

On Decoding of Reed-Muller Codes Using a Local Graph Search

Mikhail Kamenev

Abstract—We present a novel iterative decoding algorithm for Reed-Muller (RM) codes, which takes advantage of a graph representation of the code. Vertices of the considered graph correspond to codewords, with two vertices being connected by an edge if and only if the Hamming distance between the corresponding codewords equals the minimum distance of the code. The algorithm uses a greedy local search to find a node optimizing a metric, e.g. the correlation between the received vector and the corresponding codeword. In addition, the cyclic redundancy check can be used to terminate the search as soon as a valid codeword is found, leading to an improvement in the average computational complexity of the algorithm. Simulation results for both binary symmetric channel and additive white Gaussian noise channel show that the presented decoder approaches the performance of maximum likelihood decoding for RM codes of length less than 1024 and for the second-order RM codes of length less than 4096. Moreover, it is demonstrated that the considered decoding approach outperforms state-of-the-art decoding algorithms of RM codes with similar computational complexity for a wide range of block lengths and rates.

Index Terms—Reed-Muller codes, BSC channels, AWGN channels, near maximum-likelihood decoding, Fast Hadamard Transform.

I. INTRODUCTION

BINARY Reed-Muller (RM) codes were firstly discovered by Muller [2] and then by Reed, who also proposed a majority decoding algorithm for this family of error correction codes [3]. Although it has been proven that RM codes achieve the capacity on an erasure channel under maximum a posteriori (MAP) decoding [4], no MAP decoding algorithm that can be efficiently used for decoding of RM codes is known.

However, there are several decoding algorithms of RM codes that allow achieving the performance of a maximum likelihood (ML) decoder for a wide range of block lengths and rates. For instance, a recursive permutation list decoder with a list size less than or equal to 128 allows achieving near-ML decoding performance for small length (≤ 256) codes [5]. Since this algorithm processes a received codeword sequentially and uses a sorting operation, the latency of a hardware implementation of the recursive permutation list decoder is high. Permutation-based techniques improving the latency of the recursive permutation list decoder have been proposed in [6]–[8]. Unfortunately, the complexity of recursive algorithms required for near-ML decoding performance grows exponentially with the code length [9].

Another algorithm that performs close to ML decoding is a recursive projection-aggregation (RPA) decoder proposed in

[10]. This algorithm is based on projecting the code on its cosets, recursively decoding the projected codes, and aggregating the reconstructions. The decoder demonstrates near-ML decoding performance for RM codes of length less than 512 and for the second-order RM codes. Moreover, it allows for parallel implementation. Since the code is projected on all its cosets, the RPA algorithm takes a very long time to decode. Techniques aiming to improve the running time of projection-based decoding have been considered in [11]–[13].

Several other approaches for decoding of RM codes have been proposed recently. However, they demonstrate near-ML error correction performance for short length codes [14]–[19] or applicable only for a binary erasure channel [20], [21].

In this paper, we propose a new iterative decoding algorithm for RM codes. Since recursive decoding needs a large complexity to approach the ML decoding performance for RM codes of length larger than 256 [9], the proposed algorithm employs a low-complex recursive decoder [22] just to get the initial candidate codeword and then iteratively improves it using a local graph search. The nodes of the considered graph correspond to codewords. Any two nodes of the graph are connected by an edge if and only if the Hamming distance between two corresponding codewords equals the minimum distance of the code. Simulation results demonstrate that the proposed algorithm approaches the performance of the ML decoder for moderate length (≤ 512) RM codes and for the second-order RM codes of length less than 4096, with average computational complexity being reasonable. We also observe that our algorithm outperforms both recursive permutation list decoding and RPA decoding with similar computational complexity in most considered cases.

The rest of the paper is organized as follows. In Section II, we shortly introduce RM codes. In Section III, we provide a high-level description of our decoding algorithm. In Section IV, we present a greedy version of the algorithm that allows to significantly improve the running time. Numerical results are presented in Section V. We conclude the paper in Section VI.

II. RM CODES

Denote by $f(\mathbf{v}) = f(v_1, \dots, v_m)$ a Boolean function of m variables that is written in the algebraic normal form. Let \mathbf{f} be the vector of length 2^m containing values of f at all of its 2^m arguments. The binary RM code $\mathcal{R}(r, m)$ of length n , $n = 2^m$, and order r , $0 \leq r \leq m$, is the set of all vectors \mathbf{f} , where $f(\mathbf{v})$ is a Boolean function of degree at most r . Note that the minimum distance of the RM code $\mathcal{R}(r, m)$ equals 2^{m-r} [23, Sec. 13.3].

This paper has been presented in part at the 2020 IEEE Information Theory Workshop [1].

M. Kamenev is with the Moscow Research Center, Huawei Technologies Co., Ltd., Moscow, Russia. Email: kamenev.mikhail1@huawei.com

RM codes can also be considered in terms of finite geometries. Denote by $\text{EG}(m, 2)$ the Euclidean geometry of dimension m over $\text{GF}(2)$. By definition, $\text{EG}(m, 2)$ contains 2^m points, whose coordinates are all binary vectors \mathbf{v} of length m . We associate every subset \mathcal{S} of points of $\text{EG}(m, 2)$ with a binary incidence vector of length 2^m that contains one in positions $s \in \mathcal{S}$ and zeros elsewhere. Then the codewords of $\mathcal{R}(r, m)$ can be considered as incidence vectors of subsets of $\text{EG}(m, 2)$. For instance, the minimum weight codewords of $\mathcal{R}(r, m)$ are exactly the incidence vectors of $(m - r)$ -dimensional subspaces of $\text{EG}(m, 2)$ [23, Sec. 13.4].

III. GRAPH SEARCH BASED DECODING OF RM CODES

Consider an RM code $\mathcal{R}(r, m)$ of dimension k and a graph \mathcal{G} with 2^k nodes. Assign to each node in the graph \mathcal{G} a codeword of $\mathcal{R}(r, m)$. Then two nodes are connected by an edge if and only if the Hamming distance between corresponding codewords equals 2^{m-r} , i.e. the minimum distance of the code. Assume that $\mathbf{y} = (y_1, y_2, \dots, y_n)$ is a log-likelihood ratio (LLR) vector of a binary-input additive white Gaussian noise (BI-AWGN) channel. Then, for each node in the graph \mathcal{G} , we can assign the correlation metric M [24, Sec. 10.1], defined by

$$M = \sum_{i=1}^n (1 - 2c_i) y_i,$$

where (c_1, c_2, \dots, c_n) is a codeword assigned to the node. Note that the ML decoder returns a codeword with the largest correlation metric M [24, Sec. 10.1].

Using a graph traversal algorithm (breadth-first search or depth-first search), one can start at a random node, visit all reachable nodes, and return the codeword that corresponds to the node with the largest metric.

Proposition 1: The algorithm described above is equivalent to the ML decoder.

Proof: Observe that the algorithm achieves the performance of the ML decoder if all nodes in the graph \mathcal{G} can be visited. It is possible if and only if the graph \mathcal{G} is connected. Since the minimum weight codewords generate the code [23, Sec. 13.6], it follows that there is a path between the node associated with the all-zero codeword and any other node of the graph \mathcal{G} . Therefore, \mathcal{G} is a connected graph. ■

Unfortunately, the computational complexity of the algorithm grows exponentially with the code dimension that makes it infeasible for practical usage. To decrease the complexity of the search algorithm, we propose the following greedy approach. The algorithm starts at a node associated with an output of low-complex recursive decoding [22]. Then it moves to the adjacent node that has the biggest metric value M and has not been visited before. The algorithm continues till a maximum number of iterations N is reached or all adjacent nodes are already visited. The output of the algorithm is the codeword corresponding to the visited node with the biggest metric.

The maximum number of iterations is used here to limit the computational complexity of the algorithm and the amount of memory used to store visited nodes. In addition, if all adjacent nodes are visited, then the algorithm can be finished until the

maximum number of iterations is reached. Another option that can be used to decrease the average computational complexity is the termination of the algorithm if a codeword that has the best metric found so far satisfies the cyclic redundancy check (CRC). The formal description of the proposed decoding algorithm is shown in Algorithm 1.

Example 1: Consider the RM code $\mathcal{R}(2, 3)$ and an LLR vector

$$\mathbf{y} = (2.76, 5.68, -6.58, 4.42, -0.09, 3.9, 3.56, -1.91). \quad (1)$$

Let us decode this vector using the graph search algorithm with 3 iterations. Assume that the recursive decoder returns a vector $\mathbf{c} = (0, 0, 1, 1, 1, 1, 0, 0)$ with a metric of 8.44. In the first iteration, the algorithm finds a vector $\mathbf{c}_1 = (0, 0, 1, 0, 1, 0, 0, 0)$ that has the largest correlation ($M = 25.08$) among the codewords at the minimum distance from \mathbf{c} . In the next iteration, the algorithm performs a search over the codewords at the minimum distance from \mathbf{c}_1 . The result is the vector $\mathbf{c}_2 = (0, 0, 1, 0, 0, 0, 0, 1)$ with a metric of 28.72. In the last iteration, a vector with the largest metric is \mathbf{c}_1 . Note that this vector has been found before. Therefore, the algorithm chooses the vector $\mathbf{c}_3 = (1, 0, 1, 0, 1, 0, 0, 1)$ with the second largest correlation that is equal to 23.38. The output of the algorithm is a codeword with the largest correlation from the set $\{\mathbf{c}, \mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3\}$. Thus, the output of the algorithm is the codeword \mathbf{c}_2 . It can be readily verified that \mathbf{c}_2 is the ML solution for the vector \mathbf{y} .

Note that an algorithm similar to Algorithm 1 has been proposed in [25] for decoding arbitrary linear codes. This algorithm starts searching with a codeword returned by a low-complex decoding algorithm and uses decoding over a low-weight trellis instead of the `NextStep` function. Although decoding over a low-weight trellis has lower complexity compared to trellis-based ML decoding, it is computationally expensive to run this search procedure many times. As we will show in the next section, there exists a version of the `NextStep` function implementing a low-complex greedy search algorithm for RM codes. It allows running Algorithm 1 with a large N , with the average computational complexity being reasonable.

IV. A SIMPLIFIED SELECTION OF AN ADJACENT NODE

The bottleneck of Algorithm 1 is the computational complexity of the function `NextStep`. A naive implementation of this function goes through all codewords \mathbf{c} of a code $\mathcal{R}(r, m)$ such that the Hamming distance between \mathbf{c} and the current codeword equals 2^{m-r} , i.e. the minimum distance of the code. The number of minimum weight codewords in $\mathcal{R}(r, m)$ is defined as [23, Sec. 13.4]

$$A_{2^{m-r}} = 2^r \prod_{i=0}^{m-r-1} \frac{2^{m-i} - 1}{2^{m-r-i} - 1}.$$

For instance, the number of minimum weight codewords of the code $\mathcal{R}(4, 9)$ is approximately $53 \cdot 10^6$. Thus, a brute-force algorithm that goes through all adjacent nodes takes a long time to find the next codeword. To solve this issue, we propose

Algorithm 1: The GraphSearch decoding function

Input: RM code parameters r and m , a vector of LLRs \mathbf{y} , a maximum number of iterations N , a flag t whether to use the CRC to terminate the algorithm

Output: A codeword \mathbf{res}

```

1 Set  $\mathcal{C}$  to be an empty set of codewords
2 Let  $\mathbf{c}$  be the result of recursive decoding of the vector  $\mathbf{y}$ 
3 Add  $\mathbf{c}$  to  $\mathcal{C}$ 
4  $\mathbf{res} \leftarrow \mathbf{c}$ 
5  $M \leftarrow 0$ 
6 for  $i = 1, 2, \dots, 2^m$  do
7    $M \leftarrow M + (1 - 2c[i]) y[i]$  // Compute the
   metric for the codeword  $\mathbf{c}$ .
8 end
9 for  $i = 1, 2, \dots, N$  do
10   $\mathbf{c}', M' \leftarrow \text{NextStep}(r, m, \mathbf{y}, \mathbf{c}, \mathcal{C})$  // NextStep
   goes through all codewords  $\mathbf{x}$ ,  $\mathbf{x} \notin \mathcal{C}$ ,
   such that the Hamming distance between
    $\mathbf{x}$  and  $\mathbf{c}$  equals  $2^{m-r}$ , and returns the
   codeword  $\mathbf{c}'$  that has the biggest
   metric.  $M'$  is the metric of the
   output codeword. If all  $\mathbf{x} \in \mathcal{C}$ , then
    $M' = -\infty$ .
11  if  $M' = -\infty$  then
12    break // All adjacent nodes are
   already visited.
13  end
14   $\mathbf{c} \leftarrow \mathbf{c}'$  // Move to an adjacent node.
15  Add  $\mathbf{c}$  to  $\mathcal{C}$ 
16  if  $M' > M$  then
   // Save the current codeword if it
   is better than one found so far.
17     $\mathbf{res} \leftarrow \mathbf{c}$ 
18     $M \leftarrow M'$ 
19    if  $t$  and  $\mathbf{res}$  satisfies the CRC then
20      break
21    end
22  end
23 end
24 return  $\mathbf{res}$ 

```

a greedy search algorithm that has the best-case running time of $\mathcal{O}(n \log(\max\{N, n\}))$.

Consider an RM code $\mathcal{R}(r, m)$ of length n and a rooted tree \mathcal{T} constructed in the following way. The root of the tree is assigned $\text{EG}(m, 2)$. Each node at depth i , $i \in \{0, \dots, r-1\}$, has $2^{m-i+1} - 2$ children containing all possible $(m-i-1)$ -dimensional subspaces of the parent node [23, Appx. B]. By the construction of the tree \mathcal{T} , the leaf nodes contain all $(m-r)$ -dimensional subspaces of $\text{EG}(m, 2)$. Recall that the union of these subspaces contains all minimum weight codewords of $\mathcal{R}(r, m)$ [23, Sec. 13.4]. Thus, if we replace the points of subspaces with the corresponding coordinate positions in the codeword, then the leaves of the tree \mathcal{T} will store non-zero positions of all minimum weight codewords in $\mathcal{R}(r, m)$. The tree \mathcal{T} for the RM code $\mathcal{R}(2, 3)$ is illustrated in Fig. 1. Observe that different leaf nodes can contain non-zero positions of the same codeword.

The tree \mathcal{T} can be easily constructed using codewords of the first order RM codes. The root of the tree is assigned a set $[n] \triangleq \{1, 2, \dots, n\}$. Let $\mathbf{c} = (c_1, c_2, \dots, c_n)$ be a codeword of weight $n/2$ in $\mathcal{R}(1, m)$ and let $\mathcal{S}(\mathbf{c}) = \{i \in [n] \mid c_i = 1\}$ be




Fig. 1. A tree that is used to enumerate all minimum weight codewords of the RM code $\mathcal{R}(2, 3)$

a set of non-zero positions in \mathbf{c} . Since the codewords of weight $n/2$ in $\mathcal{R}(1, m)$ are incidence vectors of $(m-1)$ -dimensional subspaces [23, Sec. 13.4], the children nodes of the root contain sets $\mathcal{S}(\mathbf{c}_i)$, $i \in [2n-2]$, where $\mathbf{c}_1, \dots, \mathbf{c}_{2n-2}$ are distinct codewords of weight $n/2$ in $\mathcal{R}(1, m)$. Let us consider a child of the root with a set $\mathcal{S}(\mathbf{c}) = \{s_1, s_2, \dots, s_{n/2}\}$ and let us consider the code that is obtained by taking those codewords of the original RM code that are zero on $[n] \setminus \mathcal{S}(\mathbf{c})$ and deleting these positions. Since the set $\mathcal{S}(\mathbf{c})$ defines an $(m-1)$ -dimensional subspace, the shortened code is an RM code, whose codewords are incidence vectors of subsets in $\text{EG}(m-1, 2)$ [23, Sec. 13.4]. Therefore, we can apply a similar procedure to generate the children of the considered node as for the root node. Let $\hat{\mathbf{c}} = (\hat{c}_1, \hat{c}_2, \dots, \hat{c}_{n/2})$ be a codeword of weight $n/4$ in $\mathcal{R}(1, m-1)$ and let $\hat{\mathcal{S}}(\hat{\mathbf{c}}) = \{s_i \mid i \in [n/2], \hat{c}_i = 1\}$. Then the children nodes of the node with the set $\mathcal{S}(\mathbf{c})$ contain the sets $\hat{\mathcal{S}}(\hat{\mathbf{c}}_i)$, $i \in [n-2]$, where $\hat{\mathbf{c}}_1, \dots, \hat{\mathbf{c}}_{n-2}$ are distinct codewords of weight $n/4$ in $\mathcal{R}(1, m-1)$. This process continues until the node with a set of size 2^{m-r} , which contains non-zero positions of a minimum weight codeword in $\mathcal{R}(r, m)$, is reached.

The greedy search algorithm can be implemented using the tree \mathcal{T} . Recall that this function aims to find a codeword at the minimum distance from the current codeword \mathbf{c} that has the best metric. To decrease the complexity of the algorithm, we consider a depth-first search that is terminated when the first leaf is reached. Thus, it is required to find a leaf in the tree \mathcal{T} such that if we flip bits in the codeword \mathbf{c} at the positions in the set of this leaf, we will get the codeword with the best metric.

Consider a leaf with a set \mathcal{S} and assume that $\mathbf{c} = (c_1, c_2, \dots, c_n)$ is the current codeword. If we choose this leaf as a solution, then the metric of the output codeword is calculated as

$$\sum_{i=1}^n (1 - 2c_i) y_i - 2 \sum_{i \in \mathcal{S}} (1 - 2c_i) y_i, \quad (2)$$

where (y_1, y_2, \dots, y_n) is a vector of input LLRs. Since we are required to maximize the metric of the output codeword, the algorithm needs to find a leaf that minimizes the second sum of (2).

Consider a node of the tree \mathcal{T} that contains a set $\mathcal{S} = \{s_1, s_2, \dots, s_{\bar{n}}\}$, $1 \leq s_1 < s_2 < \dots < s_{\bar{n}} \leq n$. Assume that this node is not a leaf. Consequently, it has $2\bar{n} - 2$ children nodes. Since we consider the greedy depth-first search, the algorithm needs to move to a child node that seems best at the moment. Here, we propose to use the following heuristic: For every child node with a set $\bar{\mathcal{S}}$, we compute a metric

$$2 \sum_{i \in \bar{\mathcal{S}}} (1 - 2c_i) y_i \quad (3)$$

and move to the node with the smallest metric value. The algorithm continues until the leaf node is reached. The output of the algorithm is the codeword that is obtained from the current one by flipping bits at the positions in the set of the leaf node. Observe that the leaf with minimum metric value (3) provides the best solution in terms of (2).

A straightforward calculation of (3) for all children of a node with a set \mathcal{S} takes $\mathcal{O}(\bar{n}^2)$ time. Since a node of the tree \mathcal{T} contains a set with non-zero positions of a codeword of the first order RM code, we can improve the running time of a naive algorithm by taking advantage of the Fast Hadamard Transform (FHT) [26]. Let $\bar{\mathbf{y}} = (\bar{y}_1, \bar{y}_2, \dots, \bar{y}_{\bar{n}})$ be a vector such that $\bar{y}_i = (1 - 2c_{s_i}) y_{s_i}$ and let

$$\mathbf{H} = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}^{\otimes \log_2 \bar{n}}, \quad (4)$$

where $\mathbf{X}^{\otimes t}$ denotes t -times Kronecker product of the matrix \mathbf{X} with itself. Then the result of the FHT applied to the vector $\bar{\mathbf{y}}$ can be expressed in terms of matrix multiplication as $(h_1, h_2, \dots, h_{\bar{n}}) = \bar{\mathbf{y}} \mathbf{H}$. Note that the columns of the matrix \mathbf{H} correspond to half of the first-order RM code codewords. Namely, if $(x_1, x_2, \dots, x_{\bar{n}})^T$ is a column of the matrix \mathbf{H} associated with the codeword \mathbf{c} , then $c_i = (1 - x_i)/2$. The rest codewords are defined by the columns of $-\mathbf{H}$. Therefore, the values of (3) for the children nodes can be calculated as $(-1)^a h_i + h_1$, where $i \in \{2, \dots, \bar{n}\}$, $a \in \{0, 1\}$. The running time of this approach is $\mathcal{O}(\bar{n} \log \bar{n})$ [26].

Example 2: Consider a similar scenario as in Example 1. Namely, the RM code $\mathcal{R}(2, 3)$ and LLR vector (1). Assume that $\mathbf{c} = (0, 0, 1, 1, 1, 1, 0, 0)$ is the current codeword. The correlation metric for this codeword equals 8.44. Let us find a codeword at the minimum distance from \mathbf{c} that has the largest correlation using the depth-first algorithm described above. In the first step, the algorithm computes the metric (3) for each child of the root and moves to the node with the set $\mathcal{S} = \{2, 4, 6, 8\}$ that has the smallest metric, namely, -9.1 . Then, the algorithm applies the same procedure to the node with the set \mathcal{S} and selects a node with the set $\{4, 6\}$. The metric value for this node equals -16.64 and, from (2), it follows that if we flip bits at positions 4 and 6, we will get a codeword $\mathbf{c}' = (0, 0, 1, 0, 1, 0, 0, 0)$ with a metric of 25.04. Note that \mathbf{c}' and \mathbf{c}_1 found in Example 1 using an optimal algorithm are the same.

The version of the greedy algorithm described above is able to find only one solution. If this solution is in the set \mathcal{C} (see line 3 of Algorithm 1), then the break statement in line 12 of Algorithm 1 will terminate the for-loop in lines 9 – 23. We observe that in this case, the algorithm terminates before the

correct codeword is found, significantly decreasing the error-rate performance of the decoder. Therefore, one requires to modify the greedy approach described above in such a way that it can visit multiple leaves and select one that generates the codeword \mathbf{c} , $\mathbf{c} \notin \mathcal{C}$. For instance, one can choose l children of the root of the tree \mathcal{T} and run the depth-first search for each child independently. Then the algorithm can generate up to l different codewords and choose one that is not in \mathcal{C} . We call parameter l a search breadth of the decoding algorithm.

Note that the complexity of the decoding algorithm depends on the value of the search breadth. If the search breadth is too small, then the algorithm has a high probability of termination before reaching the correct codeword. On the other hand, a large search breadth significantly increases the running time of the algorithm. We found that a good trade-off here is to run the search with a relatively small l , for instance, $l = 8$. However, if all generated codewords are in the set \mathcal{C} , then we allow running the depth-first search for \bar{l} additional children of the root of the tree \mathcal{T} . This operation is allowed to be run only s times for the one launch of the decoder. If the algorithm is not able to find a codeword \mathbf{c} , $\mathbf{c} \notin \mathcal{C}$, for the given search breadth l and the algorithm has already run the depth-first search for \bar{l} additional children s times, then decoding is terminated. This limitation is introduced to decrease the average computational complexity of the algorithm. Note that if the search for each child node of the root runs sequentially, one can stop the search procedure as soon as a codeword \mathbf{c} , $\mathbf{c} \notin \mathcal{C}$, that improves the metric of the current best codeword is found.

Example 3: Let us consider the RM code $\mathcal{R}(2, 3)$ and let us decode LLR vector (1) using a greedy approach discussed above. Let $N = 10$, $l = 2$, $\bar{l} = 1$, $s = 1$ and assume that $\mathbf{c} = (0, 0, 1, 1, 1, 1, 0, 0)$ is the output of the recursive decoder. As in Example 2, in the first iteration, the algorithm finds a vector $\mathbf{c}_1 = (0, 0, 1, 0, 1, 0, 0, 0)$ with a metric of 25.08 by applying the depth-first search to the child node with the best metric (3). Since the metric of \mathbf{c}_1 is greater than the metric of \mathbf{c} , the algorithm does not run the depth-first search for the second-best child and moves to the next iteration. In the second iteration, the best child of the root has the set $\{1, 4, 5, 8\}$ and the depth-first search for this node returns the set $\{5, 8\}$. This set generates the codeword $\mathbf{c}_2 = (0, 0, 1, 0, 0, 0, 0, 1)$ with a metric of 28.72. As in the first iteration, the depth-first search for the best child results in a new codeword with a better metric than that of the codewords \mathbf{c} and \mathbf{c}_1 . Therefore, the algorithm terminates the second iteration and moves to the next one. In the third iteration, the two best children of the root are defined by sets $\{1, 4, 5, 8\}$ and $\{5, 6, 7, 8\}$. The depth-first search for both these nodes results in the set $\{5, 8\}$ generating the codeword \mathbf{c}_1 that has been already found in the first iteration. Since the search for the l best children does not result in a new codeword, the algorithm runs the depth-first search for the third-best child defined by the set $\{1, 2, 5, 6\}$. This search results in the set $\{1, 5\}$ and the codeword $\mathbf{c}_3 = (1, 0, 1, 0, 1, 0, 0, 1)$ with a metric of 23.38. In the fourth iteration, as in the previous one, the search for the two best children of the root defined by sets $\{1, 4, 5, 8\}$ and $\{1, 2, 5, 6\}$ results in the codeword \mathbf{c}_2 that has been already found. Since the search has been already performed

for additional \bar{l} nodes s times, the algorithm terminates here and returns a codeword with the best metric from the set $\{c, c_1, c_2, c_3\}$. Note that the result in each iteration is the same as in Example 1 that considers an optimal algorithm to find the next codeword.

The formal description of the proposed greedy search approach is shown in Algorithm 2. This algorithm uses the `NextStepGreedyRec` function presented in Algorithm 3. The `NextStepGreedyRec` function performs a depth-first search to find a leaf that minimizes metric (3). At each recursive step, the function `NextStepGreedyRec` chooses a node of the tree \mathcal{T} that minimizes (3) (see lines 12 – 22). Condition in line 1 checks whether a leaf node has been reached. If so, then the algorithm flips bits of the current codeword c in the coordinates from the vector v (see lines 3 – 5). Note that, for simplicity, we use vectors instead of sets to store positions of bits to be flipped. If the new codeword is not in the set \mathcal{C} , then the new codeword and the value of (3) are returned. Otherwise, the algorithm returns ∞ instead of the value of metric (3).

The `NextStepGreedy` function calculates metric (3) for the children of the root node of the tree \mathcal{T} (see lines 5 – 7). Then it uses the `NextStepGreedyRec` function to perform a depth-first search for l children with the smallest metric values. Note that $l + \bar{l}$ must be less than or equal to $2^{m+1} - 2$. If a new codeword has been found during the first l iterations, then the loop in lines 8 – 24 is terminated (see lines 9 – 15). If additional \bar{l} nodes are required to find a new codeword, then the algorithm sets a flag ec to one. This flag is used to count how many times the algorithm used the option to run the depth-first search for additional \bar{l} nodes. If the `NextStepGreedy` function returns ec that equals one s times during decoding of the received vector, then the graph search algorithm prohibits the `NextStepGreedy` function from visiting \bar{l} extra nodes. For instance, it can be done by setting \bar{l} to be equal to zero. Note that the loop in lines 8 – 24 is also terminated if a codeword with a better metric has been found (see line 23). If the `NextStepGreedy` function is not able to find a codeword that is not in \mathcal{C} , then the function returns the input codeword c and $-\infty$ as a metric value. The formal description of the graph search algorithm that uses the `NextStepGreedy` function is presented in Algorithm 4.

We now consider the running time and the space complexity of Algorithm 4. First, we consider the computational complexity and space requirements of the functions `NextStepGreedyRec` and `NextStepGreedy`.

Lemma 1: The running time of the `NextStepGreedyRec` function is $\mathcal{O}(n \log(\max\{N, n\}))$, $n = 2^m$.

Proof: Consider the base case in lines 2 – 10 of Algorithm 3. The algorithm flips at most 2^{m-r} bits in line 4 and checks whether a new codeword is in the set \mathcal{C} . This check can be implemented efficiently using a red-black tree guaranteeing that the worst-case running time of the search operation is $\mathcal{O}(\log N)$ [27, Sec. 13]. Since the per-bit comparison of two codewords takes linear time, the base case running time is $\mathcal{O}(n \log N)$.

Algorithm 2: The `NextStepGreedy` function

Input: RM code parameters r and m , a vector of LLRs y , a codeword c , a set of codewords \mathcal{C} , search breadth l and \bar{l}

Output: A codeword res , a metric M of the codeword res , a flag ec whether the algorithm did not find a codeword using l attempts

```

1  $res \leftarrow c$ ,  $M_c \leftarrow 0$ ,  $M \leftarrow \infty$ ,  $ec \leftarrow 0$ 
2 for  $i = 1, 2, \dots, 2^m$  do
3    $M_c \leftarrow M_c + (1 - 2c[i]) y[i]$ 
4 end
5 Let  $c^i, 1 \leq i \leq 2^{m+1} - 2$ , be all codewords of the code  $\mathcal{R}(1, m)$  with the Hamming weight  $2^{m-1}$ 
6 Let  $a$  be a vector of length  $2^{m+1} - 2$ ,
    $a[i] \leftarrow 2 \cdot \sum_{j, c^i[j]=1} (1 - 2c[j]) y[j]$ 
7 Let  $k$  be a vector that index the vector of metrics  $a$  in ascending order
8 for  $i = 1, 2, \dots, l + \bar{l}$  do
9   if  $i > l$  then
10    if  $M = \infty$  then
11       $ec \leftarrow 1$ 
12    else
13      break
14    end
15  end
16  Let  $v$  be a vector of size  $2^{m-1}$ 
17   $h \leftarrow 1$ 
18  for  $j = 1, 2, \dots, 2^m$  do
19    if  $c^{k[i]}[j] = 1$  then  $v[h] \leftarrow j$ ,  $h \leftarrow h + 1$ 
20  end
21   $c', M' \leftarrow \text{NextStepGreedyRec}(r - 1, m - 1, y, c, \mathcal{C}, v, a[k[i]])$ 
22  if  $M' < M$  then  $M \leftarrow M'$ ,  $res \leftarrow c'$ 
23  if  $M < 0$  then break
24 end
25 if  $M = \infty$  then
26    $M \leftarrow -\infty$ 
27 else
28    $M \leftarrow M_c - M'$ 
29 end
30 return  $res, M, ec$ 

```

The elements of the array a (see line 13) can be calculated using the FHT that takes time $\mathcal{O}(n \log n)$. Since the algorithm computes the array a at each recursive step, the total running time of this operation is

$$\mathcal{O}\left(\sum_{i=0}^{r-1} 2^{m-i} \log(2^{m-i})\right) = \mathcal{O}(n \log n). \quad (5)$$

Observe that the total running time of the for loop in lines 17 – 21 and the minimum value search in line 14 is $\mathcal{O}(n)$. Thus, the `NextStepGreedyRec` function takes $\mathcal{O}(n \log(\max\{N, n\}))$ time. ■

Lemma 2: The space complexity of the `NextStepGreedyRec` function is $\mathcal{O}(n^2)$, $n = 2^m$.

Proof: Observe that the base case and lines 14 – 21 of Algorithm 3 use linear space. The algorithm also stores all codewords of $\mathcal{R}(1, h)$, $h \in \{m - r + 1, m - r + 2, \dots, m\}$, with the Hamming weight 2^{h-1} to select the codeword with the smallest metric efficiently (line 14 and line 18). Since the codewords of the first order RM codes can be defined

Algorithm 3: The NextStepGreedyRec function

Input: RM code parameters r and m , a vector of LLRs \mathbf{y} , a codeword \mathbf{c} , a set of codewords \mathcal{C} , a vector of coordinates \mathbf{v} , a local metric value M'

Output: A codeword \mathbf{res} , a metric M that is used to find

```

1   $\mathbf{res} \leftarrow \mathbf{c}$ 
2  if  $r = 0$  then
3     $\mathbf{res} \leftarrow \mathbf{c}$ 
4    for  $i = 1, 2, \dots, 2^m$  do
5       $\mathbf{res}[\mathbf{v}[i]] \leftarrow \mathbf{res}[\mathbf{v}[i]] \oplus 1$ 
6    end
7    if  $\mathbf{res} \in \mathcal{C}$  then
8      return  $\mathbf{res}, \infty$ 
9    else
10     return  $\mathbf{res}, M'$ 
11  end
12 Let  $\mathbf{c}^i, 1 \leq i \leq 2^{m+1} - 2$ , be all codewords of the code
    $\mathcal{R}(1, m)$  with the Hamming weight  $2^{m-1}$ 
13 Let  $\mathbf{a}$  be a vector of length  $2^{m+1} - 2$ ,
    $\mathbf{a}[i] \leftarrow 2 \cdot \sum_{j, \mathbf{c}^j[i]=1} (1 - 2\mathbf{c}[\mathbf{v}[j]]) \mathbf{y}[\mathbf{v}[j]]$ 
14 Let  $k$  be an index of the minimum value in  $\mathbf{a}$ 
15 Let  $\hat{\mathbf{v}}$  be a vector of size  $2^{m-1}$ 
16  $h \leftarrow 1$ 
17 for  $j = 1, 2, \dots, 2^m$  do
18   if  $\mathbf{c}^k[j] = 1$  then
19      $\hat{\mathbf{v}}[h] \leftarrow \mathbf{v}[j], h \leftarrow h + 1$ 
20   end
21 end
22 return
   NextStepGreedyRec( $r - 1, m - 1, \mathbf{y}, \mathbf{c}, \mathcal{C}, \hat{\mathbf{v}}, \mathbf{a}[k]$ )

```

recursively using (4), it follows that it is enough to store only codewords of $\mathcal{R}(1, m)$ that takes $\mathcal{O}(n^2)$ space. Therefore, the algorithm uses $\mathcal{O}(n^2)$ space. ■

Lemma 3: The best-case time of the NextStepGreedy function is $\mathcal{O}(n \log(\max\{N, n\}))$. The worst-case time of the NextStepGreedy function is $\mathcal{O}((l + \bar{l}) \cdot n \log(\max\{N, n\}))$.

Proof: Consider line 6 and line 7 of Algorithm 2. As in Algorithm 3, the vector \mathbf{a} is computed in $\mathcal{O}(n \log n)$. Sorting in line 7 also takes $\mathcal{O}(n \log n)$ [27, Part II]. Observe that all the rest calculations require linear time exclusive of line 21.

In line 21, the algorithm uses the NextStepGreedyRec function to find a new codeword. In the best case, this function is called only once. Since this function runs in $\mathcal{O}(n \log(\max\{N, n\}))$, the best-case running time of the NextStepGreedy function is $\mathcal{O}(n \log(\max\{N, n\}))$. In the worst case, the NextStepGreedyRec function is called $l + \bar{l}$ times. Thus, the worst-case running time is $\mathcal{O}((l + \bar{l}) \cdot n \log(\max\{N, n\}))$. ■

Lemma 4: The space complexity of the NextStepGreedy function is $\mathcal{O}(n^2)$.

Proof: As the NextStepGreedyRec function, the NextStepGreedy function stores codewords of $\mathcal{R}(1, m)$ that take $\mathcal{O}(n^2)$ space. Observe that the auxiliary variables use linear space. Therefore, the space requirement of the NextStepGreedy function is indeed $\mathcal{O}(n^2)$. ■

Note that Algorithm 2 allows for parallel implementation. One can run the first l iterations of the for loop

Algorithm 4: A greedy version of the graph search decoding algorithm

Input: RM code parameters r and m , a vector of LLRs \mathbf{y} , a maximum number of iterations N , a flag t whether to use the CRC to terminate the algorithm, search breadth l and \bar{l} , a number of extra \bar{l} attempts to find a codeword s

Output: A codeword \mathbf{res}

```

1  Set  $\mathcal{C}$  to be an empty set of codewords
2  Let  $\mathbf{c}$  be the result of recursive decoding of the vector  $\mathbf{y}$ 
3  Add  $\mathbf{c}$  to  $\mathcal{C}$ 
4   $\mathbf{res} \leftarrow \mathbf{c}$ 
5   $M \leftarrow 0$ 
6  for  $i = 1, 2, \dots, 2^m$  do
7     $M \leftarrow M + (1 - 2\mathbf{c}[i]) \mathbf{y}[i]$  // Compute the
    metric for the codeword  $\mathbf{c}$ .
8  end
9  for  $i = 1, 2, \dots, N$  do
10    $\mathbf{c}', M', ec \leftarrow \text{NextStepGreedy}(r, m, \mathbf{y}, \mathbf{c}, \mathcal{C}, l, \bar{l})$ 
11   if  $ec = 1$  then  $s \leftarrow s - 1$ 
12   if  $s = 0$  then  $\bar{l} \leftarrow 0$ 
13   if  $M' = -\infty$  then
14     break // A codeword  $\mathbf{c}'$  such that  $\mathbf{c}' \notin \mathcal{C}$ 
    is not found.
15   end
16    $\mathbf{c} \leftarrow \mathbf{c}'$  // Move to an adjacent node.
17   Add  $\mathbf{c}$  to  $\mathcal{C}$ 
18   if  $M' > M$  then
19     // Save the current codeword if it
    is better than one found so far.
20      $\mathbf{res} \leftarrow \mathbf{c}$ 
21      $M \leftarrow M'$ 
22     if  $t$  and  $\mathbf{res}$  satisfies the CRC then
23       break
24   end
25 end
26 return  $\mathbf{res}$ 

```

in lines 8 – 24 of the algorithm in parallel and select a codeword with the best metric. If all codewords returned by l calls of the NextStepGreedyRec function are in the set \mathcal{C} , then the algorithm can run \bar{l} more iterations in parallel. Consequently, the best-case running time of the parallel implementation of the NextStepGreedy function is $\mathcal{O}(l \cdot n \log(\max\{N, n\}))$, while the worst-case time is still $\mathcal{O}((l + \bar{l}) \cdot n \log(\max\{N, n\}))$. The space complexity of the parallel implementation of the NextStepGreedy function is $\mathcal{O}(n^2)$. Indeed, the functions NextStepGreedy and NextStepGreedyRec only read information from the memory that stores codewords of $\mathcal{R}(1, h)$, $h \in \{m - r + 1, m - r + 2, \dots, m\}$. Thus, only one instance of these codewords can be stored in memory. Moreover, since $l + \bar{l} \leq 2^{m+1} - 2$, the auxiliary variables take $\mathcal{O}(n^2)$ space. Consequently, the space requirement of the sequential and the parallel versions of the algorithm is the same.

Theorem 1: The worst-case running time of Algorithm 4 is $\mathcal{O}(L \cdot n \log(\max\{N, n\}))$, where $L = N \cdot l + s \cdot \bar{l}$.

Proof: Let us consider the running time of the main loop in lines 9 – 25 of Algorithm 4. Recall that we assume that the set \mathcal{C} is implemented using a red-black tree. Therefore, the insertion operation in line 17 of Algorithm 4 takes $\mathcal{O}(n \log N)$

time [27, Sec. 13]. In addition, we assume that a reasonable length CRC is used. Under this assumption, we can conclude that line 21 of Algorithm 4 takes a shorter running time than line 10. In the worst case, the break statement does not terminate the for loop in lines 9 – 25 of Algorithm 4 and the algorithm uses the possibility to execute \bar{l} additional iterations of the for loop in lines 8 – 24 of Algorithm 2 s times. Thus, the worst-case time of the main loop is $\mathcal{O}((N \cdot l + s \cdot \bar{l}) n \log(\max\{N, n\}))$.

Since recursive decoding takes $\mathcal{O}(n \log n)$ time [22], Algorithm 4 takes $\mathcal{O}(L \cdot n \log(\max\{N, n\}))$ worst-case time. ■

Theorem 2: Algorithm 4 uses $\mathcal{O}(\max\{N, n\} \cdot n)$ space.

Proof: On the one hand, the algorithm needs to store the set of codewords \mathcal{C} . Assume that this set is implemented using a red-black tree. Thus, it takes $\mathcal{O}(N \cdot n)$ space [27, Sec. 13]. On the other hand, the algorithm calls the `NextStepGreedy` function that uses $\mathcal{O}(n^2)$ space. Since the auxiliary variables use $\mathcal{O}(1)$ space and recursive decoding uses linear space, the space complexity of Algorithm 4 is $\mathcal{O}(\max\{N, n\} \cdot n)$. ■

V. SIMULATION RESULTS

We consider RM codes of length 256 and 512 and the second-order RM codes of length 1024 and 2048. The graph search decoder (referred to as GS) is compared to RPA decoding without list [10] and the recursive permutation list decoder (referred to as RPL) [5]. We assume that the graph search decoder uses the sequential implementation of Algorithm 2 and runs with parameters $l = 8$, $\bar{l} = 8$, and $s = 5$. Note that the simulation results for the recursive permutation list algorithm are obtained using an open source project available online [28].

In Fig. 2, we present the block error rate (BLER) of the considered decoders for a BI-AWGN channel. For graph search decoding, we consider two scenarios. In the first scenario, plotted as solid blue curves, the number of iterations is chosen in such a way that the worst-case¹ running time of the proposed algorithm is similar to that of recursive permutation list decoding with the list of size L . In the second scenario, plotted as dashed orange curves, we consider the number of iterations required to perform within 0.1 dB from the ML decoding lower bound that is constructed using an approach described in [5]. Namely, we compare the metric of the output codeword of the graph search decoder with a large running time to that of the transmitted codeword. If the metric of the decoder's output is better, then the ML decoder will also return an incorrect codeword. Consequently, the ratio of such events gives a lower bound on the ML decoding performance.

We can see that, in most cases, the graph search decoder outperforms the recursive permutation list decoder with a similar running time by approximately 0.1 dB. The gap increases to 0.4 dB and 0.6 dB for $\mathcal{R}(3, 9)$ and the second-order RM codes of length larger than 512, respectively. However, the recursive permutation list algorithm performs better than

the graph search algorithm for a high-rate $\mathcal{R}(5, 8)$. Observe that for $\mathcal{R}(5, 8)$ we consider the recursive permutation list algorithm with a relatively small list of size 64, while for the rest cases the list size is at least 256. Therefore, from Fig. 2, we notice that the graph search decoder is competitive with the recursive permutation list decoder with a list of size greater than 128.

RPA decoding outperforms graph search decoding considered in the first scenario only for $\mathcal{R}(2, 11)$. However, it is done at the cost of a much longer running time. In Table I, we present the running time to decode one codeword for the recursive permutation list, the RPA, and the graph search algorithms considered in Fig. 2. Note that, for graph search decoding, results are presented for the first scenario. Since the running time of the graph search and the RPA algorithms depends on the noise level, we only present the running time at a BLER of 10^{-4} for these two cases. Recall that the worst-case running time of graph search decoding coincides with the running time of the recursive permutation list algorithm. From Table I, we can see that the running time of RPA decoding is much longer compared to the proposed algorithm. For $\mathcal{R}(2, 11)$, the graph search decoder with 1024 iterations outperforms RPA decoding by approximately 0.4 dB at a BLER of 10^{-4} , while the average-case running time of these two algorithms is the same. Thus, from Fig. 2, we can conclude that the proposed algorithm outperforms the RPA decoder with a similar average-case running time.

In [10], a simplified version of RPA decoding is proposed for high-rate RM codes. For instance, a list version of simplified RPA decoding is demonstrated to achieve the ML decoding performance for $\mathcal{R}(5, 8)$. It takes approximately 13 ms at a BLER of 10^{-4} . The graph search decoder with 128 iterations also achieves the ML decoding performance, but it takes only 3 ms on average at a BLER of 10^{-4} . Another simplified version of RPA decoding, called sparse multi-decoder RPA (SRPA), is proposed in [12]. The SRPA algorithm allows decreasing the running time of RPA decoding up to four times for $\mathcal{R}(2, 9)$ at the cost of 0.15 dB loss and up to eight times for $\mathcal{R}(3, 8)$ without any performance loss. Thus, the running time of the SRPA algorithm and the graph search algorithm for $\mathcal{R}(2, 9)$ is similar. However, due to 0.15 dB performance loss, the graph search decoding demonstrates slightly better performance. In the case of $\mathcal{R}(3, 8)$, the running time of SRPA decoding is much longer compared to the proposed algorithm. Hence, simplified versions of RPA decoding are also not competitive with the graph search algorithm.

From Fig. 2, we can see that the graph search decoder requires a large number of iterations to perform 0.1 dB from ML for RM codes of length 512 and order greater than 3. For instance, the required number of iterations equals 2^{15} for $\mathcal{R}(4, 9)$ resulting in a very long running time in the worst case. However, the average-case running time of the proposed algorithm is much smaller. In Figs. 3 and 4 we compare the average-case running time as a function of E_b/N_0 for $\mathcal{R}(3, 9)$ and $\mathcal{R}(4, 9)$, respectively. Note that the graph search decoder approaching ML performance takes less time to decode than the recursive permutation list algorithm in high signal-to-noise region. The average-case running time of the graph search

¹We consider the worst-case running time of the graph search algorithm as in Theorem 1, i.e. the break statement does not terminate the main loop of Algorithm 4 and each call of the `NextStepGreedy` function takes the worst-case running time.

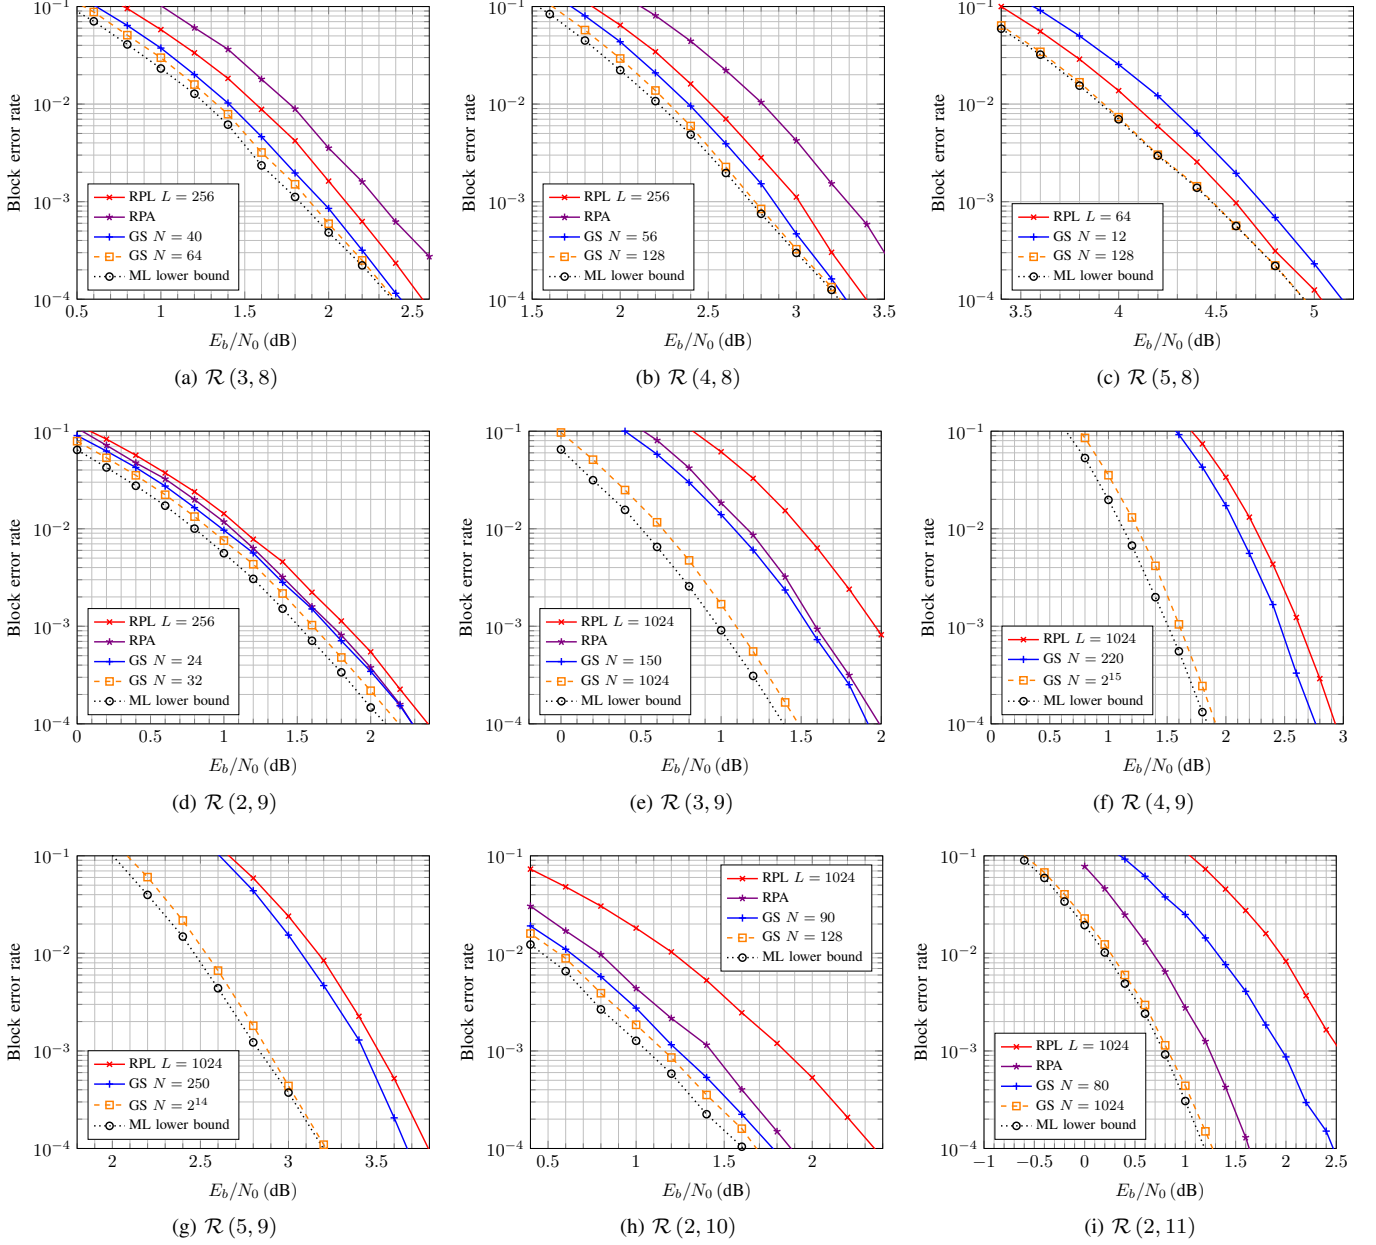


Fig. 2. The block error rate performance of RM codes on a BI-AWGN channel. For the recursive permutation list algorithm, we use L to denote the list size. For the graph search algorithm, we use N to denote the number of iterations.

decoder can be further reduced using a CRC. Recall that the proposed algorithm is terminated if a codeword with the best metric found so far satisfies the CRC. In our simulations, we consider the 24-bits CRC with the generator polynomial

$$x^{24} + x^{23} + x^{21} + x^{20} + x^{17} + x^{15} + x^{13} + x^{12} + x^8 + x^4 + x^2 + x + 1.$$

Such a long CRC guarantees that undetected errors, i.e. the codeword that satisfies the CRC is not a valid one, do not affect the performance for a BLER above 10^{-4} . Hence, the results presented in Figs. 3 and 4 can be regarded as the best possible improvement of the average-case running time if we have a genie that stops decoding immediately after a valid codeword is found. Although this technique allows reducing the average-case running time, the 24-bits CRC results in a huge rate loss

for the considered cases. Thus, a further improvement of the proposed algorithm is required to decrease the average-case running time using a reasonable length CRC with negligible performance loss.

In Fig. 5, we plot the block error probability for a binary symmetric channel (BSC). Note that we consider a simplified version of RPA decoding proposed for a BSC (see Algorithm 1 in [10]), while for graph search decoding and recursive permutation list decoding we use hard-input versions of the algorithms considered for a BI-AWGN channel. As in Fig. 2, we use solid blue curves to report results of the graph search algorithm with the worst-case running time that is similar to that of the recursive permutation list algorithm and we use dashed orange curves to report results of the graph search

TABLE I

COMPARISON OF AVERAGE-CASE RUNNING TIME TO DECODE ONE CODEWORD AMONG ALGORITHMS CONSIDERED IN FIG. 2. FOR THE GRAPH SEARCH AND THE RPA ALGORITHMS, THE RUNNING TIME IS ESTIMATED AT E_b/N_0 REQUIRED TO REACH A BLER OF 10^{-4} . THE WORST-CASE RUNNING TIME OF THE GRAPH SEARCH ALGORITHM IS SIMILAR TO THE RUNNING TIME OF THE RECURSIVE PERMUTATION LIST ALGORITHM.

Code	$\mathcal{R}(3, 8)$	$\mathcal{R}(4, 8)$	$\mathcal{R}(5, 8)$	$\mathcal{R}(2, 9)$	$\mathcal{R}(3, 9)$	$\mathcal{R}(4, 9)$	$\mathcal{R}(5, 9)$	$\mathcal{R}(2, 10)$	$\mathcal{R}(2, 11)$
RPL	3.9ms	5.5ms	1.5ms	4ms	29ms	46ms	53ms	35ms	66ms
RPA	196ms	9.1s	—	14ms	2.1s	—	—	58ms	227ms
GS	2.4ms	1ms	0.2ms	3.1ms	10ms	4.8ms	4ms	24ms	45ms




Fig. 3. The average-case running time of $\mathcal{R}(3, 9)$ decoding on a BI-AWGN channel.




Fig. 4. The average-case running time of $\mathcal{R}(4, 9)$ decoding on a BI-AWGN channel.

algorithm approaching the ML performance. In contrast to a BI-AWGN channel, the graph search decoder requires a moderate number of iterations ($N \leq 1024$) to approach the performance of ML decoding for a BSC. For RM codes of length greater than 256, the proposed algorithm outperforms the recursive permutation list algorithm with the same worst-case running time, while these decoders perform similarly for RM codes of length 256. Furthermore, the graph search algorithm outperforms RPA decoding and, at the same time, takes a shorter average-case running time. For instance, the average-case running time of RPA decoding for $\mathcal{R}(3, 8)$ at a BLER of 10^{-4} is 105 ms, while the proposed algorithm with

16 iterations takes only 0.9 ms.

VI. CONCLUSION

We presented a new decoder for RM codes, which benefits from the representation of the code as a graph. Such representation allows using a greedy local search algorithm that is able to find the transmitted codeword efficiently. In almost all considered cases, our algorithm outperforms the state-of-the-art decoders of RM codes with a similar complexity. Furthermore, the proposed algorithm allows achieving the performance of the ML decoder with a reasonable worst-case running time on a BSC. In the case of a BI-AWGN channel, we demonstrated that the ML decoder performance is achieved with a feasible average computational complexity, which can be further reduced using a CRC.

ACKNOWLEDGMENT

We thank the reviewers for valuable suggestions that have helped to significantly improve the quality of this paper. We also thank Vladimir Gritsenko and Alexey Maevskiy for useful discussions and feedback.

REFERENCES

- [1] M. Kamenov, "On decoding of Reed-Muller codes using a local graph search," in *2020 IEEE Information Theory Workshop (ITW)*, 2021, pp. 1–5.
- [2] D. E. Muller, "Application of Boolean algebra to switching circuit design and to error detection," *Transactions of the I.R.E. Professional Group on Electronic Computers*, vol. EC-3, no. 3, pp. 6–12, Sep. 1954.
- [3] I. Reed, "A class of multiple-error-correcting codes and the decoding scheme," *Transactions of the IRE Professional Group on Information Theory*, vol. 4, no. 4, pp. 38–49, Sep. 1954.
- [4] S. Kudekar, S. Kumar, M. Mondelli, H. D. Pfister, E. Şaşoğlu, and R. L. Urbanke, "Reed-Muller codes achieve capacity on erasure channels," *IEEE Transactions on Information Theory*, vol. 63, no. 7, pp. 4298–4316, July 2017.
- [5] I. Dumer and K. Shabunov, "Soft-decision decoding of Reed-Muller codes: recursive lists," *IEEE Transactions on Information Theory*, vol. 52, no. 3, pp. 1260–1266, March 2006.
- [6] N. Stolte, "Recursive codes with the Plotkin construction and their decoding," Ph.D. dissertation, Technical University of Darmstadt, Germany, 2002.
- [7] M. Geiselhart, A. Elkelesh, M. Ebada, S. Cammerer, and S. t. Brink, "Automorphism ensemble decoding of Reed-Muller codes." [Online]. Available: <https://arxiv.org/abs/2012.07635v1>
- [8] N. Doan, S. A. Hashemi, and W. J. Gross, "Successive-cancellation decoding of Reed-Muller codes with fast Hadamard transform," 2021. [Online]. Available: <https://arxiv.org/abs/2108.12550v2>
- [9] K. Ivanov and R. L. Urbanke, "On the efficiency of polar-like decoding for symmetric codes." [Online]. Available: <https://arxiv.org/abs/2104.06084v2>
- [10] M. Ye and E. Abbe, "Recursive projection-aggregation decoding of Reed-Muller codes," *IEEE Transactions on Information Theory*, vol. 66, no. 8, pp. 4948–4965, 2020.

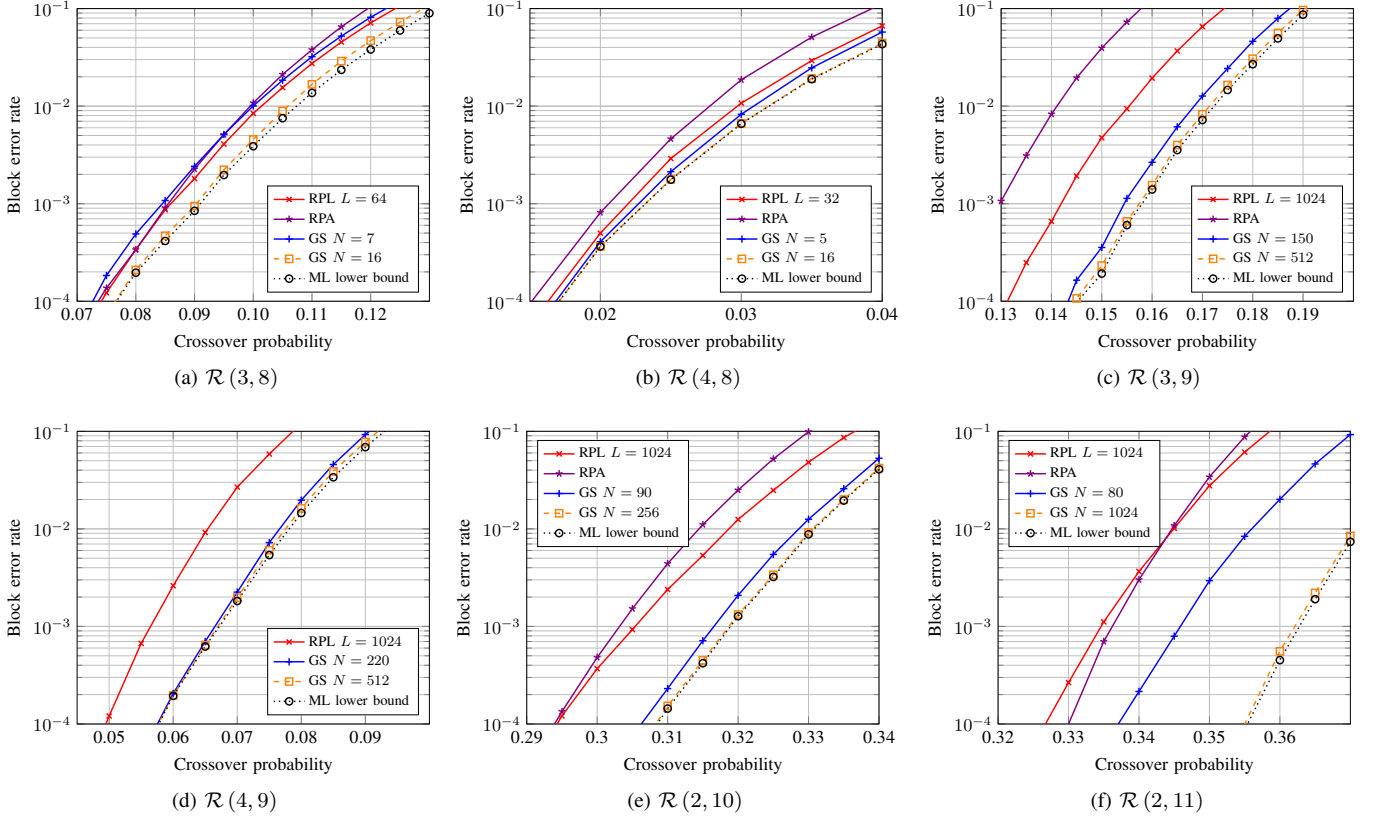


Fig. 5. The block error rate performance of RM codes on a BSC. For the recursive permutation list algorithm, we use L to denote the list size. For the graph search algorithm, we use N to denote the number of iterations.

- [11] K. Ivanov and R. Urbanke, "Improved decoding of second-order Reed-Muller codes," in *2019 IEEE Information Theory Workshop (ITW)*, 2019, pp. 1–5.
- [12] D. Fathollahi, N. Farsad, S. A. Hashemi, and M. Mondelli, "Sparse multi-decoder recursive projection aggregation for Reed-Muller codes," in *2021 IEEE International Symposium on Information Theory (ISIT)*, 2021, pp. 1082–1087.
- [13] M. V. Jamali, X. Liu, A. V. Makkua, H. MahdaviFar, S. Oh, and P. Viswanath, "Reed-Muller subcodes: machine learning-aided design of efficient soft recursive decoding," in *2021 IEEE International Symposium on Information Theory (ISIT)*, 2021, pp. 1088–1093.
- [14] S. A. Hashemi, N. Doan, M. Mondelli, and W. J. Gross, "Decoding Reed-Muller and polar codes by successive factor graph permutations," in *2018 IEEE 10th International Symposium on Turbo Codes Iterative Information Processing (ITC)*, 2018, pp. 1–5.
- [15] E. Santi, C. Hager, and H. D. Pfister, "Decoding Reed-Muller codes using minimum-weight parity checks," in *2018 IEEE International Symposium on Information Theory (ISIT)*, 2018, pp. 1296–1300.
- [16] M. Lian, C. Häger, and H. D. Pfister, "Decoding Reed-Muller codes using redundant code constraints," in *2020 IEEE International Symposium on Information Theory (ISIT)*, 2020, pp. 42–47.
- [17] A. Buchberger, C. Häger, H. D. Pfister, L. Schmalen, and A. G. Amat, "Pruning neural belief propagation decoders," in *2020 IEEE International Symposium on Information Theory (ISIT)*, 2020, pp. 338–342.
- [18] P. Yuan and M. C. Coşkun, "Complexity-adaptive maximum-likelihood decoding of modified G_N -coset codes." [Online]. Available: <https://arxiv.org/abs/2105.04048v2>
- [19] S. A. Hashemi, N. Doan, W. J. Gross, J. Cioffi, and A. Goldsmith, "A tree search approach for maximum-likelihood decoding of Reed-Muller codes." [Online]. Available: <https://arxiv.org/abs/2107.08991v1>
- [20] K. Ivanov and R. Urbanke, "Permutation-based decoding of Reed-Muller codes in binary erasure channel," in *2019 IEEE International Symposium on Information Theory (ISIT)*, 2019, pp. 21–25.
- [21] M. C. Coşkun, J. Neu, and H. D. Pfister, "Successive cancellation inactivation decoding for modified Reed-Muller and eBCH codes," in *2020 IEEE International Symposium on Information Theory (ISIT)*, 2020, pp. 437–442.
- [22] I. Dumer and K. Shabunov, "Recursive decoding of Reed-Muller codes," in *2000 IEEE International Symposium on Information Theory*, 2000, p. 63.
- [23] F. J. MacWilliams and N. J. A. Sloane, *The theory of error-correcting codes*. Amsterdam, The Netherlands: North-Holland, 1977.
- [24] S. Lin and D. J. Costello, *Error control coding*, Second ed. Upper Saddle River, NJ, USA: Pearson Prentice hall, 2004.
- [25] T. Koumoto, T. Takata, T. Kasami, and Shu Lin, "A low-weight trellis-based iterative soft-decision decoding algorithm for binary linear block codes," *IEEE Transactions on Information Theory*, vol. 45, no. 2, pp. 731–741, 1999.
- [26] R. Green, "A serial orthogonal decoder," *JPL Space Programs Summary*, vol. 37, pp. 247–253, 1966.
- [27] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to algorithms*, 3rd ed. Cambridge, MA, USA: MIT Press, 2009.
- [28] K. Shabunov, "Error correcting coding research tools," <https://github.com/kshabunov/ecclab>, 2021.