A Worst-case Analysis of an MST-based Heuristic to Construct Energy-Efficient Broadcast Trees in Wireless Networks*

Andrea E.F. Clementi† Pilu Crescenzi‡ Paolo Penna§

Gianluca Rossi¶ Paola Vocca‖

July 2001

Abstract

We consider the problem of computing an optimal range assignment in a wireless network which allows a specified source station to perform a broadcast operation. In particular, we analyze a polynomial-time algorithm proposed by Wieselthier, Nguyen, and Ephremides, which is based on the computation of a standard minimum spanning tree, and we prove that this algorithm computes a solution whose performance ratio is bounded by 20 in the case in which both the dimension d and the gradient α are equal to 2. Moreover, we show how this approximation algorithm can be generalized to the case in which $\alpha \geq d$, for any dimension d, and we prove a lower bound on its performance ratio, which is exponential with respect to d.

* **A preliminary version of this paper appeared in Proc. of STACS'01 [CCP+01].**

† Dipartimento di Matematica, Università degli Studi di Roma "Tor Vergata", Via della Ricerca Scientifica, I-00133 Roma, Italy. Email: clementi@mat.uniroma2.it

‡ Dipartimento di Sistemi e Informatica, Università degli Studi di Firenze, Via C. Lombroso 6/17, I-50134 Firenze, Italy. Email: piluc@dsi.unifi.it

§ Institut für Teoretische Informatik, Swiss Federal Institute of Technology Zürich (ETH), Clausiustrasse 49, CH-8092 Zürich, Switzerland. Email: penna@inf.ethz.ch

¶ Dipartimento di Sistemi e Informatica, Università degli Studi di Firenze, Via C. Lombroso 6/17, I-50134 Firenze, Italy. Email: rossig@dsi.unifi.it

‖ Dipartimento di Matematica, Università degli Studi di Roma "Tor Vergata", Via della Ricerca Scientifica, I-00133 Roma, Italy. Email: vocca@mat.uniroma2.it

1 Introduction

Wireless networking technology will play a key role in future communications and the choice of the network architecture model will strongly impact the effectiveness of the applications proposed for the mobile networks of the future. Broadly speaking, there are two major models for wireless networking: *single-hop* and *multi-hop*. The single-hop model [RW94], based on the cellular network model, provides one-hop wireless connectivity between mobile hosts and static nodes known as *base stations*. This type of networks relies on a fixed backbone infrastructure that interconnects all base stations by high-speed wired links. On the other hand, the multi-hop model [Lau95] requires neither fixed, wired infrastructure nor predetermined interconnectivity. *Ad hoc* networking [HT98] is the most popular type of multi-hop wireless networks because of its simplicity: Indeed, an *ad hoc* wireless network is constituted by a homogeneous system of *mobile* stations connected by wireless links. In ad hoc networks, to every station is assigned a transmission range: The overall range assignment determines a transmission (directed) graph since one station s can transmit to another station t if and only if t is within the transmission range of s. The range transmission of a station depends, in turn, on the energy power supplied to the station: In particular, the power P_s required by a station s to correctly transmit data to another station t must satisfy the inequality

$$\frac{P_s}{d(s,t)^\alpha} > \gamma \tag{1}$$

where $d(s,t)$ is the Euclidean distance between s and t, $\alpha \geq 1$ is the *distance-power gradient*, and $\gamma \geq 1$ is the *transmission-quality* parameter. In an ideal environment (i.e. in the empty space) it holds that $\alpha = 2$ but it may vary from 1 to more than 6 depending on the environment conditions of the place the network is located (see [PL95]).

The fundamental problem underlying any phase of a dynamic resource allocation algorithm in ad-hoc wireless networks is the following: Find a transmission range assignment such that (1) the corresponding transmission graph satisfies a given property π, and (2) the overall energy power required to deploy the assignment (according to Eq. 1) is minimized.

A well-studied case of the above problem consists in choosing π as follows: The transmission graph has to be strongly connected. In this case, it is known that: (a) the problem is not solvable in polynomial time (unless P=NP) [CPS99, KKKP00], (b) it is possible to compute a range assignment which is at most twice the optimal one (that is, the problem is 2-approximable), for multi-dimensional wireless networks [KKKP00], (c) there exists a constant $r > 1$ such that the problem is not r-approximable (unless P=NP), for d-dimensional networks with $d \geq 3$ [CPS99], and (d) the problem can be solved in polynomial time for one-dimensional networks [KKKP00]. Another analyzed case consists in choosing π as follows: The diameter of the transmission graph has to be at most a fixed value h. In this case, while non-trivial negative results are not known, some tight bounds (depending on h) on the minimum energy power have been proved in [CPS00], and an approximation algorithm for the one-dimensional case has been given in [CFP+00]. Other trade-offs between connectivity and energy consumption have been obtained in [MM96, Pir91, UY96].

In this paper we address the case in which π is defined as follows: *Given a source station s, the transmission graph has to contain a directed spanning tree rooted at s.* This case has been posed as an open question by Ephremides in [Eph00]: Its relevance is due to the fact that any transmission graph satisfying the above property allows the source station to perform a *broadcast* operation. Broadcast is a task initiated by the source station which transmits a message to all stations in the wireless network: This task constitutes a major part of real life multi-hop radio network [BGI92, BII93].

The optimization problem. The broadcast range assignment problem described above is a special case of the following optimization problem, called MINIMUM ENERGY CONSUMPTION BROADCAST SUBGRAPH (in short, MECBS). Given a weighted directed graph $G = (V, E)$ with edge weight function $w : E \rightarrow \mathcal{R}^+$, a *range assignment* for G is a function $r : V \rightarrow \mathcal{R}^+$: The *transmission (directed) graph* induced by G and r is defined as $G_r = (V, E')$ where

$$E' = \bigcup_{v \in V} \{(v, u) : (v, u) \in E \wedge w(v, u) \leq r(v)\}.$$

The MECBS problem is then defined as follows: Given a *source node* $s \in V$, find a range assignment r for G such that G_r contains a directed spanning tree of G rooted at s and

$$\mathsf{cost}(r) = \sum_{v \in V} r(v)$$

is minimized.

The MINIMUM ENERGY CONSUMPTION BROADCAST SUBGRAPH problem is NP-hard and, if P $\neq$ NP, it is not approximable within a sub-logarithmic factor, even when the problem is restricted to undirected graphs [GK98]. This result is obtained by providing an approximation preserving reduction from MIN SET COVER to MECBS: It is known that, unless P=NP, MIN SET COVER is not approximable within a sub-logarithmic factor [RS97] (see, also, the list of optimization problems contained in [ACG+99]).

The intractability of the general version of MECBS does not necessarily imply the same hardness result for its restriction to wireless networks. In particular, let us consider, for any $d \geq 1$ and for any $\alpha \geq 1$, the family of graphs N_d^α, called *(d-dimensional) wireless networks*, defined as follows: A complete (undirected) graph G belongs to N_d^α if it can be embedded on a d-dimensional Euclidean space such that the weight of an edge is equal to the αth power of the Euclidean distance between the two endpoints of the edge itself. The restriction of MECBS to graphs in N_d^α is then denoted by MECBS[N_d^α]: It is clear that the previously described broadcast range assignment problem in the ideal 2-dimensional environment corresponds to MECBS[N_2^2].

Observe that if $\alpha = 1$ (that is, the edge weights coincide with the Euclidean distances), then the optimal range assignment is simply obtained by assigning to s the distance from its farthest node and by assigning 0 to all other nodes. We then have that, for any $d \geq 1$, MECBS[N_d^1] is solvable in polynomial time. Moreover, by making use of dynamic programming techniques, it has also been shown that, for any $\alpha \geq 1$, MECBS[N_1^α] is solvable in polynomial time [CDS01].

On the other hand, it is possible to prove that, for any $d \geq 2$ and for any $\alpha > 1$, MECBS[N_d^α] is not solvable in polynomial time (unless P= NP). The proof of this result is an adaptation of the one given in [CPS99] to prove the NP-hardness of computing a minimum range assignment that guarantees the strong connectivity of the corresponding transmission graph: This adaptation is described in Section A. As a consequence of this intractability result, the best we can do is to look for polynomial-time algorithms that compute approximate solutions for MECBS[N_d^α] with $d \geq 2$ and $\alpha > 1$: The main contribution of this paper is to show that, whenever $\alpha \geq d$, this can be done by making use of the following heuristics based on the computation of a standard minimum spanning tree (in short, MST), which has been proposed in [WNE00].

The algorithm MST-ALG**.** Given a graph $G \in \mathsf{N}_2^\alpha$ and a specified source node s, the algorithm first computes a MST T of G (observe that this computation does not depend on the value of α). Subsequently, it makes T directed by rooting it at s. Finally, the algorithm assigns to each vertex v the maximum among the weights of all edges of T outgoing from v. Clearly, the algorithm runs in polynomial time and computes a feasible solution.

A major positive aspect of the algorithm lies on the fact that it is just based on the computation of a standard minimum spanning tree. In a network with dynamic power control, the range assigned to the stations can be modified at any time: The algorithm can thus take advantage of all known techniques to dynamically maintain MSTs (see, for example, [DRT92, Epp94, NPW00]).

Our results. In [WNE00], the performance of the above described heuristics has been evaluated by simulation and the worst-case analysis of its quality in terms of the approximation ratio has been left open: The main result of this paper is the proof of an upper bound on the algorithm's performance ratio. In particular, we first prove that, for any instance of MECBS[N_2^α] with $\alpha \geq 2$, *the cost of the range assignment computed by* MST-ALG *is at most* $5^{\alpha/2} \cdot 2^\alpha$ *times the optimal cost* and we successively show that *the approximation algorithm can be generalized* in order to deal with MECBS[N_d^α], for any $d \geq 2$ and for any $\alpha \geq d$.

Our analysis is based on computational geometry techniques and it is rather interesting by itself: Indeed, independently from our paper and successively to its conference version [CCP+01], similar techniques have been used in [WCLF01] in order to obtain an improved analysis of MST-ALG in the two-dimensional case (see Section 4, for more details on this result) and to analyze other heuristics introduced in [WNE00].

Finally, we prove that *the performance ratio of* MST-ALG *grows at least exponentially with respect to* d.

Prerequisites. We assume the reader to be familiar with the basic concepts of computational complexity theory (see, for example, [BC94, Pap94]) and with the basic concepts of the theory of approximation algorithms (see, for example, [ACG+99]).

2 The two-dimensional case

In what follows, given a graph $G \in \mathsf{N}_2^\alpha$, we denote by $G^{1/\alpha}$ the graph obtained from G by setting the weight of each edge to the αth root of the weight of the corresponding edge in G: Hence, $G^{1/\alpha} \in \mathsf{N}_2^1$, that is, there exists an embedding of $G^{1/\alpha}$ on the plane such that the Euclidean distance $d(u, v)$ between two nodes u and v coincides with the weight of the edge (u, v) in $G^{1/\alpha}$.

In this section, we prove that, for any instance $x = \langle G = (V, E), w, s\rangle$ of MECBS[N_2^α] with $\alpha \geq 2$, the range assignment r computed by MST-ALG satisfies the following inequality:

$$\mathsf{cost}(r) \leq 5^{\alpha/2} \cdot 2^{\alpha}\mathsf{opt}(x), \tag{2}$$

where $\mathsf{opt}(x)$ denotes the cost of an optimal range assignment. First notice that

$$\mathsf{cost}(r) \leq w(T),$$

where, for any subgraph G' of G, $w(G')$ denotes the sum of the weights of the edges in G'. As a consequence of the above inequality, it now suffices to show that there exists a spanning subgraph G' of G such that $w(G') \leq 5^{\alpha/2} \cdot 2^{\alpha}\mathsf{opt}(x)$. Indeed, since the weight of T is bounded by the weight of G', we have that Eq. 2 holds.

In order to prove the existence of G', we make use of the following theorem whose proof is given in Sect. 2.1

Theorem 1 *Let $G(V, E) \in \mathsf{N}_2^\alpha$ with $\alpha \geq 2$ and let* diam *be diameter of the smallest disk containing all the nodes in G. Then, for any* MST *T of G,*

$$w(T) \leq 5^{\alpha/2}\mathsf{diam}^{\alpha}.$$

Notice that, given a graph $G \in \mathsf{N}_2^\alpha$ with $\alpha \geq 2$, we will always identify the nodes of G with the points corresponding to an embedding of $G^{1/\alpha}$ on the plane. Recall that the Euclidean distance $d(u, v)$ among two points u and v coincides with the weight of the edge (u, v) in $G^{1/\alpha}$.

Let r_{opt} be an optimal assignment for x. For any $v \in V$, let

$$S(v) = \{u \in V : w(v, u) \leq r_{\mathsf{opt}}(v)\}$$

and let $T(v)$ be a MST of the subgraph of G induced by $S(v)$. From Theorem 1, it follows that $w(T(v)) \leq 5^{\alpha/2} \cdot 2^{\alpha} r_{\mathsf{opt}}(v)$. Consider the spanning subgraph $G' = (V, E')$ of G such that

$$E' = \bigcup_{v \in V} \{e \in E : e \in T(v)\}.$$

It then follows that

$$w(G') \leq \sum_{v \in V} w(T(v)) \leq 5^{\alpha/2} \cdot 2^{\alpha} \sum_{v \in V} r_{\mathsf{opt}}(v) = 5^{\alpha/2} \cdot 2^{\alpha}\mathsf{opt}(x).$$

We have thus proved the following result.

Theorem 2 *For any $\alpha \geq 2$,* MECBS[N_2^α] *is approximable within $5^{\alpha/2} \cdot 2^{\alpha}$.*

2.1 Proof of Theorem 1

Let us first consider a graph $G(V, E) \in \mathsf{N}_2^\alpha$ with $\alpha = 2$ and let $e_i = (u_i, v_i)$ be the ith edge in T, for $i = 1, \ldots, |V| - 1$ (any fixed ordering of the edges is fine). We denote by D_i the *diametral open disk* of e_i, that is, the open disk whose center c_i is on the midpoint of e_i and whose diameter is $d(u_i, v_i)$. Moreover, D_i^* will denote the open disk having the same center as D_i and whose diameter is $d(u_i, v_i)/2$. From Lemma 6.2 of [MS92], it follows that D_i contains no point from the set $V - \{u_i, v_i\}$. The following lemma, instead, states that, for any two diametral disks, the center of one disk is not contained in the other disk.

Lemma 1 *For any $i, j \in \{1, \ldots, |V| - 1\}$ with $i \neq j$, c_i is not contained in D_j.*

Proof. Suppose by contradiction that there exist two diametral disks D_i and D_j such that c_i is contained in D_j. We will show that the longest edge between e_i and e_j can be replaced by a strictly shorter one, still maintaining the connectivity of T: Since T is a MST the lemma will follow. Let us assume, without loss of generality, that $d(u_j, v_j) \geq d(u_i, v_i)$. We first prove that

$$\max\{d(u_i, u_j), d(v_i, v_j)\} < d(u_j, v_j) \tag{3}$$

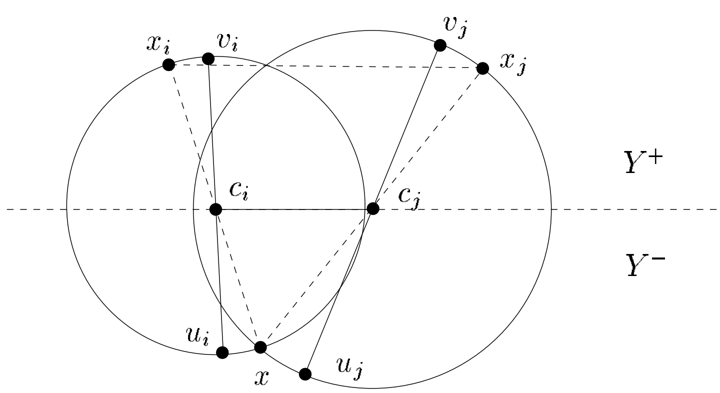

Figure 1: The proof of Lemma 1.

Let Y^+ and Y^- be the half-planes determined by the line identified by c_i and c_j: Without loss of generality, we may assume that v_i and v_j (respectively, u_i and u_j) are both contained in Y^+ (respectively, Y^-), as shown in Fig. 1. Assume also that $d(v_i, v_j) \geq d(u_i, u_j)$ (the other case can be proved in a similar way). Let x be the intersection point in Y^- between the two circumferences determined by D_i

and D_j (notice that, since D_i and D_j are open disks, neither D_i nor D_j contains x) and let x_i and x_j be the points diametrically opposite to x with respect to c_i and c_j, respectively. Clearly, $d(v_i, v_j) \leq d(x_i, x_j)$. Eq. 3 easily follows from the following

Fact 1 $d(x_i, x_j) < d(u_j, v_j)$.

Proof. By definition, c_i (respectively, c_j) is the median of the segment $\overline{xx_i}$ (respectively, $\overline{xx_j}$). Thus, the triangles $\triangle(xx_ix_j)$ and $\triangle(xc_ic_j)$ are similar. From the hypothesis that $c_i \in D_j$, it follows that $d(c_i, c_j) < d(x, c_j)$. Thus, by similarity, it must hold that

$$d(x_i, x_j) < d(x, x_j) = d(u_j, v_j)$$

and the fact follows. □

As a consequence of Eq. 3, we can replace in T, $e_j = (u_j, v_j)$ by either (u_i, u_j) or (v_i, v_j) (the choice depends on the topology of T), thus obtaining a better spanning tree. □

We now use the above lemma in order to prove that, not two disks D_i^* and D_j^* intersect each other.

Lemma 2 *For any* $i, j \in \{1, \ldots, |V| - 1\}$ *with* $i \neq j$, D_i^* *and* D_j^* *are disjoint.*

Proof. Let c_i and c_j be the centers of D_i^* and D_j^*, respectively. Also, for any $l \in \{i, j\}$, let r_l and r_l^* denote the radii of D_l and D_l^*, respectively. From Lemma 1 and from the definition of D_i^* and of D_j^*, it follows that

$$\frac{d(c_i, c_j)}{2} \geq \frac{\max\{r_i, r_j\}}{2} = \max\{r_i^*, r_j^*\}.$$

This implies that D_i^* does not intersect D_j^*. □

For any i with $1 \leq i \leq |V| - 1$, let $\overline{D_i^*}$ denote the smallest closed disk that contains D_i^*. The last lemma of this section states that the union of all $\overline{D_i^*}$s is contained in a closed disk whose diameter is comparable to the diameter of $G^{1/\alpha}$.

Lemma 3 *Let* $\mathcal{D}$ *be the smallest disk containing all the points in* V, *let* diam *be its diameter and let* $\mathcal{U} = \bigcup_{e_i \in T} \overline{D_i^*}$. *Then,* $\mathcal{U}$ *is contained into the closed disk whose diameter is equal to* $(\sqrt{5}/2)\mathsf{diam}$ *and whose center coincides with the center of* $\mathcal{D}$.

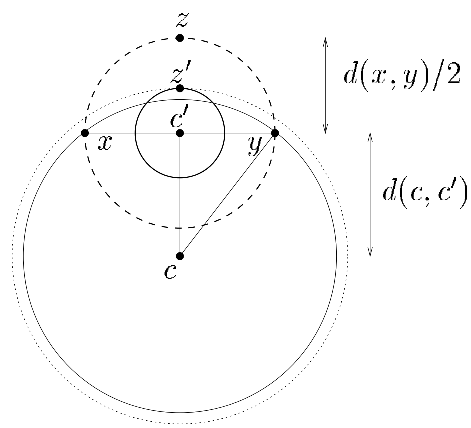

Figure 2: The proof of Lemma 3.

Proof. Consider any two points x and y within $\mathcal{D}$. Without loss of generality, we can assume that both x and y are on the boundary of $\mathcal{D}$ (indeed, for any two points x' and y' laying on the segment $\overline{xy}$, the diametral disk of x' and y' is contained in the diametral disk of x and y). Consider the closed disk whose diameter is equal to $d(x,y)$ and whose center c' is on the midpoint of the segment $\overline{xy}$, and let z be any point on its boundary (see Fig. 2). Also let z' be the middle point of $\overline{c'z}$. We now show that $d(c,z') \leq (\sqrt{5}/4)\mathsf{diam}$, where c is the center of $\mathcal{D}$. From the triangular inequality we have that

$$\begin{aligned} d(c,z') &\leq d(c,c') + d(c',z') &= d(c,c') + d(c',z)/2 \\ &= d(c,c') + d(c',y)/2 &= d(c,c') + d(x,y)/4. \end{aligned} \tag{4}$$

Moreover, since $\angle cc'y = \pi/2$,

$$d(c,c')^2 + d(c',y)^2 = d(c,y)^2 = \mathsf{diam}^2/4. \tag{5}$$

By combining Eq. 4 and Eq. 4 we get

$$d(c,z') \leq \sqrt{\frac{\mathsf{diam}^2 - d(x,y)^2}{4}} + d(x,y)/4.$$

The right hand of this equation reaches its maximum when $d(x,y) = \mathsf{diam}/\sqrt{5}$, which implies $d(c,z') \leq (\sqrt{5}/4)\mathsf{diam}$. This implies that the disk with center c and radius $(\sqrt{5}/4)\mathsf{diam}$ contains all the disks D^* given by any edge with endpoints x and y inside $\mathcal{D}$. Hence the lemma follows. □

We are now able to prove Theorem 1. In particular, for the case $\alpha = 2$ we have to prove that

$$w(T) = \sum_{i=1}^{|V|-1} d(u_i, v_i)^2 \leq 5\mathsf{diam}^2, \tag{6}$$

where (u_i, v_i) is the ith edge in T, for $i = 1, \ldots, |V| - 1$. Let $\mathsf{Area}(D_i^*)$ denote the area of $\overline{D_i^*}$. It then holds that

$$\sum_{i=1}^{|V|-1} d(u_i, v_i)^2 = \frac{16}{\pi} \sum_{i=1}^{|V|-1} \mathsf{Area}(D_i^*). \tag{7}$$

By combining Lemma 3 and 2, we have that

$$\sum_{i=1}^{|V|-1} \mathsf{Area}(D_i^*) \leq \pi \left(\frac{\sqrt{5}}{4}\mathsf{diam}\right)^2 = \frac{5\pi}{16}\mathsf{diam}^2. \tag{8}$$

By combining Eq. 7 and 8 we obtain Eq. 6, which proves the lemma for $\alpha = 2$.

Finally, we consider the case $\alpha > 2$. By simple calculations, we get

$$\begin{aligned} \mathsf{cost}(r) &= \sum_{i=1}^{|V|-1} d(u_i, v_i)^\alpha = \sum_{i=1}^{|V|-1} \left(d(u_i, v_i)^2\right)^{\alpha/2} \\ &\leq \left(\sum_{i=1}^{|V|-1} d(u_i, v_i)^2\right)^{\alpha/2} \leq 5^{\alpha/2}\mathsf{diam}^\alpha, \end{aligned}$$

where the last inequality follows from Eq. 6. This completes the proof of Theorem 1.

3 The general case

Algorithm MST-ALG achieves a constant (i.e. independent of the graph size) approximation ratio even on higher dimensions. The proof uses the same geometric arguments of those used for the two-dimensional case. Therefore, in what follows we just outline its main steps:

- Let $\alpha = d$ and let $T = \{e_1, \ldots, e_m\}$ be any MST for a set V of n points in $\mathcal{R}^d$. Let D_i the *diametral open sphere* of e_i, that is, the open sphere whose center c_i is on the midpoint of e_i and whose diameter is $d(u_i, v_i)$. Then, a positive real β_d exists such that, for any $i \neq j$ and for any point $p \in D_i \cap D_j$, the angle $\angle c_i p c_j \geq \beta_d$. (As in the proof of Lemma 1, if $\angle c_i p c_j < \beta_d$ then T would not be an MST.)

- Let $W = \{\vec{v}_1, \ldots, \vec{v}_k\}$ be any set of vectors outgoing from a fixed point in $\mathcal{R}^d$ such that, for any $i \neq j$, $\angle \vec{v}_i \vec{v}_j \geq \beta_d$. Then $k \leq n_{\beta_d}$, where n_{β_d} is a positive integer depending on d and β_d only (see the bounds on the construction of *spherical codes* in [CS88, Chapt. 2.6]).

- Let diam be the diameter of the smallest sphere $\mathcal{D}$ containing all points in V. Let $\mathcal{U} = \bigcup_{e_i \in T} \overline{D_i}$. Then, a positive integer γ_d exists such that $\mathcal{U}$ is contained into the closed sphere whose center coincides with the center of $\mathcal{D}$ and whose diameter is equal to $\gamma_d \mathsf{diam}$, with γ_d depending on d only.

- By combining the above items, we can prove that

$$w(T) = O((\gamma_d n_{\beta_d}) \cdot \mathsf{diam}^d).$$

Given any graph $G \in \mathsf{N}_d^d$, thanks to the last item, the same argument adopted to prove Theorem 2 implies the existence of a spanning subgraph G' of G such that

$$w(G') = O((\gamma_d n_{\beta_d}) \cdot \mathsf{opt}(G)).$$

Since the cost of the solution returned by MST-ALG is not larger than the cost of any spanning subgraph of G, we get the following result.

Theorem 3 *There exists a function $f : \mathcal{N} \times \mathcal{R} \to \mathcal{R}$ such that, for any $d \geq 2$ and for any $\alpha \geq d$,* MECBS[N_d^α] *is approximable within factor $f(d, \alpha)$.*

On the other hand, the following theorem shows that the function f in the statement of the previous theorem grows exponentially with respect to d.

Theorem 4 *There exists a positive constant γ such that, for any d and for any k, an instance $x_{k,d}$ of* MECBS[N_d^d] *exists such that $\mathsf{opt}(x_{k,d}) = k^d$ while the cost of the range assignment computed by* MST-ALG *is at least $k^d \cdot 2^{\gamma d}$.*

Proof. In order to prove the above theorem, we first describe the idea with an example. In the two-dimensional Euclidean space, we can construct a worst-case instance for MST-ALG by considering the star of center c and six adjacent points $c_1, \ldots, c_6$ such that $d(c, c_i) = d(c_i, c_{(i+1) \bmod 6}) = 1$, for $i = 1, \ldots, 6$. Also let c be the source node. This configuration can be viewed as an arrangement of six disks of unit diameter such that they all touch a central disk (see Figure 3).

Notice that there are (at least) two MSTs for this point set: the star with root c, and the path $c, c_1, c_2, \ldots, c_6$. The corresponding solutions have however quite

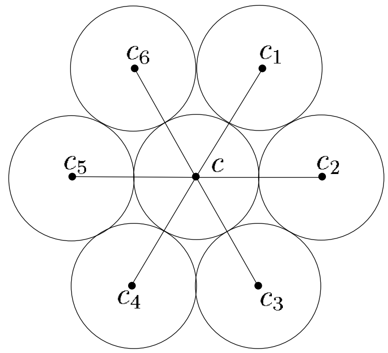

Figure 3: A worst-case instance of MST-ALG on the plane and the corresponding kissing disk arrangement.

different costs: 1 against 6. Now we modify the instance in such a way that MST-ALG is forced to "pay" all the edges in the star. Indeed, for each i with $1 \leq i \leq 6$, we place a point c'_i on the segment $\overline{cc_i}$ and at distance $\epsilon > 0$ from c. Now there is only one MST and the corresponding cost is $\epsilon^2 + 6(1-\epsilon)^2$. This implies that, on the plane, MST-ALG cannot achieve an approximation ratio better than 6.

In higher dimensions, this idea extends to *spheres* of unit diameter: a sphere in R^d of center c is defined as the set of points p such that $d(p,c) \leq 1$. Moreover, two spheres S_1, S_2 touch each other if $|S_1 \cap S_2| = 1$. We then consider the following well known problem [CS88]:

KISSING NUMBER**:** What is the maximum number of spheres that can simultaneously touch another sphere?

Let n_d be the value of the kissing number in the R^d space and let A be an arrangement of n_d unit kissing spheres. Let $\{c, c_1, \ldots, c_{n_d}\}$ be the centers of the spheres, with c be the center of the "central" sphere touching the other ones. We consider this point set as an instance of MECBS[N_d^d], with source node equal to c. Then, we modify this instance by adding $c'_1, \ldots, c'_{n_d}$, with c'_i being the point on the line passing trough c and c_i at distance $\epsilon > 0$ from c. We claim that, for this point set, there is only one MST: the one containing the n_d edges (c, c'_i) and the n_d edges (c'_i, c_i). Indeed, consider the execution of Dijkstra algorithm starting from c. First, it will add all the edges of length ϵ, and then all the ones of length $1-\epsilon$ (at any step, any other edge is strictly more costly). Then, the solution of MECBS[N_d^d] corresponding to this spanning tree has cost $\epsilon + n_d(1-\epsilon)^d$, while a solution of cost 1 exists.

This proves that MST has approximation ratio at least $\Omega(n_d)$ (we can rescale any such instance by a factor k and then define the instance $x_{k,d}$ accordingly). Hence, all we need is an exponential lower bound on n_d. We use the following well-known existential result (see also [CS88]).

Theorem 5 ([KL78, Wyn65]) *For any positive integer d,*

$$n_d \geq 2^{(1-0.5\log_2 3)d(1+o(1))}.$$

In summary, Theorem 4 has been proved.

□

4 Conclusion and open questions

In this paper we have analyzed an algorithm proposed in [WNE00] for computing an optimal range assignment in a wireless network which allows a specified source station to perform a broadcast operation. In particular, we have shown that this algorithm computes a solution whose performance ratio is bounded by $5^{\alpha/2} \cdot 2^{\alpha}$ in the two-dimensional case. Moreover, we have shown how this approximation algorithm can be generalized to the case in which $\alpha \geq d$, for any d, and we have proved a lower bound on its performance ratio, which is exponential with respect to d.

Open problems. Two main problems are left open by this paper. The first one is to improve the analysis of MST-ALG (or to develop a different algorithm with a better performance ratio). As stated in the introduction, a first step towards this direction has already been obtained in [WCLF01]: By using diamonds instead of diametral open disk, it is indeed possible to obtain a worst-case approximation ratio equal to 12 for $\alpha = d = 2$. The second open problem is to analyze the approximability properties of MECBS[N_d^α] when $\alpha < d$: Indeed, one could ask whether MST-ALG approximates MECBS[N_d^α] in the case in which $d \geq 2$ and $\alpha < d$. Unfortunately, it is not difficult to produce an instance x such that $\mathsf{opt}(x) = O(n^{\alpha/d})$ while the cost of the range assignment computed by MST-ALG is $\Omega(n)$, where n denotes the number of vertices: For example, in the case $d = 2$, we can just consider the two-dimensional grid of side $\sqrt{n}$ and the source node positioned on its center. Since the MST-based algorithm does not guarantee any approximation, it seems thus necessary to develop approximation algorithms based on different techniques.

References

[ACG+99] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi. *Complexity and Approximation – Combinatorial optimization problems and their approximability properties*. Springer Verlag, 1999.

[AK97] P. Alimonti and V. Kann. Hardness of approximating problems on cubic graphs. In *Proc. 3rd Italian Conference on Algorithms and Complexity*, volume 1203 of *LNCS*, pages 288–298, 1997.

[BC94] D. P. Bovet and P. Crescenzi. *Introduction to the Theory of Complexity*. Prentice Hall, 1994.

[BGI92] R. Bar-Yehuda, O. Goldreich, and A. Itai. On the time complexity of broadcast operations in multi-hop radio networks: an exponential gap between determinism and randomization. *J. Computer and Systems Science*, 45:104–126, 1992.

[BII93] R. Bar-Yehuda, A. Israeli, and A. Itai. Multiple communication in multi-hop radio networks. *SIAM J. on Computing*, 22:875–887, 1993.

[CCP+01] A.E.F. Clementi, P. Crescenzi, P. Penna, G. Rossi, and P. Vocca. On the complexity of computing minimum energy consumption broadcast subgraphs. In *Proc. 18th Annual Symposium on Theoretical Aspects of Computer Science*, volume 2010 of *LNCS*, pages 121–131, 2001.

[CDS01] A.E.F. Clementi, M. Di Ianni, and R. Silvestri. Private communication, 2001.

[CFP+00] A.E.F. Clementi, A. Ferreira, P. Penna, S. Perennes, and R. Silvestri. The minimum range assignment problem on linear radio networks. In *Proc. 8th Annual European Symposium on Algorithms*, volume 1879 of *LNCS*, pages 143–154, 2000.

[CPS99] A.E.F. Clementi, P. Penna, and R. Silvestri. Hardness results for the power range assignment problem in packet radio networks. In *Proc. of Randomization, Approximation and Combinatorial Optimization*, volume 1671 of *LNCS*, pages 197–208, 1999. Full version available as ECCC Report TR00-54.

[CPS00] A.E.F. Clementi, P. Penna, and R. Silvestri. The power range assignment problem in radio networks on the plane. In *Proc. 17th Annual Symposium on Theoretical Aspects of Computer Science*, volume 1770 of *LNCS*, pages 651–660, 2000. Full version available as ECCC Report TR00-54.

[CS88] J.H. Conway and N.J.A. Sloane. *Sphere Packings, Lattices and Groups*. Springer-Verlag, 1988.

[DRT92] B. Dixon, M. Rauch, and R.E. Tarjan. Verification and sensitivity analysis of minimum spanning trees in linear time. *SIAM J. Comput.*, 21:1184–1192, 1992.

[Eph00] A. Ephremides. Complicating factors for the use of distributed algorithms in wireless networks. In *1st Int. Workshop on Approximation and Randomized Algorithms in Communication Networks*, invited talk, 2000.

[Epp94] D. Eppstein. Offline algorithms for dynamic minimum spanning tree problem. *J. of Algorithms*, 17:237–250, 1994.

[GK98] S. Guha and S. Khuller. Approximation algorithms for connected dominating sets. *Algorithmica*, 20:347–387, 1998.

[HT98] Z. Haas and S. Tabrizi. On some challenges and design choices in ad hoc communications. In *Proc. IEEE MILCOM'98*, 1998.

[Kan96] G. Kant. Drawing planar graphs using the canonical ordering. *Algorithmica - Special Issue on Graph Drawing*, 16:4–32, 1996. (Extended Abstract in 33-th IEEE FOCS (1992)).

[KKKP00] L.M. Kirousis, E. Kranakis, D. Krizanc, and A. Pelc. Power consumption in packet radio networks. *Theoretical Computer Science*, 243:289–306, 2000. (Extended abstract in Proc. of STACS'97).

[KL78] G.A. Kabatiansky and V.I. Levenshtein. Bounds for packings on a sphere and in space (in russian). *Problemy Peredachi Informatsii*, 14(1):3–25, 1978. English translation: *Problems of Information Theory*, 14(1):1–17, 1978.

[Lau95] G.S. Lauer. *Packet radio routing*, chapter 11 of *Routing in communication networks*, M. Streenstrup (ed.), pages 351–396. Prentice-Hall, 1995.

[MM96] R. Mathar and J. Mattfeldt. Optimal transmission ranges for mobile communication in linear multi-hop packet radio networks. *Wireless Networks*, 2:329–342, 1996.

[MS92] C. Monma and S. Suri. Transitions in geometric minimum spanning tree. *Discrete and Computational Geometry*, 8:265–293, 1992.

[NPW00] E. Nardelli, G. Proietti, and P. Widmayer. Maintainig a minimum spanning tree under transient node failures. In *Proc. 8th Annual European Symposium on Algorithms*, volume 1879 of *LNCS*, pages 346–355, 2000.

[Pap94] C. H. Papadimitriou. *Computational Complexity*. Addison Wesley, 1994.

[Pir91] P. Piret. On the connectivity of radio networks. *IEEE Trans. on Inform. Theory*, 37:1490–1492, 1991.

[PL95] K. Pahlavan and A. Levesque. *Wireless information networks*. Wiley-Interscience, 1995.

[RS97] R. Raz and S. Safra. A sub-constant error-probability low-degree test, and sub-constant error-probability pcp characterization of np. In *Proc. 29th Ann. ACM Symp. on Theory of Comp.*, pages 784–798, 1997.

[RW94] D. Raychaudhuri and N.D. Wilson. ATM-based transport architecture for multiservices wireless personal communication networks. *IEEE J. Selected Areas in Communications*, 12:1401–1414, 1994.

[UY96] S. Ulukus and R.D. Yates. Stocastic power control for cellular radio systems. *IEEE Trans. Comm.*, 46:784–798, 1996.

[Val81] L. Valiant. Universality considerations in vlsi circuits. *IEEE Transactions on Computers*, C-30:135–140, 1981.

[WCLF01] P.-J. Wan, G. Călinescu, X.-Y. Li, and O. Frieder. Minimum-energy broadcast routing in static ad hoc wireless networks. In *Proc. of INFOCOM*, 2001.

[WNE00] J.E. Wieselthier, G.D. Nguyen, and A. Ephremides. On the construction of energy-efficient broadcast and multicast trees in wireless networks. In *Proc. of INFOCOM*, pages 585–594, 2000.

[Wyn65] A.D. Wyner. Capabilities of bounded discrepancy decoding. *BSTJ*, 44:1061–1122, 1965.

A The NP-hardness of MECBS[N_2^α]

In this section, we provide a proof of the NP-hardness of MECBS[N_2^α], for any $\alpha > 1$. The reduction is an adaptation of that in [CPS99]. In particular, we reduce the MIN VERTEX COVER problem to MECBS[N_2^α]. Recall that MIN VERTEX COVER is defined as follows: Given a undirected graph $G = (V, E)$ and a constant k, find a subset $C \subseteq V$ (a *cover*) such that for each $(u, v) \in E$ at least one of u and v belongs to C and the cardinality of C is at most k. The MIN VERTEX COVER problem is NP-hard even when restricted to planar at most cubic (i.e. 3-degree) graphs [AK97]. In order to prove the NP-hardness of MECBS[N_2^α], we will provide a reduction mapping every instance $x = \langle G = (V, E), k\rangle$ of MIN VERTEX COVER restricted to planar at most cubic graphs, to an instance of MECBS[N_2^α].

We first outline which steps have to be performed in order to derive the set of points $S(G)$ for MECBS[N_2^α] corresponding to a planar, at most cubic graph G. To this aim, we will make use of an intermediate representation of G, by means of a planar orthogonal grid drawing $D(G)$ of it. This intermediate step will make the construction of $S(G)$ simpler. The whole construction will basically take the following steps:

1. Construct a planar orthogonal grid drawing of G;
2. Add two new vertices for each bend of the drawing so to obtain a straight-line drawing $D(G)$;
3. Replace each straight-line (edge) in $D(G)$ with a suitable set of stations (gadget).

Notice that in order to obtain a polynomial time reduction we need to perform all the above steps in polynomial time. Given a planar, cubic graph $G = (V, E)$, it is always possible to derive a planar orthogonal drawing of G in which each edge is represented by a polyline having only two bends[1] [Val81, Kan96]. Moreover, in the second step, we have to preserve the optimality of the vertex cover solutions between G and the new graph represented by $D(G)$. Observe that, if $2h$ is the number of vertices added by this operation, then G has a vertex cover of size k if

[1]We use such two-bend drawings to make the reduction simpler. However, any polyline drawing algorithm, such as that in [Val81], also works: indeed the key property here is to have a drawing with a *polynomial* number of bends per edge.

and only if $D(G)$ has a vertex cover[2] of size $k+h$. Finally, in the third step, further vertices will be added in $D(G)$ still preserving the above relationship between the vertex covers of G and those of $D(G)$.

Our goal is to replace each edge (and thus both of its vertices) of $D(G)$ with a gadget of points in the Euclidean space $\mathcal{R}^2$ in order to construct an instance of the $\mathrm{MECBS}[\mathsf{N}_2^\alpha]$ problem and then show that this construction is a polynomial-time reduction.

Roadmap. In Section A.1.1 we provide the key properties of these gadgets and the reduction to $\mathrm{MECBS}[\mathsf{N}_2^\alpha]$ that relies on such properties. In Section A.1.2 we give the detailed construction of the 2-dimensional gadgets (basically the same introduced in [CPS99] but with different gadget parameters). For the sake of clarity, both those steps will be presented for the case $\alpha \geq \overline{\alpha}$, for a suitable constant $\overline{\alpha} >$. Indeed, in this case the adaptation of the reduction is simpler and easier to understand. This will prove that, for any $\alpha > \overline{\alpha}$, $\mathrm{MECBS}[\mathsf{N}_2^\alpha]$ is NP-hard. One more step is needed to deal with the case $\alpha \geq 2$: this is described in Section A.2.

A.1 $\mathrm{MECBS}[\mathsf{N}_2^\alpha]$ is NP-hard for some $\overline{\alpha} > 2$

The first step towards adapting the reduction of [CPS99] is to show that, for a suitable $\overline{\alpha} > 2$, the same gadget construction also works (with some minor changes). This case is also used to present the main ideas of the reduction, while the NP-hardness for the case $\alpha \geq 2$ requires some more technicalities and it will be given in Section A.2.

A.1.1 The Properties of the 2-Dimensional Gadgets and the Reduction

The type of gadget used to replace one edge of $D(G)$ depends on the local "situation" that occurs in the drawing (for example it depends on the degree of its endpoints). However, we can state the properties that characterize any of these gadgets.

Definition 1 (Gadget Properties) *Let $\delta, \delta', \epsilon \geq 0$ such that $\delta + \epsilon > \delta'$ and let $\beta > 1$ be a suitable parameter. For any edge (a, b) the corresponding gadget g_{ab} contains the sets of points $X_{ab} = \{x_1, \ldots, x_{l_1}\}$, $Y_{ab} = \{y_{ab}, y_{ba}\}$, $Z_{ab} = \{z_1, \ldots, z_{l_2}\}$*

[2] In what follows, we will improperly use $D(G)$ to denote both the drawing and the graph it represents.

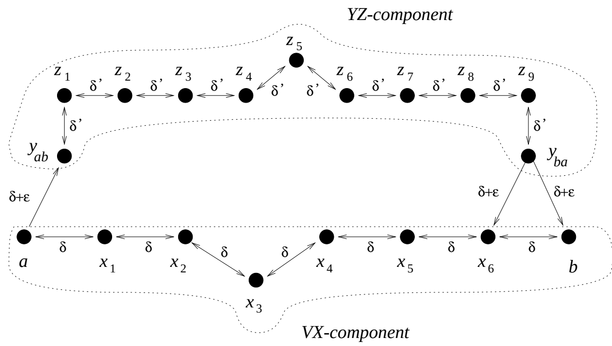

Figure 4: An example of a 2-dimensional gadget and a canonical assignment for it.

and $V_{ab} = \{a, b\}$*, where* l_1 *and* l_2 *depend on the length of the drawing of* (a, b)*. These sets of points are drawn in* $\mathcal{R}^2$ *so that the following properties hold:*

1. $d(a, y_{ab}) = d(b, y_{ba}) = \delta + \epsilon$.

2. X_{ab} *is a chain of points drawn so that*

   $$d(a, x_1) = d(x_1, x_2) = \cdots = d(x_{l_1 - 1}, x_{l_1}) = d(x_{l_1}, b) = \delta$$

   and, for any $i \neq j$, $d(x_i, x_j) \geq \delta$.

3. Z_{ab} *is a chain of points drawn so that*

   $$d(y_{ab}, z_1) = d(z_1, z_2) = \cdots = d(z_{l_2 - 1}, z_{l_2}) = d(z_{l_1}, y_{ba}) = \delta'$$

   and, for any $i \neq j$, $d(z_i, z_j) \geq \delta'$.

4. *For any* $x_i \in X_{ab}$ *and* $z_j \in Z_{ab}$, $d(x_i, z_j) > \delta + \epsilon$*. Furthermore, for any* $i = 1, \ldots, l_1$, $d(x_i, y_{ab}) \geq \delta + \epsilon$ *and* $d(x_i, y_{ba}) \geq \delta + \epsilon$.

5. *Given any two different gadgets* g_{ab} *and* g_{cd}*, for any* $u \in g_{ab} \setminus g_{cd}$ *and* $v \in g_{cd} \setminus g_{ab}$*, we have that* $d(u, v) \geq \delta$ *and if* $u \notin V_{ab} \cup X_{ab}$ *or* $v \notin V_{cd} \cup X_{cd}$ *then* $d(u, v) \geq \beta\delta$.

From the above definition, it turns out that the gadgets consist of two components whose minimum relative distance is $\delta + \epsilon$: the VX*-component* consisting of the "chain" of points in $X_{ab} \cup V_{ab}$, and the YZ*-component* consisting of the chain of points in $Y_{ab} \cup Z_{ab}$.

Let $S(G)$ be the set of points obtained by replacing each edge of $D(G)$ by one gadget having the properties described above. $S(G)$, together with any fixed node s in the VX-component of one gadget, will be our instance for MECBS[N_2^α].

Remark 1 *Let T be* any *spanning tree of $D(G)$. For every gadget g_{ab} we define the following range assignment for VX_{ab}. If (a, b) is in T, then* all *stations in VX_{ab} have range δ. If (a, b) is not in T, then we assign range δ to every stations in VX_{ab}* except one *in X_{ab} that have a null range and its always distinct from the source s. Finally, let any station in the YZ components have range δ'. Notice that this assignment is not feasible: The source s is only able to reach any other node in the union $\mathcal{U}$ of every VX-component. So, what remains to do, in order to obtain a feasible solution, is to choose some "bridge-points" from $\mathcal{U}$ to every YZ-component.*

The above observation leads us to define the following *canonical* (feasible) solutions for $S(G)$.

Definition 2 (Canonical Solutions for $S(G)$) *A range assignment r for $S(G)$ is* canonical *if, for every gadget g_{ab} of $S(G)$, the following properties hold:*

1. *For every $v \in \{a, b\}$, either $r(v) = \delta$ or $r(v) = \delta + \epsilon$. Furthermore, there exists $v \in \{a, b\}$ such that $r(v) = \delta + \epsilon$ (so, v is a radio "bridge" from the VX-component to the YZ one).*

2. *Either $r(y_{ab}) = \delta'$ and $r(y_{ba}) = 0$ (so, y_{ab} allows the broadcasting coming from to the radio "bridge" a along all the YZ one) or vice versa.*

3. *For every $x \in X_{ab}$, either $r(x) = \delta$ or $r(x) = 0$. Moreover, at most one station in X_{ab} has null range and such a station is other than the source s.*

4. *The set $\overline{E} = \{(a, b) \mid \forall x \in X_{ab} : r(x) = \delta\}$ form a spanning tree of $D(G)$.*

5. *For every $z \in Z_{ab}$, $r(z) = \delta'$.*

Informally, our reduction is based on the following facts:

1. If we minimize the number of "bridge" stations in $\mathcal{U}$ then we minimize the overall cost of any canonical solution (observe that the cost of all the X- and YZ-components is fixed);

2. The graph $D(G)$ has a vertex cover of size k' if and only if there exists a canonical solution for $S(G)$ with k' "bridge" stations of type V;

3. It is possible to choose δ, δ', ϵ, and β in order to guarantee that non-canonical feasible solutions can be transformed in polynomial time into a canonical one without paying any extra cost (notice that any canonical assignment is feasible).

In the remaining of this section we will formally prove the above statements. In particular, the most difficult part is to show that any feasible solution can be easily transformed into a canonical one without increasing the cost. Because of the slightly different definition of canonical solution used here, the proof uses a different argument from [CPS99]. Indeed, the proof in [CPS99] also applies to the case in which the cost of a connection is linear (i.e. $\alpha = 1$ in Equation 1), while here we must exploit the fact that $\alpha > 1$ (recall that MECBS[N_d^1]$\in$ P). Mainly because of this difference, we need further gadgets properties:

Definition 3 (Gadget Properties II) *Let $S(G)$ be a set of stations satisfying Definition 1. Then, for any $\alpha > \overline{\alpha}$, the following further properties must be fulfilled by the gadget set $S(G)$:*

1. *Given any two* non-adjacent *gadgets g_1 and g_2, for any $u \in g_1$ and $v \in g_2$, it holds that*

$$w(u,v) = d(u,v)^\alpha \geq \left| \bigcup_{g_{ab} \in S(G)} V_{ab} \right| (\delta+\epsilon)^\alpha + \left| \bigcup_{g_{a,b} \in S(G)} X_{ab} \right| \delta^\alpha + \left| \bigcup_{g_{a,b} \in S(G)} Y_{ab} \cup Z_{ab} \right| (\delta')^\alpha.$$

2. *If a station s has transmission range $r(s) > (\delta+\epsilon)$ and s reaches some station in an* adjacent *gadget, but no station from non-adjacent ones, then such a range contains at most $\lceil r(s)^\alpha/(\delta+\epsilon)^\alpha \rceil$ stations (including s).*

Intuition: Item 1 will be used to prove that, if the range of a station is so large to reach a non adjacent gadget, then its power is larger than the energy that would be payed a canonical (trivial) assignment (see Lemma 4). Item 2, instead, will be used to show that any solution in which a station has a (non-canonical) range reaching an adjacent gadget can be "locally" transformed into a canonical solution without increasing its cost (see the proof of Lemma 5).

Lemma 4 *Let $S(G)$ be a set of stations satisfying Definitions 1 and 3. Then, any (non canonical) assignment r_{BAD}, such that a station directly reaches another station in a non-adjacent gadget, is not optimal.*

Proof. In order to prove the lemma, we show that there exists a *canonical* assignment $\overline{r}$ that is computable in polynomial time and such that $\mathsf{cost}(\overline{r}) < \mathsf{cost}(r_{BAD})$. For each $g_{ab} \in S(G)$, the *canonical* assignment $\overline{r}$ is defined as follows: If $v \in V_{ab}$ then $\overline{r}(v) = \delta + \epsilon$; Otherwise, $\overline{r}(v)$ is chosen according to Definition 2 (notice that the overall cost of this part is fixed). Then, from the definition of $\overline{r}$ and from Item 1 of Definition 3, it follows that

$$\begin{aligned} \mathsf{cost}(\overline{r}) &< \left| \bigcup_{g_{ab} \in S(G)} V_{ab} \right| (\delta + \epsilon)^{\alpha} + \left| \bigcup_{g_{a,b} \in S(G)} X_{ab} \right| \delta^{\alpha} \\ &+ \left| \bigcup_{g_{a,b} \in S(G)} Y_{ab} \cup Z_{ab} \right| (\delta')^{\alpha} \leq \mathsf{cost}(r_{BAD}). \end{aligned}$$

□

Lemma 5 *For any feasible range assignment r for $S(G)$, there is a* canonical *range assignment r^c for $S(G)$ such that $\mathsf{cost}(r^c) \leq \mathsf{cost}(r)$.*

Proof. The proof essentially mimics that in [CPS99]. In particular, we show that if the range of some station $u \in S(G)$ does not match Definition 1 then it is possible to transform the range assignment into a canonical one without increasing its cost. From Lemma 4, we can assume that u does not reach any non-adjacent gadget. Hence, only two cases can arise:

1. u reaches only other stations of its gadget;

2. u reaches some stations in a gadget adjacent to its.

In the first case, we can perform a "local transformation" similar to that in [CPS99]. Otherwise, let $C(u)$ be the set of stations inside the transmission range of u. Since u reaches only adjacent gadgets, Item 2 of Definition 3 implies that $|C(u)| \leq \lceil r(u)^{\alpha}/(\delta + \epsilon)^{\alpha} \rceil - 1$. We now change the assignment r into r^c as follows:

1. $r^c(u) = \delta + \epsilon$;

2. For every $v \in C(u)$, $r^c(v) = \max\{\delta + \epsilon, r(u)\}$.

We now show that $\mathsf{cost}(r) \geq \mathsf{cost}(r^c)$. Indeed, we have that

$$\begin{aligned}\mathsf{cost}(r) - \mathsf{cost}(r^c) &\geq r(u)^\alpha - [r^c(u)]^\alpha - |C(u)|(\delta+\epsilon)^\alpha \\ &\geq r(u)^\alpha - (\delta+\epsilon)^\alpha - \left[\left(\frac{r(u)}{\delta+\epsilon}\right)^\alpha - 1\right](\delta+\epsilon)^\alpha = 0.\end{aligned}$$

This completes the proof. □

A.1.2 Existence of the Gadgets: Adapting the construction in [CPS99].

The following lemma states that it is possible to adapt the 2-dimensional gadgets construction in [CPS99] in order to also match Definition 3.

Lemma 6 *There exists an $\overline{\alpha} > 2$ such that the following holds. For any at most cubic graph G and for any $\alpha > \overline{\alpha}$, there exists a set of gadgets $S(G)$ satisfying Definitions 1 and 3. Moreover, $S(G)$ can be constructed in time polynomial in $|G|$.*

Proof. We first observe that the original construction in [CPS99] already satisfies Definition 1. Hence, we show how to modify any gadget set $S(G)$ yielded by that reduction in order to satisfy Definition 3 as well.
Item 1. First of all, the set of stations $S(G)$ has the following properties:

1. $S(G)$ has $m \leq n^g$ gadgets, with n being the number of vertices of G, and for some constant g (independent on G);

2. Every gadget in $S(G)$ has at most n^s stations, for some constant g (independent on G);

3. Given any two *non adjacent gadgets* g_1 and g_2, for any $u \in g_1$ and $v \in g_2$, $d(u, v) > L_{\min}$ (for a suitable constant $L_{\min} > 2(\delta + \epsilon)$).

Then, given the set of gadgets $S(G)$ with parameters δ and δ' and ϵ, we modify $S(G)$ as follows. We first observe that, for every $\delta_{new} < \delta$ and $\delta'_{new} < \delta'$, we can replace every $g \in S(G)$ with a new gadget g_{new} with parameters δ_{new}, δ'_{new} and ϵ_{new} (see the example in Figure 5). This transformation preserves the gadget properties of Definition 1 (i.e. if $S(G)$ satisfies those gadget properties, then the new set of gadget also does). So, let be $p = g + s$; we set $\delta_{new} = \delta/n^p$, $\delta'_{new} = \delta'/n^p$ and $\epsilon_{new} = \epsilon/n^p$. Hence, we are re-scaling *only the gadgets* (but not the drawing $D(G)$, that is, the position of the gadgets endpoints) by a factor $1/n^p$.

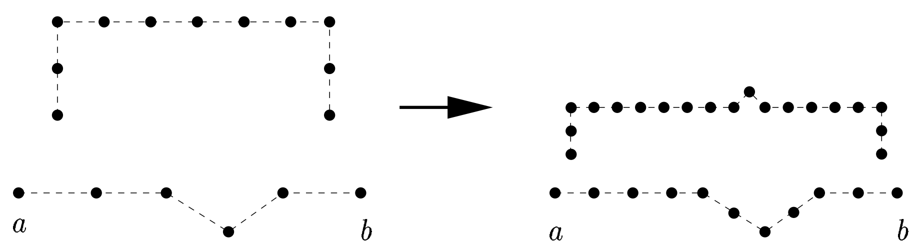

Figure 5: The proof of Lemma 6: How to rescale a gadget.

This can be done by adding, for each gadget, a polynomial number of X- and Y- points (notice that nor the position of the V-points nor their number change). Hence, the new set of gadget (also) satisfies the following property:

> Given any two *non adjacent gadgets* g_1 and g_2, for any $u \in g_1$ and $v \in g_2$, $d(u,v) > L_{\min} = 2(\delta + \epsilon) = 2n^p(\delta_{new} + \epsilon_{new})$.

By assuming $\alpha \geq \overline{\alpha} > 2$,[3] the above property implies

$$w(u,v) = d(u,v)^\alpha \geq 2^{\overline{\alpha}} n^{\overline{\alpha}p} (\delta + \epsilon)^\alpha > n^{\overline{\alpha}p} \left[(\delta + \epsilon)^\alpha + \delta^\alpha + (\delta')^\alpha \right], \qquad (9)$$

where the last inequality follows from the fact that $\delta + \epsilon \geq \delta'$ (see Definition 1). This proves Item 1 of Definition 3.

Item 2. For any station s, we consider increasing ranges such that at least one new station is covered, as long as only stations in *adjacent* gadgets are reached. Starting from its canonical range (depending on the type of s), we define $r_i(s)$ as the ith range in this sequence (with r_1 be the smallest non canonical range for which s reaches at least one station more than in the canonical assignment). Moreover, for the stations of type V, we assume its canonical range to be equal to $\delta + \epsilon$. Let also $n_i(s)$ be the number of stations (including s) covered by s when its range is equal to $r_i(s)$. Based on the fact that all the gadgets are drawn in the same way (actually, we have two sets of gadgets, those of degree two and those of degree three) we will prove that there exists a constant $\overline{\alpha}$ such that

$$n_i(s)(\delta + \epsilon)^{\overline{\alpha}} \leq r_i(s)^{\overline{\alpha}}. \qquad (10)$$

Let us first assume that $n_i(s)$ is bounded above by a function *linear* in $r_i(s)$ (let us say $c \cdot r_i(s)$). Then, it would be enough to show that

$$n_i(s)^{\overline{\alpha}} \leq c^{\overline{\alpha}} r_i(s)(\delta + \epsilon)^{\overline{\alpha}} \leq r_i(s),$$

[3] It is easy to see that the same result also holds for any $\alpha \geq 2$.

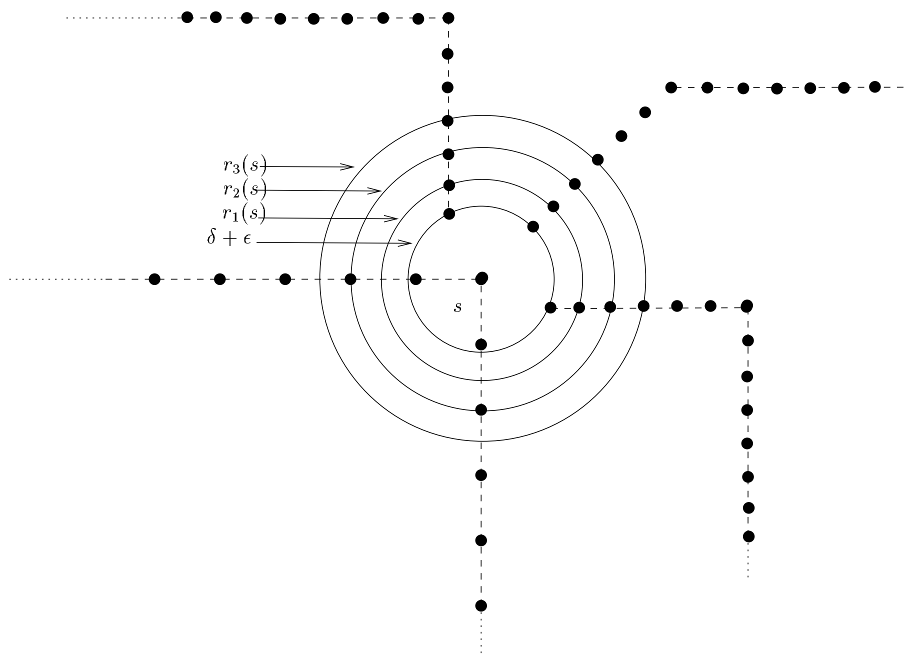

Figure 6: A 3-degree gadget from [CPS99].

which holds for a sufficiently large $\overline{\alpha}$. This linearity of $n_i(s)$ is true up to some range, let us say $\overline{r}$. (For the station s in Figure 6, this is true as long as the range of s does not reach the "bend" of some YZ-chain.) However, it is easy to see that, for a suitable $\overline{\alpha}$, any range bigger than $\overline{r}$ is more costly than assigning a range $\delta+\epsilon$ to *every* station in any (at most) three adjacent gadgets. In this case Equation 10 still hold, since we assume that s only covers stations from adjacent gadgets. This completes the proof. □

Putting things together, we get the following result.

Theorem 6 *There exists an* $\overline{\alpha} > 2$ *such that, for any* $\alpha \geq \overline{\alpha}$, MECBS[$\mathsf{N}_2^\alpha$] *is* NP-*hard.*

Proof. Let us consider the graph $D(G)$ and let us denote by V' and E' its set of vertices and edges, respectively. Recall that $D(G)$ is a planar 3-degree graph. Here we provide a polynomial time reduction from this instances of MIN VERTEX COVER to MECBS[N_2^α]. By construction, G has a vertex cover of size k if and only if $D(G)$ has a vertex cover of size $k+h$, where $2h$ is the number of new vertices added to G in the construction of $D(G)$. We can therefore consider the problem of finding an optimum vertex cover of $D(G)$. From Lemma 5, for any fixed source node s in the $\mathcal{U}$ component of the instance $S(G)$, we can restrict ourselves to canonical solutions of MECBS[N_2^α].

Given any vertex cover $K \subseteq V'$ for $D(G)$, we consider the canonical solution r_K for $S(G)$ where every $v \in K$ has range $\delta + \epsilon$ and every $w \in V' \setminus K$ has range δ. Further, the range assignment r_K for all the other points is made according to the definition of the canonical solution. Notice that the cost of this part of the assignment is *fixed*. Indeed, the overall power given to the points of type Z and Y in given by

$$M_{S(G)} = \sum_{(a,b)\in E'} \left[(|YZ_{ab}| - 1)(\delta')^\alpha\right].$$

Moreover, from items 3-4 of Definition 2 and from the fact that any spanning tree of $D(G)$ has $|V'|-1$ edges, we have that the number of stations in the VX-components that have range δ or $\delta+\epsilon$ is

$$\begin{aligned} N_{S(G)} &= |V'| + \left|\bigcup_{(a,b)\in E'} X_{ab}\right| - (|E'| - |V'| + 1) \\ &= 2|V'| + \left|\bigcup_{(a,b)\in E'} X_{ab}\right| - |E'| - 1. \end{aligned}$$

So, the *overall cost* $\mathsf{cost}(r_K)$ is given by

$$\mathsf{cost}(r_K) = |K|(\delta+\epsilon)^\alpha + (N_{S(G)} - |K|)\delta^\alpha + M_{S(G)}. \tag{11}$$

On the other hand, let r_K be any canonical solution for $S(G)$ and let κ be the number of points of type V whose range is $\delta+\epsilon$. Then, the cost of such a solution is

$$\kappa(\delta+\epsilon)^\alpha + (N_{S(G)} - \kappa)\delta^\alpha + M_{S(G)}.$$

Let K be the set of κ vertices of $D(G)$ corresponding to those stations whose solution r_K assigns range $\delta+\epsilon$. We now prove that K is a vertex cover. Suppose by contradiction that some edge (a,b) is not covered (i.e. both a and b are not in K). In the solution r_K, we should have $r_K(a) = r_K(b) = \delta$, thus contradicting the fact that r_K is canonical. □

A.2 MECBS[N_2^α] is NP-hard for any $\alpha \geq 2$

It is easy to verify that the gadget in Figure 6 does not satisfy Item 2 of Definition 3 for $\alpha = 2$. Instead of further modifying the gadget to deal with small values of α, we wish to prove Lemma 5 *without* such a property (Item 2). Notice that, once such a Lemma is proved, the remaining of the reduction still work. Moreover, in (re-)proving Lemma 5, we only have to (re-)consider the case of a non-canonical range that reaches some stations in *adjacent* gadgets. To this aim, by using simple geometric arguments, we will show that any non-canonical assignment of this type can be locally converted into a canonical one. Figure 7 shows an example of such transformation: a V-station b has a range bigger than the canonical one (in this case $\delta + \epsilon$) and it reaches several X- and YZ-stations; we can assign to any such station (including b) a canonical range thus obtaining a cheaper solution.

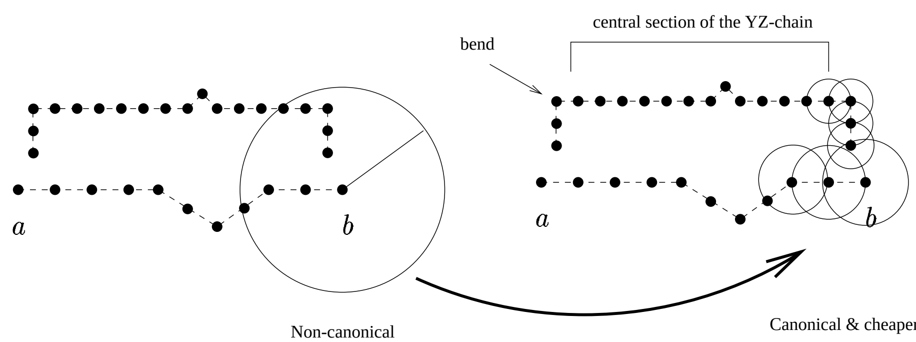

Figure 7: Locally reducing a non-canonical range.

In the rest of this section we will prove that this transformation is possible whenever a station reaches only stations in adjacent gadgets. The overall proof goes through the following main steps:

1. For every YZ-chain we define its *bend* and its *central section* as shown in Figure 7. By using a gadget rescaling technique similar to that used in the proof of Lemma 6, we can make any range from a VX-station to the central section (or the other way around) too expensive, that is, larger than the canonical cost of connecting any three (or two) adjacent gadgets. (Notice that this step uses the assumption $\alpha > 1$.)

2. By adding intermediate stations (similarly to the rescaling of the whole gadget) to the YZ-chains *only*, we can make any range covering only some new

YZ-station more expensive than some smaller range not covering such stations plus the canonical assignment to them. (Again, we use the fact that $\alpha > 1$.)

3. By virtue of Item 1 above, we need to consider YZ-stations reaching some X station only if those YZ-stations are not in the central section. As for the VX-stations, because of Item 2 above, we only have to consider a sequences of increasing ranges that include one new VX-station w.r.t. the previous one. By reffering back to Figure 6, we would not have to consider that range $r_1(s)$ and $r_3(s)$, because they only include some new YZ-station w.r.t. $r_0(s)$ and $r_2(s)$, respectively.

4. We "bend" the gadgets shape (as shown in Figure 8) so that any two adjcent VX-chains form an angle strictly bigger than $\pi/2$. Notice that, for any V-station b, where b corresponds to a vertex of degree three (i.e. its gagdet is also adjacent to other two), we can keep the angle between b and any two Y-stations in the adjacent gadgets arbitrarily close to $\pi - \pi/3 = 2\pi/3$. The latter property guarantees that we keep different YZ-chains sufficiently far apart.

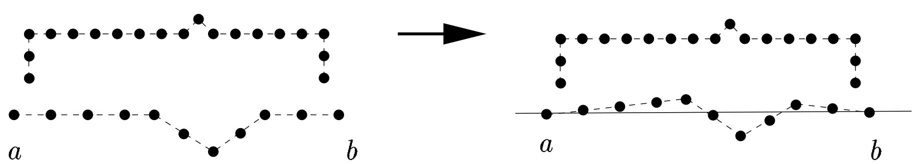

Figure 8: A further transformation for a gadget.

Since any of those steps preserves the properties of Definition 1 and Item 1 of Definition 3, we can prove the following lemma.

Lemma 7 *Let $S(G)$ be a set of gadget obtained by applying the above transformation to a gadget set satisfying Definition 1 and Item 1 of Definition 3. Then, any feasible range assignment for $S(G)$ can be trnaformed (in polynomial time) into a canonical one and whose cost in not bigger.*

As a consequence of the above lemma, we can (re-) prove[4] Theorem 6 for a more general case:

[4]The proof is actually the same, since the above lemma allows us to restrict to canonical assignments.

Theorem 7 *For any* $\alpha \geq 2$, $\mathrm{MECBS}[\mathtt{N}_2^\alpha]$ *is* NP-*hard.*