

Beam Search Decoder for Quantum Low-Density Parity-Check Codes

Min Ye[✉], Dave Wecker[✉], and Nicolas Delfosse[✉]

IonQ Inc., Bothell, Washington, USA

 (Received 23 December 2025; revised 12 May 2026; accepted 26 May 2026; published 1 July 2026)

We propose a decoder for quantum low-density parity-check (LDPC) codes based on a beam search heuristic guided by belief propagation (BP). Our beam search decoder applies to all quantum LDPC codes and achieves different speed-accuracy tradeoffs by tuning its parameters such as the beam width. We perform numerical simulations under circuit level noise for the $[[144, 12, 12]]$ bivariate bicycle (BB) code at noise rate $p = 10^{-3}$ to estimate the logical error rate and the 99.9 percentile runtime and we compare with the BP-OSD decoder which has been the default quantum LDPC decoder for the past six years. A variant of our beam search decoder with a beam width of 64 achieves a $17\times$ reduction in logical error rate. With a beam width of 8, we reach the same logical error rate as BP-OSD with a $26.2\times$ reduction in the 99.9 percentile runtime. We identify the beam search decoder with beam width of 32 as a promising candidate for trapped ion architectures because it achieves a $5.6\times$ reduction in logical error rate with a 99.9 percentile runtime per syndrome extraction round below 1 ms at $p = 5 \times 10^{-4}$. Remarkably, this is achieved in software on a single core, without any parallelization or specialized hardware (FPGA, ASIC), suggesting one might only need three 32-core CPUs to decode a trapped ion quantum computer with 1000 logical qubits.

DOI: [10.1103/6k5x-ztqt](https://doi.org/10.1103/6k5x-ztqt)

I. INTRODUCTION

Classical low-density parity-check (LDPC) codes [1–3] are widely adopted in classical information processing (WiFi, 5 G mobile, and flash memory). One of the main reasons for the success of LDPC codes is that they come with a fast and accurate decoder, the so-called belief propagation (BP) decoder [4].

The generalization of LDPC codes to the quantum setting was originally proposed by Mackay *et al.* [5]. Later several more efficient constructions of quantum LDPC codes were introduced, such as hypergraph product (HGP) codes [6], two-block codes [7], balanced product codes [8], and the recently discovered asymptotically good quantum LDPC codes [9,10]. Moreover, circuit level simulations proved that several instances of quantum LDPC codes outperform surface codes such as hyperbolic codes [11], bivariate bicycle (BB) codes [12] or their BB5 variant [13], radial code [14], or HGP codes [15,16]. Precisely, they achieve the same logical error rate as surface codes with a substantially smaller qubit overhead.

However, to make quantum LDPC codes practical, a fast and accurate decoder is needed. BP is fast but it does not

perform well in general when applied to quantum LDPC codes. The first issue comes from the definition of the decoding problem. In the classical case, BP is designed to estimate the marginal error probability of each bit. However, the concept of marginal error probability for a single qubit in a stabilizer code is not well defined. Since any error is equivalent to its product with a stabilizer, an error acting non-trivially on a specific qubit is equivalent to another error that acts trivially on that same qubit [17]. The second issue is that BP often fails to converge to a valid correction. To explain this problem, recall that BP works as a message-passing algorithm, sending data through the edges of the graph representing the code, which we call the Tanner graph [18]. For BP to be accurate, the code must be designed in such a way that its Tanner graph is cycle-free, so that a message sent from a node cannot loop back to its sender, avoiding risks of inconsistencies. For a non-trivial code, it is impossible to remove all cycles but removing short cycles is enough to ensure a high accuracy for BP [19]. Unfortunately, the structure of quantum LDPC codes makes short cycles unavoidable [5]. As a result, some marginal probabilities oscillate and BP either fails to converge, reducing its accuracy, or BP converges slowly, reducing its speed.

The BP-OSD decoder, proposed in 2019 by Panteleev and Kalachev [20] and improved in Refs. [21,22], has been the default decoder for quantum LDPC code simulations over the past six years. It is far more accurate than BP, achieving logical error rates orders of magnitude better.

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

However, it is too slow for practical applications because it relies on a matrix inversion implemented in cubic complexity by Gaussian elimination. This matrix inversion is also the bottleneck when adapting the union find decoder to quantum LDPC codes [23].

The ambiguity clustering decoder [24] and the BP-LSD decoder [25] eliminate the BP-OSD bottleneck by partitioning the matrix inversion into sub-problems that can be resolved independently, effectively reducing the average runtime, but these approaches do not improve the accuracy of BP-OSD. Moreover, the worst-case runtime of these decoders remains superlinear, which impacts the tail of the runtime distribution. For example, in Fig. 9 of [25], we see that the tail of the runtime distribution of BP-LSD reaches runtimes that are two orders of magnitude larger than the average. To avoid this issue, which could lead to a large qubit and time overhead for fault-tolerant quantum computation [26,27], we track both the average runtime and the 99.9 percentile runtime of our decoder.

Another strategy to improve BP for quantum LDPC codes is to modify the message-passing procedure by adding or removing constraints in order to suppress the impact of short cycles and the marginal oscillations. Recall that BP sends messages back and forth between error nodes, representing potential error sources, and detector nodes associated with syndrome measurement results. Poulin and Chung proposed to freeze the value of an error node neighboring an unsatisfied detector node in the Tanner graph [28]. Stabilizer inactivation (SI) removes the value of the least reliable stabilizer node [29]. Guided decimation (BP-GD) fixes the value of the most reliable qubit node [30,31]. Guided decimation is combined with backtracking in the decision-tree decoder of [32]. The ordered Tanner forest post-processing (BP-OTF) is executed in a cycle-free subgraph of the Tanner graph [33]. BP-Relay introduces a memory in BP [34]. Remarkably, SI, BP-GD, the decision-tree decoder and BP-Relay achieve a better accuracy than BP-OSD for the $[[144, 12, 12]]$ BB code with circuit level noise. The runtime distribution tail was not analyzed in these papers. When implemented on FPGA, the runtime of BP Relay seems promising for applications to superconducting qubits but the introduction of a long-term memory makes it inherently sequential which might lead to a long runtime tail at the software level [35].

Different parallel variants of BP were proposed. The automorphism-based decoder of [36] runs BP in parallel over many inputs obtained by applying an automorphism of the code. The parallel BP decoder of [37] generates several initial configurations by flipping some of the most oscillating qubits and runs BP in parallel for all these initial configurations. These decoders are embarrassingly parallel but they do not significantly improve the accuracy of BP-OSD at the circuit level.

Optimization algorithms have also been used to design decoders for quantum LDPC codes. The decision-tree

decoder of [32] explores the tree of all fault configurations starting from the empty set. Wu *et al.* treated the decoding problem as a linear program and solved it together with its dual to obtain performance guarantees [38]. An integer programming decoder and a decoder based on A^* search were considered in Ref. [39]. These four decoders can be useful to probe the optimal performance of small codes but they are too slow for real-time decoding. A recent branch and bound decoder might be faster but it was not analyzed under circuit level noise [40].

In this work, we propose a beam search decoder which is simultaneously fast, accurate, easy to parallelize, and flexible. The decoder is parametrized by the beam width and other parameters. We estimate the performance of our beam search decoder with numerical simulations for the $[[144, 12, 12]]$ bivariate bicycle (BB) code under circuit level noise with noise rate $p = 10^{-3}$. The simulations are performed on a 2023 M3 processor on a single core without any parallelization. Our beam search decoder achieves a logical error rate that is up to $17\times$ better than BP-OSD. For a beam width of 8, we achieve the same logical error rate, a $4.6\times$ reduction of the average runtime and a $26.2\times$ reduction of the 99.9 percentile runtime compared with BP-OSD. For a beam width of 32, we achieve a $5.6\times$ reduction of the logical error rate, a $2.8\times$ reduction of the average runtime and a $20.4\times$ reduction of the 99.9 percentile runtime compared with BP-OSD. This proves that the beam search decoder can be simultaneously more accurate and faster than BP-OSD.

To examine the practical application of our beam search decoder in a large-scale fault-tolerant trapped ion quantum computer, we consider its performance at lower noise rate. We pick $p = 5 \times 10^{-4}$, which is above the noise rate achievable in today's devices [41], and our goal is to design a high-accuracy decoder whose 99.9 percentile runtime per syndrome extraction round is under 1 ms, which is the expected syndrome extraction time on trapped ion machines. At this noise rate and with the $[[144, 12, 12]]$ BB code, the beam search decoder with beam width of 32 has an average runtime per syndrome extraction round of $270 \mu\text{s}$ and a 99.9 percentile runtime of $940 \mu\text{s}$, which is 24 times better than BP-OSD.

Our work proves that a software-level decoder on a single core, without any parallelization or specialized hardware (FPGA, ASIC) can simultaneously achieve a significantly better logical error rate than BP-OSD and a 99.9 percentile runtime below 1 ms. This suggests that one might only need three 32-core CPUs to decode a trapped ion quantum computer with 1000 logical qubits. For comparison, a fault-tolerant quantum computing architecture based on superconducting qubits and surface codes might need up to 1000 ASICs or FPGA decoders to correct 1000 surface code patches simultaneously.

We further demonstrate the flexibility of our decoder by performing numerical simulations showing that the beam

search decoder also outperforms BP-OSD for other BB codes and for HGP codes [16]. Moreover, it can benefit from the XYZ-decoding, which utilizes both X and Z syndrome outcomes simultaneously for decoding [35].

The rest of this paper is organized as follows. Section II introduces the main ideas of the beam search decoder. In Sec. III, we present our numerical results, followed by concluding remarks in Sec. IV. Additionally, Appendix A reviews the quantum decoding problem and the BP algorithm, while Appendix B provides a detailed algorithmic description of the beam search decoder.

II. OVERVIEW OF THE BEAM SEARCH DECODER

This section gives an overview of the beam search decoder. We first clarify the terminology used in the tree search. A path represents one partial decoding trajectory, specified by the error nodes that have already been fixed and their assigned values. A branch is obtained from a path by choosing one additional error node and fixing it to either 0 or 1, producing two child paths. The decoder maintains a set of active paths, which contains the candidate partial assignments currently being explored. After each branching step, this set is pruned so that only the most reliable paths are kept.

With this terminology, the algorithm is initialized by a standard BP run. If BP does not converge, the posterior log-likelihood ratios (LLRs) are used to identify the least reliable error node. The decoder then creates a single seed path, in which no error node has yet been fixed, but the least reliable error node has been selected as the first branching point. The initial set consists only of this seed path.

The decoder then executes multiple rounds of “masked BP.” In each round, every path in the set is expanded into two new branches by constraining the least reliable error node (identified in the previous round) to 0 and 1, respectively. This newly constrained node is now considered “masked,” meaning that it is effectively removed from the BP calculations. This constraint is added to the path’s history of masked error nodes. To maintain tractable complexity, this branching step is followed by a pruning step: the set is sorted by a reliability score, and only a fixed number of the most reliable paths are kept. This process repeats until one of two termination conditions is met: either a pre-defined number of valid solutions is collected, or the maximum number of masked BP rounds is reached. The decoder then returns the minimum-weight solution it finds. Figure 1 illustrates the workflow of the decoding algorithm, while Algorithm 1 presents a high-level sketch for the special case where the decoder terminates upon finding the first valid solution.

In our decoder, the reliability metric for each error node is the absolute value of the sum of its posterior LLRs over all BP iterations in the current round. We use this sum-based metric rather than the LLR from a single BP iteration

to ensure that oscillating error nodes are correctly identified as unreliable. The reliability score of a path, which is used to prune the set, is then calculated by summing these individual node metrics over all of its associated unmasked error nodes.

Pseudocode and a detailed description of the beam search decoder are provided in Appendix B, and the implementation source code is available on GitHub [42].

Our beam search decoder shares several similarities with the BP-GDG decoder proposed in Ref. [31]. Both algorithms, for example, employ multiple rounds of masked BP and select one error node to fix (or constrain) in each round. The BP-GDG decoder, in its initial 4 rounds, also branches by exploring both values (0 and 1) for the selected error node. However, there are three important differences between our decoder and BP-GDG. First, our decoder systematically branches in every round, whereas BP-GDG only branches for the first 4 rounds. After that, BP-GDG only performs a quick exploration of immediate side branches of the most likely decoding path. Second, and most importantly, we introduce a reliability score to predict a path’s likelihood of success before its completion. To our knowledge, this is a new approach. In contrast, standard decoders (including BP-GDG) can only evaluate a path’s quality after it terminates—either by finding a valid solution and checking its weight, or by failing to converge at the maximum iteration limit. This conventional method must run all paths to completion, even those that are unlikely to succeed. Our reliability score enables the decoder to proactively prune the set, eliminating unpromising paths early and focusing computational effort on more viable candidates, as demonstrated by our simulation results. Third, the node selection strategies differ. Our decoder branches on the least reliable error node (the one with the minimum absolute summed LLR). In contrast, BP-GDG branches on the error node most likely to be 1 (the one with the most negative LLR history).

The beam search decoder has a relatively simple and uniform structure. In every round, the decoder performs the same three operations: branch on the least reliable error node, run masked BP on the resulting paths, and prune the active set using the path reliability score. This contrasts with decoders such as BP-GDG, which combine several stages and branching rules. This uniform structure makes the algorithm easier to describe and more accessible to readers outside the specialized quantum-decoding community, while still providing a tunable accuracy-latency tradeoff through the beam width.

Our decoder is also related in spirit to Tesseract [39], a recent search-based decoder that formulates most-likely-error decoding as a shortest-path problem over the graph of error subsets and applies A^* search with several pruning heuristics, including a beam cutoff. The role of bounded search in our decoder is different. Rather than searching




FIG. 1. Overview of the beam search decoder. (a) The decoder is initialized with a small number of BP iterations. Then, the least reliable error node (red) is selected and two branches are created corresponding to the two possible values of this node. After the first step, BP is replaced by a masked BP ignoring the previously fixed error nodes. (b) The beam search decoder repeats the following three steps: (i) branching over the least reliable error node (ii) running a masked BP and (iii) pruning to reduce the number of paths to the beam width (which is 2 in this figure). The decoder is terminated once a sufficient number of solutions is found or if a maximum number of repetitions is reached.

ALGORITHM 1. Beam search decoder (high-level sketch).

```

1 Run standard BP and return if it converges.
2  $\text{path.next\_pos} \leftarrow$  the least reliable error node in BP.
3  $\text{path.pos\_val\_pairs} \leftarrow$  an empty vector.
4 Initialize  $\text{set} \leftarrow \{\text{path}\}$  with a single element path
5 for  $r = 1, 2, \dots, \text{max\_rounds}$  do
6   Initialize  $\text{next\_set}$  as an empty set.
7   for each  $\text{path} \in \text{set}$  and  $\text{val} \in \{0, 1\}$  do
8      $\text{nextp.pos\_val\_pairs} \leftarrow$ 
        $\text{path.pos\_val\_pairs} \cup (\text{path.next\_pos}, \text{val})$ 
9   Run masked BP masking the nodes in
      $\text{nextp.pos\_val\_pairs}$ .
10  If masked BP converges, return the result.
11   $\text{nextp.next\_pos} \leftarrow$  the least reliable error node
    in masked BP.
12   $\text{nextp.score} \leftarrow$  reliability score of nextp.
13  Add nextp into next_set.
14  Remove the element with the smallest score if
    the size of next_set exceeds the parameter
    beam_width.
15  $\text{set} \leftarrow \text{next\_set}$ 

```

directly over error subsets using an A^* objective, we use BP throughout the search: BP identifies the least reliable error node for branching, masked BP updates the messages associated with each partial path, and the accumulated posterior LLRs from masked BP define the reliability score used for pruning. Thus, Tesseract is a search-first decoder, while our method is a BP-guided beam search decoder. This distinction leads to a different accuracy-latency tradeoff. For example, in Fig. 2 of [39], the $[[144, 12, 12]]$ BB code corresponds to the 288-qubit data point because the horizontal axis counts both data and ancilla qubits. At physical error rate $p = 10^{-3}$, the reported mean decoding time for Tesseract is above 0.1 s per round. By comparison, our decoder runs well below 1 ms per round on average for the same code and physical error rate. This comparison illustrates that Tesseract is designed as a near-optimal search decoder, whereas our decoder is designed for real-time BP-guided decoding with controlled latency.

The comparison between beam search and decision-tree decoders [32] is more subtle, because both methods can be

TABLE I. Name and parameters for different configurations of beam search decoder. The number of iterations in name is calculated as $\text{initial_iters} + \text{max_rounds} * \text{iters_per_round}$. The parameter num_results is omitted from the name if it is equal to 1.

Name	max_rounds	beam_width	initial_iter	iters_per_round	num_results
beam8_230iters	10	8	30	20	1
beam32_340iters	10	32	40	30	1
beam64_640iters	20	64	40	30	1
beam64_32res_640iters	20	64	40	30	32

represented as searches over a tree of partial decoding decisions. However, the underlying tree and the decoding mechanism are different. In a decision-tree decoder, the decoding problem is specified by a parity-check matrix H . Given an input syndrome σ and a partial correction F , a node in the decision tree represents the residual syndrome $\sigma + HF$. The tree is expanded by selecting an unsatisfied check vertex from this residual syndrome and adding one of its neighboring fault vertices to the partial correction. The process continues until the residual syndrome becomes empty. Different decision-tree decoder variants mainly differ in the cost function used to explore this tree. By contrast, the tree in our beam search decoder is defined by BP-guided partial assignments. The decoder branches on the least reliable error node identified by BP, fixes this node to either 0 or 1, and then reruns masked BP on the resulting paths. The accumulated posterior LLRs from masked BP are then used to score and prune incomplete paths. Thus, while both approaches can be represented as searches over a tree, the branching rule, the role of BP, and the pruning mechanism are different: decision-tree decoders are syndrome-driven searches over fault additions, whereas our decoder is a BP-guided beam search over uncertain error-node assignments.

This difference is reflected in the runtime-tail behavior. In the circuit-level simulations of the $[[144, 12, 12]]$ BB code, Fig. 13 of [32] suggests that the BP-DTD runtime distribution at $p = 10^{-3}$ has an average decoding time on the order of 10 ms for the 12-round syndrome extraction circuit, while the tail extends much further, with the 99.9 percentile visually estimated to be on the order of 10^{-1} s. By comparison, for the same code and physical error rate, our decoder has total 12-round average runtime of a few milliseconds (Table II) and 99.9-percentile runtime below 15 ms (Table III). This highlights a key goal of our BP-guided beam search decoder: controlling not only the average runtime, but also the tail of the runtime distribution.

III. SIMULATION RESULTS

We provide simulation results for several configurations of our beam search decoder across three quantum LDPC codes. We compare the performance of these configurations against a baseline BP-OSD decoder [20–22].

Following [25], which compares BP-OSD and BP-LSD, we configure the baseline BP-OSD decoder with 30 min-sum BP iterations and order-10 combination-sweep OSD post-processing. Table I lists the names and corresponding parameters for each beam search decoder configuration used in our simulations.

Figure 2 plots the logical error rate of the $[[144, 12, 12]]$ BB code under circuit-level noise. At physical error rate $p = 10^{-3}$, even the most computationally efficient configuration, beam8_230iters, yields a $1.3\times$ improvement over bp30+osd. Increasing the beam width further enhances performance: beam32_340iters and beam64_640iters reduce the logical error rate by factors of $5.6\times$ and $7.0\times$, respectively. Notably, the most advanced configuration, beam64_32res_640iters, achieves a $17\times$ improvement over BP-OSD.

Tables II and III detail the decoding time statistics for the $[[144, 12, 12]]$ BB code over 12 syndrome extraction rounds, measured on a single core of a 2023 Apple M3 Pro. Table II reports the average decoding time, while Table III lists the 99.9 percentile. We evaluate two physical error rates: $p = 10^{-3}$, relevant for superconducting qubits, and $p = 5 \times 10^{-4}$, applicable to higher-fidelity hardware such




FIG. 2. Simulation results for the $[[144, 12, 12]]$ BB code under circuit-level noise. BP-OSD decoder is configured with 30 min-sum BP iterations followed by order-10 combination-sweep OSD.

TABLE II. Average decoding time in circuit-level noise simulations of the $[[144, 12, 12]]$ BB code with 12 syndrome extraction rounds.

Physical error rate	bp30+osd (ms)	beam8_230iters (ms)	beam32_340iters (ms)	beam64_640iters (ms)
5×10^{-4}	3.55	1.627	3.202	5.881
10^{-3}	10.59	2.318	3.837	6.698

TABLE III. The 99.9 percentile decoding time in circuit-level noise simulations of the $[[144, 12, 12]]$ BB code with 12 syndrome extraction rounds.

Physical error	bp30+osd (ms)	beam8_230iters (ms)	beam32_340iters (ms)	beam64_640iters (ms)
5×10^{-4}	272.5	8.704	11.26	14.53
10^{-3}	289.0	11.01	14.18	17.70

as trapped ion systems. At both error rates, all three beam search configurations improve the 99.9 percentile decoding time by at least $16\times$ compared to the bp30+osd baseline, demonstrating superior capability in mitigating worst-case decoding latency. Specifically, the fastest configuration, beam8_230iters, achieves reductions of $31\times$ and $26\times$ at $p = 5 \times 10^{-4}$ and $p = 10^{-3}$, respectively. Benefiting from the significantly reduced variance in decoding time, the beam32_340iters configuration achieves a 99.9 percentile time of less than 1 ms per syndrome extraction round at $p = 5 \times 10^{-4}$. While the reduction in average time is less pronounced, beam8_230iters still yields a $4.6\times$ speedup at $p = 10^{-3}$.

Figure 3 plots the logical error rate of the $[[90, 8, 10]]$ BB code using two strategies: standard XZ-decoding and XYZ-decoding, following the terminology in Ref. [34]. Since both BB codes and HGP codes simulated in this paper are CSS codes, X and Z errors are decoded

separately, and the total logical error rate is approximated as the sum of the X and Z logical error rates. Recall from Appendix A that the decoding problem is defined by the triple $(\mathbf{H}, \mathbf{A}, \mathbf{p})$. Both strategies utilize the same logical operator matrix $\mathbf{A}$ and error probabilities $\mathbf{p}$; they differ only in the construction of the parity-check matrix $\mathbf{H}$. When decoding X errors, both strategies employ the same logical operator matrix $\mathbf{A}$ consisting of Z logical operators; however, XZ-decoding constructs $\mathbf{H}$ using only Z stabilizers, whereas XYZ-decoding includes both X- and Z-type stabilizers. Similarly, when decoding Z errors, XZ-decoding uses only X stabilizers to form $\mathbf{H}$, while XYZ-decoding employs the full set of stabilizers.

In theory, utilizing the full set of stabilizers yields more information than restricting the decoder to a single stabilizer type. This advantage arises because the circuit-level noise model assumes that qubits undergo depolarizing noise during unitary gates and idling time. Since depolarizing noise includes a Pauli-Y component, which triggers both X- and Z-type stabilizers, employing both types simultaneously captures strictly more information.

However, a major drawback of utilizing the full set of stabilizers is the induction of length-4 cycles in the corresponding Tanner graph, which is known to degrade the performance of BP and many BP-based decoders. As shown in Fig. 3, the bp30+osd decoder performs much worse under XYZ-decoding than XZ-decoding, despite the fact that XYZ-decoding theoretically provides more information. In contrast, beam64_640iters effectively leverages this additional information, reducing the logical error rate by approximately $2\times$ compared to XZ-decoding in the range $0.002 \leq p \leq 0.004$. A likely explanation is that the short cycles in XYZ-decoding hinder BP-based decoders by causing the posterior LLRs of error nodes to oscillate between positive and negative values. However, because the reliability score in our beam search decoder is derived from the absolute sum of posterior LLRs, it can successfully identify these oscillating nodes as unreliable and exclude them from subsequent BP iterations.




FIG. 3. Simulation results for the $[[90, 8, 10]]$ BB code under circuit-level noise. The suffix _XYZ in the legend denotes XYZ-decoding, which utilizes both X and Z syndrome outcomes simultaneously for decoding.




FIG. 4. Simulation results for the $[[450, 32, 8]]$ HGP code [16].

Finally, to verify the robustness of our decoder across different quantum LDPC families, we perform simulations of the $[[450, 32, 8]]$ cyclic HGP code introduced in Ref. [16]. The logical error rates for three decoders are presented in Fig. 4. Our beam search decoders continue to exhibit lower logical error rates than the bp30+osd baseline. In particular, at a relatively high noise rate $p = 0.005$, the beam8_230iters and beam64_640iters configurations achieve improvements of 1.4× and 3.7×, respectively.

IV. CONCLUSION

In this work, we propose a simple beam search decoder for quantum LDPC codes. We show that it is simultaneously fast and accurate, significantly outperforming the BP-OSD decoder.

We demonstrate that the beam search decoder with beam width of 32 achieves a lower logical error rate than BP-OSD while satisfying the runtime requirements for trapped ion or neutral atom quantum computers, which are expected to be on the order of 1 ms. Notably, this configuration meets strict latency constraints, achieving a sub-millisecond runtime not only on average but also at the 99.9 percentile.

It is a folklore that superconducting qubits, which must be decoded within a micro-second window, might require a supercomputer to perform decoding for a large-scale fault-tolerant quantum computer. This encouraged researchers to optimize decoders at the micro-architecture level [43] and to build hardware-level decoders on FPGAs or ASICs [35,44–51]. In contrast, because the decoding time budget of trapped ions and neutral quantum computers is significantly larger, our decoder can be executed on a CPU without any parallelization or specialized hardware implementation.

The BP algorithm appeared independently in different communities. It is used in coding theory, statistical physics, machine learning and optimization and it led to fruitful exchanges of ideas between these communities [2,3,52,53]. In the quantum setting, the interface between quantum error correction and statistical physics is an active research topic [54–63]. We anticipate that the beam search decoder will play a role at this interface, similar to the connecting role of BP in the classical case.

Beyond its practical utility, the simplicity of the beam search decoder makes it a promising candidate for theoretical analysis. We hope this feature will facilitate the proof of asymptotic results for quantum LDPC codes, much like the BP decoder proved crucial in the design of capacity-achieving classical LDPC codes [3].

ACKNOWLEDGMENTS

The authors thank Aharon Brodutch, Edwin Tham, Felix Tripiier, Joe Latone and John Gamble and the whole IonQ team for insightful discussions.

DATA AVAILABILITY

The data that supports the findings of this article is openly available [42].

APPENDIX A: THE DECODING PROBLEM AND THE BP ALGORITHM

In a quantum fault-tolerant computing system, a syndrome extraction circuit periodically measures syndromes, and a decoder corrects errors based on the measurement results. Simulations of this process typically use one of two noise models. The circuit-level noise model, a realistic approach widely used in simulations, assumes every operation in the syndrome extraction circuit is a potential source of error. In contrast, the much simpler code capacity noise model only assumes idling noise on the data qubits, while all other operations are considered perfect. In this paper, we adopt the more realistic circuit-level noise model.

The decoding problem is formally defined by a parity-check matrix $\mathbf{H} \in \mathbb{F}_2^{M \times N}$, a logical operator matrix $\mathbf{A} \in \mathbb{F}_2^{K \times N}$, and a probability vector $\mathbf{p} = (p_0, \dots, p_{N-1})$. The N columns of $\mathbf{H}$ and $\mathbf{A}$ correspond to the set of all possible error sources in the circuit-level noise model, with p_i being the error probability of the i th source. The M rows of $\mathbf{H}$ correspond to the “detectors,” which are either syndrome measurement results or the XOR of two measurement results from consecutive rounds. An entry $\mathbf{H}_{ij} = 1$ if and only if the j th error source flips the i th detector. Similarly, the K rows of $\mathbf{A}$ correspond to the logical operators, and $\mathbf{A}_{ij} = 1$ if and only if the j th error source flips the i th logical operator. The parity-check matrix $\mathbf{H}$ is also referred to as the detector error model

in Stim [64]. Let $\mathbf{e} \in \mathbb{F}_2^N$ denote the binary vector of (unknown) error locations, and let $\mathbf{s} = \mathbf{H}\mathbf{e} \in \mathbb{F}_2^M$ be the observed syndrome. The decoder takes $\mathbf{s}$ as input and attempts to compute a correction $\hat{\mathbf{e}} \in \mathbb{F}_2^N$ such that it matches the syndrome ($\mathbf{H}\hat{\mathbf{e}} = \mathbf{s}$) and preserves the logical state ($\mathbf{A}\hat{\mathbf{e}} = \mathbf{A}\mathbf{e}$).

Both the new decoder proposed in this paper and our baseline for comparison, the BP-OSD decoder, are based on BP. We begin by reviewing BP decoding with the min-sum update rule. This algorithm operates on the Tanner graph, a bipartite graph constructed from the parity-check matrix $\mathbf{H}$. The graph's vertices are divided into N error nodes ($e_0, \dots, e_{N-1}$), corresponding to the columns of $\mathbf{H}$, and M detector nodes ($d_0, \dots, d_{M-1}$), corresponding to the rows of $\mathbf{H}$. An edge connects a detector node d_i to an error node e_j if and only if $\mathbf{H}_{ij} = 1$. We use $\mathcal{N}(d_i)$ and $\mathcal{N}(e_j)$ to denote the set of neighbors for d_i and e_j , respectively. The BP decoder functions by passing messages back and forth along these edges over multiple iterations. We use $D_{i \rightarrow j}(t)$ and $E_{j \rightarrow i}(t)$ to denote the detector-to-error and error-to-detector messages in the t th iteration, respectively. The decoder takes the syndrome vector $\mathbf{s} = (s_0, \dots, s_{M-1}) \in \mathbb{F}_2^M$ as input. To initialize, the decoder first calculates the prior LLR for each error node as $\Lambda_j = \log \frac{1-p_j}{p_j}$. The initial error-to-detector messages in iteration 0 are then set to these priors: $E_{j \rightarrow i}(0) = \Lambda_j$ for all connected pairs (d_i, e_j) . In each subsequent iteration $t \geq 1$, the message passing proceeds in two steps. First, the detector-to-error messages are updated using the min-sum rule:

$$D_{i \rightarrow j}(t) = (-1)^{s_i} \cdot \prod_{e_{j'} \in \mathcal{N}(d_i) \setminus \{e_j\}} \text{sign}(E_{j' \rightarrow i}(t-1)) \cdot \min_{e_{j'} \in \mathcal{N}(d_i) \setminus \{e_j\}} |E_{j' \rightarrow i}(t-1)|. \quad (\text{A1})$$

Second, the error-to-detector messages are updated by summing the prior LLR with all other incoming detector messages:

$$E_{j \rightarrow i}(t) = \Lambda_j + \sum_{d_{i'} \in \mathcal{N}(e_j) \setminus \{d_i\}} D_{i' \rightarrow j}(t). \quad (\text{A2})$$

After updating the messages, the decoder also calculates the posterior LLR $\Lambda_j(t)$ for each error node as

$$\Lambda_j(t) = \Lambda_j + \sum_{d_{i'} \in \mathcal{N}(e_j)} D_{i' \rightarrow j}(t). \quad (\text{A3})$$

Based on this, a hard decision $\hat{e}_j(t)$ is made:

$$\hat{e}_j(t) = \begin{cases} 0 & \text{if } \Lambda_j(t) > 0, \\ 1 & \text{if } \Lambda_j(t) \leq 0. \end{cases} \quad (\text{A4})$$

Define the vector $\hat{\mathbf{e}}(t) = (\hat{e}_0(t), \dots, \hat{e}_{N-1}(t))$. The iterative BP decoding stops at iteration t if the syndrome equation $\mathbf{H}\hat{\mathbf{e}}(t) = \mathbf{s}$ is satisfied, or if t exceeds a pre-fixed maximum number of iterations.

APPENDIX B: DETAILED DESCRIPTION OF THE BEAM SEARCH DECODER

In this section, we present a detailed description of our new beam search decoder. Given the extensive notation required to define the algorithm, a summary of the relevant variables is provided in Table IV to assist readers.

The algorithm maintains a set of decoding paths. It is initialized by running some standard BP iterations to generate a single “seed” path. The posterior LLRs from this initial run are used to identify the least reliable error node, which is stored in `path.next_pos` and will be masked in future BP iterations.

The decoder then iterates through multiple rounds of “masked BP.” In each round, every path in the current set is expanded into two new branches by setting its `path.next_pos` node to 0 and 1, respectively. Each new path inherits its parent’s history of masked error nodes (stored in `path.pos_val_pairs`) and appends the new (node, value) pair to the history, using the node index from `path.next_pos` and the value (0 or 1) just explored. A round of masked BP iterations is then run for each new path. The posterior LLRs in these BP iterations are used to determine two things: (1) the next least reliable unmasked node (which is stored as the new `path.next_pos`), and (2) an overall reliability score for the path. This branching process doubles the size of set each round, so to maintain tractable complexity, the set is pruned back down to a fixed parameter `beam_width` by keeping only the paths with the highest reliability scores.

We use $|\text{sum_LLR}[j]|$ as a reliability metric for each error node e_j , where $\text{sum_LLR}[j]$ is the summation of $\Lambda_j(t)$ over all BP iterations in the current round. The least reliable node to store in `path.next_pos` is the one with the smallest $|\text{sum_LLR}[j]|$. This sum-based metric is more robust than using the instantaneous magnitude $|\Lambda_j(t)|$ from the final iteration. While the instantaneous LLR reflects the decoder’s confidence at that single step, it can be misleading. For example, an unreliable node whose posterior LLR oscillates between large positive and negative values would have a (correctly) low summed reliability, even if its final $|\Lambda_j(t)|$ is large.

TABLE IV. Variables and notation used in the beam search decoder.

Notation	Meaning
<code>path</code>	An object representing a unique decoding path in the branching tree.
<code>path.pos_val_pairs</code>	An array containing indices and values of the error nodes that have been fixed for this path.
<code>path.edge_msgs</code>	A snapshot of all error-to-detector messages from the last BP iteration, used to initialize the next round of masked BP for this path.
<code>path.score</code>	The path reliability score, used for sorting and pruning.
<code>path.next_pos</code>	The index of the least reliable error node e_j chosen as the branching point for the next round of masked BP iterations.
<code>nextp</code>	The new child path object obtained by branching at <code>path.next_pos</code> from the current path
<code>set</code>	The set of active candidate paths currently being processed.
<code>next_set</code>	The set of candidate paths for the subsequent round, generated by branching every path in <code>set</code> .
<code>results</code>	A container for valid error vectors $\hat{\mathbf{e}}$ found (where $\mathbf{H}\hat{\mathbf{e}} = \mathbf{s}$).
<code>beam_width</code>	The maximum number of candidate paths allowed in <code>set</code> at each step.
<code>max_rounds</code>	The maximum depth of the branching tree (number of branching rounds) allowed.
<code>initial_iters</code>	The number of BP iterations performed on the root node before the beam search begins.
<code>iters_per_round</code>	The number of masked BP iterations performed to update a specific path after an error node is fixed.
<code>num_results</code>	The target number of valid solutions required to terminate the decoding process early.
<code>sum_LLR[j]</code>	The cumulative posterior LLR for error node e_j .

The overall reliability score for each path is then calculated by summing the per-node reliability metrics, $|\text{sum_LLR}[j]|$, over all unmasked error nodes and normalizing by the number of BP iterations in that round. This normalization is crucial because different paths may have run for a different number of iterations (due to early returns), and the division makes their scores comparable. In contrast, this normalization was not required when selecting the least reliable node (as described previously), since all nodes within a single path are guaranteed to go through the same number of iterations.

The beam search decoder has a parameter `num_results` that controls its termination logic. If `num_results` is set to 1, the decoder returns the first valid decoding solution it finds. Otherwise (for `num_results` > 1), the decoder maintains a set `results`, into which it adds every valid solution it discovers. As soon as this set reaches the target size of `num_results`, the decoder terminates and returns the minimum-weight error vector from `results`, where the weight of an error vector $\hat{\mathbf{e}} = (\hat{e}_0, \dots, \hat{e}_{N-1})$ is defined as

$$\text{wt}(\hat{\mathbf{e}}) = \sum_{j=0}^{N-1} \hat{e}_j \log \frac{1 - p_j}{p_j}$$

Finally, a key optimization in our algorithm is the order of decoding different paths. In each round, the paths in the `set` are sorted by the reliability score and explored in descending order. We prioritize high-score paths first because they are intuitively more likely to lead to a correct decoding result.

A high-level sketch of the beam search decoder is presented in Algorithm 1 for the `num_results` = 1

case. We now provide a detailed explanation of the full algorithm, shown in Algorithm 3. The decoder has five parameters: `max_rounds`, `beam_width`, `initial_iters`, `iters_per_round`, and `num_results`. We will describe the logic for the simplest case of `num_results` = 1.

The decoder begins by running `initial_iters` standard BP iterations [following Eqs. (A1)–(A4)]. If a valid solution is found (i.e., $\mathbf{H}\hat{\mathbf{e}}(t) = \mathbf{s}$), it is returned immediately. Otherwise, the decoder calculates a reliability metric $|\text{sum_LLR}[j]|$ for every error node e_j , where $\text{sum_LLR}[j]$ is the summation of $\Lambda_j(t)$ over all `initial_iters` iterations.

This information is used to build the “seed” path object, which initializes the `set` for the subsequent decoding rounds. This object contains four fields. The first, `path.edge_msgs`, stores all error-to-detector messages from the last BP iteration, which will be used to “warm-start” the next round of masked BP. The second, `path.pos_val_pairs`, stores the (position, value) pairs for all masked error nodes; for example, a pair $(j, v) \in \text{path.pos_val_pairs}$ means that error node e_j is fixed to $v \in \{0, 1\}$, and that e_j is excluded from all the calculations in the masked BP. This field is initialized as an empty vector. The third field, `path.next_pos`, stores the index of the next error node to mask, which is set to the node e_j with the minimum reliability metric $|\text{sum_LLR}[j]|$. The final field, `path.score`, is the path’s reliability score used for ranking. It is initialized to 0.

The decoder then executes up to `max_rounds` rounds of masked BP. As described in the high-level overview, each round involves expanding the `set` by a factor of two

ALGORITHM 2. Function masked_BP.

```

1 Function masked_BP(edge_msgs, pos_val_pairs, s,
  max_iters)
  Input: a vector edge_msgs used to initialize
    error-to-detector messages, a vector
    pos_val_pairs specifying the masked error
    nodes and their values, a syndrome s, and
    maximum number of iterations max_iters
  Output: A decoded error vector  $\hat{\mathbf{e}}$  or declare
    failure
2  $\Lambda_j \leftarrow \log \frac{1-p_j}{p_j}$  for  $j = 0, 1, \dots, N-1$ 
3  $[E_{j \rightarrow i}(0): (d_i, e_j) \text{ is connected}] \leftarrow \text{edge\_msgs}$ 
4 for each  $(j, v)$  pair in pos_val_pairs do
5   if  $v = 1$  then
6     Flip all  $s_i$  such that  $(d_i, e_j)$  are connected
      in the Tanner graph
7 Error node  $e_j$  is said to be masked if
   $(j, 0) \in \text{pos\_val\_pairs}$  or  $(j, 1) \in \text{pos\_val\_pairs}$ 
8 for  $t = 1, 2, \dots, \text{max\_iters}$  do
9   Update messages and posterior LLRs
    according to Eqs. (A1)–(A3) but ignore all the
    messages to or from a masked error node.
10  Calculate  $\hat{e}_j(t)$  according to Eq. (A4) for all error
    nodes that are not masked. Fill the masked positions
    with 0.
11  if  $\mathbf{H}\hat{\mathbf{e}}(t) = \mathbf{s}$  then
12    for each  $(j, v)$  pair in pos_val_pairs do
13       $\hat{e}_j(t) \leftarrow v$ 
14    return  $\hat{\mathbf{e}}(t)$ 
15 return “Failure”

```

and subsequently pruning it back to beam_width based on path.score if its size exceeds this limit. We now provide a detailed explanation of the masked_BP function, which is formally defined in Algorithm 2.

The masked_BP function takes 4 input parameters. The first, path.edge_msgs, “warm-starts” the BP run by initializing the error-to-detector messages from the previous round’s final state. The second, path.pos_val_pairs, is a vector specifying the error nodes to be masked and their fixed values. The third is the original syndrome vector, s. The final parameter, iters_per_round, specifies the maximum number of iterations for this run. Before iterating, the decoder pre-processes the syndrome based on the values of masked nodes. For every masked error node e_j that is set to 1, all syndrome bits s_i corresponding to detectors connected to e_j are flipped. This transformation creates a modified syndrome, effectively recasting the decoding problem as one where all masked error nodes are fixed to 0. The masked BP then runs for up to iters_per_round iterations, following

ALGORITHM 3. Beam search decoder.

```

Parameters: max_rounds, beam_width,
  initial_iters, iters_per_round,
  num_results
Input: A syndrome vector s.
Output: A decoded error vector  $\hat{\mathbf{e}}$ .
1 Initialize results as an empty set
2 Run BP for initial_iters iterations and break out
  whenever  $\mathbf{H}\hat{\mathbf{e}}(t) = \mathbf{s}$  is satisfied.
3 Insert the decoding result into results if BP succeeds,
  and return this result if num_results = 1.
4 for  $j = 0, 1, \dots, N-1$  do
5   sum_LLRLR[j]  $\leftarrow$  summation of  $\Lambda_j(t)$  over all BP
    iterations.
6 Build an object path with 4 fields:
7 path.edge_msgs  $\leftarrow [E_{j \rightarrow i}(t): d_i, e_j \text{ connect}],$  where  $t$ 
  is the index of the last BP iteration.
8 path.pos_val_pairs  $\leftarrow$  an empty vector
9 path.next_pos  $\leftarrow \arg \min_{0 \leq j < N} |\text{sum\_LLRLR}[j]|$ 
10 path.score  $\leftarrow 0$ 
11 Initialize set  $\leftarrow \{\text{path}\}$  with a single element path
12 for  $r = 1, 2, \dots, \text{max\_rounds}$  do
13   Initialize next_set as an empty set
14   for each path  $\in$  set and val  $\in \{0, 1\}$  do
15     Append (path.next_pos, val) to the end of
      path.pos_val_pairs
16     Run masked_BP(path.edge_msgs,
      path.pos_val_pairs, s, iters_per_
      round)
17   If masked_BP succeeds, insert the decoding
      result into results.
18   If size of results = num_results, return the
      minimum weight element in results.
19   iters  $\leftarrow$  number of iterations actually run in
      masked_BP (considering early return)
20    $\mathcal{U} \leftarrow$  set of unmasked error nodes, i.e.,
       $\{j: 0 \leq j < N, (j, v) \notin$ 
      path.pos_val_pairs for  $v = 0, 1\}$ ,
21   for  $j \in \mathcal{U}$  do
22     sum_LLRLR[j]  $\leftarrow$  summation of  $\Lambda_j(t)$  over all
      iterations in masked_BP.
23   Build an object nexttp with 4 fields:
24   nexttp.edge_msgs  $\leftarrow [E_{j \rightarrow i}(t): d_i, e_j \text{ connect}],$ 
      where  $t$  is the index of the last iteration in
      masked_BP
25   nexttp.pos_val_pairs  $\leftarrow$  path.pos_val_pairs
26   nexttp.next_pos  $\leftarrow \arg \min_{j \in \mathcal{U}} |\text{sum\_LLRLR}[j]|$ 
27   nexttp.score  $\leftarrow \sum_{j \in \mathcal{U}} |\text{sum\_LLRLR}[j]|/\text{iters}$ 
28   Insert nexttp into next_set
29   Remove last entry of path.pos_val_pairs
30   Only keep the top beam_width elements with the
      largest score in next_set if its size exceeds
      beam_width.
31   set  $\leftarrow$  next_set

```

the rules Eqs. (A1)–(A4) with a critical modification: all messages to or from masked error nodes are ignored. For example, in Eq. (A1), the message $D_{i \rightarrow j}(t)$ is not calculated if e_j is masked; if e_j is not masked, the calculation of $D_{i \rightarrow j}(t)$ will exclude any incoming messages $E_{j \rightarrow i}(t-1)$ where e_j is masked.

Finally, we note a simplification in Algorithm 3 made for clarity and brevity. The pseudocode first gathers all new paths in `next_set` (Line 28) and then applies a single, collective pruning step (Line 30). In our actual implementation, we maintain the set’s bounded size “on-the-fly” to improve efficiency. A new path (`nextp`) is inserted into `next_set` only if one of two conditions is met: (1) `next_set` is not yet full (its size is less than `beam_width`), or (2) `nextp.score` is greater than the minimum score currently in `next_set`. In the second case, the path with the smallest score is ejected as `nextp` is inserted. This optimization is significant because it avoids the costly operation of copying `nextp.edge_msgs` (a long vector) for paths that would be immediately discarded by the pruning step.

-
- [1] R. Gallager, *Low-density parity-check codes*, *IRE Trans. Inf. Theory* **8**, 21 (1962).
 - [2] D. J. MacKay, *Information Theory, Inference and Learning Algorithms* (Cambridge University Press, Cambridge, England, 2003).
 - [3] T. Richardson and R. Urbanke, *Modern Coding Theory* (Cambridge University Press, Cambridge, England, 2008).
 - [4] J. Pearl, *Reverend bayes on inference engines: A distributed hierarchical approach*, in *Proceedings of the Second AAAI Conference on Artificial Intelligence*, AAAI’82 (AAAI Press, 1982), pp. 133–136.
 - [5] D. J. C. MacKay, G. Mitchison, and P. L. McFadden, *Sparse-graph codes for quantum error correction*, *IEEE Trans. Inf. Theory* **50**, 2315 (2004).
 - [6] J.-P. Tillich and G. Zémor, *Quantum LDPC codes with positive rate and minimum distance proportional to the square root of the blocklength*, *IEEE Trans. Inf. Theory* **60**, 1193 (2013).
 - [7] A. A. Kovalev and L. P. Pryadko, *Quantum Kronecker sum-product low-density parity-check codes with finite rate*, *Phys. Rev. A* **88**, 012311 (2013).
 - [8] N. P. Breuckmann and J. N. Eberhardt, *Balanced product quantum codes*, *IEEE Trans. Inf. Theory* **67**, 6653 (2021).
 - [9] P. Panteleev and G. Kalachev, *Asymptotically good quantum and locally testable classical LDPC codes*, in *Proceedings of the 54th annual ACM SIGACT symposium on theory of computing* (2022), pp. 375–388.
 - [10] A. Leverrier and G. Zémor, *Quantum tanner codes*, in *Proceedings of the 2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, 2022), pp. 872–883.
 - [11] O. Higgott and N. P. Breuckmann, *Constructions and performance of hyperbolic and semi-hyperbolic floquet codes*, *PRX Quantum* **5**, 040327 (2024).
 - [12] S. Bravyi, A. W. Cross, J. M. Gambetta, D. Maslov, P. Rall, and T. J. Yoder, *High-threshold and low-overhead fault-tolerant quantum memory*, *Nature (London)* **627**, 778 (2024).
 - [13] M. Ye and N. Delfosse, *Quantum error correction for long chains of trapped ions*, *Quantum* **9**, 1920 (2025).
 - [14] T. R. Scruby, T. Hillmann, and J. Roffe, *High-threshold, low-overhead and single-shot decodable fault-tolerant quantum memory*, *PRX Quantum* **7**, 020310 (2026).
 - [15] M. A. Tremblay, N. Delfosse, and M. E. Beverland, *Constant-overhead quantum error correction with thin planar connectivity*, *Phys. Rev. Lett.* **129**, 050504 (2022).
 - [16] A. Aydin, N. Delfosse, and E. Tham, *Cyclic hypergraph product code*, [arXiv:2511.09683](https://arxiv.org/abs/2511.09683).
 - [17] D. Gottesman, *Stabilizer codes and quantum error correction*, Ph.D. thesis, California Institute of Technology, 1997.
 - [18] R. Tanner, *A recursive approach to low complexity codes*, *IEEE Trans. Inf. Theory* **27**, 533 (1981).
 - [19] X.-Y. Hu, E. Eleftheriou, and D.-M. Arnold, *Progressive edge-growth tanner graphs*, in *GLOBECOM’01, in Proceedings of the IEEE Global Telecommunications Conference (Cat. No. 01CH37270)* (IEEE, New York, 2001), Vol. 2, pp. 995–1001.
 - [20] P. Panteleev and G. Kalachev, *Degenerate quantum LDPC codes with good finite length performance*, *Quantum* **5**, 585 (2021).
 - [21] J. Roffe, D. R. White, S. Burton, and E. Campbell, *Decoding across the quantum low-density parity-check code landscape*, *Phys. Rev. Res.* **2**, 043423 (2020).
 - [22] J. Roffe, *LDPC: Python tools for low density parity check codes* (2022).
 - [23] N. Delfosse, V. Londe, and M. E. Beverland, *Toward a union-find decoder for quantum LDPC codes*, *IEEE Trans. Inf. Theory* **68**, 3187 (2022).
 - [24] S. Wolanski and B. Barber, *Introducing ambiguity clustering: An accurate and efficient decoder for QLDPC codes*, in *Proceedings of the 2024 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, New York, 2024), Vol. 2, pp. 402–403.
 - [25] T. Hillmann, L. Berent, A. O. Quintavalle, J. Eisert, R. Wille, and J. Roffe, *Localized statistics decoding for quantum low-density parity-check codes*, *Nat. Commun.* **16**, 8214 (2025).
 - [26] B. M. Terhal, *Quantum error correction for quantum memories*, *Rev. Mod. Phys.* **87**, 307 (2015).
 - [27] A. Khalid, A. Silva, G. A. Dagnew, T. Dvir, O. Wertheim, M. Gruda, X. Kong, M. Kramer, Z. Webb, A. Scherer *et al.*, *Impacts of decoder latency on utility-scale quantum computer architectures*, [arXiv:2511.10633](https://arxiv.org/abs/2511.10633).
 - [28] D. Poulin and Y. Chung, *On the iterative decoding of sparse quantum codes*, *Quantum Inf. Comput.* **8**, 987 (2008).
 - [29] J. Du Crest, M. Mhalla, and V. Savin, *Stabilizer inactivation for message-passing decoding of quantum LDPC codes*, in *Proceedings of the 2022 IEEE Information Theory Workshop (ITW)* (IEEE, New York, 2022), pp. 488–493.
 - [30] H. Yao, W. A. Laban, C. Häger, A. G. i Amat, and H. D. Pfister, *Belief propagation decoding of quantum LDPC codes with guided decimation*, in *Proceedings of the 2024 IEEE International Symposium on Information Theory (ISIT)* (IEEE, New York, 2024), pp. 2478–2483.

- [31] A. Gong, S. Cammerer, and J. M. Renes, *Toward low-latency iterative decoding of QLDPC codes under circuit-level noise*, [arXiv:2403.18901](#).
- [32] K. R. Ott, B. Hetényi, and M. E. Beverland, *Decision-tree decoders for general quantum LDPC codes*, [arXiv:2502.16408](#).
- [33] A. d. iOlius, I. E. Martinez, J. Roffe, and J. E. Martinez, *An almost-linear time decoding algorithm for quantum LDPC codes under circuit-level noise*, [arXiv:2409.01440](#).
- [34] T. Müller, T. Alexander, M. E. Beverland, M. Bühler, B. R. Johnson, T. Maurer, and D. Vandeth, *Improved belief propagation is sufficient for real-time decoding of quantum memory*, [arXiv:2506.01779](#).
- [35] T. Maurer, M. Bühler, M. Kröner, F. Haverkamp, T. Müller, D. Vandeth, and B. R. Johnson, *Real-time decoding of the gross code memory with FPGAs*, [arXiv:2510.21600](#).
- [36] S. Koutsoumpas, H. Sayginel, M. Webster, and D. E. Browne, *Automorphism ensemble decoding of quantum LDPC codes*, [arXiv:2503.01738](#).
- [37] M. Wang, A. Li, and F. Mueller, *Fully parallelized bp decoding for quantum LDPC codes can outperform bp-osa*, in *Proceedings of the 2026 IEEE International Symposium on High Performance Computer Architecture (HPCA)* (IEEE, New York, 2026), pp. 1–14.
- [38] Y. Wu, B. Li, K. Chang, S. Puri, and L. Zhong, *Minimum-weight parity factor decoder for quantum error correction*, [arXiv:2508.04969](#).
- [39] L. A. Beni, O. Higgott, and N. Shutty, *Tesseract: A search-based decoder for quantum error correction*, [arXiv:2503.10988](#).
- [40] L. Valentini, D. Forlivesi, A. Talarico, and M. Chiani, *Restart belief: A general quantum LDPC decoder*, *IEEE Commun. Lett.* **30**, 1185 (2026).
- [41] A. Hughes, R. Srinivas, C. Löschner, H. Knaack, R. Matt, C. Ballance, M. Malinowski, T. Harty, and R. Sutherland, *Trapped-ion two-qubit gates with > 99.99% fidelity without ground-state cooling*, [arXiv:2510.17286](#).
- [42] *Source code of beam search decoder* (2025). <https://github.com/ionq-publications/BeamSearchDecoder>.
- [43] P. Das, C. A. Pattison, S. Manne, D. M. Carmean, K. M. Svore, M. Qureshi, and N. Delfosse, *Afs: Accurate, fast, and scalable error-decoding for fault-tolerant quantum computers*, in *Proceedings of the 2022 IEEE International Symposium on High-Performance Computer Architecture (HPCA)* (IEEE, New York, 2022), pp. 259–273.
- [44] N. Liyanage, Y. Wu, A. Deters, and L. Zhong, *Scalable quantum error correction for surface codes using FPGA*, in *Proceedings of the 2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, Vol. IEEE, New York, 2023), Vol. 1, pp. 916–927.
- [45] N. Liyanage, Y. Wu, S. Tagare, and L. Zhong, *FPGA-based distributed union-find decoder for surface codes*, *IEEE Trans. Quantum Eng.* **5**, 1 (2024).
- [46] L. Caune, L. Skoric, N. S. Blunt, A. Ruban, J. McDaniel, J. A. Valery, A. D. Patterson, A. V. Gramolin, J. Majaniemi, K. M. Barnes *et al.*, *Demonstrating real-time and low-latency quantum error correction with superconducting qubits*, [arXiv:2410.05202](#).
- [47] B. Barber, K. M. Barnes, T. Bialas, O. Buğdaycı, E. T. Campbell, N. I. Gillespie, K. Johar, R. Rajan, A. W. Richardson, L. Skoric *et al.*, *A real-time, scalable, fast and resource-efficient decoder for a quantum computer*, *Nat. Electron.* **8**, 84 (2025).
- [48] A. B. Ziad, A. Zalawadiya, C. Topal, J. Camps, G. P. Gehér, M. P. Stafford, and M. L. Turner, *Local clustering decoder as a fast and adaptive hardware decoder for the surface code*, *Nat. Commun.* **16**, 11048 (2025).
- [49] Y. Wu, N. Liyanage, and L. Zhong, *Micro blossom: Accelerated minimum-weight perfect matching decoding for quantum error correction*, in *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems* (2025), Vol. 2, pp. 639–654.
- [50] F. Valentino, B. Branchini, D. Conficconi, D. Sciuto, and M. D. Santambrogio, *QUEKUF: An FPGA union find decoder for quantum error correction on the toric code*, *ACM Trans. Reconfigurable Technol. Syst.* **18**, 1 (2025).
- [51] S. Maurya, T. Maurer, M. Bühler, D. Vandeth, and M. E. Beverland, *FPGA-tailored algorithms for real-time decoding of quantum LDPC codes*, [arXiv:2511.21660](#).
- [52] H. Nishimori, *Statistical Physics of Spin Glasses and Information Processing: An Introduction* (Clarendon Press, 2001), Vol. 111.
- [53] M. Mezard and A. Montanari, *Information, Physics, and Computation* (Oxford University Press, 2009).
- [54] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, *Topological quantum memory*, *J. Math. Phys. (N.Y.)* **43**, 4452 (2002).
- [55] N. Delfosse and G. Zémor, *Quantum erasure-correcting codes and percolation on regular tilings of the hyperbolic plane*, in *2010 IEEE Information Theory Workshop* (IEEE, New York, 2010), pp. 1–5.
- [56] H. Bombin, R. S. Andrist, M. Ohzeki, H. G. Katzgraber, and M. A. Martin-Delgado, *Strong resilience of topological codes to depolarization*, *Phys. Rev. X* **2**, 021004 (2012).
- [57] A. A. Kovalev and L. P. Pryadko, *Spin glass reflection of the decoding transition for quantum error correcting codes*, *Quantum Inf. Comput.* **15**, 825 (2015).
- [58] A. Kubica, M. E. Beverland, F. Brandão, J. Preskill, and K. M. Svore, *Three-dimensional color code thresholds via statistical-mechanical mapping*, *Phys. Rev. Lett.* **120**, 180501 (2018).
- [59] Y. Li and M. P. A. Fisher, *Statistical mechanics of quantum error correcting codes*, *Phys. Rev. B* **103**, 104306 (2021).
- [60] D. Vodola, M. Rispler, S. Kim, and M. Müller, *Fundamental thresholds of realistic quantum error correction circuits from classical spin models*, *Quantum* **6**, 618 (2022).
- [61] B. Placke, T. Rakovszky, N. P. Breuckmann, and V. Khemani, *Topological quantum spin glass order and its realization in QLDPC codes*, [arXiv:2412.13248](#).
- [62] B. Placke, G. M. Sommers, N. P. Breuckmann, T. Rakovszky, and V. Khemani, *Expansion creates spin-glass order in finite-connectivity models: A rigorous and intuitive approach from the theory of LDPC codes*, [arXiv:2507.13342](#).
- [63] L. H. English, D. J. Williamson, and S. D. Bartlett, *Thresholds for postselected quantum error correction from statistical mechanics*, *Phys. Rev. Lett.* **135**, 120603 (2025).
- [64] C. Gidney, *Stim: A fast stabilizer circuit simulator*, *Quantum* **5**, 497 (2021).