

Short Length Menger's Theorem and Reliable Optical Routing

Amitabha Bagchi
Dept. of Computer and Information Science
Polytechnic University
6 Metrotech Center
Brooklyn, 11201, USA
bagchi@cis.poly.edu

Amitabh Chaudhary
Dept of Computer Science and Engineering
Notre Dame University
Notre Dame, IN 46656, USA
Amitabh.Chaudhary.1@nd.edu

Petr Kolman*
Inst. for Theoretical Computer Science
Charles University
Malostranské nám. 25
Prague, 118 00, Czech Republic
kolman@kam.mff.cuni.cz

Abstract

In the minimum path coloring problem, we are given a graph and a set of pairs of vertices of the graph and we are asked to connect the pairs by colored paths in such a way that paths of the same color are edge disjoint. In this paper we deal with a generalization of this problem where we are asked to connect each pair by k edge disjoint paths of the same color. The objective is to minimize the number of colors. The reason for multiple paths between the same pair of vertices is to ensure fault tolerance of the connections. We propose an $O(k^2 F) = O(k^2 \Delta \alpha^{-1} \log n)$ approximation algorithm for this problem where F is the flow number of the graph, Δ is the maximum degree and α is the expansion. This is an improvement even for the special case $k = 1$ where, to our knowledge, the best previously known bound is weaker by a factor of $\log n$.

The underlying problem is that of finding several disjoint paths between a given pair of vertices. Menger's theorem provides a necessary and sufficient condition for the existence of k such paths. However, it does not say anything about the length of the paths although in communication problems the number of links used is an issue. We show that any two k -connected vertices are connected by k edge disjoint paths of average length $O(kF)$ which improves an earlier result of Galil and Yu [17] for several classes of graphs. In fact, this is only a corollary of a stronger result for multicommodity flow on networks with unit edge capacities: any multicommodity flow with k units for each commodity can be rerouted such that the flow for each commodity is shipped through k -tuples of edge disjoint paths of average length $O(kF)$ without exceeding the edge capacities significantly.

1 Introduction

The goal of this paper is to design efficient and reliable (fault-tolerant) algorithms for network communication problems. We consider optical networks in which faults may appear on any link (edge),

*Partially supported by Institute for Theoretical Computer Science (project LN00A056 of MŠMT ČR), by the NSF grant CCR-0208856 and by the EU within the 6th Framework Programme under contract 001907 (DELIS).

possibly on several links. The main idea is to reserve several disjoint communication channels (paths) for a single connection (request). If a fault appears on some of the links then it is still possible to keep the connection alive along the remaining paths. The necessary condition for this strategy to work is that the paths reserved for the same connection (request) are edge disjoint.

To be more specific, we consider a generalization of the well known Minimum Path Coloring Problem (MPCP) [31], namely the *coloring k -edge disjoint path systems* problem (k -EDPCOL): An undirected graph $G = (V, E)$ and a (multi) set of requests $T = \{(s_i, t_i) | s_i, t_i \in V\}$ are given. For each request i , find k edge disjoint paths (called a *k -system*, for short) that connect s_i and t_i . Assign a color to each k -system such that no two k -systems that share an edge, have the same color. The objective is to minimize the number of colors used.

The minimum path coloring problem (i.e., 1-EDPCOL in our terminology), which is a variant of the edge disjoint paths problem (EDP), is known to be NP-complete [15]. Therefore, we deal with approximation algorithms in this paper. Because the best possible approximation for the 1-EDPCOL on directed graphs, in terms of $n = |V|$, is roughly $\Omega(\sqrt{n})$ (by a straightforward reduction from the hardness of the EDP [20]) we parametrize the performance of our algorithms by the flow number F or the expansion α of the graph, instead of the number of edges or vertices. The definitions of the expansion and of the flow number are postponed to the next section, here we just recall that the flow number is always bounded by $O(\Delta\alpha^{-1} \log n)$, where Δ is the maximum degree.

The main algorithmic result of this paper is an $O(k^2 F) = O(k^2 \Delta \alpha^{-1} \log n)$ approximation algorithm for the k -EDPCOL problem. It is worth mentioning that this is an improvement even for the special case $k = 1$: An important observation by Aumann and Rabani [3] for the path coloring problem shows that any algorithm for the EDP with approximation ratio c can be turned into an algorithm for the MPCP with approximation ratio $O(c \log n)$. Since the best approximation for EDP in terms of F is $O(F)$ [26], the resulting approximation for 1-EDPCOL is $O(F \log n)$ whereas our upper bound on the approximation ratio is $O(F)$. In a similar way, a recent $O(k^3 F)$ approximation algorithm for the k -Edge Disjoint Paths problem (see below) [6] which is a generalization of the EDP, can be turned into $O(k^3 F \log n)$ approximation algorithm for the k -EDPCOL. Compared to this the results in this paper are better by a factor of $k \log n$.

Short Disjoint Paths. The underlying problem of the k -EDPCOL is that of finding k disjoint paths between a given pair of vertices. Menger's theorem provides a necessary and sufficient condition for existence of k such paths. However, it does not say anything about the length of these paths. Even though the *length* of the paths in all-optical networks is not important with respect to the transmission time, links are rare resources [6] and the length of the paths is important with respect to the number of realized connections and also with respect to the security of the connections [4].

Galil and Yu [17] worked on short length versions of Menger's theorem and proved that for any two k -connected vertices in a graph, there are k edge-disjoint paths between them of average length $O(n/\sqrt{k})$. Henzinger et al. [21] gave a simpler proof of the same bound. Karger and Levine [22] generalized this result in two directions. First, they proved that for any two vertices in a unit-capacity graph with a flow ν between them, the ν units of flow can be sent along paths of average length $O(n/\sqrt{\nu})$. Second, they showed that this holds for *capacitated* graphs as well.

In this paper we take a slightly different approach. Since the $O(n/\sqrt{k})$ bound is very weak for many graphs (though it is the best possible in general, in terms of k and n) we try to give bounds on the path length in terms of different parameters like the expansion or the flow number F of the graph. The main result in this respect is that any two k -connected vertices are connected by k edge disjoint paths of average length $O(kF)$. This result falls out as a corollary of a general result regarding multicommodity flows on networks with unit capacities. Specifically, we show that any multicommodity flow with k

units for each commodity can be rerouted such that the flow for each commodity is shipped through k -tuples of edge disjoint paths of average length $O(kF)$ while the edge capacities are exceeded by a constant only. A weaker bound of a similar form, that is, an upper bound of $O(k^2F)$ on average path lengths, was implicitly contained in a recent paper on fault-tolerant routing [6].

The k -EDPCOL problem is related also to the k *edge-disjoint paths problem* (k -EDP) [6]. In k -EDP, we are given a graph $G = (V, E)$ and a set of requests T , and the task is to find a maximum subset of the pairs in T for which it is possible to select paths such that each pair is connected by k edge-disjoint paths and the paths for different pairs are mutually disjoint. For $k = 1$, this is the well known edge disjoint path problem. As a by-product of our results we get an improvement, by a factor of k , on the online approximation of the k -EDP. We show that there is a deterministic online algorithm with competitive ratio $O(k^2F)$ whereas the best previous bound was $O(k^3F)$ [6].

Previous work on path coloring. To fully set out the context of our work, we conclude this section with a brief discussion of relevant known results about path coloring.

The MPCP was first studied in the context of routing in optical networks [30]. As already mentioned, a simple reduction from the disjoint paths problem shows that this is an NP-hard problem on general graphs. A result of Golumbic and Jamison [19] implies that this problem is NP-complete on arbitrary tree topologies, although it was later shown that on trees of bounded degree the problem can be solved efficiently [14]. Erlebach and Jansen showed that path coloring is also NP-complete on cycles [13]. Note that there are two aspects of the MPCP: *path selection* and *coloring* of the chosen paths. As the path selection on trees (and on cycles) is a trivial task, the above result show that the coloring alone is hard enough.

One of the first papers on the MPCP was by Aggarwal et al. [1]. They showed that $O(\log n)$ colors are sufficient to route a permutation in hypercubic networks, improving a bound of Pankaj [30]. Raghavan and Upfal [33] gave several approximation algorithms for this problem including a constant factor approximation for trees, rings and trees of rings. Mihail et al. [28] proved that a constant factor approximation can also be achieved for directed versions of these topologies. Kleinberg and Tardos [25] gave an $O(\log n)$ approximation algorithm for meshes and other certain classes of planar graphs. Rabani [31] showed that a constant factor approximation is also possible for mesh topologies, improving previously known bounds [24, 3]. We refer the reader to a survey by Beauquier et al. [10] for other results about the path coloring problem.

Bartal and Leonardi [9] gave $O(\log n)$ competitive online algorithms for all the topologies mentioned above (i.e., meshes, trees, rings, ...) and showed that it was not possible to do better for meshes in the online setting. They also presented an $\Omega(\frac{\log n}{\log \log n})$ lower bound on the competitive ratio of any deterministic algorithm on trees. Later, also a randomized lower bound of $\Omega(\log d)$, for any online algorithm for the MPCP on a tree with diameter $d = O(\log n)$, was given [27]. Bartal et al. [8] showed that a polylogarithmic competitive ratio is not possible for this problem on general topologies.

Outline of the paper. Since the outlined problems are clearly related to flow problems it is not surprising that for the analysis of our algorithm we will use some results from this area, e.g., the Shortening Lemma and the Duality of so-called multiroute flows and cuts. The other techniques that are used include probabilistic arguments, namely the Chernoff bounds, Lovász Local Lemma and its algorithmic version. All these will be reviewed in the next section. Section 3 deals with the short length Menger's Theorem and Section 4 with the algorithm for the k -EDPCOL. Section 5 concludes with a few open problems.

An abstract of this paper appeared in Proceedings of the Fifteenth Annual ACM Symposium on Parallelism in Algorithms and Architectures, 2003 [5].

2 Preliminaries

A *network* is an undirected graph $G = (V, E)$ together with a *capacity* function $c : E \rightarrow R^+$. Most of the time, we consider networks with unit capacity on each edge and we call them *unit networks*. Let $n = |V|$ and $m = |E|$. We denote the degree of a vertex u in G by $\deg(u)$. If $c(U, V - U)$ denotes the sum of capacities of edges between U and $V - U$, the (edge) *expansion* of a network G is defined as

$$\alpha = \min_{U \subset V} \frac{c(U, V - U)}{\min\{|U|, |V - U|\}}.$$

A *flow* (single- or multi-commodity) is a nonnegative linear combination of unit flows along simple paths; note that a flow may possibly exceed the edge capacities. A *feasible flow* is a flow that respects the capacity constraints. We will slightly overload the term flow and we will use it also to denote the total amount of the commodity transferred.

In a *concurrent multicommodity flow problem* on a network G there are $l \geq 1$ commodities, each with two terminal vertices $s_i, t_i \in V$ and a demand d_i . The objective is to maximize the fraction of the demand that can be shipped simultaneously for all commodities, that is, to find the maximum f such that it is possible to route $f \cdot d_i$ units of commodity i between s_i and t_i for each i so that the total amount of all commodities passing through any edge is no greater than its capacity. Given an instance of the concurrent multicommodity flow problem and a feasible flow \mathcal{F} , the *flow value* of \mathcal{F} is the maximum f such that for each i , $f \cdot d_i$ units of commodity i are shipped between s_i and t_i . An instance of a *balanced* multicommodity flow problem satisfies an additional condition that for each vertex u the sum of demands between u and other vertices is equal the degree $d(u)$ [26].

Given a unit network G , let I_0 denote an instance of the concurrent multicommodity flow problem in which there is a commodity with demand $\deg(u) \cdot \deg(v)/2|E|$ for each pair of vertices (u, v) . For a feasible flow \mathcal{S} , let $D(\mathcal{S})$ be the length of the longest flow path in \mathcal{S} and let $C(\mathcal{S})$ be the inverse of the flow value of \mathcal{S} (i.e., the maximum, over all commodities, of the flow divided by the demand). Then the *flow number* F of G is the minimum of $\max\{D(\mathcal{S}), C(\mathcal{S})\}$ over all feasible solutions \mathcal{S} of I_0 [26]. We always have $F = O(\Delta\alpha^{-1} \log n)$, where Δ is the maximum degree of G , but sometimes F is smaller by a factor Δ or $\log n$ [26]. To give at least a few examples, the flow number of constant degree expanders, hypercubes, butterflies etc. is $O(\log n)$, $F(2D\text{-mesh}) = O(\sqrt{n})$, $F(3D\text{-mesh}) = O(\sqrt[3]{n})$, $F(K_n) = O(1)$.

Lemma 2.1 (Shortening Lemma [26]) *For any network with flow number F it holds: for any $\epsilon \in (0, 1]$ and any feasible flow \mathcal{S} with a flow value of f , with respect to an instance of the concurrent multicommodity flow problem, there exists a feasible flow \mathcal{S}' with a flow value $f/(1 + \epsilon)$ that uses paths of length at most $2 \cdot F(1 + 1/\epsilon)$.*

To prove the second part of Lemma 3.1 in the next section, we will use the construction from the original proof of the Shortening Lemma. Therefore, for the sake of completeness, we will provide the proof of the Shortening lemma [26].

Proof. Let \mathcal{O} denote the set of paths in the flow \mathcal{S} with the flow value f and let $\mathcal{O}' \subseteq \mathcal{O}$ consist of all paths from \mathcal{O} that are longer than L , for $L = 2 \cdot F/\epsilon$. We are going to shorten the paths in \mathcal{O}' at the cost of slightly decreasing the satisfied demand of each commodity. To do so, we need the following lemma:

Claim 2.2 [26] *For any network G with flow number F and any instance of the balanced multicommodity flow problem for G , there exists a feasible flow with flow value $1/2F$ consisting of paths of length at most $2F$.*

For a path $p \in \mathcal{O}'$ between s_p and t_p , let $a_{p,1} = s_p, a_{p,2}, \dots, a_{p,L}$ denote its first L vertices and $b_{p,1}, \dots, b_{p,L-1}, b_{p,L} = t_p$ its last L vertices and let f_p be the flow along p . Then the set $\mathcal{U} = \bigcup_{p \in \mathcal{O}'} \bigcup_{i=1}^L \{a_{p,i}, b_{p,i}, f_p\}$ is (a subset of) an instance of the balanced multicommodity flow problem. Thus, by Claim 2.2 there exists a feasible flow \mathcal{P} with flow value at least $1/(2F)$, with respect to the instance \mathcal{U} , consisting of paths of length at most $2F$. We are going to combine the initial and final parts of the long paths in \mathcal{O}' with these “shortcuts” in \mathcal{P} to obtain the desired short solution.

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

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

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

For each i , the length of each path in the flow $S_{p,i}$ is at most $L + 2 \cdot F$, and $f_p/(L(1 + \epsilon))$ units of flow are shipped along each flow $S_{p,i}$. 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 flows $S_{p,i}$ 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{O} 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{O} and \mathcal{P} sum up to at most $c(e)$ for an edge e . Therefore, the modification yields a feasible flow satisfying the desired properties. \square

Multiroute flows and cuts. For simplicity we will slightly overload the term k -system, defined in the introduction, and we will use it to also denote a flow along k edge disjoint paths from s to t , each path carrying the same amount of flow. A *unit k -system* is a k -system that carries one unit of flow along each path, in total k units of flow. A k -flow is a (single or multicommodity) flow that is a non-negative linear combination of unit k -systems ([7], cf. [23, 2]). The *size* of a k -system is the total number of edges used by the k paths. The size of a k -system is, in a manner of speaking, the length of the k -system; it has nothing to do with the amount of a flow carried by it. Figure 1 gives an example of a 3-flow.

For a single commodity, an $s - t$ cut is a partition of vertices V into two parts S and $\bar{S} = V \setminus S$ such that $s \in S$ and $t \in \bar{S}$. It will be convenient to view the cut as a set of edges $\{(u, v) \mid u \in S \wedge v \in \bar{S}\} = \{e_1, \dots, e_l\}$. The *capacity* of the cut is equal to $\sum_{i=1}^l c(e_i)$. The k -*capacity* of a cut $\{e_1, \dots, e_l\}$ is defined as the maximum k -flow in a simplified network made up just of a source and a destination vertices directly connected by l edges with respective capacities $c(e_1), \dots, c(e_l)$. Note that 1-capacity of a cut corresponds to the capacity of the cut. Sometimes we will talk about a k -cut instead of a cut to stress that we are interested in the k -capacity of the cut.

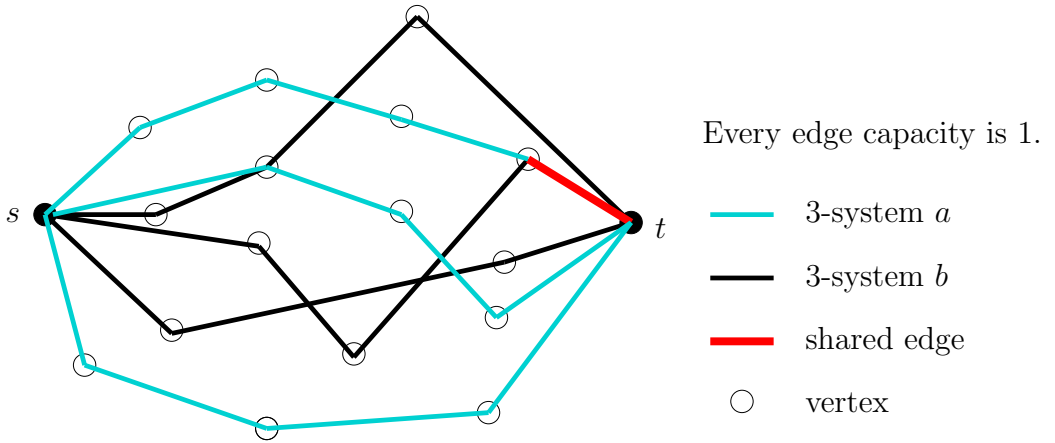


Figure 1: a and b are two 3-systems sharing one edge. The size of a is 11 and the size of b is 13. Assuming a flow of 0.3 units through each path of a and a flow of 0.4 units through each path of b , the total 3-flow between s and t is $0.3 \times 3 + 0.4 \times 3 = 2.1$ units.

Lemma 2.3 (k -capacity of a cut [7]) *Given a cut with l edges with capacities c_1, \dots, c_l such that $c_i \geq c_{i+1}$ (i.e., nonincreasing order), its k -capacity is equal to*

$$\min_{j=0, \dots, k-1} \frac{k}{k-j} \sum_{i=j+1}^l c_i .$$

The celebrated result about the duality of simple flows and cuts holds for k -flows and k -cuts, too.

Lemma 2.4 (Duality of single commodity k -flows and k -cuts [7, 23]) *The maximum feasible k -flow in G is equal to the capacity of the minimum k -cut in G .*

Probabilistic tools.

Lemma 2.5 (Chernoff Bound) *Consider any set of n independent binary random variables X_1, \dots, X_n . Let $X = \sum_{i=1}^n X_i$ and μ be chosen so that $\mu \geq E[X]$. Then it holds for all $\delta \geq 0$ that*

$$\Pr[X \geq (1 + \delta)\mu] \leq e^{-\min[\delta^2, \delta] \cdot \mu/3} .$$

Lemma 2.6 (Lovász Local Lemma) *Let A_1, \dots, A_n be “bad” events in an arbitrary probability space. Suppose that each event is mutually independent of all other events but at most d , and that $\Pr[A_i] \leq p$ for all i . If $ep(d + 1) \leq 1$, then the probability of no bad event occurring is greater than 0.*

Theorem 2.7 (Algorithmic LLL [29, 34]) *Let $\mathcal{T} = \{t_1, \dots, t_n\}$ be a set of independent random trials, and let $\mathcal{A} = \{A_1, \dots, A_m\}$ be a set of events such that each A_i is determined by the outcome of the trials in $T_i \subseteq \mathcal{T}$. For any $t_{j_1}, \dots, t_{j_k} \in T_i$ and any w_{j_1}, \dots, w_{j_k} in the domains of t_{j_1}, \dots, t_{j_k} , let $\Pr^*[A_i \mid t_{j_1} = w_{j_1}, \dots, t_{j_k} = w_{j_k}]$ be the probability of A_i conditional on the event that the outcomes of t_{j_1}, \dots, t_{j_k} are w_{j_1}, \dots, w_{j_k} .*

If we have the following:

1. for each $1 \leq i \leq m$, $\Pr[A_i] \leq p$;
2. for each T_i , there are less than d other T_j 's such that $T_i \cap T_j \neq \emptyset$;

3. for each $1 \leq i \leq m$, $|T_i| \leq w$;
4. $p \cdot d^9 \leq (1/2e)^3$;
5. $p \cdot w \leq 1$;
6. for each $1 \leq j \leq n$, we can carry out the random trial in time τ_1 ;
7. for each $1 \leq i \leq m$, $t_{j_1}, \dots, t_{j_k} \in T_i$ and w_{j_1}, \dots, w_{j_k} in the domains of t_{j_1}, \dots, t_{j_k} , we can compute every $\text{Pr}^*[A_i | *]$ in time τ_2 ;

then there is a randomized $O(n \cdot d \cdot (\tau_1 + \tau_2) + n(\tau_1 \cdot \log^{O(1)} m + 2^{(\log \log m)^{O(1)}}))$ -time algorithm which will find outcomes for t_1, \dots, t_n such that none of the events in \mathcal{A} holds.

3 A Bounded Length Version of Menger's Theorem

In this section we give an upper bound on the size of a minimum k -system connecting any two k -connected vertices. Our bound is in terms of the flow number F of a graph: we show that any two k -connected vertices are connected by k disjoint paths of average length $O(kF)$, that is, by a k -system of size $O(k^2F)$. To show this we first prove a much more general lemma which implies the bound. The lemma in its general form will serve as a crucial tool for the approximation algorithm for the k -EDPCOL in Section 4. The lemma is of a similar flavor as the Shortening lemma and is, in a sense, a generalization of it: what the Shortening lemma states about flows, Lemma 3.1 states about k -flows.

Lemma 3.1 *Given a unit network with flow number F , a set T of pairs of vertices and a feasible flow \mathcal{F} such that there are k units of flow between each pair from T , there exists a k -flow $\bar{\mathcal{F}}$ such that*

- there are k units of flow between each pair from T ,
- the flow through every edge is at most 4,
- each k -system used in $\bar{\mathcal{F}}$ has size at most $20 \cdot k^2F$.

Moreover, if the k -flow \mathcal{F} is integral (i.e., each pair from T is connected by a unit k -system), then there exists a k -flow $\bar{\mathcal{F}}$ such that

- the flow in $\bar{\mathcal{F}}$ through every edge is at most 2, and
- each k -system used in $\bar{\mathcal{F}}$ has size at most $8 \cdot k^2F$.

Proof. For $k = 1$ the claim follows immediately from the Shortening lemma. Thus, we assume that $k \geq 2$ for the rest of the proof. A k -system is called *small* if its size is at most $20 \cdot k^2F$.

Since \mathcal{F} can be viewed as a feasible solution to an instance of the concurrent multicommodity flow problem, the Shortening lemma with parameter $\epsilon = 1/(2k)$ gives a feasible flow \mathcal{F}' with flow paths of length at most $4kF + 2F$. We scale up the flow \mathcal{F}' by a factor of $1 + 1/(2k)$ to ensure that the total amount of flow between each pair of vertices from T is equal to k again. The flow in \mathcal{F}' may not be feasible but the flow in each of the edges is at most $1 + 1/(2k)$.

The goal is to transform the flow for each commodity i into a k -flow along small k -systems, while keeping the amount k units of the flow and not violating the capacity constraints “much”. For a while, we will consider the flow for each commodity i separately. Let \mathcal{F}_i denote the flow in \mathcal{F}' corresponding to the commodity i .

Claim 3.2 *The flow \mathcal{F}_i can be decomposed into a k -flow of $\frac{k+1}{2}$ units and a flow of $\frac{k-1}{2}$ units.*

Proof. Consider a network G_i that has the same set of vertices and edges as G , but the capacity of an edge e in G_i is equal to the flow through e in \mathcal{F}_i . We are going to show that the minimum k -capacity of an $s_i - t_i$ cut in G_i is at least $(k+1)/2$. By Lemma 2.4, this implies Claim 3.2,

Let $c_1 \geq c_2 \geq \dots \geq c_l$ be the capacities of edges in an $s_i - t_i$ cut with minimum k -capacity in G_i . From the characterization of the k -capacity of a cut given in Lemma 2.3 we know that its size is equal to

$$\min_{j=0, \dots, k-1} \frac{k}{k-j} \sum_{h=j+1}^l c_h .$$

By the construction of G_i , $\sum_{h=1}^l c_h \geq k$ and $c_h \leq 1 + 1/(2k)$ for every h . This implies, for every $j < k$, that

$$\frac{k}{k-j} \sum_{h=j+1}^l c_h \geq \frac{k}{k-j} \left(k - j \cdot \left(1 + \frac{1}{2k} \right) \right) = k - \frac{1}{2} \cdot \frac{j}{k-j} \geq \frac{k+1}{2} ,$$

which concludes the proof. \square

Given the k -flow from Claim 3.2 with $\frac{k+1}{2}$ units of flow, let \mathcal{F}'_i be a subset of small k -systems participating in this k -flow. So far, there is no guarantee that such small k -systems exist. We will prove that \mathcal{F}'_i is not empty and, moreover, that the k -systems in \mathcal{F}'_i cumulatively carry at least $\frac{k+2}{4}$ units of flow. This will almost complete the proof, since putting together the k -flows \mathcal{F}'_i for all i and scaling them up by four, we get the desired k -flow $\bar{\mathcal{F}}$.

Let us consider the total volume of the network G_i , that is, the sum $\sum_{e \in E} f_i(e)$, where $f_i(e)$ denotes the amount of flow belonging to commodity i through edge e in \mathcal{F}' . Due to the use of the Shortening lemma at the beginning of the proof, each path between s_i and t_i is at most $4kF + 2F$ edges long and the total flow being carried by these paths is k units. The total volume of flow between s_i and t_i is at most $4k^2F + 2kF \leq 5k^2F$ in \mathcal{F}' , which is also the total volume of G_i . Thus, the total flow through k -systems from \mathcal{F}'_i of size larger than $20k^2F$ is at most $k/4$. We conclude that at least $\frac{k+1}{2} - \frac{k}{4} = \frac{k+2}{4}$ units of flow are carried by k -systems from \mathcal{F}'_i .

At this point we put together the k -systems from all \mathcal{F}'_i , scale them up by $\frac{4k}{k+2}$ and denote the result by $\bar{\mathcal{F}}$. By construction, the resulting flow is at each edge e at most $\frac{4k}{k+2} \cdot \left(1 + \frac{1}{2k} \right) \leq 4$, the flow for each commodity is k units and only small k -systems are involved.

Better bounds for integral k -flow \mathcal{F} . To improve the bounds we dip into the proof of the Shortening lemma. We say that an edge is a *base* edge for commodity i if e is one of the first $L/2 = F/\epsilon$ or the last $L/2 = F/\epsilon$ edges on some of the k paths between s_i and t_i in \mathcal{S} . The observation is that the flow in \mathcal{F}_i on any edge that is not base is at most ϵ , every path in \mathcal{F}_i uses at most $2F$ non-base edges, and for integral \mathcal{F} , there are at most $4k^2F$ base edges in G_i . Thus, the total capacity of non-base edges in G_i is at most $2kF$, and therefore, the total flow through k -systems from \mathcal{F}_i that use more than $8kF$ non-base edges is at most $k/4$. Utilizing the fact that there are at most $4k^2F$ base edges, we conclude that at least $(k+2)/4$ units of flow from \mathcal{F}_i are carried by k -systems of size at most $8k^2F$.

The other observation from the original proof of the Shortening lemma is that for integral \mathcal{F} , the flow on every edge e in \mathcal{F}' is either at most ϵ , or at least 1 and a unit of this flow belongs to one commodity. We also note, that for every edge e and every commodity i , the flow through e in \mathcal{F}'_i is at most $1/2$. Thus, for every edge e the sum of flows in all \mathcal{F}'_i through e is at most $1/2 + 1/(2k)$, and the final flow through e in $\bar{\mathcal{F}}$ is at most $\frac{4k}{k+2} \cdot \left(\frac{1}{2} + \frac{1}{2k} \right) \leq 2$. \square

Applying Lemma 3.1 on a single pair of k -connected vertices and using just any one of the final k -systems to carry all k units of flow gives the following result:

Corollary 3.3 (Bounded length Menger’s Theorem) *Given a graph G with flow number F and a pair of k -connected vertices u and v , there are k -edge disjoint paths between u and v of average length $8k \cdot F$.*

Since there are graphs with diameter $\Omega(F)$ [26], the result of Corollary 3.3 is asymptotically the best possible for general graphs, up to a factor of k .

An improvement for the k -EDP. Let us recall the online k -Bounded Greedy Algorithm (k -BGA) for the k -EDP [5]:

k -BGA(L):

Given a request r ,

- if it is possible to realize r by a k -system q of size at most L such that q is edge disjoint with all previously selected k -systems, then accept the request r and select q for it,
- else reject r .

Note that the problem of finding k edge-disjoint paths of total length at most L between the same pair of vertices can be reduced to the classical min-cost (integral) flow problem, which can be solved by standard methods in polynomial time [12, Chapter 4]. The k -BGA can therefore also be used offline as an approximation algorithm. It is worth mentioning that if the task were to find k edge disjoint paths *each* of length at most L/k , the problem would not be tractable (cf. [11]).

The previous best competitive ratio $O(k^3F)$ was achieved by the k -BGA with parameter $L = 20k^3F$ [6]. Lemma 3.1 yields a better result.

Corollary 3.4 *Given a graph G with flow number F , the competitive ratio for the k -EDP of the k -BGA with parameter $L = 8k^2F$ is $O(k^2F)$.*

Proof. The point is that Lemma 3.1 can be used to modify the optimal offline solution into a solution that uses (fractional) k -systems of size $8k^2F$ only. Then a standard charging argument (c.f. [6]) yields the desired competitive ratio. \square

4 An Algorithm for k -EDPCOL

This section describes a randomized approximation algorithm for k -EDPCOL. We begin in Section 4.1 by a linear relaxation for a version of k -EDPCOL that restricts the size of k -systems used and we relate the optimal value of this problem to the optimal value of k -EDPCOL. Section 4.2 outlines the actual algorithm and in Section 4.3 we fill in the details of the algorithm. Throughout this section, whenever we talk about an intersection, we mean an intersection on an edge.

4.1 Relaxation

Let us start with a “relaxation” of the problem. We put relaxation in quotes since the linear program described below is a relaxation of the k -EDPCOL regarding the *integrality* of the k -systems, but is more restrictive than k -EDPCOL regarding the *size* of the k -systems.

Let $T \subseteq V \times V$ be an instance of the k -EDPCOL problem on a graph $G(V, E)$. For each $(a_j, b_j) \in T$ let \mathcal{P}_j denote the set of all k -systems of size at most $8k^2F$ between a_j and b_j and let $\mathcal{P} = \bigcup \mathcal{P}_j$. Consider the following linear program where we have a variable $x(s)$ for each k -system s from \mathcal{P} denoting the

$1/k$ fraction of flow sent along the k -system (that is, there is a flow $x(s)$ along each of the k paths of s) and a variable c denoting a half of the maximum flow through an edge in the network:

$$\text{minimize } c \quad \text{s.t.} \tag{1}$$

$$\sum_{s \in \mathcal{P}: e \in s} x(s) \leq 2 \cdot c \quad \forall e \in E \tag{2}$$

$$\sum_{s \in \mathcal{P}_j} x(s) \geq 1 \quad \forall j \tag{3}$$

$$x(s) \geq 0 \quad \forall s \in \mathcal{P} \tag{4}$$

Lemma 4.1 *The optimal fractional solution to the linear program is a lower bound for the optimal integral solution for k -EDPCOL.*

Proof. Consider the optimal integral solution to the coloring problem and let C be the number of colors used. We are going to shorten the k -systems for each color separately. Since the k -systems of the same color are disjoint, it is possible to apply Lemma 3.1 to shorten them. Now we merge together the modified flows for all C colors and we get a feasible fractional solution of LP (1)-(4), that is, a (multicommodity) k -flow consisting of only small k -systems. Since the maximum flow through an edge is $2C$ by construction, the proof is completed. \square

Solving the linear program. Since the linear program has exponentially many variables, it is not possible to solve it directly in polynomial time. It would be possible to formulate it in a different way with only polynomially many variables (which would require some effort since global properties of the flow, the disjointness of the k paths of each k -system, etc. have to be guaranteed) and then solve it as a general LP. Another option is to exploit the equivalence between the optimisation and the separation problems and to use the Ellipsoid algorithm (a violated inequality can be found in polynomial time by a minimum cost flow algorithm). Here we use a more efficient way. Since we only aim at an approximation algorithm for the k -EDPCOL, it is sufficient to start with an approximation for our linear program. A good approximation for the linear program (1)-(4) can be obtained in polynomial time using the approximation algorithms for the concurrent multicommodity flow problem by Garg and Könemann [18] or by Fleischer [16]. Basically, the only modification is that a procedure for finding a shortest path is replaced by a procedure for finding a smallest k -system, and only k -systems of size at most $8k^2F$ are considered. The smallest k -system can be efficiently computed using algorithms for minimum cost flow problems. We also note that the approximate solutions obtained by these algorithms use at most m different k -systems for each commodity. For easy later reference we state this formally as a lemma (the constant $3/2$ stated here is good enough for our purposes, even though much better approximations are possible).

Lemma 4.2 *It is possible to find a $3/2$ approximation for the linear program (1)-(4) in polynomial time. Moreover, at most m different k -systems per commodity are used in the approximation.*

We stress once again that only *short* k -systems are used in the LP (1)-(4).

4.2 Algorithm

In a high level description, the structure of our algorithm is as follows:

1. Solve the LP (1)-(4) approximately. By Lemma 4.2 this can be done efficiently, and by Lemma 4.1, the maximum flow C through an edge in the approximate solution is (roughly) a lower bound on the number of colors needed in the optimal integral solution of k -EDPCOL.

2. Round the fractional solution to an integral one. In some cases (e.g., graphs with flow number $\log n$ and more, or instances of k -EDPCOL with $|T| \geq \frac{m \log n}{k}$), this can be done directly using the rounding scheme of Raghavan and Thompson [32]. In other cases an iterative rounding method can be utilized (cf. [26]). In both cases, the aim of the rounding process is to guarantee that every k -system intersects with at most $O(Ck^2F)$ other k -systems at the end of the rounding.
3. Now, the k -systems can be easily colored with $O(Ck^2F)$ colors (e.g., greedy coloring) which is within $O(k^2F)$ from optimum.

The only missing part is the rounding procedure. We will describe it in the next subsection. Here we just summarize the main result.

Theorem 4.3 *Given a graph G with flow number F , there exists a randomized $O(k^2F)$ -approximation algorithm for the k -EDPCOL problem on G .*

4.3 Randomized rounding

Since the optimal solution has to use at least one color, we assume that the approximate fractional solution is at least one.

Lemma 4.4 *Given a feasible fractional solution to the linear program (1)-(4) from Lemma 4.2, with objective value $C \geq 1$, there exists a randomized polynomial time algorithm that rounds it to an integral solution such that each k -system in it intersects with at most $O(Ck^2F)$ other k -systems.*

Proof. Let us start with the simpler case that can be solved directly by the method of Raghavan and Thompson [32]. We assume that $F \cdot |T| \geq \frac{m \log n}{k}$ which includes both the previously mentioned cases $F \geq \log n$ and $|T| \geq \frac{m \log n}{k}$. Observing that $C \geq \frac{k|T|}{m}$ we have $CF \geq \log n$.

Consider the following random experiment. Independently for each commodity, choose exactly one of its k -systems, according to the probability distribution corresponding to the flows (i.e., a k -system s with flow $k \cdot x(s)$ is chosen with probability $x(s)$).

For an edge e , the expected number $E[n_e]$ of chosen k -systems passing through e after this experiment, is $2C$ at most. By linearity of expectation, the expected number $E[n_s]$ of chosen k -systems intersecting with a chosen k -system s is bounded by

$$E[n_s] \leq \sum_{e \in s} E[n_e] \leq 16k^2FC .$$

For a chosen k -system s , this implies, by the Chernoff bound

$$\Pr[n_s \geq 32k^2FC] \leq e^{-16k^2FC/3} .$$

We distinguish two cases. If $|T|/\log |T| \leq m$, we further bound this probability by $n^{-8k^2/3} = o(|T|^{-1})$ (using the observation that $CF \geq \log n$) and conclude that with high probability no chosen k -system intersects with more than $32k^2FC$ other chosen k -systems. Similarly, if $|T|/\log |T| > m$, the probability $\Pr[n_s \geq 32k^2FC]$ can be upper bounded by $|T|^{-\frac{16k^3F|T|}{3m \log |T|}} = o(|T|^{-1})$ and again, we conclude that with high probability no chosen k -system intersects with more than $32k^2FC$ other chosen k -systems. This completes the description of the rounding procedure in the simpler case.

In the rest of the section we assume that $F \leq n$ and $|T| \leq m \log n$ (we can actually assume more but this will be sufficient for our purposes) and we are going to describe an iterative rounding process that will result in a set of k -systems for T such that each of them intersects at most $O(k^2FC)$ others.

Initially, we have fractional k -systems with k -units of flow for each commodity. We cannot round the fractional k -systems into integral ones in one step since the probability that a chosen k -system intersects with more than $32k^2FC$ other chosen k -systems would be too large to guarantee the desired solution. The idea is to round the k -systems only partially in each step and to guarantee that the dependency between them keeps decreasing. In other words, in each iteration we want to decrease the number of k -systems that are used for each commodity while keeping a total flow of k units for each commodity, and, in the meantime decreasing the maximum number of k -systems that intersect any other k -system. To help quantify this dependency, for a k -system s that survives through iteration i , let n_s^i denote the number of other k -systems intersecting with s at the end of iteration i .

We will describe the rounding process in two steps. First, we show using the Lovász Local Lemma that there exists a way to round the fractional solution to an integral one. Then, we will explain that the Algorithmic LLL can be used instead, yielding an efficient randomized rounding algorithm.

Let $v_0 = \ln m$ and $v_{i+1} = \ln^2 v_i$, for $i > 0$. Before starting the rounding process, we split every k -system with flow $f > k/v_0$ into $\lceil f \cdot v_0/k \rceil$ k -systems, each with flow k/v_0 . We do this for technical reasons. Splitting this way increases the flow through any edge by $1 + 1/v_0$ at most, (since C is an approximation, we will ignore this as it can be hidden in the constant), and the total number of k -systems increases by v_0 , at most. Since the approximate solution from Lemma 4.2 uses at most m k -systems with non-zero flow for each request, in total there are by now at most $v_0 m^2 \log m \leq m^2 \log^2 m$ k -systems.

For simplicity, we will assume in the rest of the section that v_{i+1} divides v_i , for all i for which v_{i+1} will be defined, and that v_0 is integral (without these assumptions, additional ceilings can be used to avoid these requirements and they will only increase the constant hidden in the O notation).

Initial rounding (iteration $i = 0$): For each request from T , consider the following random experiment: for each commodity, organize its k -systems with flow less than k/v_0 into groups in such a way that the total flow in each group is exactly k/v_0 (for this purpose, possibly split some k -systems). Independently for each group, choose exactly one of its k -systems, according to the probability distribution corresponding to the flows (i.e., a k -system s with flow $k \cdot x(s)$ is chosen with probability $x(s) \cdot v_0$).

Now, we use a similar argument as before. For an edge e , the expected number $E[n_e^0]$ of chosen k -systems passing through e after the initial iteration, is $2Cv_0$. By linearity of expectation, the expected number $E[n_s^0]$ of chosen k -systems intersecting with a chosen k -system s is bounded by

$$E[n_s^0] \leq \sum_{e \in s} E[n_e^0] \leq 16k^2FCv_0 .$$

For a chosen k -system s , this implies, by the Chernoff bound (to simplify notation, we introduce a new variable $D = 16k^2FC$) that

$$\Pr[n_s^0 \geq 2Dv_0] \leq e^{-Dv_0/3} = m^{-D/3} .$$

Since $e \cdot m^{-D/3} \cdot (1 + m^2 \log^2 m) < 1$, the Lovász Local Lemma guarantees that there exists a random choice such that the maximum number of k -systems intersecting any other k -system is at most $2Dv_0$. We send k/v_0 units of flow along the chosen k -systems (that is, there is a flow $1/v_0$ along each path) and zero along the other k -systems.

Let $a_0 = 2$, and for $i \geq 1$ let $a_i = (1 + 1/\sqrt[4]{v_{i-1}})a_{i-1}$. Observe that all a_i 's are bounded by an absolute constant (using $1 + x \leq e^x$, e.g.).

Intermediate rounding (iteration $i > 0$): The input for iteration i is a set of k -systems such that the maximum number of k -systems intersecting any other k -system is at most $a_{i-1}Dv_{i-1}$, each k -system carries flow k/v_{i-1} , and there are v_{i-1} k -systems for each request.

Similarly as in the initial rounding, consider the following random experiment: for each request, organize all its k -systems into groups, each of total flow exactly k/v_i . Independently for each group, choose uniformly at random exactly one of its k -systems. Our goal is to show that there exists a choice such that at the end of iteration i (which is the beginning of iteration $i + 1$), the maximum number of k -systems intersecting any other k -system is at most $a_i Dv_i$, each k -system carries flow k/v_i , and there are v_i k -systems for each request. Figure 2 schematically outlines the intermediate rounding steps.

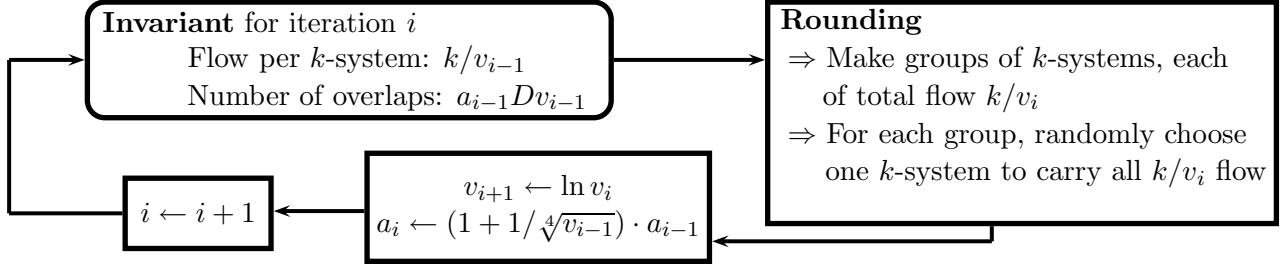


Figure 2: Schematic: Intermediate rounding

Since at the end of iteration $i - 1$ each k -system intersects with at most $a_{i-1} Dv_{i-1}$ other k -systems, and since each of the k -systems in iteration $i - 1$ survives to iteration i with probability v_i/v_{i-1} , we have

$$\mathbb{E}[n_s^i] \leq a_{i-1} Dv_i .$$

Using the Chernoff bound, it is possible to say more:

$$\Pr[n_s^i \geq (1 + 1/\sqrt[4]{v_i})a_{i-1} Dv_i] \leq e^{-\sqrt{v_i} a_{i-1} D/3} = v_{i-1}^{-a_{i-1} D/3} .$$

Now we apply the Lovász Local Lemma. For a group g of k -systems, let A_g denote the event that the k -system chosen from g in iteration i intersects with more than $(1 + 1/\sqrt[4]{v_i})a_{i-1} Dv_i$ other k -systems, at the end of iteration i . The event A_g depends only on the random outcomes in all groups that contain a k -system that intersects with a k -system from the group g . In each group there are v_{i-1}/v_i k -systems, each of them intersects with at most $a_{i-1} Dv_{i-1}$ other k -systems. In total, the event A_g depends on at most $(a_{i-1} Dv_{i-1}) \cdot v_{i-1}/v_i$ other events $A_{g'}$. Since

$$e \cdot v_{i-1}^{-a_{i-1} D/3} \cdot (a_{i-1} Dv_{i-1}^2/v_i + 1) < 1 ,$$

there exists a choice such that the maximum number of k -systems intersecting any other k -system is at most $a_i Dv_i$. We send k/v_i units of flow along each chosen k -system and nothing along all the other k -systems. Clearly, this guarantees the desired input for the next iteration.

Final rounding: In each iteration, the maximum number of k -systems intersecting any other k -system decreases exponentially. Thus, after several iterations, the maximum number of k -systems intersecting other k -system is at most $\text{poly}(kF)$ while there is still a flow of k units for each request. At this point we perform the randomized rounding once more: for each request, choose uniformly at random one of its k -systems. Since the dependency between the k -systems is bounded by $\text{poly}(kF)$, the Lovász Local Lemma guarantees existence of the desired k -systems.

To complete the proof, it is sufficient to show that the Algorithmic LLL can be used instead of the Lovász Local Lemma in each iteration. In our case in iteration i , both the random trials and the (bad) events will be indexed by the groups of k -systems: t_g is the random trial of choosing one k -system from group g and A_g is the event that the k -system chosen in group g intersects with more than $(1 + 1/\sqrt[4]{v_i})a_{i-1}k^2FCv_i$ other k -systems. Observe, that

- each random trial can be carried out time $\tau_1 = O(v_{i-1}/v_i)$
- $Pr^*[A_i|*]$ can be computed in time $\tau_2 = O(k^2FCv_{i-1}^2) = \text{poly}(m)$
- for $p = v_{i-1}^{-a_{i-1}D/3}$, $w = a_{i-1}v_{i-1}^2/v_i$ and $d = w^2$, all other assumptions of the Algorithmic LLL are satisfied.

Thus, we can use the Algorithmic LLL in each iteration. □

5 Conclusion and Open Problems

One of the main contributions of this paper is the result on “short” k -flows for multicommodity problems. There is an interesting open question here: Is it possible to generalize Lemma 3.1 to networks with nonuniform edge capacities? In our proof of this lemma (namely in the proof of Claim 3.2) the fact that all original capacities are unit is needed to prove that there exists a cut of large k -capacity for each commodity. Can this assumption be avoided?

A corollary of this general result is a new upper bound, $O(kF)$, on the average path length between two k -connected vertices a graph with flow number F . There is, however, still a gap between this upper bound and the obvious lower bound of $\Omega(F)$. We think that the gap might be an artifact of the analysis that can be removed. This would lead to an $O(kF)$ approximation ratio both for the k -EDP and k -EDPCOL problems.

The presented algorithm for k -EDPCOL problem is an offline one. This immediately brings in mind a question whether k -systems for the requests can be picked up and colored online? Is it possible to achieve a $O(k^2F)$ competitive ratio or even the $O(kF)$ ratio?

Acknowledgments. We would like to thank Ankur Bhargava and Christian Scheideler for many useful discussions, and the three anonymous referees for their several useful suggestions.

References

- [1] A. Aggarwal, A. Bar-Noy, D. Coppersmith, R. Ramaswami, B. Scheiber, and M. Sudan. Efficient routing and scheduling algorithms for optical networks. In *Proceedings of the 5th ACM-SIAM Symposium on Discrete Algorithms*, pages 412–423, 1994.
- [2] C. C. Aggarwal and J. B. Orlin. On multiroute maximum flows in networks. *Networks*, 39:43–52, 2002.
- [3] Y. Aumann and Y. Rabani. Improved bounds for all optical routing. In *Proceedings of the 6th Annual Symposium on Discrete Algorithms*, pages 567–576, 1995.
- [4] A. Bagchi, A. Chaudhary, M. T. Goodrich, and S. Xu. Constructing disjoint paths for secure communication. In *Proceedings of the 17th International Conference on Distributed Computing*, volume 2848 of *Lecture Notes in Computer Science*. Springer Verlag, 2003.

- [5] A. Bagchi, A. Chaudhary, and P. Kolman. Short length Menger's theorem and reliable optical routing. In *Proceedings of the 15th Annual ACM Symposium on Parallelism in Algorithms and Architectures*, pages 246–247, 2003.
- [6] A. Bagchi, A. Chaudhary, P. Kolman, and C. Scheideler. Algorithms for fault-tolerant routing in circuit switched networks. In *Proceedings of 14th Annual ACM Symposium on Parallel Algorithms and Architectures*, pages 265–274, 2002.
- [7] A. Bagchi, A. Chaudhary, P. Kolman, and J. Sgall. A simple combinatorial proof for the duality of multiroute flows and cuts. Technical Report 2004-662, Charles University, Prague, 2004. Available online at <http://kam.mff.cuni.cz/~kamserie>.
- [8] Y. Bartal, A. Fiat, and S. Leonardi. Lower bounds for on-line graph problems with application to on-line circuit and optical routing. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing*, pages 531–540, 1996.
- [9] Y. Bartal and S. Leonardi. On-line routing in all-optical networks. In *Proceedings of the 24th International Colloquium on Automata, Languages, and Programming*, volume 1256 of *Lecture Notes in Computer Science*, pages 516–526. Springer-Verlag, 1997.
- [10] B. Beauquier, J.-C. Bermond, L. Gargano, P. Hell, S. Perennes, and U. Vaccaro. Graph problems arising from wavelength-routing in all-optical networks. In *Proceedings of the 2nd IEEE Workshop on Optics and Computer Science*, 1997. Available online at <http://citeseer.ist.psu.edu/beauquier97graph.html>.
- [11] A. Bley. On the complexity of vertex-disjoint length-restricted path problems. Technical Report SC-98-20, Konrad-Zuse-Institut Berlin, 1998.
- [12] W. J. Cook, W. H. Cunningham, W. R. Pulleyblank, and A. Schrijver. *Combinatorial Optimization*. John Wiley, New York, 1997.
- [13] T. Erlebach and K. Jansen. Scheduling of virtual connections in fast networks. In *Proceedings of the 4th Parallel Systems and Algorithms Workshop*, pages 13–32, 1996.
- [14] T. Erlebach and K. Jansen. Call scheduling in trees, rings and meshes. In *Proceedings of the 30th Hawaii International Conference on System Sciences*, 1997.
- [15] T. Erlebach and K. Jansen. The complexity of path coloring and call scheduling. *Theoretical Computer Science*, 255:33–50, 2001.
- [16] L. K. Fleischer. Approximating fractional multicommodity flow independent of the number of commodities. *SIAM Journal on Discrete Mathematics*, 13(4):505–520, Nov. 2000.
- [17] Z. Galil and X. Yu. Short length versions of Menger's theorem. In *Proceedings of the 27th annual ACM Symposium on Theory of Computing*, 1995.
- [18] N. Garg and J. Konemann. Faster and simpler algorithms for multicommodity flow and other fractional packing problems. In *Proceedings of the 39th Annual Symposium on Foundations of Computer Science*, pages 300–309, 1998.
- [19] M. C. Golumbic and R. E. Jamison. The edge intersection graphs of paths in a tree. *J. Combinatorial Theory*, Series B 38:8–22, 1985.

- [20] V. Guruswami, S. Khanna, R. Rajaraman, B. Shepherd, and M. Yannakakis. Near-optimal hardness results and approximation algorithms for edge-disjoint paths and related problems. In *Proceedings of the 31st ACM Symposium on Theory of Computing*, pages 19–28, 1999.
- [21] M. R. Henzinger, J. Kleinberg, and S. Rao. Short-length menger theorems. SRC Technical Note 1997 - 022, DEC, Systems Research Center, Nov. 1997.
- [22] D. R. Karger and M. S. Levine. Finding maximum flows in undirected graphs seems easier than bipartite matching. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing*, pages 69–78, 1998.
- [23] W. Kishimoto. A method for obtaining the maximum multiroute flows in a network. *Networks*, 27:279–291, 1996.
- [24] J. Kleinberg and É. Tardos. Approximations for the disjoint paths problem in high-diameter planar networks. In *Proceedings of the 27th Annual ACM Symposium on the Theory of Computing*, pages 26–35, 1995.
- [25] J. Kleinberg and É. Tardos. Disjoint paths in densely embedded graphs. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, pages 52–61, Los Alamitos, Oct. 1995. IEEE Computer Society Press.
- [26] P. Kolman and C. Scheideler. Improved bounds for the unsplittable flow problem. In *Proceedings of the 13th ACM-SIAM Symposium on Discrete Algorithms*, pages 184–193, 2002.
- [27] S. Leonardi and A. Vitaletti. Randomized lower bounds for online path coloring. In *Proc. of the 2nd International Workshop on Randomization and Approximation Techniques in Computer Science*, volume 1518 of *Lecture Notes in Computer Science*, pages 232–247. Springer Verlag, 1998.
- [28] M. Mihail, C. Kaklamanis, and S. Rao. Efficient access to optical bandwidth. In *Proceedings of the 36th Annual IEEE Symposium on Foundations of Computer Science*, pages 548–557, 1995.
- [29] M. Molloy and B. Reed. Further algorithmic aspects of the Lovasz Local Lemma. In *Proceedings of the 30th Annual ACM Symposium on the Theory of Computing*, pages 524–529, 1998.
- [30] R. K. Pankaj. *Architectures for Linear Lightwave Networks*. PhD thesis, MIT, 1992.
- [31] Y. Rabani. Path coloring on the mesh. In *Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science*, pages 400–409, 1996.
- [32] P. Raghavan and C. Thompson. Randomized rounding: A technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7:365–374, 1987.
- [33] P. Raghavan and E. Upfal. Efficient routing in all-optical networks. In *Proceedings of the 26th Annual Symposium on the Theory of Computing*, pages 133–143, 1994.
- [34] C. Scheideler. *Probabilistic Methods for Coordination Problems*. Habilitation thesis, University of Paderborn, July 2000.