On the Sizes of Permutation Networks and Consequences for Efficient Simulation of Hypercube Algorithms on Bounded-Degree Networks *

(Extended Abstract)

J. Hromkovič [1,**], K. Loryś [2,***] , P. Kanarek [2,***], R. Klasing [3], W. Unger [3], H. Wagener [3,†]

[1] Institut für Informatik und Praktische Mathematik, Universität zu Kiel, D-24098 Kiel, Germany, E-Mail: jhr@informatik.uni-kiel.d400.de
[2] University of Wrocław, Institute of Computer Science, PL-51-151 Wrocław, Poland, E-Mail: lorys,pka@ii.uni.wroc.pl
[3] Fachbereich Mathematik/Informatik, Universität-GH Paderborn, D-33095 Paderborn, Germany, E-Mail: klasing@uni-paderborn.de

Classification: theory of parallel and distributed computation, parallel algorithms.

Abstract. The sizes of permutation networks for special sets of permutations are investigated. The study of the planar realization and the search for small but hard sets of permutations are also included. Several asymptotically optimal estimations for distinct subsets of the set of all permutations are established here.
The two main results are:
(i) an asymptotically optimal permutation network of size $6 \cdot N \cdot \log\log N$ for shifts of power 2.
(ii) an asymptotically optimal planar permutation network of size $\Theta(N^2 \cdot (\log\log N/\log N)^2)$ for shifts of power 2.
A consequence of our results is a construction of a 4-degree network which can simulate each communication step of any hypercube algorithm using edges from at most a constant number of different dimensions in one step in $O(\log\log N)$ communication steps. A new sorting network as well as an essential improvement of gossiping in vertex-disjoint path mode in bounded-degree networks follow.

* This work was partially supported by grants Mo 285/9-1 and Me 872/6-1 (Leibniz Award) of the German Research Association (DFG), and by the ESPRIT Basic Research Action No. 7141 (ALCOM II).
** Supported by SAV Grant 2/1138/94 and by EC Cooperation Action IC 1000 ALTEC
*** Supported by Polish Government grants KBN 2 1197 91 01, KBN 8 S503 002 07, KBN 2 P301 034 07 & Volkswagen Stiftung (Joint project of Wrocław University and Heinz–Nixdorf–Institut, University of Paderborn).
† This author was supported by the Ministerium für Wissenschaft und Forschung des Landes Nordrhein-Westfalen.

1 Introduction and Definition

The study and the comparison of the computational power of distinct interconnection networks as candidates for the use as parallel architectures for existing parallel computers is an intensively investigated research branch of current computation theory. One of the fundamental approaches helping to search for the best (most effective) structures of interconnection networks is the study of the communication facilities of networks (i.e., of the complexity (effectivity) of solving fundamental communication tasks). The basic communication tasks are routing (for an extensive survey see [Le92]), broadcasting, and gossiping. Our paper is devoted to the search for effective, realistic 1-to-1 routing (permutation) networks. The effectivity in our paper is mainly interpreted as the minimal size of a network realizing a given subset of permutations (one-to-one routing tasks). By realistic we mean that we consider only networks with bounded degree (usually with degree 4).

Our study of the sizes of permutation networks for distinct subclasses of permutations has been motivated not only by searching for efficient bounded-degree networks for special routing tasks, but mainly by looking for realistic networks in the class of bounded-degree networks with very high (or even with the highest possible) communication facilities supporting different fundamental computing and communicational tasks. For instance, the study of permutation networks for shifts 2^i is well-motivated by the simulation of hypercube algorithms in networks of bounded degree, and, in our case, especially by sorting and gossiping.

This paper is organized as follows. In this section we give the basic definitions and the overview of the achieved results. Section 2 contains the exact formulation of the results. In this extended abstract, we give some proof ideas only for theorems which do not have long technical proofs, and we completely omit the proof of theorems requiring too many technical details.

Section 2 is divided into three subsections. Section 2.1 is devoted to permutation networks for all permutations of N elements, all shifts, shifts of power two, and shifts by Fibonacci numbers. The main results of this subsection are the constructions of 4-degree permutation networks of size $6 \cdot N \cdot \log\log N$ for the shifts 2^i for $i = 1, 2, \ldots, \log_2 N$, and the construction of 4-degree permutation networks of size $O\left(N \cdot 2^{\,2\sqrt{\log\log N}}\right)$ for Fibonacci shifts. The results for shifts of power of two is asymptotically optimal.

Subsection 2.2 is devoted to planar permutation networks. Here one needs to distinguish between vertex-disjoint and edge-disjoint permutation networks. For instance, for the vertex-disjoint path mode one needs networks of size $\Theta(N^3)$ to realize all permutations [KL92], but we show that the size $O(N^2)$ is necessary and sufficient for the realization in the edge-disjoint path mode. Moreover, five random permutations are as hard as the set of all permutations for planar realization in the edge-disjoint path mode. Finally, the main result of this subsection shows that the optimal planar permutation network for shifts of power two has the size in $\Theta(N^2 \cdot (\log\log N/\log N)^2)$.

The last Subsection 2.3 shows various applications of the results of the previous

subsections. The smallest delay to simulate hypercubes of N nodes on bounded-degree networks is $\log_2 N$. A constant delay is possible if the communication steps of the hypercube algorithm consist only of communication via edges of one dimension and additionally any two consecutive communication steps correspond to communication in two consecutive dimensions [Ul84]. Here, using the permutation networks for shifts 2^i, we obtain the delay $6d \cdot \log\log N$ for the simulation of hypercube algorithms communicating via edges of d dimensions in one communication step. As a consequence of this, we obtain a sorting network of size $6 \cdot N \log\log N$ which sorts N elements in time $O(\log N \cdot (\log\log N)^3)$. Our network is neither the best in time (the Ajtai et al. network [AKS83] works in $O(\log_2 N)$ time), nor in hardware (Plaxton [CP93] and Kutyłowski [KLOW94] have quick networks using $O(N)$ hardware). But our network is an intermediate case from the periodic sorting point of view. The network of Kutyłowski et al. [KLOW94] has depth 5 and the data periodically flows $O((\log_2 N)^2)$ times via the 5 levels of the network to sort. The network of Ajtai et al. [AKS83] can be viewed as a periodic network of depth $O(\log_2 N)$ (of size $N \cdot \log N$) with only one flow of data via the $O(\log N)$ levels of the network. Our network has depth $O(\log\log N)$ and the number of repetitions is $O(\log N \cdot (\log\log N)^2)$. Another consequence are quick gossip algorithms for bounded-degree networks in vertex-disjoint path mode. This is an essential improvement of the results established in [HKS93, HKSW93, HKUW94, Kl94] for gossiping in disjoint-path modes.
Now, we give the fundamental definitions. Let Π_N denote the set of all permutations of N elements. For any set A, $|A|$ denotes the cardinality of A. $F(i)$ denotes the i-th Fibonacci number of any positive integer i.

Definition 1. Let A be a subset of Π_N, and k be a positive integer. We say that a graph $G = (V, E)$ is a **vertex-disjoint (edge-disjoint) k-permutation network for A** if the following conditions hold:

(i) the degree of G is bounded by k,
(ii) V contains $2N$ special vertices $x_1, x_2, \ldots, x_N, y_1, y_2, \ldots, y_N$,
($x_1, x_2, \ldots, x_N$ are called **inputs**, $y_1, y_2, \ldots, y_N$ are called **outputs** of G)
(iii) for every permutation $\pi = (i_1, i_2, \ldots, i_N) \in A$ there exist N vertex-disjoint (edge-disjoint) paths $x_1, \ldots, y_{i_1}; x_2, \ldots, y_{i_2}; \ldots; x_N, \ldots, y_{i_N}$ in G.

The size of G, denoted $\boldsymbol{s(G)}$, is $|V|$. The depth of G (according to A), denoted **depth(G)**, is $\max\{\min\{\text{ length of } P \mid P \text{ is a path from } x_i \text{ to } y_j\} \mid i, j \in \{1, \ldots, N\},\ \pi(i) = j \text{ for some } \pi \in A\}$. A k-permutation network $G = (V, E)$ is **leveled** if V can be partitioned into $r = \text{depth}(G) + 1$ nonempty subsets $V_1, \ldots, V_r$ such that

(a) $V_1 = \{x_1, \ldots, x_N\}$, $V_r = \{y_1, \ldots, y_N\}$;
(b) $V_i \cap V_j = \emptyset$ for $i \neq j$, $i, j \in 1, \ldots, r\}$;
(c) for every edge $e \in E$, there exists $i \in \{1, \ldots, r-1\}$ such that e connects one vertex from V_i with one vertex from V_{i+1}. □

Note that almost all permutation networks constructed here are leveled and of degree at most 4. Moreover, we use the edges only in the direction from the

i-th level to the $(i+1)$-st level. On the other hand, the lower bound proofs are general, i.e, they work for any constant-degree network and the property to be leveled is not needed.

Definition 2. Let A be a subset of Π_N, and k be a positive integer. We define

$$\begin{aligned}
\mathbf{Size\text{-}vd_k(A)} &= \min\{s(G) \mid G \text{ is a} \\
&\qquad \text{vertex-disjoint } k\text{-permutation network for } A\}, \\
\mathbf{Size\text{-}ed_k(A)} &= \min\{s(G) \mid G \text{ is an} \\
&\qquad \text{edge-disjoint } k\text{-permutation network for } A\}, \\
\mathbf{Plsize\text{-}vd_k(A)} &= \min\{s(G) \mid G \text{ is planar and } G \text{ is a} \\
&\qquad \text{vertex-disjoint } k\text{-permutation network for } A\}, \\
\mathbf{Plsize\text{-}ed_k(A)} &= \min\{s(G) \mid G \text{ is planar and } G \text{ is an} \\
&\qquad \text{edge-disjoint } k\text{-permutation network for } A\}.
\end{aligned}$$

□

In what follows we use the following notation.

$S_N \subseteq \Pi_N$ ... the set of all shifts,

$\mathrm{Pow}\,2_N \subseteq S_N$ ... the set of all shifts of the power of 2 $((1+2^i) \bmod N, (2+2^i) \bmod N, \ldots, (N+2^i) \bmod N)$ for $i = 1, 2, \ldots, \log_2 N$,

$\mathrm{Fib}_N \subseteq S_N$ ... the set of all shifts $((1+F(i)) \bmod N, (2+F(i)) \bmod N, \ldots, (N+F(i)) \bmod N)$ for all Fibonacci numbers $F(1), F(2), \ldots, F(k)$, where k is the largest integer with the property $F(k) \leq N$.

Definition 3. Let N be a positive integer, and A be a subset of Π_N. We say that B is a **kernel** of A if $B \subseteq A$ and $\text{Size-ed}_k(B) = \Omega(\text{Size-ed}_k(A))$ for any constant $k \geq 4$. We say that B is a **planar *e*-kernel** [***v*-kernel**] of A if $B \subseteq A$ and for any constant $k \geq 4$ $\text{Plsize-ed}_k(B) = \Omega(\text{Plsize-ed}_k(A))$ [$\text{Plsize-vd}_k(B) = \Omega(\text{Plsize-vd}_k(A))$]. A kernel (planar e-kernel, planar v-kernel) B of A is called **(asymptotically) minimal** if, for every kernel (planer e-kernel, planar v-kernel) C of A, $|B| = \Omega(|C|)$.

□

2 Results

2.1 On the Sizes of Permutation Networks for Special Subclasses of Permutations

Since there is no asymptotical difference between $\text{Size-vd}_k(A)$ and $\text{Size-ed}_k(A)$ for any permutation set A we shall consider only the edge-disjoint path mode in this subsection. It is well-known that

$$\text{Size-ed}_k(\Pi_N) = \Theta(N \log N) = \text{Size-ed}_k(S_N).$$

To see the upper bound Size-ed$_k$ $(\Pi_N) = O(N \log N)$ it is sufficient to take the well-known permutation network [Be64, Be65, Wa68]. To see the lower bound Size-ed$_k(S_N) = \Omega(N \log N)$ one has to apply the following lower bound method of [PV76] to S_N.

Lemma 4 [PV76], Theorem 2.2.1. *Let A be any subset of Π_N. If A fulfils the following property*

(i) Each input is assigned to $|A|$ different outputs by the $|A|$ permutations.

then Size-ed$_k(A) = \Omega(N \cdot \log_2 |A|)$. □

Thus, we obtain that S_N is a kernel of Π_N. But S_N is not the minimal kernel of Π_N because a random set A with $|A| \leq \log_2 N$ has Size-ed$_k(A) = \Omega(N \cdot |A|)$ with probability tending to 1 with growing N and $|A|$ [We94]. Obviously each $A \subseteq \Pi_N$ has Size-ed$_4(A) = O(N \cdot |A|)$, and so the random sets of $\log_2 N$ permutations are **minimal kernels** of Π_N.
In [HKS93] it has been conjectured that Pow 2_N and Fib$_N$ are kernels of S_N. We now show that this is not the case.

Theorem 5. *Size-ed$_4$(Pow 2_N) $\leq 6 \cdot N \cdot \log\log N + 3$.*

Idea of the Proof. The construction of the permutation network can be described in a recursive way. Assume we have a permutation network of $M = 2^{i/2}$ inputs for shifts $2^0, 2^1, ..., 2^{i/2}$ and one wants to construct a network for shifts $2^0, 2^1, ..., 2^i$. The idea is to put $M = 2^{i/2}$ networks for shifts $2^0, 2^1, ..., 2^{i/2}$ next to each other in the same levels and to add a few (constant) further levels to connect them. These new levels are created in such a way that they are able to establish the following two distinct kinds of connections:

(i) If one wants to realize "small" shifts of size $z \in \{2^0, 2^1, ..., 2^{i/2}\}$ for $N = 2^i$ inputs, the first levels copy the inputs to the inputs of M "small" networks (of M inputs). Then the "small" networks realize the local shifts of size z. To get the right output of the whole network, the last level shifts the outputs of the small networks which are not at the right positions because of the local mod z shifts realization.
(ii) If one wants to realize "big" shifts of size $z \in \{2^{i/2+1}, 2^{i/2+2}, ..., 2^i\}$ for $N = 2^i$ inputs, the first levels shuffle the inputs to the inputs of the small networks in such a way that the M small networks realize exactly the large shifts.

If this works and the number of additional levels is in $O(1)$, then one obtains the following recurrence for the depth of the permutation network with N inputs: $depth(N) \leq depth(\sqrt{N}) + O(1)$. Obviously, the solution is $depth(N) \leq O(\log\log N)$. Some additional technicalities are necessary to get the constant small (i.e. $depth(N) \leq 6 \cdot \log\log N + O(1)$). □

Theorem 6. *Size-ed$_4$(Fib$_N$) $\leq 13 \cdot N \cdot 2^{2\sqrt{\log\log N}}$.*

Proof. The proof is omitted in this extended abstract. □

Corollary 7. *Size-ed*$_k$*(Pow* 2_N*)* $= \Theta(N \cdot \log\log N)$ *for any constant* k *independent of* N.

Proof. The upper bound is in Theorem 5 and the lower bound follows from the fact that Pow 2_N fulfils Property (i) of Lemma 4. □

Note that especially Theorem 5 brings a crucial contribution for the simulation of hypercube algorithms in bounded-degree networks of degree 4 for several fundamental computing problems. More about this can be found in Subsection 2.3.

2.2 On the Sizes of Planar Permutation Networks

It is well-known that Plsize-vd$_k(\Pi_N) = \Theta(N^3)$ [CS78, KL92]. Our first result shows that there is an essential difference between edge-disjoint paths mode and vertex-disjoint paths mode for planar permutation networks.

Theorem 8. *For any constant* $k \geq 4$ *independent of* N

$$\text{Plsize-ed}_k(\Pi_N) = \Theta(N^2) = \text{Plsize-ed}_k(S_N).$$

Outline of the Proof. We start with the upper bound. Consider the following network $PE_N = (V, E)$ defined by $V = \{\, v_{i,j}\,;\, 1 \leq i < N \;,\; 1 \leq j \leq N\} \cup \{\, x_i, y_i\,;\, 1 \leq i \leq N\}$ and $\{v_{i,j}, v_{k,l}\} \in E$ iff $(i+1 = k$ and $j = l)$ or $(i = k$ and $j \bmod N = l-1)$. Furthermore $\{x_i, v_{1,i}\} \in E$ for $1 \leq i \leq N$ and $\{v_{N-1,i}, y_i\} \in E$ for $1 \leq i \leq N$. Observe that PE_N is planar (see Figure 1).
Furthermore, it is well-known that any permutation from Π_N can be built from $N-1$ transpositions. On each of the rows $(v_{i,1}, v_{i,2}, v_{i,3}, ..., v_{i,N})$ $(1 \leq i < N)$ one of these transpositions can be carried out. This is also shown in Figure 1: The nodes marked by a circle send the messages from the top edge to the right hand side edge and the message from the left hand side to the bottom edge. The other nodes send the messages from the top edge to the bottom edge and the message from the left hand side to the right hand side edge. In Figure 1 the permutation $(5, 4, 1, 3, 2)$ is carried out. Thus we conclude Plsize-ed$_4(\Pi_N) \leq N \cdot (N+1)$.
The lower bound proof is based on the planar separator theorem of Lipton and Tarjan [LT79]. Let, for any subset of permutations $A \subseteq \Pi_N$, $G(A) = (V, E)$, where $V = \{x_1, \ldots, x_N, y_1, \ldots, y_N\}$ and, for every $\pi = (i_1, i_2, \ldots, i_N) \in A$, E contains the edges $(x_1, y_{i_1}), \ldots, (x_N, y_{i_N})$. Let (V_1, V_2), $V_1 \cup V_2 = V$, $|V_1| = |V_2| = N$, be a bisection of $G(A)$. The bisection width of $G(A)$ according to (V_1, V_2) and π, $\underline{bw(G(A), (V_1, V_2), \pi)}$, is the number of edges defined by π and leading between V_1 and V_2. $\underline{bw(G(A), (V_1, V_2))} = \max\{bw(G(A), (V_1, V_2), \pi) \mid \pi \in A\}$. The balanced bisection width of $G(A)$ is $\underline{bw(G(A))} = \min\{bw(G(A), (V_1, V_2)) \mid (V_1, V_2)$ is a bisection of $G(A)\}$. We observe that each permutation network realizing A must have bisection width at least $bw(G(A))$. Since $bw(S_N) \geq \lceil N/2 \rceil$

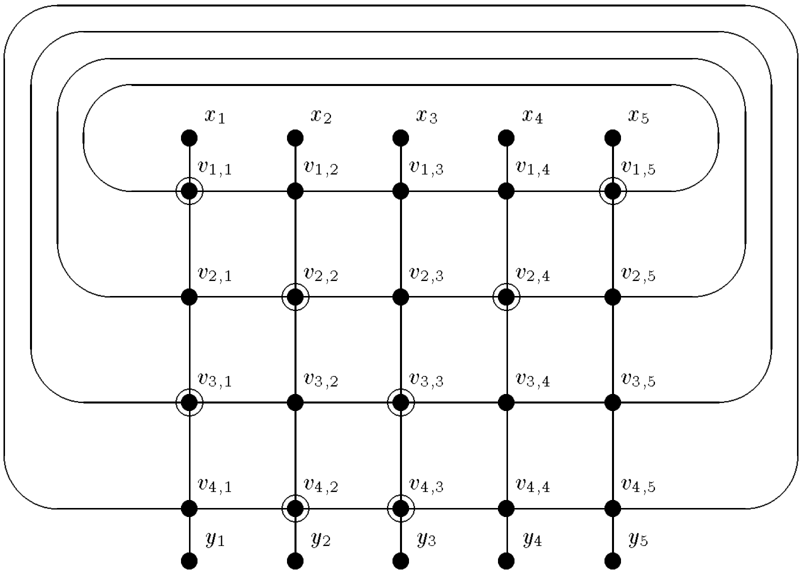

Fig. 1. PE_5

and each planar graph of m nodes has an $O(\sqrt{m})$ separator [LT79], we obtain Plsize-ed$_k(S_N) = \Omega(N^2)$. □

Again, we see that S_N is a planar e-kernel of Π_N. But the next result showing that S_N is no planar v-kernel of Π_N underlines the difference between these two communication modes.

Theorem 9. *Plsize-vd*$_4(S_N) \leq N \cdot (N+2)$.

Outline of the Proof. We define the network $PS_N = (V, E)$ by $V = \{ v_{i,j} \,;\, 0 \leq i, j < N\} \cup \{ x_i, y_i \,;\, 1 \leq i \leq N\}$ and $\{v_{i,j}, v_{k,l}\} \in E$ iff $i + 1 = k$ and $l \in \{j, (j+1) \bmod N\}$. Furthermore $\{x_i, v_{0,i-1}\} \in E$ for $1 \leq i \leq N$ and $\{v_{N-1,i-1}, y_i\} \in E$ for $1 \leq i \leq N$. Observe that PS_N is planar (see Figure 2). The messages from the nodes $\{ x_i \,;\, 1 \leq i \leq N\}$ are shifted to the nodes $\{ y_i \,;\, 1 \leq i \leq N\}$. The $i-th$ message $(1 \leq i \leq N)$ uses the path P_i^s when a shift by s $(0 \leq s < N)$ is carried out, where $P_i^s = (x_i, v_{0,i-1}, v_{1,(i+\min(1,s)-1) \bmod N}, \ldots, v_{2,(i+\min(2,s)-1) \bmod N}, \ldots, v_{j,(i+\min(j,s)-1) \bmod N}, \ldots, v_{N-1,(i+\min(N-1,s)-1) \bmod N}, y_{(i+s-1) \bmod N+1})$. Again it is easy to observe that these paths are vertex-disjoint (see Figure 2).

□

Corollary 10. *For any constant $k \geq 4$ independent of N*

$$\textit{Plsize-vd}_k(S_N) = \Theta(N^2).$$

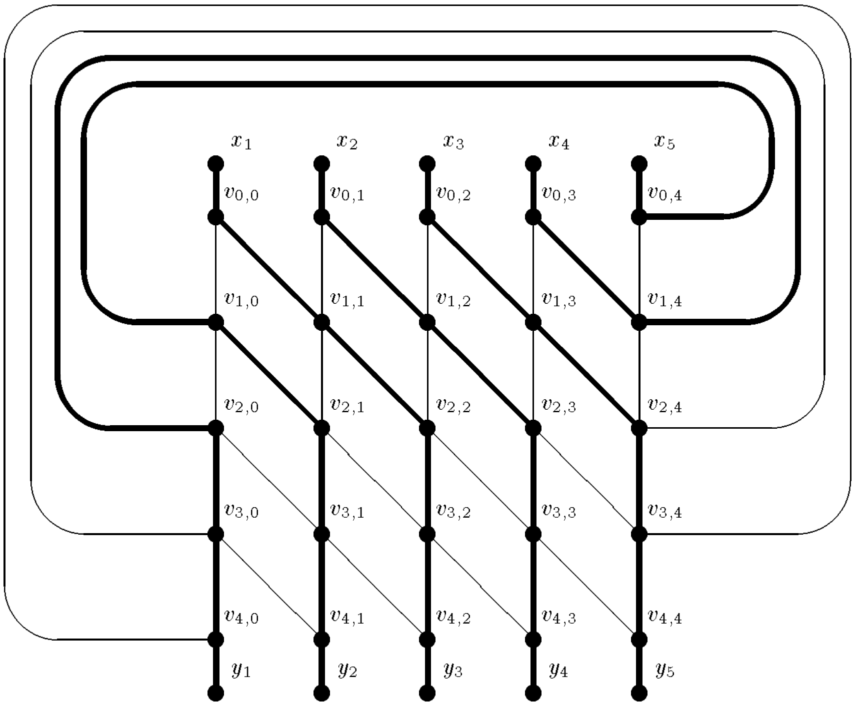

Fig. 2. a shift by 2 realized by PS_5

Proof. The result is a direct consequence of Theorem 8 and Theorem 9. □

Again, as in the general case (Subsection 2.1), S_N is no minimal e-kernel of Π_N. In the following we show that there are e-kernels of Π_N of size 5.

Theorem 11. *Let* $N = m^2$, $A_m = \{0, 1, \ldots, m-1\} \times \{0, 1, \ldots, m-1\}$, *and let* $B_N = \{\pi_1, \pi_2, \pi_3, \pi_4, \pi_5\} \subseteq \Pi_N$ *be the following permutations on* A_m*:*

$$\begin{aligned}
\pi_1(x, y) &= (x, y), \\
\pi_2(x, y) &= (x, x + y), \\
\pi_3(x, y) &= (x, x + y + 1), \\
\pi_4(x, y) &= (x + y, y), \\
\pi_5(x, y) &= (x + y + 1, y),
\end{aligned}$$

where the $+$ *is modulo* m. *Then, for any constant* k *independent of* N,

$$Plsize\text{-}ed_k(B_N) = \Theta(N^2).$$

Theorem 12. *Let* R_N *be a set of 5 random permutations from* Π_N. *Then, with probability tending to 1 with growing* N

$$Plsize\text{-}ed_k(R_N) = \Theta(N^2),$$

for any constant k independent of N.

Idea of the Proofs of Theorem 11 and 12. The proofs are based on the same idea as the lower bound of Theorem 8. Since B_N and R_N define permutation graphs which are expanders [GG81, Pi77], the planar realization of them is hard. □

Finally, let us consider the set Pow 2_N. Despite the fact that Pow 2_N is no planar e-kernel of Π_N, the following result shows that Pow 2_N requires almost the same size for planar realization as Π_N.

Theorem 13. *For any constant k independent of N*

$$\mathit{Plsize\text{-}ed}_k(\mathit{Pow}\,2_N) = \Theta(N^2 \cdot (\log\log N)^2/(\log N)^2)$$

Idea of the Proof. In this case, the lower bound is easier than the upper bound. Leighton has mentioned in [Le92] that the balanced bisection width of the hypercube of N nodes is in $\Omega(N \log\log N/\log N)$. This can be used to show that $bw(G(\text{Pow}\,2_N)) = \Omega(N \log\log N/\log N)$. Using the planar separator theorem [LT79], we obtain $\text{Plsize-ed}_k(\text{Pow}\,2_N) = \Omega(N^2 \cdot (\log\log N)^2/(\log N)^2)$.

The upper bound is based on an embedding of $G(\text{Pow}\,2_N)$ in the planar grids. There are too many technical details needed to present this proof and it does not help to show only part of them. Therefore, we completely omit this part of the proof. □

2.3 Application of Permutation Networks for Sorting and Gossiping

The aim of this subsection is to show the consequences of Theorem 5 and 6 for other tasks than routing of messages. The main idea of the use of Theorem 5 is in the simulation of hypercube algorithms in degree-bounded networks. Here, we consider the hypercube algorithms as synchronized parallel algorithms executing alternatively communication steps and computing steps. In each computing step each processor of the hypercube executes some computation on its local data. In each communication step one processor can exchange some message with one neighbour via an adjacent edge. Each processor may communicate with at most one neighbour in one communication step. We distinguish three classes of hypercube algorithms:

(1) **general hypercube algorithms**, as described above
(2) **leveled hypercube algorithms** [Le92, Wi92], which communicate in every communication step via edges of one fixed dimension of the hypercube
(3) **normal hypercube algorithms** [Ul84] are leveled ones with the additional condition that two consecutive communication steps use two consecutive dimensions of the hypercube.

Additionally, we define ***d*-leveled hypercube algorithms** as hypercube algorithms using edges of at most d different dimensions in each communication step. Any communication step of a normal hypercube algorithm running on a hypercube of size N can be simulated by two communication steps by the Shuffle-Exchange network of size N. For general algorithms, the best simulation of one communication step of the hypercube by bounded-degree networks uses $\log_2 N$ steps. Using the network from Theorem 5 we can construct a 4-degree network simulating one communication step of any leveled hypercube algorithm in $6 \cdot \log_2 \log_2 N$ steps.

Theorem 14. *There is a 4-degree network of size $6 \cdot N \cdot \log\log N$ which can simulate any communication step of any d-leveled hypercube algorithm on a hypercube of N nodes in $6 \cdot d \cdot \log\log N$ communication steps.*

Because almost all common hypercube algorithms are leveled, we obtain several efficient parallel algorithms running on bounded-degree networks as a consequence of Theorem 14. Note that our simulation beats the previous ones in the delay. All the previous simulations of the general and the d-leveled hypercube algorithms by degree-bounded networks had delay $O(\log N)$ [Le92], while we have delay $O(\log\log N)$. Another positive property of our network is its size $6 \cdot N \log\log N$, i.e. we do not need to pay too much with the increase of the size for this small delay.
We omit the formulation of all possible consequences of Theorem 14 in this extended abstract. We only present some results for sorting and gossiping.

Theorem 15. *For every $N = 2^n, n \in I\!N$, there is a leveled network $\overline{G}_N$ of size $6 \cdot N \log\log N$ sorting N elements in $O(\log N \cdot (\log\log N)^3)$ steps. This network has depth $6 \cdot \log\log N$ and the sorting algorithm periodically uses the network $O(\log N \cdot (\log\log N)^2)$ times.*

Our network is neither the best in time (the Ajtai et al. network [AKS83] works in $O(\log_2 N)$ time), nor in hardware (Cypher and Plaxton [CP93] and Kutyłowski et al. [KLOW94] have quick networks using $O(N)$ hardware). But our network is an intermediate case from the periodic sorting point of view. The network of Kutyłowski et al. [KLOW94] has depth 5 and the data periodically flows $O((\log_2 N)^2)$ times via the 5 levels of the network to sort. The network of Ajtai et al. [AKS83] can be viewed as a periodic network of depth $O(\log_2 N)$ (of size $N \cdot \log N$) with only one flow of data via the $O(\log N)$ levels of the network. Our network has depth $O(\log\log N)$ and the number of repetitions is $O(\log N \cdot (\log\log N)^2)$. If one considers the complexity measure $(\text{Depth})^2 \cdot$ (the number of repetitions) for periodic sorting networks [note that this measure is reasonable because the depth of the network influences both important parameters of the network – the size of the hardware and the time complexity], our network is even better than the networks from [AKS83, KLOW94] because we have $(\text{Depth})^2 \cdot (\text{the number of repetitions}) = O(\log N \cdot (\log\log N)^4)$ and both previous networks [AKS83, KLOW94] have $(\text{Depth})^2 \cdot (\text{the number of repetitions}) = O((\log N)^2)$.
Another application of the previous theorems is related to gossiping (all to all broadcasting) in edge-disjoint (vertex-disjoint) paths mode. Let C_N denote the

complete graph of N nodes. In [HKS93] it is shown that there are bounded-degree networks of size N whose

(i) two-way vertex-disjoint gossip complexity is $r_2(C_N) + O(\log\log N)$, where $r_2(C_N) = \log_2 N$ is the two-way gossip complexity of the complete graph C_N on N nodes, and
(ii) one-way vertex-disjoint gossip complexity is $r_1(C_N) + O(\log\log N)$, where $r_1(C_N) \approx 1.44\ldots\log_2 N$ is the one-way gossip complexity of C_N.

Thus, one can gossip in bounded-degree networks almost as quickly as in complete graphs using the vertex-disjoint paths mode. In [HKS93] it has been conjectured that the additional $\log\log N$ steps are necessary to gossip in bounded-degree networks. Theorems 5 and 6 surprisingly provide networks which can gossip even faster.

Theorem 16. *There is a 4-degree network of size $M = 6 \cdot N \cdot \log\log N$ for any positive integer $N \geq 16$ with two-way vertex-disjoint paths gossip complexity smaller than $r_2(C_M) + \log\log\log M + const.$*

Theorem 17. *There is a 4-degree network of size $M = O\left(N \cdot 2^{\,2\sqrt{\log\log N}}\right)$ for every positive integer $N \geq 16$ with one-way vertex-disjoint paths gossip complexity smaller than $r_1(C_M) + 1.12...\sqrt{\log\log M} + const.$*

Idea of the Proofs of Theorems 16 and 17. To see the connection between Theorems 5 and 6 and gossiping in the two-way mode and one-way mode resp. one has to take a look at the concept of three-phase gossip algorithms in [HKS93, HKSW93]. □

3 Conclusion

In this paper, we presented asymptotically optimal planar and non-planar permutation networks for shifts of power 2. A new sorting network as well as an essential improvement of gossiping in two-way vertex-disjoint paths mode in bounded-degree networks follow. For the set of Fibonacci shifts, Fib_N, which are crucial for the design of efficient gossip algorithms in one-way vertex-disjoint paths mode, we were not able to prove the optimality of our constructions. Thus, the main open problem left is to determine Size-ed$_k(Fib_N)$ and Plsize-ed$_k(Fib_N)$ exactly.

References

[AKS83] M. Ajtai, J. Komlós, E. Szemerédi, An $O(n\log n)$ sorting network. *Proc. 15th ACM Symposium on Theory of Computing*, pp. 1-9, 1983.

[Be64] V. Beneš, "Permutation groups, complexes, and rearrangeable multistage connecting networks", *Bell System Technical Journal*, vol. 43, pp. 1619-1640, 1964.

[Be65] V. Beneš, "Mathematical Theory of Connecting Networks and Telephone Traffic", *Academic Press*, New York, NY, 1965.

[CP93] R. Cypher, C.G. Plaxton: Deterministic Sorting in Nearly Logarithmic Time on the Hypercube and Related Computers. *Journal of Computer and System Sciences*, No. 47, 1993, pp. 501-548.

[CS78] M. Cutler, Y. Shiloach: Permutation layout. *Networks*, vol. 8 (1978) 253-278.

[GG81] O. Gabber, Z. Galil: Explicit Construction of Linear-Sized Superconcentrators. *Journal of Computer and System Sciences*, No. 22, 1981, pp. 407-420.

[HKS93] J. Hromkovič, R. Klasing, E.A. Stöhr, "Gossiping in Vertex-Disjoint Paths Mode in Interconnection Networks", Proc. of the 19th Int. Workshop on Graph-Theoretic Concepts in Computer Science (*WG'93*), Springer LNCS 790, pp. 288-300.

[HKSW93] J. Hromkovič, R. Klasing, E.A. Stöhr, H. Wagener, "Gossiping in Vertex-Disjoint Paths Mode in d-Dimensional Grids and Planar Graphs", Proc. of the First Annual European Symposium on Algorithms (*ESA'93*), Springer LNCS 726, pp. 200-211.

[HKUW94] J. Hromkovič, R. Klasing, W. Unger, H. Wagener, "Optimal Algorithms for Broadcast and Gossip in the Edge-Disjoint Path Modes", Proc. of the 4th Scandinavian Workshop on Algorithm Theory (*SWAT'94*), Springer LNCS 824, pp. 219-230.

[Kl94] R. Klasing, "The Relationship Between Gossiping in Vertex-Disjoint Paths Mode and Bisection Width", Proc. of the 19th International Symposium on Mathematical Foundations of Computer Science (MFCS'94), Springer LNCS 841, pp. 473-483.

[KL92] M. Klawe, T. Leighton: A tight lower bound on the size of planar permutation networks. *SIAM J. Disc. Math.*, Vol. 5, No. 4, pp. 558-563, November 1992.

[KLOW94] M. Kutyłowski, K. Loryś, B. Oesterdiekhoff, R. Wanka: Fast and feasible periodic sorting networks of constant depth; Proc. 35th Annual Symposium on Foundations of Computer Science (*FOCS '94*), to appear.

[Le92] F.T. Leighton: Introduction to parallel algorithms and architectures: Arrays, Trees, Hypercubes. *Morgan Kaufmann Publisher* (1992).

[LT79] R.J. Lipton, R.E. Tarjan, "A separator theorem for planar graphs", *SIAM J. Appl. Math.* 36(2), 1979, pp. 177-189.

[Pi77] N. Pippenger: Superconcentrators. *SIAM Journal of Computing*, No. 6, 1977, pp. 298-304.

[PV76] N. Pippenger, L. Valiant: Shifting graphs and their applications. *Journal of the ACM*, Vol. 23, No. 3, July 1976, pp. 423-432.

[Ul84] Ullman, J.D., Computational Aspects of VLSI *Computer Science Press, Rockville, MD*, 1984, 495 p.

[Wa68] A. Waksman, "A permutation network", *Journal of the ACM*, vol. 15, no. 1, pp. 159-163, January 1968.

[We94] R. Werchner, personal communication.

[Wi92] D.B. Wilson: Embedding leveled hypercube algorithms into hypercubes. Proc. 4th ACM Symposium on Parallel Algorithms and Architectures (*SPAA '92*), pp. 264-270.