Minimizing Uncertainty through Sensor Placement with Angle Constraints

Ioana O. Bercea* Volkan Isler† Samir Khuller*

Abstract

We study the problem of sensor placement in environments in which localization is a necessity, such as ad-hoc wireless sensor networks that allow the placement of a few anchors that know their location or sensor arrays that are tracking a target. In most of these situations, the quality of localization depends on the relative angle between the target and the pair of sensors observing it. In this paper, we consider placing a small number of sensors which ensure good angular α-coverage: given α in $[0, \pi/2]$, for each target location t, there must be at least two sensors s_1 and s_2 such that the $\angle(s_1 t s_2)$ is in the interval $[\alpha, \pi - \alpha]$. One of the main difficulties encountered in such problems is that since the constraints depend on at least two sensors, building a solution must account for the inherent dependency between selected sensors, a feature that generic Set Cover techniques do not account for.

We introduce a general framework that guarantees an angular coverage that is arbitrarily close to α for any $\alpha <= \pi/3$ and apply it to a variety of problems to get bi-criteria approximations. When the angular coverage is required to be at least a constant fraction of α, we obtain results that are strictly better than what standard geometric Set Cover methods give. When the angular coverage is required to be at least $(1 - 1/\delta) \cdot \alpha$, we obtain a $\mathcal{O}(\log \delta)$- approximation for sensor placement with α-coverage on the plane. In the presence of additional distance or visibility constraints, the framework gives a $\mathcal{O}(\log \delta \cdot \log k_{\mathsf{OPT}})$-approximation, where k_{OPT} is the size of the optimal solution. We also use our framework to give a $\mathcal{O}(\log \delta)$-approximation that ensures $(1 - 1/\delta) \cdot \alpha$-coverage and covers every target within distance $3R$.

1 Introduction

Localization is an important necessity in many mobile computing applications. In ad-hoc wireless sensor networks, it centers around the ability of nodes to self localize using little to no absolute spatial information. When mobility is considered, the problem becomes that of tracking a moving target through a sensor network in which a set of sensors must combine measurements in order to detect the location of the target.

When a large number of sensors are deployed, it becomes impractical to equip all of them with the capability of localizing themselves with respect to a global system (such as through GPS). From this perspective, a commonly used technique is to employ a small number of anchors (or beacons) that know their location and are capable of transmitting it to the other nodes seeking to localize themselves [27]. Alternatively, sensors such as cameras or microphone arrays placed in the environment can collect measurements which can then be used to estimate the locations of objects of interest. In some of the most popular scenarios, each target seeking to localize itself has access to Euclidean distances and/or angular measurements (bearing) *relative* to the sensors that are in its vicinity. When exact distances or bearings from two sensors to a target are known, localization can be easily performed through the process of triangulation. In practice, however, the inherent sensor measurements are noisy and several models of uncertainty have been proposed [4].

From a geometric perspective, a common benchmark for estimating uncertainty is the Geometric Dilution of Precision (GDOP). This benchmark investigates how the relative geometry between sensors and target nodes amplifies measurement errors and affects the localization error. Savvides et al. [23, 24] observe that the error is largest when the angle θ between two sensors and the target node is either very small or close to π. The analysis of Kelly [14] further shows that, when triangulation is used, this angle contributes to the GDOP at a fundamental level. When distance measurements are used in triangulation, the GDOP is proportional to $1/|\sin\theta|$. When angular measurements are used, the GDOP is proportional to $d_1 \cdot d_2/|\sin\theta|$, where d_1 and d_2 are the distances from the sensors to the target. Intuitively, each measurement becomes a constraint that restricts the set of possible locations of the target and the quality of localization depends on both the area and the shape of all intersections. As seen in Figure 1, when the angle θ is close to 0 or π, the feasible set becomes unconstrained and the error unbounded. In particular, when the sensors and the target are collinear, localization is impossible.

Inspired by these observations, our paper focuses on the geometry of sensor deployment and asks the question of where should the sensors be placed so as to en-

*Department of Computer Science, University of Maryland, College Park, {ioana,samir}@cs.umd.edu

†Department of Computer Science and Engineering, University of Minnesota, Minneapolis, isler@cs.umn.edu

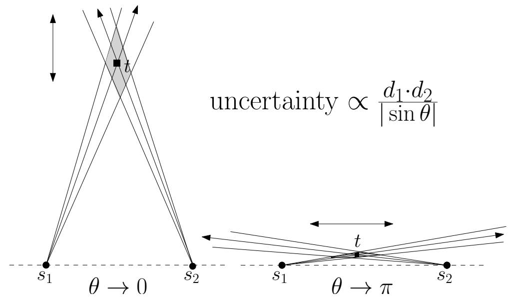

Figure 1: When bearing information is used to determine a target's location, the measurement corresponds to a cone centered at the true bearing from the sensor to the target. Given sensors s_1, s_2 and target t, let $\theta = \angle s_1 t s_2$, $d_1 = d(s_1, t)$ and $d_2 = d(s_2, t)$. When the sensors are too far from the target, the uncertainty becomes unbounded. If θ is small, the y-direction becomes unconstrained. If θ is large, the x-direction becomes unconstrained.

sure that the GDOP is below a given threshold at all target locations? To this end, we define an angular constraint which we call α-*coverage*: given a parameter $\alpha \in [0, \pi/2]$, each target at position t must be assigned two distinct sensors (at positions s_1 and s_2) such that the angle $\theta = \angle s_1 t s_2$ is in the range $[\alpha, \pi - \alpha]$ (i.e. neither too small nor too big). We then frame the problem of sensor placement as a bicriteria optimization problem: given a set of possible sensor and target locations, we wish to select the smallest number of sensors that provide α-coverage for all target locations. We address these variants from a theoretical perspective and present a general algorithmic framework that specifically addresses the angular constraint and iteratively obtains better angular guarantees at the expense of larger solution sizes.

Our model. Formally, we consider the two dimensional model in which the set of candidate sensor locations is a discrete set $X \subseteq \mathbb{R}^2$ and the set of target locations is a discrete set $T \subseteq \mathbb{R}^2$. We have chosen to discuss this discrete setting (instead of the continuous one) because we consider it to be more theoretically rich. If we could choose sensor locations anywhere in the plane, we could essentially impose a grid on the plane and use a constant number of sensors per cell to cover the targets (without violating the angular constraints). The analysis would then use k-center or Art Gallery techniques to obtain constant factor approximations.

Given a parameter $\alpha \in [0, \pi/2]$, we say that an (unordered) pair (s_1, s_2) α-covers a target t if $\theta = \angle s_1 t s_2$ is in the range $[\alpha, \pi - \alpha]$. Notice that the higher the value of α, the smaller the range of values that θ can take. Our algorithmic framework applies to several sensor coverage problems. First, we define the **Minimum Sensor Placement with α-coverage** ($\alpha-\mathsf{Ang}$) problem, which asks for the smallest set of sensors that α-covers T. We then consider its clustering variant, in which we additionally require that the sensors be within a given range R of the target (corresponding to a finite sensing range scenario). We call this the **Minimum Sensor Placement with (α, R)-coverage** ($(\alpha, R)-\mathsf{AngDist}$) problem and say that a pair of sensors (α, R)-covers a target if both sensors respect the constraints. Finally, we consider a version of the Art Gallery problem, in which the target must be visible to both sensors. This problem was first introduced by Efrat, Har-Peled and Mitchell [8] which discussed the case in which the sensors (contained in a region P) are required to α- guard targets that are contained in a smaller region $Q \subseteq P$. Two sensors s_1, s_2 α-guard a target t if they both see the target and (s_1, s_2) α-covers t. Given a sensor $s \in X$ and a target $t \in T$, we say that s sees t if the segment connecting the two does not cross the boundary of P (i.e. t is within line-of-sight of s). We henceforth refer to this problem as the **Art Gallery with α-coverage** ($\alpha-\mathsf{ArtAng}$) problem.

Related work. When it comes to sensor coverage problems, extensive work has been done although surprisingly few results discuss α-coverage. Notable exceptions are the work of Efrat, Har-Peled and Mitchell [8], Tekdas and Isler [25] and Isler, Khanna, Spletzer and Taylor [13]. As mentioned before, Efrat et al. [8] introduce the $\alpha-\mathsf{ArtAng}$ problem in which two sensors are required to α-guard a target. They present a $\mathcal{O}(\log k_{\mathsf{OPT}})$-approximation algorithm that guarantees $\alpha/2$-coverage, where k_{OPT} is the size of the optimal solution. Their main subroutine is similar to an algorithm for the Art Gallery problem that first imposes a grid Γ on P and chooses guards located at vertices of Γ. We note however, that such a step is not necessary in our case, since X is already a discrete set. Their algorithm runs in time $\mathcal{O}(nk_{\mathsf{OPT}}^4 \log^2 n \log m)$, where n is the number of vertices of P and m is the number of points in $\Gamma \cap P$ (i.e. possible sensor locations). In contrast, we present a framework that achieves $(1-1/\delta)\cdot\alpha$-coverage, approximates the number of sensors by $\mathcal{O}(\log \delta \cdot \log k_{\mathsf{OPT}})$ and runs in time $\mathcal{O}(\log \delta \cdot k_{\mathsf{OPT}} \cdot mn \log m)$, for any $\delta \geq 1$ and $\alpha \leq \pi/3$. In our case and throughout the rest of the paper, n represents the number of targets.

Tekdas and Isler [25] formalize the angle constraint by requiring that the uncertainty $d_1 \cdot d_2/|\sin\theta|$ computed by Kelly [14] be smaller than a certain threshold U. When the targets are contained in some subset of the plane and the sensors can be placed anywhere (continuous case), they present a 3-approximation with maximum uncertainty $\leq 5.5U$. We note that $(\alpha, R)-\mathsf{AngDist}$ is a generalization of the above problem in the sense in which an algorithm for $(\alpha, R)-\mathsf{AngDist}$ can be used in approximately solving the former. Finally, Isler et al [13] consider the case in which the sensor locations are already given and one must compute an assignment of sensors to targets that minimizes the total sum of errors. In addition, they require that each sensor be used in tracking only one target. The version relevant to our problem is when the error is defined as $1/\sin\theta$. In the case in which the sensors are equally spaced on a circle, they present a 1.42-approximation that also applies to minimizing the maximum error.

Our contributions. For the case of $\alpha \leq \pi/3$, we provide a general bi-criteria framework that approximates the angular coverage to arbitrary precision while guaranteeing a good approximation in the size of the solution. Specifically, for any $\delta > 1$, we propose an iterative method that guarantees $(1-1/\delta)\cdot\alpha$-coverage and approximates the solution size by $\mathcal{O}(\log\delta)$ for $\alpha-\mathsf{Ang}$ and $\mathcal{O}(\log\delta\log k_{\mathsf{OPT}})$ for $(\alpha,R)-\mathsf{AngDist}$ and $\alpha-\mathsf{ArtAng}$. When the polygon in $\alpha-\mathsf{ArtAng}$ is allowed to have h holes, we obtain a $\mathcal{O}(\log\delta\cdot\log k_{\mathsf{OPT}}\cdot\log h)$-approximation. It is worthwhile to note that the main technical theorem of the framework refers solely to the angle coverage constraint and as such, could be applied to a variety of other problems as long as the other constraints (such as distance or line-of-sight visibility) define a good set system (one with constant VC dimension, for example).

In addition, we present further approximations for $(\alpha,R)-\mathsf{AngDist}$. We relax the distance constraints from R to $3R$ and reduce the approximation factor of the solution size from $\mathcal{O}(\log\delta\cdot\log k_{\mathsf{OPT}})$ to $\mathcal{O}(\log\delta)$, while keeping the angular coverage at $(1-1/\delta)\cdot\alpha$. We achieve this by using a more involved technique of employing ϵ-nets that we believe may be of independent interest.

We also consider the case in which $\alpha = 0$ and construct a set of optimal size that covers the targets within distance $(1+\sqrt{3})\cdot R$. This particular case remains relevant since it captures the spirit of fault tolerance by requiring two distinct sensors to be assigned to a target. We achieve our result by showing a $(1+\sqrt{3})$-approximation for the more general Euclidean Fault Tolerant k-suppliers problem which improves on the existing 3-approximation by Khuller et al [15].

Discussion of existing techniques. In a more general context, we believe that angular constraints are interesting not just because they contribute to a new geometrical direction in sensor coverage problems. From a purely theoretical perspective, they also present the challenge of approximating an optimization problem whose objective function is linear in the number of chosen sensors but whose constraints depend on pairs of sensors jointly satisfying a condition.

In such a case, an algorithm that chooses pairs which satisfy the constraints might incur an overall cost that is quadratic in the objective function. For example, one natural way in which we can consider $\alpha-\mathsf{Ang}$ is as an instance of Set Cover. For each pair of sensors (s,s'), we can define $S_{(s,s')}$ to be the set of targets $t\in T$ such that (s,s') α-covers them. Set Cover asks for the smallest number of pairs whose union covers T. The generic greedy charging scheme for Set Cover gives us a solution of size at most $k^*\cdot\log n$, where k^* is the size of the optimal Set Cover solution. Notice, however, that k^* can be much larger than k_{OPT} (the size of the optimal set of sensors for $\alpha-\mathsf{Ang}$) and this could lead to a quadratic blowup in the size of the solution. In the worst case, greedy Set Cover could pick as many as $\binom{k_{\mathsf{OPT}}}{2}\cdot\log n$ sensors in its solution. This degenerate case might happen when we end up picking distinct pairs of sensors to cover each target while the optimal solution picks a much smaller set of sensors that collectively α-cover the targets.

Using the geometry of the problem, one could slightly improve the above approximation factor from $\mathcal{O}(k_{\mathsf{OPT}}\cdot\log n)$ to $\mathcal{O}(k_{\mathsf{OPT}}\cdot\log k_{\mathsf{OPT}})$. Specifically, one can show that each $S_{(s,s')}$ is induced by the symmetric difference of two circles. As such, these objects have constant Vapnik-Chervonenkis (VC) dimension [19] and allow for a $\mathcal{O}(\log k^*)$-approximation for Set Cover [10, 3]. Unfortunately, this only gets us a $\mathcal{O}(k_{\mathsf{OPT}}\cdot\log k_{\mathsf{OPT}})$-approximation guarantee. The persistent k_{OPT} factor in the approximation comes from the fact that the Set Cover framework cannot distinguish between sensors that help cover a lot of targets (locally) and sensors that, additionally, can also help cover more targets in conjunction with other sensors. In other words, it does not make use of the **global dependency** between sensors in order to get a small solution size.

In fact, such observations are more in the spirit of Label Cover type problems in which we need to assign labels to vertices of a graph but a specific labeling is considered feasible only when it satisfies certain edge constrains. Indeed, when considered in its full generality (i.e. points lie in arbitrary space and coverage is defined arbitrarily), the problem becomes a generalization of MinRep and, as such, incurs a hardness of approximation bound of $2^{\log^{1-\epsilon} n}$, for any $0<\epsilon<1$ unless NP $\subseteq$ DTIME$(n^{\mathrm{polylog}(n)})$ [17]. Such occurrence of Label Cover in a natural setting is intriguing in its own right. Furthermore, it can conceivably model other instances in which the quality of a solution depends on pairs of elements jointly satisfying a condition, such as in pairwise feature selection for Machine Learning tasks [22, 6]. We defer the rest of the discussion to the extended version of the paper [1].

2 Algorithmic Framework

Let $m = |X|$ be the number of possible sensor locations and $n = |T|$ be the number of target locations. The underlying distance function will be the Euclidean ℓ_2

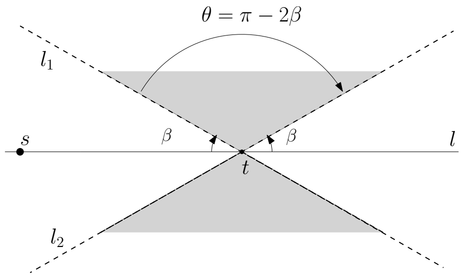

Figure 2: The set $R_t(s, \beta)$ is induced by the double-wedge generated by the lines l_1 and l_2 and has a central angle $\theta = \pi - 2\beta$.

metric. We consider (unordered) pairs of the form (s, s') where $s \neq s'$, $s \in S$, $s' \in S'$ and $S, S' \subseteq X$ are sets of sensors. We denote the set of such pairs as $S \times S'$. Formally, $S \times S' = \{(s, s') | s \neq s', s \in S, s' \in S'\}$. We say that the set of pairs $S \times S'$ α-covers t if there exists a pair $(s, s') \in S \times S', s \neq s'$, that α-covers t. When $S' = S$, we simply say that the set S α-covers t. A pair or a set of pairs α-covers a set T of targets when it α-covers each element of T. In addition, a pair or a set of pairs (α, R)-covers a set T of targets *within distance* R if, for at least one of the pairs that α-covers a target, the distance from both sensors to the target is $\leq R$.

We begin by considering a fixed set S of given sensors and asking the question of how should we pick a second set of sensors S' such that pairs of the form (s, s') with $s \in S$ and $s' \in S'$ will β-cover the set of targets (for a given β). Formally, given a target $t \in T$, a sensor $s \in X$ and angle parameter $\beta \in [0, \pi/2]$, we define the set $R_t(s, \beta)$ of feasible sensor locations that, together with s, β-cover t:

$$R_t(s, \beta) = \{s' \in X | (s, s') \ \beta\text{-covers } t\}$$

As seen in Fig. 2, this set is induced by two wedges centered around t. A *wedge* is defined as the intersection of two non-parallel half spaces in $\mathbb{R}^2$. In our case, we are interested in the two wedges defined by the lines that form angles of β with the line that passes through s and t. The union of these two wedges will be referred to as the *double-wedge* around t. The task of constructing a pair that β-covers t can now be reduced to that of finding a second sensor s' inside the double-wedge, i.e. once we find such a sensor, we are guaranteed that the pair (s, s') β-covers t.

Specifically, the set S' becomes a hitting set for the given collection of double-wedges (one for each target). Given a set system $\mathcal{F}(X, \mathcal{R})$, where X is the set of sensors and $\mathcal{R}$ is a collection of subsets (double-wedges) of X, a hitting set is a set $H \subseteq X$ that intersects every subset in $\mathcal{R}$ non-trivially. The Hitting Set problem asks for a hitting set of minimum cardinality. Our method constructs S' by approximately solving the Hitting Set problem, in which the double-wedges may be further restricted to intersect the circle of radius R centered at the targets (for $(\alpha, R)-$AngDist) or the visibility polygons of the corresponding targets (for $\alpha-$ArtAng). For this step, we employ known techniques using constant VC dimension and ϵ-net constructions.

When it comes to the analysis, the challenge is making sure that S' is not much larger than $k_{\mathsf{OPT}} = |S_{\mathsf{OPT}}|$ (the size of the optimal set of sensors) and therein lies the difficulty. We do this by showing that S_{OPT} itself is a hitting set, so the size of the optimal hitting set is smaller than k_{OPT}. Specifically, our structural theorem shows that given a set S of sensors that $(\alpha - 2\epsilon)$-covers T for some $\epsilon \in (0, \alpha/2]$, S_{OPT} must intersect the double-wedges induced by S around each target with $\beta = \alpha - \epsilon$:

Theorem 1 *Let $\epsilon > 0$ be such that $\alpha - \epsilon \leq \pi/3$ and $\epsilon \leq \alpha/2$. Given a set S that $(\alpha - 2\epsilon)$- covers T, let $T' \subseteq T$ be the set of targets that S does **not** $(\alpha - \epsilon)$-cover. Then the set of pairs $S \times S_{\mathsf{OPT}}$ $(\alpha - \epsilon)$-covers T'.*

When $\epsilon = \alpha/2$, we start with an arbitrary set S and recover the observation of Efrat et al. [8]. In order to get better than $\alpha/2$-coverage, the seed set S cannot be chosen arbitrarily. In fact, our proof crucially uses the fact that the sensors in S already $(\alpha - 2\epsilon)$-cover the targets. Given a pair (s_1, s_2) in S that $(\alpha - 2\epsilon)$-covers a target t, we show that one of the optimal sensors must be in $R_t(s_1, \alpha - \epsilon) \cup R_t(s_2, \alpha - \epsilon)$, i.e. it either makes a good pair with s_1 or with s_2. The restriction that $\alpha \leq \pi/3$ is used when $\epsilon < \alpha/2$ and guarantees that the union of the two double-wedges (for s_1 and s_2) is itself a double-wedge with a large enough central angle that it must intersect S_{OPT}.

By adding the set S' to S, we now obtain a larger set that $(\alpha - \epsilon)$-covers the targets. Iterating this procedure $\log \delta$ times, we guarantee $(1 - 1/\delta) \cdot \alpha$-coverage and approximate the size of the optimal solution by $\log \delta \cdot c$, where c is the approximation guarantee we get from solving the Hitting Set problem in each iteration (Section 2.2).

Finally, we note that while adding more sensors in order to obtain a better solution is a classical approach, the challenge is do so in a way that addresses the global dependency between sensors and does not incur a quadratic cost in each iteration. In this context, instead of looking for pairs in which both sensors are unknown, our framework looks for pairs in which one sensor is in S (i.e. already known) and the other one is in S'. This allows us to cast the task of constructing S' as a global optimization problem (Hitting Set) for which good approximations exist. Moreover, when $\alpha \leq \pi/3$, our main structural theorem allows us to bound the size of such hitting sets linearly in k_{OPT} (essentially shaving off the additional k_{OPT} factor from Set Cover).

2.1 Proof of Theorem 1

In order to prove Theorem 1, we essentially show that S_{OPT} must intersect at least one of the double-wedges generated by S around a given target t. First, notice that if a set S already $(\alpha-\epsilon)$-covers a target t, then we do not need to worry: S will continue to $(\alpha-\epsilon)-$cover t even when we add S' to S. We are therefore concerned with targets in T' that are not already $(\alpha-\epsilon)$-covered by S.

Fix such a target $t \in T'$ and let $s_1, s_2 \in S$ be any two sensors that $(\alpha-2\epsilon)$-cover t but do not $(\alpha-\epsilon)$-cover it. Our strategy will be to show that there exists $s^* \in S_{\mathsf{OPT}}$ such that either (s_1, s^*) $(\alpha-\epsilon)$-covers t or (s_2, s^*) $(\alpha-\epsilon)$-covers t, i.e. $s^* \in R_t(s_1, \alpha-\epsilon)$ or $s^* \in R_t(s_2, \alpha-\epsilon)$. The candidates will be $s_1^*, s_2^* \in S_{\mathsf{OPT}}$ where (s_1^*, s_2^*) is the optimal pair that α-covers t. We will show that s^* is either s_1^* or s_2^*. Intuitively, each of the double-wedges induced by s_1 and s_2 alone is not big enough to "capture" s_1^* or s_2^*. However, if $\angle s_1ts_2$ is in the range $[\alpha-2\epsilon, \pi-(\alpha-2\epsilon)]$, then the union D_t of these double-wedges will have a large enough central angle to guarantee that one of the optimal sensors is contained in it. In other words, at least one of the optimal sensors s_1^* or s_2^* together with either s_1 or s_2 will $(\alpha-\epsilon)$-cover t.

Let D_1 and D_2 be the double-wedges corresponding to $R_t(s_1, \alpha-\epsilon)$ and $R_t(s_2, \alpha-\epsilon)$, respectively. Notice that they have central angles $\theta_{D_1} = \theta_{D_2} = \pi - 2(\alpha-\epsilon)$. Let $\alpha' = \angle(s_1ts_2)$. We begin by first establishing that the union of these two double-wedges generated by s_1 and s_2 is a larger double-wedge.

Lemma 2 *The union of the two double-wedges D_1 and D_2 is a larger double-wedge D_t centered at t with central angle $\theta_{D_t} = \pi - 2(\alpha-\epsilon) + \alpha'$.*

Proof. We refer the reader to Figure 3 for an intuitive explanation. Formally, let l be the line that passes through s_1 and t and let l_1 and l_2 the two lines that define D_1. Since $\angle(s_1ts_2) \notin [\alpha-\epsilon, \pi-(\alpha-\epsilon)]$, it follows that s_2 is not in D_1. Assume without loss of generality that s_2 is between the lines l and l_1 in the counterclockwise direction. The same proof follows for the other possible locations of s_2.

Now consider D_2 and let l_3 and l_4 be the defining lines through t, while l' is the line that passes through s_2 and t. Notice that D_1 and D_2 are identical except that D_2 is a rotated copy of the D_1. In other words, since $\angle(l, l') = \alpha'$, we also have that $\angle(l_1, l_3) = \alpha'$ and $\angle(l_2, l_4) = \alpha'$.

We will show that $(l_1, l_3) \leq \angle(l_1, l_2)$, and hence conclude that l_3 must lie between l_1 and l_2. Since $\angle(l_1, l_3) = \alpha' \leq \alpha-\epsilon$ (by choice of s_2) and $\angle(l_1, l_2) = \pi - 2(\alpha-\epsilon)$, we have that when $\alpha-\epsilon < \pi/3$:

$$\alpha-\epsilon \leq \pi - 2(\alpha-\epsilon).$$

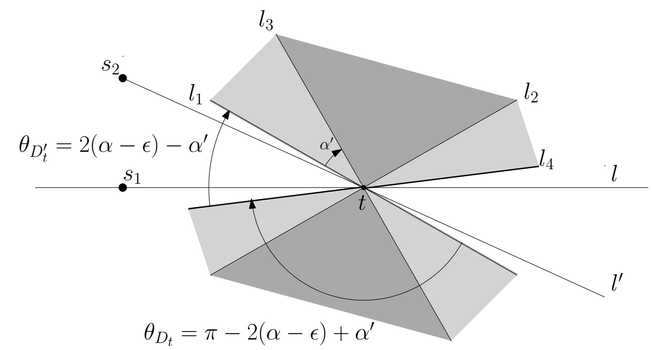

Figure 3: Since $\angle(s_1ts_2) = \alpha'$, D_2 is a rotation by α' of D_1 . Their union is another double-wedge D_t defined by l_1 and l_4 with central angle $\theta_{D_t} = \pi - 2(\alpha-\epsilon) + \alpha'$.

Therefore, the union of the two double-wedges D_1 and D_2 is a continuous double-wedge D_t determined by l_1 and l_4. It has central angle $\theta_{D_t} = \theta_{D_1} + \angle(l_2, l_4) = \pi - 2(\alpha-\epsilon) + \alpha'$.

□

The next step is to show that one of the two optimal sensors s_1^* and s_2^* must be in D_t. The intuition is that by making D_t have a large central angle, we ensure that the complement D_t' of D_t has such a small central angle that it would not be able to contain both s_1^* and s_2^*.

Lemma 3 *At least one of the two optimal sensors s_1^* and s_2^* assigned to t must be in D_t.*

Proof. Let D_t' be the complement of D_t. Notice that D_t' forms another double-wedge defined by l_1 and l_4 but that it does not actually contain points on these lines. Moreover, D_t' has a central angle $\theta_{D_t'} = \pi - \theta_D = 2(\alpha-\epsilon) - \alpha'$. Since s_1 and s_2 $(\alpha-2\epsilon)$-cover the target, and we are considering the case where s_2 is between l and l_1, we have that $\alpha' \geq \alpha - 2\epsilon$. Hence, we have that $\theta_{D_t'} \leq \alpha$. This implies that s_1^* and s_2^* cannot be both in the same wedge of D_t' without being exactly situated on the lines l_1 and l_4(i.e. in D_t). The other bad situation would be for them to be in different wedges of D'. But then the angle between them would be greater than θ_{D_t}. Since $\alpha' \geq \alpha - 2\epsilon$, we get that

$$\theta_{D_t} = \pi - 2(\alpha-\epsilon) + \alpha' \geq \pi - \alpha,$$

which would contradict the fact the $\angle(s_1^*ts_2^*) \in [\alpha, \pi - \alpha]$. In other words, at least one of the optimal sensors s_1^* and s_2^* must be in D_t. □

This concludes the proof of Theorem 1. At this point, it is worthwhile to notice that the requirement that $\alpha \leq \pi/3$ is relatively tight in this framework. When $\alpha - \epsilon > \pi/3$, both of the above claims fail. In particular, the central angle of D_1 and D_2 would be too small and their union would no longer correspond to a bigger double-wedge. Furthermore, it would no longer be true that such a union must intersect S_{OPT}.

2.2 Iterating to obtain $((1-1/\delta)\cdot\alpha)$-coverage

Given the technical lemmas from before that allow us to refine the angular coverage of a given seed set S, we can now develop a more general algorithm that constructs a new set that achieves $((1-1/\delta)\cdot\alpha)$-coverage for any $\delta > 1$. The idea is to iteratively apply the refinement step (by setting $S = S \cup S'$) $\log\delta$ times, first with $\epsilon = \alpha/2$, then with $\epsilon = \alpha/4$ etc. At the end of $\log\delta$ iterations, we have that the updated set S $((1-1/\delta)\cdot\alpha)$-covers T. The running time of the algorithm is $\log\delta$ times the time to find the appropriate hitting set plus the time it takes to find the starting set. This first set (denoted S_1) requires special care and depends on the problem at hand.

Notice that we require S_1 to 0-cover T but one can check that the proof of Theorem 1 follows in this case even when we do not have two distinct sensors 0-covering a target. Therefore, in the case of $\alpha-$Ang, it suffices to pick S_1 to consist of any sensor in X and get the following:

Theorem 4 *Given X, T, $\alpha \in [0,\pi/3]$ as above, we can find a set of sensors $S \subseteq X$ such that S $((1-1/\delta)\cdot\alpha)$-covers T and $|S| = \mathcal{O}(\log\delta)\cdot k_{OPT}$. The running time of the algorithm is $\mathcal{O}(\log\delta\cdot k_{OPT}\cdot m\log m)$.*

When it comes to the $(\alpha,R)-$AngDist problem, we require that the initial set S_1 has the property that each target is within distance R of at least one sensor in S_1. Without loss of generality, we can assume that $R=1$ and then our problem becomes an instance of the Discrete Unit Disk Cover (DUDC) problem [5]. In DUDC, we are given a set $\mathcal{P}$ of n points and a set of $\mathcal{D}$ of m unit disks in the Euclidean plane. The objective is to select a set of disks $\mathcal{D}^* \subseteq \mathcal{D}$ of minimum cardinality that covers all the points. The problem is a geometric version of Set Cover and is NP-hard [9]. Nevertheless, several constant factor approximations have been developed and all could be used to compute a good approximation while balancing the trade-off between the approximation factor and the running time. For our purposes, we use the 18-approximation by Das et al [5] that has a runtime of $\mathcal{O}(n\log n + m\log m + mn)$. We note that better approximations are known, but using them in our framework could increase the total runtime. In each iteration, we increase the size of our set by $\mathcal{O}(\log k_{\mathsf{OPT}})\cdot k_{\mathsf{OPT}}$ and since $|S_1| \le 18\cdot k_{\mathsf{OPT}}$, we get the following:

Theorem 5 *Given X, T, α and R as above, we can find a set of sensors $S \subseteq X$ such that S $((1-1/\delta)\cdot\alpha)$-covers T within distance R and $|S| = \mathcal{O}(\log\delta\cdot\log k_{OPT})\cdot k_{OPT}$. The running time of the algorithm is $\mathcal{O}(\log\delta\cdot k_{OPT}\cdot mn\log m)$.*

For $\alpha-$ArtAng, we need to find a set of sensors $S_1 \subseteq X$ that guard T. To this extent, we can again employ the fact that the set of visibility polygons has finite VC-dimension [26]. Notice that finding a small S_1 that guards T corresponds to the hitting set problem for the set system made of sensors and visibility polygons of target locations. We therefore obtain a set of size $\mathcal{O}(\log k^*)\cdot k^*$, where k^* is the size of the smallest set of sensors from X that guard T. Since S_{OPT} also guards T, we have that $k^* \le k_{\mathsf{OPT}}$, so we are guaranteed to obtain a solution of size $\mathcal{O}(\log k_{\mathsf{OPT}})\cdot k_{\mathsf{OPT}}$. In each iteration, we increase the size of our set by $\mathcal{O}(\log k_{\mathsf{OPT}})\cdot k_{\mathsf{OPT}}$ and since $|S_1| = \mathcal{O}(\log k_{\mathsf{OPT}})\cdot k_{\mathsf{OPT}}$, we get the following:

Theorem 6 *Given polygons $Q \subseteq P$,X, T, and $\alpha \in [0,\pi/3]$ as above, we can find a set of sensors $S \subseteq X$ such that S $((1-1/\delta)\cdot\alpha)$-guards T and $|S| = \mathcal{O}(\log\delta\cdot\log k_{OPT})\cdot k_{OPT}$. The running time of the algorithm is $\mathcal{O}(\log\delta\cdot k_{OPT}\cdot mn\log m)$.*

We note that, in the case of $\alpha-$ArtAng, we improve on the result of Efrat et al [8], in that we approximate the α-coverage constraint to any constant factor while maintaining the same approximation factor of $\mathcal{O}(\log k_{\mathsf{OPT}})$. Moreover, our running times are comparable: $\mathcal{O}(nk_{\mathsf{OPT}}^4\log^2 n\log m)$ in [8] versus $\mathcal{O}(k_{\mathsf{OPT}}\cdot mn\log m)$ for our approximation. We note that their running time comes from the fact they do not directly use the bounded VC dimension of the set system. Instead, they use a previous algorithm designed by Efrat et al [7] for approximating the more general Art Gallery problem when the set of targets is restricted to vertices of that grid. When angle constraints are added, they adapt this algorithm to only consider vertices of the grid that also satisfy α-coverage. In our scenario in which targets have to be chosen from a discrete set, we do not need to impose a grid and can directly apply the Brönnimann and Goodrich algorithm [3]. For the case in which the sensors can be placed anywhere, their algorithm could be employed instead while maintaining the same approximation guarantees.

3 Other bi-criteria approximations for $(\alpha,R)-$AngDist

The geometric objects at the core of our method are wedges centered at targets whose central angles depend on α and ϵ. These ranges define the set of feasible locations from which we must choose a new set of sensors and are given as input to the subsequent Hitting Set problem. In the case of $(\alpha,R)-$AngDist, the distance constraints require that the sensors we pick also be within range R of the target, so our wedges become sectors through intersection with a disk of radius R centered at the target.

When it comes to solving the Hitting Set problem, we employ the instrumental results of Haussler and Welzl [10] and Brönnimann and Goodrich [3] which are based on the existence of good ϵ-nets for a given set system. We defer the intricacies of ϵ-nets to the long for of the paper[1] but mention now that, given an algorithm that computes in polynomial time an ϵ-net of size $\mathcal{O}((1/\epsilon) \cdot g(1/\epsilon))$, the algorithm of Brönnimann and Goodrich [3] returns a hitting set of size $\mathcal{O}(\tau \cdot g(\tau))$, where τ is the size of the optimal hitting set and g is a monotonically increasing sublinear function. Several good ϵ-net constructions are known when the underlying objects are geometrical. For the case of $(\alpha, R)-\mathsf{AngDist}$, we employ the canonical ϵ-net construction of Blumer et al [2] and Komlós et al [16] which exploits the fact that our ranges have constant VC dimension. This yields a $\mathcal{O}(\log k_{\mathsf{OPT}})$-approximation for Hitting Set on sectors of radius R.

In order to reduce the approximation factor, we consider relaxing the distance constraint and allowing the chosen sensors to be within distance $3R$ of the targets. In other words, we extend the radius of our sectors from R to $3R$. Inspired by the construction of Kulkarni and Govindarajan [18] for (unbounded) wedges, we then propose a deterministic rule for picking sensors and obtain a "relaxed" ϵ-net of size $\mathcal{O}(\frac{R_I}{R} \cdot \frac{1}{\epsilon})$, where R_I is the diameter of the largest enclosing ball of all possible sensor locations. The main difference between our construction and the one in [18] is the fact that the objects the latter considers are unbounded and, as such, allow for simpler grid-based constructions. In fact, this distinction is indeed the source of the additional $\frac{R_I}{R}$ factor that we incur in our bound.

To our knowledge, this is the first ϵ-net construction whose size depends linearly on $\frac{1}{\epsilon}$ and the ratio of the diameter of the input space to the size of the ranges (that is, when size can be appropriately defined). We note that the $\mathcal{O}(\frac{1}{\epsilon})$ construction of Pach and Woeginger [21] for translates of convex polygons does implicitly depend on solving the problem for points contained inside a bounded square. It is unclear, however, how to adapt their method for the case in which the ranges are sectors of similar radius but can have arbitrary central angles and orientations.

In order to overcome the dependency on $\frac{R_I}{R}$, we further employ the shifting technique of Hochbaum and Maass [11] to first partition our space into cells of bounded width and height and apply our ϵ-net construction to obtain good hitting sets in those restricted spaces. The analysis then yields an overall hitting set of size $\mathcal{O}(k_{\mathsf{OPT}})$ that achieves the desired $(\alpha - \epsilon)$-coverage and is within distance $3R$ of the targets. The complete argument is rather involved and we defer the exact details to extended version of the paper [1]. Formally, we get that:

Theorem 7 *Given X, T, $\alpha \leq \pi/3$ and R, we can find a set of sensors $S \subseteq X$ such that S $((1 - 1/\delta) \cdot \alpha)$-covers T within distance at most $3R$ and $|S| = \mathcal{O}(\log \delta) \cdot k_{\mathsf{OPT}}$. The running time of the algorithm is $\mathcal{O}(k_{\mathsf{OPT}} \cdot mn \log m \log n)$.*

Another interesting special case of $(\alpha, R)-\mathsf{AngDist}$ is the one in which $\alpha = 0$, since it requires us to place two distinct sensors within distance R of each target. A related problem is the fault tolerant k-suppliers problem as defined by Khuller et al [15] that requires us to select k suppliers such that each client has δ suppliers within an optimal distance r^* of it. While the objective in k-suppliers is to minimize the covering radius (as opposed to the number of sensors used), we will nevertheless exploit the connection and use it for $(\alpha, R)-\mathsf{AngDist}$. Under arbitrary metrics, Khuller et al. [15] develop a 3-approximation for fault tolerant k-suppliers: they select k sensors that are guaranteed to cover the clients within $3 \cdot r^*$. Karloff and then Hochbaum et al. [12] show that this factor is tight for the general k-suppliers problem (i.e. $\delta = 1$), unless P=NP. When the underlying space is $\mathbb{R}^d$ with the ℓ_2 metric, Nagarajan et al. [20] improve this factor to $(1+\sqrt{3})$ for Euclidean k-suppliers. They crucially use the observation that three clients who are pairwise more than $\sqrt{3} \cdot r^*$ apart from each other can never be covered by the same supplier within distance r^* and reduce the problem of finding k suppliers to that of computing a minimum edge cover. A similar approach can be used in our case, except the structure of the optimal solution corresponds to a b-edge cover. In the long form of the paper [1], we present the details of the analysis and guarantee that δ suppliers are within distance $(1 + \sqrt{3})r^*$ of each client:

Theorem 8 *There exists a polynomial time $(1 + \sqrt{3})$-approximation algorithm for the Euclidean fault tolerant k-suppliers in any dimension for arbitrary $\delta \geq 1$.*

The k-suppliers technique minimizes the covering radius by relying on the existence of k suppliers that cover all the clients within a guess radius R. Given such a guess radius, the algorithm itself never picks more than k sensors (i.e. without explicitly knowing the value of k). The binary search technique of Hochbaum and Shmoys [12] is then used to obtain guarantees with respect to the optimal covering radius r^*. The algorithm hence starts with a guess R and returns a set of k sensors that cover everything within distance $(1 + \sqrt{3}) \cdot R$. In our case, $k = k_{\mathsf{OPT}}$, so we will always pick at most k_{OPT} sensors. Since we already know the value of R, we get the following result as well:

Theorem 9 *Given X, T, and R as above, we can find a set of sensors $S \subseteq X$ such that S 0-covers T within distance $R \cdot (1 + \sqrt{3})$ and $|S| = k_{\mathsf{OPT}}$, where k_{OPT} is the cardinality of the smallest set of sensors that 0-covers T*

within distance R. The running time of the algorithm is $\mathcal{O}(n^2 \log n(m + n \log n))$.

References

[1] I. O. Bercea, V. Isler, and S. Khuller. Minimizing uncertainty through sensor placement with angle constraint. *CoRR*, abs/1211.0243, 2012.

[2] A. Blumer, A. Ehrenfeucht, D. Haussler, and M. K. Warmuth. Learnability and the Vapnik-Chervonenkis dimension. *J. ACM*, 36(4), 1989.

[3] H. Brönnimann and M. Goodrich. Almost optimal set covers in finite VC-dimension. *Discrete & Computational Geometry*, 14(1):463–479, 1995.

[4] N. Cressie. *Statistics for spatial data*. John Wiley & Sons, 2015.

[5] G. K. Das, R. Fraser, A. López-Ortiz, and B. G. Nickerson. On the discrete unit disk cover problem. In *WALCOM: Algorithms and Computation*, Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2011.

[6] S. Dreiseitl and M. Osl. Feature selection based on pairwise classification performance. In *Computer Aided Systems Theory-EUROCAST 2009*, pages 769–776. Springer, 2009.

[7] A. Efrat and S. Har-Peled. Guarding galleries and terrains. *Inf. Process. Lett.*, 100(6):238–245, 2006.

[8] A. Efrat, S. Har-Peled, and J. Mitchell. Approximation algorithms for two optimal location problems in sensor networks. In *BroadNets'05*, pages 714–723 Vol. 1, Oct 2005.

[9] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York,USA, 1979.

[10] D. Haussler and E. Welzl. Epsilon-nets and simplex range queries. In *Symposium on Computational Geometry*, pages 61–71, 1986.

[11] D. S. Hochbaum and W. Maass. Approximation schemes for covering and packing problems in image processing and vlsi. *J. ACM*, 32(1):130–136, Jan. 1985.

[12] D. S. Hochbaum and D. B. Shmoys. A unified approach to approximation algorithms for bottleneck problems. *J. ACM*, 33(3):533–550, May 1986.

[13] V. Isler, S. Khanna, J. Spletzer, and C. Taylor. Target tracking with distributed sensors: The focus of attention problem. *Computer Vision and Image Understanding Journal*, (1-2):225–247, 2005.

[14] A. Kelly. Precision dilution in triangulation based mobile robot position estimation. In *Intelligent Autonomous Systems*, volume 8, pages 1046–1053, 2003.

[15] S. Khuller, R. Pless, and Y. J. Sussmann. Fault tolerant k-center problems. *Theoretical Computer Science*, 242(12):237 – 245, 2000.

[16] J. Komlós, J. Pach, and G. Woeginger. Almost tight bounds for ϵ-nets. *Discrete Comput. Geom.*, 7(2):163–173, Mar. 1992.

[17] G. Kortsarz. On the hardness of approximating spanners. *Algorithmica*, 30(3):432–450, 2001.

[18] J. Kulkarni and S. Govindarajan. New ϵ-net constructions. In *CCCG'10*.

[19] J. Matoušek. *Lectures on Discrete Geometry*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2002.

[20] V. Nagarajan, B. Schieber, and H. Shachnai. The Euclidean k-supplier problem. In *Integer Programming and Combinatorial Optimization*, volume 7801, pages 290–301. Springer Berlin Heidelberg, 2013.

[21] J. Pach and G. Woeginger. Some new bounds for epsilon-nets. In *Proceedings of the sixth annual symposium on Computational geometry*, pages 10–15. ACM, 1990.

[22] E. Pekalska, A. Harol, C. Lai, and R. P. Duin. Pairwise selection of features and prototypes. In *Computer Recognition Systems*, pages 271–278. Springer, 2005.

[23] A. Savvides, W. Garber, S. Adlakha, R. Moses, and M. B. Srivastava. On the error characteristics of multihop node localization in ad-hoc sensor networks. In *Information Processing in Sensor Networks*, pages 317–332. Springer, 2003.

[24] A. Savvides, W. Garber, R. Moses, and M. Srivastava. An analysis of error inducing parameters in multihop sensor node localization. *Mobile Computing, IEEE Transactions on*, 4(6):567–577, Nov 2005.

[25] O. Tekdas and V. Isler. Sensor placement for triangulation based localization. *IEEE Tran. Automation Science and Engineering*, 7(3):681–685, 2010.

[26] P. Valtr. Guarding galleries where no point sees a small area. *Israel Journal of Mathematics*, 104(1):1–16, 1998.

[27] F.-b. Wang, L. Shi, and F.-y. Ren. Self-localization systems and algorithms for wireless sensor networks. *Ruan Jian Xue Bao(J. Softw.)*, 16(5):857–868, 2005.