A bad example for the iterative rounding method for mincost k-connected spanning subgraphs

Ashkan Aazami[a], Joseph Cheriyan[b,*], Bundit Laekhanukit[c]

[a]Dept. of Comb. & Opt., U. Waterloo, Waterloo ON Canada N2L 3G1.
[b]Dept. of Comb. & Opt., U. Waterloo, Waterloo ON Canada N2L 3G1.
[c]Dept. of Comb. & Opt., U. Waterloo, Waterloo ON Canada N2L 3G1.

Abstract

Jain's iterative rounding method and its variants give the best approximation guarantees known for many problems in the area of network design. The method has been applied to the mincost k-connected spanning subgraph problem, but the approximation guarantees given by the method are weak. We construct a family of examples such that the standard LP relaxation has an extreme point solution with infinity norm $\leq \Omega(1)/\sqrt{k}$, thus showing that the standard iterative rounding method cannot achieve an approximation guarantee better than $\Omega(\sqrt{k})$.

Keywords: approximation algorithms, LP relaxations, iterative rounding method, graph connectivity, k-connected spanning subgraphs, extreme point solutions,
2000 MSC: 90C27,
2000 MSC: 68W25,
2000 MSC: 05C40,

1. Introduction

The topic of network design occupies a central place in Combinatorial Optimization and Theoretical Computer Science. A major theme within this topic focuses on the algorithmic problem of computing minimum-cost subgraphs: given a graph $G = (V, E)$ together with costs on the edges, find a subgraph H of G that has minimum cost and satisfies some specified connectivity requirements. An example is the well-known minimum-cost spanning tree problem. Most of these problems are NP-hard, implying that optimal solutions cannot be computed in polynomial time, modulo the P$\neq$NP conjecture. Recent research has focused on the design and analysis of approximation algorithms for these problems. Rather than computing an optimal solution, the goal changes to finding a sub-optimal solution whose cost is guaranteed to be within a known factor of the optimal cost, see [14, 15]. Jain, see [10] and also see the books [14, 15], introduced and analysed an algorithmic paradigm for solving such problems called the *iterative rounding* method.

Jain's iterative rounding method [10] works as follows. Formulate the problem as a covering integer program whose right-hand side is given by a so-called requirement function f. Then solve the LP (linear programming) relaxation to find a basic (extreme point) optimum solution x. Pick an edge e^* of highest value and add it to the solution subgraph H (initially, $E(H)$ is empty). Then update the LP and the integer program, since the variable x_{e^*} is implicitly fixed at value 1. The resulting LP is the same as the LP for the "reduced" problem where the edge e^* is pre-selected for H. Under appropriate conditions on the requirement function f, the problem turns out to be self reducible, i.e., the essential properties of the original problem

*Corresponding author
Email addresses: aazami@gmail.com (Ashkan Aazami), jcheriyan@uwaterloo.ca (Joseph Cheriyan), blaekhanukit@uwaterloo.ca (Bundit Laekhanukit)
URL: http://sites.google.com/site/aazami (Ashkan Aazami), http://www.math.uwaterloo.ca/~jcheriya (Joseph Cheriyan), http://www.math.uwaterloo.ca/~blaekhan (Bundit Laekhanukit)

are preserved in the reduced problem. Iteratively solve the reduced problem. Jain [10] applied this method to the *survivable network design problem* (SNDP), and proved that it achieves an approximation guarantee of 2 provided that the requirement function f is weakly supermodular.

In the *mincost k-connected spanning subgraph* problem, we are given a graph G together with a non-negative cost function on the edges of G. The goal is to pick a set of edges with minimum cost such that the picked edges and their incident nodes form a k-connected spanning subgraph of G. (A graph is called *k-connected* if it has at least $k+1$ nodes, and the deletion of any set of $k-1$ nodes leaves a connected graph. Alternatively, for any two nodes u and v, there exist k openly disjoint paths from u to v.) The problem has attracted research interest in the area of approximation algorithms for many years. A series of improved approximation guarantees have been obtained, see [1, 11, 4, 12]. Besides the iterative rounding method, some other algorithmic paradigms have been applied to the problem, such as, the primal-dual method, see [13, 12], and methods for covering deficient sets, see [11, 4].

Frank and Jordan [8] gave a setpair formulation for some problems in network design. A setpair (W_t, W_h) consists of two disjoint node sets W_t and W_h. Each setpair (W_t, W_h) is assigned a non-negative, integer requirement $f(W_t, W_h)$. The goal is to find a minimum-cost subgraph H that satisfies the requirement of every setpair, i.e., for each setpair (W_t, W_h), H should have at least $f(W_t, W_h)$ edges that have one end-node in W_t and the other end-node in W_h. The mincost k-connected spanning subgraph problem can be formulated via setpairs. We focus on a standard integer programming formulation and its LP relaxation, see below. Our main contribution is a family of examples such that the standard LP relaxation has an extreme-point solution with infinity norm $\leq \Omega(1)/\sqrt{k}$. The existence of such examples has been claimed in [1], but no justification has been presented till now.

The analysis of the standard iterative rounding method is based on proving a guarantee for *every* iteration; in more detail, the approximation guarantee of the method is given by a lower bound on the infinity norm of an extreme-point solution that holds for every iteration. Our example applies in the first iteration of the method, and shows that the approximation guarantee implied by that iteration is $\geq \Omega(\sqrt{k})$. Thus, in a formal sense, our example shows that the standard iterative rounding method cannot achieve an approximation guarantee better than $\Omega(\sqrt{k})$. This raises the intriguing possibility of modifying the iterative rounding method to circumvent the obstacles presented by our example.

The overall construction is somewhat complicated. Rather than giving a compact presentation that may be hard to decipher, we give a longer presentation with some redundancy that may be easier to comprehend. The construction starts with a so-called "base graph" $G_0 = (V, E_0)$ that is k-connected. There is a set of augmenting edges F'' used to increase the connectivity to $k+1$. The LP solution x'' in our construction assigns a fractional value $\leq \Omega(1)/\sqrt{k}$ to each edge in F''. We describe the set F'' in two stages, first without so-called breakpoints, and then we give the actual set F'' using breakpoints.

Our notation is as follows. Let $G = (V, E)$ be a graph. A setpair is an ordered pair of node sets $W = (W_t, W_h)$, where $W_t \subseteq V$ is called the *tail* and $W_h \subseteq V$ is called the *head*. We denote the set of all setpairs by $\mathcal{S}$. Let $e \in E$ be an edge in G. We say that e and W are *incident* (or W is incident to e) if e has one end node in W_t and has its other end node in W_h. We also say that e covers W. For a setpair W, we denote the set of edges covering W by $\delta(W)$. Given a real-valued vector on the edges, $x \in \Re^E$, and a subset F of E, we use $x(F)$ to denote $\sum_{e \in F} x_e$. In particular, $x(\delta(W))$ denotes $\sum\{x_e \mid e \in \delta(W)\}$.

Let U be a subset of nodes. We define the *neighborhood* of U as $\Gamma(U) = \{v \notin U : uv \in E, u \in U\}$. For a subset U of nodes, let $\xi(U) = V \setminus (U \cup \Gamma(U))$ denote the *node-complement* of U. We define two paths to be *openly disjoint* if any node common to both paths is an end node of both paths. We relax some of the standard notation for the sake of readability. Thus, if H denotes a subgraph, then we may also use H to denote the node set $V(H)$ of the subgraph.

We formulate the mincost k-connected spanning subgraph problem as a covering integer program by means of an integer-valued requirement function f on the setpairs. For a setpair W, we define $f(W) = k - (|V| - |W_t \cup W_h|)$; informally speaking, $f(W)$ gives the deficiency of W, the minimum number of edges from W_t to W_h required by k openly disjoint paths from a node in W_t to a node in W_h. We focus on the linear programming relaxation of the integer program. The LP has a variable x_e for each edge $e \in E$. As usual, x indicates the set of picked edges in a solution, i.e., the incidence vector χ_F of a set of picked edges

$F \subseteq E$ gives a solution $x = \chi_F$ of the LP as well as the integer program provided that F and its incident nodes form a k-connected spanning subgraph of G. (In more detail, for $F \subseteq E$ we use χ_F to denote the vector in $\{0,1\}^E$ that has $\chi_F(e) = 1$ iff the edge e is in F.)

$$\begin{aligned} (\texttt{Setpair-LP}) \quad & \min \sum_{e \in E} c_e x_e & \\ \text{subject to} \quad & x(\delta(W)) \geq f(W) & \forall W \in \mathcal{S} \\ & 1 \geq x_e \geq 0 & \forall e \in E \end{aligned}$$

Given a feasible solution x to the above LP, a setpair W is called *tight* if $x(\delta(W)) = f(W)$.

The iterative rounding method, in fact, focuses on a family of covering integer programs and their LP relaxations, namely, all of the covering integer programs obtained by fixing some of the variables at the value 1 or 0. Thus, the relevant family of LP relaxations is obtained from (`Setpair-LP`) by fixing the values of some of the variables $x_e, e \in E$ at one or zero.

Our main result is the following:

Theorem 1.1. *Let $k = 4p(p-1)$, where $p \geq 2$ is an integer. Then there exists a graph $G = (V, E)$ and an LP (Π) in the family of LPs obtained from (`SetPair-LP`) (by fixing some variables to have values of one or zero) such that (Π) has an extreme point x such that $\max_{e \in E} x_e \leq \frac{1}{p} = \frac{\Theta(1)}{\sqrt{k}}$.*

2. Extreme point example

2.1. Our construction: the base graph

Our construction is based on a construction by Ravi and Williamson [13]. See Figure 1 for an illustration.

First, we construct an (undirected) graph $G_0 = (V, E_0)$ that we call the base graph. Let $p \geq 2$ be an integer. The base graph is constructed such that its connectivity is $k = 4p(p-1)$. We use L to denote $k/2$, thus $L = 2p(p-1)$. Table 1 summarizes the parameters for our construction.

k	connectivity of $G_0 = (V, E_0)$	$k = 2L = 4p(p-1)$
L	# of setpairs per column	$L = 2p(p-1)$
p	# of columns	
	# of breakpoints per column	$p - 1$

Table 1: List of parameters

- Our construction uses p copies of a particular subgraph that we call a *column*; the columns are indexed $1, 2, \ldots, p$. Moreover, there is a clique of size $\geq 2L + 2 = k + 2$ that we call the *central clique*, denoted C_*; it is disjoint from the columns but there are some edges between every column and the central clique.

- A column in our construction consists of $L+1$ disjoint cliques with sizes $1, 2L, 2L-1, 2L-2, \ldots, L+1$. Consider the column indexed by j. We denote the $L+1$ cliques in it by $C_{i,j}, i = 0, 1, \ldots, L$, where $C_{0,j}$ consists of a single node, and for $i = 1, \ldots, L$, $C_{i,j}$ is a clique on $2L + 1 - i$ nodes. Moreover, the edge set of a complete bipartite graph is placed between every pair of consecutive cliques $C_{i,j}$ and $C_{i+1,j}$ for $i = 0, \ldots, L-1$, and also between $C_{L,j}$ and C_*. We use v_j to denote the single node in the clique $C_{0,j}, \forall j = 1, \ldots, p$ and we call these nodes the *v-nodes*. These nodes have a special purpose in the construction. Also, for each clique $C_{i,j}$, we pick an arbitrary node and call it the *designated node* of $C_{i,j}$. We also pick a node of C_* and call it the designated node of the central clique; we denote this node by r^*.

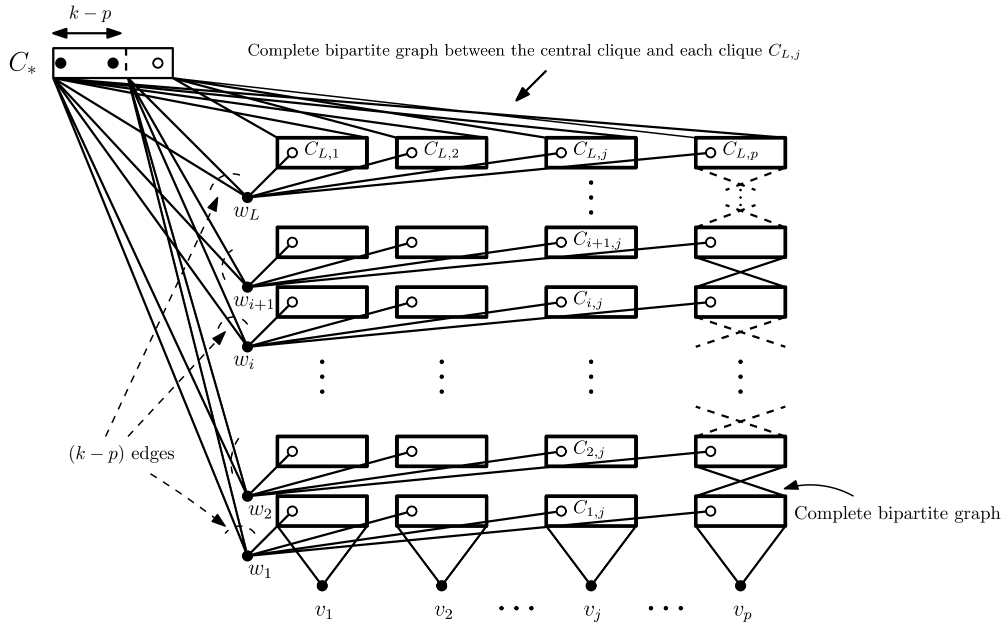

Figure 1: The base graph $G_0 = (V, E_0)$.

- Finally, we add L nodes that we denote by $w_1, w_2, \ldots, w_L$; these nodes are disjoint from the nodes in the columns and C_*. We call these nodes the *w-nodes*, and these nodes have a special purpose in the construction. For each $i = 1, \ldots, L$, the node w_i is incident to k edges; p of these edges have their other ends at the designated nodes of the cliques $C_{i,j}$ for $j = 1, \ldots, p$; the remaining $k-p$ edges incident to w_i have their other ends at distinct nodes of C_* excluding r^*. We assume that each w-node is adjacent to the *same* set of $k-p$ nodes of C_*. Note that r^* is not adjacent to any w-node.

  Thus for a fixed i, each node w_i is adjacent to the set of p designated nodes of $C_{i,1}, \ldots, C_{i,p}$. For a fixed $i = 1, \ldots, L$, we take the i-th *row* of the base graph to be the subgraph induced by the node w_i and the cliques $C_{i,j}$ for $j = 1, \ldots, p$.

By an *interior node* we mean a node that is not a v-node, or a w-node, or a node of C_*. In other words, an interior node is a node of $\bigcup_{i=1}^{L} \bigcup_{j=1}^{p} C_{i,j}$. This completes the construction of the base graph.

Claim 2.1. *The graph $G_0 = (V, E_0)$ is k-connected.*

Our proof is given in the Appendix.

2.2. Augmenting edges, preliminary version

Our construction uses a set of augmenting edges. These are the edges that may be added to the base graph $G_0 = (V, E_0)$ to make it $(k+1)$-node connected. We describe the set of augmenting edges in two stages. This subsection has a preliminary description. In the next subsection, we apply some modifications to get the actual set of augmenting edges for our construction.

In this subsection, we use F' to denote the set of augmenting edges. We have two types of edges in F'. The edges of the first type are called *long* edges and they are useful for covering the v-nodes. The edges of the second type are called *short* edges, and they are useful for covering the w-nodes.

- **Long edges:** For each $j = 1, \ldots, p$ and node v_j, we add $(p-1)$ edges from v_j to the designated node r^* of C_*.

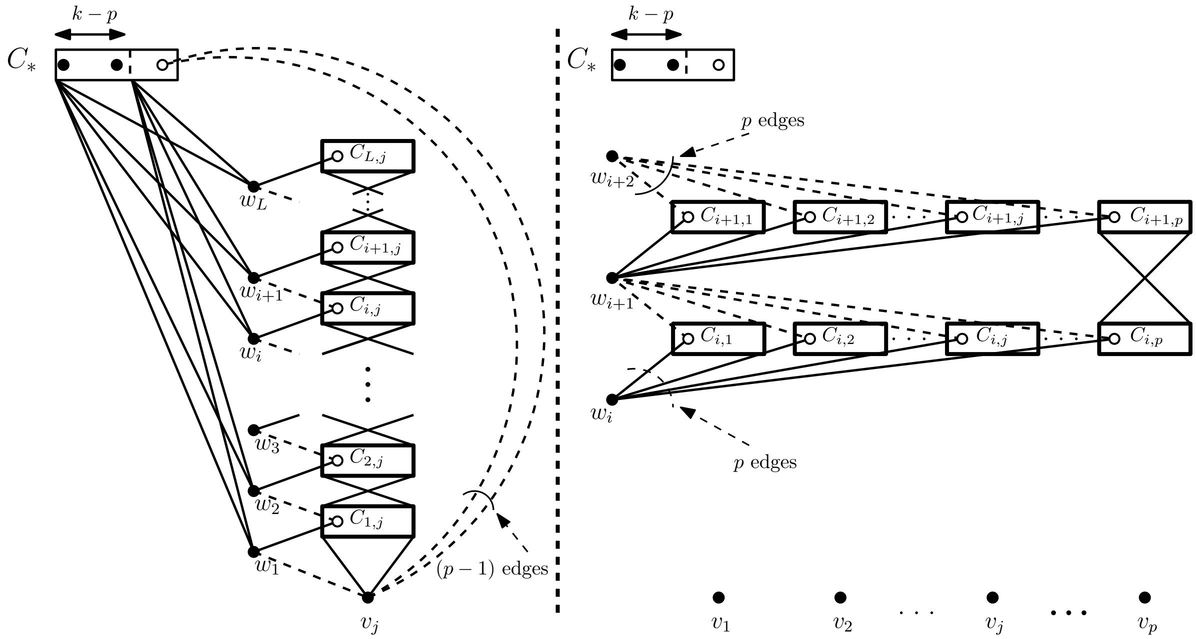

Figure 2: The augmenting edge set, preliminary version. The figure on the left shows the edges incident to the jth column, and the figure on the right shows the edges incident to two consecutive rows.

- **Short edges:** For each $i = 0, \ldots, L-1$, we add an edge between $w_{(i+1)}$ and the designated node of each of the cliques $C_{i,1}, C_{i,2}, \ldots, C_{i,p}$; thus there are p edges between $w_{(i+1)}$ and the p cliques on the i-th row of the base graph. Note that $C_{0,j} = \{v_j\}$ $(j = 1, \ldots, p)$, so there is a short edge between the node w_1 and each of the v-nodes.

This completes the description of the edges in F'. See Figure 2 for an illustration.

The following LP is obtained by relaxing an integer programming formulation of the problem of augmenting the connectivity of the base graph to $k+1$ by using the edges in F'. The LP has a variable x_e for each edge $e \in F'$, thus $x \in \Re^{F'}$. There are no LP variables for the edges of the base graph $G_0 = (V, E_0)$. Informally speaking, each edge in E_0 contributes a value of 1 to the constraints of (`Setpair-LP`); formally, we have the constraints: $x(\delta(W)) + \chi_{E_0}(\delta(W)) \geq 1 + f(W), \forall W \in \mathcal{S}$, where the two terms on the left-hand side account for the contribution of the edges in F' and E_0, respectively, and the right-hand side gives the requirement for $(k+1)$-connectivity for the augmented graph.

$$
\begin{array}{lll}
(\texttt{Augmenting-LP}) & \min \sum_{e \in F'} c_e x_e & \\
\text{subject to} & x(\delta(W)) \geq 1 + f(W) - \chi_{E_0}(\delta(W)) & \forall W \in \mathcal{S} \\
& x_e \geq 0 & \forall e \in F'
\end{array}
$$

Observe that (`Augmenting-LP`) has been obtained from (`Setpair-LP`) by fixing the values of some of the variables $x_e, e \in E$ at one or zero; in other words, (`Augmenting-LP`) belongs to the family of LP relaxations addressed by Theorem 1.1; hence, we can prove Theorem 1.1 by showing that (`Augmenting-LP`) has an extreme-point solution x such that $\max_{e \in E} x_e \leq \frac{\Theta(1)}{\sqrt{k}}$. We show this via Claims 2.3–2.6 given below.

The next result is not used in the proof of Theorem 1.1, but we build upon its proof in order to prove Claim 2.3, and our proof of Theorem 1.1 uses the latter result. The next result shows that we need a capacity of only $\frac{1}{p}$ on each edge in F' for augmenting the connectivity by 1.

Claim 2.2. *Consider the graph $G' = (V, E_0 \cup F')$. The vector given by $x'(e) = \frac{1}{p}, \quad \forall e \in F'$ is a solution for (`Augmenting-LP`).*

The proof is given in the Appendix.

2.3. Augmenting edges, final version

To prove Theorem 1.1, we need to construct a family of tight setpairs and a set of augmenting edges F'' such that the incidence matrix (of these setpairs and edges) has full row rank. Clearly, the set of edges F' given in the previous subsection is not appropriate for this purpose, because the long edges have multiplicity $(p-1)$ in F', hence, the incidence matrix cannot have full rank since the columns of the long edges are replicated with multiplicity $(p-1)$. In this subsection, we describe how to modify the set F' to get a set F'' that is appropriate for Theorem 1.1.

For each $j = 1, \dots, p$ and node v_j, we replace each of the $(p-1)$ long edges incident to v_j by a pair of long edges. We also modify some of the short edges.

Consider any column $j = 1, 2, \dots, p$. Our construction uses $(p-1)$ special row indices that we call the *breakpoints* of the column. The breakpoints of the jth column are given by the $(p-1)$ consecutive odd row indices $\ell = 2(p-1)(j-1)+1, 2(p-1)(j-1)+3, \dots, 2(p-1)(j-1)+2p-3$. For each of these breakpoints ℓ, we add a pair of long edges.

Pair of long edges for breakpoint ℓ:

1. The first long edge (of the pair) is between v_j and a non-designated node of the clique $C_{\ell+2,j}$. There is one exceptional case for the last breakpoint of the last column, and we discuss it below.
2. The second long edge (of the pair) is between a non-designated node of the clique $C_{\ell,j}$ and r^*.

Short edges: Consider a column j, $j = 1, \dots, p$. For $i = 0, \dots, L-1$, if i is a breakpoint of column j, then we place a short edge between w_{i+1} and r^*. Otherwise, if i is not a breakpoint of column j, then we place a short edge between w_{i+1} and a node of $C_{i,j}$. The choice of the node of $C_{i,j}$ depends on i, and we have two cases: if $(i-1)$ is a breakpoint (i.e., if $2(p-1)(j-1)+2 \le i \le 2(p-1)(j-1)+2p-2$) then we place a short edge between w_{i+1} and a non-designated node of $C_{i,j}$; otherwise, if $(i-1)$ is not a breakpoint then we place a short edge between w_{i+1} and the designated node of $C_{i,j}$. If both i and $i-1$ are not breakpoints, then observe that either $i < 2(p-1)(j-1)+1$ or $i \ge 2(p-1)(j-1)+2p-1$.

As mentioned above, one exceptional case comes up in the definition of the pairs of long edges. This is the case for the last breakpoint of the last column, thus we have $j = p$ and breakpoint $\ell = 2(p-1)p-1 = L-1$. Then the first long edge for this pair appears to be ill-defined, since the edge is between v_p and a node of $C_{\ell+2,p} = C_{L+1,p}$, but there is no such clique in the base graph. To avoid this difficulty, we take r^* to be the end node of the first long edge. (Informally speaking, we are taking C_* to be $C_{L+1,p}$.) Moreover, this exception needs special handling in the proof of Claim 2.3.

This completes the description of the edges in F''. See Figure 3 for an illustration.

Claim 2.3. *Consider the graph $G'' = (V, E_0 \cup F'')$. The vector given by $x''(e) = \frac{1}{p}, \quad \forall e \in F''$ is a solution for (`Augmenting-LP`).*

The Appendix has two proofs of this claim. The first proof uses general methods and is longer; the second proof was suggested by a referee, and it is shorter, but it relies on special properties of the base graph.

2.4. Rank of the incidence matrix

In this subsection, we first define a family of tight setpairs for the feasible solution x'' defined in Claim 2.3. Next, we show that the incidence matrix of these tight setpairs and the edges in F'' has full rank. This shows that x'' is an extreme-point solution of (`Augmenting-LP`), and thus completes the proof of Theorem 1.1.

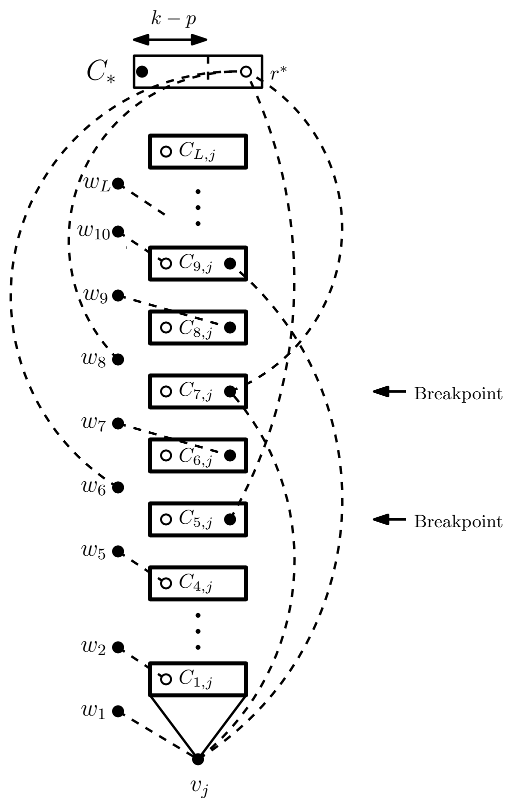

Figure 3: The augmenting edge set, final version. The figure shows the augmenting edges incident to an arbitrary column.

A family of tight setpairs:
Throughout this section, the notation for node-complement refers to the base graph $G_0 = (V, E_0)$. For example, if $S = \{v_1\} \cup C_{1,1}$, then $\xi(S) = V \setminus (\{v_1\} \cup C_{1,1} \cup C_{2,1} \cup \{w_1\})$. The word "column" may mean either a column of a matrix or a column of the base graph; the context will resolve the ambiguity. Similarly, the word "row" may mean either a row of a matrix or a row of the base graph.

Consider a column $j = 1, \ldots, p$. Recall that each column consists of a single node v_j together with a sequence of L cliques $C_{i,j}$, $i = 1, \ldots, L$. For each $i = 0, \ldots, L-1$, let $S_{i,j} = \bigcup_{\ell=0}^{\ell=i} C_{\ell,j}$; thus, $S_{i,j}$ contains v_j and the nodes of the first i (nontrivial) cliques of the j-th column. Each set $S_{i,j}$ defines a setpair $(S_{i,j}, \xi(S_{i,j}))$; for ease of notation, we may use $S_{i,j}$ to denote this setpair. Thus, we have L setpairs for each column. Also for each node w_i, $i = 1, \ldots, L$, let $W_i = \{w_i\}$. Each set W_i defines a setpair $(W_i, \xi(W_i))$; again, we may use W_i to denote this setpair. Thus, we have pL setpairs of the form $S_{i,j}$, and L setpairs of the form W_i, for a total of $(p+1)L$ setpairs. Let $\mathcal{L}$ denote the set of all these setpairs. See Figure 4 for an illustration.

Claim 2.4. *Let x'' be the feasible solution given in Claim 2.3. Then, the setpairs in $\mathcal{L}$ are tight, that is, for each $W \in \mathcal{L}$, we have $x''(\delta(W)) = 1$.*

Proof: First, consider any setpair $(W_i, \xi(W_i))$, $i = 1, \ldots, L$, where $W_i = \{w_i\}$. Observe that $\Gamma(W_i)$, the neighbourhood of w_i in the base graph, consists of $k-p$ nodes of C_* and the designated nodes from each of the cliques $C_{i,j}$, $j = 1, \ldots, p$. There are p augmenting edges incident to the setpair, namely, the p short edges incident to w_i; thus $|\delta(W_i, \xi(W_i)) \cap F''| = p$; moreover, $x''(\delta(W_i, \xi(W_i))) = 1$ because each augmenting edge e has $x''_e = \frac{1}{p}$.

Next, consider any setpair $(S_{i,j}, \xi(S_{i,j}))$, $i = 0, \ldots, L-1$ and $j = 1, \ldots, p$. Note that $\Gamma(S_{i,j})$, the

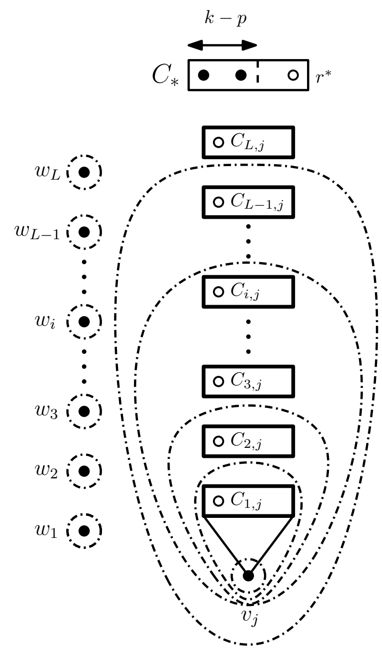

Figure 4: An illustration of the tight setpairs $(W_{i+1}, \xi(W_{i+1}))$ and $(S_{i,j}, \xi(S_{i,j}))$ for $i = 0, \ldots, L-1$ and fixed j. The tail of each of these $2L$ setpairs is indicated by a circle or oval, but the heads are not indicated.

neighbourhood of $S_{i,j}$ in the base graph, consists of k nodes, namely, $\{w_1, \ldots, w_i\} \cup C_{(i+1),j}$. We claim that each such setpair is covered by exactly p augmenting edges, either p long edges, or $p-1$ long edges and one short edge. As above, this claim implies that the setpair is tight.

Consider any column j and the ℓ-th pair of long edges in the column, for $\ell = 1, \ldots, p-1$; let $\beta(\ell)$ denote the associated breakpoint (thus $\beta(\ell) = 2(p-1)(j-1) + 2\ell - 1$). It can be seen that the setpair $(S_{\beta(\ell),j}, \xi(S_{\beta(\ell),j}))$ is covered by both long edges of the pair, and each of the other setpairs $(S_{i,j}, \xi(S_{i,j}))$, $i = 0, \ldots, L-1,\ i \neq \beta(\ell)$, is covered by exactly one of the long edges of the pair. (To verify this, recall that $\Gamma(S_{0,j}) = \Gamma(\{v_j\}) = C_{1,j}$, and $\Gamma(S_{i,j}) = C_{i+1,j} \bigcup \{w_1, \ldots, w_i\}$ for $i = 1, \ldots, L-1$.) Hence, for each $i = 0, \ldots, L-1$, if i is a breakpoint of column j, then the setpair $(S_{i,j}, \xi(S_{i,j}))$ is covered by p long edges, otherwise, the setpair is covered by $p-1$ long edges (one long edge from each of the $p-1$ pairs). Moreover, if i is a breakpoint of column j, then the setpair $(S_{i,j}, \xi(S_{i,j}))$ is covered by none of the short edges, otherwise, the setpair is covered by one short edge, namely, the short edge between $C_{i,j}$ and w_{i+1}. Our claim follows, and thus we have proved that the setpair $(S_{i,j}, \xi(S_{i,j}))$ is tight. □

The next result is not essential, but we include it since it can be used to give another proof of Claim 2.3. Let $d_{i,j}$ denote the designated node of the clique $C_{i,j}, \forall i = 0, \ldots, L,\ j = 1, \ldots, p$, and let $\mathcal{D}$ denote the set $\{d_{i,j} \mid i = 1, \ldots, L,\ j = 1, \ldots, p\}$; note that $\mathcal{D}$ does not contain any of the nodes $v_j = d_{0,j},\ j = 1, \ldots, p$.

Corollary 2.5. *Let x'' be the feasible solution given in Claim 2.3. Then, x'' covers all of the setpairs of the form $S_{i,j} - \mathcal{D}'$, where $\mathcal{D}' \subseteq \mathcal{D}$, that is, $x''(\delta((S_{i,j} - \mathcal{D}', \xi(S_{i,j} - \mathcal{D}')))) \geq 1$ holds, $\forall i = 0, \ldots, L-1,\ j = 1, \ldots, p$.*

Proof: Fix i and j, and consider any set $S = S_{i,j} - \mathcal{D}'$ and the associated setpair $(S, \xi(S))$; we may assume $\mathcal{D}' \subseteq \{d_{1,j}, \ldots, d_{i,j}\}$. First, observe that $\Gamma(S)$ has the same size as $\Gamma(S_{i,j})$, namely, k, because $\Gamma(S)$ can be obtained from $\Gamma(S_{i,j})$ by replacing the node w_ℓ by the node $d_{\ell,j}$, for each designated node in $S_{i,j} - S$. Hence, for each $\ell = 1, \ldots, i$, note that $\Gamma(S)$ has exactly one of the two nodes w_ℓ or $d_{\ell,j}$.

If $S = S_{i,j}$, then Claim 2.4 implies that $x''(\delta(S_{i,j})) = 1$, so the result follows. Otherwise, consider the smallest ℓ such that $w_\ell \notin \Gamma(S)$. If $\ell = 1$, then the short edge between v_j and w_1 together with the long edges incident to the setpair $S_{i,j}$ suffice to cover the setpair $(S, \xi(S))$. If $\ell \geq 2$, then we have $w_{\ell-1} \in \Gamma(S)$, and $d_{(\ell-1),j} \in S$, hence, the short edges between $d_{(\ell-1),j}$ and w_ℓ together with the long edges incident to the setpair $S_{i,j}$ suffice to cover the setpair $(S, \xi(S))$. The result follows. □

The incidence matrix:
Consider the incidence matrix B of setpairs in $\mathcal{L}$ and edges in F''. The rows of B are labeled by the setpairs of $\mathcal{L}$, and the columns are labeled by the edges of F''. Let $W \in \mathcal{L}$ and $e \in F''$. The entry of B corresponding to the pair (W, e) is given by:

$$B(W,e) = \begin{cases} 1 & \text{if } e \in \delta(W), \\ 0 & \text{otherwise} \end{cases}$$

The matrix B has dimension $(p+1)L \times (p+1)L$. To see this, consider the setpairs in $\mathcal{L}$ first. As noted above, we have pL setpairs of the form $S_{i,j}$ and L setpairs of the form W_i. The rows of B are partitioned into two parts; in the illustration of B (see below), the first part is above the double horizontal line, and the second part is below that line. The first part corresponds to the setpairs of the form $S_{i,j}$, and the second part corresponds to the setpairs of the form W_i. Now consider the augmenting edges. The base graph has p columns, and there are $2(p-1)$ long edges in each column; moreover, there are p short edges per w-node. In total, there are $2(p-1)p + pL = L + pL = (p+1)L$ augmenting edges. The columns of B are also partitioned into two parts; in the illustration of B, the first part is to the left of the double vertical line, and the second part is to the right of that line. The first part corresponds to the $L = 2(p-1)p$ long edges, and the second part corresponds to the pL short edges.

We show below that the rows and columns of the matrix B can be indexed such that B has the following structure:

$$B = \left(\begin{array}{c|c|c||c|c|c} Q_1 & 0 & 0 & \widehat{I}_1 & 0 & 0 \\ 0 & \ddots & 0 & 0 & \ddots & 0 \\ 0 & 0 & Q_p & 0 & 0 & \widehat{I}_p \\ \hline\hline 0 & \dots & 0 & I_{L\times L} & \dots & I_{L\times L} \end{array}\right),$$

where $I_{L\times L}$ denotes the $L \times L$ identity matrix, and 0 denotes a matrix of zeros of the appropriate dimension. The matrices Q_j and $\widehat{I}_j$ are described below.

In order to describe the indexing of the rows and columns of B, it is convenient to partition the rows and columns into a few blocks, and then describe the indexing with respect to the blocks.

The rows of the first part of B (above the double horizontal line) are partitioned into p blocks corresponding to the p columns of the base graph. The jth block, for $j = 1, \ldots, p$, has L rows for the L setpairs $S_{0,j}, S_{1,j}, \ldots, S_{L-1,j}$, and these L rows are indexed in the natural order, i.e., the ith row for the setpair $S_{i-1,j}$. The rows of the second part of B (below the double horizontal line) form one block that has L rows for the L setpairs $W_1, W_2, \ldots, W_L$; these L rows are indexed in the natural order, i.e., the ith row for the setpair W_i.

The matrix-columns of the first part of B (to the left of the double vertical line) are partitioned into p blocks corresponding to the p columns of the base graph. The jth block, for $j = 1, \ldots, p$, has $2(p-1)$ matrix-columns for the $(p-1)$ pairs of long edges incident to the jth column of the base graph, and these $2(p-1)$ matrix-columns are indexed in the natural order, i.e., the $(2\ell-1)$th and 2ℓth matrix-columns are for the first and second long edge of the pair of long edges corresponding to the ℓth breakpoint, $\ell = 1, \ldots, p-1$. The matrix-columns of the second part of B (to the right of the double vertical line) are partitioned into p blocks corresponding to the p columns of the base graph. The jth block, for $j = 1, \ldots, p$, has L matrix-columns for the L short edges associated with the jth column of the base graph, and these L matrix-columns are indexed in the natural order, i.e., the ith matrix-column for the short edge that is associated with the jth column of the base graph and is incident to w_i. Recall from Section 2.3 that for each $i = 0, \ldots, L-1$, there is a short edge associated with the jth column of the base graph. If i is not a breakpoint, then the associated short edge has one end in $S_{i,j}$ and the other end in W_{i+1}. If i is a breakpoint, then the associated short edge has one end at r^* and the other end in W_{i+1}. Thus the matrix-column for a short edge either has two nonzeros (a 1 for the incident setpair of the form W_i and another 1 for the incident setpair of the form $S_{i-1,j}$) or has one nonzero (a short edge associated with a breakpoint has a single 1 for the incident setpair of the form W_i).

Let e_i denote the i-th column of the $L \times L$ identity matrix $I_{L\times L}$. Let f_i denote a column vector of size L with a 1 in entries $1, \ldots, i$ and a 0 in entries $i+1, \ldots, L$. A column vector of size L with a 1 in all entries is denoted by $1_L = f_L$. Recall from the previous subsection that the breakpoints of the j-th column of the base graph are given by the indices $2(p-1)(j-1) + 2\ell - 1$, for $\ell = 1, \ldots, (p-1)$. Whenever we refer to breakpoint indices with respect to B or a submatrix of B, then we add an offset of 1 to these indices to account for the fact that the indices of the rows of a matrix start with 1 whereas the indices of the rows of the base graph start with 0.

Consider the j-th column of the base graph, for any $j = 1, \ldots, p$. This column is associated with the matrix Q_j, and the matrix has L rows corresponding to the L setpairs of the form $S_{i,j}$, and it has $2(p-1)$ columns corresponding to the $(p-1)$ pairs of long edges. Let h denote $2(p-1)(j-1)+2$; note that h is determined by j. For each $\ell = 1, \ldots, (p-1)$, recall that there is a pair of long edges corresponding to the ℓ-th breakpoint; the ℓ-th breakpoint has the index $2(p-1)(j-1) + 2\ell - 1 = h + 2\ell - 3$. The first long edge (of the ℓ-th pair of long edges) has end nodes at v_j and $C_{h+2\ell-1,j}$, and covers the setpairs $S_{0,j}, \ldots, S_{h+2\ell-3,j}$. The second long edge (of the pair) has end nodes at r^* and $C_{h+2\ell-3,j}$, and covers the setpairs $S_{h+2\ell-3,j}, \ldots, S_{L-1,j}$. Hence, for each $\ell = 1, \ldots, (p-1)$, the columns $2\ell - 1$ and 2ℓ of Q_j are given by the vectors $e_1 + e_2 + \ldots + e_{h+2\ell-2}$ and $e_{h+2\ell-2} + e_{h+2\ell-1} + \ldots + e_L$, respectively. See Figure 5 for an illustration. Thus, we have

$$Q_j = \left(f_h, 1_L - f_{h-1}, f_{h+2}, 1_L - f_{h+1}, \ldots, f_{h+2(p-2)}, 1_L - f_{h+2(p-2)-1}\right).$$

The matrix $\widehat{I_j}$ has L rows and L columns; the rows of $\widehat{I_j}$ correspond to the rows of Q_j and both sets of rows correspond to the setpairs $S_{i,j}, i = 0, \ldots, L-1$; each column of $\widehat{I_j}$ corresponds to a short edge incident to a w-node. Recall that a short edge connects w_{i+1} and a node of $C_{i,j}$ provided i is *not* a breakpoint of column j (of the base graph), for each $i = 0, 1, \ldots, L-1$. The matrix $\widehat{I_j}$ is a diagonal matrix whose (i,i)-entry is 0 if i is the index of a breakpoint, and the entry is 1 otherwise. See Figure 5 for an illustration.

The following claim completes the proof of Theorem 1.1.

Claim 2.6. *The matrix B has full rank.*

The proof is given in the Appendix.

Recall from Section 2.3 that an exceptional case arises for the last column p of the base graph, and the pair of long edges associated with the last breakpoint of p, namely, $L-1$. No special handling is needed for this exceptional case either in the definition of the matrix B or in the proof of Claim 2.6.

3. Conclusions

We constructed a family of examples of the mincost k-connected spanning subgraph problem such that the standard LP relaxation has an extreme point solution with infinity norm $\leq \Omega(1)/\sqrt{k}$. The number of nodes in our construction is $\Theta(p^5) = \Theta(k^{2.5})$. The example applies for the special case of the problem where a k-connected spanning subgraph of the input graph is given, and the goal is to find a mincost set of edges whose addition increases the connectivity to $k+1$. The family of tight setpairs used in our proof has the following property: if we take the smaller of the head and the tail for each setpair, then we get a laminar family of sets.

Over the past decade, the iterative rounding method and its variants have been used to achieve many remarkable results in areas such as network design for edge connectivity requirements. But these achievements have not been extended to other areas, even closely related ones such as network design for node connectivity requirements. Our main result gives some explanation for this lack of success for the iterative rounding method, but one can hope for the discovery of new algorithmic paradigms that will surmount the obstacles.

$$
Q_2 = \left(\begin{array}{cc|cc}
1 & & 1 & \\
1 & & 1 & \\
1 & & 1 & \\
1 & & 1 & \\
1 & & 1 & \\
1 & 1 & 1 & \\
 & 1 & 1 & \\
 & 1 & 1 & 1 \\
 & 1 & & 1 \\
 & 1 & & 1 \\
 & 1 & & 1 \\
 & 1 & & 1
\end{array}\right)
\qquad
\widehat{I}_2 = \begin{pmatrix}
1 & & & & & & & & & & & \\
 & 1 & & & & & & & & & & \\
 & & 1 & & & & & & & & & \\
 & & & 1 & & & & & & & & \\
 & & & & 1 & & & & & & & \\
 & & & & & 0 & & & & & & \\
 & & & & & & 1 & & & & & \\
 & & & & & & & 0 & & & & \\
 & & & & & & & & 1 & & & \\
 & & & & & & & & & 1 & & \\
 & & & & & & & & & & 1 & \\
 & & & & & & & & & & & 1
\end{pmatrix}
$$

Figure 5: Let $p = 3$ and $L = 2(p-1)p = 12$. The matrices Q_2 and $\widehat{I}_2$ corresponding to the second column in the construction are illustrated; the entries indicated by blanks are all zero; Q_2 is an $L \times 2(p-1)$ matrix, and $\widehat{I}_2$ is an $L \times L$ matrix. The breakpoints of the second column have indices $2(p-1)(2-1) + 2\ell$, for $\ell = 1, \ldots, (p-1)$, namely, $6, 8$.

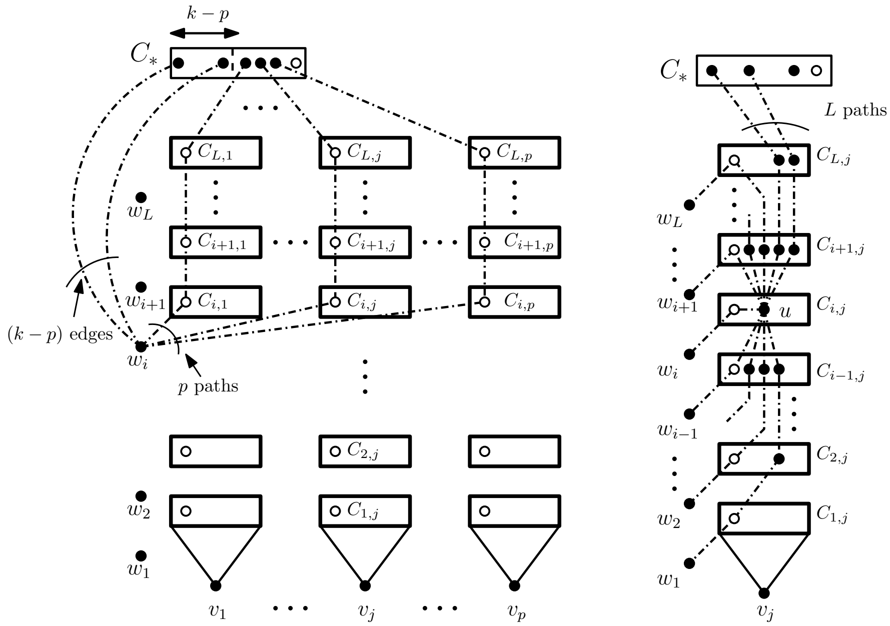

(a) k-fan from a w-node. (b) k-fan from interior node u.

Figure A.6: Illustration of k-fans for w-nodes and for interior nodes.

Appendix A. Proofs of claims

This Appendix has the proofs of the claims from Section 2.

Appendix A.1. Proof of Claim 2.1

Claim 2.1 *The graph $G_0 = (V, E_0)$ is k-connected.*

Proof: We prove that G_0 is k-connected by using a standard k-connectivity algorithm, see [2, 3]: First, pick k nodes and prove that there are k openly disjoint paths between every two of them. We say that these k nodes have been successfully "scanned." Then we scan the remaining nodes in some order; if all the nodes can be successfully scanned, then G_0 is k-connected. To scan one of the remaining nodes z, we construct a k-fan from z to the set of nodes that have been already scanned. A k-fan from a node z to a set of nodes $U, z \notin U$ means a set of k openly disjoint paths, where each path starts at z and ends at a node of U, and moreover, z is the only node that occurs in two or more of these paths. See Even [2] and Even and Tarjan [3], for a proof of correctness and further details.

For the sake of convenience, we allow some informality in the following discussion. In particular, when we say that a path P is in column j, we mean that all the interior nodes of P are in column j; thus, one or both end-nodes of P may not be in the column.

See Figure A.6 for an illustration.

To pick the initial set of k nodes, we exclude r^* from the central clique C_*, and pick any k of the remaining nodes; let $r_1, \ldots, r_k$ denote the picked nodes. Since C_* has order $\geq 2L + 2$, it contains $\geq k = 2L$ openly disjoint paths between any two of its nodes. Thus the initial set of k nodes has the required property.

We scan the remaining nodes in the following order:

1. Central clique: Let u be a node of C_*, where $u \notin \{r_1, \ldots, r_k\}$. Clearly, the central clique contains a k-fan from u to $\{r_1, \ldots, r_k\}$.

2. w-nodes: Consider any $i = 1, \ldots, L$ and the node w_i. A k-fan from w_i to $\{r_1, \ldots, r_k\}$ can be constructed by sending $k - p$ paths via the $k - p$ nodes of C_* adjacent to w_i. The remaining p paths of the k-fan are given by paths of length $2 + L - i$ in each of the p columns; the j-th of these paths starts at w_i, then uses one node from each of the cliques $C_{i,j}, C_{(i+1),j}, \ldots, C_{L,j}$.
3. Interior nodes: Consider any clique $C_{i,j}$ and any node u in it. To scan u, we construct two L-fans that have no nodes in common except u. The first L-fan is from u to a set of L nodes of C_*, call it $\{r_1, \ldots, r_L\}$; each path in this L-fan has length $1 + L - i$ and uses exactly one node of $C_{i+h,j}$ for $h = 1, 2, \ldots, L - i$. The second L-fan is from u to the w-nodes $w_1, w_2, \ldots, w_L$; each of these L paths uses one node from each of the cliques "lying between" u and the w-node. (For example, if $i < L$ and $h \in \{1, 2, \ldots, (L - i)\}$, then the path from u to w_{i+h} uses one node of $C_{i,j}, C_{i+1,j}, \ldots, C_{i+h-1,j}$ and the designated node of $C_{i+h,j}$.) Thus, for $i < L$ and each $h \in \{1, 2, \ldots, (L - i)\}$, the number of openly disjoint paths from the two L-fans incident to $C_{i+h,j}$ equals the cardinality of $C_{i+h,j}$, namely, $2L + 1 - i - h$.
   Moreover, we can augment the above k-fan to get a k-fan from u to k nodes of C_*, by adding (the edge set of) a matching between $\{w_1, \ldots, w_L\}$ and L of the neighbours of $\{w_1, \ldots, w_L\}$ in C_*, call them $r'_1, \ldots, r'_L$; such matchings exits in the base graph since it contains a complete bipartite graph on the node sets $\{w_1, \ldots, w_L\}$ and $k - p > L$ nodes of C_*. Thus, by adding the matching edges, we can augment the path from u to w_ℓ in the original k-fan to get a path from u to r'_ℓ in the new k-fan. We get a k-fan from u to $\{r_1, \ldots, r_L, r'_1, \ldots, r'_L\}$, and the latter set is contained in C_*.
4. v-nodes: Consider any $j = 1, \ldots, p$ and the node v_j. Observe that v_j is adjacent to the k nodes of $C_{1,j}$, and each of those nodes is an interior node. We scan v_j by constructing a k-fan to its k neighbours in $C_{1,j}$.

This completes the description of the scanning procedure. We successfully scanned all nodes, hence, G_0 is k-connected. ☐

Appendix A.2. Proof of Claim 2.2

Claim 2.2 *Consider the graph $G' = (V, E_0 \cup F')$. The vector given by $x'(e) = \frac{1}{p}, \quad \forall e \in F'$ is a solution for (`Augmenting-LP`).*

Proof: We prove that G' is fractionally $(k+1)$-connected using an argument similar to the one in the proof of Claim 2.1.

A *fractional $|U|$-fan* from a node z to a set of nodes $U, z \notin U$, means a flow of value $|U|$ between z and U such that the flow transiting via each node has value ≤ 1 and the flow terminating at each node has value ≤ 1. More formally, a fractional $|U|$-fan from z to U is defined in the associated directed graph, where each undirected node v is replaced by a pair of nodes v_{in}, v_{out} with a unit-capacity arc (v_{in}, v_{out}), and each undirected edge uv is replaced by a pair of arcs (u_{out}, v_{in}), (v_{out}, u_{in}) of infinite capacity; the fan refers to a flow of value $|U|$ with a single source z_{out} and a sink at each node $u_{out}, \forall u \in U$; the flow on an arc need not be integral, but note that the value of the flow transiting via any node, namely, the flow on any arc (v_{in}, v_{out}), is ≤ 1; the existence of such a fan certifies that there exists no node cut of cardinality $< |U|$ whose deletion separates z from U (i.e., results in two different connected components such that one contains z and the other contains some node of U).

We take the initial set of $k+1$ nodes to be a set of $k+1$ nodes from C_* that includes r^*; we may use $r_{(k+1)}$ to denote r^* and we denote the initial set of nodes by $r_1, \ldots, r_{(k+1)}$. Since C_* has order $\geq 2L + 2 = k + 2$, it clearly contains $\geq k + 1$ openly disjoint paths between any two of its nodes. Thus the initial set of $k + 1$ nodes has the required property.

We scan the remaining nodes in the following order:

1. Central clique: Let u be a node of C_*, where $u \notin \{r_1, \ldots, r_{(k+1)}\}$. Clearly, the central clique contains a $(k + 1)$-fan from u to $\{r_1, \ldots, r_{(k+1)}\}$.

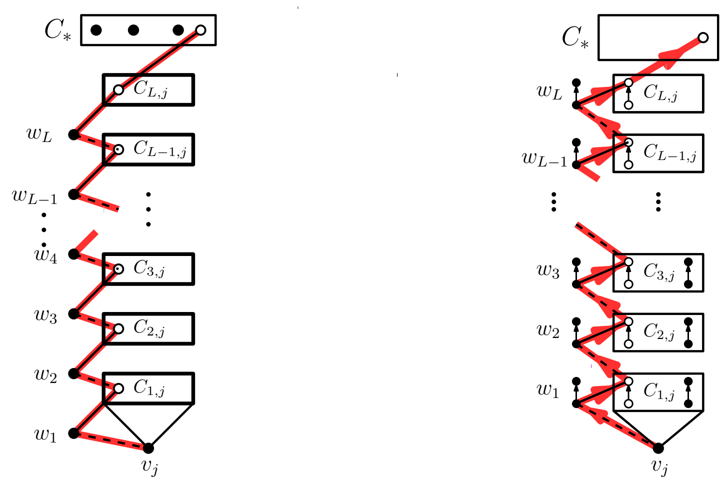

(a) Augmenting path in the undirected graph.

(b) Augmenting path in the associated directed graph.

Figure A.7: Illustration of the last augmenting path from v_j to r^*.

2. w-nodes: Consider any $i = 1, \ldots, L$ and the node w_i. We start with the k-fan from w_i to $\{r_1, \ldots, r_k\}$ constructed in the proof of Claim 2.1. Then we add p paths from w_i to $r^* = r_{(k+1)}$ such that each of these paths has a flow of value $1/p$. We call these the *fractional paths*. Moreover, the fractional paths are openly disjoint (only the end nodes w_i and r^* are common), and each integral path (from w_i to an r_ℓ, $\ell = 1, \ldots, k$) has only the node w_i in common with any other path. Each of the fractional paths has length $3 + L - i$, starts with a short edge from w_i to a node of $C_{(i-1),j}$, and uses one node from each of the cliques $C_{(i-1),j}, C_{i,j}, C_{(i+1),j}, \ldots, C_{L,j}$. Thus each column contains two disjoint subpaths of the fractional $(k+1)$-fan, an integral path and a fractional path.
3. v-nodes: For any $j = 1, \ldots, p$ and the node v_j, we construct a fractional $(k+1)$-fan from v_j. We start by constructing an (integral) k-fan from v_j to $\{w_1, \ldots, w_L\} \bigcup \{r_1, \ldots, r_L\}$ in the base graph $G_0 = (V, E_0)$. This k-fan is similar to the k-fan constructed for the interior nodes in the proof of Claim 2.1. We construct two L-fans that have no nodes in common except v_j; the first L-fan is from v_j to $\{r_1, \ldots, r_L\}$; each path in this L-fan has length $1 + L$ and uses exactly one node of $C_{i,j}$ for $i = 1, 2, \ldots, L$; the second L-fan is from v_j to the nodes $w_1, w_2, \ldots, w_L$; each of these L paths uses one node from each of the cliques "lying between" v_j and the w-node.
   Then we add p paths from v_j to $r^* = r_{(k+1)}$ such that each of these paths has a flow of value $1/p$. Moreover, the fractional paths are openly disjoint (only the end nodes v_j and r^* are common). More formally, we focus on the associated directed graph and construct a flow of value $k+1$ between v_j and $\{w_1, \ldots, w_L\} \bigcup \{r_1, \ldots, r_L, r^*\}$ such that the value of the flow transiting via any node is ≤ 1. We construct the flow of value $k+1$ by starting with the integral flow of value k (given by the k-fan), and then sending a flow of value $1/p$ on p augmenting paths. We take $p-1$ of the augmenting paths to be the paths of length one given by the long edges between v_j and r^*. The last augmenting path P^* has length $2L+1$ and is incident to all of the w-nodes and all of the cliques $C_{1,j}, C_{2,j}, \ldots, C_{L,j}$. In the undirected graph G', the path P^* has the form $v_j, w_1, u_1, w_2, u_2, \ldots, w_{L-1}, u_{L-1}, w_L, u_L, r^*$, where $u_1, \ldots, u_L$ denote the designated nodes of $C_{1,j}, \ldots, C_{L,j}$, respectively (for ease of notation, we use u_i rather than $u_{i,j}$ for the designated node of $C_{i,j}$). For the sake of notational convenience, we denote augmenting paths by their node sequence in the undirected graph, rather than in the associated directed graph.
   See Figure A.7 for an illustration of the last augmenting path in the graph, as well as in the associated directed graph.
4. Interior nodes: Consider any clique $C_{i,j}, i = 1, \ldots, L$, and any node u in it. To scan u, we first construct

a k-fan from u to $\{w_1,\ldots,w_L\}\bigcup\{r_1,\ldots,r_L\}$, in a similar way to the proof of Claim 2.1. Alternatively, we can construct a k-fan from u to k nodes of C_*, see part (3) in the proof of Claim 2.1. Next, we connect u by a path of length i to v_j that is disjoint from the nodes of the above k-fan, except for the node u; this path is easily constructed, since the cliques $C_{1,j}, C_{2,j},\ldots,C_{i-1,j}$ have cardinalities of $2L, 2L-1,\ldots,2L+2-i$ and are incident to $1,2,\ldots,i-1$ openly-disjoint paths of the k-fan. Adding this path to the k-fan gives a $(k+1)$-fan to a set of $k+1$ already scanned nodes. Observe that this $(k+1)$-fan is contained in the base graph $G_0=(V,E_0)$, that is, it does not use any augmenting edge.

□

Appendix A.3. Proof of Claim 2.3

Claim 2.3 *Consider the graph $G''=(V,E_0\cup F'')$. The vector given by $x''(e)=\frac{1}{p},\quad \forall e\in F''$ is a solution for (`Augmenting-LP`).*

Proof: The proof is similar to the proof of Claim 2.2. Here, we only mention the changes needed to complete the proof. We start with $k+1$ nodes from C_*, including the node r^*, and we denote these nodes by $r_1, r_2,\ldots,r_k, r_{(k+1)}=r^*$. The base graph $G_0=(V,E_0)$ has $\geq k+1$ openly disjoint paths between any two of these nodes.

We scan the remaining nodes in the same order as in the proof of Claim 2.2. The scanning procedure is almost the same for all of the nodes except for the v-nodes. The scanning procedure for the v-nodes needs to be modified, because the construction of the fractional $(k+1)$-fans is different; in the previous construction, we used the $(p-1)$ long edges incident to a v-node, but those edges are not present in F''. The modifications are discussed below.

The scanning procedure is the same for the nodes of the central clique, the interior nodes, and all the w-nodes, except for nodes w_i such that $i-1$ is a breakpoint of some column $j=1,\ldots,p$. Let w_i be such an exceptional node. Note that w_i has one short edge to r^*, so a flow of value $\frac{1}{p}$ can be sent directly to r^* via this short edge. The remaining flow of value $\frac{p-1}{p}$ can be sent to r^* via the remaining $(p-1)$ short edges incident to w_i as described in the proof of Claim 2.2.

Consider a node v_j, $j=1,\ldots,p$, and its fractional $(k+1)$-fan to $\{w_1,\ldots,w_L\}\bigcup\{r_1,\ldots,r_L,r^*\}$. We start with an (integral) k-fan from v_j to $\{w_1,\ldots,w_L\}\bigcup\{r_1,\ldots,r_L\}$ as in the proof of Claim 2.2. This k-fan contains an L-fan from v_j to the L nodes $r_1,\ldots,r_L$ of C_*. We impose the following requirement on this L-fan:

> There is a path $\hat{P}$ contained in the L-fan from v_j to C_* such that for each pair among the $p-1$ pairs of long edges associated with column j, $\hat{P}$ contains both interior nodes incident to the pair of long edges. Moreover, $\hat{P}$ contains all the non-designated nodes of the cliques $C_{i,j}, i=1,\ldots,L$ that are incident to short edges. (Recall from Section 2.3 that there is a short edge between a non-designated node of $C_{i,j}$ and w_{i+1} iff i is not a breakpoint and $i-1$ is a breakpoint.)

This requirement is easily satisfied since each path from v_j to C_* in the L-fan can use one arbitrary non-designated node of each of the cliques $C_{i,j}, i=1,\ldots,L$. We assume that $\hat{P}$ is the path of the L-fan from v_j to r_1. Moreover, we may view the k-fan from v_j to $\{w_1,\ldots,w_L\}\bigcup\{r_1,\ldots,r_L\}$ as an integral flow; the path $\hat{P}$ carries one unit of this flow.

We augment the integral flow (of value $k=2L$) by sending a flow of value $\frac{1}{p}$ via each of the $(p-1)$ pairs of long edges. Consider a pair of long edges, and let u' and u'' denote the interior nodes incident to the two long edges; assume that the row index of u' is less than that of u''. We send a flow of value $\frac{1}{p}$ from v_j to u'' via the first long edge, next we push back the same amount of flow from u'' to u' along the path $\hat{P}$, and finally, we send the same amount of flow from u' to r^* via the second long edge. Similarly, we send a flow of value $\frac{1}{p}$ from r^* to v_j using each of the $(p-1)$ pairs of long edges in column j. (An exceptional case is discussed below.)

Finally, we send a flow of value $\frac{1}{p}$ from v_j to r^* via the short edges, similarly to the proof of Claim 2.2. We have to modify the construction in Claim 2.2, because F'' has no edge between $w_{(i+1)}$ and $C_{i,j}$ if i is

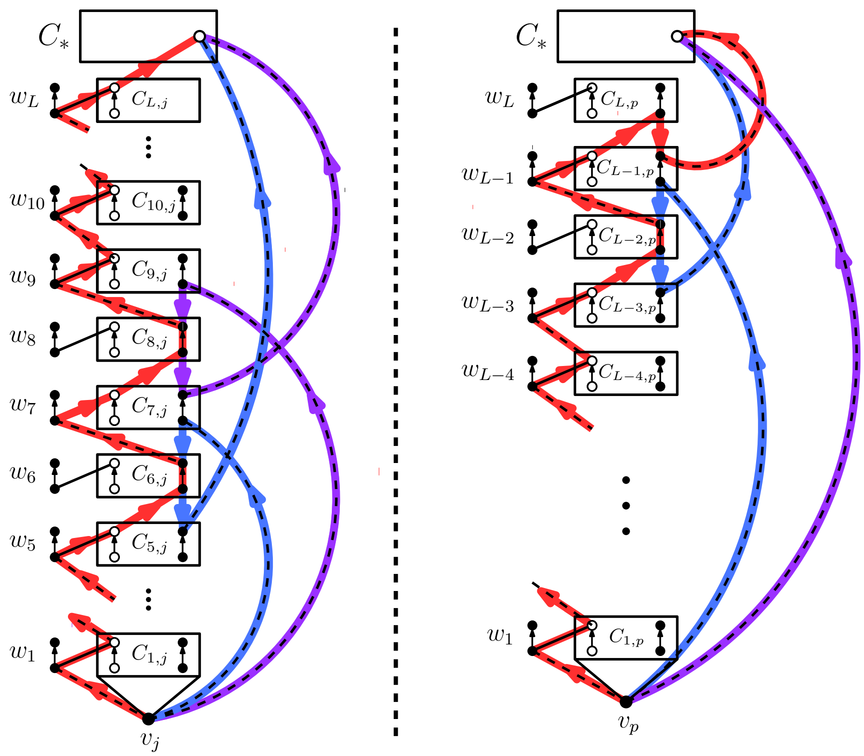

Figure A.8: The fractional augmenting paths from v_j to r^* via the long edges (there are $p-1$ such paths) and the short edges (there is 1 such path). In the figure on the left, $p = 3$ and $j = 2$, hence this (2nd) column has 2 breakpoints with indices $5, 7$. In the figure on the right, $p = 3$, $L = 2p(p-1) = 12$, and $j = 3$, hence this (3rd) column has 2 breakpoints with indices $9, 11$; note that $L - 1 = 11, L - 3 = 9$.

a breakpoint. We modify the augmenting path $P^* = v_j, w_1, u_1, w_2, u_2, \ldots, w_{L-1}, u_{L-1}, w_L, u_L, r^*$ used in the proof of Claim 2.2 to give another augmenting path P^{**}. (Note that P^* and P^{**} denote corresponding paths of $G' = (V, E_0 \cup F')$ and $G'' = (V, E_0 \cup F'')$; moreover, recall that $u_1, \ldots, u_L$ denote the designated nodes of $C_{1,j}, \ldots, C_{L,j}$.) For each of the breakpoints ℓ in column j except the last breakpoint, we replace the subpath $w_\ell, u_\ell, w_{\ell+1}, u_{\ell+1}, w_{\ell+2}$ of P^* by the subpath $w_\ell, u_\ell, q_{\ell+1}, w_{\ell+2}$, where $q_{\ell+1}$ denotes the node of $C_{\ell+1,j}$ in the path $\hat{P}$. Recall that we send a flow of value $\frac{1}{p}$ on the pair of long edges associated with the breakpoint ℓ, and in this process we push back a flow of value $\frac{1}{p}$ through a node of $\hat{P}$ that is in $C_{\ell+1,j}$; in fact, this is the node $q_{\ell+1}$; thus, it can be seen that the net flow transiting via $q_{\ell+1}$ is ≤ 1. All of the edges in the subpath are available to send the flow: we push back flow on the edge u_ℓ, w_ℓ of the base graph; the edge $u_\ell, q_{\ell+1}$ is a (so far) unused edge of the base graph; the edge $q_{\ell+1}, w_{\ell+2}$ is a (so far) unused short edge.

Recall from Section 2.3 that an exceptional case arises for the last column p of the base graph, and the pair of long edges associated with the last breakpoint of p, namely, $L-1$. Then we take the first long edge of the pair to be the edge v_p, r^*, and the second long edge of the pair to be an edge between a non-designated node of $C_{L-1,p}$ and r^*. The flow of value $\frac{1}{p}$ for this pair of long edges is sent directly on the first edge; the second edge is not used here. Instead, the second edge is used for the flow of value $\frac{1}{p}$ sent via the short edges from v_p to r^*. We take the last part of the fractional augmenting path P^{**} in column p to be $w_{L-1}, u_{L-1}, q_L, q_{L-1}, r^*$, where u_{L-1} is the designated node of $C_{L-1,p}$, and q_L, q_{L-1} are the nodes of $\hat{P}$ in $C_{L,p}, C_{L-1,p}$, respectively. All of the edges in the subpath are available to send the flow: we push back flow on the edge u_{L-1}, w_{L-1} of the base graph; the edge u_{L-1}, q_L is a (so far) unused edge of the base graph; we push back flow on the edge q_{L-1}, q_L of $\hat{P}$, and the edge q_{L-1}, r^* is the (so far unused) second edge of the exceptional pair of long edges.

Figure A.8 illustrates some of the paths in the fractional $(k+1)$-fans from $v_j, j = 1, \ldots, p-1$ and from v_p.

□

Appendix A.4. Alternative Proof of Claim 2.3

Claim 2.3 *Consider the graph $G'' = (V, E_0 \cup F'')$. The vector given by $x''(e) = \frac{1}{p}, \quad \forall e \in F''$ is a solution for (`Augmenting-LP`).*

Proof: We sketch an alternative proof that was suggested by a referee, that exploits the structure of the base graph. A *deficient set* means a nonempty set of nodes S with $\xi(S) \neq \emptyset$ and $|\Gamma(S)| = k$. A *k-node-cut* means a set of k nodes Y whose removal from the base graph results in ≥ 2 connected components; thus, $Y = \Gamma(S)$ for some deficient set S. First, we give an explicit listing of all the deficient sets of the base graph. Then we apply Claim 2.4 and Corollary 2.5 to show that each of the deficient sets in our list is covered by x'', that is, for each deficient set S, we have $x''((S, \xi(S))) \geq 1$. The second part is immediate, so we focus on the first part and on the base graph $G_0 = (V, E_0)$ for the rest of the discussion.

Recall from Section 2.4 that $d_{i,j}$ denotes the designated node of the clique $C_{i,j}, \forall i = 0, \ldots, L, \ j = 1, \ldots, p$, and $\mathcal{D}$ denote the set $\{d_{i,j} \mid i = 1, \ldots, L, \ j = 1, \ldots, p\}$; note that $\mathcal{D}$ does not contain any of the nodes $v_j = d_{0,j}, \ j = 1, \ldots, p$.

Claim *The deficient sets of the base graph are given by the sets $W_i = \{w_i\}, \ i = 1, \ldots, L$, $S_{0,j} = \{v_j\}, \ j = 1, \ldots, p$, and $S_{i,j} - \mathcal{D}', \ \mathcal{D}' \subseteq \mathcal{D}, i = 1, \ldots, L-1, j = 1, \ldots, p$.*

We sketch a proof of this claim, using several observations used in the proofs of Claims 2.1–2.3.

Let Y be any k-node-cut of the base graph. Clearly, the nodes of $(C_* \cup C_{L,1} \ldots C_{L,p}) - Y$ are in the same connected component of $G_0 - Y$ because $C_* \cup C_{L,j}$ forms a clique of G_0 with $\geq k + 2 + (k/2)$ nodes, for each $j = 1, \ldots, p$. Let this connected component be denoted by $\mathcal{CC}_*$. Moreover, note that Y contains none of the nodes of C_*, because for each node r of C_*, $G_0 - r$ is k-connected (this follows from the proof of Claim 2.1).

Next, observe that none of the nodes $v_j, \ j = 1, \ldots, p$ is contained in Y. To see this, fix $j = 1, \ldots, p$ and note that $G_0 - v_j$ is k-connected; this follows from the proof of Claim 2.1 since the openly disjoint paths and the k-fans used in the proof do not contain v_j, except for the k-fan used for scanning v_j.

Thus we have two cases. Either $\mathcal{CC}_*$ contains $\{v_1, \ldots, v_p\}$ or not.

Case 1: $v_1, \ldots, v_p$ are in $\mathcal{CC}_*$. Then we claim that every interior node of $G_0 - Y$ is in $\mathcal{CC}_*$. This follows from the proof of Claim 2.2; in part (4) of that proof we constructed a $(k+1)$-fan from any interior node u of column j to a set of $k+1$ nodes consisting of v_j and k nodes of C_*; if $u \notin Y$, then at least one path of this $(k+1)$-fan is present in $G_0 - Y$, showing that u is in $\mathcal{CC}_*$.

Thus every other connected component $\mathcal{CC}'$ of $G_0 - Y$ consists of nodes of $\{w_1, \ldots, w_L\}$; moreover, it can be seen that any such $\mathcal{CC}'$ consists of exactly one node from $\{w_1, \ldots, w_L\}$, because $|\Gamma(S)| > k$ for any set S of two or more of the nodes $w_1, \ldots, w_L$.

Case 2: there is a node $v_j,\ j = 1, \ldots, p$ that is not in $\mathcal{CC}_*$ (thus we fix $j = 1, \ldots, p$ such that v_j is not in $\mathcal{CC}_*$). Let $\mathcal{CC}_j$ denote the connected component of $G_0 - Y$ that contains v_j. Then we claim that Y contains $C_{i,j}$ for an $i \in \{1, \ldots, L\}$. This follows easily; suppose that $C_{i,j} - Y \neq \emptyset,\ \forall i \in \{1, \ldots, L\}$; then, by induction on $i = L, L-1, \ldots, 1$, it can be seen that $G_0 - Y$ has a path from each node of $C_{i,j} - Y$ to C_*; this gives a contradiction, since we get a path from v_j to C_* in $G_0 - Y$. Note that Y cannot contain more than one of the cliques $C_{i,j}$, since $|Y| = k$ and two of these cliques together contain $\geq 2k + 2 - (i + i') \geq k + 2$ nodes. Fix i such that $C_{i,j} \subseteq Y$. Clearly, every node of $\bigcup_{\ell=i+1,L} C_{\ell,j} - Y$ has a path to C_* in $G_0 - Y$, and hence is in $\mathcal{CC}_*$; similarly, every node of $\bigcup_{\ell=1,i-1} C_{\ell,j} - Y$ has a path to v_j in $G_0 - Y$, and hence is in $\mathcal{CC}_j$; moreover, $C_{\ell,j} - Y$ has at least two nodes, $\forall \ell \in \{1, \ldots, L\} - \{i\}$.

Finally, we claim that Y contains exactly one of the two nodes w_ℓ or $d_{\ell,j}$, $\forall \ell \in \{1, \ldots, i-1\}$. To see this, fix ℓ, and focus on any nondesignated node u in $C_{\ell,j} - Y$; u exists since $C_{\ell,j} - Y$ has at least two nodes. Consider the k-fan in G_0 from u to k nodes of C_* constructed in part (3) of the proof of Claim 2.1; observe that Y does not contain any node of C_*; hence, Y contains exactly one internal node of each of the k paths of this k-fan; note that the k-fan has a path of the form $u, d_{\ell,j}, w_\ell, r$, where r is in C_*; hence, Y contains exactly one of $d_{\ell,j}$ and w_ℓ.

It can be seen that $\mathcal{CC}_j$ is contained in column j, and its node set is given by $S_{i,j} - \mathcal{D}'$, where $\mathcal{D}' \subseteq \mathcal{D}$.

This completes the proof of the claim on the deficient sets of the base graph. Thus we gave an alternative proof of Claim 2.3.

□

Appendix A.5. Proof of Claim 2.6

Claim 2.6 *The matrix B has full rank.*

Proof: We show that using elementary column operations the matrix B can be transformed into a lower triangular matrix with non-zero diagonal entries. We assume that the parameter p in the construction is an integer ≥ 2. Although the proof holds for $p = 2$, some of the formulas given below (e.g., for Λ_j) apply only for $p \geq 3$. For the sake of notational convenience, we allow some informality in what follows. In particular, we may use the same symbol (e.g., $Q_j, \widehat{I}_j, B$) to denote a matrix as well as its updated version after applying elementary column operations.

Initially, before we apply any column operations, recall that

$$B = \left(\begin{array}{c|c|c||c|c|c} Q_1 & 0 & 0 & \widehat{I}_1 & 0 & 0 \\ 0 & \ddots & 0 & 0 & \ddots & 0 \\ 0 & 0 & Q_p & 0 & 0 & \widehat{I}_p \\ \hline\hline 0 & \ldots & 0 & I_{L \times L} & \ldots & I_{L \times L} \end{array} \right),$$

where $L = 2(p-1)p$; thus, B has $(p+1)L$ rows and the same number of columns; the left part of B (to the left of the vertical double line) has L columns, and the right part of B has pL columns. Moreover, recall that

$$Q_j = (f_h, 1_L - f_{h-1}, f_{h+2}, 1_L - f_{h+1}, \ldots, f_{h+2(p-2)}, 1_L - f_{h+2(p-2)-1}),$$

where $h = 2(p-1)(j-1) + 2$. Recall that e_i denote the i-th column of the $L \times L$ identity matrix, and f_i denote a column vector of size L with a 1 in entries $1, \ldots, i$ and a 0 in entries $i+1, \ldots, L$. If $i > L$, then e_i denotes a vector of zeroes of size L, and f_i denotes a vector of ones of size L; indices $i > L$ may occur in formulas pertaining to the last breakpoint of the last column of the base graph.

$$
Q_2 = \left(\begin{array}{cc|cc}
1 & & 1 & \\
1 & & 1 & \\
1 & & 1 & \\
1 & & 1 & \\
1 & & 1 & \\
1 & 1 & 1 & \\
 & 1 & 1 & \\
 & 1 & 1 & 1 \\
 & 1 & & 1 \\
 & 1 & & 1 \\
 & 1 & & 1 \\
 & 1 & & 1
\end{array}\right)
\qquad
Q_2^{(1)} = \left(\begin{array}{cc|cc}
1 & & 0 & \\
1 & & 0 & \\
1 & & 0 & \\
1 & & 0 & \\
1 & & 0 & \\
1 & 1 & 0 & \\
 & 1 & 1 & \\
 & 0 & 1 & 1 \\
 & 0 & & 1 \\
 & 0 & & 1 \\
 & 0 & & 1 \\
 & 0 & & 1
\end{array}\right)
\qquad
Q_2^{(2)} = \left(\begin{array}{cc|cc}
0 & & & \\
0 & & & \\
0 & & & \\
0 & & & \\
0 & & & \\
1 & 1 & & \\
 & 0 & 0 & \\
 & & 1 & 1 \\
 & & & 0 \\
 & & & 0 \\
 & & & 0 \\
 & & & 0
\end{array}\right)
$$

Figure A.9: Elementary column operations applied to Q_2 in the first stage.

Let Λ denote the $L \times L$ submatrix in the bottom-left corner of B (to the left of the vertical double line and below the horizontal double line). Initially, all entries in Λ are zero.

We apply elementary column operations in two stage. In the first stage, for each $j = 1, \ldots, p$, we make all entries of Q_j zero and transform $\widehat{I}_j$ into the identity matrix. These operations introduce some nonzero entries into the submatrix Λ. In the second stage, we use elementary column operations to transform Λ to a lower triangular matrix with non-zero diagonal entries. The final matrix has the following form; the matrix in the bottom right corner is not relevant in the analysis.

$$
\left(\begin{array}{c|c}
0_{pL\times L} & I_{pL\times pL} \\ \hline
\Lambda_{L\times L} &
\end{array}\right)
$$

By swapping the first column with the second column of this matrix, we get a lower triangular matrix with non-zero diagonal entries, and this completes the proof.

$$
\left(\begin{array}{c|c}
0_{pL\times L} & I_{pL\times pL} \\ \hline
\Lambda_{L\times L} &
\end{array}\right)
\quad \longrightarrow \quad
\left(\begin{array}{c|c}
I_{pL\times pL} & 0_{pL\times L} \\ \hline
 & \Lambda_{L\times L}
\end{array}\right)
$$

- **First stage:** Note that the indices of the breakpoints in the j-th column are $h, h+2, h+4, \ldots, h+2(p-2)$, where $h = 2(p-1)(j-1)+2$. First, we apply the following $2(p-2)$ elementary column operations to each matrix Q_j. We subtract the $(i+2)$-th column from the i-th column, for each $i = 2, 4, \ldots, 2(p-2)$, in this order. Next, we subtract the i-th column from the $(i+2)$-column, for each $i = 2(p-2)-1, \ldots, 3, 1$, in this order. These operations do not change other submatrices of B. After applying these operations, we get the following matrix; see Figure A.9 for an illustration.

$$
\begin{aligned}
Q_j^{(1)} &= (f_h, e_h + e_{h+1}, e_{h+1} + e_{h+2}, e_{h+2} + e_{h+3}, \ldots, e_{h+2p-5} + e_{h+2p-4}, 1_L - f_{h+2p-5}) \\
&= (e_{h-1} + e_h, e_h + e_{h+1} | e_{h+1} + e_{h+2}, e_{h+2} + e_{h+3} | \ldots | e_{h+2p-5} + e_{h+2p-4}, e_{h+2p-4} + e_{h+2p-3}) \\
&\quad + (f_{h-2}, 0, \ldots, 0, 1_L - f_{h+2p-3})
\end{aligned}
$$

Next, by subtracting the columns of the matrix $\widehat{I}_j$ from the columns of the current matrix Q_j, we can make all of the entries in Q_j zero except the entries on the rows corresponding to the breakpoints; see Figure A.9 for an illustration. After applying these operations we get the following matrix.

$$
Q_j^{(2)} = (e_h, e_h | e_{h+2}, e_{h+2} | \ldots | e_{h+2p-4}, e_{h+2p-4})
$$

These operations will change the submatrix Λ in the bottom left corner of B. We partition the columns of Λ into p blocks, denoted $\Lambda_1, \ldots, \Lambda_p$, where each block consists of $2(p-1)$ consecutive columns. Thus $\Lambda = (\Lambda_1, \Lambda_2, \ldots, \Lambda_p)$.

After applying the above operations, it can be seen that the matrix $\Lambda_j^{(1)}$ is changed to $\Lambda_j^{(2)} = Q_j^{(2)} - Q_j^{(1)}$. To see this, observe that we change an entry in the i-th row of Q_j from α to β by adding $(\beta - \alpha)$ times the i-th column of $\widehat{I}_j$, where i is *not* a breakpoint index. The column of $\widehat{I}_j$ is identical to the column of the identity matrix at the bottom (in the last L rows of B). Hence, the corresponding entry of Λ_j changes from 0 to $\beta - \alpha$. Thus, we have

$$\begin{aligned}\Lambda_j^{(2)} &= (-e_{h-1}, -e_{h+1} | -e_{h+1}, -e_{h+3} | \ldots | -e_{h+2p-5}, -e_{h+2p-3}) \\ &\quad + (-f_{h-2}, 0, \ldots, 0, -1_L + f_{h+2p-3})\end{aligned}$$

Next, we subtract each even-indexed column of Q_j from the column to the left of it. These operations change $\Lambda_j^{(2)}$ to

$$\begin{aligned}\Lambda_j^{(3)} &= (e_{h+1} - e_{h-1}, -e_{h+1} | e_{h+3} - e_{h+1}, -e_{h+3} | \ldots | e_{h+2p-3} - e_{h+2p-5}, -e_{h+2p-3}) \\ &\quad + (-f_{h-2}, 0, \ldots, 0, 1_L - f_{h+2p-3}, -1_L + f_{h+2p-3})\end{aligned}$$

Finally, we swap the column indexed by 2ℓ with the column of $\widehat{I}_j$ corresponding to the ℓ-th breakpoint, for $\ell = 1, 2, \ldots, (p-1)$. These operations make all entries of Q_j zero, and they transform the matrix $\widehat{I}_j$ into the identity matrix. The matrix $\Lambda_j^{(3)}$ is changed to the following matrix; see Figure A.10a for an illustration.

$$\begin{aligned}\Lambda_j^{(4)} &= (e_{h+1} - e_{h-1}, e_h | e_{h+3} - e_{h+1}, e_{h+2} | \ldots | e_{h+2p-3} - e_{h+2p-5}, e_{h+2p-4}) \\ &\quad + (-f_{h-2}, 0, \ldots, 0, 1_L - f_{h+2p-3}, 0)\end{aligned}$$

$$\left(\begin{array}{cccc|cccc|cccc}
-1 & & & & -1 & & & & -1 & & & \\
 & 1 & & & -1 & & & & -1 & & & \\
1 & & -1 & & -1 & & & & -1 & & & \\
 & & & 1 & -1 & & & & -1 & & & \\
 & & 1 & & -1 & & & & -1 & & & \\
 & & 1 & & & 1 & & & -1 & & & \\
 & & 1 & & 1 & & -1 & & -1 & & & \\
 & & 1 & & & & & 1 & -1 & & & \\
 & & 1 & & & & 1 & & -1 & & & \\
 & & 1 & & & & 1 & & & 1 & & \\
 & & 1 & & & & 1 & & 1 & & -1 & \\
 & & 1 & & & & 1 & & & & & 1
\end{array}\right)$$

(a) $\Lambda = (\Lambda_1 | \Lambda_2 | \Lambda_3)$ at the start of the second stage.

$$\left(\begin{array}{cccc|cccc|cccc}
-1 & & & & & & & & & & & \\
 & 1 & & & -1 & & & & & & & \\
1 & & -1 & & -2 & & & & & & & \\
 & & & 1 & -1 & & & & & & & \\
 & & 1 & & -1 & & & & & & & \\
 & & 1 & & & 1 & & & -1 & & & \\
 & & 1 & & 1 & & -1 & & -2 & & & \\
 & & 1 & & & & & 1 & -1 & & & \\
 & & 1 & & & & 1 & & -1 & & & \\
 & & 1 & & & & 1 & & & 1 & & \\
 & & 1 & & & & 1 & & 1 & & -1 & \\
 & & 1 & & & & 1 & & & & & 1
\end{array}\right)$$

(b) $\Lambda = (\Lambda_1 | \Lambda_2 | \Lambda_3)$ after the first step of the second stage.

Figure A.10: The matrix $\Lambda = (\Lambda_1 | \Lambda_2 | \Lambda_3)$ at the beginning of the second stage, and just after the first step of the second stage, which subtracts the first column of Λ_j from the first column of Λ_{j+1}. The next step obtains zeros at all the even-indexed entries in the first column of Λ_{j+1} by adding the even-indexed columns of Λ_j.

- **Second stage:** While describing this stage, we use the term diagonal to mean the diagonal of the matrix Λ. Thus for a submatrix Λ_j or for a column of Λ, the diagonal refers to the entries of the diagonal of Λ restricted to that submatrix or column.

  Note that all nonzero entries in Λ_j are below or on the diagonal, except for entries in the first column; the first column has -1 on all rows above and including the diagonal. To make Λ lower triangular, we apply some elementary column operations on the first column of each Λ_j. We do this in two steps.

  In the first step, we subtract the first column of Λ_j from the first column of Λ_{j+1}, starting from $j = (p-1)$ down to $j = 1$. This removes many consecutive -1 entries, and leaves only $2p-3$ non-zero entries above the diagonal in the first column of each Λ_j for $j = 2, \ldots, p$.

  In the second step, for each $j = 1, \ldots, p-1$, we change the remaining nonzero entries above the diagonal in Λ_{j+1} to zeros by using the columns of Λ_j. The *same* elementary column operations are applied for all Λ_j. Sequentially, consider $j = 1, \ldots, p-1$ and assume that Λ_j has no non-zero entry above the diagonal. (Note that Λ_1 has no non-zero entry above the diagonal.) First, we take each even-indexed column of Λ_j and add it to the first column of Λ_{j+1}. After this, there remain $p-2$ non-zero entries above the diagonal in the first column of Λ_{j+1}; note that the second entry above the diagonal is -2. See Figure A.10b for an illustration. Then we apply $p-2$ elementary column operations to replace the remaining $p-2$ non-zero entries above the diagonal by zeros, and in the process the diagonal entry becomes $-p$. In more detail, for each $\ell = 1, 2, \ldots, p-2$ in sequence, we multiply the $(2\ell+1)$-th column of Λ_j by $-(\ell+1)$ and add the result to the first column of Λ_{j+1}. To verify this, observe that for $\ell = 1, 2, \ldots, p-3$, the $(2\ell+1)$-th column of Λ_j is given by $e_{h+2\ell+1} - e_{h+2\ell-1}$, and the $(2p-3)$-th column is given by $e_{h+2p-3} - e_{h+2p-5} + (1_L - f_{h+2p-3})$, hence, (by induction) the topmost non-zero entry in the first column of Λ_{j+1} is $-(\ell+1)$ just before we apply the ℓ-th of these $p-2$ column operations.

  Thus, the matrix Λ is transformed into a lower triangular matrix with non-zero diagonal entries.

This completes the proof of Claim 2.6, and shows that B has full rank.

□

Acknowledgments: We thank the referees for their careful reading of a previous draft, and for suggesting an alternative proof of Claim 2.3.

References

[1] J.Cheriyan, S.Vempala and A.Vetta, *Network design via iterative rounding of setpair relaxations*, Combinatorica, **26**(3), pp. 255–275, 2006.
Preliminary version in Proc. 34th Ann. ACM STOC, pp.306–312, May 2002.

[2] S.Even, *An algorithm for determining whether the connectivity of a graph is at least k*, SIAM J. Comput. 4(3): 393-396, 1975.

[3] S.Even and R.E.Tarjan, *Network flow and testing graph connectivity*, SIAM J. Comput. 4(4): 507-518, 1975.

[4] J.Fakcharoenphol and B.Laekhanukit, *An $O(\log^2 k)$-approximation algorithm for the k-vertex connected subgraph problem*, Proc. 40th Ann. ACM STOC, May 2008.

[5] L.Fleischer, *A 2-Approximation for Minimum Cost* $\{0, 1, 2\}$ *Vertex Connectivity*, Proc. IPCO, pp.115–129, May 2001.

[6] L.Fleischer, K.Jain, and D.P.Williamson, *An iterative rounding 2-approximation algorithm for the element connectivity problem*, *Journal of Computer and System Sciences*, **72** (5), pp. 838–867, 2006.
Preliminary version in Proc. 42nd IEEE FOCS, pp.339-347, 2001.

[7] A.Frank, *Connectivity augmentation problems in network design*, in Mathematical Programming: State of the Art 1994, (Eds.J.R.Birge and K.G.Murty), pp.34–63, The University of Michigan, Ann Arbor, MI, 1994.

[8] A.Frank and T.Jordan, *Minimal edge-coverings of pairs of sets*, Journal of Combinatorial Theory, Series B, **65**, pp.73-110, 1995.

[9] A.Frank and E.Tardos, *An application of submodular flows*, Linear Algebra and its Applications, **114/115**, 1989, pp.329-348.

[10] K.Jain, *A factor* 2 *approximation algorithm for the generalized Steiner network problem*, Combinatorica, **21**, pp.39-60, 2001. Preliminary version in Proc. 39th IEEE FOCS, 1998.

[11] G.Kortsarz and Z.Nutov, *Approximating k-node connected subgraphs via critical graphs*, SIAM J. Comput. **35** (1), pp. 247–257, 2005.
Preliminary version in Proc. 36th ACM STOC, pp. 138–145, June 2004.

[12] Z.Nutov, *An almost $O(\log k)$ approximation for k-connected subgraphs*, Proc. ACM-SIAM SODA, 2009.

[13] R.Ravi and D.P.Williamson, *An approximation algorithm for minimum-cost vertex-connectivity problems*, Algorithmica, **18**, pp.21-43, 1997. Preliminary version in Proc. 6th ACM-SIAM SODA, pp.332-341, 1995.
Erratum, Algorithmica, **34**, pp.98-107, 2002.

[14] V.V.Vazirani, *Approximation Algorithms*, Springer-Verlag, Berlin, 2001.

[15] D.P.Williamson and D.B.Shmoys, *The Design of Approximation Algorithms*, Cambridge University Press, to be published in 2010.