



On Permutation Selectors and their Applications in Ad-Hoc Radio Networks Protocols

Jordan Kuschner, Yugarshi Shashwat, Sarthak Yadav, and Marek Chrobak^(✉)

University of California at Riverside, Riverside, USA
marek@cs.ucr.edu

Abstract. Selective families of sets, or selectors, are combinatorial tools used to “isolate” individual members of sets from some set family. Given a set X and an element $x \in X$, to isolate x from X , at least one of the sets in the selector must intersect X on exactly x . We study (k, N) -permutation selectors which have the property that they can isolate each element of each k -element subset of $\{0, 1, \dots, N - 1\}$ in each possible order. These selectors can be used in protocols for ad-hoc radio networks to more efficiently disseminate information along multiple hops. In 2004, Gasieniec, Radzik and Xin gave a construction of a (k, N) -permutation selector of size $O(k^2 \log^3 N)$. This paper improves this by providing a probabilistic construction of a (k, N) -permutation selector of size $O(k^2 \log N)$. Remarkably, this matches the asymptotic bound for standard strong (k, N) -selectors, that isolate each element of each set of size k , but with no restriction on the order. We then show that the use of our (k, N) -permutation selector improves the best running time for gossiping in ad-hoc radio networks by a poly-logarithmic factor.

1 Introduction

Selective families of sets, or selectors, are combinatorial tools used to “isolate” individual members of sets belonging to a given collection of sets. Given a set X and some element $x \in X$, to isolate x from X , at least one of the sets in the selector must intersect X on exactly x . Various types of selectors have been constructed in the literature and used in applications ranging from group testing to coding, bioinformatics, multiple-access channel protocols, and information dissemination in radio networks.

To illustrate this concept on a concrete example, consider the contention resolution problem in multiple-access channels (MACs) without feedback: N devices are connected to a shared broadcast channel (say, a radio frequency or ethernet), and each has a unique identifier from the set $U = \{0, 1, \dots, N - 1\}$. If exactly one device transmits in some time slot, then its message will be delivered through the channel to all other devices. However, if two or more devices transmit simultaneously, then a collision on the broadcast channel occurs. Channel collisions are indistinguishable from background noise, and the sender does not receive any feedback about the fate of its transmission.

Suppose that some unknown set of k devices initially activate and have messages to be transmitted. (See Fig. 1.) We seek a protocol that would allow these k active devices

Research supported by NSF grant CCF-2153723.

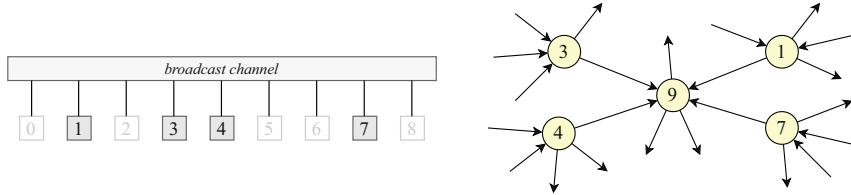


Fig. 1. On the left, an illustration of MAC contention resolution, with $N = 9$ and $k = 4$. Active devices are marked with darker colors. On the right, a node v in a radio network, with label 9 and four in-neighbors.

to transmit successfully, providing that the other devices remain idle. Without any feedback, each protocol for this model is non-adaptive, that is, it can be uniquely identified with the sequence of transmission sets $\bar{S} = S_0, S_1, \dots, S_{m-1}$, where each S_t is the set of devices allowed to transmit at time slots t , if active. A trivial protocol would avoid collisions altogether by using singleton transmission sets $\{0\}, \{1\}, \dots, \{N-1\}$, but this requires $m = N$ steps to complete k transmissions, which seems wasteful if k is small.

For an active device x to transmit successfully at a slot t of a transmission sequence \bar{S} , x must be the only active device in set S_t . Therefore \bar{S} must satisfy the following property: for any set $X \subseteq U$ of cardinality k and any $x \in X$ there is $t \in \{0, 1, \dots, m-1\}$ such that $S_t \cap X = \{x\}$. A family \bar{S} that satisfies this condition is called a *strong (k, N) -selector*. In other words, a strong (k, N) -selector isolates each element of each subset of U of cardinality k . Optimizing the time to complete all transmissions in the above MAC model is equivalent to constructing a strong (k, N) -selector of minimum size m .

Very similar challenges arise in information dissemination protocols for ad-hoc radio networks. A radio network is a directed graph whose nodes represent radio transmitters/receivers and edges represent their transmission ranges. There are n nodes, each assigned a unique label from $U = \{0, 1, \dots, N-1\}$. When a node transmits, its message is sent to all its out-neighbors; however, the possibility of collisions at the out-neighbors can result in the loss of transmitted messages. The ad-hoc property dictates that the nodes are oblivious to the network's topology at the beginning. In particular, in the scenario with a node v having k in-neighbors that attempt to transmit to v , the challenge is equivalent to the MAC contention resolution. (See Fig. 1.) A protocol that applies a strong (k, N) -selector will guarantee that in at most m steps all in-neighbors of v will successfully transmit their messages to it. This property is particularly useful in the problem of *gossiping* (full information exchange), where all nodes in the network have some information that needs to be delivered to all other nodes in the network. (The formal definition of the gossiping problem is given in Sect. 3.)

Strong (k, N) -selectors have been well studied in the literature (see the discussion at the end of this section). It is known that there are strong (k, N) -selectors of size $m = O(k^2 \log N)$, beating the earlier-mentioned trivial bound of N if k is small.

Our Contribution. We study a new type of selectors called (k, N) -permutation selectors. We earlier saw that a strong (k, N) -selector isolates all elements of each k -element

set $X \subseteq U$. The (k, N) -permutation selector guarantees an extra property that for each possible permutation of X , it isolates all $x \in X$ in the order of this permutation.

The motivation comes from radio networks, where, unlike in the MAC contention resolution problem, a message may need to travel along a path of multiple nodes. At each node v along the path, the in-neighbors of v need to overcome contention to successfully transmit to v . Consider such a path $P = v_0 v_1 \dots v_s$, and suppose that the in-neighborhood X of P (the set of nodes with edges going to P) has at most k nodes. (See Fig. 2 for illustration.) An s -fold repetition of a strong (k, N) -selector will deliver a message from v_0 to v_s in time $O(sk^2 \log N) = O(k^3 \log N)$. A (k, N) -permutation selector of size m can achieve this in time m , because it guarantees that in m steps the nodes v_0, v_1, \dots, v_{s-1} will successfully transmit, one by one, *in this order*.

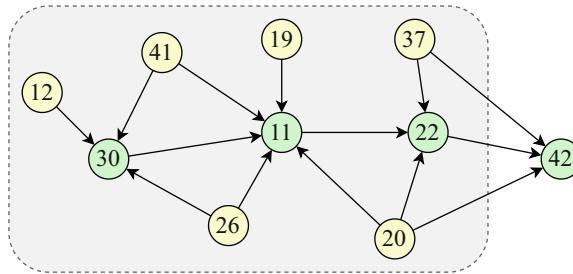


Fig. 2. An example of a path and its in-neighborhood, with nodes identified by their labels. The path is $P = 30, 11, 22, 42$ and its neighborhood (shaded) has 9 nodes.

The concept of (k, N) -permutation selectors is implicit in the work by Gasieniec, Radzik and Xin [10]. They show that a (k, N) -permutation selector can be constructed by interleaving other known types of selectors¹. Their construction is called a *path selector* and it has size $O(k^2 \log N \log^2 k)$. Thus, assuming that N is polynomial in the network size n , a protocol based on their path selector will deliver a message along a path with in-neighborhood of size at most k in time $O(k^2 \log^3 n)$. They used this idea to give an $O(n^{4/3} \log^{10/3} n)$ -time protocol for gossiping in ad-hoc radio networks. This is the best currently known bound for this problem.

The main result of our paper is an improved upper bound on the size of (k, N) -permutation selectors. Let $2 \leq k \leq N$. Using a probabilistic construction, we prove the following theorem in Sect. 2.

Theorem 1. *There is a (k, N) -permutation selector of size $m = O(k^2 \log N)$.*

This bound matches the best bound on strong (k, N) -selectors. This seems surprising, as (k, N) -permutation selectors offer more capability: they can isolate any k elements in any given order, while strong selectors will only do so in some unknown order. Theorem 1 leads to a poly-logarithmic improvement of the time complexity of

¹ We remark that in [10] the authors use a different style for specifying the parameters of selectors. The notation in our paper follows the convention from [7].

gossiping in ad-hoc radio networks. A further minor improvement can be achieved by using a faster procedure for broadcasting [6] inside the gossiping protocol. This leads to the following result:

Theorem 2. *If N is polynomial in n , then the gossiping problem in ad-hoc radio networks can be solved in time $O(n^{4/3} \log^2 n (\log \log n)^{2/3})$.*

In Sect. 4 we consider even more general structures that we named (k, q, N) -permutation selectors. These are defined by the following property: for each permutation of each k -element set X , the selector isolates some q -elements of X in the order of this permutation. This naturally extends the concept of (k, q, N) -selectors, which isolate q elements of a k -element set without restricting the order. (See the discussion below.) Extending the proof of Theorem 1, we show that there exists a (k, q, N) -permutation selector of size $O(kq \log N)$.

Related Work. Combinatorial structures closely related to selectors have been studied in different settings and under different terminology, with connections between these concepts not always obvious. Some of these structures are equivalent to strong (k, N) -selectors, while others are related to the concept of *weak (k, N) -selectors*, which only isolate any one element of each k -element set. Examples include superimposed codes used in information retrieval [12], cover-free set families [9], as well as protocols for non-adaptive group testing [11] and for MAC contention resolution [13]. The use of selectors in protocols for ad-hoc radio networks was initiated in [1, 3], and various forms of selectors have been used in essentially all deterministic protocols for information dissemination in such networks.

One classical example of applications of selectors outside of networking is group testing, a method used in fields such as medical diagnostics, quality control, or bioinformatics. The goal in group testing is to efficiently identify individuals who test positive for a specific trait, such as a disease. Instead of testing these individuals separately, group testing works by pooling individuals into groups for collective testing. With this approach, selectors can be employed to minimize the number of required tests. For more thorough discussion of applications of selectors, see [7] and the references therein.

For strong (k, N) -selectors, the upper bound of $O(k^2 \log N)$ can be established by a probabilistic construction [9, 12, 13]. An explicit construction matching this bound can be found in [16]. Note that such bounds are of interest only for $k = O(\sqrt{N} / \log N)$, because for larger values of k a trivial construction using N singletons is better. This $O(k^2 \log N)$ bound is essentially tight, given a lower bound of $\Omega(k^2 \log^{-1} k \log N)$ (for $k = O(\sqrt{N})$) established in [2, 5, 8].

De Bonis et al. [7] introduced a more general model of selectors called (k, q, N) -selectors. A (k, q, N) -selector has the property that it can isolate at least q elements from each k -element subset of U . A strong (k, N) -selector is a special case when $q = k$, and a weak (k, N) -selector is a special case when $q = 1$. They proved that there are (k, q, N) -selectors of size $O(k^2 / (k - q + 1) \log N)$.

2 Permutation Selectors

The objective in this section is to prove Theorem 1, namely to show that there is a (k, N) -permutation selector of size $m = O(k^2 \log N)$.

We start with a more formal definition of our permutation selectors. Recall that $U = \{0, 1, \dots, N-1\}$, and let $2 \leq k \leq N$. Consider a sequence $\bar{S} = S_0, S_1, \dots, S_{m-1}$ of subsets of U . We call \bar{S} a (k, N) -permutation selector if it has the following property:

(PS) For every $X \subseteq U$ with $|X| = k$ and for each permutation $\pi = x_1, x_2, \dots, x_k$ of X , there exists an increasing sequence of indices $0 \leq i_1 < i_2 < \dots < i_k \leq m-1$ such that S_{i_l} isolates x_l from X for each $l = 1, 2, \dots, k$.

If \bar{S} satisfies this property for a set X and a permutation π , we will say that it *isolates X in order π* , or simply that \bar{S} isolates π . Thus \bar{S} is a (k, N) -permutation selector if it isolates each k -permutation of U .

The proof idea is to choose each set S_i randomly, by having each label in U , independently, add itself to S_i with probability $\frac{1}{k}$. This way, for any fixed set X of cardinality k , each S_i will isolate some element of X with probability $\gamma = (1 - 1/k)^{k-1}$, so γ is in the range $\frac{1}{e} < \gamma \leq \frac{1}{2}$. (Recall that $k \geq 2$.) We then show that the probability that X is not isolated in order π is exponentially small with respect to m/k . This, using the union bound, will give us that the probability that \bar{S} is not a (k, N) -permutation selector is less than 1 if $m = \Theta(k^2 \log N)$, with a sufficiently large constant hidden behind the big-Theta. Therefore, some sequence \bar{S} of length $m = O(k^2 \log N)$ is a (k, N) -permutation selector.

We now proceed with the details. As in the first part of the proof we fix X and π , we can as well assume that $X = \{0, 1, \dots, k-1\}$ and that the desired permutation of X is $\pi = 0, 1, \dots, k-1$. Our first goal is to estimate the probability that \bar{S} does not isolate X in order π .

To this end, we start by considering an auxiliary problem, which is basically a variant of coupon collection: Suppose that we generate uniformly a random sequence $\bar{r} = r_0, r_1, \dots, r_{\ell-1}$ of elements of $\{0, 1, \dots, k-1\}$, where $\ell \geq k \geq 2$. We want to compute the probability $p_{\ell, k}$ of the event “ \bar{r} does not contain π as a subsequence”. To compute $p_{\ell, k}$, we reason as follows. For a given $j \in \{0, 1, \dots, k-1\}$, if \bar{r} contains $0, 1, \dots, j-1$ as a sub-sequence, associate with \bar{r} the unique *lexicographically-first appearance* of $0, 1, \dots, j-1$, namely the increasing sequence i_0, i_1, \dots, i_{j-1} of positions in \bar{r} , where i_0 is the position of the first 0, i_1 is the first position of 1 after i_0 , and so on. The number of \bar{r} ’s that contain $0, 1, 2, \dots, j-1$ but not $0, 1, 2, \dots, j$ can be computed by multiplying the number of choices, $\binom{\ell}{j}$, for the lexicographically-first appearance of $0, 1, \dots, j-1$, and the number of ways for filling the remaining $\ell - j$ positions, which is $(k-1)^{\ell-j}$. This yields the formula for $p_{\ell, k}$, given below, for which we then derive an upper bound estimate.

$$\begin{aligned}
 p_{\ell, k} &= \frac{1}{k^\ell} \cdot \sum_{j=0}^{k-1} \binom{\ell}{j} (k-1)^{\ell-j} \\
 &= (1 - 1/k)^\ell \cdot \sum_{j=0}^{k-1} \binom{\ell}{j} (k-1)^{-j} \\
 &< e^{-\ell/k} \cdot \sum_{j=0}^{k-1} \left(\frac{\ell}{k-1} \right)^j \leq e^{-\ell/k} (2\ell/k)^k,
 \end{aligned} \tag{1}$$

where the last step follows from $\ell/(k-1) \leq 2\ell/k$ and $\sum_{j=0}^{k-1} (2\ell/k)^j = (2\ell/k)^k/(2\ell/k-1) \leq (2\ell/k)^k$, as $\ell \geq k$.

Next, still with X and π fixed as above, we estimate the probability that \bar{S} does not isolate X in order π . Let h be the random variable representing the number of sets S_i that isolate some element of X (with repetitions counted). As stated earlier, the probability that some element of X is isolated by a given S_i is $\gamma \in (\frac{1}{e}, \frac{1}{2}]$. So h is a binomial random variable with success probability γ and mean $\mu = \gamma m$. Therefore, letting $\delta = 1 - 1/(4\gamma)$ (note that $\delta \in [\frac{1}{2}, 1 - e/4)$) and using the Chernoff bound for the lower tail of h 's distribution, we get

$$\Pr[h \leq m/4] = \Pr[h \leq (1 - \delta)\mu] \leq e^{-\delta^2\mu/2} = e^{-\delta^2\gamma m/2} = \alpha^m, \quad (2)$$

where $\alpha = e^{-\delta^2\gamma/2} \in (0, 1)$.

Consider now the length- h sequence consisting of elements of X (with repetitions) that are isolated by \bar{S} , in order in which they are isolated. For an integer $\ell \geq k$, let E_ℓ be the event that the first $\min(\ell, h)$ elements of this sequence do not contain our permutation π as a subsequence. Then $\Pr[E_\ell | h \geq \ell] = p_{\ell, k}$, so

$$\begin{aligned} \Pr[\bar{S} \text{ does not isolate } X \text{ in order } \pi] &\leq \Pr[h \leq m/4] + \Pr[E_{m/4} | h \geq m/4] \\ &\leq \alpha^m + p_{m/4, k} \\ &\leq \alpha^m + e^{-m/4k} (m/2k)^k \\ &\leq \beta^{m/k} (m/k)^k, \end{aligned} \quad (3)$$

for $\beta = \max(\alpha, e^{-1/4})$. (We can assume that $m \geq 2k$.) Clearly, $\beta \in (0, 1)$.

Let $m = ck^2 \log N$, for some constant c that will be specified later. To upper bound the probability that \bar{S} is not a (k, N) -permutation selector, we apply the union bound. We have $\binom{N}{k}$ choices of k -element sets X , and each can be permuted in $k!$ ways. So, using the bound (3), we obtain

$$\begin{aligned} \Pr[\bar{S} \text{ is not a } (k, N)\text{-permutation selector}] &\leq \binom{N}{k} \cdot k! \cdot [\beta^{m/k} (m/k)^k] \\ &\leq \binom{N}{k} \cdot k! \cdot \beta^{ck \log N} \cdot (ck \log N)^k \\ &\leq N^k \cdot N^k \cdot \beta^{ck \log N} \cdot c^{k \log N} \cdot N^k \cdot N^k \\ &= N^{4k} \cdot (c\beta^c)^{k \log N} < 1, \end{aligned} \quad (4)$$

where the last inequality holds if c is large enough so that $c\beta^c < \frac{1}{16}$. This proves that for this c and $m = ck^2 \log N$ there exists a (k, N) -permutation selector \bar{S} of length m .

3 Gossiping in Ad-Hoc Radio Networks

In this section we prove Theorem 2. We start by giving a formal description of the ad-hoc radio network model. The gossiping protocol we use is essentially identical to the one in [10], but we include its high-level description and sketch the analysis, for the sake of completeness.

Ad-Hoc Radio Network Model. A radio network can be naturally modeled as a directed graph $G = (V, E)$ whose nodes represent processing elements equipped with radio transmitters/receivers. Each node has a unique label from the set $U = \{0, 1, \dots, N - 1\}$. The directed edges represent the nodes' transmission ranges, that is $(u, v) \in E$ iff v is in the range of the transmissions from u . If $(u, v) \in E$ then v is called an *out-neighbor* of u and u is an *in-neighbor* of v .

Initially, all nodes know only their own label and the upper bound N on the number of labels. They do not have any information about the network's topology. The time is discrete, divided into equal-length time steps. If a node u transmits a message, this message is sent to all its out-neighbors at the same time step. If v is one of these out-neighbors, and u is the only in-neighbor of v transmitting at this step, then v will receive u 's message. But if some other in-neighbor of v transmits at the same step, a *collision* occurs. The model does not assume any collision detection capability, so u will not receive any collision notification and v will not know that u transmitted. There are no restrictions on message size or local computation.

The two most basic information dissemination primitives in this model are broadcasting and gossiping. In broadcasting (or one-to-all communication) the goal is to deliver a message from a designated *source* node to all other nodes. In gossiping (or all-to-all communication) each node starts with its own piece of information that we call a *rumor*, and the rumors from each node must be delivered to all other nodes. For these problems to be well defined, G must satisfy appropriate connectivity assumptions: in broadcasting all nodes must be reachable from the source node, and in gossiping the network must be strongly connected.

The Gossiping Protocol. We describe the protocol under the assumption that $N = n$, and later we will show how to extend it to the general case, when N is polynomial in n . With this assumption, a trivial gossiping protocol would broadcast the rumors from nodes labeled $0, 1, \dots, n - 1$ one by one. Denoting by $B(n)$ the running time of a broadcasting protocol, this would take time $O(nB(n))$. To speed this up, the gossiping protocols in [4, 10, 14] work by grouping rumors in some nodes so that these nodes can broadcast the collected rumors in a single message.

To reduce the number of such broadcasts, the idea is to broadcast from nodes that have many rumors. We call a rumor *active* if it has not been broadcast yet, and a node is *active* if its rumor is active and *dormant* otherwise. We use procedure $\text{DISPERSE}(\mu)$, which repeatedly chooses a node with at least μ active rumors and then broadcasts from this node. Choosing such a node can be accomplished with broadcasting and binary search [4]. The number of chosen nodes will be $O(n/\mu)$, so the total running time of $\text{DISPERSE}(\mu)$ is $O((n/\mu)B(n) \log n)$.

One clever observation in [10] is that it is sufficient to give a protocol for a task called *quasi-gossiping*, where each node needs to either become dormant itself or have its rumor delivered to a dormant node. This is because a single repetition of the transmission sequence from the quasi-gossiping protocol will in fact complete full gossiping. The pseudo-code for the quasi-gossiping protocol is given in Algorithm 1. It uses a parameter κ whose value will be determined later.

The correctness of Protocol QUASIGOSSIP is justified by focussing on *active paths*, which are paths consisting only of active nodes. For any active path, define its *active in-*

Algorithm 1. QUASI GOSSIP

```

1: for  $v = 0, 1, \dots, n - 1$  do
2:   transmit from node  $v$ 
3:   DISPERSE( $\kappa$ )
4: repeat  $\log \kappa + 1$  times
5:   the active nodes transmit according to a  $(\kappa, n)$ -permutation selector
6:   DISPERSE( $\kappa/2$ )

```

neighborhood to be the set of active in-neighbors of the nodes on this path. Let ℓ be the largest number such that each active path with ℓ nodes has its active in-neighborhood size (strictly) smaller than κ . After the execution of DISPERSE(κ) in Line 3, each node will be left with fewer than κ active in-neighbors, so at this point we have $\ell \geq 1$. Each iteration of the **repeat** loop at least doubles the value of ℓ . Since all nodes (except possibly the last) on an active path belong to the active in-neighborhood, the value of ℓ cannot exceed κ . So after $\log \kappa$ iterations there cannot be any active paths left of length at least κ , and then the last iteration will complete the quasi-gossiping task.

In line 5 we use the (κ, n) -permutation selector from Theorem 1, so the total cost of these selectors will be $O(\kappa^2 \log^2 n)$. The cost of all calls to DISPERSE() is $O((n/\kappa + \log n)B(n) \log n)$. Thus the overall running time of Protocol QUASI GOSSIP is asymptotically bounded by:

$$(n/\kappa + \log n)B(n) \log n + \kappa^2 \log^2 n$$

Letting $\kappa = (nB(n)/\log n)^{1/3}$, and using the bound $B(n) = O(n \log n \log \log n)$ from [6, 15] on the complexity of broadcasting, we obtain a gossiping protocol with running time $O(n^{4/3} \log^2 n (\log \log n)^{2/3})$.

To extend Protocol QUASI GOSSIP to the case when N is polynomial in n , we first replace n by N in all invocations of selectors. After this, the only problematic part of Protocol QUASI GOSSIP is the **for** loop in Lines 1–2 that reduces the active in-neighborhoods of the individual nodes. Instead of this loop, the protocol in [10] uses $(s, s/4, N)$ -selectors for geometrically decreasing values of s , combined with operations DISPERSE($s/4$), to gradually reduce these in-neighborhoods. (This is similar to the method in [13].) The running time of this process amortizes to $O((n/\kappa)B(n) \log n)$ —same time as for the case of small labels. So the overall running time also remains the same, completing the proof of Theorem 2.

4 (k, q, N) -Permutation Selectors

In Sect. 2 we defined (N, k) -permutation selectors, a family of sets that isolates all elements of any k -permutation π of the label set $U = \{0, 1, \dots, N - 1\}$, meaning that all elements in π are isolated in their listed order in π .

We now generalize this concept analogously to the way (k, q, N) -selectors generalize standard selectors, by requiring that some q elements of each k -permutation π are isolated in the order of π . Formally, let $\bar{S} = S_0, S_1, \dots, S_{m-1}$ be a sequence of subsets of U . \bar{S} is called a (k, q, N) -permutation selector if the following property holds:

(PS') For each $X \subseteq U$ with $|X| = k$, and for each permutation $\pi = x_1, x_2, \dots, x_k$ of X , there are increasing sequences of indices $0 \leq i_1 < i_2 < \dots < i_q \leq m-1$ and $d_1 < d_2 < \dots < d_q$ such that, for all $l = 1, 2, \dots, q$, set S_{i_l} isolates x_{d_l} from X .

This section provides a probabilistic construction of a (k, q, N) -permutation selector, proving the following theorem. Let $2 \leq k \leq N$ and $1 \leq q \leq k$.

Theorem 3. *There exists a (k, q, N) -permutation selector $\bar{S} = S_0, S_1, \dots, S_{m-1}$ of size $m = O(kq \log N)$.*

Proof. As in the proof of Theorem 1, we use a probabilistic argument. The construction is the same: for each $j = 0, 1, \dots, m-1$, we let S_j be a random subset of U obtained by each element $x \in U$ adding itself to S_j , independently, with probability $\frac{1}{k}$. We then need to prove that if c is a sufficiently large constant and $m = c \cdot qk \log N$, then

$$\Pr[\bar{S} \text{ is not a } (k, q, N)\text{-permutation selector}] < 1. \quad (5)$$

The proof's high-level strategy is the same as in the proof of Theorem 1. We use the Chernoff bound to reduce the problem to a version of the coupon collection problem. We consider the following variant of the coupon collection problem: For a random sequence $\bar{r} = r_0, r_1, \dots, r_{\ell-1}$ of coupons from $\{0, 1, \dots, k-1\}$, where $\ell \geq k \geq 2$, let $p'_{\ell, k, q}$ be the probability of the event “ \bar{r} does not contain an increasing subsequence of length q ”. To show (5), it is sufficient to prove that

$$p'_{\ell, k, q} \leq \gamma^{\ell/q} (2\ell/q)^q, \quad (6)$$

for some constant $\gamma \in (0, 1)$. Indeed, this is analogous to (1). We can then use the Chernoff-based bound (2), to obtain a bound of $\delta^{m/q} (m/q)^q$, for some $0 < \delta < 1$, on the probability that \bar{S} does not isolate q elements of π in order, analogously to (3). For the union-bound estimate, we then take $m = ckq \log N$, with large enough c . Since then $\delta^{m/q} (m/q)^q = \delta^{ck \log N} (ck \log N)^q \leq \delta^{ck \log N} (ck \log N)^k$, the derivation for the union bound (4) will be essentially the same, with δ instead of β .

It remains to estimate $p'_{\ell, k, q}$. We reason as follows. First, to avoid clutter in the calculations below, we will assume that q is a divisor of k . (If it is not, the values of k/q in the calculations below need to be rounded up or down. This does not affect our asymptotic estimate.)

We now consider a special type of increasing sub-sequences that we refer to as *q -jump sub-sequences*. Partition the set of coupons $\{0, 1, \dots, k-1\}$ into q equal size blocks B_0, B_1, \dots, B_{q-1} , each having k/q consecutive coupons. That is, the h -th block is $B_h = \left\{ h \frac{k}{q}, h \frac{k}{q} + 1, \dots, (h+1) \frac{k}{q} - 1 \right\}$, for $h = 0, 1, \dots, q-1$. For $j \leq q$, a jump sub-sequence is a sequence of j coupons a_0, a_1, \dots, a_{j-1} such that $a_h \in B_h$ for each $h = 0, \dots, j-1$.

Define now $p''_{\ell, k, q}$ to be the probability of the event “ \bar{r} does not contain a jump sub-sequence of length q ”. Since $p''_{\ell, k, q} \geq p'_{\ell, k, q}$, it is sufficient to prove that the same inequality (6) holds for $p''_{\ell, k, q}$ instead of $p'_{\ell, k, q}$. We do this by refining the argument in the proof of Theorem 1. For a given $j \in \{0, 1, \dots, q-1\}$, if \bar{r} contains a length- j jump sub-sequence then associate with \bar{r} the unique *lexicographically-first appearance*

of a length- j jump sub-sequence. The number of \bar{r} 's that contain a length- j jump sub-sequence but not a length- $(j+1)$ jump sub-sequence can be obtained by choosing the lexicographically first length- j jump sub-sequence in $\binom{\ell}{j} \cdot (k/q)^j$ ways and multiplying it by the number of ways of filling the remaining $\ell-j$ positions, which is $(k-k/q)^{\ell-j}$. We thus have

$$\begin{aligned} p'_{\ell,k,q} &\leq p''_{\ell,k,q} = \frac{1}{k^\ell} \sum_{j=0}^{q-1} \binom{\ell}{j} \cdot (k/q)^j \cdot (k-k/q)^{\ell-j} \\ &= (1-1/q)^\ell \sum_{j=0}^{q-1} \binom{\ell}{j} (q-1)^{-j} \\ &\leq (1-1/q)^\ell \sum_{j=0}^{q-1} (2\ell/q)^j \\ &\leq e^{-\ell/q} (2\ell/q)^q. \end{aligned}$$

This proves (6) (with $\gamma = 1/e$), completing the proof.

For $q = k$, the bound in Theorem 3 matches the bound from Theorem 1 and the best bound for strong (k, N) -selectors. For $q = 1$, it also matches the $O(k \log N)$ bound for weak (k, N) -selectors. We are not sure about the optimum size of (k, q, N) -permutation selectors for intermediate values of q . The case when $q = \sqrt{k}$ is particularly interesting. From the bound in [7], (k, \sqrt{k}, N) -selectors have size $O(k \log N)$. Can some probabilistic construction produce a (k, \sqrt{k}, N) -permutation selector of size $m = \tilde{O}(k)$? For such m , a random sequence of coupons from $\{0, 1, \dots, k-1\}$ can be thought of as a “noisy” permutation. We would need to show that the probability that this random sequence does not contain an increasing sub-sequence of length \sqrt{k} is exponentially small. This question is closely related to Ulam’s Problem about the distribution of the longest increasing subsequence (LIS) in a random permutation. Ulam’s Problem has been extensively studied, and it is known that for permutations of $0, 1, \dots, k-1$ the expected length of LIS is $2\sqrt{k} + O(k^{1/6})$. (See [17], for example.) To our knowledge, however, the published concentration bounds are not sufficient for refining the probabilistic construction in the proof of Theorem 3 to yield a better bound.

References

1. Basagni, S., Bruschi, D., Chlamtac, I.: A mobility-transparent deterministic broadcast mechanism for ad hoc networks. *IEEE/ACM Trans. Networking* **7**(6), 799–807 (1999)
2. Chaudhuri, S., Radhakrishnan, J.: Deterministic restrictions in circuit complexity. In: *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC 1996, pp. 30–36. Association for Computing Machinery, New York (1996)
3. Chlebus, B.S., Gasieniec, L., Gibbons, A., Pelc, A., Rytter, W.: Deterministic broadcasting in ad hoc radio networks. *Distrib. Comput.* **15**(1), 27–38 (2002)
4. Chrobak, M., Gasieniec, L., Rytter, W.: Fast broadcasting and gossiping in radio networks. *J. Algorithms* **43**(2), 177–189 (2002)

5. Clementi, A.E.F., Monti, A., Silvestri, R.: Distributed broadcast in radio networks of unknown topology. *Theor. Comput. Sci.* **302**(1–3), 337–364 (2003)
6. Czumaj, A., Davies, P.: Faster deterministic communication in radio networks. In: Proceedings of 43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016), pp. 139:1–139:14 (2016)
7. De Bonis, A., Gąsieniec, L., Vaccaro, U.: Generalized framework for selectors with applications in optimal group testing. In: Baeten, J.C.M., Lenstra, J.K., Parrow, J., Woeginger, G.J. (eds.) ICALP 2003. LNCS, vol. 2719, pp. 81–96. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-45061-0_8
8. D'yachkov, A., Rykov, V.: Bounds on the length of disjunctive codes. *Probl. Inf. Transm.* **18**, 7 (1982)
9. Erdős, P., Frankl, P., Füredi, Z.: Families of finite sets in which no set is covered by the union of r others. *Israel J. Math.* **51**(1–2), 79–89 (1985)
10. Gasieniec, L., Radzik, T., Xin, Q.: Faster deterministic gossiping in directed ad hoc radio networks. In: Proceedings of Scandinavian Workshop on Algorithm Theory (SWAT 2004), pp. 397–407 (2004)
11. Hwang, F.K., Sos, V.T.: Non-adaptive hypergeometric group testing. *Studia Scientiarum Mathematicarum Hungarica* **22**, 257–263 (1987)
12. Kautz, W.H., Singleton, R.C.: Nonrandom binary superimposed codes. *IEEE Trans. Inf. Theory* **10**(4), 363–377 (1964)
13. Komlos, J., Greenberg, A.: An asymptotically fast nonadaptive algorithm for conflict resolution in multiple-access channels. *IEEE Trans. Inf. Theory* **31**(2), 302–306 (1985)
14. Liu, D., Prabhakaran, M.: On randomized broadcasting and gossiping in radio networks. In: Proceedings of 8th Annual International Conference on Computing and Combinatorics (COCOON 2002), pp. 340–349 (2002)
15. De Marco, G.: Distributed broadcast in unknown radio networks. *SIAM J. Comput.* **39**(6), 2162–2175 (2010)
16. Porat, E., Rothschild, A.: Explicit nonadaptive combinatorial group testing schemes. *IEEE Trans. Inf. Theory* **57**(12), 7982–7989 (2011)
17. Romik, D.: The Surprising Mathematics of Longest Increasing Subsequences. Cambridge University Press (2014)