Chapter 13
Energy Consumption Minimization in Ad Hoc Wireless and Multi-interface Networks

Alfredo Navarra, Ioannis Caragiannis, Michele Flammini, Christos Kaklamanis, and Ralf Klasing

Abstract This chapter deals with energy consumption issues in wireless networks. In such networks, energy is a scarce resource and, hence, it must be used efficiently. Under these circumstances, we consider two interesting combinatorial optimization problems: *Minimum Energy Broadcast Routing* and *Cost Minimization in Multi-interface Networks*. The goal of the first problem is to perform broadcasting from a given source while minimizing the overall energy required for communication. The second problem refers to the choice of activating a set of available communication interfaces at the network nodes in order to satisfy the required connections in a wireless multi-interface network with minimum total cost. While Minimum Energy Broadcast Routing has been studied extensively during recent years, Cost Minimization in Multi-interface Networks is rather new. For both problems we survey recent complexity results and approximation algorithms under different assumptions.

Key words: energy saving, wireless networks, multi-interface networks, broadcasting, approximation algorithms

---

Alfredo Navarra
Dipartimento di Matematica e Informatica, Università degli Studi di Perugia, Via Vanvitelli 1, 06123 Perugia, Italy, e-mail: `navarra@dmi.unipg.it`

Ioannis Caragiannis · Christos Kaklamanis
Computer Technology Institute & Department of Computer Engineering and Informatics, University of Patras, 26500 Rio, Greece, e-mail: `{caragian,kakl}@ceid.upatras.gr`

Michele Flammini
Dipartimento di Informatica, Università degli Studi dell'Aquila, Via Vetoio, Loc. Coppito, 67100 L'Aquila, Italy, e-mail: `flammini@di.univaq.it`

Ralf Klasing
CNRS - LaBRI - Université Bordeaux 1, France, e-mail: `klasing@labri.fr`

13.1 Introduction

In recent years *wireless networks* have been widely deployed, mostly because of the recent drop in equipment prices and due to the features provided by the new technologies. Unlike traditional *wired networks* where signals pass from one device to another through physical cables, in wireless networks data transmissions from each node (station) occur in the open air and within a given coverage area. In this scenario, considerable attention has been devoted to the so-called *ad hoc* wireless networks, due to their potential applications in emergency disaster relief, on the battlefield, in impervious areas, and so on [29, 44]. Ad hoc wireless networks do not require any fixed infrastructure. The network is simply a collection of devices that can communicate with each other according to proximity and available common protocols and interfaces.

In this chapter we consider two important problems arising in the context of ad hoc wireless networks: *Minimum Energy Broadcast Routing* and *Cost Minimization in Multi-interface Networks*. The first problem considers the need of broadcasting information from a given source to all other network nodes when the network nodes are equipped with omnidirectional antennas. The second problem aims at establishing connections among heterogeneous nodes equipped with a set of interfaces that can be activated at a different cost.

The common objective in both the above scenarios is to minimize the total energy consumption in order to keep the network alive as long as possible. Energy is in fact a scarce resource in wireless ad hoc networks, and communication strongly depends on it.

The chapter is organized as follows. Section 13.2 is devoted to the Minimum Energy Broadcast Routing problem. First, some motivation for the problem is provided and, then, it is formally defined. Several results are surveyed and details are given for interesting analysis techniques. Section 13.3 is devoted to Cost Minimization in Multi-interface Networks. Again, the problem is first motivated, then formally defined, and a list of results is presented emphasizing some interesting techniques. Finally, Section 13.4 provides conclusive remarks and interesting directions of future research.

13.2 Minimum Energy Broadcast Routing

The study of a basic communication pattern such as broadcasting is of main interest in the context of wireless ad hoc networks. Broadcasting can in fact be used to set up the network or to rapidly spread useful information. The wireless environment allows all devices in the range of a transmitter x to receive messages sent by x. The range of transmissions basically depends on the environment where devices are distributed. According to the widely used power attenuation model [42], when a station s transmits with power P_s, a station r can receive messages from s if and only if $P_s > \beta dist(s,r)^{\alpha}$, where $dist(s,r)$ is the Euclidean distance between s and

r, α is a parameter which depends on the environment with typical values between 2 and 6, and β is a positive parameter known as the *reception quality threshold*. For the sake of simplicity, from now on we normalize parameter β to 1. Due to the nonlinearity of power attenuation, multi-hop transmission of messages through intermediate devices may result in energy conservation. The main property of wireless ad hoc networks is usually the lack of a fixed infrastructure for routing purposes. A natural issue arising in this setting is that of supporting broadcasting with minimum total energy consumption. This problem, called *Minimum Energy Broadcast Routing* (MEBR), is well-known in the literature and it has been extensively studied (see [1, 2, 5, 7–9, 12, 13, 16, 19, 20, 24–26, 31, 33, 36, 41]).

13.2.1 Definitions and Notation

Given a set of stations S, let $G(S)$ be the complete weighted directed graph whose nodes are the stations in S and in which the weight $w(x,y)$ of each edge (x,y) is the power required at x in order to transmit correctly to node y. A power assignment for S is a function $p : S \to \mathbb{R}_+$ assigning a transmission power $p(x)$ to every station x in S. A power assignment p for S yields a directed communication graph $G^p = (S,A)$ such that, for each directed edge (x,y) of $G(S)$, (x,y) belongs to A if and only if $p(x) \geq w(x,y)$, i.e., if x can correctly transmit to y. In this case, we say that y is within the range of x. The total cost of a power assignment p is then

$$cost(p) = \sum_{x \in S} p(x).$$

MEBR takes as input $G(S)$ together with a source station $s \in S$ and consists of finding a power assignment p of minimum cost such that G^p contains a directed spanning tree rooted at s (and directed towards the leaves). We call such a power assignment an *optimal power assignment* and denote its cost by $m^*(S,s)$.

The weight function $w : E \mapsto \mathbb{R}_+$ is usually symmetric (i.e., $w(x,y) = w(y,x)$ for each pair of stations x, $y \in S$). Nonsymmetric weight functions can be used to capture the irregularity of the environment or situations where stations use batteries of different types which may operate on different fixed energy levels. An important case of symmetric weight functions arises in the geometric version of MEBR. In this case, the stations of S correspond to points in a $d-$dimensional Euclidean space and the weight function is defined as $w(x,y) = dist(x,y)^\alpha$, where $dist$ is the Euclidean distance and $\alpha \geq 1$ is a positive parameter. Equivalently, in this case, we seek a range assignment $r : S \to \mathbb{R}_+$ such that the range $r(x)$ of a station x denotes the maximum Euclidean distance from x at which signals can be correctly received. Again, a range assignment r for S yields a directed communication graph $G^r = (S,A)$ such that the directed edge (x,y) belongs to A if and only if y is at distance at most $r(x)$ from x. We use the notation $G_\alpha(S)$ and $m^*_\alpha(S,s)$ to denote the input graph and the cost of the optimal range assignment in the geometric case.

13.2.2 The Geometric Version of MEBR

We first study the geometric version of MEBR. In general, the geometric version of MEBR is NP-hard, while it is solvable in polynomial time when $\alpha = 1$ or $d = 1$ [13, 20]. Many heuristics and corresponding experimental results can be found in the literature. We can find the Shortest Paths Tree (SPT), the Minimum Spanning Tree (MST), and the Broadcast Incremental Power (BIP) in [42]; the Iterative Maximum-Branch Minimization (IMBM) in [37]; the Adaptive Broadcast Consumption (ABC) in [33]; a refined BIP version in [43]; and many Integer Linear Programming approaches like the ones in [22, 30, 33, 43] (see also [2] for a comparative experimental study).

While such heuristics have been observed experimentally to perform pretty well on random instances of MEBR, the only one for which extended analytical studies were done is the MST heuristic. It is based on the idea of tuning ranges so that the communication graph contains a minimum spanning tree (see Section 1.5.1.1) of the cost graph $G(S)$. More precisely, denote by $T(S)$ a minimum spanning tree of $G(S)$. The MST heuristic considers $T(S)$ rooted at the source station s, directs the edges of $T(S)$ towards the leaves, and sets the power $p(x)$ of every internal station x of $T(S)$ with $k > 0$ children $x_1, \ldots, x_k$ in such a way that $p(x) = max_{i=1,\ldots,k} w(x, x_i)$. In other words, p is the power assignment of minimum cost inducing the directed tree derived from $T(S)$, and is such that $cost(p) \leq c(T(S))$, where $c(T(S))$ denotes the total cost of the edges in $T(S)$. Therefore, the approximation ratio of the heuristic is bounded by the ratio between the cost of a minimum spanning tree of $G(S)$ and the optimal power cost $m^*(S,s)$.

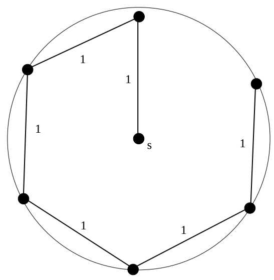

Fig. 13.1 The lower bound of 6 for the MST heuristic. On this instance, there is a source station s and six additional stations, five of which are on the circumference of a circle of radius $1+\varepsilon$ centered at s. The MST heuristic produces a tree (path) consisting of six edges of total cost 6. The optimal solution is a star connecting s to the other six nodes at a cost of $1+\varepsilon$. The approximation ratio can become arbitrarily close to 6 by selecting ε to be arbitrarily small

For the two-dimensional case (which is actually the case that has been given most attention in the literature), a lower bound of 6 on the approximation ratio of the MST heuristic was provided in [42] (see Figure 13.1). The first constant upper bound was provided in [19]. The analysis led to an approximation ratio of 40 for

MST, immediately reduced to 20 by the same authors. The general idea behind the analysis is to represent each edge of $G(S)$ chosen by the MST heuristic by a two-dimensional shape and evaluate the area occupied by all such shapes, since the cost of a solution provided by the MST heuristic (for $\alpha = 2$) is proportional to the considered area. The 40-approximation was obtained by associating with each edge *e* a circle whose diameter is the length of *e* (see Figure 13.2.a). The 20-approximation arises by associating a circle with each edge *e* whose center is the middle point of *e* and diameter is half the length of *e* (see Figure 13.2.b). A further improvement obtained by using the same technique but varying the shape was given in [33, 41]. The obtained approximation ratio is 12.15 and the shape associated with each edge *e* is a rhombus whose bigger diagonal coincides with *e*, and the angles at its endpoints formed by the sides of the rhombus are 60 degrees (see Figure 13.2.c). Note that the third shape implies an interesting property for which no overlap occurs among two shapes associated with two different edges of the minimum spanning tree. By refining the geometrical arguments but without changing the rhombus shape, a better bound of 10.86 was obtained in [8].

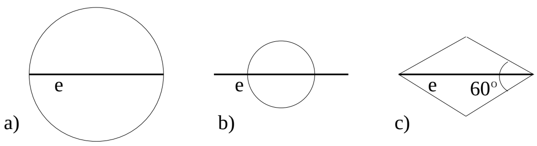

Fig. 13.2 The three shapes associated with an edge *e* of a spanning tree. The choice of shape a) leads to a 40-approximation, b) to 20, and c) to 12.15.

A different method was used in [24, 31]. A new process to evaluate the approximation factor of the MST heuristic was introduced, leading first to an upper bound of 8 and then to 6.33 by more refined geometrical arguments. The basic idea, which will be explained in detail in the next section, is to grow one circle centered at each node of the network until all the circles belong to the same connected component, that is, the union of all the circles forms one and only one delimited area. The bounds of the MST heuristic were then determined by evaluating the covered area by such a process. This method was also extended to the more general d-dimensional case for which a $(3^d - 1)$-approximation ratio has been obtained. It is worth mentioning that, for any Euclidean dimension d and power $\alpha \geq d$, the MST heuristic is lower-bounded by the so-called d-dimensional *kissing number* [19]. More precisely, the d-dimensional kissing number is the maximum number of mutually disjoint d-spheres (or hyperspheres) of a given radius r that can simultaneously touch a d-sphere of the same radius r in the d-dimensional Euclidean space [21]. For the three-dimensional Euclidean space, for instance, the kissing number is 12 while the upper bound provided by [24, 31] is 26. Such a bound has been improved in [38] to 18.8 by extending to the three-dimensional space arguments similar to the ones that led the upper bound for the two-dimensional case from 8 to 6.33. In [1] the gap between the lower bound of 6 and the upper bound of 6.33 for the MST heuristic in the two-dimensional case was finally closed by decreasing the upper bound to 6. The author provided a new analysis technique based on the Delaunay triangulation.

An interesting direction for providing worst-case scenarios and studying the performance of the MST or other heuristics was given in [26]. By means of a randomized procedure, the approach showed an almost tight 4-approximation ratio of the MST heuristic in the case of uniform high-density distributions of the radio stations.

13.2.3 An 8*-Approximation Upper Bound for the MST Heuristic*

In this section, we present the main ideas that can be used in the analysis of the MST heuristic. The particular arguments used in the following yield an upper bound of 8 on the approximation ratio and are relatively easy to follow. The improved results in [31] and [1] follow significantly more involved analysis.

Given a graph G and a weight function w defined on its edges, for any $c \in \mathbb{R}_+$, let $N(G,c)$ be the number of connected components in the graph obtained from G by keeping only the edges $e \in E$ such that $w(e) \leq c$. Then, the cost $MST(G)$ of a minimum spanning tree for G is given by the following lemma.

Lemma 13.1 (Frieze and McDiarmid [27]). $MST(G) = \int_0^\infty (N(G,c)-1)dc.$

For any subset of stations $Q \subseteq S$, let $G_\alpha(Q)$ be the subgraph of $G_\alpha(S)$ induced by Q. Also, let $n(Q,r) = N(G_\alpha(Q), r^\alpha)$, that is, the number of connected components in $G_\alpha(Q)$ obtained by maintaining only the edges between the nodes at distance at most r in Q. Recalling that each edge (x,y) has cost $w(x,y) = dist(x,y)^\alpha$ and exploiting Lemma 13.1 by substituting the variable c with r^α, we obtain the following corollary.

Corollary 13.1 (Klasing et al. [31]). *For any subset of stations* $Q \subseteq S$, $MST(G_\alpha(Q)) = \alpha \int_0^\infty (n(Q,r)-1)r^{\alpha-1}dr.$

Now, the main argument in the proof is the following. We address the case $\alpha = 2$ since the upper bound for $\alpha \geq 2$ can be directly inferred [24, 31]. Consider an optimal power assignment of cost $m^*(S,s)$ in which k stations $x_1,\dots,x_k$ are assigned nonzero power. For $i = 1,\dots,k$, denote by r_i the range of station x_i and let Q_i be the set of stations within the range of station x_i. We will show that $MST(G_\alpha(Q_i)) \leq 8r_i^\alpha$. In this way, we will obtain that

$$MST(G_\alpha(S)) \leq \sum_{i=1}^{k} MST(G_\alpha(Q_i)) \leq 8\sum_{i=1}^{k} r_i^\alpha = 8m^*(S,s),$$

thus proving the upper bound. Without loss of generality, we consider a set of stations $Q \subseteq S$ such that there exists a station $x \in Q$ with $\max_{y\in Q} dist(x,y) = 1$. Thus, all the points of Q belong to a circle of radius 1 in the plane, from now on denoted by C_1, and the cost of each edge of the weighted complete graph representing the input network is proportional to the square of the distance between its endpoints.

For simplicity of notation, for any set of stations Q, $G_2(Q)$ is denoted simply as $G(Q)$.

Let $G(Q,r)$ be the graph obtained by considering only the edges of $G(Q)$ of length at most r, or equivalently of cost at most r^2, and let $CC(Q,r)$ be the set of the connected components of $G(Q,r)$. Let r_{max} be the minimum r such that $G(Q,r)$ is connected (i.e., $n(Q,r_{max}) = |CC(Q,r_{max})| = 1$). Then, directly from Corollary 13.1,

$$MST(G) = 2\int_0^{r_{max}} (n(Q,r)-1)rdr.$$

For the sake of readability, from now Q is dropped from the notation, so that G, $G(r)$ $CC(r)$, $n(r)$, and r_{max} will denote $G(Q)$, $G(Q,r)$ $CC(Q,r)$, $n(Q,r)$ and $r_{max}(Q)$, respectively.

Theorem 13.1 ([31]). *Given any subset of stations $Q \subseteq S$ within a circle of radius 1, $MST(G(Q)) \leq 8$.*

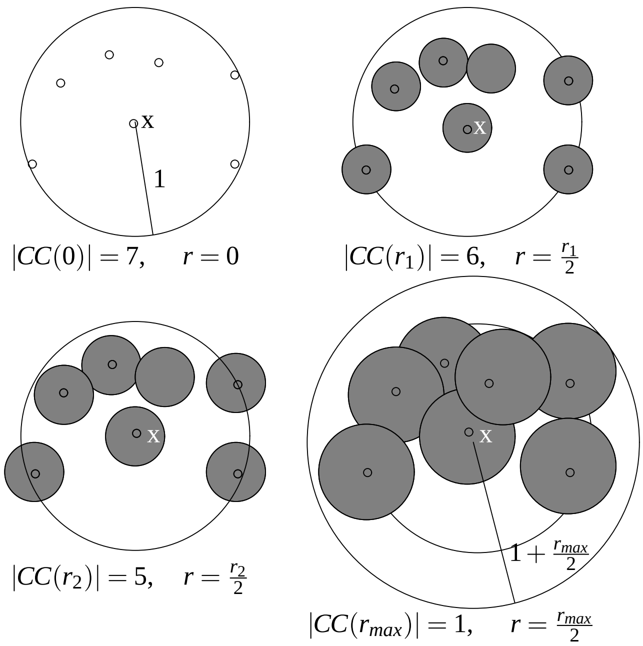

Fig. 13.3 The expanding process described in Theorem 13.1

The proof of Theorem 13.1 considers a growing process in which circles of equal radii centered at the stations of Q are synchronously grown starting from a radius $r = 0$ till $r = \frac{r_{max}}{2} \leq \frac{1}{2}$; i.e., the process ends when $G(2r)$ becomes connected. This is accomplished by increasing at any infinitesimal step the current radii, all equal to a given r, by dr.

For a set of stations P and a radius r, let $P(r)$ be the set of the points in the plane contained in the union of all the circles of radius r associated with the stations of P. Let $a(P,r)$ be the area of $P(r)$.

Suitable lower and upper bounds on the total area $a(Q, \frac{r_{max}}{2})$ covered by all the circles related to Q at the end of the process, that is, when all the radii are equal to $\frac{r_{max}}{2}$, can be determined as follows.

Consider a connected component $P \in CC(2r)$ of $G(2r)$. Then, since two circles of radius r centered in two stations at distance at most $2r$ touch each other, $P(r)$ corresponds to a closed region of the plane having perimeter $p(P,r)$ equal to at least $2\pi r$, that is, equal to at least the perimeter of a single circle of radius r centered in a station of P. Thus, the increase $p(P,r)dr$ of $a(P,r)$ when r is increased by an infinitesimal step dr is at least $2\pi r dr$.

If $P \in CC(2r)$ and $R \in CC(2r)$ are two different connected components of $G(2r)$, $P(r) \cap R(r) = \emptyset$, as any point belonging to the intersection would contradict the fact that the distance between any station in P and any station in R is strictly greater than $2r$. Therefore,

$$\begin{aligned}
a(Q, \frac{r_{max}}{2}) &= \int_0^{\frac{r_{max}}{2}} \sum_{P \in CC(2r)} p(P,r)dr \\
&\geq \int_0^{\frac{r_{max}}{2}} n(2r) 2\pi r dr \\
&= \frac{1}{4} 2\pi \int_0^{r_{max}} n(r) r dr \\
&= \frac{1}{4} 2\pi \int_0^{r_{max}} (n(r)-1) r dr + \frac{1}{4} 2\pi \int_0^{r_{max}} r dr \\
&= \frac{\pi}{4} MST(G) + \frac{\pi}{4} r_{max}^2.
\end{aligned}$$

Moreover, the total region of the plane covered by the union of all the circles related to Q of radius $\frac{r_{max}}{2}$, that is, $Q(\frac{r_{max}}{2})$, is included in a circle of radius $1 + \frac{r_{max}}{2}$ centered at the station x such that $dist(x,y) \leq 1$ for every $y \in Q$ (see Figure 13.3), so that $a(Q, \frac{r_{max}}{2}) \leq \pi + (\frac{r_{max}}{2})^2$. Therefore,

$$\frac{\pi}{4} MST(G) + \frac{\pi}{4} r_{max}^2 \leq a(Q, \frac{r_{max}}{2}) \leq \pi (1 + \frac{r_{max}}{2})^2$$

and thus, since $r_{max} \leq 1$,

$$MST(G) \leq 4(1 + \frac{r_{max}}{2})^2 - r_{max}^2 = 4(1 + r_{max}) \leq 8.$$

13.2.4 Experimental Studies with the MST Heuristic

We now show an interesting technique from [26] for obtaining "bad" instances for the MST heuristic. The goal is to maximize the cost of a possible MST inside C_1 considering its center s as the source. This was done in order to better understand the actual quality of the performance of the MST heuristic over interesting instances that are more representative of real-world applications. Starting from random instances, the maximization consists of slight movements of the nodes according to some useful properties of the MST construction. For instance, if we want to increase the cost of an edge of the MST, the easiest idea is to increase the distance of its endpoints. Let us now consider a node $v \neq s$ of a generic instance given as input. We consider the degree of such a node in the undirected tree obtained from the MST heuristic before assigning the directions. Let $N(v) = \{v_1, v_2, \ldots, v_k\}$ be the set of the neighbors of v in such a tree. We evaluate the median point $p = (x, y)$, whose coordinates are given by the average of the corresponding coordinates of the nodes in $N(v)$. The idea is then to move the node v farther from p but, of course, inside the considered circle. In general this should augment the cost of the MST on the edge connecting the node v to the rest of the tree (see Figure 13.4).

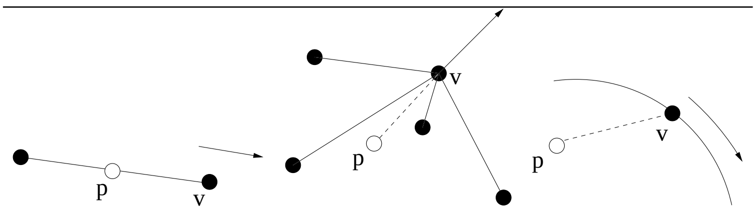

Fig. 13.4 Augmenting the edge costs when a node has one or more neighbors and when it is on the circumference of C_1

It can also happen that such a movement completely changes the structure of the MST, reducing the initial cost. In that case we do not validate the movement. Given an instance, the augmenting algorithm performs this computation for each node twisting over all the nodes but s, until no movements are allowed. Therefore, in order to give to a node a "second chance" to move, we can repeat such computations for a fixed number of rounds. Note that when a node reaches the border that is the circumference of C_1, the only allowed movement is over such a circumference.

Another way to increase the cost of the MST is to try to delete a node. The chosen candidate is the node with the highest degree. The idea behind this choice is that the highest degree node could be considered as the intermediary node to connect its neighbors, so removing it, a "big hole" will probably appear. On the one hand, this means that the distances to connect the remaining disjoint subtrees should increase the overall cost. On the other hand, we are creating more space for further movements. After a deletion, the algorithm starts again with the movements. Indeed, the deletion can be considered as a movement in which two nodes coincide. If the deletion does not increase the cost of the current MST, it is not validated. In such a case, the next step will be the deletion of the second highest degree node and so on. The whole procedure is repeated until no movements and no deletions are allowed. Note that eventually the algorithm can be repeated several times in order to obtain more accurate results. Sometimes, in fact, it can happen that the algorithm is stuck in some local maximum. Due to its randomness on the movements, the more it is executed, the higher is the probability to exit from such a situation.

The algorithm was evaluated over hundreds of instances from five up to 100 nodes. Table 13.1 shows the average and the maximum costs obtained on random instances using and not using the augmenting method (ε represents the maximum distance allowed for movements).

Table 13.1 The average and the maximum costs obtained on standard random instances and using the previous augmenting algorithm on instances of five up to 100 nodes and ε equal to 0.1 and 0.5.

n	Random		Augmented, $\varepsilon = .5$		Augmented, $\varepsilon = .1$	
	Average	Max	Average	Max	Average	Max
5	1.301	2.875	3.645	4.000	3.627	4.000
7	1.480	2.479	4.545	5.738	4.560	5.879
20	1.854	2.618	4.281	5.090	4.131	5.122
50	1.812	1.971	3.732	3.890	3.633	3.759
100	1.683	1.883	3.567	3.722	3.490	3.812

Compared to the standard random generated instances, the average costs were almost tripled while the maximum costs were almost doubled. The numerical results obtained are very interesting since they show that standard random instances are not really representative when studying the bounds of the MST heuristic for the MEBR problem. Moreover, as a "side effect" of such experiments, another very interesting property concerns the topologies obtained in the augmented instances. Whereas for instances up to around 20 nodes the method modifies the distribution of nodes, collapsing them to the hexagon shape of Figure 13.1, increasing the number of nodes makes things more interesting.

In Figure 13.5 an instance of 100 nodes is given before and after the movements and deletions. What follows from those experiments is an evident regularity on the final obtained instances. As shown in Figure 13.5, in general, after the augmentation, nodes look like they are being disposed of on some kind of regular grid, and this reflects the lower bound given by the regular hexagon shape.

13.2.5 Solving More General Instances of MEBR

In non-geometric versions of MEBR, the MST heuristic can be easily shown to have a poor approximation ratio. Better algorithms exist in both cases, where the weight

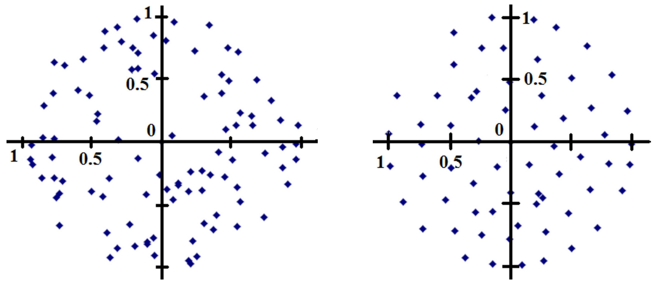

(a) before (b) after

Fig. 13.5 A random instance of 100 nodes before and after applying the augmenting method. The number of nodes decreased from 100 to 65, while the cost increased from 1.877 to 3.681

function is symmetric and nonsymmetric. For the former case, [13–15] present algorithms with a logarithmic (in the number of nodes) approximation ratio. The main idea is to reduce instances of the problem to instances of the problem Node-Weighted Connected Dominating Set [13, 14] or Node-Weighted Steiner Tree [15]. Such instances can be approximated within a logarithmic factor [28], which carries over as the approximation factor of the original instance. For nonsymmetric instances, [15] exploits a reduction due to Liang [36] of instances of MEBR to instances of Directed Steiner Tree. The obtained instances have some special properties, which allow for logarithmic approximations using techniques of Zosin and Khuller [45] for approximating special instances of Directed Steiner Tree. Similar results using different techniques have been obtained independently by Calinescu et al. [9]. All these results are the best possible; a matching inapproximability result has been presented in [19].

In the following, we discuss a recent algorithm from [12] for symmetric instances of MEBR, which has important implications for the geometric version of MEBR as well. The algorithm starts with the solution computed by the MST heuristic and gradually performs improvements on this solution according to a well-selected criterion. At the end, the solution obtained has significantly smaller cost than the initial one.

Before presenting the algorithm, we give some necessary definitions. Given a power assignment p and a station $x \in S$, let $E(p,x) = \{(x,y) | w(x,y) \leq p(x)\}$ be the set of the undirected edges induced by p at x, and $E(p) = \bigcup_{x \in S} E(p,x)$ the set of all the undirected edges induced by p. For every subset of undirected edges $F \subseteq E$ of a weighted graph $G = (S,E)$, we denote as $c(F)$ the overall cost of the edges in F, that is, the total sum of their weights. For the sake of simplicity, we will identify trees with their corresponding sets of edges. A *swap set* for a spanning tree T of an undirected graph $G(S,E)$ and a set of edges F with endpoints in S is any subset F' of edges that must be removed from the multigraph $T \cup F$ so that $T \cup F \setminus F'$ is a spanning tree of G.

We are now ready to describe the algorithm by first giving the basic underlying idea. Starting from a spanning tree T of $G(S)$, if the cost of T is significantly higher than the one of an optimal solution for performing broadcasting from a given source node $s \in S$, then there must exist a cost-efficient *contraction* of T. Namely, it must be possible to set the transmission power $p(x)$ of at least one station x in such a way that $p(x)$ is much lower than the cost of some swap set $A(p,x)$ for T and $E(p,x)$. The algorithm then repeatedly chooses at each step $p(x)$ in such a way that, starting from the current spanning tree, $c(A(p,x))/p(x)$ is maximized. The final tree will be such that, considering the correct orientation of the edges according to the final assignment p, some edges will be in the reverse direction, i.e., from the leaves towards the source s. However, the transmission powers can then be properly set with low additional cost in order to obtain the right orientation from s to the other stations.

At any intermediate step of the algorithm in which p and T are the current power assignment and maintained tree, respectively, consider a contraction at a given station x consisting of setting the transmission power of x to $p'(x)$, and let p' be the resulting power assignment. Then, a maximum cost swap set $A(p',x)$ to be attributed to the contraction can be trivially determined by letting $A(p',x)$ contain the edges that are removed when determining a minimum spanning tree in the multigraph $T \cup E(p',x)$ with the cost of all the edges in $E(p',x)$ set equal to 0. We call the ratio $\frac{c(A(p',x))}{p'(x)}$ the *cost-efficiency* of the contraction.

Formally, the algorithm performs the following steps:

- Set the transmission power $p(x)$ of every station in $x \in S$ equal to 0.
- Let $T = T(S)$ be a minimum spanning tree of $G(S)$.
- While there exists at least one contraction of cost-efficiency strictly greater than 2
  - Perform a contraction of maximum cost-efficiency, and let $p'(x)$ be the corresponding increased power at a given station x, and p' be the resulting power assignment.
  - Set the weight of all the edges in $E(p',x)$ equal to 0.
  - Let $T' = T \cup E(p',x) \setminus A(p',x)$.
  - Set $T = T'$ and $p = p'$.
- Orient all the edges of T from the source s toward all the other stations.
- Return the transmission power assignment p that induces such a set of oriented edges.

For any instance of the problem where the minimum spanning tree of the cost graph $G(S)$ is guaranteed to cost at most ρ times the cost of an optimal solution for MEBR, the algorithm achieves an approximation ratio bounded by ρ if $\rho \leq 2$ and by $2\ln\rho - 2\ln 2 + 2$ if $\rho > 2$, which exponentially improves upon the MST heuristic. Surprisingly, the algorithm and analysis do not make use of any geometric arguments, and still the results significantly improve the previously best-known approximation factor for Euclidean instances of the problem. The corresponding approximation ratio is reduced (when $\alpha \geq d$) from 6 [1] to 4.2 for $d = 2$, from 18.8 [38] to 6.49 for $d = 3$, and in general from $3^d - 1$ [24] to $2.2d + 0.61$ for $d > 3$. In the two-dimensional case, the achieved approximation is even less than the lower bound of $13/3$ on the approximation ratio of the BIP heuristic [41]. In arbitrary (i.e., non-Euclidean) cost graphs, it is not difficult to see that the cost of the minimum spanning tree is at most $n-1$ times the cost of an optimal solution for MEBR; hence, the algorithm achieves a logarithmic approximation for arbitrary symmetric weight functions matching the results in [9, 15].

13.3 Cost Minimization in Multi-interface Networks

Nowadays wireless devices hold multiple radio interfaces, allowing switching from one communication network to another according to required connectivity and related quality. The selection of the "best" radio interface for a specific connection might depend on various factors, namely, its availability in specific devices, the required communication bandwidth, the cost (in terms of energy consumption) of maintaining an active interface, the available neighbors, and so forth. While managing such connections, a lot of effort must be devoted to energy consumption issues. Devices are, in fact, usually battery-powered and network survivability might depend on their persistence in the network. This introduces a challenging and natural optimization problem that must take care of different variables at the same time. Generally speaking, given a set of k interfaces and a graph $G=(V,E)$, where V represents the set of wireless devices and E the set of required connections according to the proximity of the devices and the available interfaces that they may share, the problem can be stated as follows. What is the cheapest way, i.e., which subset of available interfaces in each node must be activated, to satisfy (cover) all the connections described by E? Note that a connection is satisfied when the endpoints of the corresponding edge share at least one active interface. Moreover, for each node $v \in V$ there is a set of available interfaces, from now on denoted as $W(v)$. $\bigcup_{v\in V} W(v)$ determines the set of all the possible interfaces available in the network whose cardinality is denoted by k. An example of a network instance is shown in Figure 13.6.

Depending on whether k is a priori bounded or not, two different problems arise. The first one is called Cost Minimization in Multi-interface Networks (k-CMI for short). The second one is called Cost Minimization in Unbounded Multi-interface Networks (CMI for short). In this section, we report results about the complexity of both k-CMI and CMI in various scenarios. The problems turn out to be very hard in general; hence, we also consider possible approximation algorithms. We deal with two main variations of the problem: the case in which the cost of activating an interface is the same for each interface (uniform case), and the more general case in which such a cost may differ (non-uniform case). Indeed, the first model is equivalent to asking for the minimum total number of activated interfaces inside the network to cover all the connections. We also consider different graph classes that are of interest from both theoretical and practical points of view, namely, graphs with bounded degree, since in real-world scenarios users are normally connected to a limited number of nodes; planar graphs, since the induced graph of joining users in a network is likely to be planar; trees, since middleware strategies are heavily based on this kind of structure (see, for instance, [10]); and complete graphs, since

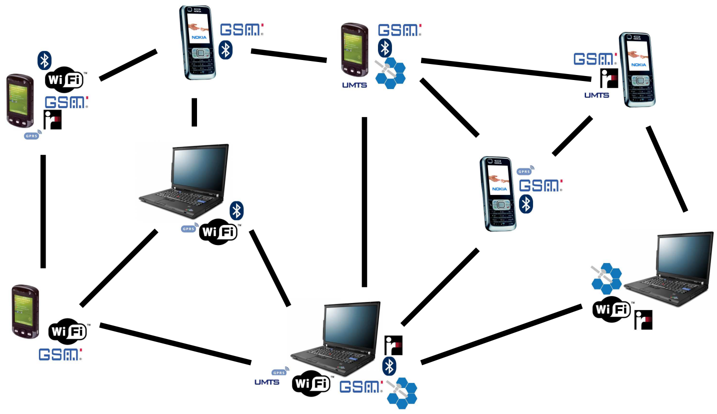

Fig. 13.6 The composed network according to available interfaces and proximities

this is one of the main structures used for modeling peer-to-peer networks (see for instance [18]).

Here we consider the bounded and the unbounded version of the problem. The two models reflect two different feasible cases, where available interfaces are either known a priori or not, respectively. Since nowadays devices support many and different interfaces, it makes sense either to assume the number of interfaces that may occur in a composed network as given, or to not. It might depend, in fact, on the number of nodes participating in the network. Regardless, k reflects the network dynamics.

The problems originated from [11], where a slightly different model of k-CMI is introduced. That model considers bandwidth constraints and also the possibility of having mutually exclusive interfaces, i.e., interfaces that, if activated, preclude the activation of some other interfaces. The motivation is quite technical. For instance, the WiFi interface can operate in different modalities: *Infrastructure* and *Ad Hoc*. If a device activates WiFi in the Infrastructure modality, it cannot satisfy connections that require Ad Hoc modality, and vice versa. This further constraint is not introduced here since the problem, although of practical interest, is not easily solvable. Other related problems were recently addressed in [23, 35] and [4], concerning connectivity and shortest path issues, respectively.

13.3.1 Definitions and Notation

Unless otherwise stated, the network graph $G = (V,E)$ is always assumed to be simple (i.e., without multiple edges), undirected, and connected. Moreover, we always denote by n and m the cardinality of the sets V and E respectively. The degree of node $v \in V$ is denoted by $deg(v)$ and the set of its neighbors by $N(v)$. The maximum node degree of graph G is denoted by $\Delta(G)$.

A global characterization of interfaces of respective nodes from V is given in terms of an appropriate interface assignment function W, according to the following definition.

Definition 13.1. A function $W\colon V \to 2^{\{1,\dots,k\}}$ is said to *cover* graph $G = (V,E)$ if for each $\{u,v\} \in E$ the set $W(u) \cap W(v) \neq \emptyset$.

The cost of activating an interface for a node is assumed to be identical for all nodes and given by cost function $c\colon \{1,\dots,k\} \to \mathbb{Z}_+$, i.e., the cost of interface i is written as c_i. The considered k-CMI optimization problem is formulated as follows.

---

k-CMI: Cost Minimization in Multi-interface Networks

Input: A graph $G = (V,E)$, a positive integer k, an allocation of available interfaces $W\colon V \to 2^{\{1,\dots,k\}}$ covering graph G, an interface cost function $c\colon \{1,\dots,k\} \to \mathbb{R}_+$.

Solution: An allocation of active interfaces $W_A\colon V \to 2^{\{1,\dots,k\}}$ covering graph G such that $W_A(v) \subseteq W(v)$ for all $v \in V$.

Goal: Minimize the total cost of the active interfaces, $c(W_A) = \sum_{v\in V} \sum_{i\in W_A(v)} c_i$.

The considered CMI optimization problem is formulated as follows.

CMI: Cost Minimization in Unbounded Multi-interface Networks

Input: A graph $G = (V,E)$, an allocation of available interfaces $W\colon V \to 2^{\{1,\dots,k\}}$ covering graph G, an interface cost function $c\colon \{1,\dots,k\} \to \mathbb{R}_+$.

Solution: An allocation of active interfaces $W_A\colon V \to 2^{\{1,\dots,k\}}$ covering graph G such that $W_A(v) \subseteq W(v)$ for all $v \in V$.

Goal: Minimize the total cost of the active interfaces, $c(W_A) = \sum_{v\in V} \sum_{i\in W_A(v)} c_i$.

13.3.2 Results for k-CMI

Table 13.2 summarizes known results for k-CMI [32]. The problem is polynomially solvable for $k = 2$ but it is already APX-hard when k grows. If the underlying graph is complete or a tree, then k-CMI is still polynomial while for planar graphs it is NP-hard but admits a PTAS.

Table 13.2 Hardness and approximability of the k-CMI problem

Graph class	Interfaces	Complexity of k-CMI	
		non-uniform costs	uniform costs
General graphs	$k = 2$	$\mathscr{O}(n^3)$	$\mathscr{O}(nm)$
	$k \geq 3$	$(k-1)$-approx, APX-hard	$\min\{\lceil \frac{k+1}{2} \rceil, \frac{2m}{n}\}$-approx, APX-hard
Graphs of bounded degree Δ	$k \geq 3$	Δ-approx, APX-hard for $\Delta \geq 5$	$\frac{\Delta+1}{2}$-approx, APX-hard for $\Delta \geq 5$
Planar graphs	$k \geq 3$	NP-hard, PTAS	NP-hard, PTAS
Trees	any k	$\mathscr{O}(n)$	$\mathscr{O}(n)$
Complete graphs	any k	$\mathscr{O}(n^2)$	$\mathscr{O}(n^2)$

The proof that provides the APX-hardness for $k \geq 3$ considers a polynomial transformation from the well-known VERTEX COVER problem on subcubic graphs[1] to k-CMI. On those instances VERTEX COVER is known to be APX-hard [39]. The transformation works as follows. Given a subcubic graph $G = (V, E)$, it is known that, in general, its chromatic number is at most 3 [6]. Nodes can then be partitioned into three subsets V_1, V_2, and V_3 according to an optimal coloring in such a way that $V_1 \bigcup V_2 \bigcup V_3 \equiv V$ and for each edge $e = \{x, y\} \in E$, x and y do not belong to the same subset V_i for every $i = 1$, 2, or 3.

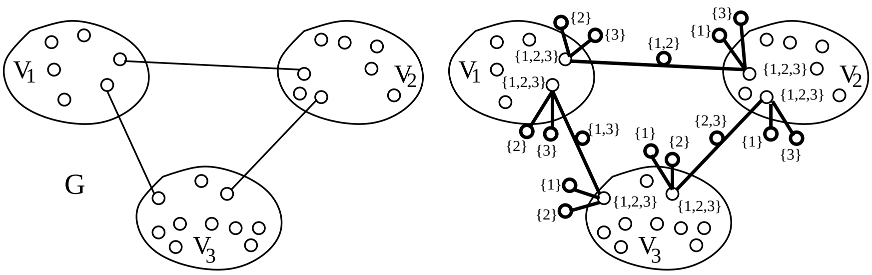

Fig. 13.7 On the left, the graph G subdivided into three node subsets according to a 3-coloring and the three possible kinds of edges. On the right are the modifications obtained for each kind of edge belonging to G and the interfaces associated with the related nodes

As illustrated in Figure 13.7, with each node $v \in V$, three interfaces, namely 1, 2, and 3 are associated. Moreover, to each $v \in V$ there are two new nodes connected. Those new nodes have only one interface: 2 and 3 (1 and 3 or 1 and 2 respectively) if $v \in V_1$ ($v \in V_2$ or $v \in V_3$). For each edge of G a further node is added. With such a node there are associated two interfaces. If the considered edge connects V_1 and V_2 (V_1 and V_3 or V_2 and V_3) then interfaces 1 and 2 (1 and 3 or 2 and 3 respectively) are associated with the added node. Considering for instance an edge $e = \{x, y\} \in E$ such that $x \in V_1$ and $y \in V_2$, in order to solve k-CMI on the new graph of maximum degree 5 built from G, a solution necessarily has to activate interfaces 2 and 3 in

[1] Graphs with maximum degree bounded by 3.

x, and 1 and 3 in y. In order for both x and y to be able to communicate with the new intermediate node, either such a node must activate both its interfaces, or one among x and y has to activate its third available interface. Both the solutions are locally equivalent. On the other hand, activating the third interface for either x or y may lead to a decrease in the number of activated interfaces in the global solution. This is implied by the fact that the neighborhood of the added intermediate node between x and y is constituted by only x and y, while both x and y may have many other connections. This implies that one can look for solutions where for each edge of the original graph at least one endpoint has all its three interfaces activated. Note that this reflects exactly the requirement of VERTEX COVER.

k-CMI can be approximated within a factor of $k-1$. A greedy algorithm activates interfaces among the nodes. It starts from the cheapest interface 1, and it activates it in each node that has a neighbor holding that interface. Let $V_1 \subseteq V$ be the set of nodes in which the algorithm activated interface 1 and let $E(V_1)$ be the corresponding set of covered edges. Note that the optimal solution restricted to $E(V_1)$ (i.e., the set of activated interfaces of an optimal solution at the endpoints of the edges belonging to $E(V_1)$) clearly costs at least as much as the cost of the algorithm. In the second step, the same is done for the next cheapest interface 2 among the remaining connections $E \setminus E(V_1)$. Again, the cost of the optimal solution restricted to $E(V_2)$ is at least the price paid by the algorithm. This is implied by the fact that any connection belonging to $E(V_2)$ cannot be covered by interface 1; otherwise, the algorithm would have covered it in the previous step. This process is continued for all the interfaces in a non-decreasing cost order, but for the last two interfaces. Referring to Table 13.2, when $k = 2$, k-CMI is polynomially solvable. Hence, when the two most expensive interfaces remain, the optimal algorithm for $k = 2$ can be applied. Since each step costs at most as much as the optimal solution, the $(k-1)$-approximation holds by observing that the whole process requires $k-1$ steps.

Concerning the uniform cost case, an easy approximation algorithm for solving k-CMI leads to a factor of $\frac{2m}{n}$. The algorithm simply chooses one interface for each edge of the input graph in order to satisfy the required connection. This means that for each edge at most one interface in each endpoint is activated. It follows that for m edges it activates at most $2m$ interfaces for n nodes. The $\lceil \frac{k+1}{2} \rceil$-approximation mentioned in Table 13.2 is instead obtained by suitably applying a hitting set algorithm.

13.3.3 Results for CMI

Table 13.3 summarizes results obtained for CMI [34]. When k depends on the instance, i.e., it is not set a priori, the problem becomes harder even for complete graphs and trees. In general, CMI is hard to approximate within a factor of $\mathcal{O}(\log k)$, even when restricted to the unit cost interface case. The proof proceeds by reduction to the MINIMUM HITTING SET problem. We recall that for a collection of non-empty subsets $C_1, C_2, \ldots, C_l \subseteq \{1, 2, \ldots, k\}$, set $S \subseteq \{1, 2, \ldots, k\}$ is called a *hitting* *set* if for all $i \in [1,\dots,\ell]$, $C_i \cap S \neq \emptyset$. The problem of minimizing the cardinality of the hitting set is as hard as the MINIMUM SET COVER problem [3], and consequently, hard to approximate within a factor of $\mathcal{O}(\log k)$ [40].

Concerning the $\sqrt{n}(1+\ln n)$-approximation factor, this is obtained by means of a polynomial transformation of the problem to the well-known WEIGHTED MINIMUM SET COVER problem. This leads to the claim that the existence of any a-approximation algorithm for WEIGHTED MINIMUM SET COVER leads to an $(a\sqrt{n})$-approximation algorithm for CMI. Since WEIGHTED MINIMUM SET COVER admits a $(1+\ln n)$-approximation [17], $\sqrt{n}(1+\ln n)$ is obtained.

Table 13.3 Hardness and approximability of the CMI problem. Entries marked by (*) follow from k-CMI results

Graph class	Complexity of CMI	
	non-uniform costs	uniform costs
General graphs	$(k-1)$-approx (*) $(\sqrt{n}(1+\ln n))$-approx not approx within $\mathcal{O}(\log k)$	$\lceil \frac{k+1}{2} \rceil$-approx (*) $\frac{2m}{n}$-approx (*) not approx within $\mathcal{O}(\log k)$
Graphs of bounded degree Δ	Δ-approx (*) APX-hard for $\Delta \geq 5$, $k \geq 3$ (*)	$\frac{\Delta+1}{2}$-approx (*) APX-hard for $\Delta \geq 5$, $k \geq 3$ (*)
Planar graphs	6-approx APX-hard	6-approx APX-hard
Trees	2-approx APX-hard	2-approx APX-hard
Complete graphs	not approx within $\mathcal{O}(\log k)$	not approx within $\mathcal{O}(\log k)$

13.4 Conclusion and Future Work

The chapter surveys recent results obtained for two interesting problems arising in the field of wireless ad hoc networks. Both the problems deal with the minimization of the overall energy needed to perform desired communication protocols. In particular, the Minimum Energy Broadcast Routing problem expresses the necessity to perform the basic broadcast pattern of communication from a given source, and the network is composed of homogeneous nodes equipped with omnidirectional radio antennas. The Cost Minimization in Multi-interface Networks expresses the need of establishing connections among heterogeneous nodes equipped with different subsets of interfaces, each associated with some activation cost. Many interesting directions for future work arise from both problems. These include the extensions of the studies to different communication protocols, to different objective functions, and to distributed environments.

Acknowledgements The research was partially funded by the project "CEPAGE" of INRIA (France), by the ANR-project "ALADDIN" (France), by the European Commission under Integrated Project "Algorithmic Principles for Building Efficient Overlay Computers" (AEOLUS), EU COST action 293 – Graphs and Algorithms in Communication Networks (GRAAL) – and EU COST action 295 – Dynamic Communication Networks (DYNAMO), by the Greek General Secretariat for Research and Technology under programme PENED, and by a "Caratheodory" research grant from the University of Patras.

References

1. Ambühl, C.: An optimal bound for the MST algorithm to compute energy efficient broadcast trees in wireless networks. In: Proceedings of the 32nd International Colloquium on Automata, Languages and Programming (ICALP), LNCS 3580, Springer, pp. 1139–1150, 2005
2. Athanassopoulos, S., Caragiannis, I., Kaklamanis, C., Kanellopoulos, P.: Experimental comparison of algorithms for energy-efficient multicasting in ad hoc networks. In: Proceedings of the 3rd International Conference on Ad Hoc Networks & Wireless (ADHOC NOW), LNCS 3158, Springer, pp. 183-196, 2004
3. Ausiello, G., D'Atri, A., Protasi, M.: Structure preserving reductions among convex optimization problems. Journal of Computer and System Sciences, 21(1):136–153, 1980
4. Barsi, F., Navarra, A., Pinotti, M. C.: Cheapest Path in Multi-Interface Networks. In: Proceedings of the 10th International Conference on Distributed Computing and Networking (ICDCN), LNCS, Springer, 2009
5. Biló, V., Flammini, M., Melideo, G., Moscardelli, L., Navarra, A.: Sharing the cost of multicast transmissions in euclidean and general wireless networks. Theoretical Computer Science, 369(1-3):269–284, 2006
6. Brooks, R. L.: On coloring the nodes of a network. In: Proceedings of Cambridge Philosophical Society, 37:194–197, 1941
7. Čagalj, M., Hubaux, J., Enz, C.: Minimum-energy broadcast in all-wireless networks: NP-completeness and distribution issues. In: Proceedings of the 8th Annual International Conference on Mobile Computing and Networking (MobiCom), ACM Press, pp. 172–182, 2002
8. Cai, H., Zhao, Y.: On approximation ratios of minimum-energy multicast routing in wireless networks. Journal of Combinatorial Optimization, 9(3):243–262, 2005
9. Calinescu, G., Kapoor, S., Olshevsky, A., Zelikovsky, A.: Network lifetime and power assignment in adhoc wireless networks. In: Proceedings of the 11th Annual European Symposium on Algorithms (ESA), LNCS 2832, Springer, pp. 114–126, 2003
10. Caporuscio, M., Carzaniga, A., Wolf, A. L.: Design and evaluation of a support service for mobile, wireless publish/subscribe applications. IEEE Transactions on Software Engineering, 29(12):1059–1071, 2003
11. Caporuscio, M., Charlet, D., Issarny, V., Navarra, A.: Energetic Performance of Service-oriented Multi-radio Networks: Issues and Perspectives. In: Proceedings of the 6th International Workshop on Software and Performance (WOSP), pp. 42–45. ACM Press, 2007
12. Caragiannis, I., Flammini, M., Moscardelli, L.: An exponential improvement on the mst heuristic for minimum energy broadcasting in ad hoc wireless networks. In: Proceedings of the 34th International Colloquium on Automata, Languages and Programming (ICALP), LNCS 4596, Springer, pp. 447–458, 2007
13. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P.: New results for energy-efficient broadcasting in wireless networks. In: Proceedings of the 13th International Symposium on Algorithms and Computation (ISAAC), LNCS 2518, Springer, pp. 332–343, 2002
14. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P.: A logarithmic approximation algorithm for the minimum energy consumption broadcast subgraph problem. Information Processing Letters, 86(3):149-154, 2003
15. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P.: Energy-efficient wireless network design. Theory of Computing Systems, 39(5), pp. 593–617, 2006

16. Cartigny, J., Simplot, D., Stojmenovic, I.: Localized minimum-energy broadcasting in ad hoc networks. In: Proceedings of the 22nd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM), Vol. 3, IEEE Computer Society Press, pp. 2210–2217, 2003
17. Chvátal, V.: A greedy heuristic for the set-covering problem. Mathematics of Operations Research, 4(3):233–235, 1979
18. Cilibrasi, R., Lotker, Z., Navarra, A., Perennes, S., Vitanyi, P.: About the lifespan of peer to peer networks. In: Proceedings of the 10th International Conference On Principles Of Distributed Systems (OPODIS), LNCS 4305, Springer, pp. 290–304, 2006
19. Clementi, A. E. F., Crescenzi, P., Penna, P., Rossi, G., Vocca, P.: On the complexity of computing minimum energy consumption broadcast subgraph. In: Proceedings of the 18th Annual Symposium on Theoretical Aspects of Computer Science (STACS), LNCS 2010, Springer, pp. 121–131, 2001
20. Clementi, A. E. F., Di Ianni, M., Silvestri, R.: The minimum broadcast range assignment problem on linear multi-hop wireless networks. Theoretical Computer Science, 299(1-3):751–761, 2003
21. Conway, J. H., Sloane, N. J. A.: "The Kissing Number Problem" and "Bounds on Kissing Numbers". Ch. 2.1 and Ch. 13 in: Sphere Packings, Lattices, and Groups. Springer-Verlag, New York, 3rd edition, 1998
22. Das, A. K., Markas, R. J., El-Sharkawai, M., Arabshahi, P., Gray, A.: Minimum energy broadcast trees for wireless networks: Integer programming formulations. In: Proceedings of the 22nd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM), IEEE Computer Society, Vol. 2, pp. 1001–1010. 2003
23. Faragó, A., Basagni, S.: The Effect of Multi-Radio Nodes on Network Connectivity—A Graph Theoretic Analysis. In: Proceedings of the 19th International IEEE Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), 2008
24. Flammini, M., Klasing, R., Navarra, A., Perennes, S.: Improved approximation results for the Minimum Energy Broadcasting Problem. In: Proceedings of ACM Joint Workshop on Foundations of Mobile Computing (DIALM-POMC), pp. 85–91, 2004
25. Flammini, M., Klasing, R., Navarra, A., Perennes, S.: Tightening the upper bound for the Minimum Energy Broadcasting. Wireless Networks, Vol. 14(5), pp. 959-669, 2008
26. Flammini, M., Navarra, A., Perennes, S.: The "real" approximation factor of the MST heuristic for the minimum energy broadcasting. ACM Journal of Experimental Algorithmics, 11, 2006. Preliminary version in: Proceedings of the 4th International Workshop on Efficient and Experimental Algorithms (WEA), LNCS 3503, Springer, pp. 22–31, 2005
27. Frieze, A. M., McDiarmid, C. J. H.: On Random Minimum Length Spanning Trees. Combinatorica, 9:363–374, 1989
28. Guha, S., Khuller, S.: Improved Methods for Approximating Node Weighted Steiner Trees and Connected Dominating Sets. Information and Computation, 150(1), pp. 57–74, 1999
29. Hac, A.: Wireless sensor network designs. John Wiley & Sons, Ltd, 2003
30. Kang, I., Poovendran, R.: Iterated local optimization for minimum energy broadcast. In: Proceedings of the 3rd International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt), pp. 332–341, 2005
31. Klasing, R., Flammini, M., Navarra, A., Perennes, S.: Improved approximation results for the Minimum Energy Broadcasting Problem. Algorithmica, 49(4):318–336, 2007
32. Klasing, R., Kosowski, A., Navarra, A.: Cost minimisation in multi-interface networks. In: Proceedings of the 1st EuroFGI International Conference on Network Control and Optimization (NET-COOP), LNCS 4465, Springer, pp. 276–285, 2007
33. Klasing, R., Navarra, A., Papadopoulos, A., Perennes, S.: Adaptive Broadcast Consumption (ABC), a new heuristic and new bounds for the minimum energy broadcast routing problem. In: Proceedings of the 3rd IFIP-TC6 International Networking Conference, LNCS 3042, Springer, pp. 866–877, 2004
34. Kosowski, A., Navarra, A.: Cost minimisation in unbounded multi-interface networks. In: Proceedings of the 2nd PPAM Workshop on Scheduling for Parallel Computing (SPC), Lecture Notes in Computer Science 4967, Springer-Verlag, pp. 1039-1047, 2007

35. Kosowski, A., Navarra, A., Pinotti, M. C.: Connectivity in Multi-Interface Networks. In: Proceedings of the 4th International Symposium on Trustworthy Global Computing (TGC), LNCS, Springer, 2008
36. Liang, W.: Constructing minimum-energy broadcast trees in wireless ad hoc networks. In: Proceedings of the 3rd ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc), ACM Press, pp. 112–122, 2002
37. Li, F., Nikolaidis, I.: On minimum-energy broadcasting in all-wireless networks. In: Proceedings of the 26th Annual IEEE Conference on Local Computer Networks (LCN), IEEE Computer Society, p. 193, 2001
38. Navarra, A.: 3-D Minimum Energy Broadcasting problem. Ad Hoc Networks, Vol. 6(5), pp. 734-743, 2008
39. Papadimitriou, C. H., Yannakakis, M.: Optimization, approximation, and complexity classes. Journal of Computer and System Sciences, 43:425–440, 1991
40. Raz, R., Safra, S.: A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In: Proceedings of the 29th Anuual ACM Symposium on Theory of Computing (STOC), pp. 475–484, 1997
41. Wan, P. J., Calinescu, G., Li, X., Frieder, O.: Minimum energy broadcasting in static ad hoc wireless networks. Wireless Networks, 8(6):607–617, 2002
42. Wieselthier, J. E., Nguyen, G. D., Ephremides, A.: On the construction of energy-efficient broadcast and multicast trees in wireless networks. In: Proceedings of the 19th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM), IEEE Computer Society, pp. 585–594, 2000
43. Yuan, D.: Computing Optimal or Near-Optimal Trees for Minimum-Energy Broadcasting in Wireless Networks. In: Proceedings of the 3rd International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt), pp. 323–331, 2005
44. Zhao, F., Guibas, L.: Wireless sensor networks: an information processing approach. Morgan Kaufmann, 2004
45. Zosin, L., Khuller, S.: On Directed Steiner Trees. In: Proceedings of the 13th Annual ACM/SIAM Symposium on Discrete Algorithms (SODA), pp. 59-63, 2002