Length-Bounded Cuts and Flows*

Georg Baier[1,**], Thomas Erlebach[2], Alexander Hall[3], Ekkehard Köhler[1], Heiko Schilling[1], and Martin Skutella[4]

[1] Institute of Mathematics, TU Berlin, Germany
[2] Department of Computer Science, U Leicester, England
[3] Institute of TCS, ETH Zurich, Switzerland
[4] Department of Mathematics, U Dortmund, Germany

Abstract. An L-length-bounded cut in a graph G with source s, and sink t is a cut that destroys all s-t-paths of length at most L. An L-length-bounded flow is a flow in which only flow paths of length at most L are used. We show that the minimum length-bounded cut problem in graphs with unit edge lengths is $\mathcal{NP}$-hard to approximate within a factor of at least 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 give approximation algorithms of ratio $\min\{L, n/L\}$ in the node case and $\min\{L, n^2/L^2, \sqrt{m}\}$ in the edge case, where n denotes the number of nodes and m denotes the number of edges. We discuss the integrality gaps of the LP relaxations of length-bounded flow and cut problems, analyze the structure of optimal solutions, and present further complexity results for special cases.

1 Introduction

In a classical article Menger [1], shows a strong relation between cuts and systems of disjoint paths: let G be a graph and s, t two nodes of G, then the maximum number of edge-/node-disjoint s-t-paths equals the minimum size of an s-t-edge-/node-cut (Menger's Theorem); see also Dantzig and Fulkerson [2] and Kotzig [3]. Ford and Fulkerson [4] and Elias, Feinstein, and Shannon [5] generalized the theorem of Menger to flows in graphs with capacities on the arcs and provided algorithms to find an s-t-flow and an s-t-cut of the same value.

Lovász, Neumann Lara, and Plummer [6] consider 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 suitable cut definition. For length-bounds greater than 4, they give upper and lower bounds for the gap between the maximum number of length-bounded node-disjoint paths and the minimum cardinality of a cut. Furthermore, they provide examples showing that some of the bounds are tight. The results were extended independently to edge-disjoint paths by Exoo [7] and Niepel and Safaríková [8].

* 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).

** Since 12/04 at Corporate Technology, Information & Communication, Siemens AG.

M. Bugliesi et al. (Eds.): ICALP 2006, Part I, LNCS 4051, pp. 679–690, 2006.

According to Bondy and Murty [9], Lovász conjectured that there is a constant C such that the size of a minimum L-length-bounded s-t-node-cut, i. e., a minimum node-set disjoint to $\{s, t\}$ which hits each L-length-bounded s-t-path, is at most a factor of $C \cdot \sqrt{L}$ larger than the cardinality of a maximum system of node-disjoint s-t-paths of length at most L. Exoo and Boyles [10] disprove this conjecture. They construct 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 $C \cdot L$ times the maximum number of node-disjoint s-t-paths of length at most L; the constant C is roughly 1/4.

Itai, Perl, and Shiloach [11] give efficient algorithms to find the maximum number of node-/edge-disjoint s-t-paths with at most 2 or 3 edges; the node-disjoint case is also solved for length-bound 4. On the complexity side they show that the node- and edge-disjoint length-bounded s-t-paths problem is $\mathcal{NP}$-complete for length-bounds greater than 4. Instead of fixing the path length, one can fix the number of paths and look for the minimal value bounding all path lengths. Again both the node- and edge-disjoint version is $\mathcal{NP}$-complete for two paths already.

Guruswami et al. [12] show that the edge-disjoint length-bounded s-t-paths problem is MAX SNP-hard even in undirected networks, and they give an $\mathcal{O}(\sqrt{m})$-approximation algorithm for it. For directed networks, they can show that the problem is hard to approximate within a factor $n^{\frac{1}{2}-\epsilon}$, for any $\epsilon > 0$.

For fractional length-bounded multi-commodity flows in graphs with edge-capacities and edge lengths Baier [13] gives a fully polynomial time approximation scheme (FPTAS). This FPTAS also yields a polynomial time algorithm for fractional length-bounded multi-commodity flows and fractional length-bounded edge-(multi-)cuts in unit-length graphs.

Mahjoub and McCormick [14] present a polynomial algorithm for the 3-length-bounded edge-cut in undirected graphs. Furthermore, they show 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 path 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, see e.g. [15,16,17,18].

Our Contribution. We present various results concerning the complexity and approximability of length-bounded cut and flow problems. After the preliminaries, in Section 3, we show that the minimum length-bounded cut problem in graphs with unit edge-lengths is $\mathcal{NP}$-hard to approximate within a factor of at least 1.1377 for $L \geq 5$ in the case of node-cuts and for $L \geq 4$ in the case of edge-cuts; see Table 1 for an overview of known and new complexity results. We also give approximation algorithms of ratio $\min\{L, n/L\}$ in the node case and $\min\{L, n^2/L^2, \sqrt{m}\}$ in the edge case. For classes of graphs such as constant degree expanders, hypercubes, and butterflies, we state an $\mathcal{O}(\log n)$-approximation algorithm pointed out by [19]. Furthermore, we give instances for which the integrality gap of the LP relaxation is $\Omega(\sqrt{n})$. Section 4 discusses the maximum

Table 1. Known and new (bold type) complexity results; $\varepsilon \in \mathbb{R}^+$ and $c \in \mathbb{N}$ are constants, ε can be arbitrarily small

L	node-cut	edge-cut
1	—	poly.
2	poly.	poly.
3	poly.	poly. [14] (undirected)
4	poly. [6] (undirected)	**inapprox. within 1.1377** (directed & undirected)
$5 \dots \lfloor n^{1-\varepsilon} \rfloor$	**inapprox. within 1.1377** (directed & undirected)	**inapprox. within 1.1377** (directed & undirected)
$n - c$	**poly.** (directed & undirected)	

length-bounded flow problem. For series-parallel graphs with unit edge lengths and unit edge-capacities, we proof a lower bound of $\Omega(\sqrt{n})$ on the integrality gap of the LP formulation. Furthermore, we show that edge- and path-flows are not polynomially equivalent for length-bounded flows: there is no polynomial algorithm to transform an edge-flow which is known to correspond to a length-bounded path-flow into a length-bounded path-flow. 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.

2 Preliminaries

We consider (directed or undirected) graphs $G = (V, E)$ with node set $V = V(G)$ and edge set $E = E(G)$. The number of nodes are denoted by n and the number of edges are denoted by m. A graph may contain *multi-edges*, i.e. parallel edges, in which case the graph will be called a *multi-graph*. Sometimes, we call an edge *simple* to distinguish it clearly from multi-edges. The graph G possesses two independent weights, an edge-capacity function $u : E \rightarrow \mathbb{Q}_{>0}$ and an edge-length function $d : E \rightarrow \mathbb{Q}_{\geq 0}$. If not stated otherwise, we assume unit-lengths and capacities.

Length-Bounded Cuts. Let $s, t \in V$ be two distinct nodes. We call a subset of edges C_e of G an *s-t-edge-cut*, if no path remains from s to t in $G \setminus C_e$. The *value* (or *capacity*) of C_e is the number of edges in C_e (or the total capacity of edges in C_e, if edge-capacities are not unit). Similarly, a node set C_n of G which separates s and t (and contains neither s nor t) is defined as an *s-t-node-cut*; its *value* is the number of nodes in C_n.

Let $\mathcal{P}_{s,t}(L)$ denote the set of all s-t-paths with length at most L. We call a subset of edges C_e^L of G an *L-length-bounded s-t-edge-cut*, if the nodes s and t have a distance greater than L in $G \setminus C_e^L$. This means that C_e^L must hit every path in $\mathcal{P}_{s,t}(L)$. Similarly, a subset C_n^L of the node set of G is called *L-length-bounded s-t-node-cut* if it destroys all paths in $\mathcal{P}_{s,t}(L)$. All of our cuts are s-t-cuts and therefore we will often omit the s-t-prefix. If the type of a cut is clear from the context, we will also omit the superscript L of C as well as the indices e and n of C. The *value* (or *capacity*) of a length-bounded cut is defined as in the standard cut case. In the *Minimum Length-Bounded Cut* problem we are looking for a length-bounded cut of minimum value.

In the linear programming relaxation of the minimum length-bounded edge-cut problem one has to assign to each edge $e \in E$ a dual length ℓ_e such that the dual length of a shortest s-t-path from $\mathcal{P}_{s,t}(L)$ is at least 1 (the LP relaxation for node-cuts is analogous):

$$\min \sum_{e \in E} u_e \ell_e \quad \text{s.t.} \sum_{e \in P} \ell_e \geq 1 \quad (P \in \mathcal{P}_{s,t}(L)), \quad \ell_e \geq 0 \quad (e \in E) . \qquad (1)$$

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 cuts* since only a fraction of an edge may contribute to the cut.

Length-Bounded Flows. Length-bounded flows are flows along paths such that the length of every path is bounded. More precisely, an *L-length-bounded s-t-flow* is a function $f : \mathcal{P}_{s,t}(L) \to \mathbb{R}_{\geq 0}$ assigning a flow value f_P to each s-t-path P in G of length at most L. The sum $\sum_{P \in \mathcal{P}_{s,t}(L)} f_P$ is called the *s-t-flow value* of f. The flow f is *feasible*, if edge-capacities are obeyed, i. e., for each edge $e \in E$ the sum of the flow values of paths containing this edge must be bounded by its capacity u_e. Since all our flows are s-t-flows, we will often omit the s-t-prefix.

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

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

We will refer to feasible solutions of this linear program as *path-flows*. Note that the dual of (2) is the linear program (1) for the minimum length-bounded cut problem. One way to prove the maximum-flow minimum-cut equality for standard flows is to apply duality theory of linear programming. In the case of multiple commodities, a source- and sink-node pair (s_i, t_i) and a length-bound $L_i \geq 0$ is given for each commodity $i = 1, \ldots, k$. An *$(L_1, \ldots, L_k)$-length-bounded multi-commodity flow* f is a set of L_i-length-bounded s_i-t_i-flows f_i, for $i = 1, \ldots, k$.

3 Length-Bounded Cuts

It follows from linear programming duality that the maximum fractional length-bounded flow value equals the minimum fractional length-bounded cut value. In the case of 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 different. This is an immediate consequence of the integrality gap that we state in the following theorem.

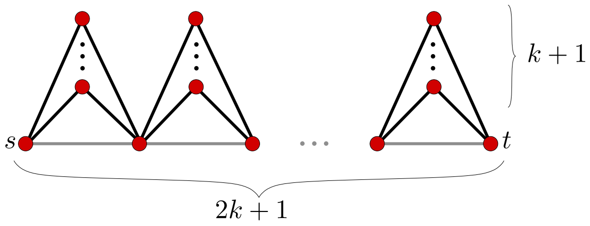

Fig. 1. Example of a large integrality gap of the linear program (1) of the minimum length-bounded cut. The straight s-t-path (*in gray*) contains $2k+1$ edges. Each of these edges is accompanied by $k+1$ parallel paths of length 2.

Theorem 1. *For (un-)directed series-parallel graphs the ratio of the minimum integral length-bounded edge-/node-cut value to the minimum fractional one can be of order $\Omega(\sqrt{n})$. In particular, the ratio of the minimum length-bounded edge-/node-cut size to the maximum number of length-bounded edge-disjoint paths can be of order $\Omega(\sqrt{n})$.*

Proof (sketch). The class of graphs depicted in Fig. 1 for $L = 3k + 1$ have a fractional length-bounded edge-cut value less than 2 but an integral length-bounded edge-cut value $k + 1 \in \theta(\sqrt{n})$. The result for node-cuts follows by considering the corresponding line graph. □

3.1 Complexity and Polynomially Solvable Cases

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

Theorem 2. *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 (un-)directed graphs.*

Proof. Enumerate all $V' \subseteq V$ with $|V'| \leq c$ and return the smallest V' which is a length-bounded node-cut, if there is any. Otherwise, any length-bounded node-cut V' contains at least $c+1$ nodes so that the longest remaining s-t-path has length at most $n - c - 1$ and therefore V' actually cuts all s-t-paths. Thus, returning a standard minimum node-cut suffices. □

Note that Theorem 2 does not carry over to the edge version of the problem, since by removing c edges one cannot guarantee that a standard cut suffices.

Theorem 3. *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 (un-)directed graphs within a factor of* 1.1377.

Proof. We first look at the case $L = 5$ and give a reduction from the well known Vertex Cover problem which has been shown to be $\mathcal{NP}$-hard to approximate within a factor ≈ 1.3606 [20]. Given a Vertex Cover instance G_{vc} with

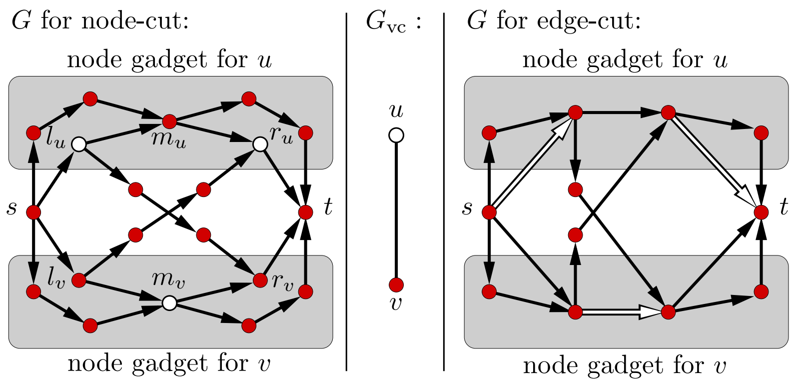

Fig. 2. Gadgets for the reduction of VERTEX COVER to length-bounded node-cut (*left*) and length-bounded edge-cut (*right*), respectively. Both correspond to two connected nodes u, v of the given VERTEX COVER instance, shown in the *middle*. The *highlighted nodes (edges)* are in the cut / vertex cover.

$n_{\text{vc}} = |V_{\text{vc}}|$ 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 $v \in V_{\text{vc}}$ we add a *node gadget* to G consisting of seven nodes which are interconnected with s, t and themselves as shown in Fig. 2 (left) – the nodes in the bottom half surrounded by a gray box. For each edge $\{u, v\} \in E_{\text{vc}}$ 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 Fig. 2 (left).

Lemma 1. *From a vertex cover V'_{vc} in G_{vc} of size x one can always construct a node-cut V' in G of size $n_{\text{vc}} + x$ and vice versa, for $x < n_{\text{vc}}$.*

We only deal with the easy direction "⇒" in Lemma 1 and omit all further details due to space limitations. Let $V'_{\text{vc}} \subseteq V_{\text{vc}}$ be a vertex cover with $|V'_{\text{vc}}| = x$. For each node $v \in V'_{\text{vc}}$ we add l_v and r_v to our cut $V' \subseteq V$ and for each node $u \in V_{\text{vc}} \setminus V'_{\text{vc}}$ we add m_u to V' (see Fig. 2 for an example). Note that $|V'| = n_{\text{vc}} + x$ and that no path of length at most 5 remains after removing V' from G.

The proof of Theorem 1.1 in [20] 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 $n_{\text{vc}} \cdot (1 - p + \varepsilon')$ exists, and the case where any vertex cover has size at least $n_{\text{vc}} \cdot (1 - 4p^3 + 3p^4 - \varepsilon')$, for any $\varepsilon' \in \mathbb{R}^+$ and $p = (3 - \sqrt{5})/2$. If we plug this into the result of Lemma 1, we have shown that the length-bounded node-cut is hard to approximate within a factor (there is an $\varepsilon' \in \mathbb{R}^+$ for which the inequality holds): $(n_{\text{vc}} + n_{\text{vc}} \cdot (1 - 4p^3 + 3p^4 - \varepsilon'))/(n_{\text{vc}} + n_{\text{vc}} \cdot (1 - p + \varepsilon')) > 1.1377$.

For other values of $L \in \{5, \ldots, \lfloor n^{1-\varepsilon} \rfloor\}$, we modify the construction of G as follows: (1) Add a path of length $L - 5$ from a new source node s' to s. Let s' be our new source. (2) Stepwise replace each node on this path after s' and until s (inclusive) by a group of $c \cdot n_{\text{vc}}$ nodes, for some constant c. For each of these groups connect all new nodes with all neighbors of the replaced node. We omit

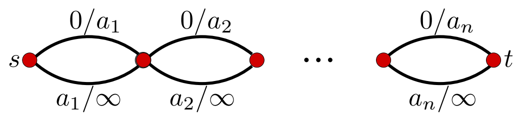

Fig. 3. Reduction of 2-PARTITION to the length-bounded cut problem. The *labels* denote length/capacity.

all further details. 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. □

The proof of the following theorem is similar to the proof of Theorem 3 with the difference that the adapted gadgets given in Fig. 2 (right) are used, which already work for length-bound $L = 4$.

Theorem 4. *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 (un-)directed (simple) graphs within a factor of* 1.1377*.*

Lemma 2. *For a series-parallel and outer-planar (un-)directed graph with edge-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 give a reduction from 2-PARTITION. Take an arbitrary 2-PARTITION instance $a_1, \dots, a_k \in \mathbb{N}$ and consider the graph in Fig. 3. Let the length-bound be $L = B - 1$. It is not difficult to see that there is an edge-cut of size at most B if and only if the instance of 2-PARTITION is a yes-instance. □

We will show in Theorem 8 that it is $\mathcal{NP}$-hard to decide whether a fractional length-bounded flow of given flow value exists even if the graph is outer-planar. Since the primal and dual programs have identical optimal objective function values, the same holds for the fractional length-bounded edge-cut problem.

3.2 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. Another extreme case is if the length-bound equals the distance between s and t, denoted by $\text{dist}(s, t)$. Lovász, Neumann Lara, and Plummer [6] show a special version of the following theorem in the context of length-bounded node-disjoint paths.

Theorem 5. *In (un-)directed multi-graphs with edge-capacities and lengths, for* $L = \text{dist}(s, t)$ *the minimum length-bounded edge-/node-cut and the maximum length-bounded flow problem can be solved efficiently. In particular, the max flow value and the min cut value coincide if* $L = \text{dist}(s, t)$*.*

The proof of this theorem is based on considering the sub-graph induced by all edges which are contained in at least one shortest s-t-path. For suitable length functions, like unit-edge-lengths, Theorem 5 yields the following approximation result for the minimum length-bounded cut problem.

Corollary 1. *In (un-)directed multi-graphs one can find an* $(L+1-\mathrm{dist}(s,t))$*-approximation to the minimum* L*-length-bounded cut.*

Proof (sketch). Repeatedly compute and remove a minimum $\mathrm{dist}(s,t)$-length-bounded cut from the graph until $\mathrm{dist}(s,t) > L$. □

It can be shown that the given performance ratio bound is tight for the sketched algorithm. The next theorem establishes bounds on the absolute difference between the sizes of standard minimum cuts and length-bounded minimum cuts.

Theorem 6. *Let* $G = (V, E)$ *be a (un-)directed multi-graph. A minimum node-cut in* G *is larger than a minimum length-bounded node-cut by at most* $\frac{n}{L}$*. If* G *is a simple graph, a minimum edge-cut is larger than a minimum length-bounded edge-cut by at most* $\mathcal{O}(\frac{n^2}{L^2})$*.*

Proof. The size of a minimum node-cut is equal to the maximum number of node-disjoint s-t-paths by Menger's theorem. Let C_1 be an optimal length-bounded node-cut. We 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 $L+1$. Thus, the number of node-disjoint s-t-paths in $G \setminus C_1$ is at most $(n-2)/L \leq n/L$. Therefore, a minimum node-cut C_2 in $G \setminus C_1$ has cardinality at most n/L. Then $C = C_1 \cup C_2$ is a node-cut in G of the desired cardinality. The proof for edge-cuts is similar. It applies a helpful lemma from [21] which states that if a (directed or undirected) simple graph contains k edge-disjoint s-t-paths, the shortest of these has length $\mathcal{O}(n/\sqrt{k})$. □

One can show that the bound of $\frac{n}{L}$ on the gap between standard and length-bounded node-cuts given in Theorem 6 is tight with a graph consisting of parallel paths of length $L+1$ except for one of them having length L. Theorem 6 leads to the following corollary.

Corollary 2. *For (un-)directed multi-graphs there exists an* $\mathcal{O}(\frac{n}{L})$*-approximation algorithm for the minimum length-bounded node-cut problem. For simple graphs (directed or undirected) there exists an* $\mathcal{O}(\frac{n^2}{L^2})$*-approximation algorithm for the minimum length-bounded edge-cut problem.*

Now we show that there are approximation algorithms with ratio $\mathcal{O}(\sqrt{n})$ for length-bounded node-cuts and with ratio $\mathcal{O}(\sqrt{m})$ for length-bounded edge-cuts.

Theorem 7. *For (un-)directed graphs there exists an* $\mathcal{O}(\min\{L, n/L, \sqrt{n}\})$*-approximation algorithm for the minimum length-bounded node-cut problem and an* $\mathcal{O}(\min\{L, n^2/L^2, \sqrt{m}\})$*-approximation algorithm for the minimum length-bounded edge-cut problem.*

Proof. The upper bounds of $\min\{L, n/L\}$ in the node case and $\min\{L, n^2/L^2\}$ in the edge case follow from Corollaries 1 and 2. Furthermore, we have $\min\{L, n/L\} \leq \sqrt{n}$, so the claimed ratio for length-bounded node cuts follows directly. It remains to show that ratio $\mathcal{O}(\sqrt{m})$ can be achieved for length-bounded edge-cuts.

Let OPT denote the size of a smallest length-bounded edge-cut. If $L \leq \sqrt{m}$, we simply apply the algorithm from Corollary 1. If $L > \sqrt{m}$, we repeatedly find an s-t-path of length at most $\lceil\sqrt{m}\rceil$, add all its edges to the cut, and remove these edges from the graph. Let C_1 denote the set of edges added to the cut in this process. Note that $|C_1| \leq \lceil\sqrt{m}\rceil \cdot \text{OPT}$.

If $G \setminus C_1$ does not contain an s-t-path of length at most L, we output C_1. Otherwise, we compute a minimum edge-cut C_2 in $G \setminus C_1$ and output $C_1 \cup C_2$. It suffices to show that $|C_2| \leq \sqrt{m}$. Let V_i denote the set of nodes at distance i from s in $G \setminus C_1$. Note that the distance from s to t in $G \setminus C_1$ is at least $\lceil\sqrt{m}\rceil + 1$. Let E_i be the set of edges in $G \setminus C_1$ with tail in V_i and head in V_{i+1}. Note that E_i is an edge-cut. Let j be such that E_j has minimum cardinality among the sets E_i for $0 \leq i \leq \lceil\sqrt{m}\rceil - 1$. Observe that $|E_j| \leq m/\lceil\sqrt{m}\rceil \leq \sqrt{m}$. □

For a large class of graphs a better approximation ratio is possible [19]: let F be the *flow number* of G, as defined in [22]. By the Shortening Lemma [22] it follows that if L is at least $4 \cdot F$, a standard minimum-cut is an $\mathcal{O}(1)$-approximation for the L-length-bounded cut. By Corollary 1 this gives an $\mathcal{O}(F)$-approximation for arbitrary L. Since $F = \mathcal{O}(\Delta\alpha^{-1}\log n)$, where Δ is the maximum degree and α the expansion, (cf. [22]) we obtain $\mathcal{O}(\log n)$-approximations for classes of graphs such as constant degree expanders, hypercubes, and butterflies.

4 Length-Bounded Flows

4.1 Edge-Based vs. Path-Based Flows: Complexity

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

Proposition 1. *Given a length-bounded (multi-commodity) path-flow in a graph with edge-capacities and lengths, and m edges. There exists a length-bounded (multi-commodity) path-flow with the same length bound and the same flow value per edge and commodity that uses at most m paths for each commodity.*

The proof of Proposition 1 follows from the fact, that the linear program in (2) has only m linear constraints. Thus, the theory of linear programming can be used to show that there is always a path-flow of maximum flow value which has a small size. Nevertheless, linear programming cannot be used to find maximum fractional length-bounded flows efficiently, unless $\mathcal{P} = \mathcal{NP}$.

Theorem 8. *For a single-commodity length-bounded flow problem in an (un-)directed outer-planar graph with edge-capacities and lengths it is $\mathcal{NP}$-complete to decide whether there is a fractional length-bounded flow of given flow value.*

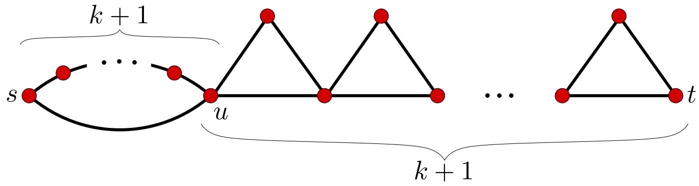

Fig. 4. Graph G_k in which the unique maximum length-bounded flow sends more than one half of the flow along paths with small flow values

Proof (sketch.). The 2-Partition problem can be reduced to the integral length-bounded flow problem for a flow of value 2 (similar to the proof of Lemma 2; see also Fig. 3). In a second step one shows that a fractional flow of value 2 in this special graph induces an integral flow of value 2. □

Finding a maximum length-bounded flow is computationally more difficult than finding a standard maximum flow. Standard flows are usually modeled as edge-flows. Each flow in a path formulation can easily be transformed into an edge-flow. For standard flows the reverse transformation is also possible. If length-bounds are present, one may try to use an edge-flow formulation, too. However, as the following theorem shows, edge- and path-flows are not polynomially equivalent for length-bounded flows. The following result is an immediate consequence of the proof of Theorem 8 and has been shown independently by Correa et al. [23, Corollary 3.4].

Corollary 3. *Unless $\mathcal{P} = \mathcal{NP}$, there is no polynomial algorithm to transform an edge-flow which is known to correspond to a length-bounded path-flow into a length-bounded path-flow, even if the graph is outer-planar.*

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 maximum flow ships a large percentage of the flow along paths with very small flow values.

Theorem 9. *There are unit-capacity outer-planar graphs of order n such that every maximum length-bounded flow ships more than one half of the total flow along paths with flow values $\mathcal{O}(1/n)$.*

Proof (sketch). Consider the family of unit-capacity and unit-length graphs depicted in Fig. 4. One can show that the unique maximum $(2k+2)$-length-bounded flow contains $k+1$ paths each with flow value $\frac{1}{k+1}$ and one path with flow value $\frac{k}{k+1}$ (using the sub-path of length $k+1$ between nodes s and u). □

For integral length-bounded flows there is a surprising structural difference between path- and edge-flows which is stated in Theorem 10.

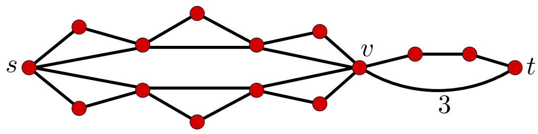

Fig. 5. A unit-length graph with an integral edge-flow of value 4 that corresponds to a maximum fractional 6-length-bounded path-flow but which has no integral 6-length-bounded path-decomposition: edge vt has capacity 3, all other edges have unit capacity

Theorem 10. *An integral (maximum) edge-flow corresponding to a (fractional) length-bounded flow in an (un-)directed graph with unit-edge-lengths does not need to have an integral length-bounded path decomposition.*

Proof (sketch). The graph in Fig. 5 has a unique maximum 6-length-bounded flow of value 4. The flow on each edge equals its capacity and is thus integral. But the unique length-bounded path decomposition is half-integral. □

In [13] it was shown that the 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.

Theorem 11. *For unit-capacity graphs with n nodes, the integrality gap of the integer program in* (2) *can be of order $\Omega(\sqrt{n})$ even for unit-edge-lengths and planar graphs. The length-bound used is of order $\Theta(\sqrt{n})$.*

The proof is based on a unit-capacity graph with n nodes, a maximum integral length-bounded flow of value 1, and a maximum half-integral flow of value $\Omega(\sqrt{n})$. The structure of this graph is a refinement of half a k by k grid. The construction is inspired by Guruswami et al. [12]. We omit all further details in this extended abstract.

The big integrality gap in Theorem 11 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 2. Indeed, the integrality gap is constant for high capacity graphs. The following result is a consequence of the randomized rounding technique of Raghavan and Thompson [24].

Theorem 12. *Consider a 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.*

Acknowledgement. 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.

References

1. Menger, K.: Zur allgemeinen Kurventheorie. Fund. Mathematicae (1927) 96–115
2. Dantzig, G.B., Fulkerson, D.R.: On the max flow min cut theorem of networks. In Kuhn, H.W., Tucker, A.W., eds.: Linear Inequalities and Related Systems. Volume 38 of Annals of Math. Studies. Princeton University Press (1956) 215–221
3. Kotzig, A.: Connectivity and Regular Connectivity of Finite Graphs. PhD thesis, Vysoká Škola Ekonomická, Bratislava, Slovakia (1956)
4. Ford, L.R., Fulkerson, D.R.: Maximal flow through a network. Canadian Journal of Mathematics **9** (1956) 399–404
5. Elias, P., Feinstein, A., Shannon, C.E.: A note on the maximum flow through a network. IRE Transactions on Information Theory **2** (1956)
6. Lovász, L., Neumann Lara, V., Plummer, M.D.: Mengerian theorems for paths of bounded length. Periodica Mathematica Hungarica **9** (1978) 269–276
7. Exoo, G.: On line disjoint paths of bounded length. Discrete Mathematics **44** (1983) 317–318
8. Niepel, L., Safaríková, D.: On a generalization of Menger's Theorem. Acta Mathematica Universitatis Comenianae **42** (1983) 275–284
9. Bondy, J.A., Murty, U.: Graph Theory with Applications. North Holland (1976)
10. Boyles, S.M., Exoo, G.: A counterexample to a conjecture on paths of bounded length. Journal of Graph Theory **6** (1982) 205–209
11. Itai, A., Perl, Y., Shiloach, Y.: The complexity of finding maximum disjoint paths with length constraints. Networks **12** (1982) 277–286
12. Guruswami, V., Khanna, S., Rajaraman, R., Shepherd, B., Yannakakis, M.: Near-optimal hardness results and approximation algorithms for edge-disjoint paths and related problems. In: Proceedings of the Symposium on Theory of Computing, ACM (1999) 19–28
13. Baier, G.: Flows with Path Restrictions. PhD thesis, TU Berlin, Germany (2003)
14. Mahjoub, A.R., McCormick, S.T.: The complexity of max flow and min cut with bounded-length paths. unpublished Manuscript (2003)
15. Perl, Y., Ronen, D.: Heuristics for finding a maximum number of disjoint bounded paths. Networks **14** (1984) 531–544
16. Brandes, U., Neyer, G., Wagner, D.: Edge-disjoint paths in planar graphs with short total length. Technical Report 19, Universität Konstanz (1996)
17. Wagner, D., Weihe, K.: A linear-time algorithm for edge-disjoint paths in planar graphs. Combinatorica **15**(1) (1995) 135–150
18. Hsu, D.: On container width and length in graphs, groups, and networks. IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences **E77-A**(4) (1994)
19. Kolman, P.: Personal communication (2006)
20. Dinur, I., Safra, S.: On the hardness of approximating minimum vertex cover. Annals of Mathematics **162**(1) (2005) 439–486
21. Galil, Z., Yu, X.: Short length versions of Menger's Theorem. In: Proceedings of the Symposium on Theory of Computing, ACM (1995) 499–508
22. Kolman, P., Scheideler, C.: Improved bounds for the unsplittable flow problem. In: Proceedings of the Symposium on Discrete Algorithms, ACM (2002) 184–193
23. Correa, J.R., Schulz, A.S., Stier Moses, N.E.: Fast, fair, and efficient flows in networks. In: Operations Research, INFORMS (2006) to appear.
24. Raghavan, P., Thompson, C.: Randomized rounding: A technique for provably good algorithms and algorithmic proofs. Combinatorica **7** (1987) 365–374