

Priority Transmission

Noga Alon* Guy Rutenberg†

October 8, 2016

Abstract

We consider a broadcast model involving multiple transmitters and receivers. Transmission is performed in rounds, where in each round any transmitter is allowed to broadcast a single message, and each receiver can receive only a single broadcast message, determined by a priority permutation over the transmitters. The message received by receiver R in a given transmission round is the one sent by the first transmitter among all those broadcasting in that round according to the permutation of R . In our model, each pair of transmitter and receiver has a unique message which the transmitter has to send to the receiver. We prove upper and lower bounds on the minimal number of rounds needed for transmitting all the messages to their respective receivers. We also consider the case where the priority permutations are determined geometrically.

1. Introduction

Consider the following broadcast model: There are k transmitters, and n receivers. Each transmitter has a (unique) message for each receiver which has to be sent. Transmission is performed in rounds of broadcasting, in each round some transmitters send messages where each transmitter can only send a single message in each broadcasting round (it can also send none). When multiple transmitters broadcast at the same time, collisions between transmissions may occur at the receiver. Such collisions are resolved by having a priority permutation associated with each receiver, determining the message that will get through according to the transmitter's priority. One possible example for such priority permutations is the one where transmitters are sorted by their respective distance to each receiver, and hence the receiver will always receive the message which was broadcasted by the closest transmitter.

*Sackler School of Mathematics and Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel and School of Mathematics, Institute for Advanced Study, Princeton, NJ 08540. Email: nogaa@tau.ac.il. Research supported in part by a USA-Israeli BSF grant, by an ISF grant, by the Israeli I-Core program and by the Oswald Veblen Fund.

†Sackler School of Mathematics, Tel Aviv University, Tel Aviv 69978, Israel. Email: guyrutenberg@mail.tau.ac.il

In this paper we are interested in finding the minimum possible number of transmission rounds $q = q(k, n)$ needed to transmit messages in a network with k transmitters and n receivers. For the case $n = 2$ of two receivers we show that $q(k, 2) = k + 1$ and determine for each two given permutations the minimum number of required rounds. For the case of networks with $n = 3$ and $n = 4$ receivers we establish lower and upper bounds on the minimum number of required rounds. We provide a general lower bound of $q(k, n) = n \log k$ and show that it is tight for $n = k!$. Finally, we provide general upper bounds realizable by permutations corresponding to distances in the plane and in higher dimensions. The higher dimension construction shows that $q(n, n) \leq n \cdot O(e^{\sqrt{\log n \log \log n}})$ for the case $n = k$ of equal number of receivers and transmitters.

In the appendix we consider those sets of permutations which can be realized as distances in d -dimensional space, and show that their number is negligible compared to the total number of sets of permutations of the same size.

2. Related work

The problem studied in this paper is closely related to the investigation of radio networks, that is, the study of propagation of messages in a network of nodes that can transmit or receive messages. There are several different models, each having different constraints on how messages can be sent and received between nodes and the type and number of messages transmitted.

The study of such networks dates back to the 1980s [13] and since then underwent extensive research. The first network model, introduced by [13], considered networks defined as directed graphs, where the vertices are nodes which both transmit and receive messages, and edges indicate to which nodes the message can propagate when it is broadcasted. The basic model is of a multi-hop network, where a single message is emitted from a distinguished node called the source node and then propagates to the entire network. In their paper, they consider collisions, which happen when multiple neighbors of a node transmit simultaneously, and optimality measures for protocols considering the maximum time and average time needed for the broadcasted message to propagate from the source node to all others.

The collision model used in [13] deals with two scenarios: If collision occurs because two neighbors of a vertex transmit simultaneously, the receiving vertex detects the collision but not its origin. However, if a collision occurs because a vertex is transmitting while a neighbor of it tries to transmit as well, none of the parties involved detects the collision. This collision model can be altered in order to further consider additional models: Collisions are not distinguishable from non-transmissions ([16, 5]) while in another model in the case of multiple transmissions, one of the signals gets through. This is the model considered in the present paper. Some papers, such as [4], use a model where collisions can be detected as noise but in practice this additional information is discarded and treated as silence.

The original radio-broadcast model considered a single message that needs to be broadcasted to the entire graph. [19] generalizes this model to include set to set broadcasting,

and introduces the concept of sets of receivers and transmitters where each transmitter has a unique message to be delivered to each receiver. [10, 11, 6] consider a different generalization where multiple messages have to be transmitted to different receivers, but some receivers have prior knowledge of messages intended to others.

A variant of the multiple-message model studied in [2, 5], considered *network-coding*: Instead of treating each message separately when routing them, the model allows taking multiple messages, combining them by means of a coding that merges information from multiple messages into a single packet. This model of broadcasting, the so-called network-coding model, sometimes performs better than the traditional routing model.

For a more extensive review of broadcasting in radio networks, see [18].

3. The model

The model of broadcasting used in this paper is as follows: We have two kinds of nodes: transmitters and receivers. We denote the number of transmitters by k and the number of receivers by n . Each pair of transmitter and receiver has a unique message that needs to be transmitted between them. The transmissions are performed in rounds of broadcasting, where the transmissions in each round are performed simultaneously. Associated with each receiver is a permutation of the transmitters, denoting their priority. This permutation is used for solving broadcast collisions in the sense that when two or more transmitters transmit in the same round, a receiver gets the message sent by the transmitter appearing first in its associated permutation. This happens regardless of whether or not that transmitter transmitted a message addressed for that particular receiver. More formally, suppose $1, \dots, k$ are the transmitters and σ is a permutation that is associated with one of the receivers. Consider a given round in which i_1, \dots, i_ℓ are the transmitters that sent messages. Then, the only message that gets through to that receiver is the one transmitted by the transmitter with minimal index in σ , that is, the message sent by i_j where j satisfies $\sigma(i_j) = \min_{m=1}^{\ell} \sigma(i_m)$.

We are interested in studying the limits on how many transmission rounds are necessary in a given network (defined by the number of transmitters, receivers and the corresponding permutations) in order for each transmitter to deliver its messages to the corresponding receivers.

4. Networks with two receivers

The case $n = 2$ is fully understood, as shown in the following theorem.

Theorem 4.1. *For any k , $q(k, 2) = k + 1$. In addition, for a given pair of permutations σ and τ of the set $[k] = \{1, 2, \dots, k\}$ of senders, the minimum number q of rounds that suffices to transmit all $2k$ messages is $k + t$, where t is the minimum number of parts in a partition of $[k]$ into disjoint parts so that each part is increasing in σ and decreasing in τ . In particular, this is easy to compute, given σ and τ , and for a pair of random permutations the expected number of rounds needed is $k + (2 + o(1))\sqrt{k}$ (the exact distribution is known as well).*

Proof. We start with the lower bound. At most two senders can transmit simultaneously in each round. Furthermore, the last two senders in each permutation can only transmit alone, as any other sender transmitting will block their transmission. Hence we need at least $k + 1$ transmission rounds.

The following 1-dimensional construction proves that the bound is tight: Place the senders S_1, \dots, S_k , in that order, between the two receivers, R_1 and R_2 . Let the permutations over the transmitters associated with each receiver be determined by the distance between that receiver and each transmitter. In the first round S_1 transmits its message to R_2 . In the i th round for $2 \leq i \leq k$, S_{i-1} transmits its message to R_1 while S_i transmits its message to R_2 . In the $(k + 1)$ th and final round, S_k transmits to R_1 . It is clear that in the transmission schedule above each transmitter succeeds in sending its messages to its two receivers.

We next consider the minimum number of rounds in an arbitrary network with two receivers. Given a transmission schedule for two permutations σ and τ , assume, without loss of generality, that $\sigma = (1, 2, 3, \dots, k)$. Construct the following bipartite graph $(A \cup B, E)$ with $A = B = \{1, 2, \dots, k\}$ where $(i, j) \in E$ for $i \in A, j \in B$, iff i transmits to the first receiver in the same round j transmits to the second receiver, or $i = j$. We may further assume that there are no redundant transmissions, that is, no sender transmits its message twice to the same receiver. It follows directly that each vertex is connected to at most one other vertex besides the copy of itself, hence, for every $v \in A \cup B$, $\deg(v) \leq 2$. Thus every connected component in the graph is either a path or a cycle. Next we show that it cannot be a cycle. Assume to the contrary that $(i_1, i_1), (i_1, i_2), (i_2, i_2), (i_2, i_3), \dots, (i_{k-1}, i_k), (i_k, i_k), (i_k, i_1)$ form a cycle in the bipartite graph. Since (i_1, i_2) is an edge, i_1 transmits to R_1 in the same round i_2 transmits to R_2 . Hence, $i_1 \leq_\sigma i_2$ and $i_2 \leq_\tau i_1$. This reasoning of course can be applied to every other edge, resulting in $i_1 \leq_\sigma i_2 \leq_\sigma \dots \leq_\sigma i_k \leq_\sigma i_1$ which is a contradiction. Therefore, every connected component is a path of the form $(i_1, i_1), (i_1, i_2), (i_2, i_2), (i_2, i_3), \dots, (i_{k-1}, i_k), (i_k, i_k)$, where $i_1 \leq_\sigma i_2 \leq_\sigma \dots \leq_\sigma i_k$ and $i_k \leq_\tau i_{k-1} \leq_\tau \dots \leq_\tau i_1$.

If a path corresponds to ℓ transmission rounds of pairs, it has $2\ell + 2$ vertices, and vice-versa. Hence, the total number of pairs transmitting simultaneously is $k - t$ where t is the number of connected components in the graph. We need two additional rounds corresponding to the transmissions of the endpoints of each path. Hence the total number of rounds is $k + t$. Furthermore, if $[k]$ can be partitioned into t parts each increasing in σ and decreasing in τ , then it is easy to see that an appropriate transmission schedule can be constructed such that the connected components in the transmission graph will match the parts of the partition.

Dilworth's Theorem ([14]) states that given a partially ordered set, the size of the maximum antichain, that is, a set where every two distinct elements are non-comparable, is equal to the minimum number of chains (sets where every two elements are comparable) that cover the set.

Consider the following partial order over $[k]$

$$x \leq y \iff x \leq_\sigma y \wedge y \leq_\tau x.$$

Every chain in this partial order corresponds to a path in the graph whose vertices




Figure 1: Consider the two permutations above. The vertices can be split into 3 parts each increasing in σ and decreasing in τ : $(1, 2, 3, 6)$, $(4, 5)$, (7) . These correspond to the graph shown above and 10 transmission rounds.

increase in σ and decrease in τ . Hence, by Dilworth's Theorem, the minimum number of such paths required to cover all vertices is equal to the size of the maximum antichain. However, an antichain set is simply a common subsequence of both permutations, σ and τ . For the case of random permutations, the length of the longest common subsequence has the well studied statistics of the length of the longest increasing subsequence of a random permutation, see [8]. In particular its expectation is $(2 + o(1))\sqrt{k}$. \square

We note that finding the largest common subsequence of two permutations can be done easily in polynomial time. Hence, we can efficiently find the optimal transmission schedule for any two given permutations.

5. Bounds on networks with three and four receivers

For any fixed $n \geq 3$ the situation already changes and $q(k, n) - k$ tends to infinity with k . We describe the cases $n = 3$, $n = 4$ separately, as these are simpler.

Theorem 5.1.

$$k + \lfloor k^{1/3} \rfloor \leq q(k, 3) \leq q(k, 4) \leq k + O(\sqrt{k})$$

Moreover, permutations achieving the upper bound can be realized by distances between senders and receivers in the plane.

Proof. Applying the Erdős-Szekeres Theorem ([15]) twice (see, e.g. [9]), we conclude that any three permutations contain two that have a common subsequence of length at least $\lfloor k^{1/3} \rfloor$. Because a schedule for three permutations induces a schedule for any two of the

permutations (not necessarily an optimal one), by Theorem 4.1 and its proof the schedule must contain at least $k + \lfloor k^{1/3} \rfloor$ rounds.

As any valid schedule for a set of permutations induces a valid schedule for any subset of the permutations we clearly have $q(k, 3) \leq q(k, 4)$. Thus, it is sufficient to provide a construction for $n = 4$ receivers that proves the upper bound. Suppose that $k = s^2$ and consider the grid in the plane

$$\{0, 1, 2, \dots, s\} \times \{0, 1, 2, \dots, s\}.$$

We place the senders on each point of the grid with non-zero coordinates, namely, in each point (i, j) with $i, j \in 1, \dots, s$ we have a sender. We place four receivers, called NE (for North-East), SE (South-East), SW (South-West) and NW (North-West). The permutation for NE is according to the projection of the location of the senders on the line $x = y$, with ties broken arbitrarily, and highest priority point being the North-East sender located at (s, s) . The permutation of SW is opposite, and those of SE and NW are according to the projections on the line $y = -x$. Now we can complete the transmissions in $(s + 1)^2$ rounds corresponding to the squares of the grid: For each square with vertices $(i, j), (i + 1, j), (i, j + 1), (i + 1, j + 1)$ ($i, j \in 0, \dots, s$) we add a transmission round where each sender sends its message to the receiver to whom it is closest (compared to the other transmitters in the square). Namely, the sender at (i, j) transmits its message to SW, that in $(i + 1, j)$ sends its message to SE, and so on. When i or j are equal to s or 0 some of the points in the squares above may not represent senders, as those are located only on $\{1, \dots, s\} \times \{1, \dots, s\}$. In that case we simply ignore these vertices and no corresponding transmission is made.

Clearly, each transmitter has a round where it transmits to each receiver and is closest to it among the transmitting senders at that transmission round. Hence, the construction yields a valid schedule with $(s + 1)^2 = k + O(\sqrt{k})$ transmission rounds.

We further note that by placing the receivers far enough away in the direction of North-East, South-West, and so on, the permutations above are simply the ones obtained by distances in the plane. \square

Remark: It is known, see [9, 12], that there are collections of $\Omega(k^{1/3})$ permutations of $[k]$, so that no two have a common subsequence of size exceeding $\Theta(k^{1/3})$. Therefore, the proof of the lower bound above does not yield a larger bound for $q(k, n)$ for any fixed n and sufficiently large k than for $n = 3$.

6. The general case: lower bound

Note, first, that the transmitters with the lowest priority for each receiver must transmit independently from any other transmitter, otherwise they will be blocked. Hence, every schedule has to contain n transmission rounds for those transmitters. Additionally, because each receiver can only receive a single transmission in each round, every round is limited to n participating transmitters. This means that the remaining $n(k - 1)$ messages will be transmitted using at least $k - 1$ transmission rounds. Therefore, we get the

following trivial lower bound:

$$q(k, n) \geq n + k - 1.$$

This trivial bound can be generalized by the following observation: Consider the set of n messages to each receiver with the $(k - 1)$ -th priority (with respect to the appropriate receiver). Each such message can only be sent with another single message (one that corresponds to a transmitter with the k th priority for the respective receiver). Furthermore, a round that includes a message which originates from a transmitter with priority i and reaches its destination, can only include $k - i$ additional messages—messages of transmitters which have lower priority for the given receiver. Hence, assuming that there are no redundant transmissions which do not reach their destinations, the round may contain up to $k - i + 1$ messages, as any additional message will necessarily be redundant.

We can use this fact to give a better lower bound on the number of transmissions rounds needed. Associate with each message a weight: The weight of a message from a transmitter with priority i to the corresponding receiver will be $1/(k - i + 1)$. Thus the sum of weights of all messages participating in a single round is at most 1. Indeed, let i be the maximum index in a priority permutation of a message in the round. Then there are at most $k - i + 1$ messages in that round, and each has a weight at most $1/(k - i + 1)$, hence the total weight of all messages in that round is 1. Summing the weights of all the messages gives us a lower bound on the number of rounds required to complete the transmission of all the messages. This proves the following theorem:

Theorem 6.1. *For every n and k ,*

$$q(k, n) \geq n \sum_{i=1}^k \frac{1}{i} \geq n \log_e k.$$

We note that for the case $k > n$ this can be slightly improved as each round cannot contain more than n messages. Therefore, if we modify the definition of the weight associated with each message to $\max\{1/(k - i + 1), 1/n\}$, we get that $q(k, n) \geq n(\log_e n - 1) + k$.

7. The general case: upper bound

The case $k = n$ appears difficult, but for k far smaller than n (specifically, for $n \geq k!$) we can show that $q(k, n) = \Theta(n \log k)$.

Theorem 7.1. *For $n = k!$, $q(k, n) = n \sum_{i=1}^k \frac{1}{i}$.*

The upper bound of $O(n \log k)$ for $n > k!$ follows by splitting the set of receivers into groups of size at most $k!$, and by handling the messages to each such group separately, using the schedule in the proof of Theorem 7.1.

Proof. Let the priority permutations of the n receivers be all the $n = k!$ possible permutations of the k transmitters. For each subset I of the transmitters, $|I| = r$, let J be the

subset of receivers whose priority permutations have the transmitters in I in the last r places. Thus $|J| = r!(k-r)!$. Note that each transmitter in I has the highest priority among the transmitters in I in $m_r = (r-1)!(k-r)!$ permutations in J .

We define the following transmission schedule: For each subset I , it is possible to split the receivers in J into m_r sets, each of size r , such that those transmitters with the highest priority among the transmitters in I are distinct. To see this, for each of the $(k-r)!$ possible permutations τ of the receivers not in I , and each cyclic permutation ν of the receivers in I , take the r cyclic shifts of ν concatenated with τ as a set of r permutations. This gives the required m_r sets. For each such r permutations, the corresponding r transmitters with the highest priority will transmit in a single round to their respective receivers (according to the permutation they have the highest priority among the transmitters in I). This ensures that each transmitter participating in the round is indeed able to deliver its message to the corresponding receiver.

Given a transmitter i and receiver j , consider the set I that includes i and every other transmitter whose priority (with regards to j 's permutation) is lower than i 's. Now $j \in J$, where J is the subset of receivers whose priority permutations have the transmitters in I in the last $r = |I|$ places. Thus the transmission schedule defined above ensures that i successfully transmits its message to j in one of the rounds associated with I and J . Therefore, every transmitter is able to send its message to every receiver in one of the rounds of the transmission schedule. The number of rounds is

$$q = \sum_{r=1}^k \binom{k}{r} m_r = \sum_{r=1}^k \binom{k}{r} (r-1)!(k-r)! = k! \sum_{r=1}^k \frac{1}{r}$$

which completes the proof. \square

For any $n > k!$, one can split the receivers to sets of size $k!$. Repeating the construction from the proof for each set shows that $q(k, n) = \Theta(n \log k)$. For $n < k!$ one can split the senders into groups of size $\log n / \log \log n$ and handle each such group separately according to the last theorem. This implies that for all $n < k!$

$$q(k, n) \leq O\left(kn \frac{(\log \log n)^2}{\log n}\right).$$

In particular, for $k = n$ the arguments here imply:

$$(1 + o(1))n \log n \leq q(n, n) \leq (1 + o(1)) \frac{n^2 (\log \log n)^2}{\log n}.$$

It is possible, however, to get a better upper bound using the basic construction in [7]. In that paper it is shown that there is a positive constant $c > 0$ and a bipartite graph with classes of vertices $A = B = [n]$ and with $n^2 - n^{2-c}$ edges, whose edges can be decomposed into at most n^{2-c} induced matchings. View A as the set of senders and B as the set of receivers. For each $b \in B$ define a permutation on A by placing all $a \in A$ connected to B before all other members of A (with the permutation being arbitrary otherwise). Now the messages corresponding to each of the induced matchings can be transmitted in

one round, for each of the matchings. The remaining messages can be sent in separate rounds. This completes the schedule in at most $2n^{2-c}$ rounds, which gives

$$q(n, n) \leq O(n^{2-c})$$

for some constant $c > 0$.

In the next section we prove a much stronger upper bound showing that in fact $q(n, n) = n^{1+o(1)}$.

7.1. A general upper bound

The construction used for getting the upper bound on $n = 3$ and $n = 4$ receivers can be generalized to an arbitrary number of receivers n , provided that the number of transmitters is sufficiently large. This is done by taking a digital convex n -gon, that is a convex n -gon whose vertices lie on lattice points, and placing the receivers on the lines perpendicular to arbitrarily chosen supporting lines of each vertex of the n -gon, while the senders will be located far enough on a square of size $\sqrt{k} \times \sqrt{k}$ centered around the origin.

We cover the senders using overlapping translates of the convex n -gon, such that each sender is covered by every vertex of the n -gon (using different translates of it). Each transmission round will correspond to one n -gon participating in the cover. Each sender covered by the n -gon will transmit its message to the corresponding receiver according to the vertex that covers it. It can be seen that such sender is the closest to the corresponding receiver among all other senders covered by that n -gon. Thus, this is a valid and complete transmission schedule.

It has been shown in [17] and [1] that such convex n -gons can be constructed so that they can be inscribed inside a square grid of side-size $\Theta(n^{3/2})$. Therefore, by using such convex n -gon, our senders can be covered, according to the construction above, using $(\sqrt{k} + O(n^{3/2}))^2 = k + O(n^{3/2}\sqrt{k} + n^3)$ n -gons, hence the schedule requires $k + O(n^{3/2}\sqrt{k} + n^3)$ transmission rounds.

We have thus proved the following theorem:

Theorem 7.2. *For any k and n , $q(k, n) \leq k + O(n^{3/2}\sqrt{k} + n^3)$ with permutations corresponding to distances in the plane.*

One such family of convex polygons with diameter $O(n^{3/2})$ is the one constructed as follows: For a given t consider the Farey sequence of order t , that is the set of irreducible fractions between 0 and 1, where the denominator is at most t , for example $F_3 = \{\frac{0}{1}, \frac{1}{3}, \frac{1}{2}, \frac{2}{3}, \frac{1}{1}\}$. Now, consider the bijection $p/q \rightarrow p/(q-p)$ applied to the Farey sequence (omitting the last term), e.g., $\bar{F}_3 = \{\frac{0}{1}, \frac{1}{2}, \frac{1}{1}, \frac{2}{1}\}$. One can verify that the bijection preserves the increasing monotonicity of the Farey sequence. Now, starting at $(0, 0)$, treat each fraction as a 2-tuple and add it coordinate by coordinate to the previous point to create the next vertex of the polygon. The last example produces the following vertex list: $(0, 0), (0, 1), (1, 3), (2, 4), (4, 5)$. This effectively constructs one quarter of the polygon. We complete the polygon by rotating and repeating the process. Because




Figure 2: Convex n -gon construction for $t = 4$. Here $n = 24$ and the n -gon is bounded by a 17×17 square.

the elements in the list were monotonely increasing, so is the slope of the corresponding edges. Therefore, the constructed polygon is indeed convex. As the sum of the numerator and denominator of each term in the modified sequence is bounded by t , the total change to both coordinates in each quarter is bounded by t times the number of terms in the modified Farey sequence of order t . The number of terms of the modified Farey sequence is one less than in the original sequence, that is $\sum_{i=1}^t \varphi(i) = \Theta(t^2)$, where $\varphi(i)$ is Euler's totient function, namely the number of positive integers smaller than i that are relatively prime to i . Therefore, the total change in both coordinates is $O(t^3)$, making the diameter of the polygon $O(t^3)$ as well. The number of vertices in the polygon is $O(t^2)$ (the same order as the number of terms in the sequence). Therefore a digital convex n -gon constructed in this manner will have a diameter of $O(n^{3/2})$. Such a construction is illustrated in Figure 2.

Theorem 7.2 can be used to derive a non-trivial upper bound for the case of equal number of receivers and senders ($n = k$), with permutations corresponding to distances in the plane. We start by partitioning the set of receivers into $n^{2/3}$ sets of $n^{1/3}$ receivers each. For each such subset of receivers in the partition and the set of all transmitters,

apply the construction of Theorem 7.2. This results in

$$n^{2/3} \cdot (k + O([n^{1/3}]^{3/2} \sqrt{k} + (n^{1/3})^3)) = O(n^{5/3})$$

rounds. This proves the following corollary:

Corollary 7.3. *For any n , $q(n, n) = O(n^{5/3})$ with permutations corresponding to distances in the plane.*

Finally, we consider a higher dimensional analogue of the construction, by using a d -dimensional sphere instead of the convex n -gon.

Theorem 7.4.

$$(1 + o(1))n \log n \leq q(n, n) \leq n \cdot O\left(e^{\sqrt{\log n \log \log n}}\right) \quad (= n^{1+o(1)}).$$

Proof. Consider the d -dimensional ball. It is well known that the volume of the unit ball in \mathbb{R}^d (for even d) is:

$$V_d = \frac{\pi^{d/2}}{(d/2)!}$$

and hence the number of lattice points in a ball of radius r is $(1 + o(1))V_d r^d$. Now if we consider the ball of radius m and all the spheres of integral square radius at most m^2 we get that by the pigeonhole principle there exists such a sphere (henceforth the "small" sphere) which has at least $(1 + o(1))V_d m^{d-2}$ lattice points on its boundary. Let the permutations of the receivers correspond to projections to lines perpendicular to the supporting hyperplanes of these lattice points. Thus we have $n = (1 + o(1))V_d m^{d-2}$.

Let the transmitters be the lattice points inside a ball of radius dm . Thus $k = (1 + o(1))V_d (dm)^d$. The transmission rounds will correspond to shifts of the "small" sphere centered in the lattice points of a ball of radius $(d+2)m$. Thus the number of transmission rounds is

$$(1 + o(1))V_d ((d+2)m)^d = (1 + o(1)) \left(1 + \frac{2}{d}\right)^d k \leq (1 + o(1))e^2 k$$

If we take $m^2 = d^d$ we get that $k = nd^{2d}$. In this case we have $d^{2d} = n^{o(1)}$ as $n > \frac{1}{(d/2)^{d/2}} d^{d(d-2)/2}$ which implies $\log n \geq (1 + o(1))d^2/2 \cdot \log d$ and hence $k \leq n \cdot O(e^{\sqrt{\log n \log \log n}})$ which gives

$$q(n, n) \leq q(k, n) \leq (1 + o(1))e^2 k \leq n \cdot O(e^{\sqrt{\log n \log \log n}}) = n^{1+o(1)}.$$

□

8. Concluding remarks

The problem of determining $q(k, n)$ or estimating it more accurately is still unresolved except for the extreme cases of $n = 2$ and $k \geq n!$. In networks with $n = 3$ and $n = 4$, it is still unclear where between $k + O(k^{1/3})$ and $k + O(\sqrt{k})$ are the real values of $q(k, n)$. For $n = k$, the gap between the lower bound $n \log n$ and the upper bound $n \cdot O(e^{\sqrt{\log n \log \log n}})$, is still substantial and it will be interesting to improve it.

One might consider the problem limited to the case where the permutations are realized as distances in the plane. In this case, the $k + O(n^{3/2}\sqrt{k} + n^3)$ bound gives good results for fixed number of receivers (fixed n). However, if we consider the special case of equal number of receivers and transmitters ($n = k$), then the bound above is trivial, but as shown in Corollary 7.3 can be improved to $O(n^{5/3})$. It will be interesting to estimate more accurately the minimum possible number of rounds that can be realized by distance permutations in the plane when $n = k$.

A natural problem that arises from constructions that appeared in the proofs of the bounds, is determining which sets of permutations are realizable as projections along lines in a d -dimensional space. In the appendix we further investigate this problem and prove upper bounds on the number of such sets of permutations whenever the dimension is smaller than the number of permutations, showing that in this case most sets of permutations are not realizable.

A. Distance permutations

The construction of the upper bound on networks with $n = 3$ and $n = 4$ receivers relies on a set of permutations that can be realized as distances in the plane. The general question of determining necessary and sufficient conditions such that a set of r permutations over n points can be realized as distances from r points in a d -dimensional space is interesting in its own right.

We focus on those permutations that can be realized as projections along arbitrary directions. That is, each permutation is determined by the projection of the points on a line through the origin in a given direction. Note that such a construction is equivalent to distances from far away placed points.

Clearly, whenever $r \leq d$, every set of permutations is realizable, as we can consider the permutations as projections along the axes, and place the points independently in each axis according to the corresponding permutation.

Consider the case of $r = 3$ and $d = 2$. Let $\vec{v}_1, \vec{v}_2, \vec{v}_3$ be the vectors representing the directions on which each permutation is projected, and let $\pi_i(p)$ be the projection of the point p on v_i . The vectors $\vec{v}_1, \vec{v}_2, \vec{v}_3$ are linearly dependent, hence there are a_1, a_2, a_3 such that $a_1\vec{v}_1 + a_2\vec{v}_2 + a_3\vec{v}_3 = 0$. Now let p_1, p_2 be two arbitrary points, and suppose, without loss of generality, that $a_1(\pi_1(p_1) - \pi_1(p_2))$ and $a_2(\pi_2(p_1) - \pi_2(p_2))$ both have the same sign. In that case $a_3(\pi_3(p_1) - \pi_3(p_2))$ must have the opposite sign. If one of the a_i 's is 0, say a_1 , the condition simplifies to requiring that $a_2(\pi_2(p_1) - \pi_2(p_2))$ and $a_3(\pi_3(p_1) - \pi_3(p_2))$ have opposing signs. This implies a necessary condition for a set of

three permutations to be realizable in the plane. Clearly, the condition can be generalized to apply for every $r > d$.

For example, the following 3 permutations over $1, \dots, 4$ cannot be realized by projections in the plane:

$$\begin{aligned}\pi_1 &= (1234) \\ \pi_2 &= (4132) \\ \pi_3 &= (3124)\end{aligned}$$

Indeed, assume to the contrary that the permutations are realizable in the plane, and let $\vec{v}_1, \vec{v}_2, \vec{v}_3$ be the directions of the projections used for defining the permutations. Let a_1, a_2, a_3 satisfy $a_1\vec{v}_1 + a_2\vec{v}_2 + a_3\vec{v}_3 = 0$. As none of the permutations is the same or an inverse of another, we have that a_1, a_2, a_3 are all non-zero. Suppose that a_1 and a_2 have the same sign then $a_1(\pi_1(1) - \pi_1(2))$, $a_2(\pi_2(1) - \pi_2(2))$, $a_1(\pi_1(1) - \pi_1(3))$ and $a_2(\pi_2(1) - \pi_2(3))$ all have the same sign. However, because $\pi_3(1) - \pi_3(2)$ and $\pi_3(1) - \pi_3(3)$ have opposing signs, the necessary condition above requires a_3 to have both an opposite sign and the same sign as a_1 and a_2 at the same time, which is a contradiction. Supposing that a_1 and a_2 have opposing signs, we may consider the pairs of points 1, 4 and 2, 3. Using the same arguments, we get that a_3 must have the same sign as a_1 and the opposite sign as well, which is, again, a contradiction. Therefore, one cannot realize the above permutations by projections in the plane.

The above condition, and its clear extension for $r > d$, presents a necessary condition for a set of permutations to be realizable using projections along lines in the d -dimensional case. However, it is not immediately clear from it how many such sets are indeed realizable. In the following subsection we show that the number of such sets is negligible compared to the total number of sets of permutations when the number of permutations is bigger than the dimension of the space used.

A.1. Upper bound on the number of realizable permutations

In order to bound the number of r -tuples of permutations realizable by projections in the plane, consider the following set of polynomials over variables in \mathbb{R}^d :

$$P_{ijk}(x_1, \dots, x_n, \alpha_1, \dots, \alpha_r) = \langle x_i, \alpha_k \rangle - \langle x_j, \alpha_k \rangle$$

where $1 \leq i < j \leq n$, $1 \leq k \leq r$. Given a set $A = \{x_1, \dots, x_n\}$ of n points in \mathbb{R}^d and a set of vectors $\alpha_1, \dots, \alpha_r \in \mathbb{R}^d$, the signs of the polynomials P_{ijk} determine the ordering of the points in A when projected on the directions $\alpha_1, \dots, \alpha_r$. That is, the pattern over $\{-1, 1\}^{d\binom{n}{2}}$ created by taking the sign of each polynomial over a given point in $\mathbb{R}^{d(n+r)}$, the so-called sign-pattern, represents a set of points and directions realizing a set of r permutations over n points using projections in \mathbb{R}^d .

Warren's Theorem ([20], see also [3]) states that given a set of m polynomials P_1, \dots, P_m over \mathbb{R}^ℓ , the number of different sign-patterns in $\{\pm 1\}^m$ created by the polynomials is bounded by $(4ekm/\ell)^\ell$ where k is the maximum degree of the polynomials and assuming

$m \geq \ell$. We can use Warren's Theorem to bound the number of different sign-patterns created by $\{P_{ijk}\}$, and thus bound the number of permutations realizable as projections.

In this case, the number of polynomials is $m = r \binom{n}{2}$ and each polynomial is of degree $k = 2$ and over $\ell = d(n+r)$ variables. Thus by Warren's Theorem (assuming $d(n+r) \leq r \binom{n}{2}$), the number of sign patterns \bar{s} satisfies

$$\bar{s} \leq \left(\frac{4ern^2}{d(n+r)} \right)^{d(n+r)} \leq (4ern/d)^{dn+dr}.$$

Thus the number of ways to pick r permutations over n points which are realizable by projections in \mathbb{R}^d is bounded by $(4ern/d)^{dn+dr}$. Compare that number with the number of ways to pick r arbitrary permutations over n points, which is $(n!)^r \geq (n/e)^{nr}$. When $d < r$ are fixed, and n tends to infinity, we get that the number of r -tuples of permutations realizable as projections in the d -space is negligible when compared to the total number of r -tuples of permutations. On the other hand, the bound becomes trivial if $d = r$, as expected, as in that case, every r permutations are realizable, for example by projecting each permutation independently on a different axis.

The above discussion considered sets of permutations realizable by projections in the d -dimensional space. A generalization of this model would be to consider sets of permutations realizable as distances from r points in the d -dimensional space (projections are distances from far away placed points). This generalization behaves similarly to the case of projections, as we can still use the signs of a set of $r \binom{n}{2}$ polynomials of degree 2 to describe the permutations defined using the distances. Indeed, here

$$P_{ijk}(x_1, \dots, x_n, \alpha_1, \dots, \alpha_r) = d^2(x_i, \alpha_k) - d^2(x_j, \alpha_k)$$

where $d^2(x_i, \alpha_j)$ is the square of the distance between x_i and α_j . Thus, the analysis is the same as before and results in proving that the number of r -tuples of permutations realizable as distances in the d -dimensional space is negligible compared to the total number of r -tuples of permutations whenever $d < r$ are fixed and n tends to infinity.

References

- [1] D. M. Acketa and J. D. Žunić. On the maximal number of edges of convex digital polygons included into an $m \times m$ -grid. *J. Combinatorial Theory, Series A*, 69(2):358–368, 1995.
- [2] R. Ahlswede, N. Cai, S. R. Li, and R. W. Yeung. Network information flow. *IEEE Trans. Inform. Theory*, 46(4):1204–1216, 2000.
- [3] N. Alon. Tools from higher algebra. In R. L. Graham, M. Grötschel, and L. Lovász, editors, *Handbook of combinatorics*, volume 2, pages 1749–1783. Elsevier, 1995.
- [4] N. Alon, A. Bar-Noy, N. Linial, and D. Peleg. A lower bound for radio broadcast. *J. Comput. System Sci.*, 43(2):290–298, 1991.

- [5] N. Alon, M. Ghaffari, B. Haeupler, and M. Khabbazian. Broadcast throughput in radio networks: Routing vs. network coding. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '14, pages 1831–1843, Philadelphia, PA, USA, 2014. Society for Industrial and Applied Mathematics.
- [6] N. Alon, E. Lubetzky, U. Stav, A. Weinstein, and A. Hassidim. Broadcasting with side information. In *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '08, pages 823–832, Washington, DC, USA, 2008. IEEE Computer Society.
- [7] N. Alon, A. Moitra, and B. Sudakov. Nearly complete graphs decomposable into large induced matchings and their applications. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing*, STOC '12, pages 1079–1090, New York, NY, USA, 2012. ACM. Also: *J. Eur. Math. Soc. (JEMS)* 15(4):1575–1596, 2013.
- [8] J. Baik, P. Deift, and K. Johansson. On the distribution of the length of the longest increasing subsequence of random permutations. *J. Amer. Math. Soc.*, 12(4):1119–1178, 1999.
- [9] P. Beame and D.-T. Huynh-Ngoc. On the value of multiple read/write streams for approximating frequency moments. In *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '08, pages 499–508, Washington, DC, USA, 2008. IEEE, IEEE Computer Society.
- [10] Y. Birk and T. Kol. Informed-source coding-on-demand (iscod) over broadcast channels. In *Proceedings of the Seventeenth Annual Joint Conference of the IEEE Computer and Communications Societies*, volume 3 of *INFOCOM '98*, pages 1257–1264. IEEE, 1998.
- [11] Y. Birk and T. Kol. Coding on demand by an informed source (iscod) for efficient broadcast of different supplemental data to caching clients. *IEEE Trans. Inform. Theory*, 52(6):2825–2830, June 2006.
- [12] B. Bukh and L. Zhou. Twins in words and long common subsequences in permutations. *Israel J. Math.*, 213(1):183–209, 2016.
- [13] I. Chlamtac and S. Kutten. On broadcasting in radio networks—problem analysis and protocol design. *IEEE Transactions on Communications*, 33(12):1240–1246, 1985.
- [14] R. P. Dilworth. A decomposition theorem for partially ordered sets. *Annals of Mathematics*, 51(1):161–166, 1950.
- [15] P. Erdős and G. Szekeres. A combinatorial problem in geometry. *Compos. Math.*, 2:463–470, 1935.
- [16] L. Gąsieniec, D. Peleg, and Q. Xin. Faster communication in known topology radio networks. *Distrib. Comput.*, 19(4):289–300, 2007.

- [17] S. Matić-Kekić, D. M. Aćketa, and J. D. Žunić. An exact construction of digital convex polygons with minimal diameter. *Discrete Math.*, 150(1):303–313, 1996.
- [18] D. Peleg. Time-efficient broadcasting in radio networks: a review. In *International Conference on Distributed Computing and Internet Technology*, pages 1–18. Springer, 2007.
- [19] D. Richards and A. L. Liestman. Generalizations of broadcasting and gossiping. *Networks*, 18(2):125–138, 1988.
- [20] H. E. Warren. Lower bounds for approximation by nonlinear manifolds. *Trans. Amer. Math. Soc.*, 133(1):167–178, 1968.