

Received 23 August 2024; revised 15 May 2025; accepted 31 May 2025; date of publication 6 June 2025;
date of current version 23 July 2025.

Digital Object Identifier 10.1109/TQE.2025.3577769

Improved Belief Propagation Decoding Algorithms for Surface Codes

JIAHAN CHEN¹, ZHENGZHONG YI¹, ZHIPENG LIANG,
AND XUAN WANG¹ (Senior Member, IEEE)

Harbin Institute of Technology, Shenzhen 518055, China

Corresponding authors: Zhengzhong Yi; Xuan Wang (e-mail: zhengzhongyi@cs.hitsz.edu.cn; wangxuan@cs.hitsz.edu.cn).

This work was supported in part by the Guangdong Provincial Key Laboratory of Novel Security Intelligence Technologies under Grant 2022B1212010005, in part by the Shenzhen Science and Technology Program, China, under Grant JCYJ20241202123906009, in part by the Colleges and Universities Stable Support Project of Shenzhen, China, under Grant GXWD20220817164856008, and in part by the Harbin Institute of Technology, Shenzhen—SpinQ quantum information Joint Research Center Project under Grant HITSZ20230111.

ABSTRACT Quantum error correction is crucial for universal fault-tolerant quantum computing. Highly accurate and low-time-complexity decoding algorithms play an indispensable role in ensuring quantum error correction works effectively. Among existing decoding algorithms, belief propagation (BP) is notable for its nearly linear time complexity and general applicability to stabilizer codes. However, BP's decoding accuracy without postprocessing is unsatisfactory in most situations. This article focuses on improving the decoding accuracy of BP over GF(4) for surface codes. Inspired by machine learning optimization techniques, we first propose Momentum-BP and AdaGrad-BP to reduce oscillations in message updating, breaking the trapping sets of surface codes. We further propose exponential weighted average initialization belief propagation (EWAInit-BP), which adaptively updates initial probabilities and provides a one to three orders of magnitude improvement over traditional BP for planar surface code, toric code, and XZZX surface code without any postprocessing method, showing high decoding accuracy even under parallel scheduling. The theoretical $O(1)$ time complexity under parallel implementation and high accuracy of EWAInit-BP make it a promising candidate for high-precision real-time decoders.

INDEX TERMS Machine learning, quantum error correction (QEC), quantum information.

I. INTRODUCTION

Information in quantum computing devices is highly susceptible to noise, leading to information loss and computational errors [1]. To address this issue, quantum error correction (QEC) introduces redundant physical qubits and encodes the information in a certain subspace of the state space of all the physical qubits to protect it. So far at least, QEC is an essential step toward universal fault-tolerant quantum computation [2], [3], [4]. To ensure that QEC works, the decoder of a QEC code should output an accurate and fast estimation of error according to the error syndromes [5].

Several decoding algorithms have achieved notable success, including minimum-weight perfect matching (MWPM) [6] algorithm, union-find (UF) algorithm [7], maximum likelihood (ML) algorithm [8], matrix-product-states-based (MPS-based) algorithm [8], renormalization group algorithm [9], and neural-network-based (NN-based) algorithm [10], [11], [12], [13], [14], [15], [16], [17]. Among these, ML, MWPM, and MPS-based algorithms offer high

accuracy but come with high complexity. In contrast, the UF algorithm provides near-linear complexity with slightly lower accuracy. NN-based decoders have to balance scalability, decoding accuracy, and time cost [14] and may struggle with generalizability across different codes. Consequently, existing decoding algorithms rarely meet the requirements of both high accuracy and low complexity in QEC decoding [18], [19].

The belief propagation (BP) algorithm [20] is applicable to any stabilizer code and can achieve a complexity of $O(j\tau N)$ [21], where j is the average weight of the rows and columns of the check matrix, and τ is the number of iterations. In addition, parallel scheduling of BP can be efficiently implemented in hardware, theoretically achieving constant-time complexity [22]. However, traditional BP exhibits low decoding accuracy for many quantum error correcting codes and even cannot achieve the code capacity threshold for surface codes, which are the focus of recent advancements in QEC implementation [3], [23].

BP can be implemented over both GF(4) [21], [24], [25], [26], [27], [28], [29] and GF(2) [30], [31], [32], [33], where the Pauli operators are represented. There have been some efforts improving decoding accuracy of BP over GF(4) and GF(2). Over GF(4), memory effect belief propagation (MBP) [21] achieves a threshold of 14.5–16% for the planar surface code by modifying the step size in the a posteriori probability updates. However, MBP is limited to serial scheduling, and the outer loop for searching step sizes increases its time complexity. Over GF(2), generalized belief propagation [33] achieves a 14% threshold by combining generalized BP (GBP) with multiple rounds of decoding. This algorithm can be implemented in parallel, but the number of outer loops increases with the code length, adding to its time complexity. Belief propagation with bias using oscillating trapping sets (BP-OTS) [34] periodically detects and breaks trapping sets using manually set parameters, allowing parallel implementation while maintaining the original complexity. It outperforms MBP for short codes but does not achieve a significant threshold.

Postprocessing methods, such as ordered statistics decoding (OSD), stabilizer inactivation [35], MWPM [36], [37], [38], and UN [3], are often used to improve BP's decoding accuracy. However, they also introduce extra time cost, and all have high time complexities, ranging from $O(N^2)$ to $O(N^3)$. Improving BP's decoding accuracy to get rid of its reliance on postprocessing algorithms remains a significant challenge.

In this article, we mainly focus on improving BP's decoding accuracy for surface codes, while we also believe that the ideas used in these improvements will also work for some other quantum low-density parity-check (LDPC) codes. There are several explanations for why surface codes are challenging for BP decoding. From the perspective of degeneracy [21], [39], stabilizers of surface codes have much lower weights compared to the code distance, leading to multiple low-weight errors corresponding to the same syndrome, thus having negative influence on BP's error convergence. From the perspective of short cycles [20], a number of four cycles over GF(4) and eight cycles over GF(2) in surface codes can cause BP's inaccurate updating. A more intuitive explanation involves trapping sets [34], [40], where surface codes have many symmetric trapping sets causing BP messages to oscillate periodically.

This article approaches BP improvements from two perspectives: optimizing message updates inspired by gradient optimization, and optimizing a priori probabilities using information of previous iterations, leading to the proposal of three algorithms.

First, inspired by the similarity between BP message update and gradient descent of the energy function, we leverage optimization methods in machine learning to optimize the message update rules, and propose Momentum-BP and AdaGrad-BP. In addition to improving decoding accuracy, experiments on trapping sets in surface codes show that Momentum-BP and AdaGrad-BP can

smooth the message update process and reduce oscillation amplitude.

Furthermore, realizing the importance of a priori probabilities involved in the iteration of BP, we propose exponential weighted average initialization belief propagation (EWAInit-BP), which dynamically updates the a priori probabilities involved in each iteration by utilizing the a posteriori probabilities output by previous iterations. Experiments on short toric codes, planar surface codes, and XZZX surface codes (under biased noise) show that EWAInit-BP outperforms existing BP improvements without postprocessing and outer loops under parallel scheduling. Notably, the time complexity of all proposed improvements remains $O(j\tau N)$. Thus, EWAInit-BP can be implemented with theoretically constant complexity in parallel while maintaining decoding accuracy.

The rest of this article is organized as follows. Section II introduces the fundamentals of surface codes and BP decoding, as well as the perspective of BP as a gradient optimization, and summarizes the challenges BP faces in decoding surface codes. Section III introduces Momentum-BP, AdaGrad-BP, and EWAInit-BP. The results of the simulation are presented in Section IV. Finally, Section V concludes this article and discusses the future challenges.

II. PRELIMINARIES

A. SURFACE CODES

A stabilizer code is defined by an abelian subgroup \mathcal{S} of the Pauli group \mathcal{P}_n excluding $-I$ and composed of operators $S \in \mathcal{P}_n$ such that $S|\psi\rangle = |\psi\rangle$ for all codewords $|\psi\rangle$ in the code space. The code space is the simultaneous $+1$ eigenspace of all elements of \mathcal{S} . The stabilizers must commute with each other, which means $[S_i, S_j] = 0 \forall S_i, S_j \in \mathcal{S}$. The dimension of the code space is determined by the number of independent generators of \mathcal{S} . For a stabilizer code with N physical qubits and K logical qubits, \mathcal{S} has $N - K$ independent generators.

The stabilizers are used to detect errors by measuring their eigenvalues. The set of measurement outcomes, known as the error syndrome, contains information about the errors on physical qubits. The logical operators of a stabilizer code are Pauli operators that commute with all elements of \mathcal{S} but are not in \mathcal{S} ; thus, they are in $\mathcal{N}(\mathcal{S}) \setminus \mathcal{S}$, where $\mathcal{N}(\mathcal{S}) = \{E : ES = SE, \forall S \in \mathcal{S}\}$. They act on the encoded logical qubits without disturbing the code space; thus, they can lead to undetectable errors. The code distance D is defined as the minimum weight of logical operators. A quantum stabilizer code can be represented by $[[N, K, D]]$.

Surface codes may be the most notable stabilizer codes in recent years. Surface codes are a class of quantum LDPC codes [41], where qubits are placed on a 2-D lattice, relying on the topology of the surface to protect quantum information [42]. Logical qubits are encoded in global degrees of freedom, and larger lattices provide larger code distances. The code distance scales with the code length as $O(\sqrt{N})$,




FIGURE 1. Lattice of the $[[2L^2 - 2L + 1, 1, L]]$ planar surface code with $L = 3$. Circles represent data qubits, yellow squares represent X-type stabilizers, green squares represent Z-type stabilizers, and arrows indicate the coupling between stabilizers and qubits.

whereas the code rate is asymptotically zero. In this article, toric codes refer to those with closed boundaries, and planar surface codes refer to those with open boundaries. Besides, we consider the unrotated surface codes, but it is clear that our proposed algorithms can be extended to rotated versions.

A $[[2L^2, 2, L]]$ toric code has a 2-D torus topology, and the closed boundaries makes it translationally invariant to the error syndrome, which is beneficial to decoding algorithms, such as BP and neural networks. The Euler characteristic of the torus is 0, but the X-type and Z-type stabilizers of the toric code both have a redundancy, resulting in a total of two encoded logical qubits.

A $[[2L^2 - 2L + 1, 1, L]]$ planar surface code is defined on a surface with open boundaries, which makes the stabilizers near the boundaries having lower weights. This reduction in stabilizer weight results in a slight decrease in the error correction capability compared to the toric code; however, it simplifies implementation in physical quantum systems. The lattice structure of the planar surface codes is illustrated in Fig. 1.

An $[[L^2, 2, L]]$ XZZX surface code is a non-CSS code designed to handle biased noise more effectively. The stabilizers of these codes have the form of XZZX. This alternating pattern of Pauli X and Z operators allows the code to reach the theoretical maximum threshold of 50% under pure Pauli X, pure Pauli Y, and pure Pauli Z noise.

B. BP DECODING

1) BP OVER GF(4)

The idea behind *ML decoding* in QEC is to find the most likely error coset given the error syndrome. This can be expressed as [24]

$$E = \arg \max_{ES} \sum_{S \in \mathcal{S}} p(ES|s) \quad (1)$$

where the term ES denotes the coset formed by the error E and the stabilizer S . However, finding this ML solution is computationally expensive and exponential [43]. The BP decoding algorithm simplifies the aforementioned objective to

$$E = \left\{ \arg \max_{E_j} p(E_j|s) \right\}_{j=1}^n \quad (2)$$

where the probability of each individual error E_j given the syndrome s is maximized independently for $j = 1$ to n . BP makes two simplifications to the objective, which compromise its performance on highly degenerate quantum codes:

- 1) assuming that the most likely error pattern is the same as the most likely error coset;
- 2) assuming that the most likely error on each qubit is the same as the most likely error pattern.

The BP algorithm passes messages on a Tanner graph, where stabilizers and qubits are represented as check nodes and variable nodes, respectively. Over GF(4) which this article mainly focuses on, each variable node corresponds to a probability vector of all Pauli errors on one qubit.

Besides, one can use log-likelihood ratio (LLR) as the message passed between variable nodes and check nodes. In this way, belief propagation (abbreviated as LLR-BP₄) receives a priori probability vectors describing each type of Pauli error on each qubit as $\{p_{v_j} = \{p_{v_j}^X, p_{v_j}^Y, p_{v_j}^Z\}\}_{j=1}^N$ and initializes the variable-to-check messages as [21]

$$\Pi_{v_j} = \{\Pi_{v_j, W}\}_{W \in \{X, Y, Z\}} = \ln \frac{1 - p_{v_j}}{p_{v_j}} \quad (3)$$

$$m_{v_j \rightarrow c_i}^{(0)} = \lambda_{H_{ij}}(\Pi_{v_j, X}, \Pi_{v_j, Y}, \Pi_{v_j, Z}) \quad (4)$$

where Π_{v_j} is the a priori probabilities of variable node j expressed by LLR. The lambda function in (4) is defined as

$$\lambda_{H_{ij}}(p^X, p^Y, p^Z) = \ln \frac{1 + e^{-p^{H_{ij}}}}{e^{-p^X} + e^{-p^Y} + e^{-p^Z} - e^{-p^{H_{ij}}}} \quad (5)$$

indicating LLR values of whether the Pauli error E_j commutes with the Pauli operator H_{ij} of the parity-check matrix H . H_{ij} is also the j th element of stabilizer S_i .

One iteration of BP consists of one horizontal update, one vertical update, and one hard decision. In the horizontal update, each check node calculates the feedback based on the messages received from the variable nodes but not from the target variable node in last vertical update, sending a check-to-variable message as follows:

$$w = 2 \tanh^{-1} \left(\prod_{v_{j'} \in \mathcal{N}(c_i) \setminus v_j} \tanh \left(\frac{m_{v_{j'} \rightarrow c_i}^{(t-1)}}{2} \right) \right) \quad (6)$$

$$m_{c_i \rightarrow v_j}^{(t)} = (-1)^{s_i} * w \quad (7)$$

where $\mathcal{N}(c_i)$ represents all variable nodes adjacent to check node i , and s_i is the syndrome corresponding to stabilizer S_i .

In the vertical update, variable nodes aggregate the messages from check nodes and, after a similar extrinsic calculation as above, send variable-to-check messages as

$$m_{v_j \rightarrow c_i}^{(t)} = \lambda_{H_{ij}} (\Pi_{v_j, W} + \sum_{\substack{c_{i'} \in \mathcal{M}(v_j) \setminus c_i \\ \langle W, H_{ij} \rangle = 1}} m_{c_{i'} \rightarrow v_j}^{(t)})_{W \in \{X, Y, Z\}} \quad (8)$$

where $\mathcal{M}(v_j)$ represents all check nodes adjacent to variable node j . $\langle W, H_{ij} \rangle$ stands for the trace inner product; here, the result is 1 when the error operator W anticommutes with the matrix element H_{ij} and 0 when they commute.

In the hard decision, variable nodes first accumulate the messages from all check nodes to calculate the a posteriori probability of each Pauli error on each qubit as

$$Q_{v_j, W}^{(t)} = \Pi_{v_j, W} + \sum_{\substack{c_i \in \mathcal{M}(v_j) \\ \langle W, H_{ij} \rangle = 1}} m_{c_i \rightarrow v_j}^{(t)}. \quad (9)$$

Then, for each variable node, if all LLRs for Pauli errors are greater than 0, then the hard decision result is the identity I ; otherwise, it is the Pauli error with the smallest LLR value. The hard decisions of all qubits combine to form the error estimate \hat{E} . If \hat{E} corrects the error syndrome, the algorithm converges; otherwise, it returns to the horizontal update for a new iteration. If the maximum number of iterations is reached, decoding has failed.

BP can be performed under either parallel or serial scheduling [44]. Under parallel scheduling, all variable nodes and check nodes update their messages simultaneously in each iteration, allowing for hardware parallelization and constant time complexity. On the other hand, the serial scheduling updates the messages of variable nodes and check nodes sequentially. This method leads to faster convergence and slightly higher accuracy, though it has a more complex hardware implementation and higher computational overhead.

One perspective on BP is that each a posteriori probabilities update in one BP iteration can be viewed as a single step of gradient descent with a simplified step size on a certain energy function, as the expanded form of the a posteriori update closely resembles the derivative of the energy function (see Appendix A). From this perspective, MBP and adaptive MBP (AMBP) [21] adjust the “learning rate” of message updates to converge more quickly to the equivalent energy minimum, which is common in degenerate quantum codes.

2) HARDNESS OF BP IN DECODING SURFACE CODES

There are several terms that can explain the unsatisfactory accuracy of BP in surface codes.

a) Cycles: Short cycles cause the joint probability calculations in BP to be dependent, which violates the independence assumption of updating variable nodes (8) and check nodes (6) [20]. Surface codes exhibit eight cycles over GF(2) and four cycles over GF(4), exacerbating this issue.

b) Degeneracy: High degeneracy leads to BP decoding objectives (2) deviating from true targets (1) [24], [39], especially when errors can combine with stabilizers to form error types that are symmetric to the original ones. In such cases, BP’s marginal probability calculations for the both error types will be completely consistent. In surface codes, the weights of the stabilizers are all no greater than 4, often much less than the weights of the logical operators. This symmetric degeneracy results in incorrect convergence of BP decoding.

c) Trapping sets: Trapping sets are specific local structures in error-correcting codes combined with corresponding error patterns, where iterative decoding algorithms can never converge [45], [46], [47]. In QEC, trapping sets are caused by the aforementioned short loops and symmetric degeneracy [40], [48]. The smallest trapping set in surface codes is shown in Fig. 2, consisting of eight cycles and four symmetric stabilizers, where the local connections of each check node and variable node are identical [34]. If any two qubits here experience errors that anticommute with the stabilizers, the hard decision result of BP will oscillate periodically between all 0’s and all 1’s.

The first two terms are directly related to the underlying principles of the BP algorithm. In contrast, trapping sets offer a more specific description, which can not only inspire potential improvements but also serve as a metric for evaluating improvements.

III. BP OPTIMIZATION

From the perspective of BP as gradient descent [21], the high degeneracy of surface codes results in many equivalent global minima, leading to BP oscillating between adjacent erroneous solutions within symmetric trapping sets. A natural idea is to explore if certain empirical methods can be employed to smooth BP’s message updates, stabilizing the a posteriori probabilities at a specific local minimum before oscillations occur, thus enabling successful convergence to degenerate errors.

A. GRADIENT OPTIMIZATION OF MESSAGE UPDATES

This section follows the gradient perspective, inspired by optimizers in machine learning, to smooth BP iterations by modifying its message update rules.

1) UPDATE GRADIENT SMOOTHING

The exponential weighted average (EWA) method updates and smooths data series by combining the current observation with a weighted average of the previous value, expressed as

$$v_t = \gamma v_{t-1} + (1 - \gamma)x_t \quad (10)$$

where x_t is the observation at time t , v_t is the smoothed value, and γ is a decay factor satisfying $0 < \gamma < 1$.

In machine learning, momentum methods utilize an idea similar to EWA, adding a momentum term to accumulate




FIGURE 2. Trapping sets with two types of error patterns in surface codes. Solid circles represent error qubits, and solid squares represent stabilizers with nontrivial error syndromes. The stabilizers in the figure can be either all X-type or all Z-type.

historical information in gradient descent as

$$m_t = \gamma m_{t-1} + (1 - \gamma) \nabla J(\theta_{t-1}) \quad (11)$$

$$\theta_t = \theta_{t-1} - \alpha m_t \quad (12)$$

where γ is the decay factor similar to that in EWA, m_t is the accumulated momentum term, $\nabla J(\theta)$ is the gradient of the loss function with respect to parameter θ , and α is the learning rate. This method integrates previous directions while updating parameters, thereby smoothing the update direction.

We incorporate the concept of momentum into BP's a posteriori probability updates to solve this issue by breaking the symmetry in BP's message updates, proposing Momentum-BP. Momentum-BP can be implemented under both parallel and serial scheduling. Here, we present the latter in Algorithm 1, where the initialization and check node update remain the same as in traditional BP. Our approach computes the a posteriori probability before performs the vertical update. The posterior update rule of Momentum-BP is shown as (13) and (14), where $g_{v_j, W}^{(t)}$ are the accumulated momentum term and $Q_{v_j, W}^{(t)}$ are the a posteriori probability for the j th variable node regarding the Pauli error W in the t th iteration. Notice that we transformed (9), so that the accumulated term is the posterior value from the previous iteration. In the variable node update step, we subtract the original message sent from the check node rather than the momentum-smoothed message to ensure stability and prevent overcompensation.

For implementation of the gradient calculation, we ignore the coefficient terms in (29), using the difference between the posterior updates calculated by traditional BP and the a posteriori probabilities from the previous iteration as the gradient for the current iteration, expanded in (16), shown at the bottom of this page. When $\alpha = 1$ and $\gamma = 0$, (16)

$$Q_{v_j, W}^{(t)} = Q_{v_j, W}^{(t-1)} - \alpha \left(\gamma g_{v_j, W}^{(t-1)} + (1 - \gamma) \left(Q_{v_j, W}^{(t-1)} - \Pi_{v_j, W} - \sum_{\substack{c_i \in \mathcal{M}(v_j) \\ (W, H_{ij})=1}} m_{c_i \rightarrow v_j}^{(t)} \right) \right) \quad (16)$$

Algorithm 1: Momentum-BP Under Serial Scheduling

Input: Parity-check matrix H of $M \times N$, syndrome vector $s = (1, -1)^M$, a priori LLR vectors $\{\Pi_{v_j}^{(0)}\}_{j=1}^N$, max iterations $iter_{\max}$, learning rate $\alpha = 1$, and momentum rate $\gamma \geq 0$.

Output: Estimated error vector \hat{E} .

Initialization:

$t = 1$

for $j = 1 \rightarrow N$, **do**

Initialize $m_{v_j \rightarrow c_i}^{(0)}$ as (4), let $g_{v_j}^{(0)} = 0$, $Q_{v_j}^{(0)} = \Pi_{v_j}$

Serial update:

for $j = 1 \rightarrow N$, $W \in \{X, Y, Z\}$, **do**

for $i \in \mathcal{N}(v_j)$, **do**

Calculate $m_{c_i \rightarrow v_j}^{(t)}$ as (6)

Momentum update

$$g_{v_j, W}^{(t)} = \gamma g_{v_j, W}^{(t-1)} + (1 - \gamma) \frac{\partial J}{\partial Q_{v_j, W}^{(t-1)}} \quad (13)$$

$$Q_{v_j, W}^{(t)} = Q_{v_j, W}^{(t-1)} - \alpha g_{v_j, W}^{(t)} \quad (14)$$

Variable node update

$$m_{v_j \rightarrow c_i}^{(t)} = Q_{v_j, W}^{(t)} - \langle W, H_{ij} \rangle m_{c_i \rightarrow v_j}^{(t)} \quad (15)$$

Hard decision:

for $j = 1 \rightarrow N$, $W \in \{X, Y, Z\}$, **do**

if $Q_{v_j, W}^{(t)} > 0$ for all W , **let** $\hat{E}_i = I$

else, let $\hat{E}_i = \arg \min_{W \in \{X, Y, Z\}} Q_{v_j, W}^{(t)}$

if $(H, \hat{E}) = s$, **return** “converge” and halt.

else if $t = iter_{\max}$, **return** “fail” and halt.

else $t = t + 1$, repeat from **serial update**.

is equivalent to the traditional BP's a posteriori update; when $0 < \alpha < 1$ and $\gamma = 0$, it is equivalent to the EWA of a posteriori updates. In practice, to avoid the malice of hyperparameter tuning, we either fix α at 1 or γ at 0 (see (17) and (18), both shown at the bottom of the next page).

The properties of Momentum-BP can be explained from two perspectives.

- 1) Since it takes into account message updates from previous iterations, the method smooths the update direction in each iteration, preventing the corresponding energy function from oscillating between two points and gradually converging to the global minimum.

- 2) Nodewise momentum updates gradually break the symmetry of message passing shown in Fig. 2, allowing BP to escape from these trapping sets.

The trapping set simulation in Section IV-A will support those explanations.

2) ADAPTIVE STEP SIZE

Oscillation of message updates in trapping sets can also be attributed to the step size. In traditional LLR-BP, all variable node posterior updates result from simply directly accumulating messages from neighboring check nodes. From the gradient perspective, if the step size is too large, the optimization algorithm is likely to oscillate; if the step size is too small, the algorithm is likely unable to update effectively. In practical decoding, the update frequency of different variable nodes may be uneven, with some requiring more rapid updates, while others need to remain stable.

Inspired by adaptive gradient methods in deep learning, we propose AdaGrad-BP. The framework of this algorithm is similar to Momentum-BP, with initialization and horizontal updates identical to traditional BP, while the posterior probability updates follow an adaptive step size rule as

$$Q_{v_j, W}^{(t)} = Q_{v_j, W}^{(t-1)} - \alpha \frac{g'_{v_j, W}^{(t)}}{\sqrt{G_{v_j, W}^{(t)} + \epsilon}} \quad (19)$$

where $g'_{v_j, W}^{(t)}$ approximates the gradient of energy function

$$g'_{v_j, W}^{(t)} = Q_{v_j, W}^{(t-1)} - \Pi_{v_j, W} - \sum_{\substack{c_i \in \mathcal{M}(v_j) \\ \langle W, H_{ij} \rangle = 1}} m_{c_i \rightarrow v_j}^{(t)} \quad (20)$$

and $G_{v_j, W}^{(t)}$ is the cumulative sum of the squared gradients as

$$G_{v_j, W}^{(t)} = G_{v_j, W}^{(t-1)} + (g'_{v_j, W}^{(t)})^2. \quad (21)$$

This method adaptively adjusts the step size by accumulating historical gradient information. If the estimate of a variable node for a certain Pauli error changes significantly, the algorithm reduces the step size to counteract oscillation. Equation (19) takes effect after the first iteration to ensure initial update stability. Here, α is the boost learning rate, providing an update impetus as the algorithm begins to accumulate gradients. α is relatively insensitive, so it is fixed at 5 in our simulations. While AdaGrad-BP may face issues with

accumulated gradients causing the step size to approach zero after many iterations, in practice, it can reduce the number of iterations, allowing most error patterns to converge before the gradients become too small.

Unfortunately, combining Momentum-BP and AdaGrad-BP using Adam optimizer or other methods will not yield satisfactory results. On the one hand, the Adam algorithm introduces additional hyperparameters whose optimal values for BP cannot be directly borrowed from deep learning, thereby complicating the algorithm. On the other hand, Momentum and AdaGrad approach similar message update optimizations from different perspectives, and a straightforward combination of the two can lead to overcompensation.

B. DYNAMIC INITIALIZATION OF PRIOR PROBABILITIES

BP decoding and gradient descent still have a fundamental difference, which limits the performance of the algorithms in Section III-A. Each update in gradient descent is based on the results of the previous step, whereas in BP decoding, each calculation of a posteriori probabilities accumulates the same a priori probabilities. Given the objective function as shown in (2), for simplicity using the binary domain for instance, BP calculates the LLR values as

$$Q_{v_j} = \ln \frac{P(e_j = 0|s)}{P(e_j = 1|s)} \quad (22)$$

$$= \ln \frac{P(e_j = 0) \cdot P(s|e_j = 0)}{P(e_j = 1) \cdot P(s|e_j = 1)} \quad (23)$$

$$= \Pi_{v_j} + \sum_{c_i \in \mathcal{M}(v_j) \setminus c_i} m_{c_i \rightarrow v_j} \quad (24)$$

where Π_{v_j} is a manually set initial value of the algorithm that participates in each iteration's a posteriori probability update. However, since the initial error probabilities are derived from a theoretical channel, which is difficult to measure and estimate in experiments, the a priori probabilities can significantly affect the BP iterations.

Existing methods for altering initial probabilities include reinitialization using information obtained from previous iterations [24], [25], [33], as well as detecting oscillations during the decoding process and modifying initial values accordingly [34]. However, they either require multiple outer loops or manually set algorithm parameters. We introduce a simple but effective method for dynamically updating a priori probabilities, employing a transformed EWA, incorporating

$$\stackrel{\text{if } \alpha=1}{=} \gamma(Q_{v_j, W}^{(t-1)} - g_{v_j, W}^{(t-1)}) + (1 - \gamma) \left(\Pi_{v_j, W} + \sum_{\substack{c_i \in \mathcal{M}(v_j) \\ \langle W, H_{ij} \rangle = 1}} m_{c_i \rightarrow v_j}^{(t)} \right) \quad (17)$$

$$\stackrel{\text{if } \gamma=0}{=} \alpha \left(\Pi_{v_j, W} + \sum_{\substack{c_i \in \mathcal{M}(v_j) \\ \langle W, H_{ij} \rangle = 1}} m_{c_i \rightarrow v_j}^{(t)} \right) + (1 - \alpha)Q_{v_j, W}^{(t-1)} \quad (18)$$

Algorithm 2: EWAInit-BP Under Parallel Scheduling.

Input: Parity-check matrix H , syndrome vector s , a priori LLR vectors $\{\Pi_{v_j}^{(0)}\}_{j=1}^N$, max iterations $iter_{\max}$, discount factor α .

Output: Estimated error vector \hat{E} .

$t = 1$, $\Pi^{(1)} = \Pi^{(0)}$, $m_{v_j \rightarrow c_i}^{(0)} =$

$\lambda_{H_{ij}}(\Pi_{v_j, X}^{(0)}, \Pi_{v_j, Y}^{(0)}, \Pi_{v_j, Z}^{(0)})$

EWA Initialization when $t > 1$:

for $j = 1 \rightarrow N$, do

$$\Pi_{v_j}^{(t)} = \alpha \Pi_{v_j}^{(0)} + (1 - \alpha) Q_{v_j}^{(t-1)} \quad (25)$$

Horizontal update:

for $i = 1 \rightarrow M$, Calculate $m_{c_i \rightarrow v_j}^{(t)}$ as (6)

Vertical update:

for $j = 1 \rightarrow N$,

Calculate $m_{v_j \rightarrow c_i}^{(t)}$ as (8), replacing $\Pi_{v_j, W}$ with

$$\Pi_{v_j, W}^{(t)}$$

Hard decision:

for $j = 1 \rightarrow N$, $W \in \{X, Y, Z\}$, do

Calculate $Q_{v_j, W}^{(t)}$ as (9), replacing $\Pi_{v_j, W}$ with

$$\Pi_{v_j, W}^{(t)}$$

if $Q_{v_j, W}^{(t)} > 0$ for all e , let $\hat{E}_i = I$

else, let $\hat{E}_i = \arg \min_{W \in \{X, Y, Z\}} Q_{v_j, W}^{(t)}$

if $\langle H, \hat{E} \rangle = s$, return “converge” and halt.

else if $t = iter_{\max}$, return “fail” and halt.

else $t = t + 1$, repeat from EWA Initialization.

components of previous a posteriori probabilities into the a priori probabilities with a decaying factor. Hence, we call this method EWAInit-BP. EWAInit-BP under parallel scheduling is shown in Algorithm 2.

One might wonder why the first term on the right-hand side (RHS) of (25) is each qubit’s list of initial a priori probabilities $\Pi_{v_j}^{(0)}$ instead of a priori probabilities from the previous iteration $\Pi_{v_j}^{(t-1)}$, as the latter aligns better with the definition and intuition of the EWA. In fact, by using the definition in (25) rather than the aforementioned approach, we achieve the exponentially weighted average of horizontal message updates for the a posteriori probabilities in EWAInit-BP, as stated in Theorem 1.

Theorem 1: The a posteriori update of EWAInit-BP (see Algorithm 2) realizes an exponentially weighted average of past horizontal message updates, specifically expressed as

$$Q_{v_j, W}^{(t)} = \Pi_{v_j, W}^{(0)} + \sum_{k=0}^{t-1} (1 - \alpha)^k \sum_{\substack{c_i \in \mathcal{M}(v_j) \\ \langle W, H_{ij} \rangle = 1}} m_{c_i \rightarrow v_j}^{(t-k)}. \quad (26)$$

Proof: This result can be derived by combining (9) and (25) and recursively expanding the update formula. The detailed process is provided in Appendix B. \square

Notably, EWAInit-BP and Momentum-BP share a similar recursive form, but the former is more concise (see Appendix B). The primary distinction between them lies in the differing placements of EWA (or momentum). Moreover, by incorporating the EWA step into the initialization phase, EWAInit-BP avoids the issue of message coefficient inconsistency caused by performing posterior updates followed by vertical updates. In fact, EWAInit-BP also demonstrates better performance in simulations.

C. COMPLEXITY

Our proposed algorithms involve the same message passing between variable nodes and check nodes as in traditional BP. The main difference lies in the additional update steps, which in Momentum-BP and AdaGrad-BP consist of recalculating the a posteriori probabilities, and in EWAInit-BP involves recalculating the a priori probabilities. These additional steps only involve simple arithmetic operations and can be executed locally at each node, without requiring extra iterations, loops, or global communication, thus not introducing any significant overhead. Therefore, the overall time complexity of the proposed algorithms remains $O(j\tau N)$, where j is the average degree of the Tanner graph, and τ is the number of iterations.

In terms of space complexity, traditional BP requires storing messages and probability values, resulting in $O(N)$ space requirements. Our algorithms only add extra storage space to accumulate gradients or maintain EWAs, which is proportional to the number of variable nodes and does not change the overall space complexity. Thus, our improved algorithms maintain the same efficiency in both time and space as traditional BP, ensuring scalability and feasibility for practical implementation.

IV. SIMULATION RESULTS

In all simulations, the criterion for successful decoding is that the error estimate satisfies $\hat{E} \in ES$, rather than $\hat{E} = E$ as in classical information. All cases not meeting the aforementioned condition are referred to as logical errors in this article. Herein, a detected error refers to $\hat{E}E$ being anticommutative with any one stabilizer $S \in \mathcal{S}$; an undetected error implies that $\hat{E}E$ commutes with all stabilizers, i.e., $\hat{E}E \in N(\mathcal{S}) \setminus \mathcal{S}$, where the result of error correction is a logical operator. Each data point on the experimental curves represents the mean of at least 10^4 simulations.

Without further specification, our simulations employ a depolarizing channel, where each qubit has an equal probability $p/3$ of experiencing Pauli X , Y , or Z errors, and a probability of $1 - p$ of no error occurring. In addition, for ease of comparison, we adopt the simple noise model, which does not consider measurement errors or error propagation in quantum circuits.

A. BREAKING THE TRAPPING SETS

We first compare the performance of improved BP on the trapping set shown in Fig. 2. In this example, qubits 1 and

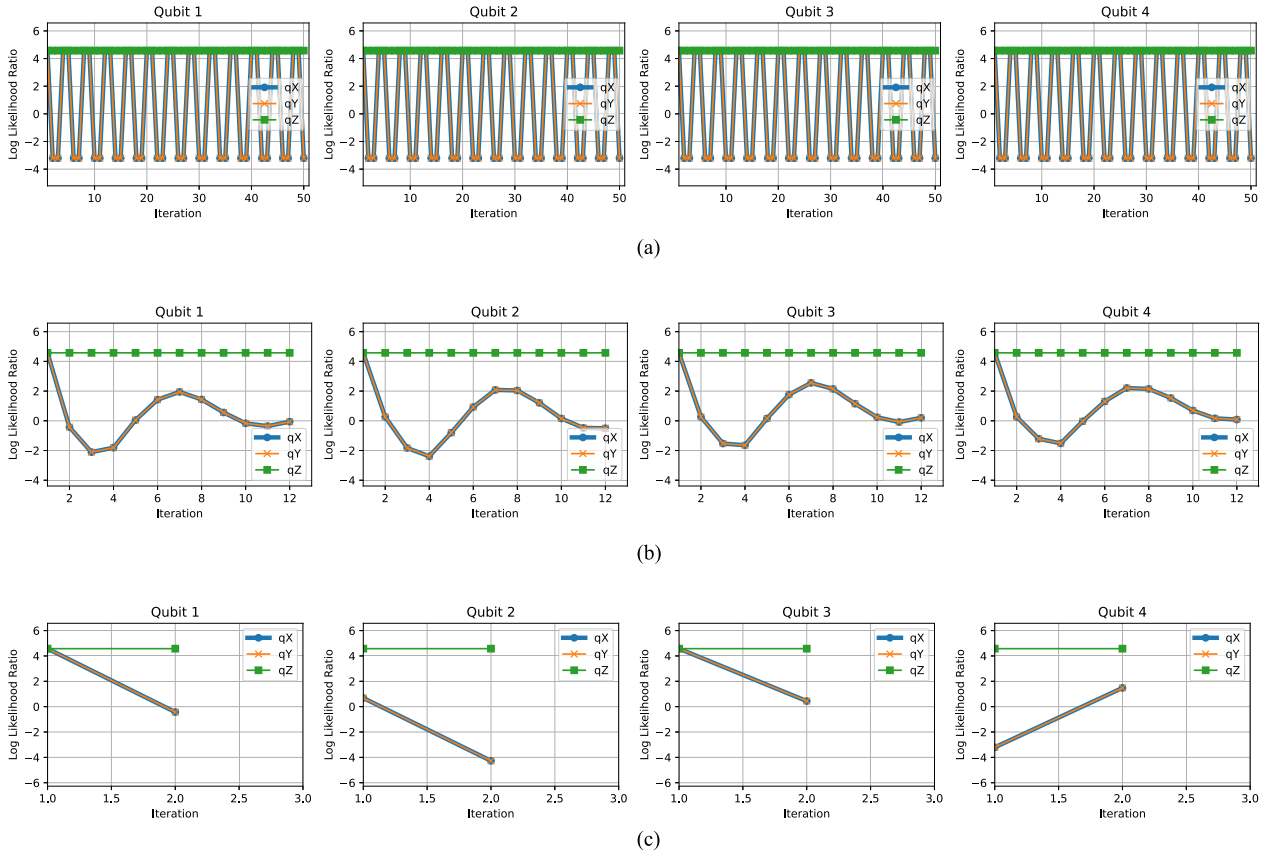


FIGURE 3. Comparison of traditional LLR-BP₄ and AdaGrad-BP under parallel scheduling and serial scheduling on the (4,0) trapping set. The improved BP algorithm break the symmetry of message updates and converge within 12 and 2 iterations separately. (a) Traditional LLR-BP₄. (b) AdaGrad-BP (parallel). (c) AdaGrad-BP (serial).

TABLE 1. Iterations Required by Different Algorithms on the (4,0) Trapping Set

Algorithm	Parallel Scheduling	Serial Scheduling
LLR-BP ₄ [31]	—	—
(A)MBP [21]	—	$\alpha = 0.5, iter = 2$
BP-OTS [34]	$iter = T + 2$	$iter = T + 2$
EWAInit-BP	—	$\alpha = 0.5, iter = 2$
AdaGrad-BP	$iter = 12$	$iter = 2$
Momentum-BP	$\alpha = 0.5, iter = 2$	$\alpha = 0.6, iter = 2$

2 experience X errors, resulting in an error syndrome of $[1,0,1,0]$. Fig. 3 shows the comparison between the traditional LLR-BP₄ and AdaGrad-BP in Section III-A under parallel and serial scheduling, with the maximum iteration set to 50 and the a priori probabilities set to 0.01. The hard decision result of the traditional BP oscillates continuously between all X and all I , whereas the improved BP algorithm shows significantly reduced oscillation amplitude and converges within few iterations.

A more detailed comparison of the different algorithms proposed in this article is presented in Table 1. As described in Section III-A, we set α of AdaGrad-BP to 5 and γ of Momentum-BP to 0. For Momentum-BP, EWAInit-BP,

and MBP, we tested ten evenly distributed values for alpha within the range $[0, 1]$ and selected the one that yielded the best performance. The same parameter selection strategy is also applied in Section IV-B. MBP (equals AMBP in this simulation) [21] and EWAInit-BP converges at the same fastest speed under serial scheduling but oscillates under parallel scheduling. BP-OTS [34] shows little dependence on scheduling method, but its convergence depends on the manually set oscillation period T , which may vary in different error correction scenarios. Algorithms proposed in Section III-A, Momentum-BP and AdaGrad-BP, converge fast under both parallel scheduling and serial scheduling.

It is worth noting that EWAInit-BP fails to decode trapping sets under parallel scheduling. This is because the recursive message update formula for each bit remains symmetric, making it unable to resolve the issue within trapping sets. In practical error-correcting code decoding, this limitation may lead to nonconvergence in very rare special cases. However, in most cases, messages originating from outside the trapping sets can assist in achieving convergence.

B. IMPROVED BP ON SURFACE CODES

In this section, we compare the performance of proposed BP improvements with other improvements for various surface




FIGURE 4. Overall performance of the proposed BP algorithms with existing algorithms on an $[[85, 1, 7]]$ planar surface code. Momentum-BP, AdaGrad-BP, MBP, and AMBP are evaluated under both parallel and serial scheduling. EWAInit-BP performs between serial MBP and serial AMBP and is almost identical to BP₂-OSD-0 in this case.

codes. For all BP algorithms, the maximum number of iterations $iter_{max}$ across all code lengths is set to 100 to ensure linear asymptotic complexity.

1) COMPARISON OF OUR BP OPTIMIZATION METHODS

A comparison of all methods proposed in this article with some existing high-accuracy BP improvements on an $L = 7$ planar surface code is illustrated in Fig. 4, using the same parameter selection strategy as described in Section IV-A. The two methods for optimizing message update rules, Momentum-BP and AdaGrad-BP, exhibit improvements over traditional LLR-BP₄ at physical error rates lower than 10%, providing evidence for the validity of the gradient perspective of BP. The prior-probability-based method, EWAInit-BP, demonstrates improvements ranging from one to three orders of magnitude over traditional BP and effectively eliminates the error floor. Compared to other algorithms, EWAInit-BP achieves performance between serial MBP and serial AMBP while using hardware-friendly parallel scheduling and, in this case, performs nearly identically to BP₂-OSD-0 [49], despite the latter having a time complexity of $O(n^3)$.

The enhancements observed with Momentum-BP and AdaGrad-BP support the validity of the gradient perspective, providing a foundation for further research in this direction. Next, we will provide detailed comparisons between our best enhancement, EWAInit-BP, and algorithms with similar time overheads (excluding postprocessing and external loops). Although AMBP [21] has an asymptotic complexity of $O(N)$, its actual runtime overhead is high, as detailed in Appendix C, which also includes the results for adaptive exponential weighted average (AEWA-BP).





FIGURE 5. Performance of LLRBP₄, MBP, and EWAInit-BP under parallel scheduling for toric codes with $L = \{4, 6, 8\}$. Only EWAInit-BP finds a reduction in the logical error rate as the lattice size increases, with the intersection of error rates occurring at 13.6%. (a) Logical error rates. (b) Average iterations.

2) TORIC CODE

We construct the $[[2n^2, 1, n]]$ toric code based on its properties as a hypergraph product of an $[n, 1]$ repetition code. The toric code features periodic boundaries; hence, all its stabilizers have a weight of four, and any error syndromes on it exhibit translational symmetry [17], which is particularly conducive to BP and neural network decoding. We compare the performance of our EWAInit-BP and other existing methods under both parallel and serial scheduling, as shown in Figs. 5 and 6, respectively. Note that in all simulations in this section, the parameter choices for different code lengths are consistent.

For the parallel schedule,¹ compared to the classic BP over GF(4) with LLR message and its enhancement, MBP, our EWAInit-BP is the only method that achieves a decrease in

¹Our implementation of BP-OTS over GF(4) under parallel scheduling did not achieve the performance reported in the original paper, necessitating further investigation into this discrepancy.


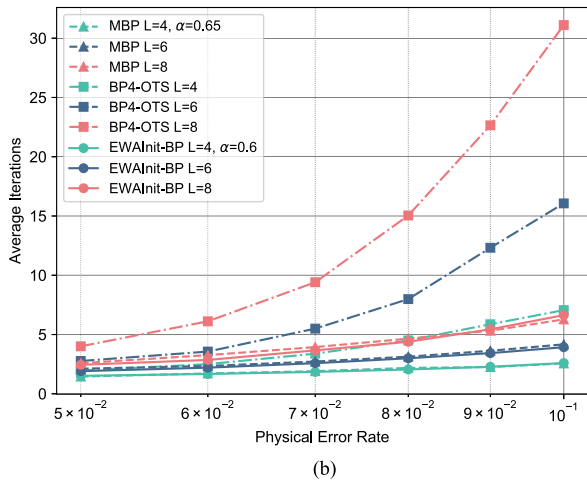




FIGURE 6. Performance of MBP, BP-OTS, and EWAInit-BP under serial scheduling for toric codes with $L = \{4, 6, 8\}$. We choose $T = 9$ and $C = 20$ for BP-OTS. EWAInit-BP outperforms other BP improvements, with the intersection of logical error rates at 15%. (a) Logical error rates. (b) Average iterations.

FIGURE 7. Performance of LLRBP4, MBP, and EWAInit-BP (EWAInit-OTS) for planar surface codes with $L = 5, 7, 9$. The intersection of logical error rates for EWAInit-BP under serial scheduling is approximately 10%. (a) LER under a parallel schedule. (b) LER under a serial schedule.

logical error rates as the lattice size increases, with significant improvements in logical error rates at the same lattice size. The intersection of the error rate curves for these lattice sizes is approximately 13.6%. In addition, the average number of iterations required for EWAInit-BP is significantly lower than the other two methods, converging in fewer than ten iterations mostly. Parallel scheduling of BP algorithms holds the potential to achieve constant complexity through hardware-based parallel computation, while EWAInit-BP ensures both the error correction capabilities and speed in decoding.

For the serial schedule, all three decoding algorithms demonstrated error correction capabilities. Among those, BP-OTS performs better than MBP at low physical error rates, but its performance is surpassed by MBP when the physical error rate exceeds 7%. The performance of EWAInit-BP, on the other hand, consistently outperforms these two methods across all tested conditions, with the logical error rate at $L = 6$ being comparable to that of the other

two algorithms at $L = 8$, and the intersection of the logical error rate curves is approximately 15%.² In terms of iteration count, both EWAInit-BP and MBP converge in remarkably few iterations (often fewer than 5).

3) PLANAR SURFACE CODE

The boundaries of the planar surface code are no longer contiguous, resulting in the emergence of some stabilizers with weight three and qubits that only connect two stabilizers at the boundaries. This configuration leads to a slight reduction in the error correction capabilities of the planar surface code compared to the Toric code. We compare the decoding logical error rates of LLR-BP₄, MBP, and EWAInit-BP under both parallel and serial scheduling, as illustrated in Fig. 7.

²We cannot yet consider the intersection of curves for short codes as the obtained threshold, as the convergence capability of EWAInit-BP degrades to some extent on larger codes. This will be discussed in Appendix C.




FIGURE 8. Performance of LLRBP4, MBP, and EWAInit-BP (EWAInit-OTS) for XZZX surface codes with $L = 5, 7, 9$ under serial scheduling. The intersection of logical error rates for EWAInit-BP is approximately 12.4%.

Under parallel scheduling, the decoding logical error rates of LLRBP4 and MBP4 increase with increasing code length, whereas EWAInit-BP achieves a reduction in logical error rates at low physical error rates. Unfortunately, at higher error rates, decoding at higher code lengths becomes more challenging, preventing EWAInit-BP from achieving a reliable threshold on the planar surface code.

Under serial scheduling, EWAInit-BP performs better than MBP at $L = 5$ and $L = 7$; however, its performance at $L = 9$ is not satisfying. Despite the dynamic updating of a priori probabilities, the method still partly relies on the symmetric initial input priors. One solution is to incorporate the OTS method as a “postprocessing” step for EWAInit, resetting the a priori probabilities of some nodes every few iterations, with the hyperparameters $T = 5$ and $C = 20$. This EWAInit-OTS method, while maintaining a time complexity of $O(n)$, shows improvements in convergence rate and logical error rates on the surface code compared to EWAInit-BP and other methods.

The patterns of average iterations of different algorithms for the planar surface code and the XZZX surface code are similar to those observed in the toric code and exhibit the same trends as the logical error rates. Therefore, we do not repeat the iteration figures here.

4) XZZX SURFACE CODE

For biased noise, the XZZX surface code exhibits superior error correction capabilities compared to the traditional planar surface code, achieving the theoretical maximum threshold of 50% across three types of pure Pauli error channels.

A comparison of decoding performance on the XZZX surface code is illustrated in Fig. 8. Our EWAInit combined with OTS method outperforms other algorithms under serial scheduling, with the logical error rate (LER) curve intersection at approximately 12.4%. Unfortunately, for odd lattice sizes, the logical error rate for higher code lengths is higher




FIGURE 9. Performance of LLRBP4 and EWAInit-BP for the XZZX surface code with $L = 5, 7, 9$ under parallel scheduling, evaluated under conditions where Pauli Z errors are more prevalent in the channel. The intersection of LERs for EWAInit-BP is approximately 10%.

than for lower code lengths at low physical error rates, while it demonstrates error-correcting capabilities at higher physical error rates.

To demonstrate the bias noise correction capability of the XZZX surface code, we chose $p_z = \frac{2}{3}$ and $p_x = p_y = \frac{1}{6}$, implying that Pauli Z errors are more likely to occur, which aligns with the conditions found in actual hardware. Fig. 9 displays the decoding performance under biased noises under parallel scheduling. It is evident that when biased noise is present, the performance of BP decoding using the XZZX surface code with parallel scheduling shows significant improvement compared to that on a depolarizing channel. Besides, EWAInit-BP offers an improvement of approximately one order of magnitude and demonstrates the error correction capability under parallel scheduling, with the intersection of logical error rates around 10%.

C. GENERALITY ON OTHER QUANTUM LDPC CODES

Although the trapping set analysis presented in Section IV-A is difficult to directly extend to arbitrary quantum LDPC codes due to the structural differences among codes, the performance of our EWAInit-BP algorithm is not limited to 2-D topological constructions. In this section, we provide empirical evidence that EWAInit-BP generalizes well to other classes of quantum LDPC codes. Two representative examples are considered: the Bivariate Bicycle codes proposed by IBM [50], and 3-D Chamon codes, which is an instance of 3-D XYZ product of three repetition codes [51].

Fig. 10 presents a comparison of logical error rates between BP₂-OSD-0 and EWAInit-BP under parallel scheduling on these two code families under depolarizing noise and code capacity noise model. For the bivariate bicycle codes, EWAInit-BP almost consistently outperforms BP₂-OSD-0 across all tested code distances, with its advantage more pronounced at lower physical error rates.

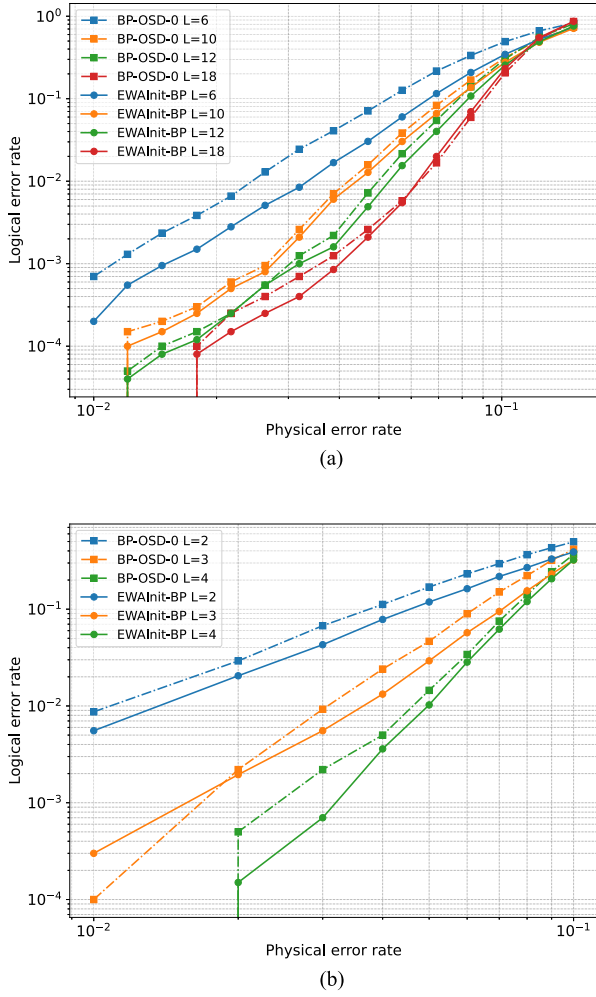


FIGURE 10. Comparison of BP₂-OSD-0 and EWAInit-BP on two non-surface-code quantum LDPC families. (a) Bivariate bicycle codes and (b) 3-D Chamon codes. EWAInit-BP almost outperforms BP-OSD-0 across all distances and physical error rates. (a) Bivariate bicycle codes. (b) 3-D Chamon codes.

Similarly, for the 3-D Chamon codes, EWAInit-BP demonstrates improved decoding performance for all tested distances. While both algorithms show typical threshold-like behavior, EWAInit-BP achieves lower logical error rates under comparable physical error rates. These results provide preliminary evidence that the exponential averaging mechanism is broadly beneficial for decoding quantum LDPC codes.

V. CONCLUSION

In this article, we proposed three improved BP algorithms: Momentum-BP, AdaGrad-BP, and EWAInit-BP.

Since the a posteriori probability update of BP can be interpreted as the gradient optimization of an energy function, we adopted optimization strategies from gradient descent in machine learning, resulting in Momentum-BP and AdaGrad-BP. These mechanisms aim to smooth the update process and reduce the amplitude of oscillations. Simulations

on trapping sets intuitively demonstrated the effectiveness of these algorithms, supporting the feasibility of leveraging gradient optimization principles for BP decoding.

From the theoretical derivation of the message update rules of the BP algorithm, we recognized the importance of a priori probabilities and transformed the message update smoothing into initial values adapting, proposing EWAInit-BP. This algorithm realizes an exponentially weighted average of past horizontal message updates and outperforms other algorithms without OSD postprocessing on various surface codes. Specifically, EWAInit-BP shows high decoding accuracy even under parallel scheduling, with logical error rates decreasing as code length increases on short blocks. Notably, all our proposed methods do not introduce any additional factors to the time complexity of BP. Thus, EWAInit-BP can be implemented with $O(1)$ complexity in fully parallelism while maintaining decoding accuracy.

Furthermore, we also evaluated the performance of EWAInit-BP on two representative quantum LDPC codes—bivariate bicycle codes and the 3-D Chamon codes. Our algorithm demonstrates competitive decoding performance on both families under parallel scheduling, showing potential adaptability to high-dimensional or irregular topologies. The results suggest that the principle behind message adaptation strategies may extend beyond the specific context of surface codes, indicating broader theoretical significance. Future work will explore how to adapt our approach to other quantum LDPC codes with higher degeneracy or more complex structures, especially focusing on maintaining decoding performance at larger code distances.

However, applying exponential weighting to messages is still insufficient to ensure BP convergence for all code lengths. It may be necessary to draw inspiration from methods like neural BP to enhance the representational power of the algorithm. Besides, our simulations only considered Pauli errors on qubits and did not account for measurement errors, error propagation in circuits, etc. Future work needs to adapt BP to the circuit-level noise model, potentially requiring specialized optimizations.

APPENDIX A BP AS GRADIENT OPTIMIZATION

Each iteration of BP can be interpreted as one step of gradient descent with a simplified step size on the following energy function [21], [52]:

$$J(\mathbf{Q}) = - \sum_{i=1}^m 2 \tanh^{-1} \left((-1)^{s_i} \prod_{v_j \in \mathcal{N}(c_i)} \tanh \left(\frac{\lambda_{H_{ij}}(Q_{v_j})}{2} \right) \right) \quad (27)$$

where \mathbf{Q} is the a posteriori probabilities of all the variable nodes in this iteration, and m is the number of check nodes. The term inside the parentheses measures the discrepancy between the result of the current iteration and the actual error syndrome of stabilizer S_i . The smaller the energy function, the closer the BP iteration result is to the correct solution.

The variables for gradient updates are the a posteriori estimates of each type of Pauli error of each variable node. Specifically, for the $W \in \{X, Y, Z\}$ type Pauli error on the n th variable node, the a posteriori probability update is

$$Q_{v_j, W}^{(t)} = Q_{v_j, W}^{(t-1)} - \alpha \frac{\partial J}{\partial Q_{v_j, W}^{(t-1)}} \quad (28)$$

where the partial derivative is

$$\begin{aligned} \frac{\partial J}{\partial Q_{v_j, W}} = & - \sum_{\substack{c_i \in \mathcal{M}(v_j) \\ \langle W, H_{ij} \rangle = 1}} \tilde{m}_{c_i \rightarrow v_j} \\ & * \frac{g_{ij}(Q) e^{-Q_{v_j, W}}}{e^{-Q_{v_j, X}} + e^{-Q_{v_j, Y}} + e^{-Q_{v_j, Z}} - e^{-Q_{v_j, H_{ij}}}} \end{aligned} \quad (29)$$

where g_{ij} is a term that is always greater than 0 (see [21]), and $\tilde{m}_{c_i \rightarrow v_j}$ is the a posteriori probability summation term, approximated using the tanh function

$$\tilde{m}_{c_i \rightarrow v_j} = (-1)^{s_i} \prod_{v_j \in \mathcal{N}(c_i) \setminus v_j} \tanh\left(\frac{\lambda_{H_{ij'}}(Q_{v_j'})}{2}\right). \quad (30)$$

If we view the last term on the RHS of (29) as a variable coefficient, then the posterior update in BP can be seen as a gradient descent step that omits this coefficient term.

APPENDIX B ITERATIVE EXPANSION OF EWA-BASED BP ALGORITHMS

Here, we present the expansion of Momentum-BP and EWAInit-BP proposed in this article to clarify their properties. To streamline calculations and maintain consistency, we set γ in Momentum-BP to 0.

A. MOMENTUM-BP

For Momentum-BP, the a posteriori probabilities are updated in the first iteration as

$$Q_{v_j, W}^{(1)} = \alpha(\Pi_{v_j, W}^{(0)} + \sum_{\substack{c_i \in \mathcal{M}(v_j) \\ \langle W, H_{ij} \rangle = 1}} m_{c_i \rightarrow v_j}^{(1)}) + (1 - \alpha)\Pi_{v_j, W}^{(0)} \quad (31)$$

$$= \Pi_{v_j, W}^{(0)} + \alpha \sum_{\substack{c_i \in \mathcal{M}(v_j) \\ \langle W, H_{ij} \rangle = 1}} m_{c_i \rightarrow v_j}^{(1)} \quad (32)$$

where the first term on the RHS represents the calculation in the current iteration. The second term represents the calculation from the previous iteration, which in the first iteration is simply the a priori probability, say $Q_{v_j, W}^{(0)} = \Pi_{v_j, W}^{(0)}$.

Now, we can expand the a posteriori update iteratively

$$Q_{v_j, W}^{(t)} = \alpha(\Pi_{v_j, W}^{(0)} + \sum_{\substack{c_i \in \mathcal{M}(v_j) \\ \langle W, H_{ij} \rangle = 1}} m_{c_i \rightarrow v_j}^{(t)}) + (1 - \alpha)Q_{v_j, W}^{(t-1)} \quad (33)$$

$$\begin{aligned} &= \alpha(\Pi_{v_j, W}^{(0)} + \sum_{\substack{c_i \in \mathcal{M}(v_j) \\ \langle W, H_{ij} \rangle = 1}} m_{c_i \rightarrow v_j}^{(t)}) + (1 - \alpha) * \\ &(\alpha(\Pi_{v_j, W}^{(0)} + \sum_{\substack{c_i \in \mathcal{M}(v_j) \\ \langle W, H_{ij} \rangle = 1}} m_{c_i \rightarrow v_j}^{(t-1)}) + (1 - \alpha)Q_{v_j, W}^{(t-2)}) \end{aligned} \quad (34)$$

$$\begin{aligned} &= \alpha \sum_{k=0}^{t-1} (1 - \alpha)^k \left(\Pi_{v_j, W}^{(0)} + \sum_{\substack{c_i \in \mathcal{M}(v_j) \\ \langle W, H_{ij} \rangle = 1}} m_{c_i \rightarrow v_j}^{(t-k)} \right) \\ &+ (1 - \alpha)^t Q_{v_j, W}^{(0)}. \end{aligned} \quad (35)$$

Using the initial condition $Q_{v_j, W}^{(0)} = \Pi_{v_j, W}^{(0)}$, we can substitute $Q_{v_j, W}^{(0)}$ into (35) to obtain

$$\begin{aligned} Q_{v_j, W}^{(t)} = & \Pi_{v_j, W}^{(0)} \left(\alpha \sum_{k=0}^{t-1} (1 - \alpha)^k + (1 - \alpha)^t \right) \\ & + \alpha \sum_{k=0}^{t-1} (1 - \alpha)^k \sum_{\substack{c_i \in \mathcal{M}(v_j) \\ \langle W, H_{ij} \rangle = 1}} m_{c_i \rightarrow v_j}^{(t-k)}. \end{aligned} \quad (36)$$

Now, let us examine the coefficient of $\Pi_{v_j, W}^{(0)}$

$$\alpha \sum_{k=0}^{t-1} (1 - \alpha)^k + (1 - \alpha)^t. \quad (37)$$

This is a finite geometric series, and we can simplify it using the formula

$$\sum_{k=0}^n r^k = \frac{1 - r^{n+1}}{1 - r}. \quad (38)$$

Substituting back into (37), we get

$$\begin{aligned} &\alpha \sum_{k=0}^{t-1} (1 - \alpha)^k + (1 - \alpha)^t \\ &= \alpha \cdot \frac{1 - (1 - \alpha)^t}{\alpha} + (1 - \alpha)^t = 1. \end{aligned} \quad (40)$$

Thus, we have shown that the coefficient of $\Pi_{v_j, W}^{(0)}$ is indeed 1. The complete general term formula is therefore

$$Q_{v_j, W}^{(t)} = \Pi_{v_j, W}^{(0)} + \sum_{k=0}^{t-1} \alpha (1 - \alpha)^k \sum_{\substack{c_i \in \mathcal{M}(v_j) \\ \langle W, H_{ij} \rangle = 1}} m_{c_i \rightarrow v_j}^{(t-k)}. \quad (41)$$

B. EWAINIT-BP

Similarly, for EWAInit-BP, we have the initial condition

$$\Pi_{v_j, W}^{(1)} = \Pi_{v_j, W}^{(0)} \quad (42)$$

$$Q_{v_j, W}^{(1)} = \Pi_{v_j, W}^{(1)} + \sum_{\substack{c_i \in \mathcal{M}(v_j) \\ \langle W, H_{ij} \rangle = 1}} m_{c_i \rightarrow v_j}^{(1)} \quad (43)$$

and the general term form

$$\Pi_{v_j, W}^{(t)} = \alpha \Pi_{v_j, W}^{(0)} + (1 - \alpha) Q_{v_j, W}^{(t-1)} \quad (44)$$

$$= \Pi_{v_j, W}^{(0)} + \sum_{k=1}^t (1 - \alpha)^k \sum_{\substack{c_i \in \mathcal{M}(v_j) \\ \langle W, H_{ij} \rangle = 1}} m_{c_i \rightarrow v_j}^{(t-k)} \quad (45)$$

$$Q_{v_j, W}^{(t)} = \Pi_{v_j, W}^{(t)} + \sum_{\substack{c_i \in \mathcal{M}(v_j) \\ \langle W, H_{ij} \rangle = 1}} m_{c_i \rightarrow v_j}^{(t)} \quad (46)$$

$$= \Pi_{v_j, W}^{(0)} + \sum_{k=0}^t (1 - \alpha)^k \sum_{\substack{c_i \in \mathcal{M}(v_j) \\ \langle W, H_{ij} \rangle = 1}} m_{c_i \rightarrow v_j}^{(t-k)} \quad (47)$$

where the coefficient of $\Pi_{v_j, W}^{(0)}$ is also 1, derived by the same approach.

Although the approaches for improvement are different, both Momentum-BP and EWAInit-BP attempt to enhance the update of the a posteriori probability $Q_{v_j, W}^{(t)}$ by applying a weighted combination of past messages. The distinction lies in the fact that the recurrence relation derived from performing EWA on the prior probability is more concise and intuitive, with only a $(1 - \alpha)$ factor applied to past messages. Interestingly, EWAInit-BP, with its simpler recurrence relation, also demonstrates better performance in simulations.

APPENDIX C COMPARED TO AMBP

AMBP [21] defines a parameter range (for surface codes, $\alpha \in [1, 0.5]$) and iterates through the possible values from large to small until MBP converges with a specific parameter α^* , at which point it returns the result. Sorting the parameters from large to small prevents overly aggressive message updates that might cause logical errors, while ensuring convergence in most cases. In this appendix, we uniformly sampled 11 values within the range $[1, 0.5]$, making the worst case time overhead for AMBP approximately ten times that of MBP. We can utilize a similar adaptive parameter selection strategy for EWAInit-BP using a parameter range of $\alpha \in [1, 0]$, where the parameter represents the strength of the EWA. We refer to the resulting algorithm as Adaptive EWAInit-BP.

We compare EWAInit-BP, AEWA-BP with parallel scheduling, and AMBP with both parallel and serial scheduling on planar surface codes of sizes $[[13, 1, 3]]$, $[[85, 1, 7]]$, and $[[221, 1, 11]]$, as shown in Fig. 11. EWAInit-BP and parallel AMBP no longer exhibit the property of decreasing




FIGURE 11. Performance of EWAInit-BP, AEWA-BP, parallel AMBP, and serial AMBP on the planar surface code with $L = 3, 7, 11$. EWAInit-BP and parallel AMBP achieve similar accuracy, but both exhibit higher error rates at $L = 11$ compared to $L = 7$. AEWA-BP achieves accuracy comparable to serial AMBP.

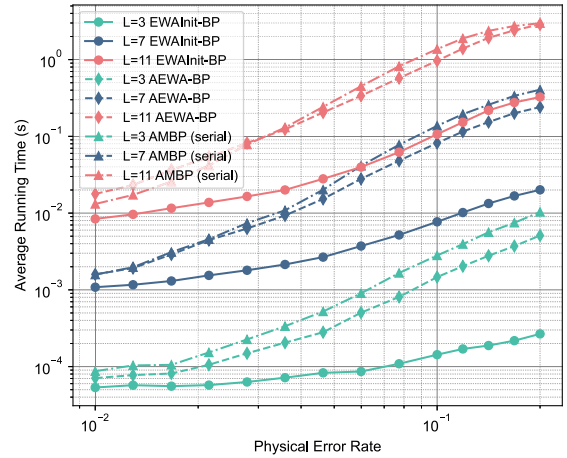


FIGURE 12. Average running time of EWAInit-BP, AEWA-BP, and serial AMBP on the planar surface code with $L = 3, 7, 11$. For the latter two algorithms that utilize adaptive methods, the running time increases significantly as the physical error rate grows.

logical error rates with increasing code length, indicating that large surface codes remain a challenge for parallel BP. However, AEWA-BP with parallel scheduling demonstrates accuracy comparable to AMBP with serial scheduling. A major issue with AEWA-BP is that its parameter range cannot cover all possible magnitudes of message updates, meaning that it theoretically cannot guarantee convergence for all errors.

Next, we compared the runtime of the aforementioned algorithms to verify the high overhead of adaptive methods discussed in the main text, as shown in Fig. 12. The runtime of parallel AMBP is very similar to that of AEWA-BP and is, therefore, omitted from the figure. As shown, EWAInit-BP exhibits the same scaling properties as traditional BP: its runtime grows linearly with code length and increases with

TABLE 2. Comparison of the Running Time of EWAInit-BP, AEWA-BP, and AMBP Under Certain Physical Error Rates

PER	EWAInit-BP (ms)	AEWA-BP (ms)	AMBP (ms)
10^{-2}	1.1	1.6	1.6
6×10^{-2}	3.7	27.9	41.3
10^{-1}	7.7	81.9	137.0
2×10^{-1}	20.1	240.6	405.6

physical error rates due to the rising number of iterations (up to a maximum of 100).

However, algorithms utilizing adaptive methods show higher scaling. As the physical error rate increases, these algorithms tend to fully traverse the parameter range, resulting in runtime scaling not only with the number of iterations but also with the parameter range length in high-error-rate scenarios. In this experiment, the parameter range is set to 11 values. It can be observed that the runtime of adaptive methods is an order of magnitude higher than that of EWAInit-BP under high error rates. Table 2 presents the runtime (ms) of the three algorithms under several error rates, offering a more intuitive perspective.

REFERENCES

- [1] Y. Wang et al., "Single-qubit quantum memory exceeding ten-minute coherence time," *Nat. Photon.*, vol. 11, no. 10, pp. 646–650, 2017, doi: [10.1038/s41566-017-0007-1](https://doi.org/10.1038/s41566-017-0007-1).
- [2] Y. Zhao et al., "Realization of an error-correcting surface code with superconducting qubits," *Phys. Rev. Lett.*, vol. 129, Jul. 2022, Art. no. 030501, doi: [10.1103/PhysRevLett.129.030501](https://doi.org/10.1103/PhysRevLett.129.030501).
- [3] Google Quantum AI, "Suppressing quantum errors by scaling a surface code logical qubit," *Nature*, vol. 614, no. 7949, pp. 676–681, 2023, doi: [10.1038/s41586-022-05434-1](https://doi.org/10.1038/s41586-022-05434-1).
- [4] P. Shor, "Fault-tolerant quantum computation," in *Proc. 37th Conf. Found. Comput. Sci.*, 1996, pp. 56–65, doi: [10.1109/SFCS.1996.548464](https://doi.org/10.1109/SFCS.1996.548464).
- [5] B. M. Terhal, "Quantum error correction for quantum memories," *Rev. Mod. Phys.*, vol. 87, pp. 307–346, Apr. 2015, doi: [10.1103/RevModPhys.87.307](https://doi.org/10.1103/RevModPhys.87.307).
- [6] J. Edmonds, "Paths, trees, and flowers," *Can. J. Math.*, vol. 17, pp. 449–467, 1965, doi: [10.4153/CJM-1965-045-4](https://doi.org/10.4153/CJM-1965-045-4).
- [7] N. Delfosse, V. Londe, and M. E. Beverland, "Toward a union-find decoder for quantum LDPC codes," *IEEE Trans. Inf. Theory*, vol. 68, no. 5, pp. 3187–3199, May 2022, doi: [10.1109/TIT.2022.3143452](https://doi.org/10.1109/TIT.2022.3143452).
- [8] S. Bravyi, M. Suchara, and A. Vargo, "Efficient algorithms for maximum likelihood decoding in the surface code," *Phys. Rev. A*, vol. 90, Sep. 2014, Art. no. 032326, doi: [10.1103/PhysRevA.90.032326](https://doi.org/10.1103/PhysRevA.90.032326).
- [9] G. Duclos-Cianci and D. Poulin, "Fast decoders for topological quantum codes," *Phys. Rev. Lett.*, vol. 104, no. 5, 2010, Art. no. 050504, doi: [10.1103/PhysRevLett.104.050504](https://doi.org/10.1103/PhysRevLett.104.050504).
- [10] G. Torlai and R. G. Melko, "Neural decoder for topological codes," *Phys. Rev. Lett.*, vol. 119, Jul. 2017, Art. no. 030501, doi: [10.1103/PhysRevLett.119.030501](https://doi.org/10.1103/PhysRevLett.119.030501).
- [11] N. P. Breuckmann and X. Ni, "Scalable neural network decoders for higher dimensional quantum codes," *Quantum*, vol. 2, 2018, Art. no. 68, doi: [10.22331/q-2018-05-24-68](https://doi.org/10.22331/q-2018-05-24-68).
- [12] Y.-H. Liu and D. Poulin, "Neural belief-propagation decoders for quantum error-correcting codes," *Phys. Rev. Lett.*, vol. 122, May 2019, Art. no. 200501, doi: [10.1103/PhysRevLett.122.200501](https://doi.org/10.1103/PhysRevLett.122.200501).
- [13] X. Ni, "Neural network decoders for large-distance 2D toric codes," *Quantum*, vol. 4, 2020, Art. no. 310, doi: [10.22331/q-2020-08-24-310](https://doi.org/10.22331/q-2020-08-24-310).
- [14] K. Meinerz, C.-Y. Park, and S. Trebst, "Scalable neural decoder for topological surface codes," *Phys. Rev. Lett.*, vol. 128, Feb. 2022, Art. no. 080505, doi: [10.1103/PhysRevLett.128.080505](https://doi.org/10.1103/PhysRevLett.128.080505).
- [15] R. W. J. Overwater, M. Babaie, and F. Sebastiano, "Neural-network decoders for quantum error correction using surface codes: A space exploration of the hardware cost-performance tradeoffs," *IEEE Trans. Quantum Eng.*, vol. 3, 2022, Art. no. 3101719, doi: [10.1109/TQE.2022.3174017](https://doi.org/10.1109/TQE.2022.3174017).
- [16] Y. Choukroun and L. Wolf, "Deep quantum error correction," in *Proc. AAAI Conf. Artif. Intell.*, 2024, vol. 38, pp. 64–72, doi: [10.1609/aaai.v38i1.27756](https://doi.org/10.1609/aaai.v38i1.27756).
- [17] P. Andreasson, J. Johansson, S. Liljestrand, and M. Granath, "Quantum error correction for the toric code using deep reinforcement learning," *Quantum*, vol. 3, 2019, Art. no. 183, doi: [10.22331/q-2019-09-02-183](https://doi.org/10.22331/q-2019-09-02-183).
- [18] 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. Archit.*, 2020, pp. 556–569, doi: [10.1109/ISCA45697.2020.00053](https://doi.org/10.1109/ISCA45697.2020.00053).
- [19] F. E. A. Battistel, "Real-time decoding for fault-tolerant quantum computing: Progress, challenges and outlook," *Nano Futures*, vol. 7, no. 3, 2023, Art. no. 032003, doi: [10.1088/2399-1984/aceba6](https://doi.org/10.1088/2399-1984/aceba6).
- [20] D. MacKay, G. Mitchison, and P. McFadden, "Sparse-graph codes for quantum error correction," *IEEE Trans. Inf. Theory*, vol. 50, no. 10, pp. 2315–2330, Oct. 2004, doi: [10.1109/TIT.2004.834737](https://doi.org/10.1109/TIT.2004.834737).
- [21] K.-Y. Kuo and C.-Y. Lai, "Exploiting degeneracy in belief propagation decoding of quantum codes," *Npj Quantum Inf.*, vol. 8, no. 1, 2022, Art. no. 111, doi: [10.1038/s41534-022-00623-2](https://doi.org/10.1038/s41534-022-00623-2).
- [22] J. Valls, F. Garcia-Herrero, N. Raveendran, and B. Vasić, "Syndrome-based min-sum vs OSD-0 decoders: FPGA implementation and analysis for quantum LDPC codes," *IEEE Access*, vol. 9, pp. 138734–138743, 2021, doi: [10.1109/ACCESS.2021.3118544](https://doi.org/10.1109/ACCESS.2021.3118544).
- [23] R. Acharya et al., "Quantum error correction below the surface code threshold," *Nature*, vol. 638, pp. 920–926, 2025, doi: [10.1038/s41586-024-08449-y](https://doi.org/10.1038/s41586-024-08449-y).
- [24] D. Poulin and Y. Chung, "On the iterative decoding of sparse quantum codes," 2008, *arXiv:0801.1241*, doi: [10.48550/arXiv.0801.1241](https://doi.org/10.48550/arXiv.0801.1241).
- [25] Y.-J. Wang, B. C. Sanders, B.-M. Bai, and X.-M. Wang, "Enhanced feedback iterative decoding of sparse quantum codes," *IEEE Trans. Inf. Theory*, vol. 58, no. 2, pp. 1231–1241, Feb. 2012, doi: [10.1109/TIT.2011.2169534](https://doi.org/10.1109/TIT.2011.2169534).
- [26] Z. Babar, P. Botsinis, D. Alanis, S. X. Ng, and L. Hanzo, "Fifteen years of quantum LDPC coding and improved decoding strategies," *IEEE Access*, vol. 3, pp. 2492–2519, 2015, doi: [10.1109/ACCESS.2015.2503267](https://doi.org/10.1109/ACCESS.2015.2503267).
- [27] N. Raveendran, M. Bahrami, and B. Vasic, "Syndrome-generalized belief propagation decoding for quantum memories," in *Proc. IEEE Int. Conf. Commun.*, May 2019, pp. 1–6, doi: [10.1109/ICC.2019.8761366](https://doi.org/10.1109/ICC.2019.8761366).
- [28] K.-Y. Kuo and C.-Y. Lai, "Refined belief propagation decoding of sparse-graph quantum codes," *IEEE J. Select. Areas Inf. Theory*, vol. 1, no. 2, pp. 487–498, Aug. 2020, doi: [10.1109/JSAIT.2020.3011758](https://doi.org/10.1109/JSAIT.2020.3011758).
- [29] K.-Y. Kuo and C.-Y. Lai, "Comparison of 2D topological codes and their decoding performances," in *Proc. IEEE Int. Symp. Inf. Theory*, 2022, pp. 186–191, doi: [10.1109/ISIT50566.2022.9834489](https://doi.org/10.1109/ISIT50566.2022.9834489).
- [30] A. Rigby, J. Olivier, and P. Jarvis, "Modified belief propagation decoders for quantum low-density parity-check codes," *Phys. Rev. A*, vol. 100, no. 1, 2019, Art. no. 012330, doi: [10.1103/PhysRevA.100.012330](https://doi.org/10.1103/PhysRevA.100.012330).
- [31] C.-Y. Lai and K.-Y. Kuo, "Log-domain decoding of quantum LDPC codes over binary finite fields," *IEEE Trans. Quantum Eng.*, vol. 2, 2021, Art. no. 2103615, doi: [10.1109/TQE.2021.3113936](https://doi.org/10.1109/TQE.2021.3113936).
- [32] Z. Yi et al., "Improved belief propagation decoding algorithm based on decoupling representation of Pauli operators for quantum stabilizer codes," *Quantum Inf. Process.*, vol. 24, no. 4, pp. 1–25, 2025, doi: [10.1007/s11128-025-04709-6](https://doi.org/10.1007/s11128-025-04709-6).
- [33] J. Old and M. Rispler, "Generalized belief propagation algorithms for decoding of surface codes," *Quantum*, vol. 7, 2023, Art. no. 1037, doi: [10.22331/q-2023-06-07-1037](https://doi.org/10.22331/q-2023-06-07-1037).
- [34] D. Chytas, M. Pacenti, N. Raveendran, M. F. Flanagan, and B. Vasić, "Enhanced message-passing decoding of degenerate quantum codes utilizing trapping set dynamics," *IEEE Commun. Lett.*, vol. 28, no. 3, pp. 444–448, Mar. 2024, doi: [10.1109/LCOMM.2024.3356312](https://doi.org/10.1109/LCOMM.2024.3356312).
- [35] J. D. Crest, M. Mhalla, and V. Savin, "Stabilizer inactivation for message-passing decoding of quantum LDPC codes," in *Proc. IEEE Inf. Theory Workshop*, 2022, pp. 488–493, doi: [10.1109/ITW54588.2022.9965902](https://doi.org/10.1109/ITW54588.2022.9965902).
- [36] B. Criger and I. Ashraf, "Multi-path summation for decoding 2D topological codes," *Quantum*, vol. 2, Oct., 2018, Art. no. 102, doi: [10.22331/q-2018-10-19-102](https://doi.org/10.22331/q-2018-10-19-102).

- [37] O. Higgott, T. C. Bohdanowicz, A. Kubica, S. T. Flammia, and E. T. Campbell, "Improved decoding of circuit noise and fragile boundaries of tailored surface codes," *PRX Quantum*, vol. 13, no. 3, 2023, Art. no. 031007, doi: [10.1103/PhysRevX.13.031007](https://doi.org/10.1103/PhysRevX.13.031007).
- [38] L. Caune, B. Reid, J. Camps, and E. Campbell, "Belief propagation as a partial decoder," 2023, *arXiv:2306.17142*, doi: [10.48550/arXiv.2306.17142](https://doi.org/10.48550/arXiv.2306.17142).
- [39] P. Fuentes, J. Etchezarreta Martinez, P. M. Crespo, and J. Garcia-Frías, "Degeneracy and its impact on the decoding of sparse quantum codes," *IEEE Access*, vol. 9, pp. 89093–89119, 2021, doi: [10.1109/ACCESS.2021.3089829](https://doi.org/10.1109/ACCESS.2021.3089829).
- [40] N. Raveendran and B. Vasić, "Trapping sets of quantum LDPC codes," *Quantum*, vol. 5, 2021, Art. no. 562, doi: [10.22331/q-2021-10-14-562](https://doi.org/10.22331/q-2021-10-14-562).
- [41] N. P. Breuckmann and J. N. Eberhardt, "Quantum low-density parity-check codes," *PRX Quantum*, vol. 2, no. 4, Art. no. 040101, 2021, doi: [10.1103/PRXQuantum.2.040101](https://doi.org/10.1103/PRXQuantum.2.040101).
- [42] H. Bombin, "An introduction to topological quantum codes," 2013, *arXiv:1311.0277*, doi: [10.48550/arXiv.1311.0277](https://doi.org/10.48550/arXiv.1311.0277).
- [43] M.-H. Hsieh and F. Le Gall, "Np-hardness of decoding quantum error-correction codes," *Phys. Rev. A*, vol. 83, no. 5, 2011, Art. no. 052331, doi: [10.1103/PhysRevA.83.052331](https://doi.org/10.1103/PhysRevA.83.052331).
- [44] D. Hocevar, "A reduced complexity decoder architecture via layered decoding of LDPC codes," in *Proc. IEEE Workshop Signal Process. Syst.*, Dec. 2004, pp. 107–112, doi: [10.1109/SIPS.2004.1363033](https://doi.org/10.1109/SIPS.2004.1363033).
- [45] M. K. Dehkordi and A. H. Banihashemi, "An efficient algorithm for finding dominant trapping sets of LDPC codes," in *Proc. 6th Int. Symp. Turbo Codes Iterative Inf. Process.*, 2010, pp. 444–448, doi: [10.1109/ISTC.2010.5613902](https://doi.org/10.1109/ISTC.2010.5613902).
- [46] S. Han, J. Oh, K. Oh, and J. Ha, "Deep-learning for breaking the trapping sets in low-density parity-check codes," *IEEE Trans. Commun.*, vol. 70, no. 5, pp. 2909–2923, May 2022, doi: [10.1109/TCOMM.2022.3157314](https://doi.org/10.1109/TCOMM.2022.3157314).
- [47] S. Kang, J. Moon, J. Ha, and J. Shin, "Breaking the trapping sets in LDPC codes: Check node removal and collaborative decoding," *IEEE Trans. Commun.*, vol. 64, no. 1, pp. 15–26, Jan. 2016, doi: [10.1109/TCOMM.2015.2498161](https://doi.org/10.1109/TCOMM.2015.2498161).
- [48] A. K. Pradhan, N. Raveendran, N. Rengaswamy, X. Xiao, and B. Vasić, "Learning to decode trapping sets in QLDPC codes," in *Proc. 12th Int. Symp. Top. Coding*, 2023, pp. 1–5, doi: [10.1109/ISTC57237.2023.10273526](https://doi.org/10.1109/ISTC57237.2023.10273526).
- [49] J. Roffe, D. R. White, S. Burton, and E. Campbell, "Decoding across the quantum low-density parity-check code landscape," *Phys. Rev. Res.*, vol. 2, no. 4, 2020, Art. no. 043423, doi: [10.1103/PhysRevResearch.2.043423](https://doi.org/10.1103/PhysRevResearch.2.043423).
- [50] 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*, vol. 627, no. 8005, pp. 778–782, 2024, doi: [10.1038/s41586-024-07107-7](https://doi.org/10.1038/s41586-024-07107-7).
- [51] A. Leverrier, S. Apers, and C. Vuillot, "Quantum XYZ product codes," *Quantum*, vol. 6, 2022, Art. no. 766, doi: [10.22331/q-2022-07-14-766](https://doi.org/10.22331/q-2022-07-14-766).
- [52] R. Lucas, M. Bossert, and M. Breitbart, "On iterative soft-decision decoding of linear binary block codes and product codes," *IEEE J. Sel. Areas Commun.*, vol. 16, no. 2, pp. 276–296, Feb. 1998, doi: [10.1109/49.661116](https://doi.org/10.1109/49.661116).