Non-Uniform Windowed Decoding For Multi-Dimensional Spatially-Coupled LDPC Codes

Lev Tauz*, Homa Esfahanizadeh†, and Lara Dolecek*
*Department of Electrical and Computer Engineering, University of California, Los Angeles, USA
†Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, USA
levtauz@ucla.edu, homaesf@mit.edu, and dolecek@ee.ucla.edu

***Abstract*—In this work, we propose a non-uniform windowed decoder for multi-dimensional spatially-coupled LDPC (MD-SC-LDPC) codes over the binary erasure channel. An MD-SC-LDPC code is constructed by connecting together several SC-LDPC codes into one larger code that provides major benefits over a variety of channel models. We propose and analyze a novel non-uniform decoder that allows for greater flexibility between latency and code reliability. Our theoretical derivations and empirical results show that our non-uniform decoder greatly improves upon the standard windowed decoder in terms of design flexibility, latency, and complexity.**

I. Introduction and Motivation

Spatially-coupled LDPC (SC-LDPC) codes are a popular choice for error-correcting codes due to their capacity-achieving performance [1] and low-latency windowed decoding [2]. Multi-Dimensional SC-LDPC (MD-SC-LDPC) codes are a class of LDPC codes [3]–[5] created by connecting several SC-LDPC codes. This class of codes has many significant benefits compared to conventional SC codes, including lower population of detrimental objects for belief propagation (BP) decoders [3], improved reliability over parallel channels [5], and robustness to burst erasures [4]. This makes MD-SC-LDPC codes a viable candidate in coding for multi-dimensional non-volatile memories. In this work, we demonstrate that another benefit of MD-SC-LDPC codes is that they can be decoded using non-uniform windowed decoders allowing for flexible decoder design.

Most MD-SC-LDPC code constructions preserve the unique chain structure of an SC-LDPC code which naturally allows for windowed decoding [2]. One way of applying windowed decoding is for each constituent SC code to have its own window and for the windows to move in unison along the coupled constituent SC chains. This approach provides the same proportional latency benefits as it does for a single SC code, relative to the block length. Conventionally, to improve latency, one can only reduce the window size uniformly across all the constituent codes at the cost of lower reliability. However, this approach does not take into account the structure of MD-SC-LDPC codes and, as a result, causes unnecessary reliability loss. By allowing non-uniform window sizes across constituent codes, we exploit the structure of MD-SC codes to provide more decoder design flexibility and better code utilization.

In this work, we define a code ensemble that captures the multi-dimensional (MD) coupling structure which can be exploited for designing a flexible decoder. Next, we propose a novel non-uniform windowed decoder that takes into account the unique structure of MD-SC codes. Finally, we demonstrate through simulations the improvements offered by our new decoder.

The main contribution of this work is presenting our non-uniform windowed decoder and demonstrating how it can provide latency and reliability trade-offs that would not be feasible for 1D SC-LDPC codes. For example, using a non-uniform windowed decoder can provide more than an order of magnitude improvement in Bit Error Rate (BER) as opposed to a uniform windowed decoder with the same latency constraints. Conversely, we also show that the non-uniform windowed decoder can greatly improve the latency for the same BER constraints.

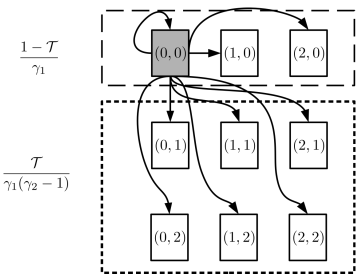

Fig. 1. Coupling from the VN perspective of section $(0, 0)$ with $\gamma_1 = \gamma_2 = 3$. The probability of connecting to a section in the top box is $(1 - \mathcal{T})/\gamma_1$ and of connecting to a section in the bottom box is $\mathcal{T}/\gamma_1(\gamma_2 - 1)$.

Due to limited space, the following sections are shortened from our long version [6]. We refer readers to the long version for detailed background and derivation of equations.

Notation: For positive integers A and B, we define the set $[A] \triangleq \{0, 1, \dots, A-1\}$ and the operation $(n)_B = n \mod B$.

II. Code Construction and Properties

In this section, we define an MD-SC code ensemble $\mathcal{C}_{\text{MD}} = \mathcal{C}(d_l, d_r, L_1, \gamma_1, L_2, \gamma_2, \mathcal{T})$. The parameters d_l and d_r denote the degrees of the variable nodes (VNs) and check nodes (CNs), respectively. We denote L_1 as the *1-dimensional (1D) coupling length* and L_2 as the *MD coupling length*. Additionally, we define $1 \le \gamma_1 \le L_1$ as the *1D coupling depth* and $1 \le \gamma_2 \le L_2$ as the *MD coupling depth*, which specify the coupling distance along a dimension. We refer to $0 \le \mathcal{T} \le 1$ as the *density* of the edges for the coupling along the second dimension.

Now, we define the MD-SC code ensemble. The code contains an $L_1 \times L_2$ grid of sections where each section contains M VNs and $M(d_l/d_r)$ CNs. Each of the d_l edges of a VN in section (i, j) is uniformly connected to a CN from sections $\{(i + k, j) : \ k \in [\gamma_1]\}$ with probability $1 - \mathcal{T}$ or sections $\{(i + k, (j + r)_{L_2}) : \ k \in [\gamma_1], r \in [\gamma_2] \setminus \{0\}\}$ with probability $\mathcal{T}$. An example of this coupling can be seen in Fig. 1. We study the performance of our ensemble for the binary erasure channel with erasure probability ϵ (BEC(ϵ)) under BP decoding which can be analyzed through density evolution (DE) techniques. The full density evolution equations can be found in [6]. Given a target erasure probability δ, we define the BP threshold as ϵ^*_δ such that for $\epsilon \le \epsilon^*_\delta$ BP is able to decode all bits to at most a target erasure probability δ after an infinite number of iterations.

In [6], we demonstrate that this ensemble has the same design rate and BP threshold as the standard 1D-SC code ensemble [1]. We also show that for finite lengths, our MD-SC ensemble is more likely to create small problematic objects for decoding due to the extra sparsity but this can be greatly mitigated with careful coupling as shown in [3]. In the next section, we demonstrate our new non-uniform windowed decoder which utilizes the structure of MD-SC codes.

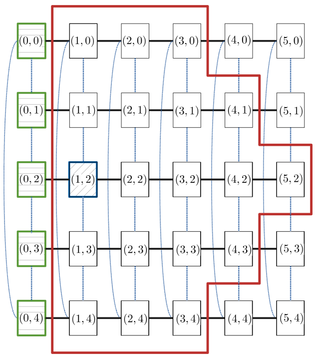

Fig. 2. Example of a non-uniform window with TVNs (blue rectangle) in section $(1, 2)$ and window shape $\mathcal{W} = [5, 4, 3, 3, 4]$. The green rectangles are the decoded VNs. Due to the TVNs being in segment 2, the window is shifted so that window size W_0 is used for segment 2.

TABLE I

THRESHOLDS FOR WINDOW SIZES WITH THE LARGEST $\epsilon^{\mathrm{WC}}_{\delta,\mathcal{W}}$ FOR $d_l = 4$, $d_r = 8$, $\gamma_1 = 2$, AND $\delta = 10^{-12}$. WINDOW SIZES ARE CONSTRAINED BETWEEN 2 AND 7, FOR $L_2 = 7$, AND BETWEEN 2 AND 5, FOR $L_2 = 9$.

L_2	γ_2	$\mathcal{T}$	C	$\mathcal{W}$	$\epsilon^{\mathrm{WC}}_{\delta,\mathcal{W}}$
7	2	0.05	28	$(5,5,4,2,3,4,5)$	≈ 0.4829
7	2	0.1	28	$(5,5,4,3,3,4,4)$	≈ 0.4722
7	3	0.05	28	$(5,4,4,3,4,4,4)$	≈ 0.4723
7	3	0.1	28	$(4,4,4,4,4,4,4)$	≈ 0.4685
9	2	0.05	36	$(5,5,4,3,2,3,4,5,5)$	≈ 0.4872
9	2	0.1	36	$(5,5,4,3,2,3,4,5,5)$	≈ 0.4806
9	3	0.05	36	$(5,5,5,2,3,3,4,4,5)$	≈ 0.4767
9	3	0.1	36	$(4,4,4,4,4,4,4,4,4)$	≈ 0.4685

III. Non-Uniform Windowed Decoding

In this section, we describe an abbreviated version of our non-uniform windowed decoder.

We define a subset of VNs for which BP will be performed over as a *window configuration (WC)*. Every WC has a unique section of VNs that are aimed to be decoded, called the *targeted VNs (TVNs)*. The TVNs of each WC are VNs of a single section (i, j) of the code. We denote $\mathcal{W} = [W_0, W_1, \ldots, W_{L_2-1}]$ to be the *vector of window sizes* of the WCs. Given (i, j), we define $S^{\mathcal{W}}_{(i,j)} = \{(i + k, j + r) : r \in [L_2], k \in [W_{(j+r)_{L_2}}]\}$ as the WC over which BP will be performed over. For any specific WC, the window sizes are cyclically shifted so that W_0 is centered on the TVNs. We define $\epsilon^{\mathrm{WC}}_{\delta,\mathcal{W}}$ as the BP threshold such that for $\epsilon \le \epsilon^{\mathrm{WC}}_{\delta,\mathcal{W}}$ the non-uniform windowed decoder is able to decode all TVNs to at most a target erasure probability δ. An example of a WC is presented in Fig. 2.

We briefly analyze the latency of the decoder. Let $s(\mathcal{W}) = \sum_{i=0}^{L_2-1} W_i$. Note that the number of VNs that need to be accessed to process a WC is at most $\mathcal{O}(s(\mathcal{W}))$ which upper bounds the latency. For a block BP decoder, the latency is $\mathcal{O}(L_1 L_2)$ since every VN has to be accessed. Thus, the latency is reduced by at least a factor of $\mathcal{O}(\frac{s(\mathcal{W})}{L_1 L_2})$. As such, we denote the latency constraint as C such that $s(\mathcal{W}) \le C$.

While there a few design choices for the non-uniform windowed decoder, we find that the most important one is identifying the best choice of $\mathcal{W}$ given the latency constraint C. In [6], we show that the best $\mathcal{W}$ that maximizes $\epsilon^{\mathrm{WC}}_{\delta,\mathcal{W}}$ while satisfying $s(\mathcal{W}) \le C$ are $\mathcal{W}$ where $s(\mathcal{W}) = C$ which significantly reduces the search space.

IV. Simulations

In this section, we demonstrate through simulations the flexibility and improvements offered by our non-uniform windowed decoder. We also provide new finite-length simulations that were not published in the full paper [6].

In Table I, we show the thresholds for different window sizes that were chosen to maximize $\epsilon^{\mathrm{WC}}_{\delta,\mathcal{W}}$ for various code parameters. We note that the uniform windowed decoder has the same threshold regardless of γ_2 and $\mathcal{T}$, which is expected since the DE equations degenerate to the 1D-SC DE equations when the window sizes are uniform. We see that for small values of γ_2 and $\mathcal{T}$, non-uniform window shapes have the highest thresholds. This demonstrates how non-uniform window shapes better utilize the code to improve the performance.

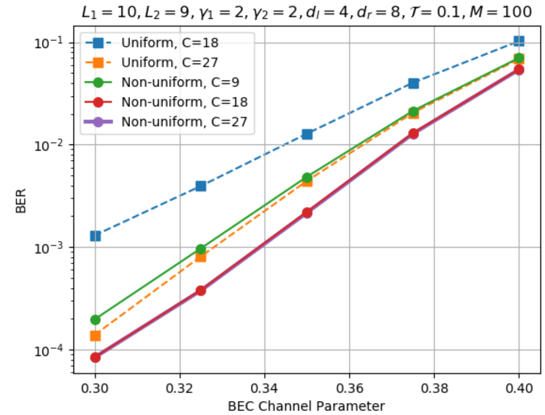

Fig. 3. Bit Error Rates for MD-SC codes with length 9000. Note that the red and purple curves overlap.

In Fig. 3, we present the BER for BP decoding of finite length realizations of the MD-SC code ensemble given different window shapes and latency constraints. The BER is averaged across many finite length realizations of the ensemble. The non-uniform window shapes were chosen to maximize the $\epsilon^{\mathrm{WC}}_{\delta,\mathcal{W}}$. We observe that a non-uniform window can provide over an order of magnitude improvement in BER in comparison to a uniform window for the same latency constraint. In addition, we note that given the same BER, the non-uniform window can achieve it with a factor of 4 reduction in latency. Finally, we observe that the performance saturates given a large enough window size but that saturation occurs faster using non-uniform window shapes.

V. Conclusion and Future Work

In this work, we defined a new variant of MD-SC-LDPC codes which offers more flexibility in designing windowed decoding. We proposed a novel windowed decoder using non-uniform window sizes which better exploits the structure of MD coupling. Future research is focused on analyzing how the threshold is affected by changes in the window shape.

VI. Acknowledgments

Research supported in part by a grant from ASRC-IDEMA and grant CCF-BSF:CIF #1718389 from NSF.

References

[1] S. Kudekar, T. J. Richardson, and R. L. Urbanke, "Threshold saturation via spatial coupling: Why convolutional LDPC ensembles perform so well over the BEC," *IEEE Transactions on Information Theory*, Feb. 2011.

[2] A. R. Iyengar, P. H. Siegel, R. L. Urbanke, and J. K. Wolf, "Windowed decoding of spatially coupled codes," *IEEE Transactions on Information Theory*, Apr. 2013.

[3] H. Esfahanizadeh, L. Tauz, and L. Dolecek, "Multi-dimensional spatially-coupled code design: Enhancing the cycle properties," *IEEE Transactions on Communications*, May 2020.

[4] R. Ohashi, K. Kasai, and K. Takeuchi, "Multi-dimensional spatially-coupled codes," in *Proc. IEEE International Symposium on Information Theory (ISIT)*, Jul. 2013.

[5] L. Schmalen and K. Mahdaviani, "Laterally connected spatially coupled code chains for transmission over unstable parallel channels," in *Proc. International Symposium on Turbo Codes & Iterative Information Processing (ISTC)*, Aug. 2014.

[6] L. Tauz, H. Esfahanizadeh, and L. Dolecek, "Non-uniform windowed decoding for multi-dimensional spatially-coupled ldpc codes," in *Proc. IEEE International Symposium on Information Theory (ISIT)*, June 2020.