An exponential improvement on the MST heuristic for the Minimum Energy Broadcasting problem

I. Caragiannis[1] M. Flammini[2] L. Moscardelli[2]

[1]Research Academic Computer Technology Institute and
Dept. of Computer Engineering and Informatics
University of Patras, Greece

[2]Dipartimento di Informatica,
Università di L'Aquila, Italy

Nice, March 9^{th}, 2007 / AEOLUS workshop on scheduling

Outline

General Characteristics

- A wireless network is a collection of transmitter/receiver stations.
- All the communication is carried over the wireless medium.
- All stations have omni-directional antennas.
- A communication is established by assigning to each station a transmitting power.
  - Power is expended for signal transmission only. No power expenditure for signal reception or processing.
- Multi-hop communication is allowed.

The model

- We are interested in the broadcast communication from a given source node s.
- Given a set of stations S, let $G(S)$ be the complete weighted graph in which the weight $w(x, y)$ of each edge between stations x and y is the power consumption needed for a communication between x and y.
- A power assignment for S is a function $p : S \rightarrow \boldsymbol{R}^+$ assigning a transmission power $p(x)$ to every station in S.
- The total cost of a power assignment is

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

The goal

- The *Minimum Energy Broadcast Routing* (MEBR) problem for a given source $s \in S$ consists in finding a power assignment p of minimum cost such that every station is able to receive the communication from s.
- A particular relevant case is when stations lie in a $d-$dimensional *Euclidean space*. Then, given an integer $\alpha \geq 1$ and a constant $\beta \in \boldsymbol{R}^+$, the power consumption needed for a correct communication between x and y is $w(x, y) = \beta \cdot dist(x, y)^\alpha$.

Previous Work (1)

- In the general case in which the weights of $G(S)$ are completely arbitrary, the problem cannot be approximated within $(1-\epsilon)\ln n$ unless $NP \subseteq DTIME(n^{O(\log\log n)})$[1], where n is the number of stations, while logarithmic (in the number of stations) approximation algorithm has been provided[2].
- When distances are induced by the positions of the stations in a d-dimensional space, for $\alpha > 1$ and $d > 1$ the MEBR problem is NP-hard, while if $\alpha = 1$ or $d = 1$ it is solvable in polynomial time[3].

[1] Clementi, Crescenzi, Penna, Rossi and Vocca, STACS 2001

[2] Calinescu, Kapoor, Olshevsky and Zelikovsky, ESA 2003; Caragiannis, Kaklamanis and Kanellopoulos, ISAAC 2002

[3] Caragiannis, Kaklamanis and Kanellopoulos, ISAAC 2002; Zagalj, Hubaux and Enz, MobiCom 2002

Previous Work (2) - Euclidean case

- The best (previously) known approximation algorithm is the *MST* heuristic[4].
  - It is based on the idea of tuning ranges so as to include a minimum spanning tree of the cost graph $G(S)$.
  - After the first approximation analysis[5], the best shown approximation ratios are 6 for $d = 2$[6], 18.8 for $d = 3$[7] and $3^d - 1$ for every $d > 3$[8].
  - A lower bound on the approximation ratio is given by the d-dimensional kissing numbers n_d (i.e. 6 for $d = 2$ and 12 for $d = 3$)[9].

[4] Wieselthier, Nguyen and Ephremides, INFOCOM 2000

[5] Clementi, Crescenzi, Penna, Rossi and Vocca, STACS 2001

[6] Ambuhl, ICALP 2005

[7] Navarra, SIROCCO 2006

[8] Flammini, Klasing, Navarra and Perennes, DIALM-POMC 2004

[9] Wan, Calinescu, Li and Frieder, Wireless Networks 2002

Our Contribution (1)

- We present a new approximation algorithm for the MEBR problem.
- For any distance metric inducing a weighting of $G(S)$ such that its minimum spanning tree is guaranteed to cost at most ρ times the cost of an optimal solution for MEBR, our algorithm achieves an approximation ratio bounded by $2\ln\rho - 2\ln 2 + 2$.
- We provide a matching lower bound, proving that the analysis is tight.

Our Contribution (2) - Euclidean case

- In the 2-dimensional case, the achieved approximation is even less than the 4.33 lower bound on the ratio of the *BIP* heuristic, the only one shown to be no worse than MST[10].

Dimensions	1	2	3	...	7	...	d
MST	2	6	18.8	...	2186	...	$3^d - 1$
Our alg.	2	4.2	6.49	...	16	...	$2.20d + 0.62$

Figure: Comparison between the approximation factors of our algorithm and the MST heuristic in Euclidean instances.

[10] Wan, Calinescu, Li and Frieder, Wireless Networks 2002

The basic idea (1)

- Starting from a spanning tree T_0 of $G(S)$, if the cost of T_0 is significantly higher than the one of an optimal solution, then there must exist a cost efficient *contraction* of T_0.
- In other words, 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 a subset of edges that can be deleted from T_0 maintaining the connectivity and eliminating cycles.
  - Let $E(p', x)$ be the set of edges induced by $p(x)$.
  - Let $A(p', x)$ be a swap set, i.e. a set of edges that can be removed maintaining the connectivity and eliminating cycles.

The basic idea (2)

- At each step, starting from the initial MST T_0, perform a maximum cost-efficiency contraction:
  - Consider a contraction at a station x consisting in setting the transmission power of x to $p'(x)$, and let p' be the resulting power assignment.
  - Then a maximal cost *swap set* $A(p', x)$ can be easily determined by considering the edges that are removed when computing a minimum spanning tree in the multigraph $T \cup E(p', x)$ with the cost of all the edges in $E(p', x)$ set to 0.
  - The ratio $\frac{c(A(p',x))}{p'(x)}$ is the cost-efficiency of the contraction.

The algorithm

- Set the transmission power $p(x)$ of every station in $x \in S$ equal to 0; set i equal to 0.
- Let T_0 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 to 0 the weight of all the edges in $E(p', x)$
    - Let $i = i + 1$ and $p = p'$
    - Let $T_i = T_{i-1} \cup E(p', x) \setminus A(p', x)$
- Orient all the edges of T_i from the source s toward all the other stations.
- Return the transmission power assignment p that induces such a set of oriented edges.

A difficult task

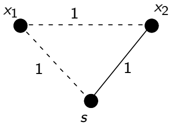

Figure: A simple network with a minimum spanning tree depicted by dashed lines.

- Consider the network in figure
  - Swap sets $A(p', x)$ are not static sets.
- Thus, we cannot statically associate edges of the initial spanning tree to the range assignments of the optimum.
- We have to ensure that at each step i of the algorithm, if the current tree T_i has a cost much grater than the optimum, a good contraction exists.

Our solution

- Consider a directed spanning tree T^* induced by the power assignment of an optimal solution.
- Given a spanning tree T with an arbitrary weighting of the edges, we want to find a one-to-one function mapping each edge of T^* to an edge of T.
- Such a mapping has to ensure that, for each station, if we consider all its outgoing edges in T^*, their images form a swap set for T with respect to such outgoing edges.

The recursive mapping (1)

- Given two trees T_A and T_B spanning the n stations, we want to find a one-to-one swap mapping $f : T_A \rightarrow T_B$ such that $f(\{u,v\})$ is an edge of T_B forming a cycle with $\{u,v\}$ in $T_B \cup \{\{u,v\}\}$.
- We consider a recursive mapping construction, in which the recursive step works on the unique edge incident to a leaf of T_A.
  - If $|V| = 2$, the base of the induction is trivially verified.
  - Now we assume that such a mapping f' exists for any T'_A and T'_B spanning $n-1$ stations, and we prove that f exists.

The recursive mapping (2)

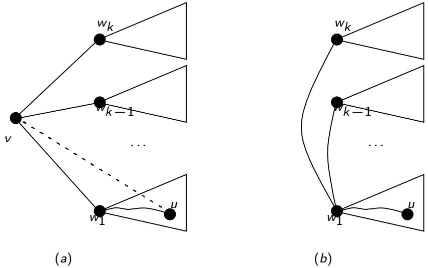

Figure: Edge $\{u, v\}$ of T_A is associated to $\{v, w_1\}$ of T_B.

- every edge $\{y, z\}$ of T_A forms a cycle with $\{w_1, w_i\}$ in T'_B, if and only if it forms also a cycle with $\{v, w_i\}$ in T_B.
- The swap mapping f' for T'_A and T'_B is a swap mapping for T_A and T_B.

Two technical lemmata (1)

Lemma 1

Given two rooted spanning trees T_A and T_B over the same set of nodes V, there exists a one-to-one mapping $f : T_A \rightarrow T_B$, called the swap mapping, such that, if $v_1, \ldots, v_k$ are all the children of a same parent node u in T_A, then the set $\{f(\{u, v_1\}), \ldots, f(\{u, v_k\})\}$ of the edges assigned to $\{u, v_1\}, \ldots, \{u, v_k\}$ by f is a swap set for T_B and $\{\{u, v_1\}, \ldots, \{u, v_k\}\}$.

- By applying the previous recursive construction in an appropriate way, i.e. considering a proper ordering of the edges, we can obtain the claimed mapping.

Two technical lemmata (2)

Lemma 2

Given any tree T, and k edges $\{u, v_1\}, \ldots, \{u, v_k\}$ not belonging to T, if $\{\{w_1, y_1\}, \ldots, \{w_k, y_k\}\}$ is a swap set for T and $\{\{u, v_1\}, \ldots, \{u, v_k\}\}$, then $\{\{w_1, y_1\}, \ldots, \{w_k, y_k\}\}$ is the subset of a swap set for T and $\{\{u, v_1\}, \ldots, \{u, v_k\}, \{u, z_1\}, \ldots, \{u, z_l\}\}$, for every set of l newly added edges $\{\{u, z_1\}, \ldots, \{u, z_l\}\}$.

- This lemma ensures that the edges in a swap set relative to a set X of edges are in a swap set relative to any set $Y \supseteq X$ of edges.

The key lemma

Lemma 3

Let T be any spanning tree for $G(S)$ with an arbitrary weighting of the edges, and let $\gamma = c(T)/m^*$ be the ratio among the cost of T and the one of an optimal transmission power assignment p^*. Then there exists a contraction of T of cost-efficiency γ.

Sketch of proof of Lemma 3

- Consider the swap mapping $f : T^* \to T$.
- By Lemma 1, f assigns to all the descending edges $D(x)$ in T^* of every station x a subset of edges $SS(x)$ forming a swap set.
- All such subsets $SS(x)$ form a partition of T, and since $\frac{c(T)}{m^*} = \frac{\sum_{x \in S} c(SS(x))}{\sum_{x \in S} p^*(x)} = \gamma$, there must exist at least one station $\bar{x}$ such that $\frac{c(SS(\bar{x}))}{p^*(\bar{x})} \geq \gamma$.
- Since $D(\bar{x}) \subseteq E(p^*, \bar{x})$, by Lemma 2, $SS(\bar{x}) \subseteq A(p^*, \bar{x})$. Therefore, there exists a contraction of T of cost-efficiency $c(A(p^*, \bar{x}))/p^*(\bar{x}) \geq c(SS(\bar{x}))/p^*(\bar{x}) = \gamma$.

$\square$

The approximation ratio of our algorithm

Theorem 4

The algorithm has approximation ratio $2\ln\rho - 2\ln 2 + 2$, where ρ is the approximation guarantee of a minimum spanning tree over $G(S)$.

Sketch of proof (1)

- Let T_0 be the minimum spanning tree for $G(S)$ computed at the beginning of the algorithm, $T_1, \ldots, T_k$ be the sequence of the trees constructed by the algorithm after the contraction steps, and $\gamma_i = c(T_i)/m^*$.
- By Lemma 3, since the algorithm always considers contractions of maximum cost-efficiency, at each step $i = 0, 1, \ldots, k-1$ it performs a contraction at node x_i having cost-efficiency at least γ_i (by assigning x_i a power p_i) and removing from the initial spanning tree edges with total cost $t_i = c(T_i) - c(T_{i+1})$.

$$\gamma_i = \frac{t_i}{p_i}.$$

Sketch of proof (2)

- In order to orient the edges of the final solution from s towards the other stations, we have at most to double the power assignment due to the contraction steps.
- The overall cost is upper bounded by

$$2\sum_{i=0}^{k-1} p_i + c(T_k) = 2\sum_{i=0}^{k-1} \frac{t_i}{\gamma_i} + c(T_k).$$

- Since $\forall i, \gamma_i > 2$, we are paying each edge of the initial spanning tree at most for its whole cost.

Sketch of proof (3)

- Recall that $\gamma_i = \frac{c(T_i)}{m^*}$.
- Since we perform contractions as long as the cost-effectiveness is grater than 2, by Lemma 3 the cost $c(T_k)$ of the final tree can be at most $2m^*$.
- Therefore, the total cost can be upper bounded as follows:

$$2\sum_{i=0}^{k-1}\frac{t_i}{\gamma_i} + c(T_k) \leq 2m^* \left(\sum_{i=0}^{k-1}\frac{t_i}{c(T_i)} + 1\right).$$

Sketch of proof (4)

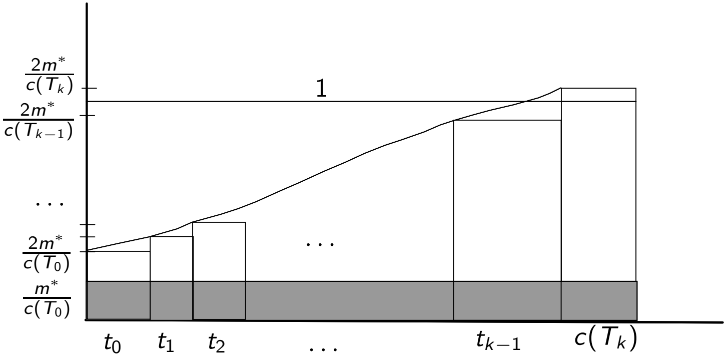

$$2m^* \left(\sum_{i=0}^{k-1} \frac{t_i}{c(T_i)} + 1 \right) = \sum_{i=0}^{k-1} \left(\frac{2m^*}{c(T_i)} t_i \right) + \frac{2m^*}{c(T_k)} c(T_k)$$

Sketch of proof (5)

- Let us observe that

$$2m^*\left(\sum_{i=0}^{k-1}\frac{t_i}{c(T_i)}+1\right)\leq$$

$$\leq 2m^*\left(\int_0^{c(T_0)\left(1-\frac{2}{\gamma_0}\right)}\frac{\mathrm{d}t}{c(T_0)-t}+1\right)=m^*(2\ln\gamma_0-2\ln 2+2).$$

- Since T_0 is a ρ-approximation of an optimal solution, we have $\gamma_0\leq\rho$ and the claim follows.

□

The building block Q_x.

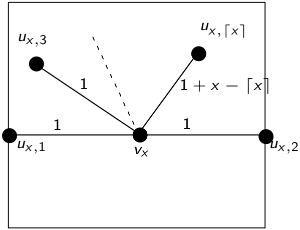

Figure: The building block Q_x.

- Notice that in Q_x there exists a contraction centered at node v_x having cost-effectiveness equal to x.

The complete construction

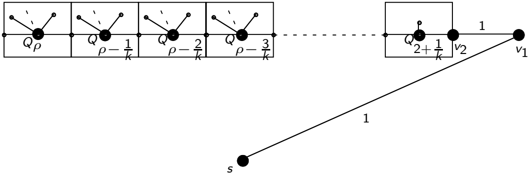

Figure: A minimum spanning tree of the lower bound instance.

- Let k be an integer parameter; the node set of the instance is obtained by sequencing $k(\rho - 2)$ building blocks $(Q_{2+\frac{1}{k}}, Q_{2+\frac{2}{k}}, \ldots, Q_3, \ldots, Q_\rho)$
- Moreover, in the instance there are other 3 nodes: the source s, and nodes v_1 and v_2, that coincides with $u_{2+\frac{1}{k},2}$.
- The weights of the edges connecting s to all the other nodes are equal to 1; moreover, $w(v_1, v_2) = 1$.

The complete construction

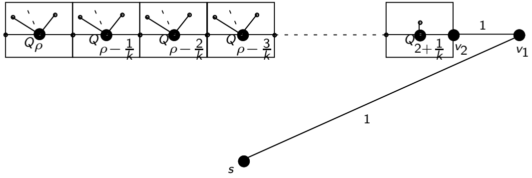

Figure: A minimum spanning tree of the lower bound instance.

- The weights of the edges contained in building block Q_x are diveded by kx, so that the sum of all the edges of each building block is equal to $\frac{1}{k}$.
- For all the other pairs of nodes, we assume that the mutual power communication cost is very high.
- Assume that the initial minimum spanning tree considered by the algorithm is the one depicted in figure, whose cost is ρ.

The lower bound (1)

- At each step, the algorithm can arbitrarily choose among two equivalent contractions, i.e. having the same cost-efficiency;
- For instance, at step 0, the first choice is the contraction centered at the source and having transmission power equal to 1, and the second choice is the contraction centered at v_ρ and having transmission power equal to $\frac{1}{\rho}$.
  - Both contractions have a cost-efficiency equal to ρ, and we assume that the algorithm chooses the contraction centered at v_ρ.
- In this way, the algorithm performs $k(\rho - 2)$ steps of contractions.
- At this point, no contraction having cost-efficiency at least 2 exists any longer.

The lower bound (2)

- Notice that the sum of the costs of the transmission powers set in the contractions is $\sum_{i=2k+1}^{k\rho} \frac{1}{i} = H_{k\rho} - H_{2k}$.
- In order to orient the edges of the final tree from the source towards the other nodes, we have to globally double the cost of the transmission powers set in the contraction steps.
- Thus, the final cost of the solution returned by the algorithm has cost $2H_{k\rho} - 2H_{2k} + 2$, while the optimal solution has cost 1.
- Letting k go to infinity, the approximation ratio tends to $2\ln\rho - 2\ln 2 + 2$.

Summary

- We have presented an approximation algorithm exponentially outperforming the MST heuristic for several specific metrics.
- Such results are particularly relevant for their consequences on Euclidean instances, for which the achieved approximation ratio has become linear in the number of dimension d instead of exponential.

Dimensions	1	2	3	...	7	...	d
MST	2	6	18.8	...	2186	...	$3^d - 1$
Our alg.	2	4.2	6.49	...	16	...	$2.20d + 0.62$

Figure: Comparison between the approximation factors of our algorithm and the MST heuristic in Euclidean instances.

Open questions

- Our analysis works for general metrics, but there might be possible improvements for specific cases, like the Euclidean, for which it would be worth to determine exact results tightening the current gap between the lower and upper bounds on the approximation ratio.
- Another interesting issue is that of determining similar contraction strategies possibly leading to better approximated solutions.