Approximating Minimum-Size k-Connected Spanning Subgraphs via Matching

Joseph Cheriyan * *Ramakrishna Thurimella* †

September 15, 1996
(Revised: November 27, 1996, March 19, 1998)

Abstract

Abstract: An efficient heuristic is presented for the problem of finding a minimum-size k-connected spanning subgraph of an (undirected or directed) *simple* graph $G = (V, E)$. There are four versions of the problem, and the approximation guarantees are as follows:

minimum-size k-**node** connected spanning subgraph of an **undirected** graph	$1 + [1/k]$,
minimum-size k-**node** connected spanning subgraph of a **directed** graph	$1 + [1/k]$,
minimum-size k-**edge** connected spanning subgraph of an **undirected** graph	$1 + [2/(k + 1)]$,
and minimum-size k-**edge** connected spanning subgraph of a **directed** graph	$1 + [4/\sqrt{k}]$.

The heuristic is based on a subroutine for the degree-constrained subgraph (b-matching) problem. It is simple, deterministic, and runs in time $O(k|E|^2)$. The analyses of the heuristics for minimum-size k-**node** connected spanning subgraphs hinge on theorems of Mader.

For undirected graphs and $k = 2$, a (deterministic) parallel **NC** version of the heuristic finds a 2-node connected (or 2-edge connected) spanning subgraph whose size is within a factor of $(1.5 + \epsilon)$ of minimum, where $\epsilon > 0$ is a constant.

1 Introduction

Given an undirected or directed *simple* graph $G = (V, E)$, an efficient approximation algorithm[1] is presented for the problem of finding a k-connected ($k = 1, 2, 3, \ldots$) spanning subgraph $G' = (V, E')$ that has the minimum number of edges. Let n and m denote $|V|$ and $|E|$, respectively. There are four versions of the problem, depending on whether G is a graph (i.e., an undirected graph) or a digraph (i.e., a directed graph), and on whether the spanning subgraph G' is required to be k-node connected or k-edge connected. All four versions of the problem are NP-hard: the two problems on graphs are NP-hard for $k \geq 2$, and the two problems on digraphs are NP-hard for $k \geq 1$, [GJ 79].

*Department of Combinatorics and Optimization, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1. Supported in part by NSERC grant no. OGP0138432 (NSERC code OGPIN 007). email: `jcheriyan@dragon.uwaterloo.ca`

†Department of Mathematics and Computer Science, University of Denver, 2360 S. Gaylord St., Denver CO 80208. Supported in part by NSF Research Initiation Award grant CCR-9210604. email: `ramki@cs.du.edu` URL: `http://www.cs.du.edu/~ramki`

[1] An *α-approximation algorithm* for a combinatorial optimization problem runs in polynomial time and delivers a solution whose value is always within the factor α of the optimum value. The quantity α is called the *approximation guarantee* of the algorithm.

	Previous results		Results in this paper	
	Undirected Graphs	Digraphs	Undirected Graphs	Digraphs
k-ECSS	$2-[1/k]$ for $k \geq 2$ [K 96] 1.85 for $k \geq 2$ [KR 96] $1+\sqrt{O(\log n)/k}$ [Ka 94]	1.61 for $k=1$ [KRY 96] 2 for $k \geq 2$	$1+[2/(k+1)]$ **improves for** $k \geq 3$	$1+[4/\sqrt{k}]$ **improves for** $k \geq 17$
k-NCSS	1.5 for $k=2$ [GSS 93] 2 for $k \geq 3$	1.61 for $k=1$ [KRY 96] 2 for $k \geq 2$	$1+[1/k]$ **improves for** $k \geq 3$	$1+[1/k]$ **improves for** $k \geq 2$

Table 1: A summary of previous & new approximation guarantees for minimum-size k-edge connected spanning subgraphs (k-ECSS), and minimum-size k-node connected spanning subgraphs (k-NCSS).

1.1 Previous work

Results of Mader [Ma 71, Ma 72] (also see [Bo 78]) imply that every *minimal*[2] k-edge connected graph has at most kn edges, and every minimal k-node connected graph has at most kn edges. Clearly, a k-connected (i.e., k-node connected or k-edge connected) graph has at least $kn/2$ edges, since each node has degree $\geq k$. Similarly, every k-connected digraph has at least kn arcs (directed edges) since each node has outdegree $\geq k$, and results of Edmonds [Ed 72] and Mader [Ma 85] imply that every minimal k-connected digraph has at most $2kn$ arcs. These facts immediately imply a 2-approximation algorithm for all four versions of the problem, since there is an easy polynomial-time algorithm to find a minimal k-edge connected (or k-node connected) spanning subgraph of a given graph or digraph. For graphs, recent algorithmic work gives another easy and efficient method for finding a k-connected spanning subgraph whose size (i.e., number of edges) is at most kn. A k-edge connected spanning subgraph (V, E') is obtained by taking $E' = F_1 \cup F_2 \cup \ldots \cup F_k$, where F_i $(1 \leq i \leq k)$ is the edge set of a maximal (but otherwise arbitrary) spanning forest of $(V, E \backslash (F_1 \cup \ldots \cup F_{i-1}))$, see [Th 89, NI 92], and a k-node connected spanning subgraph (V, E') is obtained similarly, but now each F_i is a maximal scan-first-search spanning forest, see [NI 92, FIN 93, CKT 93].

In the approximate solution of NP-hard combinatorial optimization problems, it often turns out that finding a solution within a factor of two of optimum is almost trivial, but achieving (asymptotically) better approximation guarantees needs a deeper understanding of the problem. For example, consider the metric TSP, i.e., the Traveling Salesman Problem with edge weights satisfying the triangle inequality. Finding a solution whose value is within a factor of two of optimum is trivial. The Christofides heuristic [Ch 76] broke the 2-approximation barrier by employing a powerful idea: matching.

Given a graph, consider the problem of finding a minimum-size 2-edge connected spanning subgraph (2-ECSS), or a minimum-size 2-node connected spanning subgraph (2-NCSS). Several recent papers have focused on these two problems. Khuller & Vishkin [KV 94] achieved the first significant advance by obtaining approximation guarantees of 1.5 and 1.66 for the minimum-size 2-ECSS problem and the minimum-size 2-NCSS problem. Garg et al [GSS 93] improved the approximation guarantee of the latter problem to 1.5. These algorithms are based on depth-first search (DFS), and they do not imply efficient parallel algorithms for the PRAM model. Subsequently, Chong &

[2]A graph H is called *minimal* with respect to a property $\mathcal{P}$ if H possesses $\mathcal{P}$, but for every edge e in H, $H \backslash e$ does *not* possess $\mathcal{P}$.

Lam [CL 95, CL 96] gave (deterministic) **NC** algorithms on the PRAM model with approximation guarantees of $(1.5+\epsilon)$ and $(1.66+\epsilon)$ for the minimum-size 2-ECSS problem and the minimum-size 2-NCSS problem.

For graphs and the general minimum-size k-ECSS problem, first Karger [Ka 94] used randomized rounding to improve the approximation guarantee (for k large w.r.t. $\log n$) to $1+\sqrt{[O(\log n)/k]}$; Karger's algorithm is not deterministic but Las Vegas. Then Khuller & Raghavachari [KR 96] improved the approximation guarantee (for all k) from 2 to (roughly) 1.85. They left open the problem of improving on the approximation guarantee of two for the minimum-size k-NCSS problem.

For digraphs and the problem of finding a minimum-size 1-connected (i.e., strongly connected) spanning subgraph, Khuller, Raghavachari and Young [KRY 96, KRY 95] gave a 1.61-approximation algorithm. For digraphs and $k \geq 2$, there appears to have been no previous work on achieving approximation guarantees better than two.

1.2 An illustrative example

Here is an example illustrating the difficulty in improving on the 2-approximation guarantee for the minimum-size k-connected spanning subgraph problem. Let the given graph G have n nodes, where n is even. Suppose that the edge set of G, $E(G)$, is the union of the edge set of the complete bipartite graph $K_{k,(n-k)}$ and the edge set E_{opt} of an n-node, k-regular, k-edge connected (or k-node connected) graph. For example, for $k=2$, $E(G)$ is the union of $E(K_{2,(n-2)})$ and the edge set of a Hamiltonian cycle. A naive heuristic may return $E(K_{k,(n-k)})$ which has size $k(n-k)$, roughly two times $|E_{opt}|$. A heuristic that significantly improves on the 2-approximation guarantee must somehow return many edges of E_{opt}.

1.3 Results in this paper

Heuristics and approximation guarantees. This paper first presents a simple heuristic for finding an approximately minimum-size k-NCSS of a given graph or digraph. An approximation guarantee of $1+[1/k]$ is proved. A variant of the heuristic finds a small-size k-ECSS of a given graph or digraph. For graphs and the minimum-size k-ECSS problem, the approximation guarantee is $1+[2/(k+1)]$. For digraphs and the minimum-size k-ECSS problem, the approximation guarantee is $1+[4/\sqrt{k}]$. Let $G=(V,E)$ be the given graph. The heuristic has two steps. The first step finds a minimum-size subgraph (V,M) of minimum-degree k (or $k-1$) via a subroutine for the degree-constrained subgraph (b-matching) problem. The second step adds an (inclusionwise) minimal edge set $F \subseteq E \backslash M$ such that the resulting graph $(V, M\cup F)$ is either k-node connected or k-edge connected, as required. Heuristics of this type have been considered by other researchers, but we were not aware of this when the preliminary version of this paper (Proc. IEEE FOCS'96) appeared. Subsequently, S. Khuller (personal communication, October 1996) and T. Watanabe (personal communication, October 1996) informed that they had examined or implemented heuristics of this type. One of the contributions of this paper is to refine the general heuristic to the four minimum-size k-CSS problems discussed above, and to give nearly tight analyses of the four approximation guarantees. The running time of the heuristic is $O(k|E|^2)$, and for graphs the running time improves to $O(k^3|V|^2+|E|^{1.5}(\log|V|)^2)$. The analyses on graphs/digraphs of the minimum-size k-NCSS heuristic are based on theorems of Mader [Ma 72, Ma 85]. In the context of augmenting the node connectivity of graphs and digraphs, the first application of Mader's theorems is due to Jordán [Jo 95, Jo 93]. Two key lemmas in our analyses, namely, Lemmas 3.3 and 3.19, are inspired by similar results of Jordán, namely, [Jo 95, Lemma 3.3] and the following paragraph in [Jo 95] and Lemma 2.6 and Corollary 2.7 in [Jo 93]. In the context of approximation algorithms for minimum-

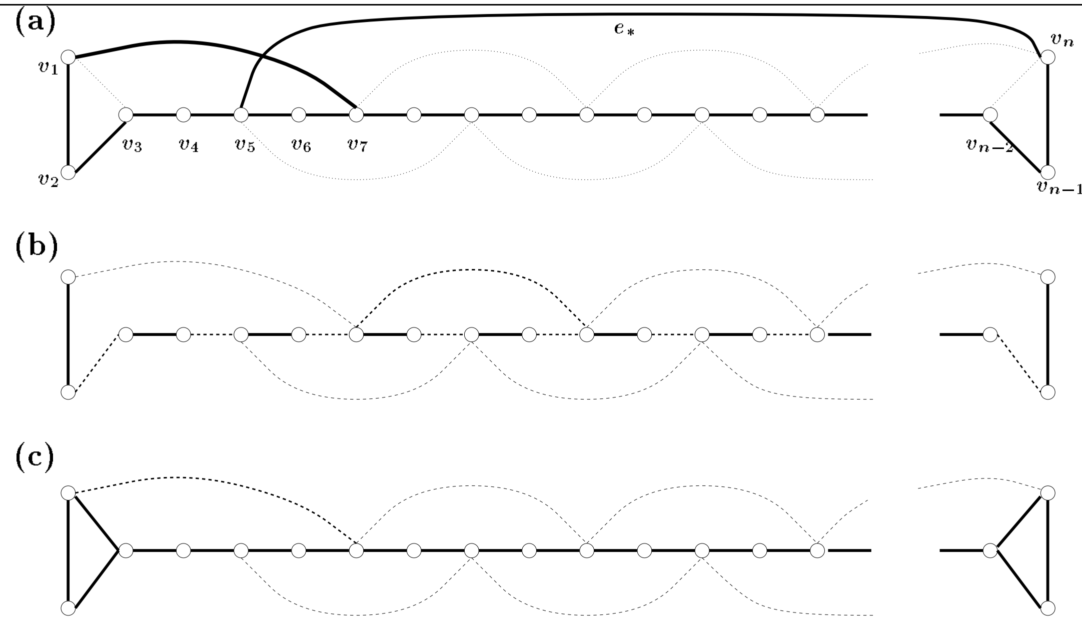

Figure 1: Illustrating the 2-NCSS heuristic on a 2-node connected graph $G = (V, E)$; $n = |V|$ is even, and $k = 2$. Adapted from Garg, Santosh & Singla [GSS 93, Figure 7].
(a) A minimum-size 2-node connected spanning subgraph has $n+1$ edges, and is indicated by thick lines (the path $v_1, v_2, \ldots, v_n$ and edges v_1v_7 and $e_* = v_5v_n$).
(b) The first step of the heuristic in Section 3.1 finds a minimum-size $M \subseteq E$ such that every node is incident to $\geq (k-1) = 1$ edges of M. The thick lines indicate M; it is a perfect matching. The second step of the heuristic finds an (inclusionwise) minimal edge set $F \subseteq E$ such that $(V, M \cup F)$ is 2-node connected. F is indicated by dashed lines – the "key edge" e_* is not chosen in F. $|M \cup F| = 1.5n - 5$.
(c) Another variant of the heuristic first finds a minimum-size $M \subseteq E$ such that every node is incident to $\geq k = 2$ edges of M. The thick lines indicate M (M is the path $v_1, v_2, \ldots, v_n$ and edges v_1v_3, $v_{n-2}v_n$). The second step of the heuristic finds the edge set $F \subseteq E$ indicated by dashed lines – the "key edge" e_* is not chosen in F. $(V, M \cup F)$ is 2-node connected, and for every edge vw in F, $(V, M \cup F) \backslash vw$ is *not* 2-node connected. $|M \cup F| = 1.5n - 3$.

size k-connected spanning subgraph problems, Chong & Lam [CL 95] appear to be the first to use matching.

For graphs, the heuristic finds a 2-node connected or 2-edge connected spanning subgraph whose size is within a factor of 1.5 of the minimum size. A parallel (deterministic) version gives a $(1.5+\epsilon)$-approximation **NC** algorithm. Similarly, a sequential *linear-time* version gives an approximation guarantee of $(1.5 + \epsilon)$.

Independently of this paper, and using different methods, Chong and Lam [CL 96b] have also obtained a parallel (deterministic) $(1.5+\epsilon)$-approximation **NC** algorithm for the minimum-size 2-NCSS problem on graphs. Recently, Fernandes [Fe 97, Theorem 5.1] showed that the minimum-size 2-ECSS problem on graphs is MAX SNP-hard.

Table 1 summarizes the approximation guarantees obtained in this paper for the four versions of the problem, and compares these with the previous best approximation guarantees. Figure 1 illustrates the working of the heuristic on an example.

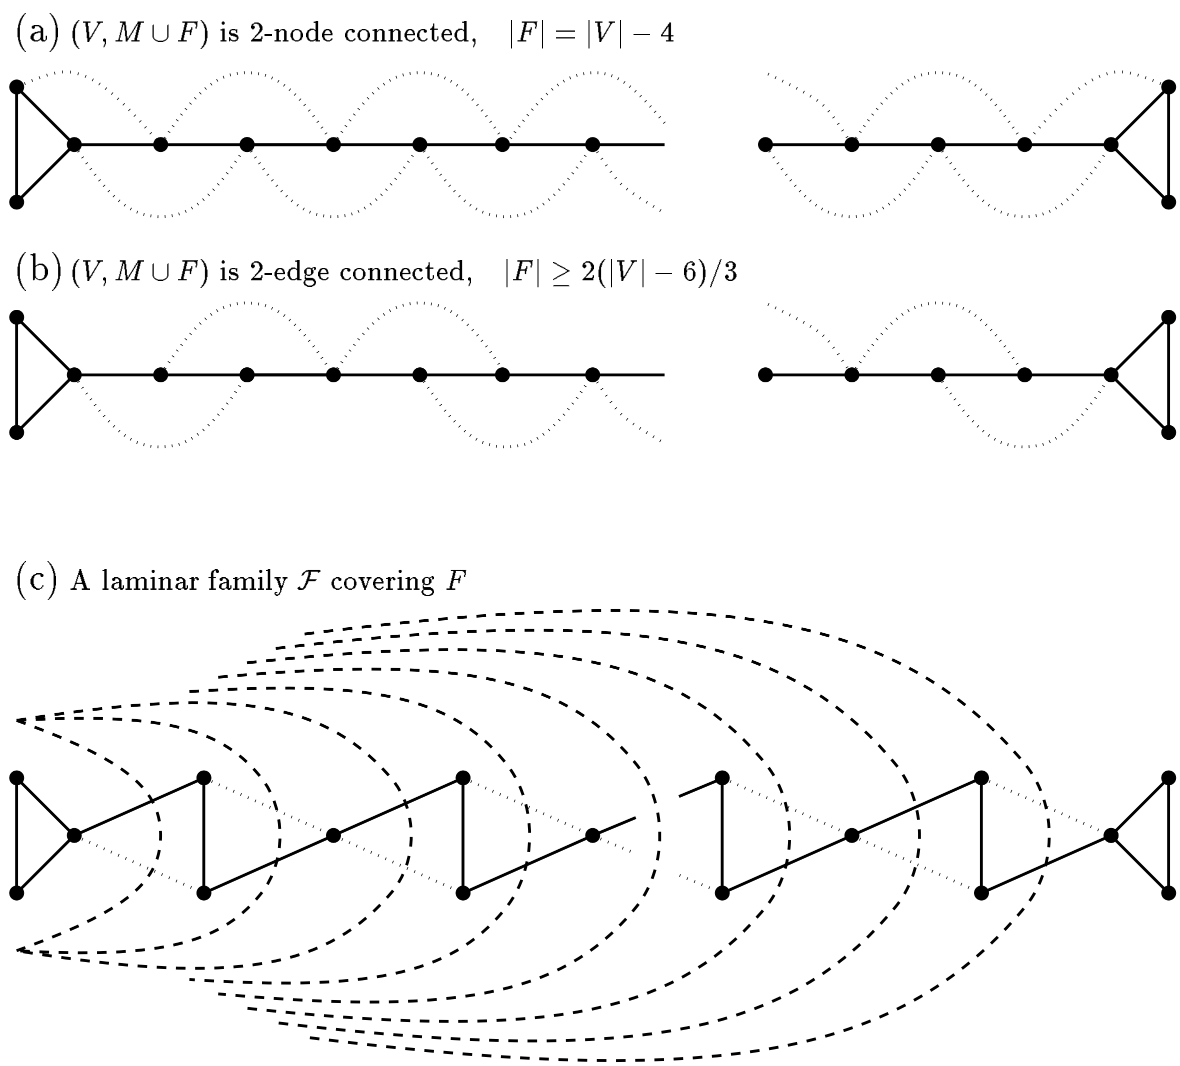

Figure 2: An illustration of Lemma 3.3 (a corollary of Mader's theorem, Theorem 3.2) and of Theorem 4.3. An n-node graph of minimum degree $k = 2$, (V, M), is indicated by solid lines.
(a) The dotted lines indicate an (inclusionwise) minimal edge set F such that $(V, M \cup F)$ is 2-node connected. F has size $n - 4$, for $n \geq 4$. By Lemma 3.3, the maximum size of F over all possible M is $\leq n - 1$.
(b) The dotted lines indicate an (inclusionwise) minimal edge set F such that $(V, M \cup F)$ is 2-edge connected. F has size $\geq 2(n - 6)/3$, for $n \geq 6$. By Theorem 4.3, the maximum size of F over all possible M is $\leq 2(n - 1)/3$.
(c) The dashed lines indicate a laminar family of tight node sets $\mathcal{F}$ covering the F-edges of the 2-edge connected graph in (b). The proof of Theorem 4.3 is based on examining M, F and $\mathcal{F}$.

Contributions to approximation algorithms for "uniform" network design. As discussed above, the subarea of network design with uniform edge costs and uniform connectivity requirements has attracted a fair amount of recent interest in theoretical computer science, e.g., the references cite ten papers from this subarea, as well as a survey paper [K 96]. This paper takes up four central questions from this subarea, and settles them in the sense that reasonably good approximation guarantees are derived based on a simple heuristic. (For the sake of referees who may have to make an extrinsic comparison, we mention that this paper subsumes some of the main results in eight of the recent papers cited in the references.) To achieve the approximation guarantees, the paper has to rely on some deep areas of graph theory and combinatorial optimization.

Combinatorial contributions. The paper has two combinatorial results that may be of independent interest. The first is Theorem 3.5 that gives a new lower bound on the size of a k-edge connected spanning subgraph. The proof relies on the Gallai-Edmonds decomposition theorem of matching theory. Theorem 3.5 is related to a result of R. P. Gupta: a bipartite graph of minimum degree k has k edge-disjoint edge covers. Theorem 3.5 implies some interesting results in matching theory such as Petersen's theorem (see Corollary 3.11), and the following: "except for cycles of odd length, every 2-edge connected graph has two edge-disjoint edge covers" (see Corollary 3.12). The second combinatorial result of independent interest is Theorem 4.3. This theorem gives an asymptotically tight upper bound of $k|V|/(k+1)$ on the size of an (inclusionwise) minimal edge set F such that $(V, M \cup F)$ is a k-edge connected (simple) graph, where (V, M) is a graph of minimum degree $\geq k$. The proof makes use of a laminar family of tight node sets that covers F. The proof is long, and at several points, novel arguments have to be developed. Theorem 4.3 is related to a theorem of Mader on "critical cycles" in a k-node connected graph, see Theorem 3.2. Apparently, Mader's theorem has *no* analogue for k-edge connected graphs; for $k = 2$, this can be seen from the example in Figure 5; the example generalizes to all $k \geq 2$. However, there is one implication of Mader's theorem that is an analogue of Theorem 4.3: If (V, M) is as above, and F is an (inclusionwise) minimal edge set such that $(V, M \cup F)$ is a k-node connected graph, then $|F| \leq |V| - 1$ (see Lemma 3.3). Both the bounds ($k|V|/(k+1)$ in Theorem 4.3, and $|V| - 1$ in Lemma 3.3) are tight up to an additive term of $(k+1)$, for all $k \geq 2$. Figure 2 has relevant examples for $k = 2$, and these examples generalize for all $k \geq 2$. Although Theorem 4.3 and Lemma 3.3 are analogous, the two results seem to be focusing on two essentially different combinatorial structures, and neither result implies the other one.

Organization of the paper. The rest of the paper is organized as follows. Section 2 has definitions and notation. Section 3 presents the heuristic for approximating a minimum-size k-node connected spanning subgraph of a graph or a digraph, and separately analyzes the approximation guarantees on graphs and digraphs. Section 4 describes and analyzes the heuristic for approximating a minimum-size k-edge connected spanning subgraph of a graph or a digraph. Section 5 has conclusions, including a discussion of the relationship to extremal graph theory.

2 Definitions and notation

For a subset S' of a set S, $S \backslash S'$ denotes the set $\{x \in S : x \notin S'\}$.

This paper considers finite *simple* graphs and digraphs, i.e., the graphs/digraphs have no loops nor multiedges. (But, Propositions 3.9 and 3.10 do allow multiedges.) Let $G = (V, E)$ be a graph or a digraph. $V(G)$ and $E(G)$ stand for the node set and the edge set of G. By the *size* of G we mean $|E(G)|$. First, suppose that G is a graph. An edge incident to nodes v and w is denoted by vw. For a subset M of E and a node v, we use $\deg_M(v)$ to denote the number of edges of M

incident to v; $\deg(v)$ denotes $\deg_E(v)$.

A node is said to be *covered* by an edge set M if the node is incident to at least one edge of M; otherwise, the node is *uncovered* by M. An *edge cover* is a set of edges that covers all the nodes. A *matching* of a graph $G = (V, E)$ is an edge set $M \subseteq E$ such that $\deg_M(v) \leq 1$, $\forall v \in V$; furthermore, if every node $v \in V$ has $\deg_M(v) = 1$, then M is called a *perfect matching*. A graph G is called *factor-critical* if for every node $v \in V$, there is a perfect matching in $G \backslash v$, see [LP 86].

An $x \leftrightarrow y$ path refers to a path whose end nodes are x and y. We call two paths *openly disjoint* if every node common to both paths is an end node of both paths. Hence, two (distinct) openly disjoint paths have no edges in common, and possibly, have no nodes in common. A set of $k \geq 2$ paths is called openly disjoint if the paths are pairwise openly disjoint. For a node set $S \subseteq V(G)$, $\delta_G(S)$ denotes the set of all edges in $E(G)$ that have one end node in S and the other end node in $V(G) \backslash S$ (when there is no danger of confusion, the notation is abbreviated to $\delta(S)$); $\delta(S)$ is called a *cut*, and by a *k-cut* we mean a cut that has exactly k edges.

A graph $G = (V, E)$ is said to be k-edge connected if $|V| \geq k + 1$ and the deletion of any set of $< k$ edges leaves a connected graph. A graph $G = (V, E)$ is said to be *k-node connected* if $|V| \geq k + 1$, and the deletion of any set of $< k$ nodes leaves a connected graph.

Let $G = (V, E)$ be a digraph. An arc (directed edge) with start node v and end node w is denoted (v, w). For $M \subseteq E$ and a node v, $\deg_{M,out}(v)$ ($\deg_{M,in}(v)$) denotes the number of arcs of M with start node v (end node v). For a node set $S \subseteq V$, $\delta_{out}(S)$ ($\delta_{in}(S)$) denotes the set of arcs with start nodes in S and end nodes in $V \backslash S$ (end nodes in S and start nodes in $V \backslash S$). The digraph is called *strongly connected* (1-connected) if for every (ordered) pair of nodes v,w, there exists a directed path from v to w. The digraph is called *k-edge connected* if $|V| \geq k + 1$, and the deletion of any set of $< k$ arcs leaves a strongly-connected digraph. The digraph is called *k-node connected* if $|V| \geq k + 1$, and the deletion of any set of $< k$ nodes leaves a strongly-connected digraph.

An edge vw (arc (v, w)) of a k-node connected graph G (digraph G) is called *critical* w.r.t. k-node connectivity if $G \backslash vw$ ($G \backslash (v, w)$) is *not* k-node connected. Similarly, we have the notion of critical edges (arcs) w.r.t. k-edge connectivity.

Let $G = (V, E)$ be a graph, and let $b : V \rightarrow \mathbb{Z}_+$ assign a nonnegative integer b_v to each node $v \in V$. The perfect b-matching (or perfect degree-constrained subgraph) problem is to find an edge set $M \subseteq E$ such that each node v has $\deg_M(v) = b_v$. The maximum b-matching (or maximum degree-constrained subgraph) problem is to find a maximum-cardinality $M \subseteq E$ such that each node v has $\deg_M(v) \leq b_v$. The b-matching problem can be solved in time $O(|E|^{1.5}(\log |V|)^{1.5}\sqrt{\alpha(|E|, |E|)})$, see [GaTa 91, Section 11] (for our version of the problem, note that each edge has unit cost and unit capacity, and each node v may be assumed to have $0 \leq b_v \leq \deg(v)$). Also, see [Ge 95, Section 7.3] and [Ga 85].

3 A $(1 + \frac{1}{k})$-approximation algorithm for minimum-size k-node connected spanning subgraphs

This section presents the heuristic for finding an approximately minimum-size k-node connected spanning subgraph (abbreviated k-NCSS), and proves an approximation guarantee of $1 + [1/k]$. First, we focus on graphs, and then turn to digraphs. The analysis of the heuristic for graphs hinges on a deep theorem of Mader [Ma 72, Theorem 1]. Given a graph $G = (V, E)$, a straightforward application of Mader's theorem shows that the number of edges in the k-NCSS returned by the heuristic is at most

$$(|V| - 1) + \min\{|M| \,:\, M \subseteq E \text{ and } \deg_M(v) \geq (k - 1),\, \forall v \in V\},$$

see Lemma 3.3 below. An approximation guarantee of $1+[2/k]$ on the heuristic follows, since the number of edges in a k-node connected graph is at least $k|V|/2$, by the "degree lower bound", see Proposition 3.4. Often, the key to proving improved approximation guarantees for (minimizing) heuristics is a nontrivial lower bound on the value of every solution. We improve the approximation guarantee from $1+[2/k]$ to $1+[1/k]$ by exploiting a new lower bound on the size of a k-edge connected spanning subgraph, see Theorem 3.5:

> The number of edges in a k-edge connected spanning subgraph of a graph $G = (V, E)$ is at least $\lfloor |V|/2 \rfloor + \min\{|M| \,:\, M \subseteq E \text{ and } \deg_M(v) \geq (k-1), \forall v \in V\}$.

The analysis of the heuristic for digraphs is similar, and hinges on another theorem of Mader [Ma 85, Theorem 1], which may be regarded as the generalization of [Ma 72, Theorem 1] to digraphs. An approximation guarantee of $1+[1/k]$ is proved on the digraph heuristic by employing a simpler version of Theorem 3.5, namely Proposition 3.8, to give a lower bound on the number of edges in a solution.

Assume that the given graph or digraph $G = (V, E)$ is k-node connected, otherwise, the heuristic will detect this and report failure.

3.1 Undirected graphs

Let $E^* \subseteq E$ denote a minimum-cardinality edge-set such that the spanning subgraph (V, E^*) is k-*edge* connected. Note that every k-node connected spanning subgraph (V, E') (such as the optimal solution) is necessarily k-edge connected, and so has $|E'| \geq |E^*|$.

The heuristic has two steps. The first finds a *minimum-size* spanning subgraph (V, M), $M \subseteq E$, whose minimum degree is $(k-1)$, i.e., each node is incident to $\geq (k-1)$ edges of M. Clearly, $|M| \leq |E^*|$, because (V, E^*) has minimum degree k, i.e., every node is incident to $\geq k$ edges of E^*. To find M efficiently, we use the algorithm for the maximum degree-constrained subgraph (b-matching) problem. Our problem is:

$$\min\{|M| : \deg_M(v) \geq (k-1),\ \forall v \in V,\ \text{and}\ M \subseteq E\}.$$

To see that this is a b-matching problem, consider the equivalent problem of finding the complement $\overline{M}$ of M w.r.t. E, where $\overline{M} = E \backslash M$:

$$\max\{|\overline{M}| : \deg_{\overline{M}}(v) \leq \deg(v) + 1 - k,\ \forall v \in V,\ \text{and}\ \overline{M} \subseteq E\}.$$

The b-matching problem can be solved in time $O(|E|^{1.5}(\log|V|)^2)$ see [GaTa 91], hence this running time suffices to find M.

The second step is equally simple. We find an (inclusionwise) minimal edge set $F \subseteq E \backslash M$ such that $M \cup F$ gives a k-node connected spanning subgraph, i.e., $(V, M \cup F)$ is k-node connected and for each edge $vw \in F$, $(V, M \cup F) \backslash vw$ is *not* k-node connected. Recall that an edge vw of a k-node connected graph H is *critical* (w.r.t. k-node connectivity) if $H \backslash vw$ is *not* k-node connected. The next result characterizes critical edges.

Proposition 3.1 *An edge vw of a k-node connected graph H is not critical iff there are at least $k+1$ openly disjoint $v \leftrightarrow w$ paths in H (including the path vw).*

To find F efficiently, we start with $F = \emptyset$ and take the current subgraph to be $G = (V, E)$ (which is k-node connected). We examine the edges of $E \backslash M$ in an arbitrary order, say, $e_1, e_2, \ldots, e_\ell$ ($\ell = |E \backslash M|$). For each edge $e_i = v_i w_i$, we attempt to find $(k+1)$ openly disjoint $v_i \leftrightarrow w_i$ paths in the current subgraph. If we succeed, then we remove the edge e_i from the current subgraph (since e_i is not critical), otherwise, we retain e_i in the current subgraph and add e_i to F (since e_i is critical). At termination, the current subgraph with edge set $M \cup F$ is k-node connected, and every edge $vw \in F$ is critical. The running time for the second step is $O(k|E|^2)$.

The proof of the next lemma hinges on a theorem of Mader [Ma 72, Theorem 1]. For an English translation of the proof of Mader's theorem see Lemma I.4.4 and Theorem I.4.5 in [Bo 78].

Theorem 3.2 (Mader [Ma 72, Theorem 1]) *In a k-node connected graph, a cycle consisting of critical edges must be incident to at least one node of degree k.*

Lemma 3.3 $|F| \leq |V| - 1$.

Proof: Consider the k-node connected subgraph returned by the heuristic, $G' = (V, E')$, where $E' = M \cup F$. Suppose that F contains a cycle C. Note that every edge in the cycle is critical, since every edge in F is critical. Moreover, every node v incident to the cycle C has degree $\geq (k+1)$ in G', because v is incident to two edges of C, as well as to at least $(k-1)$ edges of $M = E' \backslash F$. But this contradicts Mader's theorem. We conclude that F is acyclic, and so has $\leq |V| - 1$ edges. The proof is done. □

Proposition 3.4 *Let $G = (V, E)$ be a graph of node connectivity $\geq k$. The heuristic above finds a k-node connected spanning subgraph (V, E') such that $|E'| \leq (1 + [2/k])|E_{opt}|$, where $|E_{opt}|$ denotes the cardinality of an optimal solution. The running time is $O(k^3|V|^2 + |E|^{1.5}(\log|V|)^2)$.*

Proof: The approximation guarantee follows because $|E_{opt}| \geq (k|V|/2)$, so

$$\frac{|M| + |F|}{|E_{opt}|} = \frac{|M|}{|E_{opt}|} + \frac{|F|}{|E_{opt}|} \leq 1 + \frac{|V|}{(k|V|/2)} = 1 + [2/k].$$

We have already seen that M can be found in time $O(|E|^{1.5}(\log|V|)^2)$, and F can be found in time $O(k|E|^2)$. The running time of the second step can be improved to $O(k^3|V|^2)$ as follows: we run a linear-time preprocessing step to compute a sparse certificate $\widetilde{E}$ of G for k-node connectivity, i.e., $\widetilde{E} \subseteq E$, $|\widetilde{E}| \leq k|V|$, and for all nodes v, w, $(V, \widetilde{E})$ has k openly disjoint $v \leftrightarrow w$ paths iff G has k openly disjoint $v \leftrightarrow w$ paths, see [NI 92, FIN 93, CKT 93]. We compute M as before, by running the first step on G. To find the set $F \subseteq E \backslash M$, we run the second step on $\widetilde{E} \cup M$ rather than on E, and for each edge $v_i w_i \in \widetilde{E} \backslash M$, we attempt to find $(k+1)$ openly disjoint $v_i \leftrightarrow w_i$ paths in the current subgraph of $(V, \widetilde{E} \cup M)$. The second step runs in time $O(k|\widetilde{E} \cup M|^2) = O(k^3|V|^2)$, since $|\widetilde{E} \cup M| = O(k|V|)$. □

To improve the approximation guarantee to $1 + [1/k]$, we present an improved lower bound on $|E^*|$, where E^* denotes a minimum-cardinality edge set such that $G^* = (V, E^*)$ is k-edge connected. Suppose that E^* contains a perfect matching P_0 (so $|P_0| = n/2$). Then $|E^*| \geq (n/2) + \min\{|M^*| \ : \ M^* \subseteq E, \ \deg_{M^*}(v) \geq (k-1), \forall v \in V\}$. To see this, focus on the edge set $M' = E^* \backslash P_0$. Clearly, every node $v \in V$ is incident to at least $(k-1)$ edges of M', because $\deg_{E^*}(v) \geq k$ and $\deg_{P_0}(v) = 1$. Since M^* is a minimum-size edge set with $\deg_{M^*}(v) \geq (k-1)$, $\forall v \in V$, we have $|M^*| \leq |M'| = |E^*| - (n/2)$. The next theorem generalizes this lower bound to the case when E^* has no perfect matching. The proof is given in the next subsection (Section 3.2), after developing some preliminaries.

Theorem 3.5 *Let $G^* = (V, E^*)$ be a graph of edge connectivity $\geq k \geq 1$, and let n denote $|V|$. Let $M^* \subseteq E^*$ be a minimum-size edge set such that every node $v \in V$ is incident to $\geq (k-1)$ edges of M^*. Then $|E^*| \geq |M^*| + \lfloor n/2 \rfloor$.*

Theorem 3.6 *Let $G = (V, E)$ be a graph of node connectivity $\geq k$. The heuristic described above finds a k-node connected spanning subgraph (V, E') such that $|E'| \leq (1 + [1/k])|E_{opt}|$, where $|E_{opt}|$ denotes the cardinality of an optimal solution. The running time is $O(k^3|V|^2 + |E|^{1.5}(\log|V|)^2)$.*

Proof: The approximation guarantee of $1 + [1/k]$ follows easily from Theorem 3.5, using an argument similar to Proposition 3.4. We have $E' = M \cup F$, where $|F| \leq (n-1)$. Moreover, since M is a minimum-size edge set with $\deg_M(v) \geq (k-1)$, $\forall v \in V$, Theorem 3.5 implies that $|M| \leq |E_{opt}| - \lfloor n/2 \rfloor \leq |E_{opt}| - (n-1)/2$. Hence,

$$\frac{|M| + |F|}{|E_{opt}|} \leq \frac{|E_{opt}| - (n-1)/2 + (n-1)}{|E_{opt}|} \leq 1 + \frac{n/2}{|E_{opt}|} \leq 1 + [1/k],$$

where the last inequality uses the "degree lower bound", $|E_{opt}| \geq kn/2$.

The running time analysis is the same as that in Proposition 3.4. □

3.2 A lower bound for the size of a k-connected spanning subgraph and Gupta's theorem on bipartite graphs

This subsection gives a proof of Theorem 3.5. This theorem is used in the previous subsection to prove an approximation guarantee of $1 + [1/k]$ for a minimum-size k-NCSS. Theorem 3.5 gives the following new lower bound on the size of a k-ECSS:

> Let $G^* = (V, E^*)$ be a k-edge connected graph ($k \geq 1$), and let n denote $|V|$. Let $M^* \subseteq E^*$ be a minimum-size edge set such that every node $v \in V$ is incident to $\geq (k-1)$ edges of M^*. Then $|E^*| \geq |M^*| + \lfloor n/2 \rfloor$.

First, a theorem of R. P. Gupta on bipartite graphs is recalled. For the special case of bipartite graphs, (a stronger form of) the lower bound in Theorem 3.5 follows easily from Gupta's theorem, see Proposition 3.8. This proposition is used in Section 3.4 to prove an approximation guarantee of $1 + [1/k]$ for a minimum-size k-NCSS of a digraph. Gupta's theorem does not apply to non-bipartite graphs. The proof of Theorem 3.5 (for arbitrary graphs) relies on the Gallai-Edmonds decomposition theorem of matching theory. When the Gallai-Edmonds decomposition of the graph is "nontrivial", one can define a bipartite graph B that partially represents the decomposition. The proof of Theorem 3.5 is completed by examining B. One way is to prove a variant of Gupta's theorem (see Proposition 3.9), and then apply it to B. This is described below. Readers interested in a detailed study of the proofs in this subsection may find it useful to review two results in matching theory, namely, the Gallai-Edmonds decomposition theorem [LP 86, Theorem 3.2.1], and the Hungarian method for bipartite matching [LP 86, Lemma 1.2.2].

Theorem 3.7 (Gupta [Gu 67]) *Let $B = (X \cup Y, E)$ be a bipartite graph with minimum degree k. Then there exists a partition of the edge set of B, namely E, into k sets $E_1, E_2, \ldots, E_k$ such that each node $v \in X \cup Y$ is incident to at least one edge from each set E_i, $1 \leq i \leq k$.*

For an elegant proof, see the solutions to Problems 10–12 in [L 93, Chapter 7]. Also, see [BM 76, Problem 6.1.6]. The next result strengthens Theorem 3.5 for bipartite graphs. The proof is via Gupta's theorem. Another brief proof follows from Proposition 3.10.

Proposition 3.8 *Let $B^* = (X \cup Y, E^*)$ be a bipartite graph with minimum degree $\geq k$. Let $M^* \subseteq E^*$ be a minimum-size edge set such that every node $v \in X \cup Y$ is incident to $\geq k-1$ edges of M^*. Then $|E^*| \geq |M^*| + (|X \cup Y|/2)$.*

Proof: Apply Gupta's theorem to E^*, and let $E_1, E_2, \ldots, E_k$ be the partition of E^*. Focus on the set, say E_k, that has the maximum cardinality. Clearly, $|E_k| \geq |E^*|/k \geq |X \cup Y|/2$. Now, consider $M' = E^* \backslash E_k$, and observe that each node $v \in X \cup Y$ is incident to $\geq (k-1)$ edges of M', because Gupta's result shows that v is incident to some edge from each of the remaining $(k-1)$ sets $E_1, E_2, \ldots, E_{k-1}$. The proof is done since $|E^*| - (|X \cup Y|/2) \geq |M'|$ and $|M'| \geq |M^*|$. □

Proposition 3.8 does *not* generalize to nonbipartite graphs B^*, even if we strengthen the condition "B^* has minimum degree $\geq k$" to "B^* is k-edge connected". For example, let $k = 2$, and let $B^* = K_3$, the complete graph on three nodes. Then M^* is a minimum edge cover of K_3, and has size two. But then $|E^*| = |M^*| + 1 < |M^*| + (|V|/2)$. The generalization of Proposition 3.8 fails because B^* is a 2-edge connected, 2-regular graph such that for every edge cover M^*, the edge-complement of M^* in B^*, $(V, E^* - M^*)$, has an isolated node, so it does not have an edge cover. For every *even* integer $k \geq 2$, there is an infinite family of nonbipartite graphs such that the generalization of Proposition 3.8 fails. Take B^* to be a k-edge connected, k-regular graph with an *odd* number of nodes n. Then M^* has size at least $(1 + (k-1)n)/2$, so $(V, E^* - M^*)$ has an isolated node, and hence has size $< n/2$. It can be seen that the examples in this paragraph are factor-critical graphs.

The next proposition may be regarded as a variant of Gupta's theorem. Note that the bipartite graph B in the next proposition may have minimum degree one, and B may have multiple copies of an edge.

Proposition 3.9 *Let $B = (X \cup Y, E)$ be a bipartite (loopless) multigraph with node bipartition $X \cup Y$. Let each node $y \in Y$ have $\deg(y) \geq k$, and let B have a matching of size $|X|$. Then B has an edge cover J such that each node $y \in Y$ is incident to exactly one edge of J, and each node $x \in X$ is incident to either exactly one edge of J or at least $(k-1)$ edges of $E \backslash J$.*

Proof: See Figure 3(b) for an illustration. Let J_0 be a matching of size $|X|$. The edge cover J is constructed iteratively, starting with $J' = J_0$ and $J'' = \emptyset$. Throughout, J' is a matching of the current B, and at the end of the construction, $J' \cup J''$ is an edge cover of the original B that satisfies the proposition.

If $J' \cup J''$ is an edge cover, i.e., if J' is a perfect matching, then the proof is completed by taking $J = J' \cup J''$. Clearly, the degree requirements in the proposition hold. Otherwise, if $J' \cup J''$ is not an edge cover, the size of $J' \cup J''$ is increased by one such that one more Y-node is covered and the degree requirements in the proposition are maintained. Let $v \in Y$ be a node that is *not* covered by $J' \cup J''$. Let T be the node set of the maximal J'-alternating tree that contains v. That is, a node w is in T iff there exists a J'-alternating path between v and w. (For a matching J', recall that a J'-alternating path means a path whose edges are alternately in J' and not in J'.)

Claim: There is a node $x \in T \cap X$ with $\deg(x) \geq k + 1$.

To prove this claim, note that (i) $|T \cap Y| = |T \cap X| + 1$ (since each node $y \in T \cap Y$ except v is incident to an edge of J'), and (ii) for every node $y \in T \cap Y$, every incident edge xy has the other end node x in $T \cap X$ (otherwise, x can be added to T, and so T is not maximal). By assumption, each node $y \in T \cap Y$ has $\deg(y) \geq k$, hence, (i), (ii), and the pigeon-hole principle guarantee that there is a node $x \in T \cap X$ with $\deg(x) > k$. This proves the claim.

Let xz be the J'-edge incident to x, i.e., x is matched to z by J'. This edge is (permanently) added to the edge cover J by taking $J'' = J'' \cup \{xz\}$. The node z is deleted from B. Since $x \in T$, there exists a J'-alternating path between v and x (by definition of T). Let this path be P'. The matching J' is updated by switching alternate edges along P', i.e., J' is replaced by the symmetric difference of J' and $E(P')$. Note that the current B (with node z deleted) has a matching of size

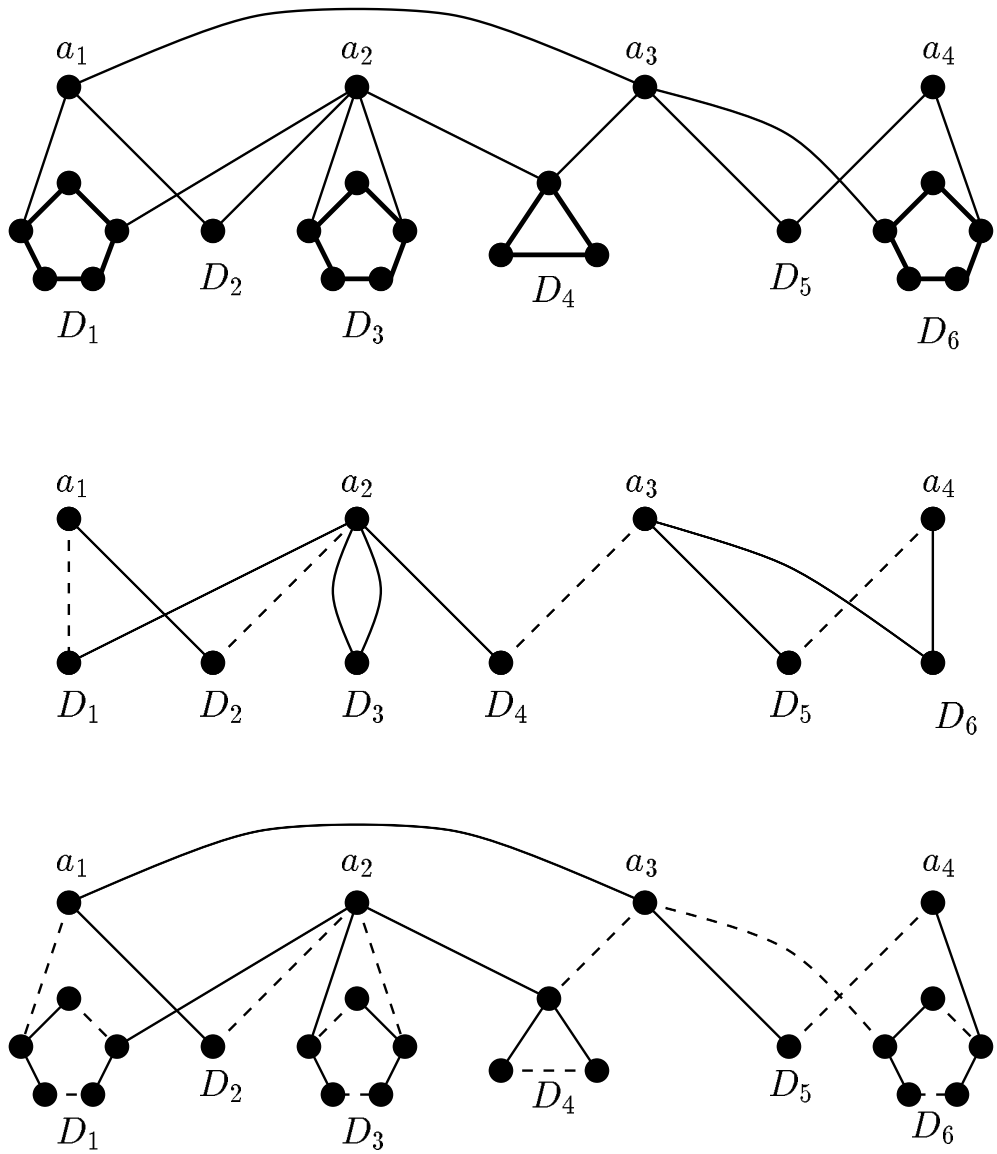

Figure 3: An illustration of the proofs of Theorem 3.5 and Propositions 3.9, 3.10.
(a) $G = (V, E)$ is a 2-edge connected graph ($k = 2$), and the Gallai-Edmonds decomposition is given by $A = A(G) = \{a_1, a_2, a_3, a_4\}$, and $D = D(G) = V(D_1) \cup V(D_2) \cup V(D_3) \cup V(D_4) \cup V(D_5) \cup V(D_6)$. The odd (factor-critical) components of $G \backslash A$ are $D_1, \ldots, D_6$.
(b) The bipartite multigraph B in the proofs of Propositions 3.9, 3.10. In Proposition 3.10, B is obtained from G by deleting the nodes in $V \backslash (A \cup D)$ and the edges in $E(A)$, and shrinking $D_1, \ldots, D_6$ into single nodes. In B, note that $\deg(D_1), \ldots, \deg(D_6) \geq k = 2$, and there is a matching J' of size $|A| = 4$. J' is indicated by dashed lines, $J' = \{a_1 D_1, a_2 D_2, a_3 D_4, a_4 D_5\}$.
In the construction of Proposition 3.9, the 1st iteration chooses, say $v = D_3$. Then $T = \{D_3, a_2, D_2, a_1, D_1\}$, and $x = a_2 \in T \cap A$ has degree $\geq k+1 = 3$. The edge $a_2 D_2$ is added to J'', the node D_2 is deleted, and in J', $a_2 D_2$ is replaced by $a_2 D_3$. Finally, $J' = \{a_1 D_1, a_2 D_3, a_3 D_6, a_4 D_5\}$, $J'' = \{a_2 D_2, a_3 D_4\}$, and $J = J' \cup J''$ is the required edge cover.
(c) In G, J maps to an edge set $\widetilde{J}$. $\widetilde{J}$ is extended to the required edge cover P of G by adding a perfect matching on the nodes of G *not* incident to $\widetilde{J}$. P is indicated by dashed lines.

$|X|$, namely J', and has $\deg(y) \geq k$, for all nodes $y \in V(B)\backslash X$. Therefore, the hypothesis of the proposition continues to hold.

The above step is repeated till $J' \cup J''$ covers all nodes of B. Finally, J is taken to be $J' \cup J''$. The construction guarantees that J satisfies the degree requirements in the proposition. $\square$

Recall the Gallai-Edmonds decomposition theorem of matching theory, [LP 86, Theorem 3.2.1]. For every graph H, there is a partition of $V(H)$ into a set of (matching) noncritical nodes $D(H)$ and a set of (matching) critical nodes $V\backslash D(H)$ (i.e., $D(H)$ consists of all nodes that are left uncovered by some maximum matching of H). The partition is "trivial" if either H has a perfect matching, or if H is factor-critical: in the first case, $D(H) = \emptyset$, and in the second case, $D(H) = V(H)$. Let $A(H)$ be the set of critical nodes of H that are adjacent to one or more noncritical nodes of H. Possibly, $A(H)$ is the empty set. When there is no danger of confusion, we use A and D instead of $A(H)$ and $D(H)$. Let $\mathrm{def}(H)$ denote the deficiency of H, i.e., the number of nodes that are not covered by a maximum matching of H. (So, $\mathrm{def}(H) = |V(H)| - 2|P_0|$, where P_0 is a maximum matching of H.) The Gallai-Edmonds decomposition theorem shows that in the graph $H\backslash A$, the noncritical nodes D form $q = |A(H)| + \mathrm{def}(H)$ odd components $D_1, D_2, \ldots, D_q$, i.e., each D_i $(i = 1, \ldots, q)$ is a connected component of $H\backslash A$ with $V(D_i) \subseteq D(H)$ and $|V(D_i)|$ odd. Moreover, every one of these odd components D_i is factor-critical.

The next result is a generalization of Proposition 3.9.

Proposition 3.10 *Let G be a graph, and let $D = D(G)$ and $A = A(G)$ be the node sets in the Gallai-Edmonds decomposition. Let $q = |A(G)| + \mathrm{def}(G)$, and let $D_1, D_2, \ldots, D_q$ be the odd components of $G\backslash A$. If every D_i gives a cut containing at least k edges, i.e., if $\delta(V(D_i))$ has size $\geq k$ for $i = 1, \ldots, q$, then G has an edge cover P such that each node in $V(G)\backslash A$ is incident to exactly one edge of P, and each node in A is incident to either exactly one edge of P or at least $(k-1)$ edges of $E(G)\backslash P$.*

Proof: See Figure 3 for an illustration. The proof follows easily by applying Proposition 3.9 to a bipartite graph associated with the Gallai-Edmonds decomposition.

If $\mathrm{def}(G) = 0$, then the proof is done: take P to be a perfect matching of G. Otherwise, $\mathrm{def}(G) > 0$, and so $D \neq \emptyset$. Suppose that $A = \emptyset$. Then every component D_i of G is factor-critical, but this violates the condition on $|\delta(V(D_i))|$. Hence, A is nonempty. Clearly, every edge in $\delta(V(D_i))$ $(i = 1, \ldots, q)$ has one end node in A and the other in D_i. Let $G[A \cup D]$ be the subgraph of G induced by $A \cup D$. Let $B = (X \cup Y, E')$, $X = A$, be the bipartite (loopless) multigraph obtained from $G[A \cup D]$ by deleting all edges with both end nodes in A and by shrinking the components $D_1, D_2, \ldots, D_q$ of $G[A \cup D]\backslash A$ to single nodes. The shrunk nodes are also called $D_1, D_2, \ldots, D_q$, and so $Y = \{D_1, D_2, \ldots, D_q\}$. B has $\geq k$ edges incident to each of the shrunk nodes $D_1, D_2, \ldots, D_q$, since in G each of the cuts $\delta(V(D_i))$ $(i = 1, \ldots, q)$ has $\geq k$ edges. Moreover, B has a matching of size $|X| = |A|$, by the Gallai-Edmonds decomposition theorem. Therefore, B satisfies the conditions in Proposition 3.9. By the proposition, B has an edge cover J satisfying the degree requirements in the proposition; note that each node $D_i \in Y$ is incident to exactly one edge of J. Let $\widetilde{J}$ denote a set edges of G that corresponds to J, i.e., for each edge $a_h D_i \in J$ with $a_h \in X = A$, $D_i \in Y$, there is an edge $a_h w_i \in \widetilde{J}$ such that (in G) w_i is a node in D_i and w_i is adjacent to a_h. Let $V(\widetilde{J})$ be the set of nodes of G incident to edges in $\widetilde{J}$, i.e., $V(\widetilde{J}) = A \cup \{w_i \in V(D_i) : i = 1, \ldots, q\}$. By the Gallai-Edmonds decomposition theorem, $G\backslash V(\widetilde{J})$ has a perfect matching $\widetilde{P}$. To see this, note that each component of $G\backslash V(\widetilde{J})$ is either an even component of $G\backslash A$ or is obtained by deleting one node from an odd (factor-critical) component of $G\backslash A$; in either case, the component has a perfect matching.

Take $P = \widetilde{J} \cup \widetilde{P}$. Clearly, P is an edge cover of G such that each node $v \in V \backslash A$ is incident to exactly one edge of P. Moreover, by Proposition 3.9, every node in A is incident to either exactly one edge of P or to $\geq (k-1)$ edges of $E \backslash P$. $\square$

Proof: **(Theorem 3.5)** See Figure 3 for an illustration. We construct an appropriate edge set P^* such that $|P^*| \geq \lfloor n/2 \rfloor$ and every node $v \in V$ is incident to $\geq (k-1)$ edges of $E^* \backslash P^*$. In the statement of Theorem 3.5, note that M^* is a minimum-size edge set such that (V, M^*) has minimum degree $(k-1)$. Hence, $|E^* \backslash P^*| \geq |M^*|$. The theorem follows immediately from the existence of the edge set P^*, because $|E^*| = |E^* \backslash P^*| + |P^*| \geq |E^* \backslash P^*| + \lfloor n/2 \rfloor \geq |M^*| + \lfloor n/2 \rfloor$.

If the size of a maximum matching of G^* is $\geq (n-1)/2$, i.e., if G^* has a matching that leaves at most one node uncovered, then we take P^* to be a maximum matching. (This handles the case when G^* is a factor-critical graph.)

To handle the case when $\operatorname{def}(G^*) \geq 2$, we apply Proposition 3.10 to G^*, noting that G^* satisfies the conditions in the proposition. (Since G^* is k-edge connected, $\deg(v) \geq k$, $\forall v \in V$, and every node set $S \subseteq V$, $\emptyset \neq S \neq V$, has $|\delta(S)| \geq k$.) We take P^* to be the edge cover P guaranteed by the proposition. Since P^* is an edge cover of G^*, $|P^*| \geq n/2$. Moreover, $(V, E^* \backslash P^*)$ has minimum degree $\geq k-1$ by the proposition and the fact that G^* has minimum degree $\geq k$. The theorem follows. $\square$

We mention two corollaries of Theorem 3.5/Proposition 3.10, though these are not relevant to the main theme of the paper.

Corollary 3.11 (Petersen's Theorem) *A 3-regular graph without cut edges has a perfect matching.*

Proof: Let $G^* = (V, E^*)$ be the graph, and let $n = |V|$. Clearly, n is even, and $|E^*| = 3n/2$. The key point is that every node set S of odd cardinality (i.e., $S \subset V$ and $|S|$ odd) has $|\delta(S)| \geq 3$ since $|\delta(S)|$ is *odd* (since $3|S| - 2|E(S)|$ is odd) and is ≥ 2. Suppose that G^* has no perfect matching. Then $\operatorname{def}(G^*) > 0$, and so in the Gallai-Edmonds decomposition we have $D(G^*) \neq \emptyset$; moreover, G^* is not factor-critical (n is even) so $A(G^*) \neq \emptyset$. Applying Proposition 3.10 with $k = 3$ shows that G^* has an edge cover P such that every node is incident to $\geq (k-1) = 2$ edges of $M = E^* \backslash P$. Clearly, $|P| \geq n/2$, since P is an edge cover, and $|M| = |E^* \backslash P| \geq n$, since (V, M) has minimum degree 2. Since $|E^*| = |P| + |M| = 3n/2$, we have $|P| = n/2$ and $|M| = n$. Therefore, P is a perfect matching of G^*. $\square$

Corollary 3.12 *Let $G = (V, E)$ be a 2-edge connected graph. G has two edge-disjoint edge covers iff G is not a cycle of odd length.*

Proof: If G is an odd-length cycle, then it does not have two edge-disjoint edge covers.

Suppose that G is not a cycle of odd length. If G has a perfect matching P, then clearly P and $E \backslash P$ are edge-disjoint edge covers of G. Suppose that G is factor-critical and has a node v with $\deg(v) \geq 3$. Let w be a neighbour of v. Now $G \backslash w$ has a perfect matching, say P_0. Then $P = P_0 \cup \{vw\}$ is an edge cover of G such that $(V, E \backslash P)$ has an edge cover. Otherwise, G is not factor-critical and has no perfect matching. Then Proposition 3.10 gives an edge cover P such that $E \backslash P$ is an edge cover. $\square$

3.3 Minimum-size 2-connected spanning subgraphs of undirected graphs: a parallel $(1.5+\epsilon)$-approximation algorithm

This subsection focuses on the design of an efficient parallel algorithm and a linear-time sequential algorithm for the problem of finding a minimum-size 2-node connected (2-edge connected) spanning subgraph of a graph. Let $\epsilon > 0$ be a constant, independent of $|V(G)|$. A deterministic parallel version of the main heuristic runs in **NC** and achieves an approximation guarantee of $(1.5+\epsilon)$, whereas a randomized **NC** version achieves an approximation guarantee of 1.5. A sequential *linear-time* version of the main heuristic achieves an approximation guarantee of $(1.5+\epsilon)$. The proof of the 1.5 approximation guarantee in this subsection again hinges on Mader's theorem (Theorem 3.2), but instead of employing the lower bound in Theorem 3.5, we employ a nice lower bound result due to Chong and Lam (Proposition 3.14).

The heuristic for a minimum-size 2-NCSS described below can be used to find a 1.5-approximation of a minimum-size 2-ECSS. For this, we run a preprocessing step on the given graph $G = (V, E)$, which is assumed to be 2-edge connected, to partition the edge set into blocks (maximal 2-node connected subgraphs). Then separately for each block, we run our heuristic for a minimum-size 2-NCSS. For a block, the optimal 2-ECSS may *not* be 2-node connected, nevertheless, the lower bound used by the 2-NCSS heuristic applies to 2-ECSS too, so the edge set found by our algorithm will have size within 1.5 times the minimum size of a 2-ECSS.

Consider the problem of approximating a minimum-size 2-NCSS. Assume that the given graph $G = (V, E)$ is 2-node connected. The heuristic consists of two steps. The first finds a *minimum edge cover* $M \subseteq E$ of G, i.e., a minimum-cardinality edge set such that every node is incident to at least one edge of M. One way of finding M is to start with a *maximum matching* $\widetilde{M}$ of G, and then to add one edge incident to each node that is not matched by $\widetilde{M}$. Recall that $\text{def}(G)$ denotes the number of nodes not matched by a maximum matching of G, i.e., $\text{def}(G) = |V| - 2|\widetilde{M}|$. Then we have $|M| = |\widetilde{M}| + \text{def}(G)$. (It is easily seen that no edge cover of G has smaller cardinality than $|\widetilde{M}| + \text{def}(G)$.) The second step of the heuristic finds an (inclusionwise) minimal edge set $F \subseteq E \backslash M$ such that $M \cup F$ gives a 2-NCSS. In other words, $(V, M \cup F)$ is 2-node connected, but for each edge $vw \in F$, $(V, M \cup F) \backslash vw$ is *not* 2-node connected. Let E' denote $M \cup F$, and let $E_{opt} \subseteq E$ denote a minimum-cardinality edge set such that (V, E_{opt}) is 2-*edge* connected.

Lemma 3.13 $|E'| = |M| + |F| \leq 1.5|V| + \text{def}(G) - 1$.

Proof: By Mader's theorem (Theorem 3.2), F is acyclic, so $|F| \leq |V| - 1$. A minimum edge cover M of G has size $|M| = |\widetilde{M}| + \text{def}(G)$, where $\widetilde{M}$ is a maximum matching of G. Obviously, $|\widetilde{M}| \leq |V|/2$. The result follows. □

The next result, due to Chong and Lam, gives a lower bound on the size of a 2-ECSS. Proposition 3.15 generalizes Chong and Lam's lower bound to k-edge connected spanning subgraphs, $k \geq 1$.

Proposition 3.14 (Chong & Lam [CL 95, Lemma 3]) *Let $G = (V, E)$ be a graph of edge connectivity ≥ 2, and let $|E_{opt}|$ denote the minimum size of a 2-edge connected spanning subgraph. Then $|E_{opt}| \geq \max(|V| + \text{def}(G) - 1, |V|)$.*

Proposition 3.15 *Let $G = (V, E)$ be a graph of edge connectivity $\geq k \geq 1$, and let $|E_{opt}|$ denote the minimum size of a k-edge connected spanning subgraph. If G is not factor-critical, then $|E_{opt}| \geq \frac{k}{2}(|V| + \text{def}(G))$. In general, $|E_{opt}| \geq \frac{k}{2}\max(|V| + \text{def}(G) - 1,\ |V|)$.*

Proof: Suppose that G is *not* factor-critical and $\mathrm{def}(G)$ is ≥ 1. Then, by the Gallai-Edmonds decomposition theorem of matching theory [LP 86, Theorem 3.2.1], there is a *nonempty* node set A such that $G\backslash A$ has $|A| + \mathrm{def}(G)$ odd components ($G\backslash A$ may have some even components too). Focus on an (odd or even) component D_i of $G\backslash A$. The number of edges of E_{opt} such that either one or both end nodes are in D_i is at least $(|V(D_i)| + 1)k/2$, because every node $v \in V(D_i)$ is incident to $\geq k$ edges of E_{opt}, and moreover, $\delta(V(D_i))$ has at least k edges of E_{opt}. Summing over all components D_i of $G\backslash A$ proves the proposition. □

Theorem 3.16 *Let $G = (V, E)$ be a graph of node (edge) connectivity ≥ 2. Let $\epsilon > 0$ be a constant. The heuristic described above finds a 2-node connected (2-edge connected) spanning subgraph (V, E') such that $|E'| \leq 1.5|E_{opt}|$, where $|E_{opt}|$ denotes the minimum size of a 2-ECSS.*

A randomized parallel version of the heuristic runs in **RNC** *and achieves an approximation guarantee of* 1.5. *A deterministic parallel version of the heuristic runs in* **NC** *and achieves an approximation guarantee of* $(1.5 + \epsilon)$.

The sequential running time is $O(\sqrt{|V|}|E|)$. A sequential linear-time version of the heuristic achieves an approximation guarantee of $(1.5 + \epsilon)$.

Proof: The approximation guarantee follows from Lemma 3.13 and Proposition 3.14, since

$$\frac{|E'|}{|E_{opt}|} \leq \frac{1.5|V| + \mathrm{def}(G) - 1}{\max(|V| + \mathrm{def}(G) - 1, |V|)} \leq 1 + \frac{0.5|V|}{|V|} \leq 1.5.$$

Consider the deterministic parallel version of the heuristic. Let $\widetilde{M}$ denote a maximum matching of G. For Step 1, we find an approximately maximum matching in **NC** using the algorithm of [FGHP 93]: for a constant ϵ, $0 < \epsilon < 0.5$, the algorithm finds a matching M' with $|M'| \geq (1-2\epsilon)|\widetilde{M}|$ in parallel time $O(\epsilon^{-4}(\log |V|)^3)$ using $O(\epsilon^{-1}|V|^{2+(2/\epsilon)})$ processors. We obtain an (inclusionwise) minimal edge cover M of size $\leq (1 + 2\epsilon)|\widetilde{M}| + \mathrm{def}(G)$ by adding to M' one edge incident to every node that is not matched by M'. For Step 2, we use a variant of the **NC** algorithm of [HKe+ 95, KeR 95], see Algorithm 2 and Lemma 2 in Kelsen & Ramachandran [KeR 95]. Let G' be a 2-node connected spanning subgraph of G such that $E(G')$ contains the minimal edge cover M. Call an edge vw of G' *essential* if either vw is in M or $G'\backslash vw$ is not 2-node connected (i.e., an edge of G' is nonessential if it is not in M and it is not critical w.r.t. the 2-node connectivity of G'). Algorithm 2 of [KeR 95] starts by taking the current subgraph G' to be G, and repeatedly finds a spanning tree T of G' that has the minimum number of nonessential edges, *minimally* augments T to obtain a 2-node connected spanning subgraph G'' of G', and then replaces the current subgraph G' by G''. Finding the spanning tree T is easy: we compute a minimum spanning tree of G' where the cost of each edge in M is taken to be (-1), the cost of each remaining essential edge of G' is zero, and the cost of each nonessential edge of G' is one. The parallel complexity of the whole algorithm is in **NC**, see [HKe+ 95, KeR 95]. Now, the approximation guarantee is $(1.5 + \epsilon)$.

For the sequential linear-time version of the heuristic, note that a matching M' with $|M'| \geq (1 - 2\epsilon)|\widetilde{M}|$ can be found in time $O((|V| + |E|)/\epsilon)$. Moreover, in linear time, we can find a minimal 2-node connected spanning subgraph whose edge set contains the minimal edge cover $M \subseteq E$ obtained by adding edges to M', see [HKe+ 95]. □

3.4 Directed graphs

The main heuristic extends to digraphs. The key tool in the analysis of the approximation guarantee is another theorem of Mader, [Ma 85, Theorem 1]. Given a digraph $G = (V, E)$ that is assumed to have node connectivity at least k, the first step of the heuristic finds an arc set $M \subseteq E$ of minimum cardinality such that for every node v, there are $\geq (k-1)$ arcs of M going out of v and $\geq (k-1)$ arcs of M coming into v. Clearly, $|M| \leq |E_{opt}|$, where $E_{opt} \subseteq E$ denotes a minimum-cardinality arc set such that (V, E_{opt}) is k-node connected. The second step of the heuristic is as in Section 3.1: we find an (inclusionwise) minimal arc set $F \subseteq E \backslash M$ such that $M \cup F$ is the arc set of a k-node connected spanning subgraph. The key point is that $|F| \leq 2|V|-1$, by Mader's digraph theorem (Theorem 3.17).

Consider the first step in more detail. To find the arc set M, we transform the digraph problem to a b-matching problem on the bipartite graph $B(G)$ associated with G. For each node $v \in V(G)$, there is a pair of nodes v_-, v_+ in the bipartite graph $B(G)$, and for each arc (v, w) of G, there is one edge v_+w_- in the bipartite graph. Our problem of finding a minimum-cardinality $M \subseteq E$ with $\deg_{M,in}(v) \geq (k-1)$, $\deg_{M,out}(v) \geq (k-1)$, $\forall v \in V$, corresponds to the problem of finding a minimum-cardinality edge set M' of the bipartite graph such that each node of the bipartite graph is incident to $\geq (k-1)$ edges of M'. As in Section 3.1, this is a b-matching problem.

An *alternating cycle* of a digraph is a nonempty, even-length sequence of distinct arcs $C = e_1, e_2, \ldots, e_{2\ell-1}, e_{2\ell}$, $\ell \geq 1$, such that (using indices modulo 2ℓ) for each $i = 0, 1, \ldots$, the arcs e_{2i} and e_{2i+1} have the same start node, and the arcs e_{2i+1} and e_{2i+2} have the same end node. In other words, the set of undirected edges corresponding to an alternating cycle C is a union of cycles, and moreover, alternate occurrences of nodes have two C-arcs coming out or two C-arcs going in. See Figure 4 for an illustration. For an alternating cycle C, a C-out node is a node having two outgoing arcs of C, and a C-in node is a node having two incoming arcs of C. Recall that an arc e of a k-node connected digraph H is called critical if $H \backslash e$ is not k-node connected. Here is Mader's theorem on the critical arcs of a k-node connected digraph; see Figure 4 for an illustration.

Theorem 3.17 (Mader [Ma 85, Theorem 1]) *In a k-node connected digraph, if there is an alternating cycle C each of whose arcs is critical, then there is either a C-out node of outdegree k or a C-in node of indegree k.*

Fact 3.18 (Mader [Ma 85, Lemma 2]) *Let H be a digraph, and let $B(H)$ be the associated bipartite graph. There is a cycle in $B(H)$ iff there is an alternating cycle in H.*

Remarks: Mader [Ma 85] states the theorem for minimal k-node connected digraphs, but in fact, his proof needs only the fact that every arc in the alternating cycle is critical. Now, consider a digraph H_0 that is obtained from an arbitrary strongly connected digraph by subdividing every arc at least once (i.e., an arc is replaced by ≥ 1 new nodes and a directed path of ≥ 2 arcs). Note that H_0 contains no alternating cycle. Mader [Ma 85, p. 104] shows that there exists a minimal k-node connected digraph G such that H_0 is contained in the subgraph of G induced by arcs whose start nodes have outdegrees $> k$ and whose end nodes have indegrees $> k$.

Lemma 3.19 *Let $F \subseteq E \backslash M$ be the set of critical arcs found by the second step of the heuristic. Then $|F| \leq 2|V|-1$.*

Proof: Let $G' = (V, E')$, where $E' = M \cup F$. We claim that F contains no alternating cycle. By way of contradiction, suppose that $C \subseteq F$ is an alternating cycle. Observe that every C-out node v has $\geq (k+1)$ outgoing arcs of E', since there are $\geq (k-1)$ arcs of M outgoing from v, and there are two arcs of C outgoing from v. Similarly, every C-in node has $\geq (k+1)$ incoming arcs of E'. This contradicts Mader's digraph theorem. Hence, F contains no alternating cycle. Then $|F| \leq 2|V|-1$, because the bipartite graph associated with (V, F) is acyclic. $\square$

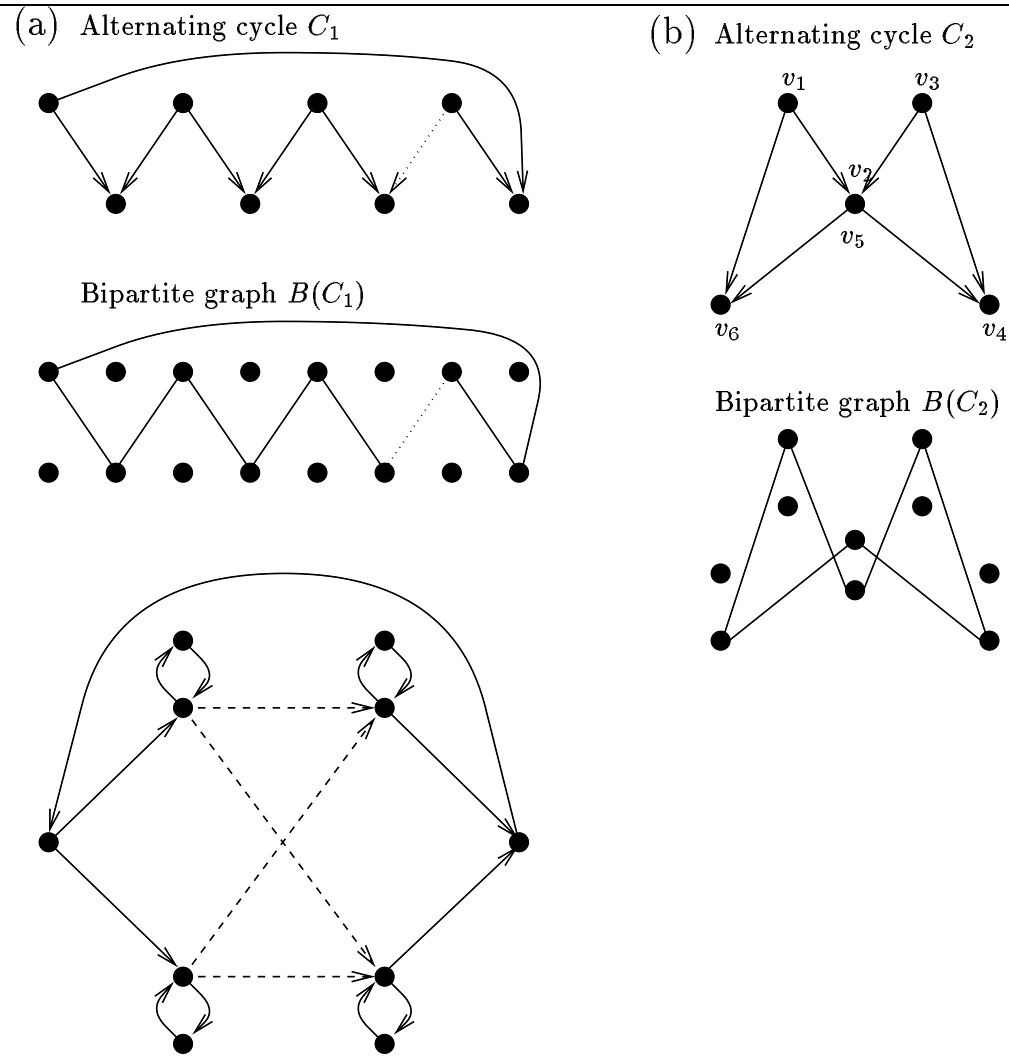

(c) An alternating cycle in a strongly connected digraph

Figure 4: An illustration of an alternating cycle in a digraph, and of Mader's theorem on critical alternating cycles in a k-node connected digraph, see Theorem 3.17.
(a) An alternating cycle C_1, and its bipartite graph $B(C_1)$.
(b) Another alternating cycle $C_2 = (v_1, v_2), (v_3, v_2), (v_3, v_4), (v_5, v_4), (v_5, v_6), (v_1, v_6)$ and its bipartite graph $B(C_2)$. For an alternating cycle, the undirected version may not be a cycle, but the bipartite graph has at least one cycle.
(c) An alternating cycle C of a 1-connected (strongly connected) digraph is indicated by dashed lines. Every C-out node has outdegree $> k = 1$, and every C-in node has indegree $> k = 1$. None of the arcs in the alternating cycle is critical for 1-connectivity. This example is modified from an example of Mader [Ma 85].

The previous lemma immediately gives an approximation guarantee of $1+[2/k]$ for a minimum-size k-NCSS of a digraph, because the "degree lower bound" implies that a digraph k-NCSS has $\geq k|V|$ arcs. The approximation guarantee can be improved to $1+[1/k]$ via the lower bound on the size of a digraph k-NCSS implied by Proposition 3.8.

Proposition 3.20 *Let $G=(V,E)$ be a digraph of node connectivity $\geq k$. The heuristic above finds a k-node connected spanning subgraph (V,E') such that $|E'| \leq (1+[2/k])|E_{opt}|$, where $|E_{opt}|$ denotes the cardinality of an optimal solution.*

Theorem 3.21 *Let $G=(V,E)$ be a digraph of node connectivity $\geq k$. The heuristic described above finds a k-node connected spanning subgraph (V,E') such that $|E'| \leq (1+[1/k])|E_{opt}|$, where $E_{opt} \subseteq E$ denotes a minimum-cardinality arc set such that (V,E_{opt}) is k-node connected. The running time is $O(k|E|^2)$.*

Proof: The proof of the approximation guarantee is similar to the proof for undirected graphs in Theorem 3.6. Let $G_{opt}=(V,E_{opt})$ be a k-node connected spanning subgraph of minimum size. Apply Proposition 3.8 to the bipartite graph $B(G_{opt})$ of G_{opt} to deduce that $|M^*| \leq |E(B(G_{opt}))| - |V(B(G_{opt}))|/2$, where $M^* \subseteq E(B(G_{opt}))$ is a minimum-size edge set such that every node of $B(G_{opt})$ is incident to $\geq k-1$ edges of M^*. Since the arc set $M \subseteq E(G)$ found by the heuristic has $|M| \leq |M^*|$ (since M comes from a supergraph of E_{opt}), it follows that $|M| \leq |E(B(G_{opt}))| - |V(B(G_{opt}))|/2 = |E_{opt}| - |V(G)|$. Consequently, since $|E'| = |M| + |F|$ and $|F| \leq 2|V(G)|-1$,

$$\frac{|E'|}{|E_{opt}|} \leq \frac{|E_{opt}| - |V(G)| + (2|V(G)|-1)}{|E_{opt}|} \leq 1+\frac{1}{k},$$

where the last inequality uses the "degree lower bound", $|E_{opt}| \geq k|V(G)|$. The running time analysis is similar to that for the heuristic for graphs, see Section 3.1. □

4 Approximating minimum-size k-edge connected spanning subgraphs

The heuristic can be modified to find an approximately minimum-size k-edge connected spanning subgraph (abbreviated k-ECSS) of a graph or a digraph. First, we focus on graphs, and prove a $(1+[2/(k+1)])$-approximation guarantee for finding a minimum-size k-ECSS. The analysis hinges on Theorem 4.3 which may be regarded as an analogue of Mader's theorem [Ma 72, Theorem 1] for k-edge connected graphs. Then we turn to digraphs, and prove an approximation guarantee of $1+[4/\sqrt{k}]$ for the k-ECSS heuristic.

In this section, an edge e (arc e) of a k-edge connected graph (digraph) H is called *critical* if $H \backslash e$ is *not* k-edge connected. Assume that the given graph or digraph $G=(V,E)$ is k-edge connected, otherwise, the heuristic will detect this and report failure.

4.1 Undirected graphs

In this subsection, $G=(V,E)$ is a graph. The first step of the heuristic finds an edge set $M \subseteq E$ of minimum cardinality such that every node in V is incident to $\geq k$ edges of M. Clearly, $|M| \leq |E_{opt}|$, where $E_{opt} \subseteq E$ denotes a minimum-cardinality edge set such that (V,E_{opt}) is k-edge connected. The second step of the heuristic finds an (inclusionwise) minimal edge set $F \subseteq E \backslash M$ such that $M \cup F$ is the edge set of a k-edge connected spanning subgraph. In detail, the second step starts with $F = \emptyset$ and $E' = E$. Note that $G' = (V, E')$ is k-edge connected at the start. We examine the edges of $E \backslash M$ in an arbitrary order $e_1, e_2, \ldots$. For each edge $e_i = v_i w_i$ (where $1 \leq i \leq |E \backslash M|$), we determine whether or not $v_i w_i$ is critical for the current graph by finding the maximum number of edge-disjoint $v_i \leftrightarrow w_i$ paths in G'.

Proposition 4.1 *An edge $v_i w_i$ of a k-edge connected graph is not critical iff there exist at least $k+1$ edge-disjoint $v_i \leftrightarrow w_i$ paths (including the path $v_i w_i$).*

If $v_i w_i$ is noncritical, then we delete it from E' and G', otherwise, we retain it in E' and G', and also, we add it to F. At termination of the heuristic, $G' = (V, E')$, $E' = M \cup F$, is k-edge connected and every edge $vw \in F$ is critical, i.e., $G' \backslash vw$ is *not* k-edge connected. Theorem 4.3 below shows that $|F| \leq k|V|/(k+1)$ for $k \geq 1$. Since $|E_{opt}| \geq k|V|/2$, the minimum-size k-ECSS heuristic achieves an approximation guarantee of $1 + [2/(k+1)]$ for $k \geq 1$.

The next lemma turns out to be quite useful. A straightforward counting argument gives the proof, see Mader [Ma 71, Lemma 1], or Cai [Ca 93, Claim 3].

Lemma 4.2 *Let $G = (V, M)$ be a simple graph of minimum degree $k \geq 1$.*
(i) Then for every node set $S \subseteq V$ with $1 \leq |S| \leq k$, the number of edges with exactly one end node in S, $|\delta(S)|$, is at least k.
(ii) If a node set $S \subseteq V$ with $1 \leq |S| \leq k$ contains at least one node of degree $\geq (k+1)$, then $|\delta(S)|$ is at least $k+1$.

The goal of Theorem 4.3 is to give an upper bound on the number of critical edges in the edge-complement of a spanning subgraph of minimum degree k in an arbitrary k-edge connected graph H. Clearly, every critical edge $e \in E(H)$ is in some k-cut $\delta(A_e)$, $A_e \subseteq V(H)$. By a *tight* node set S of a k-edge connected graph H we mean a set $S \subset V(H)$ with $|\delta_H(S)| = k$, i.e., a node set S such that $\delta_H(S)$ is a k-cut. As usual, a family of sets $\{S_i\}$ is called *laminar* if for any two sets in the family, either the two sets are disjoint, or one set is contained in the other. For an arbitrary subset F' of the critical edges of H, it is well known that there exists a laminar family $\mathcal{F}$ of tight node sets *covering* F', i.e., there exists $\mathcal{F} = \{A_1, A_2, \ldots, A_\ell\}$, where $A_i \subseteq V(H)$ and $\delta(A_i)$ is a k-cut, for $1 \leq i \leq \ell$, such that each edge $e \in F'$ is in some $\delta(A_i)$, $1 \leq i \leq \ell$. (For details, see [Fr 93, Section 5] or [Ca 93, Lemma 3], but in the latter reference note that the associated family (of a collection of k-cuts) should be laminar rather than crossing-free.) It is convenient to define a tree T corresponding to $\mathcal{F} \cup \{V(H)\}$: there is a T-node corresponding to each set $A_i \in \mathcal{F}$ and to $V(H)$, and there is a T-edge $A_i A_j$ (or $V(H) A_j$) iff $A_j \subset A_i$ and no other node set in $\mathcal{F}$ contains A_j and is contained in A_i. Note that the T-node corresponding to the node set A_i of the laminar family $\mathcal{F}$ is denoted by A_i, and the T-node corresponding to the node set $V(H)$ is denoted by $V(H)$. Each T-edge corresponds to a k-cut of H. Suppose that the tree T is rooted at the T-node $V(H)$. We associate another node set $\phi_i \subseteq V(H)$ with each node set A_i of $\mathcal{F}$:

$$\phi_i = A_i \backslash \bigcup \{A \in \mathcal{F} : A \subset A_i, A \neq A_i\}.$$

In other words, a T-node $A_i \in \mathcal{F}$ that is a leaf node of T has $\phi_i = A_i$, otherwise, ϕ_i consists of those H-nodes of A_i that are not in the node sets $A', A'', \ldots$, where $A', A'', \ldots \in \mathcal{F}$ correspond to the children of A_i in the tree T. For distinct T-nodes A_i and A_j, note that ϕ_i and ϕ_j are disjoint. See Figure 5 for an illustration of $\mathcal{F} = \{A_i\}$, the family of node sets $\{\phi_i\}$, and the tree T for a particular graph.

The proof of Theorem 4.3 is long and nontrivial. Readers interested in a detailed study of the proof may be helped by: (i) an examination of the examples in Figure 2(c) and Figure 5, (ii) the

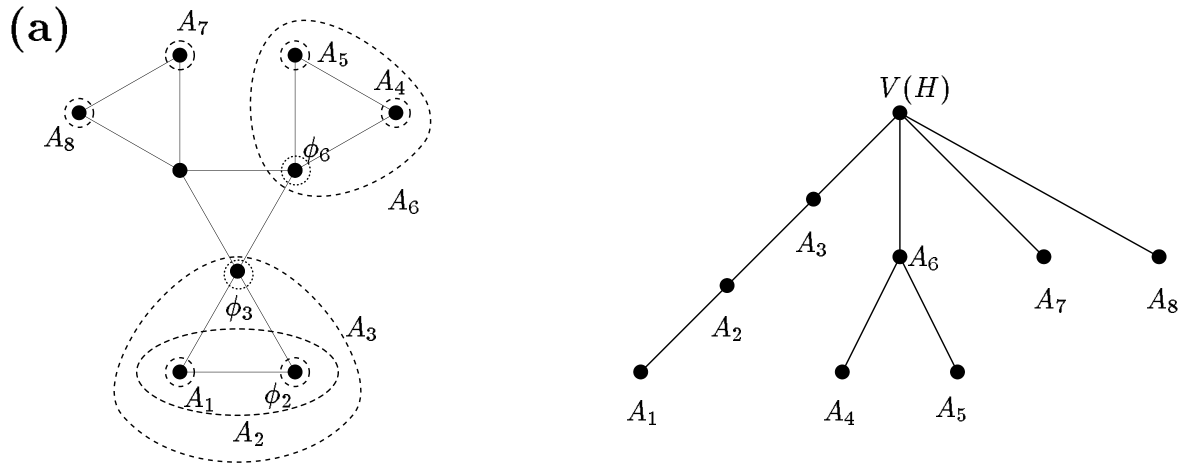

Laminar family $\mathcal{F}$ of tight node sets

Tree T of $\mathcal{F}$

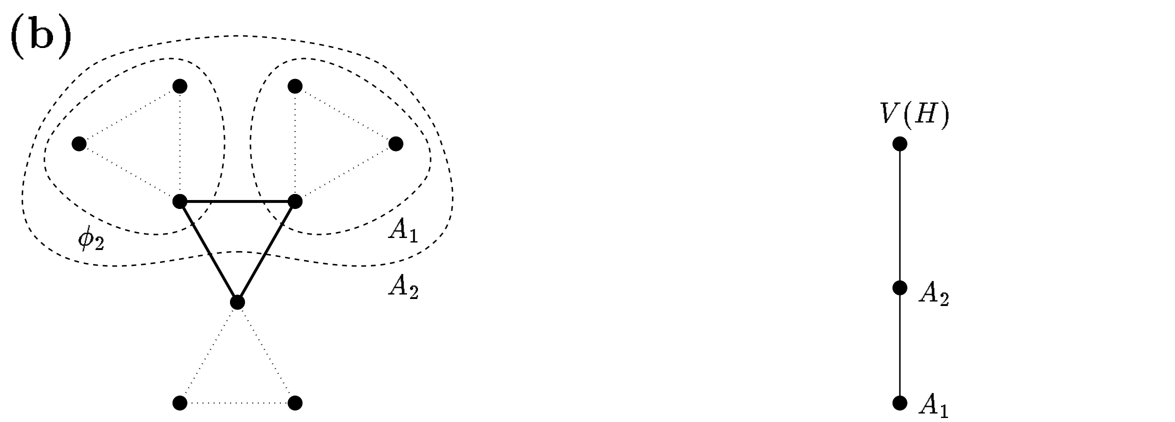

Laminar family $\mathcal{F}'$ of tight node sets

Tree T' of $\mathcal{F}'$

Figure 5: Two laminar families of tight node sets for a 2-edge connected graph H ($k = 2$).
(a) The laminar family $\mathcal{F}$ covers *all* critical edges of H. $\mathcal{F}$ consists of the node sets $A_1, \ldots, A_8$, where each A_i is tight since $|\delta(A_i)| = 2 = k$. For a node set A_i, ϕ_i is the node set $A_i \backslash \bigcup\{A_j \in \mathcal{F} : A_j \subset A_i,\ A_j \neq A_i\}$. Note that $\phi_i = A_i$ for the inclusionwise minimal A_i, i.e., for $i = 1, 4, 5, 7, 8$. Also, the tree T corresponding to $\mathcal{F} \cup \{V(H)\}$ is illustrated.
(b) The laminar family $\mathcal{F}'$ covers all critical edges of $E(H)\backslash M$, where $M \subset E(H)$ is such that every node is incident to at least $k = 2$ edges of M. M is indicated by dotted lines. All edges of $E(H)\backslash M$ are critical. $\mathcal{F}'$ consists of the tight node sets A_1, A_2. Also, the node sets ϕ_1, ϕ_2 are indicated ($\phi_1 = A_1$), and the tree T' representing $\mathcal{F}' \cup \{V(H)\}$ is illustrated.

illustration of the proof in Figure 6(a)–(d), (iii) a study of the proof of Theorem 4.9, which is an analogous but weaker result for k-edge connected digraphs, and (iv) a study of the relevant parts of the papers by A. Frank [Fr 93] and by M. Cai [Ca 93].

Theorem 4.3 *Let $H=(V,E)$ be a k-edge connected, n-node graph ($k \geq 1$). Let $M \subseteq E$ be an edge set such that the spanning subgraph (V,M) has minimum degree $\geq k$. Let F be the set consisting of edges of $E \backslash M$ that are in some k-cut of H. Let $\mathcal{F}=\{A_1,\ldots,A_\ell\}$ be a laminar family of tight node sets that covers F, i.e., for each $e \in F$, there is an $A_i \in \mathcal{F}$ such that $e \in \delta(A_i)$. Then*

$$|F| \leq \frac{k}{k+1}\left|\bigcup_{i=1}^{\ell} A_i\right| \leq \frac{k}{k+1}(n-1). \tag{1}$$

Some key preliminaries are discussed, before delving into the proof. The upper bound on $|F|$ is asymptotically tight. Consider the k-edge connected graph G obtained as follows: take $\ell+1$ copies of the $(k+1)$-clique, $C_0, C_1, \ldots, C_\ell$, and for each $i=1,\ldots,\ell$, choose an arbitrary node v_i in C_i and add k (nonparallel) edges between v_i and C_0. Take $M=\bigcup_{i=0}^{\ell} E(C_i)$, and $F=E(G)\backslash M$. Observe that $|F|=k(n-(k+1))/(k+1)$.

Fact 4.4 *For a laminar family of tight node sets $\mathcal{F}=\{A_1,\ldots,A_\ell\}$, $\bigcup_{i=1}^{\ell}\delta(A_i)=\bigcup_{i=1}^{\ell}\delta(\phi_i)$.*

Proof: For each $i=1,\ldots,\ell$, an edge in $\delta(\phi_i)$ is either in $\delta(A_i)$ or in $\delta(A'), \delta(A''),\ldots$, where $A', A'', \ldots \in \mathcal{F}$ correspond to the children of A_i in the tree T. Hence, the set on the left side contains the set on the right side.

To see that the set on the left side is contained in the set on the right side, note that for every edge e in the left side set, there is an (inclusionwise) minimal tight node set $A_{i(e)}$ such that $e \in \delta(A_{i(e)})$, and the associated node set $\phi_{i(e)}$ has $e \in \delta(\phi_{i(e)})$. □

Fact 4.5 *Let H, M, F and $\mathcal{F}=\{A_1,\ldots,A_\ell\}$ be as in Theorem 4.3. The inequality in the theorem*

$$|F| \leq \frac{k}{k+1}\left|\bigcup_{i=1}^{\ell} A_i\right|$$

is implied by the inequality

$$\left|\bigcup_{i=1}^{\ell}\delta(A_i)\right| \leq \frac{k}{k+1}\sum_{i=1}^{\ell}|\phi_i| + \frac{1}{2}\sum_{i=1}^{\ell}|M \cap \delta(\phi_i)|.$$

Proof: Let $M_c \subseteq M$ denote the set of M-edges that are covered by the laminar family $\mathcal{F}$, i.e.,

$$M_c = \bigcup_{i=1}^{\ell}[M \cap \delta(A_i)] = M \cap \left[\bigcup_{i=1}^{\ell}\delta(A_i)\right] = M \cap \left[\bigcup_{i=1}^{\ell}\delta(\phi_i)\right] = \bigcup_{i=1}^{\ell}[M \cap \delta(\phi_i)].$$

Consider an arbitrary edge $e=vw$ that is in M_c. If $e \in \delta(\phi_i)$ $(i=1,\ldots,\ell)$, then either $v \in \phi_i, w \notin \phi_i$ or $w \in \phi_i, v \notin \phi_i$. Since the node sets ϕ_i $(i=1,\ldots,\ell)$ are mutually disjoint, there are at most two tight node sets $A_i \in \mathcal{F}$ such that $e \in \delta(\phi_i)$. (E.g., if there are tight node sets $A_g, A_h \in \mathcal{F}$,

$g \neq h$, with $v \in \phi_g, w \in \phi_h$, then $e \in \delta(\phi_g), e \in \delta(\phi_h)$, and $e \notin \delta(\phi_i)$ for $i = 1, \ldots, \ell$, $i \neq g$, $i \neq h$.) Then

$$|M_c| = \left| \bigcup_{i=1}^{\ell} [M \cap \delta(\phi_i)] \right| \geq \frac{1}{2} \sum_{i=1}^{\ell} |M \cap \delta(\phi_i)|, \tag{2}$$

since we are counting the cardinality of a union of sets such that each element occurs in at most two of these sets.

Now note that $\bigcup_{i=1}^{\ell} \delta(A_i) = F \cup M_c$, hence

$$\left| \bigcup_{i=1}^{\ell} \delta(A_i) \right| = |F| + |M_c|. \tag{3}$$

Also, $\bigcup_{i=1}^{\ell} A_i = \bigcup_{i=1}^{\ell} \phi_i$, hence

$$\frac{k}{k+1} \left| \bigcup_{i=1}^{\ell} A_i \right| = \frac{k}{k+1} \left| \bigcup_{i=1}^{\ell} \phi_i \right| = \frac{k}{k+1} \sum_{i=1}^{\ell} |\phi_i|. \tag{4}$$

Substituting inequalities (2), (3) and (4) into the second inequality in the fact gives

$$|F| + |M_c| \leq \frac{k}{k+1} \left| \bigcup_{i=1}^{\ell} A_i \right| + |M_c|,$$

which is the inequality in Theorem 4.3. □

Most of the complications in the proof of Theorem 4.3 seem to be caused by the presence of tight node sets $A_i \in \mathcal{F}$ such that $|\phi_i| = 1$. To illustrate the main ideas in the proof, we first prove (in 15 lines) a weaker version of Theorem 4.3. In the weaker version, the required upper bound of $k(n-1)/(k+1)$ is relaxed to $(n-1)$, and the laminar family of tight node sets $\mathcal{F} = \{A_1, \ldots, A_\ell\}$ is restricted such that every $A_i \in \mathcal{F}$ has $|\phi_i| \geq 2$. (The motivation for putting the restriction on $\mathcal{F}$ is expository. Such restricted laminar families $\mathcal{F}$ do not seem to be of mathematical interest.)

Proposition 4.6 *Let H, M, F and $\mathcal{F}$ be as in Theorem 4.3, and moreover, suppose that each tight node set $A_i \in \mathcal{F}$ has $|\phi_i| \geq 2$. Then*

$$|F| \leq \left| \bigcup_{i=1}^{\ell} A_i \right| \leq n - 1.$$

Proof: For an arbitrary $i = 1, \ldots, \ell$, consider A_i, ϕ_i, and let p denote $|\phi_i|$. By assumption, $p \geq 2$. Suppose that $p \leq k$ (the other case $p \geq k+1$ turns out to be easy). Then

$$|M \cap \delta(\phi_i)| \geq p(k - (p-1)), \tag{5}$$

since for every node $v \in \phi_i$, there are at most $(p-1)$ incident edges $vw \in E(H)$ with $w \in \phi_i$. Adding $2|\phi_i|$ to both sides of inequality (5) gives

$$2|\phi_i| + |M \cap \delta(\phi_i)| \geq 2p + p(k - (p-1)) \geq -p^2 + (k+2)p. \tag{6}$$

Subtracting $2k$ from both sides of inequality (6) gives

$$2|\phi_i| + |M \cap \delta(\phi_i)| - 2k \geq -p^2 + (k+2)p - 2k = -(p-k)(p-2) \geq 0, \tag{7}$$

where the last inequality $-(p-k)(p-2) \geq 0$ holds because $2 \leq p \leq k$. Inequality (7) implies

$$|\phi_i| + \frac{1}{2}|M \cap \delta(\phi_i)| \geq k = |\delta(A_i)|. \tag{8}$$

If $|\phi_i| \geq (k+1)$, then obviously inequality (8) holds.

Summing up inequality (8) over $i = 1, \ldots, \ell$ gives

$$\left|\bigcup_{i=1}^{\ell} \delta(A_i)\right| \leq \sum_{i=1}^{\ell} |\delta(A_i)| = k \cdot \ell \leq \sum_{i=1}^{\ell} |\phi_i| + \frac{1}{2}\sum_{i=1}^{\ell} |M \cap \delta(\phi_i)|. \tag{9}$$

The proof of Fact 4.5 shows that inequality (9) implies the inequality in the proposition, $|F| \leq \left|\bigcup_{i=1}^{\ell} A_i\right| \leq n-1$. □

Proof: (Theorem 4.3)

W.l.o.g. assume that $\mathcal{F}$ is minimal, i.e., for every $A_i \in \mathcal{F}$ there is an edge $e_i \in F$ such that $e_i \in \delta(A_i)$ and $e_i \notin \delta(A)$ for all $A \in \mathcal{F}$, $A \neq A_i$. Since $\mathcal{F}$ is minimal, every $A_i \in \mathcal{F}$ has $|\phi_i| \geq 1$. Let T be the tree representing $\mathcal{F} \cup \{V(H)\}$. The proof examines the node sets $A_i \in \mathcal{F}$, ϕ_i, but the node set $V(H) \backslash \bigcup\{A_i : A_i \in \mathcal{F}\}$ is not relevant for the proof. Every inclusionwise minimal $A_i \in \mathcal{F}$ has $|A_i| \geq (k+1)$, since $\delta(A_i) \cap F \neq \emptyset$ implies that A_i contains a node v with $\deg_H(v) \geq (k+1)$, so Lemma 4.2 implies this bound on $|A_i|$. Hence, every $A_i \in \mathcal{F}$ with $|\phi_i| = 1$ has at least one child in the tree T.

Two key assumptions are needed to complete the proof.

Assumption 1: For $1 \leq i \leq \ell$, every ϕ_i induces a complete subgraph of H, and moreover, every edge of this complete subgraph is in M, i.e., for $i = 1, \ldots, \ell$, $\forall v, w \in \phi_i, vw \in E(H)$ and $vw \in M$.

Assumption 2: For every $A_i \in \mathcal{F}$ with $|\phi_i| = 1$, there is an $A_j \in \mathcal{F}$ such that $|\phi_j| \leq k$ and A_j is a child of A_i in the tree T.

Claim 1: Assumption 1 causes no loss of generality.

Here is the proof of Claim 1. For an arbitrary $i = 1, \ldots, \ell$, consider ϕ_i and $E(\phi_i)$, the set of edges of H with both end nodes in ϕ_i. Clearly, an edge $vw \in E(\phi_i)$ is *not* in F, since vw is in none of the k-cuts $\delta(A_j)$ $(j = 1, \ldots, \ell)$. Therefore, all the missing edges vw with $v \in \phi_i, w \in \phi_i$ can be added to H (say, vw is first added to $E \backslash (M \cup F)$) such that ϕ_i induces a clique, and this will keep M, F and $\mathcal{F}$ unchanged. Moreover, every edge $vw \in E(\phi_i)$ can be placed in M, and the minimum degree requirement on (V, M) will continue to hold. By repeating this for each $i = 1, \ldots, \ell$, we obtain $H', M', F' = F$ and $\mathcal{F}' = \mathcal{F}$ that satisfy Assumption 1 and the conditions in the theorem. Clearly, if the inequality in the theorem holds for $H', M', F', \mathcal{F}'$, then it also holds for $H, M, F, \mathcal{F}$.

Claim 2: Assumption 2 causes no loss of generality.

Here is the proof of Claim 2. Consider an $A_i \in \mathcal{F}$ $(i = 1, \ldots, \ell)$ such that $|\phi_i| = 1$ and in the tree T every child $A_j \in \mathcal{F}$ of A_i has $|\phi_j| \geq (k+1)$. Let $\phi_i = \{v^*\}$. Let $A_j \in \mathcal{F}$ be an arbitrary T-child of A_i with $|\phi_j| \geq (k+1)$. Clearly, by Assumption 1, the subgraph of H induced by ϕ_j is a clique, and every edge in the clique is in M. Suppose that H has an edge wv^* such that $w \in A_j \backslash \phi_j$,

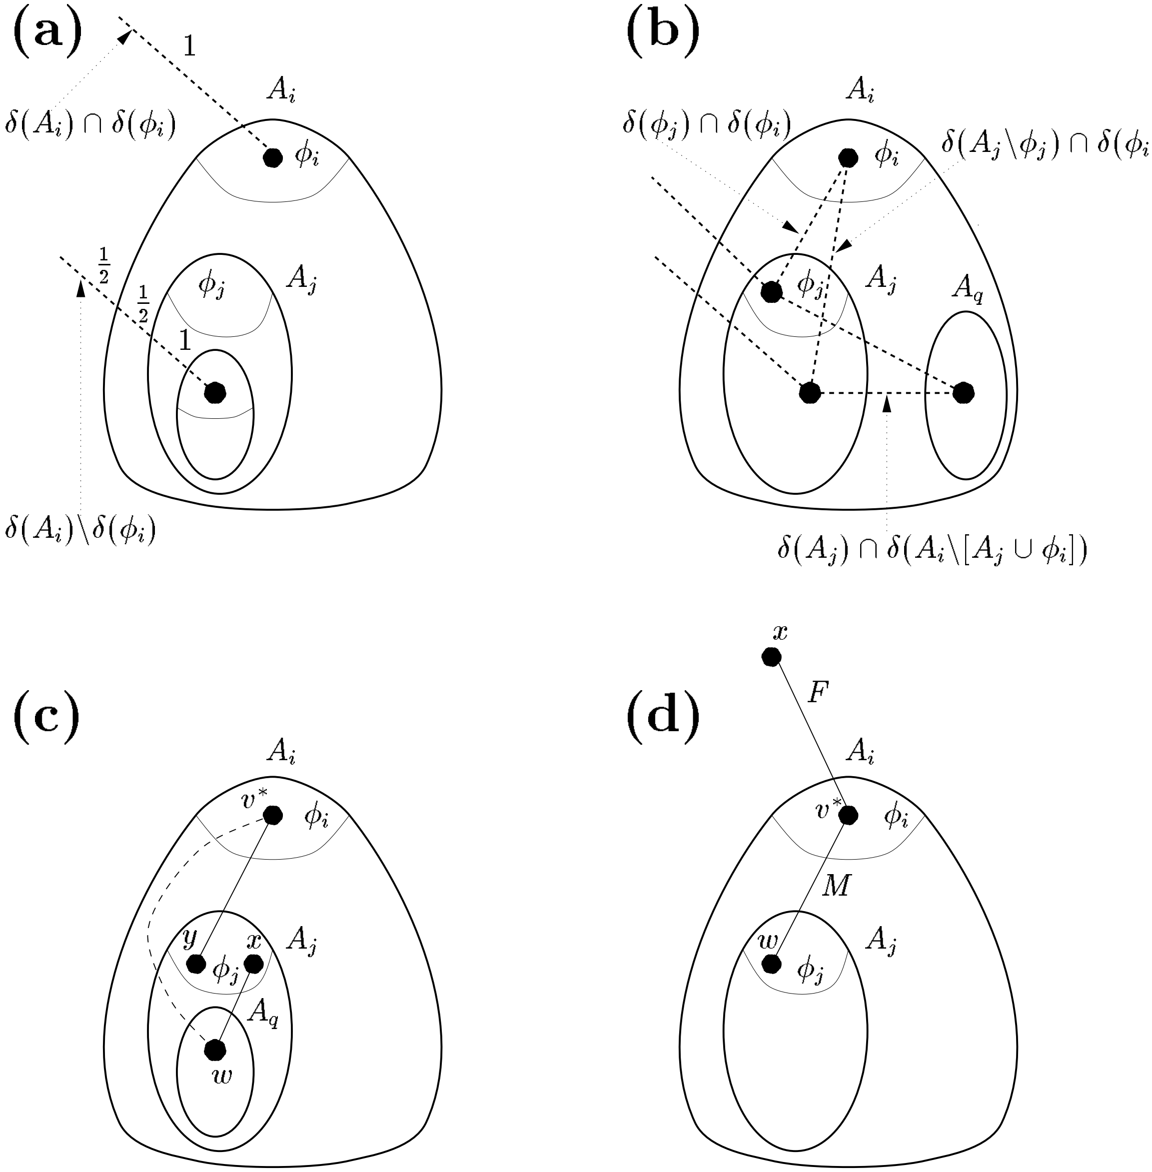

Figure 6: An illustration of the proof of Theorem 4.3.
(a) Every edge in $\delta(A_i) \cap \delta(\phi_i)$ contributes ≥ 1 to the l.h.s. of inequality (σ), and every edge in $\delta(A_i) \backslash \delta(\phi_i)$ contributes $\geq \frac{1}{2}$.
(b) The tight node set A_i is shown, together with two tight node sets A_j, A_q contained in A_i. The node sets ϕ_i and ϕ_j are also shown. The three kinds of edges arising in the proof are illustrated.
(c) In Claim 2, $\phi_i = \{v^*\}$ and $|\phi_j| \geq (k+1)$. An edge wv^* with $w \in A_j \backslash \phi_j$ is replaced by a pair of new edges wx and yv^*, where $x \in \phi_j$, $y \in \phi_j$.
(d) In Claim 2, $\phi_i = \{v^*\}$. If an edge v^*x with $x \notin A_i$ is in F (so $v^*x \notin M$), then there is an edge v^*w in M with $w \in \phi_j$, where $\phi_j \subseteq A_j \subset A_i$ and $|\phi_j| \geq (k+1)$. Edges v^*x and v^*w are swapped between M and F.

i.e., $wv^* \in \delta(A_j)\backslash\delta(\phi_j)$. (Figure 6(c) illustrates this.) Then we replace wv^* by a pair of new edges wx, yv^* with $x \in \phi_j, y \in \phi_j$ (possibly, $x = y$) such that the resulting graph H' is simple (i.e., H' has no multiedges); this can be done always, since $|\phi_j| \geq (k+1)$ and both $\delta(A_j)$ and $\delta(A_q)$ are k-cuts, where $A_q \in \mathcal{F}$ is the T-child of A_j containing node w. The resulting graph H' is k-edge connected. To see this, note that H is k-edge connected, and H' is obtained from H by replacing one edge wv^* by two edges wx, yv^*, where the nodes x and y are contained in the $(k+1)$-clique induced by ϕ_j. The formal proof of the k-edge connectivity of H' is easy, and is left to the reader. If $wv^* \in M$, then we take $M' = (M\backslash\{wv^*\}) \cup \{wx, yv^*\}$, $F' = F$, otherwise, we take $M' = M$, $F' = (F\backslash\{wv^*\}) \cup \{wx, yv^*\}$. In either case $\mathcal{F}$ covers F'. By repeating this manoeuvre for all relevant $i = 1, \ldots, \ell$, we obtain H', M', F' and $\mathcal{F}' = \mathcal{F}$ with $|F'| \geq |F|$ that satisfy the conditions in the theorem. Clearly, if the inequality in the theorem holds for $H', M', F', \mathcal{F}'$, then it also holds for $H, M, F, \mathcal{F}$. Moreover, the following condition $(*)$ holds:

for every $A_i \in \mathcal{F}$ with $|\phi_i| = 1$, for every T-child $A_j \in \mathcal{F}$ of A_i with $|\phi_j| \geq (k+1)$, every edge in $\delta(A_j) \cap \delta(\phi_i)$ is in $\delta(\phi_j)$. $\quad (*)$

Now w.l.o.g. suppose that H, M, F and $\mathcal{F}$ satisfy condition $(*)$. Call an $A_i \in \mathcal{F}$ *bad* if $|\phi_i| = 1$ and every T-child $A_j \in \mathcal{F}$ of A_i has $|\phi_j| \geq (k+1)$. Suppose that there is a bad $A_i \in \mathcal{F}$ with $\phi_i = \{v^*\}$ such that one of the edges $v^*x \in \delta(A_i) \cap \delta(\phi_i)$ is not in M. (Figure 6(d) illustrates this.) Then since $|\delta(A_i)| = k$, $|\delta(A_i) \cap F| \geq 1$, and $|M \cap \delta(\phi_i)| \geq k$, there must be an M-edge wv^* in $\delta(\phi_i)\backslash\delta(A_i)$. Let $A_j \in \mathcal{F}$ be the T-child of A_i such that $w \in A_j$. Since A_i is bad, $|\phi_j| \geq (k+1)$, therefore condition $(*)$ applies and ensures that the node w is in ϕ_j. Moreover, by Assumption 1, w is incident to $\geq k$ edges of M that have both end nodes in ϕ_j. Take $M' = (M\backslash\{wv^*\}) \cup \{v^*x\}$, $F' = (F\backslash\{v^*x\}) \cup \{wv^*\}$, and observe that $|M| = |M'|$, $|F| = |F'|$, every node $v \in V(H)$ is incident to $\geq k$ edges of M', F' consists of critical edges in $E(H)\backslash M'$, and $\mathcal{F}$ covers F'. By repeating this manoeuvre for all relevant $i = 1, \ldots, \ell$, we obtain H, M', F' and $\mathcal{F}$ that satisfy the conditions in the theorem such that $|F'| = |F|$, and for every bad $A_i \in \mathcal{F}$, no edge in $\delta(A_i) \cap \delta(\phi_i)$ is in F'. Then we can start with $\mathcal{F}$, and remove each bad A_i from $\mathcal{F}$ to obtain another laminar family $\mathcal{F}'$ covering F' such that $|\bigcup_{A \in \mathcal{F}'} A| \leq |\bigcup_{A \in \mathcal{F}} A|$, and $\mathcal{F}'$ satisfies Assumption 2. Clearly, if the inequality in the theorem holds for $H', M', F', \mathcal{F}'$, then it also holds for $H, M, F, \mathcal{F}$. This completes the proof of Claim 2.

Instead of proving that $F, \mathcal{F}$ satisfy inequality (1), we prove that under Assumption 2, M, F and $\mathcal{F} = \{A_1, \ldots, A_\ell\}$ satisfy the following sharper inequality (see Fact 4.5):

$$\left|\bigcup_{i=1}^{\ell} \delta(A_i)\right| \leq \frac{k}{k+1} \sum_{i=1}^{\ell} |\phi_i| + \frac{1}{2} \sum_{i=1}^{\ell} |M \cap \delta(\phi_i)|. \tag{10}$$

Clearly, every $A_i \in \mathcal{F}$ with $|\phi_i| \geq (k+1)$ satisfies the inequality

$$|\delta(A_i)| \leq \frac{k}{k+1} |\phi_i|. \tag{11}$$

From the proof of Proposition 4.6 (see inequalities (5), (6), (7), (8)), it follows that every $A_i \in \mathcal{F}$ with $2 \leq |\phi_i| \leq k$ satisfies the inequality

$$|\delta(A_i)| + \frac{k-1}{2(k+1)} |\phi_i| \leq \frac{k}{k+1} |\phi_i| + \frac{1}{2} |M \cap \delta(\phi_i)|, \tag{12}$$

where the surplus term on the left hand side (l.h.s.) is the difference between $k|\phi_i|/(k+1)$ and $|\phi_i|/2$. Every $A_i \in \mathcal{F}$ with $|\phi_i| = 1$ satisfies the inequality

$$|\delta(A_i) \cap \delta(\phi_i)| + \frac{1}{2} |\delta(A_i)\backslash\delta(\phi_i)| + \frac{k}{k+1} - \frac{1}{2} |\delta(A_i) \cap \delta(\phi_i)| \leq \frac{k}{k+1} |\phi_i| + \frac{1}{2} |M \cap \delta(\phi_i)|, \tag{13}$$

because $|\delta(A_i)\cap\delta(\phi_i)| + |\delta(A_i)\backslash\delta(\phi_i)| = |\delta(A_i)| = k \le |M\cap\delta(\phi_i)|$.

Claim 3: Under Assumption 2, the inequality (σ) obtained by summing up over all $A_i \in \mathcal{F}$ the appropriate one of inequalities (11), (12), (13) implies inequality (10), i.e., the l.h.s. of inequality (σ) is $\ge$ the l.h.s. of inequality (10), and the r.h.s. of inequality (σ) is $\le$ the r.h.s. of inequality (10).

Here is the proof of Claim 3. Clearly, inequality (σ) will imply inequality (10) if for every $A_i \in \mathcal{F}$, every edge in $\delta(A_i)\cap\delta(\phi_i)$ contributes ≥ 1 to the l.h.s. of inequality (σ). This property holds for $A_i \in \mathcal{F}$ with $|\phi_i| \ge 2$ by inequalities (11),(12), but for $A_i \in \mathcal{F}$ with $|\phi_i| = 1$ the property fails to hold (see inequality (13)). Fortunately, there is a way around this difficulty. For $A_i \in \mathcal{F}$ with $|\phi_i| = 1$, we allow A_i, ϕ_i to contribute a deficit of $\frac{1}{2}|\delta(A_i)\cap\delta(\phi_i)|$ on the l.h.s. of inequality (σ); using this deficit, we can ensure that every edge in $\delta(A_i)\cap\delta(\phi_i)$ (in $\delta(A_i)\backslash\delta(\phi_i)$) contributes ≥ 1 ($\ge 1/2$) to the l.h.s. of inequality (σ), see inequality (13). (Figure 6(a) illustrates the general scheme.) For each $A_i \in \mathcal{F}$ with $|\phi_i| = 1$, let $A_{c(i)} \in \mathcal{F}$ be an arbitrary T-child of A_i such that $1 \le |\phi_{c(i)}| \le k$; $A_{c(i)}$ exists by Assumption 2. Inequality (σ) implies inequality (10) because the deficit contributed by each $A_i \in \mathcal{F}$ with $|\phi_i| = 1$ is compensated by the surplus contributed by $A_{c(i)}, \phi_{c(i)}$. To see this, focus on an arbitrary $A_i \in \mathcal{F}$ with $|\phi_i| = 1$, and let $j = c(i)$. First observe that if an edge $vw \in \delta(A_j)$ with $v \in A_j$ is not in $\delta(A_i)$, then there are three possibilities: (i) $v \in \phi_j$, $w \in \phi_i$, (ii) $v \notin \phi_j$, $w \in \phi_i$, i.e., $v \in A_g$, where $A_g \in \mathcal{F}$ corresponds to a child of A_j in the tree T, and (iii) $v \in A_j$, $w \in A_i\backslash[A_j \cup \phi_i]$, i.e., $w \in A_q$, where $A_q \in \mathcal{F}$ corresponds to a sibling of A_j in the tree T. (Figure 6(b) illustrates the three possibilities.) Second, observe that

$$|\delta(A_i)\cap\delta(\phi_i)| \le |\delta(A_j)\backslash\delta(A_i)| = |\delta(A_j\backslash\phi_j)\cap\delta(\phi_i)| + |\delta(A_j)\cap\delta(A_i\backslash[A_j\cup\phi_i])| + |\delta(\phi_j)\cap\delta(\phi_i)|.$$

For each of the first two terms t on the right hand side, A_j, ϕ_j contributes a surplus of at least $t/2$ to the l.h.s. of inequality (σ), because (i) every edge in two distinct k-cuts $\delta(A_g)$ and $\delta(A_j)$, $A_g \in \mathcal{F}$, $A_j \in \mathcal{F}$, $A_g \subset A_j$, contributes a surplus of $1/2$ or more, since $A_h \in \mathcal{F}$ such that $\delta(\phi_h)\cap\delta(A_h)$ contains the edge contributes one for the edge, and every other $A \in \mathcal{F}$ such that $\delta(A)$ contains the edge contributes $\ge 1/2$ for the edge; (ii) every edge in two distinct k-cuts $\delta(A_q)$ and $\delta(A_j)$, $A_q \in \mathcal{F}$ disjoint from $A_j \in \mathcal{F}$, contributes a surplus of one or more.

Focus on the third term $|\delta(\phi_j)\cap\delta(\phi_i)|$, and note that its value is $\le |\phi_j|$, since $|\phi_i| = 1$ and the graph is simple. If $|\phi_j| = 1$, then the deficit contributed by A_i, ϕ_i (to the l.h.s. of inequality (σ)) is compensated, because the surplus of $\frac{k}{k+1}$ (on the l.h.s. of A_i's inequality) is $\ge \frac{1}{2}$ (for $k \ge 1$), hence

$$\frac{1}{2}|\delta(A_i)\cap\delta(\phi_i)| \quad \le \quad \frac{1}{2}|\delta(A_j\backslash\phi_j)\cap\delta(\phi_i)| + \frac{1}{2}|\delta(A_j)\cap\delta(A_i\backslash[A_j\cup\phi_i])| + \frac{k}{k+1}.$$

If $2 \le |\phi_j| \le k$, then the deficit contributed by A_i, ϕ_i (to the l.h.s. of inequality (σ)) is compensated, because the surplus of $\frac{k-1}{2(k+1)}|\phi_j| + \frac{k}{k+1}$ (on the l.h.s. of A_j's and A_i's inequalities) is $\ge |\phi_j|/2$ (for $k \ge |\phi_j| \ge 1$), hence

$$\frac{1}{2}|\delta(A_i)\cap\delta(\phi_i)| \quad \le \quad \frac{1}{2}|\delta(A_j\backslash\phi_j)\cap\delta(\phi_i)| + \frac{1}{2}|\delta(A_j)\cap\delta(A_i\backslash[A_j\cup\phi_i])| + \frac{k-1}{2(k+1)}|\phi_j| + \frac{k}{k+1}.$$

This completes the proof of Claim 3 and the proof of the theorem. □

Theorem 4.7 *Let $G = (V, E)$ be a graph of edge connectivity $\ge k \ge 1$. The heuristic described above finds a k-edge connected spanning subgraph (V, E') such that $|E'| \le (1 + [2/(k+1)])|E_{opt}|$, where $|E_{opt}|$ denotes the cardinality of an optimal solution. The running time is $O(k^3|V|^2 + |E|^{1.5}(\log|V|)^2)$.*

The next result is not relevant for the analysis of the heuristics in this paper, but may be of interest in graph theory. Given a k-edge connected graph H, let us call a critical edge of H *special* if both end nodes have degree at least $(k+1)$ in H. The number of special edges is at most $k|V(H)|$, since by Mader's result [Ma 72], the maximum number of critical edges in a k-edge connected graph H is at most $k|V(H)|$. Based on theorems of Cai [Ca 93], we give a bound of $6|V(H)|$ (independent of k) on the number of special edges in H, see Proposition 4.8.

Proposition 4.8 *The number of special edges in a k-edge connected, n-node graph H is at most $6n$ for odd $k \geq 1$, and at most $4n$ for even $k \geq 2$.*

Proof: Let $\mathcal{F}$ be a laminar family of tight node sets that covers all the special edges such that every $A \in \mathcal{F}$ has at least one special edge in $\delta(A)$. Let T be the tree representing $\mathcal{F} \cup \{V(H)\}$. Each special edge of H is in some k-cut that corresponds to a T-edge. Hence, the number of special edges is at most $k \cdot |E(T)| \leq k \cdot |V(T)|$. To estimate $|V(T)|$, we apply Theorems 5 and 6 of Cai [Ca 93], with slight modifications. One point to note is that Theorems 5,6 of [Ca 93] are stated for minimal k-edge connected graphs, but an examination of the proofs shows that these theorems apply to *all* k-edge connected (undirected) graphs. There are two cases:

- if $4 \leq k$ and k is even, then $|V(T)| \leq (4n/(k+4)) - (5k/(k+4)) \leq 4n/k$, and
- if $7 \leq k$ and k is odd, then $|V(T)| \leq (6n/(k+6)) - (8k/(k+6)) \leq 6n/k$.

Hence, the number of special edges in H is at most $6n$ for odd k, and at most $4n$ for even k. Note that for $k = 2$ (or $k = 1, 3$ or 5), the number of special edges is at most kn, which is $\leq 4n$ (or $\leq 6n$), since by results in [Ma 72], the number of critical edges is at most kn, and every special edge is a critical edge. $\square$

4.2 Directed graphs

The heuristic for finding an approximately minimum-size k-edge connected spanning subgraph of a digraph has two steps. Similarly to Section 3.4, the first step finds a minimum-cardinality arc set $M \subseteq E$ such that for every node v, there are $\geq k$ arcs of M going out of v and $\geq k$ arcs of M coming into v. Clearly, $|M| \leq |E_{opt}|$, where $E_{opt} \subseteq E$ denotes a minimum-cardinality arc set such that (V, E_{opt}) is k-edge connected. The second step of the heuristic finds an (inclusionwise) minimal arc set $F \subseteq E \backslash M$ such that $E' = M \cup F$ is the arc set of a k-edge connected spanning subgraph. To prove the approximation guarantee, we need to estimate $|F|$. We use the notion of special arcs to estimate $|F|$. Call an arc (v, w) of a k-edge connected digraph *special* if the arc is critical, and in addition, $\deg_{out}(v) \geq (k+1)$ and $\deg_{in}(w) \geq (k+1)$. Clearly, every arc in F is a special arc of the digraph $G' = (V, E')$, $E' = M \cup F$, returned by the heuristic. We can deduce a bound of $O(\sqrt{k}|V|)$ on the number of special arcs in G' by examining chains of tight node sets $S_1 \subset S_2 \subset \ldots \subset S_q$, where a node set S_i is called tight if G' has exactly k arcs in $\delta_{out}(S_i)$.

Theorem 4.9 *Let $k \geq 1$ be an integer, and let H be a k-edge connected, n-node digraph. The number of special arcs in H is at most $4\sqrt{k} \cdot n$.*

Proof: Let V denote $V(H)$ for this proof. Each special arc e is in a k-dicut $\delta_{out}(A_e) = \delta_{in}(V \backslash A_e)$, where $2 \leq |A_e| \leq n-2$. As in Section 4.1, we obtain *two* laminar families of tight node sets $\mathcal{F}_{out}$ and $\mathcal{F}_{in}$ that cover all the special arcs: that is, for each $A_i \in \mathcal{F}_{out}$ ($A_i \in \mathcal{F}_{in}$), A_i is a set of H-nodes, $\delta_{out}(A_i)$ ($\delta_{in}(A_i)$) has k arcs including at least one special arc, and each special arc is in some $\delta_{out}(A_i)$, $A_i \in \mathcal{F}_{out}$, or is in some $\delta_{in}(A_i)$, $A_i \in \mathcal{F}_{in}$. Focus on $\mathcal{F}_{out}$; the analysis is symmetric for $\mathcal{F}_{in}$. Let $\mathcal{F}_{out} = \{A_1, A_2, \ldots, A_\ell\}$. To estimate the number of special arcs, we need to examine the tree T corresponding to $\mathcal{F}_{out} \cup \{V(H)\}$. For $i = 1, \ldots, \ell$, recall that the T-node corresponding to a node set $A_i \in \mathcal{F}_{out}$ is also denoted A_i (the T-node corresponding to $V(H)$ is denoted by V), and recall that ϕ_i denotes $A_i \backslash \bigcup\{A \in \mathcal{F}_{out} : A \subset A_i,\ A \neq A_i\}$. Partition the set $\{A_1, \ldots, A_\ell\}$ of T-nodes into two sets R_1 and R_2, where R_2 consists of the T-nodes incident to precisely two T-edges, and $R_1 = \{A_1, \ldots, A_\ell\} \backslash R_2$. Note that $V \notin R_1$ and $V \notin R_2$.

Claim 1: $|R_1| \leq 2|V_1|/(k+1)$, where V_1 denotes the set of H-nodes in $\bigcup\{\phi_i : A_i \in R_1\}$.

Here is the proof of Claim 1. Let T_1 be the tree obtained from the tree T by "unsubdividing" all the T-nodes in R_2, i.e., by repeatedly replacing a degree-two T-node in R_2 and its two incident edges by an edge between the two neighbours. Then T_1 is a tree whose nonleaf T-nodes in R_1 have T_1-degree ≥ 3, whereas the T-node V may have T_1-degree $1, 2$ or ≥ 3. Let ℓ_1 be the number of leaf nodes (degree-1 nodes) of T_1 in R_1. Then, $|R_1| \leq \ell_1 + (\ell_1 + 1) - 2 \leq 2\ell_1$. Now, Claim 1 follows because $\ell_1 \leq |V_1|/(k+1)$, because for each (inclusionwise) minimal $A_i \in \mathcal{F}_{out}$, the set $\phi_i = A_i$ of H-nodes has cardinality at least $(k+1)$ by the digraph version of Lemma 4.2(ii). (A_i contains a node v with $\deg_{out}(v) \geq (k+1)$ since $\delta_{out}(A_i)$ contains a special arc.)

Now focus on a maximal path $P = A_0, A_1, \ldots, A_{q+1}$ of T such that every T-node A_i with $1 \leq i \leq q$ is in R_2. In H, the node sets $A_0, A_1, \ldots, A_{q+1}$ satisfy $A_0 \subset A_1 \subset \ldots \subset A_{q+1}$, and for $i = 1, \ldots, q$, if $A' \in \mathcal{F}_{out}$ is contained in A_i, then either $A' = A_{i-1}$ or $A' \subset A_{i-1}$. Let V_P denote the set of H-nodes $\phi_1 \cup \phi_2 \cup \ldots \cup \phi_q$. Also, note that for $i = 1, 2, \ldots, q$, $A_i = A_0 \cup \phi_1 \cup \phi_2 \cup \ldots \cup \phi_i$.

Claim 2: The number of arcs (v, w) such that $v \in V_P$ and $(v, w) \in \bigcup\{\delta_{out}(A_i) : 1 \leq i \leq q\}$ is at most $k + 2\sqrt{k} \cdot |V_P|$.

Here is the proof of the Claim 2; see Figure 7 for an illustration. The additional term of k in the upper bound accounts for the arcs with start nodes in A_q and end nodes in $V \backslash A_q$; there are at most k such arcs, since each such arc is in $\delta_{out}(A_q)$. Now ignore the arcs in $\delta_{out}(A_q)$. Linearly order the H-nodes in V_P such that for each i, $1 \leq i < q$, the H-nodes in ϕ_i precede the H-nodes in ϕ_{i+1}. Let v be an arbitrary node in V_P. Let $\Gamma_v \subseteq V_P$ denote the set of end nodes w_j of the arcs (v, w_j) outgoing from v such that $w_j \in V_P$ and $(v, w_j) \in \bigcup\{\delta_{out}(A_i) : 1 \leq i \leq q\}$. Let the linear ordering of the nodes in Γ_v be $w_1, w_2, \ldots, w_{|\Gamma_v|}$. Call an arc (v, w_j) *short* if $j \leq \sqrt{k}$, otherwise, call the arc *long*. We "charge" each long arc (v, w_j) to the first $\sqrt{k}$ nodes $w_1, w_2, \ldots, w_{\sqrt{k}}$ in Γ_v, i.e., each of these nodes is charged $1/\sqrt{k}$ for each arc (v, w_j), $w_j \in \Gamma_v$ and $j > \sqrt{k}$. Now consider the total charge on an arbitrary node $w_a \in V_P$ due to all long arcs $(x, y) \in \bigcup\{\delta_{out}(A_i) : 1 \leq i \leq q\}$ with $x \in V_P$ and $y \in V_P$. The key fact is this: the total charge on w_a is at most $\sqrt{k}$. To see this suppose that $w_a \in \phi_i$, where $1 \leq i \leq q$. Then for every arc (v, w_j) charged to w_a, $(v, w_j) \in \delta_{out}(A_{i-1})$, because $v \in A_i \backslash \phi_i$ (if $v \in V \backslash A_i$ or $v \in \phi_i$, then clearly Γ_v does not contain a node of ϕ_i such as w_a). Furthermore, by the linear ordering of Γ_v, $w_j \in \phi_i \cup \phi_{i+1} \cup \ldots \cup \phi_q$, i.e., $w_j \notin A_{i-1}$. Since $\delta_{out}(A_{i-1})$ has k arcs, the total charge to w_a is at most $k \cdot (1/\sqrt{k}) = \sqrt{k}$. Finally, consider the total number, m_P, of short arcs $(x, y) \in \bigcup\{\delta_{out}(A_i) : 1 \leq i \leq q\}$ with $x \in V_P$ and $y \in V_P$. Obviously, m_P is at most $\sqrt{k}|V_P|$. Claim 2 is completed by summing up the three terms: k (for arcs in $\delta_{out}(A_q)$), $\sqrt{k}|V_P|$ (for the total charge on nodes $w \in V_P$), and $\sqrt{k}|V_P|$ (for m_P).

We account for the special arcs in $\delta_{out}(A_q)$ by "charging" the additional term of k to the "unsubdivided edge" $A_0 A_{q+1}$ of the tree T_1 in the proof of Claim 1. Thus each edge $A_i A_{i+q+1}$, $A_i \subset A_{i+q+1}$, of T_1 is "charged" for $\leq 2k$ special arcs (these are the special arcs in $\delta_{out}(A_i) \cup \delta_{out}(A_{i+q})$). Since the number of edges in T_1 is $\leq |R_1|$, the number of special arcs contributed by the T-nodes in R_1 is $\leq 2k|R_1|$. We "charge" $2\sqrt{k}$ to each H-node v such that $v \in \phi_i$ for a T-node $A_i \in R_2$. Combining the contributions of special arcs from the T-nodes in R_1 and R_2 and applying Claim 1,

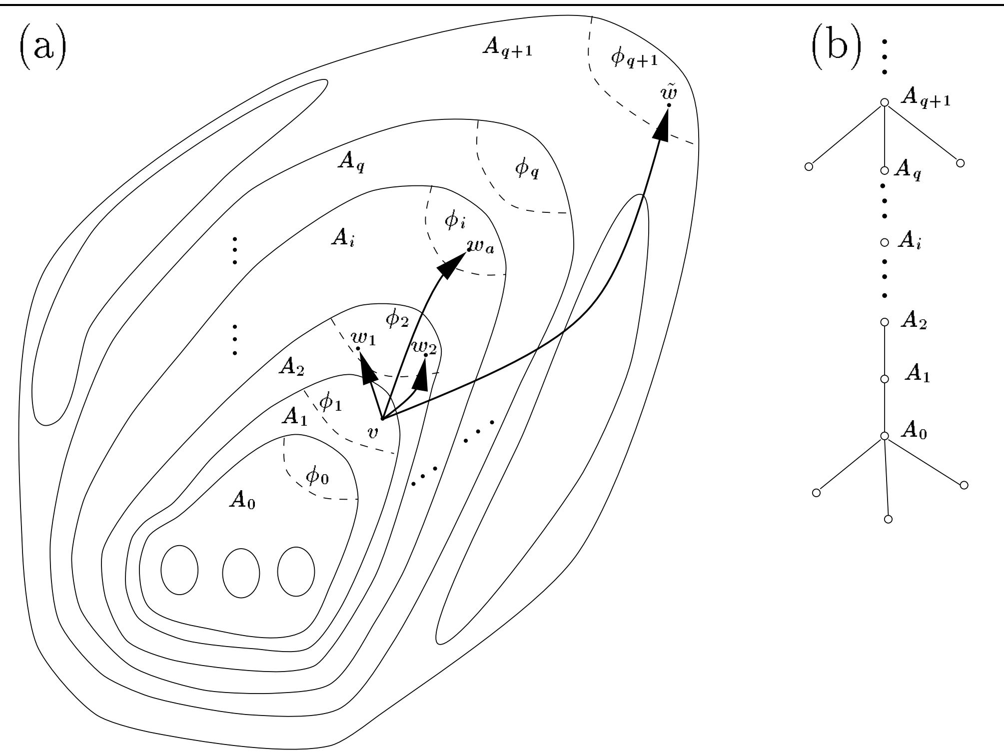

Figure 7: An illustration of Claim 2 in the proof of Theorem 4.9.
(a) A subfamily of the laminar family of tight node sets $\mathcal{F}_{out}$ that covers (some of) the special arcs.
(b) The subtree corresponding to the subfamily of $\mathcal{F}_{out}$ in (a). Each of the T-nodes $A_1, A_2, \ldots, A_q$ is incident to exactly two edges of T, where T is the tree corresponding to $\mathcal{F}_{out}$.

we see that the number of special arcs is at most

$$2k|R_1| + 2\sqrt{k} \cdot n_2 \leq \frac{4kn_1}{(k+1)} + 2\sqrt{k} \cdot n_2$$

where n_1 and n_2 denote the cardinalities of $V_1 = \bigcup\{\phi_i : A_i \in R_1\}$ and $V_2 = \bigcup\{\phi_i : A_i \in R_2\}$, respectively. For $k \geq 1$, the number of special arcs is maximized when n_2 is maximum possible and n_1 is minimum possible. Since the tree T has at least two leafs, n_2 is at most $n - (2k+2)$. Hence, the number of special arcs contributed by $\mathcal{F}_{out}$ is at most $4k(2k+2)/(k+1) + 2\sqrt{k}(n - (2k+2))$. The total number of special arcs in H is at most $16k + 4\sqrt{k}(n - (2k+2)) \leq 4\sqrt{k}n$. □

The heuristic clearly runs in time $O(k|E|^2)$. This can be improved by implementing the second step to run in time $O(k^3|V|^2)$. We run Gabow's algorithm [Ga 95] as a preprocessing step to compute a sparse certificate $\widetilde{E}$ of G for k-edge connectivity, i.e., $\widetilde{E} \subseteq E$, $|\widetilde{E}| \leq 2k|V|$, and for all nodes v, w, $(V, \widetilde{E})$ has k arc-disjoint $v \rightarrow w$ directed paths iff G has k arc-disjoint $v \rightarrow w$ directed paths. In detail, we fix a node $a \in V(G)$ and take $\widetilde{E} = \widetilde{E}_{out} \cup \widetilde{E}_{in}$, where $\widetilde{E}_{out}$ ($\widetilde{E}_{in}$) is the union of k arc-disjoint out-branchings (in-branchings) rooted at a. Gabow's algorithm [Ga 95] runs in time $O(k|V|^2)$, and the second step runs in time $O(k|\widetilde{E} \cup M|^2) = O(k^3|V|^2)$.

Theorem 4.10 *Let $G = (V, E)$ be a digraph of edge connectivity $\geq k$. The heuristic described above finds a k-edge connected spanning subgraph (V, E') such that $|E'| \leq (1 + [4/\sqrt{k}])|E_{opt}|$, where $|E_{opt}|$ denotes the cardinality of an optimal solution. The running time is $O(k^3|V|^2 + |E|^{1.5}(\log|V|)^2)$.*

The upper bound on the number of special arcs in Theorem 4.9 is not tight, but is within a factor of (roughly) three of the tight bound for $n \gg k$. To see this, take $n \geq 3k+2$ and consider the following k-edge connected, n-node digraph $\hat{G}$ with at least $\beta n - 2\beta(k+1) + k$ special arcs, where β is the maximum integer such that $\beta(\beta+1)/2 \leq k$, i.e., $\beta = \lfloor\sqrt{2k+0.25} - 0.5\rfloor$. See Figure 8 for an illustration of $\hat{G}$. $\hat{G}$ has a "left" $(k+1)$-directed clique K_L and a "right" $(k+1)$-directed clique K_R. Let $v_1, v_2, \ldots, v_\ell$ be a linear ordering of the remaining nodes, where $\ell = n - 2(k+1) \geq k$. There is one arc from v_i $(1 \leq i \leq \ell)$ to each of the next β nodes $v_{i+1}, \ldots, v_{i+\beta}$; hence, each node v_i has one arc coming in from each of the previous β nodes $v_{i-1}, \ldots, v_{i-\beta}$. (Take $v_0, v_{-1}, v_{-2} \ldots, v_{-\beta+1}$ to mean nodes in K_L, and take $v_{\ell+1}, v_{\ell+2}, \ldots, v_{\ell+\beta}$ to mean nodes in K_R.) These β left-to-right arcs starting from v_i will turn out to be special arcs. Additionally, there are $(k+1-\beta)$ arcs from K_R to each of the nodes $v_1, v_2, \ldots, v_\ell$, and there are $(k+1-\beta)$ arcs from each of the nodes $v_1, v_2, \ldots, v_\ell$ to K_L. Finally, there are $(k - \beta(\beta+1)/2)$ arcs from K_L to K_R. This completes the construction of $\hat{G}$. It can be checked that $\hat{G}$ is k-edge connected. (Note that besides the $(k - \beta(\beta+1)/2)$ arcs from K_L to K_R, there are $\beta(\beta+1)/2$ arc-disjoint directed paths from K_L to K_R, such that there is one "one-hop" directed path, two "two-hop" directed paths, ..., β "β-hop" directed paths). For each node set A in the laminar family of node sets $\{K_L, (K_L \cup \{v_1\}), \ldots, (K_L \cup \{v_1, v_2, \ldots, v_\ell\})\}$, the out-directed cut $\delta_{out}(A)$ has cardinality k, and every arc in $\delta_{out}(A)$ is a special arc.

5 Conclusions

Our analyses of the heuristics exploit results from extremal graph theory, such as Mader's remarkable theorem [Ma 72, Theorem 1], and raise new problems in the areas of approximation algorithms and extremal graph theory.

For a graph G and an integer $k \geq 1$, let $\mu(k, G)$ denote the minimum number of edges in a spanning subgraph of minimum degree k. For a digraph G and integer $k \geq 1$, define $\mu(k, G)$ similarly. For a graph (or digraph) G and integer $k \geq 1$, let $\mu'(k, G)$ denote the minimum number

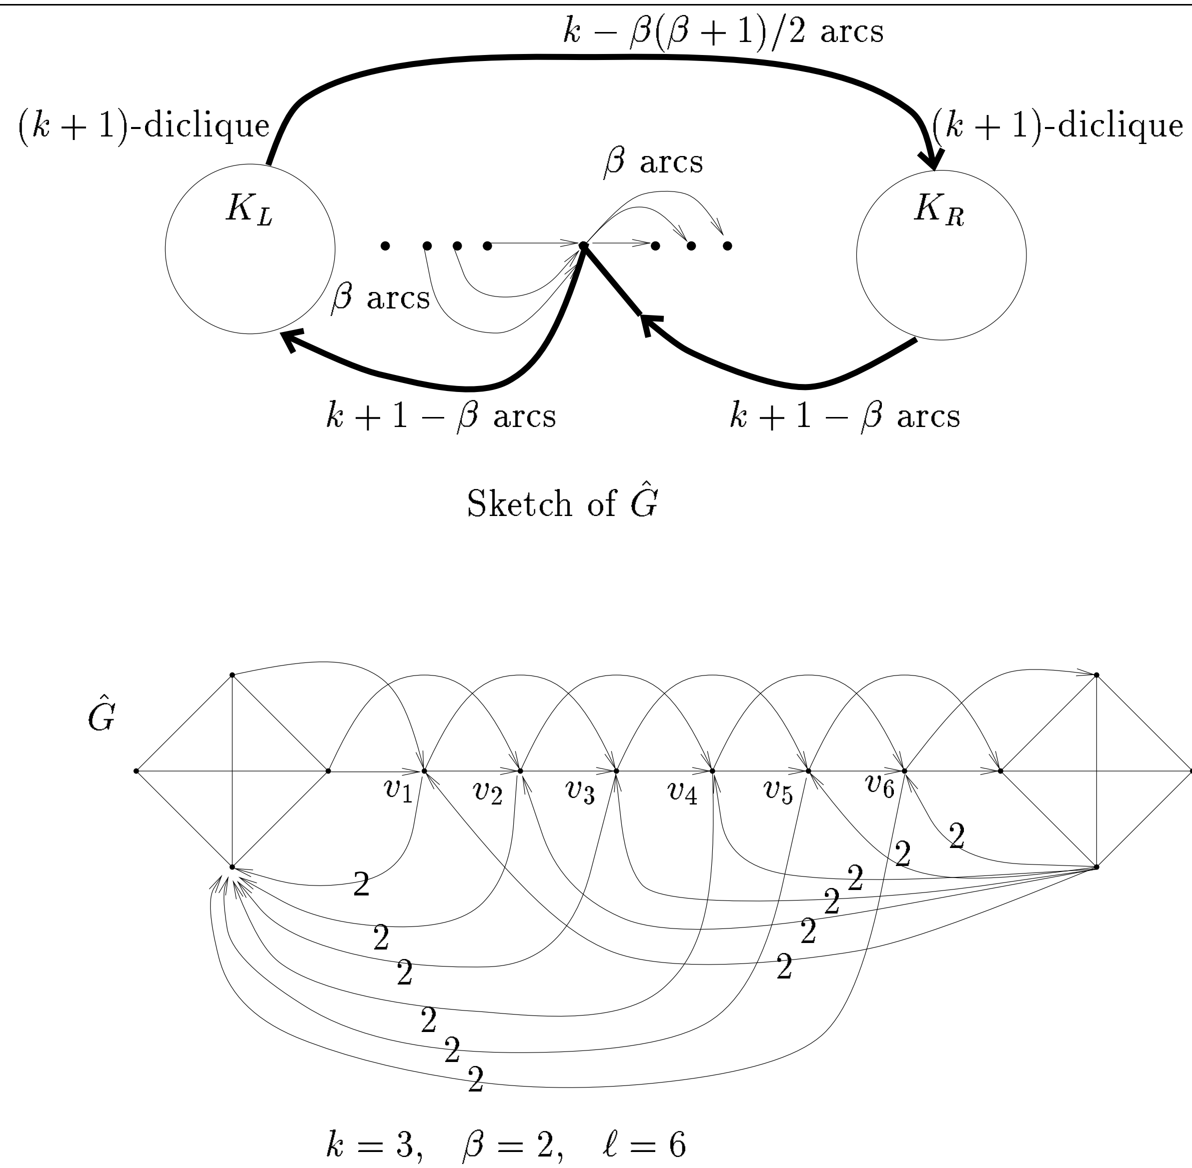

Figure 8: The digraph $\hat{G}$ described in the last paragraph of Section 4.2. $\hat{G}$ has $n \geq 3k+2$ nodes, and has $\geq \beta(n - 2(k+1)) + k$ special arcs, $\sqrt{k} \leq \beta < \sqrt{2k}$, showing that the upper bound on the number of special arcs in Theorem 4.9 is within a small constant factor of being tight for $n \gg k$.

of edges (arcs) in a k-edge connected spanning subgraph (k-ECSS), and let $\mu''(k,G)$ denote the minimum number of edges (arcs) in a k-node connected spanning subgraph (k-NCSS). While $\mu(k,G)$ can be computed efficiently via b-matchings, computing either $\mu'(k,G)$ or $\mu''(k,G)$ is NP-hard. This paper shows that (i) by computing $\mu(k-1,G)$, we can efficiently approximate $\mu''(k,G)$ to within a factor of $1+[1/k]$ for both graphs and digraphs, and (ii) by computing $\mu(k,G)$, we can efficiently approximate $\mu'(k,G)$ to within a factor of $1+[2/(k+1)]$ for graphs, and a factor of $1+[4/\sqrt{k}]$ for digraphs. Theorem 3.6 shows that for a k-node connected graph G, $\frac{\mu''(k,G)}{\mu'(k,G)} \le \frac{k+1}{k}$, and Theorem 3.21 shows that for a k-node connected digraph G, $\frac{\mu''(k,G)}{\mu(k,G)} \le \frac{k+1}{k}$. Propositions 3.4 and 3.20 show that for a k-node connected graph or digraph G, $\frac{\mu''(k,G)}{\mu(k-1,G)} \le \frac{k+1}{k-1}$. Theorem 4.7 shows that for a k-edge connected graph G, $\frac{\mu'(k,G)}{\mu(k,G)} \le \frac{k+3}{k+1}$.

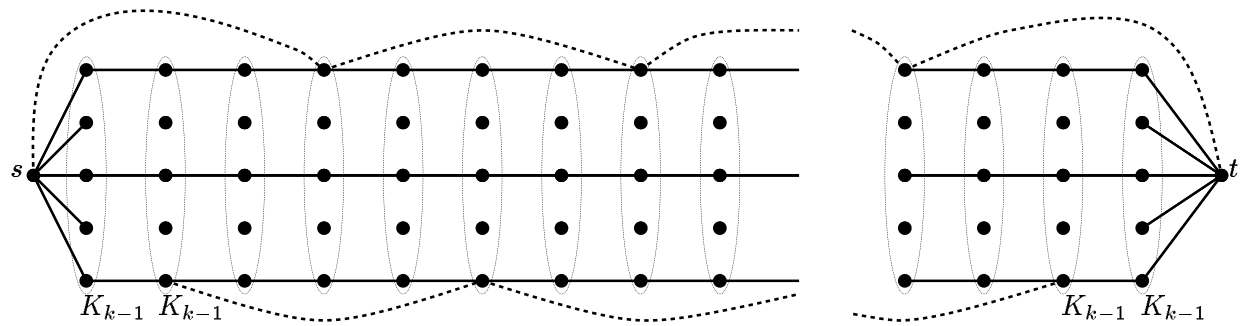

Figure 9: A k-node connected graph $G=(V,E)$ (with $k \ge 2$) such that the minimum size μ'' of a k-node connected spanning subgraph decreases by $(n-3k+1)/(2k-2)$ on adding one edge.
G consists of nodes s, t, and ℓ copies of the $(k-1)$-clique, and has $k-1$ openly disjoint $s{\leftrightarrow}t$ paths such that each path uses exactly one node from each $(k-1)$-clique; also, G has $(\ell-1)/2$ dashed edges. Every edge in G is critical w.r.t. k-node connectivity. Adding the edge st to G, and then removing all the dashed edges leaves a k-node connected graph, so μ'' decreases from $|E|$ to $|E|+1-(\ell-1)/2$.
A k-edge connected (and k-node connected) graph $\widetilde{G}$ such that the minimum size μ' of a k-edge connected spanning subgraph decreases by $\frac{|V(\widetilde{G})|-4k+2}{3k-3}$ on adding one edge can be obtained by modifying G as follows: "split" every $(k-1)$-clique incident with a dashed edge into a pair of $(k-1)$-cliques connected by a matching of size $(k-1)$.

For minimum-size k-ECSS (k-NCSS) problems, there appears to be a difficulty in achieving approximation guarantees of $1+\frac{\omega(1)}{k^2}$. A graph theoretic function g is said to satisfy the edge Lipschitz condition if whenever graphs H and H' differ in only one edge, then $|g(H)-g(H')| \le 1$, see [AS 92, p. 86]. Observe that $\mu(k,G)$ satisfies the edge Lipschitz condition. Moreover, the optimal size of a b-matching satisfies the edge Lipschitz condition, and so do most functions related to matchings of graphs. In contrast, both $\mu'(k,G)$ and $\mu''(k,G)$ violate this condition. First, focus on $\mu'(k,G)$ for graphs G and $k \ge 2$. Let G be the minimal k-edge connected graph obtained by "stringing" ℓ copies of the $(k+1)$-clique, i.e., take ℓ copies of the $(k+1)$-clique, and for each copy i, $1 \le i \le \ell$, designate a pair of distinct nodes as s_i and t_i, and then identify t_i and s_{i+1} for

$i = 1, 2, \ldots, \ell - 1$. Adding the edge $s_1 t_\ell$ decreases μ' by $\ell = (|V(G)| - 1)/k$, since removing all the edges $s_i t_i$, $1 \leq i \leq \ell$, leaves a k-edge connected graph. Now consider $\mu''(k, G)$ for graphs G and $k \geq 2$. For each $k \geq 2$, there exists a k-node connected graph G such that adding a particular new edge decreases μ'' by $\dfrac{|V(G)| - 3k + 1}{2k - 2}$; see Figure 9 for an illustration. For $k = 2$ and the graph G in Figure 1, observe that μ'' decreases from $1.5|V| - 5$ to $|V| + 1$ upon adding the edge e_*. A k-edge connected (and k-node connected) graph $\widetilde{G}$ such that adding a particular new edge decreases μ' by $\dfrac{|V(\widetilde{G})| - 4k + 2}{3k - 3}$ can be obtained by modifying the graph in Figure 9 as indicated in the figure caption. Garg et al [GSS 93] discuss similar issues for the minimum-size 2-NCSS problem on graphs.

Another drawback of the analysis of the k-NCSS heuristic for graphs in Section 3.1 is that the size of the edge set $E' = M \cup F$ returned by the heuristic is compared against $\mu'(k, G)$, the minimum size of a k-ECSS. Given an integer $k \geq 2$, for each integer $n = 2k(i + k) + k$, where $i = 0, 1, 2, \ldots$, there exists a k-node connected, n-node graph $\hat{G}$ such that

$$\frac{\mu''(k, \hat{G})}{\mu'(k, \hat{G})} = 1 + \frac{(k - 2)}{(2k^2 + k)}.$$

In view of this, for large k, a sharper lower bound will have to be employed for proving approximation guarantees substantially better than $1 + [1/2k]$ for the minimum-size k-NCSS problem. For $k = 2$ or $k = 3$, larger values of $\mu''(k, G)/\mu'(k, G)$ are given by the graph G in Figure 9 with the parameter k fixed at 2 or 3 and with $|V(G)| \gg k$: for $k = 2$, the ratio approaches $6/5$, and for $k = 3$, the ratio approaches $14/13$.

Here is another consequence of R. P. Gupta's result, see the proof of Proposition 3.8: For a bipartite graph G with minimum degree $\geq k$,

$$\frac{\mu(k - 1, G)}{\mu(k, G)} \leq \frac{(k - 1)}{k}.$$

This inequality does not hold for nonbipartite graphs, since for $G = K_{(k+1)}$, $\mu(k - 1, G)/\mu(k, G)$ equals $(k - 1)/k$ for k odd, and equals $k/(k + 1)$ for k even. Another result of Gupta, see [BM 76, Problem 6.2.8], shows that $\mu(k - 2, G)/\mu(k, G) \leq (k - 2)/(k - 1)$ for all graphs G of minimum degree $\geq k$.

Acknowledgments: A jointly authored and preliminary version of this paper, containing Propositions 3.4, 3.20, Lemmas 3.3, 3.13, 3.19, and Theorem 3.16, was submitted to the IEEE FOCS 1996 program committee. The other results are due to the first author.
Thanks to W. H. Cunningham, H. R. Hind, A. V. Kotlov, U. S. R. Murty and A. Sebo for helpful discussions. U. S. R. Murty suggested the use of Gupta's result for proving Proposition 3.8.

References

[AS 92] N. Alon and J. H. Spencer, *The Probabilistic Method*, John Wiley & Sons Inc., New York, 1992.

[Bo 78] B. Bollobás, *Extremal Graph Theory*, Academic Press, London, 1978.

[BM 76] J. A. Bondy and U. S. R. Murty, *Graph Theory with Applications*, American Elsevier Publishing Co., New York, 1976.

[Ca 93] M. Cai, "The number of vertices of degree k in a minimally k-edge-connected graph," *J. Combinatorial Theory Series B* **58** (1993), 225–239.

[CKT 93] J. Cheriyan, M. Y. Kao and R. Thurimella, "Scan-first search and sparse certificates: An improved parallel algorithm for k-vertex connectivity," *SIAM J. Computing* **22** (1993), 157–174.

[CL 95] K. W. Chong and T. W. Lam, "Approximating biconnectivity in parallel," Proc. 7th Annual ACM SPAA 1995, 224–233. To appear in *Algorithmica*.

[CL 96] K. W. Chong and T. W. Lam, "Improving biconnectivity approximation via local optimization," Proc. 7th Annual ACM-SIAM Symposium on Discrete Algorithms 1996, 26–35.

[CL 96b] K. W. Chong and T. W. Lam, "Towards more precise parallel biconnectivity approximation," Proc. 7th International Symposium on Algorithms and Computation, December 1996.

[Ch 76] N. Christofides, "Worst-case analysis of a new heuristic for the travelling salesman problem," Technical report, G.S.I.A., Carnegie-Mellon Univ., Pittsburgh, PA, 1976.

[Ed 72] J. Edmonds, "Edge-disjoint branchings," in *Combinatorial Algorithms*, Ed. R. Rustin, Algorithmics Press, New York, 1972, 91–96.

[Fe 97] C. G. Fernandes, "A better approximation ratio for the minimum k-edge-connected spanning subgraph problem," Proc. 8th Annual ACM-SIAM Symposium on Discrete Algorithms 1997, 629–638.

[FGHP 93] T. Fischer, A. V. Goldberg, D. J. Haglin and S. Plotkin, "Approximating matchings in parallel," *Information Processing Letters* **46** (1993), 115-118.

[Fr 93] A. Frank, "Submodular functions in graph theory," *Discrete Mathematics* **111** (1993), 231–243.

[FIN 93] A. Frank, T. Ibaraki and H. Nagamochi, "On sparse subgraphs preserving connectivity properties," *J. Graph Theory* **17** (1993), 275–281.

[Ga 85] H. N. Gabow, "A scaling algorithm for weighted matching on general graphs," Proc. 26th Annual IEEE FOCS 1985, 90–100.

[Ga 95] H. N. Gabow, "A matroid approach to finding edge connectivity and packing arborescences," *Journal of Computer and System Sciences* **50** (1995), 259–273.

[GaTa 91] H. N. Gabow and R. E. Tarjan, "Faster scaling algorithms for general graph matching problems," *Journal of the ACM* **38** (1991), 815–853.

[GJ 79] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman, San Francisco, 1979.

[GSS 93] N. Garg, V. S. Santosh, and A. Singla, "Improved approximation algorithms for biconnected subgraphs via better lower bounding techniques," Proc. 4th Annual ACM-SIAM Symposium on Discrete Algorithms 1993, 103–111.

[Ge 95] A. M. H. Gerards, "Matching," in Handbook of Operations Research and Management Science (Eds. M. O. Ball, T. L. Magnanti, C. L. Monma, and G. L. Nemhauser), North-Holland, Amsterdam, 1995.

[Gu 67] R. P. Gupta, "A decomposition theorem for bipartite graphs," in *Théorie des Graphes Rome I. C. C.*, Ed. P. Rosenstiehl, Dunod, Paris, 1967, 135–138.

[HKe+ 95] X. Han, P. Kelsen, V. Ramachandran and R. Tarjan, "Computing minimal spanning subgraphs in linear time," *SIAM J. Computing* **24** (1995), 1332-1358.

[Jo 93] T. Jordán, "Increasing the vertex-connectivity in directed graphs," *Proc. Algorithms — ESA '93, 1st Annual European Symposium*, **LNCS 726**, Springer, New York, (1993), 236–247.

[Jo 95] T. Jordán, "On the optimal vertex-connectivity augmentation," *J. Combinatorial Theory Series B* **63** (1995), 8–20. Preliminary version in *Proc. 3rd I.P.C.O.* (1993), 75-88.

[Ka 94] D. R. Karger, "Random sampling in cut, flow, and network design problems," Proc. 26th Annual ACM STOC 1994, 648–657. To appear in *Mathematics of Operations Research.*

[KeR 95] P. Kelsen and V. Ramachandran, "On finding minimal two-connected subgraphs," *Journal of Algorithms* **18** (1995), 1–49.

[K 96] S. Khuller, "Approximation algorithms for finding highly connected subgraphs," in *Approximation algorithms for NP-hard problems*, Ed. D. S. Hochbaum, PWS publishing co., Boston, 1996.

[KR 96] S. Khuller and B. Raghavachari, "Improved approximation algorithms for uniform connectivity problems," *Journal of Algorithms* **21** (1996), 434–450. Preliminary version in: Proc. 27th Annual ACM STOC 1995, 1–10.

[KRY 95] S. Khuller, B. Raghavachari and N. Young, "Approximating the minimum equivalent digraph," *SIAM J. Computing* **24** (1995), 859-872. Also in Proc. 5th Annual ACM-SIAM Symposium on Discrete Algorithms 1994, 177-186.

[KRY 96] S. Khuller, B. Raghavachari and N. Young, "On strongly connected digraphs with bounded cycle length," *Discrete Applied Mathematics* **69** (1996), 281–289.

[KV 94] S. Khuller and U. Vishkin, "Biconnectivity approximations and graph carvings," *Journal of the ACM* **41** (1994), 214–235. Preliminary version in: Proc. 24th Annual ACM STOC 1992, 759–770.

[L 93] L. Lovász, *Combinatorial Problems and Exercises*, North-Holland, Amsterdam, 1993.

[LP 86] L. Lovász and M. D. Plummer, *Matching Theory*, North-Holland, Amsterdam, 1986.

[Ma 71] W. Mader, "Minimale n-fach kantenzusammenhängende Graphen," *Math. Ann.* **191** (1971), 21–28.

[Ma 72] W. Mader, "Ecken vom Grad n in minimalen n-fach zusammenhängenden Graphen," *Archive der Mathematik* **23** (1972), 219–224.

[Ma 85] W. Mader, "Minimal n-fach zusammenhängenden Digraphen," *J. Combinatorial Theory Series B* **38** (1985), 102–117.

[NI 92] H. Nagamochi and T. Ibaraki, "A linear-time algorithm for finding a sparse k-connected spanning subgraph of a k-connected graph," *Algorithmica* **7** (1992), 583–596.

[Th 89] R. Thurimella, *Techniques for the Design of Parallel Graph Algorithms*, Ph.D. Thesis, The University of Texas at Austin, August 1989.