$ for each training. Fig. 14 plots the logical error rate during training per iteration of 2000 batches on the largest used neural network, showing the performance saturation after 150 iterations. The testing after training is done similarly to training. We again run 2000 batches to obtain the desired statistical accuracy. This is done for several logarithmically spaced values in a range between 0.03 and 0.3. As can be seen in Fig. 7, this range includes the $p_{\rm th}$, the decoder threshold, and clearly shows the slope difference. To obtain the slope, a fit is performed using the model in (1), and to obtain the p_{th} , we interpolate the two values above and below the ler = per line in the logarithmic domain. The reported variance used in the confidence interval is the sum of the variances of these two points. For the simulations that include quantization, this process is repeated for every combination of quantization levels and regularization levels.

All training and testing were done on custom-written code in C++ and Cuda, which was run on NVIDIA Tesla K40 GPUs over a span of a couple of months. All code is available at [1].

C. COST FUNCTION FOR QUANTIZATION

Due to the targeted hardware implementation, some additional regularization terms are added to the typical meansquared-error cost function used during the NN training. The process is illustrated in Fig. 13. Usually, the weights are randomly initialized in a certain range [see Fig. 13(a)]. During training, the weights expand outward [see Fig. 13(b)] [52], which usually is not a problem for weights using the floatingpoint representation. However, if we want to quantize those weights to a set of discrete levels [see Fig. 13(c)], issues arise when limiting the number of bits used in quantization. If, for example, all the weights are quantized between -1and +1, all weights outside this region are clipped to -1 and +1 [see Fig. 13(d)]. To push the weights toward 0 during training, we can add the sum of all squared weights $|w|^2$ to the cost function [see Fig. 13(e)] [52]. This will push less important weights to zero and decrease the average size of all weights. Another problem is that before quantization all the weights are uniformly distributed between the ± 1 range. To minimize the quantization error, we can also try to push the weights toward certain quantization levels during training [see Fig. 13(f)]. This can be done by adding the sum of the squares of the difference between every weight and the nearest quantization level $|w - w_q|^2$. Combining the two [see Fig. 13(g)] results in the following cost function:

$$\sum (y-t)^{2} + r \cdot \left(\sum |w|^{2} + \sum |w-w_{q}|^{2} \right)$$
(10)



FIGURE 13. Method adopted for training to optimize the representation of the weight using fixed point rather than floating point. This involves using weight regularization to push the weights toward zero and toward the nearest quantization level. The weights and outputs are then quantized after training.



FIGURE 14. Logical error rate during training for a distance 9 code with the maximum neural network used in this work (hidden layer sizes: 256 and 64), comparing to the MWPM decoder the final error rate for different transferfunctions (TanH and SQNL) and for the use of the rotational symmetry (0 or 1). The error bars show a confidence interval of 99.9%.

where y is the output, t is the target output, w is the value of the weight, and w_q is the quantized weight. The additional scaling term r decreases the influence of the weight regularization compared to the output error.

To further decrease the quantization error, a different number of quantization bits can be used to sample the weights than is used for the regularization. An example where we sample with an additional bit is shown in Fig. 13(h). The

3101719

final reported performance is the optimum over all possible regularization bits. Finally, because we use two's complement signed fixed-point numbers, the discussion mentioned above should be in the range $[-1, 1 - 1/2^{b-1}]$, where *b* is the number of bits.

V. DECODING PERFORMANCE RESULTS

The main objective of this work is to minimize the complexity of the neural network used in the HLD, while still obtaining a competitive decoding performance. Reducing the complexity implies a reduction in the number of free parameters. For instance, this can be done by reducing the size of the neural network, or by constraining the architecture and weights. Constraining should be done with care [53], but can reduce the size while still improving the performance [54]. An example is CNNs, where the connectivity is limited and weights are reused.

In this work, we focus on four tuning knobs that influence both the decoding performance and the hardware cost: the rotational symmetry, the transfer functions, the layer sizes, and the number of bits used for quantization. When looking at these parameters, we will compare the obtained decoder slope and the p_{th} . First, the influence of rotational symmetry is determined. Next, different transfer functions are compared. These results are then used in a layer-size sweep and in a bit-width sweep for the quantization.

A. ROTATIONAL SYMMETRY

This work focuses on fully connected neural networks, thus leaving exploration of other SC symmetries using CNNs for future work. However, the rotational symmetry of the SC and our PED can be investigated. The weights of the neural network can be copied and rotated four times. If we then use an initial neural network with a quarter of the size, the number of independent weights is divided by four, while still keeping the same total amount of weights and connectivity. This will both reduce the hardware cost by reducing on-chip memory, and ease the optimization during training.

In order to compare the performance difference between rotating the neural network or not, we ran simulations for all different distances, transfer functions (TanH, ReLU, SQNL, see following section), and layer sizes (from 4 to 256). For all these simulations, both a version with and without rotational symmetry is trained. The p_{th} and slope have been extracted and the averages over the different transfer functions and layer sizes are presented in the top two rows of Table II. The average is calculated by taking the geometric mean of the ratio between including and excluding symmetry for every configuration. If the ratio is larger than 1, an improvement is found by rotating. Since the performance for some configurations was too low, resulting in an undefined slope or p_{th} , those configurations were excluded from the computed average.

We find that including rotational symmetry has a positive effect on both the p_{th} and the slope for all distances. This is mainly attributed to the training having an easier time finding

Performance Distance 5 9 3 All comparison 7 Rotated ^a 1.0133 1.0320 1 0563 1.0850 1 0411 $p_{\rm th}$ Slope 1.0042 1.0107 1.0290 1.0546 1.0207 Unrotated SQNL^b 1.0291 1.1637 1.2830 1.3154 1.1703 p_{th} TanH Slope 1.0021 1.0552 1.1378 1.2137 1.0822 ReLU^b 0.9771 0.9957 1.0177 1.0387 1.0008 p_{th} Slope 0.9936 0.9978 1.0127 1.0605 1.0101 TanH

TABLE II Average ρ_{th} and Slope Improvements When Comparing the Use of Rotational Symmetry and the Different Transfer Functions

^aAveraged over all transfer functions.

^bAveraged over both rotated and unrotated configurations.

The ratio of the geometric means of the performance over different layer sizes (excluding cases when the parameter is undefined) is reported.

an optimum and is in line with the results in [55] for a Toric code. The improvement increases for larger distances, which is likely due to the larger neural networks performing better and benefiting more from the weight regularization. This is in line with the results found on layer sizes later in this section.

B. TRANSFER FUNCTIONS

Different transferfunctions will have different hardware costs and decoding performance, with this work will focus on, in descending hardware cost, hyperbolic tangent (TanH), squared nonlinearity (SQNL), and rectified linear unit (ReLU), defined as

$$TanH(x) = \frac{e^{x} - e^{-x}}{e^{x} + e^{-x}}$$
(11)

$$\operatorname{ReLU}(x) = \begin{cases} 0, & \text{for } x < 0\\ x, & \text{for } x \ge 0 \end{cases}$$
(12)

$$SQNL(x) = \begin{cases} -1, & \text{for } x < -1\\ 2x + x^2, & \text{for } -1 \le x < 0\\ 2x - x^2, & \text{for } 0 \le x \le 1\\ 1, & \text{for } x > 1 \end{cases}$$
(13)

For simplicity, every node will use the same transfer function.

The lower half of Table II compares the performance for different transfer functions similarly to the comparison for the rotational symmetry. The more computationally expensive TanH is taken as a reference. The average performance difference of the ReLU over the TanH is not significant. Since the ReLU is much cheaper in hardware implementation (not requiring any exponential), it is attractive even for roughly equal performance.

The SQNL does show a large improvement, up to 31%. This, in combination with the simplicity of the required hardware, make SQNL the preferred choice.

The transfer functions show the same performance increase for larger distances, similar to the rotational symmetry results. This strengthens our belief that this trend is due to the need for larger neural networks for larger distances.

Based on these results, we limit the search space for the follow-up analyzes by adopting rotated neural networks with



FIGURE 15. p_{th} performance for the smallest four SC distances as a function of the number of nodes used in the first hidden layer. For each distance, we report the results for a second hidden layer with 64 nodes (line with higher p_{th}) and with 4 nodes (line with lower p_{th}). The dashed lines show the MWPM p_{th} . The error bars represent a confidence interval of 99.9%.

the SQNL function. The results with the other options are available in the supplementary materials [1]. We also limit ourselves to the p_{th} performance.

C. LAYER SIZES

As discussed in Section III, this work uses two hidden layers and an output layer with two nodes, as there are two outputs, thus using the neural network shown in Fig. 11 for l = 3. We assume that reducing the number of nodes in any of the hidden layers will reduce the hardware cost but, as shown in the following, will degrade the decoding performance due to the reduced computational power of the neural network.

A summary of the layer size sweep is shown in Fig. 15 by plotting p_{th} as a function of the number of nodes in the first hidden layer. The achieved slope is not shown in this section and the next, but the correlation between p_{th} and the slope is shown at the end of this section in Fig. 17. For each distance, we report the results for a second hidden layer with 64 nodes (line with higher p_{th}) and with 4 nodes (line with lower p_{th}). All other investigated second layer sizes lie in between these two cases.

All lines in Fig. 15 show a similar trend, saturating to maximum performance with increasing layer sizes. This maximum p_{th} increases for larger distances, and, as expected before, a larger distance requires a larger neural network for the same p_{th} . The largest tested neural network decoder (256 and 64 nodes in the first- and second-hidden layer, respectively) is enough to outperform MWPM, whose performance is indicated by the dashed lines. For the smallest distance of 3, almost any neural network reaches maximum performance. For distance 9, the maximum is not yet visible, indicating that future research should include larger neural networks.



FIGURE 16. $p_{\rm th}$ performance degradation of the quantized neural networks with respect to their floating-point counterparts. All configurations are shown in lighter colours. The thicker lines represent the average performance degradation.

The effect of the second layer size is shown as the difference between the higher and lower line for each distance. More detailed data are included in [1]. The first hidden layer size has a stronger impact on performance than the second layer, although this is stronger for larger distances.

D. QUANTIZATION

Using a fixed-point representation for the data in the NN instead of a floating-point representation can significantly save on hardware costs, at the price of the decoding performance. Since there is no strategy to determine the optimal number of bits in the fixed-point representation [56], we explore how performance varies for different quantization levels. Ideally, the optimum number of bits for each node depends on the number of input nodes and the domain of the transfer function, and in principle can be different for node output, weights, and biases. However, for simplicity of the study, we adopt the same number of bits to represent the outputs, the weights, and the biases in all layers, except for the 1-bit global inputs and outputs of the NN.

To see the effect of quantization, we compare the performance of the neural network before and after quantization. This is done by dividing the quantized (fixed-point) p_{th} by the floating-point p_{th} . The floating-point performance before quantization does include the regularization term that attracts the weights to certain quantization levels. These results are shown in Fig. 16, plotting the performance degradation due to the quantization for all different configurations of layer sizes and quantization levels used in weight regularization. For layer sizes, this is between 4 and 256 in the first hidden layer and between 4 and 64 in the second hidden layer. The weight regularization levels are varied between 4 and 256 levels (2 and 8 bits). The *x*-axis shows the number of bits used for representing the data during evaluation.

TABLE III Decoding Performance	of This Work Compared to the MWPM
Decoder and the Work of [33]	

Performance	Used	Distance			
parameter	Decoder	3	5	7	9
p_{th}	MWPM	0.08251	0.10372	0.11368	0.11932
	[33]	0.09815	0.12191	0.12721	0.12447
	Float	0.09769	0.12657	0.12917	0.12490
	Fixed	0.09781	0.12637	0.12934	0.12430
Slope	MWPM	1.856	2.723	3.601	4.496
	Float	1.886	2.869	3.812	4.663
	Fixed	1.894	2.866	3.820	4.667
Min. Bits	Fixed	3	4	5	7

Results are limited to the use of the rotational symmetry and the SQNL transfer function, for both the best floating-point performance and the best performing fixed point neural network. We also report the minimum number of bits needed to obtain a $p_{\rm th}$ higher than that of the MWPM decoder.

A couple of trends can be observed. First, 9 bits are enough for most configurations to reach floating-point performance (a value of 1). Next, larger distances need more bits for the same performance degradation. These trends are captured in the thick lines as cumulative average performance degradation. These lines represent the average degradation in the performance when going down to fewer bits. First, all points at 9 bits are averaged for a certain distance. Then, every following step down in bits, we subtract the average degradation of all lines in that distance.

From the data [1], we also see that we need more bits to reach the MWPM performance for larger distances. For increasing distances we need 3, 4, 5, and 7 bits. One explanation could be due to the relation between layer sizes and quantization discussed at the start of this section. This would indicate that more research into quantization dependent on layer size is needed. However, another option would be to look into CNNs, which effectively decreases the number of inputs per node.

E. COMPARING DECODING PERFORMANCES

Table III combines all the results discussed so far. It compares the maximum p_{th} and slope found in this work to the neural network decoders in [33] and the MWPM decoder. The data in the table are limited to the use of the rotational symmetry and the SQNL activation function, and to a layer size of 256 and 64 for the first and second-hidden layer, respectively. Even with neural networks smaller than in [33], we obtain slightly higher p_{th} , thus confirming the validity of our training method. It also means that our choice for a simpler transfer function does not affect the performance. Finally, we see that quantizing (at least 8 or more bits) does not result in any significant degradation in the performance.

Interestingly, a distance-3 decoder requires only a very small neural network and a few bits. Because, as shown before, the added symmetry improves the performance, an interesting possibility is using CNNs [39] or distributed decoders [40] based on distance-3 kernels, thus requiring



FIGURE 17. Correlation between the slope and p_{th} for the four distances. Every dot represents a different neural network configuration (rotated with SQNL function and all layer sizes). The dashed lines indicate the slope and p_{th} of the MWPM decoder.

simpler hardware, allowing the decoding of larger distances and even increasing the decoding performance.

Finally, an overview of all decoders is shown in Fig. 17. In this figure, every decoder is depicted with a dot in the color of the corresponding distance, and the dashed lines represent the performance of the MWPM decoder. The correlation between the slope and the $p_{\rm th}$ is quite apparent. The main takeaway from this plot is that a decoder at a larger distance can have a worse slope than a decoder at a smaller distance. Thus, if the decoder cannot achieve the sufficient performance, it might not be economical to go for a larger distance. For example, a decoder could be too large to be cointegrated with the qubits, consume too much power, or have too large a delay to keep up with the generated data. For this reason, the following section will analyze the estimated hardware cost.

VI. HARDWARE COSTS

A. HARDWARE ESTIMATE

As we want to find out the minimal achievable delay and we do not need any memory for recurrency inside of the neural network, we chose a fully parallel combinatorial implementation. Some flip-flops would be needed to store the input and outputs of the neural network during inference, and the weights and the biases would be stored in some external RAM. However, for illustrating the effects of solely the neural-network logic on the hardware, the cost in area and power of those memories and the memory access are not included in the following estimates.

Although one could expect that the absence of a clock and flip-flops would decrease the power consumption, this is not always the case [57]. A fully combinatorial circuit will propagate any glitch to the outputs, thereby increasing the



FIGURE 18. Schematic of the hardware in one node. First, (a) *m n*-bit inputs are combined with their corresponding weights into (b) multiplicants needed for the Baugh–Wooley (BW) multiplication scheme. (c) These multiplicates are then added in a carry-save adder (CSA) tree. Finally, (d) result is split and passed through the SQNL block to obtain (e) *n*-bit output.

amount of charging and discharging of the capacitance of the digital cells and the interconnect parasitics. Instead, these glitches would not be able to propagate in a more pipelined approach. In combination with clock gating, this could even decrease the power consumption [57], although pipelining would increase the area and delay.

The data flow in each node are depicted in Fig. 18. First, all the multiplicands are calculated using the Modified Baugh– Wooley 2's complement method [57], [58] (b). These are then summed in a Wallace carry-save adder tree [59] (c). Finally, to implement the SQNL function [see (13)], the resulting fractional part of the sum (d) is passed through a squaring unit [57] and added to or subtracted from a bit-shifted sum depending on the sign (e). This process is different for input nodes and output nodes. Each input is only one unsigned bit, and thus, the Baugh–Wooley method is replaced by a simple AND operation. The outputs are also just one bit, thus, only the sign of the sum is needed.

A hardware description of the resulting digital circuit is fed to the GENUS synthesizer using the CMOS standard static library form the adopted TSMC 40-nm CMOS process to obtain the circuit schematic. To efficiently obtain the hardware estimates for all the neural network configurations, only the individual nodes are synthesized. The resulting delay, area, and power from the individual nodes are then summed to obtain an estimate for the total neural network.

To estimate the delay of the NN, we extract from the synthesis the critical path of the node, which is equal to the critical path of the respective layer as every node in a layer is the same. The total delay is simply obtained by summing the critical path delays of the layers. This estimate does not include the additional delay due to the interconnect between layers.

The area of each node is estimated as the sum of the area of the digital cells in the node, scaled by a fixed fill factor of 1. Since all the nodes are equal, the total area is just the sum of the node areas, as we assume that the top-level layout can be quite efficient and does not consume any additional area.

The power is defined as the average energy needed per decoding cycle of 440 ns. To get a more accurate power estimate, we perform transient simulations in Cadence Virtuoso



FIGURE 19. *p*th of the quantized neural network decoders versus their estimated hardware cost. The figure shows the Pareto front of the different distances and compares them to the MWPM decoder. The error bars represent a confidence interval of 99.9%.

of the nodes. This will include the extra power consumption due to the propagation of glitches. The energy is then averaged over 100 cycles. The simulations have an input activity factor of 50%.

B. RESULTS

Figs. 19 and 20 show the hardware cost estimates (delay, area, and power) of an ASIC design versus the obtained p_{th} and slope, respectively. For both figures, the hardware cost is divided into three plots. The dots in these plots, colored according to the corresponding distance, represent all the quantized configurations, as shown in Fig. 16. The solid lines show the Pareto front for every distance, i.e., the set of the neural networks that perform better for that particular trade-off. The horizontal dashed lines represent the performance of the MWPM decoder. The hardware estimates of the individual nodes are included in [1].

As shown in Fig. 19, a larger circuit is needed to obtain the same p_{th} at a larger distance. This trend holds until the Pareto fronts start to flatten out. The performance saturation is very clearly visible for distances 3 and 5. For distances 7 and 9, the curves also seem to saturate but this is only visible in the Pareto front and not in the whole dataset. We would also expect p_{th} to go on further, especially for distance 9, as larger distances are expected to have a larger p_{th} . The saturation of the Pareto front is, therefore, attributed to the finite size of our neural networks, which limits their performance at larger distances. Similar trends appear in Fig. 20 on the slope data. Here, the data points also show less saturation for the larger distances. However, increasing the neural network sizes would be necessary to prove this in future work.

The data on the decoder slope in Fig. 20 also indicate a correlation between the hardware cost and the obtained performance. In contrast to the p_{th} , the Pareto fronts for the different distances seem to follow the same basic trend, and differentiate only by a different saturation value.

The delays of the neural networks are all smaller than 30 ns, i.e., an order of magnitude lower than the required 440 ns, thus indicating that less parallelism in the hardware implementation is feasible. Lower parallelism could help reduce both the area and power, which are quite large and scale exponentially with the performance. However, when using non-Clifford gates, a decoding time faster than 440 ns might be preferred. In case of errors, we either must correct the state before applying the non-Clifford gate [44], or we must update the logical Pauli frame and correct the state just after applying the non-Clifford gate and before applying the next gate to prevent the errors from spreading into a complex multiqubit error [60], [61]. Both methods require the decoder to keep up with the error syndrome generation, i.e., keep the throughput faster than 440 ns, but a faster decoder would be preferred so as not to limit the execution speed of the quantum algorithm.

Not being limited by the delay, there should be enough time to use RNN. RNNs are needed for decoding measurement errors and would increase the number of inputs per



FIGURE 20. Slopes of the quantized neural network decoders versus their estimated hardware cost. The figure shows the Pareto front of the different distances and compares them to the MWPM decoder.

node, as their own outputs in the previous cycle are concatenated with the outputs of the previous layer. It also opens the possibility to use the more complex recurrent cells of long-short term memory that typically have longer delays, consume more hardware, but also have shown a better decoding performance. Finally, more (and thus perhaps smaller) layers could be used as required by CNNs, which might need a deeper neural network, especially for larger distances.

Looking at the p_{th} and slope plot together, two considerations can be drawn as follows.

- If the highest *p*_{th} is preferred given certain hardware constraints, one can best choose a smaller SC distance. As a high *p*_{th} is mainly preferred when the physical error rates are higher, this means the qubits probably could not be integrated into larger distances anyway.
- 2) If the highest slope is preferred, it does not make sense to choose a distance where the slope already starts to saturate. Thus, the largest distance is the best choice. Similar to the previous point, if a high slope is preferred, probably the qubits are performing quite well at a low physical error rate. This means that the qubits can probably be used together to form larger logical qubits.

Fig. 21 supports these conclusions. In this figure, the optimum distance is shown for certain area constraints. Every line shows a different area constraint. This does not mean that this configuration will fully occupy this area but it will never exceed it. The markers on the lines change color depending



FIGURE 21. Optimal distance to obtain the lowest logical error rate for a certain physical error rate. The different lines indicate the maximum allowed area for the decoder.

on what distance obtains the lowest logical error rate for that physical error rate. As discussed, lower distances perform better when the physical error rate is around the p_{th} . In that region, as also Fig. 19 illustrates, the best performance for a restricted area is always for lower distances. However, when the physical error rates become low enough, the slope starts to dominate. This is seen by a shift toward larger distances for lower physical error rates.

TABLE IV Three Designs That Fit on the Artix-7 FPGA and Lie on the $p_{\rm th}$ Pareto Front

]	Distance	3	3	5
Layer 1	Size	8	16	64
Layer 2	Size	4	4	64
•	Bits	3	5	4
	$p_{ m th}$ Slope	0.0823 1.8641	$0.0976 \\ 1.8868$	0.1037 2.6641
FPGA	Delay Area Power	17.9 ns 351 LUT < 1 mW	71.9 ns 2942 LUT 6 mW	87.6 ns 44670 LUT 132 mW
ASIC	Delay Area Power	7.3 ns 0.0031 mm ² 10.7 μW	12.3 ns 0.0114 mm ² 43.2 μ W	14.3 ns 0.3937 mm ² 1.0 mW

The table shows the distances and configuration, along with the decoding performance and hardware cost for both the FPGA and the ASIC design. All designs have $d^2 - 1$ inputs for layer 1 and two nodes in the output layer.

C. COMPARING ASIC TO FGPA

We demonstrated that for near-term QEC, smaller distances are optimal. They obtain a higher p_{th} given the same hardware. Near-term QEC will also be a lot more experimental, requiring more frequent changes to the configuration of the neural network decoder and the corresponding weights and biases. As ASIC designs are optimized for integration given a certain configuration, they might not be suited for this task due to very limited reconfigurability. On the contrary, FPGAs can be reconfigured very easily and can even be synthesized to incorporate certain weights and biases, thus even optimizing away unnecessary hardware. A downside to FPGAs is that their hardware is more generic and, thus, less optimized in terms of delay and power. Furthermore, their adoption imposes a strict limit in the area, as only a fixed amount of hardware primitives (look-up tables, flip-flops) can be used.

To compare the FPGA to the ASIC designs, the same hardware was implemented on the Xilinx Artix-7 FPGA. However, instead of synthesizing all the nodes individually, the whole design was synthesized and implemented as a whole, thereby demonstrating if the design would actually fit on the selected FPGA. The most promising designs were chosen, either the ones lying on the Pareto front that would just barely obtain the MWPM performance or the ones for which the curves in Figs. 19 and 20 start to saturate.

After running synthesis and implementation using Xilinx Vivado, only the designs in Table IV did actually fit on the FPGA. These are the distance-3 decoders achieving the maximum performance and performance comparable to MPWM and the distance-5 decoder with the MPWM performance. All FPGA values are postimplementation room-temperature estimates. The power estimate is the reported dynamic power for a 440-ns clock period. The table shows that all decoders still fit the delay requirements by a large margin. This means that if this hardware would be optimized for area, e.g., by reducing the parallelism, probably larger neural networks could fit as well.

Despite those limitations, the results are very promising, as even the FPGA designs can meet the required delay. The delay and hardware costs would be even lower if the synthesis was run with hard programmed weights but this was omitted for a fair comparison to the ASIC designs.

D. MOVING TO CRYOGENIC TEMPERATURE

The values reported for ASICs are extracted from simulations at 300 K. From these, we can draw some conclusions about the performance at 4.2 K. The work of [62, Fig. 10.1], [63] found that at cryogenic temperatures the delay of digital cells decreases by up to 50% for mature CMOS technologies thanks to the increase in mobility. However, the speed-up will be much less significant in advanced commercial technologies as the increase in threshold voltage combined with the reduction of supply voltages mitigates those effects. As a result, the delay estimates can be assumed to approximately hold also at 4.2 K. Due to the increased subthreshold slope [64], the leakage power at cryogenic temperatures is greatly reduced. Thus, the power at cryogenic temperature is estimated to be lower than at 300 K. Finally, due to the increase in mismatch [64] and latch-up [65], [66], a larger area might be needed at 4.2 K to decrease the mismatch and to increase the number of well-taps to combat latch-up [67].

VII. DISCUSSION

We have proposed a new PED with more symmetries than previous works. These symmetries were also incorporated into the neural network resulting in the improved performance. For future works, however, even more symmetries could be exploited, for example, using a convolutional neural network. Another benefit of the novel PED is the equal delay of every chain.

A fully connected feed-forward neural network with two hidden layers was used for the HLD. Two hidden layers were chosen to minimize the delay and because this is enough to fit any possible function given enough nodes. The hardware estimates show, however that the delay is small enough to allow for more layers.

For the space exploration, first the TanH transfer function was approximated by the SQNL function, which is significantly cheaper in hardware cost and also outperforms the TanH.

Next, the layer sizes of the two hidden layers were swept and compared to the MWPM decoder and previous work [19], [33]. Even though our layers were significantly smaller, the obtained performance was on par with or better than previous research. The results also showed that, especially for distances 7 and 9, the performance could be further improved by increasing the layer sizes or the neural network depth. These results for a feed-forward neural network can be extended and compared to recurrent and CNNs in future work.

To reduce the hardware cost, all weights and outputs were quantized between 3 and 9 bits using a fractional fixed-point two's complement representation. For 9 bits, all distances performed on par with the floating-point results before quantization. Even fewer bits were needed to keep the performance above that of MWPM. More bits are needed for larger distances, possibly pointing to the need for some rescaling of the transferfunction domain depending on the layer size. Another solution might be to use a sparsely connected or CNN. This solution could also exploit the translational symmetry of the SC and our PED. The final option to move toward fewer bits is to either use a binary neural network or to train using a different methodology than was used in this article.

Finally, the ASIC hardware cost was estimated for every configuration in terms of delay, power, and area. The data show clear trends for both the p_{th} and the decoding slope. The hardware cost for a certain p_{th} is a lot higher for larger distances. On the one hand, this means that, if the main objective is to obtain the lowest-cost hardware for physical qubits with a high error rate, the lowest distance that can obtain that p_{th} should be chosen. On the other hand, if the physical qubits perform well below the p_{th} , the data on the decoding slope clearly show that there is a correlation between the steepest slope and the hardware costs. It also indicates that with an increasing distance, the maximum slope increases as well. If the qubits perform well enough and a large enough SC can be made, it is desirable to choose the larger distance.

The fully parallelized implementation chosen in this work makes the needed hardware larger than necessary, as the 440 ns that are required to avoid a data backlog are met by a large margin with by the ASIC and the FPGA designs. This would justify designs with more hardware reuse to tradeoff the extra delay for smaller and lower-power solutions, which is needed to get the area well below 10 mm² and the power below 1 W. Due to the reconfigurability of FPGAs, the weights could also be synthesized into the design, removing a lot of unnecessary hardware. Both these optimizations could significantly reduce the area and power. Combined with the exploration of CNNs, this is a relevant direction for future research. The hardware tradeoffs for RNNs should also be explored when using a more realistic error model including measurement errors, such as circuit noise. Employing recurrency and using the circuit noise error model would allow the comparison with other competitive hardware decoders, such as [30], [31]. Currently, our decoder shows competitive delays, while still performing better than the MWPM and Union-Find algorithms under depolarizing error models. Previous work has shown that this performance advantage also holds for RNNs under a circuit noise model [33], [38]. Furthermore, our current work focuses on fully connected neural networks, which are not yet scalable for future large-scale quantum computers. However, previous work [39] has shown scalable neural networks based on CNNs performing on par with MWPM up to distance 64 Toric codes. Extrapolating these results, hardware implementations of convolutional and/or RNNs represent a very promising alternative to be investigated in future research.

VIII. CONCLUSION

This work presents an extensive space exploration of a HLD for SCs consisting of a PED and a fully connected feedforward neural network with two hidden layers. The results show that the decoder can be optimized for hardware simplicity, while still obtaining state-of-the-art decoding performance. The resulting hardware implementation allows the decoder to obtain decoding times less than 30 ns in ASIC and less than 90 ns in FPGA implementations, which is significantly lower than the required 440 ns needed to keep up with the SC cycles of current solid-state qubit technologies. The required area and power dissipation for the ASICs are realistic for a practical implementation at SC distances up to 9 and for decoding performance well superior to the MWPM algorithm. This paves the way for QEC hardware that can be cointegrated with the qubits at cryogenic temperatures.

ACKNOWLEDGMENT

The authors would like to thank S. Varsamopoulos and P. Padalia for their useful discussions.

REFERENCES

- R. W. J. Overwater, "Data for: Neural network decoders for surface codes," Data for: Neural network decoders for surface codes, 4TU.ResearchData, Dataset, 2021, doi: 10.4121/16539786.
- [2] A. W. Harrow and A. Montanaro, "Quantum computational supremacy," *Nature*, vol. 549, no. 7671, pp. 203–209, Sep. 2017, doi: 10.1038/nature23458.
- [3] A. Montanaro, "Quantum algorithms: An overview," NPJ Quantum Inf., vol. 2, Jan. 2016, doi: 10.1038/npjqi.2015.23.
- [4] P. W. Shor, "Polynomial-time algorithms for prime factorizaiton and discrte logarithms on a quantum computer," *SIAM J. Comput.*, vol. 26, 1997, Art. no. 1484, doi: 10.1137/S0097539795293172.
- [5] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proc. 28th Annu. ACM Symp. Theory Comput.*, 1996, pp. 212–219, doi: 10.1145/237814.237866.
- [6] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*. Cambridge, U.K.: Cambridge Univ. Press, 2000.
- [7] R. P. Feynman, "Simulating physics with computers," *Int. J. Theor. Phys.*, vol. 21, pp. 467–488, 1982, doi: 10.1007/BF02650179.
- [8] J. Park, "The concept of transition in quantum mechanics," *Found. Phys.*, vol. 1, pp. 23–33, 1970, doi: 10.1007/BF00708652.
- [9] W. Wootters and W. Zurek, "A single quantum cannot be cloned," *Nature*, vol. 299, no. 5886, pp. 802–803, 1982, doi: 10.1038/299802a0.
- [10] A. Y. Kitaev, "Quantum computations: Algorithms and error correction," in Proc. 3rd Int. Conf. Quantum Commun., Comput. Meas., pp. 1191– 1249, 1997, doi: 10.1070/rm1997v052n06abeh002155.
- [11] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, "Surface codes: Towards practical large-scale quantum computation," *Phys. Rev. A*, vol. 86, 2012, Art. no. 032324, doi: 10.1103/PhysRevA.86.032324.
- [12] J. Koch et al., "Charge insensitive qubit design derived from the Cooper pair box," Feb. 2007, arXiv:cond-mat/0703002, doi: 10.48550/arXiv.cond-mat/0703002.
- [13] J. A. Schreier *et al.*, "Suppressing charge noise decoherence in superconducting charge qubits," *Phys. Rev. B*, vol. 77, no. 18, May 2008, Art. no. 180502, doi: 10.1103/PhysRevB.77.180502.
- [14] M. Veldhorst *et al.*, "An addressable quantum dot qubit with fault-tolerant control-fidelity," *Nature Nanotechnol.*, vol. 9, no. 12, pp. 981–985, Dec. 2014, doi: 10.1038/nnano.2014.216.
- [15] J. Edmonds, "Paths, trees and flowers," Can. J. Math., vol. 17, pp. 449–467, 1965, doi: 10.4153/CJM-1965-045-4.
- [16] S. Bravyi, M. Suchara, and A. Vargo, "Efficient algorithms for maximum likelihood decoding in the surface code," *Phys. Rev. A*, vol. 90, no. 3, Sep. 2014, Art. no. 032326, doi: 10.1103/PhysRevA.90.032326.

- [17] G. Duclos-Cianci and D. Poulin, "Fast decoders for topological quantum codes," *Phys. Rev. Lett.*, vol. 104, Feb. 2010, Art. no. 050504, doi: 10.1103/PhysRevLett.104.050504.
- [18] N. Delfosse and N. H. Nickerson, "Almost-linear time decoding algorithm for topological codes," *Quantum*, vol. 5, p. 595, Dec. 2021, doi: 10.22331/q-2021-12-02-595.
- [19] S. Varsamopoulos, B. Criger, and K. Bertels, "Decoding small surface codes with feedforward neural networks," *Quantum Sci. Technol.*, vol. 3, no. 1, 2017, Art. no. 0 15004, doi: 10.1088/2058-9565/aa955a.
- [20] B. Patra *et al.*, "Cryo-CMOS circuits and systems for quantum computing applications," *IEEE J. Solid-State Circuits*, vol. 53, no. 1, pp. 309–321, Jan. 2018, doi: 10.1109/JSSC.2017.2737549.
- [21] F. Sebastiano et al., "Cryo-CMOS electronic control for scalable quantum computing: Invited," in Proc. 54th Annu. Des. Autom. Conf., 2017, pp. 13:1–13:6, doi: 10.1145/3061639.3072948.
- [22] B. Patra et al., "19.1 a scalable cryo-CMOS 2-to-20ghz digitally intensive controller for 4×32 frequency multiplexed spin qubits/transmons in 22 nm finfet technology for quantum computers," in Proc. IEEE Int. Solid- State Circuits Conf., 2020, pp. 304–306, doi: 10.1109/ISCAS.2019.8702442.
- [23] C. Degenhardt *et al.*, "Systems engineering of cryogenic CMOS electronics for scalable quantum computers," in *Proc. IEEE Int. Symp. Circuits Syst.*, 2019, pp. 1–5, doi: 10.1109/ISCAS.2019.8702442.
- [24] T. Lehmann, "Cryogenic support circuits and systems for silicon quantum computers," in *Proc. IEEE Int. Symp. Circuits Syst.*, 2019, pp. 1–5, doi: 10.1109/ISCAS.2019.8702413.
- [25] S. Bonen *et al.*, "Cryogenic characterization of 22-nm FDSOI CMOS technology for quantum computing ICs," *IEEE Electron Device Lett.*, vol. 40, no. 1, pp. 127–130, Jan. 2019, doi: 10.1109/LED.2018.2880303.
- [26] A. Ruffino, Y. Peng, and E. Charbon, "Interfacing qubits via Cryo-CMOS front ends," in *Proc. IEEE Int. Conf. Integr. Circuits, Technol. Appl.*, 2018, pp. 42–44, doi: 10.1109/CICTA.2018.8705712.
- [27] J. M. Hornibrook *et al.*, "Cryogenic control architecture for large-scale quantum computing," *Phys. Rev. Appl.*, vol. 3, Feb. 2015, Art. no. 24010, doi: 10.1103/PhysRevApplied.3.024010.
- [28] H. Bohuslavskyi et al., "Cryogenic characterization of 28-nm FD-SOI ring oscillators with energy efficiency optimization," *IEEE Trans. Electron Devices*, vol. 65, no. 9, pp. 3682–3688, Sep. 2018, doi: 10.1109/TED.2018.2859636.
- [29] A. Holmes, M. R. Jokar, G. Pasandi, Y. Ding, M. Pedram, and F. T. Chong, "NISQ: Boosting quantum computing power by approximating quantum error correction," in *Proc. ACM/IEEE 47th Annu. Int. Symp. Comput. Architecture*, 2020, pp. 556–569, doi: 10.1109/ISCA45697.2020.00053.
- [30] P. Das et al., "A scalable decoder micro-architecture for fault-tolerant quantum computing," 2020, arXiv:2001.06598, doi: 10.48550/arXiv.2001.06598.
- [31] P. Das, A. Locharla, and C. Jones, "LILLIPUT: A lightweight low-latency lookup-table based decoder for near-term quantum error correction," 2021, arXiv:2108.06569, doi: 10.48550/arXiv.2108.06569.
- [32] Y. Ueno, M. Kondo, M. Tanaka, Y. Suzuki, and Y. Tabuchi, "Qecool: Online quantum error correction with a superconducting decoder for surface code," in *Proc. 58th ACM/IEEE Des. Autom. Conf.*, 2021, pp. 451–456, doi: 10.1109/DAC18074.2021.9586326.
- [33] S. Varsamopoulos, K. Bertels, and C. G. Almudever, "Comparing neural network based decoders for surface codes," *IEEE Trans. Comput.*, vol. 69, no. 2, pp. 300–311, Feb. 1, 2020.
- [34] T. Fösel, P. Tighineanu, T. Weiss, and F. Marquardt, "Reinforcement learning with neural networks for quantum feedback," *Phys. Rev. X*, vol. 8, no. 3, Jul. 2018, Art. no. 0 31084, doi: 10.1103/PhysRevX.8.031084.
- [35] P. Baireuther, T. E. O'Brien, B. Tarasinski, and C. W. J. Beenakker, "Machine-learning-assisted correction of correlated qubit errors in a topological code," *Quantum*, vol. 2, p. 48, Jan. 2018, doi: 10.22331/q-2018-01-29-48.
- [36] G. Torlai and R. G. Melko, "Neural decoder for topological codes," *Phys. Rev. Lett.*, vol. 119, no. 3, Jul. 2017, Art. no. 30501, doi: 10.1103/Phys-RevLett.119.030501.
- [37] S. Krastanov and L. Jiang, "Deep neural network probabilistic decoder for stabilizer codes," *Sci. Rep.*, vol. 7, Sep. 2017, Art. no. 11003, doi: 10.1038/s41598-017-11266-1.
- [38] C. Chamberland and P. Ronagh, "Deep neural decoders for near term faulttolerant experiments," *Quantum Sci. Technol.*, vol. 3, no. 4, Oct. 2018, Art. no. 44002, doi: 10.1088/2058-9565/aad1f7.

- [39] X. Ni, "Neural network decoders for large-distance 2D toric codes," *Quantum*, vol. 4, p. 310, Aug. 2020, doi: 10.22331/q-2020-08-24-310.
- [40] S. Varsamopoulos, K. Bertels, and C. G. Almudever, "Decoding surface code with a distributed neural network–based decoder," *Quantum Mach. Intell.*, vol. 2, no. 1, pp. 1–12, 2020, doi: 10.1007/s42484-020-00015-9.
- [41] C. Horsman, A. G. Fowler, S. Devitt, and R. V. Meter, "Surface code quantum computing by lattice surgery," *New J. Phys.*, vol. 14, no. 12, Dec. 2012, Art. no. 123011, doi: 10.1088/1367-2630/14/12/123011.
- [42] R. Versluis *et al.*, "Scalable quantum circuit and control for a superconducting surface code," *Phys. Rev. Appl.*, vol. 8, Sep. 2017, Art. no. 34021, doi: 10.1103/PhysRevApplied.8.034021.
- [43] D. Poulin, "Optimal and efficient decoding of concatenated quantum block codes," *Phys. Rev. A*, vol. 74, no. 5, Nov. 2006, Art. no. 52333, doi: 10.1103/PhysRevA.74.052333.
- [44] L. Riesebos, X. Fu, C. G. Almudever, and K. Bertels, "Pauli frames for quantum computer architectures," in *Proc. 54th Annu. Des. Autom. Conf.*, 2017, Art. no. 76, doi: 10.1145/3061639.3062300.
- [45] T. E. O'Brien, B. Tarasinski, and L. DiCarlo, "Density-matrix simulation of small surface codes under current and projected experimental noise," *Nature Partner J. Quantum Inf.*, vol. 3, Sep. 2017,Art. no. 39, doi: 10.1038/s41534-017-0039-x.
- [46] J. M. Boter *et al.*, "The spider-web array—A sparse spin qubit array," 2021, arXiv:2110.00189, doi: 10.48550/arXiv.2110.00189.
- [47] J. P. Van Dijk, E. Charbon, and F. Sebastiano, "The electronic interface for quantum processors," *Microprocessors Microsyst.*, vol. 66, pp. 90–101, 2019, doi: 10.1016/j.micpro.2019.02.004.
- [48] E. Jeffrey et al., "Fast accurate state measurement with superconducting qubits," Phys. Rev. Lett., vol. 112, May 2014, Art. no. 190504, doi: 10.1103/PhysRevLett.112.190504.
- [49] C. Barthel *et al.*, "Fast sensing of double-dot charge arrangement and spin state with a radio-frequency sensor quantum dot," *Phys. Rev. B*, vol. 81, Apr. 2010, Art. no. 161308, doi: 10.1103/PhysRevB.81.161308.
- [50] K. Hornik, M. Stinchcombe, and H. White, "Multilayer feedforward networks are universal approximators," *Neural Netw.*, vol. 2, no. 5, pp. 359–366, 1989, doi: 10.1016/0893-6080(89)90020-8.
- [51] D. P. Kingma and J. Ba, "ADAM: A method for stochastic optimization," Dec. 2014, arXiv:1412.6980, doi: 10.48550/arXiv.1412.6980.
- [52] C. Bishop, Neural Networks for Pattern Recognition. Oxford, U.K.: Clarendon Press, 1995.
- [53] R. Sutton, "The bitter lesson," Mar. 2019. [Online]. Available: http://www.incompleteideas.net/IncIdeas/BitterLesson.html
- [54] Y. le Cun, "Generalization and network design strategies," Dept. Comput. Sci., Univ. Toronto, Toronto, ON, Canada, Tech. Rep. CRG-TR-89-4, 1989.
- [55] T. Wagner, H. Kampermann, and D. Bruß, "Symmetries for a high-level neural decoder on the toric code," *Phys. Rev. A*, vol. 102, Oct. 2020, Art. no. 42411, doi: 10.1103/PhysRevA.102.042411.
- [56] M. Verhelst, "Deep learning processor survey," 2020. [Online] Available: http://www.esat.kuleuven.be/ mverhels/DLICsurvey.html
- [57] B. Parhami, Algorithms and Design Methods for Digital Computer Arithmetic. New York, NY, USA: Oxford Univ. Press, 2012.
- [58] C. R. Baugh and B. A. Wooley, "A two's complement parallel array multiplication algorithm," *IEEE Trans. Comput.*, vol. 22, no. 12, pp. 1045–1047, Dec. 1973, doi: 10.1109/T-C.1973.223648.
- [59] C. S. Wallace, "A suggestion for a fast multiplier," *IEEE Trans. Electron. Comput.*, vol. EC-13, no. 1, pp. 14–17, Feb. 1964, doi: 10.1109/PGEC.1964.263830.
- [60] B. M. Terhal, "Quantum error correction for quantum memories," *Rev. Mod. Phys.*, vol. 87, pp. 307–346, Apr. 2015, doi: 10.1103/RevMod-Phys.87.307.
- [61] C. Chamberland, P. Iyer, and D. Poulin, "Fault-tolerant quantum computing in the pauli or clifford frame with slow error diagnostics," *Quantum*, vol. 2, p. 43, Jan. 2018, doi: 10.22331/q-2018-01-04-43.
- [62] H. Homulle, "Cryogenic electronics for the read-out of quantum processors," Ph.D. dissertation, Dept. Quantum Comput. Eng., Delft Univ. Technology, Delft, The Netherlands, Jun. 2019, doi: 10.4233/uuid:e833f394-c8b1-46e2-86b8-da0c71559538.
- [63] J. van Dijk et al., "Cryo-CMOS for analog/mixed-signal circuits and systems," in Proc. IEEE Custom Integr. Circuits Conf., 2020, pp. 1–8, doi: 10.1109/CICC48029.2020.9075882.

- [64] P. A. 't Hart, J. P. G. van Dijk, M. Babaie, E. Charbon, A. Vladimircscu, and F. Sebastiano, "Characterization and model validation of mismatch in nanometer CMOS at cryogenic temperatures," in *Proc. 48th Eur. Solid-State Device Res. Conf.*, 2018, pp. 246–249, doi: 10.1109/ESS-DERC.2018.8486859.
- [65] L. Deferm, E. Simoen, B. Dierickx, and C. Claeys, "Anomalous latch-up behaviour of CMOS at liquid helium temperatures," *Cryogenics*, vol. 30, no. 12, pp. 1051–1055, 1990, doi: 10.1016/0011-2275(90)90206-R.
- [66] C. J. Marshall *et al.*, "Mechanisms and temperature dependence of single event latchup observed in a CMOS readout integrated circuit from 16– 300 K," *IEEE Trans. Nucl. Sci.*, vol. 57, no. 6, pp. 3078–3086, Dec. 2010, doi: 10.1109/TNS.2010.2085018.
- [67] E. Schriek, "A low-power standard cell library for cryogenic operation," Master's thesis, Dept. Quantum Comput. Eng., Delft Univ. Technology, Delft, The Netherlands, 2018.



Ramon W. J. Overwater received the B.S. degree in electrical engineering from the Delft University of Technology, Delft, The Netherlands, in 2016 and the double M.S. degrees in microelectronics (*cum laude*) and computer engineering in 2019 from the Delft University of Technology, where he is currently working toward the Ph.D. degree in cryogenic electrical engineering.

His research interests include cryogenic electronic characterization, mixed-signal design, and high-performance computing.



Masoud Babaie (Member, IEEE) received the B.Sc. (Hons.) degree in electrical engineering from the Amirkabir University of Technology, Tehran, Iran, in 2004, the M.Sc. degree in electrical engineering from the Sharif University of Technology, Tehran, Iran, in 2006, and the Ph.D. degree (*cum laude*) in electrical engineering from the Delft University of Technology, Delft, The Netherlands, in 2016.

From 2006 to 2011, he was with the Kavoshcom Research and Development Group, Tehran,

Iran, where he was involved in designing wireless communication systems. From 2014 to 2015, he was a Visiting Scholar Researcher with the Berkeley Wireless Research Center, Berkeley, CA, USA. In 2016, he joined the Delft University of Technology, where he is currently a tenured Assistant Professor. He has authored or coauthored one book, three book chapters, 11 patents, and more than 70 technical articles. His research interests include RF/millimeter-wave integrated circuits and systems for wireless communications and cryogenic electronics for quantum computation.

Dr. Babaie was a co-recipient of the 2015–2016 IEEE Solid-State Circuits Society Pre-Doctoral Achievement Award, the 2019 IEEE ISSCC Demonstration Session Certificate of Recognition, the 2020 IEEE ISSCC Jan Van Vessem Award for Outstanding European Paper, and the 2022 IEEE CICC Best Paper Award. He was the recipient of the Veni Award from the Netherlands Organization for Scientific Research (NWO) in 2019. He also serves on the Technical Program Committee of the IEEE International Solid-State Circuits Conference (ISSCC) and the Co-Chair for the Emerging Computing Devices and Circuits Subcommittee of the IEEE European Solid-State Circuits Conference (ESSCIRC).



Fabio Sebastiano (Senior Member, IEEE) received the B.Sc. (*cum laude*) and M.Sc. (*cum laude*) degrees in electrical engineering from the University of Pisa, Pisa, Italy, in 2003 and 2005, respectively, the M.Sc. degree (*cum laude*) in electrical engineering from the Sant–Anna School of Advanced Studies, Pisa, Italy, in 2006, and the Ph.D. degree in electrical engineering from the Delft University of Technology, Delft, The Netherlands, in 2011.

From 2006 to 2013, he was with NXP Semi-

conductors Research, Eindhoven, The Netherlands, where he researched fully integrated CMOS frequency references, nanometer temperature sensors, and area-efficient interfaces for magnetic sensors. In 2013, he joined the Delft University of Technology, where he is currently an Associate Professor and the Research Lead of the Quantum Computing Division of QuTech. He has authored or coauthored one book, 11 patents, and more than 90 technical publications. His main research interests are cryogenic electronics, quantum computing, sensor read-outs, and fully integrated frequency references.

Dr. Sebastiano is on the technical program committee of the ISSCC, the RFIC Symposium, and the IMS, and he is currently an Associate Editor for the IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION. He was co-recipient of the 2008 ISCAS Best Student Paper Award, the 2017 DATE best IP award, the ISSCC 2020 Jan van Vessem Award for Outstanding European Paper, and the 2022 CICC Best Paper Award. He was the Distinguished Lecturer of the IEEE Solid-State Circuit Society.