

Length-Bounded Cuts and Flows

GEORG BAIER

Siemens

THOMAS ERLEBACH

University of Leicester

ALEXANDER HALL

Google

EKKEHARD KÖHLER

BTU Cottbus

PETR KOLMAN AND ONDŘEJ PANGRÁC

Charles University in Prague

HEIKO SCHILLING

TomTom

AND

MARTIN SKUTELLA

TU Berlin

Some of the results described in this article were presented at ICALP'06 (Baier et al. 2006).

This work was partly supported by the Federal Ministry of Education and Research (BMBF grant 03-MOM4B1), by the European Commission - Fet Open project DELIS IST-001907 (SBF grant 03.0378-1), and by the German Research Foundation (DFG grants MO 446/5-2 and SK 58/5-3). P. Kolman was partially supported by grant 201/09/0197 of GA ČR and by project 1M0021620808 (ITI) of the Ministry of Education of the Czech Republic. M. Skutella is supported by DFG Research Center MATHEON in Berlin.

Authors' addresses: G. Baier, Siemens AG, Corporate Technology, Otto-Hahn-Ring 6, 81739 Munich, Germany, e-mail: georg.baier@gmail.com; T. Erlebach, Department of computer Science, University of Leicester, University Road, Leicester LE1 7RH, U.K., e-mail: t.erlebach@mcs.le.ac.uk; A. Hall, Google Switzerland GmbH, Brandschenkestrasse 110, 8002 Zürich, Switzerland, e-mail: alex.hall@gmail.com; E. Köhler, Institute of Mathematics, BTU Cottbus, Postfach 10 13 44, 03013 cottbus, Germany, e-mail: ekoehler@math.tu-cottbus.de; P. Kolman and O. Pangrác, Department of Applied Mathematics, Faculty of Mathematics and Physics, Charles University in Prague, Malostranské nám. 25, 118 00, Prague 1, Czech Republic, e-mail {kolman, pangrac}@kam.mff.cuni.cz; H. Schilling, TomTom International B.V., Oosterdokstraat 114, 1001 DK Amsterdam, The Netherlands + Rudower Chaussee 29, 12489 Berlin, Germany + 20th Floor Euston Tower, 286 Euston Road, London, NW1 3AS, U.K., e-mail: heiko.schilling@gmail.com; M. Skutella, Institute of Mathematics, TU Berlin, Strasse des 17. Juni 136, 10623 Berlin, Germany, e-mail: skutella@math.tu-berlin.de.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute

to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2010 ACM 1549-6325/2010/11-ART4 \$10.00

DOI 10.1145/1868237.1868241 <http://doi.acm.org/10.1145/1868237.1868241>

Abstract. For a given number L , an L -length-bounded edge-cut (node-cut, respectively) in a graph G with source s and sink t is a set C of edges (nodes, respectively) such that no s - t -path of length at most L remains in the graph after removing the edges (nodes, respectively) in C . An L -length-bounded flow is a flow that can be decomposed into flow paths of length at most L . In contrast to classical flow theory, we describe instances for which the minimum L -length-bounded edge-cut (node-cut, respectively) is $\Theta(n^{2/3})$ -times ($\Theta(\sqrt{n})$ -times, respectively) larger than the maximum L -length-bounded flow, where n denotes the number of nodes; this is the worst case. We show that the minimum length-bounded cut problem is \mathcal{NP} -hard to approximate within a factor of 1.1377 for $L \geq 5$ in the case of node-cuts and for $L \geq 4$ in the case of edge-cuts. We also describe algorithms with approximation ratio $\mathcal{O}(\min\{L, n/L\}) \subseteq \mathcal{O}(\sqrt{n})$ in the node case and $\mathcal{O}(\min\{L, n^2/L^2, \sqrt{m}\}) \subseteq \mathcal{O}(n^{2/3})$ in the edge case, where m denotes the number of edges. Concerning L -length-bounded flows, we show that in graphs with unit-capacities and general edge lengths it is \mathcal{NP} -complete to decide whether there is a fractional length-bounded flow of a given value. We analyze the structure of optimal solutions and present further complexity results.

Categories and Subject Descriptors: G.2.2 [Discrete Mathematics]: Graph Theory—*Graph algorithms; network problems*; F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—*Routing and layout*; G.2.1 [Discrete Mathematics]: Combinatorics—*Combinatorial algorithms*

General Terms: Algorithms, Theory

Additional Key Words and Phrases: Approximation algorithms, flows and cuts in graphs, graphs

ACM Reference Format:

Baier, G., Erhlebach, T., Hall, A., Köhler, E., Kolman, P., Pangrác, O., Schilling, H., and Skutella, M. 2010. Length-bounded cuts and flows. *ACM Trans. Algor.* 7, 1, Article 4 (November 2010), 27 pages. DOI = 10.1145/1868237.1868241 <http://doi.acm.org/10.1145/1868237.1868241>

1. Introduction

In a classical article Menger [1927] showed a strong relation between cuts and systems of disjoint paths (Menger’s Theorem): given a graph G and two nodes s, t in G , the maximum number of edge- and node-disjoint s - t -paths equals the minimum size of an s - t -edge- and node-cut, respectively (cf. Dantzig and Fulkerson [1956]; Kotzig [1956]). Ford and Fulkerson [1956] and Elias et al. [1956] generalized the theorem of Menger to flows in graphs with capacities on the edges and provided algorithms to find an s - t -flow and an s - t -cut of the same value. (All results mentioned in this introduction hold both for directed and undirected graphs, unless stated otherwise.)

As far as we know, the problem of length-bounded flows was first considered by Adámek and Koubek [1971], who observed that a natural generalization of the max-flow min-cut theorem does not hold for length-bounded flows and gave some estimations on the value of a maximum length-bounded flow. Independently, Lovász et al. [1978] studied the maximum length-bounded node-disjoint s - t -paths problem. For length-bounds 2, 3, and 4, a relation holds that is analogous to Menger’s theorem but with a new cut definition. For length-bounds greater than 4, they gave upper and lower bounds for the gap between the maximum number of length-bounded node-disjoint paths and the cardinality of a minimum cut. Furthermore, they provided examples showing that some of the bounds are tight.

The results were extended independently to edge-disjoint paths by Exoo [1983] and Niepel and Safariková [1983]. Pyber and Tuza [1993] proved the following theorem, improving an earlier result of Lovász et al. [1978]: if the size of a minimum L -length bounded s - t -node-cut is k , then the number of node-disjoint s - t -paths of length at most $\binom{k+L-2}{L-2} + \binom{k+L-3}{L-2}$ is at least k .

According to Bondy and Murty [1976], Lovász conjectured that there is a constant γ such that the size of a minimum L -length-bounded s - t -node-cut is at most a factor of $\gamma \cdot \sqrt{L}$ larger than the cardinality of a maximum system of node-disjoint s - t -paths of length at most L . Boyles and Exoo [1982] disproved this conjecture. They constructed, for each length-bound $L > 0$, a graph and a node pair s, t such that the minimum L -length-bounded s - t -node-cut has size greater than $\gamma \cdot L$ times the maximum number of node-disjoint s - t -paths of length at most L , where the constant γ is roughly $1/4$. The ratio between the maximum number of node-disjoint s - t -paths and the size of a minimum length-bounded s - t -cut was also studied by Ben-Ameur [2000].

Itai et al. [1982] gave efficient algorithms to find the maximum number of node- and edge-disjoint s - t -paths with at most two or three edges; the node-disjoint case was also solved for length-bound 4. On the complexity side they showed that the maximum node- and edge-disjoint length-bounded s - t -paths problem is \mathcal{NP} -hard for length-bounds greater than 4. Van der Holst and de Pina [2002] proved that the problem is \mathcal{NP} -hard in planar graphs. Guruswami et al. [2003] showed that the edge-disjoint 6-length-bounded s - t -paths problem is MAX \mathcal{SNP} -hard, and for any length-bound they gave an $\mathcal{O}(\sqrt{m})$ -approximation algorithm, where m denotes the number of edges. Bley [2003] proved that both the node- and the edge-disjoint maximum 5-length-bounded s - t -paths problem are \mathcal{APX} -complete. For directed networks, Guruswami et al. [2003] showed that the problem is \mathcal{NP} -hard to approximate within a factor of $n^{\frac{1}{2}-\epsilon}$ for any $\epsilon > 0$, where n denotes the number of nodes.

For maximum fractional length-bounded multicommodity flows in unit-length graphs with general capacities, Baier [2003] proved that the maximum fractional length-bounded multicommodity flow can be computed exactly in polynomial time using linear programming methods. Independently, another polynomial-time exact algorithm for the same setting was given by Kolman and Scheideler [2006]; again, the algorithm exploits linear programming methods. For maximum fractional length-bounded multicommodity flows in graphs with general edge lengths and general capacities, Baier [2003] gave a fully polynomial time approximation scheme (FPTAS).

Mahjoub and McCormick [2003] presented a polynomial algorithm for the 3-length-bounded edge-cut in undirected graphs. Furthermore, they showed that the fractional versions of the length-bounded flow- and cut problem are polynomial even if L is part of the input but that the integral versions are strongly \mathcal{NP} -hard even if L is fixed.

Length-bounded paths problems arise naturally in a variety of real-world optimization problems and therefore many heuristics for finding large systems of length-bounded paths have been developed [Perl and Ronen 1984; 1996 Brandes et al. 1996; Wagner and Weihe 1995; Hsu 1994].

1.1. OUR CONTRIBUTION. In contrast to the classical flow theory, we describe instances for which the minimum L -length-bounded edge-cut is $\Theta(n^{2/3})$ -times

TABLE I. KNOWN AND NEW (BOLD TYPE) COMPLEXITY AND (IN)APPROXIMABILITY RESULTS

L	Node cut	Edge cut
1	—	poly.
2	poly.	poly.
3	poly.	poly. (undirected) [Mahjoub and McCormick 2003]
4	poly. (undirected) [Lovász et al. 1978]	inapprox. within 1.1377
$5 \dots \lfloor n^{1-\varepsilon} \rfloor$	inapprox. within 1.1377	inapprox. within 1.1377
arbitrary	$\mathcal{O}(\min\{L, n/L\})$ -approx. $\mathcal{O}(\sqrt{n})$ -approx.	$\mathcal{O}(\min\{L, n^2/L^2, \sqrt{m}\})$ -approx. $\mathcal{O}(n^{2/3})$ -approx.
$2 \cdot F(1 + 1/\varepsilon) \dots n$	$(1 + \varepsilon)$ -approx.	$(1 + \varepsilon)$ -approx.
$n - c$	poly.	

$\varepsilon \in \mathbb{R}_{>0}$ and $c \in \mathbb{N}$ are constants, ε can be arbitrarily small. F is the flow number of a graph. All results hold for the directed and undirected cases, unless stated otherwise.

larger than the maximum L -length-bounded flow, and instances for which the minimum L -length-bounded node cut is $\Theta(\sqrt{n})$ -times larger than the maximum L -length-bounded flow. In both cases we prove that this is the worst case, and we explain how this corresponds to the integrality gap of a natural linear programming formulation of the L -length-bounded cut problem. Further, we show that the minimum length-bounded cut problem is \mathcal{NP} -hard to approximate within a factor of 1.1377 for $L \geq 5$ in the case of node cuts and for $L \geq 4$ in the case of edge cuts; Table I provides an overview of known and new complexity results. We also give approximation algorithms of ratio $\mathcal{O}(\min\{L, n/L\}) \subseteq \mathcal{O}(\sqrt{n})$ in the node case and $\mathcal{O}(\min\{L, n^2/L^2, \sqrt{m}\}) \subseteq \mathcal{O}(n^{2/3})$ in the edge case. For instances with the length bound L larger than $2 \cdot F(1 + 1/\varepsilon)$ where F is the flow number of the graph [Kolman and Scheideler 2006] and ε is any constant larger than zero, we give $(1 + \varepsilon)$ -approximation algorithms for both the node cuts and edge cuts (e.g., for hypercubes and $L \geq 3F = \mathcal{O}(\log n)$ this yields a constant approximation). For length bounds $L = n - c$, where $c \in \mathbb{N}$ is a constant, we provide a polynomial time algorithm for the minimum L -length-bounded node-cut problem.

Concerning L -length-bounded flows, we show that in graphs with unit capacities and general edge lengths it is \mathcal{NP} -complete to decide whether there is a fractional length-bounded flow of a given flow value. Even worse, the edge representation and the path representation of an L -length-bounded (fractional) flow are not polynomially equivalent. In particular, for graphs with general edge lengths we prove that there is no polynomial algorithm which transforms an edge representation of an L -length-bounded flow into a path representation, unless $\mathcal{P} = \mathcal{NP}$. We analyze the structure of optimal solutions and give instances where each maximum flow ships a large percentage of the flow along paths with an arbitrarily small flow value. We also provide a lower bound of $\Omega(\sqrt{n})$ on the integrality gap of the linear programming formulation of the maximum L -length-bounded flow (we remark that for some instances the size of the linear program is exponential in the size of the graph). The integrality gap applies even for planar graphs.

2. Preliminaries

2.1. GRAPHS. We consider both directed and undirected graphs; the number of nodes of a graph is denoted by n and the number of edges is denoted by m . There

are two independent functions associated with each graph G , an *edge-capacity* function $u : E \rightarrow \mathbb{Q}_{>0}$ and an *edge-length* function $d : E \rightarrow \mathbb{Q}_{\geq 0}$. We denote the capacity and length of an edge $e \in E$ by u_e and d_e , respectively. Unless stated otherwise, the length of each edge is 1 (*unit lengths*) and the capacity of each edge is 1 (*unit capacities*). The *length of a path* is the sum of the lengths of the edges on the path. The *distance* between two nodes u and v , denoted $\text{dist}(u, v)$, is the length of a shortest path from u to v .

A *multigraph* is a graph that may contain several edges between two vertices; such edges are called *multi-edges* or *parallel edges*. Occasionally, we use the term *simple graph* to stress that we are dealing with a graph that does not have multiedges. Unless stated otherwise, our results apply to simple graphs.

2.2. LENGTH-BOUNDED CUTS. Let $s, t \in V$ be two distinct nodes in a graph $G = (V, E)$. We call a subset of edges $C \subseteq E$ of G an *s-t-edge-cut*, if no path remains from s to t in the graph $(V, E \setminus C)$. The *value* of C is the sum of the capacities of the edges in C , that is, $\sum_{e \in C} u_e$. In the case of unit capacities, the value of a cut is also called its *size*. Similarly, a node set $C \subseteq V$ of G that separates s and t (and contains neither s nor t) is an *s-t-node-cut*; its *value* (or *size*) is the number of nodes in C .

Let $\mathcal{P}_{s,t}(L)$ denote the set of all *s-t*-paths of length at most L . We call a subset of edges $C \subseteq E$ of G an *L-length-bounded s-t-edge-cut* if the nodes s and t have a distance greater than L in the graph $(V, E \setminus C)$. This means that C must hit every path in $\mathcal{P}_{s,t}(L)$. Similarly, a subset C of the node set of G is called an *L-length-bounded s-t-node-cut* if it hits all paths in $\mathcal{P}_{s,t}(L)$. All of our cuts are *s-t*-cuts (for some nodes s, t) and therefore we will often omit the *s-t*-prefix when talking about cuts. The *value* of a length-bounded cut is defined in the same way as in the standard cut case. In the *minimum length-bounded cut* problem we are looking for an *L*-length-bounded cut of minimum value.

In a linear programming relaxation of the minimum length-bounded edge-cut problem, one has to assign to each edge $e \in E$ a nonnegative value ℓ_e , called its *dual length*, such that the dual length of every path from $\mathcal{P}_{s,t}(L)$ is at least one (the linear programming relaxation for node-cuts is analogous):

$$\begin{aligned} \min \quad & \sum_{e \in E} u_e \ell_e & (1) \\ & \sum_{e \in P} \ell_e \geq 1 \quad \forall P \in \mathcal{P}_{s,t}(L) \\ & \ell_e \geq 0 \quad \forall e \in E \end{aligned}$$

An integral solution to this linear program corresponds to a length-bounded *s-t*-cut, and vice versa. In particular, the minimum length-bounded *s-t*-cut value and the value of a minimum integral solution are equal. We will refer to feasible solutions of (1) as *fractional edge-cuts*; *fractional node-cuts* are defined similarly as feasible solutions of the analogous linear programming relaxation of the length-bounded node-cut problem.

2.3. LENGTH-BOUNDED FLOWS. An *s-t-flow* in a directed graph $G = (V, E)$ is a function $f : E \rightarrow \mathbb{Q}_{\geq 0}$ satisfying for each vertex $v \in V \setminus \{s, t\}$ the flow conservation constraint $\sum_{(u,v) \in E} f(u, v) = \sum_{(v,u) \in E} f(v, u)$ [Ahuja et al. 1993]; we call the function f an *edge representation* of the flow. A classical theorem by

Ford and Fulkerson [1962] states that every flow can be represented as a nonnegative linear combination of unit flows along cycles and s - t -paths in G ; we call such a linear combination a *path representation* of the flow. The flow f is *feasible*, if all capacity constraints are satisfied, that is, $f(e) \leq u_e$ for each edge $e \in E$. The *value* (or *size*) of the flow f is the quantity $\sum_{(u,t) \in E} f(u,t) - \sum_{(t,u) \in E} f(t,u)$. For undirected graphs, the definitions are analogous. If $f(e)$ is integral for every $e \in E$, the flow f is said to be *integral with respect to the edge representation*. If f can be decomposed into integral flows along cycles and s - t -paths, it is said to be *integral with respect to the path representation*.

Length-bounded flows are flows that can be decomposed into flows along paths of bounded length. More precisely, an L -length-bounded s - t -flow is an s - t -flow f such that there exists a decomposition of f into flows along s - t -paths of length at most L ; such a decomposition is called an L -length-bounded *path representation* of the flow.

A natural optimization problem is to find a feasible length-bounded s - t -flow of maximum value. We can formulate this problem as a linear program:

$$\begin{aligned} \max \quad & \sum_{P \in \mathcal{P}_{s,t}(L)} f_P & (2) \\ & \sum_{P: e \in P} f_P \leq u_e \quad \forall e \in E \\ & f_P \geq 0 \quad \forall P \in \mathcal{P}_{s,t}(L). \end{aligned}$$

Note that the dual of (2) is the linear program (1). One way to prove the maximum-flow minimum-cut equality for standard flows is to apply the duality of linear programming and to observe that there always exists an *integral* optimal solution (which does not hold in the length-bounded case).

In the case of multiple commodities, we are given k source-sink node pairs $(s_1, t_1), \dots, (s_k, t_k)$ called *commodities*. A *multicommodity flow* f is a set of s_i - t_i -flows f_i , for $i = 1, \dots, k$. The multicommodity flow f is *feasible* if for each edge $e \in E$ the capacity constraint holds, that is, $\sum_{i=1}^k f_i(e) \leq u_e$. An L -length-bounded *multi-commodity flow* f is a multicommodity flow such that the flow of each commodity i is an L -length-bounded s_i - t_i -flow.

2.4. SERIES-PARALLEL AND OUTERPLANAR GRAPHS. In this article we deal several times with series-parallel graphs and outerplanar graphs and therefore their definitions are provided. A directed acyclic graph G with two dedicated and distinct nodes $s, t \in V$, the source and the sink, is *series-parallel* if and only if one of the following holds (an equivalent definition can be found in Brucker [2001]):

- (1) (*Base case*). G consists only of the nodes s, t and the edge (s, t) .
- (2) (*Parallel decomposition*). G can be obtained from two series-parallel graphs G_1 and G_2 , with source-sink pairs s_1, t_1 and s_2, t_2 , by taking the disjoint union of G_1 and G_2 and identifying s_1 with s_2 and t_1 with t_2 , which gives the source s and sink t of G , respectively.
- (3) (*Series decomposition*). G is obtained analogously to the parallel composition from two series-parallel graphs G_1 and G_2 , except that in this case t_1 is identified with s_2 and $s = s_1, t = t_2$.

An undirected graph is series-parallel, if it can be derived from a series-parallel directed graph by removing the edge directions.

An undirected graph is *outerplanar*, if it has a planar embedding such that all vertices are on the same face. A directed graph is called *outerplanar* if its underlying undirected graph is outerplanar.

2.5. FLOW NUMBER. At the end of Section 3.4 we deal with a graph parameter called *flow number*. Here we recall the definition of the flow number for an undirected graph. First, we define a few auxiliary terms. In a *concurrent multicommodity flow problem*, there are k commodities, each specified by a pair of nodes (s_i, t_i) and a demand d_i . A *feasible solution* for this problem is a multicommodity flow that obeys the capacity constraints but need not meet the specified demands. The *flow value of a feasible solution* is the maximum value f such that at least $f \cdot d_i$ units of commodity i are routed for each i . The *max-flow* for a multicommodity flow problem is defined as the maximum flow value over all feasible solutions. Given a concurrent multicommodity flow problem with feasible solution \mathcal{S} , the *dilation* $D(\mathcal{S})$ of \mathcal{S} is the length of the longest flow path in \mathcal{S} and the *congestion* $C(\mathcal{S})$ of \mathcal{S} is the inverse of its flow value. Let I be the instance of the concurrent multicommodity flow problem with a commodity for every ordered pair of nodes such that the demand for the pair (u, v) is $c(u)c(v) / \sum_{w \in V} c(w)$, where $c(w) = \sum_{e=\{w,z\} \in E} u_e$ for each vertex $w \in V$. The *flow number* $F(G)$ of a graph G is the minimum over all feasible solutions \mathcal{S} for the instance I of $\max\{C(\mathcal{S}), D(\mathcal{S})\}$ [Kolman and Scheideler 2006]. The flow number F of a graph is closely related to the expansion α of the graph: $F = \Omega(\alpha^{-1})$ and $F = \mathcal{O}(\Delta \alpha^{-1} \log n)$, where Δ is the maximum degree in the graph [Kolman and Scheideler 2006].

3. Length-Bounded Cuts

3.1. LENGTH-BOUNDED FLOWS VERSUS LENGTH-BOUNDED CUTS

3.1.1. Edge-Cuts. It follows from linear programming duality that the maximum (fractional) length-bounded flow value equals the minimum fractional length-bounded cut value. For standard flows, this equality holds for integral cuts as well. In the presence of a length-bound, the maximum flow value and the minimum cut value may be very different. This is in intimate relationship with the integrality gap of the linear program (1).

THEOREM 3.1. *There exist infinite families of directed and undirected graphs for which the ratio of the minimum integral length-bounded edge-cut value to the minimum fractional length-bounded edge-cut value is of order $\Theta(n^{2/3})$ for a graph with n vertices, and this is the worst possible ratio.*

PROOF. We describe the family of undirected graphs; the directed graphs are obtained from the undirected graphs by a natural orientation of the edges (“from left to right”). Let n be an integer such that $n^{1/3}$ is integral and let $k = n^{2/3}$. We describe a graph G on $\Theta(n)$ vertices for which the ratio between the two cuts is of order $\Theta(n^{2/3})$. The core of G is a layered graph G' consisting of $4k + 1$ layers. The first layer contains only a single vertex s and the last layer contains only a single vertex t . The second layer and the last but one layer consist of $n^{2/3}$ vertices each. Every other layer consists of $n^{1/3}$ vertices. Except for the pair 2, 3 and the




FIG. 1. An instance ($L = 3k$ where $k = n^{2/3}$) with integrality gap $\Omega(n^{2/3})$. Every L -length-bounded path has to use at least k shortcut edges (the topmost edges in the figure).

pair $4k - 1, 4k$, each pair of consecutive layers of G' forms a complete bipartite graph. The vertices of layer 2 are partitioned into $n^{1/3}$ groups of size $n^{1/3}$, and all vertices from each group are connected to a unique vertex (different from the vertices chosen for the other groups) from the third layer. Vertices in layers $4k - 1$ and $4k$ are connected in an analogous way.

The desired graph $G = (V, E)$ is obtained from G' by choosing a single vertex from each odd layer, say a vertex v_i from the odd layer i , and connecting v_i with v_{i+2} by an edge for each odd i ; these edges are called *shortcut edges*. The graph G is depicted in Figure 1. We note that a similar graph was used in a lower bound proof by Galil and Yu [1995] in the context of bounds on the average path length of edge disjoint s - t -paths and by Chekuri and Khanna [2007] in the analysis of the performance of a greedy algorithm for the edge disjoint paths problem.

For length bound $L = 3k$, we show that the minimum length-bounded cut has size $\Omega(n^{2/3})$. Note that any s - t -path of length at most L must use at least k shortcut edges. Given a subset $F \subseteq E$ of edges, we say that a shortcut edge $\{v_i, v_{i+2}\}$ from $E \setminus F$ is *substantial* if $5 \leq i \leq 4k - 5$, the vertex v_i is connected in $(V, E \setminus F)$ to more than $n^{1/3}/2$ vertices from layer $i - 1$, and the vertex v_{i+2} is connected in $(V, E \setminus F)$ to more than $n^{1/3}/2$ vertices from layer $i + 3$. Note that $2k - 4$ of the $2k$ shortcut edges can potentially be substantial. We claim that if the number of substantial edges in $(V, E \setminus F)$ is at least k and the size of F is at most $k/2$, then there exists an L -length-bounded path between s and t in $(V, E \setminus F)$. To see this, note that G contains more than $n^{2/3}/2$ edge-disjoint paths from s to any set of more than $n^{1/3}/2$ vertices in a layer i , $3 \leq i \leq 4k - 1$; more than $n^{2/3}/2$ edge-disjoint paths from any set of more than $n^{1/3}/2$ vertices in a layer i to any set of more than

$n^{1/3}/2$ vertices in a layer j , for $3 \leq i \leq 4k - 3$ and $i + 2 \leq j \leq 4k - 1$; and more than $n^{2/3}/2$ edge-disjoint paths from any set of more than $n^{1/3}/2$ vertices in a layer i , $3 \leq i \leq 4k - 1$, to t . In each of the three cases, a set F of at most $k/2$ edges cannot hit all the more than $n^{2/3}/2$ edge-disjoint paths. Therefore, we can construct an L -length-bounded path from s to t in $(V, E \setminus F)$ by concatenating shortest paths from s to the first substantial edge, between any two consecutive substantial edges, and from the last substantial edge to t . Hence, if F is an L -length-bounded edge-cut, we must have that $(V, E \setminus F)$ contains less than k substantial edges (implying that $|F| > k - 4$ or $|F| > k/2$; in both cases, F has size $\Omega(k) = \Omega(n^{2/3})$, as claimed).

On the other hand, assigning each shortcut edge e a dual length $l_e = 1/k$ ensures that the dual length of every s - t -path from $\mathcal{P}_{s,t}(L)$ is at least 1. Thus, the integrality gap is $\Theta(k) = \Theta(n^{2/3})$ in this instance. It remains to show that this is the worst case.

For length bounds $L \leq n^{2/3}$, a simple rounding scheme proves that the ratio between fractional and integral minimum cuts in every graph is at most $n^{2/3}$: given a minimum fractional L -length-bounded cut, round every $l_e \geq 1/L$ to 1 and all other l_e to zero; clearly this yields an integral L -length-bounded cut that is at most L -times larger than the fractional one. For $L > n^{2/3}$, we start with a similar rounding scheme and round every $l_e \geq 1/n^{2/3}$ to 1 and all other l_e to zero. If we remove at this point all edges with $l_e = 1$ from the graph, the distance between s and t will be at least $n^{2/3} + 1$ and thus, by a theorem of Even and Tarjan [1975], the minimum cut between s and t has size $\mathcal{O}(n^{2/3})$. Putting together edges with $l_e = 1$ and edges from the minimum cut in the reduced graph gives an integral L -length-bounded cut that is at most $\mathcal{O}(n^{2/3})$ -times larger than the fractional one. \square

As mentioned earlier, by duality of linear programming, the size of a minimum fractional L -length-bounded cut equals the size of a maximum fractional L -length-bounded flow, which implies the following corollary.

COROLLARY 3.2. *There exist infinite families of directed and undirected graphs for which the ratio of the minimum integral length-bounded edge-cut value to the maximum fractional length-bounded flow is of order $\Theta(n^{2/3})$ for a graph with n vertices, and this is the worst case.*

Remarkably, asymptotically the same ratio applies for the minimum integral L -length-bounded edge-cut and the maximum *integral* L -length-bounded flow (i.e., the maximum number of edge disjoint L -length-bounded paths between s and t).

COROLLARY 3.3. *There exist infinite families of directed and undirected graphs for which the ratio of the minimum integral length-bounded edge-cut value to the maximum number of length-bounded edge disjoint paths is of order $\Theta(n^{2/3})$ for a graph with n vertices, and this is the worst case.*

PROOF. The graphs described in the proof of Theorem 3.1 provide again the lower bound. For the upper bound, we describe an L -length-bounded cut $C \subseteq E$ of size at most $\mathcal{O}(n^{2/3})$ -times larger than the maximum number of L -length-bounded edge disjoint paths between s and t , for a given instance of the problem. We argue similarly as in the proof of Theorem 3.1. For $L \leq n^{2/3}$, let h be the maximum number of edge-disjoint L -length-bounded paths between s and t ; the cut simply consists of the edges of these h paths. For $L > n^{2/3}$, let h be the maximum




FIG. 2. Example of a large integrality gap of the linear program (1) of the minimum length-bounded cut. The straight s - t -path (in gray) contains $2k + 1$ edges. Each of these edges is accompanied by $k + 1$ parallel paths of length 2 and the length bound is $L = 3k + 1$.

number of edge-disjoint $n^{2/3}$ -length-bounded paths between s and t . Consider the cut consisting of all edges of the h edge-disjoint paths of length at most $n^{2/3}$ between s and t . By removing these edges, the distance between s and t increases to $n^{2/3} + 1$ at least, and by the theorem of Even and Tarjan [1975], the minimum cut between s and t has size $\mathcal{O}(n^{2/3})$. Altogether, we have again an L -length-bounded cut that is at most $\mathcal{O}(n^{2/3})$ -times larger than the maximum number of edge disjoint L -length-bounded s - t -paths. \square

For the instance used in the above proofs (cf. Figure 1), the maximum fractional and integral L -length-bounded flows have the same size, namely, 2. It is worth stressing that this is generally not the case (cf. the construction by Guruswami et al. [2003] for the lower bound $n^{1/2-\varepsilon}$ on the approximation ratio of the L -length-bounded edge disjoint paths problem).

3.1.2. Node Cuts. Analogous results hold for node cuts; the integrality gap is smaller in this case. Since the same bounds apply for edge cuts on series-parallel graphs, in the next theorems we mention both node and edge cuts.

THEOREM 3.4. *There exist infinite families of directed and undirected graphs for which the ratio of the minimum integral length-bounded node-cut value to the minimum fractional one is of order $\Theta(\sqrt{n})$ for a graph with n vertices, and this is the worst possible ratio. The same bound applies for edge-cuts on series-parallel graphs.*

PROOF. We start by giving the construction for edge cuts and then describe how to adapt it to node cuts. The construction is very similar to the construction in the proof of Theorem 3.1. For every $k \in \mathbb{N}$ we construct a graph G_k on $2k^2 + 5k + 3$ vertices with a fractional length-bounded edge-cut value less than 2 and an integral length-bounded cut value $k + 1$. The graph G_k is generated from an s - t -path containing $2k + 1$ edges; we call these edges *ground edges*. Parallel to each ground edge we add $k + 1$ paths of length 2; see Figure 2 for the undirected case. Note that the proof works for directed edges as well (direct edges from left to right).

Consider a graph G_k for arbitrary k and let $L = 3k + 1$. A minimum fractional edge cut has a value less than 2. This can be seen as follows. For a fractional edge cut, we have to assign a dual edge-length to each edge, such that the dual length of each s - t -path with at most $L = 3k + 1$ edges is not less than 1. An s - t -path with at most $3k + 1$ edges must contain at least $k + 1$ ground edges. Thus, assigning each ground edge a dual length of $\frac{1}{k+1}$ and assigning zero to the remaining edges yields a fractional cut of value $\frac{2k+1}{k+1} < 2$.

Now we give a lower bound of $k + 1$ on the size of an (integral) edge cut. If we take a nonground edge, we must take at least $k + 1$ nonground edges. Otherwise for any nonground edge in the cut there would always be another equivalent length 2 path which is not cut and thus the nonground edges could be removed from the cut without invalidating it. A cut containing only ground edges must have size greater than k ; otherwise an s - t -path of length $L = 3k + 1$ remains. Since $k + 1$ is in $\Theta(\sqrt{n})$, this completes the proof.

For node-cuts one can simply take the line graph (replace each edge by a node, connect two nodes, if the corresponding edges shared a node) of the above construction. This gives the $\Omega(\sqrt{n})$ lower bound on the integrality gap in undirected and directed (direct edges from left to right) graphs.

The proof of optimality of the bound for node cuts follows the same argument as the proof of Theorem 3.1. Given a fractional node cut and $L \leq \sqrt{n}$, round every $l_v \geq 1/L$ to 1 and all others to 0. For $L > \sqrt{n}$, first round every $l_v \geq 1/\sqrt{n}$ to 1 and remove these nodes from the graph; at this point, the distance between s and t is at least $\sqrt{n} + 1$ and thus the minimum cut has size at most \sqrt{n} . Putting together nodes from this cut and the nodes with $l_v = 1$, we get an L -bounded node cut that is at most $\Theta(\sqrt{n})$ -times larger than the size of the fractional cut. \square

Similarly as Theorem 3.1 implies Corollaries 3.2 and 3.3, Theorem 3.4 implies the next two corollaries.

COROLLARY 3.5. *There exist infinite families of directed and undirected graphs for which the ratio of the minimum integral length-bounded node-cut size to the maximum fractional length-bounded flow is of order $\Theta(\sqrt{n})$ for a graph with n vertices, and this is the worst possible ratio. The same bound applies for edge cuts on series-parallel graphs.*

COROLLARY 3.6. *There exist infinite families of directed and undirected graphs for which the ratio of the minimum integral length-bounded node-cut size to the maximum number of length-bounded node disjoint paths is of order $\Theta(\sqrt{n})$ for a graph with n vertices, and this is the worst possible ratio. The same bound applies for edge cuts and edge-disjoint paths on series-parallel graphs.*

We note that one of the main questions posed in the paper by Lovász et al. [1978] was exactly about the ratio of the minimum integral length-bounded node cut size and the maximum number of node disjoint length-bounded paths.

3.2. CUTS VERSUS LENGTH-BOUNDED CUTS. In this subsection we establish bounds on differences between the sizes of standard minimum cuts and length-bounded minimum cuts.

THEOREM 3.7. *Let $G = (V, E)$ be a directed or undirected multigraph. The difference between the size of a minimum node cut in G and the size of a minimum L -length-bounded node cut is at most $\mathcal{O}(\frac{n}{L})$. If G is a simple graph, the difference between the size of a minimum edge cut and the size of a minimum L -length-bounded edge cut is at most $\mathcal{O}(\frac{n^2}{L^2})$; if $L \geq \sqrt{m}$, then that difference is at most $\mathcal{O}(\sqrt{m})$, even for multigraphs.*

PROOF. Our arguments apply to directed or undirected graphs in the same way; therefore we describe the proof for undirected graphs only.




FIG. 3. Example of the $\frac{n}{L}$ gap between the standard and the length-bounded cut.

First, consider the case of node cuts. Let C_1 be a minimum length-bounded node cut. We will construct a node cut C of size at most $|C_1| + \frac{n}{L}$. In $G \setminus C_1$, all s - t -paths have length at least $L + 1$. The number of node-disjoint s - t -paths in $G \setminus C_1$ is at most $(n - 2)/L \leq n/L$, as each such path contains at least L internal nodes and no two such paths contain the same node. Therefore, a minimum node cut in $G \setminus C_1$ has size at most n/L . Let C_2 be such a cut. Then $C = C_1 \cup C_2$ is a node cut in G of the desired size.

The proof for edge cuts follows along the same lines. Let C_1 be a minimum length-bounded edge cut. We will construct an edge cut C of size at most $|C_1| + \mathcal{O}(\frac{n^2}{L^2})$, and for $L \geq \sqrt{m}$ we also give an upper bound $|C_1| + \sqrt{m}$ on the size of C . In $G \setminus C_1$, all s - t -paths have length at least $L + 1$. A result of Even and Tarjan [1975] (cf. Galil and Yu [1995]; Chekuri and Khanna [2007]) implies that the number of edge-disjoint s - t -paths in $G \setminus C_1$ is $\mathcal{O}(n^2/L^2)$ and thus a minimum edge cut in $G \setminus C_1$ has size $\mathcal{O}(n^2/L^2)$. If $L \geq \sqrt{m}$, then the number of s - t -paths in $G \setminus C_1$ is at most \sqrt{m} since each of these paths uses at least $\sqrt{m} + 1$ edges; again, this implies that the minimum edge cut in $G \setminus C_1$ has size at most \sqrt{m} . Let C_2 be a minimum cut in $G \setminus C_1$. Then $C = C_1 \cup C_2$ is an edge cut in G of size $|C_1| + \mathcal{O}(\frac{n^2}{L^2})$ and, if $L \geq \sqrt{m}$, the size of C is at most $|C_1| + \sqrt{m}$. \square

Figure 3 gives an example showing that the bound of $\frac{n}{L}$ on the gap between standard and length-bounded node cuts given in Theorem 3.7 is tight. In this example, s and t are connected by one path of length L and by $\frac{n-L-1}{L} \approx \frac{n}{L}$ paths of length $L + 1$. A minimum L -length-bounded node cut has size one while the minimum standard node cut needs to cut all paths and has size approximately $\frac{n}{L}$.

3.3. HARDNESS OF APPROXIMATION. Table I provides an overview of known and new results concerning the complexity, inapproximability, and polynomially solvable cases of length-bounded cut problems. Furthermore, we give an \mathcal{NP} -hardness proof for the edge version in weighted series-parallel and outerplanar graphs. Note that the polynomial algorithms for L equals 2, 3 and 1, 2 for the node and edge version, respectively, are easy exercises for both directed and undirected graphs (for the case $L = 3$ node cut or $L = 2$ edge cut: after directly cutting length 2 or length 1 paths, respectively Theorem 3.13 can be applied).

3.3.1. Node Cuts. We first present a simple polynomial time algorithm for length-bounded node cuts with $L = n - c$, where $c \in \mathbb{N}$ is an arbitrary constant.




FIG. 4. Gadgets for the reduction of VERTEX COVER to length-bounded node-cut (left) and length-bounded edge-cut (right), respectively. Both correspond to two connected nodes u, v of the given VERTEX COVER instance, shown in the middle. The highlighted nodes (edges) are in the cut and in the vertex cover.

Then we come to the main result of this section, which is the inapproximability result.

THEOREM 3.8. *If $c \in \mathbb{N}$ is constant and $L = n - c$, then a minimum length-bounded node cut can be computed in polynomial time in directed and undirected graphs.*

PROOF. Enumerate all $C \subseteq V$ with $|C| \leq c$ and return the smallest C which is a length-bounded node cut, if there is any. Otherwise, any length-bounded node cut C contains at least $c + 1$ nodes so that the longest remaining s - t -path has a length at most $n - c - 1$ and therefore C actually cuts all s - t -paths. Thus, returning a standard minimum node cut suffices. \square

THEOREM 3.9. *For any $\varepsilon > 0$ and $L \in \{5, \dots, \lfloor n^{1-\varepsilon} \rfloor\}$, it is \mathcal{NP} -hard to approximate the minimum length-bounded node cut in directed and undirected graphs within a factor of 1.1377.*

PROOF. We first look at the case $L = 5$ in directed graphs and give a reduction from the well known VERTEX COVER problem. A *vertex cover* for an undirected graph $G_{vc} = (V_{vc}, E_{vc})$ is a subset V'_{vc} of the nodes V_{vc} such that for each edge $\{u, v\} \in E_{vc}$ at least one of the nodes u, v is in V'_{vc} . The problem of finding a minimum vertex cover has been shown to be \mathcal{NP} -hard to approximate within a factor ≈ 1.3606 [Dinur and Safra 2005]. Given a VERTEX COVER instance G_{vc} with $n_{vc} = |V_{vc}|$ nodes, we construct a length-bounded node-cut instance $G = (V, E)$ as follows: start with $V = \{s, t\}$ and no edges. For each node $v \in V_{vc}$ we add a *node gadget* to G consisting of seven nodes which are interconnected with s, t and themselves as shown in Figure 4 (left)—the nodes in the top half surrounded by a gray box. For each edge $\{u, v\} \in E_{vc}$ we add an *edge gadget* consisting of four nodes and six edges connecting them to the node gadgets corresponding to u and v as shown in Figure 4 (left).

LEMMA 3.10. *From a vertex cover V'_{vc} in G_{vc} of size x one can always construct in polynomial time a length-bounded node cut C in G of size $n_{vc} + x$ and vice versa, for $x \leq n_{vc}$.*

PROOF. We start with the easier direction “ \Rightarrow ”: let $V'_{vc} \subseteq V_{vc}$ be a vertex cover with $|V'_{vc}| = x$. For each node $v \in V'_{vc}$ we add two nodes l_v and r_v to our node cut $C \subseteq V$ and for each node $u \in V_{vc} \setminus V'_{vc}$ we add m_u to C (see Figure 4 for an example). Clearly this ensures $|C| = n_{vc} + x$.


To see that no 5-length-bounded path between s and t remains after removing the nodes C from G , first consider paths that use vertices of the node gadgets only, apart from s and t . With respect to a node gadget for a node v , there are two cases to distinguish. In the case that m_v was added to the cut, no length-bounded path remains in the gadget. In the case that l_v and r_v were added, the only remaining path (via m_v) has length 6, which is greater than the length-bound 5.

Now assume that after removing nodes in C there remains an s - t -path of length at most five that uses vertices from an edge gadget for some edge $\{u, v\} \in E_{vc}$. Then either l_v and r_u are not in the cut C or l_u and r_v are not in the cut C . By construction this means that both u and v are not in V'_{vc} which is a contradiction to V'_{vc} being a vertex cover. We conclude that no length-bounded s - t -path remains in G after removing C .

Now we come to the direction “ \Leftarrow ”: let $C \subseteq V$ be a length-bounded node cut of size $|C| = n_{vc} + x$. With two simple transformations we ensure that C contains nodes from the node gadgets only, and that for each node gadget either the m -type node or both the l - and r -type nodes are contained in C .

- No nodes from edge gadgets in C .* Consider an edge gadget, say for an edge $\{u, v\} \in E_{vc}$, for which at least one of its four nodes is in C . The edge gadget consists of two paths, one from l_v to r_u and one from l_u to r_v . If an inner node of the l_v - r_u -path or the l_u - r_v -path is in C , we replace it by l_v or l_u , respectively. This introduces no new length-bounded paths and does not increase the size of the node cut.
- Node gadget: either m -type node or both l - and r -type nodes in C .* Consider a node gadget for a node $v \in V_{vc}$. First note that at least one node of the gadget must be in C ; otherwise there exist three s - t -paths of length at most 5 via this gadget. If only one node of the gadget is in C , it must be m_v ; otherwise there exists at least one length-bounded path via this gadget. If two or more nodes of the gadget are in C , we replace them by l_v and r_v ; this guarantees that no 5-length-bounded path via this node gadget exists. The transformation clearly does not increase the size of the node-cut C .

A vertex cover V'_{vc} of size $|C| - n_{vc}$ can easily be derived from the transformed C as follows: include in V'_{vc} all nodes $v \in V_{vc}$ for which both nodes l_v and r_v are in C . Assume, for a contradiction, that V'_{vc} is not a vertex cover, that is, there exists an edge $\{u, v\} \in E_{vc}$ that is uncovered by V'_{vc} . Then both m_u and m_v are in the cut C (and no other nodes of the node gadgets for u and v) which implies that there exist two 5-length-bounded s - t -paths via the edge gadget for $\{u, v\}$, a contradiction to C being a node cut. Concerning the size of V'_{vc} , we observe that every node gadget after the two transformations contains either one or two nodes from C ; thus, by construction, the size of V'_{vc} is $|C| - n_{vc}$.

FIG. 5. Replacing s by a path of length $L - 5$.

If the size of C was decreased by the two transformations and the resulting vertex cover V'_{vc} has therefore smaller size than claimed in the lemma, we artificially increase the size of the vertex cover by adding other vertices to it. \square

The proof of Theorem 1.1 in Dinur and Safra [2005] gives the following gap. There are graphs G_{vc} for which it is \mathcal{NP} -hard to distinguish between two cases: the case where a vertex cover of size $n_{vc} \cdot (1 - p + \varepsilon')$ exists, and the case where any vertex cover has size at least $n_{vc} \cdot (1 - 4p^3 + 3p^4 - \varepsilon')$, for any $\varepsilon' \in \mathbb{R}_{>0}$ and $p = (3 - \sqrt{5})/2$. If we plug this into the result of Lemma 3.10, we have shown that the length-bounded node cut is hard to approximate within a factor (there is an $\varepsilon' \in \mathbb{R}_{>0}$ for which the inequality holds): $(n_{vc} + n_{vc} \cdot (1 - 4p^3 + 3p^4 - \varepsilon')) / (n_{vc} + n_{vc} \cdot (1 - p + \varepsilon')) > 1.1377$.

For other values of $L \in \{5, \dots, \lfloor n^{1-\varepsilon} \rfloor\}$, we modify the construction of G as follows: (1) add a path of length $L - 5$ from a new source node s' to s . (2) Stepwise replace each node on this path after s' and until s (inclusive) by a *group* of $2n_{vc}$ nodes, and connect each new node with all neighbors of the replaced node (see Figure 5).

We need to verify that Lemma 3.10 applies for the new construction too. The proof of the left to right implication goes through without any modification. For the right to left implication, we first observe that, if only a few nodes of a group are in the cut, we easily find a smaller cut by removing these nodes from the cut. Further, for every cut that contains all nodes from a group, we can find a smaller cut by taking the l - and r -type nodes of all node gadgets but one and taking the m -type node of the last node gadget in the cut. Thus, we can assume that none of the new nodes appear in the length-bounded cut which makes it possible to transform again a cut of size $n_{vc} + x$ in a vertex cover of size x . The total number of nodes in G depends on L and is $n = \Theta(L \cdot n_{vc} + n_{vc} + m_{vc})$, where $m_{vc} = |E_{vc}|$. Therefore, we can create instances for which L is as large as $\lfloor n^{1-\varepsilon} \rfloor$, for arbitrarily small $\varepsilon \in \mathbb{R}_{>0}$.

To see that the reduction also works for undirected graphs, observe that, by removing the edge directions in the gadgets, no new undirected paths of length less than L are introduced. \square

3.3.2. Edge Cuts. The polynomial time algorithm for node cuts with length-bound $n - c$ does not carry over for the edge version of the problem since by removing c edges one cannot guarantee that computing a standard cut suffices. The inapproximability result does carry over, as stated in the following theorem. The proof is a straightforward modification of the proof of Theorem 3.9 and therefore we omit it; the difference is that other gadgets (described in Figure 4, right) are used.




FIG. 6. Reduction of 2-PARTITION to the length-bounded cut problem. The *labels* denote length/capacity.

THEOREM 3.11. *For any $\varepsilon > 0$ and $L \in \{4, \dots, \lfloor n^{1-\varepsilon} \rfloor\}$, it is \mathcal{NP} -hard to approximate the length-bounded edge-cut in directed and undirected graphs within a factor of 1.1377.*

If we allow the edges to have different lengths and capacities, the length-bounded cut problem is \mathcal{NP} -hard even for the classes of series-parallel and outerplanar graphs.

LEMMA 3.12. *For series-parallel and outerplanar directed and undirected graphs with general capacities and lengths, it is \mathcal{NP} -hard to decide whether there is a length-bounded edge cut of size less than a given value.*

PROOF. We will show a reduction of 2-PARTITION to the length-bounded cut problem. We are given an arbitrary 2-PARTITION instance $a_1, \dots, a_k \in \mathbb{N}$. We have to decide if there exists a partition A_1, A_2 of the ground set $A_1 \cup A_2 = \{a_1, \dots, a_k\}$ such that $\sum_{i \in A_1} a_i = \sum_{i \in A_2} a_i =: B$ holds.

Graph G is a single s - t -path with k multiedges; each multiedge consists of two parallel edges; see Figure 6. All k upper edges have length zero and successively a_1, \dots, a_k as capacity. The lower edges get successively a_1, \dots, a_k as length and capacity ∞ (or any finite capacity larger than B). Note that, to obtain a simple graph, we can simply subdivide the parallel edges which still yields a series-parallel and outerplanar graph. For the directed version simply direct the edges from left to right.

Let the length bound be $L = B - 1$. We will show that there is an L -length-bounded edge-cut of value at most B if and only if the instance of 2-PARTITION is a yes-instance.

“ \Leftarrow ”: Given a solution A_1, A_2 to the 2-PARTITION instance, we take the upper edges corresponding to set A_1 as our edge cut. Clearly only s - t -paths of length at least B remain and the cut has value B .

“ \Rightarrow ”: We start by showing that any L -length-bounded edge cut must have value at least B . Assume C is a cut of value less than B ; then the path which takes the upper edges complementary to C will have length less than B , which gives a contradiction. Thus, a given edge cut of value at most B has value exactly B and yields a 2-PARTITION in the obvious way. \square

We will show in Theorem 4.2 that it is \mathcal{NP} -hard to decide whether a fractional length-bounded flow of a given flow value exists in a graph with edge lengths even if the graph is series-parallel and outerplanar. Since the primal and dual programs have identical optimal objective function values, the same holds for the fractional length-bounded edge-cut problem.

3.4. APPROXIMATION ALGORITHMS. If the length-bound L is so large that the system of L -length-bounded s - t -paths contains the set of all s - t -paths, then length-bounded cuts and flows reduce to standard cuts and flows. The maximum-flow

minimum-cut equality holds and there are many efficient algorithms to compute minimum cuts and maximum flows exactly. If the length-bound L equals the distance between s and t , we get another case solvable in polynomial time. Lovász et al. [1978] showed a special version of the following theorem in the context of length-bounded node-disjoint paths.

THEOREM 3.13. *In directed and undirected multigraphs with general capacities and edge lengths, for $L = \text{dist}(s, t)$ the minimum length-bounded edge- and node-cut problem and the maximum length-bounded flow problem can be solved in polynomial time. In particular, the maximum flow value and the minimum cut value coincide if $L = \text{dist}(s, t)$.*

PROOF. We first consider directed graphs. Let G be such a graph with edge capacities and nonnegative edge lengths and let $L = \text{dist}(s, t)$. First we generate the subgraph \overline{G} induced by all edges which are contained in at least one shortest s - t -path in G . This subgraph can be found with a slightly modified Dijkstra-labeling algorithm; one has to remember for each node *all* incoming edges generating the smallest label at this node. The edges in \overline{G} that have positive length form a directed acyclic graph, and the edges in \overline{G} with length zero connect nodes that have the same distance to t in G . Each s - t -path in \overline{G} is a shortest s - t -path in G . Therefore, a standard minimum cut and a maximum flow in \overline{G} correspond to a minimum length-bounded cut and a maximum length-bounded flow in G . The theorem follows from standard flow theory. For undirected graphs we replace each edge by two antiparallel directed edges with the capacity and length of the original edge. The subgraph \overline{G} is then constructed in the same way, and any cut or flow in \overline{G} directly translates into a length-bounded cut or flow in the original graph. \square

For graphs with unit-lengths, Theorem 3.13 yields the following approximation result for the minimum length-bounded cut problem. A similar result for node cuts appears implicitly in Ben-Ameur [2000].

COROLLARY 3.14. *In directed and undirected multigraphs with general capacities and unit lengths, one can find in polynomial time an $(L + 1 - \text{dist}(s, t))$ -approximation to the minimum L -length-bounded edge or node cut by at most $L + 1 - \text{dist}(s, t)$ standard minimum cut calculations.*

PROOF. Removing a minimum $\text{dist}(s, t)$ -length-bounded cut from the graph increases the distance of s and t by at least 1. Repeating this iteratively increases the s - t -distance to $L + 1$ within at most $L + 1 - \text{dist}(s, t)$ iterations. The value of each intermediate minimum cut is a lower bound on the value of the minimum L -length-bounded cut. Thus, their union has value at most $L + 1 - \text{dist}(s, t)$ times the minimum value of an L -length-bounded cut. \square

In Figure 7 we provide an instance showing that the performance bound in the above corollary is tight in the worst case. The core of the graph is a path on $n - 2$ vertices, the last of which is a node t , and two other vertices v and s . Each of the $n - 2$ vertices on the path is connected to the vertex v , and the vertex v is connected to the vertex s . All edges have unit length and unit capacity. The length bound is $L = n - 1$. If the algorithm breaks ties in favor of edges that are incident with v but not incident with s , the algorithm finds a cut of value $n - 2$ while the minimum cut has value 1.




FIG. 7. An example showing that the performance bound is asymptotically tight. The length bound is $L = n - 1$.

Another way to obtain an L -approximation for the L -length-bounded edge cut is the following: while $\text{dist}(s, t) \leq L$, remove edges on the shortest s - t -path. Since all removed paths are edge disjoint, their number is a lower bound on the L -length-bounded edge-cut value which, together with the fact that the removed paths have length at most L , implies the approximation ratio.

Consider the following combination of the two approaches described above. First, while there exists an s - t -path of length at most $L/2$, remove all edges on such a path. Then, while there exists an s - t -path of length at most L , find a minimum cut in the subgraph consisting of edges on the shortest s - t -paths in the current graph, and remove edges in the cut. We observe that the number of iterations of the first phase (i.e., the number of edge disjoint paths of length at most $L/2$ that the algorithm finds), plus the size of any cut from the second iteration, is a lower bound on the size of the minimum L -bounded cut. Since the length of each path deleted in the first phase is at most $L/2$ and since there are at most $L/2$ iterations of the second phase, the algorithm computes an $L/2$ -approximation.

For the sake of completeness we mention again yet another L -approximation algorithm that appeared already in the proof of Theorem 3.1: given a fractional L -length-bounded cut, round every dual length with $l_e \geq 1/L$ to 1 and all others to zero.

The algorithms described above for edge cuts can be adapted to node cuts in a straightforward way.

For large values of the length-bound L , the $\mathcal{O}(L)$ -approximations are not very satisfying. In such cases, a combination of Corollary 3.14 and Theorem 3.7 yields the following theorem; we exploit the fact that minimum cuts can be computed in polynomial time.

THEOREM 3.15. *For directed and undirected graphs there exists an $\mathcal{O}(\min\{L, n/L\}) \subseteq \mathcal{O}(\sqrt{n})$ -approximation algorithm for the minimum L -length-bounded node-cut problem and an $\mathcal{O}(\min\{L, n^2/L^2, \sqrt{m}\}) \subseteq \mathcal{O}(n^{2/3})$ -approximation algorithm for the minimum L -length-bounded edge-cut problem; the algorithms have polynomial running times.*

For a large class of graphs, a better approximation ratio is often possible. Let F be the *flow number* of G , as defined by Kolman and Scheideler [2006]. The following lemma from the same paper will be useful.

LEMMA 3.16 (SHORTENING LEMMA [KOLMAN AND SCHEIDELER 2006]). *Let an undirected network with flow number F be given. Then, for any $\varepsilon \in (0, 1]$ and any feasible multicommodity flow with a flow value of f , there exists a feasible multicommodity flow with a flow value of $f/(1 + \varepsilon)$ that can be decomposed into paths of length at most $2 \cdot F(1 + 1/\varepsilon)$.*

THEOREM 3.17. *For undirected graphs and $L \geq 2 \cdot F(1 + 1/\varepsilon)$, there exists a polynomial-time $(1 + \varepsilon)$ -approximation algorithm for the minimum L -length-bounded node cut and minimum L -length-bounded edge cut problems, where F is the flow number of the given graph.*

PROOF. The proofs for node and edge cuts are identical. Consider a graph G with flow number F . Let f denote the size of the standard minimum cut between two vertices s and t ; by the duality of flows and cuts, f is also the size of the standard maximum flow between s and t . By Lemma 3.16, there exists an L -length-bounded flow between s and t of size at least $f/(1 + \varepsilon)$. Since the size of a length-bounded flow between s and t is a lower bound on the size of the minimum length-bounded cut, we conclude that a standard minimum-cut is a $(1 + \varepsilon)$ -approximation for the L -length-bounded cut. \square

As the flow number of hypercubes or expander graphs is $\mathcal{O}(\log n)$, the theorem implies polynomial-time $\mathcal{O}(1)$ -approximation algorithms for $L \geq 3F = \mathcal{O}(\log n)$ for these graphs. We remark that the definition of the flow number can be adapted to directed graphs in which for every node the total capacity of the in-edges equals the total capacity of the out-edges [Kolman and Scheideler 2006]; the corresponding version of Theorem 3.17 holds for such graphs.

We conclude this section with a simple observation that for a graph G with a source s and a sink t one can easily obtain an $\mathcal{O}(\max\{\text{degree}(s), \text{degree}(t)\})$ -approximation by removing all edges adjacent to the source or the sink. Thus, we have $\mathcal{O}(1)$ -approximations for constant-degree graphs.

4. Length-Bounded Flows

4.1. COMPLEXITY. Choosing infinity as a length bound reduces the length-bounded flow problem to the standard flow problem. For computing maximum standard flows, in most cases one does not use the linear program (2), since the number of paths and thus the number of variables may be exponential in the input size. It is more common to use an edge-based formulation since it always uses a polynomially bounded number of variables. For length-bounded (multicommodity) flows in unit-length graphs (and general capacities), Kolman and Scheideler [2006] described an edge-based linear programming formulation of polynomial size. The main result of this section, Theorem 4.2, implies that there is no linear programming formulation of polynomial size for the problem in graphs with general lengths, unless $\mathcal{P} = \mathcal{NP}$; the formulation of Kolman and Scheideler [2006] has size $\Omega(L(n+m))$. We also note that, for graphs with unit lengths (and general capacities), Baier [2003] described how to solve the linear program (2) in polynomial time. However, we are not aware of a combinatorial algorithm with polynomial running time for the maximum length-bounded flow problem in graphs with unit lengths.

When looking at a given length-bounded flow, from linear programming theory we can infer the existence of a path decomposition of small size, where all paths fulfill the length bound.

THEOREM 4.1. *For every L -length-bounded (multicommodity) flow, there exists an L -length-bounded path representation of the flow that uses at most m paths for each commodity.*




FIG. 8. The depicted graph possesses no small multicommodity pathrepresentation corresponding to the prescribed edge representation of the flow. As set of commodities, we choose all 9 pairs (a_i, b_j) , $i, j = 1, 2, 3$. The capacity of all 6 edges is set to 1. For each commodity (a_i, b_j) there is a flow of size $1/3$ (going along the only path between a_i and b_j). This is a feasible multi-commodity flow. However, the only corresponding path representation has to use altogether all 9 paths, which is greater than $|E| = 6$.

PROOF. The theorem follows from the fact that the linear program in (2) has only m linear constraints. Therefore, the rank of the linear program for a single commodity is at most m . Consequently, there has to be a solution using no more than m paths. We can modify the edge capacities appropriately and apply this argument to each commodity one after another. \square

The argument in the proof of Theorem 4.1 can be applied simultaneously to all commodities of a length-bounded flow. Hence, there always exists an optimal solution (maximizing the sum of the flow values of all the commodities) that uses no more than $|E|$ paths in total. However, in general this transformation may change the flow on some of the edges (and also the total flow value) for individual commodities. An example is given in Figure 8.

We see that the theory of linear programming guarantees that there is always a path representation of maximum flow value that has a small size. Nevertheless, for graphs with general edge lengths, linear programming is unable to find maximum fractional length-bounded flows efficiently, unless $\mathcal{P} = \mathcal{NP}$. We formalize this statement in the following theorem.

THEOREM 4.2. *For the length-bounded flow problem in directed and undirected series-parallel and outerplanar graphs with unit capacities and general lengths it is \mathcal{NP} -complete to decide whether there is a fractional length-bounded flow of given flow value.*

PROOF. First of all we observe that this decision problem is in \mathcal{NP} . Theorem 4.1 guarantees the existence of a polynomially sized path representation of a flow. This is a certificate which certainly can be checked in polynomial time.

To prove \mathcal{NP} -hardness, we proceed similarly as in the proof of Lemma 3.12 by giving a reduction of 2-PARTITION. We are given an arbitrary 2-PARTITION instance $a_1, \dots, a_k \in \mathbb{N}$. We have to decide if there exists a partition A_1, A_2 of the ground set $A_1 \cup A_2 = \{a_1, \dots, a_k\}$ such that $\sum_{i \in A_1} a_i = \sum_{i \in A_2} a_i =: B$ holds.

Graph G is a single s - t -path with k multiedges; each multiedge consists of two parallel edges. All k upper edges have length zero and the lower edges are successively assigned a_1, \dots, a_k as lengths. All capacities are set to 1. The constructed graph is the same as the one shown in Figure 6 except for the choice of capacities. Note that, to obtain a simple graph, we subdivide each of the parallel edges, which still yields a series-parallel and outerplanar graph. For the directed version, we direct the edges from left to right. Let the length-bound be $L = B$.

Let us first consider a maximum *integral* flow (integral with respect to the path-representation) respecting the length bound L . Obviously, the value of this flow is bounded by 2 from above. Assume that the maximum flow has value 2. By the integrality of the flow, there have to be exactly two edge-disjoint s - t -paths with flow value 1 each. The total length of these two paths is $2L$; thus both must have length exactly L . The edges of length greater than zero in one path define a feasible partition for our 2-PARTITION instance. On the other hand, each feasible partition for the 2-PARTITION instance describes two s - t -paths in G , each of length exactly L . Thus, there is an integral L -length-bounded flow of value 2 in G if and only if the 2-PARTITION problem is a *yes* instance.

To complete the proof, we have to show that there is an integral L -length-bounded flow of value 2 if there is a fractional L -length-bounded flow of value 2. If the fractional solution contains a single path of length exactly L we are done since this path describes a feasible partition and thus an integral flow. We now show that each path with flow value greater than zero in a maximum fractional path flow f must have length exactly L . Assume that there is at least one path with nonzero flow value and length strictly less than L . Let $\mathcal{P}_<$ and \mathcal{P}_\geq denote the sets of s - t -paths with nonzero flow value and length strictly less than L and greater or equal to L , respectively. Since the fractional flow has a total flow value of 2, there is a total of one unit of flow on each edge. Thus, for the “flow-weighted” sum of path lengths we get $2L = \sum_{P \in \mathcal{P}_<} f_P |P| + \sum_{P \in \mathcal{P}_\geq} f_P |P|$, where $|P|$ denotes the length of a path P . Since $\sum_{P \in \mathcal{P}_< \cup \mathcal{P}_\geq} f_P = 2$, by averaging, there must be a path of length greater than L if there is a path of length less than L . This contradicts the condition that the length of all flow paths is bounded by L . \square

Finding a maximum length-bounded flow is computationally more difficult than finding a standard maximum flow. For standard flows, the edge representation of a flow is usually used. Each path representation of a flow can be easily transformed into an edge representation. For standard flows, the reverse transformation can also be efficiently computed. If length bounds are present, one may use the edge representation, too. However, as the following corollary shows, edge and path representations are not polynomially equivalent for length-bounded flows. The following result is an immediate consequence of the proof of Theorem 4.2 and has been shown independently by Correa et al. [2007, Corollary 3.4].

COROLLARY 4.3. *Unless $\mathcal{P} = \mathcal{NP}$, there is no polynomial algorithm to transform an edge representation of a length-bounded flow in a graph with unit edge capacities and general edge lengths, into a path representation, even if the graph is series-parallel and outerplanar.*

4.2. STRUCTURE OF OPTIMAL SOLUTIONS AND INTEGRALITY GAP. For standard single-commodity flows with integral capacities, there is always an integral maximum flow. The situation is completely different in the presence of length constraints. We will not only show that there need not exist an integral maximum flow but also that there are instances where each fractional maximum flow ships a large percentage of the flow along paths with very small flow values.

THEOREM 4.4. *There exist directed and undirected series-parallel and outerplanar graphs with n vertices such that every maximum fractional length-bounded flow ships more than one half of the total flow along paths with flow values $\mathcal{O}(1/n)$.*




FIG. 9. Graph G_k in which the unique maximum length-bounded flow sends more than one-half of the flow along paths with small flow values.

PROOF. We construct an infinite family $\{G_k\}_{k \in \mathbb{N}}$ of series-parallel and outerplanar graphs such that G_k has $3k + 4$ vertices and a maximum fractional L_k -length-bounded s - t -flow of value less than 2, for a certain length-bound $L_k \in \Theta(k)$. The unique maximum L_k -length-bounded flow in G_k contains $k + 1$ paths each with flow value $\frac{1}{k+1}$. We describe the construction for undirected graphs; for the directed case, simply direct edges from left to right.

The graph G_k consists of a sequence of $k + 1$ triangles preceded by a path of length $k + 1$ and a single edge that is parallel to the path; see Figure 9.

In G_k we consider a maximum fractional $(2k + 2)$ -length-bounded s - t -flow, that is, $L_k = 2k + 2$. There is only one s - t -path \tilde{P} of length at most $2k + 2$ that contains the s - u -path of length $k + 1$. Indeed, this path has length exactly $2k + 2$ and contains the unique shortest u - t -path. To obtain a total flow value larger than 1, path \tilde{P} has to be used. We call the edges in the shortest u - t -path *ground edges*.

All s - t -paths of length at most $2k + 2$ except \tilde{P} contain the edge su and at least one of the ground edges. Consider the s - t -paths of length exactly $2k + 2$ that contain edge su . There are $k + 1$ of those paths, one corresponding to each ground edge. Routing a fraction of $\frac{1}{k+1}$ units along each of them yields a feasible flow of value 1. Each ground edge is contained in exactly one of these paths and has therefore a residual capacity of $1 - \frac{1}{k+1}$. Thus, along path \tilde{P} we can route further $1 - \frac{1}{k+1}$ units of flow and obtain a feasible $(2k + 2)$ -length-bounded s - t -flow of value $2 - \frac{1}{k+1}$. We claim that this flow is maximum and unique.

Sending 1 unit of flow along path \tilde{P} blocks each other path containing a ground edge, i. e., each further feasible s - t -path. Assume $1 - \delta$ units of flow are sent along path \tilde{P} , for an arbitrary $0 < \delta < 1$. Then all remaining paths have a flow value not greater than δ each and thus altogether at most $\min\{1, (k + 1)\delta\}$. Therefore, the maximum flow value dependent on δ is bounded by $1 - \delta + \min\{1, (k + 1)\delta\}$. This expression, viewed as a function of δ , reaches its unique maximum for $0 < \delta < 1$ at $\delta = \frac{1}{k+1}$. Hence, $2 - \frac{1}{k+1}$ is the maximum fractional s - t -flow value for the given length-bound and the above constructed flow is unique. \square

For length-bounded flows, there is a surprising structural difference between integrality of path and edge representations, stated in the next theorem.

THEOREM 4.5. *There exist directed and undirected series-parallel graphs such that a (maximum) fractional length-bounded flow has an integral edge representation but does not have an integral length-bounded path representation.*

PROOF. For the sake of simplicity, we start by proving an analogous result for graphs with general edge capacities, and then describe how to modify the construction for unit-capacity graphs. Consider the undirected graph in Figure 10




FIG. 10. A unit-length graph with an integral edge flow of value 4 that corresponds to a maximum fractional 6-length-bounded path flow but which has no integral 6-length-bounded path decomposition: edge vt has capacity 3; all other edges have unit capacity.

and a length-bound $L = 6$. (For the directed case, direct edges from left to right.) We show that there exists a 6-length-bounded s - t -flow of value 4 with an integral edge representation and that no 6-length-bounded s - t -flow of value 4 has an integral path representation.

We send half a unit of flow along each of the two 6-length-bounded paths avoiding the edge vt . All remaining 6-length-bounded paths contain the edge vt and may use up to two detours from the shortest s - v -paths. Consider only paths using exactly two detours. There are three of them using the upper-left-side part and three using the lower-left-side part of the graph. In each triple, every two of them share a detour, and no two of them share an edge from the shortest s - v -paths. We send half a unit of flow along each of these six paths. Altogether we get a feasible 6-length-bounded s - t -flow of value $\frac{8}{2} = 4$ with integral edge representation; note that the path representation is half-integral.

Assume now that there is a 6-length-bounded s - t -flow of value 4 that has an integral path representation. All edges must have flow value 1 in such a flow. Since the shortest s - v -path has length 3, each s - t -path not using edge vt must go along one of the two shortest s - v -paths. An integral 6-length-bounded s - t -flow of value 4 must send one unit of flow along this path. Assume that this path uses the upper half of the graph. Then each additional path in the upper half of the graph has length 7 and is therefore infeasible. Thus, no integral 6-length-bounded s - t -flow has value 4.

To prove the theorem for unit-capacity graphs, we replace the edge vt by three paths of length 2 and the v - t -path of length 3 by a path of length 4, and we increase the length bound to 7. \square

Baier [2003] showed that the fractional length-bounded flow problem can be approximated within arbitrary precision. Having this in mind, it is interesting how far the value of such a fractional solution is away from the maximum integral solution. We note that the hardness results in Guruswami et al. [2003] imply for *directed* graphs an integrality gap of order $\Omega(n^{1/2-\varepsilon})$, for every $\varepsilon > 0$.

THEOREM 4.6. *The integrality gap of the linear program (2) is of order $\Omega(\sqrt{n})$ even for directed or undirected planar graphs with n nodes.*

PROOF. For each k we describe an undirected graph G_k on $n = \Theta(k^2)$ vertices such that the maximum integral length-bounded flow (integral with respect to the path-representation) has value 1 while the maximum fractional length-bounded flow has value $\Omega(\sqrt{n})$. For ease of presentation, we first allow integral edge lengths and then at the end of the proof we explain how to modify the construction for unit lengths. The construction is inspired by Guruswami et al. [2003] and Kleinberg [1996].




FIG. 11. A graph with a large integrality gap for the maximum length-bounded s - t -flow.

The basic structure is half of a $k - 1$ by $k - 1$ grid (see Figure 11). The top row has $k - 1$ columns, and each subsequent row has one column less than the previous one. Each element of the grid (gray ellipse) consists of two nodes (a *left* node and a *right* node). In addition, there are k nodes s_1, \dots, s_k arranged vertically, and k nodes t_1, \dots, t_k arranged horizontally, and two nodes s and t . The node s is connected to all s_i nodes and the node t is connected to all t_i nodes. For each grid element, its left and right nodes are connected by a *horizontal* edge (drawn dashed). Further horizontal edges connect the right node of a grid element to the left node of the grid element to the right (if any), and *vertical* edges connect the right node of a grid element to the left node of the grid element above (if any). Furthermore, for $2 \leq i \leq k$, s_i is connected by a horizontal edge to the left node of the first grid element in the corresponding row, and, for $1 \leq i \leq k - 1$, t_i is connected by a vertical edge to the right node of the top grid element in the corresponding column. Finally, there are *diagonal* edges connecting the right node of the rightmost grid element in each row (or s_1) to the left node of the rightmost grid element in the row above (or to t_k).

All edges have capacity 1. All edges except those adjacent to s or t have length 1. The edges $\{s, s_i\}$ and $\{t, t_{k+1-i}\}$ are assigned length $2i$, for $i = 1, \dots, k$. As length bound we choose $L = 4k + 1$.

Consider an s - t -path P of length at most L . Assume that P does not contain a diagonal edge. Let $\{s, s_i\}$ and $\{t_j, t\}$ be the first and last edge of P , respectively. Because of the grid structure of the graph G_k , it is easy to see that, between s_i and

t_j , the path P must use at least $k - i + 1$ vertical edges and at least $2j + (k - i)$ horizontal edges. Together with the lengths of the edges $\{s, s_i\}$ and $\{t_j, t\}$, path P has length at least $k - i + 1 + 2j + (k - i) + 2i + 2(k + 1 - j) = 4k + 3 > L$, a contradiction. Thus, each L -length-bounded s - t -path must contain a diagonal edge. It is easy to see that any two such paths must share an edge. Therefore, there are no two edge disjoint s - t -paths of length at most L .

The i th canonical path, for $1 \leq i \leq k$, is the path starting in s , going to s_i and then horizontally to the right end of the row, then using one diagonal edge, and then going vertically up to t via $\{t_i, t\}$ (using one horizontal dashed edge in each row it passes through). The length of the i th canonical path is exactly L . Since each pair of such canonical paths shares a different single edge, we can feasibly send half a unit of flow along each of them. That is, there is a fractional L -length-bounded s - t -flow of value $k/2$.

Since the maximum number of edge-disjoint L -length-bounded s - t -paths is 1, the gap between a maximum integral and a maximum fractional flow is at least $k/2$. Since k is of size $\Theta(\sqrt{n})$, this shows the lemma for integer edge lengths. If we subdivide each edge of length ℓ into a path with ℓ unit-length edges, we increase the number of nodes by a constant factor only and obtain the same result for unit lengths.

Directed graphs with the same integrality gap can be obtained by directing the edges from the undirected construction above in the following way: horizontal edges toward the right, vertical and diagonal edges toward the top. \square

The big integrality gap in Theorem 4.6 is tied to the unit-capacities of the graph used in the proof. Raising the edge capacities in this graph to 2 brings the integrality gap down to a constant. Indeed, the integrality gap is always constant for high-capacity graphs. The following result is a consequence of the randomized rounding technique of Raghavan and Thompson [1987].

THEOREM 4.7. *Consider a directed or undirected graph with minimal edge capacity at least $c \log n$, for a suitable constant c . Using randomized rounding, one can convert any fractional length-bounded flow into an integral length-bounded flow whose value is at most a constant factor smaller (with high probability). In particular, the integrality gap is constant for high-capacity graphs.*

5. Open Problems

In the Introduction and in Section 4.1 we mentioned that for graphs with unit lengths it is possible to compute a maximum fractional length-bounded flow in polynomial time using linear programming. However, we are not aware of a combinatorial algorithm for this problem. Even worse, a combinatorial algorithm that would decide in polynomial time whether a given length-bounded flow with a path representation is maximum has not been found either.

Another problem is to design an approximation algorithm for the minimum length-bounded cut with an approximation factor better than $L/2$.

ACKNOWLEDGMENT. The authors residing in the lowlands would like to thank the author residing in the highlands for his sedulous commitment and readiness to make sacrifices for the sake of this work. The authors would like to thank the reviewers for helpful comments and their careful proofreading.

REFERENCES

- ADÁMEK, J., AND KOUBEK, V. 1971. Remarks on flows in network with short paths. *Commentationes Mathematicae Universitatis Carolinae* 12, 4, 661–667.
- AHUJA, R. K., MAGNANTI, T. L., AND ORLIN, J. B. 1993. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, Englewood Cliffs, NJ.
- BAIER, G. 2003. Flows with path restrictions. Ph.D. dissertation, TU Berlin, Germany.
- BAIER, G., ERLEBACH, T., HALL, A., KÖHLER, E., SCHILLING, H., AND SKUTELLA, M. 2006. Length-bounded cuts and flows. In *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP)*. Lecture Notes in Computer Science, vol. 4051. Springer-Verlag, Berlin, 679–690.
- BEN-AMEUR, W. 2000. Constrained length connectivity and survivable networks. *Networks* 36, 1, 17–33.
- BLEY, A. 2003. On the complexity of vertex-disjoint length-restricted path problems. *Comput. Complex.* 12, 3-4, 131–149.
- BONDY, J. A., AND MURTY, U. 1976. *Graph Theory with Applications*. North Holland, New York, NY.
- BOYLES, S. M., AND EXOO, G. 1982. A counterexample to a conjecture on paths of bounded length. *J. Graph Theor.* 6, 205–209.
- BRANDES, U., NEYER, G., AND WAGNER, D. 1996. Edge-disjoint paths in planar graphs with short total length. Tech. rep. 19. Universität Konstanz, Koustanz, Germany.
- BRUCKER, P. 2001. *Scheduling Algorithms*, 3rd ed. Springer-Verlag, Berlin, Germany.
- CHEKURI, C., AND KHANNA, S. 2007. Edge-disjoint paths revisited. *ACM Trans. Algo.* 3, 4, Article No. 46.
- CORREA, J. R., SCHULZ, A. S., AND STIER MOSES, N. E. 2007. Fast, fair, and efficient flows in networks. *Oper. Res.* 55, 2, 215–225.
- DANTZIG, G. B., AND FULKERSON, D. R. 1956. On the max-flow min-cut theorem of networks. In *Linear Inequalities and Related Systems*, H. W. Kuhn and A. W. Tucker, Eds. Number 38 in Annals of Mathematics Studies. Princeton University Press, Princeton, NJ, 215–221.
- DINUR, I., AND SAFRA, S. 2005. On the hardness of approximating minimum vertex cover. *Ann. Math.* 162, 1, 439–486.
- ELIAS, P., FEINSTEIN, A., AND SHANNON, C. E. 1956. A note on the maximum flow through a network. *IRE Trans. Inform. Theor.* 2, 117–119.
- EVEN, S., AND TARJAN, R. E. 1975. Network flow and testing graph connectivity. *SIAM J. Comput.* 4, 507–518.
- EXOO, G. 1983. On line disjoint paths of bounded length. *Discrete Math.* 44, 317–318.
- FORD, L. R., AND FULKERSON, D. R. 1956. Maximal flow through a network. *Canad. J. Math.* 8, 399–404.
- FORD, L. R., AND FULKERSON, D. R. 1962. *Flows in Networks*. Princeton University Press, Princeton, NJ.
- GALIL, Z., AND YU, X. 1995. Short length versions of Menger’s Theorem. In *Proceedings of the 27th Annual ACM Symposium on Theory of Computing (STOC)*. ACM, New York, NY 499–508.
- GURUSWAMI, V., KHANNA, S., RAJARAMAN, R., SHEPHERD, B., AND YANNAKAKIS, M. 2003. Near-optimal hardness results and approximation algorithms for edge-disjoint paths and related problems. *J. Comput. Syst. Sci.* 67, 3, 473–496.
- HSU, D. F. 1994. On container width and length in graphs, groups, and networks. *IEICE Trans. Fund. Electron. Commun. Comput. Sci.* E77-A, 4, 668–680.
- ITAI, A., PERL, Y., AND SHILOACH, Y. 1982. The complexity of finding maximum disjoint paths with length constraints. *Networks* 12, 277–286.
- KLEINBERG, J. 1996. Approximation algorithms for disjoint paths problems. Ph.D. dissertation. Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA.
- KOLMAN, P., AND SCHEIDELER, C. 2006. Improved bounds for the unsplittable flow problem. *J. Algor.* 61, 1, 20–44.
- KOTZIG, A. 1956. Connectivity and regular connectivity of finite graphs. Ph.D. dissertation, Vysoká Škola Ekonomická, Bratislava, Slovakia.
- LOVÁSZ, L., NEUMANN-LARA, V., AND PLUMMER, M. D. 1978. Mengerian theorems for paths of bounded length. *Period. Math. Hungar.* 9, 269–276.
- MAHJOUB, A. R., AND MCCORMICK, S. T. 2003. The complexity of max-flow and min-cut with bounded-length paths. Unpublished Manuscript.
- MENGER, K. 1927. Zur allgemeinen Kurventheorie. *Fund. Math.* 10, 96–115.

- NIEPEL, L., AND SAFARÍKOVÁ, D. 1983. On a generalization of Menger's Theorem. *Acta Mathematica Universitatis Comenianae* 42, 275–284.
- PERL, Y., AND RONEN, D. 1984. Heuristics for finding a maximum number of disjoint bounded paths. *Networks* 14, 531–544.
- PYBER, L., AND TUZA, Z. 1993. Menger-type theorems with restrictions on path lengths. *Discrete Math.* 120, 1-3, 161–174.
- RAGHAVAN, P., AND THOMPSON, C. 1987. Randomized rounding: A technique for provably good algorithms and algorithmic proofs. *Combinatorica* 7, 4, 365–374.
- VAN DER HOLST, H., AND DE PINA, J. C. 2002. Length-bounded disjoint paths in planar graphs. *Discrete Appl. Math.* 120, 1-3, 251–261.
- WAGNER, D., AND WEIHE, K. 1995. A linear-time algorithm for edge-disjoint paths in planar graphs. *Combinatorica* 15, 1, 135–150.

RECEIVED DECEMBER 2007; REVISED JUNE 2009; ACCEPTED JULY 2009