Convex Combination Approximation for the Min-Cost WSN Point Coverage Problem

Zheng Fang and Jie Wang

Department of Computer Science
University of Massachusetts, Lowell, MA 01854
{zfang,wang}@cs.uml.edu
http://www.cs.uml.edu

Abstract. This paper presents a new algorithm for finding better approximation solutions to the min-cost point coverage problem in wireless sensor networks. The problem is to compute a deterministic sensor deployment plan, with minimum monetary cost on sensors, to cover the set of targets spread across a geographical region such that each target is covered by multiple sensors. This is a Max-SNP-complete problem. Our approximation algorithm, called alpha-beta approximation, is a convex combination of greedy LP-rounding and greedy set-cover selection. We show that, through a large number of numerical simulations on randomly generated targets and sites, alpha-beta approximation produces efficiently better approximation results than the best approximation algorithm previously known. In particular, the alpha-beta approximation in our experiments never exceeds an approximation ratio of 1.07, providing up to 14.86% improvement over previous approximation algorithms.

Keywords: sensor deployment, point coverage, minimum set multicover, LP-rounding, approximation algorithm.

1 Introduction

The min-cost point coverage (MCPC) problem is a classic sensor coverage problem in wireless sensor networks. This problem has been studied intensively in recent years (see, e.g., [Vaz01, CIQ$^+$02, SX05, CW04, WZ06, YW07, WZ08]). Given a set of targets in a 3D (or 2D) geographical region, a set of sensor sites in the proximity of targets, and multiple types of sensors with different monetary costs, the problem is to select a set of sensors with minimum monetary cost on sensors, a set of sites, and a mapping of the selected sensors to the selected sites, so that each target under this mapping is covered by multiple sensors.

The MCPC problem is Max-SNP-complete, and so it does not have polynomial-time approximation schemes unless P = NP. On the other hand, there are approximation algorithms for this problem with proven approximation guarantees. These algorithms can be characterized as greedy set-cover selection,

Y. Li et al. (Eds.): WASA 2008, LNCS 5258, pp. 188–199, 2008.

greedy LP-rounding, and randomized LP-rounding. The approximation ratios of these algorithms are based on different attributes that are not directly comparable, and so they provide little insights as how well each of these approximation algorithm would perform in practical applications.

To obtain better insights, it is desirable to compare these algorithms using numerical experiments and this paper takes up this task. In particular, we design and carry out a large number of numerical experiments on randomly generated sensors and sites with various densities. we show that greedy LP-rounding provides better approximation results than both randomized LP-rounding and greedy set-cover selection. We also show that the solutions produced by randomized LP-rounding are unstable. That is, running the randomized LP-rounding algorithm at different times on the same set of data will produce fluctuating results of wide spans, where the cost differences can be as large as 20%.

Next, we devise a new approximation algorithm using a convex combination of greedy LP-rounding and greedy set-cover selection. We call our new algorithm *alpha-beta approximation*. We show that alpha-beta approximation provides feasible solution to the MCPC problem. We then show that, through a large number of experiments on randomly-generated targets and sites with different densities, alpha-beta approximation provides better approximation results than all previous approximation algorithms. In particular, the actual approximation ratios of alpha-beta approximation in our experiments never exceed 1.07, and it provides up to 14.86% improvement over the best approximation using previous algorithms.

This paper is structured as follows. Section 2 describes the MCPC problem, greed set-cover selection, greedy LP-rounding, and randomized LP-rounding. Section 3 presents alpha-beta approximation. Section 4 describes experiment settings and provides numerical results. Section 5 presents final remarks and open problems.

2 Preliminaries

Let R denote a set of targets spread across a 3D (or 2D) geographical region and S a set of sites to place sensors in the proximity of targets, where R and S may or may not intersect. Targets and sites are represented as 3D (or 2D) points.

Let $\langle t_1, \ldots, t_\ell \rangle$ be ℓ types of sensors with sensing radius $\langle r_1, \ldots, r_\ell \rangle$ and monetary costs $\langle c_1, \ldots, c_\ell \rangle$, where $r_1 < \cdots < r_\ell$. We assume that a sensor can only be placed on a point, and each point can only be occupied by at most one sensor. We also assume that there is an unlimited supply for each type of sensor. Moreover, we assume that S is fully usable to R; that is, for every site j there is at least one target i and one sensor type t_v such that i falls in the sensing range of a type-t_v sensor placed at site j.

The basic form of the MCPC problem is to select sensors and sites to place these sensors, such that every target in R is covered by at least σ sensors and that the total monetary cost of the selected sensors is minimum, where $\sigma \geq 1$ is a given integer.

Denote by S_j^v the set of targets that can be covered by a type-t_v sensor placed at site j. Let c_v be the cost of set S_j^v. Then solving the MCPC problem is equivalent to solving the weighted set multicover problem, which is known to be Max-SNP-complete.

Let $|R| = n$ and $|S| = m$. For simplicity, we label targets and sites as $R = \{1, 2, \ldots, n\}$ and $S = \{1, 2, \ldots, m\}$. Denote by $d(i, j)$ the Euclidean distance between target $i \in R$ and site $j \in S$. Let

$$\begin{aligned} E_v &= \{(i, j) \mid 0 \leq d(i, j) \leq r_v,\ i \in R,\ j \in S\},\ v = 1, \ldots, \ell. \\ E_v[i] &= \{j \mid (i, j) \in E_v\},\ i = 1, \ldots, n. \\ E'_v[j] &= \{i \mid (i, j) \in E_v\},\ j = 1, \ldots, m. \end{aligned}$$

Note that $S_j^v = E'_v[j]$. Solving the MCPC problem is equivalent to solving the following ILP problem, where x_j^v is a 0-1 variable indicating the number of type-t_v sensor placed at site j:

$$\text{Minimize } \sum_{j \in S} \sum_{v=1}^{\ell} c_v \cdot x_j^v$$

$$\text{Subject to } (\forall\, i \in R) \sum_{v=1}^{\ell} \sum_{j \in E_v[i]} x_j^v \geq \sigma, \tag{1}$$

$$(\forall j \in S) \sum_{v=1}^{\ell} x_j^v \leq 1. \tag{2}$$

A feasible solution of this ILP model is also referred to as a σ-*cover*.

Greedy set-cover selection, greedy LP-rounding, and randomized LP-rounding approximation algorithms have the following approximation ratio upper bounds r:

1. For greedy set-cover selection [RV99],

$$r = H_d = \sum_{i=1}^{d} \frac{1}{i} = 1 + O(\log d),$$

   where d is the largest number of targets that can be covered by a sensor, namely, $d = \max_{j \in S,\ 1 \leq v \leq \ell} |E'_v[j]|$.

2. For greedy LP-rounding [WZ06, WZ08],

$$r = f - \sigma + 1,$$

   where f is the largest number of sensors that cover a target, namely, $f = \max_{i \in R}(\sum_{v=1}^{\ell} |E_v[i]|)$.

3. For randomized LP-rounding [Vaz01],

$$r = O(\log n).$$

Greedy Set-Cover Selection

We say that a target is σ-*covered* if it is covered by at least σ sensors. The greedy set-cover selection selects the largest set with the smallest cost at each step as follows:

1. Set $C \leftarrow \emptyset$, $S' \leftarrow \{S_j^v \mid 1 \leq j \leq m \text{ and } 1 \leq v \leq \ell\}$, and $A \leftarrow R$.
2. Choose an S_j^v from S' such that $c(S_j^v)/|a(S_j^v)| = \min_{S_i^u \in S'}\{c(S_i^u)/|a(S_i^u)|\}$, where $a(S_j^v) = S_j^v \cap A$ is the set of targets in S_j^v that are still active; namely, these targets are not σ-covered yet at this point.
3. Set $C \leftarrow C \cup S_j^v$, remove from A all the targets that are σ-covered, and remove S_j^v from S'.
4. If $A \neq \emptyset$, go back to Step 2.
5. Output C.

Greedy LP-Rounding

A simple form of greedy LP-rounding is as follows:

1. Solve the LP-relaxation of the ILP model by allowing variables x_j^v to take real values between 0 and 1.
2. Let $S^* = \{x_j^{v,*} \mid x_j^{v,*} > 0, 1 \leq j \leq m, \text{ and } 1 \leq v \leq \ell\}$ be the optimal solution to the LP model. Sort S^* in non-increasing order to produce a sorted list L^*.
3. Select a variable from L^* one at a time, round it to 1, until a σ-cover is obtained, where variables not selected are set to 0.

Randomized LP-Rounding

Treat each value $x_j^{v,*} \in (0,1]$ in the optimal solution S^* to the LP model as a probability. We use a biased coin to select set S_j^v with probability $x_j^{v,*}$ for all j and v; i.e. set $x_j^v = 1$ if its corresponding biased coin toss turns to head, and set $x_j^v = 0$ otherwise. This forms a sub-collection of sets. Repeat this process independently $k \log n$ times and compute the union of all the sub-collections of sets, where k is a constant such that $(1/e)^{k \log n} \leq 1/(4n)$. When $\sigma = 1$, it can be shown [Vaz01] that the resulting collection of sets from the union is, with high probability, a σ-cover with an approximation ratio of $O(\log n)$.

We will show in Section 4.2 that randomized LP-rounding does not produce stable results. In particular, our numerical experiments show that it produces fluctuating results on the same set of data with as much as 20% difference from different executions. Thus, randomized LP-rounding may only have theoretical interests.

3 Alpha-Beta Approximation

We observe that in greedy LP-rounding, we can obtain a better approximation by also considering greedy set-cover selection. In other words, in addition to considering the value of $x_j^{v,*}$, we also consider how many targets a type-t_v sensor placed at site j can cover. Since $|a(S_j^v)|$ may be much larger than 1 and $x_j^{v,*} \le 1$, we will consider $|a(S_j^v)|/K$ instead of $|a(S_j^v)|$ to balance two greedy strategies, where K is the largest number of targets a sensor can cover. That is,

$$K = \max_{1 \le j \le m, 1 \le v \le \ell} \{|E'_v[j]|\}.$$

We consider the convex combination of x_j^v and $|a(S_j^v)|/K$, namely, let

$$\begin{aligned} h(x_j^v) &= \alpha \cdot x_j^v + \beta \cdot \frac{|a(S_j^v)|}{K}, \\ \alpha + \beta &= 1, \\ \alpha &\ge 0, \\ \beta &\ge 0. \end{aligned}$$

We will then select S_j^v if $x_j^{v,*} > 0$ and $h(x_j^{v,*})$ is large.

Alpha-Beta Approximation Algorithm

1. Select values of α and β.
2. Set $C \leftarrow \emptyset$, $A \leftarrow R$ (the set of all targets), and $G \leftarrow \{x_j^v \mid x_j^{v,*} > 0\}$.
3. Let $x_j^v \in G$ and $h(x_j^{v,*}) = \max_{x_i^u \in G}\{h(x_i^{u,*})\}$.
4. Set $C \leftarrow C \cup S_j^v$ and $G \leftarrow G - \{x_j^v\}$.
5. Remove from A all the targets that become σ-covered, and compute $a(S_j^v) = S_j^v \cap A$ for all $1 \le j \le m$ and $1 \le v \le \ell$.
6. If $A \ne \emptyset$, goto Step 3.
7. Output C.

Clearly, setting x_j^v to 1 for all $x_j^{v,*} > 0$ provides a σ-cover. Thus, the alpha-beta approximation algorithm guarantees a σ-cover.

We note that the alpha-beta approximation is the same as greedy set-cover selection when $\alpha = 0$, and is the same as greedy LP-rounding when $\beta = 0$. When $\alpha \ne 0$ and $\beta \ne 0$, we note that the alpha-beta approximation algorithm may select a variable x_j^v that will not be selected by greedy LP-rounding or greedy set-cover selection. In our numerical experiments, this phenomenon have happened. Thus, it is possible that, by selecting α and β appropriately, the alpha-beta approximation can produce better results than greedy LP-rounding and greedy set-cover selection. Our numerical experiments confirm this observation.

4 Numerical Experiments and Performance Analysis

We describe our settings for numerical experiments and present numerical results on ILP, alpha-beta approximation, greedy LP-rounding, greedy set-cover selection, and randomized LP-rounding.

4.1 Experiment Settings

In our experiments, we use three types of sensors A, B, and C, where $\langle r_A, r_B, r_C \rangle = \langle 15, 25, 40 \rangle$ and $\langle c_A, c_B, c_C \rangle = \langle 200, 350, 580 \rangle$. Targets and sites are generated in a 300×300 area, where targets are generated uniformly and independently at random. For each target generated, 5 sites are generated uniformly and independently at random within the radius of 20 of the target. This allows all three types of sensors to be evenly selected in the optimal solutions. Experiments are carried out with $\sigma = 1, 2, 3$ on a variety of target density for each value of σ, from sparse to dense, with n = 100, 200, 300, 400, 500, and 600.

4.2 Fluctuating Behavior of Randomized LP-Rounding

Figure 1 is a sample of fluctuating behavior of randomized LP-rounding from 20 executions on the same set of 600 targets, where $\sigma = 3$. Our experiments show that, when $\sigma \geq 2$ and $n \geq 200$, there is an up to 20% difference in the results produced by randomized LP-rounding from different executions on the same set of data.

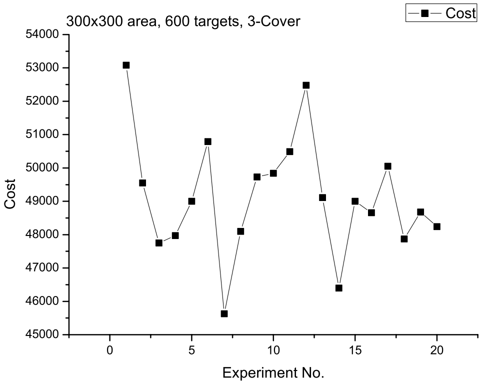

Fig. 1. Fluctuating behavior of randomized LP-rounding

4.3 Cost Curves in Terms of α

The values of α and β determine how much the LP solution and set-cover selection will influence the final decision. We show that, through numerical experiments, the value of α should be roughly reciprocal of K. For example, we observe that, when $K \leq 20$, choosing $\alpha = 0.6$ is better. When $K \geq 25$, choosing $\alpha = 0.2$ is better. Fig. 2 shows two examples of cost curves in terms of α.

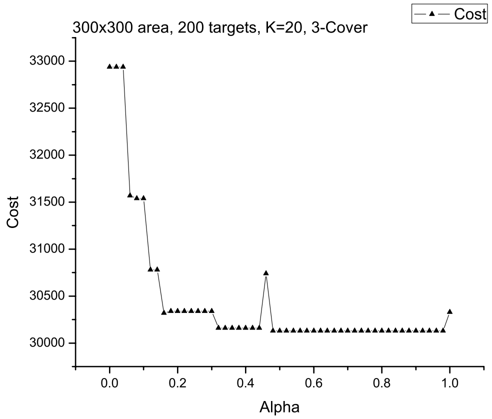

(a) $K = 20$

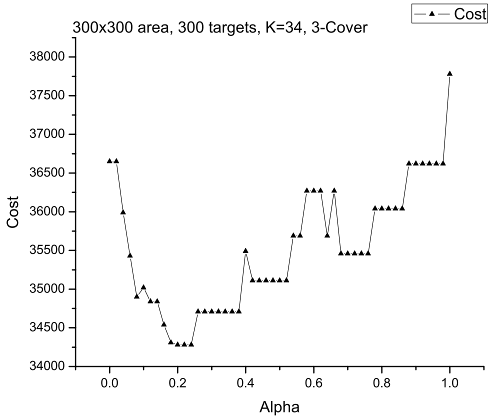

(b) $K = 34$

Fig. 2. Cost curves vs. α values at different target density

4.4 Variation of Alpha-Beta Approximation

The alpha-beta approximation algorithm runs in polynomial time, but it is slower than greedy LP-rounding. This is because calculating

Table 1. Comparisons of running time and results between original alpha-beta approximation and its variation

number of targets	σ	*original time* (10^{-3} *s*)	*variation time* (10^{-3} *s*)	*original result*	*variation result*
100	1	40	10	8680	8680
200	1	90	30	10150	10150
300	1	420	410	11950	11950
400	1	920	910	12900	12900
500	1	1720	1720	12940	12940
600	1	2110	2110	12960	12960
100	2	90	10	15940	15940
200	2	470	300	21150	21150
300	2	1400	1340	23370	23370
400	2	1560	1470	24190	23610
500	2	3100	2900	24820	24240
600	2	4570	4420	25200	25200
100	3	200	20	26680	26680
200	3	600	490	30130	30130
300	3	1690	1440	34280	34280
400	3	2720	2510	35500	35500
500	3	5410	4740	37980	37800
600	3	5620	5380	38210	38210

$$h(x_j^{v,*}) = \max_{x_i^u \in G}\{h(x_i^{u,*})\}$$

involves counting active targets for each set. If we can avoid doing this calculation as much as we can without losing accuracy, we can reduce computing time of the algorithm. For example, when $x_j^{v,*}$ is closer to 1 in an LP solution, S_j^v is more likely to be selected in the ILP solution. Thus, we may choose an appropriate threshold value t and select S_j^v directly when $x_j^{v,*} \geq t$.

Our numerical experiments show that when $t = 0.9$ and $\sigma > 1$, the variation of alpha-beta approximation can reduce much running time while producing almost the same result as (and at time even slightly better than) the original alpha-beta approximation. The variation is more effective when the density of targets is low, for there would be more $x_j^{v,*}$ close to 1 in the LP solution. Table 1 compares the running time and results generated by the original alpha-beta algorithm and its variation with $t = 0.9$.

4.5 Performance Comparisons

In all of our experiments, except for greedy set-cover selection, all approximation algorithms generate result is less than 7 seconds. Computing optimal solutions take much longer time and it takes a number of days of running time when n is large. When $\sigma > 1$, the solution produced by alpha-beta approximation has the actual approximation ratio in the range between 1.01 to 1.07.

Table 2. Performance Comparison of all algorithms with $\sigma = 2$

n	*LP*	*OPT*	*set-cover*	r_{sc}	*LP-round.*	r_{lp}	
100	15760	15760	17680	1.121827	16340	1.036802	
200	19791.61	19880	23800	1.197183	22710	1.142354	
300	21593.5	21980	26810	1.219745	25810	1.174249	
400	21513	22450	27830	1.239644	27730	1.235189	
500	21928.7	23020	28490	1.237619	27810	1.20808	
600	22307.7	23920	29410	1.229515	28840	1.205686	
	Rand-LP rounding	r_{rlp}	*alpha-beta original*	$r_{\alpha\beta}$	*alpha-beta variation*	$r'_{\alpha\beta}$	*Imp*
100	16140	1.024112	15940	1.011421	15940	1.011421	1.24%
200	23640	1.189135	21150	1.063883	21150	1.063883	6.87%
300	29010	1.319836	23370	1.063239	23370	1.063239	9.45%
400	29980	1.335412	24190	1.077506	23610	1.05167	14.86%
500	31060	1.349262	24820	1.078193	24240	1.052997	12.84%
600	32220	1.346990	25200	1.053512	25200	1.053512	12.62%

Table 2 shows the optimal solution and the approximation solution, for $\sigma = 2$, of greedy set-cover selection, greedy LP-rounding, randomized LP-rounding, original alpha-beta approximation, variation of alpha-beta approximation, and the improvement of alpha-beta over LP-rounding. Here r_{sc}, r_{lp}, r_{rlp}, $r_{\alpha\beta}$, and $r'_{\alpha\beta}$ denote, respectively, the approximation ratio of the greedy set-cover selection, greedy LP-rounding, randomized LP-rounding, alpha-beta approximation, and alpha-beta variation. We use *Imp* to denote the improvement of alpha-beta variation over greedy LP-rounding.

Figure 3 compares the cost from solutions generated by each algorithm, i.e., the ILP model, greedy set-cover selection, the greedy LP-rounding, randomized LP-rounding, and alpha-beta approximation with $\sigma = 1, 2, 3$, respectively. It shows that the alpha-beta approximation provides the best solution than any other approximation algorithm.

Figure 4 shows visualized solutions generated by the ILP model and the alpha-beta approximation for $n = 100$ targets.

5 Final Remarks and Open Problems

This paper presents alpha-beta approximation algorithm for the MCPC problem and provides numerical results that compares performance of the alpha-beta algorithm with other approximation algorithms. We show that alpha-beta approximation algorithm provides better solutions with up to 14.86% improvement over the best approximation algorithm previously known, and it does so efficiently. This indicates that, alpha-beta approximation, i.e., a convex combination of greedy LP-rounding and greedy set-cover selection is a promising new approach. There are a number of issues that warrant a further investigation.

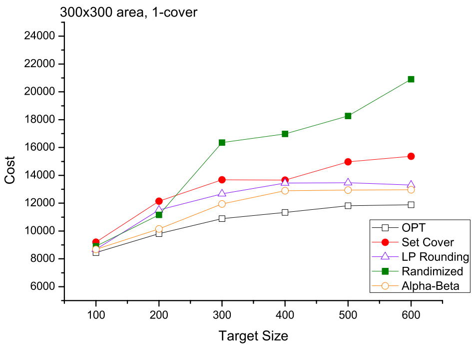

(a) $\sigma = 1$

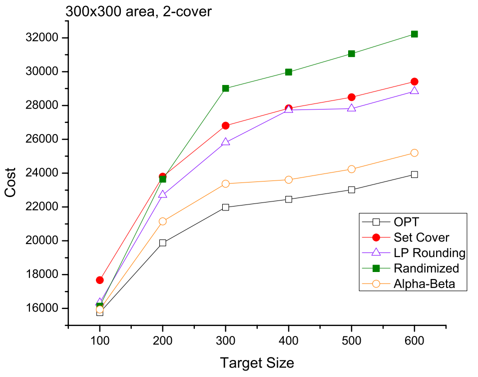

(b) $\sigma = 2$

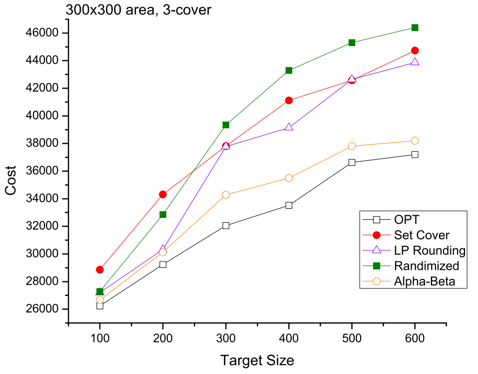

(c) $\sigma = 3$

Fig. 3. Performance comparison of all algorithms

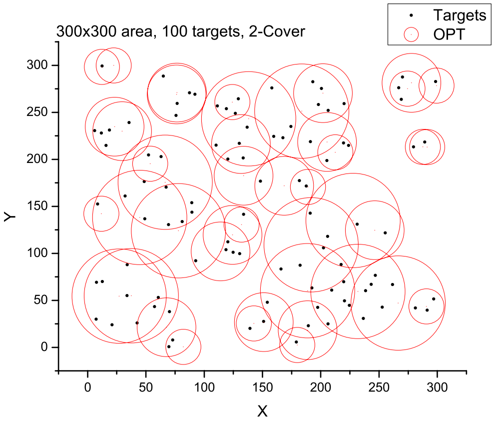

(a) OPT

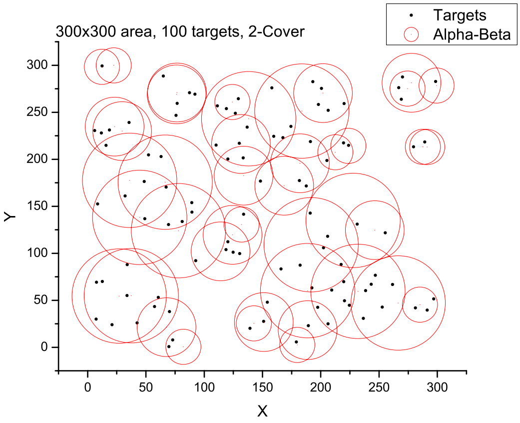

(b) alpha-beta (ratio = 1.01)

Fig. 4. Visualized approximation solutions

For example, we suspect that the approximation ratio of alpha-beta approximation would depend on the values of α and the approximation ratios of greedy LP-rounding and greedy set-cover selection.

Acknowledgements

This work was supported in part by NSF under grant CNS-0709001. Points of view in this document are those of the authors and do not necessarily represent the official position of NSF.

References

[CW04] Cardei, M., Wu, J.: Coverage in Wireless Sensor Networks. In: Ilyas, M., Magboub, I. (eds.) Handbook of Sensor Networks, CRC Press, Boca Raton (2004)

[CIQ$^+$02] Chakrabarty, K., Iyengar, S.S., Qi, H., Cho, E.: Grid coverage for surveillance and target location in distributed sensor networks. IEEE Trans. on Comput. 51(12), 1448–1453 (2002)

[Hoc97] Hochbaum, D.: Approximating covering and packing problems: set cover, vertex cover, independent set, and related problems. In: Hochbaum, D. (ed.) Approximation Algorithms for NP-Hard Problems, PWS Publishing Company (1997)

[RV99] Rajagopalan, S., Vazirani, V.V.: Primal-dual RNC approximation algorithms for set cover and covering integer programs. SIAM J. on Computing 28, 526–541 (1999)

[SX05] Sahni, S., Xu, X.: Algorithms for wireless sensor networks. Intl. Jr. on Distr. Sensor Networks 1(1), 35–56 (2005)

[Vaz01] Vazrani, V.V.: Approximation Algorithms. Springer, Heidelberg (2001)

[WZ06] Wang, J., Zhong, N.: Efficient point coverage in wireless sensor networks. Journal of Combinatorial Optimization 11, 291–305 (2006)

[WZ08] Wang, J., Zhong, N.: Minimum-cost sensor arrangement for achieving wanted coverage lifetime. International Journal on Sensor Networks (in press)

[YW07] Yu, Z., Wang, J.: Reliable sensor arrangement for achieving wanted coverage lifetime with minimum cost. In: Proc. of the 2nd International Conference on Wireless Algorithms, Systems and Applications (WASA 2007), Chicago, August 2007, pp. 95–102. IEEE Computer Society Press, Los Alamitos (2007)