Sampling Based Sensor-Network Deployment

Volkan Isler, Sampath Kannan and Kostas Daniilidis
GRASP Laboratory
Department of Computer and Information Science
University of Pennsylvania
Philadelphia, PA, USA
Email: {isleri,kannan,kostas}@cis.upenn.edu

***Abstract*— In this paper, we consider the problem of placing networked sensors in a way that guarantees coverage and connectivity. We focus on sampling based deployment and present algorithms that guarantee coverage and connectivity with a small number of sensors. We consider two different scenarios based on the flexibility of deployment. If deployment has to be accomplished in one step, like airborne deployment, then the main question becomes how many sensors are needed. If deployment can be implemented in multiple steps, then awareness of coverage and connectivity can be updated. For this case, we present incremental deployment algorithms which consider the current placement to adjust the sampling domain. The algorithms are simple, easy to implement, and require a small number of sensors. We believe the concepts and algorithms presented in this paper will provide a unifying framework for existing and future deployment algorithms which consider many practical issues not considered in the present work.**

I. Introduction

Recent achievements in low power processors, wireless networking, and sensor technology gave rise to the new field of wireless sensor networks. The technological and scientific issues underlying their development and deployment are very challenging and crosscut most of the areas in electrical engineering and computer science. The reader is referred to two of the main centers of activity, at the University of California at Berkeley [24] and Los Angeles [16], for an overview of the state of the art.

From a robotics perspective, sensor-network technology is very important for two reasons. First, we can view teams of robots with communication capabilities as sensor-networks. Therefore, the results in this paper directly apply to the deployment of mobile robots. Second, sensor-networks can help in many mobile robotics tasks [6]. For example, imagine for a foreign terrain mapping task, we could deploy sensor networks on a field before we deploy our robots. The robots then could use these sensors to help in many issues such as localization.

In this paper, we address the problems of coverage and connectivity using as few as possible static wireless sensor units, similar to the micro-routers of [35]. Examples include fire sensors deployed in a forest to detect forest fires or sensor fields for surveillance. By coverage we mean that every point in the environment is within the range of at least one sensor. Connectivity, on the other hand, is the requirement that every sensor can communicate with every other sensor. This is simply the connectivity requirement of the underlying graph whose two nodes (sensors) are adjacent only if their distance is less than a given reachability threshold.

The assumptions are very general because, both for coverage and connectivity, we do not apply any weighting of distances or any quality measurement. However, this simplified model enables us to answer the following question: *What is the minimum number of sensors needed for guaranteed coverage and connectivity?* Though terms like “smart dust” might imply that there is infinite availability of such sensors, there are several reasons why we need to estimate a minimal necessary number of them: Deployment of “infinite” sensors causes interference [9], the system becomes easily detectable by adversaries in surveillance tasks and placing many sensors increases non-invasiveness in natural and urban environments.

Our approach brings together results from sampling theory in machine learning and from the theory of geometric random graphs. When considering large sets of data, it is natural to ask what the most representative samples of them are. Let us, for now, be sloppy and associate samples with sensor units. For continuous signals, *representative* means being able to recover the original signal from the sampled ones. For recognition, it means that a classifier trained on these data will be accurate. In this paper, we use the idea of random sampling in geometric sets [27], [2]. The question we address is: How many samples (sensor units or robots) must be drawn such that every point in a possibly unknown scene is covered by least one sensor. Learning theory provides us a theorem which places an upper bound on the number of sensors for a given relative coverage. The critical concept here is that of Vapnik-Chervonenkis (VC) dimension [39], described in Section II. In wireless ad hoc networks of sensors, coverage is meaningless if sensors can not convey this information to each other. When the sensors are placed randomly, it is not clear whether the underlying graph will be connected. Here we employ the theory of random graph connectivity. Given a range specification, we can estimate the number of sensors sufficient for graph connectivity. Similarly, we can estimate the communication range required to ensure connectivity for a given number of sensors.

If the system can obtain coverage information during deployment, placing sensors to only uncovered areas results in a more efficient sampling scheme in terms of number of sensors. A second algorithm guarantees connectivity by adjusting the sampling domain with respect to not only the uncovered area but also the current communication range.

The problem of sensor placement has been studied extensively for non-connected sensors in robotics, computer vision, and computational geometry. Perhaps the most theoretically solid study of sensor placement considers the art gallery theorem –given a known environment, modeled by a simple polygon, the problem asks for guards to be placed so that the gallery is completely guarded [29], [36], [30]. It is known that $\lfloor \frac{n}{3} \rfloor$ guards are always sufficient to cover a polygon with n vertices but the problem of finding the minimum number of guards is NP-hard. There are several proposed approximations and a recent survey of them can be found in [15]. The study of art galleries has direct consequences for sensor networks that employ line-of-sight communication [1] and is also important for localization tasks [34], [40], [9]. View planning and sensor placement has also been addressed in the field of robotics [11], [18], [25] and computer vision [37].

In sensor networks literature, the term coverage has several meanings. Usually, however, coverage refers to the quality of service that can be offered by a particular network [26], [28]. Deployment, on the other hand, refers to the placement of sensors. The main issues when considering sensor deployment are maximizing network coverage and maintaining the connectivity of the network [9], [22], [31], [12], [40], [33].

For many sensor network applications, sampling has been the choice placement strategy [7], [17] for a variety of reasons. Generally, it is either impossible or very difficult to place sensors precisely. A typical example is the dissemination of sensors from an airplane onto foreign territory for surveillance purposes [10], under severe weather conditions.

It is also known that random geometric networks are robust to network failures [14]. In the case of mobile robot networks, the nodes of the network can indeed be at random positions at any given time. Assuming random node locations can help in the analysis of various issues as [17] discusses. In addition, sampling usually leads to simple, distributed, and low cost (both in running time and communication complexity) algorithms [9]. Sampling based algorithms provide a unifying framework for many environments, as opposed to more specialized algorithms that make many assumptions about the underlying environment. For example, art galleries require a polygonal representation of the environment. Moreover, sampling approaches are robust to uncertainties in the environment. An additional benefit is the availability of a well established theory of sampling and randomization [27], [2] which enables us to analytically evaluate the performance and design of the network.

In the next section, we present the basics of sampling theory, VC-dimension and ϵ-nets. Further, in Section III, we model the limitations of sensors using geometric objects and show how these concepts can be used for sensor placement. In addition, we address the connectivity, and present the trade-offs between connectivity and coverage. In Section V, we present two algorithms that incorporate the coverage and connectivity information obtained during placement and conclude the paper with a discussion of results presented.

II. Preliminaries

In this section, we introduce the basic concepts used in this paper: set-systems, VC-dimension, and ε-nets. A set system is a pair $(X, \mathcal{R})$ where X is a set and $\mathcal{R}$ is a collection of subsets of X.

Definition 1 ([39]): Given a set system $(X, \mathcal{R})$, let A be a subset of X. We say A is shattered by $\mathcal{R}$ if for all $Y \subseteq A$, $\exists R \in \mathcal{R}$ such that $R \cap A = Y$. The VC-dimension of $(X, \mathcal{R})$ is the cardinality of the largest set that can be shattered by $\mathcal{R}$.

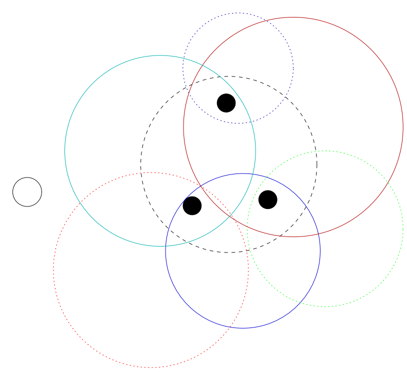

Fig. 1. 3 points shattered by 8 disks: for any subset of these points, there exists a disk that contains only those points and none other.

For example, consider the set system $(X, \mathcal{R})$, where X is a set of points on the plane and the subsets in $\mathcal{R}$ are obtained by intersecting the points in X with the set of *all* disks in the plane. Figure 1 shows 3 points shattered by 8 disks. In other words, for any subset of these points, there exists a disk that contains only those points and none other. However, this is not possible for 4 points as illustrated in the Figure 2. It is easy to verify that, if points are on their convex hull, two points opposite from each other can not be separated from the other two (Figure 2 left). If, on the other hand, there is a point inside the convex hull, it is not possible to separate the ones on the convex hull from the one inside (Figure 2 right). Therefore, no 4 points can be shattered and the VC-dimension of the disk set system is 3. In fact, this result generalizes to higher dimensions: The VC-dimension of the system that contains points in $\mathcal{R}^d$ and subsets intersected by balls in $\mathcal{R}^d$ has a VC-dimension $d + 1$ [20].

We proceed with defining the ε-net and the fundamental theorem justifying our approach.

Definition 2: $N \subseteq X$ is an ε-net for $(X, \mathcal{R})$, if for all $R \in \mathcal{R}$ with $|R| \geq \varepsilon$, we have $N \cap R \neq \emptyset$.
In other words, an $\varepsilon-$net is a set of points that meet every large set in $\mathcal{R}$. To give some preliminary idea of what is to come, the ε-net corresponds to the sampled sensor locations. For example, suppose our sensors have a sensing range and are therefore modeled with disks. If our sensors form an ε-net, this means that any activity within a

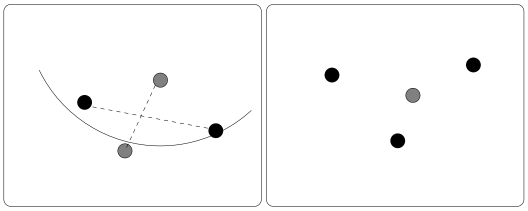

Fig. 2. 4 points can not be shattered: There are no disks that contain only the black points.

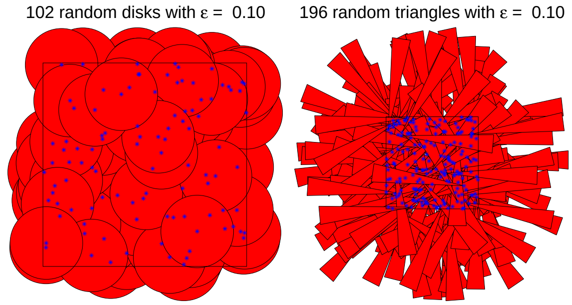

Fig. 3. Coverage with sampling

specified range (determined by ϵ) will be intersected (and hence detected) by one of our sensors. If we obtain our sensor locations with sampling, the following theorem tells us how many samples we will need:

Theorem 3 ([21]): Let d be the VC-dimension of $(X, \mathcal{R})$ and $\delta > 0$. Then, $O(\frac{d}{\varepsilon} \log \frac{d}{\varepsilon} + \frac{1}{\varepsilon} \log \frac{1}{\delta})$ points drawn at random from $\mathcal{R}$ is an ε-net with probability at least $1-\delta$.

Theorem 3 is a remarkable result which states that if a set system has a finite (i.e. constant) VC-dimension, then a small number of points (independent of the size of X) sampled from X is an ε-net with high probability.

III. Deployment with Sampling

In this section, we consider a scenario where deployment has to be accomplished in one step and no intermediate information becomes available during deployment. We describe how the results of the previous section can be used to design a sensor network for this case. The first step is to model the device constraints, typically using geometric objects. For example, one can model range constraints by representing sensors with disks or field of view constrains using triangles. Line-of-sight or visibility type of constraints can be modeled with star shaped polygons.[1]

Now suppose we would like to cover a square region A of unit area with disks of area ε. Consider the set system (X, R), where X is the (infinite) set of all disks with centers inside A and for every point p in A, there is a subset $R(p)$ in R that contains the disks that intersect p. Note that $R(p)$ itself can be represented by a disk of area ε. Therefore an ε-net for (X, R) intersects all the subsets in R, meaning that all points in A are covered by the ε-net.

It is worth emphasizing that all the sets mentioned above are infinite. The crucial property here is that the set system itself has a *constant* VC-dimension. Therefore using Theorem 3, we can sample a constant number of points (for a given ε), independent of the size of X (which is infinite most of the time) and obtain an $\varepsilon-$net with high probability.

Figure 3 shows some examples of coverage with random samples. In the left figure a unit area is covered with disks of radius $\varepsilon = 0.1$. In this simulation $d/\varepsilon \log d/\varepsilon = 102$ disks (with $d = 3$) were placed at locations chosen uniformly at random from $[0, 1]^2$. In the right figure, a similar experiment was performed for identical triangles that cover the same area. However in this case, 196 sensors are required as the triangle set system is a more complicated one with VC-dimension 5.

We conclude this section by giving examples of set systems with bounded VC-dimension: Half-plane set systems have VC-dimension 3, arbitrary convex k-gons have VC-dimension $2k + 1$ which can be obtained easily. However, determining the VC-dimension of set system can be rather tricky. A good example is the VC-dimension of visibility set systems that arise from simply-connected polygons [38], [23].

IV. Connectivity of Random Graphs

Random graphs are graphs whose properties such as the number of vertices or edges are determined randomly [8]. In this section we investigate the connectivity of sensors placed randomly using the theory of random graphs. Consider the following model for geometric random graphs: Given n, the number of vertices and ρ, range; the vertices of $G_n(\rho)$ are n points sampled uniformly from $[0, 1]^2$ and there is an edge between two vertices if the distance between them is less than ρ ([4], [3], see also [13]). The graph $G_n(\rho)$ can be used to model the properties of sensor networks obtained by sampling. Specifically, we would like to answer the following question: *Given n sensors deployed with uniform sampling, what is the communication range needed so that the underlying communication graph is connected?* Here, we assume that two sensors can communicate if the distance between them is less than or equal to some threshold value. Connectivity of sensor networks under this model has been studied previously in [17], [7], [19]. The theoretical study of the connectivity of geometric random graphs is fairly recent [32], [5]. The following theorem is adapted from [7] where it is also possible to find an extensive experimental study of similar properties of geometric random graphs:

Theorem 4: $G_n(\rho)$ is connected with probability p if $\rho \geq \sqrt{\frac{-\ln(1-p^{1/n})}{\pi n}}$.

Note that it is also possible to compute the number of sensors that ensure connectivity for a given communication range simply by holding ρ constant and solving for n in Theorem 4.

[1]A star-shaped polygon is a polygon which contains at least one point from which the entire boundary is visible.

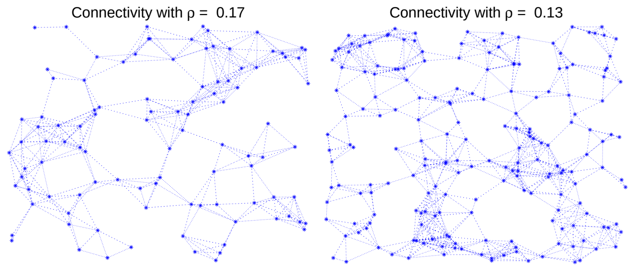

Fig. 4. Connectivity of the sensors in Figure 3 where the communication range ρ was chosen according to Theorem 4.

In Figure 4, the connectivity of sensors obtained in Figure 3 and communication ranges obtained by Theorem 4 (with $p = 0.99$) is shown.

In the important special case of range sensors, Theorems 3 and 4 enable us to establish the trade-off between the sensing range r and communication range ρ: Given a sensing range r, we can determine the area covered by the sensor ($\varepsilon = \pi r^2$) and we can apply Theorem 3 to obtain the number of sensors n required for coverage. We can then apply Theorem 4 to obtain the required communication range ρ.

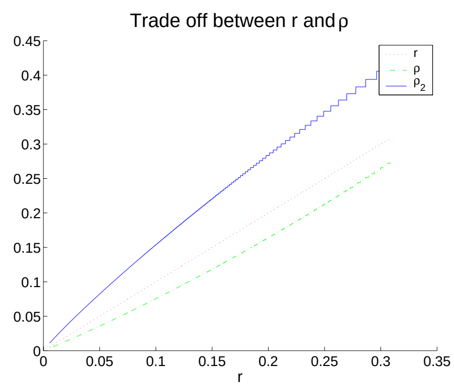

Fig. 5. The trade-off between the sensing range r and communication range ρ: For a given r value, ρ is obtained by using $n = \frac{d}{\epsilon}\log\frac{d}{\epsilon}$ in Theorem 4 and ρ_2 was obtained by $n = \frac{d}{\epsilon}$ with $\epsilon = \pi r^2$.

Figure 5 demonstrates another utility of sampling based placement. For a given $\epsilon = \pi r^2$, if we place $\frac{d}{\epsilon}\log\frac{d}{\epsilon}$ sensors and compute the communication range ρ, we see that the communication range required is smaller than r. Since typically communication range is much larger that the sensing range, this means the underlying network will be connected in such cases. However, if we had placed a smaller number say, $\frac{d}{\epsilon}$, of sensors and computed the communication range ρ_2, then we would need a bigger communication range. This is due to the decrease in the number of regions that are covered by multiple sensors which may cause loss of connectivity.

V. INCREMENTAL DEPLOYMENT

Usually sensor placement with sampling uses more sensors than the minimum required number. For example, for covering a unit area with disks of area ε, it is easy to see that $O(\frac{1}{\varepsilon})$ sensors are sufficient, whereas sampling requires $O(\frac{1}{\varepsilon}\log\frac{1}{\varepsilon})$ disks (see Figure 3).

The reason for the need for many sensors is that each sample is in a sense blind to the previous samples. In fact, if k discs have already covered an area B of the square, then the next random disc is expected to cover only $\frac{1-B}{\varepsilon}$ additional area.

Sometimes this redundancy is desirable, for example in order to have lots of overlapping observations for robust sensor fusion. Other times it is mandatory, for example when the sensors are disseminated from a plane under severe weather conditions. However, if the sensors are costly and we have more control over the placement (as in the case of mobile robots), an alternative approach based on sampling without replacement can be used to overcome this redundancy [2]. The general idea is to remove the regions that have already been covered from the sampling domain before choosing the next sensor position.

Suppose we would like to cover region A and i samples have already covered a region $B \subseteq A$. We select the next sample from $A \setminus B$. In other words, we no longer place sensors to regions that have already been covered. This algorithm is given in Table I. Figure 6-left shows a simulation of this deployment for the same conditions of the simulation illustrated in Figure 3-left. In this simulation, the number of sensors required decreased from 102 to 22 when sampled without replacement.

However, the algorithm SampleWithoutReplacement has a serious drawback. For the sensing range $r = 0.18$, if we choose the communication range $\rho = 0.21$, 20% more than r we get the communication graph given in Figure 6, middle-left. The drop in the number of sensors results in losing the connectivity of the graph!

It is possible to ensure connectivity by embedding it into the sampling process. Similar to the previous scenario, suppose we would like to cover region A and at any given time, the samples have already covered the region B and they are accessible from a region C. We simply select the next sample from $C \setminus B$ instead of $A \setminus B$. That is, we restrict the sampling domain to sensor locations that are accessible from the current sensors before we remove the regions that have already been covered from it. This algorithm is summarized in Table II. Figure 6 (middle-right, right) illustrates the placement with algorithm SampleWithoutReplacementwithComm and ensured connectivity for the same communication range $\rho = 0.21$. Another simulation where the area A that we would like to cover has a slightly more complicated structure can be found in Figure 7.

Even though they result in smaller number of sensors, the algorithms based on sampling without replacement have some drawbacks: First, they need more sophisticated data structures to keep track of covered areas. In the

[2]Consider drawing two balls from an urn one at a time. The term sampling without replacement is used to express the sampling scheme where the ball drawn first is not placed back into the urn prior to drawing the second one.

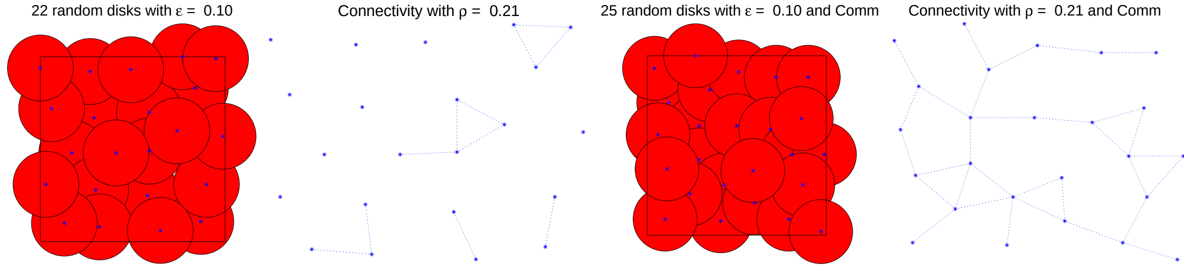

Fig. 6. Disks sampled without replacement. **: Left, Middle-left:**Even though removing already covered regions from the sampling domain results in fewer sensors, it also causes loss of connectivity. **Middle-right, Right:** Algorithm *SampleWithoutReplacementwithComm* ensures connectivity.

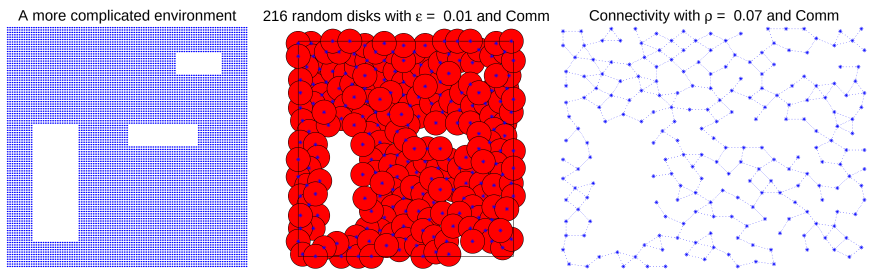

Fig. 7. Disks sampled without replacement with assured connectivity – A complex environment.

simulations of this section we solved this problem by putting a dense grid on the environment and keeping track of grid points covered. This results in decreased precision and increased running time proportional to the density of the grid. Second, it is expected that the placement obtained by these algorithms is more sensitive to errors in the actual positioning with respect to coverage and sensitivity due to the decrease in the number of sensors. Finally, sometimes it may be impossible to adjust the sampling range for some applications such as deploying sensors from an airplane under severe weather conditions.

SampleWithoutReplacement(A: Area to be covered)

area-Covered $\leftarrow \emptyset$
while *area-Covered* $\neq A$
  $s \leftarrow$ random sample from $A\setminus$ *area-Covered*
  area-Covered $\leftarrow$ *area-Covered* $\cup$ Range(s)

TABLE I

SAMPLING WITHOUT REPLACEMENT

SampleWithoutReplacementwithComm(A: Area to be covered)

s $\leftarrow$ random sample from A
area-Covered $\leftarrow$ Range(s)
sampling-Domain $\leftarrow$ CommunicationRange(s) $\setminus$ Range(s)
while *area-Covered* $\neq A$
  s $\leftarrow$ random sample from *sampling-Domain*
  area-Covered $\leftarrow$ *area-Covered* $\cup$ Range(s)
  sampling-Domain $\leftarrow$ *sampling-Domain* $\cup$ CommunicationRange(s)
  sampling-Domain $\leftarrow$ *sampling-Domain* $\setminus$ Range(s)

TABLE II

SAMPLING WITHOUT REPLACEMENT WITH COMMUNICATION CONSTRAINTS: RANGE(s) DENOTES THE SENSING RANGE OF THE SENSOR LOCATED AT s WHEREAS COMMUNICATIONRANGE(s) IS THE COMMUNICATION RANGE. CONNECTIVITY IS ENSURED BY EMBEDDING IT INTO THE SAMPLING PROCESS.

VI. CONCLUSIONS AND DISCUSSION

In this paper, we proposed a sampling based approach for the deployment of networked sensors. Our said goal was to guarantee coverage and connectivity using as few sensors as possible. Our model was very general: sensing range represented by simple shapes like disks or triangles, and connectivity depending on distance. However, this simple model enabled us to present bounds on the number of sensors required for coverage and connectivity.

Two scenarios have been addressed: concurrent deployment and incremental deployment. In the latter scenario, the deploying system receives information regarding coverage and connectivity after each deployed sensor. Afterwards, it applies sampling without replacement. The former scenario of concurrent deployment addresses deployment tasks for which no intermediate information is available. It is limiting in that it forces the system to decide on the number and location of the sensors a-priori. For this scenario, we have formulated coverage as a covering problem for set-systems. Given this, the ϵ-net theorem guarantees coverage using a small number of sensors, provided that the underlying set system has a finite VC-dimension.

Even though we have not addressed the quality of coverage, it may be possible to extend the algorithms of section V to tackle this issue using weighted sampling. Suppose for each point p, there is a weighting function $w(p)$ that returns how well that point is covered. If p is covered from a single sensor, $w(p)$ can be a function that is inversely proportional to the distance from the sensor. For coverage from multiple sensors, an averaging scheme can be used. If such a function is available, we can bias the sampling scheme such that points that are likely to increase the total weight more, are chosen with higher probabilities. In fact, at this point of the algorithm any information about the current state of the environment can be incorporated into the sampling by introducing bias. Further, these algorithms can be derandomized by considering new location candidates at each step and choosing the best among them. However, not only does such a locally greedy deployment scheme come with an additional computational cost, it is also unclear that it will result in a globally better placement, especially for complex environments.

With this paper, we would like to initiate a methodology in deployment where geometry of the sensing range meets sampling theory. It can be extended to several types of sensors with more complex coverage definitions [28]. We regard our method complementary to mobility approaches because it can provide an initial deployment guaranteeing coverage. This can allow for mobility to be used for tasks like tracking after initial deployment.

REFERENCES

[1] A. Acampora. Last mile by laser. *Scientific American*, 287:48–53, July 2002.

[2] P. K. Agarwal and S. Sen. Randomized algorithms for geometric optimization problems. *Handbook of Randomization (P. Pardalos, S. Rajasekaran, J. Reif, and J. Rolim, eds.), Kluwer Academic Publishers*, pages 151–201, 2001.

[3] M. J. B. Appel and R. P. Russo. The maximum vertex degree of a graph on uniform points in $[0, 1]^d$. *Advances in Applied Probability*, 29(3):567–581, 1997.

[4] M. J. B. Appel and R. P. Russo. The minimum vertex degree of a graph on uniform points in $[0, 1]^d$. *Advances in Applied Probability*, 29(3):582–594, 1997.

[5] M. J. B. Appel and R. P. Russo. The connectivity of a graph on uniform points on $[0, 1]^d$. *Statistics and Probability Letters*, 60:351–377, 2002.

[6] Maxim Batalin and Gaurav S. Sukhatme. Coverage, exploration and deployment by a mobile robot and communication network. *Telecommunication Systems*, 2004.

[7] C. Bettstetter. On the minimum node degree and connectivity of a wireless multihop network. In *Proceedings of the third ACM international symposium on Mobile ad hoc networking and computing*, pages 80–91. ACM Press, 2002.

[8] B. Bollabas. *Random Graphs*. Cambridge University Press, 2001.

[9] N. Bulusu, J. Heidemann, and D. Estrin. Adaptive beacon placement. In *Twenty-first International Conference on Distributed Computing Systems (ICDCS-21)*, page xxx. University of California, Los Angeles, IEEE Computer Society, April 2001.

[10] Peter I. Corke, Stefan E. Hrabar, Ron Peterson, Daniela Rus, Srikanth Saripalli, and Gaurav S. Sukhatme. Autonomous deployment and repair of a sensor network using an unmanned aerial vehicle. In *IEEE International Conference on Robotics and Automation*, 2004.

[11] J. Cortés, S. Martínez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1327–1332, Arlington, VA, May 2002.

[12] S. S. Dhillon and K. Chakrabarty. A fault-tolerant approach to sensor deployment in distributed sensor networks. *Proc. Army Science Conference, Paper ID: JP-05*, 2002.

[13] J. Diaz, J. Petit, and M. Serna. Random geometric problems on $[0, 1]^2$. In J. Rolim, M. Luby, and M. Serna, editors, *Randomization and Approximation Techniques in Computer Science*, volume 1518 of *Lecture Notes in Computer Science*, pages 294–306. Springer-Verlag, 1998.

[14] J. Diaz, J. Petit, and M. J. Serna. Faulty random geometric networks. *Parallel Processing Letters*, 10(4):343–358, 2000.

[15] S. Eidenbenz, C. Stamm, and P. Widmayer. Inapproximability results for guarding polygons and terrains. *Algorithmica*, 31:79–113, 2001.

[16] D. Estrin, R. Govindan, J. S. Heidemann, and S. Kumar. Scalable coordination in sensor networks. In *ACM Mobile Computing and Networking (MOBICOM)*, pages 263–270, 1999.

[17] A. Farag. Scalable analysis and design of ad hoc networks via random graph theory. In *Proceedings of the 6th international workshop on Discrete algorithms and methods for mobile computing and communications*, pages 43–50. ACM Press, 2002.

[18] H.H. Gonzalez-Banos and J.C. Latombe. A randomized art-gallery algorithm for sensor placement. *Proc. ACM Symp. on Computational Geometry (SoCG'01)*, 2001.

[19] P. Gupta and P. R. Kumar. Critical power for asymptotic connectivity in wireless networks. In *Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of W.H. Fleming*. Birkhauser, Boston, 1998.

[20] P. Gupta and P. R. Kumar. Internets in the sky: The capacity of three dimensional wireless networks. *Communications in Information and Systems*, 1:33–49, 2001.

[21] D. Haussler and E. Welzl. ϵ-nets and simplex range queries. *Discrete Comput. Geom.*, 2:127–151, 1987.

[22] A. Howard, M. J. Mataric, and G.S. Sukhatme. An incremental self-deployment algorithm for mobile sensor networks. *Autonomous Robots, Special Issue on Intelligent Embedded Systems*, 13(2):113–126, 2002.

[23] V. Isler, S. Kannan, K. Daniilidis, and P.Valtr. VC-dimension of exterior visibility. *"IEEE Trans. Pattern Analysis and Machine Intelligence"*, 26(5):667–671, 2004.

[24] J. M. Kahn, R. H. Katz, and K. S. J. Pister. Mobile networking for "smart dust". In *ACM Mobile Computing and Networking (MOBICOM)*, pages 271–278, 1999.

[25] L. E. Kavraki, J. C. Latombe, R. Motwani, and Pr. Raghavan. Randomized query processing in robot path planning. In *ACM Symposium on Theory of Computing*, pages 353–362, 1995.

[26] X. Y. Li, P. J. Wan, and O. Frieder. Coverage in wireless ad-hoc sensor networks, April 2002.

[27] J. Matousek. Geometric computation and the art of sampling. In *IEEE Symposium on Foundations of Computer Science*, 1998.

[28] S. Meguerdichian, F. Koushanfar, M. Potkonjak, and M. B. Srivastava. Coverage problems in wireless ad-hoc sensor networks. In *INFOCOM*, pages 1380–1387, 2001.

[29] J. O'Rourke. *Art Gallery Theorems And Algorithms*. Oxford University Press, 1987.

[30] J. O'Rourke. Visibility. In Jacob E. Goodman and Joseph O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 25, pages 467–480. CRC Press LLC, 1997.

[31] D. Payton, M. Daily, M. Howard, and C. Lee. Pheromone robotics. *Autonomous Robots*, 11(3):319–324, November 2001.

[32] M. D. Penrose. On k-connectivity for a geometric random graph. *Random Structures and Algorithms*, 15(2):145–164, 1999.

[33] Sameera Poduri and Gaurav S. Sukhatme. Constrained coverage for mobile sensor networks. In *IEEE International Conference on Robotics and Automation*, 2004.

[34] A. Savvides, C. Han, and M. Strivastav. Dynamic fine-grained localization in ad-hoc networks of sensor s. In *ACM SIGMOBILE*, Rome, Italy, July 2001.

[35] S. D. Servetto and G. Barrenechea. Constrained random walks on random graphs: Routing algorithms for large scale wireless sensor networks. In *1st ACM International Workshop on Wireless Sensor Networks and Applications,*. ACM MobiCom, September 2002.

[36] T. Shermer. Recent results in art galleries. *Proceedings of the IEEE*, 80:1384–1399, 1992.

[37] K.A. Tarabanis, P.K. Allen, and R.Y. Tsai. A survey of sensor planning in computer vision. *IEEE Trans. Robotics and Automation*, 11:86–104, 1995.

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

[39] V. N. Vapnik and A. Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264–280, 1971.

[40] Y. Zou and K. Chakrabarty. Sensor deployment and target localization in distributed sensor networks. *accepted for publication in ACM Transactions on Embedded Computing Systems*, 2003.