

Relay placement for two-connectivity [★]

Gruia Calinescu¹

Illinois Institute of Technology, Chicago IL 60616, USA,
calinescu@iit.edu,
WWW home page: <http://www.cs.iit.edu/~calinesc>

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 must 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.

Kashyap, Khuller, and Shayman (Infocom 2006) present algorithms for the two variants of the problem: two-edge-connectivity and biconnectivity. For both they prove an approximation ratio of at most $2d_{MST}$, where d_{MST} is the maximum degree of a minimum-degree Minimum Spanning Tree in the normed space. In the Euclidean two and three dimensional spaces, $d_{MST} = 5$, and $d_{MST} = 13$ respectively. We give a tight analysis of the same algorithms, obtaining approximation ratios of d_{MST} for biconnectivity and $2d_{MST} - 1$ for two-edge-connectivity respectively.

Keywords: wireless and sensor networks, approximation algorithms, Steiner nodes

1 Introduction

A wireless sensor network is composed of a large number of sensors, which can be densely deployed to monitor the targeted environment. Sensors may have a short transmission range since long transmission consumes more energy, and the sensors normally have limited power. Therefore, network partitions may 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. A normed space is a metric space (X, d) , given by a set X and a symmetric function $d : X \times X \rightarrow \mathbb{R}^+$ 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 [2], a normed space also has the following property (and others that we don't 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 $d(x, z) = \alpha d(x, y)$. In

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

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).

We must 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 (we normalize to 1 the transmission range of the sensors). For a multiset of points P , let $U(P)$ be the unit-disk graph induced by P . Also, we call two vertices U -adjacent, or U -neighbors, if their distance is at most 1.

Kashyap, Khuller, and Shayman [12] introduce the two variants of this problem: TWO-EDGE-CONNECTED RELAY PLACEMENT ($U(S \cup Q)$ 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 [21, 18] that d_{MST} is the Hadwiger number of the normed space, defined as follows: the maximum size of an independent set in $U(N_x)$, taken over all points x of the space, with N_x being the points, other than x , within distance 1 of x . It is known that $d_{MST} = 5$ in the Euclidean two-dimensional space, and $d_{MST} = 13$ in three dimensions.

[12] presents two algorithms, based on the Khuller and Vishkin (**Algorithm KV**[14]) and the Khuller and Raghavachari [13] (**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 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, [12] proves that each of the two algorithms has approximation ratio of at most $2d_{MST}$. [12] also presents simulation results.

We give a tight analysis 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, for two-edge-connectivity and 5, instead of 10 ([12]), for biconnectivity. Assuming that no post-processing removes redundant relay nodes, the ratios given in this paragraph are tight for these algorithms. We are not able to analyze the effect of removing useless beads, a step applicable after both **Algorithm KR** and **Algorithm KV**.

For the ratio of $2d_{MST} - 1$, we use a more careful accounting and look inside **Algorithm KV**. Due to space limitations, we fully omit presenting this result, as well as several other proofs. 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.

Precisely, we claim the following “parsimony” 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 S be a subset of $V(H)$. Then there exists a set of arc-disjoint paths P_i of H , all starting and ending at a vertex of S and without interior vertices from S , such that, if we replace each P_i by an arc e_i joining the start and end vertex of P_i , we obtain a biconnected digraph on vertex set S . This property allows one to “bypass” Steiner vertices (“parsimony”) and in some sense eliminate them.

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 [9, 8, 24, 17], and we use some of their results and techniques for “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 [5], who proved that it is integral (see also [6]). Thus, there exists a 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 we describe later is (almost) derived from an integral polytope optimum solution.

Also in previous work, Wang, Thai, and Du [22] and Bredin, Demaine, Hajiaghayi, and Rus [1] also gave constant factor algorithms, those of [1] achieving a $O(k^4)$ approximation for k -connectivity. 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 , [1] obtains a $O(k^3)$ -approximation ratio, improved to $O(k^2)$ by Kamra and Nutov [10].

MSPT (Minimum Number of Steiner Points Tree with bounded edge length) is the following related 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 [15] 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 [19] give a tight analysis of 4 for the MST-based algorithm described above, and generalize to arbitrary normed spaces obtaining a ratio of $d_{MST} - 1$. Chen, Du, Hu, Lin, Wang, and Xue also prove in [3] the same ratio of




Fig. 1. Left, an optimum solution for two-edge-connectivity. The nodes of S are black disks, and the relay nodes are empty circles. Right, an optimum bead solution.

4 but with a different approach, and present a 3-approximation algorithm. Later, Cheng, Du, Wang, and Xu [4] improve the running time of the algorithms found in [3] while the approximation ratio is unchanged. They also present a randomized algorithm with approximation ratio 2.5 for the same problem. In arbitrary normed spaces, Nutov and Yaroshevitch [20] obtain a $\lfloor (d_{MST} + 1)/2 \rfloor + 1 + \epsilon$ -approximation.

2 Biconnectivity

For any graph G , we use \vec{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 \vec{G} is defined as $w(D) = \sum_{e \in E(D)} w(e)$.

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

For any subdigraph D of \vec{G} , we use \bar{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. Call a feasible solution Q of a TWO-CONNECTED RELAY PLACEMENT instance a *bead-solution* if $U(Q \cup S)$ contains a two-edge-connected graph (or biconnected, respectively) H where each node of Q has degree exactly two. The Kashyap et. al. [12] algorithm produces a bead solution - see for example Figure 1, borrowed from the thesis of Kashyap [11]. In a bead-solution, we may call the relay nodes *beads*.

As [12], we use the approximation algorithm of Khuller and Raghavachari [13], which we refer to as **Algorithm KR**. We use a variant of **Algorithm KR** and go deeper in the algorithm to obtain a better approximation ratio.

For $x, y \in S$, define $w(x, y) = \max(0, \lceil d(x, y) \rceil - 1)$, 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 $d(x, y)$. Note that w is not a metric.

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 [5]. Gabow [7] has given a faster implementation of the Frank-Tardos algorithm. Suppose that D is a 2-outconnected digraph from a vertex s in which s has exactly two outgoing neighbors x and y . Then the graph $(V(\overline{D}) \setminus \{s\}, E(\overline{D}) \cup \{xy\} \setminus \{sx, sy\})$ is biconnected [13]. **Algorithm KR** constructs a biconnected spanning subgraph of a given weighted graph G as follows.

1. Let xy be an edge of G and s be a vertex not in V . Add to \overrightarrow{G} two arcs sx and sy of weight 0. The resulting digraph is denoted by G^+ .
2. Find the minimum-weighted spanning subdigraph D of G^+ which is 2-outconnected from s .
3. Output the edge set $E(\overline{D}) \cup \{xy\} \setminus \{sx, sy\}$.

Using as G the complete graph on vertex set S with w as the weight defined above, both us and [12] use the slightly modified algorithm that tries as xy above all the edges of G , and picks the minimum of all output solutions. Replace each edge of positive weight by new beads (that is, every such edge has its own distinct beads); this is the output. It is known that this modified version is a 2-approximation for MINIMUM-WEIGHT SPANNING BICONNECTED SUBGRAPH. The approximation ratio of $2d_{MST}$ obtained by [12] 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. Our approximation ratio of d_{MST} follows from Theorem 2, proven after preliminary lemmas.

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 path P in H between u and v that, except for u and v , shares neither vertices nor edges with C .

Lemma 1. *Let J be a biconnected simple undirected graph and A be a subset of $V(J)$ with $|A| > 1$. Assume no proper biconnected subgraph J' of J exists such that $A \subseteq V(J')$. Then for every cycle in J , any chord-path has in its interior a vertex of A . Every simple cycle of J contains two vertices of A .*

Proof omitted. The property above is proved by Luebke and Pravan [17] (see also [16]) with a slightly different hypothesis and their approach works here as well. From here it is immediate to deduce the following:

Corollary 1. *Let J be a biconnected simple undirected graph and A be a subset of $V(J)$ with $|A| > 1$. Assume no proper biconnected subgraph J' of J exists such that $A \subseteq V(J')$. Let Q_i be a connected component of the subgraph of J induced by $V(J) \setminus A$. Let A_i be the set of vertices of A adjacent to some vertex in Q_i , and let T_i be the the subgraph of J with vertex set $A_i \cup Q_i$ and containing all the*

edges of J with at least one endpoint in Q_i . Then T_i is a tree (called full Steiner component).

We also need a maximum degree condition that is claimed and used in [12]. The condition is stated in the lemma below whose proof we omit.

Lemma 2. *Assume Q is minimal such that $U(S \cup Q)$ is biconnected. $U(S \cup Q)$ contains a biconnected subgraph such that every vertex of Q has degree at most d_{MST} .*

Now, it will be nice if we could use the ‘‘parsimony’’ property for planar graphs mentioned in the introduction, and whose proof we omit due to space limitations. However we are unable to prove or disprove this property for non-planar graphs, and in three dimensions we cannot count on planarity. We do have Lemma 3 below, weaker in two respects: the solution is ‘‘fractional’’, and 2-outconnectivity replaces biconnectivity. It will be enough for our purpose.

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

$$0 \leq \beta_e \leq 1 \quad \forall e \in E(L) \quad (1)$$

$$\sum_{e \in \Lambda(V \setminus X, X)} \beta_e \geq 2 \quad \forall \emptyset \neq X \subseteq (V(L) \setminus \{s\}) \quad (2)$$

$$\sum_{e \in \Lambda(V \setminus (\{z\} \cup X), X)} \beta_e \geq 1 \quad \forall z \neq s \quad \forall \emptyset \neq X \subseteq (V(L) \setminus \{s, z\}) \quad (3)$$

Using Menger’s theorem, one can check that, for an integral vector β valid for $\mathcal{P}(L, s)$, the set A of arcs e of $E(L)$ with $\beta_e = 1$ is such that the digraph $(V(L), A)$ is 2-outconnected from s . Thus one can think of a valid vector β as being ‘‘fractional-2-outconnected’’. Theorem 17.1.14 of [6], (given there with more complicated notation as it solves k -outconnectivity), is given below:

Theorem 1. (originally [5]) *The system giving \mathcal{P} is Total Dual Integral, which implies that for any $c : E(L) \rightarrow \mathbb{N}$, if the linear program [Minimize $\sum_{e \in E(L)} c_e \beta_e$ subject to $\beta \in \mathcal{P}(L, s)$] 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 prove our main structural property (the ‘‘fractional outconnected parsimony’’):

Lemma 3. *Let J be a biconnected undirected graph, and replace every edge by two anti-parallel directed arcs. Let A be a subset of $V(J)$. Then there exist vertices $x, y \in A$, and there exist positive reals α_i and a set of paths P_i of J , all starting and ending at a vertex of A and without interior vertices from A , with the following properties. P_0 starts at x and ends at y and $\alpha_0 = 1/2$. P_1 starts at y and ends at x and $\alpha_1 = 1/2$. For every arc of $e \in E(J)$,*

$$\sum_{i \geq 0 \mid e \in E(P_i)} \alpha_i \leq 1$$

For $i \geq 2$, replace each P_i by an arc e_i joining the start and end vertex of P_i , obtaining a directed graph H with vertex set A . Add new vertex s and two arcs sx and sy to H , resulting in digraph H^+ . Let $\alpha_{sx} = \alpha_{sy} = 1$, and for $i \geq 2$, let $\alpha(e_i) = \alpha_i$. Then the vector α is feasible for $\mathcal{P}(H^+, s)$.

Proof. Remove edges and vertices not in A from J until it satisfies the conditions of Lemma 1 and Corollary 1. Let T_j ($j \geq 0$) be the full Steiner components (all our full Steiner components have at least one vertex not in A , and no edge with both endpoints in A) given by this corollary. Do an Eulerian traversal of each bidirected T_j (as in Christofides' algorithm). Recall that the vertices of $A \cap V(T_j)$ are leaves, and thus each is visited exactly once. If vertices u, v of A appear in this traversal such that, after u, v is the next vertex of A (thus skipping the vertices not in A), have two paths P_j and P_k one from u to v and one from v to u , both with $\alpha_i = 1/2$. P_j follows the traversal, while P_k is the reverse of P_j . Arbitrarily pick a T_0 and x and y consecutive in the Eulerian traversal of T_0 , and renumber the paths such that P_0 starts at x and ends at y and P_1 starts at y and ends at x . All α values are still $1/2$. For two vertices u and v of A adjacent in J , make (one-arc) paths P_j and P_k , one from u to v and one from v to u , both with $\alpha_i = 1$. One can immediately check that for every arc $e \in E(J)$, $\sum_{i \geq 0 \mid e \in E(P_i)} \alpha_i \leq 1$, as indeed, for an arc e of a full Steiner component, e appears in two paths P_i : one is a part of the Eulerian traversal, and one is the reverse of a path P in the Eulerian traversal, precisely P that contains the arc antiparallel to e . Incidentally to this proof, we remark that Kashyap et al. [12] also do this Eulerian traversal (though they don't 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 $i \geq 2$, replace each P_i by an arc e_i joining the start and end vertex of P_i , obtaining a directed graph H with vertex set A ; note that H does not include arcs given by P_0 and P_1 . Add new vertex s and two arcs sx and sy to H , resulting in digraph H^+ . Let $\alpha_{sx} = \alpha_{sy} = 1$. It remains to show that the vector α is feasible for $\mathcal{P}(H^+, s)$. Once again incidentally, we mention that [12] implicitly obtain the same H^+ but put $\alpha = 1$ on all the arcs, while we use $1/2$ for all arcs with at least one Steiner endpoint. This is where we improve the ratio - and this also explains why our proof is much longer and complicated.

Constraints (1) are immediate. We proceed to Constraints (2). Pick an arbitrary $X \subseteq (V(L) \setminus \{s\})$. If $\{x, y\} \in X$, then the two arcs sx and sy with $\alpha = 1$ satisfy Constraint (2) for X .

Consider now the case $x \in X$ and $y \notin X$. We have $\alpha_{sx} = 1$. Going back to the undirected J , there are two internally-disjoint paths P and P' from y to x . We claim that one of them must be P_1 , the path inside the full Steiner component T_0 . Indeed, assume otherwise, and let C be the cycle obtained from putting together P and P' . If P_1 has edges not C , then it has a subpath that is, in J , a chord-path for C with no internal vertices of A , contradicting Lemma 1. Thus without loss of generality we assume $P = P_1$. As P_1 does not give an arc in H^+ , we concentrate on P' . Note that the lack of chord-paths without internal

vertices of A also shows that not internal vertex of P' appears in T_0 . Let v be the first vertex of X on P' ($v = x$ possible), and v' be the vertex before v on P' ($v' = y$ possible). If $v' \in A$ (note that $v \in A$), then we have a path P_j from v' to v in the bidirected J with $\alpha_j = 1$, and then in H we have an arc e_j from v' to v with $\alpha_j = 1$. Then the arcs sx and $v'v$ together satisfy Constraint (2). If $v' \notin A$, then there is a full Steiner component T_i that contains v' and that has endpoints both in X and outside X (the last vertex of A on P' before v is not in X ; recall $v \in X$). Note that $i \neq 0$. The Eulerian traversal of T_i gives two paths, P_j and P_k , one entering X and one exiting X . Then $\alpha_j = 1/2$, and also there is another path P_q , the reversal of P_k , that also enters X and has $\alpha_q = 1/2$. Then the arcs sx , e_j , and e_q together satisfy Constraint (2).

The case $x \notin X$ and $y \in X$ is symmetric. The last case for verifying Constraint (2), considered next, has $x \notin X$ and $y \notin X$; let $v \in X$. Going back to the undirected J , there are two internally-disjoint paths P'_1 and P'_2 from x to v . Let C be the cycle obtained from putting together P'_1 and P'_2 . Let v_1 be the first vertex of X on P'_1 ($v_1 = v$ possible), and v_2 be the first vertex of X on P'_2 ($v_2 = v$ possible). Let v'_1 be the vertex before v_1 on P'_1 ($v'_1 = x$ possible), and let v'_2 be the vertex before v_2 on P'_2 ($v'_2 = x$ possible). If $v'_1 \in A$ (note that $v_1 \in A$), then we have a path P_{r_1} from v'_1 to v_1 in the bidirected J with $\alpha_{r_1} = 1$, and then in H we have an arc e_{j_1} from v'_1 to v_1 with $\alpha_{j_1} = 1$. Similarly, if $v'_2 \in A$ (note that $v_2 \in A$), then we have a path P_{r_2} from v'_2 to v_2 in the bidirected J with $\alpha_{r_2} = 1$, and then in H we have an arc e_{j_2} from v'_2 to v_2 with $\alpha_{j_2} = 1$. If $v'_1 \notin A$, then there is a full Steiner component T_{i_1} that contains v'_1 and that has endpoints both in X and outside X . If $v'_2 \notin A$, then there is a full Steiner component T_{i_2} that contains v'_2 and that has endpoints 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 A . The Eulerian traversal of T_{i_1} gives two 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 T_{i_2} gives two 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 path P_{q_2} , the reversal of P_{k_2} , that also enters X and has $\alpha_{q_2} = 1/2$. Note that neither of j_1 , q_1 , j_2 , and q_2 could be 0 or 1, as both P_0 and P_1 have their endpoints outside X . Thus H contains either e_{r_1} with $\alpha_{j_1} = 1$, or both e_{j_1} and e_{q_1} with $\alpha_{j_1} = \alpha_{q_1} = 1/2$. Also H contains either e_{r_2} with $\alpha_{j_2} = 1$, or both e_{j_2} and e_{q_2} with $\alpha_{j_2} = \alpha_{q_2} = 1/2$. In all four subcases, Constraint (2) is satisfied.

We proceed to Constraints (3), which must hold $\forall z \neq s \quad \forall \emptyset \neq X \subseteq (V(L) \setminus \{s, z\})$. If $x \in X$, regardless of y and z , the arc sx with $\alpha_{sx} = 1$ satisfies the constraint. Similarly, sy satisfies the constraint if $y \in X$. Consider now the case $x \notin X$ and $y \notin X$. The argument is the same whether $z = y$ or $z = x$ or $z \notin \{x, y\}$; we will assume by symmetry $x \neq z$, and let $v \in X$. Going back to the undirected J , there are two internally-disjoint paths P'_1 and P'_2 from x to v ; assume by renaming P'_1 and P'_2 that z is not a vertex of P'_1 . Let C be the cycle obtained from putting together P'_1 and P'_2 . Let v_1 be the first vertex of X on

P'_1 ($v_1 = v$ possible), and v_2 be the first vertex of X on P'_2 ($v_2 = v$ possible). Let v'_1 be the vertex before v_1 on P'_1 ($v'_1 = x$ possible), and let v'_2 be the vertex before v_2 on P'_2 ($v'_2 = x$ or $v'_2 = z$ possible). If $v'_1 \in A$ (note that $v_1 \in A$), then we have a path P_{r_1} from v'_1 to v_1 in the bidirected J with $\alpha_{r_1} = 1$, and then in H we have an arc e_{j_1} from v'_1 to v_1 with $\alpha_{j_1} = 1$. So, if $v'_1 \in A$, Constraint (3) is satisfied.

Assume from now on that $v'_1 \notin A$; therefore there is a full Steiner component T_{i_1} that contains v'_1 and that has endpoints both in X and outside X . Consider the case when z is an interior vertex of 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 A . The Eulerian traversal of T_{i_1} gives two 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 $\alpha_{q_1} = 1/2$. 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} . Also, $\{j_1, q_1\} \cap \{0, 1\} = \emptyset$, since both P_{j_1} and P_{q_1} have one end in X , and P_0 and P_1 have endpoints x and y which are not in X . The two arcs of H : e_{j_1} and e_{q_1} satisfy Constraint (3).

From now on, z is not an interior vertex of P'_2 . If $v'_2 \in A$ (note that $v_2 \in A$ and $v'_2 \neq z$), then we have a path P_{r_2} from v'_2 to v_2 in the bidirected J with $\alpha_{r_2} = 1$, and then in H we have an arc e_{j_2} from v'_2 to v_2 with $\alpha_{j_2} = 1$. So, if $v'_2 \in A$, Constraint (3) is satisfied.

We are left with the case $v'_1 \notin A$, z not on P'_2 , and $v'_2 \notin A$; recall that z is not on P'_1 . 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 endpoints 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 A .

Let v''_1 be the last vertex of A before v_1 on P_1 ; then $v''_1 \in V(H^+) \setminus (X \cup \{z\})$. Consider the Eulerian traversal of T_{i_1} ; it passes through each vertex of $A \cap V(T_{i_1})$ 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 path P_{j_1} of this traversal goes from $V(H^+) \setminus (X \cup \{z\})$ to X , or goes from X to $V(H^+) \setminus (X \cup \{z\})$. In the second case, P_{q_1} , the reverse of P_{j_1} goes from $V(H^+) \setminus (X \cup \{z\})$ to X . Let P_{k_1} be either P_{j_1} or P_{q_1} , such that it goes from $V(H^+) \setminus (X \cup \{z\})$ to X . P_{k_1} cannot be P_0 or P_1 , since it ends in X , and P_0 and P_1 end in x or y , both being outside X . Thus e_{k_1} exists in H^+ ; also $\alpha_{k_1} = 1/2$. We repeat the argument for T_{i_2} , to get another arc $e_{k_2} \in E(H^+)$ going from $V(H^+) \setminus (X \cup \{z\})$ to X , and with $\alpha_{k_2} = 1/2$. These two arcs of H : e_{k_1} and e_{k_2} satisfy Constraint (3).

In all cases, Constraint (3) is satisfied. \square

Theorem 2. *Let S be an instance and Q an optimum feasible solution for BI-CONNECTED RELAY PLACEMENT. Let G be the weighted graph on S . Then there exists $xy \in E(G)$ such that, after bidirecting G and adding to \vec{G} a vertex s and arcs sx and sy of weight 0, resulting in graph G^+ , G^+ contains a 2-outconnected subgraph D from s such that $w(D) + w(xy) \leq d_{MST}|Q|$.*

Proof. Define, for graph L , $d_L(v)$ to be the degree of vertex v in L . For path P in L , define \mathring{P} to be the set of vertices in its interior, and define $l'_L(P) = |\mathring{P}|$.

Assume Q is minimal such that $U(Q \cup S)$ is biconnected, and $|S| \geq 2$. Choose a biconnected spanning subgraph K of $U(Q \cup S)$ as in Lemma 2. Apply Lemma 3 with K as J , and S as A , obtaining vertices $x, y \in S$, paths $(P_i)_{i \geq 0}$, nonnegative numbers α_i , arcs e_i (for $i \geq 2$) giving digraph H , and then digraph H^+ after adding vertex s and arcs sx and sy . Use these x, y as the two vertices of G required by the theorem.

For arcs e of H^+ , define $c_e = 0$ if $e = sx$ or $e = sy$, and $c_e = l'_K(P_i)$ if $e = e_i$ is obtained from P_i (for $i \geq 2$). Consider the linear program **LP**: Minimize $\sum_{e \in E(H^+)} c_e \beta_e$ subject to $\beta \in \mathcal{P}(H^+, s)$, and note that α from Lemma 3 gives a feasible solution. Apply Theorem 1 to get an integral solution for **LP**, and therefore a digraph D , subgraph of H^+ , 2-outconnected from s . It remains to check the weight condition. Note that for any edge e of G with endpoints u and v , and for any path P from u to v in K , $w(e) \leq l'_K(P)$, as beads can be placed on the vertices of \check{P} . Also, every arc e_i of D of non-zero weight comes from a path P_i (with $i \geq 2$) of K . Thus (using Theorem 1 for the second inequality):

$$w(D) = \sum_{e \in D} w(e) \leq \sum_{i \geq 2 : e_i \in D} l'_K(P_i) \leq \sum_{i \geq 2} \alpha_i l'_K(P_i).$$

Write $e \diamond v$ if edge e is adjacent to vertex v , and $next(P, v, e)$ if on path P , edge e is used to leave v . We have:

$$\begin{aligned} w(D) + w(xy) &\leq \sum_{i \geq 0} \alpha_i l'_K(P_i) = \sum_{i \geq 0} \alpha_i \sum_{v \in \check{P}_i} 1 = \sum_{v \in Q} \sum_{i \geq 0 \mid v \in \check{P}_i} \alpha_i \\ &= \sum_{v \in Q} \sum_{e \mid e \diamond v} \sum_{i \geq 0 \mid next(P_i, v, e)} \alpha_i \leq \sum_{v \in Q} \sum_{e \mid e \diamond v} 1 \\ &= \sum_{v \in Q} d_K(v) \leq |Q| \cdot d_{MST} \end{aligned}$$

where the first inequality follows from the previous equation and the facts that P_0 is a path of K that starts at x and ends at y with $\alpha_0 = 1/2$, and that P_1 is a path of K that starts at y and ends at x with $\alpha_1 = 1/2$. The second inequality comes from Lemma 3, and the last inequality from Lemma 2. \square

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 2. It has two *sea stars* (one relay node, the *star's center*, 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, and we connect their tentacles as those of u, v are in Figure 2 – this can always be done 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




Fig. 2. The nodes of S are black disks. Optimum uses the relay nodes u and v . If we start **Algorithm KR** with x and y 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 3 are represented, except for edges sx and sy , by dotted and solid arcs, respectively. One could get only nine beads by starting with x, y not U -adjacent. However, in a larger example, one or two beads saved still results in a ratio of five.

connected component of $U(S)$ (one for each curve), such that $m = 5q/2 - o(q)$. Optimum is q (and let $q \rightarrow \infty$). We use the following theorem of Whitty [23]:

Theorem 3. [23] *Suppose that, given a directed graph $D = (V, A)$ 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 $v \in V - s$ the two paths to s from v uniquely determined by the arborescences are internally vertex-disjoint.*

Wherever we start with x and y in **Algorithm KR**, each of the two arborescences from the theorem above needs $m - 1$ arcs of weight 1 to enter each of the m paths/connected components of $U(S)$, except those containing x or y . Thus **Algorithm KR** produces a solution of weight at least $2(m - 1) = 5q - o(q)$.

A similar (but non-planar) construction can be made for the three dimensional Euclidean space, using far-apart sea stars with 13 tentacles each.

References

1. 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, 2010.
2. Victor Bryant. *Metric Spaces: Iteration and Application*. Cambridge University Press, 1985.

3. 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.
4. Xiuzhen Cheng, Ding-Zhu Du, Lusheng Wang, and Baogang Xu. Relay sensor placement in wireless sensor networks. *Wirel. Netw.*, 14(3):347–355, 2008.
5. A. Frank and E. Tardos. An application of submodular flows. *Linear Algebra and its Applications*, 114/115:320–348, 1989.
6. Andras Frank. *Connections in Combinatorial Optimization*. Oxford University Press, 2011.
7. Harold N. Gabow. A representation for crossing set families with applications to submodular flow problems. In *Proc. SODA*, pages 202–211, 1993.
8. K. Hvam, L. Reinhardt, P. Winter, and M. Zachariassen. Bounding component sizes of two-connected Steiner networks. *Inf. Process. Lett.*, 104(5):159–163, 2007.
9. K. Hvam, L. Reinhardt, P. Winter, and M. Zachariassen. Some structural and geometric properties of two-connected Steiner networks. In *Proc. CATS*, volume 65 of *CRPIT*, pages 85–90. Australian Computer Society, 2007.
10. Lior Kamma and Zeev Nutov. Approximating survivable networks with minimum number of Steiner points. In *Proc. WAOA*, volume 6534 of *Lecture Notes in Computer Science*, pages 154–165, 2010.
11. A. Kashyap. Robust design of wireless networks, *Ph.D. thesis*, University of Maryland, 2006.
12. 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.
13. S. Khuller and B. Raghavachari. Improved approximation algorithms for uniform connectivity problems. *Journal of Algorithms*, 21:433–450, 1996.
14. S. Khuller and U. Vishkin. Biconnectivity approximation and graph carvings. *J. ACM*, 41:214–235, 1994.
15. 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.
16. E. L. Luebke. *k*-Connected Steiner Network Problems, *Ph.D. thesis*, University of North Carolina, 2002.
17. 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.
18. Horst Martini and Konrad J. Swanepoel. Low-degree minimal spanning trees in normed spaces. *Applied Mathematics Letters*, 19(2):122 – 125, 2006.
19. Ion I. Măndoiu and Alexander 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.
20. Zeev Nutov and Ariel Yaroshevitch. Wireless network design via 3-decompositions. *Inf. Process. Lett.*, 109(19):1136–1140, 2009.
21. Gabriel Robins and Jeffrey S. Salowe. Low-degree minimum spanning trees. *Discrete & Computational Geometry*, 14(2):151–165, 1995.
22. Feng Wang, M.T. Thai, and Ding-Zhu Du. On the construction of 2-connected virtual backbone in wireless networks. *Wireless Communications, IEEE Transactions on*, 8(3):1230–1237, 2009.
23. R.W. Whitty. Vertex-disjoint paths and edge-disjoint branchings in directed graphs. *J. Graph Theory*, 11(3):349–358, 1987.
24. Pawel Winter and Martin Zachariassen. Two-connected Steiner networks: structural properties. *Oper. Res. Lett.*, 33(4):395–402, 2005.