The Wake Up Dominating Set Problem

Amir Bannoura, Christian Ortolf, Christian Schindelhauer, and Leonhard Reindl

Faculty of Engineering, University of Freiburg, Germany
{ortolf, schindel}@informatik.uni-freiburg.de
{amir.bannoura, reindl}@imtek.de

Abstract. Recently developed wake-up receivers pose a viable alternative for duty-cycling in wireless sensor networks. Here, a special radio signal can wake up close-by nodes. We model the wake-up range by the unit-disk graph. Such wake-up radio signals are very energy expensive and limited in range. Therefore, the number of signals must be minimized. So, we revisit the Connected Dominating Set (CDS) problem for unit-disk graphs and consider an online variant, where starting from an initial node all nodes need to be woken up, while the online algorithm knows only the nodes woken up so far and has no information about the number and location of the sleeping nodes.

We show that in general this problem cannot be solved effectively, since a worst-case setting exists where the competitive ratio, i.e. the number of wake-up signals divided by the size of the minimum CDS, is $\Theta(n)$ for n nodes. For dense random uniform placements, this problem can be solved within a constant factor competitive ratio with high probability, i.e. $1 - n^{-c}$.

For a restricted adversary with a reduced wake-up range of $1 - \epsilon$ we present a deterministic wake-up algorithm with a competitive ratio of $O(\epsilon^{-\frac{1}{2}})$ for the general problem in two dimensions.

In the case of random placement without any explicit position information we present an $O(\log n)$-competitive epidemic algorithm with high probability to wake up all nodes. Simulations show that a simplified version of this oblivious online algorithm already produces reasonable results, that allows its application in the real world.

1 Introduction

Energy is the driving problem of wireless sensor networks (WSNs), since sensor nodes usually operate for long periods and the only source of energy is battery cells which are difficult to be exchanged. The functionality of WSNs can be extended through the use of low power microprocessors, sensors, and radio transceivers. The availability of low power hardware components provided a technological break-through of wake-up receivers. These receivers interact only when a special wake-up signal is addressed to them. When a wake-up signal is received, the wake-up receiver triggers an interrupt to wake the sensor node. In addition, any sensor node has the capability of transmitting a wake-up signal to wake up all other close-by sensor nodes. Recent research [9] has decreased the energy consumption of sensor nodes when no activity is required to less than 9 μW, whereas sensor nodes that are not equiped with wake-up receivers who uses duty cycle would be spending around 51 mW checking the medium from time to time. Fig 1 shows a sample board of a designed wake-up receiver integrated with a wireless sensor node.

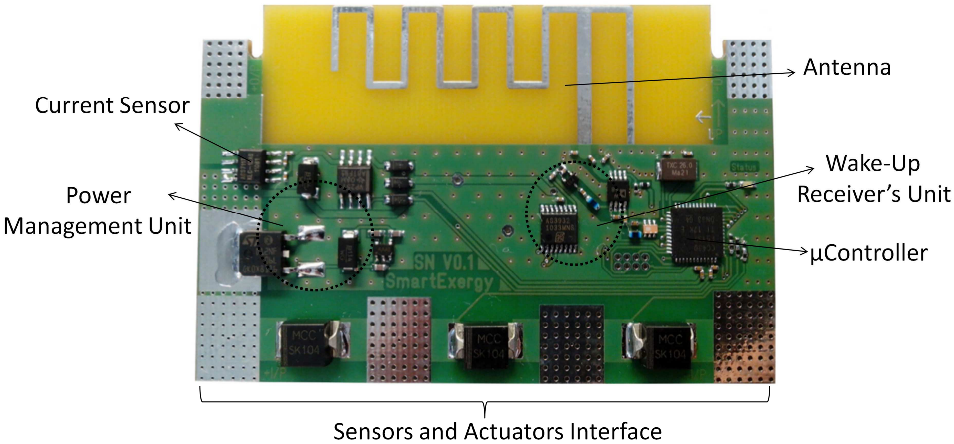

Fig. 1. Wireless Sensor Node with Integrated Wake-Up Receiver

Wake-up receivers integrated in sensor nodes constitute a paradigm shift for wireless sensor protocols. In which, sensor nodes interact with the surrounding neighbors only when they are required to receive and send information. The duty cycle process for periodically checking to find out whether messages need to be received or synchronizing with other sensor nodes may not be required anymore.

Despite that this technology provides a new solution for the energy consumption problem, new problems arise. Sensor nodes are required to produce a wake-up signal, these signals are energy expensive compared with the signals that are required for normal data communication. Furthermore, the communication range of a wake-up signal is smaller than the normal data communication range, which requires a multi-hop wake-up signals to wake up sensor nodes that are located in the normal data communication range. Our aim is to minimize the number of wake-up singals transmitted as much as possible and maximize the covered area to reduce the energy required to wake up sensor nodes.

A straight-forward solution is to establish minimum set of sensors which are able to wake up all sensor nodes in case some data need to be collected or distributed. This is the well known Connected Dominating Set (CDS) problem, where one tries to compute the Minimum CDS (MCDS). This problem plays an important role in wireless networks and it is known to be NP-complete.

For simplicity we assume a unit-disk range model and we are interested in computing the minimum dominating connected set in unit disk graphs, i.e.

geometric graphs where an undirected edge exists between nodes, if their distance is at most 1.

However, our problem is somehow different. When the sensor nodes are placed, no positions are known before the first wake-up signal. Also, the nodes may be moving, sensor nodes may fail, and possible persistent memory might not be available. All these are reasons to build up a CDS from scratch regularly.

So, we face an online version of the MCDS problem in the context of wake-up receivers. At the beginning, one sensor node wakes up, e.g. because of new sensor data. It sends a wake-up signal and receives responses from all next neighbors. Then, a decision needs to be taken which of the neighbored sensors is allowed to send the next wake-up signal. Since normal data communication consumes only little energy compared to the wake-up signal, we can assume that all active nodes are aware of each other. Furthermore, the information which sensor received a wake-up signal is available to us, even if the sensor has already been woken up. The question is now, can we wake up all nodes with minimal number of wake-up signals. This is what we address as the *wake-up minimum connected dominating set problem in unit disk graphs*. In this variant the positions of the woken up nodes become available as soon as they are awake. For the *wake-up position-aware minimum connected dominating set problem in unit disk graphs* positions are not known at all.

2 Related Work

The new wake-up receivers developed by Gamm et. al. in [9] give us an alternative to the concept of duty cycles for awaiting incoming messages in wireless sensor networks.

A perfectly efficient online wake-up would use a minimum connected dominating set of nodes to wake up all the nodes. Finding such a MCDS was already shown to be NP-hard for general graphs, as well as for unit disc graphs [13, 4]. For the general (non unit-disk graph) problem no polynomial time approximation exists unless $NP \subseteq DTIME[n^{O(\log \log n)}]$ [10], yet for MCDS with unit disc graphs a PTAS has been presented in [3].

Movement of sensor nodes and maintaining an existing MCDS in their presense was discussed before by Das et al in [5].

Of course the wake-up problem is an online version of MCDS, because of the differences to the online version presented by Eidenbenz [8] we are referring to it as the wake-up problem. Eidenbenz models the online problem by node added every round by an adversary, while the online algorithm has to present a CDS, but may never remove nodes once added to the CDS. He shows a competitive ratio of $\Theta(n)$ for the CDS size.

Another online MCDS problem closer related to the wake-up problem is discussed as broadcast problem by Bar-Yehuda et al. [1] and leads to the same lower bound of $\Omega(n)$, so does the reactive routing problem [14]. This motivates why comparing to an adversary with the same radio range is pointless, as asymptotically no online algorithm can beat trivial flooding, i.e. using the whole graph as CDS.

Further research discusses Minimum Routing Cost CDS (MOC-CDS) [6] which requires the hop distance, between any pair of nodes to be minimal. This problem is also NP hard, but for the unit disc graph a PTAS exists [7]. A more generalized version of the problem is called α-MOC-CDS, where the dominating set must have an α-spanner property additionally. For $\alpha = 1$ these problems are the same.

For distributed generation of MCDS approximations are discussed in [15]. Our problem differs, since not all nodes are awake in the beginning, but have to discover all other nodes from the starting node, leading to different time and message complexity.

The motivation for proving an algorithm for a dense random network stems from the requirement for density in a random unit disk graph to guarantee connectivity of the network [11].

In the position oblivious case we use a push-based epidemic rumor spreading algorithm with a simple counter mechanism, similar to the one in [12]. Rumor spreading turns out to be simple and robust. For other applications epidemic algorithms have been already proposed. For a survey of epidemic algorithms in wireless sensor networks we refer to the chapter 3, Epidemic Models, Algorithms and Protocols in Wireless Sensor and Ad-hoc Networks in [2] by Das and Prabib.

3 Preliminaries

We assume that points are in general positions, i.e. that neither three points are on a line nor four points on a circle in two dimensions. For points in two dimensions or three dimensions the unit-disk graph (UDG) of a given point set V is an undirected graph with edge set $E := \{\{u, v\} \mid u, v \in V \ : |u, v| \leq 1\}$. Later on we refer also to UDGs with different radius r.

We consider the following problems.

Definition 1. *Given an undirected graph $G = (V, E)$ a* connected dominating set (CDS) *S has the following properties*

1. *S is connected in G, i.e. for all $u, v \in S$ there exists a path from u to v in G using only nodes of S.*
2. *S is dominating all nodes in V, i.e. for all $u \in V$ there exists a node $v \in S$ such that $\{u, v\} \in E$.*

Definition 2. *The* wake-up position oblivious Minimum Connected Dominating Set Problem in Unit Disk Graphs (Wake-Up-PO-MCDS-UDG) *is to construct a CDS where the algorithm works in rounds and starts with a node s_0.*

1. *At the beginning only the nodes $V_1 = \{u \in V : \{u, s_0\} \in E\}$ and edges $E_1 = \{\{u, s_0\} : u \in V_1\}$ are known.*
2. *In round i a node $s_i \in V_1$ may be selected by the algorithm and then the nodes $V_{i+1} = \{u \in V : \{u, s_i\} \in E\}$ and edges $E_{i+1} = \{\{u, s_i\} : u \in V_1\}$ are added to the knowledge base of the algorithm.*

In the wake-up position-aware minimum connected dominating set problem in unit disk graphs *also the position of the known nodes is available to the algorithm.*

4 Lower Bounds

The problem of computing the minimum connected dominating set for unit-disk graphs (MCD-UD) has been proven to be NP-complete by Lichtenstein [13]. Lower approximation bounds are not known, while the best approximation factor so far has a bound of 3.8 [16].

For the Wake-up version there is a trivial, but hard computational lower bound for the competitive ratio, i.e. the number of nodes of a connected dominating set woken up by an algorithm divided by the number of nodes of the MCDS.

Proposition 1. *The competitive ratio of all deterministic algorithms for Wake-Up-MCD-UD is at least* $\frac{n}{2} - \frac{1}{2}$*. For probabilistic algorithms the expected competitive ratio is at least* $\frac{n}{4}$*.*

Proof. We use a variant of the construction presented in [1, 8, 14], see Fig. 2. The optimal solution uses wake-up calls from the start node s and the node u_i connected to t. Any deterministic algorithm can be fooled to use $n-1$ wake-up calls of nodes $u_1, \ldots, u_{n-2}$ such that the final wake up call reaches t.

If the connected node u_i is chosen randomly, then any randomized algorithm needs in the expectation $1 + \frac{n-2}{2} = \frac{n}{2}$ calls to launch a wake-up at u_i.

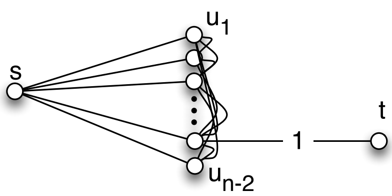

Fig. 2. Lower Bound Construction with competitive ratio $n/2 - \frac{1}{2}$

5 Algorithms

While the wake-up problem can not be efficiently solved in general, for high node density a straight-forward grid based algorithm already achieves constant approximation ratio.

5.1 A grid-based online algorithm

We partition the area into a grid of a square size of $\frac{1}{\sqrt{5}}$ for two dimensions and $\frac{1}{\sqrt{6}}$ for three dimensions, see Fig. 3. This grid size guarantees that any node in a cell can reach all nodes in (orthogonally) neighbored cells with a unit-distance wake-up call. We assume that each node is aware of its grid position and let denote cell(u) the grid cell of u.

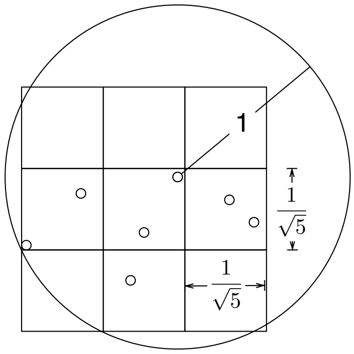

Fig. 3. Grid construction in two dimensions.

The grid based wake-up algorithm 1 chooses a representative for each cell and performs a flooding on the grid structure. In particular, it solves the problem if all grid cells are non-empty. Note that such a $m \times m$ grid can be covered only by a CDS of size of at least $(\frac{m}{\sqrt{5}} - 1)^2 = \frac{1}{5}m^2 - \frac{2}{\sqrt{5}}m + 1$ in two dimensions, while the number of nodes who perform wake up calls is bounded by m^2. Hence, in a square we have a competitive ratio of $5 + o(1)$, and in a cube a competitive ratio of $6 + o(1)$ by an analogous calculation.

Algorithm 1: Grid based wake-up algorithm

Send wake up from s
$G_{\text{done}} \leftarrow \{\text{cell}(s)\}$
$G_{\text{to-do}} \leftarrow \{\text{cell}(u) \ : \ \{u,s\} \in E\} \setminus \{\text{cell}(s)\}$
while $G_{to\text{-}do} \neq \emptyset$ **do**
  Pick a node w such that $\text{cell}(w) \in G_{\text{to-do}}$
  Send wake up from w
  $G_{\text{done}} \leftarrow G_{\text{done}} \cup \{\text{cell}(w)\}$
  $G_{\text{to-do}} \leftarrow G_{\text{to-do}} \cup \{\text{cell}(u) \ : \ \{u,w\} \in E\} \setminus \text{cell}(s) \setminus G_{\text{done}}$
end

If enough nodes are placed randomly, then every cell is occupied with high probability, i.e. $1 - n^{-c}$ for a constant c where $c \in (0, 1)$.

Theorem 1. *If n nodes are place randomly in a $m \times m$-grid with $m \leq \sqrt{\frac{n}{c \ln n}}$ for some constant c. Then the grid based wake-up algorithm computes a CDS with a constant competitive ratio with high probability, i.e. $1 - \frac{1}{n^{c+1}}$ for any $c > 1$.*

Proof. A node is placed in one of the m^2 cells with probability $\frac{1}{m^2} \geq \frac{c \ln n}{n}$. Therefore, the chance that it is not placed in a cell is $1 - \frac{c \ln n}{n}$. So, the probability that a cell is empty can be upperbounded as follows.

$$\left(1 - \frac{c \ln n}{n}\right)^n \leq e^{-c \ln n} = n^{-c} \tag{1}$$

We have used $(1 - 1/m)^m \leq 1/e$ for $m > 0$. By the union bound the probability that any of cell is empty is therefore at most n^{-c+1}. □

Note that when the node density is decreased only by a constant factor, that the unit disk graphs becomes disconnected [11] and there is no solution for CDS.

5.2 A competitive algorithm with respect to a weaker adversary

We have seen that for non-randomized placement the wake-up variant cannot compete with the offline version, which can be seen as an adversary which places the sleeping nodes in the area outside of the wake up signals. However, if we compare the wake-up algorithm with a weaker adversary, we can show some interesting results.

For this we consider two unit-disk graphs with radius 1 and $1 - \epsilon$ for some $\epsilon \in (0, 1)$. The wake-up algorithm solves the wake-up problem for CDS in unit disk graphs, i.e. with radius 1. We compare its performance to the size of the minimum connected dominating set of the unit disk graph with radius $1 - \epsilon$ of the same event. In this way, counter-examples cannot occur as shown in Fig. 2.

Theorem 2. *In two dimensions there is a wake-up algorithm which produces at most $O(\epsilon^{-1/2})$ more wake-up calls than the number of nodes of the CDS of the $1 - \epsilon$ unit disk graph for $\epsilon \in (0, 1)$.*

Proof. The key idea is for a node u to cover the two-hop neighborhood of the $1 - \epsilon$ unit disk graph with a set of nodes in the neighborhood of the unit-disk graph. We prove that the size of this set of this $\left\lceil \frac{32\pi}{\sqrt{\epsilon}} \right\rceil$-coverage boundary set nodes is bounded by $O(\frac{1}{\sqrt{\epsilon}})$.

Using this observation we use a grid based approach like in [14]. If in two dimensions we choose the grid size of $\frac{1}{\sqrt{2}}(1 - \epsilon)$, then any node in a cell of size u can reach with two hops all nodes which any node of its cell can reach in one hop, since the diagonal of the cell is $1 - \epsilon$. So, it suffices to broadcast a message in a grid, which can be done in constant factor overhead.

It remains to construct the coverage of the two-hop neighborhood of the $1 - \epsilon$ unit disk graphs with $O(\sqrt{\epsilon})$ nodes. For this, we need to investigate some properties of the two-hop neighborhood.

Algorithm 2: $(1-\epsilon)$ cover two hop wake-up grid algorithm

Send wake up from s
$G_{\text{done}} \leftarrow \{\text{cell}(s)\}$
$G_{\text{to-do}} \leftarrow \{\text{cell}(u) \ : \ \{u,s\} \in E\} \setminus \{\text{cell}(s)\}$
while $G_{to\text{-}do} \neq \emptyset$ **do**
  Pick a node w such that $\text{cell}(w) \in G_{\text{to-do}}$
  Send wake up from w
  Compute the $\left\lceil \frac{32\pi}{\sqrt{\epsilon}} \right\rceil$-coverage boundary set S of the neighbored nodes of w
  forall the $v \in S$ **do**
    $G_{\text{done}} \leftarrow G_{\text{done}} \cup \{\text{cell}(v)\}$
    $G_{\text{to-do}} \leftarrow G_{\text{to-do}} \cup \{\text{cell}(v) \ : \ \{v,w\} \in E\} \setminus G_{\text{done}}$
  end
end

Definition 3. *The two-hop covering node set of a unit-disk-graph of radius r starting with node s is the set of nodes $S \subseteq V$ with the following properties:*

1. *All nodes of S are within distance r of s.*
2. *Each node of S is necessary, i.e. for all nodes of $u \in S$ there exists a point p within distance $[r, 2r]$ from s with $|u,p| \leq r$ and $|v,p| > r$ for all $v \in S \setminus \{u\}$.*

We call the nodes of $S = \{c_1, \ldots, c_m\}$ the *cover nodes*. The *outer ring* is the disk of radius $2r$ without the disk of radius r with center s. The *coverage area* is the union of all disks of radius r and center points of S. The *coverage boundary* is the boundary of this coverage area. By definition it consists of arcs with radius r and center points of S. When two points u, v of S have distance r to the same point w the coverage boundary, we call this point a *boundary point*, see Fig. 4. Now the following geometric observations can be made.

Lemma 1. 1. *Given a point p of the coverage boundary, which is not a boundary point, and the tangent T of the boundary, then the angle δ between (p,s) and T is in the range $\delta \in [\frac{1}{6}\pi, \frac{5}{6}\pi]$.*
2. *The length of the coverage boundary is at most $8\pi r$.*

Proof. The first statement comes from a simple geometric observation which is based on possible placement of nodes c_i, see Fig. 5.

For the second statement, note that each distance d traveled on the boundary region corresponds to an angle difference of at least $\frac{d}{z}\frac{1}{\sin\delta}$, where z is the distance to s which is in the range $[r, 2r]$. The angle of the tangent is δ. Since $\frac{1}{2} \leq \sin\delta \leq 1$ and since the total angle difference is bounded by 2π, it follows that the maximum length of the boundary is bounded by $8\pi r$. □

We neglect the case, where the coverage boundary intersects with the inner ring. It is straight-forward that the following claims also hold for this case.

Using the observation of Lemma 1 we can order all cover points according to their direction seen from s as $c_1, \ldots, c_m$. Only neighbored nodes share a boundary point. We name the angles according to Fig. 4. We denote b_i as the boundary point between c_i and c_{i+1}, and b_m as the boundary node between c_1 and c_m. Let $\alpha_i = \angle c_i b_i c_{i+1}$, $\beta_i = \angle b_{i-1} c_i b_i$ and $\gamma_i = \pi - \angle c_{i-1} c_i c_{i+1}$.

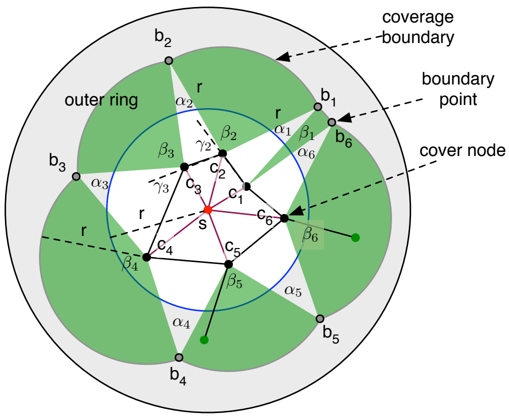

Fig. 4. Definition of coverage boundary, boundary points $b_1, \ldots, b_m$ and cover nodes $S = \{c_1, \ldots, c_m\}$ labeling

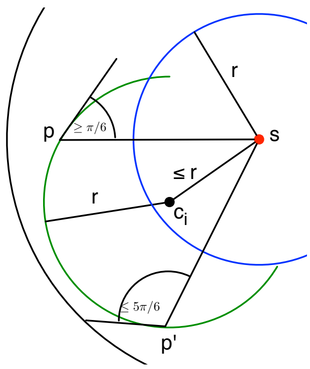

Fig. 5. Angle property of the boundary region

From the definition of the angles we derive the following equalities for all $i \in [1, m]$:

$$\gamma_i = \beta_i - \frac{\alpha_i + \alpha_{i-1}}{2} \ , \tag{2}$$

$$\sum_{i=1}^{m} \gamma_i = 2\pi \ . \tag{3}$$

Since the boundary region is defined by the arcs of angles $\alpha_1, \ldots, \alpha_n$ with radius r now Lemma 1 implies

$$\sum_{i=1}^{m} \beta_i \leq 8\pi \ .$$

So,

$$\sum_{i=1}^{m} \alpha_i = \left(\sum_{i=1}^{m} \alpha_i + \gamma_i \right) - 2\pi = \left(\sum_{i=1}^{m} \beta_i \right) - 2\pi \leq 6\pi$$

While the cover points do not necessarily form a convex hull, its form is quite well behaved. For a large number of cover points we need to find groups of near points, i.e. with small angles α_i and β_i.

Lemma 2. *Given m boundary points then, there exist at least $k \leq m$ interval indices $i_1, \dots, i_k$ such that for all $\nu \in \{1, \dots k\}$:*

$$\sum_{j=i_\nu}^{i_{\nu+1}-1} \alpha_j \leq \frac{12\pi}{k} \quad \textit{and} \quad \sum_{j=i_\nu+1}^{i_{\nu+1}-1} \beta_j \leq \frac{16\pi}{k}$$

Proof. Start with $i_1 = 1$. Now for $\nu = 1, 2, \dots$ choose the largest q such that

$$\sum_{j=i_\nu}^{q} \alpha_j \leq \frac{12\pi}{k} \quad \text{and} \quad \sum_{j=i_\nu+1}^{q} \beta_j \leq \frac{16\pi}{k}$$

and set $i_{\nu+1} := q + 1$.

By definition

$$\sum_{j=i_\nu}^{i_{\nu+1}} \alpha_j > \frac{12\pi}{k} \quad \text{or} \quad \sum_{j=i_\nu}^{i_{\nu+1}} \beta_j > \frac{16\pi}{k}$$

One of this property must be violated more than $k/2$ times. This would imply

$$\sum_{j=1}^{m} \alpha_j > \frac{12\pi}{k}\frac{k}{2} = 6\pi \quad \text{or} \quad \sum_{j=i_\nu}^{i_{\nu+1}} \beta_j > \frac{16\pi}{k}\frac{k}{2} = 8\pi$$

which contradicts Lemma 1 □

Lemma 3. *For given boundary nodes $c_1, \dots, c_m$ with*

$$\sum_{i=1}^{m} \alpha_i \leq \frac{1}{2}\sqrt{\epsilon} \quad \textit{and} \quad \sum_{i=2}^{m} \beta_i \leq \frac{1}{2}\sqrt{\epsilon}\ ,$$

the disks with center $c_1, \dots, c_m$ with radius $1 - \epsilon$ are covered by the two disks with center c_1 and c_m with radius 1.

Proof. From Equation 2 we get for all $\ell \leq m$:

$$\left|\sum_{i=2}^{\ell} \gamma_i\right| = \left|\sum_{i=2}^{\ell} \beta_i - \frac{1}{2}(\alpha_1 + \alpha_\ell) - \sum_{i=2}^{\ell-1} \alpha_i\right| \leq \left|\sum_{i=2}^{\ell} \beta_i - \sum_{i=1}^{\ell} \alpha_i\right| \leq \frac{1}{2}\sqrt{\epsilon}\ .$$

Furthermore, we assume $\epsilon < 1$ and use $\tan \frac{1}{2}\alpha_i = |c_i, c_{i+1}|/2(1 - \epsilon)$ and $\tan(x) \leq 2x$ for $x \in [0, 1]$.

$$|c_1, c_m| \leq \sum_{i=1}^{m-1} |c_i, c_{i+1}| \leq \sum_{i=1}^{m-1} 2(1 - \epsilon) \tan \frac{1}{2}\alpha_i \leq \sum_{i=1}^{m-1} 2(1 - \epsilon)\alpha_i \leq \sqrt{\epsilon}$$

Since all sums $\left|\sum_{i=2}^{\ell} \gamma_i\right| \leq \frac{1}{2}\sqrt{\epsilon}$ and the maximum distance $|c_1, c_m| \leq \frac{1}{2}\sqrt{\epsilon}$ we can conclude that $c_1, \dots, c_m$ fits into a rectangle of length $\frac{1}{2}\sqrt{\epsilon}$ and width $\frac{1}{2}\epsilon$. The rest follows by the following geometric argument.

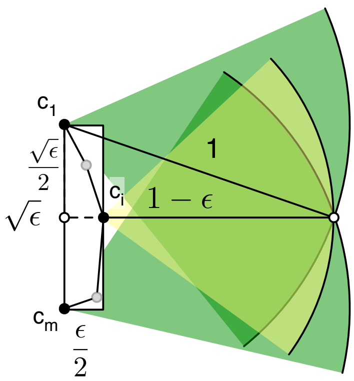

Fig. 6. If the cover nodes are inside a $\sqrt{\epsilon} \times \frac{1}{2}\epsilon$-rectangle, then two disks with radius 1 of the two outermost nodes cover all disks with radius $1-\epsilon$

The worst case placement for the outer nodes c_1 and c_m and some inner node c_i is depicted in Fig. 6. Note that

$$\left(1-\frac{1}{2}\epsilon\right)^2 + \left(\frac{\sqrt{\epsilon}}{4}\right)^2 = 1-\epsilon+\frac{1}{16}\epsilon+\frac{1}{4}\epsilon^2 = 1-\frac{15}{16}\epsilon+\frac{1}{4}\epsilon^2 \leq 1 \ .$$

Therefore the two outermost nodes with radius 1 always cover the disks with radius $1-\epsilon$. □

Putting all pieces together, given a start node s it wakes up all neighbor nodes. They report their position and the algorithms chooses $k = \left\lceil \frac{32\pi}{\sqrt{\epsilon}} \right\rceil$ intervals according to Lemma 2. From these k intervals surrounding s we start wake-up calls only from the nodes at the interval borders, i.e. $c_{i_1}, \ldots, c_{i_k}$. These nodes can cover also the area covered by all cover nodes. So, we need $k = O(\epsilon^{-1/2})$ wake-up calls to wake up all nodes in the two-hop $(1-\epsilon)$-neighborhood. For a grid size of $\frac{1}{\sqrt{2}}(1-\epsilon)$ this includes all points in the neighborhood of any given point. Therefore, the cell based approach will inform at least the same node set as the adversary.

Another implication from the (constant size) cell based approach is that if we only count cells instead of nodes, there are straight-forward linear upper and lower bounds for the number of wake-up calls. From this observation the competitive ratio of $O(\epsilon^{-1/2})$ follows. □

5.3 A position oblivious wake up algorithm

It seems natural and necessary that the positions of the nodes is used by the wake up algorithms. There are a lot of reasons why the positions might not be known. It is expensive and time-consuming to measure the coordinates and store it on each sensor node. Some sensor nodes might have no persistent memory and cannot store such information. And most important, since the communication range is dependent to the environment it is not clear what the position means in comparison to the unit disk range.

The following oblivious wake-up algorithms come into mind: Flooding, random walk, and epidemic algorithms. While flooding reaches all nodes, it is the worst with respect to energy. Random walks neither reduce the number of wake-up signals nor does it give any delivery guarantees. Epidemic algorithms appear to be the most reasonable solution. The question, however, is how to stop the epidemic wake-up of nodes. We use a push-based epidemic rumor spreading algorithm [12] which will be combined with a simple counter mechanism, which stops the wake-up if k other wake-ups have been received.

However, it is possible to exploit some position information given by the graph structure itself. The random k-covered wake-up of Figure 3 distinguish between covered and uncovered nodes. A node is covered, if it has received two wake-up signals. The algorithm starts with one node and picks in each round a node which has not been covered twice.

This idea generalizes to the random k-covered wake-up algorithm, where nodes continue to send until each node has been covered $k + 1$ times or has send a wake-up signal.

Algorithm 3: Random k-covered wake-up

```
Input graph G = (V, E), start node s ∈ V
forall the u ∈ V do
  | counter(u) ← 0
end
W ← {s}
Node s sends wake-up signal
forall the u ∈ V : {u, s} ∈ E do
  | counter(u) ← counter(u) + 1
end
while ∃u ∈ V \ W : 0 < counter(u) ≤ k do
  | Pick a random node w ∈ V \ W with 0 < counter(u) ≤ k
  | W ← W ∪ {w}
  | Node w sends wake-up signal
  | forall the u ∈ V : {u, w} ∈ E do
  |   | counter(u) ← counter(u) + 1
  | end
end
```

While we show in Section 5.3 that for $k = 1$ the algorithms performs very well, the generalization is necessary since there are situations where the algorithm fails from the start. In Fig. 7 such case is depicted. The node s wakes up four nodes and the near-by node u continues. This node wakes up the same set of nodes as s and so the algorithm stops, since all nodes are covered twice.

With the help of the position information this case could have been clearly avoided. But even without it one could increase k. However, the factor k increases the message complexity. So, a compromise between error rate and message complexity needs to be made.

For the dense case, one can show that the error rate can be reduced to any polynomial if the density is large enough and k is chosen to be logarithmic in n.

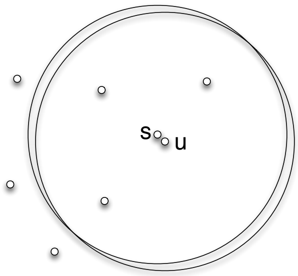

Fig. 7. A Counter-Example for the Random-k-Covered-Wake-Up

Theorem 3. *If n nodes are place randomly in a $m \times m$-grid with $m \leq \sqrt{\frac{n}{c \ln n}}$ for some constant c. Then the Random $O(\log n)$ covered wake-up algorithm computes a CDS with a competitive ratio of $O(\log n)$ high probability, i.e. $1 - n^{-c}$ for some $c \geq 1$.*

Proof. The expected number of nodes in each cell is $\frac{n}{m^2} \geq c \ln n$. Using Chernoff bounds it is possible to prove that this amount is in the range $[\frac{1}{2}\frac{n}{m^2}, 2\frac{n}{m^2}]$ with high probability. The expected number of nodes in the communication range is upper bounded by $5\pi\frac{n}{m^2}$ since the cell length is $\frac{1}{\sqrt{5}}$. Again Chernoff bound can provide with high probability that the number of nodes reachable in one hop is at most $10\pi\frac{n}{m^2}$ with high probability.

So, the probability that a node of a neighbor cell is activated is at least $\frac{1}{10\pi}$. Therefore, if each node is randomly activated until $10(c+1)\ln n$ wake up calls have been reached, this results in a probability of $1 - n^{-c}$ that each cell starts at least one wake-up call, when a neighbor cell has been activated before. □

6 Simulations

We have simulated the epidemic random k-covered algorithms to evaluate the efficency of the covered nodes and the wake-up signals needed from the source node to reach every single node in the area. The grid based flooding algorithm and the competitive algorithm (with $(1-\epsilon)$ unit disk graphs) have high constant factors involved such that they clearly cannot compete.

We randomly deployed varying number of nodes in a square area of 100m edge length. The middle node is woken up first and the wake-up communication range is limited to 10m based on real-world data. Fig. 8 and 9 show how the CDS is constructed in a network with 1,000 randomly deployed nodes using the random 1-covered, resp. 2-covered, wake-up algorithm.

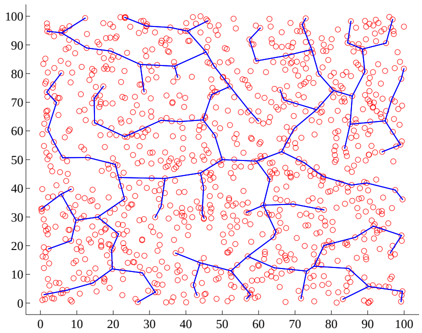

Fig. 8. Random-1-Covered-Wake-Up

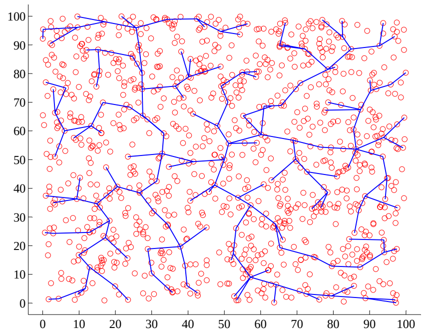

Fig. 9. Random-2-Covered-Wake-Up

Starting with a source node at a position (50, 50), the algorithm randomly picks the next node to be woken in order to cover the rest of the nodes that are found in the area. The nodes transmitting wake-up signal form a tree from the source node to each covered nodes. Only edges where a new wake-up call is initiated are depicted. In case of k = 1, a node is considered to be covered when it is covered by two wake-up signals or if it transmits one. When the algorithm considers $k > 1$ then a possible intersection in the tree can be formed as it appears in Fig. 9.

The quality of the algorithms are measured according to their coverage, i.e. the ratio of uncovered nodes after the algorithm has terminated, and the complexity, i.e. the number of wake-up calls sent. We have simulated this for the above parameters for increasing density. For this, we increase the number of nodes from 1 to $2,000$.

Fig. 10 shows that for $k = 1$ the ratio of uncovered nodes is relatively high compared to the set of nodes which can be reached, this percentage is displayed as the result of the flooding algorithm. Increasing k ameliorates this behavior. For high node density all algorithms reach nearly a full coverage. For $k = 1$ a coverage of 95% happens when 350 nodes are participating, for $k \geq 2$ this already happens for 250 nodes.

Fig. 11 indicates that the message complexity grows linearly with k and converges for increasing node density. Surprisingly, the complexity increases from $k = 1$ to $k = 2$ only by around 40%. So, $k = 2$ appears to be a good compromise between coverage and complexity.

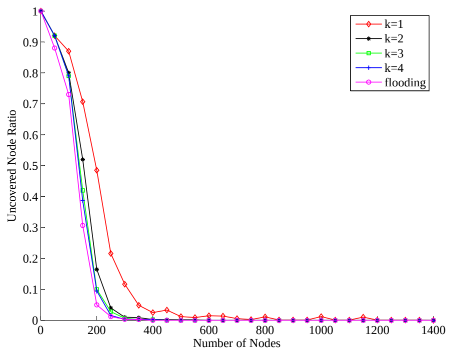

Fig. 10. Uncovered Node Ratio

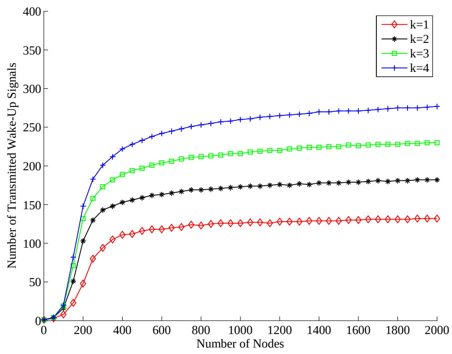

Fig. 11. Wake-Up Signal Transmitted for n nodes

7 Conclusions and Future Work

The improved efficiency of wake-up receivers down to the microwatt range implies a paradigm shift for wireless sensor networks. Now, no busy idling implemented as duty-cycling is necessary until the first sensor information or the first message arrives. However, when the nodes need to be woken up without any prior knowledge we face the wake-up connected dominating set problem, presented here.

For this problem, we provide theoretical and practical solutions. Our algorithms work for random placement and an adversarial setting, where we needed to reduce the power of the adversary, otherwise no efficient algorithms can be found. It turns out that for the random placement the position information is not necessary to find an efficient algorithm with a $O(\log n)$-competitive ratio. We have simulated a simplified variant of this and have seen that it reaches nearly all nodes with small number of wake-up calls.

This raises the hope that duty-cycling might soon be a technique of the past. However, with the available transceiver technology a wake-up call is orders of magnitudes more energy-consuming than standard operation. Taking this into account, it does not make sense to wake up the network from scratch every time a sensor reading appears. At this moment, it is more efficient to put the full network into sleep after some thousand communication cycles. So with the current hardware, a hybrid solution of wake-up calls and duty-cycling is the optimal solution.

Another available technique is the use of IDs for wake-up calls. It is possible to program sensor nodes to be woken up only on a special signal, which are given by a programmable ID. This may help protocols to build up a wake-up infrastructure, where the wake-up signals may trigger different nodes or paths. At this point, it is not clear how this feature can be used in future protocols and what can be achieved with it.

Since, the wake-up transceiver have only become available recently, we are in the process of implementing the given protocols and further research will show, how well these wake-up algorithms behave in real world.

Acknowledgements This work has been supported by the Ministry for Education and Research (Bundesministerium für Bildung und Forschung, BMBF) within the Research project AURIS (13N11746).
The authors would like to thank SmartExergy for providing the wake-up receiver's board.

References

1. Reuven Bar-Yehuda, Oded Goldreich, and Alon Itai. On the time-complexity of broadcast in radio networks: an exponential gap between determinism randomization. In *Proceedings of the sixth annual ACM Symposium on Principles of distributed computing*, PODC '87, pages 98–108, New York, NY, USA, 1987. ACM.
2. Azzedine Boukerche. *Algorithms and protocols for wireless sensor networks*, volume 62. Wiley-IEEE press, 2008.
3. Xiuzhen Cheng, Xiao Huang, Deying Li, Weili Wu, and Ding-Zhu Du. A polynomial-time approximation scheme for the minimum-connected dominating set in ad hoc wireless networks. *Networks*, 42(4):202–208, 2003.
4. Brent N. Clark, Charles J. Colbourn, and David S. Johnson. Unit disk graphs. *Discrete Math.*, 86(1-3):165–177, January 1991.
5. B. Das and V. Bharghavan. Routing in ad-hoc networks using minimum connected dominating sets. In *Communications, 1997. ICC '97 Montreal, Towards the Knowledge Millennium. 1997 IEEE International Conference on*, volume 1, pages 376–380 vol.1, 1997.
6. Ling Ding, Xiaofeng Gao, Weili Wu, Wonjun Lee, Xu Zhu, and Ding-Zhu Du. Distributed construction of connected dominating sets with minimum routing cost in wireless networks. In *Proceedings of the 2010 IEEE 30th International Conference on Distributed Computing Systems*, ICDCS '10, pages 448–457, Washington, DC, USA, 2010. IEEE Computer Society.
7. Hongwei Du, Qiang Ye, Jioafei Zhong, Yuexuan Wang, Wonjun Lee, and Haesun Park. Ptas for minimum connected dominating set with routing cost constraint in wireless sensor networks. In Weili Wu and Ovidiu Daescu, editors, *Combinatorial Optimization and Applications*, volume 6508 of *Lecture Notes in Computer Science*, pages 252–259. Springer Berlin Heidelberg, 2010.
8. Stephan Eidenbenz. Online Dominating Set and Variations on Restricted Graph Classes. Technical report, Institue of Theoretical Computer Science , ETH Zürich, 2002.
9. G.U. Gamm, M. Sippel, M. Kostic, and L.M. Reindl. Low power wake-up receiver for wireless sensor nodes. In *Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP), 2010 Sixth International Conference on*, pages 121–126, 2010.
10. S. Guha and S. Khuller. Approximation algorithms for connected dominating sets. *Algorithmica*, 20(4):374–387, April 1998.
11. P. Gupta and P.R. Kumar. Critical power for asymptotic connectivity. In *Decision and Control, 1998. Proceedings of the 37th IEEE Conference on*, volume 1, pages 1106–1110 vol.1, 1998.

12. R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized rumor spreading. In *In IEEE Symposium on Foundations of Computer Science*, pages 565–574, 2000.
13. D. Lichtenstein. Planar formulae and their uses. *SIAM Journal on Computing*, 11(2):329–343, 1982.
14. S. Rührup and C. Schindelhauer. Competitive time and traffic analysis of position-based routing using a cell structure. In *Parallel and Distributed Processing Symposium, 2005. Proceedings. 19th IEEE International*, pages 8 pp.–, 2005.
15. Peng-Jun Wan, K.M. Alzoubi, and O. Frieder. Distributed construction of connected dominating set in wireless ad hoc networks. In *INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 3, pages 1597–1604 vol.3, 2002.
16. Weili Wu, Hongwei Du, Xiaohua Jia, Yingshu Li, and Scott C.-H. Huang. Minimum connected dominating sets and maximal independent sets in unit disk graphs. *Theor. Comput. Sci.*, 352(1):1–7, March 2006.