

Snake-in-the-Box Codes for Rank Modulation

Yonatan Yehezkeally and Moshe Schwartz, *Senior Member, IEEE*

Abstract—Motivated by the rank-modulation scheme with applications to flash memory, we consider Gray codes capable of detecting a single error, also known as snake-in-the-box codes. We study two error metrics: Kendall’s τ -metric, which applies to charge-constrained errors, and the ℓ_∞ -metric, which is useful in the case of limited-magnitude errors. In both cases, we construct snake-in-the-box codes with rate asymptotically tending to 1. We also provide efficient successor-calculation functions, as well as ranking and unranking functions. Finally, we also study bounds on the parameters of such codes.

Index Terms—Flash memory, permutations, rank modulation, snake-in-the-box codes.

I. INTRODUCTION

FLASH memory is a nonvolatile storage medium which is electrically programmable and erasable. Its current wide use is motivated by its high storage density and relative low cost. Among the chief disadvantages of flash memories is their inherent asymmetry between cell programming (injecting cells with charge) and cell erasure (removing charge from cells). While single cells can be programmed with relative ease, in the current architecture, the process of erasure can only be performed by completely depleting large blocks of cells of their charge. Moreover, the removal of charge from cells physically damages them over time.

This issue is exacerbated as a result of the ever-present demand for denser memory: smaller cells are more delicate, and are damaged faster during erasure. They also contain less charge and are, therefore, more prone to error. In addition, flash memories, at present, use multilevel cells, where charge levels are quantized to simulate a finite alphabet—the more levels, the less safety margins are left, and data integrity is compromised. Thus, overprogramming (increasing a cell’s charge level above the designated mark) is a real problem, requiring a costly and damaging erasure cycle. Hence, in a programming cycle, charge levels are usually made to gradually approach the desirable amount, making for lengthier programming cycles as well (see [3]).

In an effort to counter these effects, a different modulation scheme has recently been suggested for flash memories—rank modulation [12]. This scheme calls for the representation of the

data stored in a group of cells in the permutation suggested by their relative charge levels. That is, if $c_1, c_2, \dots, c_n \in \mathbb{R}$ represent the charge levels of $n \in \mathbb{N}$ cells, then that group of cells is said to encode that permutation $\sigma \in S_n$ such that

$$c_{\sigma(1)} > c_{\sigma(2)} > \dots > c_{\sigma(n)} > 0.$$

This scheme eliminates the need for discretization of charge levels. Furthermore, it was suggested in [12] that programming could be restricted to “push-to-the-top” operations. In this scheme, one only programs a group of cells by increasing the charge level of a single cell above that of all others. In this manner, overprogramming is no longer an issue.

In addition, storing data using this scheme also improves the memory’s robustness against other noise types. Retention, the process of slow charge leakage from cells, tends to affect all cells with a similar trend [3]. Since rank modulation stores information in the differences between charge levels rather than their absolute values, it offers more resilience against that type of noise. It is also worth noting that the advantages of rank modulation have been experimentally applied to phase-change memory (see [18]).

Gray codes using “push-to-the-top” operations and spanning the entire space of permutations were also studied in [12]. The Gray code [9] was first introduced as a sequence of distinct binary vectors of fixed length, where every adjacent pair differs in a single coordinate. It has since been generalized to sequences of distinct states $s_1, s_2, \dots, s_k \in S$ such that, for every $i < k$, there exists a function in a predetermined set of transitions $t \in T$ such that $s_{i+1} = t(s_i)$ (see [19] for an excellent survey). When the states one considers are permutations on $n \in \mathbb{N}$ elements and the allowed transitions are “push-to-the-top” operations, Jiang *et al.* [12] referred to such Gray codes as n -length rank-modulation Gray codes (n -RMGC’s), and it presented such codes traversing the entire set of permutations. In this fashion, a set of n rank-modulation cells could implement a single logical multilevel cell with $n!$ levels, where increasing the logical cell’s level by 1 corresponds to a single transition in the n -RMGC. This allows for a natural integration of rank modulation with other multilevel approaches such as rewriting schemes [4], [10], [11], [24].

Other recent works have explored error-correcting codes for rank modulation, where different types of errors are addressed by a careful choice of metric. In [2], [13], and [17], Kendall’s τ -metric was considered, since a small charge-constrained error translates into a small distance in the metric. In contrast, the ℓ_∞ -metric was used in [15] and [22], as small distances in this metric correspond to small limited-magnitude errors.

In this paper, we explore Gray codes for rank modulation which detect a single error, under both metrics mentioned previously. Such codes are known as snake-in-the-box codes, and

Manuscript received July 18, 2011; revised January 25, 2012 and April 16, 2012; accepted April 16, 2012. Date of publication April 27, 2012; date of current version July 10, 2012. This work was supported in part by the Israel Science Foundation under Grant 134/10. This paper was presented in part at the 2012 IEEE International Symposium on Information Theory.

The authors are with the Department of Electrical and Computer Engineering, Ben-Gurion University of the Negev, Beer Sheva 84105, Israel (e-mail: yonatan@bgu.ac.il; schwartz@ee.bgu.ac.il).

Communicated by N. Kashyap, Associate Editor for Coding Theory.

Digital Object Identifier 10.1109/TIT.2012.2196755

have been studied extensively for binary vectors with the Hamming metric and with single-bit flips as allowable transitions (see [1] and references therein).

This paper is organized as follows. In Section II, we present basic notation and definitions. In Section III, we review properties of Kendall's τ -metric, present a recursive construction of snake-in-the-box codes over the alternating groups of odd orders with rate asymptotically tending to 1, then present some upper bounds on the size of such snake-in-the-box codes in general, and conclude by presenting auxiliary functions needed for the use of codes generated by this construction. In Section IV, we present a direct construction of snake-in-the-box codes of every order in the ℓ_∞ -metric based on results from [12], with rates that asymptotically tend to 1. We conclude in Section V with some ad hoc results, as well as some open questions.

II. PRELIMINARIES

Given a permutation σ on n elements (i.e., a bijection from and into the set $[n] = \{1, 2, \dots, n\}$), we shall denote it by $\sigma = [\sigma(1), \sigma(2), \dots, \sigma(n)]$. This form is called the *vector notation* for permutations. We let S_n be the symmetric group on $[n]$ (i.e., the group of all permutations on $[n]$). For $\sigma, \tau \in S_n$, their composition, denoted $\sigma\tau$, is the permutation for which $\sigma\tau(i) = \sigma(\tau(i))$ for all $i \in [n]$. It is well known that $|S_n| = n!$.

Example 1: One has precisely six ways of organizing the elements of [3] in a row. These are

$$[1, 2, 3], [1, 3, 2], [2, 3, 1], [2, 1, 3], [3, 1, 2], [3, 2, 1].$$

These six permutation form the group S_3 . \square

A cycle, denoted (a_1, a_2, \dots, a_k) , is a permutation mapping $a_i \mapsto a_{i+1}$ for all $i \in [k-1]$, as well as $a_k \mapsto a_1$. We shall occasionally use *cycle notation* in which a permutation is described as a composition of cycles. We also recall that any permutation may be represented as a composition of cycles of size two (known as *transpositions*), and that the parity of the number of transpositions does not depend on the decomposition. Thus, we have *even* and *odd* permutations, with positive and negative *signs*, respectively. We let A_n be the subgroup of all even permutations on $[n]$, called the alternating group of order n . Again, it is well known that $|A_n| = \frac{1}{2}|S_n|$.

Example 2: Of the permutations presented in Example 1, only the following are even:

$$[1, 2, 3], [2, 3, 1], [3, 1, 2].$$

They form the group A_3 . Put in cycle notation, they are

$$\text{id}, (1, 2, 3), (1, 3, 2)$$

where id denotes the identity permutation. \square

Definition 3: Given a set S and a subset of transformations $T \subseteq S^S = \{f \mid f: S \rightarrow S\}$, a *Gray code* over S , using transitions T , of size $M \in \mathbb{N}$, is a sequence $C = (c_0, c_1, \dots, c_{M-1})$ of M distinct elements of S , called *codewords*, such that for all $j \in [M-1]$ there exists $t \in T$ such that $c_j = t(c_{j-1})$.

Alternatively, when the original permutation c_0 is known (or irrelevant), we use a slight abuse of notation in referring to

the sequence of transformations $(t_{k_1}, \dots, t_{k_{M-1}})$ generating the code (i.e., $c_j = t_{k_j}(c_{j-1})$) as the code itself.

In the aforementioned definition, when $M = |S|$, the Gray code is called *complete*. If there exists $t \in T$ such that $t(c_{M-1}) = c_0$, the Gray code is called *cyclic*, M is called its *period*, and we shall, when listing the code by its sequence of transformations, include $t_{k_M} = t$ at the end of the list. The *rate* of C , denoted $R(C)$, is defined as

$$R(C) = \frac{\log_2 M}{\log_2 |S|}.$$

In the context of rank modulation for flash memories, the set of transformations T comprises of “push-to-the-top” operations, first used in [12], and later also in [8] and [23]. We denote by $t_i: S_n \rightarrow S_n$ the “push-to-the-top” operation on index i , i.e.,

$$\begin{aligned} t_i[a_1, a_2, \dots, a_{i-1}, a_i, a_{i+1}, \dots, a_n] \\ = [a_i, a_1, a_2, \dots, a_{i-1}, a_{i+1}, \dots, a_n] \end{aligned}$$

and throughout the paper we set $T = \{t_2, t_3, \dots, t_n\}$. We also note that in cycle notation

$$t_i\sigma = \sigma(i, i-1, \dots, 1). \quad (1)$$

For ease of presentation only, we also denote by \underline{t}_i the “push-to-the-bottom” operation on index $n+1-i$, i.e.,

$$\begin{aligned} \underline{t}_i[a_1, a_2, \dots, a_{n-i}, a_{n+1-i}, a_{n+2-i}, \dots, a_n] \\ = [a_1, a_2, \dots, a_{n-i}, a_{n+2-i}, \dots, a_n, a_{n+1-i}]. \end{aligned}$$

Restricting the transformations to “push-to-the-top” operations allows fast cell programming, and eliminates overshoots (see [12]). In the context of flash memory, “push-to-the-top” operations have also been used in [6] and [7]. We also note that generating permutations using “push-to-the-top” operations is of independent interest, called “nested cycling” in [21] (see also references therein), motivated by a fast “push-to-the-top” operation¹ (cycling) available on some computer architectures.

Let $d: S \times S \rightarrow \mathbb{N} \cup \{0\}$ be a distance function inducing a metric \mathcal{M} over S . Given a transmitted codeword $c \in C$ and its received version $\tilde{c} \in S$, we say a single error occurred if $d(c, \tilde{c}) = 1$. We are interested in Gray codes capable of detecting single errors, which we now define.

Definition 4: Let \mathcal{M} be a metric over S induced by a distance measure d . A *snake-in-the-box code* over \mathcal{M} and S , using transitions T , is a Gray code C over S and using T , in which for every pair of distinct elements $c, c' \in C$, $c \neq c'$, one has $d(c, c') \geq 2$.

Since throughout the paper our ambient space is S_n , and the transformations we use are the “push-to-the-top” operations T , we shall abbreviate our notation and call the snake-in-the-box code of size M an (n, M, \mathcal{M}) -snake, or an \mathcal{M} -snake. We will be considering two metrics in the next sections: Kendall's τ -metric \mathcal{K} and the ℓ_∞ -metric, with their respective \mathcal{K} -snakes and ℓ_∞ -snakes.

¹The operations described in [21] are actually mirror images of “push-to-the-top.” Furthermore, in the permutation-generation scheme, there is not a Gray code since it repeats some of the previously generated permutations, also making it inefficient.

It is interesting to note that the classical definition of snake-in-the-box codes (see the survey [1]) is slightly weaker in the sense that $d(c, c') \geq 2$ is required for distinct $c, c' \in C$, unless c and c' are adjacent in C . This, however, is a compromise due to the fact that in the classical codes over binary vectors, the transformations (which flip a single bit) always create adjacent codewords at distance 1 apart. This compromise is unnecessary in our case since, as we shall later see, the “push-to-the-top” operations allow adjacent words at distance 2 or more apart.

III. KENDALL’S τ -METRIC AND \mathcal{K} -SNAKES

Kendall’s τ -metric [14], denoted \mathcal{K} , is induced by the bubble-sort distance which measures the minimal amount of adjacent transpositions required to transform one permutation into the other. For example, the distance between the permutations $[2, 1, 4, 3]$ and $[2, 4, 3, 1]$ is 2, as

$$[2, 1, 4, 3] \rightarrow [2, 4, 1, 3] \rightarrow [2, 4, 3, 1]$$

is a shortest sequence of adjacent transpositions between the two. More formally, for $\alpha, \beta \in S_n$, as noted in [13]

$$d_K(\alpha, \beta) = |\{(i, j) \mid \alpha(i) < \alpha(j) \wedge \beta(i) > \beta(j)\}|.$$

The metric \mathcal{K} was first introduced by Kendall and Gibbons [14] in the study of ranking in statistics. It was observed in [13] that a bounded distance in Kendall’s τ -metric models errors caused by bounded changes in charge levels of cells in the flash memory. Error-correcting codes for this metric were studied in [2], [13], and [17].

We let Kendall’s τ adjacency graph of order $n \in \mathbb{N}$ be the graph $G_n = (V_n, E_n)$ whose vertices are the elements of the symmetric group $V_n = S_n$, and $\{\alpha, \beta\} \in E_n$ if and only if $d_K(\alpha, \beta) = 1$. It is well known that Kendall’s τ -metric is *graphic* [5], i.e., for every $\alpha, \beta \in S_n$, $d_K(\alpha, \beta)$ equals the length of the shortest path between the two in the adjacency graph, G_n .

A. Construction

We begin the construction process by restricting ourselves to Gray codes using only “push-to-the-top” operations on odd indices. The following lemma provides the motivation for this restriction.

Lemma 5: A Gray code over S_n using only “push-to-the-top” operations on odd indices is a \mathcal{K} -snake.

Proof: According to (1), a “push-to-the-top” operation on an odd index is a composition with an odd-length cycle (which is an even permutation). Thus, the codewords in a Gray code using only such operations are all with the same sign.

On the other hand, an adjacent transposition is an odd permutation, thus, flipping the sign of the permutation it acts on. It follows that in a list of codewords, all with the same sign, there are no two codewords which are adjacent in G_n , i.e., the Gray code is a \mathcal{K} -snake. ■

Lemma 5 saves us the need to check whether a Gray code is in fact a \mathcal{K} -snake, at the cost of restricting the set of allowed tran-

sitions (and the size of the resulting code, although Theorems 15 and 17, presented as follows, work to mitigate this concern). In particular, if n is even, the last element cannot be moved.

By starting with an even permutation, and using only “push-to-the-top” operations on odd indices, we get a sequence of even permutations. Thus, throughout this part, the context of the alternating group A_{2n+1} is assumed, where $n \in \mathbb{N}$.

The construction we are about to present is recursive in nature. As a base for the recursion, we note that three consecutive “push-to-the-top” operations on the third index of permutations in A_3 constitute a complete cyclic $(3, 3, \mathcal{K})$ -snake

$$C_3 = ([1, 2, 3], [3, 1, 2], [2, 3, 1]).$$

We shall extend C_3 to the next order as a running example alongside the general construction in the following.

Now, assume that there exists a cyclic $(2n - 1, M_{2n-1}, \mathcal{K})$ -snake C_{2n-1} , and let

$$t_{k_1}, t_{k_2}, \dots, t_{k_{M_{2n-1}}}$$

be the sequence of transformations generating it, where k_j is odd for all $j \in [M_{2n-1}]$. We also assume that $k_1 = 2n - 1$ (this requirement, while perhaps appearing arbitrary, is actually quite easily satisfied. Indeed, every sufficiently large cyclic \mathcal{K} -snake over S_{2n-1} must, without loss of generality, satisfy it. We shall make it a point to demonstrate that this holds for our construction).

We fix arbitrary values for $a_0, a_1, \dots, a_{2n-2}$ such that

$$\{a_0, a_1, \dots, a_{2n-2}\} = [2n + 1] \setminus \{1, 3\}. \quad (2)$$

For all $i \in [2n - 1]$, we define

$$\sigma_0^{(i)} = [1, a_i, 3, a_{i+1}, \dots, a_{i+2n-2}]$$

where the indices are taken modulo $2n - 1$, and such that we indeed have $\sigma_0^{(i)} \in A_{2n+1}$, i.e., $\sigma_0^{(i)}$ is an even permutation (one simple way of achieving this is to choose them in ascending order).

Example 6: We recall that C_3 is generated by the operations (t_3, t_3, t_3) , which satisfy our requirement. As suggested previously, we order $[5] \setminus \{1, 3\}$ in ascending order, i.e.,

$$(a_0, a_1, a_2) = (2, 4, 5).$$

We define the following permutations as starting points for our construction:

$$\begin{aligned} \sigma_0^{(0)} &= \sigma_0^{(3)} = [1, 2, 3, 4, 5] \\ \sigma_0^{(1)} &= [1, 4, 3, 5, 2] \\ \sigma_0^{(2)} &= [1, 5, 3, 2, 4] \end{aligned}$$

and readily verify that they are all even. □

We now define, for all $i \in [2n - 1]$ and $j \in [M_{2n-1}]$, the permutation

$$\sigma_{j(2n+1)}^{(i)} = t_{k_j} \left(\sigma_{(j-1)(2n+1)}^{(i)} \right)$$

i.e., we construct cycles corresponding to a mirror view of C_{2n-1} on all but the two first elements of $\sigma_0^{(i)}$ (which, as we recall, are $(1, a_i)$).

Example 7: In our running example, we define the following permutations:

$$\begin{aligned} \sigma_5^{(0)} = t_3 \sigma_0^{(0)} &= [1, 2, 4, 5, 3] & \sigma_5^{(1)} = t_3 \sigma_0^{(1)} &= [1, 4, 5, 2, 3] \\ \sigma_{10}^{(0)} = t_3 \sigma_5^{(0)} &= [1, 2, 5, 3, 4] & \sigma_{10}^{(1)} = t_3 \sigma_5^{(1)} &= [1, 4, 2, 3, 5] \\ \sigma_{15}^{(0)} = t_3 \sigma_{10}^{(0)} &= [1, 2, 3, 4, 5] & \sigma_{15}^{(1)} = t_3 \sigma_{10}^{(1)} &= [1, 4, 3, 5, 2] \\ \sigma_5^{(2)} = t_3 \sigma_0^{(2)} &= [1, 5, 2, 4, 3] \\ \sigma_{10}^{(2)} = t_3 \sigma_5^{(2)} &= [1, 5, 4, 3, 2] \\ \sigma_{15}^{(2)} = t_3 \sigma_{10}^{(2)} &= [1, 5, 3, 2, 4] \end{aligned}$$

and resume our construction. \square

We now note the following properties of our construction.

Lemma 8: Let $i, k \in [2n - 1]$, and $j, l \in [M_{2n-1}]$. The following are equivalent.

- 1) The permutations $\sigma_{j(2n+1)}^{(i)}$ and $\sigma_{l(2n+1)}^{(k)}$ are cyclic shifts of each other.
- 2) $\sigma_{j(2n+1)}^{(i)} = \sigma_{l(2n+1)}^{(k)}$.
- 3) $i = k$ and $j = l$.

Proof: First, if $\sigma_{j(2n+1)}^{(i)}$ is a cyclic shift of $\sigma_{l(2n+1)}^{(k)}$, since

$$\sigma_{j(2n+1)}^{(i)}(1) = 1 = \sigma_{l(2n+1)}^{(k)}(1)$$

then necessarily

$$\sigma_{j(2n+1)}^{(i)} = \sigma_{l(2n+1)}^{(k)}.$$

It then follows that

$$a_i = \sigma_{j(2n+1)}^{(i)}(2) = \sigma_{l(2n+1)}^{(k)}(2) = a_k$$

hence $i = k$. Moreover, since the two permutations' last $n - 1$ elements agree, and $t_{k_1}, t_{k_2}, \dots, t_{k_{M_{2n-1}}}$ induce a Gray code, we necessarily have $j = l$.

Finally, the last statement implies that the first is trivial. \blacksquare

Lemma 9: For all $i \in [2n - 1]$, it holds that

$$\sigma_{M_{2n-1}(2n+1)}^{(i)} = \sigma_0^{(i)}.$$

Proof: The transformations $t_{k_1}, t_{k_2}, \dots, t_{k_{M_{2n-1}}}$ induce a cyclic code, and the claim follows directly. \blacksquare

Therefore, we have constructed $2n - 1$ cycles comprised of cyclically nonequivalent permutations (although, at this point, they are not generated by ‘‘push-to-the-top’’ operations).

It shall now be noted that

$$t_k = t_{2n+1}^{2n} t_{2n+2-k}.$$

Hence, if we define for all $i \in [2n - 1]$, $0 \leq j < M_{2n-1}$, and $1 < m \leq 2n$, the permutations

$$\begin{aligned} \sigma_{j(2n+1)+1}^{(i)} &= t_{2n+2-k_{j+1}} \sigma_{j(2n+1)}^{(i)} \\ \sigma_{j(2n+1)+m}^{(i)} &= t_{2n+1}^{m-1} \sigma_{j(2n+1)+1}^{(i)} \end{aligned}$$

then it holds that

$$\sigma_{(j+1)(2n+1)}^{(i)} = t_{2n+1} \sigma_{j(2n+1)+2n}^{(i)}.$$

Our observation from the earlier paragraph means that, at this point, we have $2n - 1$ disjoint cycles, which we conveniently denote

$$C_{2n+1}^{(i)} = \left(\sigma_0^{(i)}, \sigma_1^{(i)}, \dots, \sigma_{M_{2n-1}(2n+1)-1}^{(i)} \right)$$

for all $i \in [2n - 1]$ (for ease of notation, we let $C_{2n+1}^{(0)} = C_{2n+1}^{(2n-1)}$).

Example 10: In our construction, the cycles we produced are

$$\begin{aligned} \sigma_0^{(0)} = t_5 \sigma_{14}^{(0)} &= [1, 2, 3, 4, 5] & \sigma_0^{(1)} = t_5 \sigma_{14}^{(1)} &= [1, 4, 3, 5, 2] \\ \sigma_1^{(0)} = t_3 \sigma_0^{(0)} &= [3, 1, 2, 4, 5] & \sigma_1^{(1)} = t_3 \sigma_0^{(1)} &= [3, 1, 4, 5, 2] \\ \sigma_2^{(0)} = t_5 \sigma_1^{(0)} &= [5, 3, 1, 2, 4] & \sigma_2^{(1)} = t_5 \sigma_1^{(1)} &= [2, 3, 1, 4, 5] \\ \sigma_3^{(0)} = t_5 \sigma_2^{(0)} &= [4, 5, 3, 1, 2] & \sigma_3^{(1)} = t_5 \sigma_2^{(1)} &= [5, 2, 3, 1, 4] \\ \sigma_4^{(0)} = t_5 \sigma_3^{(0)} &= [2, 4, 5, 3, 1] & \sigma_4^{(1)} = t_5 \sigma_3^{(1)} &= [4, 5, 2, 3, 1] \\ \sigma_5^{(0)} = t_5 \sigma_4^{(0)} &= [1, 2, 4, 5, 3] & \sigma_5^{(1)} = t_5 \sigma_4^{(1)} &= [1, 4, 5, 2, 3] \\ \sigma_6^{(0)} = t_3 \sigma_5^{(0)} &= [4, 1, 2, 5, 3] & \sigma_6^{(1)} = t_3 \sigma_5^{(1)} &= [5, 1, 4, 2, 3] \\ \sigma_7^{(0)} = t_5 \sigma_6^{(0)} &= [3, 4, 1, 2, 5] & \sigma_7^{(1)} = t_5 \sigma_6^{(1)} &= [3, 5, 1, 4, 2] \\ \sigma_8^{(0)} = t_5 \sigma_7^{(0)} &= [5, 3, 4, 1, 2] & \sigma_8^{(1)} = t_5 \sigma_7^{(1)} &= [2, 3, 5, 1, 4] \\ \sigma_9^{(0)} = t_5 \sigma_8^{(0)} &= [2, 5, 3, 4, 1] & \sigma_9^{(1)} = t_5 \sigma_8^{(1)} &= [4, 2, 3, 5, 1] \\ \sigma_{10}^{(0)} = t_5 \sigma_9^{(0)} &= [1, 2, 5, 3, 4] & \sigma_{10}^{(1)} = t_5 \sigma_9^{(1)} &= [1, 4, 2, 3, 5] \\ \sigma_{11}^{(0)} = t_3 \sigma_{10}^{(0)} &= [5, 1, 2, 3, 4] & \sigma_{11}^{(1)} = t_3 \sigma_{10}^{(1)} &= [2, 1, 4, 3, 5] \\ \sigma_{12}^{(0)} = t_5 \sigma_{11}^{(0)} &= [4, 5, 1, 2, 3] & \sigma_{12}^{(1)} = t_5 \sigma_{11}^{(1)} &= [5, 2, 1, 4, 3] \\ \sigma_{13}^{(0)} = t_5 \sigma_{12}^{(0)} &= [3, 4, 5, 1, 2] & \sigma_{13}^{(1)} = t_5 \sigma_{12}^{(1)} &= [3, 5, 2, 1, 4] \\ \sigma_{14}^{(0)} = t_5 \sigma_{13}^{(0)} &= [2, 3, 4, 5, 1] & \sigma_{14}^{(1)} = t_5 \sigma_{13}^{(1)} &= [4, 3, 5, 2, 1] \end{aligned}$$

$$\begin{aligned} \sigma_0^{(2)} = t_5 \sigma_{14}^{(2)} &= [1, 5, 3, 2, 5] \\ \sigma_1^{(2)} = t_3 \sigma_0^{(2)} &= [3, 1, 5, 2, 4] \\ \sigma_2^{(2)} = t_5 \sigma_1^{(2)} &= [4, 3, 1, 5, 2] \\ \sigma_3^{(2)} = t_5 \sigma_2^{(2)} &= [2, 4, 3, 1, 5] \\ \sigma_4^{(2)} = t_5 \sigma_3^{(2)} &= [5, 2, 4, 3, 1] \\ \sigma_5^{(2)} = t_5 \sigma_4^{(2)} &= [1, 5, 2, 4, 3] \\ \sigma_6^{(2)} = t_3 \sigma_5^{(2)} &= [2, 1, 5, 4, 3] \\ \sigma_7^{(2)} = t_5 \sigma_6^{(2)} &= [3, 2, 1, 5, 4] \\ \sigma_8^{(2)} = t_5 \sigma_7^{(2)} &= [4, 3, 2, 1, 5] \\ \sigma_9^{(2)} = t_5 \sigma_8^{(2)} &= [5, 4, 3, 2, 1] \\ \sigma_{10}^{(2)} = t_5 \sigma_9^{(2)} &= [1, 5, 4, 3, 2] \\ \sigma_{11}^{(2)} = t_3 \sigma_{10}^{(2)} &= [4, 1, 5, 3, 2] \\ \sigma_{12}^{(2)} = t_5 \sigma_{11}^{(2)} &= [2, 4, 1, 5, 3] \\ \sigma_{13}^{(2)} = t_5 \sigma_{12}^{(2)} &= [3, 2, 4, 1, 5] \\ \sigma_{14}^{(2)} = t_5 \sigma_{13}^{(2)} &= [5, 3, 2, 4, 1] \end{aligned}$$

where the permutations in bold are those from Example 7. \square

Each of the cycles is of size $(2n + 1)M_{2n-1}$, is generated by ‘‘push-to-the-top’’ operations, and contains all cyclic shifts of elements present in our previous version of that cycle. We merge these cycles into a single cycle in the following theorem.

Theorem 11: Given a cyclic $(2n - 1, M_{2n-1}, \mathcal{K})$ -snake using only ‘‘push-to-the-top’’ operations on odd indices, and such that its first transformation is t_{2n-1} , there exists a cyclic $(2n + 1, M_{2n+1}, \mathcal{K})$ -snake with the same properties, whose size is

$$M_{2n+1} = (2n - 1)(2n + 1)M_{2n-1}.$$

Proof: Since $k_1 = 2n - 1$, it holds for all $i \in [2n - 1]$ that $\sigma_1^{(i)} = t_3 \sigma_0^{(i)}$, and we recall $\sigma_2^{(i)} = t_{2n+1} \sigma_1^{(i)}$. More explicitly

$$\begin{aligned}\sigma_1^{(i)} &= [3, 1, a_i, a_{i+1}, \dots, a_{i+2n-2}] \\ \sigma_2^{(i)} &= [a_{i+2n-2}, 3, 1, a_i, a_{i+1}, \dots, a_{i+2n-3}]\end{aligned}$$

where, again, the indices are taken modulo $2n - 1$. Thus, for all $i \in [2n - 2]$, we have

$$t_3 \sigma_1^{(i)} = [a_i, 3, 1, a_{i+1}, \dots, a_{i+2n-2}] = \sigma_2^{(i+1)}$$

and $t_3 \sigma_1^{(2n-1)} = \sigma_2^{(1)}$.

Let E denote the left-shift operator, and so

$$E^2 C_{2n+1}^{(i)} = \left(\sigma_2^{(i)}, \sigma_3^{(i)}, \dots, \sigma_{M_{2n+1}(2n+1)-1}^{(i)}, \sigma_0^{(i)}, \sigma_1^{(i)} \right).$$

By the aforementioned observations, we conclude that

$$C_{2n+1} = E^2 C_{2n+1}^{(0)}, E^2 C_{2n+1}^{(1)}, \dots, E^2 C_{2n+1}^{(2n-2)}$$

is a cyclic $(2n + 1, M_{2n+1}, \mathcal{K})$ -snake, consisting of

$$M_{2n+1} = (2n - 1)(2n + 1)M_{2n-1}$$

permutations. The code C_{2n+1} obviously uses t_{2n+1} , and so some cyclic shift of it has it as its first transition (in fact, for every $i \in [2n - 1]$, one has $\sigma_3^{(i)} = t_{2n+1} \sigma_2^{(i)}$, and in particular, $E^2 C_{2n+1}^{(0)}$ has t_{2n+1} as its first transition, and so does C_{2n+1}). Finally, it is easily verifiable that all ‘‘push-to-the-top’’ operations are on odd indices. ■

Example 12: Our running example ends with the full construction of a $(5, 45, \mathcal{K})$ -snake C_5 from Theorem 11. The down arrows stand for an omitted sequence of t_5 transformations. The transition from column to column uses a single t_3 transformation

[5,3,1,2,4]	$\sigma_2^{(0)}$	[2,3,1,4,5]	$\sigma_2^{(1)}$	[4,3,1,5,2]	$\sigma_2^{(2)}$
↓	↓	↓	↓	↓	↓
[1,2,4,5,3]	$\sigma_5^{(0)}$	[1,4,5,2,3]	$\sigma_5^{(1)}$	[1,5,2,4,3]	$\sigma_5^{(2)}$
[4,1,2,5,3]	$\sigma_6^{(0)}$	[5,1,4,2,3]	$\sigma_6^{(1)}$	[2,1,5,4,3]	$\sigma_6^{(2)}$
↓	↓	↓	↓	↓	↓
[1,2,5,3,4]	$\sigma_{10}^{(0)}$	[1,4,2,3,5]	$\sigma_{10}^{(1)}$	[1,5,4,3,2]	$\sigma_{10}^{(2)}$
[5,1,2,3,4]	$\sigma_{11}^{(0)}$	[2,1,4,3,5]	$\sigma_{11}^{(1)}$	[4,1,5,3,2]	$\sigma_{11}^{(2)}$
↓	↓	↓	↓	↓	↓
[1,2,3,4,5]	$\sigma_0^{(0)}$	[1,4,3,5,2]	$\sigma_0^{(1)}$	[1,5,3,2,4]	$\sigma_0^{(2)}$
[3,1,2,4,5]	$\sigma_1^{(0)}$	[3,1,4,5,2]	$\sigma_1^{(1)}$	[3,1,5,2,4]	$\sigma_1^{(2)}$

□

We now turn to consider the size and rate of the constructed codes, and show that their rate asymptotically tends to 1.

Theorem 13: The size of \mathcal{K} -snakes constructed in Theorem 11 behaves asymptotically as

$$|C_{2n+1}| = M_{2n+1} = \frac{(2n)!(2n+1)!}{n!^2 \cdot 2^{2n}} \sim \frac{1}{\sqrt{\pi n}} |S_{2n+1}|$$

which leads to an asymptotic rate of 1.

Proof: Starting from our base case of a complete cyclic $(3, 3, \mathcal{K})$ -snake, we define for all $n \in \mathbb{N}$ the ratio

$$D_{2n+1} = \frac{M_{2n+1}}{(2n+1)!}$$

which is the size of our constructed code over the total size of S_{2n+1} . We note that

$$\frac{D_{2n+1}}{D_{2n-1}} = \frac{M_{2n+1} \cdot (2n-1)!}{(2n+1)! \cdot M_{2n-1}} = \frac{2n-1}{2n}.$$

Therefore, since $D_3 = \frac{1}{2}$, we have for all $2 \leq n \in \mathbb{N}$ that

$$D_{2n+1} = \frac{1}{2} \prod_{m=2}^n \frac{2m-1}{2m} = \frac{(2n)!}{n!^2 \cdot 2^{2n}}.$$

Using Stirling’s approximation, one observes

$$\begin{aligned}\lim_{n \rightarrow \infty} D_{2n+1} \sqrt{\pi n} &= \lim_{n \rightarrow \infty} \frac{(2n)! \sqrt{\pi n}}{n!^2 \cdot 2^{2n}} \\ &= \lim_{n \rightarrow \infty} \frac{\sqrt{4\pi n} \left(\frac{2n}{e}\right)^{2n} \sqrt{\pi n}}{\left(\sqrt{2\pi n} \left(\frac{n}{e}\right)^n\right)^2 \cdot 2^{2n}} = 1.\end{aligned}$$

Moreover, one can now readily verify that

$$\lim_{n \rightarrow \infty} R(C_{2n+1}) = \lim_{n \rightarrow \infty} \frac{\log_2 M_{2n+1}}{\log_2 |S_{2n+1}|} = 1.$$

■

Section III-B will focus in exploring the possible size of \mathcal{K} -snakes in general.

Before we conclude this part, we recall that Flash memory cells suffer long-time damage from erasure cycles, and therefore, it is desirable to minimize the number of times such cycles are required.

A property of rank-modulation cell programming is that an erasure of an entire cell block is required only when a specific cell is to exceed its maximal permitted charge level. It is therefore of interest to analyze the rate with which our constructed codes increase the charge level of any given cell.

Repeated ‘‘push-to-the-top’’ operations on a given cell will result in a fast increase in that cell’s charge level, and growing gaps between it and the charge levels of other cells. It is therefore most cost economic, in the sense that it delays the need for a time-consuming erasure and reprogramming cycle to employ a programming strategy which retains the charge levels of individual cells as balanced as possible. Such balanced Gray codes were constructed in [12].

In this part’s context, this goal is achieved if and only if every two subsequent incidents in a cyclic $(2n + 1, M, \mathcal{K})$ -snake where a ‘‘push-to-the-top’’ operation is applied to a certain cell are separated by at most $2n + 1$ operations on other cells. Our family of codes nearly achieves this goal.

Proposition 14: For every permutation $\sigma \in C_{2n+1}$, in the \mathcal{K} -snake constructed in Theorem 11, there exists another $\sigma' \in C_{2n+1}$ such that $\sigma(1) = \sigma'(1)$, following it by no more than $2n + 3$ steps.

Proof: Recall that

$$C_{2n+1} = E^2 C_{2n+1}^{(0)}, E^2 C_{2n+1}^{(1)}, \dots, E^2 C_{2n+1}^{(2n-2)}.$$

By the nature of our construction, for $n \geq 2$, every “push-to-the-top” operation, on all but the last rank in the code, appears either as part of the pattern

$$\dots, \underbrace{t_{2n+1}, \dots, t_{2n+1}}_{2n}, t_i, \underbrace{t_{2n+1}, \dots, t_{2n+1}}_{2n}, \dots$$

or as

$$\dots, \underbrace{t_{2n+1}, \dots, t_{2n+1}}_{2n}, t_3, t_3, \underbrace{t_{2n+1}, \dots, t_{2n+1}}_{2n}, \dots$$

It is therefore the case that there exist $0 \leq k \leq 2n$ and $j \in [n]$ such that the transformations used in C_{2n+1} after σ are of the following two forms.

$$\begin{aligned} 1) & \underbrace{t_{2n+1}, \dots, t_{2n+1}}_k, t_{2j+1}, \underbrace{t_{2n+1}, \dots, t_{2n+1}}_{2n} \\ 2) & \underbrace{t_{2n+1}, \dots, t_{2n+1}}_k, t_3, t_3, \underbrace{t_{2n+1}, \dots, t_{2n+1}}_{2n} \end{aligned}$$

In the second case, one notes

$$\sigma(1) = \begin{cases} t_{2n+1}^{2n-1} t_3^2 \sigma(1), & k = 0 \\ t_3^2 t_{2n+1}^2 \sigma(1), & k = 1 \\ t_3 t_{2n+1}^2 \sigma(1), & k = 2 \\ t_{2n+1}^{2n+1-k} t_3^2 t_{2n+1}^k \sigma(1), & k > 2. \end{cases}$$

Finally, in the first case, we note that

$$\sigma(1) = \begin{cases} t_{2n+1}^{2n-k} t_{2j+1} t_{2n+1}^k \sigma(1), & k < 2j+1 \\ t_{2j+1} t_{2n+1}^k \sigma(1), & k = 2j+1 \\ t_{2n+1}^{2n+1-k} t_{2j+1} t_{2n+1}^k \sigma(1), & k > 2j+1. \end{cases}$$

It is of interest to note that, of all cases discussed in the last proof, the second case where $k > 2$ is the only situation in which another instance of programming to the specific cell fails to occur in $2n+2$ steps, i.e., for the large majority of cases (in all but $\frac{2n-1}{M_{2n+1}}$ of them), the construction of Theorem 11 yields optimally behaving codes in this respect.

B. Bounds on \mathcal{K} -Snakes

We now turn our attention to bounding the parameters of \mathcal{K} -snakes. We begin by noting a simple upper bound on the size of \mathcal{K} -snakes.

Theorem 15: If C is an (n, M, \mathcal{K}) -snake, then:

- 1) $M \leq \frac{1}{2} |S_n|$;
- 2) $M = \frac{1}{2} |S_n|$ if and only if, for all $\{\alpha, \beta\} \in E_n$, it holds that $\alpha \in C$ or $\beta \in C$.

Proof: Every $\alpha \in S_n$ has exactly $(n-1)$ neighbors in G_n . When we sum the edges for every vertex in G_n , each edge in E_n is counted precisely twice, hence

$$|E_n| = \frac{n-1}{2} \cdot |S_n| = \frac{n!(n-1)}{2}.$$

On the other hand, for every $\alpha, \beta \in C$ and $e_1, e_2 \in E_n$ such that $\alpha \in e_1$ and $\beta \in e_2$, clearly $e_1 \neq e_2$. It follows that there are no less than $M(n-1)$ distinct edges in E_n . Hence

$$M \leq \frac{1}{2} |S_n|.$$

Finally, we note that $M = \frac{1}{2} |S_n|$ iff $M(n-1) = |E_n|$, iff every edge in E_n contains a (unique) element of C . ■

It is worth mentioning, at this point, that this upper bound might not be tight. Indeed, we know by Theorem 13 that

$$\frac{M_{2n+1}}{\frac{1}{2} |S_{2n+1}|} \sim \frac{2}{\sqrt{\pi n}}$$

and no constructions are currently known which attain the upper bound, except for the trivial case of C_3 .

The codes we constructed in the previous part use only “push-to-the-top” operations on odd indices. We would now like to show that using even a single “push-to-the-top” operation on an even index can never result in a code attaining the bound of Theorem 15 with equality. We first require a simple lemma.

Lemma 16: Let C be a \mathcal{K} -snake over S_n . If $\sigma, \sigma' \in C$ and there exists a path in G_n of odd length between them, then that path contains an edge both of whose endpoints are not in C .

Proof: Consider such a path of odd length in G_n , connecting σ and σ' . Now, color the vertices of C black, and those of $S_n \setminus C$ white. Since C is a \mathcal{K} -snake, no edge in E_n has both its ends colored black. In the aforementioned path, the vertices cannot alternate in color since σ and σ' are colored black and the path has odd length. It follows that there is an edge in the path with both ends colored white, as claimed. ■

With this lemma in hand, we can now further bound the size of \mathcal{K} -snakes employing a “push-to-the-top” operation on an even index.

Theorem 17: If an (n, M, \mathcal{K}) -snake C contains a “push-to-the-top” operation on an even index, then

$$M \leq \frac{1}{2} |S_n| - \Theta(n) < \frac{1}{2} |S_n|.$$

Proof: Let $C = (\sigma_1, \dots, \sigma_M)$. We take $i \in [M-1]$ such that $\sigma_{i+1} = t_{2m}(\sigma_i)$, where $2m \in [n]$. Then, σ_i and σ_{i+1} have different signs. We will also find it convenient to denote

$$r = \sigma_i(2m) \in [n].$$

We shall construct as many distinct paths in Kendall’s τ adjacency graph G_n connecting σ_i with σ_{i+1} , knowing they must all have odd lengths, and therefore, by Lemma 16 they each contain an edge completely disjoint from C . We will then show that these edges are all distinct, allowing us to improve upon the bound of Theorem 15.

One natural such path is generated by subsequently applying to σ_i the adjacent transpositions $(j, j+1)$ for $j = 2m-1, 2m-2, \dots, 1$. By taking more care before applying these transpositions, we shall arrive at more paths.

Consider the set of adjacent transpositions that do not involve the index $2m$, namely

$$T = \{(j, j+1) \mid j \in [n-1] \setminus \{2m-1, 2m\}\}.$$

For every subset $B \subseteq T$ of size $|B| \leq 2$, we generate a new permutation ω^B by applying to σ_i the elements of B (in some arbitrary order, say from smallest to largest indices). Naturally, for two distinct such subsets, B and B' , we have $\omega^B \neq \omega^{B'}$, but still $\omega^B(2m) = r = \omega^{B'}(2m)$.

We can now apply to ω^B the aforementioned transpositions in the following way:

$$\begin{aligned} \omega_0^B &= \omega^B \\ \omega_j^B &= \omega_{j-1}^B(2m - j, 2m - j + 1); \quad j \in [2m - 1] \end{aligned}$$

and, for every choice of subset, we have $\omega_{2m-1}^B(1) = r$. Clearly, we can generate σ_{i+1} from ω_{2m-1}^B by reversing the effect of B 's elements (the actual transpositions required are altered by the change of index for r , but all other elements retain their relative positions with respect to each other. Formally, we need to apply the elements of B in the reverse order, but whenever $(j, j + 1) \in B$ such that $j < 2m$ we instead apply the adjacent transposition $(j + 1, j + 2)$).

Now, note that if $\omega_k^B = \omega_l^{B'}$, then in particular

$$2m - k = (\omega_k^B)^{-1}(r) = (\omega_l^{B'})^{-1}(r) = 2m - l$$

hence $k = l$. Therefore, the induced permutation on $[n] \setminus \{r\}$ agrees as well. This, however, is impossible unless $B = B'$. Hence, any two paths of this sort can only intersect in the first step of obtaining ω^B from σ_i (or the last step from ω_{2m-1}^B to σ_{i+1}), i.e., in the first (or last) edge of the path.

Finally, by Lemma 16, each path hereby described contains an edge disjoint from C . Note that it cannot be its first or last edge (since $\sigma_i, \sigma_{i+1} \in C$); hence, these edges are all distinct. It follows (in the same manner used in the proof of Theorem 15) that, where N denotes the number of subsets of T with cardinality 2 or less, we have

$$M(n - 1) \leq |E_n| - N = \frac{n - 1}{2} |S_n| - N$$

and naturally $N = \binom{n-3}{2} + (n - 3) + 1 = \Theta(n^2)$. ■

Before concluding this section, we note that the upper bound of Theorem 17 is still higher than M_{2n+1} , the size of codes generated by the construction of Theorem 11. See Section V for some ad hoc results of codes with optimal sizes.

C. Successor Calculation and Ranking Algorithms

We now turn to present algorithms associated with the codes we constructed in the previous sections. The algorithms are brought here for completeness of presentation, and are straightforward derivations from the construction. We shall, therefore, only provide an intuitive sketch of correctness for them, as we shall later do in the section corresponding to l_∞ -snakes.

In order to use the codes described in Theorem 11 in the implementation of a logic cell (with M_{2n+1} levels), importance is known to the ability of efficiently increasing the cell's level. That is, one needs to know, for every given permutation in the code, the appropriate “push-to-the-top” operation required to produce the subsequent permutation.

For the code C_{2n+1} from Theorem 11, the function $\text{Successor}_{\mathcal{K}}(n, [b_1, \dots, b_{2n+1}])$ takes as input a permutation in the code, and returns as output the index i of the required transformation t_i . It is assumed throughout this part that the elements $\{a_i\}_{i=0}^{2n-2}$ from (2), used in our construction, are known, and we will denote them with superscript (n) to indicate order when it is not clear from context. Furthermore, we require a function

$$\text{Ind}_n(b) : [2n + 1] \setminus \{1, 3\} \rightarrow [0, 2n - 2]$$

which returns the unique index such that $a_{\text{Ind}_n(b)} = b$. We assume Ind_n runs in $O(1)$ time.² ne possible way, among many, of achieving this is by defining

$$a_i^{(n)} = \begin{cases} 2, & i = 0 \\ i + 3, & i \geq 1 \end{cases} \quad \text{Ind}_n(b) = \begin{cases} 0, & b = 2 \\ b - 3, & b \geq 4. \end{cases}$$

Finally, we naturally assume validity of the input in all procedures.

Our strategy will be to identify the vertices in C_{2n+1} which require a transformation other than t_{2n+1} . Those are either permutations with leading 1's (those on which we initially performed “push-to-the-bottom” operations in our construction), or the last permutation in each $E^2 C_{2n+1}^{(j)}$. In the latter case, we need only apply t_3 , where the former requires translation of the $a_i^{(n)}$'s according to their respective positions in the originating permutation of each $C_{2n+1}^{(j)}$, and a recursive run of $\text{Successor}_{\mathcal{K}}$ to determine the correct “push-to-the-bottom” operation to be performed.

It shall be noted at this point that a degree of freedom exists in the cyclic shift of C_{2n-1} one applies to construct each $C_{2n+1}^{(j)}$ (one only needs to confirm that the first “push-to-the-top” operation shall be on the last index). This shift shall be denoted by the following bijection for every order $n \in \mathbb{N}$ and index $j \in [2n - 1]$:

$${}^n_j \downarrow : \{3\} \cup \{a_i^{(n)}\}_{i \neq j} \longrightarrow [2n - 1]$$

defined such that the “push-to-the-bottom” operation applied to

$$\left[1, a_j^{(n)}, b_1, \dots, b_{2n-1}\right] \in C_{2n+1}^{(j)}$$

matches the “push-to-the-top” operation applied in C_{2n-1} to

$$\left[{}^n_j \downarrow b_{2n-1}, {}^n_j \downarrow b_{2n-2}, \dots, {}^n_j \downarrow b_1\right].$$

We shall further denote its inverse as ${}^n_j \uparrow$. These two bijections can be implemented in $O(1)$ time, for example, by taking as a starting point C_{2n-1} 's $(2n - 4)$ -ranked permutation

$$\left[a_0^{(n-1)}, \dots, a_{2n-4}^{(n-1)}, 3, 1\right]$$

²Though the integers used throughout are of magnitude $O(n)$, and so may require $O(\log n)$ bits to represent, we tacitly assume (as in [12]) all simple integer operations, e.g., assignment, comparison, addition, etc., to take $O(1)$ time.

Function $\text{Successor}_{\mathcal{K}}(n, [b_1, \dots, b_{2n+1}])$	
input	: $n \in \mathbb{N}$, A permutation $[b_1, \dots, b_{2n+1}] \in C_{2n+1}$
output	: An odd $i \in \{3, \dots, 2n+1\}$ that determines the transition t_i to the next permutation in C_{2n+1}
1	if $n = 1$ then
2	return 3
3	if $b_1 = 3$ and $b_2 = 1$ and $\forall 3 \leq i \leq 2n$: $(\text{Ind}_n(b_{i+1}) - \text{Ind}_n(b_i)) \equiv 1 \pmod{2n-1}$ then
4	return 3
5	if $b_1 = 1$ then
6	$j \leftarrow \text{Ind}_n(b_2)$
7	$i \leftarrow \text{Successor}_{\mathcal{K}}(n-1, [{}^n_j \downarrow b_{2n+1}, {}^n_j \downarrow b_{2n}, \dots, {}^n_j \downarrow b_3])$
8	return $2n+2-i$
9	return $2n+1$

Function $\text{Rank}_{\mathcal{K}}([b_1, \dots, b_{2n+1}])$	
input	: A permutation $[b_1, \dots, b_{2n+1}] \in C_{2n+1}$
output	: The rank $k \in \{0, \dots, M_{2n+1} - 1\}$ associated with the given permutation in C_{2n+1}
1	if $n = 1$ then
2	return $3 - b_2$
3	$i \leftarrow \min \{l \in [2n+1] \mid b_l = 1\}$
4	$j \leftarrow \text{Ind}_n(b_{(i \bmod (2n+1))+1})$
5	for $l \leftarrow 1$ to $2n-1$ do
6	$c_l \leftarrow {}^n_j \downarrow b_{((i-1) \bmod (2n+1))+1}$
7	$r \leftarrow (\text{Rank}_{\mathcal{K}}([c_1, \dots, c_{2n-1}]) - r_{2n+1}^{(j)}) \bmod M_{2n-1}$
8	$rn \leftarrow ((2n+1)(r-1) - 1 + ((i-2) \bmod (2n+1))) \bmod ((2n+1)M_{2n-1})$
9	return $(2n+1)M_{2n-1} \cdot j + rn$

and defining accordingly

$${}^n_j \downarrow b = \begin{cases} 1, & b = 3 \\ 3, & \text{Ind}_n(b) = j + 1 \\ a_{(j - \text{Ind}_n(b) - 1) \bmod (2n-1)}^{(n-1)}, & \text{otherwise} \end{cases} \quad (3)$$

where $\text{Ind}_n(b) = j + 1$ is checked modulo $2n - 1$, as well as

$${}^n_j \uparrow b = \begin{cases} 3, & b = 1 \\ a_{(j+1) \bmod (2n-1)}^{(n)}, & b = 3 \\ a_{(j - \text{Ind}_n(b) - 1) \bmod (2n-1)}^{(n)}, & \text{otherwise.} \end{cases} \quad (4)$$

Lemma 18: $\text{Successor}_{\mathcal{K}}$ runs in $O(1)$ amortized time.

Proof: We first note that by the nature of our construction, the element 1 appears in the leading index precisely $(2n-1) \cdot M_{2n-1}$ times, which constitutes $\frac{1}{2n+1}$ of the code's size. The pair (3, 1) leads to no more (and in fact strictly less) permutations.

Therefore, if we let E_n denote the expected number of steps performed by $\text{Successor}_{\mathcal{K}}$ when called on input of length $2n+1$, then we note the recursive connection

$$\begin{aligned} E_n &\leq O(1) + \frac{1}{2n+1}O(n) + \frac{1}{2n+1}(O(n) + E_{n-1}) \\ &= O(1) + \frac{1}{2n+1}E_{n-1}. \end{aligned}$$

Developing this inequality recursively, there exists $L \in \mathbb{N}$ such that

$$\begin{aligned} E_n &\leq L + \frac{1}{2n-1}E_{n-1} \\ &\leq \left(1 + \frac{1}{2n-1}\right)L + \frac{1}{(2n-1)(2n-3)}E_{n-2} \leq \\ &\vdots \\ &\leq \left(1 + \frac{1}{2n-1} + \frac{n-2}{(2n-1)(2n-3)}\right)L + \frac{n!2^n}{(2n)!}E_1 \end{aligned}$$

and so $E_n = O(1)$. \blacksquare

To use C_{2n+1} in the implementation of a logic cell, one also needs a method of computing a given permutation's rank in the code. We implement the function $\text{Rank}_{\mathcal{K}}([b_1, \dots, b_{2n+1}])$ which receives as input a permutation $[b_1, \dots, b_{2n+1}] \in C_{2n+1}$ and returns its rank in

$$C_{2n+1} = E^2 C_{2n+1}^{(0)}, E^2 C_{2n+1}^{(1)}, \dots, E^2 C_{2n+1}^{(2n-2)}$$

Function $\text{Unrank}_{\mathcal{K}}(n, k)$	
input	: $n \in \mathbb{N}$; rank $k \in [0, M_{2n+1} - 1]$
output	: The permutation $[b_1, \dots, b_{2n+1}]$ which is k th in C_{2n+1}
1	if $n = 0$ then
2	return [1]
3	$j \leftarrow \lfloor \frac{k}{(2n+1) \cdot M_{2n-1}} \rfloor$
4	$pos \leftarrow k \bmod ((2n+1)M_{2n-1})$
5	$perm \leftarrow \left(\lfloor \frac{pos+1}{2n+1} \rfloor + 1 + r_{2n+1}^{(j)}\right) \bmod M_{2n-1}$
6	$shift \leftarrow (pos+2) \bmod (2n+1)$
7	$[c_1, \dots, c_{2n-1}] \leftarrow \text{Unrank}_{\mathcal{K}}(n-1, perm)$
8	return $t_{2n+1}^{shift} [1, a_j^{(n)}, {}^n_j \uparrow c_{2n-1}, {}^n_j \uparrow c_{2n-2}, \dots, {}^n_j \uparrow c_1]$

in the order indicated by that notation. The assumptions made in the previous part are still in effect. Moreover, we will require the knowledge of the cyclic shift of C_{2n-1} used in the construction of each $C_{2n+1}^{(j)}$, which we retain in the form of $r_{2n+1}^{(j)}$, the rank of permutation in C_{2n-1} which was chosen as a starting point. For example, in the method suggested in (3) and (4), we have

$$r_{2n+1}^{(j)} = 2n - 4$$

for all $j \in [2n-1]$.

We use the following method: first identify the position of 1 in the permutation, and the following element, which gives us both the subcode the permutation belongs to and the cyclic shift in our mock "push-to-the-bottom" operation. Armed with that information, we then scan the permutation backward and translate the $a_j^{(n)}$'s indices according to the subcode in the same way we did in $\text{Successor}_{\mathcal{K}}$. After that, a recursive run of $\text{Rank}_{\mathcal{K}}$ will give us the permutation's position in its subcode, which we will combine with the cyclic shift to produce the correct rank, taking $r_{2n+1}^{(j)}$ into account and remembering that C_{2n+1} is constructed of the $E^2 C_{2n+1}^{(j)}$'s rather than the $C_{2n+1}^{(j)}$'s.

Lemma 19: The function $\text{Rank}_{\mathcal{K}}$ operates in $O(n^2)$ steps.

Proof: We note that $\text{Rank}_{\mathcal{K}}$ performs $O(n)$ operations before calling upon itself with an order reduced by one. It therefore operates in $O(n^2)$ time. \blacksquare

Unranking permutations, i.e., the process of assigning to a given rank in $[0, M_{2n+1} - 1]$ the corresponding permutation in the C_{2n+1} , might also be needed if one requires the logic cell to perform as more than a counter. We implement a function $\text{Unrank}_{\mathcal{K}}(n, k)$ which returns as output the k -ranked permutation in C_{2n+1} .

Naturally, all the aforementioned assumptions made still hold. We will follow the same general method used for $\text{Rank}_{\mathcal{K}}$, i.e., we shall compute $j \in [2n-1]$ such that the given rank

belongs to $\sigma \in E^2 C_{2n+1}^{(j)}$, then adjust the rank to indicate the correct position in $C_{2n+1}^{(j)}$. It will then remain to compute the correct permutation in the “push-to-the-bottom” cycle using a recursive run, and shift it the required number of times.

Lemma 20: The function Unrank_K operates in $O(n^2)$ steps as well.

Proof: Follows exactly the same lines as our proof to Lemma 19. ■

IV. ℓ_∞ -METRIC AND ℓ_∞ -SNAKES

The ℓ_∞ -metric is induced on S_n by the embedding in \mathbb{Z}^n implied by the vector notation. More precisely, for $\alpha, \beta \in S_n$, one defines

$$d_\infty(\alpha, \beta) = \max_{i \in [n]} |\alpha(i) - \beta(i)|.$$

We use the ℓ_∞ -metric to model a different kind of noise mechanism than that modeled by Kendall’s τ -metric, namely spike noise. In this model, the rank of each memory cell is assumed to have been changed by a bounded amount (see [22]).

Error-correcting and -detecting codes in S_n for the ℓ_∞ -metric are referred in [22] as limited-magnitude rank-modulation codes (LMRM codes). In that paper, constructions of such codes achieving nonvanishing normalized distance and rate are presented. Moreover, bounds on the size of optimal LMRM codes are proven. In particular, it has been shown [22, Th. 20] that if C is an $(n, M, 2)$ -LMRM, then

$$M \leq \frac{n!}{2^{\lfloor n/2 \rfloor}}.$$

Using a simple translation to an extremal problem involving permanents of $(0, 1)$ -matrices (see [20]), this is also the best possible bound using the set-antiset method. For our needs, it follows that the size of every n -length ℓ_∞ -snake is bounded by this term. We shall present a construction of ℓ_∞ -snakes achieving this upper bound by a factor of $\lfloor \frac{n}{2} \rfloor 2^{\lfloor n/2 \rfloor}$, which we will show achieves an asymptotic rate of 1.

A. Construction

In order to use the code constructions presented in [12], we first prove the following lemma.

Lemma 21: Both constructions in [12, Ths. 4 and 7], when applied recursively, yield complete cyclic n -RMGC’s containing both “push-to-the-top” operations t_2 and t_n .

Proof: The proposition was, while not fully stated, actually proven in [12, Th. 4].

For [12, Th. 7], we shall assume that the recursive process was applied to a length- $(n-1)$ Gray code satisfying these conditions (as is the case with the base example given in that article). The resulting code uses t_n by definition. Moreover, since the original code used t_{n-1} , the resulting code uses $t_{n-(n-1)+1} = t_2$. ■

This lemma now allows for the construction of a basic building block which we will later use.

Lemma 22: Let $\{a_j\}_{j=1}^n$, $n \geq 2$, be a set of integers of the same parity. Let

$$\sigma = [x, a_1, a_2, \dots, a_n, b_{n+2}, b_{n+3}, \dots, b_m] \in S_m$$

be a permutation such that the parity of x differs from that of the elements of $\{a_j\}_{j=1}^n$. Then, there exists a (noncyclic) $(m, n + (n-1)!, \ell_\infty)$ -snake starting with σ and ending with the permutation

$$t_2 t_{n+1}^n(\sigma) = [a_2, a_1, a_3, a_4, \dots, a_n, x, b_{n+2}, b_{n+3}, \dots, b_m].$$

Proof: Let $\sigma_0, \dots, \sigma_{n+(n-1)!-1}$ denote the codewords of the claimed code, and denote by $t_{k_1}, \dots, t_{k_{n+(n-1)!-1}}$ the list of transformations generating it.

We set $\sigma_0 = \sigma$. For all $i \in [n]$, we let $\sigma_i = t_{n+1}^i(\sigma)$, i.e., $t_{k_i} = t_{n+1}$. Quite clearly, any two of these $n+1$ permutations are at ℓ_∞ -distance at least 2 apart, since the a_j ’s share parity.

Now, by Lemma 21, there exists a complete cyclic $(n-1)$ -RMGC starting with σ_n , with its last operation being t_2 . We, therefore, let $t_{k_{n+i}}$ for $i \in [(n-1)!]$ represent that code; hence, $t_{k_{n+(n-1)!}} = t_2$ and $\sigma_{n+(n-1)!} = \sigma_n$ (we then, obviously, omit the last transformation as well as the repeated codeword $\sigma_{n+(n-1)!}$). These $(n-1)!$ permutations, $\sigma_n, \dots, \sigma_{n+(n-1)!-1}$, also represent an ℓ_∞ -snake, for the same reason.

Finally, take $0 \leq k < n$ and $0 \leq l < (n-1)!$, and observe σ_k and σ_{n+l} . Suppose $d_\infty(\sigma_k, \sigma_{n+l}) \leq 1$. Then, in particular, $|a_{n-k} - x| = 1$. Moreover, if $k = n-1$, then $|x - a_n| = 1$, but then a_n ’s position in σ_k correlates to that of $\{a_j\}_{j=1}^{n-1}$ in σ_{n+l} , in contradiction. Therefore, $k \leq n-2$, but then a_n ’s position in σ_{n+l} (n th from left) correlates to that of a_{n-k-1} in σ_k , where $1 \leq n-k-1 \leq n-1$, again in contradiction. This concludes our proof. ■

Example 23: We shall start this example with the permutation

$$\sigma = [1, 2, 4, 6, 3, 5].$$

We will also require a complete 2-RMGC, which clearly comprises of two subsequent t_2 operations. We are now ready to present a (noncyclic) $(6, 3 + 2!, \ell_\infty)$ -snake

$$\begin{array}{l} [1, 2, 4, 6, 3, 5] \\ \downarrow \quad t_4 \\ [6, 1, 2, 4, 3, 5] \\ \downarrow \quad t_4 \\ [4, 6, 1, 2, 3, 5] \\ \downarrow \quad t_4 \\ [2, 4, 6, 1, 3, 5] \\ \downarrow \quad t_2 \\ [4, 2, 6, 1, 3, 5]. \end{array}$$

An additional t_2 operation is called for by our complete 2-RMGC, but we omit it. We also note that any permutation on the *odd* element in this example will not change its properties as an ℓ_∞ -snake. □

Having this building block in hand, we continue to describe a construction of a cyclic ℓ_∞ -snake. The construction follows by dividing the ranks in a length- n permutation into even and odd elements, and covering permutations on each half separately.

Theorem 24: For all $4 \leq n \in \mathbb{N}$, there exists an (n, M, ℓ_∞) -snake of size

$$M = \left\lfloor \frac{n}{2} \right\rfloor! \left(\left\lfloor \frac{n}{2} \right\rfloor + \left(\left\lfloor \frac{n}{2} \right\rfloor - 1 \right)! \right).$$

Proof: To simplify notations, we start by noting that $[n]$ has $p = \lceil \frac{n}{2} \rceil$ odd elements and $q = \lfloor \frac{n}{2} \rfloor$ has even ones. We shall use that notation throughout this proof.

Using [12, Ths. 4 and 7], we take a complete cyclic p -RMGC using the operations

$$t_{\alpha(1)}, t_{\alpha(2)}, \dots, t_{\alpha(p!)}.$$

Moreover, we use Lemma 22 to come by a (q, M_q, ℓ_∞) -snake of size $M_q = q + (q-1)!$ given by the operations

$$t_{\beta(1)}, t_{\beta(2)}, \dots, t_{\beta(q+(q-1)!-1)}.$$

As the origin for the code we construct, we use

$$\sigma_0 = [1, 2, 4, \dots, 2q, 3, \dots, 2p-1].$$

For all $i \in [p!]$ and $j \in [q + (q-1)! - 1]$, we define the sequence of transformations generating the code as

$$\begin{aligned} t_{k_{(i-1)(q+(q-1)!)+j}} &= t_{\beta(j)} \\ t_{k_{i(q+(q-1)!)}} &= t_{\alpha(i)+q} \end{aligned}$$

and where, naturally, the codewords satisfy $\sigma_i = t_{k_i}(\sigma_{i-1})$.

We start by noting that, for all $i \in [p!]$, the permutation $\sigma_{(i-1)(q+(q-1)!)}$ satisfies the requirements of Lemma 22 as a simple matter of induction. It follows that, for all $i \in [p!]$, the permutations

$$\{\sigma_{(i-1)(q+(q-1)!)+1}, \sigma_{(i-1)(q+(q-1)!)+2}, \dots, \sigma_{i(q+(q-1)!)-1}\}$$

are at ℓ_∞ -distance of at least 2 apart.

Furthermore, for $i, i' \in [p!]$, $i < i'$, since the code generated by $t_{\alpha(1)}, t_{\alpha(2)}, \dots, t_{\alpha(p!)}$ is indeed a Gray code, we are assured that, for all $0 \leq j, j' \leq q + (q-1)! - 1$, the last $p-1$ elements of both $\sigma_{(i-1)(q+(q-1)!)+j}$ and $\sigma_{(i'-1)(q+(q-1)!)+j'}$ are all odd and represent two distinct permutations; hence

$$d_\infty(\sigma_{(i-1)(q+(q-1)!)+j}, \sigma_{(i'-1)(q+(q-1)!)+j'}) \geq 2.$$

Finally, we note that

$$t_{\alpha(p!)}(\sigma_{p!(q+(q-1)!)-1}) = \sigma_0$$

since the code provided by $t_{\alpha(1)}, t_{\alpha(2)}, \dots, t_{\alpha(p!)}$ is cyclic and $o(t_2) = 2$ divides $p!$. ■

Example 25: For this example, using an order of six as the last example (i.e., $p = q = 3$), we refer to [12, Ths. 4 and 7] for a complete cyclic 3-RMGC as well as the aforementioned 2-RMGC. One such is created by the transitions

$$t_2, t_3, t_3, t_2, t_3, t_3.$$

Moreover, we have our $(6, 3+2!, \ell_\infty)$ -snake from the last example, generated by (recall that it's *not* cyclic)

$$t_4, t_4, t_4, t_2.$$

We start our cyclic $(6, 3!(3+2!), \ell_\infty)$ -snake in the same permutation as we did the last example, and use the generating tran-

sitions of the code presented in the last example as a building block

σ_0	[1, 2, 4, 6, 3, 5]	
	⇓	
σ_4	[4, 2, 6, 1, 3, 5]	
	↓	$t_5 = t_{3+2}$
σ_5	[3, 4, 2, 6, 1, 5]	
	⇓	
σ_9	[2, 4, 6, 3, 1, 5]	
	↓	$t_6 = t_{3+3}$
σ_{10}	[5, 2, 4, 6, 3, 1]	
	⇓	
σ_{14}	[4, 2, 6, 5, 3, 1]	
	↓	$t_6 = t_{3+3}$
σ_{15}	[1, 4, 2, 6, 5, 3]	
	⇓	
σ_{19}	[2, 4, 6, 1, 5, 3]	
	↓	$t_5 = t_{3+2}$
σ_{20}	[5, 2, 4, 6, 1, 3]	
	⇓	
σ_{24}	[4, 2, 6, 5, 1, 3]	
	↓	$t_6 = t_{3+3}$
σ_{25}	[3, 4, 2, 6, 5, 1]	
	⇓	
σ_{29}	[2, 4, 6, 3, 5, 1]	
	↓	$t_6 = t_{3+3}$
σ_{30}	[1, 2, 4, 6, 3, 5].	

Down double arrows stand for an omitted sequence of transitions generating Example 23, i.e., t_4, t_4, t_4, t_2 . One sees that we indeed have a cyclic ℓ_∞ -snake of size 30. □

We note that by switching the roles of odd and even numbers in Theorem 24, we can construct an (n, M, ℓ_∞) -snake of size

$$M = \left\lfloor \frac{n}{2} \right\rfloor! \left(\left\lceil \frac{n}{2} \right\rceil + \left(\left\lceil \frac{n}{2} \right\rceil - 1 \right)! \right).$$

However, the resulting code is strictly smaller for odd n .

Theorem 26: The ℓ_∞ -snakes constructed in Theorem 24 have an asymptotic rate of 1.

Proof: Let C_n denote the ℓ_∞ -snake of length n constructed by Theorem 24. Using the crude

$$\left(\frac{n}{e} \right)^n \leq n! \leq n^n$$

the proof is a matter of simple calculation

$$\begin{aligned} \lim_{n \rightarrow \infty} R(C_n) &= \lim_{n \rightarrow \infty} \frac{\log_2 \left(\left\lceil \frac{n}{2} \right\rceil! \left(\left\lfloor \frac{n}{2} \right\rfloor + \left(\left\lfloor \frac{n}{2} \right\rfloor - 1 \right)! \right) \right)}{\log_2(n!)} \\ &\geq \lim_{n \rightarrow \infty} \frac{2 \log_2 \left(\left(\left\lfloor \frac{n}{2} \right\rfloor - 1 \right)! \right)}{\log_2(n!)} \\ &\geq \lim_{n \rightarrow \infty} \frac{(n-4) \log_2 \left(\frac{n-4}{2e} \right)}{n \log_2 n} = 1. \end{aligned}$$

B. Successor Calculation and Ranking Algorithms

Finding the correct “push-to-the-top” operation to propagate a given permutation to the following one is naturally

Function $\text{Successor}_\infty([a_1, \dots, a_n])$	
input	: A permutation $[a_1, a_2, \dots, a_n]$
output	: $i \in \{2, 3, \dots, n\}$ that determines the transition t_i to the next permutation in the ℓ_∞ -snake from Theorem 24
1	$q \leftarrow \lfloor \frac{n}{2} \rfloor; p \leftarrow \lceil \frac{n}{2} \rceil$
2	if $a_{q+1} \equiv 0 \pmod{2}$ then
3	return $q + 1$
4	if $\text{Rn}(\lfloor \frac{a_{q+1}+1}{2}, \dots, \frac{a_{q+1}}{2} \rfloor) \equiv 0 \pmod{2}$ then
5	if $[a_1, \dots, a_q] = [4, 2, 6, \dots, 2q]$ then
6	return $q + \text{Succ}(\lfloor \frac{a_{q+1}+1}{2}, \dots, \frac{a_{q+1}}{2} \rfloor)$
7	return $\text{Succ}(\lfloor \frac{a_1}{2}, \dots, \frac{a_q}{2} \rfloor)$
8	if $[a_1, \dots, a_q] = [2, 4, \dots, 2q]$ then
9	return $q + \text{Succ}(\lfloor \frac{a_{q+1}+1}{2}, \dots, \frac{a_{q+1}}{2} \rfloor)$
10	return $\text{Succ}(\text{sw}(\lfloor \frac{a_1}{2}, \dots, \frac{a_{q-1}}{2} \rfloor))$

dependent upon one's ability to do the same with the $\lfloor \frac{n}{2} \rfloor$ - and $(\lfloor \frac{n}{2} \rfloor - 1)$ -RMGC's used in our construction. We, therefore, assume to have the function $\text{Succ}([a_1, a_2, \dots, a_n])$ which accepts as input a permutation $[a_1, a_2, \dots, a_n] \in S_n$ and returns the correct transformation used in the codes we used. Furthermore, we assume to have the function $\text{Rn}([a_1, a_2, \dots, a_n])$ which returns the respective rank of the input permutation in that code, where the identity permutation is assumed to have rank zero. Finally, we shall use an auxiliary function $\text{sw} : S_n \rightarrow S_n$ defined by $\text{sw}(\sigma) = (1, 2) \circ \sigma$ (which naturally operates in $O(n)$ steps).

The function $\text{Successor}_\infty([a_1, \dots, a_n])$ then returns as output the index i of the required transformation t_i to produce the subsequent permutation in the code from $[a_1, \dots, a_n]$. It operates by considering the following cases: in each block of Lemma 22, one computes the proper index by propagating the leading element of odd rank as long as that is needed, then applying Succ to the permutation on the elements of even ranks (where one distinguishes between blocks in which 2, 4 were switched). Only the last permutation of each block calls for applying Succ to the permutation on the elements of odd ranks.

Lemma 27: If the functions Succ , Rn operate in L_n , M_n steps, respectively, in the average case, then Successor_∞ has an average runtime of $O(n + L_{q-1} + M_p)$.

Proof: We partition our proof by return cases. Successor_∞ exits at line 3 in precisely $\frac{q}{q+(q-1)!}$ of cases, in which case it returns within a fixed number of operations.

It exits at lines 6 and 9 in $\frac{1}{q+(q-1)!}$ of cases, in which case it operates in at most (depending on the data structures in use) $O(n) + M_p + L_p$ steps in the average case.

Finally, Successor_∞ returns from lines 7 and 10 in $\frac{(q-1)!-1}{q+(q-1)!}$ of cases, after performing $O(n) + M_p + L_{q-1}$ steps.

In every sensible implementation of Succ (i.e., where we assume $\frac{L_p - L_{q-1}}{q+(q-1)!} \rightarrow 0$), we then have an amortized runtime of $O(n + L_{q-1} + M_p)$. ■

We now note that by [12, Ths. 7 and 10], we may assume Succ to operate in $O(1)$ steps in the average case, and, by [12, Part III-C] (which also relies on [16]), we assume Rn runs in $O(n)$ steps, yielding an average runtime of $O(n)$ for Successor_∞ .

We shall also present the function $\text{Rank}_\infty(n, [a_1, \dots, a_n])$ that, given a permutation in the ℓ_∞ -snake presented in

Function $\text{Rank}_\infty([a_1, \dots, a_n])$	
input	: A permutation $[a_1, a_2, \dots, a_n]$ in the ℓ_∞ -snake from Theorem 24
output	: $k \in \mathbb{N}$ that represents the given permutation's rank in the code
1	$q \leftarrow \lfloor \frac{n}{2} \rfloor; p \leftarrow \lceil \frac{n}{2} \rceil$
2	if $a_{q+1} \equiv 0 \pmod{2}$ then
3	$i \leftarrow \min\{j \in [n] \mid a_j \not\equiv 0 \pmod{2}\}$
4	return $i - 1 + (q + (q - 1)!) \cdot \text{Rn}(\lfloor \frac{a_{q+1}}{2}, \frac{a_{q+2}+1}{2}, \dots, \frac{a_{q+1}}{2} \rfloor)$
5	$R \leftarrow \text{Rn}(\lfloor \frac{a_{q+1}+1}{2}, \dots, \frac{a_{q+1}}{2} \rfloor)$
6	if $R \equiv 0 \pmod{2}$ then
7	return $q + (q + (q - 1)!) \cdot R + \text{Rn}(\lfloor \frac{a_1}{2}, \dots, \frac{a_{q-1}}{2} \rfloor)$
8	return $q + (q + (q - 1)!) \cdot R + \text{Rn}(\text{sw}(\lfloor \frac{a_1}{2}, \dots, \frac{a_{q-1}}{2} \rfloor))$

Function $\text{Unrank}_\infty(n, k)$	
input	: $4 \leq n \in \mathbb{N}$; rank $k \in \mathbb{N}$
output	: The permutation $[a_1, a_2, \dots, a_n]$ which is k th in the (n, M, ℓ_∞) -snake from Theorem 24
1	$q \leftarrow \lfloor \frac{n}{2} \rfloor; p \leftarrow \lceil \frac{n}{2} \rceil$
2	$R \leftarrow \lfloor \frac{k}{q+(q-1)!} \rfloor; r \leftarrow (k \bmod (q + (q - 1)!))$
3	$[b_1, \dots, b_p] \leftarrow \text{UnR}(p, R)$
4	if $r \geq q$ then
5	$[a_1, \dots, a_{q-1}] \leftarrow \text{UnR}(q - 1, r - q)$
6	if $R \equiv 1 \pmod{2}$ then
7	$[a_1, \dots, a_{q-1}] \leftarrow \text{sw}([a_1, \dots, a_{q-1}])$
8	return $[2a_1, \dots, 2a_{q-1}, 2q, 2b_1 - 1, \dots, 2b_p - 1]$
9	if $R \equiv 0 \pmod{2}$ then
10	return $[2, 4, 6, \dots, 2r, 2b_1 - 1, 2(r+1), \dots, 2q, 2b_2 - 1, \dots, 2b_p - 1]$
11	return $[4, 2, 6, \dots, 2r, 2b_1 - 1, 2(r+1), \dots, 2q, 2b_2 - 1, \dots, 2b_p - 1]$

Section IV-A, returns that permutation's rank in the code. This function uses the function Rn discussed earlier as well, and works by considering the same cases discussed earlier.

Lemma 28: If the function Rn operates in M_n steps, then Rank_∞ has a runtime of $O(n + M_p)$ (in the average or worst case, respectively).

Proof: We partition our proof by return condition once more. If the program exits from 4, then it performed $O(q) + M_p$ steps.

If it exits from 7 or 8, then it performed $O(1) + M_p + M_{q-1}$ steps. ■

Again, by results discussed previously, we note that Rank_∞ runs in $O(n)$ steps in the average case.

As mentioned previously, unranking permutations in the code might also be required. For that purpose, we implement the function $\text{Unrank}_\infty(n, k)$, accepting as input the length of the code and a specific rank and returning the implied permutation. We will assume the existence of a similar function UnR for the construction used in Section IV-A, where again we assume the unit permutation to have rank zero.

Once more, our implementation and estimate of Unrank_∞ 's runtime relies heavily on that of its auxiliary functions.

Lemma 29: If the function UnR operates in N_n steps, then Unrank_∞ runs in $O(n + N_p)$ steps.

Proof: One notes that the only operations in Unrank_∞ that take more than a fixed number of steps are calls for sw (taking $O(n)$), calls for UnR , and, depending on the data structures in use, concatenation of indices (at most $O(n)$ as well). The claim follows. ■

Again, it shall be noted that, relying on Lemma 21 and [12, Part III-C], Unrank_∞ can be performed in $O(n^2)$ operations.

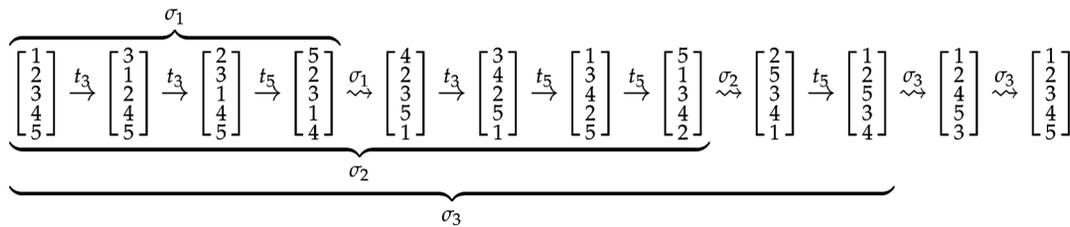


Fig. 1. $(5, 57, \mathcal{K})$ -snake generated by a computer search. Squiggly arrows stand for a repetition of the transitions defined by the braces.

V. CONCLUSION

In this paper, we explored rank-modulation snake-in-the-box codes under both Kendall’s τ -metric and the ℓ_∞ -metric. In both cases, we presented a construction yielding codes with rates asymptotically tending to 1, and implemented auxiliary functions for the production of the successor permutation, as well as ranking and unranking for permutations in such codes. We also proved upper bounds on the size of \mathcal{K} -snakes.

However, it is not presently known whether the upper bounds presented and referenced in this paper are achievable, and therefore, we were unable to show how close the codes generated by our constructions come to being optimal with respect to their sizes (rather than their asymptotic rates). A computer search for *cyclic* codes, performed on S_5 , yielded $(5, M, \mathcal{K})$ -snakes of maximal size $M = 57$ (for comparison, the construction from Theorem 11 yields a $(5, 45, \mathcal{K})$ -snake). While an abundance of such codes were found (well over 500 nonequivalent codes), they all were in fact codes over A_5 . For completeness, we present one of those codes in Fig. 1.

Searches of a higher order appear to be infeasible, but we include one more peculiar result: every maximal code we tested, skipped three permutations who all agree on 4, 5, i.e., it skipped a coset of S_3 . While we have no optimal codes of a higher order to test this phenomenon on, the codes generated by Theorem 11 of lengths 7 and 9 display it as well—several cosets of S_5 and S_7 were absent, respectively.

It shall be noted that a complete (but not cyclic) $(5, 60, \mathcal{K})$ -snake over A_5 can easily be constructed from each cyclic code we tested by generating the skipped coset of S_3 with two t_3 operations, followed by a t_5 operation and the given code, in order. However, we do not currently know whether $(2n + 1, \frac{(2n+1)!}{2}, \mathcal{K})$ -snakes over A_{2n+1} exist for every length.

These results, along with the bounds we showed in Lemma 17 and Theorem 15, give rise to the following conjecture: for all $n \in \mathbb{N}$, a \mathcal{K} -snake exists over A_n whose size is no less than that of every \mathcal{K} -snake over S_n .

In addition, searches done in a computer for ℓ_∞ -snakes for lengths 4, 5, and 6 returned codes of size 6, 30, and 90, respectively, suggesting that perhaps the upper bound of [22, Th. 20] is achievable. Moreover, in these cases, we were able to find codes generated only by “push-to-the-top” operations on the last two indices. A code for each length is presented in Fig. 2 in binary representation (conveniently written in octal notation), where zeroes stand for t_n ’s and ones for t_{n-1} ’s. Searches for higher lengths again seem infeasible.

n	Defining Transitions
4	55
5	0212206063
6	010204410222042124446130162347

Fig. 2. $(4, 6, \ell_\infty)$ -, $(5, 30, \ell_\infty)$ -, and $(6, 90, \ell_\infty)$ -snakes generated by a computer search. All codes represented by a sequence of “push-to-the-top” operations, applied in order to the identity permutation, where zeroes stand for t_n ’s and ones for t_{n-1} ’s. The binary strings are given in octal notation and should be read from left to right.

ACKNOWLEDGMENT

The authors would like to thank the two anonymous reviewers, whose comments improved the presentation of this paper.

REFERENCES

- [1] H. L. Abbot and M. Katchalski, “On the construction of snake in the box codes,” *Utilitas Math.*, vol. 40, pp. 97–116, 1991.
- [2] A. Barg and A. Mazumdar, “Codes in permutations and error correction for rank modulation,” *IEEE Trans. Inf. Theory*, vol. 56, no. 7, pp. 3158–3165, Jul. 2010.
- [3] J. Brewer and M. Gill, *Nonvolatile Memory Technologies With Emphasis on Flash*. New York: Wiley, 2008.
- [4] F. Chierichetti, H. Finucane, Z. Liu, and M. Mitzenmacher, “Designing floating codes for expected performance,” *IEEE Trans. Inf. Theory*, vol. 56, no. 3, pp. 968–978, Mar. 2010.
- [5] M. Deza and H. Huang, “Metrics on permutations, a survey,” *J. Comb. Inf. Sys. Sci.*, vol. 23, pp. 173–185, 1998.
- [6] E. E. Gad, A. Jiang, and J. Bruck, “Compressed encoding for rank modulation,” in *Proc. IEEE Int. Symp. Inf. Theory*, St. Petersburg, Russia, Aug. 2011, pp. 884–888.
- [7] E. E. Gad, M. Langberg, M. Schwartz, and J. Bruck, “Constant-weight Gray codes for local rank modulation,” *IEEE Trans. Inf. Theory*, vol. 57, no. 11, pp. 7431–7442, Nov. 2011.
- [8] E. E. Gad, M. Langberg, M. Schwartz, and J. Bruck, “Generalized Gray codes for local rank modulation,” in *Proc. IEEE Int. Symp. Inf. Theory*, St. Petersburg, Russia, Aug. 2011, pp. 839–843.
- [9] F. Gray, Pulse Code Communication U.S. 2 632 058, Mar. 17, 1953.
- [10] A. Jiang, V. Bohossian, and J. Bruck, “Rewriting codes for joint information storage in flash memories,” *IEEE Trans. Inf. Theory*, vol. 56, no. 10, pp. 5300–5313, Oct. 2010.
- [11] A. Jiang, M. Langberg, M. Schwartz, and J. Bruck, “Universal rewriting in constrained memories,” in *Proc. IEEE Int. Symp. Inf. Theory*, Seoul, Korea, Jun. 2009, pp. 1219–1223.
- [12] A. Jiang, R. Mateescu, M. Schwartz, and J. Bruck, “Rank modulation for flash memories,” *IEEE Trans. Inf. Theory*, vol. 55, no. 6, pp. 2659–2673, Jun. 2009.
- [13] A. Jiang, M. Schwartz, and J. Bruck, “Correcting charge-constrained errors in the rank-modulation scheme,” *IEEE Trans. Inf. Theory*, vol. 56, no. 5, pp. 2112–2120, May 2010.
- [14] M. Kendall and J. D. Gibbons, *Rank Correlation Methods*. New York: Oxford Univ. Press, 1990.
- [15] T. Klöve, T.-T. Lin, S.-C. Tsai, and W.-G. Tzeng, “Permutation arrays under the Chebyshev distance,” *IEEE Trans. Inf. Theory*, vol. 56, no. 6, pp. 2611–2617, Jun. 2010.
- [16] M. Mares and M. Straka, “Linear-time ranking of permutations,” in *Proc. 15th Annu. Eur. Conf. Algorithms*, 2007, pp. 187–193.

- [17] A. Mazumdar, A. Barg, and G. Zémor, "Constructions of rank modulation codes," in *Proc. IEEE Int. Symp. Inf. Theory*, St. Petersburg, Russia, Aug. 2011, pp. 834–838.
- [18] N. Papandreou, H. Pozidis, T. Mittelholzer, G. F. Close, M. Breiwisch, C. Lam, and E. Eleftheriou, "Drift-tolerant multilevel phase-change memory," in *Proc. 3rd IEEE Int. Memory Workshop*, Monterey, CA, May 2011, pp. 22–25.
- [19] C. D. Savage, "A survey of combinatorial Gray codes," *SIAM Rev.*, vol. 39, no. 4, pp. 605–629, Dec. 1997.
- [20] M. Schwartz and I. Tamo, "Optimal permutation anticodes with the infinity norm via permanents of $(0, 1)$ -matrices," *J. Combin. Theory Ser. A*, vol. 118, pp. 1761–1774, 2011.
- [21] R. Sedgewick, "Permutation generation methods," *ACM Comput. Surv.*, vol. 9, no. 2, pp. 137–164, Jun. 1977.
- [22] I. Tamo and M. Schwartz, "Correcting limited-magnitude errors in the rank-modulation scheme," *IEEE Trans. Inf. Theory*, vol. 56, no. 6, pp. 2551–2560, Jun. 2010.
- [23] Z. Wang and J. Bruck, "Partial rank modulation for flash memories," in *Proc. IEEE Int. Symp. Inf. Theory*, Austin, TX, Jun. 2010, pp. 864–868.
- [24] E. Yaakobi, A. Vardy, P. H. Siegel, and J. K. Wolf, "Multidimensional flash codes," presented at the presented at the 46th Annu. Allerton Conf. Commun. Control Comput., Monticello, IL, 2008.

Yonatan Yehezkeally was born in Israel in 1983. He received a double B.Sc. degree from the Mathematics Department and the Electrical and Computer Engineering Department, Ben-Gurion University of the Negev, Israel. He is now a graduate student with the Department of Electrical and Computer Engineering, Ben-Gurion University of the Negev. His research interests include algebraic coding, combinatorial structures, and finite group theory.

Moshe Schwartz (M'03–SM'10) was born in Israel in 1975. He received the B.A. (*summa cum laude*), M.Sc., and Ph.D. degrees from the Technion—Israel Institute of Technology, Haifa, Israel, in 1997, 1998, and 2004 respectively, all from the Computer Science Department.

He was a Fulbright post-doctoral researcher in the Department of Electrical and Computer Engineering, University of California San Diego, and a post-doctoral researcher in the Department of Electrical Engineering, California Institute of Technology. He now holds a position with the Department of Electrical and Computer Engineering, Ben-Gurion University of the Negev, Israel.

He received the 2009 IEEE Communications Society Best Paper Award in Signal Processing and Coding for Data Storage, and the 2010 IEEE Communications Society Best Student Paper Award in Signal Processing and Coding for Data Storage. His research interests include algebraic coding, combinatorial structures, and digital sequences.