

Network Algorithms

Prof. Dr. Christian Scheideler

Technische Universität München, April 17, 2007

1 Introduction

The goal of this lecture is to give an introduction to the state in network algorithms, particularly algorithms for overlay networks. It is a widely accepted fact that algorithmic advances in the area of computer science are only useful to society if they are based on models that truthfully reflect the restrictions and requirements of the corresponding applications. This is certainly also true for network communication. For example, in general any infrastructure connecting processing units with each other may be called a network, but certain infrastructures such as a dedicated line between any pair of nodes are certainly unrealistic, because they are too expensive to build. Thus, messages may have to traverse several units to reach their destination, causing (among other problems) route selection and scheduling problems. Also, we will not just study network communication out of context, but we will also look at various topics that require or support efficient network communication such as distributed data management or design strategies for overlay networks.

In this section we will first give a basic introduction to graph theory and will then introduce some popular families of networks and investigate their structural properties.

1.1 Graph theory

A *graph* $G = (V, E)$ consists of a set of *nodes* (or *vertices*) V and a set of *edges* (or *arcs*) E . The nodes represent the processing units and the edges represent the communication links between the units. Often, we will set $n = |V|$ (the size of V) and $m = |E|$. The *size* of G is defined as the number of nodes it contains. For all $v, w \in V$, (v, w) denotes a *directed* edge from v to w , and $\{v, w\}$ denotes an *undirected* edge from v to w . G is called *undirected* if $E \subseteq \{\{v, w\} \mid v, w \in V\}$ and *directed* if $E \subseteq \{(v, w) \mid v, w \in V\}$. Unless explicitly mentioned, we assume for the rest of this lecture that G is undirected and that each undirected edge $\{v, w\}$ represents two independent, directed edges (v, w) and (w, v) .

A sequence of contiguous edges in G is called a *path*. The *length* of the path is defined as the number of edges it contains. A path is called *node-simple* if it visits every node in G at most once. Similarly, it is called *edge-simple* (or *simple*) if it contains every edge in G at most once. G is called *connected* if, for any pair of nodes $v, w \in V$, there is a path in G from v to w . We call a simple path a *cycle* if it starts and ends at the same node. The *girth* of a graph G is defined as the length of the shortest cycle G contains. G is called a *tree* if it is connected and contains no cycle. A graph $T = (V', E')$ is called a *spanning tree* of G if $V' = V$, $E' \subseteq E$, and T is a tree. G is called *bipartite* if its node set can be partitioned into two node sets V_1 and V_2 such that $E \subseteq \{\{v, w\} \mid v \in V_1, w \in V_2\}$.

For any pair of nodes $v, w \in V$, let $\delta(v, w)$ denote the *distance* of v and w in G , that is, the length of a shortest path from v to w . The *diameter* D of G is defined as $\max\{\delta(v, w) \mid v, w \in V\}$. If $\{v, w\} \in E$ then v is called a *neighbor* of w . For any subset $U \subseteq V$, the *neighborhood* of U is defined as

$$\Gamma(U) = \{v \in V \setminus U \mid \exists u \in U : \{u, v\} \in E\}.$$

The number of neighbors of v is called the *degree* of v and denoted by d_v . The degree of G is defined as $d = \max\{d_v \mid v \in V\}$. If all nodes in G have the same degree, then G is called *regular*.




Figure 1: An example of an undirected graph with diameter 4.

A family of graphs $\mathcal{G} = \{G_n \mid n \in \mathbb{N}\}$ has degree $d(n)$ if for all $n \in \mathbb{N}$ the degree of G_n is $d(n)$. If it is clear to which family a graph belongs, we say that this graph has constant (or bounded) degree if and only if its family has constant degree.

A *network* is specified by a graph $G = (V, E)$ with edge capacities given by a function $c : E \rightarrow \mathbb{R}^+$. Given a graph G with capacities c , let the capacity of a node $v \in V$ be defined as

$$c(v) = \sum_{w \in V} c(v, w)$$

and the capacity of any node set or edge set U be defined as $c(U) = \sum_{u \in U} c(u)$. Given a subset $U \subseteq V$, (U, \bar{U}) denotes the set of all edges $(u, v) \in E$ (or $\{u, v\} \in E$ if G is undirected) with $u \in U$ and $v \in \bar{U}$. So $c(U, \bar{U})$ is the sum of the capacities of all edges in (U, \bar{U}) . The *conductance* α of a network G with capacities c is defined as

$$\alpha = \min_{\emptyset \neq U \subset V} \frac{c(U, \bar{U})}{\min\{c(U), c(\bar{U})\}}.$$

1.2 Basic network topologies

The most basic network topologies used in practice are trees, cycles, grids and tori. Many other suggested networks are simply combinations or derivatives of these. The advantage of trees is that the path selection problem is very easy: for every source-destination pair there is only one possible simple path. However, since the root of a tree is usually a severe bottleneck, so-called *fat trees* have been used. These trees have the property that higher-level edges have a (much) larger capacity than lower-level edges. See Figure 2 for an example.




Figure 2: The structure of a fat tree.

Fat trees belong to a family of networks that require edges of non-uniform capacity to be efficient. Easier to build are networks with edges of uniform capacity. This is usually the case for grids and tori. Unless explicitly mentioned, we will treat all edges in the following to be of capacity 1. In the following, $[x]$ means the set $\{0, 1, \dots, x - 1\}$.

Definition 1.1 (Torus, Mesh) Let $m, d \in \mathbb{N}$. The (m, d) -mesh $M(m, d)$ is a graph with node set $V = [m]^d$ and edge set

$$E = \left\{ \{(a_{d-1} \dots a_0), (b_{d-1} \dots b_0)\} \mid a_i, b_i \in [m], \sum_{i=0}^{d-1} |a_i - b_i| = 1 \right\} .$$

The (m, d) -torus $T(m, d)$ is a graph that consists of an (m, d) -mesh and additionally wrap-around edges from $(a_{d-1} \dots a_{i+1} (m - 1) a_{i-1} \dots a_0)$ to $(a_{d-1} \dots a_{i+1} 0 a_{i-1} \dots a_0)$ for all $i \in [d]$ and all $a_j \in [m]$ with $j \neq i$. $M(m, 1)$ is also called a line, $T(m, 1)$ a cycle, and $M(2, d) = T(2, d)$ a d -dimensional hypercube.

Figure 3 presents a linear array, a torus, and a hypercube.




Figure 3: The structure of $M(m, 1)$, $T(4, 2)$, and $M(2, 3)$.

The hypercube is a very important class of networks, and many derivatives, the so-called *hyper-cubic networks*, have been suggested for it. Among these are the butterfly, cube-connected-cycles, shuffle-exchange, and de Bruijn graph. We start with the butterfly, which is basically a rolled out version of a hypercube.

Definition 1.2 (Butterfly) Let $d \in \mathbb{N}$. The d -dimensional butterfly $BF(d)$ is a graph with node set $V = [d + 1] \times [2]^d$ and an edge set $E = E_1 \cup E_2$ with

$$E_1 = \{ \{ (i, \alpha), (i + 1, \alpha) \} \mid i \in [d], \alpha \in [2]^d \}$$

and

$$E_2 = \{ \{ (i, \alpha), (i + 1, \beta) \} \mid i \in [d], \alpha, \beta \in [2]^d, \alpha \text{ and } \beta \text{ differ only at the } i\text{th position} \} .$$

The node set $\{ (i, \alpha) \mid \alpha \in [2]^d \}$ represents level i of the butterfly. The d -dimensional wrap-around butterfly $W-BF(d)$ is defined by taking the $BF(d)$ and identifying level d with level 0 .

Figure 4 shows the 3-dimensional butterfly $BF(3)$. The $BF(d)$ has $(d + 1)2^d$ nodes, $2d \cdot 2^d$ edges and degree 4. It is not difficult to check that combining the node sets $\{ (i, \alpha) \mid i \in [d] \}$ into a single node results in the hypercube.




Figure 4: The structure of $BF(3)$.

Next we define the cube-connected-cycles network. It only has a degree of 3 and it results from the hypercube by replacing the corners by cycles.

Definition 1.3 (Cube-Connected-Cycles) Let $d \in \mathbb{N}$. The cube-connected-cycles network $CCC(d)$ is a graph with node set $V = \{ (a, p) \mid a \in [2]^d, p \in [d] \}$ and edge set

$$E = \{ \{ (a, p), (a, (p + 1) \bmod d) \} \mid a \in [2]^d, p \in [d] \} \\ \cup \{ \{ (a, p), (b, p) \} \mid a, b \in [2]^d, p \in [d], a = b \text{ except for } a_p \}$$

Two possible representations of a CCC can be found in Figure 5.

The shuffle-exchange is yet another way of transforming the hypercubic interconnection structure into a constant degree network.

Definition 1.4 (Shuffle-Exchange) Let $d \in \mathbb{N}$. The d -dimensional shuffle-exchange $SE(d)$ is defined as an undirected graph with node set $V = [2]^d$ and an edge set $E = E_1 \cup E_2$ with

$$E_1 = \{ \{ (a_{d-1} \dots a_0), (a_{d-1} \dots \bar{a}_0) \} \mid (a_{d-1} \dots a_0) \in [2]^d, \bar{a}_0 = 1 - a_0 \}$$

and

$$E_2 = \{ \{ (a_{d-1} \dots a_0), (a_0 a_{d-1} \dots a_1) \} \mid (a_{d-1} \dots a_0) \in [2]^d \} .$$




Figure 5: The structure of $CCC(3)$.

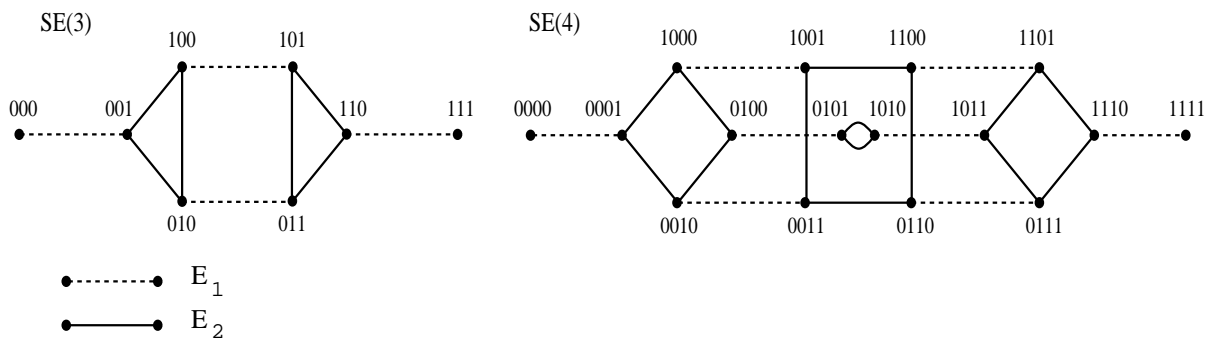


Figure 6: The structure of $SE(3)$ and $SE(4)$.

Figure 6 shows the 3- and 4-dimensional shuffle-exchange graph.

Definition 1.5 (de Bruijn) *The b -ary de Bruijn graph of dimension d $DB(b, d)$ is an undirected graph $G = (V, E)$ with node set $V = \{v \in [b]^d\}$ and edge set E that contains all edges $\{v, w\}$ with the property that $w \in \{(x, v_{d-1}, \dots, v_1) : x \in [b]\}$, where $v = (v_{d-1}, \dots, v_0)$.*

Two examples of a de Bruijn graph can be found in Figure 7.

1.3 Direct and indirect networks

Networks are usually separated into *direct* and *indirect* networks. Direct networks are networks in which every node represents a processing unit that can inject and absorb packets, whereas in indirect




Figure 7: The structure of $DB(2, 2)$ and $DB(2, 3)$.

networks only certain nodes (the so-called *input nodes*) can inject packets and certain nodes (the so-called *output nodes*) can absorb packets. An important subclass of indirect networks are the so-called leveled graphs.

Definition 1.6 (Leveled Graph) A graph $G = (V, E)$ is called leveled with depth D if the nodes of G can be partitioned into $D + 1$ levels L_0, \dots, L_D such that every edge in E connects nodes of consecutive levels. Nodes in level 0 are called inputs, and nodes in level D are called outputs. If, in addition, $|L_0| = |L_D|$ and L_0 is identified with L_D , then G is called a wrapped leveled graph with depth D .

Examples of leveled graphs are the fat tree and the butterfly, and an example of a wrapped leveled graph is the wrap-around butterfly. In a butterfly it is usually assumed that the nodes in L_0 represent the input nodes and the nodes in level L_D represent the output nodes. In a fat tree the nodes in level L_D are usually both input and output nodes.

1.4 The diameter

Recall the definition of the diameter in Section 1.1. One important goal in choosing a topology for a network is that it has a small diameter. The following theorem presents a lower bound for this.

Theorem 1.7 Every graph of maximum degree $d \geq 3$ and size n must have a diameter of at least $\lfloor (\log n) / (\log(d - 1)) \rfloor - 1$.

Proof. Suppose we have a graph $G = (V, E)$ of maximum degree d and size n . Start from any node $v \in V$. In a first step at most d other nodes can be reached. In two steps at most $d \cdot (d - 1)$ additional nodes can be reached. Thus, in general, in at most k steps at most

$$1 + \sum_{i=0}^{k-1} d \cdot (d - 1)^i = 1 + d \cdot \frac{(d - 1)^k - 1}{(d - 1) - 1} \leq \frac{d \cdot (d - 1)^k}{d - 2}$$

nodes (including v) can be reached. This has to be at least n to ensure that v can reach all other nodes in V within k steps. Hence,

$$(d - 1)^k \geq \frac{(d - 2) \cdot n}{d} \Leftrightarrow k \geq \log_{d-1} \left(\frac{(d - 2) \cdot n}{d} \right) \Leftrightarrow k \geq \log_{d-1} n + \log_{d-1} \left(\frac{d - 2}{d} \right).$$

Since $\log_{d-1}((d-2)/d) > -2$ for all $d \geq 3$, this is true only if $k \geq \lfloor \log_{d-1} n \rfloor - 1$. □

Theorem 1.7 uses as a construction for the lower bound a complete $(d-1)$ -ary tree with a root of degree d . However, it is easy to see that in this tree there are two nodes (see the leaves v and w in Figure 8) with a distance of approximately $2 \log_{d-1} n$, which is by a factor of 2 larger than the lower bound. Can networks with a better diameter be constructed? The next theorem gives an answer to this.




Figure 8: Nodes with highest distance in a tree.

Theorem 1.8 For every even $d > 2$ there is an infinite family of graphs G_n of maximum degree d and size n with a diameter of at most $(\log n)/(\log d - 1)$.

Proof. The proof is part of the assignment. □

1.5 The conductance

Recall the definition of the conductance in Section 1.1. We start with an upper bound on the conductance that must hold for all networks.

Theorem 1.9 For every network $G = (V, E)$ with non-negative edge capacities, the conductance can be at most 1.

Proof. For every set $U \subseteq V$ let $E_U = \{\{v, w\} \in E \mid v \in U\}$, where an edge appears twice in E_U if both v and w are in U . Certainly, $(U, \bar{U}) \subseteq E_U$. Since $c(U) = c(E_U)$ it must therefore hold that $c(U, \bar{U}) \leq c(U)$. Equivalently, it must also hold that $c(U, \bar{U}) = c(\bar{U}, U) \leq c(\bar{U})$. Hence, $c(U, \bar{U}) \leq \min\{c(U), c(\bar{U})\}$ and therefore

$$\alpha(G) = \min_{U \subseteq V} \frac{c(U, \bar{U})}{\min\{c(U), c(\bar{U})\}} \leq 1.$$

□

Interestingly, for any $d \geq 3$ there are graphs that can achieve a constant conductance. These are the so-called expanders. One explicit construction is known as the Gabber-Galil graph [4]:

Definition 1.10 Let $n \in \mathbb{N}$. The Gabber-Galil graph $GG(n)$ is a graph with node set $V = [n]^2$ and edge set E consisting of all edges $((x, y), (x', y'))$ with

$$(x', y') \in \{(x, x + y), (x, x + y + 1), (x + y, y), (x + y + 1, y)\} \pmod{n}$$

Other explicit constructions of expanders can be found in [8, 9, 10]. Also random regular graphs are known to be expanders, with high probability. For the classes of graphs we presented above the conductance is quite complicated to compute. Therefore, we just list some results here.

Theorem 1.11 The d -dimensional hypercube, cube-connected-cycles, butterfly, shuffle-exchange, and de Bruijn graph with uniform edge capacities all have an conductance of $\Theta(1/d)$.

Using the fact that for these networks $d = \Theta(\log n)$, where n is the number of nodes in the network, it follows that all of these networks have an conductance of $\Theta(1/\log n)$.

1.6 The flow number

In order to define the flow number, we first have to introduce the concept of multicommodity flows. Consider any (bi-)directed network $G = (V, E)$ with non-negative edge capacities given by c . A *multicommodity flow instance* on G is a set of ordered pairs of vertices $(s_1, t_1), (s_2, t_2), \dots, (s_k, t_k)$. Each pair (s_i, t_i) represents a *commodity* with source s_i and target t_i . A *multicommodity flow* for that instance is a flow $f : E \times \{1, \dots, k\} \rightarrow \mathbb{R}^+$ with the following properties:

- For all edges $e = (v, w)$, $\sum_{i=1}^k f(e, i) \leq c(e)$, and
- for all nodes v and commodities i with $v \notin \{s_i, t_i\}$, $\sum_{w:(u,v) \in E} f((u, v), i) = \sum_{w:(v,w) \in E} f((v, w), i)$.

The objective is to maximize the amount of flow traveling from the sources to the corresponding destinations, subject to the capacity constraints. The problem comes in two flavors. In the first, called the *maximum multicommodity flow* problem, the total flow, summed over all commodities, is to be maximized. The second is called the *concurrent multicommodity flow* problem. Here, for each commodity (s_i, t_i) a non-negative demand d_i is specified. The objective is to maximize the *fraction* of the demand that can be shipped simultaneously for all commodities. In other words, we want to find the maximum φ so that a flow of $\varphi \cdot d_i$ can be shipped for every commodity i without exceeding the capacities of the edges. φ is called the *concurrent max-flow*. A *balanced multicommodity flow problem* (BMFP) is a concurrent multicommodity flow problem in which the sum of the demands of the commodities originating and the commodities terminating in a node v is equal to $c(v)$ for every $v \in V$. Both the maximum throughput problem and the maximum concurrent flow problem can be solved in polynomial time using linear programming.

Given a concurrent multicommodity flow problem with feasible solution S , let the *dilation* $D(S)$ of S be defined as the length of the longest flow path in S and the *congestion* $C(S)$ of S be defined as the inverse of its concurrent flow value (i.e., the congestion says how many times the edge capacities would have to be increased in order to satisfy the demands of all commodities when using the same set of flow paths). Let \mathcal{B} be the special BMFP in which each pair of nodes (v, w) has a commodity of demand $c(v) \cdot c(w)/c(V)$. The *flow number* $F(G)$ of a network G is defined as the minimum over all feasible solutions S of \mathcal{B} of $\max\{C(S), D(S)\}$ [6]. In the case that there is no risk of confusion, we




Figure 9: Solution to a 2-commodity flow problem (a). The routing of the first commodity is shown in (b) and the second commodity is shown in (c).

will simply write F instead of $F(G)$. Note that the flow number of a network is invariant to a scaling of the capacities.

The flow number of a network G can be computed in polynomial time. Another advantage of the flow number is that, as shown by the next theorem, it can be applied to much more general multicommodity flow problems than just the one that defines it.

Theorem 1.12 *For any network G with flow number F and any instance I of the BMFP for G , there is a feasible solution for I with congestion and dilation at most $2F$.*

Proof. The idea is to decompose I into two multicommodity flow problems: for every commodity i with source s_i and destination t_i , the first problem I_1 has commodities i_u from s_i to u for all $u \in V$ with demands $d_{i_u} = d_i \cdot c(u)/c(V)$, and the second problem I_2 has commodities i'_u from u to t_i for all $u \in V$ with demands $d_{i'_u} = d_i \cdot c(u)/c(V)$. For every commodity i from the original problem, the total demand of corresponding commodities in I_1 is d_i and is d_i in I_2 as well. Moreover, for every node $u \in V$ the amount of commodity i shipped to u in I_1 is equal to the amount of commodity i shipped from u in I_2 .

Interestingly, both of the flow problems I_1 and I_2 are equal to the special flow problem \mathcal{B} because for any pair $v, w \in V$, the total demand of the commodities with source v and destination w in I_1 is equal to

$$\sum_{i: s_i=v} \frac{d_i \cdot c(w)}{c(V)} = \frac{c(v) \cdot c(w)}{c(V)},$$

and in I_2 it is also equal to

$$\sum_{i: t_i=w} \frac{d_i \cdot c(v)}{c(V)} = \frac{c(v) \cdot c(w)}{c(V)}.$$

Thus, according to the definition of the flow number, both I_1 and I_2 have a feasible solution with congestion and dilation at most F . Hence, the original problem I has a feasible solution with congestion and dilation at most $2F$, which proves the claim. \square

With techniques similar to those used in the proof of Theorem 5.0.3 in [11] one can also prove the following result.

Theorem 1.13 *On average over all BMFPs I , the minimum $\max\{C(\mathcal{S}), D(\mathcal{S})\}$ over all feasible solutions \mathcal{S} of I is $\Omega(F)$.*

Hence, the flow number truthfully captures the problem of routing BMFPs in networks. Using Theorem 1.12, we prove another powerful result, called *Shortening Lemma*, that shows that the flow number allows one to convert arbitrary multicommodity flow solutions into solutions with short flow paths.

Theorem 1.14 (Shortening Lemma [6]) *Suppose we are given a network with flow number F . Then, for any $\epsilon \in (0, 1]$ and any feasible multicommodity flow f , there exists a feasible multicommodity flow f' with flow values $|f'_i|$ of at least $|f_i|/(1 + \epsilon)$ for every commodity i that uses paths of length at most $2 \cdot F(1 + 1/\epsilon)$.*

Proof. Given a flow solution \mathcal{S} , let $\mathcal{S}' \subseteq \mathcal{S}$ consist of all paths from \mathcal{S} that are longer than L , for $L = 2 \cdot F/\epsilon$. We are going to shorten the paths in \mathcal{S}' at the cost of slightly decreasing the satisfied demand of each commodity.

For a path $p \in \mathcal{S}'$ between s_p and t_p , let $a_{p,1} = s_p, a_{p,2}, \dots, a_{p,L}$ denote its first L nodes and $b_{p,1}, \dots, b_{p,L-1}, b_{p,L} = t_p$ its last L nodes and let f_p be the flow value along p . Then the set $\mathcal{U} = \bigcup_{p \in \mathcal{S}'} \bigcup_{i=1}^L \{a_{p,i}, b_{p,i}, f_p\}$ is (a subset of) an instance of the BMFP. By Theorem 1.12, there exists a feasible solution \mathcal{P} to \mathcal{U} with flow value at least $1/(2F)$ consisting of paths of length at most $2F$. We are going to combine the initial and final parts of the long paths in \mathcal{S}' with these “shortcuts” in \mathcal{P} to obtain the desired short solution.

First, decrease the flows along all paths $p \in \mathcal{S}$ by a factor of $1/(1 + \epsilon)$ so that we have room to accommodate new, short paths for the paths in \mathcal{S}' . These short paths are constructed in the following way:

For every path $p \in \mathcal{S}'$, we replace p by L flow systems $S_{p,i}$, $i = 1, \dots, L$. Each flow system $S_{p,i}$ consists of two parts:

1. the flow paths between $a_{p,i}$ and $b_{p,i}$ in \mathcal{P} corresponding to the request $\{a_{p,i}, b_{p,i}, f_p\}$ from \mathcal{U} , now with a flow of $f_p/(L(1 + \epsilon))$, and
2. $f_p/(L(1 + \epsilon))$ units of flow between $a_{p,1}$ and $a_{p,i}$ along p , and $f_p/(L(1 + \epsilon))$ units of flow between $b_{p,i}$ and $b_{p,L}$ along p .

For each i , the length of each path in the subsystem $S_{p,i}$ is at most $L + 2 \cdot F$, and $f_p/(L(1 + \epsilon))$ units of flow are shipped along each path system $S_{p,i}$. Summed over all $i = 1 \dots L$, we have $f_p/(1 + \epsilon)$ units of flow between $s_p = a_{p,1}$ and $t_p = b_{p,L}$, which is as high as the original flow through p reduced by $1/(1 + \epsilon)$. Hence, we can replace p by the systems $S_{p,i}$ without changing the amount of flow from s_p to t_p .

Now, it holds for every edge e that the flow traversing e due to the paths in \mathcal{S} is at most $c(e)/(1 + \epsilon)$, and due to the shortcuts in \mathcal{P} is at most

$$\sum_{p \in \mathcal{P}: e \in p} \frac{f_p}{L(1 + \epsilon)} \leq \frac{2F}{L(1 + \epsilon)} \cdot c(e) = \frac{\epsilon \cdot c(e)}{1 + \epsilon},$$

since

$$\sum_{p \in \mathcal{P}: e \in p} \frac{f_p}{2F} \leq c(e).$$

Thus, the flows in \mathcal{S} and \mathcal{P} sum up to at most $c(e)$ for an edge e . Therefore, the modification yields a feasible solution satisfying the desired properties. \square

Next, we explore the relationship of the flow number with the diameter and the conductance of a network. The first result immediately follows from the definition of F .

Fact 1.15 *For every network with diameter D and flow number F , it holds that $F \geq D$.*

The next result reveals a very close relationship between the conductance and the flow number of a network.

Theorem 1.16 *For any network G with conductance α and flow number F it holds that*

$$\alpha^{-1} \leq F \leq c \cdot \alpha^{-1} \log n$$

for some constant c .

Proof. We only prove here that $F \geq \alpha^{-1}$. (The entire proof can be found in [6].) For this we need some notation. Given a concurrent multicommodity flow problem, the *cut ratio* of a cut (U, \bar{U}) is defined as

$$R_U = \frac{c(U, \bar{U})}{d(U, \bar{U})} \quad \text{where} \quad d(U, \bar{U}) = \sum_{(s_i, t_i) \in (U \times \bar{U}) \cup (\bar{U} \times U)} d_i.$$

Now, let f be the concurrent max-flow of the problem \mathcal{B} used for the definition of F . Consider any cut (U, \bar{U}) and let i_1, i_2, \dots, i_r denote the commodities whose source and target are separated by this cut. Since all flows for these commodities must cross (U, \bar{U}) , we know that

$$\sum_{j=1}^r f \cdot d_{i_j} \leq c(U, \bar{U}).$$

Since $\sum_{j=1}^r d_{i_j} = d(U, \bar{U})$, this means that

$$f \leq \frac{c(U, \bar{U})}{d(U, \bar{U})}.$$

For \mathcal{B} it holds that

$$d(U, \bar{U}) = \sum_{(u,v) \in (U \times \bar{U}) \cup (\bar{U} \times U)} \frac{c(u) \cdot c(v)}{c(V)} = \frac{2c(U) \cdot c(\bar{U})}{c(V)}.$$

We distinguish between two cases. If $c(U) \geq c(V)/2$, then $c(U) \cdot c(\bar{U})/c(V) \geq c(\bar{U})/2$. Thus,

$$f \leq \frac{c(U, \bar{U})}{2 \cdot c(\bar{U})/2} = \frac{c(U, \bar{U})}{\min\{c(U), c(\bar{U})\}}.$$

If $c(\bar{U}) \geq c(V)/2$, then $c(U) \cdot c(\bar{U})/c(V) \geq c(U)/2$ and therefore

$$f \leq \frac{c(U, \bar{U})}{2 \cdot c(U)/2} = \frac{c(U, \bar{U})}{\min\{c(U), c(\bar{U})\}}.$$

Hence, in both cases,

$$f \leq \frac{c(U, \bar{U})}{\min\{c(U), c(\bar{U})\}}$$

and therefore $f \leq \alpha$ or $1/f \geq \alpha^{-1}$. Since according to the definition of F , $F \geq 1/f$, it follows that $F \geq \alpha^{-1}$. \square

Since the flow number of a network is an upper bound on its diameter, it follows from Theorem 1.16:

Corollary 1.17 *For every network with conductance α the diameter is at most $O(\alpha^{-1} \log n)$.*

From Theorem 1.7 it follows that this bound is exact for constant degree expanders. Do there exist networks where the flow number is in $O(\max\{D, \alpha^{-1}\})$? The next theorem lists some.

Theorem 1.18 *The d -dimensional hypercube, cube-connected-cycles, butterfly, shuffle-exchange, and de Bruijn graph with uniform edge capacities all have a flow number of $\Theta(d)$.*

For proofs see, for example, [7] or [11]. Thus, for these networks it actually holds that $F = \Theta(\alpha^{-1})$, i.e. the conductance describes very well the routing ability of the network. It also follows from the bound that all networks must have a diameter of $O(\log n)$.

1.7 Expansion and span

We end this section with the definition of two more parameters: the *expansion* and the *span* of a graph $G = (V, E)$. The expansion β measures how well a graph can sustain adversarial node faults, and the span σ measures how well a graph can sustain random node faults.

In the past, researchers have mostly studied the problem up to which point a network can sustain faults so that the size of its largest connected component is still a constant fraction of its original size. However, for network theory, such a question is not too useful as it only gives a qualitative statement about the state of the graphs after faults and not a quantitative statement. Hence, a better question would be:

Up to which fault probability does a network still contain a network of at least a constant fraction of its original size that still has approximately the same expansion?

Knowing an answer to this question would have many useful consequences for distributed data management, routing, and distributed computing. Research on load balancing has shown that if the expansion basically stays the same, the ability of a network to balance single-commodity or multi-commodity load basically stays the same, and this ability can be exploited through simple local algorithms [5, 2, 1]. Hence, we will study the two parameters under the question above.

Expansion

Recall the definition of the neighbor set $\Gamma(U)$. The expansion of a graph $G = (V, E)$ is defined as

$$\beta = \min_{\emptyset \neq U \subseteq V, |U| \leq |V|/2} \frac{|\Gamma(U)|}{|U|}$$

In words, the expansion measures the impact node failures can have on disconnecting intact nodes from the rest of the network. In [3] the following result was shown.

Theorem 1.19 *Let G be any graph with n nodes, maximum degree δ and expansion β . Suppose that the adversary can select up to $f = \frac{\beta n}{4\delta k^2}$ faulty nodes for some constant $k > 1$. Then there is a subgraph H in G of size at least $n - \frac{f \cdot k}{\beta}$ with expansion at least $(1 - \frac{1}{k}) \cdot \beta$.*

Hence, for constant degree graphs, f adversarial faults still leave a subgraph H of $n - O(f)$ nodes that essentially has the same expansion as the original graph.

Span

Interestingly, the expansion of a network cannot be used to predict how well a network can sustain random faults. In fact, networks with expansion $1/\sqrt{n}$ (e.g., 2-dimensional meshes) are known that can sustain a constant fault probability whereas other networks with expansion $1/\sqrt{n}$ (e.g., n -node expanders in which every edge is replaced by a path of length n) are known that can only sustain a fault probability of $O(1/\sqrt{n})$. Hence, a new parameter is needed. A very promising parameter appears to be the span of a graph, which is defined as follows [3]:

Consider a graph $G = (V, E)$. Let $U \subseteq V$ be any subset of nodes. U is defined to be *compact* if and only if U and $V \setminus U$ are connected in G . Let \mathcal{U} be the set of all compact sets of G . Let $P(U)$ be the smallest tree in G which connects every node in $\Gamma(U)$ (i.e., it essentially spans the boundary of U). Note that the set of nodes in $P(U)$ need not be from U alone or from $V \setminus U$ alone. Then the *span* of a graph is defined as:

$$\sigma = \max_{U \in \mathcal{U}} \left\{ \frac{|P(U)|}{|\Gamma(U)|} \right\}$$

Using this parameter, the following theorem was recently shown [3]:

Theorem 1.20 *Consider any graph G with maximum degree δ , span σ , and expansion $\beta \geq (\gamma \delta \ln^3 n)/n$ for some sufficiently large constant γ and $|\Gamma(U)| \geq \log_\delta |U|$ for every node set U in G with $|U| \leq |V|/2$. Then, with high probability, provided the fault probability $p \leq 1/(16e \cdot \delta^{8\sigma})$ and $\epsilon \leq 1/4$, there is a non-faulty subgraph $H \subseteq G$ of size $|H| \geq n/3$ with expansion at least $(\epsilon/\delta) \cdot \beta$.*

Since it is not difficult to see that all d -dimensional meshes with constant d have a constant span, this means that all of these can sustain a constant fault probability and still have a large connected component of essentially the same expansion. We believe that also the hypercubic networks have a constant span and that Theorem 1.20 is far from being tight, so investigating the span further is an ongoing, interesting research issue.

References

- [1] A. Anagnostopoulos, A. Kirsch, and E. Upfal. Load balancing in arbitrary network topologies with stochastic adversarial input. *SIAM Journal on Computing*, 34(3):616–639, 2005.
- [2] E. Anshelevich, D. Kempe, and J. Kleinberg. Stability of load balancing algorithms in dynamic adversarial systems. In *Proc. of the 34th ACM Symp. on Theory of Computing (STOC)*, 2002.
- [3] A. Bagchi, A. Bhargava, A. Chaudhary, D. Eppstein, and C. Scheideler. The effect of faults on network expansion. In *Proc. of the 16th ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, pages 286–293, 2004.
- [4] O. Gabber and Z. Galil. Explicit constructions of linear-sized superconcentrators. *Journal of Computer and System Sciences*, 22:407–420, 1981.
- [5] B. Ghosh, F. T. Leighton, B. M. Maggs, S. Muthukrishnan, C. G. Plaxton, R. Rajaraman, A. W. Richa, R. E. Tarjan, and D. Zuckermann. Tight analyses of two local load balancing algorithms. *SIAM Journal on Computing*, 29(1):29–64, 1999.
- [6] P. Kolman and C. Scheideler. Improved bounds for the unsplittable flow problem. In *ACM/SIAM Symp. on Discrete Algorithms (SODA)*, 2002.
- [7] F. Leighton. *Introduction to Parallel Algorithms and Architectures: Arrays · Trees · Hypercubes*. Morgan Kaufmann Publishers (San Mateo, CA), 1992.
- [8] A. Lubotzky, R. Phillips, and R. Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.
- [9] G. Margulis. Explicit group theoretical constructions of combinatorial schemes and their application to the design of expanders and superconcentrators. *Problems Inform. Transmission*, 11:39–46, 1988.
- [10] M. Morgenstern. Existence and explicit constructions of $q + 1$ regular Ramanujan graphs for every prime power q . *Journal of Combinatorial Theory, Series B*, 62:44–62, 1994.
- [11] C. Scheideler. *Universal Routing Strategies for Interconnection Networks*, volume 1390 of *Lecture Notes in Computer Science*. Springer, 1998.