Constructions of Permutation Arrays for Certain Scheduling Cost Measures

Joseph Naor*　　Ron M. Roth*

Abstract

Constructions of permutation arrays are presented that are optimal or nearly-optimal with respect to two cost measures: the so-called longest-jump measure and the longest-monotone-greedy-subsequence measure. These measures arise in the context of scheduling problems in asynchronous, shared memory, multiprocessor machines.

*Computer Science Department, Technion — Israel Institute of Technology, Haifa 32000, Israel.

1 Introduction

Let $[n]$ denote the set $\{1, 2, \ldots, n\}$ and let π and σ be two permutations over $[n]$. By a *cost measure of* π *with respect to* σ we refer to any real-valued function $(\pi, \sigma) \mapsto \text{cost}(\pi, \sigma)$. Given a cost measure $\text{cost}(\pi, \sigma)$, we define the cost of an $m \times n$ array P of m permutations $P_1, P_2, \ldots, P_m$ on $[n]$ with respect to σ by

$$\text{cost}(P, \sigma) \triangleq \sum_{i=1}^{m} \text{cost}(P_i, \sigma) .$$

The *maximal cost* of an $m \times n$ permutation array P is defined as $\text{cost}(P) \triangleq \max_\sigma \text{cost}(P, \sigma)$.

For a permutation π on $[n]$, denote by π_j the jth element of π, $j = 1, 2, \ldots, n$. Also, let π^{i} denote the inverse permutation of π. We extend these notations to $m \times n$ permutation arrays P by letting $P_{i,j}$ denote the jth element in the ith row of P, and P^{i} be the $m \times n$ permutation array for which $(P^{\text{i}})_i = (P_i)^{\text{i}}$, $i = 1, 2, \ldots, n$.

In this paper, we obtain constructions for $n \times n$ permutation arrays P with small maximal cost for each of the following two cost measures:

1. *The longest-jump measure,* defined by

   $$\text{cost}_1(\pi, \sigma) \triangleq n - \max_{j=1}^{n} \left((\sigma^{\text{i}})_j - (\pi^{\text{i}})_j \right) .$$

   That is, $\text{cost}_1(\pi, \sigma)$ is n minus the largest (signed) difference between the location indexes of the same element of $[n]$ in both σ and π. For example, for $\pi = (2, 1, 5, 3, 4)$ and $\sigma = (3, 1, 4, 2, 5)$ we have $(\sigma^{\text{i}})_1 - (\pi^{\text{i}})_1 = 0$, $(\sigma^{\text{i}})_2 - (\pi^{\text{i}})_2 = 3$, $(\sigma^{\text{i}})_3 - (\pi^{\text{i}})_3 = -3$, $(\sigma^{\text{i}})_4 - (\pi^{\text{i}})_4 = -2$, and $(\sigma^{\text{i}})_5 - (\pi^{\text{i}})_5 = 2$. Hence, in this case, $\text{cost}_1(\pi, \sigma) = 5 - 3 = 2$.

2. *The longest-monotone-greedy-subsequence measure* (in short, the longest-m.g.s. measure), which we will denote by $\text{cost}_2(\pi, \sigma)$ and define next. For a permutation π, construct the *longest monotone greedy subsequence* $j_1 < j_2 < \ldots < j_r$, of length r, with respect to σ as follows: Set $j_1 \triangleq 1$. Now, assume that $j = j_s$ has been determined. If $\pi_j = \sigma_n$ then $r = s \triangleq \text{cost}_2(\pi, \sigma)$. Otherwise, set j_{s+1} to be the smallest integer $k > j_s$ such that π_k appears after π_j in σ; i.e., k is the smallest integer greater than j such that $(\sigma^{\text{i}}\pi)_k > (\sigma^{\text{i}}\pi)_j$, where multiplication stands for permutation composition.

For example, the longest m.g.s. of $\pi = (2, 1, 5, 3, 4)$ with respect to $\sigma = (3, 1, 4, 2, 5)$ is $j_1 = 1$, $j_2 = 3$.

The following proposition is easily verified.

Proposition 1. *For every two permutations* π *and* σ *over* $[n]$, $\text{cost}_2(\pi, \sigma) \leq \text{cost}_1(\pi, \sigma)$.

Our treatment of these two cost measures has been motivated, in part, by previous work on task scheduling in asynchronous, shared-memory, faulty multiprocessor environments [5], [6], [7], [9], [11], [13], [14], [17]. Our model follows that of Anderson and Woll [5] and can be formulated as follows. Let P be an $n \times n$ permutation array whose ith row, P_i, is a permutation on $[n]$ that defines the order in which tasks, numbered $1, 2, \ldots, n$, in the system should be executed by processor i, $i = 1, 2, \ldots, n$. The order in which tasks are completed is determined by a permutation σ. The objective is to obtain a worst-case estimate of performance where the permutation σ maximizes a cost function that measures the work carried out by the processors.

Anderson and Woll used in [5] the longest-m.g.s. cost function for measuring that work. They showed that if an array is chosen uniformly at random from an ensemble of all $N \times N$ permutation arrays, then, with high probability, there exists an $N \times N$ permutation array Q for which $\text{cost}_2(Q) \leq c\, N \log N$, where c is an absolute constant. Up to a scalar multiplier, this maximal cost is the smallest possible for any $N \times N$ permutation array [8]. An $N \times N$ permutation array Q with the smallest maximal cost was then used by Anderson and Woll as a building block for constructing $n \times n$ permutation arrays $\Pi(n)$ with $\text{cost}_2(\Pi(n)) \leq n^{1+o_n(1)}$, where $o_n(1)$ stands for an expression which tends to zero as n goes to infinity. Taking n to be sufficiently large relative to N, an exhaustive search for the permutation array Q — and, subsequently, the overall time required to construct $\Pi(n)$ — can be made at most polynomially large in n.

In Section 2 we provide a construction of $n \times n$ permutation arrays $\mathcal{M}$ with $\text{cost}_1(\mathcal{M}) = O(n^{3/2})$. The optimality (up to a scalar multiplier) of this construction is established by proving a lower bound of $\Omega(n^{3/2})$ on the maximal cost of any $n \times n$ permutation array P with respect to the longest-jump measure.

In Section 3 we address the longest-m.g.s. measure. We first provide a construction for $n \times n$ permutation arrays with $\mathrm{cost}_2(\Pi(n)) \le n^{1+o_n(1)}$. Our basic construction is somewhat similar to the one of Anderson and Woll; however the presentation is different than the one in [5] and involves the definition of Kronecker product of permutation arrays. Then, we present a variation of the basic construction that avoids the need for a search for the aforementioned best $N \times N$ permutation array Q. To this end, we employ a technique similar to the one used in coding theory to obtain the well-known Justesen codes [10].

It should be noted that permutation arrays with other optimality properties were considered elsewhere. See, for example, Alon [3].

2 Constructions for the longest-jump measure

In this section we provide an optimal construction for $n \times n$ permutation arrays for the longest-jump measure. We start by deriving a construction for $n \times n$ permutation arrays with maximal cost $O(n^{5/3})$. Then, we apply this construction recursively to obtain permutation arrays of maximal cost $O(n^{3/2})$. We conclude this section by proving an $\Omega(n^{3/2})$ lower bound on the maximal cost according to the longest-jump measure, thus establishing the optimality of the construction.

Throughout this section, 'cost' means a cost with respect to the longest-jump measure.

2.1 A construction of maximal cost $O(n^{5/3})$

Let π be a permutation on $[n]$ and let k be a positive integer not greater than n. We define the *k-support* of π by the set $\mathcal{S}(\pi, k) \triangleq \{\pi_1, \pi_2, \ldots, \pi_k\}$ and its complement set $\overline{\mathcal{S}}(\pi, k)$ by $[n] - \mathcal{S}(\pi, k)$.

Let π and σ be two permutations on $[n]$. We say that π is *k-intersecting* with σ if $\mathcal{S}(\pi, k) \cap \overline{\mathcal{S}}(\sigma, n-k) \neq \emptyset$. Note that when π is k-intersecting with σ, then there exists an element $j \in [n]$ for which $(\sigma^{\mathrm{i}})_j \ge n-k+1$ and $(\pi^{\mathrm{i}})_j \le k$, implying $\mathrm{cost}_1(\pi, \sigma) \le n - ((\sigma^{\mathrm{i}})_j - (\pi^{\mathrm{i}})_j) \le 2k-1$.

Our goal is to construct a class of $n \times n$ permutation arrays $\mathcal{Q}(n)$, such that for *any* permutation σ over $[n]$, 'most' of the rows in $\mathcal{Q}(n)$ are k-intersecting with σ for a relatively 'small' k. To this end, we associate a permutation array $\mathcal{Q}(n,k)$ with a certain k-regular n-vertex directed graph $\mathcal{G}(n,k)$. Under such an association, a permutation σ defines a set S of k vertices in $\mathcal{G}(n,k)$, and the ith row in $\mathcal{Q}(n,k)$ is k-intersecting with σ if vertex i in $\mathcal{G}(n,k)$ is adjacent to some vertex in S. We choose the graph $\mathcal{G}(n,k)$ so as to guarantee that for any set S of k vertices in $\mathcal{G}(n,k)$, at least $n-k$ vertices in $\mathcal{G}(n,k)$ are adjacent to some vertex in S. We then take $\mathcal{Q}(n)$ to be a permutation array $\mathcal{Q}(n,k)$ with the smallest k for which such graphs $\mathcal{G}(n,k)$ exist.

We now describe the construction of $\mathcal{Q}(n,k)$ in detail. For a directed (respectively undirected) graph $G=(V,E)$ and a subset $S \subseteq V$, let $\mu(G,S)$ denote the number of vertices in G with no edges outgoing (respectively connecting) to any of the vertices of S. For a positive integer $r \le |V|$, let $\mu(G;r)$ be defined by

$$\mu(G;r) \stackrel{\triangle}{=} \max_{S \subseteq V \,:\, |S|=r} \mu(G,S) \; .$$

Let $\mathcal{G}(n,k)$ denote a k-regular n-vertex graph for which $\mu(\mathcal{G}(n,k);k) \le k$; a construction of such graphs is provided later on. Given $\mathcal{G}(n,k)$, we construct a permutation array $\mathcal{Q}(n,k)$ out of $\mathcal{G}(n,k)$ for which $\mathcal{S}\Big((\mathcal{Q}(n,k))_i\Big)$ is the set of terminal vertices of the outgoing edges from vertex i in $\mathcal{G}(n,k)$ for all $i=1,2,\ldots,n$. In other words, an edge $i \to j$ in $\mathcal{G}(n,k)$ implies that the element j appears in the k-prefix of $(\mathcal{Q}(n,k))_i$ (i.e., among the first k elements in the ith row of $\mathcal{Q}(n,k)$). The order of elements within the k-prefix, and respectively within the $(n-k)$-suffix, of each row of $\mathcal{Q}(n,k)$ is arbitrary and, therefore, the construction of $\mathcal{Q}(n,k)$ out of $\mathcal{G}(n,k)$ is not unique. Given a permutation σ, the number of rows in $\mathcal{Q}(n,k)$ which are not k-intersecting with σ is at most $\mu\Big(\mathcal{G}(n,k), \mathcal{S}(\sigma, n-k)\Big) \le \mu(\mathcal{G}(n,k);k) \le k$ and, therefore, the total cost with respect to σ of these non-k-intersecting rows is at most $k \cdot n$. On the other hand, the cost with respect to σ of a row in $\mathcal{Q}(n,k)$, which is k-intersecting with σ, is at most $2k-1$. Hence,

$$\mathrm{cost}(\mathcal{Q}(n,k)) \;\le k \cdot n \;+\; (2k-1)\, n \;=\; O(k \cdot n) \; . \tag{1}$$

To construct $\mathcal{G}(n,k)$ we make use of a generalization of the expander construction given by Lubotzky, Phillips, and Sarnak [15]. The next theorem, proved in [4], will be used to bound from above the value of $\mu(\mathcal{G}(n,k);k)$.

Theorem 1. [4]. *Let $H = (V, E)$ be an undirected Δ-regular graph with $|V| = n$ and let λ denote the second largest absolute value of any eigenvalue of the adjacency matrix of H. Let $S \subseteq V$ be of size $|S| = r$. Then,*

$$\mu(H, S) \le \frac{\lambda^2 (n-r) n}{r\Delta^2} \le \frac{\lambda^2 n^2}{r\Delta^2} .$$

In the construction of $\mathcal{G}(n, k)$ we use the generalization of the Lubotzky-Phillips-Sarnak construction as it appears in [4] and which is summarized here for the sake of completeness. Let p and q be two primes congruent to 1 modulo 4 such that p is a quadratic residue modulo q. Let h be a positive integer and denote by Z_{q^h} the ring of integers modulo q^h. Also, let $\mathcal{X}$ denote all 2×2 matrices over Z_{q^h} of determinant 1, where both matrices A and $-A$ are regarded as the same element $\pm A$. It can be readily verified that $\mathcal{X}$ contains $\frac{1}{2}(q^{3h} - q^{3h-2})$ elements.

Given p, q, and h, we now define the undirected graphs $H_{p,q,h}$ as Cayley graphs of $\mathcal{X}$ in the following manner: the vertices of $H_{p,q,h}$ are all $n = \frac{1}{2}(q^{3h} - q^{3h-2})$ elements of $\mathcal{X}$, and two such elements A and B are adjacent if and only if AB^{-1} is a matrix of the form

$$\pm\frac{1}{\sqrt{p}} \begin{bmatrix} a_0 + \imath a_1 & a_2 + \imath a_3 \\ -a_2 + \imath a_3 & a_0 - \imath a_1 \end{bmatrix} ,$$

where $\imath$ is an integer satisfying $\imath^2 \equiv -1 \pmod{q^h}$, a_0 is an odd positive integer, and a_1, a_2, a_3 are even integers satisfying $a_0^2 + a_1^2 + a_2^2 + a_3^2 = p$. (Note that the square root of p modulo q^h does exist; see [4] for details.) The resulting graphs $H_{p,q,h}$ turn out to be $(p+1)$-regular graphs with $n = \frac{1}{2}(q^{3h} - q^{3h-2})$ vertices and $\lambda(H_{p,q,h}) \le 2\sqrt{p}$ [15], [4].

Now fix q (to, say, 5) and let n be an integer of the form $n = \frac{1}{2}(q^{3h} - q^{3h-2})$ for some integer h. For such integers n, we construct the k-regular directed graphs $\mathcal{G}(n, k)$ as follows: let $p(k)$ be the largest prime smaller than k which is congruent to 1 modulo 4 and which is a quadratic residue modulo q (e.g., when $q = 5$, we take $p(k)$ to be the largest prime smaller than k which is congruent to 1 or 9 modulo 20). Now, to obtain $\mathcal{G}(n, k)$, make each edge in $H_{p(k),q,h}$ into two anti-parallel edges and add $k - p(k) - 1$ outgoing edges from each vertex in $H_{p(k),q,h}$ to obtain out-degree k at each vertex. Note that, by the distribution of prime numbers, we have $p(k) = k\,(1 - o_k(1))$, where $o_k(1)$ stands for an expression which tends to zero as k goes to infinity.

Lemma 1. *For every graph $\mathcal{G}(n,k)$ and for any positive integer $r \le n$,*

$$\mu(\mathcal{G}(n,k);r) \;\le\; \frac{4\,n^2}{r \cdot k\,(1-o_k(1))}\,.$$

Proof. By Theorem 1 we have, for each subset S of vertices of $H_{p,q,h}$,

$$\mu(H_{p,q,h},S) \;\le\; \frac{(\lambda(H_{p,q,h}))^2 n^2}{r\,(p+1)^2} \;\le\; \frac{4\,p\,n^2}{r\,(p+1)^2} \;<\; \frac{4\,n^2}{r\cdot k\,(1-o_k(1))}\,.$$

The lemma now follows by the inequality $\mu(\mathcal{G}(n,k),S) \le \mu(H_{p,q,h},S)$. □

Returning now to the discussion that led to Equation (1), we require that $\mu(\mathcal{G}(n,k);k)$ be at most k. By Lemma 1, we can guarantee that by requiring $k^3(1-o_k(1)) \ge 4\,n^2$. Hence, we can take $k = k(n) = O(n^{2/3})$, thus yielding an $n \times n$ permutation array $\mathcal{Q}(n) \triangleq \mathcal{Q}(n,k(n))$ with maximal cost $O(n^{5/3})$.

2.2 A construction of maximal cost $O(n^{3/2})$

In this section we obtain a construction of $n \times n$ permutation arrays with maximal cost $O(n^{3/2})$ by applying the construction of Section 2.1 recursively. More specifically, we construct a family of $n \times n$ permutation arrays $\mathcal{M}(n,\ell)$, $\ell \ge 0$, such that the maximal cost of any $\lfloor n/4^\ell \rfloor \times n$ sub-array of $\mathcal{M}(n,\ell)$ is $O(n^{3/2}/2^\ell)$. In particular, the maximal cost of $\mathcal{M}(n,0)$ will be $O(n^{3/2})$.

The arrays $\mathcal{M}(n,\ell)$ are obtained in decreasing order of ℓ as follows. Let $k_{n,\ell}$ denote the integer $2^{\ell+3}\lceil\sqrt{n}\,\rceil$. When $k_{n,\ell} \ge n$, we take $\mathcal{M}(n,\ell)$ to be an arbitrary $n \times n$ permutation array. As for smaller values of ℓ, assume that $\mathcal{M}(n,\ell+1)$ has been constructed. To obtain $\mathcal{M}(n,\ell)$, we first apply the construction of Section 2.1, using $\mathcal{G}(n,k_{n,\ell})$, to obtain an $n \times n$ permutation array $\mathcal{Q}(n,k_{n,\ell})$. Note that this requires that n be a value for which the construction $\mathcal{G}(n,k_{n,\ell})$ exists (e.g., $n = \frac{1}{2}(5^{3h} - 5^{3h-2})$). For $i = 1,2,\ldots,n$, we take the $k_{n,\ell}$-prefix of $(\mathcal{M}(n,\ell))_i$ to be the $k_{n,\ell}$-prefix of $(\mathcal{Q}(n,k_{n,\ell}))_i$. The $(n-k_{n,\ell})$-suffix of $(\mathcal{M}(n,\ell))_i$ is obtained by taking the last $n-k_{n,\ell}$ entries of $(\mathcal{Q}(n,k_{n,\ell}))_i$ permuted according to their order of appearance in $(\mathcal{M}(n,\ell+1))_i$.

Proposition 2. *Let n and ℓ denote integers for which the construction $\mathcal{G}(n,k_{n,\ell})$ exists.*

Then, the maximal cost of any $\lfloor n/4^\ell \rfloor \times n$ *sub-array of* $\mathcal{M}(n,\ell)$ *is* $O(n^{3/2}/2^\ell)$*. In particular,* $\mathrm{cost}_1(\mathcal{M}(n,0)) = O(n^{3/2})$.

Proof. Let $k \triangleq k_{n,\ell}$ and $m_\ell \triangleq \lfloor n/4^\ell \rfloor$. We prove by a backward induction on ℓ that the maximal cost of any $m_\ell \times n$ sub-array of $\mathcal{M}(n,\ell)$ is at most $28\, m_\ell \cdot 2^\ell \lceil \sqrt{n}\, \rceil$, whenever n is greater than some absolute constant N which is independent of ℓ. (We remark that the constant 28 is not the smallest possible, but it makes the proof relatively simple.)

The claim is trivial when $k \ge n$. We now consider smaller values of ℓ for which $k < n$. The next two observations can be easily verified:

- For $i = 1, 2, \ldots, n$ and for any permutation σ,

$$\mathrm{cost}_1\Big((\mathcal{M}(n,\ell))_i, \sigma\Big) \;\le\; \mathrm{cost}_1\Big((\mathcal{M}(n,\ell+1))_i, \sigma\Big) \;+\; k \;, \tag{2}$$

  since each element of $(\mathcal{M}(n,\ell+1))_i$ is "shifted" in $(\mathcal{M}(n,\ell))_i$ at most k positions to the right.

- If $(\mathcal{Q}(n,k))_i$ is k-intersecting with σ then $\mathrm{cost}_1\Big((\mathcal{M}(n,\ell))_i, \sigma\Big) \le 2k-1$.

Now fix some permutation σ. By Lemma 1, the number of rows in $\mathcal{Q}(n,k)$ that are non-k-intersecting with σ is at most $4n^2/(k^2(1-o_k(1))) \le n/4^{\ell+1}$ for every $n > N$, where N is an absolute integer independent of ℓ. In particular, the number of rows in $\mathcal{Q}(n,k)$ that are non-k-intersecting with σ is at most $m_{\ell+1}$. Note also that $m_{\ell+1} \le m_\ell/4$.

Let M be an $m_\ell \times n$ sub-array of $\mathcal{M}(n,\ell)$ and let M' be an $m_{\ell+1} \times n$ sub-array of M that includes all rows of M whose k-prefixes have been originated, while constructing $\mathcal{M}(n,\ell)$, from rows of $\mathcal{Q}(n,k)$ that are non-k-intersecting with σ. Also, denote by M'' the $m_{\ell+1} \times n$ sub-array of $\mathcal{M}(n,\ell+1)$ consisting of rows from which the $(n-k)$-suffixes of the rows of M' have been originated. By the induction hypothesis,

$$\mathrm{cost}_1(M'', \sigma) \;\le\; 28\, m_{\ell+1} \cdot 2^{\ell+1} \lceil \sqrt{n}\, \rceil \;.$$

However, by Equation (2) we have

$$\mathrm{cost}_1(M', \sigma) \;\le\; \mathrm{cost}_1(M'', \sigma) \;+\; m_{\ell+1} \cdot k \;.$$

Hence,

$$\begin{aligned}
\mathrm{cost}_1(M,\sigma) &< 28\, m_{\ell+1} \cdot 2^{\ell+1} \lceil \sqrt{n}\, \rceil \;+\; m_{\ell+1} \cdot k \;+\; (m_\ell - m_{\ell+1}) \cdot 2\, k \\
&= (16\, m_\ell + 48\, m_{\ell+1}) \cdot 2^\ell \lceil \sqrt{n}\, \rceil \\
&\le \Big(16\, m_\ell + 48\, (m_\ell/4)\Big) \cdot 2^\ell \lceil \sqrt{n}\, \rceil \\
&= 28\, m_\ell \cdot 2^\ell \lceil \sqrt{n}\, \rceil \;,
\end{aligned}$$

as claimed. □

It is easy to verify that the construction of $\mathcal{M}(n,0)$ can be carried out in polynomial time. Nevertheless, the description — and consequently the construction — of $\mathcal{M}(n,0)$ can be considerably simplified if we replace the graph sequence $\{\mathcal{G}(n,k_{n,\ell})\}_\ell$ by a sequence of n-vertex graphs $\{\mathcal{H}(n,\ell)\}_\ell$ which are nested; that is, $\mathcal{H}(n,\ell_1)$ is a subgraph of $\mathcal{H}(n,\ell_2)$ whenever $\ell_1 < \ell_2$.

One easy way to obtain such a graph sequence is as follows: Let $\ell_{\max}$ denote the largest integer ℓ for which $k_{n,\ell} \le n$. Set $\mathcal{H}(n,0) \triangleq \mathcal{G}(n,k_{n,0})$ and, for $0 < \ell \le \ell_{\max}$, let

$$\mathcal{H}(n,\ell) \;\triangleq\; \mathcal{H}(n,\ell-1) \cup \mathcal{G}(n,\ell-1) \;=\; \bigcup_{r=0}^{\ell-1} \mathcal{G}(n,k_{n,r}) \;,$$

where union of graphs should read as union of their edges (resulting parallel edges are merged into one). The out-degree of each vertex in $\mathcal{H}(n,\ell)$, $\ell > 0$, is bounded from above by $\sum_{r=0}^{\ell-1} k_{n,r} < k_{n,\ell} \triangleq k$ and, so, we can add edges to $\mathcal{H}(n,\ell)$ to have out-degree exactly k at each vertex. Furthermore, by Lemma 1 we have, for $\ell > 0$,

$$\mu(\mathcal{H}(n,\ell);k) \;\le\; \mu(\mathcal{G}(n,k_{n,\ell-1});k) \;\le\; \frac{8\, n^2}{k^2\,(1-o_k(1))} \;.$$

It can be readily verified that Proposition 2 still holds when we construct $\mathcal{M}(n,\ell)$ out of $\mathcal{H}(n,\ell)$, instead of $\mathcal{G}(n,k_{n,\ell})$. Once we construct $\mathcal{M}(n,0)$ this way, the description of each row $\pi = (\mathcal{M}(n,0))_i$ becomes quite straightforward:

- For the index range $1 \le j \le 8\lceil \sqrt{n}\, \rceil$, the elements π_j are the terminal vertices of the edges emanating from vertex i in $\mathcal{H}(n,0)$.
- For $0 < \ell \le \ell_{\max}$ and the index range $2^{\ell+2}\lceil \sqrt{n}\, \rceil < j \le 2^{\ell+3}\lceil \sqrt{n}\, \rceil$, the elements π_j are the terminal vertices of all edges which emanate from vertex i in $\mathcal{H}(n,\ell)$ and which do not appear in $\mathcal{H}(n,\ell-1)$.

- For the index range $2^{\ell_{\max}+3}\lceil\sqrt{n}\,\rceil < j \le n$ (if nonempty), the elements π_j can be set arbitrarily to complete row π into a permutation.

2.3 A lower bound of $\Omega(n^{3/2})$

Theorem 2. *For every $n \times n$ permutation array P,*

$$\mathrm{cost}_1(P) \;\ge\; \tfrac{1}{4}\, n^{3/2} \,.$$

Proof. Let P be an $n \times n$ permutation array and let $k \triangleq \lfloor \frac{1}{2}\sqrt{n} \rfloor$. Denote by A the $n \times n$ matrix over $\{0,1\}$ whose ith row is the characteristic vector of $\mathcal{S}(P_i,k)$; that is, $A_{i,j} = 1$ if $j \in \mathcal{S}(P_i,k)$ and $A_{i,j} = 0$ otherwise. Clearly, each row of A contains exactly k nonzero entries and, therefore, the total number of nonzero entries in A is $n \cdot k$.

Let X denote the set of indexes of those columns in A that contain at most $2k$ nonzero entries. It can be readily verified that $|X| \ge n/2$, or else there would have been more than $n \cdot k$ nonzero entries in A. Let X' be any subset of X of size k. The $n \times k$ sub-matrix $A_{X'}$ of A defined by the columns indexed by X' contains at most $2k^2$ nonzero entries. Hence, there exist at least $n - 2k^2 \ge n/2$ all-zero rows in $A_{X'}$. Denote by Y the set of indexes of the all-zero rows in $A_{X'}$. Note that for each $i \in Y$ we have

$$\mathcal{S}(P_i,k) \cap X' \;=\; \emptyset \,. \tag{3}$$

Now, take σ to be a permutation with $\overline{\mathcal{S}}(\sigma, n-k) = X'$. By (3) it follows that for every $i \in Y$, row P_i is non-k-intersecting with σ and, therefore, $\mathrm{cost}_1(P_i,\sigma) \ge k+1$. Hence,

$$\mathrm{cost}_1(P) \;\ge\; \mathrm{cost}_1(P,\sigma) \;\ge\; \sum_{i \in Y} \mathrm{cost}_1(P_i,\sigma) \;\ge\; |Y| \cdot (k+1) \;>\; \tfrac{1}{4}\, n^{3/2} \,. \qquad \square$$

We remark that by optimizing over the constants involved in the last proof, we can obtain a lower bound $\mathrm{cost}_1(P) \ge \frac{2+\epsilon}{3\sqrt{3}} n^{3/2}$ for any fixed $\epsilon > 0$ and sufficiently large n. However, for the sake of simplicity, we chose to present Theorem 2 in its current form.

3 The longest-monotone-greedy-subsequence case

In this section we first present a construction for $n \times n$ permutation arrays $\mathcal{P}(n)$ whose maximal cost with respect to the longest-m.g.s. measure is $O(n^{1+o_n(1)})$. We then discuss a variation of the basic construction that avoids the exhaustive search which is required in the basic scheme.

Throughout this section, 'cost' means a cost with respect to the longest m.g.s. measure.

3.1 Kronecker product of permutation arrays

Let π be a permutation on $[m]$ and χ be a permutation on $[n]$. The *Kronecker product* of π and χ is a permutation $\pi \otimes \chi$ on $[m\,n]$ defined by

$$(\pi \otimes \chi)_{(j-1)n+\ell} \triangleq (\pi_j - 1) \cdot n + \chi_\ell \, , \quad j = 1, 2, \ldots, m \, , \quad \ell = 1, 2, \ldots, n \, .$$

Similarly, for an $m \times m$ permutation array P and an $n \times n$ permutation array Q we define the Kronecker product of P and Q as the $(m\,n) \times (m\,n)$ permutation array $P \otimes Q$ given by

$$(P \otimes Q)_{(i-1)n+k,\,(j-1)n+\ell} = (P_i \otimes Q_k)_{(j-1)n+\ell} \, , \quad i, j = 1, 2, \ldots, m \, , \quad k, \ell = 1, 2, \ldots, n \, .$$

For example,

$$\begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix} \otimes \begin{pmatrix} 1 & 2 & 3 \\ 2 & 1 & 3 \\ 1 & 3 & 2 \end{pmatrix} = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 2 & 1 & 3 & 5 & 4 & 6 \\ 1 & 3 & 2 & 4 & 6 & 5 \\ 4 & 5 & 6 & 1 & 2 & 3 \\ 5 & 4 & 6 & 2 & 1 & 3 \\ 4 & 6 & 5 & 1 & 3 & 2 \end{pmatrix} .$$

Proposition 3. *For every two permutation arrays P and Q,*

$$\mathrm{cost}_2(P \otimes Q) \leq \mathrm{cost}_2(P) \cdot \mathrm{cost}_2(Q) \, .$$

Proof. Let σ be any permutation on $[m\,n]$. We construct a permutation τ on $[m]$ from σ as follows:

$$(\tau^{\mathrm{i}})_u < (\tau^{\mathrm{i}})_v \quad \text{if and only if} \quad \max_{\ell=1}^{n} \Big((\sigma^{\mathrm{i}})_{(u-1)n+\ell} \Big) < \max_{\ell=1}^{n} \Big((\sigma^{\mathrm{i}})_{(v-1)n+\ell} \Big) \, .$$

That is, an integer u has a smaller location index in τ than that of an integer v if and only if there exists an integer $a \in [n]$ such that $(u-1)n+1, (u-1)n+2, \ldots, un$ all have smaller location indexes in σ than that of $(v-1)n+a$.

Let $\mathcal{J} \subseteq [m]$ be the set of indexes in the longest m.g.s. of a row $\pi = P_i$ with respect to τ and let s be an integer in $[m]$.

Case 1: $s \notin \mathcal{J}$. We claim that none of the integers $(s-1)n+1, (s-1)n+2, \ldots, sn$ appears in the longest m.g.s. of a row $\pi \otimes Q_k$, with respect to σ, for all $k = 1, 2, \ldots, n$.

To see this, note that since s is not in the longest m.g.s. of π with respect to τ, then there exists some index $t < s$ such that

$$(\tau^{\mathrm{i}})_{\pi_s} \;=\; (\tau^{\mathrm{i}}\pi)_s \;<\; (\tau^{\mathrm{i}}\pi)_t \;=\; (\tau^{\mathrm{i}})_{\pi_t} \;.$$

Hence, by construction, there exists an integer $a \in [n]$ such that

$$\max_{\ell=1}^{n} \Big((\sigma^{\mathrm{i}})_{(\pi_s-1)n+\ell} \Big) \;<\; (\sigma^{\mathrm{i}})_{(\pi_t-1)n+a} \;. \tag{4}$$

Letting b_k denote the location index, $((Q_k)^{\mathrm{i}})_a$, of a in the row Q_k, we can rewrite (4) as

$$\Big(\sigma^{\mathrm{i}}\,(\pi \otimes Q_k)\Big)_{(s-1)n+\ell} \;<\; \Big(\sigma^{\mathrm{i}}\,(\pi \otimes Q_k)\Big)_{(t-1)n+b_k} \quad \text{for all } k, \ell = 1, 2, \ldots, n \;.$$

On the other hand, noting that $(t-1)n + b_k < (s-1)n + \ell$ for all $k, \ell = 1, 2, \ldots, n$, we conclude that none of the indexes $(s-1)n+\ell$ appears in the longest m.g.s. of a row $\pi \otimes Q_k$ with respect to σ for all $k, \ell = 1, 2, \ldots, n$.

Case 2: $s \in \mathcal{J}$. For $k = 1, 2, \ldots, n$, let $\mathcal{L}_k$ denote the set of indexes in the longest m.g.s. of $\pi \otimes Q_k$ with respect to σ and let

$$\mathcal{L}_k(s) \;\triangleq\; \{\; \ell \in [n] \;:\; (s-1)n+\ell \in \mathcal{L}_k \;\} \;.$$

Also, let χ be a permutation on $[n]$ defined by

$$(\chi^{\mathrm{i}})_u \;<\; (\chi^{\mathrm{i}})_v \quad \text{if and only if} \quad (\sigma^{\mathrm{i}})_{(\pi_s-1)n+u} \;<\; (\sigma^{\mathrm{i}})_{(\pi_s-1)n+v} \;.$$

It can be readily verified that $\mathcal{L}_k(s)$ is a set of indexes of a (not-necessarily longest) m.g.s. of Q_k with respect to χ and, therefore, $|\mathcal{L}_k(s)| \le \mathrm{cost}_2(Q_k, \chi)$. Hence,

$$\sum_{k=1}^{n} |\mathcal{L}_k(s)| \;\le\; \mathrm{cost}_2(Q, \chi) \;\le\; \mathrm{cost}_2(Q) \;.$$

Ranging over all values $s \in \mathcal{J}$, we thus obtain,

$$\begin{aligned}\mathrm{cost}_2(\pi \otimes Q, \sigma) &= \sum_{k=1}^{n} |\mathcal{L}_k| = \sum_{s \in \mathcal{J}} \sum_{k=1}^{n} |\mathcal{L}_k(s)| \\ &\leq |\mathcal{J}| \cdot \mathrm{cost}_2(Q) = \mathrm{cost}_2(\pi, \tau) \cdot \mathrm{cost}_2(Q) .\end{aligned}$$

Finally, summing over all rows π of P yields the desired inequality

$$\mathrm{cost}_2(P \otimes Q, \sigma) \leq \mathrm{cost}_2(P, \tau) \cdot \mathrm{cost}_2(Q) \leq \mathrm{cost}_2(P) \cdot \mathrm{cost}_2(Q) . \qquad \square$$

We point out the analogy between Proposition 3 and the technique suggested by Abbot in [1] to construct Ramsey graphs through graph products.

3.2 The basic construction

The basic construction of $n \times n$ permutation arrays with respect to the longest m.g.s. measure is obtained by applying Proposition 3 recursively. Let Q be an $N \times N$ permutation array and, for $n = N^m$, let $\mathcal{P}(n)$ denote the $n \times n$ permutation array obtained by taking the Kronecker product $Q \otimes Q \otimes \cdots \otimes Q$ of m copies of Q. Clearly,

$$\mathrm{cost}_2(\mathcal{P}(n)) \leq (\mathrm{cost}_2(Q))^m = n^{(\log \mathrm{cost}_2(Q))/\log N} . \tag{5}$$

Now, assume that Q has the smallest maximal cost among all $N \times N$ permutation arrays. By the probabilistic proof given in [5], there exists an absolute constant c such that $\mathrm{cost}_2(Q) \leq c\, N \log N$. Hence, by (5),

$$\mathrm{cost}_2(\mathcal{P}(n)) \leq n^{1+\epsilon(N)} , \tag{6}$$

where $\epsilon(N) = (\log(c \log N))/(\log N) = o_N(1)$. Taking m to be sufficiently large, we can afford an exhaustive search over all $(N!)^N$ permutation arrays for the optimal permutation array Q in time complexity which is polynomial in n. To this end, we can take N to be as large as $O(\sqrt{m}/\sqrt{\log m}) = O(\sqrt{\log n}/\log\log n)$, in which case we obtain $\epsilon(N) = O((\log\log\log n)/\log\log n) = o_n(1)$.

A stricter complexity requirement might be that of "explicitness", where the computation of each entry in the permutation array should be polynomial in the presentation of the indexes of that entry. In such a case, the exhaustive search for Q should be polylogarithmic in n and, therefore, N can be taken to be as large as $O(\sqrt{\log\log n}/\log\log\log n)$.

The resulting permutation array $\mathcal{P}(n)$ obtained by the above construction resembles the one given in [5]. Yet, our characterization of $\mathcal{P}(n)$ and the analysis of its cost through Kronecker product of arrays differ from the treatment in [5]. (For the sake of comparison, note that for an array P with the smallest maximal cost among all $n \times n$ permutation arrays we have $\mathrm{cost}_2(P) \le n^{1+\epsilon(n)}$, where $\epsilon(n) = ((\log\log n)/\log n) + O(1)$.)

3.3 Avoiding the search

To avoid the search for the $N \times N$ permutation array Q in the construction of Section 3.2, we apply a technique that has been used in coding theory in several applications, e.g., in constructing Justesen codes [10] that attain the Zyablov bound [18]: Instead of taking a Kronecker product of m copies of the *same* $N \times N$ permutation array with the smallest maximal cost, we take the Kronecker product of m *distinct* $N \times N$ permutation arrays. In fact, setting $m = (N!)^N$, we can take the Kronecker product of *all* $N \times N$ permutation arrays. Now, the probabilistic proof given in [5] implies the existence of not only one $N \times N$ permutation array Q with $\mathrm{cost}_2(Q) \le c\,N\log N$, but rather the existence of at least $(1 - \delta_N)(N!)^N$ such permutation arrays, where δ_N tends to zero at least as fast as $1/\log N$ when N goes to infinity. Taking $m = (N!)^N$, and letting $\mathcal{P}(N^m)$ be the Kronecker product of all $N \times N$ permutation arrays (according to some lexicographic order), we obtain,

$$\begin{aligned}\log\,\mathrm{cost}_2(\mathcal{P}(N^m)) &\le \sum_{N\times N \text{ arrays } P} \log\,\mathrm{cost}_2(P) \\ &\le m\cdot(1-\delta_N)\cdot\log\left(cN\log\,N\right) \;+\; m\cdot\delta_N\log(N^2)\;,\end{aligned}$$

as the maximal cost of those permutation arrays P with $\mathrm{cost}_2(P) \ge c\,N\log N$ is bounded from above by N^2. Hence, for $m = (N!)^N$ and $n = N^m$,

$$\begin{aligned}\frac{\log\,\mathrm{cost}_2(\mathcal{P}(n))}{\log\,n} &= \frac{m\,(1-\delta_N)\cdot\log\left(cN\log\,N\right)}{m\cdot\log\,N} \;+\; \frac{2m\,\delta_N\log N}{m\cdot\log\,N} \\ &\le 1 \;+\; \frac{O(1)+\log\log N}{\log N} \;+\; \delta_N \\ &= 1+\epsilon(N)\;,\end{aligned}$$

where $\epsilon(N)$ is essentially the same as in (6). Note that the computation of any entry in $\mathcal{P}(n)$ takes time which is polylogarithmic in n and that $n = N^{(N!)^N}$ implies $N = O(\sqrt{\log\log n}/\log\log\log n)$, as in the 'explicit' version of the construction of Section 3.2.

Acknowledgment

The authors thank Yehuda Afek, Noga Alon, and Gadi Taubenfeld for the many helpful discussions.

References

[1] H.L. Abbot, *A note on Ramsey's theorem, Canad. Math. Bull.,* 15 (1972), 9–10.

[2] Y. Afek, D. Greenberg, M. Merritt, and G. Taubenfeld, Private communication (1992).

[3] N. Alon, *Generating pseudo-random permutations and maximum flow algorithms, Inform. Proc. Lett.,* 35 (1990), 201–204.

[4] N. Alon, J. Bruck, J. Naor, M. Naor, and R.M. Roth, *Construction of asymptotically good low-rate error-correcting codes through pseudo-random graphs, IEEE Trans. Inform. Theory,* IT-38 (1992), 509–516.

[5] R.J. Anderson and H.S. Woll, *Wait-free algorithms for the union-find problem, Proc. 23rd ACM Symp. Theory of Computation* (1991), 370–380.

[6] Y. Aumann and M.O. Rabin, *Fast* PRAM *simulation on a fully asynchronous parallel system, Proc. 33rd IEEE Symp. Foundations of Computer Science* (1992), 147–156.

[7] M.F. Bridgland and R.J. Watro, *Fault tolerant decision making in totally asynchronous distributed systems, Proc. 6th Annual ACM Symp. Principles of Distributed Computing* (1987), 52–63.

[8] J. Buss and P. Ragde, *Certified write-all on a strongly asynchronous* PRAM, Manuscript (1990).

[9] C. Dwork, J. Halperin, and O. Waartz, *Performing work efficiently in the presence of faults, Proc. 11th Annual ACM Symp. Principles of Distributed Computing* (1992), 91–102.

[10] J. Justesen, *A class of constructive asymptotically good algebraic codes, IEEE Trans. Inform. Theory,* IT-18 (1972), 652–656.

[11] P. Kanellakis and A. Shvartsman, *Efficient parallel algorithms can be made robust, Proc. 8th Annual ACM Symp. Principles of Distributed Computing* (1989), 211–222.

[12] Z.M. Kedem, K.V. Palem, A. Raghunathan, and P.G. Spirakis, *Combining tentative and definite executions for very fast dependable parallel computing, Proc. 23rd ACM Symp. Theory of Computation* (1991), 381–390.

[13] Z.M. Kedem, K.V. Palem, M.O. Rabin, and A. Raghunathan, *Efficient program transformations for resilient parallel computation via randomization, Proc. 24th ACM Symp. Theory of Computation* (1992), 306–317.

[14] Z.M. Kedem, K.V. Palem, and P.G. Spirakis, *Efficient robust parallel computations, Proc. 22nd ACM Symp. Theory of Computation* (1990), 138–148.

[15] A. Lubotzky, R. Phillips, and P. Sarnak, *Explicit expanders and the Ramanujan conjecture, Proc. 18th Annual ACM Symp. Theory of Computation* (1986), 240–246. (See also: *Ramanujan graphs, Combinatorica,* 8 (1988), 261–277.)

[16] G.A. Margulis, *Explicit group-theoretical constructions of combinatorial schemes and their application to the design of expanders and superconcentrators, Probl. Inform. Trans.* 24 (1988), 39–46.

[17] C. Martel, R. Subramonian, and A. Park, *Asynchronous* PRAMs *are (almost) as good as synchronous* PRAMs*, Proc. 31st IEEE Symp. Foundations of Computer Science* (1990), 590–599.

[18] V.V. Zyablov, *An estimate of the complexity of constructing binary linear cascade codes, Probl. Inform. Trans.,* 7 (1971), 3–10.