More Is More: The Benefits of Denser Sensor Deployment

MATTHEW P. JOHNSON, DENIZ SARIÖZ, AMOTZ BAR-NOY,
and THEODORE BROWN, City University of New York
DINESH VERMA and CHAI WAH WU, IBM T. J. Watson Research Center

Positioning disk-shaped sensors to optimize certain coverage parameters is a fundamental problem in ad hoc sensor networks. The hexagon lattice arrangement is known to be optimally efficient in the plane, even though 20.9% of the area is unnecessarily covered twice, however, the arrangement is very rigid—any movement of a sensor from its designated grid position (due to, e.g., placement error or obstacle avoidance) leaves some region uncovered, as would the failure of any one sensor. In this article, we consider how to arrange sensors in order to guarantee multiple coverage, that is, k-coverage for some value $k > 1$. A naive approach is to superimpose multiple hexagon lattices, but for robustness reasons, we may wish to space sensors evenly apart.

We present two arrangement methods for k-coverage: (1) optimizing a Riesz energy function in order to evenly distribute nodes, and (2) simply shrinking the hexagon lattice and making it denser. The first method often approximates the second, and so we focus on the latter. We show that a density increase tantamount to k copies of the lattice can yield k'-coverage, for $k' > k$ (e.g., $k = 11, k' = 12$ and $k = 21, k' = 24$), by exploiting the double-coverage regions. Our examples' savings provably converge in the limit to the $\approx 20.9\%$ maximum. We also provide analogous results for the square lattice and its $\approx 57\%$ inefficiency (e.g., $k = 3, k' = 4$ and $k = 5, k' = 7$) and show that for multi-coverage for some values of k', the square lattice can actually be more efficient than the hexagon lattice.

We also explore other benefits of shrinking the lattice: Doing so allows all sensors to move about their intended positions independently while nonetheless guaranteeing full coverage and can also allow us to tolerate probabilistic sensor failure when providing 1-coverage or k-coverage. We conclude by construing the shrinking factor as a budget to be divided among these three benefits.

Categories and Subject Descriptors: C.2.1 [**Computer Communication Networks**]: Network Architecture and Design

General Terms: Theory, Design

Additional Key Words and Phrases: Sensor networks, sensor coverage, k-coverage, characterization

ACM Reference Format:
Jahnson, M. P., Sariöz, D., Bar-Noy, A., Brown, T., Verma, D., and Wu, C. W. 2012. More is more: The benefits of denser sensor deployment. ACM Trans. Sensor Netw. 8, 3, Article 22 (July 2012), 19 pages.
DOI = 10.1145/2240092.2240096 http://doi.acm.org/10.1145/2240092.2240096

1. INTRODUCTION

An ad hoc sensor network is composed of sensing devices which can measure or detect features of their environment and can communicate with one other and possibly with other devices that perform data fusion. One of the problems motivated by ad hoc sensor

Portions of this work appeared in preliminary form in Wu and Verma [2008] and Johnson et al. [2009].
Authors' addresses: M. P. Johnson, D. Sariöz, A. Bar-Noy, and T. Brown, Department of Computer Science, The Graduate Center, City University of New York; A. Bar-Noy, Department of Computer and Information Science, Brooklyn College, City University of New York; D. Verma and C. W. Wu, IBM T. J. Watson Research Center; corresponding author's email: mpjohnson@gmail.com.

DOI 10.1145/2240092.2240096 http://doi.acm.org/10.1145/2240092.2240096

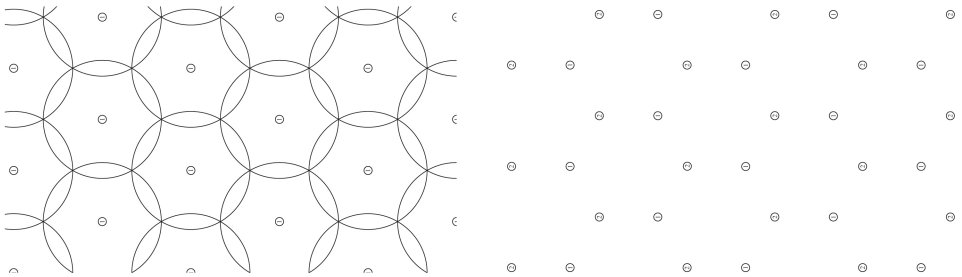

(a) A rigid 6-grid (hexagon lattice). (b) A 2-covering of the plane.

Fig. 1. Two grid arrangements.

networks is how to position sensors in order to maximize coverage, or equivalently, to minimize the number of sensors required to cover a given area. In such problems, for example, in monitoring applications, sensors must be positioned so that every point in the region of interest is observable by at least one sensor. In the boolean model of coverage, a sensor is able to observe all points within a certain distance r of its location, that is, a disk of radius r. The problem them becomes where to place sensors so that disks centered on them will cover the area. More generally, we may require k-coverage in which each point must be covered by at least k sensors. This may be required for robustness reasons, in order to perform triangulation, or to obtain multimodal data. Even in applications requiring only 1-coverage, it may be desirable for power management reasons to deploy multiple lattices' worth of sensors and then alternate the sensors between active and dormant states (duty cycling [Kansal et al. 2007]), since wireless sensors are typically powered by finite-lifetime batteries.

For the 1-coverage problem, placing homogeneous sensors in a hexagon grid pattern (see Figure 1(a)) is known to be optimal [Kershner 1939; Brass et al. 2005]. Every point is covered at least once, and no other covering structure can achieve this property with fewer sensors. Notwithstanding this optimality, there are aspects of this configuration that could be regarded as deficiencies. Since the sensors' sensing regions are circular and therefore not tileable, there inevitably are regions (lens-shaped areas in Figure 1, totaling about 20.9% of the total area) that are superfluously covered by two sensors. Second, the configuration is very sensitive. Any movement of one sensor from its designated grid position leaves some region uncovered. Third, any sensor failure will similarly cause a gap in coverage.

A 2-covering obtained by overlaying two translated copies of the optimal 1-covering is shown in Figure 1(b). We give similar coverings for $k > 2$. For robustness reasons, we obtain a k-covering by arranging k translated copies of the hexagonal lattice to maximize the minimum pairwise distance d_{min} between sensors. Such a k-fold covering has the additional advantage of being easy to understand and construct by virtue of being based on the familiar 1-coverage lattice. We show that this reduces to distributing k points evenly on the torus. This problem can be approached by minimizing a Riesz s-energy objective function with a sufficiently large parameter s. We show examples in which the s value needed varies based on the value of k. It turns out that for larger values of k, the solution provided by optimizing Riesz energy is often simply a denser hexagon lattice or is close to this. The second method we consider is therefore simply shrinking the lattice, that is, deploying sensors in a more densely packed hexagon arrangement, thus avoiding the optimization calculation.

In the optimal hexagonal covering, approximately 20.9% of the overlap offers no benefit in terms of providing 1-coverage. When the goal is multi-coverage, however, it may be possible to exploit this inherent redundancy. We show that for many values

of k, shrinking to increase the number of sensors by a factor of k (often corresponding to combining k separate 1-covers as previously mentioned) yields k'-coverage for some $k' > k$. We also show that the efficiency ratio k'/k approaches the maximal ≈ 1.209 as $k \to \infty$, for example, $12/11 \approx 1.09,\ 24/21 \approx 1.14,\ 58/49 \approx 1.18$. In this sense, the $(k' - k)/k$ ratio is an efficiency gain over the standard hexagon lattice. Perhaps counterintuitively, we provide similar results for any regular lattice, such as the square lattice in which all of its $\approx 57\%$ overlap is recovered with efficiency ratios of $4/3$ and $7/5$ for small k. This suggests a sense in which the square lattice *is* competitive with the hexagon lattice. For some multi-coverage requirements, that is, a fixed area and a desired k', a shrunken square arrangement will actually use fewer sensors than the corresponding shrunken hexagon arrangement.

For example, in the case of $k = 5$, we obtain $k' = 7$, which yields $k'/k = 7/5 = 1.4$ and $1.57/(k'/k) \approx 1.12$, which is better than the efficiency 1.209 of the shrunken hexagon lattice.

We also consider other benefits of the shrunken lattice arrangement. First, it enjoys greater robustness against placement error, which has practical advantages. If the designated position for the sensor is not ideal because of difficult terrain, for example, or because of man-made obstructions, it is natural to place the sensor nearby and hope to continue to provide full coverage. Or, if (some of) the sensors are mobile, they may take advantage of any available wiggle room to move about and avoid detection. By shrinking the lattice, the sensors would still be deployed according to a hexagonal structure but with a denser structure allowing the sensors some freedom of movement without invalidating the full coverage guarantee. We call the area in which a sensor is free to move its *wiggle area* and explore its shape when, for fairness reasons, each sensor is obliged to cover the points in its Voronoi cell. In this case, each sensor is free to move independently within its wiggle area. We note that this area need not be a disk, though the maximum *wiggle disk*,[1] corresponding to a sensor's maximum allowable displacement in any direction would be of particular practical interest. Conversely, we compute the lattice density (degree of shrinking) required in order to support a certain desired amount of allowable placement error.

Another benefit of shrinking concerns sensor reliability. Depending on sensor type, we may expect a certain percentage of deployed sensors to fail. Alternatively, in event detection applications, there may be a known false negative probability. These considerations provide a third kind of benefit to shrinking the sensor network. By sufficiently shrinking the lattice (corresponding to k times as many sensors), we can provide full coverage (in expectation) in the face of a nonzero failure probability, that is, fault tolerance, and even spend the $k \to k'$ efficiency gain on an allowable failure probability.

Model. We assume a boolean sensor coverage model with sensing coverage areas as disks. While this is a simplified model, it can be used to provide conservative bounds in applications with more complex anisotropic or probabilistic sensor models [Bai et al. 2006] and is interesting enough to be widely studied on its own. Throughout the article, we assume a large convex area of interest which allows us to ignore the edges. If connectivity is also required, then we assume that $r_c \geq 2r_s$, that is, the communication range is at least double the sensing range. This is known to suffice for inferring that full coverage implies connectivity [Zhang and Hou 2005], and so we do not further discuss connectivity in this article.

Organization. In Section 2, we present related work. In Section 3, we explore a Riesz energy technique for finding robust k-coverage arrangements before turning to the

[1]Defined as the largest disk centered at the sensor that is contained within the wiggle area.

simpler method of simply shrinking the lattice. We demonstrate the efficiency of the shrunk lattice for multi-covering, showing that the equivalent of k copies can provide k'-coverage. We give a series of examples, converging on the limit, in which the percentage of recovered area approaches the $\approx 20.9\%$ maximum. Second, we quantify the flexibility benefits of shrinking the hexagon lattice, characterize the resulting wiggle area, and prescribe the amount of shrinking required to support a given flexibility requirement (Section 4). Third, we study the fault tolerance provided by shrinking (Section 5). We conclude by construing the shrinking factor as a budget divided between these benefits (Section 6).

2. RELATED WORK

There have been efforts in various communities to characterize, find approximate solutions to, and analyze the intrinsic properties of the coverage problem based on the boolean model. Brass [2007] focused on theoretical bounds for various coverage requirements using static as well as mobile sensors, all under the boolean model. Zhang and Hou [2005] gave a distributed algorithm that attempts to minimize the number of sensors needed to obtain full coverage by minimizing the amount of intersection areas of sensor disks. Gupta et al. [2006] gave an $O(\log n)$ set cover-style greedy algorithm for the setting in which sensor locations must be chosen for a given discrete set, which was extended to k-coverage in Zhou et al. [2004]. Funke et al. [2007] proved a matching $\Omega(\log n)$ approximation ratio lower bound. The setting in which only a discrete set of clients must be covered is also well studied [Alt et al. 2006]. In this article, however, our focus is on the setting of full coverage with unrestricted sensor placement.

Some research combines coverage goals with connectivity or other constraints. Bai et al. [2006], for example, provided sensor deployments yielding 1-coverage and either 1-connectivity or 2-connectivity. Potential field methods have been used to distribute sensors so that in the resulting graphs, sensor nodes have a minimum degree [Poduri and Sukhatme 2004].

Bar-Noy et al. [2009] discussed some other problems related to inexact sensor placement. A connection was also drawn between probabilistic placement error and probabilistic (nonbinary) sensing models. Zou and Chakrabarty [2004] combined probabilistic placement error and probabilistic sensing models in seeking a high-confidence full cover of grid points. They presented experimental results of a heuristic algorithm that greedily selects grid points with low-current coverage probability.

There is a large literature on k-coverage. Wang and Tseng [2008] provided k-coverage arrangements for binary and probabilistic sensors. They noted the presence of overlap among disks but take advantage of it only in the setting of small r_c. For large r_c, which is our focus here, they simply duplicate the 1-covering k times. Bartolini et al. [2008] gave distributed algorithms for assembling sensors into a hexagon lattice, whereas we take the ability to position sensors as a given and study the benefits of such arrangements. Focusing on probabilistic sensors, Zhang and Hou [2006] gave bounds on density for guaranteeing full k-coverage with high probability with various arrangements. This non-local guarantee requires increasing sensor density by a logarithmic factor. Our focus in Section 5, however, is purchasing a local coverage guarantee using only a constant factor increase in density. Huang and Tseng [2005], for example, gave efficient algorithms for determining the maximum value k for which a given sensor configuration provides k-coverage, supporting both unit-disk sensors and non-unit-disk sensors.

Finally, Pach and Tóth [2007] studied the problem of decomposing a k-cover into a maximum number of disjoint 1-covers. They noted that when the region of interest is a disk, the only positive result that has been claimed is that any 33-cover may be decomposed into two disjoint 1-covers. Abrams et al. [2004] solved a related but relaxed problem in which a set of sensors is partitioned into subsets with the objective

of maximizing the sum of the (partial) covers of the sets. We remark that our second motivation for shrinking the lattice is the reverse of this, that is, of combining multiple 1-covers.

3. *K*-COVERAGE

Under k-coverage, each point in the region R is covered by at least k different sensors. A 2-covering (see Figure 1(b)) can be obtained by concatenating two optimal 1-coverings where there is a translation between the two 1-coverings. This approach enjoys a robustness property, namely, the sensors are placed maximally far away from one another. We now explore how k-coverings for higher k can be achieved that also satisfy this property. Sufficiently increasing the density of the sensor network lattice will yield a k-fold coverage. Furthermore, we will see how to exploit the overlap in a single lattice and improve the efficiency of the obtained coverage relative to the number of sensors used.

3.1. Riesz Energy

We construct a k-covering by superimposing k translated copies of the hexagonal 1-covering. Because of the lattice structure of the optimal 1-covering, each sensor within the 1-covering can be construed as lying on the surface of a torus, that it, the 1-covering is reduced to a single sensor located on the surface of a torus. This can be seen as follows. The lattice can be deformed into a regular square lattice and each sensor is located in the same place in each lattice square. This means that if we glue the right and left edges of the square and the top and bottom edges of the square (i.e., we form a quotient of the square lattice by identifying the left with the right edge and the top with the bottom edge), we obtain a torus with a single sensor.

Since this lattice is the same for all the 1-coverings, the tori for the different 1-coverings can be thought of as the same object. The distance metric d on the torus is induced by the distance metric d_R in the region R. Since each point on the torus corresponds to a lattice in the plane, when unwrapped onto R, d is the Hausdorff distance between two lattices. (This follows readily from the definition of the Hausdorff distance [Rockafellar and Wets 2009].) In this case, for robustness reasons, that is, to maximize $d_{\min} = \min_{1\le i<j\le k} d(a_i, a_j)$, the problem is reduced to dispersing k points on a torus.[2] This problem has been studied using the approach of Riesz energy minimization [Hardin and Saff 2004]. The s-Riesz energy of a set of points $\{a_i\}$ is defined as

$$E_s = \sum_{i\neq j} \frac{1}{d(a_i, a_j)^s}$$

for a distance metric d and exponent $s > 0$. For $s = 0$, the Riesz energy is defined as

$$E_0 = \sum_{i\neq j} \log \frac{1}{d(a_i, a_j)}$$

Womersley [Womersley] reports experiments in best-effort (local search) optimization algorithms to minimize E_s for $s \to \infty$, which experimentally appear to often produce solutions with maximum $d_{\min}$. He also reports that Quasi-Newton-based optimization algorithms have been developed for cases in which the number of points is in the thousands.

[2]The problem of maximizing the pairwise distance between a fixed number of points, known as a *best packing* of the points [Hardin and Saff 2004], has previously been studied on the sphere, where it is known as Fejes Tóth's problem [Brass et al. 2005]. Exact solutions are known for $k \le 12$ and $k = 24$, but no general solutions are known.

(a) A 3-covering of the plane with $d_{\min} = \frac{1}{\sqrt{3}}$. Note that the union of 3 hexagonal 1-coverings forms a hexagonal covering.

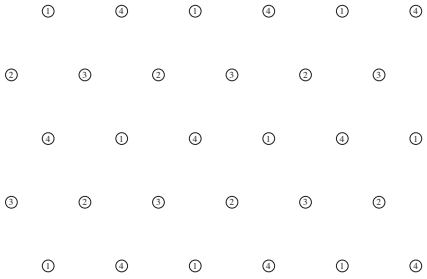

(b) A 4-covering of the plane with $d_{\min} = \frac{1}{2}$. Note that the union of the 4 hexagonal 1-coverings is again a hexagonal covering.

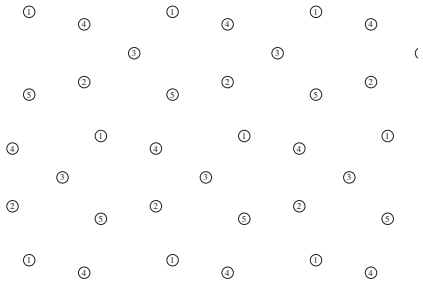

(c) A 5-covering of the plane with $d_{\min} = 0.3937$.

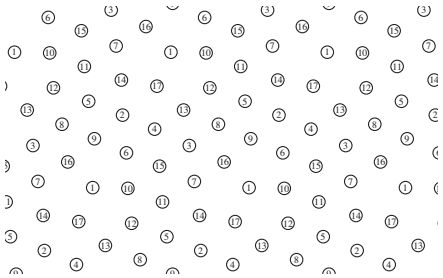

(d) A 17-covering of the plane with $d_{\min} = 0.2198$.

Fig. 2. Computed k-coverings for various values of k.

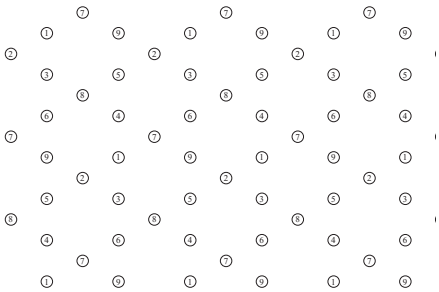

(a) Optimized using $s = 1$. $d_{\min} = \frac{1}{2\sqrt{3}}$.

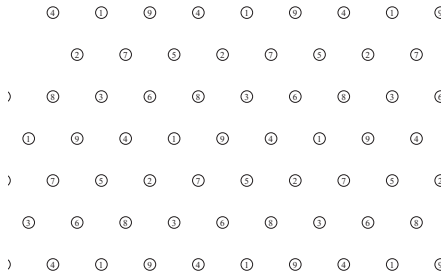

(b) Optimized using $s = 2$. $d_{\min} = \frac{1}{3}$.

Fig. 3. Two 9-coverings of the plane.

3.2. Experimental Results

By minimizing the Riesz energy E_s, we obtain configurations of k-coverings for various values of k. In our experiments, we choose $s = 10$ and $d_R(a_i, a_j) = \|a_i - a_j\|_2$ to be the Euclidean distance on the plane. For 2-coverings, we obtain the same result as shown in Figure 1(b). The results for k-coverings for various values of $k > 2$ are shown in Figure 2. In these figures, sensors labeled with the same number belong to the same hexagonal 1-covering. Note that the union of all the sensors can form a regular geometric structure and, in some instances, a hexagonal lattice. For instance, for $k = 3, 4, 7, 9$ (see Figures 2(a), 2(b), and 3(b)), the union of all the sensors form a hexagonal

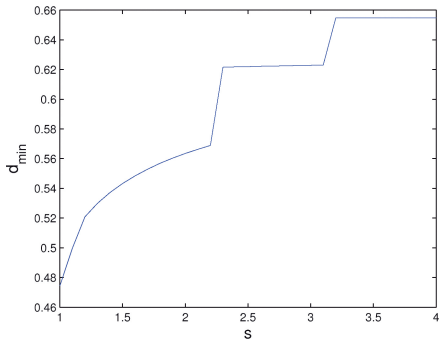

Fig. 4. $d_{\min}$ obtained in a 7-covering by minimizing E_s for various s.

lattice—the optimal 1-covering. (The configuration for $k = 3$ is used in Sun et al. [2005].) Equivalently, the hexagon lattice can be decomposed into k disjoint hexagon lattices scaled by factor $\sqrt{k}$. Such values of k are said to "divide the lattice".

Let Q' be the set of all such hexagon lattice divisors. Observe that by repeated decomposition, it follows that for any $k = \Pi_{q_i \in Q} q_i^{m_i}$ for $m_i \in \mathbb{N}$, we also have $k \in Q'$. It is known [Bernstein et al. 1997] that $Q' = \{a^2 + ab + b^2 | a, b \in \mathbb{N}_1\} = \{3, 4, 7, 9, 12, 13, 16, \ldots\}$ (the Loeschian numbers other than 0,1),[3] several of which we saw in the figures.[4]

For arbitrary integers k', we can round up to the first $k = \Pi_{q_i \in Q'} q_i^{m_i}$ bounding k' and use the resulting optimal k-covering (in general suboptimally) as the k'-covering. Then, the optimal $d_{\min}$ for a k'-covering is lower bounded by $\frac{1}{\sqrt{k}}$ and upper bounded by $\frac{1}{\sqrt{k'}}$ (due to the optimality of the hexagonal lattice).

3.2.1. Choosing the Exponent s. How large should the value of the exponent s be when computing the Riesz energy E_s? In order to approximate the energy function that maximizes d_{min}, the parameter s should be chosen to be large. However, a large s leads to large gradients and can cause problems in the numerical computations. On the other hand, in our experiments, we found that choosing small values of s can produce suboptimal results. For instance, we show in the subfigures of Figure 3 the 9-coverings produced by optimized Riesz energy with $s = 1$ and $s = 2$, respectively. According to the a preceding discussion, the latter subfigure illustrates the optimal configuration maximizing d_{min}. Numerical experiments show that there is a transition from one configuration to the other configuration around $s = 1.2$.

In fact, by the so-called Poppy-seed Bagel Theorem [Hardin and Saff 2004], as the number of points k approaches infinity, minimizing the s-energy for any value $s \geq 2$ will distribute the points uniformly on the surface of the torus, thus maximizing $d_{\min}$. This result only refers to k in the limit, however, and may not hold for small k. For small values of k, exponent s may need to be large in order to maximize $d_{\min}$. In practice, it could be difficult to know for certain whether the s value chosen is large enough for the chosen k. For example, for $k = 7$, Figure 4 shows $d_{\min}$ as E_s is minimized for various values of s. We see that $d_{\min}$ reaches the optimal value of $\sqrt{3}/\sqrt{7}$ at $s \approx 3.2$ after passing through several other configurations. On the other hand, for $k = 9$, the only transition occurs at $s \approx 1.2$. As we noted before, we cannot simply increase s until we obtain a hexagonal solution, since the optimal configuration need not be hexagonal.

[3]http://oeis.org/A003136.

[4]Members of Q' not expressible as products of members of Q' are known as *prime divisors* of the lattice: $\{3, 4, 7, 13, 19, 25, \ldots\}$ (the norms of Eisenstein-Jacobi primes; http://oeis.org/A055664).

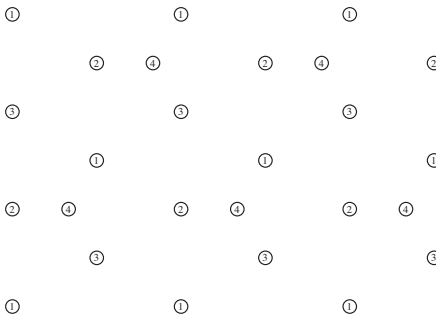

(a) $d_{\min} = \frac{1}{3}$. This covering is obtained by adding an additional 1-covering to the 3-covering in Figure 2(a).

(b) The base 1-covering is the square grid lattice and $d_{\min} = \frac{1}{2\sqrt{2}}$.

Fig. 5. Two 4-coverings of the plane.

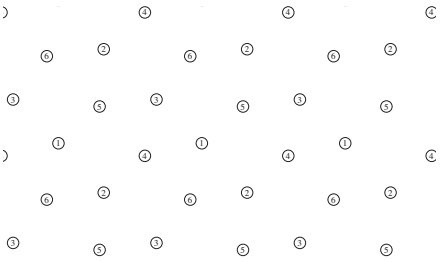

(a) A 6-covering of the plane. The base 1-covering is the square grid lattice and $d_{\min} = 0.2708$.

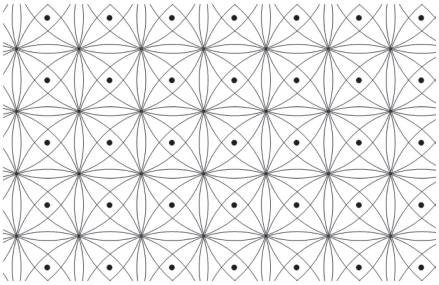

(b) The square lattice shrunk by a factor of five, providing 7-covering.

Fig. 6. Two grid arrangements.

3.3. Incremental *k*-Coverings and Other Base Coverings

Consider the scenario where a k-covering is in place and we want to expand the sensor network to a k'-covering ($k' > k$) by adding more sensors: where should the new sensors go. Again, the same optimization approach can be used with the additional constraint that the sensors in the original k-covering remain fixed. In Figure 5(a), we show a 4-covering obtained by extending the 3-covering in Figure 2(a). Compare this with the optimal 4-covering in Figure 2(b).

So far we have applied this algorithm to the case where the underlying 1-covering (i.e., the *base* covering) is the hexagonal lattice. This approach can also be applied to other base 1-coverings. For instance, Figure 5(b) shows a 4-covering obtained by translating four square grid 1-coverings and optimizing the arrangement by minimizing E_{10}. We see that the totality of sensors form a hexagonal lattice. A more irregular-looking 6-covering based on six square grid lattices is shown in Figure 6(a).

There is a close relationship between these k-coverings and dispersed-dither digital halftoning [Ulichney 1997]. In dispersed dither, a pattern is tiled on the plane such that it forms a visually pleasing pattern. The dispersed-ordered dither patterns (e.g., blue noise dispersed dither [Spaulding et al. 1997]) are regular patterns that are repeated and similar to k-coverings, such as in Figure 6(a). The stacking constraint in these dither patterns dictates that a pattern with more dots is a superset of a

pattern with fewer dots—the same constraint as with the incremental k-coverings discussed previously. Indeed, many algorithms for generating blue noise dither patterns are based on energy minimization [Wu et al. 2003]. In a k-covering, however, the sensors can be positioned at any point on the plane, resulting in a linear programming problem, whereas in a dispersed-dither pattern, the dots are constrained to lie on the printer addressability grid (e.g., a 600 dpi or a 1,200 dpi grid), resulting in an integer programming problem.

3.4. Definitions

Definition 1. Let r be the sensing radius. Let the *granularity* δ of a lattice be the minimum distance between two nodes. Let A indicate the size of the area to cover.

When the hexagon lattice is used to provide 1-coverage, we have $\delta = r\sqrt{3}$.

Definition 2. Let n be the number of sensors used in a given deployment. Let the *coverage density* ρ of the deployment be the ratio of the *sum* of all sensors' coverage region sizes to the area size A. Let $n_{opt} = n_{hex}$ and $\rho_{opt} = \rho_{hex}$ indicate these values in the case of the hexagon lattice, and let n_{sq} and ρ_{sq} indicate them in that of the square lattice. Let $g = n/n_{opt}$ be the number of *hexagon lattices' worth of sensors* used in the deployment.

The density ρ_{opt} is simply the ratio of the areas of circle and hexagon.

$$\rho_{opt} = \frac{2\pi}{3\sqrt{3}} \approx 1.209. \tag{1}$$

Hence, the number of sensors n_{opt} needed to cover a region of area A is given by

$$n_{opt} = \frac{A \cdot \rho_{opt}}{\pi r^2} = \frac{2A}{r^2 3\sqrt{3}} \approx \frac{0.385A}{r^2}. \tag{2}$$

Although the hexagon lattice uses fewer sensors than any other, many other grid formations could be considered, some of which have practical advantages. One example is the square lattice which has the following coverage density and sensor count, respectively.

$$\rho_{sq} = \frac{\pi}{2} \approx 1.57; \tag{3}$$

$$n_{sq} = \frac{A}{2r^2}. \tag{4}$$

Definition 3. Let (a, b) and (c, d) indicate the two basis vectors generating the lattice.

We have $(a, b) = (\frac{r\sqrt{3}}{2}, \frac{r3}{2})$, $(c, d) = (r\sqrt{3}, 0)$ for the hexagonal lattice and $(a, b) = (r, 0)$, $(c, d) = (0, r)$ for the square lattice. Then the coverage density is equal to $\frac{\pi r^2}{|ad-bc|}$, which generalizes the two densities just stated.

3.5. Analysis

The observation that in some k-coverings generated by the Riesz energy method all the sensors themselves form a denser hexagon lattice suggests another scheme for obtaining k coverage, that is, take the standard hexagon lattice initially drawn on a scale where $r = 1$ and $\delta = \sqrt{3}$, and rather than minimizing Riesz energy, simply increase its density (or equivalently increase the value r) until k-coverage has been achieved. Notice that there is an equivalence between increasing the density and increasing r: We always have $g = r^2$. For example, in the unscaled hexagon lattice, we have $g = 1 = r^2$.

Unlike in single coverage, a hexagon lattice is not the most efficient configuration for every value k, for example, consider the 2-coverage configuration shown in Figure 1(b). By what factor must the hexagon lattice be shrunk in order to provide, for example, 2-coverage? Consider a neighborhood $N_\epsilon(x)$ centered on point x which is a sensor location in the hexagon lattice. The first cover of this area comes from the sensor located at x. The second cover must be provided by (perhaps some of) the sensors surrounding it, that is, x's six neighbors in the hexagon lattice. Unfortunately, any radius assigned to those points that yields an additional cover for $N_\epsilon(x)$ will actually provide two more covers for it. That is, the required shrinking factor corresponds to three times as many sensors as a single lattice (i.e., the hexagon lattice with $g = 3$ or $r = \sqrt{3}$). The corresponding arrangements of disks is known as the Flower of Life. Naively shrinking until we double the number of sensors (i.e., until $g = 2$) would fail to produce 2-coverage. Riesz energy minimization provides a 2-cover using two copies of the hexagon lattice with $\delta = r = 1$ (see Figure 1(b)). Nonetheless, for many values (including but not limited to $k \in Q$), shrinking performs well.

The proposed k-covering satisfies a stronger property of providing k-coverage, using one sensor from each of the underlying k lattices, which may be useful in multimodal data fusion applications, where there are k types of sensors, and each point in the region of interest needs to be covered by at least one sensor of each type. On the other hand, there is overlap in the coverage area, that is, there are areas in a 1-covering that are covered by more than one sensor (e.g., Figure 1(a)). We observed previously that the overlap density ρ is equal to $\frac{\pi r_c^2}{|ad-bc|}$, which is $\frac{2\pi}{3\sqrt{3}} \approx 1.209$ and $\frac{\pi}{2} \approx 1.57$ for the hexagonal and square lattices, respectively.

In the absence of such data fusion applications, we can take advantage of this overlap, that is, for sufficiently large k, the k-covering constructed is actually a k'-covering with $k'/k \approx \rho_{opt}$.

Proposition 4. For all $\epsilon > 0$, there exists k such that the k-covering constructed by concatenating k 1-coverings is also a k'-covering, and $k'/k > \rho_{opt} - \epsilon$.

PROOF. Let k be chosen such that the lattice can be decomposed as k identical sublattices, for example, we pick $k \in Q$ for the hexagonal lattice. The distance between the sensors is $1/\sqrt{k}$ times the corresponding distance in the sublattice. For each point p in the region, consider the neighborhood $N_r(p)$. The number of lattice points in $N_r(p)$—all of which cover point p—is approximately $k \cdot \frac{\pi r^2}{|ad-bc|}$ as $k \to \infty$. Therefore, this is a k'-covering where $k' \approx k \cdot \frac{\pi r_c^2}{|ad-bc|} = k\rho_{opt}$. □

We now discuss the speed of convergence of k'/k to ρ_{opt} in the preceding proposition. It follows from Huxley [1995] that for k'/k approaching ρ_{opt} as $k \to \infty$, we have $k'/k > \rho_{opt} - \epsilon$, where $\epsilon = O(k^{-\frac{2}{3}+\nu})$ for any $\nu > 0$. First consider the easier case of a unit square grid in which case $|ad - bc| = 1$. The number of lattice points in the circular disk is k', and the area of the disk is $k\pi$. A corollary in Huxley [1995] (substituting $k^{1/2}$ in for M in the corollary of Huxley [1995]) shows that the discrepancy in the number of lattice points in the disk, that is, $k' - k\pi$, is $O(k^{\frac{1}{3}}(\log k)^{\frac{1}{3}}) = O(k^{\frac{1}{3}+\nu})$. Therefore, $k'/k = \pi + O(k^{-\frac{2}{3}+\nu})$. For a general lattice, an affine transformation transforms the lattice into a unit square lattice, and the area of the unit disk is transformed into an ellipse of area $\frac{\pi}{|ad-bc|}$, but the conclusion still holds.

In fact, ρ_{opt} upper-bounds the maximum achievable gain.

Proposition 5. k'/k cannot exceed ρ_{opt}.

Table I. $k \to k'$ Values for $k \in [1, 27]$ such that $\frac{k'}{k} > 1$

Sensing Range	k	k'	Gain Ratio
3.31662	11	12	1.09091
3.60555	13	14	1.07692
3.87298	15	16	1.06667
4.00000	16	18	1.12500
4.12311	17	19	1.11765
4.24264	18	19	1.05556
4.58258	21	24	1.14286
4.69042	22	25	1.13636
4.79583	23	26	1.13043
4.89898	24	27	1.12500
5.00000	25	28	1.12000
5.09902	26	28	1.07692
5.19615	27	30	1.11111

PROOF. We consider the hexagon lattice case here, but the same argument is valid for other lattices. Let n_{opt} be the number of sensors appearing in the hexagon lattice covering a certain area of size A (see Definition 2). Suppose kn_{opt} sensors (i.e., k hexagon lattices' worth) provide a k'-coverage. A hexagon lattice covers the entire area at least once and approximately 20.9% of it twice, that is, $n_{opt}S \approx 1.209A$, where S is the size of a single sensor's sensing range (with radius r). Then, the total amount of area covered by k lattice's worth of sensors is $kn_{opt}S \approx k1.209A$. Therefore, with kn_{opt} sensors, the best we could hope for is that the full area A is covered $\approx 1.209k$ times, that is, $k'/k \leq \rho_{opt}$. □

We show the following by machine computation. We verified the results by hand for several cases.

Proposition 6. The coverage multiplicity k' provided by the shrinking factor corresponding to a k-fold duplication of the hexagon lattice can be computed in time $O(k^2 \log k)$.

PROOF. (sketch) The multiplicity of coverage provided by a given arrangement can be found efficiently by applying the algorithm of Huang and Tseng [2005]. Their algorithm computes the minimum value k' for which all sensor disks are *k'-perimeter covered*, which means that all points on the perimeter are covered by k' sensors other than itself, on the assumption that no two sensors lie at the same location. Let r be the sensing range, then computing the perimeter coverage of one sensor s can be done in time $O(c \log c)$, where c is the number of sensors within distance $2r$ of s. Due to symmetry, we only need to compute the perimeter coverage of a single sensor in the hexagon lattice (in fact, for only 1/6 of its perimeter). Since there are $O(k^2)$ sensors in the hexagon lattice within distance $2k$ of any point, the total running time will be $O(k^2 \log k)$. □

We used the method just described to compute the coverage multiplicity k' provided by a number of small duplication values $k = 1, \ldots, 27$. The values k in this range yielding an improved multiplicity (i.e., those with $k' > k$) are shown in Table I). In particular, 11 copies of the hexagon grid suffice to provide 12-coverage.

In Figure 7(a) we show the obtained efficiency gain k'/k for various values of k. We see that it approaches $\rho_{opt} = 1.209$ for large k, as promised by Proposition 4. Analogously, Figure 7(b) shows the obtained efficiency gain for a shrunken square lattice, which correspondingly approaches a recovery of the $\approx 57\%$ inefficiency of the square grid. We emphasize here that although the square lattice is typically considered inferior to the hexagon lattice, these results argue that the hexagon's advantage disappears

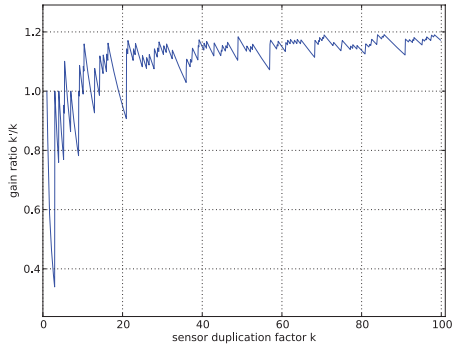

(a) Shrunken hexagon lattice efficiency gain $\frac{k'}{k}$ v. k.

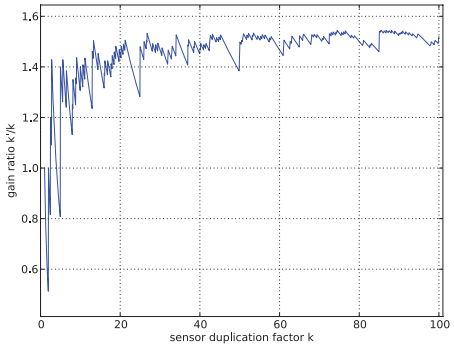

(b) Shrunken square lattice efficiency gain $\frac{k'}{k}$ v. k.

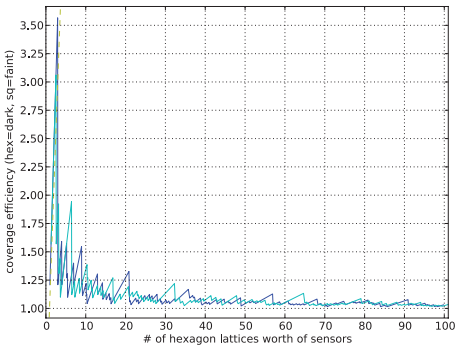

(c) Coverage efficiency (coverage size / sum of disk sizes)

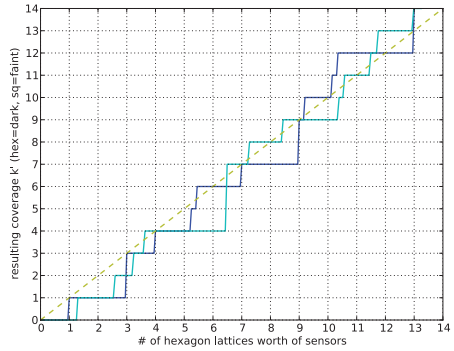

(d) Coverage provided vs. cost

Fig. 7. Computational results.

when comparing shrunken lattices. In the limit, all the single lattice's redundancy is recovered.

Figure 7(c) plots the coverage efficiency of the two lattices as we vary the number of sensors used. For both curves in this figure, the X coordinate indicates the number of hexagon lattice's worth used, which is a $1/n_{opt}$ fraction (ignoring remainders) of the total number of sensors used (see Definition 2). The Y coordinate indicates the efficiency, that is, $\rho_{opt}/(k'/k)$. This efficiency measure is equivalent to the coverage area size divided by the total coverage paid for, that is, the number of sensors times the coverage area size of one sensor. Following the known discrepancies in this efficiency for values $k = 1$ and $k = 2$, we observe as expected that both efficiency ratios converge to unity in the limit. Finally, Figure 7(d) illustrates a portion of this data differently, plotting the coverage provided by hexagon and square lattices as a function of the number of sensors used (varying from 1 to 14 hexagon lattices' worth). Note that each time one of these curves supersedes the dashed $y = x$ line, it corresponds to a particular shrunken lattice arrangement that is more efficient than the base hexagon lattice.

Again, we find potential grounds for using the square lattice. Suppose, for example, that we wish to cover a given area with 4-coverage. Shrinking the hexagon lattice by a factor of four will provide 4-coverage, yielding the same efficiency ratio as simply using four copies of the hexagon lattice. Shrinking the square lattice by a factor of

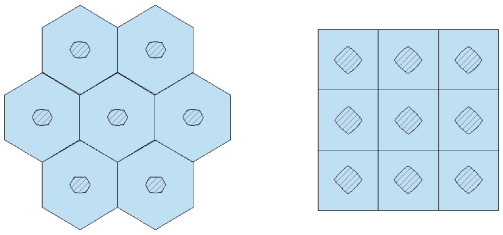

Fig. 8. In hexagon and square lattices, wiggle regions such that sensors cover their responsibility regions.

three, however, will provide 4-coverage while using $3 \cdot \rho_{sq}/\rho_{opt} \approx 3 \cdot 1.57/1.209 \approx 3.896$ times as many sensors as one hexagon lattice. Or, if our goal is, for example, 7-coverage (see Figure 6(b)), we find that shrinking the square lattice by a factor of five will provide 7-coverage for the cost of $\approx 5 \cdot 1.57/1.209 \approx 6.493$ hexagon lattices. That is, the hexagon lattice shrunken by a factor of seven provides 7-coverage using approximately $7 \cdot 1.209 \approx 8.463$ sensors per unit area, whereas the square lattice shrunken by a factor of five provides the same coverage with a per-unit area sensor count of only about $5 \cdot 1.57 \approx 7.85$.

4. WIGGLE ROOM

As we noted previously, the optimal hexagon lattice is very rigid in the sense that the movement of a single sensor will result in uncovered space. In this section, we will explore a trade-off between density and position flexibility.

Although our main focus is on the hexagon lattice, the square lattice is popular because of practical convenience, and so we sometimes include calculations for the square lattice for comparison. By a *rigid* lattice, we mean a lattice (hexagon or square) with the smallest number of sensors among lattice arrangements (of the same kind, hexagon or square) providing full coverage. An operational difficulty of these rigid solutions is that sensors must be placed exactly, with no room for placement error. After all, any room for error would imply that the lattice was not tight and so used more sensors, or sensors with larger sensing radii, than necessary. Conversely, if we do use more or larger-radius sensors than the absolute minimum needed, there may be flexibility in allowable sensor placement while maintaining full sensing coverage over the region. We refer to the region within which a sensor can be placed while full coverage is maintained as the sensor's *wiggle region*. Given a particular feasible configuration (meaning one which covers the entire region of interest), we may be interested in the wiggle regions of all the individual sensors. Wiggle regions are such that full coverage is maintained as sensors move about, perhaps simultaneously, as long as each sensor remains in its wiggle region. For sensors in any lattice configuration, the wiggle regions would be congruent but translated; for an arbitrary arrangement, different sensors may have different sorts of wiggle regions. Given a placement of sensors, we may compute the shape and size of the wiggle region for each.

A sensor's *intended location* is the location where it lies if there is zero placement error. A simple and practical way of characterizing the size of a wiggle region is by the *wiggle radius* w, that is, the radius of the largest circle centered on the intended sensor location inscribable in it. We call this circle the *wiggle disk*. The wiggle disk will be of practical importance when placement error is isotropic, that is, the same for all directions. For some configurations, such as the hexagon lattice, the difference between the wiggle disk and the full wiggle area will be quite small (see Figure 8 (left), where the shaded wiggle area is itself close to disk-shaped), though this need not be the case

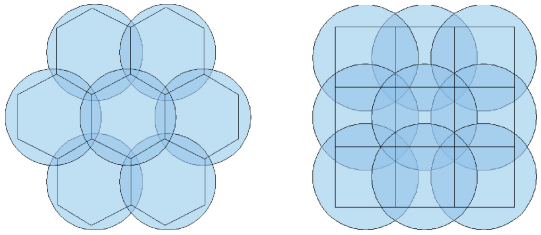

Fig. 9. In hexagon and square lattices, congruent responsibility regions that tile the region.

in general (see Figure 8 (right), where the four-pointed shaded wiggle area is far from disk-shaped).

If (conversely) we are given desired amounts of allowable placement error, we could compute a configuration that is robust enough to continue to guarantee full coverage in spite of this error. While lattice-based solutions have desirable properties and are straightforward to apply, it may in reality be impractical to place sensors exactly on grid points. A dual problem is to find an efficient configuration that compensates for the placement error that we wish to allow and, therefore, assume to occur. It is for this second problem that we provide the denser hexagon lattice as a solution.

4.1. From Arrangements to Wiggle Regions

In this section, we examine the scenario in which all the sensors can be repositioned from their intended position. Locally, this means that a sensor and all its neighbors may be inexactly positioned. There must be some overlap among sensors' coverage disks in any arrangement fully covering a region. We now characterize the wiggle region, comprising all those points on which this sensor could be repositioned such that full coverage is maintained, for these two settings. We focus on lattice-based solutions with congruent wiggle regions.

We introduce the notion of the *responsibility region* (RR) for a sensor s, which is the region that s must fully cover no matter where the other sensors lie. By achieving coverage for each RR, we therefore achieve full coverage. We say that s is responsible for any point p whose distance to the initially designated position of sensor s is no greater than its distance to the initially designated position of any other sensor, in which case s's responsibility region is simply its Voronoi neighborhood [Pach and Agarwal 1995]. Figure 9 shows the boundaries of the Voronoi neighborhoods of sensors in lattice arrangements. Due to symmetry in lattice arrangements (and due to the nature of Voronoi diagrams), these RRs will tile the plane, as shown in Figure 9.

The wiggle regions based on these congruent responsibility regions can once again be constructively obtained. Let R be the bounded responsibility region of a particular sensor s with sensing radius r, and let P be the set of extreme points of R. Then full coverage is maintained if and only if s lies in the intersection of closed disks with radius r centered at the points in P. We take the extreme points to be the Voronoi points in the resulting Voronoi diagram.

Each responsibility region R in a lattice-based arrangement is its own convex hull and will be covered by placing a sensor anywhere within the intersection of radius-r disks centered at R's extreme points. We show the wiggle regions for a hexagon lattice in Figure 8. Since each sensor can be placed within its wiggle region independently of the choice of where others are placed within their wiggle regions, the placement policies must naturally be more conservative. Treating Voronoi neighborhoods in arbitrary arrangements as responsibility regions implies corresponding wiggle regions as well.

Since Voronoi neighborhoods are always convex, the construction in Proposition 1 can be applied in every case in order to obtain the wiggle region.

In a general arrangement of sensors, given a designated sensor with position p, let X be the set of extreme points of its responsibility region, and for each $q \in X$, let $s(q)$ denote the distance between p and q. Let $d(q)$ denote the distance between q and the boundary of the same sensor's coverage disk. That is, $d(q) = r - s(q)$ for $0 \leq s(q) \leq r$. Then the disk centered at p with radius $\min_{q \in X} d(q)$ necessarily fits within the wiggle region. In the field, it may be easier to work with this smaller disk instead of the whole wiggle region, since it is easier to calculate and specify. Note that for non-lattice arrangements, this need not be the largest disk that can be inscribed in the wiggle region.

4.2. From Desired Flexibility to Required Density

We use the preceding characterization in calculating the lattice granularity (and thus the number of sensors) required for covering a region with a lattice, in accordance with a minimum required wiggle radius.

Suppose we are given an upper bound on placement error w, and we would like to find a lattice arrangement that attains full coverage. Observe that it is necessary and sufficient to position sensors in a lattice such that a circle of radius w could be inscribed in the wiggle region of every hypothetical sensor placed exactly at grid points. In other words, no matter where the sensor is placed inside this disk, its distance to every extreme point of the (Voronoi) responsibility region must be at most r. The best arrangement, that is, one with the largest granularity, under these additional constraints is one in which the distance from every critical point to the center of a Voronoi region is exactly $r - w$.

By definition, in a rigid lattice arrangement, the distance from each sensor to the extreme points of its Voronoi region is r; we can allow for wiggle room w by scaling the rigid lattice by a factor of $(r - w)/r$. Because a disk with radius r and placement error w can be treated as an exactly placed disk with radius $r - w$ [Bar-Noy et al. 2009], we can compensate for the inexact placement by scaling the lattice.

Since in the rigid hexagon lattice, the granularity is $r\sqrt{3}$, the granularity δ_{hex} of the hexagon lattice accommodating w-bounded placement error is given by $\delta_{hex} = (r-w)\sqrt{3}$. Conversely, if all sensors are free to move in a hexagon lattice with granularity δ_{hex} and coverage radius r, solving for w indicates how far the sensors may be allowed to stray from their original positions without violating the coverage guarantee.

For a hexagon lattice, or indeed for any sensor configuration in general, the number of sensors used when we allow wiggle room w is in general scaled by a factor $r^2/(r-w)^2$.

Since given r and w the number of sensors required is affected by the same factor in all cases, and since the rigid hexagon lattice is optimal, the best solution in the case of inexact placement with bounded placement error is the 6-grid previously discussed. From Eq. (1), the density of sensors $\rho_{hex}(w)$ in a hexagon lattice which permits a placement error of w is given by

$$\rho_{hex}(w) = \frac{2\pi r^2}{3(r-w)^2\sqrt{3}} \approx \frac{1.209r^2}{(r-w)^2}. \tag{5}$$

Recalling Eq. (2), where A is the area of the entire region of interest, the number of sensors $n_{hex}(w)$ needed in this hexagon lattice is found to be

$$n_{hex}(w) = \frac{2A}{3(r-w)^2\sqrt{3}} \approx \frac{0.385A}{(r-w)^2}. \tag{6}$$

Similarly, if it were an additional requirement to use a 4-grid, recalling Eq. (3), the density of sensors $\rho_{sq}(w)$ in such an arrangement allowing a placement error of w is given by

$$\rho_{sq}(w) = \frac{\pi r^2}{2(r-w)^2} \approx \frac{1.57r^2}{(r-w)^2}. \tag{7}$$

Recalling Eq. (4), the number of sensors $n_{sq}(w)$ required in this square lattice is found to be

$$n_{sq}(w) = \frac{A}{2(r-w)^2} = \frac{0.5A}{(r-w)^2}. \tag{8}$$

We note that the ratio of required densities for coverage by square or hexagon lattices is invariant with w, that is, is the same as the efficiency ratio for the two underlying (exactly placed) lattices.

5. ROBUSTNESS TO SENSOR FAILURE

A third benefit of shrinking is robustness to sensor failure. Suppose each deployed sensor fails with probability p. Clearly this will tend to break the full coverage of the hexagon lattice, but shrinking will decrease this tendency. Computing the probability of full coverage in this case is problematic, since the probability of full coverage will depend heavily on the size of the relative coverage area. One alternative measure is the probability that a single *local triangle* (i.e., a triangle formed by three mutually adjacent sensors) is covered (or k''-covered for each $k'' \in [1, k']$). Restricting our attention to the disks that fully cover the triangle (in the case of $k = 49$, there are 48 such disks), we obtain the following lower bound on the expected coverage multiplicity (shown with a faint dashed line) through linearity of expectation.

Proposition 7. The expected coverage multiplicity k'' for a local triangle T is at least $(1-p) \cdot n_f$, where n_f is the number of sensors fully covering T.

We performed simulations to estimate the coverage obtained when sensors fail (independently) with a fixed probability, as well as the most generous failure probabilities that will yield coverage. Figure 10(a) illustrates the expected multiplicity k'' of coverage for a deployment with $k = 49$ for varying failure probabilities p. Two faint bounding lines are shown, indicating that k'' is bounded below by the stated Bernoulli distribution mean but is strictly convex. Conversely, Figure 10(b) shows the minimum success probabilities required, when the lattice has been shrunk by a factor k, in order to maintain 1-coverage or k-coverage (in expectation).

6. DISCUSSION

We have examined three separate motivations for shrinking the hexagon lattice (or increasing the sensor radii): obtaining flexibility in sensor placement, increasing the sensor efficiency in obtaining k-coverage, and providing fault tolerance. Increasing the radius by value w allows for amount w of placement error, as discussed in Section 4. For values $k \in Q'$, this is equivalent to k-covering where the distance between sensors is shrunk by a factor of $1/\sqrt{k}$. The third benefit to shrinking is to maintain coverage when sensors fail.

Thus, increasing the radius r from 1 to a larger value can be construed as providing a redundancy budget which may be spent in multiple ways. If the enlarged radius r is spent entirely on wiggle room $w = r - 1$, then we may only rely upon 1-coverage. Conversely, if r is spent entirely on getting $k' \approx 1.209k$ coverage, then no flexibility in sensor placement is provided. Finally, it can be spent entirely on failure probability (as in the lower curve in Figure 10(a)). First of all, we can partition the benefit of r as

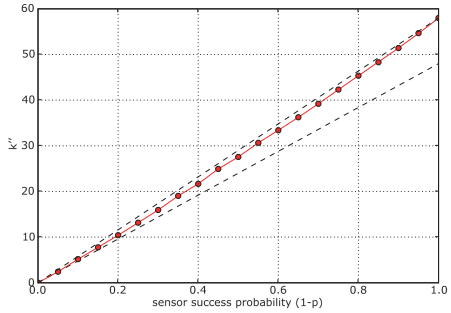

(a) Expected coverage multiplicity (k'') with $k = 49$).

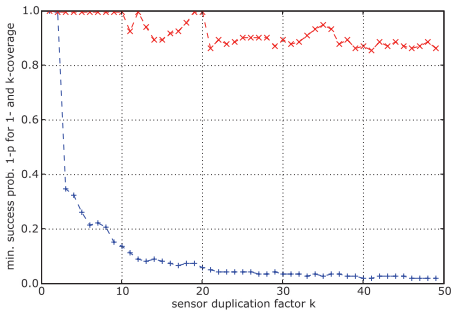

(b) Success probabilities for 1- and k-coverage.

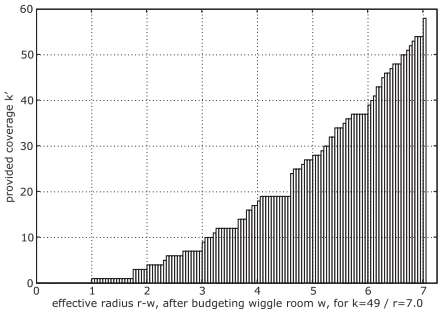

(c) Efficiency gain after allocating wiggle room budget.

Fig. 10. Computational results.

$r = r_w + r_c$, yielding wiggle room $w = r_w$ and at least k-coverage for $k \geq r_c^2$ (which may, as previously indicated, be made probabilistic). A single shrinking factor could provide a range of possible trade-offs between benefits, with the degree of coverage and the amount of allowable placement error varying depending on how the shrunken lattice is interpreted. Figure 10(c) illustrates the amount of efficiency gain received after allocating some amount $w < r$ to wiggle room, starting with the radius r corresponding to $k = 49$. Notice that in varying the effective radius, we revisit some of the values k shown in Table I. For example, for $k = 27$ and $w = 0.613$, we have $r = \sqrt{27} - w \geq \sqrt{21}$, which implies that 24-coverage is obtained.

Second, after allocating wiggle room, there are various potential ways to further divide the shrinking benefit between multi-coverage and failure robustness. Figure 10(a), for example, can be interpreted as characterizing the possible ways of allocating the shrinking factor 7 (corresponding to $k = 49$) between allowable sensor failure probability and expected multiplicity of coverage. For a particular value k, two natural trade-off choices are those in which the maximum allowable failure probabilities are chosen that yield either 1-coverage or k-coverage in expectation (see Figure 10(b) for an estimation of these curves based on random trials). In the case of $k = 49$, these two choices can be extracted from the points on the curve in Figure 10(a) with $x = 1$ or $x = 49$.

In conclusion, we have examined three kinds of benefits between which the shrinking budget can be divided. We emphasize that this budget need not be merely metaphorical. As we saw in Section 4, the size of the coverage region and the lattice density together determine the number of sensors used; their cost can be significant in terms

of equipment costs, placement, and energy. In this article, we have considered three reasons to pay this price.

REFERENCES

ABRAMS, Z., GOEL, A., AND PLOTKIN, S. A. 2004. Set k-cover algorithms for energy efficient monitoring in wireless sensor networks. In *Proceedings of the International Conference on Information Processing in Sensor Networks*.

ALT, H., ARKIN, E. M., BRÖNNIMANN, H., ERICKSON, J., FEKETE, S. P., KNAUER, C., LENCHNER, J., MITCHELL, J. S. B., AND WHITTLESEY, K. 2006. Minimum-cost coverage of point sets by disks. In *Proceedings of the Symposium on Computational Geometry*. 449–458.

BAI, X., KUMAR, S., XUAN, D., YUN, Z., AND LAI, T.-H. 2006. Deploying wireless sensors to achieve both coverage and connectivity. In *Proceedings of the International Symposium on Mobile Ad Hoc Networking and Computing.*

BAR-NOY, A., BROWN, T., JOHNSON, M. P., AND LIU, O. 2009. Cheap or flexible sensor coverage. In *Proceedings of the International Conference on Distributed Computing in Sensor Systems*.

BARTOLINI, N., CALAMONERI, T., FUSCO, E. G., MASSINI, A., AND SILVESTRI, S. 2008. Snap and spread: A self-deployment algorithm for mobile sensor networks. In *Proceedings of the International Conference on Distributed Computing in Sensor Systems*.

BERNSTEIN, M., SLOANE, N. J. A., AND WRIGHT, P. E. 1997. On sublattices of the hexagonal lattice. *Discrete Math. 170*, 29–39.

BRASS, P. 2007. Bounds on coverage and target detection capabilities for models of networks of mobile sensors. *ACM Trans. Sen. Netw. 3,* 2.

BRASS, P., MOSER, W. O. J., AND PACH, J. 2005. *Research Problems in Discrete Geometry*. Springer, Berlin-Heidelberg.

BROWN, T., SARIÖZ, D., BAR-NOY, A., T. LA PORTA, VERMA, D., JOHNSON, M. P., AND ROWAIHY, H. 2007. Geometric considerations for distribution of sensors in ad-hoc sensor networks. In *Proceedings of the SPIE Defense, Security, & Sensing Conference*.

FUNKE, S., KESSELMAN, A., KUHN, F., LOTKER, Z., AND SEGAL, M. 2007. Improved approximation algorithms for connected sensor cover. *Wirel. Netw. 13,* 2, 153–164.

GUPTA, H., ZHOU, Z., DAS, S. R., AND GU, Q. 2006. Connected sensor cover: Self-organization of sensor networks for efficient query execution. *IEEE/ACM Trans. Netw. 14*, 55–67.

HARDIN, D. P. AND SAFF, E. B. 2004. Discretizing manifolds via minimum energy points. *Not. AMS 51,* 10, 1186–1194.

HUANG, C.-F. AND TSENG, Y.-C. 2005. The coverage problem in a wireless sensor network. *Mob. Netw. Appl. 10,* 4, 519–528.

HUXLEY, M. N. 1995. The mean lattice point discrepancy. In *Proceedings of the Edinburgh Mathematical Society*, vol. 38. 523–531.

JOHNSON, M. P., SARIOZ, D., BAR-NOY, A., BROWN, T., WU, C. W., AND VERMA, D. 2009. More is more: The benefits of dense sensor assignment. In *Proceedings of the Annual Joint Conference of the IEEE Computer and Communications Societies*.

KANSAL, A., HSU, J., ZAHEDI, S., AND SRIVASTAVA, M. B. 2007. Power management in energy harvesting sensor networks. *ACM Trans. Embed. Comput. Syst.* 6, 2.

KERSHNER, R. 1939. The number of circles covering a set. *Am. J. Math. 61*, 665–671.

PACH, J. AND AGARWAL, P. 1995. *Combinatorial Geometry*, 3rd Ed. Wiley-Interscience, New York, NY.

PACH, J. AND TÓTH, G. 2007. Decomposition of multiple coverings into many parts. In *Proceedings of the Annual ACM Symposium on Computational Germetry*. 133–137.

PODURI, S. AND SUKHATME, G. S. 2004. Constrained coverage for mobile sensor networks. In *Proceedings of the IEEE International Conference on Robotics and Automation*. 165–171.

ROCKAFELLAR, R. T. AND WETS, R. J.-B. 2009. *Variational Analysis*. Springer, Berlin-Heidelburg.

SPAULDING, K. E., MILLER, R. L., AND SCHILDKRAUT, J. 1997. Methods for generating blue-noise dither matrices for digital halftoning. *J. Electron. Imag. 6,* 2, 208–230.

SUN, T., CHEN, L.-J., HAN, C.-C., , AND GERLA, M. 2005. Reliable sensor networks for planet exploration. In *Proceedings of the IEEE International Conference on Networking, Sensing, and Control*. 373–378.

ULICHNEY, R. 1997. *Digital Halftoning*, 3rd Ed. the MIT Press, Cambridge, MA.

WANG, Y.-C. AND TSENG, Y.-C. 2008. Distributed deployment schemes for mobile wireless sensor networks to ensure multilevel coverage. *IEEE Trans. Parallel Distrib. Syst. 19*.

WOMERSLEY, R. Visualization of minimum energy points on the torus. http://web.maths.unsw.edu.au/~rsw/Torus/.

WU, C. W., THOMPSON, G., AND STANICH, M. 2003. A unified framework for digital halftoning and dither mask construction: Variations on a theme and implementation issues. In *Proceedings of the International Conference on Digital Printing Technologies*. 793–796.

WU, C. W. AND VERMA, D. 2008. A sensor placement algorithm for redundant covering based on riesz energy minimization. In *Proceedings of the IEEE. Circuits and Systems*. 2074–2077.

ZHANG, H. AND HOU, J. C. 2005. Maintaining sensing coverage and connectivity in large sensor networks. *Ad Hoc Sen. Wirel. Netw. 1,* 1-2, 89–124.

ZHANG, H. AND HOU, J. C. 2006. Is deterministic deployment worse than random deployment for wireless sensor networks? In *Proceedings of the Annual Joint Conference of the IEEE Computer and Communications Societies*.

ZHOU, Z., DAS, S., AND GUPTA, H. 2004. Connected K-coverage problem in sensor networks. In *Proceedings of the International Conference on Computer Communications Networks*. 373–378.

ZOU, Y. AND CHAKRABARTY, K. 2004. Uncertainty-aware and coverage-oriented deployment for sensor networks. *J. Parallel Distrib. Comput. 64,* 7, 788–798.

Received August 2010; revised December 2010, May 2011; accepted May 2011