On k-Guarding Polygons

Daniel Busto* William Evans* David Kirkpatrick*

Abstract

We describe a polynomial time $O(k \log\log \mathrm{OPT}_k(P))$-approximation algorithm for the k-guarding problem of finding a minimum number, $\mathrm{OPT}_k(P)$, of vertex guards of an n-vertex simple polygon P so that for every point $p \in P$, the number of guards that see p is at least the minimum of k and the number of vertices that see p. Our approach finds $O\left(\frac{k}{\varepsilon} \log\log \frac{1}{\varepsilon}\right)$ size (k, ε)-*nets* for instances of the k-hitting set problem arising from the k-guarding problem. These nets contain k distinct elements (or the entire set if it has fewer than k elements) from any set that has at least an ε fraction of the total weight of all elements. To find a nearly optimal k-guarding, we slightly modify the technique of Brönnimann and Goodrich [4] so that the weights of all elements remain small, which is necessary for our (k, ε)-net finder. Our approach, generalizes, simplifies, and corrects a subtle flaw in the technique introduced by King and Kirkpatrick [11] to find small ε-nets for set systems arising from 1-guarding instances.

1 Introduction

In the classic art gallery problem one is given a simple polygon P in the plane and asked to find a smallest subset G of P, called guards, so that every point $p \in P$ is seen by at least one guard $g \in G$ (i.e., $\overline{gp} \subseteq P$). In this paper, we allow only *vertex guards*, meaning G is a subset of V, the vertices of polygon P. The original results on this problem addressed the extremal question of how many guards are needed to guard a simple polygon with n vertices. Chvátal [6], answering a question of Klee, showed that $\lfloor n/3 \rfloor$ guards are occasionally necessary and always sufficient to guard such polygons. O'Rourke's book [16] and Urrutia's chapter [18] describe the history and subsequent flurry of results in this area.

In some applications, we would like every point in P to be seen by more than one guard. For example, to use triangulation to locate an intruder, his position must be seen by at least two guards (whose locations are different). Surprisingly, the generalization of Klee's question to this form of 2-guarding[1] seems to have taken over 30 years to appear, though, earlier, Belleville et al. [2] addressed a different form of k-guarding where each guard must be an interior point of a distinct edge of P. Salleh [17] showed that $\lfloor 2n/3 \rfloor$ guards are occasionally necessary and always sufficient to 2-guard a simple n-gon; and that for 3-guarding convexly quadrilateralizable n-gons, the bound is $\lfloor 3n/4 \rfloor$ guards. A simple proof [3, 15] follows from Fisk's triangulation colouring proof [8] of Chvátal's result. If we insist that guards must be at different vertices then there are simple polygons that cannot be k-guarded in this way for $k \geq 4$ because some points in P are seen by fewer than k vertices. While our original motivation concerned 2-guarding, which is always possible, our results apply to k-guarding in general, if we only require that the guarding do as well as it can for un-k-guardable points. Thus we say that a subset G of the vertices of P is a *k-guarding* of P if for every point $p \in P$ the number of guards that see p is at least the minimum of k and the number of vertices that see p.

Another option is to allow multiple guards at the same vertex. A multiset G of vertices of P is a *multi-k-guarding* of P if every point $p \in P$ is seen by at least k guards in G. Since k copies of any 1-guarding is a multi-k-guarding, $k\lfloor n/3 \rfloor$ guards are always sufficient to multi-k-guard any n-gon, and Chvátal's "comb" 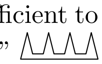 shows they are occasionally necessary. While the smallest multi-2-guarding is at most twice the smallest 1-guarding, a smallest 2-guarding may be much larger. An s-spiked fan 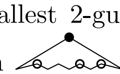 admits a 1-guarding (black dot) of size one and requires at least $\lceil s/2 \rceil + 1$ guards (all dots) to 2-guard. On the other hand, a 2-guarding may be smaller than twice the smallest 1-guarding . We address the problem of finding the smallest number of vertex guards, $\mathrm{OPT}_k(P)$, needed to k-guard a given simple polygon P. In Section 3, we show that the problem is NP-hard for $k > 1$. Note, however, that for certain classes of polygons such as spiral polygons, a smallest 2-guarding can be found efficiently [3]. The hardness of the problem motivates consideration of approximation algorithms.

For multi-k-guarding, one option is to use k copies of a ρ-approximation of a smallest 1-guarding to produce a $k\rho$-approximation of a smallest multi-k-guarding. For k-guarding, this is not an option since a k-guarding cannot place multiple guards at a single vertex. Even a k-step

*Department of Computer Science, University of British Columbia, {busto,will,kirk}@cs.ubc.ca

[1] Belleville [1] uses the term "two-guarding" to mean guarding the entire polygon using only two guards, which is different from our usage.

approach that approximates an optimal 1-guarding using only vertices that don't appear as guards in previous steps is difficult to make work; the relationship between the size of the resulting k-guarding and a smallest k-guarding is not clear.

Fusco and Gupta [9] introduced the notion of (k, ε)-nets to find k-covers of sets in the context of sensor networks. Their result implies an $O(\log \mathrm{OPT}_k(P))$-approximation algorithm for minimally k-guarding P. Chekuri, Clarkson, and Har-Peled [5] described a technique for set multicover with potentially different coverage requirements for each set that also implies an $O(\log \mathrm{OPT}_k(P))$-approximation algorithm. We describe a polynomial time $O(k \log\log \mathrm{OPT}_k(P))$-approximation algorithm that generalizes the technique King and Kirkpatrick [11] used to show an $O(\log\log \mathrm{OPT}_1(P))$-approximation algorithm for 1-guarding P.

2 k-guarding as k-hitting

A *k-hitting set* of a set $\mathcal{R}$ of sets is a set H such that $|H \cap R| \geq \min\{k, |R|\}$ for all $R \in \mathcal{R}$. Let V_p be the set of potential guards that can see p. Let $\mathcal{R} = \{V_p \mid p \in P\}$. A k-hitting set of $\mathcal{R}$ is a k-guarding of P.

Let V be the set of n potential guard locations. Let $w(v)$ be the weight of potential guard location v and $w(R) = \sum_{v \in R} w(v)$. Rather than k-hitting every set in $\mathcal{R}$, we will be happy with k-hitting the heavy sets, those with weight at least $\varepsilon w(V)$. Such a k-hitting set is called a *(k, ε)-net* for the weighted set system $(V, \mathcal{R}, w)$. The algorithm for finding a k-hitting set initially sets these weights to 1 and updates them so that a (k, ε)-net for the weighted set system will be a k-hitting set for $\mathcal{R}$.

Suppose we can find a (k, ε)-net for $(V, \mathcal{R}, w)$ of size $s(1/\varepsilon)$ for some function s. We slightly modify the technique of Brönnimann and Goodrich [4] to obtain a k-hitting set of size $s(4c)$ where c is the smallest k-hitting set. Assume we know that c is the size of the smallest k-hitting set. We find a $(k, 1/2c)$-net N of size $s(2c)$. If there is a set $R \in \mathcal{R}$ (called a *witness*) that is not k-hit, we double the weight of the guards in R *that are not in the net* N. (The italics indicate our modification.) Since $w(R) \leq w(V)/2c$, the weight of V is multiplied by at most a factor $1 + 1/2c$. If H is a k-hitting set of size c, $(H \cap R) \setminus N$ is not empty and the weight of some element in H is doubled. It follows from the proof of Lemma 3.4 [4] that:

Lemma 1 *If there is a k-hitting set of size c, the weight-doubling process iterates at most $4c \lg(n/c)$ times before finding a k-hitting set.*

The rest of their algorithm is a doubling search for the correct value of c. Since the search will succeed when the guess is at most $2c$, the size of the k-hitting set is at most $s(4c)$. The running time is dominated by the time for the last weight-doubling process which involves $O(c \log(n/c))$ invocations of the (k, ε)-net finder (and witness finder).

King and Kirkpatrick [11] show how to find a $(1, \varepsilon)$-net of size $O\left(\frac{1}{\varepsilon} \log\log \frac{1}{\varepsilon}\right)$ for the weighted set systems arising from 1-guarding. To approximately k-guard, we require small (k, ε)-nets for the weighted set systems arising from k-guarding. In Section 6, we show how to find such (k, ε)-nets of size $O\left(\frac{k}{\varepsilon} \log\log \frac{1}{\varepsilon}\right)$. It will be important that the weight of each element remains small. Our modified weight-doubling process insures this. A possible alternative is to solve a linear program to obtain the weights before invoking an ε-net finder once [14, 7]. However, it seems difficult to modify the program to insure that the weights remain small enough that our (k, ε)-net finder will produce a near-optimal k-hitting set.

3 Hardness of k-guarding

Lee and Lin [13] show that minimally 1-guarding a simple polygon is NP-hard. We extend this result to k-guarding for $k > 1$.

Theorem 2 *Finding a minimum k-guarding of a simple polygon is NP-hard for all $k > 1$.*

Proof. We reduce from the problem of minimally 1-guarding a terrain (an x-monotone polygonal curve), which is known to be NP-hard [12]. Given a terrain T, place a concave, x-monotone path $v_1, v_2, \ldots, v_{k-1}$ sufficiently far above the terrain so that each v_i can see all of T. This can be done in polynomial time by intersecting the positive halfplanes coincident with edges of T. Create a polygon P by connecting v_1 to the leftmost and v_{k-1} to the rightmost vertex of T, as shown in Fig. 1. Now P contains a k-guarding of size h if and only if T contains a 1-guarding of size $h - k + 1$. Clearly, adding $v_1, v_2, \ldots, v_{k-1}$ to any 1-guarding of T creates a k-guarding of P. Let G be a k-guarding of P. Let ℓ be the number of v_i in G. Removing all v_i and any $k-1-\ell$ terrain vertices from G creates a 1-guarding of T. $\square$

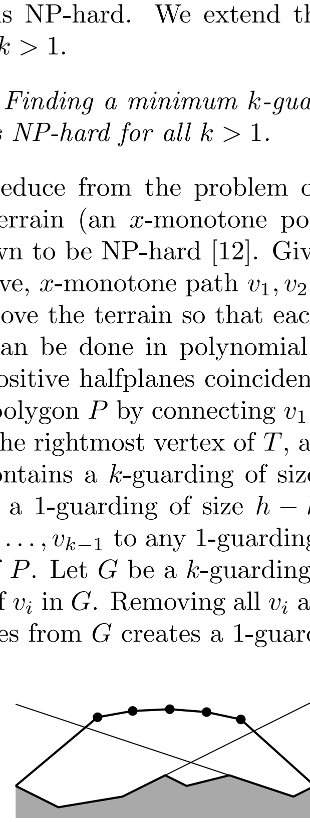

Figure 1: The construction used to show k-guarding is NP-hard (for $k = 6$).

A simple extension of Lin and Lee's original proof shows that minimally multi-k-guarding a simple polygon is NP-hard, but only for odd k. Their proof is based on "critical locations" for guards; a guard at one location is equivalent to assigning a variable true and at the other to false. With odd k whichever location is given more guards can be interpreted as the choice for the corresponding variable. With even k an optimal guard set may have the same number of guards at each location.

4 Distinguished vertices

We describe in this section a method for choosing vertex guards in a vertex-weighted polygon P so that any vertex that sees a large fraction of the total weight will be seen by k guards. The basic idea is to partition the boundary of P into fragments of consecutive edges so that all fragments have approximately equal weight, and to place guards at extreme points of visibility between fragments. Lemma 3 implies that this is a good choice of guards. In the following section, we adopt a hierarchical fragmentation scheme to insure that the size of the guard set is small.

Let ∂P denote the boundary of polygon P. We refer to any sequence of consecutive edges on ∂P as a *(boundary) fragment* of P. Two fragments are *adjacent* if they share an endpoint. We say that a fragment f is *visible* from a point p of P if there is a point q on f such that the open line segment pq lies in the interior of P. The *extreme points of visibility* of fragment f from a set of points S are the first point in f visible from some point in S and the last point in f visible from some point in S. It follows from the simplicity of P that if f is visible from p then the points on f that are visible from p span a contiguous sector of angles centered at p determined by the extreme points of visibility of f from p. If the sectors of one or more visible fragments together have a span greater than π then the corresponding fragments are said to *surround* p; clearly at most one fragment has this property in isolation. Two fragments a and b, visible from p, are *clockwise consecutive from* p if the clockwise extreme visibility angle of a coincides with the counterclockwise extreme visibility angle of b. Suppose that a and b are clockwise consecutive from p. Then, if the clockwise extreme visibility point of a is no further from p than the counterclockwise extreme visibility point of b, then a is said to support a right tangent from p; otherwise b is said to support a left tangent from p, Fig. 2(a).

Following King and Kirkpatrick [11], we consider various partitions of ∂P into fragments, based in part on the weights associated with the vertices of P. [2] For any such fragmentation F of ∂P, we *distinguish* those vertices of P that coincide with extreme points of visibility between some pair of fragments in F, Fig. 2(b).

[2] As indicated earlier, there is a subtle, though significant, flaw in the construction described in [11]. This is discussed in more detail in Appendix A.

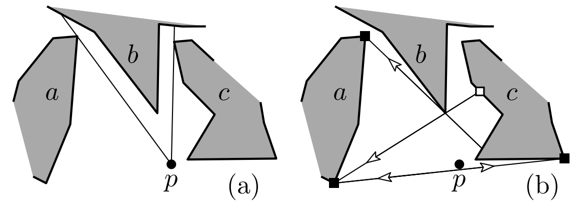

Figure 2: Fragments a, b, and c are clockwise consecutive from p. (a) Fragment a supports a right tangent and c supports a left tangent from p. (b) Squares are extreme points of visibility between fragments a and c. Filled squares are distinguished vertices.

The intuition that these distinguished vertices serve as a good choice for guard locations is based on the following geometric lemma:

Lemma 3 (Lemma 2 [11]) *Let a and b be clockwise consecutive fragments of a fragmentation F, visible from a point p, that do not surround p. If a supports a left tangent from p (resp. if b supports a right tangent from p) then p sees at least one of the distinguished vertices on a (resp. on b).*

The preceding lemma is enough to show that if p sees many visible fragments then it must see many different distinguished vertices, which is a generalization of King and Kirkpatrick's Lemma 1 for $k \geq 1$:

Lemma 4 *Any point p of P that sees at least $2k+3$ fragments of F in total must see a distinguished vertex on at least k fragments.*

Proof. If p sees at least $2k+3$ fragments then it sees a set T of at least $2k$ consecutive fragments each of which cannot pair with a consecutive fragment to surround p, otherwise there would be two disjoint pairs that both span more than π. If a fragment in T doesn't have a distinguished vertex then it has no tangent from p and its consecutive fragment(s) in T have tangents and thus distinguished vertices, by Lemma 3. Hence, at least $|T|/2 \geq k$ fragments have distinguished vertices. □

Remark. If we divide ∂P into $\frac{2k+2}{\varepsilon}$ equal weight fragments where each vertex has weight 1 then it follows that every point p that sees more than a fraction ε of the total weight $w(V) = n$, sees more than $2k+2$ fragments, and by Lemma 4, must see at least k among the set of $O\left(\frac{k^2}{\varepsilon^2}\right)$ distinguished vertices associated with this fragmentation. Thus the distinguished vertices form a (k, ε)-net for the weighted set system $(V, \mathcal{R}, w(v) = 1)$. It remains to find other fragmentations that will work for more weight functions and whose associated distinguished vertices will provide a smaller net.

5 A net-finder based on hierarchical fragmentation

In the last section we showed that by placing guards at a set of $O\big(\frac{k^2}{\varepsilon^2}\big)$ distinguished vertices associated with a flat fragmentation of ∂P, we can ensure that any point that sees fewer than k of these guards sees less than an ε fraction of the total weight. In this section we discuss how hierarchical fragmentation can be used to reduce the number of guards required to $O\big(\frac{k}{\varepsilon}\log\log\frac{1}{\varepsilon}\big)$. It will suffice to prove the result for $k \le \frac{\lg(1/\varepsilon)}{r\lg\lg(1/\varepsilon)}$, for a sufficiently large constant r, since otherwise the (k,ε)-net-finder of Fusco and Gupta [9] or Chekuri, Clarkson, and Har-Peled [5] can be used. It will be clear that the construction of our (k,ε)-net takes polynomial time. (In fact, with some care, it can be constructed in $O(n^3)$ time [10].)

We can represent the hierarchical fragmentation as a tree. At the root there is a single fragment representing the entire boundary ∂P. This root fragment is broken up into a certain number of child fragments. Fragmentation continues recursively until a specified depth t is reached. The integer t is chosen so that

$$(t-1)+2^{t-1} < \lg\frac{1}{\varepsilon} \le t+2^t. \qquad (1)$$

Note that $\lg\lg\frac{1}{\varepsilon}-1 < t < \lg\lg\frac{1}{\varepsilon}+1$, which implies $2^t/t > \frac{\lg(1/\varepsilon)}{2(\lg\lg(1/\varepsilon)+1)} > \frac{\lg(1/\varepsilon)}{3\lg\lg(1/\varepsilon)}$, (provided $\varepsilon < 1/16$), which by our assumption is at least $rk/3$.

The number of children of a fragment f depends on both t and the level of f in the tree. Specifically, if $e = t+2^t-\lceil\lg\frac{1}{\varepsilon}\rceil$ (note $0 \le e \le 2^{t-1}$), then b_i, the number of children of fragments at level $i-1$, is:

$$b_i = \begin{cases} \beta_k t\cdot 2^{2^{t-1}+1}\cdot 2^{1-e} & , \quad i=1 \\ 2^{2^{t-i}+1} & , \quad 1<i\le t\ , \end{cases}$$

where β_k is a linear function of k that will be specified later. Let ϕ_i be the number of fragments at level i.

$$\phi_i = \begin{cases} 1 & , \quad i=0 \\ \beta_k t\cdot 2^{2^t-2^{t-i}-e+i+1} & , \quad 0<i\le t \end{cases}$$

since $\phi_i = \prod_{j=1}^{i} b_j = \beta_k t\cdot 2^{1-e}\cdot\prod_{j=1}^{i}2^{2^{t-j}+1} = \beta_k t\cdot 2^{1-e+i+\sum_{j=1}^{i}2^{t-j}} = \beta_k t\cdot 2^{2^t-2^{t-i}-e+i+1}$. Note that

$$\phi_t = \beta_k t\cdot 2^{2^t+t-e} = \beta_k t 2^{\lceil\lg\frac{1}{\varepsilon}\rceil} \ge \beta_k t\frac{1}{\varepsilon}. \qquad (2)$$

Each collection of child fragments with the same parent fragment f, together with the complement $\overline{f}$ of f, defines a fragmentation F_f of ∂P. The guard set, D_{HF}, is the union, over all parent fragments f in the tree, of the set of vertices that are distinguished by fragmentation F_f. The total number of guards chosen is

$$|D_{HF}| \ \le\ 4\sum_{i=1}^{t}\binom{b_i+1}{2}\phi_{i-1} \ \le\ b_1^2+4\sum_{i=2}^{t}b_i^2\phi_{i-1}\ .$$

Since

$$b_1^2 = (\beta_k t\cdot 2^{2^{t-1}+1}\cdot 2^{1-e})^2 = \beta_k t\cdot 2^{2^t+t-e}(\beta_k t\cdot 2^{4-t-e})$$
$$< \beta_k t\cdot 2^{\lceil\lg\frac{1}{\varepsilon}\rceil}\cdot(\beta_k t\cdot 2^{4-t-e}) \le 2\beta_k t\cdot\frac{1}{\varepsilon}\ ,$$

for $2^t/t \ge 16\beta_k$ (which holds when $rk/3 \ge 16\beta_k$), and

$$\sum_{i=2}^{t}b_i^2\phi_{i-1} = \sum_{i=2}^{t}b_i\phi_i = \sum_{i=2}^{t}(2^{2^{t-i}+1})(\beta_k t 2^{2^t-2^{t-i}-e+i+1})$$
$$= \beta_k t\cdot 2^{2^t-e+2}\sum_{i=2}^{t}2^i < \beta_k t\cdot 2^{2^t+t-e+3}$$
$$\le 8\beta_k t\cdot 2^{\lceil\lg\frac{1}{\varepsilon}\rceil} \le 16\beta_k t\cdot\frac{1}{\varepsilon}\ ,$$

we know,

$$|D_{HF}| = O\bigg(\beta_k\frac{1}{\varepsilon}\log\log\frac{1}{\varepsilon}\bigg). \qquad (3)$$

We must now provide a generalization of Lemma 4 that works with our hierarchical fragmentation.

Lemma 5 *Any point p in P that sees fewer that k vertices in D_{HF} sees no more than $(8k+4)i+1$ fragments at level i of the hierarchical fragmentation.*

The proof of Lemma 5 uses a potential function defined on fragments. A *narrow* fragment is a visible fragment with sector at most π. A *wide* fragment is a visible fragment that isn't narrow. The *potential* of a visible fragment f (wide or narrow) is given by $2g+1-t$, where g is the number of guards in D_{HF} from fragment f that see p, and t is the number of tangents from p to f. Note that g counts all the guards on f in D_{HF}, which includes distinguished vertices in fragmentations F_a for all descendants a of f in the tree. We call a visible fragment *potent* if it is wide or it is narrow and has positive potential; otherwise we call it *impotent*. In addition to potent and impotent fragments, both of which are visible to p, there are also non-visible fragments in the tree, which we regard as having potential zero.

Lemma 6 *The potential of f is at least the total potential of its children.*

Proof. Let $f_1, f_2, \ldots, f_c$ be the visible children of f. Let g_i be the number of guards in D_{HF} from f_i that see p, and t_i be the number of tangents from p to f_i. We want to show that $2g+1-t \ge \sum_{i=1}^{c}(2g_i+1-t_i)$. The number of tangents to $\overline{f}$ is $2-t$ and the total number of tangents to visible fragments in F_f is $c+1$. Thus $2-t+\sum_{i=1}^{c}t_i = c+1$, which implies $\sum_{i=1}^{c}(1-t_i) = 1-t$. The lemma follows since $2g \ge \sum_{i=1}^{c}2g_i$. □

Lemma 7 *An impotent fragment f has at most one visible child and that child is impotent.*

Proof. Let f be an impotent fragment with potential $2g + 1 - t \leq 0$, which implies $g = 0$ and $t \geq 1$. If f has two visible children then the (at least one) child with f's tangent contains a visible distinguished vertex by Lemma 3, which contradicts $g = 0$. Thus f has at most one visible child and, since it shares all tangents with f, its potential is at most $1 - t \leq 0$. □

Lemma 8 *The number of potent fragments at any level is at most the total potential at that level plus two.*

Proof. Potent fragments have positive potential, except for the (at most one) wide fragment that has potential at least -1. □

Proof of Lemma 5. If p sees fewer than k vertices in D_{HF} then the potential of the root is at most $2(k-1)+1$. By Lemmas 6 and 8, the number of potent fragments at any level is at most $2(k-1)+3$. By Lemma 7, every impotent fragment has at most one visible child, which is impotent. Since an impotent child cannot contain a distinguished vertex, any fragment can have at most four impotent children (by Lemma 4 with $k = 1$). Thus every potent fragment has at most four impotent children and its other visible children all have positive potential. Since the total number of potent fragments remains at most $2(k-1)+3$, the number of visible fragments increases by at most $4 \times \#$ potent fragments $\leq 8k + 4$ at each level. Thus the number of visible fragments at level i is no more than $(8k+4)i + 1$. □

6 Near optimal k-guarding

The discussion of the hierarchical fragmentation scheme in the previous section did not take into account the weight of the vertices V of P in the weighted set system $(V, \mathcal{R}, w)$. The weights play a role in the selection of the children of a parent fragment. We choose the children to have approximately equal weight and to contain an integral number of vertices. Both these requirements may be impossible to satisfy if the weights of vertices are very different. The potential problem is that a fragment at level i may have fewer than b_i vertices and thus cannot produce b_i child fragments in the hierarchy. To address this problem, we keep the weights of all vertices small and build the tree bottom up.

Let $w_{\max}$ be the maximum weight of a vertex $v \in V$. Let $\gamma = w(V)/\phi_t$ be the target weight of a leaf, i.e. a fragment at level t in the tree. We create the tree from the bottom up by fragmenting the perimeter of the polygon P into ϕ_t leaf fragments where each leaf fragment f has weight $w(f)$ that satisfies $\gamma - w_{\max} \leq w(f) \leq \gamma + w_{\max}$. One way to do this is to imagine fragmenting the perimeter into ϕ_t equal weight pieces, which may split some vertices in two, and then putting any vertex that is split into the last (clockwise-most) fragment in which it occurs. As long as $w_{\max} < \gamma$, the resulting fragmentation has at least one vertex in every fragment. We then combine each collection of b_i adjacent fragments ($i = t$) to form their parent fragment at level $i - 1$ and repeat for all i down to $i = 1$. This creates the tree from which we extract the set of guards D_{HF}. Applying Lemma 5 with $i = t$, choosing $\beta_k = 8k + 5$, and using equations (2) and (3), we get

Lemma 9 *Given a weighted set system $(V, \mathcal{R}, w)$ such that $w(v) < \varepsilon w(V)/(\beta_k t)$ for all $v \in V$ where t satisfies equation (1), the set D_{HF} is a (k, ε)-net for $(V, \mathcal{R}, w)$ of size $O\big(\frac{k}{\varepsilon} \log\log \frac{1}{\varepsilon}\big)$.*

To keep our slightly modified version of the technique of Brönnimann and Goodrich [4] from increasing the weight of any vertex to γ or more, we add all vertices v of weight $w(v) \geq \gamma/2$ to D_{HF}. Since no such vertex has its weight doubled, because of our modification, the maximum weight that any vertex can attain is less than γ. The number of such vertices is at most $2w(V)/\gamma \leq 2\beta_k t \frac{1}{\varepsilon}$, which increases the size of D_{HF} by only a (small) constant factor. Our main theorem follows:

Theorem 10 *For a simple polygon P with n vertices, we can find a k-guarding set of vertices for P of size $O(k\,OPT_k(P) \log\log OPT_k(P))$, where $OPT_k(P)$ is the size of the minimum k-guarding set of vertices for P, in time polynomial in n.*

7 Extensions and open questions

We have shown how to get an $O(k \log\log \mathrm{OPT}_k(P))$-approximation to the minimum k-guarding of a simple polygon P. This improves the $O(\log \mathrm{OPT}_k(P)$-approximation algorithms of Fusco and Gupta [9] and Chekuri, Clarkson, and Har-Peled [5], when $k = o\Big(\frac{\log \mathrm{OPT}_k(P)}{\log\log \mathrm{OPT}_k(P)}\Big)$. It would be interesting to know if the dependence on k in our algorithm can be eliminated or reduced.

Our version of k-guarding models situations in which the guarding requirement of every point of P is the same. We addressed the non-uniformity that arises simply because points of P may not see as many as k vertices of P, by reducing the guarding requirement of a point p to the minimum of k and the number of vertices that see p. In general, a uniform guarding requirement might be undesirable for other reasons as well; some (perhaps most) points may not need to be guarded at all, whereas others may need extraordinary attention. To capture this variety, we suppose that all points $p \in P$ have an associated non-negative guarding demand $d(p)$, and we seek a *demand-guarding* of P, a guard set G that satisfies the individual guarding demand of every point in P. The associated optimization problem takes as input a pair (P, d) and asks for a minimum size demand-guarding of P.

In fact, the $O(\log \mathrm{OPT})$-approximation algorithm of Chekuri, Clarkson, and Har-Peled [5] addresses this more general demand-guarding problem, where OPT denotes the size of the optimal demand-guarding set. Obviously, our algorithm could be used to provide an $O(d_{\max} \log\log \mathrm{OPT}_{d_{\max}})$-approximation (by increasing all demands to $d_{\max}$). However, since our algorithm never exploits the uniformity of guarding demands, it is straightforward to confirm that, essentially without modification, it provides an $O(d_{\max} \log\log \mathrm{OPT})$-approximation. Once again, it would be interesting to know if the dependence on $d_{\max}$ can be eliminated or reduced.

Our k-guarding problem assumes that (i) all of the vertices of the polygon are potential guard locations, and (ii) the polygon has no holes. It would be interesting to determine if our techniques can be used to get good approximation algorithms when these constraints are relaxed.

Existing NP-hardness proofs for minimally 1-guarding polygons extend easily to show that minimally multi-k-guarding simple polygons is NP-hard when k is odd. It seems likely that minimally multi-k-guarding is NP-hard when k is even as well.

There is no known non-trivial lower bound on the ratio between the sizes of a minimum k-guarding and a minimum 1-guarding of a polygon. What is the precise relationship between these two values? Are there classes of polygons where the size of a minimum 2-guarding is not much larger than a minimum 1-guarding?

References

[1] P. Belleville. Two-guarding simple polygons. In *Proc. 4th Canadian Conference on Computational Geometry*, pages 103–108, 1992.

[2] P. Belleville, P. Bose, J. Czyzowicz, J. Urrutia, and J. Zaks. K-guarding polygons on the plane. In *Proc. 6th Canadian Conference on Computational Geometry*, pages 381–386, 1994.

[3] S. Bereg. On k-vertex guarding simple polygons. In *Computational Geometry and Discrete Mathematics*, volume 1641 of *Kôkyûroku*, pages 106–113. Research Institute for Mathematical Sciences, Kyoto University, October 2008.

[4] H. Brönnimann and M. T. Goodrich. Almost optimal set covers in finite VC-dimension. *Discrete & Computational Geometry*, 14(4):463–479, 1995.

[5] C. Chekuri, K. L. Clarkson, and S. Har-Peled. On the set multicover problem in geometric settings. *ACM Trans. Algorithms*, 9(1):9:1–9:17, Dec. 2012.

[6] V. Chvátal. A combinatorial theorem in plane geometry. *Journal of Combinatorial Theory, Series B*, 18(1):39–41, 1975.

[7] G. Even, D. Rawitz, and S. Shahar. Hitting sets when the VC-dimension is small. *Inf. Process. Lett.*, 95(2):358–362, 2005.

[8] S. Fisk. A short proof of Chvátal's Watchman Theorem. *Journal of Combinatorial Theory, Series B*, 24(3):374, 1978.

[9] G. Fusco and H. Gupta. ε-net approach to sensor k-coverage. In B. Liu, A. Bestavros, D.-Z. Du, and J. Wang, editors, *Wireless Algorithms, Systems, and Applications*, volume 5682 of *Lecture Notes in Computer Science*, pages 104–114. Springer Berlin Heidelberg, 2009.

[10] J. King. Fast vertex guarding for polygons with and without holes. *Computational Geometry*, 46(3):219–231, 2013.

[11] J. King and D. G. Kirkpatrick. Improved approximation for guarding simple galleries from the perimeter. *Discrete & Computational Geometry*, 46(2):252–269, 2011.

[12] J. King and E. Krohn. Terrain guarding is NP-hard. *SIAM Journal on Computing*, 40(5):1316–1339, 2011.

[13] D. Lee and A. Lin. Computational complexity of art gallery problems. *IEEE Transactions on Information Theory*, 32(2):276–282, 1986.

[14] P. M. Long. Using the pseudo-dimension to analyze approximation algorithms for integer programming. In F. K. H. A. Dehne, J.-R. Sack, and R. Tamassia, editors, *WADS*, volume 2125 of *Lecture Notes in Computer Science*, pages 26–37. Springer, 2001.

[15] K. Mehlhorn, J. Sack, and J. Zaks. Note on the paper "K-vertex guarding simple polygons" [Computational Geometry 42 (4) (May 2009) 352–361]. *Computational Geometry*, 42(6–7):722, 2009.

[16] J. O'Rourke. *Art Gallery Theorems and Algorithms*. The International Series of Monographs on Computer Science. Oxford University Press, New York, NY, 1987.

[17] I. Salleh. K-vertex guarding simple polygons. *Computational Geometry*, 42(4):352–361, 2009.

[18] J. Urrutia. Art gallery and illumination problems. In J.-R. Sack and J. Urrutia, editors, *Handbook of computational geometry*, pages 973–1027. North-Holland Publishing Co., 2000.

A Remark on the hierarchical fragmentation construction of King and Kirkpatrick [11]

A very similar hierarchical fragmentation (differing from ours only in the definition of t, and in the level-1 fragmentation factor b_1) was described by King and Kirkpatrick [11] in developing their approximation bound for optimal 1-guarding. Unfortunately, the choice of α (which, together with t determines b_1) given in their equation (3) does not always guarantee that their equation (1) holds. In particular, consider the case when $1/\varepsilon = 2^{2^{t-1}+1}$ (so $t = \lceil \log\log(1/\varepsilon) \rceil$, as specified). In this case, $\alpha = 1/(4t2^{2^{t-1}+1-t})$ and so $t\alpha 2^{2^t}$ (the bound on $|S_{HF}|$, the size of their guard set) is essentially $2^t \cdot 1/\varepsilon$, which is $\Theta((1/\varepsilon)\log(1/\varepsilon))$, not $O((1/\varepsilon)\log\log(1/\varepsilon))$, as claimed in their equation (1).