Sparse asymmetric connectors in communication networks

R. Ahlswede and H. Aydinian

Abstract

An (n, N, d)–connector (also referred to as "rearrangeable" or "permutation" network) is a digraph with n inputs and N outputs in which for any injective mapping of input vertices into output vertices there exists n vertex disjoint paths of length d joining each input to its corresponding output. We consider the problem of construction of $(n, N, 2)$–connectors (depth–2 connectors) with few edges when $n \ll N$. The first nontrivial construction of $(n, N, 2)$–connectors was given Baltz, Jäger and Srivastav [1]. For $n \leq \sqrt{N}$ their construction requires at most $\big(2 + o(1)\big)N\sqrt{n}$ edges. In this paper we give a simple construction of $(n, N, 2)$–connectors with less than $\big(1 + o(1)\big)N \log_2 n$ edges for any given $n \geq 16$ and $N \geq N_0(n)$. We also consider fault–tolerant depth–2 connectors under t or less arc or node failures giving optimal constructions for certain values of t.

1 Introduction

The problem of designing optimal connectors (rearangeable networks) goes back to pioneering works by Shannon [11], Slepian [12], Clos [4] and Benes [2].

Connector Problem

Given $n, N \in \mathbb{N}$ $(n \leq N)$, let $G = (V, E)$ be a digraph, whose vertex set $V = I \cup L \cup O$ is partitioned into input vertices I, link vertices L and output vertices O such that $|I| = n$, $|O| = N$ and for every injective mapping $f : I \to O$ there exist vertex disjoint paths connecting i to $f(i)$ for all $i \in I$.

Then G is called an (n, N)–**connector** (or a rearrangeable network).

The problem is to construct an (n, N)–connector with small number of edges $|E|$.

An (n, N)–connector is called an (n, N, d)–connector (or (n, N)–connector of **depth** d) if any output can be reached by any input via a path of length d.

Define

$$e(n, N, k) = \min\big\{|E| : G = (V, E) \text{ is an } (n, N, d)\text{–connector}\big\}.$$

Symmetric connectors, i.e. (n, n, d)–connectors are well studied in works by Pippenger and Yao [8], Pippenger [9], Hwang and Richards [5]. Oruc [7] was the first who constructed (n, N, d)–(asymmetric) connectors for arbitrary n and N.

Using the recent result by Radhakrishnan and Ta–Shma [10] on optimal depth–2 superconcentrators Baltz, Jäger and Srivastav [1] showed that there exist "good" (n, N, d)–connectors for $d \geq 3$. It was shown that for $n \leq \sqrt{N}$

$$e(n, N, d) = \begin{cases} O\left(N \frac{\log^2 n}{\log \log n}\right), & \text{if } d = 3 \\ O(N \log^* n), & \text{if } d \geq 4, \end{cases}$$

where $\log^* n$ is $\log \log n$ for $d = 4$ and tends to a constant for large d. Thus the most challenging case is $d = 2$.

In [1] Baltz, Jäger and Srivastav considered the following model of (n, N, d)–connectors

(c1) I and L are completely connected.

(c2) $\forall x \in O$ the degree $\deg(x) = k$.

(c3) $\forall x, x' \in O$ $(x \neq x')$
$|\Gamma(x) \cap \Gamma(x')| \leq 1$ ($\Gamma(x) \triangleq$ the set of vertices adjacent to x).

Remark 1. We note that the condition (c1) makes sense only if $L < N$, since in this model we have $|E| \geq e(n, N, 2) > nL$ and the trivial upper bound is $e(n, N, 2) \leq nN$. The condition (c2) seems to be natural since all outputs have the same significance. Condition (c3) looks somewhat unmotivated, however under this additional restriction Baltz et.al [1] found a sufficient condition which makes it possible to construct $(n, N, 2)$–connectors.

Note also that we are interested in the case where $n \ll N$. Such connectors (as was mentioned in [1]) with $k = 2$ or 3 are of interest in the design of sparse electronic switches.

The next observation directly follows from Hall's theorem (see e.g. [3]).

Connector condition [1]:

Suppose the vertex set V of a graph $G = (V, E)$ is partitioned into sets I, L and O with $|I| = n$, $|O| = N$. Moreover let I and L be completely connected, (i.e. the condition (c1) is satisfied). Then G is an $(n, N, 2)$–connector if and only if $|\Gamma(S)| \geq |S|$ for all $S \subseteq O$ with $|S| \leq n$, (that is the Hall condition is satisfied for every n–subset of O).

In [1] the following sufficient condition for a connector was established.

Lemma [1]. *Suppose a graph G satisfies the conditions (c1), (c2) and (c3). Let also $k \geq \sqrt{n}$. Then G satisfies the connector condition.*

Using this lemma three different methods of construction of $(n, N, 2)$ connectors were proposed in [1]. In all cases the number of edges required to build an $(n, N, 2)$–connector is at most

$$\big(1 + o(1)\big) \cdot \left(Nn^{\frac{1}{2}} + 2N^{\frac{1}{2}}n^{\frac{3}{2}}\right).$$

In particular this gives the folowing constructive upperbound.

Theorem BJS [1]. *For $n \leq \sqrt{N}$*

$$e(n, N, 2) \leq \big(2 + o(1)\big)N\sqrt{n}.$$

Our main result is the following.
Given n let k be an integer such that

$$\binom{2k-3}{k-1} < n \leq \binom{2k-1}{k}. \tag{1.1}$$

Then for $N > N_0(n)$ the construction method, which we call ∂–construction, defined in Section 2 gives an $(n, N, 2)$–connector $G = (V, E)$ with

$$|E| = \big(1 + o(1)\big)Nk. \tag{1.2}$$

Note that for $n \geq 16$, with (1.1) we have $n \geq 2^k$ and hence $k \leq \log_2 n$. Thus we get the following constructive upper bound in terms of n.

Theorem 1. *Given $n \geq 16$ and $N > N_0(n)$ we have*

$$e(n, N, 2) \leq \big(1 + o(1)\big)N \log_2 n. \tag{1.3}$$

In Section 3 we consider reliable or fault–tolerant connectors. An $(n, N, 2)$–connector is called t–fault–tolerant under t or less edge failures (also called $(E, t) - (n, N, 2)$–connector) if in spite of any t or less edge failures it still remains an $(n, N, 2)$–connector. Correspondingly we define a $(V, t) - (n, N, 2)$–connector if t or less vertex failures are admissible in the connector.

We show that any ∂–construction of an $(n, N, 2)$–connector gives a $(E, t) - (n, N, 2)$–connector with $t = k - 1$, where k is the degree of output vertices of the connector graph G (and defined by (1.1)). Clearly for any $(E, t) - (n, N, 2)$–connector $t < \min\{\deg(x) : x \in O\}$. Therefore for the model described above (by c1, c2, c3) the ∂–construction gives an optimal edge fault–tolerant connector. The same construction gives a $(k-1)$–fault–tolerant connector under vertex failures if $\binom{2k-3}{k-1} < n \leq \binom{2k-1}{k} - k + 1$.

2 Proof of Theorem 1

We need some notation and definitions from extremal set theory. Set $[n] = \{1, \ldots, n\}$ and $[i, j] = \{i, \ldots, j\}$ for $1 \leq i < j \in \mathbb{N}$. We also use the notation

$$2^{[n]} \triangleq \{A : A \subset [n]\},$$
$$\binom{[n]}{k} \triangleq \{A \in 2^{[n]} : |A| = k\}.$$

For a family $\mathcal{A} \subset \binom{[n]}{k}$ define its r–shadow $\partial_r(\mathcal{A})$ by

$$\partial_r(\mathcal{A}) = \left\{F \subset \binom{[n]}{k-r} : \exists A \in \mathcal{A} : F \subset A\right\}.$$

For $r = 1$ we just write $\partial(\mathcal{A})$. Let us define the colex order for the k–subsets of $\mathbb{N}$. We say that $A \prec B$ iff $\max\big((A \smallsetminus B) \cup (B \smallsetminus A)\big) \in B$.

For a positive integer m let

$$m = \binom{a_k}{k} + \binom{a_{k-1}}{k-1} + \cdots + \binom{a_t}{t} \tag{2.1}$$

with $a_k > a_{k-1} > \cdots > a_t \geq t \geq 1$.

Then we say that (2.1) is a k–cascade representation of m. The folowing is well–known (see [3]).

Proposition 1. Every positive integer m has a unique k–cascade representation (2.1) with $a_k > \cdots > a_t \geq t \geq 1$.

Next we recall the well–known Kruskal–Katona Theorem (see [3]). Let $\mathtt{L}(k, m)$ denote the first m members of k–sets of $\mathbb{N}$ in the colex ordering.

Theorem KK. *Let $\mathcal{A}$ be a family of k–sets with $|\mathcal{A}| = m$ and let m have the cascade representation (2.1).*

Then for all r, $1 \le r \le k$ holds

$$|\partial_r(\mathcal{A})| \ge |\partial_r\big(\mathtt{L}(k,m)\big)|, \tag{2.2}$$

or equivalently

$$|\partial_r(\mathcal{A})| \ge \binom{a_k}{k-r} + \cdots + \binom{a_t}{t-r}$$

(for $a < b$ $\binom{a}{b} \triangleq 0$).

We are prepared now for the proof of Theorem 1.

Proof: We are given the numbers $n < N \in \mathbb{N}$. For some integer k we have

$$\binom{2k-3}{k-1} < n \le \binom{2k-1}{k}. \tag{2.3}$$

Next consider the k–cascade representation of N

$$N = \binom{a_k}{k} + \cdots + \binom{a_t}{t} \tag{2.4}$$

and take the first N members of $\binom{\mathbb{N}}{k}$ in colex order, i.e. $\mathtt{L}(k,N)$.

It is known [3] that

$$\partial\big(\mathtt{L}(k,N)\big) = \mathtt{L}\big(k-1, |\partial\big(\mathtt{L}(k,N)\big)|\big). \tag{2.5}$$

Let us denote $|\partial\big(\mathtt{L}(k,N)\big)| = \ell$. By Theorem KK

$$\ell = \binom{a_k}{k-1} + \cdots + \binom{a_t}{t-1}. \tag{2.6}$$

∂–Construction

We consider the following incidence matrix $H(k,k-1)$ of size $\ell \times N$.

The columns of $H(k,k-1)$ are labelled by elements of $\mathtt{L}(k,N)$ and the rows by elements of $\partial\big(\mathtt{L}(k,N)\big)$. For $A \in \mathtt{L}(k,N)$ and $B \in \partial\big(\mathtt{L}(k,N)\big)$ the entry $h(A,B)$ of $H(k,k-1)$ is defined by

$$h(A,B) = \begin{cases} 1, & \text{if } A \subset B \\ 0, & \text{if } A \not\subset B. \end{cases}$$

Claim: The boolean sum of any s columns $h_{i_1},\dots,h_{i_s}$ of $H(k,k-1)$ with $s \le \binom{2k-1}{k}$ contains at least s ones.

Proof: The number of ones in $\bigvee\limits_{j=1}^{s} h_{i_j}$ indicates the size of the shadow of the k–sets $A_{i_1},\dots,A_{i_s}$, where A_{i_j} is the i_j–th member of $\mathtt{L}(k,\mathbb{N})$. On the other hand the Theorem KK implies that for any family $\mathcal{A} \subset \binom{\mathbb{N}}{k}$ with $|\mathcal{A}| \le \binom{2k-1}{k}$ we have $|\partial(\mathcal{A})| \ge |\mathcal{A}|$. $\square$

We build now our graph $G = (V, E)$, with $V = I \cup L \cup O$, $|I| = n$, $|O| = N$ as follows.

Associate the vertices of L with $\partial\big(\mathtt{L}(k, N)\big)$ and the vertices of O with $\mathtt{L}(k, N)$.

Then we join a vertex $A \in L$ with a vertex $B \in O$ if and only if $h(A, B) = 1$ in the incidence matrix $H(k, k-1)$. The vertices of I and L are considered to be completely connected. In view of the claim above the graph G is an $(n, N, 2)$–connector. Now we can easily calculate the number of edges $|E|$. The number of edges between L and O is equal to the number of ones in the incidence matrix $H(k, k-1)$. Each column of H contains exactly k ones. Therefore we have

$$|E| = n|L| + k|O| = n\ell + kN.$$

Hence with (2.3) we have

$$|E| \le \binom{2k-1}{k}\ell + kN. \tag{2.7}$$

In view of (2.4), (2.6) we finally get for $N > N_0(k)$ (suitable)

$$|E| = \big(1 + o(1)\big)Nk.$$

This completes the proof of the theorem. □

3 Fault–Tolerant Connectors

In this section we will prove that the ∂–construction of $(n, N, 2)$–connectors gives us also t–fault–tolerant connectors.

Let $G = (V, E)$ be an (n, N)–connector given by ∂–construction. Given injective mapping $f : I \to O$ let $M = f(I) \subset O$. Define $L' = \big\{x \in L : x \in \Gamma(M)\big\}$. By the construction any subset $S \subset O$, with $|S| \le n$ satisfies the Hall condition.

Consider now the subgraph $G' \subset G$ with vertex set $V' = I \cup L' \cup M$ and the edge set $E' = E_1 \cup E_2$. The edges E_2 between L' and M are the same as in G, i.e. the bipartite graph $L' \cup M, E_2)$ is an induced subgraph of G. The edges $E_2 \subset E$ between I and L' are defined as follows. For $x \in I$ and $y \in L'$

$$(x, y) \in E_1 \text{ only if } \big(\varphi(x), y\big) \in E_2.$$

In other words for any $x \in I$

$$\Gamma(x) = \Gamma\big(\varphi(x)\big) \text{ and } \Gamma(x) = k.$$

Suppose now $t \le k-1$ edge failures have occured in the network. That is t or less edges in G are deleted. From each vertex $x \in I$ we have k disjoint paths to $\varphi(x)$, and deletion of any t edges in $E_1 \cup E_2$ (independently whether these edges are in E_1 or E_2) destroys t or less paths from input vertices to $M \subset O$.

Thus with every deleted edge $(x, u) \in E_1$ the "useless" edge $\big(\varphi(x), u\big) \in E_2$ can be also considered to be deleted. This implies that w.l.o.g. we may assume that the edge failures occur only in E_2. Hence to show that the graph G is edge t–fault–tolerant, it is sufficient to verify that any subset $S \subset M$ satisfies the condition $|\Gamma(S)| \geq S$, in spite of removing t or less edges in bipartite graph $(L' \cup M, E_2)$. We distinguish two cases

Case $|S| \leq \binom{2k-1}{k} - k + 1$.

Consider the cascade representation of $|S|$. We have

$$|S| = \binom{a_k}{k} + \cdots + \binom{a_t}{t}$$

where $2k - 1 > a_k > \cdots > a_t \geq 1$.

In view of our construction and Theorem KK we have

$$|\Gamma(S)| \geq \binom{a_k}{k-1} + \cdots + \binom{a_t}{t-1}.$$

Claim 1.

$$|\Gamma(S)| \geq |S| + k - 1. \tag{3.1}$$

This can be easily seen by simple calculations.

From (3.1) we infer that removing the faulty edges with the vertices in L' which are incident to some faulty edge the resulting graph still satisfies $|\Gamma(S)| \geq |S|$ for any $S \subset M$ with $|S| \leq \binom{2k-1}{k} - k + 1$.

Case $\binom{2k-1}{k} - k + 1 < |S| \leq \binom{2k-1}{k}$.

Here we have

$$|S| = \binom{2k-2}{k} + \binom{2k-3}{k-1} + \cdots + \binom{a_1}{1}$$

(or $|S| = \binom{2k-1}{k}$), where $1 \leq a_1 \leq k - 1$.

Hence by Theorem KK

$$|\Gamma(S)| \geq \binom{2k-2}{k-1} + \binom{2k-3}{k-2} + \cdots + \binom{k}{1} + \binom{a_1}{0} = \binom{2k-1}{k}.$$

Define the size of the support of $\mathcal{A}$ by $\|\mathcal{A}\| = \left|\bigcup_{i=r}^{|S|} A_i\right|$. Clearly we have $\|\mathcal{A}\| \geq 2k - 1$.

Suppose now $\|\mathcal{A}\| > 2k - 1$. Then we claim that

$$|\partial(\mathcal{A})| \geq \binom{2k-1}{k} + k - 1. \tag{3.2}$$

This follows from a result due to Moers [6] which generalizes the Kruskal–Katona Theorem. In our special case it says that if $\|\mathcal{A}\| > 2k - 1$ then the minimum for $|\partial(\mathcal{A})|$ is achieved for $\|\mathcal{A}\| = 2k$. Moreover an optimal family achieving the minimum is

$$\mathcal{A} = L(k, |S| - 1) \cup \big(\{1, \ldots, k - 1, 2k\}\big).$$

Hence if $\|\mathcal{A}\| > 2k-1$ one has (3.2) and we come to the first case when $|\partial(\mathcal{A})| - |\mathcal{A}| \geq k-1$.

Thus is remains to consider the case $\|\mathcal{A}\| = 2k-1$. Without loss of generality we may assume that $\mathcal{A} \subset \binom{[2k-1]}{k-1}$. Suppose first that $\mathcal{A} = \binom{[2k-1]}{k-1}$ and consider the corresponding incidence matrix H of size $\binom{2k-1}{k} \times \binom{2k-1}{k}$.

Each column and each row of H contains exactly k ones. Delete now any $k-i-1$ $(1 \leq i \leq k-1)$ columns in H.

Claim 2.

(i) Any row in the resulting matrix H' contains at least $i+1$ ones.

(ii) Any two rows of H' contain together at least $k+i$ ones.

Proof. (i) is obvious. (ii) follows from the fact that any two rows have at most one 1 in common and hence two rows in H' have at least $2k-(k-i) = k+i$ ones. Thus any family $\{A_1, \dots, A_{|S|}\}$ with $|S| = \binom{2k-1}{k} - k + 1 + i$ has an incidence matrix H' with the properties (i) and (ii).

Exchange now any $t \leq k-1$ ones in H' into zeros (which corresponds to a deletion of t edges in E_2), obtaining a new matrix H^*. In view of (i) and (ii) H^* contains at most one all–zero row. Thus the sum of all columns of H^* contains at most one zero coordinate if $i < k-1$ and no zero coordinates if $i = k-1$. This implies that the corresponding bipartite graph G^* with t edge–failures has the property $|\Gamma(S)| \geq |S|$ for any set of output vertices S with $|S| \leq \binom{2k-1}{k}$.

Thus we have proved the following

Theorem 2. *A graph G given by the ∂–construction is $(k-1)$–fault–tolerant under edge–failures.*

Remark 2. It follows from the proof of the theorem that G is $(k-1)$–fault–tolerant under vertex failures if

$$\binom{2k-3}{k-1} < n \leq \binom{2k-1}{k} - k + 1.$$

Clearly if G is a $(V,t)-(n,N)$–connector then it is also an $(E,t)-(n,N)$–connector. Thus in general we can speak about t–fault–tolerant connectors under edge–vertex failures.

References

[1] A. Baltz, G. Jäger and A. Srivastav, Elementary Constructions of Sparse Asymmetric Connectors, preprint.

[2] V.E. Benes, Mathematical theory of connecting networks and telephone traffic, New York Acad. Pr., 1965.

[3] B. Bollobas, Combinatorics, Cambridge University Press, 1986.

[4] C. Clos, A study of non–blocking switching networks, Bell System Tech. J. 32, 406–424, 1953.

[5] F.K. Hwang, G.W. Richards, A two–stage network with dual partial concentrators, Network, 23, 53–58, 1992.

[6] M. Moers, A generalization of a theorem of Kruskal, Graphs and Combinatorics 1, 167–183, 1985.

[7] A.Y. Oruc, A study of permutation networks: some generalizations and tradeoffs, J. of Paral. and Distr. Comput., 359–366, 1994.

[8] N. Pippenger, A.C. Yao, On rearangeable networks with limited depth, SIAM J. Algebraic Discrete Methods, 3, 411–417, 1982.

[9] N. Pippenger, On rearangeable and nonblocking switching networks, J. Comput. System Sci., 17, 145–162, 1987.

[10] J. Radhakrishuan, A. Ta–Shma, Bound for dispensers, extractors and depth–two–superconcentrators, SIAM J. Discrete Math., 13, 2–24, 2000.

[11] C.E. Shannon, Memmory Requirements in a Telephone Exchange, Bell System Tech. J., 343–349, 1950.

[12] D. Slepian, Two Theorems on a Particular Crossbar Switching Network, unpublished manuscript, 1952.