Short reachability networks

Carla Groenland* Tom Johnston†

Jamie Radcliffe‡ Alex Scott†§

August 13, 2022

Abstract

We investigate a generalisation of permutation networks. We say a sequence $T = (T_1, \ldots, T_\ell)$ of transpositions in S_n forms a t-reachability network if for every choice of t distinct points $x_1, \ldots, x_t \in \{1, \ldots, n\}$, there is a subsequence of T whose composition maps j to x_j for every $1 \leq j \leq t$. When $t = n$, then any permutation in S_n can be created, and T is a permutation network. Waksman [*JACM*, 1968] showed that the shortest permutation networks have length about $n \log_2 n$. In this paper, we investigate the shortest t-reachability networks. Our main result settles the case of $t = 2$: the shortest 2-reachability network has length $\lceil 3n/2 \rceil - 2$. For fixed t, we give a simple randomised construction which shows there exist t-reachability networks using $(2 + o_t(n))n$ transpositions. We also study the case where all transpositions are of the form $(1, \cdot)$, separating 2-reachability from the related probabilistic variant of 2-uniformity. Many interesting questions are left open.

*Utrecht University, Utrecht, The Netherlands, c.e.groenland@uu.nl. This project has received funding from the European Union's Horizon 2020 research and innovation programme under the ERC grant CRACKNP (number 853234) and the Marie Skłodowska-Curie grant GRAPHCOSY (number 101063180). Views and opinions expressed are however those of the author(s) only.

†School of Mathematics, University of Bristol, Bristol, BS8 1UG, UK and Heilbronn Institute for Mathematical Research, Bristol, UK, tom.johnston@bristol.ac.uk.

‡University of Nebraska-Lincoln, USA,jamie.radcliffe@unl.edu.

§Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK, scott@maths.ox.ac.uk. Supported by EPSRC grant EP/V007327/1.

1 Introduction

Let S_n be the symmetric group on n elements, and write $(a, b) \in S_n$ for the transposition that swaps a and b, and $1 \in S_n$ for the identity permutation. Let $T_1, \ldots, T_\ell$ be a sequence of transpositions with $T_i = (a_i, b_i)$ for $i = 1, \ldots, \ell$. The sequence forms a *permutation network* if for every choice of $\sigma \in S_n$, there is some subsequence $T_{i_1}, \ldots, T_{i_{\ell'}}$ such that

$$\sigma = (a_{i_1}, b_{i_1}) \cdots (a_{i_{\ell'}}, b_{i_{\ell'}}).$$

Permutation networks are also called *non-blocking networks* and have been well-studied due to their usefulness in communication networks and distributed computing (see e.g. [4]). Of particular relevance is the following result which shows that permutation networks can be constructed with at most $\sum_{i=1}^{n} \lceil \log_2(i) \rceil$ transpositions.

Theorem 1 (Waksman [6], Beauquier and Darrot [2])**.** *There is a permutation network on n elements using $\sum_{i=1}^{n} \lceil \log_2(i) \rceil$ transpositions.*

Each transposition in a permutation network can only increase the number of possible final states by a factor of two, and so there must be at least $\log_2(n!) = \sum_{i=1}^{n} \log_2(i)$ transpositions.[1]

Permutation networks can also be described in terms of rearrangements of counters. If counters $\{1, \ldots, n\}$ are placed on the vertices of a complete graph on n vertices (with each vertex having exactly one counter), then a permutation network can reach any configuration of counters by following the switches prescribed by a suitable subsequence of $T_1, \ldots, T_\ell$. But what happens with fewer counters? How many transpositions are needed to ensure a specific set of counters can be moved to any configuration? We say a sequence of transpositions $T_1, \ldots, T_\ell \in S_n$ is a *t-reachability network* if, for every choice of t distinct ordered points $x_1, \ldots, x_t \in [n] = \{1, \ldots, n\}$, there is some subsequence $T_{i_1}, \ldots, T_{i_{\ell'}}$ such that the composition $(a_{i_1}, b_{i_1}) \cdots (a_{i_{\ell'}}, b_{i_{\ell'}})$ maps $j \mapsto x_j$ for $1 \leq j \leq t$.

Let us first consider the case of 1-reachability where we want to shuffle one counter across n positions. By noting there are at most 2^ℓ possible subsequences of a sequence of length ℓ, we immediately get a lower bound $\ell \geq \log_2 n$, but this is far from optimal. Indeed, a transposition (a, b) can only move the counter if either a or b is a position where there may already be a counter, and the number of 'reachable positions' can only increase by

[1] It was claimed by Waksman [6] that their construction is best possible, although there does not appear to be a published proof of the corresponding lower bound.

one for each transposition. This implies $n-1$ transpositions are needed, and it is easy to find a tight example, e.g. $(1,2),(1,3),\ldots,(1,n)$.

Our main result is an exact bound for 2-reachability.

Theorem 2. *Let $n \geq 2$. The shortest 2-reachability network on n elements contains $\lceil 3n/2 \rceil - 2$ transpositions.*

We also give a simple randomised construction which shows that one can achieve t-reachability using $(2+o_t(1))n$ transpositions. Surprisingly, the coefficient of the leading term is independent of t.

Theorem 3. *Let $t \geq 3$. There is a t-reachability network on n elements of length at most $(2+o_t(1))n$.*

Reachability questions are related to the problem of generating uniform random permutations. A *lazy transposition* $T = (a,b,p)$ is the random permutation

$$T = \begin{cases} (a,b) & \text{with probability } p, \\ 1 & \text{otherwise.} \end{cases}$$

The composition of an independent sequence of lazy transpositions is also a random permutation. A *transposition shuffle* is an independent sequence of lazy transpositions $T_1,\ldots,T_\ell$ such that $T_1 \cdots T_\ell \sim \text{Uniform}(S_n)$, the uniform distribution over S_n. Let $U(n)$ denote the minimum ℓ for which there exists a transposition shuffle on n elements of length ℓ. Angel and Holroyd [1] asked whether $U(n) = \binom{n}{2}$. This is disproved in [5], which shows that $U(n) \leq \frac{2}{3}\binom{n}{2} + O(n \log n)$ (which is currently the best upper bound). For a lower bound, note that by ignoring the probabilities, any transposition shuffle gives a permutation network, and so $U(n) = \Omega(n \log n)$.[2]

In the case when all transpositions are of the form $(i, i+1)$, Angel and Holroyd [1] showed that $\binom{n}{2}$ is best possible. However, with this restriction there is also a lower bound of $\binom{n}{2}$ from the 'reachability' point of view: this is needed to reach the 'reverse permutation' that maps i to $(n+1)-i$. It is an interesting open problem to prove better lower bounds than $\Omega(n \log n)$ in the general case.

Similar to how we defined t-reachability, we say a sequence of lazy transpositions $T_1,\ldots,T_\ell$ is a *t-uniformity network* if the composition $T_1 \ldots T_\ell$

[2] Conversely, any sorting network can be used to give a sequence which achieves every permutation with non-zero probability, but not necessarily the uniform distribution. It is important here that the sequence of permutations achieves the uniform distribution exactly: Czumaj [3] showed that there are sequences of lazy transpositions of length $O(n \log n)$ which are very close to uniform.

maps the tuple $(1, \ldots, t)$ to $(x_1, \ldots, x_t)$ with equal probability for each tuple $(x_1, \ldots, x_t)$ of t distinct elements from $[n]$. Of course, any t-uniformity network is also a t-reachability network.

In this paper, we also provide a separation between 2-uniformity and 2-reachability when restricted to transpositions of the form $(1, \cdot)$, which we call *star transpositions* (as they match the edges of a star graph $K_{1,n-1}$). This does not affect the order of magnitude in the uniformity set-up, since any lazy transposition (i, j, p) can be simulated by the three star transpositions $(1, i, 1), (1, j, p), (1, i, 1)$ (so, in particular, showing that the minimum number of star transpositions is $\Omega(n^2)$ would give a lower bound of $\Omega(n^2)$ for the general case as well).

A modification of Theorem 2 shows that the minimum number of transpositions needed for 2-reachability is again around $3n/2$ even when restricting to star transpositions.

Theorem 4. *For $n \geq 3$, the shortest 2-reachability network on n elements in which each transposition is a star transposition, contains $\lceil 3(n-1)/2 \rceil$ transpositions.*

Using the restrictive nature of the transpositions, we are able to show that achieving 2-uniformity is strictly harder than achieving 2-reachability.

Theorem 5. *For some $C > 0$, any 2-uniformity network on n elements in which each transposition is a star transposition has length at least $1.6n - C$.*

This is the only setting in which a separation between reachability and uniformity is known. In the setting of star transpositions, we also provide a lower bound on the length of t-reachability networks (see Proposition 7).

The remainder of the paper is organised as follows. We settle the case of 2-reachability (Theorem 2) in Section 2 and prove Theorem 3 in Section 3. In Section 4 we study star transpositions, providing a separation between 2-reachability and 2-uniformity (Theorem 4 and Theorem 5). We finish with some open problems in Section 5.

2 Exact bound for 2-reachability

We prove Theorem 2: the smallest 2-reachability network on n elements has length $\lceil 3n/2 \rceil - 2$.

Proof of Theorem 2. We first give an upper bound construction. Let $n \geq 2$ be given. Our sequence of transpositions starts with the following transpositions, in order.

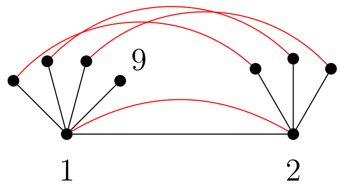

Figure 1: The multigraph corresponding to the minimum 2-reachable sequence given in the proof of Theorem 2. All vertices except 1 and 9 have deficiency 0.

1. $(1, 2)$.
2. $(1, x)$ for all odd $3 \leq x \leq n$.
3. $(2, y)$ for all even $4 \leq y \leq n$.
4. $(x, x + 1)$ for all odd $3 \leq x \leq n - 1$.

If n is odd, we also add the transposition $(1, 2)$ at the end, which will be needed to reach the position $(1, n)$. This defines a sequence of transpositions of length $\lceil 3n/2 \rceil - 2$, and it is straightforward to check that this sequence forms a 2-reachability network.

The main work in this proof is in showing the lower bound. Suppose that $\sigma = (\sigma_1, \ldots, \sigma_\ell)$ is a shortest 2-reachability network on $n \geq 2$ elements. Given two (distinct) counters on positions 1 and 2, a subsequence of σ defines a permutation that moves the counters to new (distinct) positions $x, y \in \{1, \ldots, n\}$. By the definition of 2-reachability, it must be possible to reach any such pair of positions (x, y).

We will consider a process on an auxiliary multigraph G with vertex set $\{1, \ldots, n\}$. An example is drawn in Figure 1 for the sequence of transpositions given earlier for the upper bound (for $n = 9$). The process begins by growing two trees of black edges from the vertices 1 and 2, which we denote T_1 and T_2 respectively. We initialise G as the empty graph and process the transpositions in order, adding an edge for each. We start with a set of active vertices $\{1, 2\}$. Any other vertex becomes active when it first appears in a transposition with an active vertex. Note that no counter can be sitting on an inactive vertex, and (by minimality) no transposition joins two inactive vertices. When we process a transposition (a, b), we add an edge ab to G. We colour the edge black if either a or b was inactive and red otherwise. The black edges always form a forest consisting of at most two trees each containing one of the starting positions.

We now modify our sequence to obtain a sequence with a nicer form (and at most the same length). Let c_1 and c_2 denote the counters that start on 1 and 2 respectively, and let (a, b) where $a \in V(T_1)$ and $b \in V(T_2)$ be the transposition at the first time the two trees meet. Before this time, the counter c_1 is contained in $V(T_1)$ and the counter c_2 is contained in $V(T_2)$, so immediately after (a, b) the first counter c_1 is contained in $V(T_1) \cup \{b\}$ and the second counter c_2 is contained in $V(T_2) \cup \{a\}$. We replace the initial sequence of transpositions, up to and including (a, b), by the following sequence. Start with the transpositions $(1, v)$ for all $v \in V(T_1) \setminus \{a\}$, followed by the transpositions $(1, u)$ for all $u \in V(T_2) \setminus \{b\}$. Lastly, we do $(1, 2)$, $(1, a)$ and $(2, b)$. The new sequence can reach every pair (x, y) that the original sequence could reach, and uses at most as many transpositions. We colour the edge $(1, 2)$ black, so the black edges form a spanning tree which we will view as rooted at the pair of vertices $\{1, 2\}$.

We now prove that the sum of the red degrees (counted at the end of the process) is at least $n - 2$, which shows that the number of edges in G (and transpositions in our sequence) is at least $n - 1 + \lceil (n - 2)/2 \rceil = \lceil 3n/2 \rceil - 2$, as desired. We first need some more definitions. The *subtree* T_v rooted at v is the set of vertices u (including v) such that the unique path from u to $\{1, 2\}$ in the black tree passes through v. Any vertex in the subtree rooted at v is said to be *above* v. The *deficit* of a vertex v is defined as

$$\mathrm{def}(v) = |V(T_v)| - \sum_{u \in V(T_v)} d_R(u),$$

where $d_R(u)$ denotes the number of red edges incident to u. Note that the red edges may join vertices of T_v to vertices outside T_v, so may not sit inside the tree. In order to bound the sum of the red degrees, we will inductively bound the deficit of vertices.

Each vertex $v \notin \{1, 2\}$ has a unique vertex adjacent to v on the unique path in the black tree from $\{1, 2\}$ to v. We call this the *parent* of v and we denote it by $p(v)$ (which could equal 1 or 2). The *children* $\mathcal{C}(v)$ of a vertex v are the vertices u for which $p(u) = v$. We may rewrite the formula for the deficit of v in the following inductive manner.

$$\mathrm{def}(v) = 1 - d_R(v) + \sum_{u \in \mathcal{C}(v)} \mathrm{def}(u).$$

Note that the first edge that can carry a counter to $v \notin \{1, 2\}$ is the black edge to its parent; if v is not incident to a red edge, then this is the only way to get a counter to v.

We show the following claim inductively, which we will extend to $v \in \{1, 2\}$ to finish the proof.

Claim 6. *Every vertex $v \notin \{1,2\}$ has deficit at most 1. Moreover, if the deficit of v is equal to 1, then there is a vertex ℓ such that the only way for a counter to reach ℓ is to enter v using the black edge $p(v)v$ and then to follow the path of black edges from v to ℓ.*

Proof. We prove the claim by induction on the height of the subtree rooted at $v \notin \{1,2\}$. If the height is 0, that is, v is a leaf of the black tree, then the deficit is at most 1 and equals 1 if and only if v can only be reached via the black edge from its parent $p(v)$.

Now suppose that the claim has been shown up to height $h \geq 0$, and suppose the subtree rooted at v has height $h+1$. We first show that there are at least as many red edges incident to v as there are children of v with deficit 1, which implies that deficit $\sum_{u \in V(T_v)}(1 - d_R(u))$ of v is at most

$$1 - d_R(v) + |\{u \in V(T_v) : d_R(u) = 1 \text{ and } u \neq v\}| \leq 1.$$

If a vertex u of deficit 1 is above v and connected to v via a black edge, then by induction there exists a vertex $\ell(u)$ in the subtree rooted at u which can only be reached using black edges from v. Let vz be the last black edge from v to a child of v with deficit 1. Then it is not possible to get counters into both v and $\ell(z)$ unless there is a red edge incident with v after (v,z). Similarly, if v has black edges to vertices $x, y \in \mathcal{C}(v)$ that have have deficit 1, then there has to be a red edge incident with v, produced by some transposition that occurs between the transpositions (v,x) and (v,y) else there is no way to reach $(\ell(x), \ell(y))$.

Suppose now that v has deficit exactly 1. We need to find a vertex ℓ which can only be reached using a path of black edges from the parent of v. If none of the children of v have deficit 1, then the deficit of v is $1 - d_R(v)$ and there are no red edges incident to v. Hence, we can take $\ell = v$. Otherwise, let x be the first neighbour above v which has deficit 1. Since, v has deficit exactly 1, our argument above shows that there cannot be a red edge incident to v which corresponds to a transposition before (v,x). Hence, the only way of getting a counter into $\ell(x)$ is using the black edge to v, the edge (v,x) and the path of black edges from x to $\ell(x)$, and we take $\ell = \ell(x)$. This completes the proof of Claim 6. □

We now show that the vertices 1 and 2 have deficit at most 1 as well. The difficulty in this case is that it is possible to put a counter on 1 (say) using the black edge $(1,2)$. This means that we can put a counter on $\ell(x)$ and $\ell(y)$ without requiring a red edge between them. However, this does not take into account the fact that the counters are distinguishable, and we only need to modify the argument from the claim slightly to make use of this.

Let the counters starting in positions 1 and 2 be c_1 and c_2 respectively, and suppose there is a child x of 1 which has deficit 1. Then there is a vertex $\ell(x)$ such that the only way for a counter to reach $\ell(x)$ is to follow the path of black edges from 1 to x to $\ell(x)$. If the transposition $(1, x)$ occurs before the black edge $(1, 2)$, then there is no way to move the counter c_2 to $\ell(x)$, and x must have deficit 0. The black edge $(1, 2)$ is the only way to put a counter on 1 without using a red edge, so after this transposition we can apply a proof similar to that of the claim above to see the deficit of 1 is at most 1. Indeed, if x and y are children of 1 with deficit 1, then there has to be some transposition that can place a counter on 1 in between $(1, x)$ and $(1, y)$ in order to reach $(\ell(x), \ell(y))$. We have just argued that this is not the transposition $(1, 2)$, and so it must be a red transposition. Likewise, there must be a red transposition after the last black edge to a vertex z of deficit 1 in order to end with the counters in 1 and $\ell(z)$.

Since every vertex is in either the subtree rooted at 1 or the subtree rooted at 2 and the deficit of these two subtrees is at most 1, the sum of the red degrees must be at least $n - 2$. Hence, there are at least $\lceil (n-2)/2 \rceil$ red edges and at least $\lceil 3n/2 \rceil - 2$ transpositions in the sequence. $\square$

3 Upper bound for t-reachability

We now give a probabilistic construction which shows that there are t-reachability networks on n elements of length $(2 + o_t(1))n$. This is much smaller than the best known constructions for uniformity and, surprisingly, the coefficient of the leading term is bounded. We remark that all the transpositions used in this construction are star transpositions.

Proof of Theorem 3. Let $t \geq 3$ be a natural number and choose $\epsilon \in (0, \frac{1}{t+1})$. Let $n > t$ and set $L = \lfloor n^{1-\varepsilon} \rfloor$. Recall that the counters start in positions $\{1, \ldots, t\}$.

Let G be a random bipartite graph with vertices $A \cup B = \{a_{t+1}, \ldots, a_n\} \cup \{b_1, \ldots, b_L\}$ constructed by adding two uniformly random edges from a_j to $\{b_1, \ldots, b_L\}$ for each j. The random transposition sequence begins with L phases. In the ith phase, we add the transpositions $(1, j)$ for each $j \in \{2, \ldots, t\}$ followed by $(1, j)$ for the $j \in \{t+1, \ldots n\}$ such that a_j is adjacent to b_i. Finally, the sequence ends with a t-reachable network over the first t positions.

Fix positions $x_1, \ldots, x_t \in \{1, \ldots, n\}$ for which we need to find a subsequence that puts our t counters into those positions. Let the x_i which are in $\{t+1, \ldots, n\}$ be $x_{i_1}, \ldots, x_{i_m} = y_1, \ldots, y_m$.

Suppose that there is a matching $\{a_{y_j}b_{s_j} : j \in [m]\}$ in G containing the vertices $a_{y_1}, \ldots, a_{y_m}$. For $j \in [t]$, we can put the jth counter in position $y_j = x_{i_j}$ during phase s_j using the transpositions $(1, j)$ and $(1, y_j)$. The remaining counters can easily be positioned using the t-reachable sequence at the end.

It remains to show that there is some choice for G such that every set of $s \leq t$ vertices from A is in a matching. The probability that a given set of s vertices from A has at most $s-1$ neighbours in B is at most

$$\binom{L}{s-1}\left(\frac{s-1}{L}\right)^{2s} \leq \frac{(s-1)^{2s}}{(s-1)!}L^{-(1+s)}.$$

Hence, by the union bound, the probability that there is a set $A' \subseteq A$ of size at most t and with at most $|A'|-1$ neighbours is at most

$$\sum_{s=3}^{t}\binom{n}{s}\frac{(s-1)^{2s}}{(s-1)!}L^{-(1+s)} \leq \sum_{s=3}^{t}\frac{(s-1)^{2s}}{s!(s-1)!}\frac{n^s}{L^{s+1}} = O_t\left(n^{-1+(t+1)\varepsilon}\right).$$

Since $\varepsilon < 1/(t+1)$, this probability is less than 1 for n sufficiently large. This implies there exists a suitable choice for the graph G.

Note that for $j \in \{t+1, \ldots, n\}$, the transposition $(1, j)$ appears exactly twice in the sequence since the vertex a_j has degree exactly 2. We can create a t-reachable network on the first t positions using $O(t \log t)$ star transpositions. Hence, there exists a t-reachability network using at most

$$(t-1)L + 2(n-t) + O(t \log t) = 2n + o_t(n)$$

transpositions. □

4 Star transpositions

We now turn our attention to the setting where all transpositions are of the form $(1, \cdot)$ and prove Theorem 4, which shows that restricting to star transpositions leads to only a small difference in the number of transpositions needed for 2-reachability. In fact, there is no difference when n is odd, and they differ by only 1 when n is even.

Let us first consider the upper bound. The idea is simple: we will start with the transposition $(1, 2)$ so the two counters are indistinguishable, then we will sweep through the even numbers and potentially load one of them with a counter. Finally, we will sweep through every position. That is,

1. $(1, 2)$,

2. $(1,4),(1,6),\ldots,(1,2\lfloor n/2\rfloor)$,

3. $(1,2),(1,3),(1,4),\ldots,(1,n)$.

It is easy to see that the counters can reach (x,y) when x and y are not both odd: simply place one counter in an even position in the first sweep and place the other counter in the second sweep. To put a counter on x and y when both are odd and $1 < x < y$, we load $(x+1)$ in the first sweep and use this to "reload" 1 before the transposition $(1,y)$. When $x = 1$, this sequence works when n is even and we can use $(1,n)$ to leave a counter on 1, but it breaks down when n is odd. This could be fixing by appending the transposition $(1,2)$ say, but it is possible to do slightly better with the following "twisted" sequence. Let $n = 2m+1$. The following sequence of transpositions is 2-reachable.

1. $(1,2)$,

2. $(1,4),(1,6),\ldots,(1,2m)$,

3. $(1,3),(1,2),(1,5),(1,4),\ldots,(1,2m+1),(1,2m)$.

Proof of Theorem 4. The upper bound is given in the discussion above, so we only need to prove a matching lower bound. For this we proceed much as in the proof of Theorem 2. We colour the transposition $(1,x)$ black if it is the only occurrence, red if it has occurred previously and blue otherwise (i.e. if it is a first occurrence, but will occur later as well). There are $n-1$ transpositions that are either black or blue, and it suffices to show that there are at least $(n-1)/2$ red transpositions.

We claim that the number of red transpositions is at least the number of black transpositions. Let the two counters be c_1 and c_2. Suppose $(1,x)$ is the last black transposition where $x \geq 3$. Then there needs to be a red transposition after $(1,x)$ in order to leave the counters c_1 and c_2 in 1 and x respectively. Similarly, in between any pair of black transpositions $(1,x)$ and $(1,y)$ where $x,y \geq 3$ there has to be a red transposition in order to leave the counters c_1 and c_2 in y and x respectively. This shows the claim when $(1,2)$ is used multiple times, and there is no black $(1,2)$.

If $(1,2)$ is used exactly once, then it must be the first black transposition. Indeed, if $(1,x)$ is a black transposition where $x \geq 3$, then there is no way to put the counter c_2 in x unless $(1,2)$ has already been. If $(1,2)$ were to be the only black transposition, then $(1,x)$ must appear twice for every $x \geq 3$ and there are $2(n-2)+1$ transpositions. This is at least $\lceil 3(n-1)/2\rceil$ for $n \geq 3$ and we would therefore be done. Hence, if $(1,2)$ is used exactly once, it must be the first but not the last black transposition. The above argument shows there is a red transposition after the last black transposition and in between every pair of black transpositions $(1, x)$ and $(1, y)$ where $x, y \geq 3$. It is enough to show there is a red transposition between $(1, 2)$ and the first black transposition $(1, x)$ where $x \geq 3$, and this is easy to argue: there must be a red transposition between $(1, 2)$ and $(1, x)$ in order to leave c_1 in x and c_2 in 1.

By definition the number of red transpositions is at least the number of blue transpositions. Since there is a total of $n - 1$ black and blue transpositions, there must be at least $(n-1)/2$ of one of them and there are at least $(n-1)/2$ red transpositions. $\square$

The lower bound for 2-reachability above immediately gives a lower bound for the more difficult problem of constructing a 2-uniformity network with star transpositions but, using the restrictive nature of the transpositions, we can show that a 2-uniformity network has length at least $1.6n - C$, a constant factor higher. This confirms, for this specific case, that the problem of uniformity is strictly more difficult than reachability.

Proof of Theorem 5. We use a discharging argument to show that, after disregarding a constant number of the transpositions, 1.6 is a lower bound on the average number of times that a star transposition is used.

Let $T = (T_1, \ldots, T_\ell)$ be a sequence of star transpositions that forms a 2-uniformity network on n elements. We assign each transposition $(1, a)$ a weight equal to the number of times $(1, a)$ is used in T. Let $\sigma_1, \ldots \sigma_m$ be the transpositions of T which are used exactly once and are of the form $(1, x)$ for $x \geq 3$. We transfer weight from the transpositions which are used multiple times to the transpositions which are used exactly once according to the following rules.

- If the last appearance of $(1, a)$ and $(1, b)$ is between σ_i and σ_{i+1}, then they each transfer transfer 0.3 to σ_i and 0.1 to σ_{i+1}.
- If there is only one transposition which appears for the last time between σ_i and σ_{i+1}, it transfers 0.4 to σ_i.
- Each transposition which appears between σ_i and σ_{i+1} for neither the first time nor the last time transfers 0.6 to σ_i and 0.2 to σ_{i+1}.

If $(1, a)$ is used exactly twice, then it transfers 0.4 and ends with weight 1.6. A transposition used more than twice transfers $0.8(i-2) + 0.4$ so ends with weight $1.6 + 0.2(i-2)$. Hence, we only need to show that all but a constant number of the transpositions that are used once (the σ_i) end with weight at least 1.6.

Let $\sigma_{i-1} = (1, a)$ and $\sigma_i = (1, b)$. In order to end with counters in both a and b, there must a transposition between σ_{i-1} and σ_i which can place a counter in position 1. In particular, there is either a transposition which is not appearing for the first time, or $(1, 2)$ appearing for the first time. This immediately shows that all but one of $\sigma_1, \dots, \sigma_{m-1}$ have weight at least 1.4. We claim that in between either σ_{i-1} and σ_i or between σ_i and σ_{i+1} one of the following must hold:

1. there are at least two transpositions appearing for the last time,
2. there is a transposition appearing for neither the first nor the last time,
3. there is the transposition $(1, 2)$ appearing for the first time.

If one of the first two cases occurs, then σ_i ends with weight at least 1.6, while the last case can only occur twice. This analysis does not apply to σ_1 and σ_m, showing that the number of σ_i which ends up with weight less than 1.6 is at most four.

We now prove the above claim. Let $\sigma_{i-1} = (1, a)$, $\sigma_i = (1, b)$ and $\sigma_{i+1} = (1, c)$ and assume that none of the conditions above hold. In between σ_{i-1} and σ_i, there can be only transpositions used for the first time and a single transposition $(1, \ell)$ used for the last time. We may write the sequence as

$$(1, a), (1, f_1), \dots, (1, f_k), (1, \ell), (1, f_{k+1}), \dots, (1, f_s), (1, b) \tag{1}$$

where the f_j are distinct, not equal to 2 and are appearing for the first time.

The only way to end the sequence of lazy transpositions with the counters in $\{a, b\}$, $\{a, \ell\}$ or $\{b, \ell\}$ is to start the sequence (1) with the two counters in $\{1, \ell\}$. Let p_j be the probability associated with the transposition $(1, j)$ in (1). The probability that this sequence ends with counters in $\{a, b\}$ and $\{a, \ell\}$ must be equal, so

$$p_a(1 - p_\ell) = p_a p_\ell (1 - p_{f_{k+1}}) \cdots (1 - p_{f_s}) p_b.$$

In particular, if $q = (1 - p_{f_{k+1}}) \cdots (1 - p_{f_s})$, then $1 - p_\ell = p_\ell q p_b$ which implies $p_\ell = 1/(qp_b + 1) \geq 1/2$. We now claim that the probability a counter ends in ℓ is strictly higher than the probability a counter ends in b unless $qp_b = 1$. We condition on which of 1 and ℓ contain a counter before the transposition $(1, \ell)$ and calculate the probability a counter ends in each of the positions ℓ and b.

Starting positions	{1,ℓ}	{ℓ}	{1}	$\emptyset$
Counter ends in ℓ	1	$1-p_\ell$	p_ℓ	0
Counter ends in b	qp_b	$p_\ell qp_b$	$(1-p_\ell)qp_b$	0

In every case, the probability that a counter ends in ℓ is at least the probability that a counter ends in b, and the first one is strict unless $qp_b = 1$.

Now we consider the transpositions between $(1, b)$ and $(1, c)$ below.

$$(1, b), (1, f'_1), \dots, (1, f'_k), (1, \ell'), (1, f'_{k+1}), \dots, (1, f'_s), (1, c)$$

Since the transposition $(1, b)$ fires with probability 1, there is no way for the counters to end in both ℓ' and c, giving a contradiction and proving the claim. □

We have seen in Theorem 4 that the shortest 2-reachability network has length approximately $3n/2$. Moreover, it is possible to construct a 3-reachable network using $5n/3 + C$ such transpositions. Using a similar argument to the proof of Theorem 5, we offer the following lower bound for arbitrary t when using star transpositions, which is tight up to additive constants when $t = 2, 3$.

Proposition 7. *For each $t \geq 1$, there is a constant C_t such that any t-reachability network on n elements using only star transpositions has length at least $(2 - 1/t)n - C_t$.*

Proof. We will again apply a discharging argument. Let T be a given t-reachability network consisting of star transpositions. Similar to before, we assign the transposition $(1, a)$ a weight equal to the number of times $(1, a)$ is used and let $\sigma_1, \dots, \sigma_m$ be the subsequence of transpositions that are used exactly once and of the form $(1, x)$ for $x \notin [t]$. We transfer weight from a transposition used multiple times according to the σ_i using the following simple rules.

- For each transposition between σ_i and σ_{i+1} used for the last time (but not the first), transfer $1/t$ to σ_i.
- For each transposition between σ_i and σ_{i+1} used for neither the first time nor the last time, transfer 1 to σ_i.

We will show all but at most $2t$ transpositions end up with weight at least $2 - 1/t$. Clearly, any transposition which is used multiple times ends with weight $2 - 1/t$ as required. We will ignore any σ_i, σ_{i+1} for which some transposition $(1, x)$ with $x \in [t]$ occurs for the first time between them. This disregards at most $2(t-1)$ of the σ_i.

We show all the other σ_i (except σ_m) get weight at least $2-1/t$. Consider the section between $\sigma_i = (1,a)$ and $\sigma_{i+1} = (1,b)$

$$(1,a)(1,s_1)\cdots(1,s_\ell)(1,b).$$

If there is a transposition between σ_i and σ_{i+1} used for neither the first nor the last time, then σ_i has weight at least 2. So we assume all $(1,s_j)$ are used for either the first or last time. Let $s'_1,\dots,s'_{\ell'}$ be the transpositions used for the last time. If $\ell' \geq t-1$, then σ_i receives weight at least $2-1/t$ as desired, and we show that this is always the case.

We show this by proving that we cannot simultaneously reach the positions $\{a, s'_1,\dots,s'_{\ell'}, b\}$ (which we should be able to if it has at most t elements). The transposition $\sigma_i = (1,a)$ must be used (as this is the only way to put a counter on a) and then the only way for a counter to 'reload' position 1 before σ_{i+1} is using a transposition $(1,s'_j)$ that has been used before. (Here we use our assumption that no $(1,x)$ appears for the first time in our segment for $x \in [t]$.) When we use $(1,s'_j)$ to 'reload', then no counter can end in position s'_j since this is the last use of this transposition. Hence, we cannot 'reload' the position 1 to leave a counter in b. $\square$

5 Open problems

We determined exactly the minimum number of transpositions needed in a 2-reachable network. We also gave upper and lower bounds for t-reachable networks using star transpositions, although there is still a small gap between them.

Problem 1. *What is the minimum number of transpositions needed in a t-reachable network? What if all transpositions are of the form $(1,\cdot)$?*

We proved in Theorem 5 that there is a gap between 2-uniformity and 2-reachability when restricting to star transpositions. There are many 2-uniformity networks of length $2n-3$, e.g. consider the sequence of lazy transpositions

$$\left(1,2,\tfrac{1}{2}\right),\left(1,3,\tfrac{2}{n}\right),\left(1,2,\tfrac{1}{2}\right),\left(1,4,\tfrac{2}{n-1}\right),\dots,\left(1,2,\tfrac{1}{2}\right),\left(1,n,\tfrac{2}{3}\right),\left(1,2,\tfrac{1}{2}\right),$$

and we conjecture that no smaller sequences exist, i.e. $U_2(n) = 2n-3$. We remark that, if this were true, it would match nicely with selection networks.

Conjecture 1. *For $n \geq 2$, $U_2(n) = 2n-3$.*

As a first case, it would be interesting to close the gap when restricting to star transpositions and confirm the conjecture in this case.

References

[1] O. Angel and A. E. Holroyd. Perfect shuffling by lazy swaps. *Electronic Communications in Probability*, **23**, 2018.

[2] B. Beauquier and E. Darrot. On arbitrary size Waksman networks and their vulnerability. *Parallel Processing Letters*, **12**(03n04):287–296, 2002.

[3] A. Czumaj. Random permutations using switching networks. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, STOC '15, pages 703–712. Association for Computing Machinery, 2015.

[4] W. J. Dally and B. P. Towles. *Principles and Practices of Interconnection Networks*. Elsevier, 2004.

[5] C. Groenland, T. Johnston, J. Radcliffe and A. Scott. Perfect shuffling with fewer lazy transpositions. *In preparation*, 2022.

[6] A. Waksman. A permutation network. *Journal of the ACM (JACM)*, **15**(1):159–163, 1968.