Relay placement for two-connectivity

Gruia Calinescu[a,1]

[a] *Department of Computer Science, Illinois Institute of Technology, 10 W. 31st St., Chicago, IL 60616, USA*

Abstract

Motivated by applications to wireless sensor networks, we study the following problem. We are given a set S of wireless sensor nodes, given as a multiset of points in a normed space. We seek to place a minimum-size (multi)set Q of wireless relay nodes in the normed space such that the unit-disk graph induced by $Q \cup S$ is two-connected. The unit-disk graph of a set of points has an edge between two points if their distance is at most 1.

In Infocom 2006, Kashyap, Khuller, and Shayman present two algorithms, for the two variants of the problem: two-edge-connectivity and biconnectivity. For both they prove an approximation ratio of $2d_{MST}$, where d_{MST} is the maximum degree of a minimum-degree Minimum Spanning Tree in the normed space. It is known that in the Euclidean two-dimensional space, $d_{MST} = 5$, and in the three dimensional space, $d_{MST} = 12$.

We give a tight analysis of variants of the same algorithms, obtaining approximation ratios of d_{MST} for biconnectivity and $2d_{MST} - 1$ for two-edge-connectivity respectively. To do so we prove additional structural properties regarding bypassing Steiner nodes in biconnected graphs.

Keywords: approximation algorithm, wireless network, Steiner points, submodular flows, two connectivity, parsimony

1. Introduction

A wireless sensor network is composed of a large number of sensors, which can be densely deployed to monitor the targeted environment. Some of the most important application areas of sensor networks include military, natural calamities such as forest fire detection and tornado motion, and different sorts of surveillance. When compared to traditional ad hoc networks, the most noticeable point about sensor networks is that they are limited in power, computational capacities, and memory.

Sensors may have a short transmission range, since long transmission consumes more energy, and the sensors normally have limited power. Therefore, network partitions may

☆An extended abstract appeared in Proc. IFIP Networking 2012

Email address: `calinescu@iit.edu` (Gruia Calinescu)

[1]Research supported in part by NSF grant NeTS-0916743.

occur or more sensors must be placed to maintain connectivity. Higher connectivity may be desired to ensure fault-tolerance.

Formally, in the TWO-CONNECTED RELAY PLACEMENT problem, we are given a set S of wireless sensor nodes, given as a multiset of points in a finite-dimensional normed space (”multiset” meaning two nodes may be placed at the same location). We seek to place a minimum-size (multi)set Q of wireless relay nodes in the normed space such that the unit-disk graph induced by $Q \cup S$ is two-connected.

A normed space is a metric space (X, d), given by a set X (of points) and a symmetric function (distance) $d : X \times X \to \mathbb{R}^+$ that obeys the triangle inequality: $\forall x, y, z \in X$, $d(x, y) \leq d(x, z) + d(z, y)$, and the property that $d(x, y) = 0$ if and only if $x = y$. As defined in the literature (3), a normed space also has the following property (and others that we do not use): $\forall x, y \in X$ and $\forall \alpha \in [0, 1]$, there exists $z \in X$ such that $d(x, y) = d(x, z) + d(z, y)$ and $d(x, z) = \alpha \cdot d(x, y)$. In other words, the normed space contains all the *Steiner* points. Normed spaces of interest to wireless networks are the two and three dimensional Euclidean space, with d being the Euclidean distance (the l_2 norm).

A multiset allows several nodes to be placed at the same location. The unit-disk graph of a set of points has an edge between two points if their distance is at most 1 (we normalize to 1 the transmission range of the sensors). For a multiset of points Z, let $U(Z)$ be the unit-disk graph induced by Z. Also, we call two vertices U-adjacent, or U-neighbors, if their distance is at most 1.

Kashyap, Khuller, and Shayman (17; 18) introduce the two variants of this problem: TWO-EDGE-CONNECTED RELAY PLACEMENT (when $U(S \cup Q)$ must be two-edge-connected, that is, have between any two vertices two edge-disjoint paths) and BICONNECTED RELAY PLACEMENT ($U(S \cup Q)$ must be biconnected, that is, have between any two vertices two internally vertex-disjoint paths). Two paths are internally vertex-disjoint if they only have the endpoints in common. Biconnectivity also goes by the name of two-vertex-connectivity, or two-connectivity.

Let d_{MST} be the maximum degree of a minimum-degree Minimum Spanning Tree in the normed space. It is known (27; 25) that d_{MST} is the strict Hadwiger number of the unit ball in the normed space, defined as follows: the maximum size of an independent set in $U(N_x)$, taken over all the points x of the space, with N_x being the points, other than x, within distance 1 of x. Is is known that $d_{MST} = 5$ in the Euclidean two-dimensional space, and $d_{MST} = 12$ in three dimensions (25).

(17) presents two algorithms, based on the Khuller and Vishkin (20) (**Algorithm KV**) and the Khuller and Raghavachari (19) (**Algorithm KR**) algorithms for MINIMUM-WEIGHT SPANNING TWO-EDGE-CONNECTED SUBGRAPH, and MINIMUM-WEIGHT SPANNING BICONNECTED SUBGRAPH, respectively. For these problems, a weighted graph $G = (V, E, w)$ is given as an input, and one must select a minimum weight set of edges F such that (V, F) is two-edge-connected, or biconnected respectively. For TWO-CONNECTED RELAY PLACEMENT, (17) proves that variants of the two algorithms have each approximation ratio of $2d_{MST}$.

We give a tight analysis of variants of the same algorithms, obtaining approximation ratios of d_{MST} for biconnectivity and $2d_{MST} - 1$ for two-edge-connectivity respectively. Thus, in the two-dimensional Euclidean plane, we get a ratio of 9, instead of 10 ((17)), for two-edge-connectivity and 5, instead of 10 ((17)), for biconnectivity. Assuming that no post-processing removes redundant relay nodes, the ratios given in this paragraph are tight for these algorithms.

For the ratio of $2d_{MST} - 1$, we use a more careful accounting and look inside **Algorithm KV**. For the ratio of d_{MST} for biconnectivity, we look inside **Algorithm KR**, and prove a property of biconnected graphs that may be of independent interest.

This property is technical and we only describe here a simpler version that is not used in proving our main results. We prove the following result, new to the best of our knowledge. Let H be a biconnected planar undirected graph, and replace every edge by two anti-parallel directed arcs. Let R be a subset of $V(H)$. Then there exists a set of arc-disjoint directed paths P_i of H, all starting and ending at a vertex of R and without interior vertices from R, such that, if we replace each P_i by an arc e_i joining the start and the end vertex of P_i, we obtain a biconnected digraph on R. This property allows one to "bypass" Steiner vertices ("parsimony") and in some sense eliminate them. This parsimony differs from the classical concept as given in (13) since it applies to combinatorial (and not linear programming) solutions.

For graphs in general, we prove a "fractional outconnected" variant of the property above, and use it together with **Algorithm KR** to obtain an approximation ratio of d_{MST} for biconnectivity in arbitrary normed spaces. Structural properties of biconnected Steiner networks were also studied by (15; 14; 30; 23), and we use some of their results and techniques for our "outconnected parsimony". Using these structural properties, we construct from the optimum solution of an arbitrary Biconnected Relay Placement instance a fractional solution to a certain polytope. This polytope was proposed by Frank and Tardos (11), who proved that it is integral (see also (10)). Thus, there exists an integral solution with cost at most this fractional solution, for any non-negative cost function. We define costs that relate the objective function to an optimum relay solution, and notice that the output of **Algorithm KR** in a weighted graph that we describe later is (almost) derived from an integral polytope optimum solution.

As an example of its possible applications, this new fractional outconnected parsimony property can be used to prove that a variant of **Algorithm KR** has approximation ratio of 2 for the following network design problem: we are given a normed space and a set of terminals S. We must choose a set Z of points and a set of edges F of minimum total distance such that the graph $(Z \cup S, F)$ is biconnected. This would not be an improvement, since an approximation algorithm with a ratio of 2 is already known in finite graphs (even without a metric cost function) from the paper by Fleischer, Jain, and Williamson (9). In Euclidean spaces, a PTAS (for any $\epsilon > 0$, there is an algorithm with approximation ratio of $1 + \epsilon$; the running time being polynomial for any *fixed* ϵ) was announced by Czumaj and Lingas (7); their algorithms have running time exponential in the dimension and in $1/\epsilon$. Also, no fully PTAS exists (23).

In other previous works, Wang, Thai, and Du (28) and Bredin, Demaine, Hajiaghayi, and

Rus (2) also gave constant factor algorithms, those of (2) achieving an $O(k^4 \log k)$ approximation for k-connectivity of $U(S \cup Q)$ in the Euclidean plane (it seems from their proof that a d_{MST} factor would apply in other normed spaces). We remark that in a wireless setting, one only needs k-connectivity between the vertices of S, i.e. k edge-disjoint (or internally vertex-disjoint) paths between any two vertices of S. For $k = 1$ or $k = 2$, by eliminating redundancy from any solution, one can see that k-connectivity or k-edge-connectivity between the vertices of S implies k-connectivity or k-edge-connectivity, respectively, of $U(S \cup Q)$. We only present the argument for 2-connectivity: If $U(S \cup Q)$ is not biconnected, it has a vertex v such that $U((S \cup Q) \setminus \{v\})$ has at least two connected components, and one of these two components contains no vertex of S, since we have two-connectivity between the vertices of S; the removal of this component does not decrease the connectivity between the vertices of S. This argument fails for $k > 2$. When requiring only k-connectivity between the vertices of S, the best known approximation ratio is obtained by Kamma and Nutov (16): $O(d_{MST} k^2 \log k)$.

1.1. Related Work

MSPT (Minimum Number of Steiner Points Tree with bounded edge length) is the following problem: Given S in the plane, find minimum Q such that $U(S \cup Q)$ is connected. This problem was introduced by Lin and Xue (21) and proven NP-hard. They also prove that taking a Euclidean minimum spanning tree, and placing a minimum number of relay nodes on each edge of the tree to connect the endpoints of the edge, achieves an approximation ratio of 5. Mandoiu and Zelikovsky (24) give a tight analysis of 4 for the MST-based algorithm described above, and generalize the proof to arbitrary normed spaces obtaining a ratio of $d_{MST} - 1$. Chen, Du, Hu, Lin, Wang, and Xue also prove in (4) the same ratio of 4 but with a different approach, and present a 3-approximation algorithm. Later, Cheng, Du, Wang, and Xu (5) improve the running time of some of the algorithms found in (4) and present a randomized algorithm with approximation ratio 2.5. In arbitrary normed spaces, Nutov and Yaroshevitch (26) obtain a $\lfloor (d_{MST} + 1)/2 \rfloor + 1 + \epsilon$-approximation.

2. Two-edge-connectivity

We start with notation. For any graph G, we use $\overrightarrow{G}$ to represent the bidirected version of G, that is the weighted digraph obtained from G by replacing every edge uv of G with two oppositely oriented arcs uv and vu with the same weight as the edge uv in G. As usual, the *weight* of a subgraph H of G is defined as $w(H) = \sum_{e \in E(H)} w(e)$, and the *weight* of a subdigraph D of $\overrightarrow{G}$ is defined as $w(D) = \sum_{e \in E(D)} w(e)$.

A spanning subdigraph A of $\overrightarrow{G}$ is said to be an *arborescence* rooted at some vertex $s \in V(G)$ if A contains exactly $|V(G)| - 1$ arcs and there is a path in A from s to any other vertex. In other words, arborescences in directed graphs are directed analogs of spanning trees in undirected graphs.

Call a feasible solution Q of a Two-Connected Relay Placement problem instance a *bead-solution* if $U(Q \cup S)$ contains a two-edge-connected graph (or biconnected, respectively)

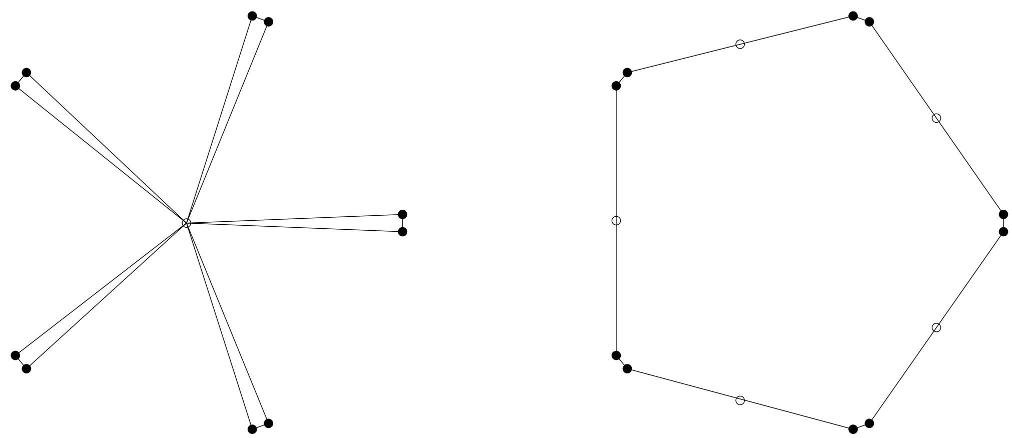

Figure 1: On the left, an optimum solution for two-edge-connectivity. The nodes of S are black disks, and the relay nodes are empty circles. On the right, an optimum bead solution.

H where each node of Q has degree exactly two. The Kashyap et al. (17) algorithms produce bead solutions - see for example Figure 1, borrowed from (17). In a bead-solution, we may call the relay nodes *beads*.

For two-edge connectivity, we describe the analysis using Euclidean distance in two dimensions. It is straightforward to extend the analysis below to any normed space, using d_{MST} instead of 5 and $d_{MST} - 1$ instead of 4, as the only time 4 comes in below is as $d_{MST} - 1$ in the MST-based algorithm for MSPT, and we can use the result of (24) in arbitrary normed spaces.

The Kashyap, Khuller, and Shayman algorithm works as follows. For $x, y \in S$, define $w(x, y) = \max(0, \lceil ||x, y|| \rceil - 1)$, where $||u, v||$ denotes the Euclidean distance from u to v. One can easily verify that $w(x, y)$ is the minimum number of relay nodes required to connect x and y, and that $w(x, y)$ is an increasing function of $||x, y||$. However, w is not a metric.

If $w(x, y) > 0$, allow two parallel edges of weight $w(x, y)$ between x and y; otherwise allow only one edge of weight 0, plus parallel edges of weight 1 (thus a bead may be placed to increase connectivity among U-adjacent nodes of S), creating an edge-weighted multigraph G. Use **Algorithm KV** to compute in G a set of edges F, attempting to minimize $w(F)$ while (V, F) is two-edge-connected. Replace each edge of positive weight by new beads (that is, every such edge has its own distinct beads); this is the output. **Algorithm KV** is a 2-approximation for Minimum Weight Spanning Two-Edge-Connected Subgraph, and the approximation ratio of (17) is based on showing that G has a two-edge-connected subgraph of weight at most $5opt$, where opt is the value of an optimum relay solution.

The constant 5 above is known to be tight (17), as in Figure 1.

Theorem 1. ***Algorithm KV*** *has approximation ratio* 9.

Proof: We look deeper into **Algorithm KV**, which works as follows. Pick in $V(G)$ an arbitrary root r. Use the polynomial-time algorithm of Gabow (12) to compute two arc-disjoint arborescences A and B of $\overrightarrow{G}$, rooted at r, such that $w(A) + w(B)$ is minimized. Output an edge xy if either xy or yx are in $A \cup B$. It is shown in (20) (and not hard to see) that the output graph is two-edge-connected.

To prove the approximation ratio, it suffices to construct, from an optimal solution, two arc-disjoint r-rooted arborescences A and B in $\overrightarrow{G}$ satisfying

$$w(A) + w(B) \leq 9\,opt. \tag{1}$$

This process is illustrated in Figure 2, with $S = \{r, x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_8, x_9, x_{10}\}$.

We assume $opt \geq 1$, as we can easily check if $U(S)$ is two-edge-connected. Consider OPT, a multiset of points giving an optimal solution. Partition OPT into connected components $Q_1, Q_2, \ldots, Q_k$ of $U(OPT)$. In the example of the Figure 2, we have $k = 2$, $Q_1 = \{u_1, u_2\}$ and $Q_2 = \{u_3\}$. Let S_i be the set of U-neighbors of Q_i in S; these S_i are not necessarily disjoint sets. In Figure 2, we have $S_1 = \{x_3, x_4, x_5, x_{10}\}$ and $S_2 = \{x_2, x_5, x_6, x_7, x_8, x_9\}$.

There exist two r-rooted arc-disjoint arborescences A_0 and B_0 in D_0, the bidirected $U(OPT \cup S)$, as $U(OPT \cup S)$ is two-edge-connected (the existence of the two arborescences is stated and used in (20), and is not hard to see). An example of A_0 and B_0 appears in Figure 2 (b). The edges of $U(OPT \cup S)$ have weight 0 and therefore

$$w(A_0) + w(B_0) = 0. \tag{2}$$

We do the following for $i = 1, 2, \ldots, k$. From A_{i-1} and B_{i-1}, we create A_i and B_i, arc-disjoint r-rooted arborescences in the digraph D_i, where D_i consists of the bidirected version of $U\left(S \cup \bigcup_{j=i+1}^{k} Q_i\right)$, plus some arcs of $\overrightarrow{G}$ that have positive weight. Moreover, we will ensure that

$$w(A_i) + w(B_i) \leq w(A_{i-1}) + w(B_{i-1}) + 9|Q_i| \tag{3}$$

and therefore, by adding up for all i, and using Equation 2, we get

$$w(A_k) + w(B_k) \leq 9\,opt, \tag{4}$$

thus obtaining Equation 1 with $A = A_k$ and $B = B_k$. In the example of Figure 2, A_1 and B_1 appear in subfigure (d) and A_2 and B_2 appear in subfigure (f).

The arcs of $\overrightarrow{G}$ that are used in A_i and B_i and that are not from $A_{i-1} \cup B_{i-1}$, have both endpoints in S_i. Let us fix an $i \in \{1, 2, \ldots, k\}$ from now on. We look at S_i, and choose in it r_A and r_B such that the paths in A_{i-1} and B_{i-1} from r to r_A and r_B respectively do not contain other vertices in S_i; strictly speaking, $r_A = r_{A,i}$ and r_A is obtained by picking an arbitrary vertex of $x \in S_i$ and following the path of A_{i-1} from r to x until we meet the first vertex of S_i (so, if $r \in S_i$, $r_A = r$). The existence of r_B is obtained similarly. Start with A_i being A_{i-1} after removing all the arcs entering S_i other than, if $r \notin S_i$, the arc entering r_A. Similarly, start with B_i being B_{i-1} after removing all the arcs entering S_i other than, if $r \notin S_i$, the arc entering r_B. Further remove from A_i and from B_i all the arcs that have both endpoints in $S_i \cup Q_i$. Note also that no arc of $A_{i-1} \cup B_{i-1}$ enters Q_i from outside S_i, by the way Q_i and S_i are constructed. For $i = 1$, the example in Figure 2 could have $r_A = x_4$ and

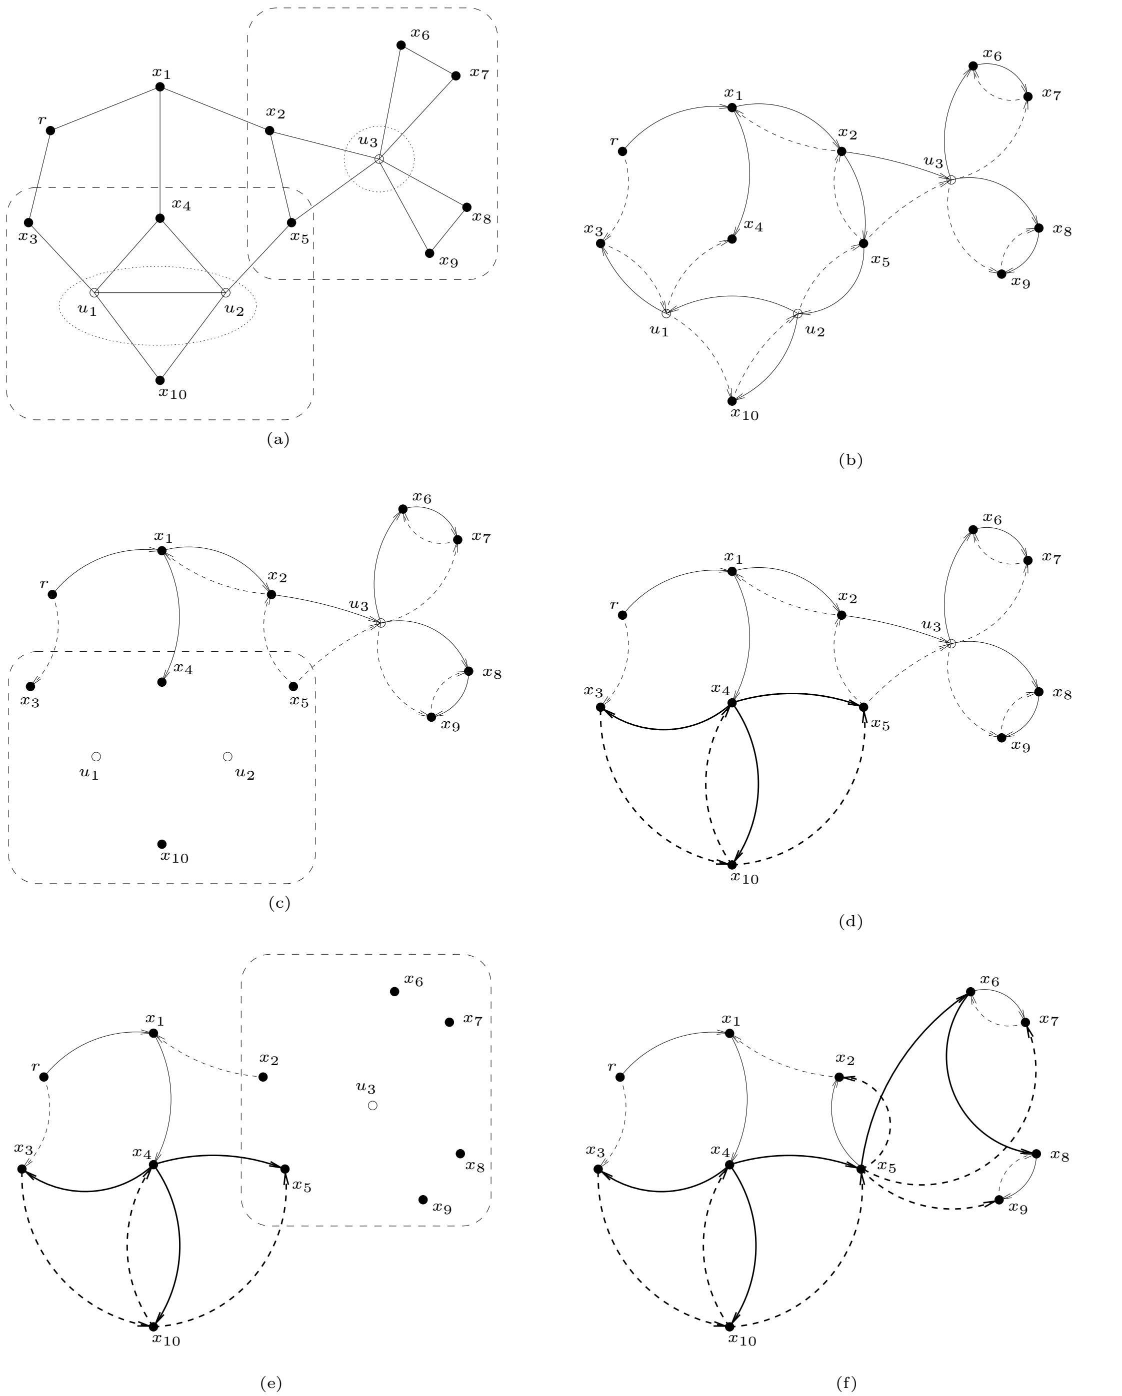

Figure 2: An example giving the construction of arborescences A (solid arc segments) and B (dashed arc segments) in the proof of Theorem 1. Figure (a) shows a (possible) optimum solution, where the input set S is given by the black disks, and the solution OPT is given by empty circles, with line segments giving the graph $U(OPT \cup S)$. The sets Q_i and S_i are given by dashed ellipses and rounded rectangles minus the dashed ellipses, respectively. The thin and thick arc segments represent some of the arcs of the bidirected version of G that have weight 0 and 1 respectively. Further explanation is in the text.

$r_B = x_3$, with subfigure (c) depicting A_i and B_i at this moment in the proof. For $i = 2$, we may have $r_A = r_B = x_5$ with subfigure (e) depicting A_i and B_i at this moment in the proof.

Construct sets X, Y, Z as follows. Initially, $X = Y = \emptyset$, and $Z = S_i \setminus \{r_A, r_B\}$. If $r_A \neq r_B$, set $r'_A = r_A$ and $r'_B = r_B$. This is the case in the example of Figure 2 for $i = 1$. If $r_A = r_B$ and there are at least two distinct nodes of Z that are U-adjacent to r_A, then arbitrarily pick two such nodes r'_A and r'_B and remove them from Z. If $r_A = r_B$, and there is at most one node of $v \in Z$ that is U-adjacent to r_A, then remove v (if it exists) from Z, and set $r'_A = r_A$ and $r'_B = r_B$. This is the case in the example of Figure 2 for $i = 2$, with $v = x_2$. Regardless of what case was before, if r'_A has U-neighbors in Z, remove them from Z and place them in Y; otherwise place r'_A in Y. If r'_B has U-neighbors in Z, remove them from Z and place them in X; otherwise place r'_B in X. Assume two vertices $x, y \in Z$ are U-adjacent. Put x in X and y in Y, taking both out of Z. Do this as long as possible, and note that Z at the end is an independent set of $U(S)$, and that X, Y, Z are disjoint. In the example of Figure 2 for $i = 1$, we end up with $X = \{x_3\}$, $Y = \{x_4\}$ and $Z = \{x_5, x_{10}\}$. For $i = 2$, we may end up with $X = \{x_6, x_8\}$, $Y = \{x_7, x_9\}$, and $Z = \emptyset$.

As $U(Q_i)$ is connected, and every vertex in S_i is U-adjacent to a vertex of Q_i, Q_i is a MSPT feasible solution for input $\{r'_A\} \cup X \cup Z$. Use the tight analysis for MSPT proposed by (24) to get a tree T_A connecting $\{r'_A\} \cup X \cup Z$, using edges of G of total weight at most $4|Q_i|$. Orient the edges of T_A away from r'_A. If $r'_A \neq r_A$, add the arc $r_A r'_A$. Every vertex $y \in (Y \setminus \{r'_A\})$ has a U-neighbor $x \in (X \cup \{r'_A\})$; add to T_A the arc xy of weight 0, using $x = r'_A$ whenever possible. If $r'_B \neq r'_A$ and $r'_B \notin X$, then r'_B has a U-neighbor $x \in (X \cup \{r'_A\})$; add to T_A the arc $x r'_B$ of weight 0. In the case (mentioned earlier when X, Y, Z were constructed) that $r_A = r_B$ and there was exactly one node of Z that is U-adjacent to r_A (this node is called v and was removed from Z), add to T_A the arc $r_A v$ of weight 0. In the example of the Figure 2, compare subfigures (c) and (d) with T_A possibly having the arcs x_4x_3, x_4x_5, and x_4x_{10}, and compare subfigures (e) and (f) with T_A possibly starting with the arcs x_5x_6 and x_6x_8, and to which later were added the three arcs x_6x_7, x_8x_9, and x_5x_2, all three of weight 0.

Similarly, note that Q_i is a MSPT feasible solution for input $\{r'_B\} \cup Y \cup Z$. Obtain T_B connecting $\{r'_B\} \cup Y \cup Z$ of weight at most $4|Q_i|$. Orient the edges of T_B away from r'_B. If $r'_B \neq r_B$, add the arc $r_B r'_B$. Every vertex $x \in (X \setminus \{r'_B\})$ has a U-neighbor $y \in (Y \cup \{r'_B\})$; add to T_B the arc yx of weight 0. If $r'_B \neq r'_A$ and $r'_A \notin Y$, then r'_A has a U-neighbor $y \in (Y \cup \{r'_B\})$; add to T_A the arc $y r'_A$ of weight 0. In the case (mentioned earlier when X, Y, Z were constructed) that $r_A = r_B$ and there was exactly one node of Z that is U-adjacent to r_A (this node is called v and was removed from Z), add to T_B the arc (of weight 1): $r_B v$. The total weight of the arcs added is at most $8|Q_i| + 1 \leq 9|Q_i|$. In the example of the Figure 2, compare subfigures (c) and (d) with T_B possibly having the arcs x_3x_{10}, $x_{10}x_4$, and $x_{10}x_5$, and compare subfigures (e) and (f) with T_A possibly starting with the arcs x_5x_7 and x_5x_9, and to which later were added the arcs x_7x_6, x_9x_8, both of weight 0, and the arc x_5x_2 of weight 1. One may have avoided the "+1" in this last example by picking $r_A = x_2$ for $i = 2$, but there are instances where the "+1" cannot be avoided, as it can be seen by running the algorithm on the example from Figure 1, with the output shown in Figure 3.

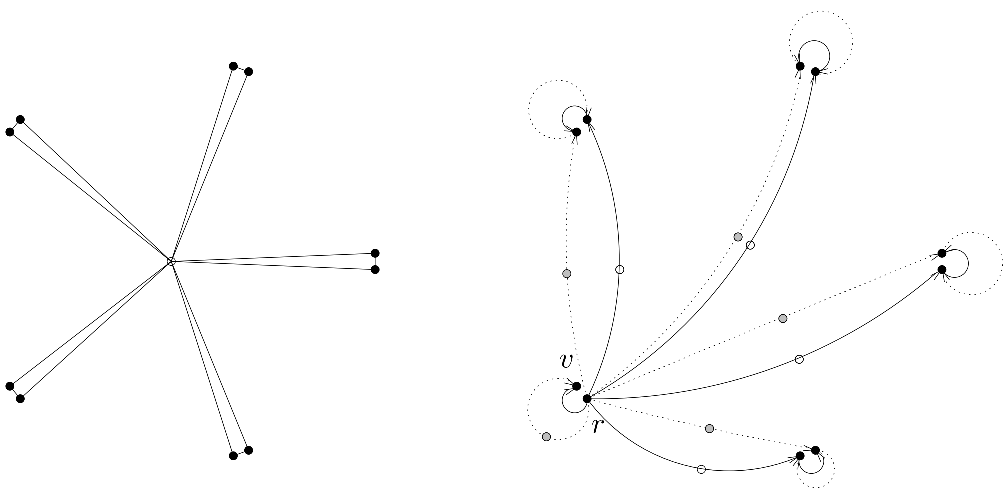

Figure 3: On the left, we have the optimal solution of Figure 1. As there, S is given by the black disks. On the right, there is a possible output of **Algorithm KV**, with the two arborescences given by solid and dotted arcs, and beads empty or dotted. One arc of weight $w = 1$ from r to v is required by one of the two arborescences.

Add T_A to A_i and T_B to B_i. Note that T_A indeed reaches every vertex of S_i, and thus A_i indeed contains a r-rooted arborescence in D_i. Indeed, any directed path P in A_{i-1} from r to some vertex $x \notin Q_i$ can be replaced by the following path in A_i (assuming P uses vertices of $Q_i \cup S_i$): if y is the last vertex of P in $Q_i \cup S_i$, then we use A_{i-1} to get from r to r_A, then in T_A from r_A to y (as $y \in S_i$ and $y \notin Q_i$, by the method Q_i and S_i were constructed), followed by the part of P from y to x. Similarly, note that T_B indeed reaches every vertex of S_i, and thus B_i indeed contains a r-rooted arborescence in D_i.

The theorem follows from the following claim:

Claim 2. *T_A and T_B are arc disjoint.*

Proof: Only the arcs of weight 0 are an issue; for the others arcs, parallel arcs are allowed. With Z independent in $U(S)$ (indeed, $U(Z)$ has no edges), and X, Y, Z disjoint, all the 0-weight arcs in T_A are in one of the following categories:

1. $r_A r'_A$, if $r_A \neq r'_A$
2. from Z to X
3. from X to $Z \cup X \cup Y$
4. from r'_A to Y (in this case, $r'_B \neq r'_A$ and $r'_A \notin Y$)
5. from r_A to v, if v exists
6. from $X \cup \{r'_A\}$ to r'_B (in this case, $r'_B \neq r'_A$ and $r'_B \notin X$)

(note that $r'_B \in X$ is possible and T_A can have arcs out of r'_B; however these arcs, other than $r_A v$, would have positive weight, as argued next. The case $r'_B \in X$ only happens if we had that r'_B had no U-neighbors in Z, and if r'_B does have a U-neighbor $y \in Y$, then y is also a U-neighbor of r'_A and the construction of T_A uses the arc $r'_A y$ instead of $r'_B y$). The 0-weight arcs in T_B are in one of the following categories:

1. $r_B r'_B$, if $r_B \neq r'_B$
2. from Z to Y
3. from Y to $Z \cup Y \cup X$
4. from r'_B to X (in this case, $r'_B \neq r'_A$ and $r'_B \notin X$)
5. from $Y \cup \{r'_B\}$ to r'_A (in this case, $r'_B \neq r'_A$ and $r'_A \notin Y$)

(note that $r'_A \in Y$ is possible and T_B can have arcs out of r'_A; however these would have positive weight, since if $r'_A \in Y$, then we had that r'_A had no U-neighbors in Z, which at that time contained all the nodes of $S_i \setminus \{r_A, r_B, r'_A, r'_B, v\}$). □

Thus T_A and T_B are indeed arc disjoint, and therefore we conclude that the approximation ratio of **Algorithm KV** is at most 9. Without a post-processing phase removing redundant nodes, the example of Figure 3 shows that the approximation ratio is exactly 9. □

3. Biconnectivity

Here (17) use the approximation algorithm of Khuller and Raghavachari (19). We could use this algorithm, but prefer a variant of the Khuller and Raghavachari algorithm proposed by Auletta, Dinitz, Nutov, and Parente (1), since this variant has a somewhat nicer analysis, and also is faster by an order of $|S|$. It is this variant that we call **Algorithm KR**, and go deeper in the algorithm to obtain a better approximation ratio.

A digraph is said to be *k-outconnected* (short for *k-vertex-outconnected*) from a vertex s if it contains k internally vertex-disjoint paths from s to any other vertex. The min-weight spanning subdigraph of a given weighted digraph which is k-outconnected from a specified vertex, if such a digraph exists, can be found in polynomial time by an algorithm of Frank and Tardos (11).

For any digraph D, we use $\overline{D}$ to represent the undirected graph obtained from D by ignoring the orientations of the arcs and then removing multiple edges between any pair of nodes. Suppose that D is a 2-outconnected digraph from a vertex s in which s has exactly two outgoing neighbors. Then the graph $\overline{D}$ is biconnected (1). **Algorithm KR** constructs a biconnected spanning subgraph of a given complete weighted graph G as follows.

1. For all $s \in V$
   (a) Bidirect G, and add to the weight of the arcs leaving s an integer $M > 2\sum_{uv \in G} w(u,v)$. The resulting digraph is denoted by $G^+(s)$.
   (b) Find a minimum-weighted spanning subdigraph $D(s)$ of $G^+(s)$ which is 2-outconnected from s (We can assume no arc of $D(s)$ enters s. Moreover, M is large enough that s is incident to exactly two arcs in $D(s)$).
   (c) Store the graph $\overline{D}(s)$
2. output the $\overline{D}(s)$ that has minimum weight in G

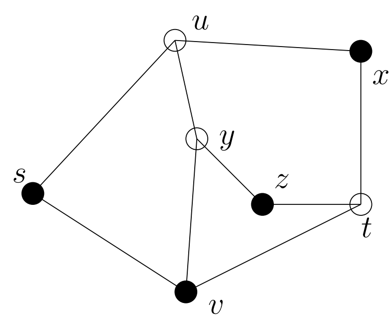

Figure 4: u, y, z, t is a chord-path between u and t for the cycle s, u, x, t, v. As an example for Lemma 4, let vertices of R be represented by black disks, and the vertices not in R by empty small circles. Then the cycle s, u, x, t, v has in u, y, v a chord-path without a vertex of R.

(1) prove that this modified version is a 2-approximation for MINIMUM-WEIGHT SPANNING BICONNECTED SUBGRAPH. We use as the input of **Algorithm KR** (the graph G, above) the complete simple graph on S with weight w defined by (similar to the previous section) $w(x, y) = \max(0, \lceil d(x, y) \rceil - 1)$ (for all $x \neq y \in S$), where $d(u, v)$ denotes the distance from u to v. One can easily verify that $w(x, y)$ is the minimum number of relay nodes required to connect x and y, and that $w(x, y)$ is an increasing function of $d(x, y)$. As noted before, w is not a metric.

After running **Algorithm KR** , replace each edge of positive weight of $\overline{D}(s)$ by new beads (that is, every such edge has its own distinct beads); this is the output of the algorithm for BICONNECTED RELAY PLACEMENT. The approximation ratio of $2d_{MST}$ obtained by (17) is based on showing that G has a biconnected subgraph of weight at most $d_{MST} \cdot opt$, where *opt* is the value of an optimum relay solution, and the fact that **Algorithm KR** is a 2-approximation for MINIMUM-WEIGHT SPANNING BICONNECTED SUBGRAPH. For us, the approximation ratio of d_{MST} follows from Theorem 10, proven in the next subsection after preliminary lemmas.

3.1. Analysis for arbitrary normed spaces

First, a corollary of Menger's Theorem:

Theorem 3 (Fan Lemma). *(see, for example, (8)) Suppose that G is a k-vertex connected undirected graph and X is a proper subset of its vertices with $|X| = k$. Then for any vertex v not in X, there are k paths that link v to vertices of X, and the only vertex appearing on more than one path is v.*

Luebke (22) and Luebke and Provan (23) proved an equivalent of the following property when minimizing the total distance. We use their approach below. Given a cycle C in an undirected graph H and two distinct vertices u and v on C, a *chord-path* between u and v is a path $\bar{P}$ in H between u and v that, except for u and v, shares neither vertices nor edges with C. See Figure 4 for an illustration.

Lemma 4. *Let J be a biconnected simple undirected graph and R be a subset of $V(J)$ with $|R| > 2$. Assume no proper biconnected subgraph J' of J exists such that $R \subseteq V(J')$. Then for every cycle in J, any chord-path has in its interior a vertex of R. Every simple cycle of J contains at least two vertices of R, and if it contains exactly two vertices, then each of the two paths contained in the cycle between these two vertices has an interior vertex of degree greater than two. Moreover, there exists a vertex of R of degree 2 in J.*

Proof: First assume, for a contradiction, that a cycle C has a chord-path $\bar{P}$ with no vertices of R in its interior. A biconnected graph J has an *ear-decomposition* ((10), Theorem 2.1.6). That is, J can be built up from a simple cycle by sequentially adjoining edges (loops are not allowed) and subdividing edges (in any order). We call the paths added (the subdivided edges) *ears*. By following the proof, we can see that one can choose the starting cycle and a number of ears arbitrarily, and after this continue to finish the ear decomposition.

We start an ear decomposition of J with C, followed by $\bar{P}$, and we obtain that there exist ear decompositions of J that have ears without vertices of R in the interior. Among all such ear decompositions, choose a decomposition with a shortest possible ear without vertices of R in the interior. Let C' be the first cycle of the ear decomposition, and $\bar{P}_1, \ldots, \bar{P}_k$ be the ears added. Let $\bar{P}_i$ be the ear without vertices of R. If no further ear uses internal vertices of $\bar{P}_i$, then we simply do not add $\bar{P}_i$, adding all the other ears, and obtain a smaller graph, biconnected and spanning R, a contradiction.

If ear $\bar{P}_j$, for $j > i$, is with the smallest j such that it has as one endpoint a vertex x interior of $\bar{P}_i$, (note that $x \notin R$), then we have two cases. If vertex y, the other end of $\bar{P}_j$, is not on $\bar{P}_i$ (including the ends), then we split $\bar{P}_i$ in two paths $\bar{P}'$ and $\bar{P}''$ at x. We do not use $\bar{P}_i$ in the ear decomposition, and have $\bar{P}_{i+1}$ follow either $\bar{P}_{i-1}$, or C' if $i = 1$. We join $\bar{P}'$ with $\bar{P}_j$ and use it as an ear instead of $\bar{P}_j$. Immediately after this new $\bar{P}_j$, and before $\bar{P}_{j+1}$, we use $\bar{P}''$ as an ear. We get another ear decomposition of the same graph that has in $\bar{P}''$ an ear shorter than $\bar{P}_i$ without vertices of R in its interior, contradicting the choice of the original ear decomposition. If y, the other end of $\bar{P}_j$, is on $\bar{P}_i$ (including its ends), then we split $\bar{P}_i$ in three paths $\bar{P}'$, $\bar{P}^*$, and $\bar{P}''$, at x and y. $\bar{P}^*$ is not trivial since $x \neq y$ (loops are not allowed in the ear decomposition), and one of $\bar{P}'$ and $\bar{P}''$ is not trivial (the other path could have no edges). We use instead of $\bar{P}_i$ the concatenation of $\bar{P}'$, $\bar{P}_j$, and $\bar{P}''$, and instead of $\bar{P}_j$ the non-trivial path $\bar{P}^*$. We get another ear decomposition of the same graph that has in $\bar{P}^*$ an ear shorter than $\bar{P}_i$ without vertices of R in its interior, contradicting the choice of the original ear decomposition. Thus every ear must contain in its interior a vertex of R, and so does every chord-path of every cycle of J.

Second assume, for a contradiction, that C' is a cycle with one vertex of R. Let us choose another vertex of R and two internally disjoint paths between these two vertices of R as the starting simple cycle in the ear decomposition. Let $\bar{P}_i$ be the last ear in the ear decomposition to contain an edge of C'. Then $\bar{P}_i$ must be a part of C', as its ends must be vertices of C' either from the first cycle of the ear decomposition or from an ear added before $\bar{P}_i$. Thus $\bar{P}_i$ is an ear without interior vertices from R. We argued above that no such ears exist and therefore we reached a contradiction. If C' has no vertex of R, a similar argument works, after starting with another arbitrary cycle containing vertices of R.

If C' has exactly two vertices of R, say, x, y, and a path P contained in C' from x to y that does not contain a vertex of degree greater than two, denote by P' the other path of C' between x and y. Choose another vertex of v of R, and based on the Fan Lemma (Theorem 3), there are two paths of J from v to x and to y that only share vertex v. These two paths are disjoint from the internal vertices of P (if any) since these vertices only have degree two, and the two paths cannot both contain x or y. One can then start the ear decomposition with the cycle made by these two paths together with P (note that this cycle is not C' since v is not a vertex of C') followed by a non-trivial part of P' as an ear without interior vertices of R. We argued above that no such ears exist and therefore we reached a contradiction.

Finally, note that the last ear added in any ear decomposition must contain in its interior a vertex of R, and this vertex has degree 2 in J. □

From here it is immediate to deduce the following:

Corollary 5. *Let J be a biconnected simple undirected graph and R be a subset of $V(J)$ with $|R| > 2$. Assume no proper biconnected subgraph J' of J exists such that $R \subseteq V(J')$. Let X_i be a connected component of the subgraph of J induced by $V(J) \setminus R$. Let R_i be the set of vertices of R adjacent to some vertex in X_i, and let T_i be the the subgraph of J induced by $R_i \cup X_i$. Then T_i is a tree (called* full Steiner component*) with all the leafs in R_i.*

We also need the next lemma, giving a maximum degree condition that is claimed and used in (17) (Property 4.3). The journal version of (17) sidesteps this property and thus we include its proof.

Lemma 6. *Assume Q is minimal such that $U(S \cup Q)$ is biconnected. In $U(S \cup Q)$, there is a biconnected subgraph such that every vertex of Q has degree at most d_{MST}.*

Proof: We follow the proof from Robins and Salowe (27), with extra work since biconnected subgraphs are harder than trees. Assume for the moment that no two edges of $U(S \cup Q)$ have exactly the same distance. Let H be a biconnected subgraph of $U(S \cup Q)$ of minimum total distance, defined as $\sum_{xy \in E(H)} d(x, y)$.

Assume now that H does have a vertex $x \in Q$ of degree more than d_{MST}. We will obtain a contradiction, as follows. First, we show (whole argument taken from (27)) that x has neighbors in H: y and z, such that $d(y, z) < d(x, z)$. Let $y_1, y_2, \ldots, y_k$ be the neighbors of x in H. Draw a ball of radius ϵ (using distance d) around x, where $\epsilon < d(x, y_i)$ for all i. Let y'_i be the intersection of a segment xy_i with the boundary of this ball. Since k exceeds the Hadwiger number of the unit ball in the normed space, there exist i, j with $d(y'_i, y'_j) \leq \epsilon$. Assume by symmetry that $d(x, y_i) \leq d(x, y_j)$. Drawing the ball of radius $d(x, y_i)$ around x; and let y''_j be the point where the segment xy_j used for finding y'_j intersects the boundary of this bigger ball. Then we also have $d(y_i, y''_j) \leq d(x, y_i)$, and therefore

$$d(x, y_j) = d(x, y''_j) + d(y''_j, y_j) = d(x, y_i) + d(y''_j, y_j) \geq d(y_i, y''_j) + d(y''_j, y_j) \geq d(y_i, y_j),$$

and in fact (there are no ties) $d(x, y_j) > d(y_i, y_j)$; use y_i as y and y_j as z, and then $d(y, z) < d(x, z)$ as desired.

Following standard notation (10), a block of a graph L is a maximal biconnected subgraph of L (assuming that the complete simple graph with two vertices is biconnected). Let $W_1, W_2, \dots, W_k$ be the blocks of L. Form a bipartite graph $T = (V, N, F)$, where $V = V(L)$, the elements of $N = \{b_1, b_2, \dots, b_k\}$ correspond to the blocks of L, and F connects v_i to b_j if in L $v_i \in W_j$.

Then we have ((10), Theorem 2.1.13): The blocks of the graph $L = (V, E)$ partition the set E of edges. Two edges belong to the same block if and only if there is a simple cycle containing both. Any two blocks have at most one node in common, and the nodes belonging to more than one block are cut nodes. The graph T is a tree, called the *block-vertex* tree of L.

Remove xz from H obtaining the graph H', and let $W_1, W_2, \dots, W_k$ be the blocks of H'. Construct the block-vertex tree $T' = (V', N, F)$ of H'. In fact, all the vertices of N (corresponding to the blocks of H') are on the path from x to z in T, or otherwise adding xz to H' will not result in a single block for H. There is a block W' of H' that contains edge xy; renumber such that $W' = W_1$. The vertex of b_1 of T (corresponding to W_1) is on the path from x to z in T.

In fact, b_1 must be the first vertex on this path as it is adjacent to x in the tree T. If W_1 only has x and y, then the path from x to z in T starts with x, b_1, y, followed by the vertex b corresponding to some other block, say, W_2. We can see that removing x (who is not in S) and adding the edge yz again merges together all the blocks of H' (except W_1); thus one can get a biconnected subgraph of H contradicting the minimality of Q.

If W_1 has more than two vertices, and y is not on the path in T from z to x, then adding to H' the edge yz results in a single block again, and since $d(y, z) < d(x, z)$, we obtain a biconnected graph of shorter total distance than H, a contradiction.

We are left with the case that W_1 has more than two vertices, and y is on the path in T from z to x. Remove xy and W_1 will split into blocks $W'_1, W'_2, \dots, W'_{k'}$, all with corresponding vertices on the path from x to y in T', the block-vertex tree of H' with edge xy removed. The other blocks of H' are not affected, and therefore all the vertices of T' corresponding to all the blocks of H' with edge xy removed are on the path from x to z in T'. Adding xz back results in a biconnected graph of shorter total distance than H, a contradiction.

It is only left to justify how we assumed that no two distances among vertices of $U(S \cup Q)$ are equal (even if initially $S \cup Q$ is a multiset). First, if $1 + \epsilon$ is the smallest $d(x, y)$ for those $x, y \in S \cup Q$ with $xy \notin E(U(S \cup Q))$, replace ϵ by $\min(\epsilon, 1)$, and then shrink all the distances by $1 + \epsilon/2$, obtaining distance function d'. Let $U'(Z)$ be the unit-distance graph of set of points Z with respect to d'. Note that $U'(S \cup Q)$ is isomorphic to $U(S \cup Q)$, and also, all $x, y \in S \cup Q$ with $xy \in E(U(S \cup Q))$ have $d'(x, y) \le 1 - \epsilon/3$. Part of the following argument is taken directly from (27). Make a small random perturbation, replacing point x with point x' such that $d'(x, x') \le \epsilon/6$. Note that $U'(S' \cup Q')$ is isomorphic to $U'(S \cup Q)$, where $Z' = \{x' | x \in Z\}$. With respect to $S' \cup Q'$, all pairwise d'-values are distinct with probability 1. Use d' and $U'(S' \cup Q')$ at the beginning of this proof, as we only need that $U'(S' \cup Q') = U(S \cup Q)$, and that we have a normed space with distance d'. $\square$

Now, it will be nice if we had "parsimony", described in the introduction and proven later in Theorem 12. However we are unable to prove or disprove Theorem 12 without using planarity, and in three dimensions surely we cannot count on planarity. We do have Lemma 8 below, weaker than Theorem 12 in two respects: the solution is "fractional", and 2-outconnectivity replaces biconnectivity. It will be enough for our purpose.

Given a digraph D and X, Y disjoint sets of $V(D)$, define $E(X, Y)$ to be the set of arcs in $E(D)$ with tail in X and head in Y. Given a digraph D and $s \in V(D)$, consider the polytope $\mathcal{P}(D, s)$ in $\mathbb{R}^{|E(D)|}$ (with vectors β having entries $\beta(e)$ for all e arcs of D) defined by the constraints:

$$0 \le \beta(e) \le 1 \qquad \forall\, e \in E(D) \tag{5}$$

$$\sum_{e \in E(V(D)\setminus X, X)} \beta(e) \ge 2 \qquad \forall\, \emptyset \ne X \subseteq (V(D) \setminus \{s\}) \tag{6}$$

$$\sum_{e \in E(V(D)\setminus(\{z\}\cup X), X)} \beta(e) \ge 1 \qquad \forall\, z \ne s \;\; \forall\, \emptyset \ne X \subseteq (V(D) \setminus \{s, z\}) \tag{7}$$

Using Menger's theorem, one can check that, for a simple digraph D and for an integral vector β valid for $\mathcal{P}(D, s)$, the set E' of arcs e of $E(D)$ with $\beta(e) = 1$ is such that the digraph $(V(D), E')$ is 2-outconnected from s. Thus one can think of a valid vector β as being "fractional-2-outconnected".

Our big hammer is Theorem 17.1.14 of (10), (given there with more complicated notation as it solves k-outconnectivity), given below and originally from Frank and Tardos (11).

Theorem 7. *(11) For a simple digraph D, the linear system above giving $\mathcal{P}$ is Total Dual Integral, which implies that for any $c : E(D) \to \mathbb{N}$, if the linear program*

$$[\textit{Minimize} \sum_{e \in E(D)} c_e \beta(e) \ \textit{subject to}\ \beta \in \mathcal{P}(D, s)]$$

has a valid optimum, it has an integer-valued optimum.

To use this deep theorem, which is also at the basis of **Algorithm KR**, we need the following "outconnected fractional parsimony" lemma:

Lemma 8. *Let J be a biconnected undirected graph, and $\overrightarrow{J}$ be its bidirected version. Let R be a subset of $V(J)$ with $|R| > 2$. Then there exists vertex $s \in R$, and there exist positive real numbers α_i and a set of directed paths P_i of $\overrightarrow{J}$, all starting and ending at a vertex of R and without interior vertices from R, with the following properties. For every arc of $e \in E(\overrightarrow{J})$,*

$$\sum_{i \mid e \in E(P_i)} \alpha_i \le 1$$

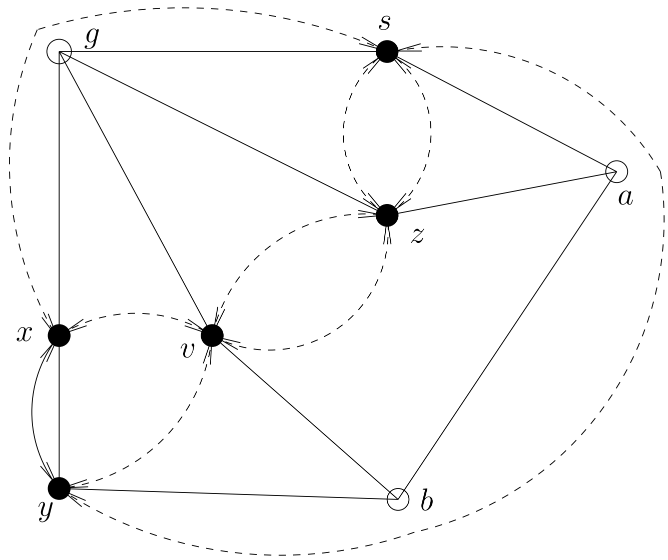

Figure 5: Illustrating Lemma 8 (this graph, being planar, also allows Theorem 12). J is shown with straight segments, with the vertices of R represented by black disks, and the vertices not in R by empty small circles. We have two full Steiner components, one with interior vertex g and leafs (which all must be vertices of R): s, z, v, x, and a second one with interior vertices a and b and leafs s, z, v, y. Some of the paths P_i are following some Eulerian traversal of the bidirected version of each full Steiner component. For example, for the component with interior vertex g, the Eulerian traversal can be $s, g, z, g, v, g, x, g, s$, giving the directed paths s, g, z, z, g, v, v, g, x, and x, g, s, as well as the reverse of each path, that is, z, g, s, v, g, z, x, g, v, and s, g, x. For the component with interior vertices a and b, we could have directed paths z, a, b, v and v, b, y and y, b, a, s and s, a, z, as well as the reverse each path, v, b, a, z etc. Each of these directed paths is replaced by an arc of K, which is shown in the figure by a (curved) arc. Other paths P_i are obtained from edges of J with both endpoints in R. Here, from the edge xy with both endpoints in R, we get two paths, each with one arc: x, y, and y, x. In the proof, we always get the antiparallel arc of every arc of K; therefore we put double arrows on our curved arcs. Antiparallel arcs get the same α-value, and we used dashed arcs when $\alpha = 1/2$ and solid arcs when $\alpha = 1$.

For all i, replace each P_i by an arc e_i joining the start and the end vertex of P_i, obtaining a directed multigraph K with vertex set R (and an edge set we call E). Let $\alpha(e_i) = \alpha_i$. Then the vector α is feasible for $\mathcal{P}(K, s)$ and $\sum_{e \in E(\{s\}, R\setminus\{s\})} \alpha(e) = 2$. Moreover, for all $x, y \in R$, $\sum_{e \in E(\{x\},\{y\})} \alpha(e) \leq 1$.

Figure 5 illustrates the lemma and its proof idea.

Proof: Remove edges and vertices not in R from J until it satisfies the conditions of Lemma 4 and Corollary 5. We pick s to be some vertex of R of degree 2 in J, whose existence is guaranteed by Lemma 4. Let T_i be the full Steiner components (all of our full Steiner components have at least one vertex not in R, and no edge with both endpoints in R) given by Corollary 5. Do an Eulerian traversal of each bidirected $\overrightarrow{T}_i$ (as in Christofides' algorithm). Recall that the vertices of $R \cap V(\overrightarrow{T}_i)$ are leafs, and thus each such vertex is visited exactly once by this traversal, assuming we start at an interior vertex. If vertices u,

v of R appear in this traversal such that, after u, v is the next vertex of R (thus skipping the vertices not in R), have two directed paths P_j and P_l one from u to v and one from v to u, where P_j follows the traversal, while P_l is the reverse of P_j. Set $\alpha_l = \alpha_j = 1/2$.

For two vertices u and v of R that are adjacent in J, make (one-arc) directed paths P_j and P_l, one from u to v and one from v to u, with $\alpha_l = \alpha_j = 1$. One can check that for every arc $e \in E(\overrightarrow{J})$,

$$\sum_{i \mid e \in E(P_i)} \alpha_i \leq 1,$$

as argued next. Indeed, for an arc e of a bidirected full Steiner component, e appears in exactly two directed paths P_i: one path is a part of the Eulerian traversal, and one path is the reverse of a directed path P in the Eulerian traversal (precisely P that contains the arc antiparallel to e). Both these two paths have α-value of 1/2. Also, for an arc e of the bidirected J connecting two vertices of R, we get exactly one arc of K (with α-value of 1). Incidentally to this proof, we remark that Kashyap et al. (17) also do this Eulerian traversal (though they do not call it Eulerian, look at their Figure 2), but implicitly set $\alpha_i = 1$ for all i and then the equation above only holds with 2 as the RHS. Here is where we improve the approximation ratio by a factor of two.

For all i, replace each P_i by an arc e_i joining the start and the end vertex of P_i, obtaining a directed multigraph K with vertex set R (and an edge set we call E). Since s has degree 2 in J, then irrespective on how many of its neighbors in J belong to R, we have $\sum_{e \in E(\{s\}, R \setminus \{s\})} \alpha(e) = 2$ as required.

Next we prove that, for all $x, y \in R$, $\sum_{e \in E(\{x\},\{y\})} \alpha(e) \leq 1$. Let arbitrarily $x \neq y \in R$. There can be several arcs of K from x to y, and we have several cases. If x, y are adjacent in J, then one arc e with tail x and head y is assigned $\alpha(e) = 1$. In this case no full Steiner component can have both x, y among its leafs, according to Corollary 5. Therefore no arc parallel to e can exists in K.

In a second case, x and y are not adjacent in J and there is one full Steiner component with leafs only x and y. This component gives us exactly two arcs with tail x and head y, and each is assigned $\alpha = 1/2$. No other full Steiner component can have both x, y among its leafs, since then we would have a cycle of J (the paths between x and y in the two full Steiner components) with exactly two vertices of R and a path between these two vertices (namely, the path in the full Steiner component with only x, y as leafs) with no interior vertex of degree three or more in J, contradicting Lemma 4. Therefore K cannot contain other arcs with tail x and head y.

In a third case, x and y are not adjacent in J, and all the full Steiner component that have both x and y as leafs have at least three leafs. Then each such component can produce at most one arc with tail x and head y, and all these arcs are assigned $\alpha = 1/2$. However, we cannot have three such components, since from two of them we can obtain a cycle going through x and y, and from the third we obtain a chord-path without any vertex of R, contradicting Lemma 4. Thus in all cases $\sum_{e \in E(\{x\},\{y\})} \alpha(e) \leq 1$.

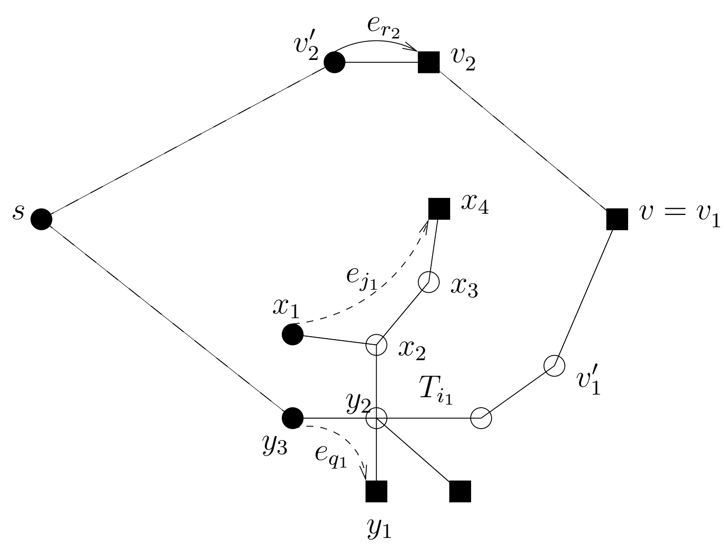

Figure 6: A subgraph of J is shown with straight segments, with vertices of R represented by either black filled squares (if in X) or by black disks (if not in X), and the vertices not in R represented by empty small circles. We use dashed segments to represent paths of J, with their interior vertices not depicted. In this case, the path P_{r_2} only has the edge $v_2'v_2$. The path P_{j_1} has vertices x_1, x_2, x_3, x_4, the path P_{k_1} has vertices y_1, y_2, y_3 and P_{q_1} has vertices y_3, y_2, y_1. The arcs of K used by the proof are represented by curved arrows, and have the values $\alpha(e_{r_2}) = 1$ and $\alpha(e_{j_1}) = \alpha(e_{q_1}) = 1/2$.

Once again incidentally, we mention that (17) implicitly obtain a similar K but put $\alpha = 1$ on all the arcs, while we use $1/2$ for all the arcs obtained from full Steiner components. This is where we improve the ratio - and this also explains why our proof is longer and more complicated.

Claim 9. *The vector α is feasible for $\mathcal{P}(K, s)$.*

Proof: Constraints (5) are immediate. We proceed to Constraint (6). Let $X \subseteq (V(D) \backslash \{s\})$ be arbitrary; arbitrarily pick $v \in X$. Going back to the undirected J, there are two internally-disjoint paths $\bar{P}_1'$ and $\bar{P}_2'$ from s to v. Let C be the cycle obtained from putting together $\bar{P}_1'$ and $\bar{P}_2'$. Let v_1 be the first vertex of X on $\bar{P}_1'$ ($v_1 = v$ possible), and v_2 be the first vertex of X on $\bar{P}_2'$ ($v_2 = v$ possible). Let v_1' be the vertex before v_1 on $\bar{P}_1'$ ($v_1' = s$ possible), and let v_2' be the vertex before v_2 on $\bar{P}_2'$ ($v_2' = s$ possible). Figure 6 provides an illustration.

If $v_1' \in R$ (note that $v_1 \in R$), then we have a directed path P_{r_1} from v_1' to v_1 in $\overrightarrow{J}$ with $\alpha_{r_1} = 1$, and then in K we have an arc e_{r_1} from v_1' to v_1 with $\alpha_{r_1} = 1$. Similarly, if $v_2' \in R$ (note that $v_2 \in R$), then we have a directed path P_{r_2} from v_2' to v_2 in $\overrightarrow{J}$ with $\alpha_{r_2} = 1$, and then in K we have an arc e_{r_2} from v_2' to v_2 with $\alpha_{r_2} = 1$. If $v_1' \notin R$, then there is a full Steiner component T_{i_1} that contains v_1' and that has leafs both in X and outside X. If $v_2' \notin R$, then there is a full Steiner component T_{i_2} that contains v_2' and that has leafs both in X and outside X. If we have both T_{i_1} and T_{i_2}, we remark that $i_1 \neq i_2$ since otherwise we obtain in J a chord-path for C with no internal vertex in R, contradicting Lemma 4.

The Eulerian traversal of $\overrightarrow{T}_{i_1}$ gives two directed paths, P_{j_1} and P_{k_1}, one entering X and one exiting X. Then $\alpha_{j_1} = 1/2$, and also there is another path P_{q_1}, the reversal of P_{k_1}, that also enters X and has $\alpha_{q_1} = 1/2$. The Eulerian traversal of $\overrightarrow{T}_{i_2}$ gives two directed paths, P_{j_2} and P_{k_2}, one entering X and one exiting X. Then $\alpha_{j_2} = 1/2$, and also there is another directed path P_{q_2}, the reversal of P_{k_2}, that also enters X and has $\alpha_{q_2} = 1/2$. Thus K contains either e_{r_1} with $\alpha_{r_1} = 1$, or both e_{j_1} and e_{q_1} with $\alpha_{j_1} = \alpha_{q_1} = 1/2$. Also K contains either e_{r_2} with $\alpha_{r_2} = 1$, or both e_{j_2} and e_{q_2} with $\alpha_{j_2} = \alpha_{q_2} = 1/2$. In all four subcases, Constraint (6) is satisfied.

We proceed to Constraints (7), which must hold $\forall\ z \neq s\ \ \forall\ \emptyset \neq X \subseteq (V(D) \setminus \{s, z\})$. Arbitrarily pick $v \in X$. Going back to the undirected J, there are two internally-disjoint paths $\bar{P}'_1$ and $\bar{P}'_2$ from s to v; assume by renaming $\bar{P}'_1$ and $\bar{P}'_2$ that z is not a vertex of $\bar{P}'_1$. Let C be the cycle obtained from putting together $\bar{P}'_1$ and $\bar{P}'_2$. Let v_1 be the first vertex of X on $\bar{P}'_1$ ($v_1 = v$ possible), and v_2 be the first vertex of X on $\bar{P}'_2$ ($v_2 = v$ possible). Let v'_1 be the vertex before v_1 on $\bar{P}'_1$ ($v'_1 = s$ possible), and let v'_2 be the vertex before v_2 on $\bar{P}'_2$ ($v'_2 = s$ or $v'_2 = z$ possible). If $v'_1 \in R$ (note that $v_1 \in R$), then we have a directed path P_{r_1} from v'_1 to v_1 in $\overrightarrow{J}$ with $\alpha_{r_1} = 1$, and then in K we have an arc e_{r_1} from v'_1 to v_1 with $\alpha_{r_1} = 1$. So, if $v'_1 \in R$, Constraint (7) is satisfied.

Assume from now on that $v'_1 \notin R$; therefore there is a full Steiner component T_{i_1} that contains v'_1 and that has leafs both in X and outside X. Consider the case when z is an interior vertex of $\bar{P}'_2$; then we cannot have that T_{i_1} has z as a vertex, since otherwise, in J, we get a chord-path of C with no internal vertex in R. One can look at Figure 6 for intuition, with $z = v'_2$. The Eulerian traversal of $\overrightarrow{T}_{i_1}$ gives two directed paths, P_{j_1} and P_{k_1}, one entering X and one exiting X. Then $\alpha_{j_1} = 1/2$, and also there is another directed path P_{q_1}, the reversal of P_{k_1}, that also enters X and has $\alpha_{q_1} = 1/2$. None of P_{j_1} and P_{k_1} and P_{q_1} start or end at z, since z is not a vertex of T_{i_1}. Constraint (7) is satisfied by $\alpha(e_{j_1})$ and $\alpha(e_{q_1})$.

From now on, z is not an interior vertex of $\bar{P}'_2$. If $v'_2 \in R$ (note that $v_2 \in R$ and $v'_2 \neq z$), then we have a directed path P_{r_2} from v'_2 to v_2 in $\overrightarrow{J}$ with $\alpha_{r_2} = 1$, and then in K we have an arc e_{r_2} from v'_2 to v_2 with $\alpha_{r_2} = 1$. So, if $v'_2 \in R$, Constraint (7) is satisfied.

We are left with the case $v'_1 \notin R$, z not on $\bar{P}'_2$, and $v'_2 \notin R$; recall that z is not on $\bar{P}'_1$. Please refer to Figure 7 for an illustration of this case. We have the full Steiner component T_{i_1} as above, and the full Steiner component T_{i_2} that contains v'_2 and that has leafs in both X and outside X. Note that $i_1 \neq i_2$ since otherwise we obtain, in J, a chord-path for C with no internal vertex in R.

Let v''_1 be the last vertex of R before v_1 on $\bar{P}'_1$ ($v''_1 = s$ is possible); then $v''_1 \in V(K) \setminus (X \cup \{z\})$. Consider the Eulerian traversal of $\overrightarrow{T}_{i_1}$; it passes through each vertex of $R \cap V(T_{i_1})$ exactly once (as these are the leafs of T_{i_1}). Then, in this traversal, we can get from v''_1 to v_1, or from v_1 to v''_1, without passing through z (which can be a leaf of T_{i_1}). Thus, we have that either a directed path P_{j_1} of this traversal goes from $V(K) \setminus (X \cup \{z\})$ to X, or goes from X to $V(K) \setminus (X \cup \{z\})$. In the second case, P_{q_1}, the reverse of P_{j_1} goes from $V(K) \setminus (X \cup \{z\})$ to X. Let P_{k_1} be either P_{j_1} or P_{q_1}, such that it goes from $V(K) \setminus (X \cup \{z\})$ to X. Thus e_{k_1}

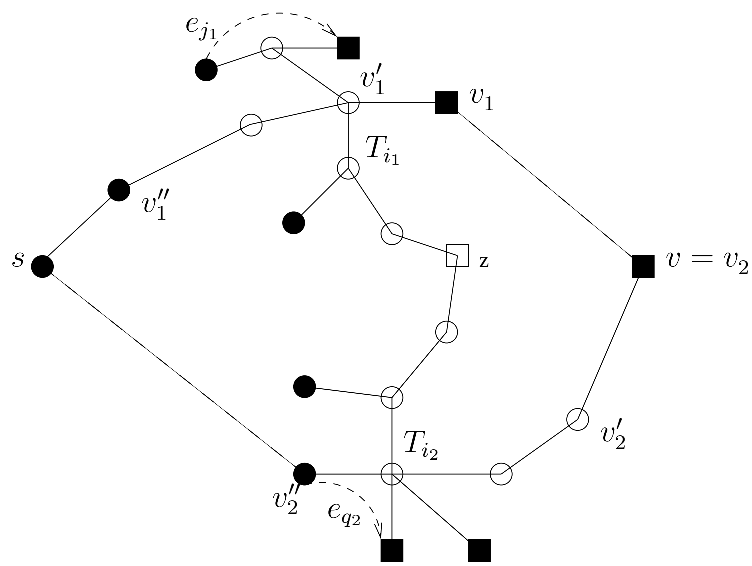

Figure 7: A subgraph of J is shown with straight segments, with vertices of R represented by black filled squares (if in X), by black disks (if not in $X \cup \{z\}$), and z is the empty square. The vertices not in R are represented by empty small circles. We use dashed segments to represent paths of J, with their interior vertices not depicted. The arcs of K used by the proof are represented by curved arrows, and $\alpha(e_{j_1}) = \alpha(e_{q_2}) = 1/2$.

exists in K; also $\alpha_{k_1} = 1/2$. We repeat the argument for T_{i_2}, to get another arc $e_{k_2} \in E(K)$ going from $V(K) \setminus (X \cup \{z\})$ to X, and with $\alpha_{k_2} = 1/2$. Constraint (7) is satisfied by $\alpha(e_{k_1})$ and $\alpha(e_{k_2})$, since $e_{k_1} \neq e_{k_2}$ (true even if they share endpoints).

In all cases, Constraints (7) are satisfied, finishing the proof of Claim 9. □

With this claim, the proof of Lemma 8 is complete. □

Theorem 10. *Let S be a multiset of at least three points making an instance of* Biconnected Relay Placement, *and Q be an optimum feasible solution. Let G be the weighted complete graph on S with weight w defined in the description of our algorithm. Then there exists $s \in V(G)$ such that $w(\overline{D}(s)) \leq d_{MST} \cdot |Q|$, where $\overline{D}(s)$ is computed by the algorithm in Line 1.c.*

Proof: Q is optimum and thus minimal such that $U(Q \cup S)$ is biconnected. Choose a biconnected spanning subgraph L of $U(Q \cup S)$ satisfying the hypothesis of Lemma 6. Apply Lemma 8 with L as J, and S as R, obtaining vertex $s \in S$, directed paths (P_i), non-negative numbers α_i, and arcs e_i giving multidigraph K. Use this s as the vertex of G required by the theorem.

For directed path P, define $\breve{P}$ to be the set of vertices in its interior. For an arc e of K with tail x and head y, define

$$c(xy) = \begin{cases} \min_{i \mid P_i \text{ starts at } x \text{ and ends at } y} |\breve{P}_i| & \text{if } x \neq s \\ M + \min_{i \mid P_i \text{ starts at } s \text{ and ends at } y} |\breve{P}_i| & \text{if } x = s \end{cases}$$

(this and the next definition give the same result for parallel arcs). For an arc e of K with tail x and head y, define $\alpha'(xy) = \sum_{i \mid P_i \text{ starts at } x \text{ and ends at } y} \alpha_i$.

Let K' be the simple digraph on S obtained from K by removing parallel arcs. Consider the linear program **LP**:

$$\text{Minimize} \sum_{xy \in E(K')} c(xy)\beta(xy) \quad \text{subject to } \beta \in \mathcal{P}(K', s).$$

Based on the fact that α from Lemma 8 belongs in the polytope $\mathcal{P}(K, s)$ and that $\sum_{e \in E(K)(\{x\},\{y\})} \alpha(e) \leq 1$, we obtain that α' belongs in the polytope $\mathcal{P}(K', s)$. Note that costs c are integral. Apply Theorem 7 to get an integral solution for **LP**, and therefore a digraph D', subgraph of K', 2-outconnected from s, satisfying $c(D') \leq \sum_{xy \in E(K')} c(xy)\alpha'(xy)$. Lemma 8 gives $\sum_{y \in S \setminus \{s\}} \alpha'(sy) = 2$ and therefore, taking the definition of $c()$ into account, we obtain:

$$c(D') \leq \left(\sum_i \alpha_i |\breve{P}_i| \right) + 2M. \tag{8}$$

Let $D(s)$ be a minimum-weighted subgraph of $G^+(s)$ which is 2-outconnected from s, as computed by the algorithm. Let w^+ be the weights of $G^+(s)$, that is

$$w^+(x, y) = \begin{cases} w(x, y) & \text{if } x \neq s \\ M + w(x, y) & \text{if } x = s \end{cases}$$

Note that for any edge e of G with endpoints u and v, and for any directed path P from u to v in $\overrightarrow{L}$, $w(e) \leq |\breve{P}|$, as beads can be placed on the vertices of $\breve{P}$. Thus, for any arc e of K with tail x and head y, we have $w^+(x, y) \leq c(xy)$.

Recall that $D(s)$ has no arc entering s and exactly two arcs leaving s. Thus

$$w(\overline{D}(s)) \leq w(D(s)) \leq w^+(D(s)) - 2M \leq w^+(D') - 2M \leq c(D') - 2M. \tag{9}$$

Write $e \diamond v$ if edge e is incident to vertex v, and write $next(P, v, e)$ if, on directed path P, one arc obtained from bidirecting edge e is used to leave v. We have:

$$\begin{aligned} w(\overline{D}(s)) &\leq \sum_i \alpha_i |\breve{P}_i| \\ &= \sum_i \quad \alpha_i \sum_{v \in \breve{P}_i} 1 \\ &= \sum_{v \in Q} \quad \sum_{i \mid v \in \breve{P}_i} \alpha_i \\ &= \sum_{v \in Q} \quad \sum_{e \mid e \diamond v} \quad \sum_{i \mid next(P_i, v, e)} \alpha_i \\ &\leq \sum_{v \in Q} \quad \sum_{e \mid e \diamond v} 1 \\ &\leq |Q| \cdot d_{MST} \end{aligned}$$

where the first inequality follows from Equations 9 and 8. The second inequality comes from Lemma 8, and the last inequality from Lemma 6. This finishes the proof. □

The analysis given above is tight. Precisely, in the two-dimensional Euclidean plane, the ratio of the biconnectivity algorithm above is indeed $5 - o(1)$, assuming all ties are broken in worst-case manner, and no post-processing removes redundant relay nodes. First look at the example in Figure 8. It has two *sea stars* (one relay node, the *star's center*, that is U-adjacent to five U-independent nodes of S, called *tentacles*) with u and v in the center. In general, we are going to use q spread-out sea stars (as, for example, in Figure 9), and we connect their tentacles as those of u, v are in Figure 8 - this can always be done while maintaining planarity to create a biconnected graph. Precisely, plane curves connect tentacles of different sea stars such that no two points on distinct curves are at distance at most 1. Each curve is subdivided such that only consecutive nodes on the curve are U-adjacent; the nodes used for subdivision are put in S. Done carefully, we end up with m paths, each giving a connected component of $U(S)$ (one for each curve), such that $m = 5q/2$ (assume q is even). Optimum is q. We use the following theorem of Whitty (29):

Theorem 11. *(29) Suppose that, given a directed graph $D = (V', E')$ and a specified vertex $s \in V'$, there are two internally vertex-disjoint paths from s to any other vertex of D. Then D has two arc-disjoint outgoing arborescences rooted at s such that for any vertex $v \in V' \setminus \{s\}$ the two paths to s from v uniquely determined by the arborescences are internally vertex-disjoint.*

Wherever we start with s in **Algorithm KR**, each of the two arborescences from the theorem above needs $m-1$ arcs of weight 1 to enter all of the m paths/connected components of $U(S)$, with the exception of the component containing s. Thus **Algorithm KR** produces a solution of weight at least $2(m-1)$, and with q large, this converges to $5q$. Other variants of **Algorithm KR** also produce solutions with $2m - o(1)$ relay nodes on this example.

In the three dimensional Euclidean space, one cannot assume any planar structure of the optimum, as explained below. Consider an even number q of far-apart sea stars in three dimensions, each with 12 tentacles. For each star, the set of first vertices of each of the 12 tentacles is an independent set of the unit-disk graph. These tentacles can be arbitrarily connected two-by-two in three dimensions while ensuring that the connected components of $U(S)$ are each one path; these paths are obtained by joining two tentacles of different sea stars. Precisely, curves in three dimensions connect tentacles of different sea stars such that no two points on distinct curves are at distance at most 1. From now on we use the same argument as in the two dimensional case. Thus in three dimensions, without a preprocessing step eliminating redundant nodes, the approximation ratio of **Algorithm KR** is at least $d_{MST} - o(1)$.

4. Stronger version of parsimony for planar graphs

Theorem 12. *Let J be a biconnected plane undirected graph, and $\overrightarrow{J}$ be its bidirected version. Let R be a subset of $V(J)$ with $R| > 1$. Then there exists a set of arc-disjoint directed paths*

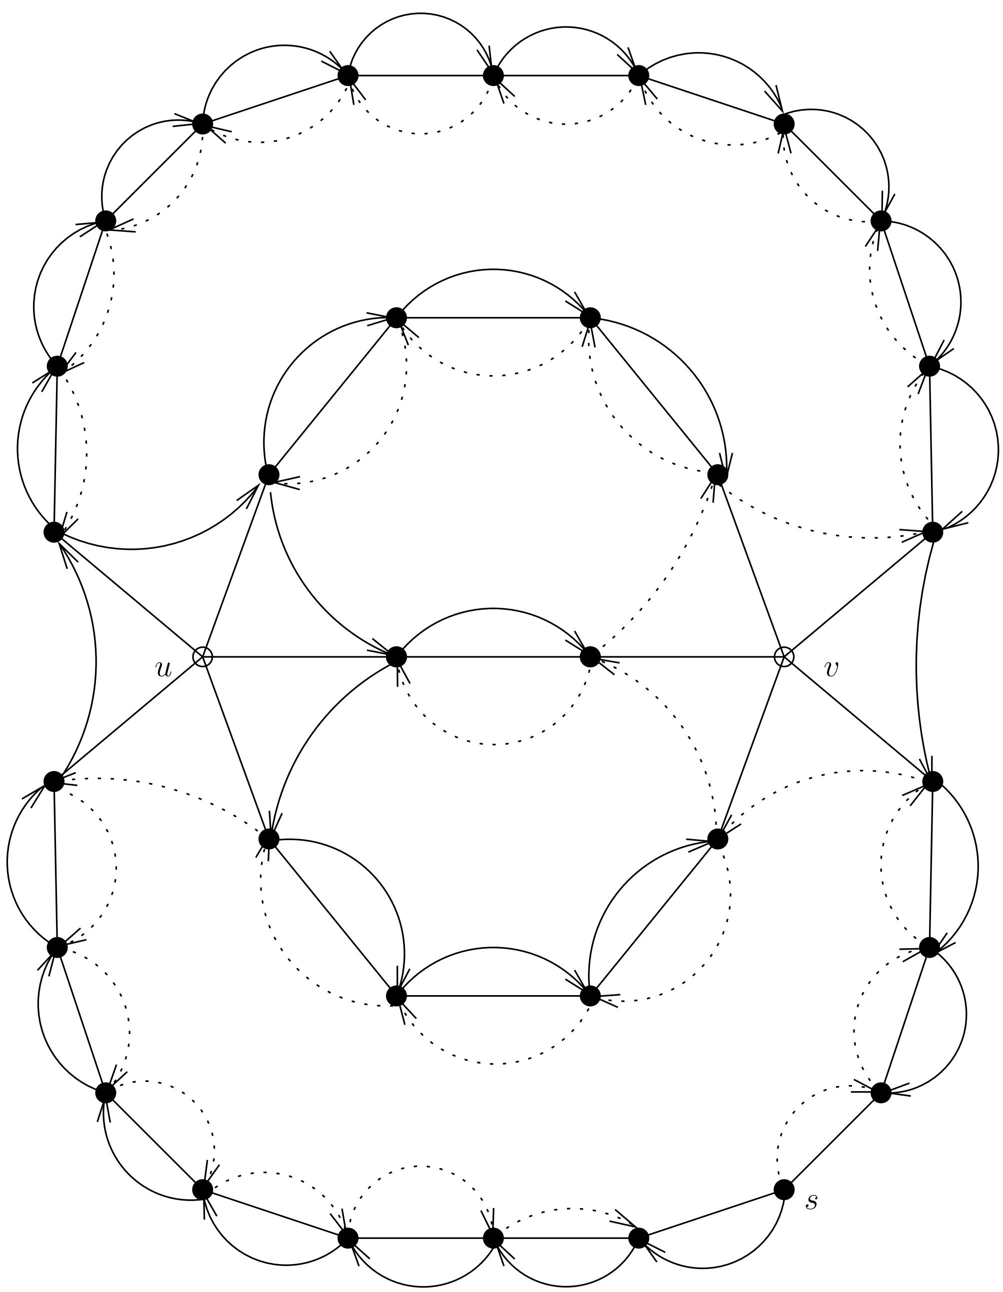

Figure 8: The nodes of S are black disks. Optimum uses the relay nodes u and v. If we start **Algorithm KR** with s as in the figure, ten edges of weight one would be chosen by the algorithm (precisely, the arcs passing "around" each of u and v, each arc needing a bead node). The two arborescences from Theorem 11 are represented by dotted and solid arcs, respectively.

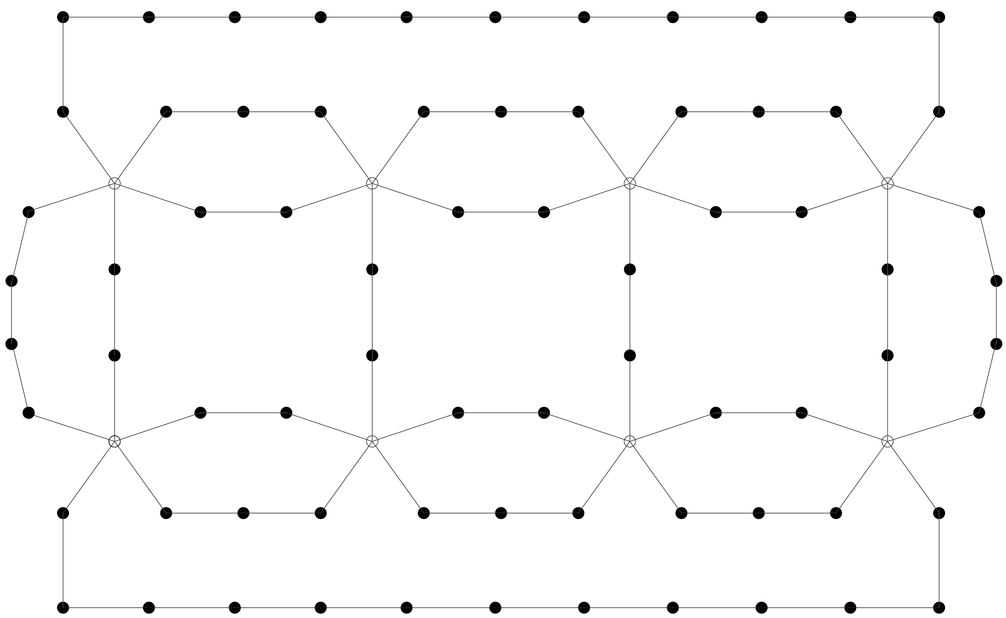

Figure 9: The nodes of S are black disks, while the sea-stars have empty circles as their center (meant to be the relay nodes of an optimum solution).

P_i of $\overrightarrow{J}$, all starting and ending at a vertex of R and without interior vertices from R, such that, if we replace each P_i by an arc e_i joining the start and the end vertex of P_i, we obtain a biconnected digraph on vertex set R. These directed paths P_i are obtained by removing vertices and edges of J, and routing counter-clockwise on all the remaining inner faces of J, and clockwise on the outer face of J.

Proof: The proof proceeds by induction on $|E(J)| + |V(J)|$. The base case is when J is an edge-minimal biconnected planar undirected graph such that every edge has at least one endpoint in R. Get the directed paths P_i by walking counter-clockwise on the inner faces, and clockwise on the outer face, and start/finish a path when encountering vertices of R. It is immediate that each arc of $\overrightarrow{J}$ is used exactly once when computing the paths. See Figure 10 for an example. Another example appears in Figure 8.

Now we show that the resulting digraph, which we call here D, is biconnected. Pick arbitrary $x, y \in R$. As J is biconnected, we have two internally vertex disjoint $x - y$ paths $\bar{P}_1$ and $\bar{P}_2$. Rename the paths such that $\bar{P}_1$ is followed by the reverse of $\bar{P}_2$ gives the counter-clockwise orientation of a bounded region.

For intuition, we refer to Figure 10 again. The first directed $x - y$ path in D is obtained as described below. Let $\bar{P}_1$ have the vertices $x = v_0, v_1, \ldots, v_k = y$. We give an arc of D from each $v_i \in R$ to either v_{i+1} or, if $v_{i+1} \notin R$, to v_{i+2} (as every edge of J has an endpoint in R, we cannot have that both $v_{i+1} \notin R$ and $v_{i+2} \notin R$). If $v_{i+1} \in R$, then there is an arc in D from v_i to v_{i+1}, going counterclockwise on one of the two faces bordered by the edge of J

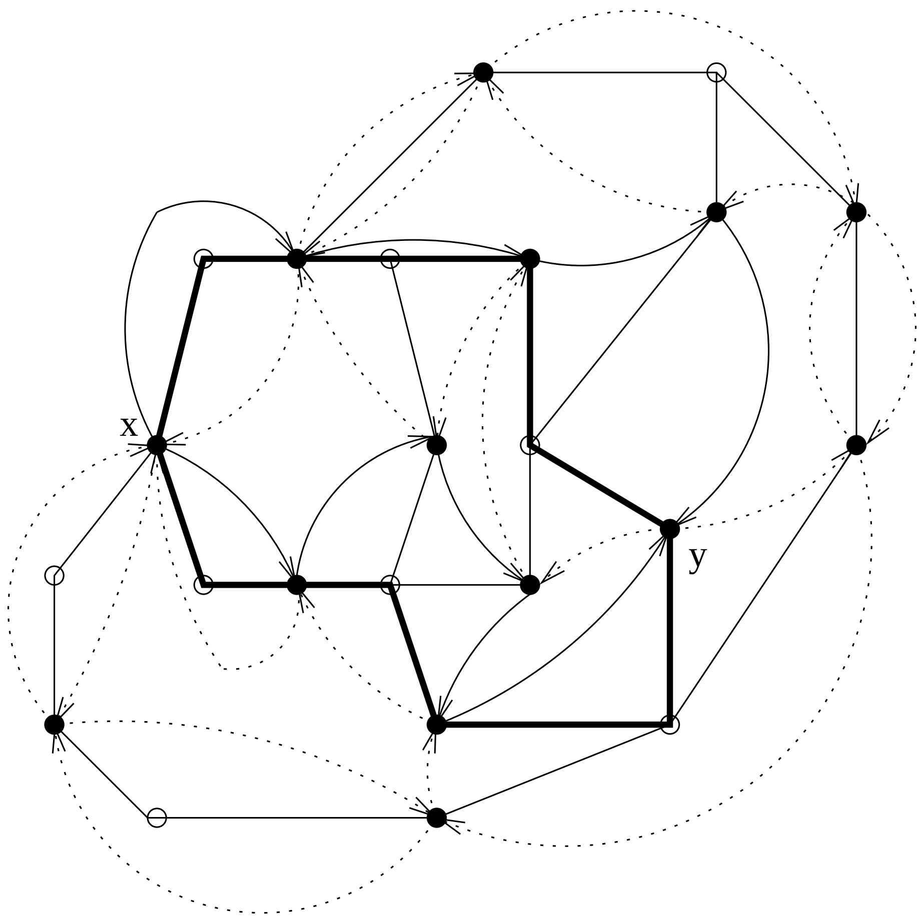

Figure 10: An edge-minimal biconnected planar graph J is given by the (straight) segments without arrows. The nodes of R are black disks, and the other nodes of J are circles. The directed paths P_i of Theorem 12 are given by all the arcs (solid or dotted). Given x, y as in the figure, two internally vertex disjoint $x - y$ paths are given by the thick segments. From these two path, we can obtain two directed paths from x to y in the resulting digraph as given by the solid arcs.

$v_i v_{i+1}$.

If $v_{i+1} \notin R$, all its J-neighbors are in R, and let $z_1, \ldots, z_l$ be the vertices of R that are adjacent to v_{i+1} and such that the l edges of J: $v_{i+1}z_1, \ldots, v_{i+1}z_l$ leave clockwise v_{i+1} and such that $v_i = z_1$ and $v_{i+2} = z_l$. The following arcs are obtained in D by walking counter-clockwise on the borders of the faces bordered by v_{i+1}: $z_1z_2, \ldots, z_{l-1}z_l$; we use all these arcs to find a directed path in D from $z_1 = v_i$ to $z_l = v_{i+2}$. We call this process the *bypass* of v_{i+1}.

We claim that none of the $z_k \in V(\bar{P}_2)$, except for the cases $v_i = z_1 = x$ and/or $v_{i+2} = z_k = y$, since otherwise the cycle of J obtained from $\bar{P}_1$ and $\bar{P}_2$ has a chord with no vertex of R, contradicting Lemma 4 (one does need $|R| > 2$ in proving this part of Lemma 4).

Doing this for all i with $v_i \in R$ let us obtain a first path in D from x to y. All the vertices of this path are either on $\bar{P}_1$ or embedded strictly in the bounded plane region bordered by $\bar{P}_1$ and the reverse of $\bar{P}_2$.

For the second directed path from x to y, apply the same procedure to the path $\bar{P}_2$, and note that all the vertices of the directed path obtained are either on $\bar{P}_2$ or embedded strictly in the unbounded plane region bordered by $\bar{P}_1$ and the reverse of $\bar{P}_2$. This finishes the base cases.

For the induction/recursion step, if J is not edge-minimal, then remove edges and apply recursion to the resulting graph. The solution for the smaller graph is good for J as well. If J has two vertices u, v which are adjacent and such that $u, v \notin R$ and such that the removal of edge uv leaves us with a graph that is not biconnected, then proceed as argued below.

Contract uv, obtaining plane graph J' with the same set R. J' is biconnected - since ((10), Proposition 2.1.9) any edge can either be contracted or removed while preserving biconnectivity. If J' has parallel edges, remove one of them from J' (this does not affect biconnectivity) and remove from J the corresponding edge incident to v. Let u' be the vertex obtained from u, v. Note that, going clockwise, the edges in J' of u' are the edges of u, followed by the edges of v. Thus, if in J' we have, in clockwise order, the edges $e_1, \ldots, e_k$ incident to u', then we can renumber them such that $e_1, \ldots, e_l$ are incident to u and follow each-other in clockwise order, and $e_{l+1}, \ldots, e_k$ are incident to v, and follow each-other in clockwise order.

Apply induction/recursion on J'. Note also that, in our routing (after possibly removing some edges of J' - in which case we remove those edges from J as well), every time we bypass a vertex z not in R, a directed path P_i enters trough an edge of z and exits trough the next clockwise edge of z (this statement is equivalent to the counter-clockwise routing on the inner faces, clockwise on the outer face). Whenever a bypass enters u' through edge e_i, it exits through edge e_{i+1} (except that we exit through e_1 if entering through e_k). If either $l = 0$ or $l = k$ (all the edges used in J' come from edges of J incident to only v, or to only u), then we use exactly the same edges and the same bypass in J as in J'. Otherwise, we use the same edges/ordering of edges in J as in J', except that when entering u through e_l, we exit through uv, thus entering v, from which we exit on e_{l+1}, and that when entering v through e_k, we exit through vu, thus entering u, from which we exit on e_1.

These changes result in exactly the same digraph D for J as for J', and keep the clockwise/counter-clockwise property. Two paths are becoming longer, by incorporating the two arcs of the bidirected J: uv and vu, but all the paths connect exactly the same vertices of R in J as they did in J'. And the two arcs resulting from bidirecting uv are used exactly once each. This completes the induction step of the proof of Theorem 12. □

5. Conclusions

Using variants of previously proposed algorithms, we improved the approximation ratio of Two-Connected Relay Placement for biconnectivity from $2d_{MST}$ to d_{MST}, and for two-edge-connectivity from $2d_{MST}$ to $2d_{MST} - 1$. In the Euclidean two-dimensional space, $d_{MST} = 5$, and in the three dimensional space, $d_{MST} = 12$. Assuming that no post-processing removes redundant relay nodes, these ratios are tight, including the ratio of 5 for biconnectivity in two dimensions. We are not able to analyze the effect of removing useless beads, a step applicable after both **Algorithm KR** and **Algorithm KV**.

It may be unrealistic to place two relay nodes at exactly the same location, or one at the same location with a sensor. Two segments, each connecting two sensors, that intersect, have different slope, and place relay nodes at the same location can be replaced, using an "uncrossing" procedure, by two segments that do not intersect (this is true in three dimensions as well). We do not have a clean solution for the case when two such segments do have the same slope (so in effect one is on top of the other); in the case of biconnectivity this can happen only if we have three sensors that are collinear and their pairwise distance is integral. Previous work also does not handle such degenerate cases.

Our result for biconnectivity has been generalized by Cohen and Nutov (6), who obtain the same approximation ratio when the problem statement includes, for every pair $\{u, v\}$ of nodes of S, a connectivity requirement $r_{uv} \in \{0, 1, 2\}$ (in our case, all the requirements are 2), and $U(S \cup Q)$ must have r_{uv} internally disjoint uv paths, for all $u, v \in S$. Their algorithm also can handle "unstable terminals". (6) uses structural results from this paper.

[1] V. Auletta, Y. Dinitz, Z. Nutov, and D. Parente. A 2-approximation algorithm for finding an optimum 3-vertex-connected spanning subgraph. *J. Algorithms*, 32:21–30, 1999.

[2] J.L. Bredin, E.D. Demaine, M.T. Hajiaghayi, and D. Rus. Deploying sensor networks with guaranteed fault tolerance. *IEEE/ACM Trans. Netw.*, 18:216–228, February 2010.

[3] V. Bryant. *Metric Spaces: Iteration and Application*. Cambridge University Press, 1985.

[4] D. Chen, D.-Z. Du, X. Hu, G. Lin, L. Wang, and G. Xue. Approximation for Steiner trees with minimum number of Steiner points. *Journal of Global Optimization*, 18:17–33, 2000.

[5] X. Cheng, D.-Z. Du, L. Wang, and B. Xu. Relay sensor placement in wireless sensor networks. *Wirel. Netw.*, 14(3):347–355, 2008.

[6] N. Cohen and Z. Nutov. Approximating {0,1,2}-survivable networks with minimum number of Steiner points. *CoRR*, abs/1304.7571, 2013.

[7] A. Czumaj and A. Lingas. A polynomial time approximation scheme for Euclidean minimum cost k-connectivity. In Kim Guldstrand Larsen, Sven Skyum, and Glynn Winskel, editors, *ICALP*, volume 1443 of *Lecture Notes in Computer Science*, pages 682–694. Springer, 1998.

[8] R. Diestel. *Graph Theory*. Springer, 2. edition, 2000.

[9] L. Fleischer, K. Jain, and D.P. Williamson. Iterative rounding 2-approximation algorithms for minimum-cost vertex connectivity problems. *J. Comput. Syst. Sci.*, 72:838–867, August 2006.

[10] A. Frank. *Connections in Combinatorial Optimization*. Oxford University Press, 2011.

[11] A. Frank and E. Tardos. An application of submodular flows. *Linear Algebra and its Applications*, 114/115:320–348, 1989.

[12] H. N. Gabow. A matroid approach to finding edge connectivity and packing arborescences. *Journal of Computer and System Sciences*, 50(2):259 – 273, 1995.

[13] M.X. Goemans and D.J. Bertsimas. Survivable networks, linear programming relaxations and the parsimonious property. *Mathematical Programming*, 60:145–166, 1993.

[14] K. Hvam, L. Reinhardt, P. Winter, and M. Zachariasen. Bounding component sizes of two-connected Steiner networks. *Inf. Process. Lett.*, 104(5):159–163, 2007.

[15] K. Hvam, L. Reinhardt, P. Winter, and M. Zachariasen. Some structural and geometric properties of two-connected Steiner networks. In Joachim Gudmundsson and C. Barry Jay, editors, *CATS*, volume 65 of *CRPIT*, pages 85–90. Australian Computer Society, 2007.

[16] L. Kamma and Z. Nutov. Approximating survivable networks with minimum number of Steiner points. In Klaus Jansen and Roberto Solis-Oba, editors, *WAOA*, volume 6534 of *Lecture Notes in Computer Science*, pages 154–165. Springer, 2010.

[17] A. Kashyap, S. Khuller, and M. Shayman. Relay placement for higher order connectivity in wireless sensor networks. *INFOCOM 2006. 25th IEEE International Conference on Computer Communications. Proceedings*, pages 1–12, April 2006.

[18] A. Kashyap, S. Khuller, and M.A. Shayman. Relay placement for fault tolerance in wireless networks in higher dimensions. *Comput. Geom.*, 44(4):206–215, 2011.

[19] S. Khuller and B. Raghavachari. Improved approximation algorithms for uniform connectivity problems. *Journal of Algorithms*, 21:433–450, 1996.

[20] S. Khuller and U. Vishkin. Biconnectivity approximation and graph carvings. *J. ACM*, 41:214–235, 1994.

[21] G. Lin and G. Xue. Steiner tree problem with minimum number of Steiner points and bounded edge-length. *Information Processing Letters*, 69:53–57, 1999.

[22] E.L. Luebke. k-Connected Steiner Network Problems, 2002.

[23] E.L. Luebke and J.S. Provan. On the structure and complexity of the 2-connected Steiner network problem in the plane. *Oper. Res. Lett.*, 26(3):111–116, 2000.

[24] I.I. Mandoiu and A.Z. Zelikovsky. A note on the MST heuristic for bounded edge-length Steiner trees with minimum number of Steiner points. *Inf. Process. Lett.*, 75(4):165–167, 2000.

[25] H. Martini and K.J. Swanepoel. Low-degree minimal spanning trees in normed spaces. *Applied Mathematics Letters*, 19(2):122 – 125, 2006.

[26] Z. Nutov and A. Yaroshevitch. Wireless network design via 3-decompositions. *Inf. Process. Lett.*, 109(19):1136–1140, 2009.

[27] G. Robins and J.S. Salowe. Low-degree minimum spanning trees. *Discrete & Computational Geometry*, 14(2):151–165, 1995.

[28] F. Wang, M.T. Thai, and D.-Z. Du. On the construction of 2-connected virtual backbone in wireless networks. *Wireless Communications, IEEE Transactions on*, 8(3):1230 – 1237, 2009.

[29] R.W. Whitty. Vertex-disjoint paths and edge-disjoint branchings in directed graphs. *J. Graph Theory*, 11(3):349–358, 1987.

[30] P. Winter and M. Zachariasen. Two-connected Steiner networks: structural properties. *Oper. Res. Lett.*, 33(4):395–402, 2005.