Approximating barrier resilience in wireless sensor networks

Sergey Bereg[1] and David Kirkpatrick[2]

[1] Department of Computer Science, University of Texas at Dallas, USA.
[2] Department of Computer Science, University of British Columbia, Canada.

Abstract. Barrier coverage in a sensor network has the goal of ensuring that all paths through the surveillance domain joining points in some start region S to some target region T will intersect the coverage region associated with at least one sensor. In this paper, we revisit a notion of redundant barrier coverage known as k-barrier coverage.
We describe two different notions of width, or impermeability, of the barrier provided by the sensors in $\mathcal{A}$ to paths joining two arbitrary regions S to T. The first, what we refer to as the *thickness* of the barrier, counts the minimum number of sensor region intersections, over all paths from S to T. The second, what we refer to as the *resilience* of the barrier, counts the minimum number of sensors whose removal permits a path from S to T with no sensor region intersections. Of course, a configuration of sensors with resilience k has thickness at least k and constitutes a k-barrier for S and T.
One of our two main results demonstrates that any (Euclidean) shortest path from S to T that intersects a fixed number of distinct sensors, never intersects any one sensor more than three times. It follows that the resilience of $\mathcal{A}$ (with respect to S and T) is at most three times the thickness of $\mathcal{A}$ (with respect to S and T). (Furthermore, if points in S and T are moderately separated (relative to the radius of individual sensor coverage) then no shortest path intersects any one sensor more than two times, and hence the resilience of $\mathcal{A}$ is at most two times the thickness of $\mathcal{A}$.)
Our second main result, which we are only able to sketch here, shows that the approximation bounds can be tightened (to 1.666 in the case of moderately separated S and T) by exploiting topological properties of simple paths that make double visits to a collection of disks.

1 Introduction

Various notions of *coverage* provided by wireless sensor networks have attracted considerable attention over the past few years. (D.W. Gage [1] initiated the formal study of sensor coverage, and the recent survey papers of Meguerdichian *et al.* [2] and Cardei and Wu [3], as well as the Ph.D. thesis of S. Kumar [4] provide comprehensive overviews of work on the topic). A fundamental concern is the design of networks that achieve high quality of coverage. Central to this endeavor is the evaluation of the quality of coverage of a given sensor network.

In general, coverage can be expressed geometrically, by relating the positions, and associated *coverage regions*, of individual sensors to some underlying surveillance domain. However, different applications motivate different notions of coverage. Three concepts that have received a significant amount of study are *area coverage*, where the goal is to achieve coverage for *all points* in the surveillance domain by a static collection of sensors, *sweep coverage*, where the goal is to ensure that any point moving continuously within the surveillance domain will be detected at some point in time by a collection of moving sensors, and *barrier coverage*, where the goal is to ensure that all paths through the surveillance domain joining points in some start region S to some target region T will intersect the coverage region associated with at least one member of some static collection of sensors.

Barrier coverage has the attractive feature of guaranteeing the absence of undetected transitions between critical subsets of the surveillance domain (for example, between unsecured entry and exit points) without the high (and in many cases, unwarranted) cost of full area coverage. However, as has been observed in several papers, barrier coverage, in its simplest formulation, does not adequately capture the robustness requirements of typical applications; for example, a configuration of sensors could provide a barrier cover between S and T which, on the failure of even a single sensor would disintegrate into something that does not even provide a reasonable approximation to barrier coverage.

Several proposals have been made to increase the robustness of the barrier coverage concept. Some [5, 6] are based on probabilistic assumptions about the distribution of sensors or paths. Other proposals retain the deterministic/worst-case nature of basic barrier coverage. Meguerdichian *et al.* [2, 7] suggest measuring the quality of barrier coverage in terms of what they call *maximal breach paths* (paths that maximize the distance to their closest sensor) and *minimum exposure paths* (paths that minimize the total degree of exposure to sensors, measured in terms of both proximity and duration). Jiang and Chen [8] consider double barrier coverage, which holds when every path from S to T must, at some point, be simultaneously covered by at least two distinct sensors. Kumar *et al.* [9, 10] introduce a different notion of multiple coverage for paths connecting certain highly constrained regions S and T. Specifically, they define a configuration of sensors to provide k-barrier coverage if every path joining a point in S to a point in T must intersect at least k distinct sensor regions. Most recently, Chen *et al.* [11, 12] have studied *localized* notions of k-barrier coverage that constrain the space of feasible paths that need to be covered.

In this paper, we revisit the notion of k-barrier coverage in a more general context than has been previously studied. Specifically, we consider an arbitrary arrangement $\mathcal{A}$ of sensors and two arbitrary regions S and T within the surveillance domain. We describe two different notions of impermeability of the barrier provided by the sensors to paths from S to T. The first, what we refer to as the *thickness* of the barrier, counts the minimum number of sensor region intersections, over all paths from S to T. The second, what we refer to as the *resilience* of the barrier, counts the minimum number of sensors whose removal permits a path from S to T with no sensor region intersections. The critical distinction between these notions is the fact that thickness counts multiple encounters of the same sensor, while resilience counts only the first encounter. It follows that any arrangement of sensors with resilience k has thickness at least k, and constitutes a k-barrier for S and T.

Figure 1 illustrates an arrangement $\mathcal{A}$ of *disk sensors* (sensors whose coverage regions are unit disks) and a path joining a point $s \in S$ to a point $t \in T$. The dual of $\mathcal{A}$, denoted $\hat{\mathcal{A}}$ is a directed graph whose vertices are the faces of $\mathcal{A}$ and whose arcs connect vertices corresponding to adjacent faces in $\mathcal{A}$. If we assign weight 1 to arcs that correspond to a transition entering a disk region, and weight 0 to arcs that correspond to a transition exiting a disk region, then it is easy to see that geometric paths that intersect k (non-necessarily distinct) sensor regions while traversing the arrangement $\mathcal{A}$ correspond to combinatorial paths in the dual graph $\hat{\mathcal{A}}$ with path length k. Thus, the thickness of $\mathcal{A}$ with respect to the points s and t corresponds to the length of the shortest path from the face containing s to the face containing t in $\hat{\mathcal{A}}$. (Of course, such a path can be computed efficiently using standard graph algorithms). Note that the same reduction extends to arbitrary regions S and T. (It suffices to add an artificial source node s^* with an edge (s^*, f) to every face f that intersects S, with weight equal to the number of disks that cover f). The exact resilience of $\mathcal{A}$ seems much more difficult to compute in general.

Our motivation is to extend the analysis of what we call barrier resilience beyond the restricted contexts (regions separated by either open or closed belts) examined by Kumar *et al.* [10]. While there is evidence to suggest that determining the exact resilience of an arbitrary sensor configuration with arbitrary regions S and T is hard, we show that for configurations of sensors with identical disk coverage regions there is a

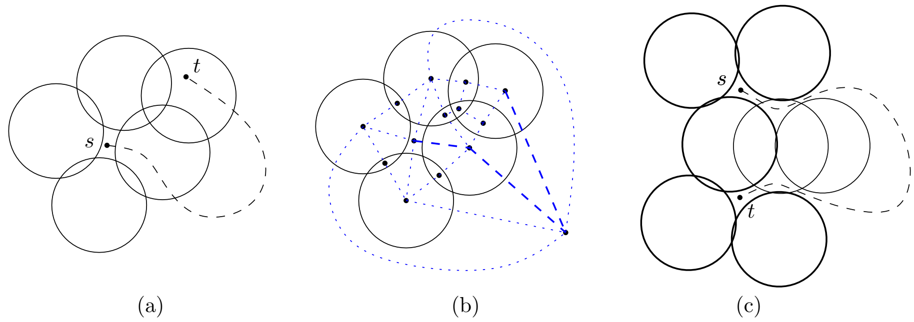

Fig. 1. (a) An st-path in a disk arrangement $\mathcal{A}$. (b) Dual of $\mathcal{A}$ with corresponding path highlighted. (c) An arrangement with thickness 2 and resilience 1 (bold disks correspond to double sensors).

close relationship between thickness and resilience. Since, as we have seen, thickness can be computed efficiently, it follows that there is an efficient algorithm to approximate thickness,

Arrangements with thickness one are exactly the same as those with resilience one (for all types of sensor coverage regions). However, for arrangements of disk sensors with thickness greater than one, the relationship between thickness and resilience is non-trivial. We note that for line-sensors (sensors whose coverage regions are unbounded lines) the notions of thickness and resilience coincide (since paths need never intersect a line-sensor more than once). On the other hand, for line-segment-sensors (sensors whose coverage regions are unit length line segments) – as studied by Kloder and Hutchinson [13], among others – there are sensor arrangements whose thickness is arbitrarily larger than their resilience.

We show, in section 3, that for an arbitrary arrangement $\mathcal{A}$ of unit disk sensors, and arbitrary points s and t in $\mathcal{A}$, any (Euclidean) shortest path from s to t that intersects a fixed number of distinct sensors, never intersects any one sensor more than three times. Furthermore, if s and t are moderately separated, more specifically if they do not co-reside on the fringe of some disk of $\mathcal{A}$, then any (Euclidean) shortest path from s to t that intersects a fixed number of distinct sensors, never intersects any one sensor more than two times. (Both of these bounds are tight, in the worst case.)

It follows immediately that the resilience of $\mathcal{A}$ (with respect to s and t) is at least one-third of the thickness of $\mathcal{A}$ (with respect to s and t). Furthermore, if s and t are moderately separated (as above) then the resilience of $\mathcal{A}$ is at least one-half of the thickness of $\mathcal{A}$. Thus any algorithm that computes the thickness of $\mathcal{A}$ provides a 3-approximation (or, under mild restrictions, a 2-approximation) of the resilience of $\mathcal{A}$.

It is natural to ask if these approximation bounds can be tightened by somehow recognizing when a path in the dual graph $\hat{A}$ re-enters a sensor disk (and can thus have its length discounted). In the next section we recall some related work that shows that such a discounting scheme, operating in polynomial time, is unlikely to exist for general graphs, even if one is only looking for a partial discount. Despite this, we are able to tighten the approximation bound (to 1.666 in the case of moderately separated s and t) by exploiting topological properties of simple paths that make double visits to a collection of disks. These results are described in section 4.

2 Background and related work

2.1 k-barrier coverage of belt regions

Kumar *et al.* [9, 10] introduced the notions of k-coverage of paths and k-barrier coverage of *belt regions*. Belt regions are defined by two uniformly separated boundaries and are either open, in which case the boundaries define the opposite sides of a strip in the plane, or closed, in which case they form the sides of a fixed width ring. In either case, the deployment of sensors within the belt is intended to cover all possible paths joining one boundary to the other.

Kumar *et al.* showed that the problem of determining if a given configuration of sensors provides a k-barrier cover for an open belt region can be reduced to the problem of determining if a given graph admits k vertex-disjoint paths between two specified vertices. Their reduction represents the sensor configuration by its intersection graph IG and exploits a well-known result in graph theory (Menger's theorem) to relate the size of minimum separating sets and maximum sets of vertex disjoint paths. Since efficient algorithms (based on maximum flows) exist for determining the existence of a maximum set of vertex disjoint paths between two specified vertices, the k-barrier question is essentially settled in this instance.

In the preliminary version [9] of their work Kumar *et al.* claimed that similar results hold for closed belts, based on the existence of non-contractable cycles in IG. This claim was subsequently retracted [10] and, consequently, the k-barrier question, even for this constrained setting, remains open.

2.2 The minimum colour single path problem

We described a reduction of the problem of determining the thickness of a sensor arrangement $\mathcal{A}$ with respect to the points s and t to that of determining the length of the shortest path between two specified faces in $\hat{\mathcal{A}}$. It is natural to ask if a similar reduction might hold for the problem of determining the resilience of $\mathcal{A}$. A promising approach in this direction is choose unique colours for each sensor and then colour the directed edges of the $\hat{\mathcal{A}}$ with colour i if that edge corresponds to a crossing into the region covered by sensor i. With this coloured-dual representation of $\mathcal{A}$ the problem of determining the resilience of $\mathcal{A}$ with respect to the points s and t reduces to that of finding a path from s to t, in this coloured dual, that uses the minimum number of distinct colours.

This *minimum colour single path problem*, for general edge-coloured graphs, was apparently first mentioned in [14]. Unfortunately, it has been shown to be NP-hard in general [15, 16]. In fact, the NP-hardness of Yuan *et al.* [16], which describes a simple reduction from the well-known set cover problem, can be easily strengthened to show that the show that the minimum colour single path problem remains NP-hard even if the underlying graph is planar and no colour appears on more than two edges.

Since the minimum set cover problem is hard to approximate to within a logarithmic factor [17], it follows from the reduction of Yuan *et al.* that the minimum colour single path problem is also hard to approximate to within a logarithmic factor. (Of course, if no colour is used more than d times then there is an obvious d-approximation algorithm.)

3 Relating resilience and thickness for unit disk sensors

In this section we establish a close connection between resilience and thickness for arrangements of disk sensors by proving that minimum (Euclidean) length paths, among all s, t paths that intersect at most k sensors, have the property that they intersect any fixed sensor at most a small constant number of times. The intuition behind the result is quite straightforward: (i) if a path π visits some disk D too many times then there must exist a shortcut; (ii) the absence of a shortcut would require that D be intersected by a large number of other pairwise-disjoint disks, whose avoidance is what forces π to repeatedly intersect D; and (iii) the existence of such a collection of disks is impossible, by standard packing arguments. Unfortunately, the details are lengthy and somewhat intricate, especially if one wants to produce the tightest possible bounds.

Let π be any s, t path through $\mathcal{A}$. We find it useful to consider the entire family $\mathcal{P}_\pi$ of s, t paths that avoid all of the sensors avoided by π. We can think of such paths as geometric paths that avoid a fixed set of disk obstacles corresponding to the avoided sensors. Clearly, there is nothing lost by restricting attention to minimum (Euclidean) length paths in $\mathcal{P}_\pi$. Thus, we begin by establishing some properties of Euclidean shortest disk-obstacle-avoiding paths.

3.1 Shortest paths avoiding disk obstacles

Given a collection unit disk obstacles $D_1, \ldots, D_n$ and two points s and t in the plane, we say that a path from s to t is *legal* if it avoids the interior of every obstacle. Any shortest legal path has the property that it is a sequence $a_1, \ldots, a_k$ where (i) each element is either a straight line segment or an arc of the boundary of an obstacle, and (ii) successive elements have common tangent and direction at the common endpoint. As is common in the bounded-curvature motion planning literature, we refer to such paths as *Dubins paths*.

Dubins paths have have many interesting local properties, some of which we develop in this section. Our broader goal, however, is to establish the following global property.

Lemma 1. *Let D be an arbitrary unit disk that does not contain either s or t. Any shortest legal st-path π crosses the boundary of D at most six times. Furthermore if either s or t has distance at least $\sqrt{3} - 1$ from the boundary of D then the path π crosses the boundary of D at most four times.*

Since it is straightforward to account for coverage of the endpoints s and t, the following theorem is an immediate consequence:

Theorem 1. *Let $\mathcal{A}$ be an arrangement of disks and s and t be two points in the plane. If s and t are well-separated then the resilience of $\mathcal{A}$ is at least one-half of the thickness of $\mathcal{A}$. For general s and t the resilience of $\mathcal{A}$ is at least one-third of the thickness of $\mathcal{A}$.*

As indicated above, we begin by developing some structural properties of shortest legal paths.

Lemma 2. *Let a and b be two points on the boundary of a unit disk D. If there exists a legal path from a to b within D, then the shortest legal path from a to b within D is no longer than any path from a to b outside D.*

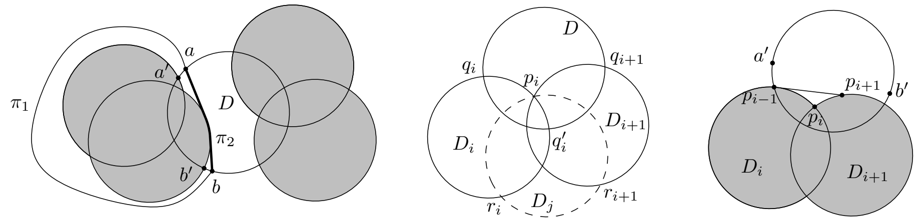

Fig. 2. Illustrations for the proof of Lemma 2.

Proof. Let π_1 be a path from a to b outside D and let π_2 be the shortest path from a to b within D. Let $\overset{\frown}{\mathrm{ab}}$ be the shortest arc of D with endpoints a and b. π_1 is not shorter than $\overset{\frown}{\mathrm{ab}}$. It suffices to prove that $\overset{\frown}{\mathrm{ab}}$ is not shorter that π_2. Let R be the region between $\overset{\frown}{\mathrm{ab}}$ and π_2. Let $\hat{D}$ be the union of disks intersecting R. Let C be any connected component of $R \cap \hat{D}$. The boundary of C contains an arc $\overset{\frown}{a'b'}$ of the boundary of D. We call $\overset{\frown}{a'b'}$ the external boundary of C and the remainder the internal boundary of C. It suffices to prove that the length of $\overset{\frown}{a'b'}$ is no less than the length of the internal boundary of C.

Fix any k and consider a set of k obstacles such that the length of the internal boundary of C is largest among all sets of k obstacles. Let $a' = p_0, p_1, p_2, \ldots, p_m = b'$ be the endpoints of arcs on the internal boundary of C. Let D_i be the obstacle with arc $p_{i-1}p_i$ on its boundary. Let $S = \{D_1, \ldots, D_{m-1}\}$.

First we show that every $p_i, i = 1, \ldots, m-1$ lies on the boundary of only two obstacles from S. Suppose to the contrary that p_i lies on the boundary of some other disk D_j from S. We specify circular arcs by their endpoints with the assumption that the arc is traced counterclockwise from its first to its second endpoint. The length of arc p_iq_i is at most the length of arc $q_i'q_i$ which is less than π (since the length of arc $a'b'$ is less than π). However the arc p_ir_i on D_i must have length greater than π (since it is outer arc of $D_i \cup D_j$). Similarly r_{i+1} is not in D. It follows that $D_j - (D_i \cup D_{i+1})$ is disjoint from D contradicting that $D_j \in S$.

If all p_i lie on the boundary of D then it is obvious that the sum of arc lengths is the same as the length of arc $\overset{\frown}{a'b'}$ on the boundary of D. Otherwise, suppose that p_i is the first point that does not lie on the boundary of D. In this case we show that the disks D_i and D_{i+1} can be perturbed to increase the length of internal boundary contradicting our maximality assumption. Rotate D_i about p_{i-1} and D_{i+1} about p_{i+1} in such a way that the intersection point p_i moves perpendicular to and away from the line $p_{i-1}p_{i+1}$. Since p_i is not on the boundary of other disks there must be a sufficiently small such motion that preserves the structure of the internal boundary while increasing the length of both $p_{i-1}p_i$ and p_ip_{i+1} (both line segments and arcs). The lemma follows. □

It follows from lemma 2 that no three successive crossing points on the boundary of D can be joined by legal paths within D. Obstacles that together block any legal path within D joining two successive boundary crossing points a to b form what we call an *obstruction* for a and b. Obstructions consisting of just two disks are called 2-*obstructions* for a and b.

Lemma 3. *Let a and b be two points on the boundary of a unit disk D. If there does not exist a legal path from a to b within D, then there must be a 2-obstruction for a and b.*

Proof. It follows from the fact that every unit disk intersecting D intersects one of the arcs of D with endpoints a and b. □

The next lemma shows that every 2-obstruction in D must cover at least half of the boundary of D. It follows that no disk D contains two disjoint 2-obstructions.

Lemma 4. *Let D be a unit disk and $\hat{D} = A \cup B$ be the union of two unit disks that form a 2-obstruction for points a and b on the boundary of D. Then at least half of the boundary of D lies inside $\hat{D}$.*

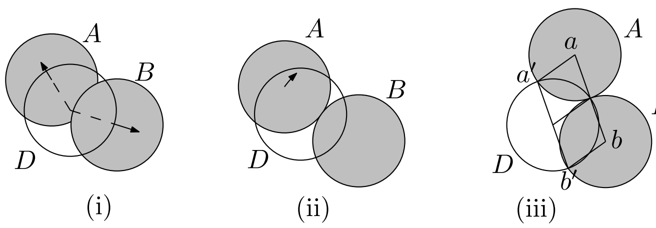

Fig. 3. Illustrations for Lemma 4.

Proof. We move disks A and B such that the total length of the boundary of D inside $\hat{D}$ decreases: (i) move A and B away from D until A and B are tangent, (ii) rotate A about center of B, so that the distance between centers of A and D increases, until (iii) the tangent point of A and B lies on the boundary of D. Then $\hat{D}$ covers half of the boundary of D since $abb'a'$ is a parallelogram, see Fig. 3. □

We exploit the bounded curvature of Dubins paths to show that if path π approaches disk D by passing between two obstacles D_1 and D_2 that intersect D then either (i) π terminates in or near D, or (ii) π is constrained in terms of the depth of its approach to D.

Lemma 5 (Access lemma). *Let D_1, D_2 be unit disks such that D properly intersects both D_1 and D_2. Assume that the centers of D_1 and D_2 are on the x-axis as in Figure 4. Let D' be the unit disc tangent to both D_1 and D_2 and with center below x-axis. If a Dubins path does not cross the top arc of D then it crosses the top arc of D' at most once.*

Proof (See Figure 4.). Suppose a Dubins path π' crosses the top arc of D' twice. Consider the highest point a of π' between these two crossings. We can place a unit disk A tangent to π' at a. π' cannot cross the top arc of A without crossing the horizontal line through the center of A. Then the top arc of A acts as a barrier for π'. Any way of placing of that barrier must intersect either D_1, D_2 or the top arc of D. Consequently it is impossible either to reach a from the first crossing point or continue from a to the second crossing point. □

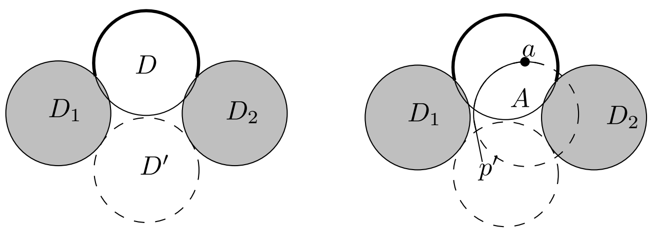

Fig. 4. The top arc of D (bold) is an obstacle.

We are now prepared to prove the main result of this section.

3.2 Proof of Lemma 1

Suppose that π crosses the boundary of D k times. For any two crossing points a and b, if there is a legal ab-path within D then a and b are set to be *D-connected*; otherwise they are *D-obstructed.* Crossing points partition the boundary of D into k arcs. By Lemma 2 if two crossing points a and b are D-connected then we can assume that π contains the shortest legal ab-path within D as a subpath. It follows that any crossing point a is D-connected to at most one other crossing point otherwise π is not a shortest legal st-path.

Let a and b be D-connected crossing points and let π' be the subpath of π connecting a and b. One of the arcs with endpoints a and b must be free of crossing points. Suppose otherwise. Then there must exist two crossing points c and d on either side of π' that are both D-obstructed from a. By Lemma 3 there are two obstructions for pairs a, c and a, d. These obstructions are disjoint within D since they do not obstruct π'. By Lemma 3 two obstructions cover the entire boundary length of D, a contradiction.

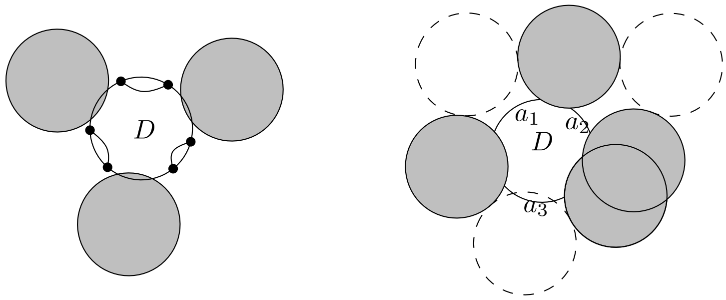

Fig. 5. (a) D-connections. (b) Narrowly exposed arcs a_1 and a_2 and widely exposed arc a_3.

Suppose that s, t are not in D. The crossings of the boundary of D by π form successive pairs. Consider any 3 such pairs. The associated 6 crossings must be connected in D as shown in Fig. 5 (a). Consider disks intersecting the 3 arcs between connections. Each connection corresponds to an arc on the boundary of D. There two types of arcs: narrowly exposed and widely exposed arcs, see Fig. 5 (b). An arc is *widely exposed* if we can place a unit circle tangent to D and disjoint from obstacles.

We first show that only one arc can be widely exposed. Suppose that there are two widely exposed arcs. We assume that the third arc is narrowly exposed. We use capital letters for disks and small letters for their centers. Let E, F be the disks corresponding to the widely exposed arcs. By Lemma 3 there are two different D-obstructions. One exposed arc corresponds to two disks of these D-obstructions. By pushing them toward D (rotation about e or f), we assume that they coincide. Let A, B, C be the three disks in the obstructions.

Case 1. The narrowly exposed arc is obstructed by A and B as in Fig. 6 (a). ec is parallel to df in parallelogram $ecfd$. If A is tangent to both B and E as shown in Fig. 6 (b), then ab, ec and df are parallel and $|ad| = 2$

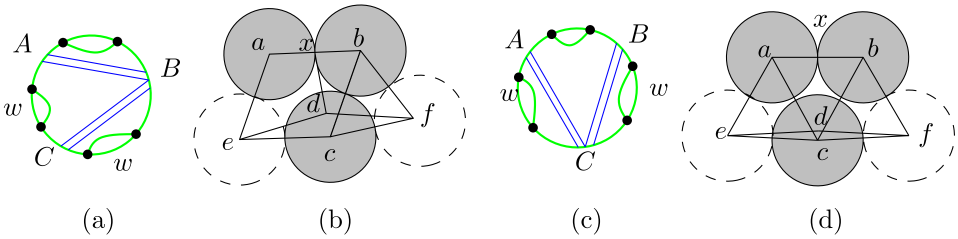

Fig. 6. Two widely exposed arcs are labeled w. D-obstructions are shown by straight line segments.

in parallelogram $abfd$. Let x be the tangent point of A and B. Then $|xd| > |ad| - |ax| = 1$ (since a, x and d are not collinear). If A rotates clockwise about e then the top crossing point x of A and B rotates about b clockwise increasing the distance $|xd|$. Thus $|xd| > 1$ in any case. It contradicts the condition that $A \cap B \subseteq D$.

Case 2. D-obstructions and widely exposed arcs are as in Fig. 6 (c). If disks A and B intersect in D then it is Case 1. They also do not intersect outside D (then the exposed arc is trapped). Rotate disk A about c clockwise until A and B are tangent, see Fig. 6 (d). We assume that ef is horizontal and c is below ef. Since $|ac| \leq 2$ then slope of ab is larger than slope of cf (in quadrilateral $abfc$). On the other hand slope of ab is smaller than slope of ec (in quadrilateral $abce$). Contradiction.

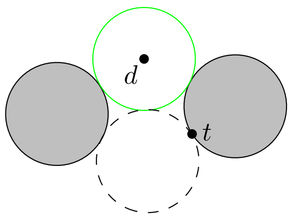

Fig. 7. Narrow pocket.

It follows from Lemma 5 that if π intersects a narrowly exposed arc then one of its endpoints must be in a pocket bounded by disk D'. Thus, at most two arcs can be narrowly exposed. It follows that there are at most 6 crossings in total and, if there are exactly 6 crossings, then both s and t are nearby d. The longest distance from d to t (or s) in a narrow pocket is $\sqrt{3}$, see Fig. 7. Thus if D has three crossings by π both endpoints of π must lie within distance $\sqrt{3}$ of the center of D.

4 Tightening the approximation factor

The algorithm implicit in the last section computes an approximation of the resilience of the sensor arrangement $\mathcal{A}$ by simply finding a shortest path π in the dual graph $\hat{\mathcal{A}}$. In this section we show how the approximation factor for any path π' can be improved by identifying a large collection of subpaths of π' whose individual subpath lengths all overcount the number of distinct sensor intersections. We record these subpaths as *shortcut* edges and combine as many shortcuts as possible to provide a discounted path length (and hence a tighter resilience estimate) for the pair s, t.

The details of this refined algorithm, particularly its analysis, are quite involved. First we develop two types of easily identifiable shortcut edges. Next we argue that we can find, among these shortcut edges, one edge associated with each doubly visited disk, that together form a *weakly compatible* set. Finally, we show that any weakly compatible set S of shortcut edges has a subset of size at least $|S|/3$ that forms a *strongly compatible* set. Since the discount achieved by our algorithm equals or exceeds the size of the largest strongly compatible subset of shortcut edges, it follows that if the minimum resilience path π' makes double visits to d disks (i.e. its resilience estimate provided by the unmodified shortest path algorithm exceeds its true resilience by d) then our modified algorithm provides a resilience estimate that exceeds the true value by at most $2d/3$.

5 Extensions

As we explained in the introduction, our results rely heavily on the assumption that the sensor regions associated with individual sensors are disks. However total congruence of sensor regions is not essential. Although it would result in weaker bounds, our arguments could still be applied if coverage regions were disks with radius between r_0 and r_1.

We have implicitly assumed that the underlying surveillance domain is structured in such a way that it does not restrict the feasibility of paths through the sensor arrangement. In general, it might be of interest to model physical obstacles, either to paths or sensor coverage, in trying to more accurately evaluate the barrier resilience of sensor networks. Of course, if obstacles are disks (or can be expressed as the union of disks) then it is straightforward to extend our existing results. More generally, obstacles can have a significant impact on our multiple visitation bounds. It is not hard to construct examples in which long thin obstacles force multiple crossings of one sensor. Even if we assume that thin obstacles are "walls" and the coverage region of a sensor is reduced to the region which is the connected component of D (minus the obstacles) that contains the center of D, this is not sufficient to prevent multiple crossings.

References

1. M. Cardei and J. Wu. Coverage in wireless sensor networks. In M. Ilyas and I. Mahgoub, editors, *Handbook of Sensor Networks: Compact Wireless and Wired Sensing Systems*, chapter 19, pages 432–446. CRC Press, 2005.
2. R. D. Carr, S. Doddi, G. Konjevod, and M. Marathe. On the red-blue set cover problem. In *Proc. of the 11th Annual ACM-SIAM Symposium on Discrete algorithms (SODA'00))*, pages 345–353, 2000.
3. A. Chen, S. Kumar, and T. H. Lai. Designing localized algorithms for barrier coverage. In *Proc. of the 13th Annual ACM International Conference on Mobile Computing and Networking (MobiCom'07)*, pages 63–74, 2007.
4. Ai Chen, Ten H. Lai, and Dong Xuan. Measuring and guaranteeing quality of barrier-coverage in wireless sensor networks. In *MobiHoc '08: Proceedings of the 9th ACM international symposium on Mobile ad hoc networking and computing*, pages 421–430, New York, NY, USA, 2008. ACM.
5. U. Feige. A threshold of $\ln n$ for approximating set cover. *J. ACM*, 45(4):634–652, 1998.
6. D. W. Gage. Command control for many-robot systems. In *Proc. 19th Annu. AUVS Tech. Symp., AUVS- 92, Hunstville AL*, pages 28–34, 1992.
7. R. Jacob, G. Koniedov, S. Krumke, M. Marathe, R. Ravi, and H. Wirth. The minimum label path problem. In *Unpublished manuscript, Los Alamos National Laboratory*, 1999.
8. C-D Jiang and G-L Chen. Double barrier coverage in dense sensor networks. *J. Comput. Sci. and Tech.*, 23(1):154–165, 2008.
9. S. Kloder and S. Hutchinson. Barrier coverage for variable bounded-range line-of-sight guards. In *Proc. of IEEE International Conference on Robotics and Automation (ICRA'07)*, pages 391–396, 2007.
10. S. Kumar, T. H. Lai, and A. Arora. Barrier coverage with wireless sensors. In *Proc. of the 11th Annual International Conference on Mobile Computing and Networking (MobiCom'05)*, pages 284–298, 2005.
11. S. Kumar, T. H. Lai, and A. Arora. Barrier coverage with wireless sensors. *Wireless Networks*, 13(6), 2007.
12. Santosh Kumar. *Foundations of coverage in wireless sensor networks.* PhD thesis, Columbus, OH, USA, 2006. Adviser-Lai, Ten H.
13. Loukas Lazos, Radha Poovendran, and James A. Ritcey. Probabilistic detection of mobile targets in heterogeneous sensor networks. In *IPSN '07: Proceedings of the 6th international conference on Information processing in sensor networks*, pages 519–528, New York, NY, USA, 2007. ACM.
14. Benyuan Liu, Olivier Dousse, Jie Wang, and Anwar Saipulla. Strong barrier coverage of wireless sensor networks. In *MobiHoc '08: Proceedings of the 9th ACM international symposium on Mobile ad hoc networking and computing*, pages 411–420, New York, NY, USA, 2008. ACM.
15. A. Meguerdichian, F. Koushanfar, G. Qu, and N. Potkonjak. Exposure in wireless ad-hoc sensor networks. In *Proc. of the 7th Annual International Conference on Mobile Computing and Networking (MobiCom'01)*, pages 139–150, 2001.
16. S. Meguerdichian, F. Koushanfar, M. Potkonjak, and M. Srivastava. Coverage problems in wireless ad-hoc sensor networks. In *Proc. of the 20th Annual Joint Conference of the IEEE Conference on Computer and Communications (INFO-COM'01)*, volume 3, pages 1380–1387, 2001.

17. S. Yuan, S. Vannat, and J. P. Juex. Minimum-color path problems for reliability in mesh networks. In *Proc. of the 24th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'05)*, volume 4, pages 2658–2669, 2005.