## Title

Selecting forwarding neighbors in wireless ad hoc networks
Permalink
https://escholarship.org/uc/item/5cq136vm

## Journal

Mobile Networks \& Applications, 9(2)
ISSN
1383-469X

## Authors

Calinescu, G
Mandoiu, Ion I
Wan, P J
et al.

## Publication Date

2004-04-01

Peer reviewed

# Selecting Forwarding Neighbors in Wireless Ad Hoc Networks* 

Gruia Călinescu ${ }^{\dagger} \quad$ Ion I. Măndoiu ${ }^{\ddagger} \quad$ Peng-Jun Wan ${ }^{\S} \quad$ Alexander Z. Zelikovsky ${ }^{〔}$


#### Abstract

Broadcasting is a fundamental operation which is frequent in wireless ad hoc networks. A simple broadcasting mechanism, known as flooding, is to let every node retransmit the message to all its 1-hop neighbors when receiving the first copy of the message. Despite its simplicity, flooding is very inefficient and can result in high redundancy, contention, and collision. One approach to reducing the redundancy is to let each node forward the message only to a small subset of 1-hop neighbors that cover all of the node's 2-hop neighbors. In this paper we propose two practical heuristics for selecting the minimum number of forwarding neighbors: an $O(n \log n)$ time algorithm that selects at most 6 times more forwarding neighbors than the optimum, and an $O\left(n \log ^{2} n\right)$ time algorithm with an improved approximation ratio of 3 , where $n$ is the number of 1- and 2-hop neighbors. The best previously known algorithm, due to Bronnimann and Goodrich [2], guarantees $O(1)$ approximation in $O\left(n^{3} \log n\right)$ time.


Keywords: wireless ad hoc networks, broadcast, approximation algorithms, unit-disk graphs, disk cover.

## 1 Introduction

Wireless ad hoc networks can be flexibly and quickly deployed for many applications such as automated battlefield, search and rescue, and disaster relief. Unlike wired networks or cellular networks, no wired backbone infrastructure is installed in wireless ad hoc networks. A communication session is achieved either through a single-hop radio transmission if the communication parties are close enough, or through relaying by intermediate nodes otherwise. In this paper, we assume that all nodes in a wireless ad hoc network are distributed in a two-dimensional plane and have an equal maximum transmission range of one unit.
*GC's research was partially performed while visiting the Department of Combinatorics and Optimization at University of Waterloo, where supported by an NSERC research grant. The work of AZZ was partially supported by NSF Grant CCR-9988331, Award No. MM2-3018 of the Moldovan Research and Development Association (MRDA) and the U.S. Civilian Research \& Development Foundation for the Independent States
of the Former Soviet Union (CRDF), and by the State of Georgia's Yamacraw Initiative.
${ }^{\dagger}$ Computer Science Department, Illinois Institute of Technology, Chicago, IL 60616.
$\ddagger$ Department of Computer Science and Engineering, University of California at San Diego, La Jolla, CA 92093-0114
${ }^{\S}$ Computer Science Department, Illinois Institute of Technology, Chicago, IL 60616
『 Department of Computer Science, Georgia State University, University Plaza, Atlanta, Georgia 30303

Broadcasting is a fundamental networking operation in wireless ad hoc networks. It is widely and frequently performed in many networking tasks such as paging a particular host, sending an alarm signal, and finding a route to a particular host [1][9][17]. A simple broadcasting mechanism, known as flooding, is to let every node retransmit the message to all its 1-hop neighbors when receiving the first copy of the message. Despite its simplicity, flooding has a serious drawback, known as the broadcast storm [16]. First, because the radio propagation is omnidirectional and a physical location may be covered by the transmission ranges of several nodes, many retransmissions are redundant. Second, heavy contention could exist because retransmitting nodes are probably close to each other. Third, collisions are more likely to occur because the RTS/CTS dialogue is inapplicable and the timing of retransmissions is highly correlated.

The following simple technique was recently exploited in [13] (see also [12]) and [19] to reduce redundant retransmissions: By virtue of beaconing, each node maintains a local topology of its 2-hop neighborhood, and relays the message only to a small subset of 1-hop neighbors which cover (in terms of radio range) all nodes that are two hops away. The subset of 1-hop neighbors selected by each node is referred to as forwarding set [19] or multipoint relaying set [12][13]. In this paper we consider the problem of finding a forwarding set of minimum size.

Minimum Forwarding Set Problem: Given a source $A$, let $\mathcal{D}$ and $\mathcal{P}$ be the sets of 1-and 2-hop neighbors of $A$. Find a minimum-size subset $\mathcal{F}$ of $\mathcal{D}$ such that every node in $\mathcal{P}$ is within the coverage area of at least one node from $\mathcal{F}$.


Figure 1: An instance for which the size of the solution computed by the greedy algorithm, $\left\{g_{1}, \ldots, g_{\log k}\right\}$, is larger than the optimum solution, $\left\{o p t_{1}, o p t_{2}\right\}$, by a logarithmic factor.

### 1.1 Previous work

Laouiti et al. [13] (see also Jacquet et al. [12]) and Sinha et al. [19] considered the Minimum Forwarding Set problem assuming no knowledge of the geographic location of the nodes. In this case, the Minimum Forwarding Set problem is essentially the well-studied Set Cover problem. Not surprisingly, the heuristic proposed in [13] is a translation of Chvátal's greedy algorithm [4] for Set Cover, and thus guarantees an approximation factor of $O(\log m)$, where $m$ is the maximum neighborhood size. The greedy algorithm iteratively selects a 1-hop neighbor covering the maximum number of 2-hop neighbors not yet covered, and terminates when all 2-hop neighbors have been covered. The greedy algorithm does not take into account the geometric properties of the Minimum Forwarding Set problem, and in fact Figure 1 shows a family of instances for which the size of the solution found by the greedy algorithm is larger than the optimum by a logarithmic factor.

Under the assumption that the nodes in the wireless network are distributed in a two-dimensional plane and each node has unit transmission range, the topology of the network is modeled as a unit-disk graph [5]. In this graph, there is an edge between two nodes if and only if their distance is at most one. The Minimum Forwarding Set problem for a given source node $s$ asks for a minimum size set of 1-hop neighbors of $s$ dominating 2-hop neighbors of $s$ in the unit-disk graph. The related Dominating Set problem in unit-disk graphs [5] asks for a subset of nodes dominating (i.e., adjacent to) all the other nodes. The Dominating Set problem in unit-disk graphs is NP-hard [5] but admits a PTAS [11]. The Minimum Forwarding Set problem does not reduce to the Dominating Set problem in unit-disk graphs since dominators are restricted to the set of 1-hop neighbors.

The Minimum Forwarding Set problem is also related to the Unit-Disk Cover problem [10], which asks for the minimum number of unit disks covering a given set of points in the plane. The Unit-Disk Cover problem is also NP-hard [5] and admits a PTAS [10]. Since in the Unit-Disk Cover problem disk centers can be chosen arbitrarily in the plane, the algorithms for this problem do not apply to the Minimum Forwarding Set problem where disks must be centered at 1-hop neighbors only.

The Minimum Forwarding Set problem is a special case of the NP-Hard Disk Cover problem [2], which asks for a minimum size subset of a given set of disks covering a given set of points. The complexity of Minimum Forwarding Set problems is not known. A constant-ratio approximation algorithm for Disk Cover, and therefore also for Minimum Forwarding Set, was given by Bronnimann and Goodrich [2] However, their algorithm - which is a special case of a sophisticated algorithm for spaces with bounded VC-dimension - has impractical running-time and its proven approximation ratio is a very large constant.

### 1.2 Our contributions

- An exact $O\left(n \log ^{2} n\right)$ time, and a 2-approximation $O(n \log n)$ time algorithm for the special case of the Minimum Forwarding Set problem when all 2-hop neighbors are in the same quadrant with respect to the source node.
- A 6-approximation algorithm for the Minimum Forwarding Set problem running in $O(n \log n)$ time, where $n$ is the total number of 1- and 2-hop neighbors.
- A 3-approximation algorithm for the Minimum Forwarding Set problem running in $O\left(n \log ^{2} n\right)$ time.
- A constant-factor approximation for the Disk Cover problem with disks of the same radius, based on rounding the optimal solution of a linear programming relaxation.

A preliminary version of this paper [3] presents the same results, except that the exact algorithm for the quadrant case and the 3-approximation algorithm given in [3] run in $O\left(n^{2}\right)$ time.

The paper is organized as follows. In next section we reformulate the Minimum Forwarding Set problem in geometric terms, give a high-level algorithm based on decomposition into quadrants, and establish basic geometric properties of the partitioned sets of 1- and 2-hop neighbors. The next three sections deal with covering 2-hop neighbors in a quadrant: we first describe an $O(n \log n)$ 2-approximation algorithm (Section 3), then we give an $O\left(n^{2}\right)$ exact algorithm (Section 4) and finally describe details of data structures needed to obtain the $O\left(n \log ^{2} n\right)$ implementation of the exact algorithm (Section 5). In Section 6 we give an extension of our techniques to the Disk Cover problem of [2], and conclude in Section 7.

## 2 Partition based algorithm

Throughout this paper a unit disk, or just disk for short, refers to a closed disk of radius 1 . The boundary of a region $R$ of the Euclidean plane is denoted by $\partial R$, e.g., the boundary circle of a disk $D$ is denoted by $\partial D$. Under the assumption that each network node has unit transmission range, we reformulate the Minimum Forwarding Set problem as follows.

1-Hop Disk Cover Problem: Given a unit-disk $A$, a set $\mathcal{D}$ of unit disks centered inside $A$, and a set of points $\mathcal{P}$ outside $A$ such that $\mathcal{P} \subseteq \cup\{D \in \mathcal{D}\}$, find a minimum-size subset $\mathcal{F}$ of $\mathcal{D}$ such that $\mathcal{P} \subseteq \cup\{D \in \mathcal{F}\}$.

Our high-level algorithm (Algorithm 1) partitions the points of $\mathcal{P}$ according to the four quadrants defined by two orthogonal lines through the center of $A$, and then independently solves the 1-Hop Disk Cover problem for each quadrant. The union of these four disk covers is then a disk cover for all the points in $\mathcal{P}$. As usual, the approximation ratio of an algorithm $\mathcal{A}$ for a minimization problem $\Pi$ is the supremum, over all instances of $\Pi$, of the ratio between

## Algorithm 1: 1-Hop Disk Cover

Input: Unit-disk $A$, set of unit disks $\mathcal{D}$ centered inside $A$, set of points $\mathcal{P}$ outside $A$ such that $\mathcal{P} \subseteq \cup\{D \in \mathcal{D}\}$
Output: Subset $\mathcal{F} \subseteq \mathcal{D}$ such that $\mathcal{P} \subseteq \cup\{D \in \mathcal{F}\}$

1. Partition the exterior of $A$ into four quadrants $Q_{1}-Q_{4}$ by two orthogonal lines, not containing points in $\mathcal{P}$, through the center of $A$ (see Figure 2).
2. For $q=1, \ldots, 4$, compute a disk cover, $\mathcal{F}_{q}$, for the points in $\mathcal{P} \cap Q_{q}$.
3. Output $\mathcal{F}=\mathcal{F}_{1} \cup \mathcal{F}_{2} \cup \mathcal{F}_{3} \cup \mathcal{F}_{4}$.
the output value of $\mathcal{A}$ and the optimal value. The following theorem relates the approximation ratio of Algorithm 1 to the approximation ratio that can be guaranteed for the 1-Hop Disk Cover restricted to points in a single quadrant.

Theorem 1 If disk covers $\mathcal{F}_{q}$ computed in Step 2 are within a factor of $\alpha$ of optimum, then Algorithm 1 has an approximation ratio of at most $3 \alpha$ for the 1-Hop Disk Cover problem.

Proof. Let $O P T$ be the optimal set of disks, and denote by $O P T_{q}, q=1,2,3,4$, the subset of disks in $O P T$ having centers in the $q^{\text {th }}$ sector of disk $A$. The key observation is that points in quadrant $Q_{q}$ cannot be covered by disks in $O P T_{q+2(\bmod 4)}$. Therefore, points in $\mathcal{P} \cap Q_{1}$ must be covered by disks in $O P T_{4} \cup O P T_{1} \cup O P T_{2}$, and thus, by the assumption that $\mathcal{F}_{q}$ 's are within a factor of $\alpha$ of the respective optimum solutions,

$$
\left|\mathcal{F}_{1}\right| \leq \alpha\left(\left|O P T_{4}\right|+\left|O P T_{1}\right|+\left|O P T_{2}\right|\right)
$$

Similarly,

$$
\begin{aligned}
\left|\mathcal{F}_{2}\right| & \leq \alpha\left(\left|O P T_{1}\right|+\left|O P T_{2}\right|+\left|O P T_{3}\right|\right) \\
\left|\mathcal{F}_{3}\right| & \leq \alpha\left(\left|O P T_{2}\right|+\left|O P T_{3}\right|+\left|O P T_{4}\right|\right) \\
\left|\mathcal{F}_{4}\right| & \leq \alpha\left(\left|O P T_{3}\right|+\left|O P T_{4}\right|+\left|O P T_{1}\right|\right)
\end{aligned}
$$

Thus, the output of the algorithm has size

$$
\begin{aligned}
& \left|\mathcal{F}_{1}\right|+\left|\mathcal{F}_{2}\right|+\left|\mathcal{F}_{3}\right|+\left|\mathcal{F}_{4}\right| \\
\leq & 3 \alpha\left(\left|O P T_{1}\right|+\left|O P T_{2}\right|+\left|O P T_{3}\right|+\left|O P T_{4}\right|\right) \\
= & 3 \alpha|O P T|
\end{aligned}
$$



Figure 2: The four quadrants in Algorithm 1.

We will show that $\alpha=2$ can be achieved in $O(n \log n)$ time (see Section 3), and $\alpha=1$ can be achieved in $O\left(n \log ^{2} n\right)$ time (see Section 4). Hence, Algorithm 1 achieves an approximation factor of 6 , respectively 3, within the same time bounds. It is natural to ask if these approximation ratios can be improved by partitioning the set of points according to $k<4$ equal sectors defined by half-lines starting at the center of $A$. The proof of Theorem 1 can be generalized to show that partitioning into $k$ sectors gives an approximation ratio of $(\lceil k / 2\rceil+1) \alpha$ for the 1-Hop Disk Cover problem if the disk cover for each sector is approximated within a factor of $\alpha$. Thus, using decomposition into 3 equal sectors does not lead to an approximation ratio better than that obtained by decomposition into quadrants. Improvements using decomposition into 2 equal sectors are possible provided that we can find an algorithm for covering the points in a $180^{\circ}$ sector with an approximation ratio of less than $3 / 2$. The ideas used in Section 4 to solve exactly the problem for a quadrant do not extend to $180^{\circ}$ sectors, since these lack the second of the essential topological properties established for the quadrants in the following lemma.

Lemma 2 Let $Q$ be an exterior quadrant of $A, J=\partial Q$ be its border, and $\mathcal{D}$ be a set of disks intersecting the interior of $Q$. Then:
(a) For any disk $D \in \mathcal{D},|\partial D \cap J|=2$.
(b) For any two disks $D, D^{\prime} \in \mathcal{D},\left|\partial D \cap \partial D^{\prime} \cap Q\right| \leq 1$.
(c) No two disks in $\mathcal{D}$ are tangent in $Q$.

Proof. Without loss of generality, we may assume that the unit-disk $A$ is centered at the origin and that $Q$ is defined by the positive $x$ - and $y$-axes. Then, the boundary of quadrant $Q, J$, consists of the two half-lines from $(0,1)$ to $(0, \infty)$, and from $(1,0)$ to $(1, \infty)$, together with a quarter-circle of $\partial A$. Let $a, b, c$ be the points with coordinates $(0,0),(1,0)$,


Figure 3: The extreme configuration in the proof of Lemma 2(b).
and $(0,1)$, respectively. We will use $\widehat{b c}$ to denote the quarter-circle of $A$ enclosed in $J$.
(a) Since every $D \in \mathcal{D}$ has non-empty intersection with the interior of $Q$, every circle $\partial D$ has at least two intersection points with $J$. The closed simple Jordan curve $\partial D$ and the infinite simple Jordan curve $J$ must intersect an even number of times (unless they are tangent, but this cannot happen), and thus cannot intersect three times. To complete the proof of part (a) we need to show that $\partial D$ does not intersect $J$ four or more times.

Let $d$ denote the center of disk $D$. Then $0<|d a| \leq 1$, since $d$ is inside $A$. Note that $\partial D$ can intersect the $x$-axis in at most two points, of which only one can have $x$-coordinate bigger than 1 . Similarly, $D$ can intersect the $y$-axis in at most two points, of which only one can have $y$-coordinate bigger than 1. Furthermore, $D$ intersects $\widehat{b c}$ at most once. Indeed, when two unit-circles with centers within distance of at most 1 intersect, the two intersection points are at least $2 \pi / 3$ apart on each of the circles, and hence a quarter-circle may contain only one of them.
(b) Assume, for a contradiction, that $D$ and $D^{\prime}$ are two distinct disks in $\mathcal{D}$ that intersect at points $h$ and $l$, with both $h$ and $l$ in $Q \cup J$. Let $d$ and $d^{\prime}$ be the centers of $D$ and $D^{\prime}$, respectively. We will change the configuration a bit, to obtain a more extreme case. First, translate $d, d^{\prime}, h$, and $l$ to the right until $d$ or $d^{\prime}$ hits $\widehat{b c}$, and assume, by symmetry, that $d$ is on $\widehat{b c}$. We still have $h, l \in R \cup J$. Assume also that $h$ is to the right of $l$. Now start rotating the rhombus $h d^{\prime} l d$ clockwise around $d$ until $h$ hits either the $x$-axis or $\overparen{b c}$, whichever happens first (see Figure 3). This procedure also keeps $d^{\prime}$ inside the unit-disk $A$ and $l$ in $Q$. Let $m$ be the point where the line $h d$ intersects the $y$-axis. As $\left|d^{\prime} h\right|=|d h|=1, d^{\prime}$ must lie on the same side of line $h m$ as $a$. As the angle $\widehat{h a m} \leq \frac{\pi}{2}, m$ must be within the diameter of the unit-disk centered at $d$ that contains $h$. Therefore $|d m| \leq 1$. Thus $l$, which is in $Q$, must be outside the triangle $a h m$, and consequently $d^{\prime}$ must be on the other side of line $h m$ than $a$, which is a contradiction.
(c) Let $D$ and $D^{\prime}$ be two disks from $\mathcal{D}$. Then $\partial D$ and $\partial D^{\prime}$ cannot be tangent from the interior since they have the same radius. If $\partial D$ and $\partial D$ are tangent from the exterior, then the distance between their centers is 2 , and the common


Figure 4: The skyline of a set of disks in a quadrant.
point can only be the origin $a$, which is not in $Q$.

## 3 Fast geometric disk covering in a quadrant

In this section we give a fast 2-approximation algorithm for the 1-Hop Disk Cover problem with all points of $\mathcal{P}$ coming from an exterior quadrant $Q$ of the unit disk $A$.

The skyline $S=\left(x_{0}, x_{1}, \ldots, x_{k}\right)$ of $\mathcal{D}$ is the upper envelope of $Q \cap(\cup\{D \in \mathcal{D}\} \cup A)$ (see Figure 4). The skyline consists of arcs $\widehat{x_{i-1} x_{i}}$ on the border of disks $D_{i} \in \mathcal{D} \cup\{A\}, i=1, \ldots, k$, such that $x_{0} \in \partial Q \cap \partial D_{1}$, $x_{i} \in \partial D_{i-1} \cap \partial D_{i}(i=1, \ldots, k-1)$, and $x_{k} \in \partial D_{k} \cap \partial Q$. The algorithm (Algorithm 2) starts by computing the skyline $S$ with $x_{i}$ 's numbered in counter-clockwise order, i.e., with polar coordinates $\left(\rho_{i}, r_{i}\right)$ of points $x_{i}$ satisfying $\rho_{0} \leq \rho_{1} \leq \rho_{2} \leq \cdots \leq \rho_{k}$. As established in Lemma 8 below, the skyline disks $D_{i}$ covering a point $p \in \mathcal{P}$ form an interval in the sequence $D_{1}, \ldots, D_{k}$. The algorithm computes these intervals for each point of $\mathcal{P}$, then outputs a minimum size set $\mathcal{F}$ of skyline disks hitting all intervals, i.e., containing at least one disk $D_{i}$ from each interval. Computing the minimum size hitting set (Step 3) is done using a simple greedy algorithm, similar, e.g., to the algorithm in [18] for finding a minimum weight dominating set in an interval graph given. Clearly, the hitting set $\mathcal{F}$ computed by Algorithm 2 is a disk cover for the points in $\mathcal{P}$. Furthermore, we have:

Theorem 3 Algorithm 2 runs in $O(n \log n)$ time, and has an approximation ratio of 2 for the 1-Hop Disk Cover problem in a quadrant.

Theorems 1 and 3 immediately give:

## Algorithm 2: Geometric 1-Hop Disk Covering in a quadrant

Input: Unit-disk $A$, set of unit disks $\mathcal{D}$ centered inside $A$, set of points $\mathcal{P}$ in the exterior quadrant $Q$ of $A$ such that $\mathcal{P} \subseteq \cup\{D \in \mathcal{D}\}$

Output: A subset $\mathcal{F} \subseteq \mathcal{D}$ such that $\mathcal{P} \subseteq \cup\{D \in \mathcal{F}\}$

1. Find the skyline $S=\left(x_{0}, x_{1}, x_{2}, \ldots, x_{k}\right)$ of $\mathcal{D}$, where the polar coordinates of $x_{i}$ are $\left(\rho_{i}, r_{i}\right)$ and $\rho_{0} \leq \rho_{1} \leq \rho_{2} \leq$ $\cdots \leq \rho_{k}$. Let $D_{i}$ be the disk containing arc $\overparen{x_{i-1} x_{i}}$.
2. For each $p \in \mathcal{P}$ with polar coordinates $(\rho, r)$, find the interval $\left[D_{f i r s t(p)}, D_{\text {last }(p)}\right]$ of skyline disks $D_{i}$ that cover $p$, via three binary searches:
(a) find $i \in\{1, \ldots, k\}$, such that $\rho \in\left[\rho_{i-1}, \rho_{i}\right]$
(b) $\operatorname{first}(p) \leftarrow \min \left\{j: 1 \leq j \leq i, p \in D_{j}\right\}$
(c) $\operatorname{last}(p) \leftarrow \max \left\{j: i \leq j \leq k, p \in D_{j}\right\}$
3. Find the minimum set $\mathcal{F}$ of skyline disks hitting each interval $\left[D_{\text {first }(p)}, D_{\text {last }(p)}\right], p \in \mathcal{P}$, using the following Interval Hitting Algorithm:
(a) Sort the set of all intervals $I=\left\{\left[D_{\text {first }(p)}, D_{\text {last }(p)}\right], p \in \mathcal{P}\right\}$ in ascending order of their right end, $D_{\text {last }(p)}$
(b) $\mathcal{F} \leftarrow \emptyset$
(c) While $I \neq \emptyset$ do

Add to $\mathcal{F}$ the right end disk $D_{\text {last }(p)}$ of the first interval
Remove from $I$ all intervals hit by $D_{\text {last }(p)}$
4. Output $\mathcal{F}$

Corollary 4 Combined with Algorithm 2, Algorithm 1 runs in $O(n \log n)$ time and has an approximation ratio of 6 for the Minimum Forwarding Set problem.

The rest of the section is devoted to the proof of Theorem 3.

Lemma 5 A point $q \in Q$ belongs to a disk $D \in \mathcal{D}$ if and only if the half-line $L$ from the center a of $A$ through a point $q$ intersects $\partial D \cap Q$ at a point $q^{\prime}$ such that $q$ belongs to the segment $\left[a, q^{\prime}\right]$.

Proof. Every disk $D \in \mathcal{D}$ contains $a$. Thus, the segment $\left[a, q^{\prime}\right]$ is fully contained in $D$, and every point of $L$ outside this segment is in the exterior of $D$.

Lemma 6 If point $p \in \mathcal{P}$ has polar coordinates $(\rho, r)$ such that $\rho \in\left[\rho_{i-1}, \rho_{i}\right]$, then $p \in D_{i}$.


Figure 5: The skyline of $\left\{D_{1}, D_{2}, D_{3}\right\}$ in Lemma 7.

Proof. Follows immediately from Lemma 5.

Lemma 7 Let $D_{1}, D_{2}, D_{3}$ be three disks of $\mathcal{D}$ appearing in this order in the skyline of $\left\{D_{1}, D_{2}, D_{3}\right\}$. Then $D_{1} \cap$ $D_{3} \cap Q \subseteq D_{2} \cap Q$.

Proof. Assume that $D_{1} \cap D_{3} \cap Q \neq \emptyset$, and let $S^{\prime}=\left(y_{0}, y_{1}, y_{2}, y_{3}\right)$ be the skyline of $\left\{D_{1}, D_{2}, D_{3}\right\}$ (see Figure 5). Since $y_{1}=\partial D_{1} \cap \partial D_{2} \cap Q, y_{2}=\partial D_{2} \cap \partial D_{3} \cap Q$, and $y_{1}, y_{2} \notin D_{1} \cap D_{3}$, Lemma 2(b) implies that $\partial D_{2} \cap \partial\left(D_{1} \cap\right.$ $\left.D_{3} \cap Q\right)=\emptyset$.

To complete the proof, it suffices to show that $D_{2}$ contains some point of $D_{1} \cap D_{3} \cap Q$. Let $x=\partial D_{1} \cap \partial D_{3} \cap Q$, and let $L$ be the half-line from $a$ through $x$. Since $a \in D_{1}, L$ intersects $\partial D_{1}$ exactly once, at $x$. Thus, $L$ does not intersect the arc $\overparen{y_{0} y_{1}}$ of the skyline. Similarly, $L$ does not intersect $\overparen{y_{2} y_{3}}$. It follows that $L$ intersects $\overparen{y_{1} y_{2}}$, and, by Lemma 5, $x \in D_{2}$.

The following is a straightforward corollary of Lemma 7:

Lemma 8 For every $p \in \mathcal{P}$, the skyline disks $D_{i}$ covering $p$ form an interval $\left[D_{\text {first }(p)}, D_{\text {last }(p)}\right]$ in the sequence $D_{1}, \ldots, D_{k}$.

Lemma 9 The optimum cover of $\mathcal{P}$ with disks from the set $\left\{D_{1}, \ldots, D_{k}\right\}$ of skyline disks contains at most 2 times more disks than the optimum cover of $\mathcal{P}$ with disks from $\mathcal{D}$.

Proof. It suffices to prove that, for every $D \in \mathcal{D}, D \cap Q$ is covered by at most two skyline disks. Furthermore, since Lemma 5 implies that any set of disks covering $\partial D \cap Q$ fully covers $D \cap Q$, we only need to show that $\partial D \cap Q$ is covered by at most two skyline disks.


Figure 6: A tight example for the approximation ratio of Algorithm 2.

Let $d_{1}$ and $d_{2}$ be the two points of intersection of $\partial D$ with the boundary of the central disk $A$. By Lemma 2(b), any skyline disk $D_{i}$ intersecting $D \cap Q$ contains at least one of the points $d_{1}$ and $d_{2}$. The key observation is that, for any two skyline disks $D_{i}$ and $D_{j}$ both containing $d_{1}$ (or both containing $d_{2}$ ), the arc $\partial D \cap \partial D_{i} \cap Q$ is contained in the arc $\partial D \cap \partial D_{j} \cap Q$ or vice versa. Therefore the minimal set of skyline disks covering $\partial D \cap Q$ has at most two disks.

Proof of Theorem 3. The approximation ratio of Algorithm 2 follows from Lemma 9. Step 1 of the algorithm can be implemented in $O(n \log n)$ time using, e.g., an adaptation of the divide-and-conquer algorithm in [15] for computing the Manhattan skyline. The binary searches in Step 2 also take $O(n \log n)$ time.

Finally, the Interval Hitting Algorithm can be implemented in $O(n \log n)$. Indeed, an interval is hit by the disk $D_{\text {last }(p)}$ if its left covering disk is before $D_{\text {last }(p)}$ in the skyline $S$. Therefore, by traversing all intervals sorted in ascending order of their left ends $D_{\text {first }(p)}$, we can delete each hit interval in constant time per interval.

Remark. The approximation ratio of 2 in Theorem 3 is tight: Figure 6 gives an instance when the optimum disk cover consisting of skyline disks has size 2 , while there is a single (non-skyline) disk covering the two points of $\mathcal{P}$.

## 4 Exact combinatorial disk covering in a quadrant

In this section we present an $O\left(n^{2}\right)$ exact algorithm for the 1 -Hop Disk Cover problem with all points of $\mathcal{P}$ coming from an exterior quadrant $Q$ of the unit disk $A$. In the next section we describe a faster $O\left(n \log ^{2} n\right)$ implementation of this algorithm based on efficient data structures. In [3], a different $O\left(n^{2}\right)$ algorithm is presented for the 1-Hop Disk Cover problem with all points of $\mathcal{P}$ coming from $Q$. That algorithm is somehow easier to implement, but has a more involved correctness proof. Furthermore, the algorithm in [3] does not appear to have an implementation faster than
$O\left(n^{2}\right)$.
The new exact algorithm (see Algorithm 3) starts by sorting and renumbering all disks $D_{i}$ with respect to the intersection points of $\partial D_{i}$ with $\partial Q$ (Step 1). The first and the last disk in this order covering each point are determined. Finally, a modified version of the Interval Hitting Algorithm (see Step 3 of Algorithm 2 ) finds the minimum disk cover.

## Algorithm 3: Combinatorial 1-Hop Disk Covering in a quadrant

Input: Unit-disk $A$, set of unit disks $\mathcal{D}$ centered inside $A$, set of points $\mathcal{P}$ in the exterior quadrant $Q$ of $A$ such that $\mathcal{P} \subseteq \cup\{D \in \mathcal{D}\}$

Output: Minimum size subset $\mathcal{F} \subseteq \mathcal{D}$ such that $\mathcal{P} \subseteq \cup\{D \in \mathcal{F}\}$

1. For each $D_{i} \in \mathcal{D}$ find $l_{i}$ and $r_{i}$, the two points of intersection between the boundaries $\partial D_{i}$ with $\partial Q$. We assume that $l_{j}<r_{j}$ in a fixed orientation of $\partial Q$. Renumber the disks in $\mathcal{D}$ such that either $l_{i}<l_{i+1}$ or $l_{i}=l_{i+1}$ and $r_{i}<r_{i+1}$ for every $i=1, \ldots, m-1$. Further, $D<D^{\prime}$ denotes that the disk $D$ has smaller index than the disk $D^{\prime}$
2. For every point $p \in \mathcal{P}$, compute $D_{f}^{p}$ and $D_{q}^{p}$, the first and last disks containing $p$.
3. Set $\overline{\mathcal{P}} \leftarrow \mathcal{P}$ and $\mathcal{F} \leftarrow \emptyset$
4. While $\overline{\mathcal{P}} \neq \emptyset$
a) Find $p \in \overline{\mathcal{P}}$ with the minimum last disk $D=D_{q}^{p}$
b) While there exists $p^{\prime} \in \overline{\mathcal{P}}$ with $D_{f}^{p^{\prime}}<D$ and $p^{\prime} \notin D$

Replace $D$ by the the last disk before $D$ containing $p$
c) Set $\mathcal{F} \leftarrow \mathcal{F} \cup\{D\}$
d) Remove from $\overline{\mathcal{P}}$ all the points $p^{\prime}$ with $D_{f}^{p^{\prime}} \leq D$
5. Output $\mathcal{F}$

We start the proof of correctness of the Algorithm 3 with the following definition and a crucial topological lemma. For $i<j<s$, we say that disks $D_{i}$ and $D_{s}$ supercover disk $D_{j}$ if there is a point $p$ in $\mathcal{P}$ such that $p \notin D_{j}$ and $p \in D_{i} \cap D_{s}$.

Lemma 10 There is an optimum solution which uses no supercovered disk.

Proof. Assume, for a contradiction, that $O P T$ is an optimum solution such that the area of $\cup_{D \in O P T}(Q \cap D)$ is maximum, and that there is a disk $D_{j} \in O P T$ which is supercovered, by say $D_{i}$ and $D_{s}$, where $i<j<s$. This means that there is a point $p \in \mathcal{P} \cap D_{i} \cap D_{s}$ which is not in $D_{j}$. There is a disk $D_{k} \in O P T$ such that $p \in D_{k}$. In the following we consider only the case when $k>j$, the other case being symmetric.


Figure 7: The existence of a nonempty region in $Q \cap\left(\left(D_{i} \cup D_{k}\right) \backslash D_{j}\right)$ implies that $\left(D_{j} \cap Q\right) \subseteq\left(D_{i} \cap D_{k}\right)$.

Lemma 2 implies that $\left(D_{j} \cap Q\right) \subseteq\left(D_{i} \cup D_{k}\right)$. See Figure 7 for an illustration.
Consider the walk on $\partial D_{i}$ starting from $l_{i}$ towards $r_{i}$. Let $p_{1}$ be the last point in this walk ( $p_{1}$ might not be in $\mathcal{P}$ ) which is inside a disk $D_{g}$ with $g<i$ and $D_{g} \in O P T$. Let $p_{2}$ be the first point in this walk ( $p_{2}$ might not be in $\mathcal{P}$ ) which is inside a disk $D_{r}$ with $i<r$ and $D_{r} \in O P T$. We have two cases.

If $p_{1}$ comes before $p_{2}$ on the walk on $\partial D_{i}$, then $O P T \backslash\left\{D_{j}\right\} \cup\left\{D_{i}\right\}$ is a solution of the same size as $O P T$, but of bigger total area, as it includes the arc of $\partial D_{i}$ in between $p_{1}$ and $p_{2}$. This contradicts the assumption that $O P T$ is an optimum solution with maximum area.

Now assume $p_{2}$ comes no later than $p_{1}$ on the walk on $\partial D_{i}$. Let $D_{g}<D_{i}$ and $D_{r}>D_{i}$ be the disks in OPT such that $p_{1}=\partial D_{i} \cap \partial D_{g}$ and $p_{2}=\partial D_{i} \cap \partial D_{r}$. The fact that there is a point in $\left(Q \cap D_{i} \cap D_{k}\right) \backslash D_{j}$ and Lemma 2 imply that $\partial D_{i} \cap \partial D_{k}$ comes before $\partial D_{i} \cap \partial D_{j}$ on the walk, and therefore $r \neq j$. As $\partial D_{i} \cap \partial D_{r}$ comes no later than $\partial D_{i} \cap \partial D_{g}$, Lemma 2 implies that $\left(D_{i} \cap Q\right) \subseteq\left(D_{g} \cup D_{r}\right)$. Then $\left(D_{j} \cap Q\right) \subseteq\left(D_{g} \cup D_{r} \cup D_{k}\right)$, implying that $O P T \backslash\left\{D_{j}\right\}$ is a solution. This contradicts the fact that $O P T$ is an optimum solution, and completes the proof.

From now on, we assume that $O P T$ is a fixed optimum solution which contains no supercovered disk.
Based on the previous lemma, an $O\left(n^{4}\right)$ exact algorithm for disk cover is immediate: eliminate all the supercovered disks. After that, for every point $p \in \mathcal{P}$, the set of disks which cover $p$ forms an interval. Then the Interval Hitting Algorithm finds the optimum solution.

To prove that the runtime of the Algorithm 3 is $O\left(n^{2}\right)$ we need to show that the total time spent in the inner while loop (Step 4.b) is $O(n)$. This fact follows from the following lemma.

Lemma 11 A disk cannot appear as $D$ inside the inner while loop (Step 4.b) for two different points.

Proof. Assume, for a contradiction, that two points, $p_{1}$ and $p_{2}$, selected in this order in Step 4.a, use the $D$ in Step
4.b. Since $p_{2}$ is not removed from $\overline{\mathcal{P}}$ while processing $p_{1}$, we have $D_{f}^{p_{1}}<D_{f}^{p_{2}}$. There is a disk $\bar{D}$ which is not supercovered and contains $p_{2}$. We have $\bar{D} \leq D$, since otherwise $\bar{D}$ would be selected by $p_{2}$ in Step 4.c, and $D$ would not be processed by $p_{2}$. Thus $D_{f}^{p_{1}}<D_{f}^{p_{2}} \leq \bar{D} \leq D \leq D_{q}^{p_{1}}$, and since $\bar{D}$ is not supercovered, we deduce that $\bar{D}$ contains $p_{1}$.

Let $D_{1}$ be the disk selected in Step 4.c while processing $p_{1}$. We cannot have $D_{1}<\bar{D}$, since otherwise, while processing $p_{1}, \bar{D}$ would have been considered as $D^{\prime}$ before $D_{1}$, and since $\bar{D}$ is not supercovered, $\bar{D}$ would be selected by $p_{1}$ in Step 4.c. So we have $D_{f}^{p_{2}} \leq \bar{D} \leq D_{1} \leq D \leq D_{q}^{p_{2}}$. As $D_{1}$ is selected while processing $p_{1}$, the check in the if statement ensures that $D_{1}$ contains $p_{2}$. So $p_{2}$ is eliminated from $\overline{\mathcal{P}}$ in Step 4 .d while processing $p_{1}$, a contradiction.

Lemma 12 The runtime of Algorithm 3 is $O\left(n^{2}\right)$.

Proof. Using straightforward for loops, every inner step, except for the inner while loop (Step 4.b), of the algorithm takes time $O(n)$. The replacement of a disk by a previous disk (inside the inner while loop) takes time $O(n)$ for a given point $p$. Checking the condition in the inner while loop (Step 4.b) takes $O(n)$ time, and Lemma 11 ensures that such a check is done at most once per disk.

Before proving the correctness of Algorithm 3, we give the intuition which motivates its steps. $\overline{\mathcal{P}}$ is the current set of uncovered points. In each iteration of the outer while loop in Step 4, as in the Interval Hitting Algorithm, we pick the point with the first last-covering disk. Then we select a disk which includes this point, as close as possible to the last-covering disk (in order to include a larger number of points) subject to excluding supercovered disks. Note that the inner while loop in Step 4.b simply jumps over disks which are supercovered.

Theorem 13 Algorithm 3 gives an optimal solution for the 1-hop Disk Cover problem in a quadrant.

Proof. First we show that every point is covered. We look at the current situation at the beginning of each execution of the outer while loop in Step 4. By the time Step 4.b (the inner while loop) is finished, $p$ is removed from $\overline{\mathcal{P}}$. Indeed, in $O P T$, there is a disk $D_{o p t}^{p}$ which is not supercovered and which covers $p$, and $D_{f}^{p} \leq D_{o p t}^{p} \leq D_{q}^{p}$. Every disk ignored by the algorithm in Step 4.b is supercovered, since by the choice of $p$ in Step 4.a, we have $D_{f}^{p^{\prime}} \leq D^{\prime} \leq D_{q}^{p} \leq D_{q}^{p^{\prime}}$. Therefore the inner while loop will stop at $D_{o p t}^{p}$, if not earlier.

To prove optimality, let $p_{1}, p_{2}, \ldots, p_{s}$ be the set of points selected by the algorithm in Step 4.a. We claim that no disk $D$ of $O P T$ can include $p_{i}$ and $p_{j}$, where $i<j$. Let $D_{i}$ be the disk selected to cover $p_{i}$ in Step 4.c. Note that while processing $p_{i}$, the algorithm cannot ignore $D$ in Step 4.b, as $D$ is not supercovered, and therefore $D \leq D_{i}$. Assuming by contradiction that $D$ contains $p_{j}$, we obtain that $D_{f}^{p_{j}} \leq D$. But then $D_{f}^{p_{j}} \leq D_{i}$, and $p_{j}$ should have been removed from $\overline{\mathcal{P}}$ in Step 4.d.

It follows that $|O P T| \geq s$, and the theorem follows from the fact that the algorithm also selects $s$ disks.
Theorems 1, 13, and Lemma 12 imply:

Theorem 14 Combined with Algorithm 3, Algorithm 1 runs in $O\left(n^{2}\right)$ time and has an approximation ratio of 3 for the Minimum Forwarding Set problem.

## 5 An $O\left(n \log ^{2} n\right)$ implementation of the combinatorial disk covering algorithm

In this section we describe an enhanced data structure based on Voronoi diagrams which allows to implement Algorithm 3 in time $O\left(n \log ^{2} n\right)$

Step 1 can obviously be executed in $O(n \log n)$. Next we describe the data structure which we use for Step 2. Construct a balanced binary search tree $T$, with the centers of the individual disks as leaves, in the sorted order obtained in Step 1. For a node $v \in T$, denote by $A_{v}$ the set of the centers in the subtree rooted at $v$. For every node $v$ in $T$, construct the Voronoi diagram of $A_{v}$, and preprocess the diagram for membership queries. A membership query returns, for a give point $p$, the face of the Voronoi diagram where $p$ lies. The preprocessing time for constructing the Voronoi diagrams is $O(n \log n)$ per level in $T$ (see, for example, Chapter 20.2 of [8]), for a total of $O\left(n \log ^{2} n\right)$. The space requirement is $O(n)$ per level of $T$, for a total of $O(n \log n)$. Preprocessing a node $v$ for membership queries also takes time $O\left(\left|A_{v}\right| \log \left|A_{v}\right|\right)$ and space $O\left(\left|A_{v}\right|\right)$, and a query can be answered in time $O\left(\log \left|A_{v}\right|\right)$ (see, for example, Chapter 30.3 of [8]). The total time is $O\left(n \log ^{2} n\right)$, and the total space is $O(n \log n)$.

With this data structure, given any node $v$ in $T$, and a point $p$, one can find whether a disk in the subtree rooted at $v$ covers $p$ in time $O(\log n)$, by finding out in which cell of the Voronoi diagram of $A_{v} p$ lies, and computing the distance from $p$ to the center of that cell. Using this observation, a binary search in $T$ can find for a point $p$ the disks $D_{f}^{p}$ and $D_{q}^{p}$ in time $O\left(\log ^{2} n\right)$.

Before proceeding to Step 3, in time $O(n \log n)$, sort the points of $\mathcal{P}$ with respect to $D_{f}^{p}$ and put them in a list $L_{1}$. To represent $\overline{\mathcal{P}}$, we only use the position in $L_{1}$ of the first point of $\overline{\mathcal{P}}$; this is enough since the points removed from $\overline{\mathcal{P}}$ in Step 4.d are consecutive in $L_{1}$, and include the first uncovered point of $L_{1}$. This representation allows the implementation of Step 4.d in total time $O(n)$ : to remove points from $\overline{\mathcal{P}}$, simply move forward in $L_{1}$, checking at every step if $D_{f}^{p^{\prime}} \leq D$.

Also, in time $O(n \log n)$, sort the points of $\mathcal{P}$ with respect to $D_{q}^{p}$ and put them in a separate list $L_{2}$. This allows us to implement Step 4.a to run in total time $O(n)$, by keeping track of where point $p$ is in $L_{2}$, and moving only forward in $L_{2}$, ignoring the points which are not in the current $\overline{\mathcal{P}}$.

In the following, we describe how to use the data structure $T$ to implement one replace statement inside the inner while loop (Step 4.b) in time $O\left(\log ^{2} n\right)$. Start at the leaf of $T$ which has the center of $D$. Let $v$ be the current node and $p(v)$ be its parent in $T$. If $v$ is a left child, then replace $v$ by its parent, and repeat. Assume now that $v$ is a right child, and let $v^{\prime}$ be its sibling. In time $O(\log n)$, we can check whether there is a center in $A_{v^{\prime}}$ whose disk includes $p$. If yes, starting from $v^{\prime}$, find the last center in $A_{v^{\prime}}$ whose disk covers $p$ (this is the same binary search procedure used to compute $D_{q}^{p}$ ). If no, then replace $v$ by its parent, and repeat. Theorem 13 implies that there always exists a disk before $D$ which contains $p$, and thus the above procedure is correct.

The condition of the while loop in Step 4.b is checked by using a data structure $T^{\prime}$ which we describe below. $T^{\prime}$ is a balanced binary tree, whose leaves are the points of $\mathcal{P}$, sorted as in $L_{1}$ (the smaller $D_{f}$ is first). For a node $v \in T$, denote by $A_{v}^{\prime}$ the set of points in the subtree rooted at $v$. For every node $v$ in $T$, construct the furthestsite Voronoi diagram of $A_{v}^{\prime}$, and preprocess the diagram for membership queries. A cell of the furthest-site Voronoi diagram contains points which have the same furthest site in $A_{v}^{\prime}$. The furthest-site Voronoi diagram can be constructed in $O\left(\left|A_{v}^{\prime}\right| \log \left|A_{v}^{\prime}\right|\right)$ and uses $O\left(\left|A_{v}^{\prime}\right|\right)$ space (see, for example, Chapter 20.3 in [8], and with the same time and space bounds, it can be preprocessed for membership queries (Chapter 30.3 of [8]). The total preprocessing time is $O\left(n \log ^{2} n\right)$ and the total space is $O(n \log n)$. Given the center of a disk $D$, and a node $v \in T^{\prime}$, finding if all the points of $A_{v}^{\prime}$ are contained in $D$ can be done in time $O(\log n)$ by a membership query in the furthest-site Voronoi diagram of $A_{v}^{\prime}$.

Now we describe how exactly to check if there is a point $p^{\prime} \in \overline{\mathcal{P}}$ with $D_{f}^{p^{\prime}}<D$ and $p^{\prime} \notin D$. Given $D$, the set of points of $\overline{\mathcal{P}}$ with $D_{f}^{p^{\prime}}<D$ are consecutive in $L_{1}$. We denote by $p_{1}$ and $p_{2}$ the first, respectively last such point. Note that $p_{1}$ is the first point in the remaining part of $L_{1}$, and that $p_{2}$ can be found by binary search in $O(\log n)$ time. Checking the condition for $p_{1}$ and $p_{2}$ takes constant time. In $O(\log n)$ we can also locate $p_{1}$ and $p_{2}$ in the leaves $v_{1}$ and $v_{2}$ of $T^{\prime}$ (by binary search), and find the least common ancestor $v$ of $v_{1}$ and $v_{2}$ in $T^{\prime}$. Let $v^{\prime}$ be $v_{1}$, and $p\left(v^{\prime}\right)$ be its parent in $T^{\prime}$. As long as $p\left(v^{\prime}\right) \neq v$, do the following: if $v^{\prime}$ is a left child, and if $v^{\prime \prime}$ is its sibling, check if $D$ contains $A_{v^{\prime \prime}}^{\prime}$. Then let $v^{\prime}$ be $p\left(v^{\prime}\right)$.

Similarly, let $v^{\prime}$ be $v_{2}$. As long as $p\left(v^{\prime}\right) \neq v$, do the following: if $v^{\prime}$ is a right child, and if $v^{\prime \prime}$ is its sibling, check if $D$ contains $A_{v^{\prime \prime}}^{\prime}$. Then let $v^{\prime}$ be $p\left(v^{\prime}\right)$. If any of the checks above fails, then there is a point $p^{\prime} \in \overline{\mathcal{P}}$ with with $D_{f}^{p^{\prime}}<D$ and $p^{\prime} \notin D$; otherwise there is no such point. The total time for one disk $D$ is $O\left(\log ^{2} n\right)$, as there are at most $2 \log n$ queries of the type: check whether all the points of $A_{v}^{\prime}$ are contained in $D$.

Finally, Lemma 11 implies that each disk $D$ will need at most once such processing. Thus Step 4.b takes $O\left(n \log ^{2} n\right)$ time overall. Based on the discussion above, we have:

Theorem 15 Algorithm 3 can be implemented to run in time $O\left(n \log ^{2} n\right)$.

## 6 The general minimum disk cover problem

In this section we describe a constant-factor approximation algorithm for the following
Minimum Disk Cover Problem. Given a set of unit disks $\mathcal{D}$ and a set of points $\mathcal{P}$ in the Euclidean plane, find a minimum-size subset $\mathcal{F} \subseteq \mathcal{D}$, such that $\mathcal{P} \subseteq \cup\{D \in \mathcal{F}\}$.

This problem is NP-Hard since it contains as a special case Dominating Set in unit-disk graphs, a problem shown to be NP-Hard in [5]. A polynomial-time algorithm with constant approximation ratio for Minimum Disk Cover was first provided by [2].

If we can obtain a constant ratio for covering an unit-side equilateral triangle, we can obtain a constant ratio for the whole plane, by tiling the plane into triangles and separately covering all the triangles, and using the fact that one disk in the optimum can only cover points in a constant number of triangles.

Let $A B C$ be such a triangle. If no point of $\mathcal{P}$ is in the triangle, there is nothing to be done. Also, if there is a disk $D \in \mathcal{D}$ whose center is in the triangle, then $D$ covers all the triangle. So, in the following, we assume all the points are in the triangle, and all the centers of disks in $\mathcal{D}$ are outside the triangle.

The algorithm has four phases:

1. After removing those disks that do not intersect the triangle, partition the remaining disks into three sets $\mathcal{D}_{1}, \mathcal{D}_{2}$, and $\mathcal{D}_{3}$, such that all the centers of the disks in $\mathcal{D}_{1}$ are on the other side of the line $A B$ than $C$, all the centers of the disks in $\mathcal{D}_{2}$ are on the other side of the line $B C$ than $A$, and all the centers of the disks in $\mathcal{D}_{3}$ are on the other side of the line $A C$ than $B$. If a disk could be put in more than one $\mathcal{D}_{i}$, pick one arbitrarily.
2. For $i=1,2,3$, let $Q_{i}$ be the triangle $A B C$ and let $J_{i}$ be a line which separates the centers of the disks of $\mathcal{D}_{i}$ from the interior of the triangle. Find the skyline as in Section 3, and compute $\mathcal{F}_{i}$, the set of disks containing some arc of the skyline.
3. Write the natural Integer Programming formulation involving only the disks in $\mathcal{F}_{1} \cup \mathcal{F}_{2} \cup \mathcal{F}_{3}$. Solve the Linear Programming relaxation.
4. Round the linear programming optimum to an integer solution, as described in Subsection 6.1.

Later we prove Theorem 16, which claims that the algorithm described above has approximation ratio at most 6 for the problem of covering the points inside the triangle.

First, we note that Lemma 2 holds easily when $J_{i}$ is a straight line. For each $\mathcal{F}_{i}$, Lemma 8 also holds. Let $\mathcal{F}=\mathcal{F}_{1} \cup \mathcal{F}_{2} \cup \mathcal{F}_{3}$, and assume $\mathcal{F}$ is sorted with $\mathcal{F}_{1}$ (which is sorted) followed by the sorted $\mathcal{F}_{2}$, and followed by
the sorted $\mathcal{F}_{3}$. Lemma 9 also holds, and therefore $\mathcal{F}$ contains a solution at most twice opt, the size of an optimum solution.

### 6.1 Rounding

We use the natural IP, with variables $x_{D}$, for $D \in \mathcal{F}$ :

$$
\begin{gather*}
\text { minimize } \sum_{D \in \mathcal{F}} x_{D} \\
\text { subject to } \sum_{D: P \in D} x_{D} \geq 1 \quad \forall P \in \mathcal{P}  \tag{1}\\
x_{D} \in\{0,1\} \quad \forall D \in \mathcal{D} . \tag{2}
\end{gather*}
$$

Let LP be the linear programming relaxation of IP, obtain by replacing the constraints 2 by

$$
\begin{equation*}
x_{D} \geq 0 \quad \forall D \in \mathcal{D} \tag{3}
\end{equation*}
$$

Let $Z_{I P}^{*}$ the value of the IP optimum. As argued above, we have $Z_{I P}^{*} \leq 2$ opt.
Let $y$ be a (fractional) solution to LP. For a point $P \in \mathcal{P}$, the set of disks covering it consists of at most three intervals, say $I_{1}^{P}, I_{2}^{P}$, and $I_{3}^{P}$. For one of the three intervals, which we call simply $I^{P}$, we have: $\sum_{D \in I^{P}} y_{D} \geq 1 / 3$.

We introduce a second integer program, which we are able to solve in polynomial time exactly (see details below), and which approximates well LP, the linear programming relaxation described above. Precisely, consider the integer program IP', with variables $x_{D}$, for $D \in \mathcal{F}$ :

$$
\begin{gather*}
\text { minimize } \sum_{D \in \mathcal{F}} x_{D} \\
\text { subject to } \sum_{D \in I^{P}} x_{D} \geq 1 \quad \forall P \in \mathcal{P}  \tag{4}\\
x_{D} \in\{0,1\} \quad \forall D \in \mathcal{D} . \tag{5}
\end{gather*}
$$

Let LP' be the linear programming relaxation of IP'. The matrix of IP' is totally unimodular (see [6], Theorem 6.28, page 223, and Example 3 on the next page), and $3 y$ is a solution to LP'. Therefore IP' has a solution of size at most $3 \sum_{D \in \mathcal{F}} y_{D}$, and an optimum for IP' can be found easily by the greedy algorithm, as described at the end of the proof of Theorem 3. Now, if $y$ is an optimum solution to LP, then $\sum_{D \in \mathcal{F}} y_{D} \leq Z_{I P}^{*} \leq 2$ opt, and therefore the solution found by the greedy algorithm has size at most 6opt.

Rounding consists of finding for each point $P \in B$ the interval $I^{P}$, and then using the greedy algorithm to hit each $I^{P}$ with elements of $\mathcal{F}$. In conclusion, we proved:

Theorem 16 The algorithm described in this section has approximation ratio at most 6 for covering points inside a unit-size equilateral triangle with sizes equal to 1 with unit-disks from a fixed set $\mathcal{D}$.

Since a single disk from the optimum solution can cover points in at most 17 triangles of the tiling we conclude

Corollary 17 There is a 102-approximation algorithm for the Minimum Disk Cover problem.

## 7 Conclusions

In this paper we presented a geometric $O(n \log n)$ 6-approximation algorithm and a combinatorial $O\left(n \log ^{2} n\right)$ 3approximation algorithm for selecting forwarding neighbors in wireless ad-hoc networks, significantly improving both the running time and the approximation ratio of the best previously known algorithm. An extension of our method can be used to obtain an alternative constant-ratio polynomial-time algorithm for the Minimum Disk Cover problem.

We mention that Theorem 13 is true in the following more general setting. Let $J$ be an infinite simple Jordan curve which separates the plane into exactly two regions, and let $B$ be one of these two regions. Let all points $\mathcal{P}$ be in $B$, and each $D_{j}$ be a region bordered by a simple closed Jordan curve $\partial \mathcal{D}_{j}$. Assume that each $\partial \mathcal{D}_{j}$ intersects the infinite curve $J$ in exactly two points, and, for any two regions $D_{j}$ and $D_{k}, \partial \mathcal{D}_{j} \cap \partial \mathcal{D}_{k} \cap B$ has at most one point. Moreover, assume that, whenever two of the curves $\partial \mathcal{D}_{j}$ intersect, they cross each other. Then Lemma 10 holds, and therefore a polynomial-time exact disk covering algorithm exists.

On the other side, if the disks have arbitrary radii, the boundaries of two disks in the region $B$ (as defined in the paragraph above) can cross in two distinct points. Our arguments rely implicitly on the assumption that boundaries intersect in at most one point inside of the region $B$ where all the points lie, and all the presented algorithms fail when applied to arbitrary disks.

When $\mathcal{P}$ and the centers of $\mathcal{D}$ are separated by a straight line, we can apply the techniques from this paper to obtain a rounding procedure with ratio of 2 to the natural linear program LP. Then, as in Section 6 , it follows that the linear program LP has constant integrality ratio for the general Disk Cover problem. However, when the disks in $\mathcal{D}$ are weighted, we do not know the integrality ratio of the corresponding integer and linear programs. The linear program is given below:

$$
\begin{gather*}
\text { minimize } \sum_{D \in \mathcal{D}} w_{D} x_{D} \\
\text { subject to } \sum_{D: P \in D} x_{D} \geq 1 \quad \forall P \in \mathcal{P}  \tag{6}\\
x_{D} \geq 0 \quad \forall D \in \mathcal{D} \tag{7}
\end{gather*}
$$

This research work assumed that the wireless nodes are not able to adjust the range of transmission. When nodes are able to adjust the range of transmission, it is possible that congestion can be further reduced. We leave for further research the design of forwarding algorithms in the variable transmission range setting.

We also leave for further research the formulation of the overall broadcast problem, even if the algorithms one obtains after doing it are centralized, and thus not practical. A comparison with the classical broadcasting problems in the telephone and post office models would then be beneficial.

Student Yuchen Wu, whom we wish to thank, implemented the geometric 1-hop covering algorithm in a quadrant (Algorithm 2), and the exact $O\left(n^{2}\right)$ combinatorial algorithm for the quadrant given in [3]. On random instances, the much simpler geometric algorithm finds solutions which are on the average only $17-44 \%$ larger than the optimum found by the exact combinatorial algorithm.

## References

[1] J. Broch, D. B. Johnson, and D. A. Maltz, The Dynamic Source Routing Protocol for Mobile Ad Hoc Networks, IETF Internet Draft, draft-ietf-manet-dsr-05.txt, March 2001.
[2] H. Bronnimann and M.T. Goodrich, Almost Optimal Set Covers in Finite VC-Dimension. Proc. 1Oth ACM Symp. on Computational Geometry (SCG), 1994, 293-302.
[3] G. Calinescu, I. Mandoiu, P-J. Wan, and A. Zelikovsky, Selecting Forwarding Neighbors in Wireless Ad Hoc Networks, Proc. $5^{\text {th }}$ International Workshop on Discrete Algorithms and Methods for Mobility, 2001, 34-43.
[4] V. Chvátal. A greedy heuristic for the set-covering problem. Mathematics of Operation Research, 4(3):233-235, 1979.
[5] B. N. Clark, C. J. Colbourn, and D. S. Johnson, Unit Disk Graphs, Discrete Mathematics, 86:165-177, 1990.
[6] W. J. Cook, W. H. Cunningham, W. R. Pulleyblank, and A. Schrijver Combinatorial Optimization, Wiley-Interscience, 1998.
[7] T. H. Cormen, C. E. Leiserson, and R. L. Rivest Introduction to Algorithms, McGrow Hill, 1990.
[8] J. E. Goodman and J. O'Rourke (editors), Handbook of Discrete and Computational Geometry, CRC Press, 1997.
[9] Z. J. Haas, M. R. Pearlman, and P. Samar. The Interzone Routing Protocol (IERP) for Ad Hoc Networks, IETF Internet Draft, draft-ietf-manet-zone-ierp-00.txt, January 2001.
[10] D. S. Hochbaum and W. Maass, Approximation schemes for covering and packing problems in imageprocessing and VLSI, Journal of the ACM, 32(1): 130-136, 1985.
[11] H. B. Hunt III, M. V. Marathe, V. Radhakrishnan, S. S. Ravi, D. J. Rosenkrantz, and R. E. Stearns, NC-Approximation schemes for NP- and PSPACE-hard problems for geometric graphs, Journal of Algorithms, 26(2):238-274, 1998.
[12] P. Jacquet, P. Muhlethaler, A. Qayyum, A. Laouiti, L. Viennot, and T. Clausen, Optimized Link State Routing Protocol, IETF Internet Draft, draft-ietf-manet-olsr-04.txt, March 2001.
[13] A. Qayyum, L. Viennot, and A. Laouiti, Multipoint Relaying: An Efficient Technique for Flooding in Mobile Wireless Networks, Proc. 35th Annual Hawaii International Conference on System Sciences (HICSS 2002).
[14] M. V. Marathe, H. Breu, H. B. Hunt III, S. S. Ravi and D. J. Rosenkrantz, Simple Heuristics for Unit Disk Graphs, Networks, Vol. 25, 1995, pp. 59-68.
[15] B.M.E. Moret and H.D. Shapiro. Algorithms from P to NP, Volume I: Design and Efficiency. Benjamin/Cummings, 1991.
[16] S.-Y. Ni, Y.-C. Tseng, Yuh-Shyan Chen, and J.-P. Sheu, The Broadcast Storm Problem in a Mobile Ad Hoc Network, Proceedings of the fifth annual ACM/IEEE international conference on Mobile computing and networking, 1999, Pages 151-162.
[17] C. E. Perkins, E. M. Royer, and S. Das, Ad Hoc On Demand Distance Vector (AODV) Routing, IETF Internet Draft, draft-ietf-manet-aodv-08.txt, March 2001.
[18] G. Ramalingam and C. Pandu Rangan, A unified approach to domination problems in interval graphs, Information Processing Letters 27 (1988), pp. 271-274.
[19] P. Sinha, R. Sivakumar, and B. Vaduvur, Enhancing Ad Hoc Routing with Dynamic Virtual Infrastructures, Proc. IEEE INFOCOM 2001, pp. 1763-1772.

## Author biographies

Gruia Călinescu is an Assistant Professor of Computer Science at the Illinois Institute of Technology. He has a Diploma from University of Bucharest and a PhD from Georgia Institute of Technology. His research interests are in the area of algorithms.

E-mail: calinesc@cs.iit.edu.
Ion I. Măndoiu received the M.S. degree from Bucharest University in 1992 and the Ph.D. degree from Georgia Institute of Technology in 2000, both in Computer Science. He worked as a Research Assistant at Bucharest University (1992-1995), Instructor at Georgia Institute of Technology (2000-2001), and Postgraduate Researcher at University of California, Los Angeles (2000-2001). Since 2001 he has been a Research Scientist at University of California, San Diego. His research interests include approximation algorithms, VLSI physical layout design, and combinatorial optimization.

E-mail: mandoiu@cs.ucsd.edu.
Peng-Jun Wan has been an Assistant Professor in Computer Science at Illinois Institute of Technology since 1997. He received his PhD in Computer Science from University of Minnesota in 1997, MS in Operations Research and Control Theory from Chinese Academy of Science in 1993, and BS in Applied Mathematics from Qsinghua University in 1990. His research interests include optical networks and wireless networks.

E-mail: wan@cs.iit.edu.
Alexander Zelikovsky received the Ph.D. degree in Computer Science from the Institute of Mathematics of the Belorussian Academy of Sciences in Minsk, Belarus, in 1989. He worked at the Institute of Mathematics in Kishinev as a senior research scholar from 1989 to 1995. Between 1992 and 1995 he visited Bonn University and the Institut fur Informatik in Saarbrueken (Germany). Dr. Zelikovsky was a Research Scientist at Virginia University (19951997) and a Postdoctoral Scholar at UCLA (1997-1998). Since 1999 he has been an Assistant Professor at the Computer Science Department of Georgia State University. He is the author of more than 60 refereed publications. Dr. Zelikovsky's research interests include VLSI physical layout design, discrete and approximation algorithms, combinatorial optimization, and ad-hoc wireless networks.

E-mail: alexz@cs.gsu.edu

