

Generalized Gray Codes for Local Rank Modulation

Eyal En Gad

Elec. Eng.
Caltech
Pasadena, CA 91125, U.S.A.
eengad@caltech.edu

Michael Langberg

Comp. Sci. Division
Open University of Israel
Raanana 43107, Israel
mikel@openu.ac.il

Moshe Schwartz

Elec. and Comp. Eng.
Ben-Gurion University
Beer Sheva 84105, Israel
schwartz@ee.bgu.ac.il

Jehoshua Bruck

Elec. Eng.
Caltech
Pasadena, CA 91125, U.S.A.
bruck@paradise.caltech.edu

Abstract—We consider the local rank-modulation scheme in which a sliding window going over a sequence of real-valued variables induces a sequence of permutations. Local rank-modulation is a generalization of the rank-modulation scheme, which has been recently suggested as a way of storing information in flash memory.

We study Gray codes for the local rank-modulation scheme in order to simulate conventional multi-level flash cells while retaining the benefits of rank modulation. Unlike the limited scope of previous works, we consider code constructions for the entire range of parameters including the code length, sliding window size, and overlap between adjacent windows. We show our constructed codes have asymptotically-optimal rate. We also provide efficient encoding, decoding, and next-state algorithms.

I. INTRODUCTION

With the recent application to flash memories, the rank-modulation scheme has gained renewed interest as evident in the recent series of papers [9], [15], [16], [18], [20], [22]. In the conventional modulation scheme used in flash-memory cells, the absolute charge level of each cell is quantized to one of q levels, resulting in a single demodulated symbol from an alphabet of size q . In contrast, in the rank modulation scheme a group of n flash cells comprise a single virtual cell storing a symbol from an alphabet of size $n!$, where each symbol is assigned a distinct configuration of relative charge levels in the n cells. Thus, there is no more need for threshold values to distinguish between various stored symbols, which mitigates the effects of retention in flash cells (slow charge leakage). In addition, if we allow only a simple programming (charge-injection) mechanism called “push-to-the-top”, whereby a single cell is driven above all others in terms of charge level, then no over-programming can occur, a problem which considerably slows down programming in conventional multi-level flash cells.

Rank modulation has been studied intermittently since the early works of Slepian [19] (later extended in [1]), in which permutations were used to digitize vectors from a time-discrete memoryless Gaussian source, and Chadwick and Kurz [6], in which permutations were used in the context of signal detection over channels with non-Gaussian noise (especially

impulse noise). Other works on the subject include [1]–[3], [5], [7], [8]. More recently, permutations were used for communicating over powerlines (for example, see [21]), and for modulation schemes for flash memory [15], [16], [20], [22].

One drawback to the rank-modulation scheme is the fact that we need to reconstruct the permutation induced by the relative charge levels of the participating cells. If n cells are involved, at least $\Omega(n \log n)$ comparisons are needed, which might be too high for some applications. It was therefore suggested in [9], [18], [22] that only *local* comparisons be made, creating a sequence of small induced permutations instead of a single all-encompassing permutation. This obviously restricts the number of distinct configurations, and thus, reduces the size of the resulting alphabet as well. In the simplest case, requiring the least amount of comparisons, the cells are located in a one-dimensional array and each cell is compared with its two immediate neighbors requiring a single comparator between every two adjacent cells [9], [18].

Yet another drawback of the rank-modulation scheme is the fact that distinct n charge levels are required for a group of n physical flash cells. Therefore, restricted reading resolution prohibits the use of large values of n . However, when only local views are considered, distinct values are required only within a small local set of cells, thus enabling the use of large groups of cells with local rank modulation.

An important application for rank-modulation in the context of flash memory was described in [15]: A set of n cells, over which the rank-modulation scheme is applied, is used to simulate a single conventional multi-level flash cell with $n!$ levels corresponding to the alphabet $\{0, 1, \dots, n! - 1\}$. The simulated cell supports an operation which raises its value by 1 modulo $n!$. This is the only required operation in many rewriting schemes for flash memories (see [4], [12]–[14], [23]), and it is realized in [15] by a Gray code traversing the $n!$ states where, physically, the transition between two adjacent states in the Gray code is achieved by using a single “push-to-the-top” operation. In the context of local rank modulation, Gray codes for the local scheme were studied in [9], [18], where necessary conditions as well as constructions were provided.

Having considered the two extremes: full rank modulation with a single permutation of n cells, and extreme local

This work was supported in part by ISF grant 134/10, ISF grant 480/08, the Open University of Israel’s research fund (grant no. 46114), the NSF grant ECCS-0802107, and an NSF-NRI award.

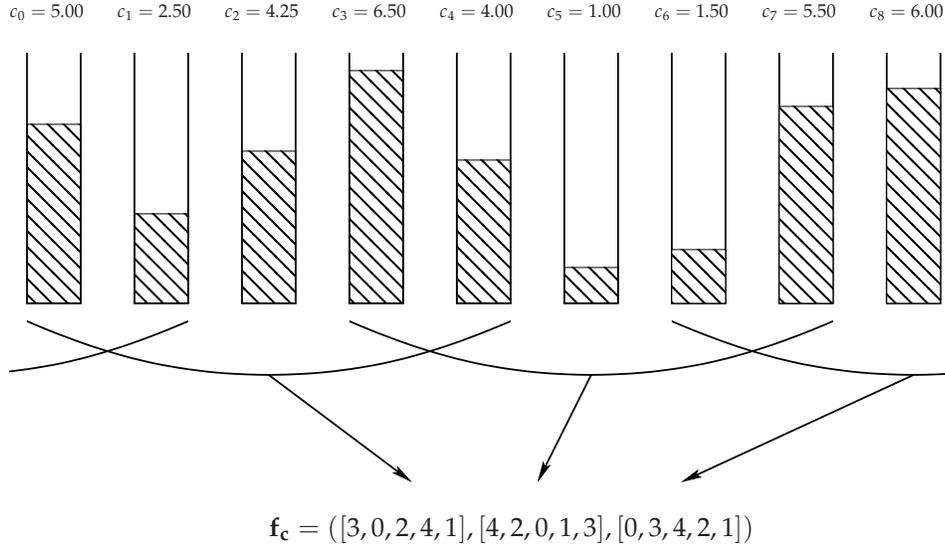


Figure 1. Demodulating a (3,5,9)-locally rank-modulated signal.

rank modulation with a sequence of n permutations over 2 elements, the question of whether any middle-road solutions exist remains open. We address this question in this paper by considering the generalized local rank modulation scheme in which a sequence of several permutations of a given size provide the local views into ranking of the cells. We construct Gray codes for this scheme which asymptotically achieve the maximum possible rate, and consider efficient encoding/decoding algorithms, as well as efficient next-state computation.

The rest of the paper is organized as follows. In Section II we give preliminary definitions and notation. In Section III we present our construction for optimal local rank modulation for general degrees of locality. We conclude with a discussion in Section IV.

II. DEFINITIONS AND NOTATION

We shall now proceed to introduce the notation and definitions pertaining to local rank modulation and Gray codes. We will generally follow the notation introduced in [9], [18].

A. Local Rank Modulation

Let us consider a sequence of t real-valued variables, $\mathbf{c} = (c_0, c_1, \dots, c_{t-1}) \in \mathbb{R}^t$, where we further assume $c_i \neq c_j$ for all $i \neq j$. The t variables induce a permutation $f_{\mathbf{c}} \in S_t$, where S_t denotes the set of all permutations over $[t] = \{0, 1, 2, \dots, t-1\}$. The permutation $f_{\mathbf{c}}$ is defined as

$$f_{\mathbf{c}}(i) = |\{j \mid c_j < c_i\}|.$$

Loosely speaking, $f_{\mathbf{c}}(i)$ is the rank of the i th cell in ascending order. This ranking is equivalent to the permutation described in [9], [18], though different.

Given a sequence of n variables, $\mathbf{c} = (c_0, c_1, \dots, c_{n-1})$, we define a window of size t at position p to be

$$\mathbf{c}_{p,t} = (c_p, c_{p+1}, \dots, c_{p+t-1})$$

where the indices are taken modulo n , and also $0 \leq p \leq n-1$, and $1 \leq t \leq n$. We now define the (s,t,n) -local rank-modulation (LRM) scheme, which we do by defining the demodulation process. Let $s \leq t \leq n$ be positive integers, with $s|n$. Given a sequence of n distinct real-valued variables, $\mathbf{c} = (c_0, c_1, \dots, c_{n-1})$, the demodulation maps \mathbf{c} to the sequence of n/s permutations from S_t as follows:

$$\mathbf{f}_{\mathbf{c}} = (f_{\mathbf{c}_{0,t}}, f_{\mathbf{c}_{s,t}}, f_{\mathbf{c}_{2s,t}}, \dots, f_{\mathbf{c}_{n-s,t}}). \quad (1)$$

Loosely speaking, we scan the n variables using windows of size t positioned at multiples of s and write down the permutations from S_t induced by the local views of the sequence.

In the context of flash-memory storage devices, we shall consider the n variables, $\mathbf{c} = (c_0, c_1, \dots, c_{n-1})$, to be the charge-level readings from n flash cells. The demodulated sequence, $\mathbf{f}_{\mathbf{c}}$, will stand for the original information which was stored in the n cells. This approach will serve as the main motivation for this paper, as it was also for [9], [15], [16], [18], [20], [22]. See Figure 1 for an example.

We say a sequence \mathbf{f} of n/s permutations over S_t is (s,t,n) -LRM realizable if there exists $\mathbf{c} \in \mathbb{R}^n$ such that $\mathbf{f} = \mathbf{f}_{\mathbf{c}}$, i.e., it is the demodulated sequence of \mathbf{c} under the (s,t,n) -LRM scheme. Except for the degenerate case of $s = t$, not every sequence is realizable. We denote the set of all (s,t,n) -LRM realizable permutation sequences as $\mathcal{R}(s,t,n)$. In a later part of this section, we show that the number of states representable by an (s,t,n) -LRM scheme, i.e., the size of $\mathcal{R}(s,t,n)$, is roughly $(t \cdot (t-1) \cdot \dots \cdot (t-s+1))^{n/s}$ (this fact is also stated in [22]).

While any $\mathbf{f} \in \mathcal{R}(s,t,n)$ may be represented as a sequence of n/s permutations over S_t , a more succinct representation is possible based on the (mixed-radix) factoradic notation system (see [17] for the earliest-known definition, and [15] for a related use): We can represent any permutation

$f = [f(0), \dots, f(t-1)] \in S_t$ with a sequence of digits $d_{t-1}, d_{t-2}, \dots, d_1, d_0$, where $d_i \in \mathbb{Z}_i$, and d_i counts the number of elements to the right of $f(i)$ which are of lower value. We call d_{t-1} the *most significant digit* and d_0 the *least significant digit*. If $f = f_{\mathbf{c}}$ for some $\mathbf{c} \in \mathbb{R}^t$, then the factoradic representation is easily seen to be equivalent to counting the number of cells to the right of the i th cell which are with lower charge levels.

Continuing with our succinct representation, we now contend that due to the overlap between local views, we can then represent each of the local permutations $f_{c_{i-s,t}}$ using only the s most-significant digits in their factoradic notation. We denote this (partial) representation as $\bar{f}_{c_{i-s,t}}$. Accordingly, we define,

$$\bar{\mathbf{f}}_{\mathbf{c}} = (\bar{f}_{c_{0,t}}, \bar{f}_{c_{s,t}}, \bar{f}_{c_{2s,t}}, \dots, \bar{f}_{c_{n-s,t}}),$$

and the set of all such presentations as $\bar{\mathcal{R}}(s, t, n)$. Thus, for example, the configuration of Figure 1 would be represented by $((3, 0, 1), (4, 2, 0), (0, 2, 2))$.

Lemma 1. For all $1 \leq s \leq t \leq n$,

$$|\bar{\mathcal{R}}(s, t, n)| \leq |\mathcal{R}(s, t, n)| \leq (t-s)! \cdot \left(\frac{t!}{(t-s)!} \right)^{\frac{n}{s}}.$$

Proof: That $|\bar{\mathcal{R}}(s, t, n)| \leq |\mathcal{R}(s, t, n)|$ is trivial, since any $\mathbf{f} \in \mathcal{R}(s, t, n)$ results in one $\bar{\mathbf{f}} \in \bar{\mathcal{R}}(s, t, n)$. For the other inequality, assume we fix the permutation induced by the first $t-s$ cells, where there are $(t-s)!$ ways of doing so. It follows that there are $t!/(t-s)!$ ways of choosing $f_{c_{n-s,t}}$, and then the same amount of ways of choosing $f_{c_{n-2s,t}}$, and continuing all the way up to $f_{c_{0,t}}$ we get the desired result. ■

When $s = t = n$, the (n, n, n) -LRM scheme degenerates into a single permutation from S_n . This was the case in most of the previous works using permutations for modulation purposes. A slightly more general case, $s = t < n$ was discussed by Ferreira *et al.* [10] in the context of permutation trellis codes, where a binary codeword was translated tuple-wise into a sequence of permutation with no overlap between the tuples. An even more general case was defined by Wang *et al.* [22] (though in a slightly different manner where indices are not taken modulo n , i.e., with no wrap-around). In [22], the sequence of permutations was studied under a charge-difference constraint called *bounded rank-modulation*, and mostly with parameters $s = t - 1$, i.e., an overlap of one position between adjacent windows. Finally, using the same terminology as this paper, the case of $(1, 2, n)$ -LRM was considered in [9], [18].

B. Gray Codes

Generally speaking, a *Gray code*, G , is a sequence of distinct states (codewords), $G = g_0, g_1, \dots, g_{N-1}$, from an ambient state space, $g_i \in S$, such that adjacent states in the sequence differ by a “small” change. What constitutes a “small” change usually depends on the code’s application.

Since we are interested in building Gray codes for flash memory devices with the (s, t, n) -LRM scheme, our ambient

space is $\mathcal{R}(s, t, n)$, which is the set of all realizable sequences under (s, t, n) -LRM.

The transition between adjacent states in the Gray code is directly motivated by the flash memory application, and was first described and used in [15], and later also used in [9], [18]. This transition is the “push-to-the-top” operation, which takes a single flash cell and raises its charge level above all others.

In our case, however, since we are considering a *local* rank-modulation scheme, the “push-to-the-top” operation merely raises the charge level of the selected cell above those cells which are comparable with it. Thus, we define the set of allowed transitions as $T = \{\tau_0, \tau_1, \dots, \tau_{n-1}\}$, which is a set of functions, $\tau_j : \mathcal{R}(s, t, n) \rightarrow \mathcal{R}(s, t, n)$, where τ_j represents a “push-to-the-top” operation performed on the j -th cell. More precisely, let \mathbf{f} be an (s, t, n) -LRM realizable sequence of permutations, i.e., there exists $\mathbf{c} \in \mathbb{R}^n$ such that $\mathbf{f} = \mathbf{f}_{\mathbf{c}}$. Now define the transition τ_j acting on \mathbf{f} as $\mathbf{f}' = \mathbf{f}'_{\mathbf{c}'}$, realizable by the variables $\mathbf{c}' = (c'_0, \dots, c'_{n-1}) \in \mathbb{R}^n$ such that c'_j is pushed to a value higher than all of its comparable cells. We denote $r(j)$ as the rightmost index (cyclically) among the cells that share a window with c'_j , and $l(j)$ as the leftmost index (cyclically) among them. We can find $r(j)$ and $l(j)$ by the following expressions:

$$l(j) = s \left\lceil \frac{j-t+1}{s} \right\rceil \bmod n,$$

$$r(j) = \left(s \left\lfloor \frac{j}{s} \right\rfloor + (t-1) \right) \bmod n.$$

Now \mathbf{c}' is given by the following expression:

$$c'_i = \begin{cases} c_i & i \neq j, \\ \max \{c_{l(i)}, \dots, c_{r(i)}\} + 1 & i = j. \end{cases}$$

Definition 2. A Gray code G for (s, t, n) -LRM (denoted (s, t, n) -LRMGC) is a sequence of distinct codewords, $G = g_0, g_1, \dots, g_{N-1}$, where $g_i \in \mathcal{R}(s, t, n)$. For all $0 \leq i \leq N-2$, we further require that $g_{i+1} = \tau_j(g_i)$ for some j . If $g_0 = \tau_j(g_{N-1})$ for some j , then we say the code is cyclic. We call N the size of the code, and say G is optimal if $N = |\mathcal{R}(s, t, n)|$.

Definition 3. We say a family of codes, $\{G_i\}_{i=1}^{\infty}$, where G_i is an (s, t, n_i) -LRMGC of size N_i , $n_{i+1} > n_i$, is asymptotically optimal if

$$\lim_{i \rightarrow \infty} \frac{\log_2 N_i}{\log_2 |\mathcal{R}(s, t, n_i)|} = 1.$$

III. GRAY CODES FOR (s, t, n) -LRM

In this section we present *efficiently* encodable and decodable *asymptotically-optimal* Gray codes for (s, t, n) -LRM. A rough description of our construction follows. First we partition the n cells into m blocks each of size m/n . To simplify our presentation we set $m = \sqrt{n}$, implying that we have m blocks, each of size m . Denote the cells in block i by \mathbf{c}_i . For each block \mathbf{c}_i we will use the factoradic representation $\bar{\mathbf{f}}_{\mathbf{c}_i}$ to represent permutations in $\bar{\mathcal{R}}(s, t, m)$. Namely, each and

every block can be thought of an element of an alphabet $\Sigma = \{v_0, \dots, v_{V-1}\}$ of size V .

Now, consider any de-Bruijn sequence S of order $m-1$ over Σ (of period V^{m-1}). Namely, S will consist of a sequence of V^{m-1} elements $v_{s_0}, v_{s_1}, \dots, v_{s_{V^{m-1}-1}}$ over Σ such that the subsequences $v_{s_i}, \dots, v_{s_{i+m-2}}$ of S cover all $(m-1)$ -tuples of Σ exactly once, sub-indices of s taken modulo V^{m-1} . Here, $s_i \in [V]$. Such sequences S exist, e.g., [11].

We are now ready to construct our Gray code G . The construction will have two phases. First we construct so-called *anchor* elements in G , denoted as $\bar{G} = \{g_0, \dots, g_{L-1}\}$. The elements of \bar{G} will consist of a cyclic Gray code over Σ^m . That is, the difference between each g_i and g_{i+1} in \bar{G} will be in only one out of the m characters (from Σ) in g_i . Specifically, the code \bar{G} will be derived from the de-Bruijn sequence S as follows: we set g_0 to be the first m elements of S , and in the transition from g_i to g_{i+1} we change v_{s_i} to $v_{s_{i+m}}$. The code \bar{G} is detailed below:

$$\begin{aligned} g_0 &= v_{s_{m-1}} \ v_{s_{m-2}} \ \dots \ v_{s_1} \ \underline{v_{s_0}} \\ g_1 &= v_{s_{m-1}} \ v_{s_{m-2}} \ \dots \ \underline{v_{s_1}} \ v_{s_m} \\ g_2 &= v_{s_{m-1}} \ v_{s_{m-2}} \ \dots \ v_{s_{m+1}} \ v_{s_m} \\ &\vdots \\ g_{L-2} &= v_{s_{L-1}} \ \underline{v_{s_{L-2}}} \ \dots \ v_{s_1} \ v_{s_0} \\ g_{L-1} &= \underline{v_{s_{L-1}}} \ v_{s_{m-2}} \ \dots \ v_{s_1} \ v_{s_0} \end{aligned}$$

where $L = \text{lcm}(m, V^{m-1})$, the sub-indices of s are taken modulo V^{m-1} , and the underline is an imaginary marking distinguishing the block which is about to change.

With the imaginary marking of the underline, the code \bar{G} is clearly a Gray code over Σ^m due to the properties of the de-Bruijn sequence S . However \bar{G} does not suffice for our construction as the transitions between the anchors g_i and g_{i+1} involve changing the entries of an entire block, which may involve many push-to-the-top operations. We thus refine \bar{G} by adding additional elements between each pair of adjacent anchors from \bar{G} that allow us to move from the block configuration in g_i to that in g_{i+1} by a series of push-to-the-top operations. Our construction is summarized below formally.

Construction 1. We consider the (s, t, n) -LRM, n be a square, $m = \sqrt{n} \geq t+2$, and require that $s|m$. Let $\{v_0, v_1, \dots, v_{V-1}\}$ be a set of V distinct mixed-radix vectors of length m taken from $([t] \times [t-1] \times \dots \times [t-s])^{m/s}$. The values of the last $s(\lceil (t+2)/s \rceil - 1)$ digits of each v_i do not play a role in the representation of the stored data and are called non-information digits, so by abuse of notation, a mixed-radix vector $(r_0, r_1, \dots, r_{m-1})$ actually represents the value $(r_0, r_1, \dots, r_{m-1-s(\lceil (t+2)/s \rceil - 1)})$ regardless of the value of the last $s(\lceil (t+2)/s \rceil - 1)$ elements. Therefore, we get

$$V = \left(\frac{t!}{(t-s)!} \right)^{\frac{m}{s} - \lceil \frac{t+2}{s} \rceil + 1}.$$

We also denote $L = \text{lcm}(m, V^{m-1})$.

Consider a de-Bruijn sequence S of order $m-1$ over the alphabet $\{0, 1, \dots, V-1\}$. The Gray code \bar{G} of anchor vectors

is a sequence g_0, g_1, \dots, g_{L-1} of L mixed-radix vectors of length $m^2 = n$. Each vector is formed by a concatenation of m blocks of length m . Between the anchors g_i and g_{i+1} , the block v_{s_i} is transformed into the block $v_{s_{i+m}}$.

Within each of the m blocks comprising any single anchor, the $(m-2)$ nd digit (the second-from-right digit – a non-information digit) is set to 1 in all blocks except for the underlined block. For brevity, we call this digit the underline digit.

Between any two anchors, g_i and g_{i+1} , a sequence of vectors called auxiliary vectors and denoted $g_i^0, g_i^1, \dots, g_i^{\ell_i}$, is formed by a sequence of push-to-the-top operations on the cells of the changing block. The auxiliary vectors are determined by Algorithm 1 described shortly.

In what follows we present Algorithm 1 that specifies the sequence $g_i^0, g_i^1, \dots, g_i^{\ell_i}$ that allow us to move from anchor state g_i to state g_{i+1} . As g_i and g_{i+1} differ only in a single block (and this block is changed from v_{s_i} to $v_{s_{i+m}}$), the same will hold for the sequence $g_i^0, g_i^1, \dots, g_i^{\ell_i}$, i.e., g_i^j and $g_i^{j'}$ will only differ in the block in which g_i and g_{i+1} differ. Thus, it suffices to define in Algorithm 1 how to change a block of length m with cell values that represent v_{s_i} into a block that represents $v_{s_{i+m}}$ using push-to-the-top operations. However, we call the attention of the reader to the fact that while the change in represented value affects only one block, for administrative reasons we also push cells of the block to the left (cyclically).

We now present Algorithm 1 and describe some of its properties. We then prove that indeed the resulting code G is an asymptotically-optimal cyclic (s, t, n) -LRMGC. We assume that the following algorithm is applied to positions $\{0, 1, \dots, m-1\}$. We further assume $(r_0, r_1, \dots, r_m) \in ([t] \times [t-1] \times \dots \times [t-s])^{m/s}$ represents the value v_ℓ , then we say the j th digit of v_ℓ is

$$v_\ell(j) = \begin{cases} r_j & 0 \leq j < m - s(\lceil (t+2)/s \rceil - 1) \\ 0 & \text{otherwise.} \end{cases}$$

Finally, we restrict $l(\cdot)$ and $r(\cdot)$ by defining

$$\begin{aligned} l'(j) &= \begin{cases} l(j) & 0 \leq l(j) \leq m-3 \\ 0 & \text{otherwise} \end{cases} \\ r'(j) &= \begin{cases} r(j) & 0 \leq r(j) \leq m-3 \\ m-3 & \text{otherwise} \end{cases} \end{aligned}$$

Our algorithm changes a block of length m with cell values that represent v_{s_i} into one that represents $v_{s_{i+m}}$ using push-to-the-top operations. It is strongly based on the factoradic representation of $v_{s_{i+m}}$. Let $v_{s_{i+m}}(j)$ be the j th entry in this representation. Namely, if $\mathbf{c} = (c_1, \dots, c_m)$ is a cell configuration that corresponds to $v_{s_{i+m}}$, then for each index $j \in [m]$ the number of entries in the window corresponding to j that are to the right of j and are of value lower than c_j equal $v_{s_{i+m}}(j)$. Roughly speaking, to obtain such a configuration \mathbf{c} , our algorithm, for $j \in [m]$, pushes each cell c_j in \mathbf{c} to the top

Algorithm 1 Transform block v_{s_i} to block $v_{s_{i+m}}$

Push the rightmost cell of the block to the left (cyclically)

$a_j \leftarrow 0$ for all $j = 0, 1, \dots, m-3$

$j \leftarrow 0$

repeat

if $v_{s_{i+m}}(j) = \sum_{i=j+1}^{r'(j)} a_i$ and $a_j = 0$ **then**

Push the j th cell of current block.

$a_j \leftarrow 1$

$j \leftarrow l'(j)$

else

$j \leftarrow j + 1$

end if

until $j = m - 2$

Push the next-to-last cell of current block.

exactly once and only after exactly j cells to the right of c_j (and participating in the window corresponding to j) have been pushed to the top. As each time a cell is changed it is pushed to the top, this will ensure that the resulting cell configuration \mathbf{c} will have a factoradic representation corresponding to $v_{s_{i+m}}$.

A few remarks are in place. In order to keep track of which cells were pushed during our algorithm, we save an array of bits a_j for each cell in the block (initialized to 0), indicating whether the cell c_j has been pushed before. We note that in order to be able to decode a state, we need to have some way to know which block is being currently changed, i.e., the imaginary underline in the anchor. We use the last two cells of each block for that purpose.

Example 4. Take the case of $(1, 2, 16)$ -LRM with $m = 4$, $V = 2$, and a de-Bruijn sequence of order 3 and alphabet of size 2 is $S = 00010111$. The list of anchors is

$$\begin{aligned} g_0 &= 1010 & 0010 & 0010 & \underline{0000} \\ g_1 &= 1010 & 0010 & \underline{0000} & 0010 \\ g_2 &= 1010 & \underline{0000} & 1010 & 0010 \\ g_3 &= \underline{1000} & 1010 & 1010 & 0010 \\ g_4 &= 1010 & 1010 & 1010 & \underline{0000} \\ g_5 &= 1010 & 1010 & \underline{1000} & 0010 \\ g_6 &= 1010 & \underline{1000} & 0010 & 0010 \\ g_7 &= \underline{1000} & 0010 & 0010 & 0010 \end{aligned}$$

The bold bit (the leftmost bit in each group of four) denotes the information bit, while the rest are non-information bits. We note that the underlined vectors are easily recognizable by next-to-right bit being 0.

Notice that in this example the information bit is dominated in size by the remaining bits of each block. This is an artifact of our example in which we take n be be small. For large values of n the overhead in each block is negligible with respect to the information bits.

As an example, the transition between g_1 and g_2 is (the

changed positions are underlined)

$$\begin{aligned} g_1 &= 1010 & 0010 & 0000 & 0010 \\ g_1^0 &= 1010 & \underline{0001} & 0000 & 0010 \\ g_1^1 &= 1010 & 0001 & \underline{0100} & 0010 \\ g_1^2 &= 1010 & 0000 & \underline{1100} & 0010 \\ g_2 &= 1010 & 0000 & \underline{1010} & 0010 \end{aligned}$$

□

We now address the analysis of Algorithm 1.

Lemma 5. Assuming the underline is known, all anchors used in Construction 1 are distinct.

Proof: Proof follows directly from the properties of the de-Bruijn sequence S and the fact that we are taking L to be the $\text{lcm}(m, V^{m-1})$. ■

Lemma 6. Algorithm 1 maintains the correctness of the underline digit in anchors. In addition, between any two adjacent anchors, Algorithm 1 guarantees the underline digits of the changing block and the block to its left (cyclically), are both not maximal.

Proof: Proof is by induction. The base case follows from our construction of the first anchor element g_0 . Assume g_i satisfies the inductive claim. When applying Algorithm 1 to move from anchor g_i to g_{i+1} , we start by pushing the rightmost cell of the block to the left of that being changed. This implies that the value of the underline cell in both the block being changed and that to its left are now not maximal. This state of affairs remains until the end of Algorithm 1, in which we push the next-to-the last cell in the changed block. At that point in time, the underline cell in the changed block obtains it's maximal value, while the block to its left (that to be changed in the next application of Algorithm 1) is of non-maximal value. All the other underline cells remain unchanged throughout the execution of Algorithm 1. ■

Lemma 7. Algorithm 1 terminates, and when it does, all of the cells are pushed exactly once.

Proof: That Algorithm 1 terminates is easy to see. For convenience, we denote $z = s(\lceil (t+2)/s \rceil - 1)$. For each $k \in \{m-z, \dots, m-3\}$, $v_{s_{i+m}}(k) = 0$, and therefore each of those cells is pushed the first time that $j = k$. Now we assume by induction that for each $k \leq m-z$, all of the cells with indices j , $k \leq j \leq m-3$, are pushed before the algorithm terminates.

The base case, $k = m-z$, was already proved above. For the induction step, by the induction assumption, we know that all the cells in $\{k, \dots, r'(k)\}$ are pushed. At the point where exactly $v_{s_{i+m}}(k-1)$ of them are pushed, cell $k-1$ is pushed in the next visit. Since also the algorithm never pushes a cell more than once, the claim is proved. ■

Theorem 8. Algorithm 1 changes a block representing v_{s_i} into a block representing $v_{s_{i+m}}$.

Proof: Before cell j is being pushed, exactly $v_{s_{i+m}}(j)$ cells from $\{j+1, \dots, r'(j)\}$ have been pushed already. The

rest will be pushed after and above it, and therefore its rank is exactly $v_{s_i+m}(j)$, as desired. ■

One drawback of Algorithm 1 is that it may visit a codeword multiple times. For example, assume a $(1, 2, 25)$ -LRM scheme, with $v_{s_i} = 11XXX$ and $v_{s_i+5} = 10XXX$, where X is the “don’t care” symbol. The algorithm would, after an initial push of a cell on the adjacent block to the left, first push cell 1, changing the block state to $01XXX$. Afterwards, the algorithm would push cell 0, changing the state back to v_{s_i} .

To solve that problem, we suggest to simulate the entire remaining execution of the algorithm every time we push a cell. If the resulting configuration after the planned push appears another time in the future, we change the algorithm’s inner state to that of the latest such repeat appearance. That way we make sure that each codeword appears only once in the Gray code. We call the revised algorithm the *repetition-avoiding algorithm*.

Lemma 9. *The time complexity of the repetition-avoiding algorithm is $O(tn)$.*

Proof: Each cell is visited by the algorithm at most t times, once during the first visit of the algorithm, and once following each of the $t - 1$ cells immediately to its right being pushed. Since each cell is pushed exactly once, a full execution of the algorithm takes $O(tm)$ steps. For the repetition-avoiding algorithm, simulating a full execution after each push results in total time complexity $O(tm + tm^2) = O(tn)$. ■

Combining all of our observations up to now, we are able to summarize with the following theorem for G from Construction 1.

Theorem 10. *G is a cyclic gray code of size at least L .*

Corollary 11. *For all constants $1 \leq s < t$, there exists an asymptotically-optimal family of codes, $\{G_i\}_{i=t+2}^{\infty}$, where G_i is an (s, t, n_i) -LRMGC of size N_i , $n_{i+1} > n_i$, with*

$$\lim_{i \rightarrow \infty} \frac{\log_2 N_i}{\log_2 |\mathcal{R}(s, t, n_i)|} = 1.$$

Proof: We set $n_i = s^2 i^2$ for all $i \geq t + 2$. Then $N_i \geq L_i \geq V_i^{s^2 i^2}$. It follows that

$$\begin{aligned} & \lim_{i \rightarrow \infty} \frac{\log_2 N_i}{\log_2 |\mathcal{R}(s, t, n_i)|} \geq \\ & \geq \lim_{i \rightarrow \infty} \frac{(si - 1) \log_2 V_i}{\log_2 \left((t-s)! \cdot \left(\frac{t!}{(t-s)!} \right)^{si^2} \right)} \\ & = \lim_{i \rightarrow \infty} \frac{(si - 1) \log_2 \left(\frac{t!}{(t-s)!} \right)^{i - \lceil \frac{t+2}{s} \rceil + 1}}{\log_2 \left((t-s)! \cdot \left(\frac{t!}{(t-s)!} \right)^{si^2} \right)} \\ & = 1. \end{aligned}$$

IV. CONCLUSIONS

We presented the framework for (s, t, n) -local rank modulation, and studied Gray codes for the most general case. The codes we present are asymptotically optimal.

Several questions remain open. For the case of $(1, 2, n)$ -LRM, a previous work describes asymptotically-optimal codes for which the weight of the codewords is constant and approaches $\frac{n}{2}$ [9]. That property guarantees a bounded charge difference in any “push-to-the-top” operation. Constant-weight codes for the general case are still missing. Of more general interest is the study of codes that cover a constant fraction of the space.

REFERENCES

- [1] T. Berger, F. Jelinek, and J. K. Wolf, “Permutation codes for sources,” *IEEE Trans. on Inform. Theory*, vol. IT-18, no. 1, pp. 160–169, Jan. 1972.
- [2] I. F. Blake, “Permutation codes for discrete channels,” *IEEE Trans. on Inform. Theory*, vol. 20, pp. 138–140, 1974.
- [3] I. F. Blake, G. Cohen, and M. Deza, “Coding with permutations,” *Inform. and Control*, vol. 43, pp. 1–19, 1979.
- [4] V. Bohossian, A. Jiang, and J. Bruck, “Buffer coding for asymmetric multi-level memory,” in *Proceedings of the 2007 IEEE International Symposium on Information Theory (ISIT2007), Nice, France, Jun. 2007*, pp. 1186–1190.
- [5] H. Chadwick and I. Reed, “The equivalence of rank permutation codes to a new class of binary codes,” *IEEE Trans. on Inform. Theory*, vol. 16, no. 5, pp. 640–641, 1970.
- [6] H. D. Chadwick and L. Kurz, “Rank permutation group codes based on Kendall’s correlation statistic,” *IEEE Trans. on Inform. Theory*, vol. IT-15, no. 2, pp. 306–315, Mar. 1969.
- [7] G. Cohen and M. Deza, “Decoding of permutation codes,” in *Intl. CNRS Colloquium, July, France, 1977*.
- [8] M. Deza and P. Frankl, “On maximal numbers of permutations with given maximal or minimal distance,” *J. Combin. Theory Ser. A*, vol. 22, 1977.
- [9] E. En Gad, M. Langberg, M. Schwartz, and J. Bruck, “On a construction for constant-weight gray codes for local rank modulation,” in *Proceedings of the 2010 IEEE 26-th Convention of Electrical and Electronic Engineers in Israel (IEEEI2010), Eilat, Israel, Nov. 2010*, p. 996.
- [10] H. C. Ferreira, A. J. H. Vinck, T. G. Swart, and I. de Beer, “Permutation trellis codes,” *IEEE Trans. on Communications*, pp. 1782–1789, Nov. 2005.
- [11] S. W. Golomb, *Shift Register Sequences*. Holden-Day, San Francisco, 1967.
- [12] A. Jiang, V. Bohossian, and J. Bruck, “Floating codes for joint information storage in write asymmetric memories,” in *Proceedings of the 2007 IEEE International Symposium on Information Theory (ISIT2007), Nice, France, Jun. 2007*, pp. 1166–1170.
- [13] A. Jiang and J. Bruck, “Joint coding for flash memory storage,” in *Proceedings of the 2008 IEEE International Symposium on Information Theory (ISIT2008), Toronto, Canada, Jul. 2008*, pp. 1741–1745.
- [14] A. Jiang, M. Langberg, M. Schwartz, and J. Bruck, “Universal rewriting in constrained memories,” in *Proceedings of the 2009 IEEE International Symposium on Information Theory (ISIT2009), Seoul, Korea, Jun. 2009*, pp. 1219–1223.
- [15] A. Jiang, R. Mateescu, M. Schwartz, and J. Bruck, “Rank modulation for flash memories,” *IEEE Trans. on Inform. Theory*, vol. 55, no. 6, pp. 2659–2673, Jun. 2009.
- [16] A. Jiang, M. Schwartz, and J. Bruck, “Correcting charge-constrained errors in the rank-modulation scheme,” *IEEE Trans. on Inform. Theory*, vol. 56, no. 5, pp. 2112–2120, May 2010.
- [17] C. A. Laisant, “Sur la numération factorielle, application aux permutations,” *Bulletin de la Société Mathématique de France*, vol. 16, pp. 176–183, 1888.
- [18] M. Schwartz, “Constant-weight Gray codes for local rank modulation,” in *Proceedings of the 2010 IEEE International Symposium on Information Theory (ISIT2010), Austin, TX, U.S.A., Jun. 2010*, pp. 869–873.

- [19] D. Slepian, "Permutation modulation," in *Proc. of the IEEE*, vol. 53, no. 3, 1965, pp. 228–236.
- [20] I. Tamo and M. Schwartz, "Correcting limited-magnitude errors in the rank-modulation scheme," *IEEE Trans. on Inform. Theory*, vol. 56, no. 6, pp. 2551–2560, Jun. 2010.
- [21] H. Vinck, J. Haering, and T. Wadayama, "Coded M-FSK for power line communications," in *Proceedings of the 2000 IEEE International Symposium on Information Theory (ISIT2000), Sorrento, Italy, 2000*, p. 137.
- [22] Z. Wang, A. Jiang, and J. Bruck, "On the capacity of bounded rank modulation for flash memories," in *Proceedings of the 2009 IEEE International Symposium on Information Theory (ISIT2009), Seoul, Korea, Jun. 2009*, pp. 1234–1238.
- [23] E. Yaakobi, P. H. Siegel, and J. K. Wolf, "Buffer codes for multi-level flash memory," in *Proceedings of the 2008 IEEE International Symposium on Information Theory (ISIT2008), Toronto, Canada, 2008*, poster.