

Erasure Decoding for Quantum LDPC Codes via Belief Propagation with Guided Decimation

Mert Gökduman^{*†}, Hanwen Yao^{*†}, and Henry D. Pfister^{*†§}

^{*}Duke Quantum Center, Duke University, Durham, NC, USA

[†]Department of Electrical and Computer Engineering, Duke University, Durham, NC, USA

[‡]Department of Mathematics, Duke University, Durham, NC, USA

Abstract—Quantum low-density parity-check (LDPC) codes are a promising family of quantum error-correcting codes for fault tolerant quantum computing with low overhead. Decoding quantum LDPC codes on quantum erasure channels has received more attention recently due to advances in erasure conversion for various types of qubits including neutral atoms, trapped ions, and superconducting qubits. Belief propagation with guided decimation (BPGD) decoding of quantum LDPC codes has demonstrated good performance in bit-flip and depolarizing noise. In this work, we apply BPGD decoding to quantum erasure channels. Using a natural modification, we show that BPGD offers competitive performance on quantum erasure channels for multiple families of quantum LDPC codes. Furthermore, we show that the performance of BPGD decoding on erasure channels can sometimes be improved significantly by either adding damping or adjusting the initial channel log-likelihood ratio for bits that are not erased. More generally, our results demonstrate BPGD is an effective general-purpose solution for erasure decoding across the quantum LDPC landscape.

I. INTRODUCTION

For the development of scalable and fault-tolerant quantum computation, quantum error correction is a crucial component that protects quantum information against noise. Among the proposed error correction schemes, quantum low-density parity-check (LDPC) codes stand out as strong candidates because they promise lower overhead [1], [2] when compared to topological codes such as surface codes [3], [4] and color codes [5]. Belief-propagation (BP) decoding of quantum LDPC codes was first introduced in [6] and considered further in [7]. Recent breakthrough results have introduced constructions of asymptotically *good* quantum LDPC codes with constant rate and linear minimum distance [8]–[10]. In terms of practical implementation, results by Bravyi et al. [11] have shown how certain quantum LDPC codes can be embedded into a bilayer hardware architecture.

In this paper, we focus on decoding quantum LDPC codes over the quantum erasure channel [12]. This model has received more attention recently due to proposals and demonstrations that erased qubits can be realized in several architectures including neutral atom [13]–[15], trapped ions [16], and superconducting qubits [17], [18]. Moreover, it has been shown that quantum error correction schemes based on erased qubits achieve better finite-length performance, and higher thresholds compared to those designed for Pauli noise [14], [19]–[21]. In the erasure model we use, a random subset of the coded qubits is chosen and then subjected to uniform

random Pauli errors. While the subset is known to the decoder, the error values are not.

Several decoding algorithms have been proposed for quantum erasure correction, targeting various classes of quantum codes in the code. In [22], a linear-time decoder was introduced for surface codes, achieving maximum-likelihood (ML) performance by peeling on a spanning tree of erasures on the surface code lattice. This was later extended into the union-find decoder, initially for topological codes [21] and subsequently for more general quantum LDPC codes [23], which is capable of correcting both Pauli errors and erasures with a higher complexity. In [24], a trimming decoder was proposed for the erasure decoding of color codes, that combines peeling on a spanning tree with erasure set extension or vertex inactivation. Erasure decoding of subsystem color codes has also been studied using a combination of techniques, including peeling, clustering, and gauge fixing [25], [26].

Two erasure decoding algorithms were proposed by Connolly et al. in [27]: pruned peeling and vertical-horizontal (VH) decoding. The pruned peeling decoder can be applied to any code and it combines peeling with a greedy search for stabilizers contained wholly within the erased qubits. The VH decoder, which can only be applied to hypergraph product (HGP) codes, integrates pruned peeling with an iterative procedure that mitigates vertical and horizontal stopping sets. For the sub-threshold regime, simulation of the VH decoder has demonstrated performance close to ML with a computational complexity of $O(n^2)$ for codes of length n .

In this work, we evaluate belief propagation with guided decimation (BPGD) decoding of quantum LDPC codes for the quantum erasure channel. A recent study by Yao et al. demonstrated that the BPGD decoder has excellent performance for quantum LDPC codes under bit-flip noise and depolarizing noise in the code capacity model [28]. Over circuit-level noise, a recent study also shows that a variant of BPGD applied as the inner decoder of sliding window decoding achieves performance on par with BP+OSD for quantum LDPC codes [29]. Here, we show that BPGD outperforms both the peeling and pruned peeling decoders when applied to erasures. While its performance lags slightly behind the VH decoder for HGP codes, BPGD offers a significant computational advantage with its reduced complexity. We further show that tuning the prior log-likelihood ratios (LLRs) and introducing damping techniques significantly improves performance and enables

BPGD to approach the performance of VH decoding with reduced complexity.

II. PRELIMINARIES

A. Classical Erasure Correction

Error-correcting codes protect information against noise by introducing redundancy. Let $\mathbb{F}_2 = \{0, 1\}$ denote the Galois field with 2 elements. A binary linear code $\mathcal{C} \subseteq \mathbb{F}_2^n$ is a subspace of binary vectors satisfying $x_1 + x_2 \in \mathcal{C}$ for all $x_1, x_2 \in \mathcal{C}$. The code can be specified either by a generator matrix $G \in \mathbb{F}_2^{k \times n}$ whose rows span the code or by a parity-check matrix $H \in \mathbb{F}_2^{m \times n}$ whose rows are orthogonal to all codewords.

When a codeword is transmitted over the classical binary erasure channel (BEC), each bit is either received correctly or erased with some probability. The received vector $y \in \{0, 1, ?\}^n$ identifies the positions of the erasures. The decoder's task is to recover the original codeword x from the observation y . In *erasure syndrome decoding*, the decoder only sees the locations of the erasures and the syndrome $s = H\tilde{y}$, where $\tilde{y} \in \mathbb{F}_2^n$ is a binary version of y whose erasures are replaced by uniform random bits. This allows the decoder to recover x if a unique solution exists. For further details, see [27, Section 2].

B. Stabilizer Formalism

For a single qubit, a pure state is defined by a unit vector in \mathbb{C}^2 and the Pauli matrices I, X, Y, Z generate a unitary subgroup that acts on the qubit, where

$$I = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad Y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \quad Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}.$$

For an n -qubit system, the Hilbert space $\mathcal{H} = (\mathbb{C}^2)^{\otimes n}$ is used to represent all pure quantum states. The Pauli group \mathcal{P}_n is the subgroup of unitary transformations on \mathcal{H} that consists of all n -fold tensor products of the Pauli matrices along with a coefficient in $\{\pm 1, \pm i\}$. An element in \mathcal{P}_n can be written as

$$P = \alpha P_1 \otimes P_2 \otimes \cdots \otimes P_n, \quad (1)$$

where $\alpha \in \{\pm 1, \pm i\}$ and $P_1, P_2, \dots, P_n \in \{I, X, Y, Z\}$.

An $[[n, k]]$ stabilizer code is a quantum error-correcting code that protects k logical qubits encoded into n physical qubits. A stabilizer group \mathcal{S} is an abelian subgroup of the Pauli group \mathcal{P}_n that does not include $-I^{\otimes n}$. The stabilizer code \mathcal{C} defined by the stabilizer group \mathcal{S} is the subspace $\mathcal{C} \subseteq \mathcal{H}$ that is invariant under the action of all operators in \mathcal{S} . For a stabilizer code, the *stabilizer generators* are a set of operators that generate the stabilizer group \mathcal{S} .

Thus, for any state $|\psi\rangle \in \mathcal{C}$ and $S \in \mathcal{S}$, we have:

$$S|\psi\rangle = |\psi\rangle. \quad (2)$$

The *weight* of a Pauli operator $P \in \mathcal{P}_n$ is the number of qubits on which the Pauli operator acts non-trivially. For instance, the weight of $X \otimes I \otimes Z$ in \mathcal{P}_3 is 2. The *distance* d of a stabilizer code is defined as the minimum weight of a Pauli operator that is not in \mathcal{S} but which commutes with all elements of the

stabilizer group \mathcal{S} . Such an operator is referred to as a *logical operator*, as it acts non-trivially on the encoded qubits.

In this paper, we focus on the *Calderbank-Shor-Steane (CSS) code*, a class of stabilizer code constructed using two classical linear codes C_1 and C_2 satisfying $C_2^\perp \subseteq C_1$. The stabilizer generators of a CSS code can be divided into X -type and Z -type operators:

$$S_X = \{X^v : v \in C_2\}, \quad S_Z = \{Z^u : u \in C_1^\perp\},$$

where $X^v := X_1^{v_1} \otimes \cdots \otimes X_n^{v_n}$ and Z^u is defined similarly. The binary representations of these stabilizers are generated by a pair of parity check matrices, H_X and H_Z , satisfying

$$H_X H_Z^T = 0. \quad (3)$$

This condition is the specialization to CSS codes of the more general *commutativity constraint* required by all stabilizer codes (i.e., that the stabilizer group is commutative). For CSS codes, the rows of H_X and H_Z correspond to X -type and Z -type stabilizer generators. Since stabilizers of the same type automatically commute, the above condition ensures that the stabilizer generators corresponding to H_X commute with those defined by H_Z .

For noise model where the X - and Z -type errors are independent, this structure simplifies the error correction procedure by leveraging the separation of X - and Z -type errors. In this context, we cover the correction of X -type errors using classical decoding of C_2 , with the understanding that the correction of Z -type errors follows the same approach using classical decoding of C_1^\perp , thus effectively reducing the quantum error correction problem to two separate classical decoding tasks.

C. Syndrome Decoding for the Quantum Erasure Channel

Erasure channels provide a simplified model for information loss by specifying the locations of likely errors. This allows for efficient decoding, especially with LDPC codes, which perform exceptionally well on these channels. In classical coding theory, LDPC codes first demonstrated their potential to approach channel capacity through iterative decoding algorithms like belief propagation [30]–[32]. As the study of erasure channels helped advance classical coding, we expect similar benefits for quantum codes in terms of improving decoding strategies and understanding performance limits.

In this work, we consider the quantum erasure channel [12], [19] as our noise model, where each qubit in the encoded state of a stabilizer code is independently erased with probability p . When a qubit is erased, it is affected by a Pauli operator in $\{I, X, Y, Z\}$ chosen uniformly at random. The locations of the erasures are given to the decoder as side information.

To correct the Pauli error $E \in \mathcal{P}_n$ affecting the erased qubits, we measure the syndrome of the stabilizers. The goal of the decoder is to identify a Pauli error estimate \hat{E} that matches the syndrome and acts trivially outside of the erased locations. This decoding process is successful if the decoded error either

exactly matches the actual Pauli error or differs from it by a stabilizer. In other words, the decoding is successful if:

$$E\mathcal{S} = \hat{E}\mathcal{S}, \quad (4)$$

where \mathcal{S} denotes the stabilizer group, and $E\mathcal{S}$ denotes its coset shifted by E . This condition implies that the identified error is logically equivalent to the actual error up to a stabilizer transformation.

The non-trivial elements of E are restricted to the locations of the erased qubits. Since the locations of the erasures are known, the decoding problem is constrained to finding an error in the support of the known erasures that matches the syndrome. There are four possible outcomes for this decoding process:

- 1) **Exact Match:** The decoder identifies an error \hat{E} that exactly matches the actual error E , meaning $\hat{E} = E$.
- 2) **Degenerate Match:** The identified error \hat{E} differs from the error E by a stabilizer, i.e., $\hat{E} \cdot \mathcal{S} = E \cdot \mathcal{S}$. In this case, the identified error is logically equivalent to the actual error but does not affect the logical state.
- 3) **Logical Error:** The identified error \hat{E} differs from the actual error E by a logical operator, i.e., $\hat{E} \cdot \mathcal{S} \neq E \cdot \mathcal{S}$, resulting in a logical error that changes the logical state.
- 4) **Decoder Failure:** The decoder fails to find any error in the support of the erasures that matches the syndrome, resulting in no valid decoding solution.

The probability of each Pauli operator I , X , Y , or Z on the erased qubits is $1/4$, as previously stated. Importantly, for CSS codes, these probabilities translate to independent $1/2$ probabilities for X -type and Z -type errors. This is because a Pauli- Y error corresponds to both an X -type error and a Z -type error occurring on the same qubit.

Thus, the syndrome decoding process is separated into independent corrections of X -type and Z -type errors. The syndrome vector s_X corresponding to Z -stabilizers provides information about X -errors, and the syndrome vector s_Z from X -stabilizers provides information about Z -errors. The decoding problem reduces to solving:

$$s_X = H_Z e_X^T, \quad s_Z = H_X e_Z^T, \quad (5)$$

where H_X and H_Z are the parity-check matrices corresponding to the X - and Z -stabilizers, and e_X and e_Z are the error vectors for X - and Z -type errors, respectively.

D. Quantum LDPC Codes

Quantum Low-Density Parity-Check (QLDPC) codes are a class of quantum error-correcting codes characterized by sparse parity check matrices. These codes are an extension of classical LDPC codes [33], [34] and they offer promising performance for fault-tolerant quantum computing due to their low overhead relative to topological codes [1], [11], [35].

A CSS quantum LDPC code can be represented by a pair of sparse matrices, H_X and H_Z , whose rows and columns each have only a small number of non-zero entries. This sparsity is crucial both for the practical implementation of

stabilizer measurement and for decoding algorithms like belief propagation.

Here we describe an important class of quantum LDPC codes called hypergraph product (HGP) codes. Constructed from two classical binary linear codes with parity-check matrices $H_1 \in \mathbb{F}_2^{m_1 \times n_1}$ and $H_2 \in \mathbb{F}_2^{m_2 \times n_2}$, the HGP code is defined by the check matrices

$$H_X = (H_1 \otimes I_{n_2} \quad I_{m_1} \otimes H_2^T), \quad (6)$$

$$H_Z = (I_{n_1} \otimes H_2 \quad H_1^T \otimes I_{m_2}). \quad (7)$$

An extension of this construction called the *lifted product construction* generalizes the HGP framework [36]. While HGP codes have binary entries in their parity-check matrices, lifted product codes replace these scalar entries with higher-dimensional objects such as *circulant matrices*. This process is called *lifting*, as it effectively "lifts" scalar entries to matrix entries, increasing the number of qubits and stabilizers

In the lifted product construction, the parity-check matrices become:

$$H_X = (\tilde{H}_1 \otimes I_{n_2} \quad I_{m_1} \otimes \tilde{H}_2^T), \quad (8)$$

$$H_Z = (I_{n_1} \otimes \tilde{H}_2 \quad \tilde{H}_1^T \otimes I_{m_2}), \quad (9)$$

where \tilde{H}_1 and \tilde{H}_2 are $m_1 \times n_1$ and $m_2 \times n_2$ matrices, respectively, whose elements are $L \times L$ binary circulant matrices. The circulant matrices form a commutative ring under matrix multiplication, which plays an essential role for maintaining the commutation relations required for the code to satisfy the commutativity constraint.

In this work, we use quantum LDPC codes from the two aforementioned code families to evaluate the decoding performance of BPGD over the quantum erasure channel.

III. DECODING ERASURES WITH BPGD

A. Belief Propagation Decoding

BP is an iterative message-passing algorithm that operates on the Tanner graph of an LDPC code [37]. For CSS codes, the correction of independent X - and Z -type errors can be treated separately. Thus, we focus on the correction of X -type errors using the classical decoding with the parity-check matrix H_Z . For correction of Z -type errors, one can follow the same approach using classical decoding with the parity-check matrix H_X . BP decoding of quantum LDPC codes was first introduced in [6] and considered further in [7].

In this context, the Tanner graph has variable nodes $V = \{v_1, \dots, v_n\}$ representing elements of the X -error vector $e = (e_1, \dots, e_n)$ and check nodes $C = \{c_1, \dots, c_m\}$ representing the Z -stabilizers in the matrix H_Z . A variable node v_i is connected to a check node c_j if $H_Z(i, j) = 1$, meaning that the corresponding j -th qubit participates in the parity-check defined by the i -th Z -stabilizer. The BP algorithm estimates the most likely values for the error bits by exchanging messages between the variable nodes and the check nodes (representing stabilizers).

Let $m_{v \rightarrow c}^{(t)}$ denote the BP message passed from variable node v to check node c during the t -th iteration and let $m_{c \rightarrow v}^{(t)}$

denote the BP message passed from check node c to variable node v during the t -th iteration. The messages are typically represented by log-likelihood ratios (LLRs). For a channel where X -errors occur with probability p , the variable nodes are initialized to the channel LLR defined by

$$\lambda(p) = \ln \frac{1-p}{p}. \quad (10)$$

This LLR represents the initial belief that a qubit has an X -error.

At the start of the BP algorithm, these initial LLRs are used to set the messages sent from each variable node v to all its connected check nodes c . Specifically, at iteration $t = 0$, the message from variable node v to check node c is set to

$$m_{v \rightarrow c}^{(0)} = \mu_v, \quad (11)$$

where $\mu_v = \lambda(p_v)$ and p_v is the prior probability that qubit v has an X -error.

Check nodes subsequently update their messages based on the incoming messages from all connected variable nodes and the syndrome bits $s = \{s_1, s_2, \dots, s_m\}$, using the equation

$$m_{c \rightarrow v}^{(t)} = (-1)^s \tanh^{-1} \left(\prod_{v' \in \mathcal{N}(c) \setminus v} \tanh \left(m_{v' \rightarrow c}^{(t)} \right) \right), \quad (12)$$

where $\mathcal{N}(c)$ represents the set of variable nodes connected to the check node c , and $m_{v' \rightarrow c}$ is the message from variable node v' to check node c .

Then, the variable nodes update their beliefs using the messages received from check nodes:

$$m_{v \rightarrow c}^{(t+1)} = \mu_v + \sum_{c' \in \mathcal{N}(v) \setminus c} m_{c' \rightarrow v}^{(t)}, \quad (13)$$

where μ_v is the initial LLR for variable node v , and $\mathcal{N}(v)$ is the set of check nodes connected to v . Finally, the variable node output message is computed with

$$m_v^{(t+1)} = \mu_v + \sum_{c' \in \mathcal{N}(v)} m_{c' \rightarrow v}^{(t)}. \quad (14)$$

B. BP Convergence Issues for Quantum LDPC Codes

For stabilizer codes, the commutativity constraint requires that all stabilizers commute. When applied to CSS codes, this implies that the binary representations of all X -stabilizers must be codewords of the binary linear code with parity-check matrix H_Z . In terms of erasure decoding, this means that the variable nodes in the support of an X -stabilizer must be a stopping set for the peeling decoder based on H_Z . Moreover, the subgraph of the Tanner graph induced by a stopping set must contain a cycle. Thus, short cycles in the Tanner graph arise from the inherent degeneracy of CSS LDPC codes due to low-weight stabilizers and this affects BP convergence.

While BP convergence suffers, there is no reason that decoding performance must also suffer. This is because the resulting uncertainty in the X -error pattern is exactly equal to an X -stabilizer. This observation was highlighted in [27] and used to motivate the pruned peeling decoder.

The key idea of pruned peeling is that, if a X -stabilizer completely is covered by erasures, then any bit in its support can be fixed to 0 (i.e., no error) without risk because applying the stabilizer can always flip its value to 1. We note that this operation “uses up” that stabilizer. Pruned peeling makes use of this idea by searching for linear combinations of at most M stabilizer generators which are completely covered by erasures and fixing one of their variables to 0.

C. Belief Propagation with Guided Decimation

Building on BP, the BPGD decoding algorithm operates by sequentially running BP and fixing values of variable nodes based on the resulting BP beliefs. Message-passing algorithms that incorporate “decimation” were first introduced in the classical context for constraint satisfaction problems [38], [39]. For quantum LDPC codes, BPGD decoding seems to mitigate the non-convergence issue of BP due to the commutativity constraint of stabilizer codes. In [28], the BPGD decoding algorithm is shown to be effective for correcting bit-flip errors in quantum CSS codes.

BPGD begins by initializing the LLRs for each variable node on the Tanner graph, similar to the standard BP algorithm. The process proceeds iteratively where, in each round, BP is run for a fixed number of iterations T using the sum-product algorithm to compute the estimated beliefs for the variable nodes.

If BP converges and the resulting error estimate matches the observed syndrome, the decoding process terminates, and the hard values of the variable nodes are returned as the estimated error. However, if convergence is not achieved, the algorithm identifies the variable node with the highest reliability, defined by the magnitude of the LLR, and decimates that bit. Decimation involves fixing the value of this variable node based on the sign of its current bias and marking the bit as decimated, effectively reducing the complexity of subsequent BP iterations.

The decimation step is controlled by a parameter llr_{\max} , which is set to a large value (typically $llr_{\max} = 25$) to ensure numerical stability in practical implementations. This fixed value serves as a strong bias, guiding the remaining variable nodes toward a solution in the subsequent rounds of BP. The process repeats, with each round consisting of BP followed by decimation, until either all variable nodes have been decimated or BP successfully converges.

While BPGD often enhances the probability of convergence compared to standard BP, there is a potential for *non-convergence failure* if the final set of hard-decimated values does not match the syndrome. Nonetheless, BPGD seems to be a valuable approach, particularly in scenarios where traditional BP struggles to find a solution. The effectiveness of BPGD in such cases underscores its utility as a robust decoder for quantum error correction.

D. BPGD Decoding for the Quantum Erasure Channel

Here, we apply the BPGD algorithm to quantum erasure channels by choosing the initial LLRs correctly. Specifically,

Algorithm 1 BPGD over erasures

Input: erasure locations V_e , block length n , H_Z Tanner graph $G = (V, C, E)$, syndrome s , BP iterations per round T

Output: estimated \hat{x} or non-convergence

1 **for** $i = 1$ to n **do**

2 **if** $v_i \in V_e$ **then**

3 $\mu_{v_i} = llr_{\min}$

4 **else**

5 $\mu_{v_i} = llr_{\max}$

6 $m_{v_i \rightarrow c_j}^{(0)} \leftarrow \mu_{v_i}$ for all $c_j \in \mathcal{N}(v_i)$

7 $V_u = V$

8 **for** $r = 1$ to n **do**

9 run BP on Tanner graph G for T iterations

10 $\hat{x} \leftarrow$ hard values of the variable nodes

11 **if** $\hat{x}H_1^T = s$ **then**

12 **return** \hat{x}

13 **else**

14 $v_i = \arg \max_{v \in V_u} \gamma(v_i)$

15 **if** $m_{v_i}^{(rT)} \geq 0$ **then**

16 $\mu_{v_i} = llr_{\max}$

17 **else**

18 $\mu_{v_i} = -llr_{\max}$

19 $V_u = V_u \setminus \{v_i\}$

20 **return** non-convergence

when a bit is erased, the implied error rate is $1/2$ and the corresponding LLR value is $\lambda(1/2) = 0$. One can also use a very small value, $llr_{\min} \approx 0$, to avoid numerical issues.

Conversely, when a qubit is not erased, the decoder is very confident that the qubit is not in error. While the formal LLR value should be $\lambda(0) = \infty$, a large finite value is used for numerical reasons. We define $llr_{\max} = 25$ to be this value.

These modifications align the decoder with the characteristics of erasure channels and result in reasonable decoding performance. The pseudo-code for the BPGD algorithm, along with further technical details, can be found in Algorithm 1.

To compare BPGD with pruned peeling, consider the case where $llr_{\max} = \infty$ and BP iterates to a fixed point. This is equivalent to the peeling decoder and always outputs a stopping set. At this point, the pruned peeling decoder will make a guess only if there is a stabilizer that can absorb the possible error. On the other hand, the BPGD decoder simply guesses bits in the stopping set randomly and continues decoding. While this approach has no guarantees, some of the guesses will be correct and the others can hopefully be absorbed by the stabilizer. Overall, performance is improved over pruned peeling.

IV. BPGD WITH DAMPING AND ADJUSTED LLRs

The performance of BPGD for quantum LDPC codes with erasures can be improved by adjusting the initial LLRs and incorporating damping into the belief propagation process. In particular, we observed experimentally that the convergence rate of BPGD on quantum LDPC codes over erasures can be

TABLE 1
OPTIMIZED VALUES OF c_{opt} FOR SIMULATIONS

[[1600,64]] HGP code						
Erasure Rate	≤ 0.08	0.10	0.12	0.14-0.24	0.26	≥ 0.28
c_{opt}	0.1	0.2	0.1	0.2	0.3	0.5
[[2025,81]] HGP code						
Erasure Rate	≤ 0.12	0.14-0.22	≥ 0.24			
c_{opt}	0.1	0.15	0.3			
B1 code						
Erasure Rate	≤ 0.34	0.36-0.46	0.48	0.50		
c_{opt}	0.3	0.5	0.8	0.9		

TABLE 2
OPTIMIZED VALUES OF γ FOR SIMULATIONS

[[1600,64]] HGP code						
Erasure Rate	≤ 0.06	0.08-0.20	0.22-0.24	0.26	0.28	0.30
γ	0.87	0.88	0.90	0.93	0.94	0.95
[[2025,81]] HGP code						
Erasure Rate	≤ 0.08	0.10-0.16	0.18-0.22	0.24	0.26	0.28
γ	0.86	0.87	0.88	0.90	0.92	0.95
B1 code						
Erasure Rate	0.32					
γ	0.97					
Erasure Rate						
Erasure Rate	≤ 0.36	0.38-0.40	0.42-0.48	0.50		
γ	0.90	0.95	0.94	0.95		

significantly improved by these modifications. This, in turn, reduces the failure rate especially at erasure rates where the rate of non-convergence dominates the rate of logical error.

A. BPGD with Adjusted LLRs

It has been observed that, for BP decoding over both classical [40] and quantum LDPC codes [41], adjusting the initialization of BP on the variable nodes can potentially improve performance. In the context of BPGD decoding of quantum LDPC codes over erasures, we have also observed performance improvement upon adjusting the priors on the variable nodes.

In particular, we found that reducing the initial confidence in the known bits can improve the convergence of the decoder. To achieve this, we introduce a scaling factor c_{opt} and adjust the initial LLRs for non-erased bits to be

$$llr'_{\max} = c_{opt} llr_{\max}. \quad (15)$$

By choosing $c_{opt} < 1$, we reduce the initial LLR magnitude for non-erased bits and this decreases the number of decoder failures due to non-convergence.

To select the optimal value of c_{opt} , we performed experiments over different ranges of values for different erasure rates, fine-tuning to narrower intervals when necessary. We observed that lower erasure rates required lower values of c_{opt} than higher erasure rates for better convergence and lower failure rates. The optimized values of c_{opt} used in our simulations is provided in Table 1.

B. BPGD with Damping

During the BPGD decoding process, BP is applied to compute approximate marginals for the variable nodes that guide the decimation. This approximation is exact if the underlying Tanner graph is a tree [37]. However, the Tanner

TABLE 3
OPTIMIZED VALUES OF c_{opt} WITH γ VALUES FROM TABLE 2

[[2025, 81]] HGP code							
Erasure Rate	0.10	0.12	0.14	0.16	0.18	0.20	0.22
c_{opt}	0.065	0.077	0.077	0.089	0.101	0.11	0.115
Erasure Rate	0.24	0.26	0.28	0.30	0.32		
c_{opt}	0.135	0.14	0.19	0.23	0.25		

graphs of quantum LDPC codes have many short cycles due to low-weight stabilizers and BP may approximate the marginal poorly. If the main problem is that BP is not converging to a fixed point, then one technique that can be used to improve the BP performance is called *damping* [42]–[44].

At each step of the BP iteration in equation (13), the evaluation of $m_{v \rightarrow c}^{(t+1)}$ is taken to be a weighted average between the old estimate and the new estimate:

$$\tilde{m}_{v \rightarrow c}^{(t+1)} = \mu_v + \sum_{c' \in N(v) \setminus c} m_{c' \rightarrow v}^{(t)}. \quad (16)$$

The damped variable-to-check messages are computed with

$$m_{v \rightarrow c}^{(t+1)} = (1 - \gamma) \cdot m_{v \rightarrow c}^{(t)} + \gamma \cdot \tilde{m}_{v \rightarrow c}^{(t+1)}, \quad (17)$$

where $\gamma \in [0, 1]$ is the damping factor.

By choosing an appropriate value for the damping factor γ we can affect the influence of new messages versus old messages. A smaller γ provides more damping to help prevent oscillations and improves convergence of BP especially for lower erasure rates. The optimized values of γ used in our simulations are provided in Table 2.

C. BPGD with Combined Adjustments

For the [[2025, 81]] HGP code, neither adjusting the initial LLR nor the damping factor were enough by themselves for the BPGD performance to match the VH decoder. But, jointly optimizing LLR adjustment and damping factor provided further improvement. First, we performed a grid search over both parameters and found that the damping coefficient γ did not change much for a given erasure rate. Thus, we found suitable parameters by fixing γ for each erasure rate and searching over possible values of c_{opt} . The chosen values minimized the failure rates in our experiments.

Table 3 summarizes the optimized values of c_{opt} used in our simulations for different erasure rates with the γ values being same as in Table 2.

V. NUMERICAL RESULTS

In Figures 1 and 2, we show simulation results for the BPGD decoder on the quantum erasure channel for two HGP QLDPC codes. These figures present the performance of peeling, pruned peeling, and VH decoders for comparison. The data points for the latter decoders were obtained by executing the code available in the publicly accessible GitHub repository referenced by [27]. BPGD decoding performs better than the peeling decoder for both codes but performs worse than pruned peeling for the [[2025, 81]] code. BPGD with adjustments outperforms pruned peeling decisively, achieving proximity

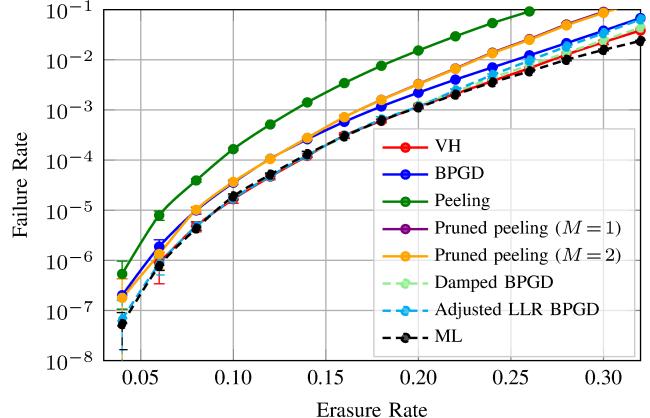


Fig. 1. Comparison the quantum erasure channel of BPGD decoding with peeling, pruned peeling, VH, and ML decoding. The plot presents the decoder failure rates for the HGP QLDPC code [[1600, 64]] from [27], with convergence to a degenerate codeword considered a success. The number of simulations per data point varies and was chosen to ensure short error bars.

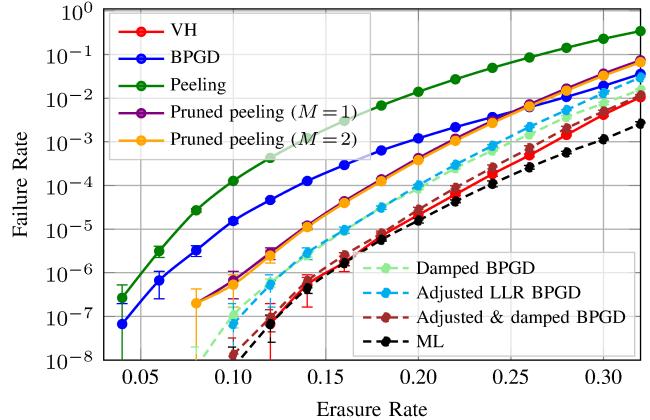


Fig. 2. Comparison on the quantum erasure channel of BPGD decoding with peeling, pruned peeling, VH, and ML decoding. The plot presents the decoder failure rates for the HGP QLDPC code [[2025, 81]] from [27].

to VH decoder which itself is close to the ML decoding performance for most data points.

In Figure 3, we show the simulation results for the [[882, 24, 18 ≤ d ≤ 24]] B1 lifted-product code¹ from [45]. Adjusted versions of BPGD decoding are compared with peeling, pruned peeling, and ML decoding, all on the quantum erasure channel. While the original BPGD decoder does not uniformly outperform pruned peeling due to convergence issues, BPGD with the proposed modifications is much closer to ML decoding than all other decoders. We also want to highlight the sharp threshold-like behavior of both BPGD and ML decoding for the B1 code with erasures. While this type of performance is typical for classical LDPC codes, the performance of quantum LDPC codes usually has the slower decay rate shown, for example, by pruned peeling with $M = 2$.

VI. CONCLUSION

In this paper, we present and analyze the performance of the belief propagation with guided decimation (BPGD) decoder

¹This was called a generalized hypergraph product code in the original paper but the literature now uses the term lifted-product code.

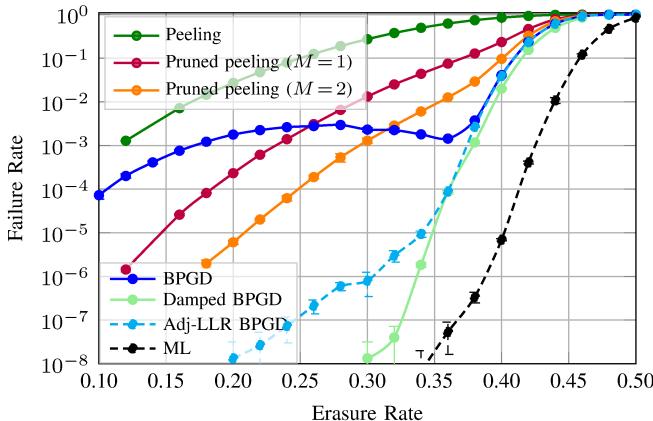


Fig. 3. Comparison of BPGD decoding over the quantum erasure channel with peeling, pruned peeling, and ML decoding. The curves show the decoder failure rates for the $[[882, 24, 18 \leq d \leq 24]]$ B1 lifted-product QLDPC code from [45].

on the quantum erasure channel. Specifically, we compare the performance of BPGD decoding with other established erasure decoders via simulation for a lifted product QLDPC code and two hypergraph product (HGP) QLDPC codes. The impressive performance of BPGD decoding for the $[[882, 24, 18 \leq d \leq 24]]$ B1 QLDPC code demonstrates the potential of BPGD as a generic decoder for QLDPC codes.

For the tested HGP codes, our results show that the BPGD decoder outperforms both peeling and pruned peeling decoders when correcting erasures. Although the BPGD decoder without adjustment is slightly inferior to the Vertical-Horizontal (VH) decoder for these codes, its computational complexity is substantially lower.

Moreover, we observe that small adjustments to the BP algorithm can enhance BPGD decoding performance by encouraging the decoder to converge. In particular, by adjusting the initial LLRs and applying damping techniques, we find that BPGD decoding can approach the performance of the VH decoder.

Finally, we note that quantum codes with erasure conversion are likely to have both erasures and errors. Due to its good performance for channels with erasures and channels with errors, the BPGD decoder is an ideal candidate for channels with both errors and erasures.

ACKNOWLEDGEMENTS

This material is based on work supported by the NSF under Grants 2106213 and 2120757. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

The decoder performance curves from [27] have been reproduced in this work based on a larger simulation study. We also acknowledge the use of generative AI tools for assistance in coding-related tasks during the development of our decoding algorithms.

REFERENCES

- [1] D. Gottesman, “Fault-tolerant quantum computation with constant overhead,” *Quantum Info. Comput.*, vol. 14, no. 15–16, p. 1338–1372, nov 2014.
- [2] O. Fawzi, A. Gaspeller, and A. Leverrier, “Constant overhead quantum fault tolerance with quantum expander codes,” *Communications of the ACM*, vol. 64, no. 1, pp. 106–114, 2020.
- [3] A. Kitaev, “Fault-tolerant quantum computation by anyons,” *Annals of Physics*, vol. 303, no. 1, pp. 2–30, 2003. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0003491602000180>
- [4] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, “Topological quantum memory,” *Journal of Mathematical Physics*, vol. 43, no. 9, pp. 4452–4505, 2002.
- [5] H. Bombin and M. A. Martin-Delgado, “Topological quantum distillation,” *Phys. Rev. Lett.*, vol. 97, p. 180501, Oct 2006. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.97.180501>
- [6] D. J. MacKay, G. Mitchison, and P. L. McFadden, “Sparse-graph codes for quantum error correction,” *IEEE Transactions on Information Theory*, vol. 50, no. 10, pp. 2315–2330, 2004.
- [7] D. Poulin and Y. Chung, “On the iterative decoding of sparse quantum codes,” *Quantum Info. Comput.*, vol. 8, no. 10, p. 987–1000, Nov. 2008.
- [8] P. Panchev 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.
- [9] A. Leverrier and G. Zémor, “Quantum Tanner codes,” in *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2022, pp. 872–883.
- [10] I. Dinur, M.-H. Hsieh, T.-C. Lin, and T. Vidick, “Good quantum LDPC codes with linear time decoders,” in *Proceedings of the 55th annual ACM symposium on theory of computing*, 2023, pp. 905–918.
- [11] 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.
- [12] C. H. Bennett, D. P. DiVincenzo, and J. A. Smolin, “Capacities of quantum erasure channels,” *Physical Review Letters*, vol. 78, no. 16, p. 3217, 1997.
- [13] Y. Wu, S. Kolkowitz, S. Puri, and J. D. Thompson, “Erasure conversion for fault-tolerant quantum computing in alkaline earth Rydberg atom arrays,” *Nature communications*, vol. 13, no. 1, p. 4657, 2022.
- [14] K. Sahay, J. Jin, J. Claes, J. D. Thompson, and S. Puri, “High-threshold codes for neutral-atom qubits with biased erasure errors,” *Physical Review X*, vol. 13, no. 4, p. 041013, 2023.
- [15] S. Ma, G. Liu, P. Peng, B. Zhang, S. Jandura, J. Claes, A. P. Burgers, G. Pupillo, S. Puri, and J. D. Thompson, “High-fidelity gates and mid-circuit erasure conversion in an atomic qubit,” *Nature*, vol. 622, no. 7982, pp. 279–284, 2023.
- [16] M. Kang, W. C. Campbell, and K. R. Brown, “Quantum error correction with metastable states of trapped ions using erasure conversion,” *PRX Quantum*, vol. 4, no. 2, p. 020358, 2023.
- [17] A. Kubica, A. Haim, Y. Vakinin, H. Levine, F. Brandão, and A. Retzker, “Erasure qubits: Overcoming the T_1 limit in superconducting circuits,” *Physical Review X*, vol. 13, no. 4, p. 041022, 2023.
- [18] J. D. Teoh, P. Winkel, H. K. Babla, B. J. Chapman, J. Claes, S. J. de Graaf, J. W. Garmon, W. D. Kalfus, Y. Lu, A. Maiti *et al.*, “Dual-rail encoding with superconducting cavities,” *Proceedings of the National Academy of Sciences*, vol. 120, no. 41, p. e2221736120, 2023.
- [19] M. Grassl, T. Beth, and T. Pellizzari, “Codes for the quantum erasure channel,” *Physical Review A*, vol. 56, no. 1, p. 33, 1997.
- [20] T. M. Stace, S. D. Barrett, and A. C. Doherty, “Thresholds for topological codes in the presence of loss,” *Physical review letters*, vol. 102, no. 20, p. 200501, 2009.
- [21] N. Delfosse and N. H. Nickerson, “Almost-linear time decoding algorithm for topological codes,” *Quantum*, vol. 5, p. 595, 2021.
- [22] N. Delfosse and G. Zémor, “Linear-time maximum likelihood decoding of surface codes over the quantum erasure channel,” *Physical Review Research*, vol. 2, no. 3, p. 033042, 2020.
- [23] N. Delfosse, V. Londe, and M. E. Beverland, “Toward a union-find decoder for quantum LDPC codes,” *IEEE Transactions on Information Theory*, vol. 68, no. 5, pp. 3187–3199, 2022.
- [24] S. Lee, M. Mhalla, and V. Savin, “Trimming decoding of color codes over the quantum erasure channel,” in *2020 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2020, pp. 1886–1890.

[25] H. M. Solanki and P. K. Sarvepalli, "Correcting erasures with topological subsystem color codes," in *2020 IEEE Information Theory Workshop (ITW)*. IEEE, 2021, pp. 1–5.

[26] ———, "Decoding topological subsystem color codes over the erasure channel using gauge fixing," *IEEE Transactions on Communications*, vol. 71, no. 7, pp. 4181–4192, 2023.

[27] N. Connolly, V. Londe, A. Leverrier, and N. Delfosse, "Fast erasure decoder for hypergraph product codes," *Quantum*, vol. 8, p. 1450, 2024.

[28] H. Yao, W. A. Laban, C. Häger, H. D. Pfister *et al.*, "Belief propagation decoding of quantum LDPC codes with guided decimation," *arXiv preprint arXiv:2312.10950*, 2023.

[29] A. Gong, S. Cammerer, and J. M. Renes, "Toward low-latency iterative decoding of QLDPC codes under circuit-level noise," 2024. [Online]. Available: <https://arxiv.org/abs/2403.18901>

[30] C. Di, D. Proietti, I. Telatar, T. Richardson, and R. Urbanke, "Finite-length analysis of low-density parity-check codes on the binary erasure channel," *IEEE Transactions on Information Theory*, vol. 48, no. 6, pp. 1570–1579, 2002.

[31] M. Luby, M. Mitzenmacher, M. Shokrollahi, and D. Spielman, "Efficient erasure correcting codes," *IEEE Transactions on Information Theory*, vol. 47, no. 2, pp. 569–584, 2001.

[32] H. Pfister, I. Sason, and R. Urbanke, "Capacity-achieving ensembles for the binary erasure channel with bounded complexity," in *2004 23rd IEEE Convention of Electrical and Electronics Engineers in Israel*, 2004, pp. 110–113.

[33] R. Gallager, "Low-density parity-check codes," *IRE Transactions on Information Theory*, vol. 8, no. 1, pp. 21–28, 1962.

[34] D. Spielman, "Linear-time encodable and decodable error-correcting codes," *IEEE Transactions on Information Theory*, vol. 42, no. 6, pp. 1723–1731, 1996.

[35] Q. Xu, J. P. Bonilla Ataides, C. A. Pattison, N. Raveendran, D. Bluvstein, J. Wurtz, B. Vasić, M. D. Lukin, L. Jiang, and H. Zhou, "Constant-overhead fault-tolerant quantum computation with reconfigurable atom arrays," *Nature Physics*, vol. 20, no. 7, pp. 1084–1090, Jul 2024. [Online]. Available: <https://doi.org/10.1038/s41567-024-02479-z>

[36] P. Panteleev and G. Kalachev, "Quantum LDPC codes with almost linear minimum distance," *IEEE Transactions on Information Theory*, vol. 68, no. 1, pp. 213–229, 2022.

[37] F. Kschischang, B. Frey, and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Transactions on Information Theory*, vol. 47, no. 2, pp. 498–519, 2001.

[38] M. Mézard, G. Parisi, and R. Zecchina, "Analytic and algorithmic solution of random satisfiability problems," *Science*, vol. 297, no. 5582, pp. 812–815, 2002. [Online]. Available: <https://www.science.org/doi/abs/10.1126/science.1073287>

[39] A. Montanari, F. Ricci-Tersenghi, and G. Semerjian, "Solving constraint satisfaction problems through belief propagation-guided decimation," *arXiv preprint arXiv:0709.1667*, 2007.

[40] M. Hagiwara, M. P. Fossorier, and H. Imai, "Fixed initialization decoding of LDPC codes over a binary symmetric channel," *IEEE transactions on information theory*, vol. 58, no. 4, pp. 2321–2329, 2012.

[41] S. Miao, A. Schnerring, H. Li, and L. Schmalen, "Quaternary neural belief propagation decoding of quantum LDPC codes with overcomplete check matrices," *arXiv preprint arXiv:2308.08208*, 2023.

[42] M. Pretti, "A message-passing algorithm with damping," *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2005, no. 11, p. P11008, 2005.

[43] E. Nachmani, E. Marciano, L. Lugosch, W. J. Gross, D. Burshtein, and Y. Be'ery, "Deep learning methods for improved decoding of linear codes," *IEEE Journal of Selected Topics in Signal Processing*, vol. 12, no. 1, pp. 119–131, 2018.

[44] M. Lian, F. Carpi, C. Häger, and H. D. Pfister, "Learned belief-propagation decoding with simple scaling and SNR adaptation," in *2019 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2019, pp. 161–165.

[45] P. Panteleev and G. Kalachev, "Degenerate Quantum LDPC Codes With Good Finite Length Performance," *Quantum*, vol. 5, p. 585, Nov. 2021. [Online]. Available: <https://doi.org/10.22331/q-2021-11-22-585>