On the Total$_k$-Diameter of Connection Networks

Yefim Dinitz * Tamar Eilam Shlomo Moran †
Shmuel Zaks ‡

September 20, 1998

Department of Computer Science
Technion, Haifa, Israel

Abstract

We study *connection networks* in which certain pairs of nodes have to be connected by k edge-disjoint paths, and study bounds for the minimal sum of lengths of such k paths. We define the related notions of *total$_k$-distance* for a pair of nodes and *total$_k$-diameter* of a connection network, and study the value $TD_k(d)$ which is the maximal such total$_k$-diameter of a network with diameter d. These notions have applications in fault-tolerant routing problems, in ATM networks, and in compact routing in networks. We prove an upper bound on $TD_k(d)$ and a lower bound on the growth of $TD_k(d)$ as functions of k and d; those bounds are tight, $\theta(d^k)$, when k is fixed. Specifically, we prove that $TD_k(d) \leq 2^{k-1}d^k$, with the exceptions $TD_2(1) = 3$, $TD_3(1) = 5$, and that for every $k, d_0 > 0$, there exists (a) an integer $d \geq d_0$ such that $TD_k(d) \geq d^k/k^k$, and (b) a k-connected simple graph G with diameter d such that $d \geq d_0$, and whose total$_k$-diameter is at least $(d-2)^k/k^k$.

---

*Up to 1990, E. A. Dinic, Moscow.

†Research supported by the Bernard Elkin Chair for Computer Science and by US-Israel BSF grant 95-00238

‡Contact author; *Tel*: (+972)-4-8294368, *Fax*: (+972)-4-8221128, *Email*: zaks@cs.technion.ac.il

1 Introduction

A connection network consists of a network (modeled by a graph) and a set of pairs of its nodes between which a reliable communication is required. This new concept is a generalization of a communication network, in the sense that some pairs of nodes are distinguished. Define the total$_k$-distance for a pair of nodes as the minimal sum of the lengths of k edge-disjoint paths between them. This quantity corresponds to the cost of a reliable routing of a message between those nodes in a network with less than k undetected faults of links. We give upper and lower bounds on the worst case value of this quantity in a connection network, as a function of its diameter, which are tight up to a multiplicative constant when k is fixed. Besides the "natural" implications of our results in the field of fault tolerant routing, our work has implications also in the field of ATM networks and compact routing methods.

Previous works in this field ([LSX95, HL94, MP88]) concentrate on node connectivity. They define a *k-container* for s and t as a set of k node-disjoint paths between s and t, and its *length* as the maximal length of a path in it; the *k-distance* for two nodes is the minimal length of a k-container for them, and the *k-diameter* is the maximal k-distance among all pairs of nodes in the graph. Upper and lower bounds for the k-diameter of k-regular k-connected graphs are given in [HL94] (as a function of k and the number of nodes in the graph). In [MP88], a family of graphs, termed *flip-trees*, is introduced; it is shown that in a flip-tree with degree d and diameter $2l-1$ there exists a container of width d and length $\leq 2l+1$ for any two nodes. The 2-diameter of de-Bruijn graphs is discussed in [LSX95]; it is shown that the 2-diameter of the de-Bruijn graph of dimension n is exactly n.

1.1 Definitions and Notations

We define a new concept of a *connection network* $S = (G, \zeta)$ which consists of (i) a graph $G = (V, E)$, where V is the set of nodes and E is the set of edges, and (ii) a *connection demand system* ζ, which is a set of pairs of distinct nodes in V. This concept is a generalization of a communication network in the sense that it models a network in which a reliable communication is required for every pair of nodes in ζ but not necessarily for every pair of nodes in the underlying graph (see [D94, M89]). (Note that when ζ includes all pairs of nodes in G, then our model coincides with the more familiar model of a communication network.) In this paper we restrict ourselves to the undirected case, i.e, when the graph G is undirected and the pairs in ζ are unordered.

Let s and t be two nodes in G. An *edge-container* for s and t is a set of edge-disjoint paths between s and t. The *width* of an edge-container is the number of paths in it. We abbreviate an edge-container of width k to a *k-container*. The *total-length* of a k-container $\mathcal{A}$, denoted by $|\mathcal{A}|$, is the sum of the lengths of all of its paths. A k-container for s and t is *minimum* if there is no k-container for s and t with a smaller total-length. The *total$_k$-distance for s and t*, $td_k(s,t)$, is the total-length of a minimum k-container for them if one exists, otherwise $td_k(s,t) = \infty$ (evidently, $td_k(s,t) = td_k(t,s)$). Two nodes s and t are *k-edge-connected* if the deletion of any set of $k-1$ edges from G does not disconnect them. It is well known (see, e.g., [FF62]) that s and t are k-edge-connected if and only if there exist k-edge-disjoint paths between them; it follows that $td_k(s,t) < \infty$ if and only if s and t are k-edge-connected.

We note that the triangle inequality holds for total$_k$-distances as for ordinary distances. A sketch of the proof is as follows. The union of any two minimum k-containers for s and u and for u and t is a subgraph, say, G', where s and t are k-connected. Indeed, neither s and u nor u and t can be disconnected in G' by at most $k-1$ edges. Hence, this is true for s and t as well (since otherwise such an edge subset would disconnect u from either s or t). Therefore, there exists an s,t-container in G'; evidently, it has a total-length of at most the sum of the total-lengths of the two above containers. Interestingly, the triangle inequality does not hold for the "k-distance" of [HL94, LSX95, MP88].

The *total$_k$-diameter*, $td_k(S)$, of a connection network $S = (G, \zeta)$, is the maximum $td_k(s,t)$ over all pairs of nodes in ζ. A connection network S is *k-edge-connected* if s and t are k-edge-connected for every pair $\{s,t\} \in \zeta$. (Note that $td_k(S) < \infty$ if and only if S is k-edge-connected.) We denote by $TD_k(d)$ the maximal total$_k$-diameter over all k-connected connection networks whose underlying graph has diameter d.

We consider two special cases. The first one is when the connection network (G, ζ) is such that ζ includes all pairs of distinct nodes in the graph; in this case we refer to the total$_k$-diameter of the graph G, denoted $td_k(G)$. The second one is when $\zeta = \{\{s,t\}\}$ (i.e., it includes only one unordered pair of nodes); in this case we abbreviate the notation $(G, \{\{s,t\}\})$ to (G, s, t).

There are two well-known routing paradigms. The first is the *one-to-one routing* paradigm ([CC95, DMS91, HL94, JLD91, LSX95, MP88]), which finds disjoint paths between two nodes in the network; we concentrate on this paradigm. The second paradigm is the *one-to-many routing* paradigm ([DMS91, HLN89, MS90, R89]) which aims to find k disjoint paths between one node and each of the nodes in a given set of k nodes. We extend our upper and lower bound results also for this case.

We summarize the differences in the model and definitions between our work and the works mentioned above as follows.

1. We concentrate on edge-connectivity (rather than node-connectivity). In the sequel k-*connected* always means k-edge-connected and *disjoint* always means edge-disjoint.

2. We consider the total-length of a k-container, i.e., the sum of lengths of the paths in it, rather than its length (which is the maximal length of a path in it). Note that the total-length and the length of any 1-container are equal.

3. Previous works examined k-node-connected graphs. We generalize the term "graph" (or "network") to "connection network" and examine k-connected connection networks (G, ζ) for which the underlying graph G is not necessarily k-connected.

4. We examine the relations between $TD_k(d)$ and the diameter d of the underlying graph G, whereas other works discuss the relationship between the k-diameter and other parameters such as the degree of nodes and the number of nodes.

1.2 Our Results

We prove an upper bound on $TD_k(d)$ and a lower bound on the growth of $TD_k(d)$ as functions of k and d, which are tight when k is fixed. Specifically, we prove:

1. $TD_k(d) \leq 2^{k-1}d^k$, with the exceptions $TD_2(1) = 3$, $TD_3(1) = 5$.

2. For every $k, d_0 > 0$, there exists an integer $d \geq d_0$ such that $TD_k(d) \geq d^k/k^k$.

2'. For every $k, d_0 > 0$, there exists a k-connected simple graph G [1] with diameter d such that $d \geq d_0$ and $td_k(G) \geq (d-2)^k/k^k$.

[1] A graph is termed *simple* if it has no multiple edges and no loops.

(It is worth noting that our lower bounds are reached on connection networks of type (G, s, t); hence, evidently, these bounds hold for any structural type of the connection demand system ζ.)

These results are extended also to the *one-to-many* routing paradigm [DMS91, HLN89, MS90, R89].

Let k be fixed (but arbitrary). Then **2** above implies that there are infinitely many values of d such that $TD_k(d) \geq d^k/k^k$. On the other hand, by **1** above, $TD_k(d) \leq 2^{k-1}d^k+1$. Thus, for any fixed k, the upper and lower bounds on the worst case growth of TD_k differ by a multiplicative constant and are both of order d^k.

1.3 Applications and Related Results

One possible application of our results is in the field of fault-tolerant routing methods. Consider a network $N = (V, E)$ in which some edges may fail. We say that the communication between two nodes s and t is $(k-1)$-reliable in N if there exists a routing scheme, termed $(k-1)$-*reliable*, that guarantees that in the presence of at most $k-1$ edge failures, a message sent from s to t will eventually arrive. Assuming that edge-failures cannot be detected, it is not hard to see that a $(k-1)$-reliable routing scheme from s to t requires that—in the case that all edges are non-faulty—each message sent from s to t is delivered on a set of edges which forms a k-connected subgraph, i.e., includes a k-container for s and t. On the other hand, if a k-container $\mathcal{A}$ for s and t exists, then a $(k-1)$-reliable routing scheme between s and t is achieved by the algorithm that delivers k copies of any message sent from s to t over the k edge-disjoint paths connecting s and t in $\mathcal{A}$.

Let the cost of a routing scheme between s and t be the maximal number of hops performed by a message sent from s to t ("hop" denotes the delivery of a message from a node to its neighbor over an edge connecting them). Then the discussion above implies that the cost of a most efficient $(k-1)$-reliable routing scheme between s and t is equal to the total$_k$-distance between s and t.

The definition of $(k-1)$-reliable communication is easily generalized to an arbitrary set of distinct pairs of nodes ζ. Specifically, a network is called $(k-1)$-*reliable w.r.t.* ζ if and only if it is $(k-1)$-reliable for each pair $\{s, t\} \in \zeta$. By the discussion above, $N = (V, E)$ is $(k-1)$-reliable w.r.t. ζ if and only if the connection network $S = (N, \zeta)$ is k-edge-connected, and the maximal cost of a delivery of a message between a pair of nodes in ζ, when using a most efficient $(k-1)$-reliable routing scheme, is equal to $td_k(S)$.

A few special cases are worth consideration. The familiar one is when a reliable communication is required between every pair of nodes in the network, which is the case where ζ is the set of all pairs of distinct nodes in the network. Another special case, introduced and discussed in [D94], is when a subset of important nodes ("terminals"), is distinguished, and a reliable communication is required for all pairs of terminals. Following [D94], we argue that such a reliability model is more realistic than the usual "all-to-all" model. Yet another case is when a reliable communication is required between one "special" node and every one in a set of other nodes, which does not necessarily include all the nodes of the network.

One conclusion of our result, which holds for each of the special cases considered above, is that a small diameter of the underlying graph of a k-connected connection network does not guarantee, in general, existence of an efficient fault-tolerant routing scheme for it. In fact, the cost of an optimal $(k-1)$-reliable routing scheme could be as large as the diameter to the power k but not larger than that (ignoring factors depending on k only).

The fault-tolerant routing application above corresponds to the one-to-one routing paradigm in which the goal is to construct disjoint paths between pairs of given nodes. Another well-known paradigm is the one-to-many routing that aims to construct disjoint paths between a given node and each of the nodes in a given set. To extend our model for the one-to-many paradigm, given a node $s \in V$ and a set of nodes $T \subset V$ of size k (which does not include s), we define $td_k(s,T)$ as the minimal sum of lengths of a set of k edge-disjoint paths between s and each node in T. The quantity $td_k(s,T)$ could be seen as the cost of a reliable routing of a message in the presence of at most $k-1$ undetected faults, where the interpretation of "reliable communication" is a communication that guarantees that at least one node in T will eventually receive the message. We extend our results to similar upper and lower bounds for $td_k(s,T)$.

Observe that our general worst case bounds are in contrast with previous results for special settings. For example, for one-to-n routing in the n-dimensional hypercube (whose diameter is n), there always exists a set of node-disjoint paths of length at most $n+1$ each (see [R89]). It could be interesting to study the worst case behavior of the total$_k$-diameter for other network topologies which are used in practice.

In a related work in the area of hop-congestion trade-off for ATM networks, [KKP95] studies the following problem, termed the *augmented path* *problem*: "Determine the minimal diameter $D_N(k)$, of an augmented path of length N and cutwidth k." (The definitions of augmented path and cutwidth will be given in Section 5). In [KKP95] upper and lower bounds are given for $D_N(k)$. We show that by applying our upper bound on the total$_k$-distance of two nodes we can easily derive the same (and even slightly better) lower bound on $D_N(k)$ as a special case. The construction which we use for proving the lower bound on the total$_k$-distance is actually similar to the construction used in [KKP95] for proving the upper bound for $D_N(k)$.

Another related work is in the field of space-efficient routing methods (termed *compact routing*). In [EMZ96] we show a strong connection between the efficiency of *linear interval routing* (which is one such method) for a network N and the total$_2$-diameter of N. Specifically, we present a lower bound on the efficiency of linear interval routing for a network in terms of its total$_2$-diameter.

The rest of the paper is organized as follows: In Section 2 we present and prove the lower bound. The upper bound is proved in Section 3. Our results are extended to the one-to-many paradigm in Section 4. In Section 5 we discuss the connection to the augmented path problem.

2 Lower Bound

In this section we discuss lower bounds on the growth of $TD_k(d)$, specifically, we prove the following theorem.

Theorem 2.1

a. *For every $k, d_0 > 0$, there exists a graph G with diameter d, $d_0 \le d \le d_0 + 3k$, and a k-connected connection network $S = (G, s, t)$ such that*

$$td_k(S) \ge d^k/k^k;$$

thus, $TD_k(d) \ge d^k/k^k$. Hence, for any fixed k, $TD_k(d) = \Omega(d^k)$.

b. *For every $k, d_0 > 0$, there exists a k-connected simple graph G with diameter d, where $d_0 \le d \le d_0 + 3k + 2$, such that*

$$td_k(G) \ge (d-2)^k/k^k.$$

For any $k > 0$, we construct (Section 2.1) an infinite series of graphs $G_x(k)$, where x is any positive integer. We show that for every positive diameter d_0, there exists a graph $G_{x_0}(k)$, with diameter d, $d_0 \leq d \leq d_0 + 3k$, and a pair of its nodes s and t, such that the connection network $S_{x_0}(k) = (G_{x_0}(k), s, t)$ satisfies $d^k/k^k \leq td_k(S_{x_0}(k)) < \infty$. The graph $G_{x_0}(k)$ is 2-connected for $k > 1$. Then we transform $G_{x_0}(k)$ (Section 2.2) to a k-connected simple graph $G'_{x_0}(k)$ with diameter $d' = d + O(1)$ such that $(d'-2)^k/k^k \leq td_k(G'_{x_0}(k)) < \infty$.

2.1 The Connection Network $S_x(k)$

Let $x \geq 2$ be an arbitrary positive integer. We construct a family of connection networks $\{S_x(k) | k \in \mathcal{N}^+\}$, where $S_x(k) = (G_x(k), s, t)$. The nodes s and t are termed the *endpoints* of $G_x(k)$. The graph $G_x(1)$ is simply a path of length x, and the nodes s and t are its endpoints. The graph $G_x(k)$ is constructed from x copies of $G_x(k-1)$; the ith copy, $1 \leq i \leq x$, is denoted $G_x^i(k-1)$ and its endpoints are s^i and t^i. To each copy i we add an edge $\{s^i, t^i\}$ termed a *k-edge*, and we combine the copies by identifying the node s^i with the node t^{i-1}, for every $2 \leq i \leq x$. The endpoints s and t of the new graph $G_x(k)$ are the nodes s^1 and t^x, respectively. Figure 1(a) demonstrates the construction. Figure 1(b) shows the graphs $G_3(1), G_3(2)$, and $G_3(3)$. It is worthwhile to mention that a similar construction is also used in [KKP95].

The following three propositions form the proof of Theorem 2.1(a).

Proposition 2.2 *The diameter $d_x(k)$ of $G_x(k)$ satisfies: $d_x(k) \leq kx$.*

Proof: We prove the following two facts by induction on k: (a) For any node u in $G_x(k)$, its distance from either s or t is at most $k\lfloor x/2 \rfloor$ and, (b) $d_x(k) \leq kx$. The basis ($k = 1$) is trivial. Now consider the graph $G_x(k)$, for $k > 1$. By the induction hypothesis (a), the distance from any node u in $G_x^i(k-1)$ to s^i or to t^i is at most $(k-1)\lfloor x/2 \rfloor$. Consider the following path between any pair of nodes u in $G_x^i(k-1)$ and v in $G_x^j(k-1)$. The path starts in u, continues in the shortest way to the nearest endpoint of $G_x^i(k-1)$ (s^i or t^i), then continues on k-edges to the endpoint of $G_x^j(k-1)$ which is the nearest to v, and then continues in the shortest way to v. Since there are exactly x k-edges, and by the induction hypothesis (a), the length of this path is at most $(k-1)\lfloor x/2 \rfloor + x + (k-1)\lfloor x/2 \rfloor \leq kx$ and thus, $d_x(k) \leq kx$. By using similar arguments for any node u, one of the distances to s or to t is at most $(k-1)\lfloor x/2 \rfloor + \lfloor x/2 \rfloor \leq k\lfloor x/2 \rfloor$. □

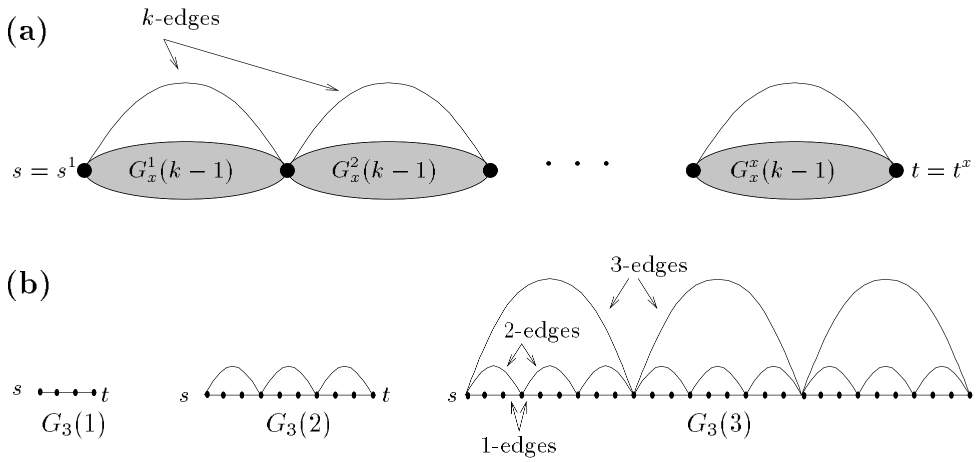

Figure 1: **(a)** A recursive description of the graph $G_x(k)$. **(b)** The graphs $G_3(1)$, $G_3(2)$, and $G_3(3)$.

Proposition 2.3 *The diameter $d_x(k)$ of $G_x(k)$ satisfies $d_x(k) \geq k(x-2)$.*

Proof: We first prove the following claim: there exists a node at distance at least $k\lceil (x-2)/2 \rceil$ from both s and t (such a node is termed *a middle node*). The basis ($k=1$) is trivial. Consider the $\lceil x/2 \rceil$-th copy $G_x^{\lceil x/2 \rceil}(k-1)$ of $G_x(k-1)$ in $G_x(k)$. By the induction hypothesis, there is a middle node u in $G_x^{\lceil x/2 \rceil}(k-1)$ with distances at least $(k-1)\lceil (x-2)/2 \rceil$ from both $s^{\lceil x/2 \rceil}$ and $t^{\lceil x/2 \rceil}$. From the construction of the graph $G_x(k)$ and by the induction hypothesis, it is clear that the distance of u from s and t is at least $(k-1)\lceil (x-2)/2 \rceil + \lceil (x-2)/2 \rceil \geq k\lceil (x-2)/2 \rceil$. Thus, u is a middle node in $G_x(k)$.

Now consider the middle nodes in $G_x^1(k-1)$ and $G_x^x(k-1)$. Since t^1 and s^x are separation nodes in $G_x(k)$ at distance $x-2$ and by the above claim, the distance between the two middle nodes is at least $(k-1)\lceil (x-2)/2 \rceil + (x-2) + (k-1)\lceil (x-2)/2 \rceil \geq k(x-2)$ and thus $d_x(k) \geq k(x-2)$. $\Box$

Proposition 2.4 *The connection network $S_x(k)$ satisfies: $d^k/k^k \leq td_k(S_x(k)) < \infty$, where d is the diameter of $G_x(k)$.*

Proof: Consider the graph $G_x(k)$ and its endpoints s and t. We first show that there exists a k-container for s and t, thus $S_x(k)$ is k-connected (so, $td_k(S_x(k)) < \infty$), and then we show that the total-length of any k-container for s and t is at least d^k/k^k. Let the path P_j between s and t, $1 \leq j \leq k$, be the one consisting exactly of the set of all j-edges (i.e., the edges which we add when constructing $G_x(j)$ from x copies of $G_x(j-1)$). The paths $P_1, P_2, ..., P_k$ are edge-disjoint and, clearly, they form a k-container for s and t.

Now, since every edge in the graph is in a cut of size k between s and t, it must be contained in any set of k edge-disjoint paths between s and t (by [FF62]); thus, any k-container for them contains all the edges of the graph. Since $G_x(k)$ contains x^k 1-edges, $td_k(S_x(k)) \geq x^k = (kx)^k/k^k \geq d^k/k^k$ (the last inequality is by Proposition 2.2). $\Box$

Remark 2.5 Note that we actually proved in Proposition 2.4 that the number of 1-edges in $G_x(k)$ is at least d^k/k^k. (We will use this observation in Section 5.)

Proof [of Theorem 1(a)]: Given any positive constant d_0, consider the connection network $S_x(k)$, where $x = \lceil (d_0 + 2k)/k \rceil$. By Proposition 2.2, the diameter d of $G_x(k)$ satisfies: $d \leq kx = k\lceil (d_0 + 2k)/k \rceil \leq d_0 + 3k$. By Proposition 2.3, $d \geq k(x-2) \geq k\lceil (d_0 + 2k)/k - 2 \rceil \geq d_0$. Thus, $d_0 \leq d \leq d_0 + 3k$ and, by Proposition 2.4, $td_k(S_x(k)) \geq d^k/k^k$. $\Box$

2.2 The k-Connected Graph $G'_x(k)$

Note that if we just replace every 1-edge (i.e., an edge in a copy of $G_x(1)$) by k parallel edges, then we get a (non-simple) k-connected graph with the same diameter as that of $G_x(k)$ and the total$_k$-diameter at least as large as $td_k(S_x(k))$. We show a k-connected simple graph that satisfies a slightly weaker inequality.

For any $k, d_0 > 0$, we construct a k-connected graph $G'_x(k)$ with diameter $d' \geq d_0$, for which $td_k(G') \geq (d'-2)^k/k^k$. Consider the graph $G_x(k)$. To construct $G'_x(k)$ from $G_x(k)$, we replace each 1-edge in $G_x(k)$ by a complete graph of size $k+1$, i.e., for each such edge $\{v_i, v_{i+1}\}$ we add the nodes $v_i^1, ..., v_i^{k-1}$, an edge between each v_i^m and v_i^n, and the edges $\{v_i, v_i^m\}$ and $\{v_{i+1}, v_i^m\}$ $1 \leq m \neq n \leq k-1$. The graph $G_x(k)$ is a subgraph of $G'_x(k)$; we term the nodes and edges of $G_x(k)$ *original*, and the new nodes and edges *new*.

Note that $G'_x(k)$ is k-connected; even if we remove all the j-edges, $j \geq 2$, then we get x^k copies of a $(k+1)$-clique such that the ith copy and the $(i+1)$st copy have one node in common, for every $1 \leq i \leq x^k - 1$. Clearly, such a graph is k-connected.

In the construction of $G'_x(k)$ from $G_x(k)$, we didn't remove any edge, thus the distance between any two original nodes did not increase. Moreover, any new node is at distance 1 from an original node. It follows that the diameter of $G'_x(k)$, d', satisfies: $d' \leq d + 2 \leq kx + 2$.

Proposition 2.3 holds also for $G'_x(k)$ (i.e., $d' \geq k(x-2)$), since a shortest path between two original nodes never passes through a new node and there are no new edges between original nodes, thus the distances between any two original nodes did not decrease.

Finally, it is easy to see that the same k-container for s and t as described in the proof of Proposition 2.4 is a minimum k-container for s and t also in $G'_x(k)$ (since a passage through a new node never shortens the distance). It follows that $td_k(G') \geq x^k \geq (xk)^k/k^k \geq (d'-2)^k/k^k$.

Given any positive integer d_0, we consider the graph $G'_x(k)$, where $x = \lceil (d_0 + 2k)/k \rceil$. We have shown that the diameter d' of $G'_x(k)$ satisfies $d_0 \leq k(x-2) \leq d' \leq kx + 2 \leq d_0 + 3k + 2$ and $(d'-2)^k/k^k \leq td_k(G'_x(k)) < \infty$. This completes the proof of Theorem 2.1(b).

3 Upper Bound

The upper bound is stated in the following theorem.

Theorem 3.1 *Every k-connected system $S = (G, \zeta)$, the underlying graph G of which has diameter d, satisfies:*

$$td_k(S) \leq 2^{k-1}d^k,$$

except for the cases $d = 1$, $k = 2, 3$, where

$$td_k(S) \leq 2^{k-1}d^k + 1.$$

Thus, for any fixed k, $td_k(S) = O(d^k)$.

Clearly, this theorem is implied by the following lemma, which is proven in the next subsection.

Lemma 3.2 *The total$_k$-distance of any two k-connected nodes in an arbitrary graph with diameter d is at most $2^{k-1}d^k$, except for the cases $d = 1$, $k = 2$ and $d = 1$, $k = 3$, when it can exceed this bound by one, i.e., be equal to 3 and 5, respectively.*

3.1 Proof of Lemma 3.2

The mathematical foundations of the considered field are closely related to network flow theory. As mentioned in Section 1.1, the total$_k$-distance $td_k(s,t)$ of two nodes s and t is finite if and only if s and t are k-edge-connected. For any graph $G = (V, E)$ and its two nodes s and t, let $\vec{G}(s,t) = (V, \vec{E}, s, t)$ be the directed flow network with source s, sink t, and $\vec{E}$ defined as follows: each edge $\{u, v\}$ in E defines a pair of anti-parallel arcs (u, v) and (v, u) in $\vec{E}$ with capacities 1 (we use the term "arc" for a directed edge). In what follows, we mean by "flow" a flow in $\vec{G}(s,t)$ with flows in arcs equal to 0 or 1 only and without flows both equal to 1 in any two anti-parallel arcs. We will shorten the notation $\vec{G}(s,t)$ to $\vec{G}$.

It is well known (see [FF62]) that s and t are k-connected in G (i.e., a k-container for them exists) if and only if there is a flow with value k from s to t in $\vec{G}$. For a flow f, let us denote by $\#f$ the number of edges of G such that f is 1 in any of the two arcs defined by it; henceforth, we call such edges "used by f". It can be easily shown that for any k-container for s and t there exists a flow from s to t with value k, which uses exactly the edges of that k-container, and that for any flow f from s to t with value k there exists a k-container for s and t, all of whose edges are used by f. It follows that $td_k(s,t)$ is equal to the number of edges used by a flow f with minimal $\#f$ among all flows from s to t with value k. [2]

Consider any two k-connected nodes s and t in a graph G with diameter d. By the above, the value of a maximum flow is at least k. For our proof it is sufficient to show that there exists a flow f_k with value k with $\#f_k$ bounded as in the statement of the lemma. We establish the existence of such a flow by induction. Let $\delta(x,y)$ denote the distance between nodes x and y in G. Trivially, a flow f_1 as required can be defined as the unit flow along a shortest path from s to t, with $\delta(s,t) \leq d$ used edges. For the induction step we need the following lemma.

[2] Such a flow, in the usual network flow terminology, is a min-cost (s,t)-flow with value k, where the cost 1 is assigned to each arc of $\vec{G}$.

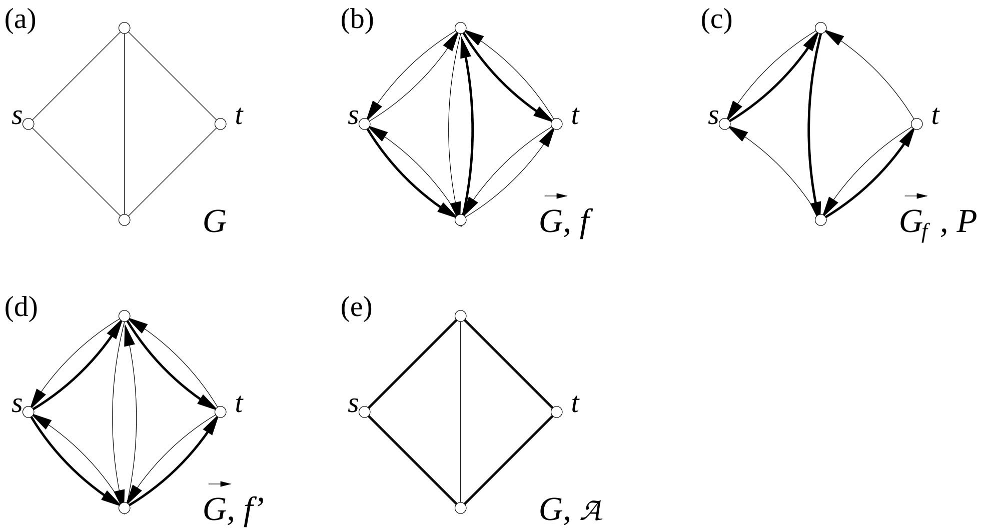

Figure 2: **(a)** A graph G and 2-connected vertices s and t in it. **(b)** The network $\vec{G}$ and a flow f with value 1 in it (the arcs with flow 1 are shown bold). **(c)** The residual network $\vec{G}_f$ and an augmenting path P in it (its arcs are shown bold). **(d)** The flow f' in $\vec{G}$ obtained via augmenting f by P. **(e)** The corresponding 2-container $\mathcal{A}$ for s and t in G (its edges are shown bold).

Lemma 3.3 *For any non-maximum flow f with value at least 1, there exists a flow f' with value greater by one with $\#f' \leq \max\{2d \cdot \#f,\ \#f + 2d\}$.*

Proof: As is widely known (see, for example, [CLR90]), to transform any non-maximum flow f in $\vec{G}$ into a flow with value greater by one, it is sufficient to find a path P from s to t in the *residual network* $\vec{G}_f = (V, \vec{E}_f)$ (an *augmenting path*[3]), and to *augment* f by the unit flow along P (henceforth, for simplicity, "to augment f by P"). Details are as follows (for illustration see Figure 2(b-d)). The arcs in $\vec{E}_f$ are defined as the arcs a in $\vec{E}$ for which $f(a) = 0$; for our purposes, all edge capacities in $\vec{G}_f$ maybe assumed 1. The augmentation of a flow f in $\vec{G}$ by a flow g in $\vec{G}_f$ is almost the usual vector addition, with the unique exception, in our case: if for some two anti-parallel arcs a and $\bar{a}$, $f(a) = 1$, $f(\bar{a}) = 0$, $g(a) = 0$, $g(\bar{a}) = 1$, then the sum flow in a and $\bar{a}$ is set to be 0 and 0, respectively (instead of 1 and 1).

Let us denote the length of a path Q by $|Q|$. Let a flow f' be the flow f augmented by a path P. It is straightforward to see that $\#f'$ is at

[3]To distinguish from the notion of *augmented path* of [KKP95].

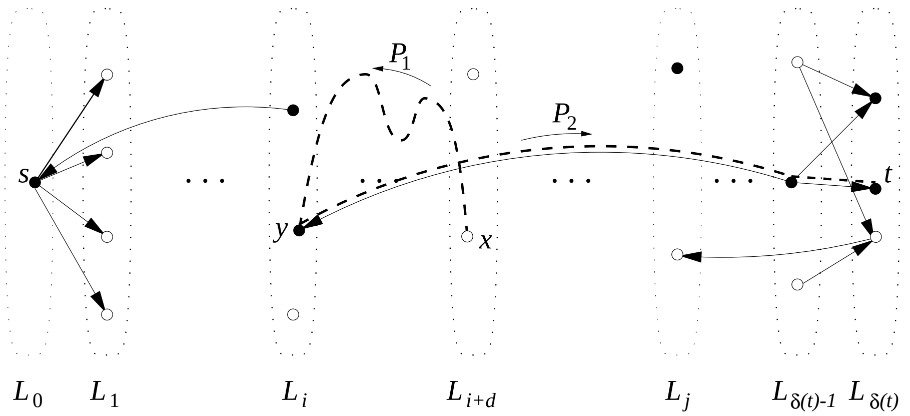

Figure 3: An example of the distribution of used vertices in a network $\vec{G}_f$ and of the path $P = P_1 + P_2$, where $y \in L_i$ and $|P_2| = 2$. (Vertices and arcs are shown partly, layers are encircled by dotted lines, used vertices are shown black, and paths P_1 and P_2 in G are shown by dashed lines.)

most $\#f + |P|$ (for illustration see Figure 2, where $\#f = 3$, $|P| = 3$, and $\#f' = 4$). Thus, it is sufficient to prove that for any non-maximum flow f with value at least 1, the minimum length of an augmenting path is at most $\max\{(2d-1)\cdot\#f,\ 2d\}$.

Let us denote by $\delta(v)$ the distance from s to any node v in $\vec{G}_f$ (or ∞, when v is not reachable from s). By non-maximality of f, $\delta(t)$ is finite, and we prove that it satisfies the above bound. Let us consider, in the BFS-structure of $\vec{G}_f$ with the origin s, the *layers* $L_i = \{v : \delta(v) = i\}$, $1 \le i \le \delta(t)$ (see Figure 3). Clearly, there is no arc of $\vec{G}_f$ from L_i to L_j with $j \ge i+2$.

A node v is termed *used by f* if there is an edge incident to v used by f; clearly, s and t are used by f. Let us denote by N the number of nodes used by f; clearly, $N \le \#f + 1$.

We first prove that $\delta(t)$ is at most $2d(N-1) \le 2d\cdot\#f$ (observe that this implies the weaker upper bounds $(2d+1)\cdot\#f$ for $\#f'$ and thus, by induction, $d(2d+1)^{k-1}$ for $td_k(s,t)$). We say that a layer is *marked* if it contains a node used by f. Assuming, to the contrary, that $\delta(t) \ge 2d(N-1)+1$, we obtain that there exist two marked layers L_i and L_j, $j \ge i+2d+1$, and $i, j \le \delta(t)$, with no marked layers between them. Let us consider a node x in layer L_{i+d}; by assumption, x is not used by f. Let $P(x)$ be a shortest path from x to t in the original graph G; its length is at most d (the diameter of G). Let y be the first node on $P(x)$ used by f, and let P_1, P_2 be the sub-paths of $P(x)$

from x to y and from y to t, respectively (see Figure 3).

Observe that, for any edge $\{u, v\}$ in P_1, the flow f cannot be 1 in any of the arcs (u, v) and (v, u) (otherwise, y would not be the first used vertex in P_1 and, thus, in P); therefore, for each such edge, both these arcs belong to $\vec{G}_f$. We obtain now that the entire P_1 is contained in $\vec{G}_f$, hence, $\delta(y)$ is finite and $|P_1| \geq |\delta(x) - \delta(y)|$.

By the choice of x and y, $\delta(x) \neq \delta(y)$. Let us consider two cases.

Case 1: $\delta(y) < \delta(x)$. By our assumptions, $\delta(y) \leq i$ and so y cannot coincide with t. We obtain $|P(x)| = |P_1| + |P_2| \geq (\delta(x) - \delta(y)) + 1 \geq ((i + d) - i) + 1 = d + 1$, a contradiction.

Case 2: $\delta(y) > \delta(x)$. Now $\delta(y) \geq j \geq i + 2d + 1$, hence $|P(x)| \geq |P_1| \geq \delta(y) - \delta(x) \geq (i + 2d + 1) - (i + d) = d + 1$, a contradiction.

It follows that $\delta(t) \leq 2d \cdot \#f$, as required.

Let us enhance this method to achieve the bound of the lemma. For this, we show that if $\delta(t) > (2d - 1) \cdot \#f$, then $\delta(t) = 2d$. The above inequality implies that $\delta(t) > (2d - 1)(N - 1)$, and hence that there exist two marked layers L_i and L_j with no marked layers between them, with $j \geq i + 2d$ and $i, j \leq \delta(t)$.

As before, let x be any vertex in layer $i + d$ and $P(x)$ a shortest path in G from x to t. Going along the same lines, we obtain that the only case that does not produce the contradiction $|P(x)| \geq d + 1$ is when t lies in L_j and $j = i + 2d$, i.e. $\delta(t) = i + 2d$.

Let now $Q(x)$ be a shortest path from x to s in G. Similar considerations as above, using the fact that x is in layer $L_{j-d} = L_{i+d}$, show that the only case without the contradiction $|Q(x)| \geq d + 1$ is when s lies in L_i, i.e., $i = 0$. Hence, our assumption $\delta(t) > (2d-1)(N-1)$ leads to $\delta(t) = j = 0 + 2d = 2d$, as claimed. □

We now return to the inductive proof of Lemma 3.2. We prove that f_k, the flow with value k constructed from f_1 by a sequence of augmentations as in the proof of Lemma 3.3, satisfies the inequality $\#f_k \leq 2^{k-1}d^k$, with exceptions as in Lemma 3.2. We assume the inequality holds for $k = l$, i.e., $\#f_l \leq 2^{l-1}d^l$, and prove it for $k = l+1$, $l \geq 1$. Let us denote $2^l d^{l+1} - 2^{l-1}d^l = (2d - 1)2^{l-1}d^l$ by Δ_l. Then, by the induction hypothesis and Lemma 3.3, $\#f_{l+1} \leq \#f_l + \max\{(2d - 1) \cdot \#f_l, 2d\} \leq \#f_l + \max\{\Delta_l, 2d\}$. Assume first that d is at least 2. Then $\Delta_l = (2d - 1)2^{l-1}d^l \geq 3 \cdot 1 \cdot d = 3d > 2d$. Hence, in this case $\#f_{l+1} \leq \#f_l + \Delta_l \leq 2^l d^{l+1}$, as required.

Assume now that $d = 1$, i.e., G is a complete graph with, maybe, multiple edges. In this case, $\#f_1$ is equal to 1. In the considered sequential construction of f_k, let δ_l denote the length of the l-th augmenting path (i.e., $\#f_{l+1} \leq \#f_l + \delta_l$). Since in G there are at least $|V| - 2$ edge-disjoint paths of length 2 between s and t, the first $|V| - 2$ shortest augmenting paths must be each of length at most 2. Clearly, the length of any augmenting path is at most $|V| - 1$. Hence, for $|V| = 2$, $\delta_l = 1$ for all l. For $|V| = 3$, $\delta_l \leq 2$ for all l. For $|V| = 4$, $\delta_1, \delta_2 \leq 2$ and $\delta_l \leq 3$ for $l > 2$. For $|V| > 4, \delta_l \leq 2$ for $l = 1, 2, 3$. Thus, we get that for any graph with diameter 1, $\#f_1 = 1$, $\#f_2 \leq 3$, $\#f_3 \leq 5$ and $\#f_4 \leq 8$. In particular, $\#f_4 \leq 2^3 1^4$ is within the bound $2^{k-1}d^k$. Also, for $l \geq 4$ it always holds that $\Delta_l = (2d-1)2^{l-1}d^l \geq 2d$. Therefore, the bound $2^{k-1}d^k$ for the case $d = 1$ and $k \geq 4$ can be proved by induction. Thus, the cases $d = 1$, $k = 2, 3$ are the only possible exceptions, and the bounds for them are as in Lemma 3.2.

This completes the proof of Lemma 3.2 and, thus, of Theorem 3.1.

3.2 Discussion

Observe that the minimization of the length of the augmenting path does not imply exact minimization of the number of edges used by the resulted flow. Indeed, assume that we wish to compute the exact minimum. Then, we must take into account that using in an augmentation path a "reverse" arc (u, v) (i.e., with the flow in (v, u) equal to 1) *decreases*, not increases, the number of used edges by 1 (for example see Figure 2). In fact, the exact minimization of the number of edges used by a flow is the minimization of the flow cost when the price of pushing a unit of flow through any arc is 1. The theory of min-cost flows is well known (see [FF62] and many others), but is rather complicated. Observe, however, that in our induction step the total number of reverse arcs is $O(\#f)$, which is of smaller order than the upper bound obtained, and hence their influence is negligible in our order-of-magnitude calculations. Therefore, we preferred to use the simpler proof tools ignoring a smaller-order term.

The thesis [E98] presents another approach to achieve upper bounds, which by contrast, is based on analysis of an *exact* minimum solution. The technique used is more complicated, and the result obtained is an upper bound for $td_2(s, t)$ of the same order as the bound presented above. However, that technique is much deeper, and we suppose that it might serve as a basis for achieving a tight upper bound for the case of general (non fixed) k.

4 Extensions

In this section we extend our model to the one-to-many routing paradigm. Given a node $s \in V$ and a set of k nodes $T \subset V$ (which does not include s), we define a *k-container for s and T* as a set of k edge-disjoint paths between s and each node in T. The definitions of *total-length* (of a k-container), *minimality* (of a k-container) and *total$_k$-distance of s and T* (denoted by $td_k(s,T)$) are similar to the definitions for the one-to-one paradigm with respect to the modified definition of a k-container. We present and prove similar upper and lower bounds on $td_k(s,T)$.

Theorem 4.1 *Let s be a node and T a set of k nodes (which does not include s) in a graph G with diameter d. If $td_k(s,T) < \infty$ then*

$$td_k(s,T) \leq 2^{k-1}(d+1)^k - k + 1.$$

Proof: Given G, s, and T as in Theorem 4.1, construct a new graph G' by adding to G a node t and k edges: one between t and each node in T. Note that, by the construction of G', any k-container for s and t contains k disjoint paths, each of which starts in s and terminates with the unique edge from a node in the set T to the node t. Therefore, there is a natural bijective correspondence between k-containers for s and t in G' and for s and T in G, given by removing/adding the new edges from/to the paths of containers, respectively. Let $\mathcal{A}$ be a minimal k-container for s and t in G'. Since clearly the diameter of G' is at most $d+1$, by Theorem 3.1, $|\mathcal{A}| \leq 2^{k-1}(d+1)^k + 1$. Hence, the corresponding k-container $\mathcal{A}'$ for s and T in G satisfies $|\mathcal{A}'| \leq 2^{k-1}(d+1)^k - k + 1$. Theorem 4.1 follows. $\Box$

Now, let us adjust our proof of the lower bound for the total$_k$-distance for two nodes to the one-to-many case.

Theorem 4.2 *For every $k, d_0 > 0$, there exists a graph $\bar{G}$, with a node s and a set T of k nodes (which does not include s) such that the diameter d of $\bar{G}$ satisfies $d_0 \leq d \leq d_0 + 3k + 2$ and*

$$(d-2)^k/k^k \leq td_k(s,T) < \infty.$$

Proof: Let us consider the graph $G_x(k)$, where $x = \lceil (d_0+2k)/k \rceil$, as defined in Section 2. We construct a graph $\bar{G}$ from it by adding a node breaking each of the k edges that meet node t. Formally, every edge $\{u,t\}$ that meets

t in $G_x(k)$, is replaced by the two edges $\{u, w\}$ and $\{w, t\}$ the node w is termed *new*.

Every path P in $G_x(k)$ corresponds to a path in $\bar{G}$ in a natural way (replace each edge in P that meets t by the corresponding path of length 2 in $\bar{G}$). Since every shortest path in $G_x(k)$ is simple, it contains at most two edges that meet t, and therefore the corresponding path in $\bar{G}$ has length which is at most greater by 2. Also, a path from a new node of $\bar{G}$ can be composed of a new edge and a path in $G_x(k)$. It follows that the diameter d of $\bar{G}$ is at most the diameter of $G_x(k)$ plus 2, i.e., the latter is at least $d-2$.

Consider the node s and the set T of the new nodes in $\bar{G}$ (of size k). By the construction of $\bar{G}$ and by the same arguments as in the proof of Theorem 2.1, the minimal k-container for s and T consists of exactly all edges in $\bar{G}$, except for the edges that meet t. It follows that $(d-2)^k/k^k \leq td_k(s, T) < \infty$, as required. □

By Theorem 4.1 and Theorem 4.2, when extending our model to the one-to-many paradigm, similar upper and lower bounds on total$_k$-distances hold.

5 The connection to the augmented path problem

The *augmented path problem*, studied in [KKP95] in relation with hop-congestion tradeoff in ATM networks, is informally described as follows: An augmented path of length N and cutwidth k, denoted $P_N(k)$, is a chain of $N+1$ nodes drawn on a plane, to which some auxiliary edges are added, such that each chain-edge is bypassed by at most $k-1$ auxiliary edges (not including the chain-edge itself). The problem is to determine the minimum diameter, $D_N(k)$, of an augmented path of length N and cutwidth k. In [KKP95] it is shown that $\frac{1}{2}N^{\frac{1}{k}} \leq D_N(k) \leq kN^{\frac{1}{k}}$.

We show that our upper bound on the total$_k$-distance for two nodes in a graph can be used to slightly improve the lower bound on $D_N(k)$ in [KKP95]: Let $P_N(k)$ be an augmented path and let the endpoints of the path be the nodes s and t. It is easy to see that if there are no k edge-disjoint paths between s and t, then one could add to $P_N(k)$ edges in such a way that the cutwidth remains k and the diameter does not increase. Therefore, we can assume, w.l.o.g., that there are k edge-disjoint paths between s and t in $P_N(k)$. Now, let us consider the system $S = (P_N(k), s, t)$. By Theorem 3.1, $td_k(S) \leq 2^{k-1}d^k$ (except for uninteresting cases), and we get $N \leq td_k(S) \leq 2^{k-1}d^k$. It follows that for every augmented path $P_N(k)$ with diameter d $(d > 1)$, $d \geq 2^{-\frac{k-1}{k}} N^{\frac{1}{k}}$, specifically, $D_N(k) \geq 2^{-\frac{k-1}{k}} N^{\frac{1}{k}}$, which is slightly better than the lower bound in [KKP95].

Note also that our graphs $G_x(k)$ are similar to the ones used in the proof of the upper bound of [KKP95], and by Remark 2.5, we actually provide an identical upper bound, since we show that $N \geq d^k/k^k$, where N is the number of nodes and $D_N(k) \leq d$.

Though we do not significantly improve any known result, we show that there is a strong connection between the two considered problems (actually, the augmented path problem could be seen as a special case of the total$_k$-diameter problem). We believe that improvements of our results will imply with no additional effort corresponding improvements of the results in [KKP95] and other related problems (such as the problem of augmented ring networks in [ABCRS96]).

References

[ABCRS96] Aiello, W., Bhatt, S.N., Chung, F.R.K., Rosenberg, A.L., Sitaraman, R.K.: Augmented Ring Networks. J. Math. Modelling and Scientific Computing 8, to appear.

[CC95] Chen, C., Chen, J.: Vertex-Disjoint Routing in Star Graphs. The 1st IEEE Intl. Conf. on Algorithms and Architecture for Parallel Processing, (1995), pp. 460-464.

[CLR90] Cormen, T. H., Leiserson, C. E., and Rivest, R. L. Introduction to Algorithms. The MIT Press, 1990.

[D94] Dinitz, Ye.: Subset Connectivity and Edge-Disjoint Trees, with an Application to Target Broadcasting. TR #822, Dept. of Computer Science, Technion, Haifa, Israel (1994).

[DMS91] Dietzfelbinger, M., Madhavapeddy S., Sudborough, I. H.: Three Disjoint Path Paradigms in Star Networks. Proc. 3nd IEEE Symp. on Parallel and Distributed Processing, (1991), pp. 400–406.

[E98] Eilam, T. Master Thesis. Department of Computer Science, Technion, Israel, in preparation.

[EMZ96] Eilam, T., Moran, S., Zaks, S.: A Lower Bound for Linear Interval Routing. Proceedings of 10th Intl. Workshop, WDAG 96, (1996), pp. 191–205.

[FF62] Ford, L.R., Jr and Fulkerson, D.R. Flows in Networks. Princeton Univ. Press, Princeton, NJ, 1962.

[HLN89] Hastad, J., Leighton, F. T., Newman, M.: Fast Computation Using Fault Hypercubes. Proc. 21st ACM Symp. on Theory of Computing, (1989), pp. 251–263.

[HL94] Hsu, D. F., Łuczak, T.: Note on the k-diameter of k-regular k-connected graphs. Discrete Mathematics 133 (1994) 291–296.

[JLD91] Jwo, J., Lakshmivarahan, S., Dhall, S. K.: Characterizations of node disjoint (parallel) paths in star graphs. Proc. 5th Intl. Parallel Processing Symp., (1991), pp. 404-409.

[KKP95] Kranakis, E., Krizanc, D., Pelc, A.: Hop-Congestion Tradeoffs for High-Speed Networks. 7th IEEE Symp. on parallel and distr. Processing, 662–668.

[LSX95] Li, Q., Sotteau, D., Xu, J.: 2-Diameter of de-Bruijn Networks. Rapport de Recherche 950, Universite de Paris Sud, centre d'Orsay, Laboratoire de Recherche en Informatique, 91405 Orsay, France (1995).

[M89] S. Moran, Message Complexity Versus Space Complexity in Fault Tolerant Broadcast Protocols, Networks, Vol. 19 (1989), pp. 505-519.

[MP88] Meyer, F. J., Pradhan, D. K.: Flip-Trees: Fault-Tolerant Graphs with Wide Containers. IEEE Transactions on Computers. Vol 37, No. 4, (1988).

[MS90] Madhavapeddy S., Sudborough, I. H.: Deterministic Message Routing in Faulty Hypercubes. Lecture Notes in Computer Science 484, (1990), pp. 154–169.

[R89] Rabin, M. O.: Efficient Dispersal of Information for Security, Load Balancing, and Fault Tolerance. Journal of ACM 36, (1989), pp. 335–348.