

Optimally Guarding Perimeters and Regions with Mobile Range Sensors

Si Wei Feng

Jingjin Yu

Abstract— We investigate the problem of using mobile robots equipped with 2D range sensors to optimally guard perimeters or regions. Given a bounded set in \mathbb{R}^2 to be guarded, and k mobile sensors where the i -th sensor can cover a circular region with a variable radius r_i , we seek the optimal strategy to deploy the k sensors to fully cover the set such that $\max r_i$ is minimized. On the side of computational complexity, we show that computing a 1.152-optimal solution for guarding a perimeter or a region is NP-hard even when the set is a simple polygon or the boundary of a simple polygon, i.e., the problem is hard to approximate. The hardness result on perimeter guarding holds when each sensor may guard at most two disjoint perimeter segments. On the side of computational methods, for the guarding perimeters, we develop a fully polynomial time approximation scheme (FPTAS) for the special setting where each sensor may only guard a single continuous perimeter segment, suggesting that the aforementioned hard-to-approximate result on the two-disjoint-segment sensing model is tight. For the general problem, we first describe a polynomial-time $(2 + \varepsilon)$ -approximation algorithm as an upper bound, applicable to both perimeter guarding and region guarding. This is followed by a high-performance integer linear programming (ILP) based method that computes near-optimal solutions. Thorough computational benchmarks as well as evaluation on potential application scenarios demonstrate the effectiveness of these algorithmic solutions.

I. INTRODUCTION

In this paper, we consider the problem of using mobile robots equipped with range sensors to guard (1D) perimeters or (2D) regions. Given a bounded polygonal one- or two-dimensional set to be secured, and k mobile robots where robot i 's sensor covers a circular region of radius r_i , we seek a deployment of the robots so that $\max r_i$ is minimized. That is, we would like to minimize the maximum single-sensor coverage across all sensors. We denote this multi-sensor coverage problem under the umbrella term *optimal set guarding with 2D sensors*, or OSG_{2D} .¹ The specific problem for guarding perimeters (resp., regions) is denoted as *optimal perimeter (resp., region) guarding with 2D sensors*, abbreviated as OPG_{2D} (resp., ORG_{2D}). Beside direct relevance to sensing, surveillance, and monitoring applications using mobile sensors [1]–[3], OSG_{2D} applies to other robotics related problem domains, e.g., the deployment of ad-hoc mobile wireless networks [4], [5], in which case an optimal solution to OSG_{2D} provides a lower bound on the guaranteed network strength over the targeted 2D region.

S.-W. Feng and J. Yu are with the Department of Computer Science, Rutgers, the State University of New Jersey, Piscataway, NJ, USA. E-Mails: {siwei.feng, jingjin.yu} @rutgers.edu. This work is supported by NSF awards IIS-1734419 and IIS-1845888.

¹The subscript is placed here to distinguish our setup from the OPG problem studied in [1], which assumes a 1D sensing model.

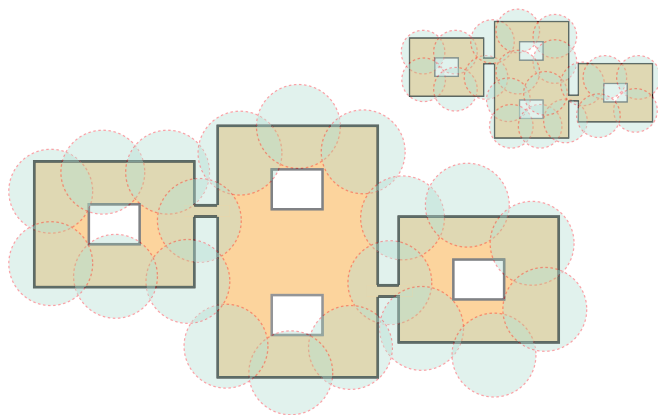


Fig. 1: An illustration of the OSG_{2D} setup and sample solutions. [center] The background shows the footprint of a building, e.g., an apartment complex. Scenarios may arise that a dangerous criminal might be hiding in the building and we would like to closely monitor the outer boundary of the building. For the setting, the shaded discs provide a near-optimal cover with minimum radii for 20 mobile sensors that fully encloses the outer perimeter, computed using algorithms presented in this work with optimality guarantees. [upper right] A near-optimal solution for guarding the interior of the building footprint minus the four holes.

As a summary of the study, on the side of computational complexity, we establish that OPG_{2D} is hard to approximate within a factor of 1.152 even when the perimeter is a simple closed polygonal chain whose length is bounded by the input size, through a reduction from vertex cover on planar 3-regular graphs. A unique property of our reduction is that it shows the inapproximability gap remains when each sensor can cover at most two disjoint perimeter segments. The proof also shows that ORG_{2D} is at least as hard to approximate. Therefore, no polynomial time algorithm may exist that solves OSG_{2D} to better than the 1.152-optimal lower bound, unless $P=NP$. On the algorithmic side, we begin by providing an efficient $(1 + \varepsilon)$ approximation algorithm for a specific class of OPG_{2D} problems in which each mobile sensor must cover a continuous perimeter segment. This implies that the aforementioned inapproximability result on OPG_{2D} under the two-disjoint-segment sensing model is tight. For the general OSG_{2D} problem, we first describe a polynomial time $(2 + \varepsilon)$ approximation algorithm as a reasonable approximability upper bound. Then, an integer linear programming (ILP) model is devised that allows the fast computation of highly optimal solutions for fairly large problem instances. Results described in this paragraph,

together with the introduction of $\text{OSG}_{2\text{D}}$ as a practical multi-robot deployment problem focusing on global optimality, constitute the main contributions of this work.

As an intermediate result toward showing the hardness of the simple polygon coverage problem, we also supply a hardness proof of vertex cover on planar bridgeless² 3-regular graphs, which may be of independent interest.

Related work. Our work on optimal perimeter and region guarding draws inspiration from a long line of multi-robot coverage planning and control research, e.g., [3], [6]–[10]. In an influential body of work on coverage control [3], [6], a gradient based iterative method is shown to drive one or multiple mobile sensors to a locally optimal configuration with convergence guarantees. Whereas [3], [6] assume that the distribution of sensory information is available *a priori*, it is shown that such information can be effectively learned [9]. Subsequently, the control method is further extended to allow the coverage of non-convex and disjoint 2D domains [7] and to work for mobile robots with varying sensing or actuation capabilities [10]. In contrast to these control-based approaches, which produce iterative locally optimal solutions, $\text{OSG}_{2\text{D}}$ emphasizes the direct computation of globally optimal deployment solutions and supports arbitrarily shaped bounded (1D) perimeters and (2D) regions.

Recently, the problems of globally optimally covering perimeters using one-dimensional sensors have been studied in much detail [1], [11]. It is shown that when the sensors are homogeneous, the optimal deployment of sensors can be computed very efficiently, even for highly complex perimeters [1]. On the other hand, the problem becomes immediately intractable, sometimes strongly NP-hard, when sensors are heterogeneous [11]. Our research is distinct from [1], [11] in that we employ a (two-dimensional) range sensing model and work on the coverage of both perimeters and regions, which has much broader applicability.

As pointed out in [3], [9], distributed sensor coverage, as well as $\text{OSG}_{2\text{D}}$, has roots in the study of the facility location optimization problem [12], [13], which examines the selection of facility (e.g., warehouses) locations that minimize the cost of delivery of supplies to spatially distributed customers. In theoretical computer science and operations research, these are known as the k -center, k -means, and k -median clustering problems [14], the differences among which are induced by the cost structure. Our investigation of $\text{OSG}_{2\text{D}}$ benefits from the vast literature on the study of k -center clustering and related problems, e.g., [15]–[19]. These clustering problems are in turn related to packing [20], tiling [21], and the well-studied art gallery problems [22], [23].

Organization. The rest of the paper is organized as follows. In Section II, we introduce the $\text{OSG}_{2\text{D}}$ formulation. Section III is devoted to establishing that $\text{OSG}_{2\text{D}}$ is hard to approximate to better than 1.152-optimal, providing a theoretical lower bound. In Section IV, focusing on the upper bound, we describe algorithms that for $\text{OSG}_{2\text{D}}$ and the

special $\text{OPG}_{2\text{D}}$ variant where a sensor is allowed to cover a continuous perimeter segment. In Section V, we benchmark the algorithms and illustrate two potential applications. We discuss and conclude the work in Section VI.

II. PRELIMINARIES

Let $\mathcal{W} \subset \mathbb{R}^2$ be a polygonal workspace, which may contain one or multiple connected components. A *critical subset* of \mathcal{W} needs to be guarded by k indistinguishable point guards with range sensing capabilities. For example, the workspace may be a forest reserve and the critical subset may be its boundary. Or, the workspace may be a high-security facility, e.g., a prison, and the critical subset the prison yard. The i th guard, $1 \leq i \leq k$, located at $c_i \in \mathbb{R}^2$, can monitor a circular area of radius r_i centered at c_i with r_i being a variable. For example, the guard may be a watchtower equipped with a vision sensor that can detect intruders. As the watchtower’s altitude increases, its sensing range also increases; but its monitoring quality will decrease at the same time due to resolution loss. In this study, we seek to compute the optimal strategy to deploy these k guards so that the required sensing range, $\max_i r_i$, could be minimized.

More formally, we model a connected component of \mathcal{W} as some 2D polygonal region containing zero or more simple polygonal obstacles. For a bounded set $D \subset \mathbb{R}^2$, we define

$$\text{size}(k, D) = \min_{c_1, \dots, c_k \in \mathbb{R}^2} \max_{p \in D} \min_{1 \leq i \leq k} \|c_i - p\|_2$$

and use $B(c, r)$ to denote the disc of radius r centered at a point $c \in \mathbb{R}^2$ (the definition of $\text{size}(k, D)$ is used extensively in later sections). Intuitively, $\text{size}(k, D)$ represents the minimum radius needed such that there exists k circles with radius $\text{size}(k, D)$ that can cover the 2D bounded region D entirely. The main problem studied in this work is:

Problem II.1 (Optimal Set Guarding with 2D Sensors). *Given a polygonal workspace $\mathcal{W} \subset \mathbb{R}^2$, let $D \subset \mathcal{W}$ be a critical subset to be guarded by k robots each with a variable coverage radius of r . Find the smallest r and corresponding robot locations $c_1, \dots, c_k \in \mathbb{R}^2$, such that $D \subset \cup_i B(c_i, r)$.*

For making accurate statements about computational complexity, we make the assumption that the length of $\partial\mathcal{W}$ is bounded by a polynomial with respect to the complexity of \mathcal{W} , (i.e. the number of vertices of the polygon).

For convenience, we give specific names to these optimal guarding problems based critical subset types. If the critical subset belongs to $\partial\mathcal{W}$, we denote the problem as *optimal perimeter guarding with 2D sensors* or $\text{OPG}_{2\text{D}}$. If the critical subset is \mathcal{W} , we denote the problem as *optimal region guarding with 2D sensors* or $\text{ORG}_{2\text{D}}$. When there is no need to distinguish, the problem is denoted as *optimal set guarding with 2D sensors* ($\text{OSG}_{2\text{D}}$).

As an example, to guard the boundary of a plus-shaped polygon with 5 robots, an optimal solution could be Fig. 2 where the inner circle covers 4 disconnected boundary segments, such pattern in the optimal solution also renders $\text{OPG}_{2\text{D}}$ much more difficult than the simplified 1D sensing

²That is, the deletion of any edge does not disconnect the graph.

model studied in [1] (indeed, $\text{OPG}_{2\text{D}}$ becomes hard to approximate, as will be shown shortly). The solution is also optimal under the $\text{ORG}_{2\text{D}}$ formulation.

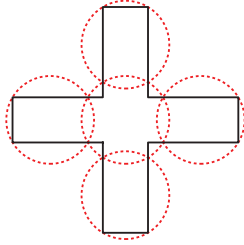


Fig. 2: An example showing an optimal solution of using five discs to cover the plus-shaped polygon. The solution is optimal for both $\text{OPG}_{2\text{D}}$ and $\text{ORG}_{2\text{D}}$ formulations.

III. INTRACTABILITY OF APPROXIMATE OPTIMAL GUARDING OF SIMPLE POLYGON

In this section, we prove that $\text{OSG}_{2\text{D}}$ with the set being a simple polygon is strongly NP-hard to approximate within a factor of $\alpha \approx 1.152$, through a sequence of auxiliary NP-hardness results. First, in Section III-A, we prove an intermediate result that the vertex cover problem is NP-complete on planar bridgeless 3-regular graphs. Next, in Section III-B, starting from a planar bridgeless 3-regular graph, we construct a structure which we call *3-net* and prove the the problem of finding the minimum coverage radius of the 3-net is NP-hard to approximate within α . Then, in Section III-C, we apply a straightforward reduction to transform the 3-net into a simple polygon to complete the hard-to-approximate proof for $\text{OSG}_{2\text{D}}$ for a simple polygon.

We then further show the inapproximability of the special $\text{OPG}_{2\text{D}}$ setup when each robot can only guard at most two disjoint perimeter segments (Section III-D), contrasting the FPTAS for the special $\text{OPG}_{2\text{D}}$ setup when each robot can only guard a continuous perimeter segment in Section IV-A.

A. Vertex Cover on Planar Bridgeless 3-Regular Graph

Our reduction uses the hardness result on the vertex cover problem for planar graphs with maximum degree 3 [24]. Such a vertex cover problem can be fully specified with a 2-tuple (G, k) where $G = (V, E)$ is a planar graph with max degree 3 and k is an integer specifying the allowed number of vertices in a vertex cover. We note that the result has been suggested implicitly in [25]; we provide an explicit account with a simple proof.

Lemma III.1. *Vertex cover on planar bridgeless 3-regular graph is NP-complete.*

Proof. For a given planar graph G with max degree 3 and an integer k , we construct a planar bridgeless 3-regular graph G'' and provide an integer k'' such that G has a vertex cover of size k if and only if G'' has a vertex cover of size k'' .

The reduction first makes G 3-regular by attaching (one or two of) the gadget shown in Fig. 3 to $v \in G$ that are

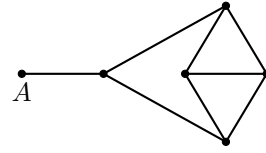


Fig. 3: A gadget that can be attached to a degree one or two vertices (at the point A) in a max degree 3 graph to make all vertices have degree 3. With each addition of the gadget, we increase the vertex cover by a size of 3, regardless of whether A is part of a vertex cover.

not degree 3. This results in a 3-regular graph G' . For each attached gadget, k is bumped up by 3, i.e., we let k' for G' be $k' = k + 3(3|V(G)| - 2|E(G)|)$. It is straightforward to see that G has a vertex cover size of k if and only if G' has a vertex cover size of k' .

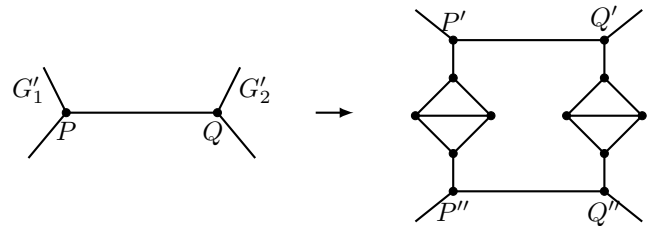


Fig. 4: Transformation that removes bridge PQ and does not introduce new bridges. The minimum vertex cover number is increased by 6 after each transformation.

In the second and last step, we remove bridges in G' . As in Fig. 4, for a bridge PQ that divides G' into G'_1 (containing P) and G'_2 (containing Q), we split the bridge edge PQ using the illustrated transformation, which yields a new graph G'' that is planar, bridgeless, and 3-regular, after all bridges are removed this way. For each such augmentation, the size of the vertex cover is bumped up by 6. Let $br(G')$ be the number of bridges in G' , G' has a vertex cover of size k' if and only if G'' has a vertex cover of size $k'' = k' + 6br(G')$. This completes the proof. \square

B. Hardness on Optimally Guarding A 3-Net

Starting from a planar cubic graph G , we construct a structure that we call 3-net, T_G , as follows. First, similar to [15], to embed G into the plane, an edge $uw \in E(G)$ is converted to an odd length path $uv_1, v_1v_2, \dots, v_{2m}w$ where $m > 3$ is an integer. We note that m is different in general for different edges of G . Denote such a path as $u \cdots w$; each edge along $u \cdots w$ is straight and has unit edge length. We also require that each path is nearly straight locally. For a vertex of G with degree 3, e.g., a vertex $u \in V(G)$ neighboring $w, x, y \in V(G)$, we choose proper configurations and lengths for paths, $u \cdots w$, $u \cdots x$, and $u \cdots y$ such that these paths meet at u forming pairwise angles of $2\pi/3$. We denote the resulting graph as G' , which becomes the *backbone* of the 3-net T_G .

From here, a second modification is made which completes the construction of T_G . In each previously constructed

path $u \cdots w = uv_1 \dots v_{2m}w$, for each $v_i v_{i+1}$, $1 \leq i \leq 2m - 1$, we add a line segment of length $\sqrt{3}$ that is perpendicular to $v_i v_{i+1}$ such that $v_i v_{i+1}$ and the line segment divide each other in the middle. A graphical illustration is given in Fig. 5. G' and the bars form the 3-net, which we denote as T_G . An example of transforming K_4 into a 3-net is given in Fig. 6.

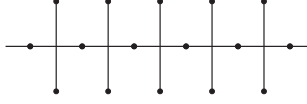


Fig. 5: Structure within the odd length path and attached perpendicular “bars” with length $\sqrt{3}$. Regarding the representation of such non-integral coordinates in the problem input, we may scale the coordinates to some certain extent and round them to integers so that the relative distance between each other is precise enough for the proof.

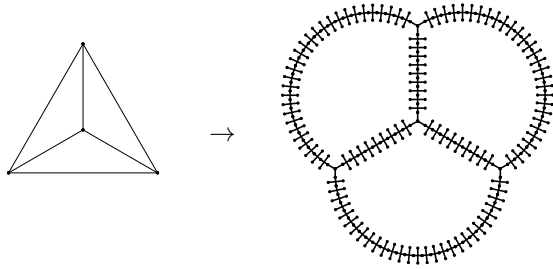


Fig. 6: Illustration of a 3-net obtained from K_4 , the complete graph on 4 vertices.

Let L be the number of (unit length) edges of G' (i.e., $L = \sum_{uv \in E(G')} \text{len}(u \cdots w)$).

Lemma III.2. *A planar bridgeless 3-regular graph G has a vertex cover of size k if and only if its transformed 3-net T_G can be covered by $K = k + (L - |E(G)|)/2$ circles of radius approximately $\alpha = 1.152$.*

We refer the readers to [26] for the lengthy and technical proof. From the proof, it is also clear that Lemma III.2 holds for discs with radius in $[1, \alpha)$. Thus, approximating $\text{size}(T_G, K)$ to less than a factor of α will decide whether G has a vertex cover of size k , yielding the hard-to-approximate result. Also, it can be observed that all lengths are polynomial with respect to the problem input size, which implies strongly NP-hardness.

Theorem III.1. *The minimum radius for cover a 3-net using k circular discs is strongly NP-hard to approximate within a factor of $\alpha \approx 1.152$.*

C. From 3-Net to A Simple Polygon

We proceed to show that OSG_{2D} is hard to approximate for a simple polygon by converting a 3-net into one. Along the backbone G' of a 3-net T_G , we first expand the line segments by δ to get a 2D region (see Fig. 7(a)). We may describe the interior of the resulting polygon as

$$P = \{p \in \mathbb{R}^2 \mid \min_{q \in T_G} (\|p - q\|_1) \leq \delta/2\}$$

For small enough δ , it's clear that P is a polygon with holes. Let $K = ((L - |E|)2 + k)$, it holds that

$$\begin{aligned} \text{size}(K, T_G) &\leq \text{size}(K, P) \leq \text{size}(K, T_G) + \delta, \\ \text{size}(K, T_G) &\leq \text{size}(K, \partial P) \leq \text{size}(K, T_G) + \delta. \end{aligned}$$

To convert the structure into a simple polygon, we can open “doors” of width δ on the structure to get rid of the holes (see Fig. 7(b)). Each opening removes one hole from P . This is straightforward to check; we omit the details.

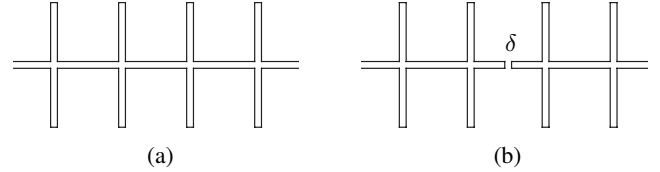


Fig. 7: (a) A 3-net T_G maybe readily converted into a simple polygon P with holes by expanding along its backbone. (b) Creating a “door” of width δ will remove one hole from P .

Denoting the resulting simple polygon as P' , we have

$$\begin{aligned} \text{size}(K, P) - \delta &\leq \text{size}(K, P') \leq \text{size}(K, P), \\ \text{size}(K, \partial P) - \delta &\leq \text{size}(K, \partial P') \leq \text{size}(K, \partial P). \end{aligned}$$

Therefore, both $\text{size}(k, P')$ and $\text{size}(k, \partial P')$ are between $\text{size}(k, T_G) - \delta$ and $\text{size}(k, T_G) + \delta$. Suppose the OSG_{2D} for $\partial P'$ or P' has a polynomial approximation algorithm with approximation ratio $1.152 - \varepsilon$ where $\varepsilon > 0$, let $\delta = \varepsilon/2$, then the optimal guarding problem for the T_G can be approximated within 1.152 disobeying the inapproximability gap by Theorem III.1. Therefore,

Theorem III.2. *OSG_{2D} is NP-hard and does not admit a polynomial time approximation within a factor of α with $\alpha \approx 1.152$, unless $P=NP$.*

D. OPG_{2D} with Sensor Guarding Limitations

The inapproximability gap from Theorem III.2 prompts us to further consider limitations on the setup with the hope that meaningful yet more tractable problems may arise. One natural limitation is to limit the number of continuous segments a mobile sensor may cover. As will be shown in Section IV-A, if a mobile sensor may only guard a single continuous perimeter segment, a $(1 + \varepsilon)$ -optimal solution can be computed efficiently. On the other hand, it turns out that if a sensor can guard up to two continuous perimeter segments, OPG_{2D} remains hard to approximate.

Theorem III.3. *OPG_{2D} of a simple polygon cannot be approximated within $\alpha \approx 1.152$ even when each robot can guard no more than two continuous boundary segments, unless $P=NP$.*

Proof. Due to [27], every bridgeless 3-regular graph G has a perfect matching. We can obtain such a perfect matching of the 3-regular graph using Edmonds Blossom algorithm in polynomial time [28]. Doubling the edges in the perfect matching, we can then obtain a 4-regular graph G' .

With each vertex's degree even on G' , a Eulerian tour exists on G' and can be efficiently computed. For the 3-net T_G , we may incorporate the bars into the Eulerian tour, corresponding to that for G' , as illustrated in Fig. 8.

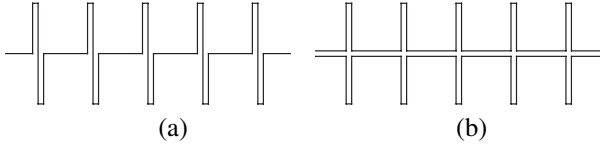


Fig. 8: (a) Part of the augmented Eulerian path for non-doubled paths. (b) Part of the augmented Eulerian path for doubled paths.

The Eulerian tour on T_G may have self-intersections, which will prevent the tour from being a simple polygon. To address this, we may use one of two possible solutions outlined in Fig. 9 to eliminate the self-intersections.

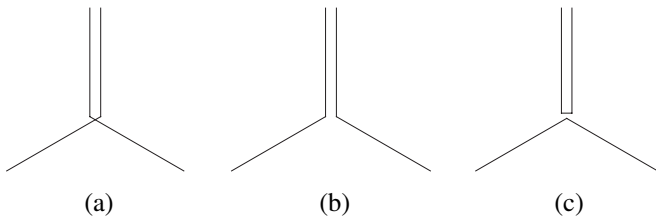


Fig. 9: In order to eliminate possible self-intersections in (a), we may transform it into one of the solutions given in (b) and (c) to make the Eulerian tour remain connected (one of the two solutions will satisfy this).

At this point, we readily observe that Theorem III.2 applies. Furthermore, an optimal solution always allows each mobile sensor to cover only two continuous perimeter segments. This is clear in the middle of any paths of T_G ; at junctions, the polygon boundary will be either one of two possibilities shown in Fig. 10, where a sensor again covers at most two continuous segments of the simple polygon. \square

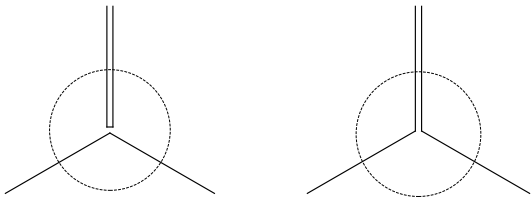


Fig. 10: The figure shows two possible types of boundaries near a vertex with degree of 4. A robot near the vertex will only be able to cover two disjoint but individually continuous boundary segments with sensing radius less than α if the solution is to be optimal.

IV. EFFECTIVE ALGORITHMIC SOLUTIONS FOR OSG_{2D}

In this section, we present several algorithmic solutions for OSG_{2D} . First, a fully polynomial approximation scheme (FPTAS) is presented that solve OPG_{2D} with the additional requirement that each sensor is responsible for a continuous perimeter segment. This contrasts Theorem III.3. Then, we

show that there exist polynomial time algorithms that readily guarantee a $(2 + \varepsilon)$ -approximation for OSG_{2D} . This is followed by an integer linear programming (ILP) method that delivers high-quality solutions (as compared with the $(2 + \varepsilon)$ -approximate one) and has good scalability.

In preparation for introducing the result, we first describe a method that is used for discretizing the problem. For a simple polygon P , we can approximately represent its boundary ∂P as a set of balls with radius ε along ∂P , by splitting ∂P into $N = \lceil \text{len}(\partial P)/(2\varepsilon) \rceil$ continuous pieces of length at most 2ε and putting the balls' centers at their midpoints. Denote set of ε -balls as S_B , and the set of their centers as $S_O = \{o_1, \dots, o_N\}$. Since it holds that $\text{size}(k, \partial P) \leq \text{size}(k, S_B) \leq \text{size}(k, S_O) + \varepsilon \leq \text{size}(k, \partial P) + \varepsilon$, the minimum coverage radius of the discretized version of covering S_O will differ no more than ε from the original problem of covering ∂P . Similarly, for covering the interior of P , we can put P into a grid with cell side length ε , and set the center of the grid cells intersecting with P as S_O , creating at most $N = O((\text{len}(\partial P)/\varepsilon)^2)$ samples. The discretization process converts guarding P or ∂P to guarding S_O .

A. OPG_{2D} with Single Segment Guarding Limitation

By Theorem. III.3, if a mobile sensor can guard up to two continuous perimeter segments, OSG_{2D} is hard to approximate within 1.152-optimal. Translating this into guarding elements of S_O , this means that a sensor can guard two chains of elements from S_O , where each chain contains some m elements o_1, \dots, o_m that are neighbors along ∂P . Interestingly, if each sensor may only guard a single chain of elements from S_O , we may compute an optimal cover for S_O using $O(N^2 \log N)$ time. This readily turns into a fully polynomial time approximation scheme (FPTAS) for OPG_{2D} . The algorithm operates by checking multiple times whether a given radius r is sufficient for k discs of the given radius to cover elements of S_O where each disc covers only a single chain of elements.

A single feasibility check is outlined in Algorithm 1. In the pseudo code, it is assumed that the indices are modulo N , e.g. $M[N + 1] = M[1]$, $o_{N+1} = o_1$. Algorithm 1 is based on an efficient implementation of a subroutine `MIN_ENCLOSE_DISC` (from e.g., [29], [30]) that computes the disc with minimum radius to enclose a given set of points in expected linear time. With this, a sliding window can be applied to find the rightmost end for each $1 \leq i \leq N$ such that o_i, \dots, o_{end} can be enclosed in a circle of radius r . The length of this sequence is stored in $M[i]$.

As o_{end} cannot come around and meet o_i , the total call to `MIN_ENCLOSE_DISC` is no more than $2N$. After this, the algorithm simply tries to put discs from each o_i to cover as many centers as possible to see whether S_O can be enclosed with k discs. An optimization can be made by only examining starting point as $o_1, \dots, o_{M[1]+1}$, since there is no circle of radius of r that can cover them together by the definition of M . The apparent complexity of Algorithm 1 is $O(N^2)$. Since there are a total of N points and k robots, in

a majority of cases a circle would enclose about N/k points, which effectively lowers the time complexity to $O(N^2/k)$.

Algorithm 1: OPG_2D_CONT_FEASIBLE

Data: $S_O = \{o_1, \dots, o_N\}$, sample points in circular order
 k , the number of robots
 r , the candidate sensing radius
Result: **true** or **false**, indicating whether S_O can be covered with k discs with radius r

```

1 if MIN_ENCLOSE_DISC( $o_1, \dots, o_N$ )  $\leq r$  then
2   | return true
3 end

%Phase 1: find the maximum number of
consecutive points a disc of radius  $r$ 
can enclose from each  $c_i$ .
4  $M \leftarrow$  an array of length  $N$ ;  $end \leftarrow 1$ ;
5 for  $i = 1$  to  $N$  do
6   | while MIN_ENCLOSE_DISC( $o_i, \dots, o_{end+1}$ )  $\leq r$  do
7     |    $end \leftarrow end + 1$ ;
8     | end
9     |  $M[i] \leftarrow end - i + 1$ ;
10 end

%Phase 2: try to tile from each  $o_i$ .
11 for  $i = 1$  to  $N$  do
12   |  $j \leftarrow i$ ,  $cnt \leftarrow k$ ;
13   | while  $cnt > 0$  do
14     |    $j \leftarrow j + M[j]$ ;
15     |   if  $j - i \geq N$  then
16     |     | return true
17     |   end
18     |    $cnt \leftarrow cnt - 1$ 
19   | end
20 end
21 return false

```

Note that for the optimal coverage radius r^* , it holds that $r_{min} = 0 < r^* \leq \text{len}(\partial P)/(2k) = r_{max}$. Recall that $N = \lceil \text{len}(\partial P)/(2\varepsilon) \rceil$. Hence, after at most

$$\log \frac{r_{max} - r_{min}}{\varepsilon} = \log \left(\frac{\text{len}(\partial P)}{2k\varepsilon} \right) = O(\log \frac{N}{k})$$

times of binary search on the optimal radius r^* by calling OPG_2D_CONT_FEASIBLE, the search range of r^* or the gap between r_{max} and r_{min} will be reduced to within ε . So, it takes expected $O(N^2 \log(N/k))$ time in total to get an approximate solution with radius at most ε more than $\text{size}(k, S_O)$ or $\text{size}(k, \partial P)$.

Theorem IV.1. *Under the rule of continuous coverage, OPG_{2D} for a simple polygon can be approximated to $(1 + \varepsilon)$ -optimal in expected $O(N^2 \log N)$ time, and $O((N^2/k) \log(N/k))$ in most cases, where $N = \lceil \text{len}(\partial P)/(2\varepsilon) \rceil$.*

Remark. In the running time complexity analysis, we implicitly used the assumption that $\text{len}(\partial P)$ is polynomial to problem input size (see Section II). Also, the algorithm given above computes an $OPT + \varepsilon$ optimal solution. However, it can be naturally assumed that the optimal sensing radius OPT is lower bounded in realistic scenarios. So, an $(OPT +$

$\varepsilon)$ solution directly translates into a $(1 + \varepsilon)$ -optimal solution. Lastly, using techniques similar to those from [1], [11], we mention that results in this subsection readily extends to multiple simple polygons with gaps along the boundary. These arguments continue to apply throughout the rest of this section.

Regarding the choice in implementation, the minimum enclosing disc problem (1-center problem) also has deterministic solution [31] in linear time, but a randomized algorithm is considered to be more efficient [29] and easier to implement.

B. $(2 + \varepsilon)$ Approximation

In dealing with Euclidean k -clustering problems, two seminal methods are often brought out, both of which compute 2-approximation solutions for k -center problem in polynomial time. This is fairly close to the inapproximability gap of 1.822 for Euclidean k -center problem [15]. The first [16], [32] transforms the clustering problem to a dominating set problem and then applies parametric search on the cluster size (radius), resulting in a 2-approximation in time $O(n^2 \log n)$ with n being the number of points to cover. A second method [17] takes a simpler farthest clustering approach by iteratively choosing the furthest point from the current centers as the new center. The method runs in $O(nk)$ but is subsequently improved to $O(n \log k)$ in [15]. So, by applying either of them on S_O , we have

Proposition IV.1. OSG_{2D} can be approximated to $(2 + \varepsilon)$ -optimal in polynomial time with $N = O(\text{len}(\partial P)/\varepsilon)$ samples for perimeter guarding and $N = O((\text{len}(\partial P)/\varepsilon)^2)$ samples for region guarding.

For evaluation, we implemented the farthest clustering approach [17].

C. Grid and Integer Programming-based Algorithm

Approximation using grids [14] often exhibits good optimality guarantees and bounded time complexity. Seeing that and knowing that OSG_{2D} is hard in general, we attempted grid-based integer linear programming (ILP) methods for solving OSG_{2D} with good success. Our ILP model construction is done as follows.

Consider bounding the polygon P of interest by an $m \times n$ square grid where each cell is $\varepsilon \times \varepsilon$, and denote g_{ij} as the center of the cell at row i and column j . If we limit the possible locations of each robot to the center of some grid cell, the optimal radius with this limitation will only be at most $\sqrt{2}\varepsilon/2$ away from $\text{size}(k, S_O)$. This could be seen by moving the robot locations in the optimal deployment to their nearest grid centers respectively and applying triangle inequality.

So, given a candidate radius r , to check the feasibility of whether ∂P can be covered by k circles of radius r , we adapt an approach for solving the k -center problem [18] with integer linear programming. Specifically, we create $m \times n$ boolean variables y_{ij} , $1 \leq i \leq m$, $1 \leq j \leq n$, indicating

whether there is a robot at g_{ij} , then start to check the feasibility of following integer programming model.

$$\sum_{1 \leq i \leq m} \sum_{1 \leq j \leq n} y_{ij} \leq k \quad (1)$$

$$\sum_{i, j \text{ s.t. } \|g_{ij} - o_\ell\|_2 \leq r} y_{ij} \geq 1 \text{ for each } 1 \leq \ell \leq N \quad (2)$$

$$y_{ij} \in \{0, 1\} \quad 1 \leq i \leq m, 1 \leq j \leq n \quad (3)$$

The first constraint says the number of locations is no more than k , and the second ensures each o_ℓ can be covered by at least one circle with radius r illustrated in Fig. 11.

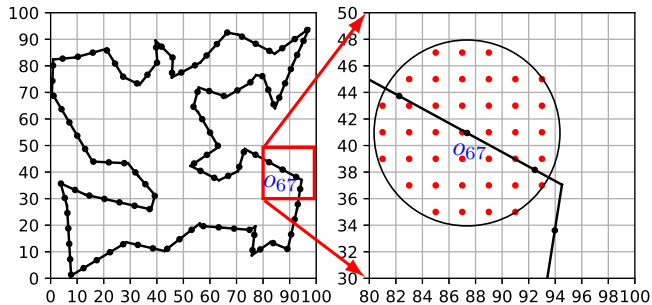


Fig. 11: This perimeter guarding example illustrates constraint (2) for o_{67} with $r = 7$. The black dots are the sampled $S_O = \{o_1, \dots, o_{100}\}$. In order to cover o_{67} , at least one among the red color grid cell centers need to be selected as robot location.

When the ILP model has a feasible solution, $r^* = \text{size}(k, S_O) \leq r$ and $r \leq r^* = \text{size}(k, S_O)$ otherwise. This means that we can do a binary search on r^* , from an initial range of $r^* = \text{size}(k, S_O)$: $r_{min}^i = 0 < r^* \leq \text{len}(\partial P)/(2k) = r_{max}^i$, until finally $r_{max}^f - r_{min}^f$ is reduced to the selected granularity of ε .

Remark. With minor modifications, the ILP model applies to 2D region guarding, where the number of constraint (2) will then be $O(mn)$ with one for each grid that intersects with the polygon in an $m \times n$ grid. The initial upper bound set as r^* be $\text{len}(\partial P)$ and lower bound set as $\sqrt{\text{area}(P)/(k\pi)}$. It is also possible to apply the $(2 + \varepsilon)$ -approximation algorithm and set the result as the initial upper bound with the half of it as the initial lower bound.

V. EVALUATION AND APPLICATION SCENARIOS

For the three algorithms described in Section IV, we developed implementations in C++ and evaluated them on an Intel Core i7 PC with a boost clock of 4.2GHz and 16GB RAM. For solving ILP models, Gurobi solver [33] is used. To evaluate the algorithms, we first generate a set of performance benchmarks obtained by subjecting these algorithms through a large set of benchmark cases. Following the synthetic benchmarks, we applied the algorithms on two potential application scenarios: guarding the outer perimeter of the Warwick Castle and monitoring a building for potential fire eruption points.

A. Performance Benchmarks

For creating synthetic benchmarks, to generate the test set \mathcal{W} , we created simple polygons with the number of vertices ranging between 10 and 200. For each instance of the tested polygon, vertices are picked uniform at random from $[0, 1] \times [0, 1]$ and the TSP tour among these vertices are used for generating a simple polygon of a reasonable shape. An example is given in Fig. 11.

We first evaluate the computational performance of the special OPG_{2D} algorithm where each sensor may cover a single continuous perimeter segment; denote this algorithm as AL_OPG_2D_CONT. Table I lists the running time in seconds for various N (number of discretized samples) and k (number of guards). Various values of N suggest the choices of ε according to the setup of N in Section IV, in this case $N = \lceil \text{len}(\partial P)/(2\varepsilon) \rceil$. Each data point is an average of 100 examples. As we can observe, the method has very good scalability. It also demonstrates the behavior that running time is inverse proportional to the number of guards, conforming with the statement about time complexity in Section IV-A. The normalized average standard deviation is about 0.06, which is pretty small.

$N \backslash k$	5	10	20	30	50	100
500	0.097	0.044	0.019	0.013	0.007	0.004
800	0.257	0.118	0.054	0.036	0.019	0.011
1000	0.385	0.183	0.082	0.055	0.029	0.016
1500	0.912	0.436	0.203	0.120	0.073	0.039
2000	1.597	0.743	0.345	0.225	0.123	0.062

TABLE I: Running time (seconds) for AL_OPG_2D_CONT.

Since the $(2 + \varepsilon)$ -optimal algorithm is extremely efficient, we do not report its running time. For the ILP methods, Table II and Table III provide the running times for solving OPG_{2D} and ORG_{2D}, respectively (for conve-

$GS \backslash k$	10	15	20	30	50	100
50×50	0.219	0.127	0.092	0.051	0.023	0.009
100×100	0.686	0.383	0.250	0.141	0.089	0.033
200×200	1.915	1.132	0.792	0.444	0.281	0.115
300×300	7.782	4.201	2.613	1.513	0.814	0.435
400×400	21.23	11.63	7.275	3.827	2.231	1.318

TABLE II: Running time (seconds) for AL_OPG_2D_ILP.

nience, denote these two methods as AL_OPG_2D_ILP and AL_ORG_2D_ILP). Each data point is an average over 10 cases. GS denotes the discrete grid size, suggesting the choice of the grid granularity ε and the single grid cell size $\varepsilon \times \varepsilon$. We observe that the ILP method is highly effective for solving OPG_{2D} and fairly good for solving ORG_{2D}. The normalized average standard deviation is about 0.125 for AL_OPG_2D_ILP (which is reasonable) and 0.545 for AL_ORG_2D_ILP (which is relatively large).

For solution quality, we compare AL_OPG_2D_CONT, AL_OPG_2D_ILP, and AL_ORG_2D_ILP with the $(2 + \varepsilon)$ -optimal solution. For example, given a test case, let the

$GS \backslash k$	10	15	20	30	50	100
20×20	0.252	0.245	0.200	0.170	0.136	0.094
30×30	1.413	1.064	0.886	0.799	0.858	0.576
40×40	5.048	3.598	3.055	2.252	6.114	1.156
50×50	7.003	5.617	4.984	5.836	10.91	0.925
80×80	87.14	84.18	82.09	423.5	>2e3	>2e3

TABLE III: Running time (seconds) for AL_ORG_2D_ILP.

resulting radius for AL_OPG_2D_CONT be r_1 and that for the $(2 + \varepsilon)$ -optimal algorithm be r_2 , we compute the optimality gain as the reduce of coverage radius over r_2 in percentage, that is $(r_2 - r_1)/r_2 \cdot 100$. These are then averaged over 10 cases. Selected representative results (only three out of a total of 18 rows) are given in Table IV. In the table, m denotes the method where **1** = AL_OPG_2D_CONT, **2** = AL_OPG_2D_ILP, and **3** = AL_ORG_2D_ILP. Number of samples for AL_OPG_2D_CONT is set to 2000. Grid size for AL_OPG_2D_ILP is 200×200 . Grid size for AL_ORG_2D_ILP is set to 40×40 . For each method, we used polygons with 200 vertices. We observe that these algorithms do significantly better than 2-optimal with AL_OPG_2D_ILP getting very close to being 1-optimal (whose optimality gain is no more than around 50).

$m \backslash k$	5	10	20	30	50	100
1	22.34	23.89	27.07	29.14	32.32	34.18
2	36.29	34.82	36.22	36.98	37.69	38.29
3	35.69	32.58	30.06	25.22	21.99	15.46

TABLE IV: Optimality gain of AL_OPG_2D_CONT, AL_OPG_2D_ILP, and AL_ORG_2D_ILP over the $(2 + \varepsilon)$ -optimal method.

B. Two Application Scenarios

Next, we demonstrate the solutions computed by our algorithms on two potential application scenarios. For the first one, we apply algorithms for OPG_{2D} on the outer boundary of the Warwick Castle in England (data retrieved from openstreetmap.org [34]). Fig. 12 shows the solution for 15 guards computed by the $(2 + \varepsilon)$ -optimal algorithm, AL_OPG_2D_CONT, and AL_OPG_2D_ILP, respectively. Both AL_OPG_2D_CONT and AL_OPG_2D_ILP do about 40% better when compared with the $(2 + \varepsilon)$ -optimal algorithm. AL_OPG_2D_CONT does 3% better than AL_OPG_2D_ILP since the perimeter is suitable for continuous guarding while the ILP method is slightly limited by the chosen resolution.

In a second application, we took the footprint of the Brazil National Museum and use 40 mobile robots to monitor it. The solution, shown in Fig. 13, is computed using AL_ORG_2D_ILP. This could be useful when a building is on fire and drones equipped with heat sensors can monitoring “hot spots” on top of the building to prioritize fire extinguishing effort. There are also many other similar application scenarios.

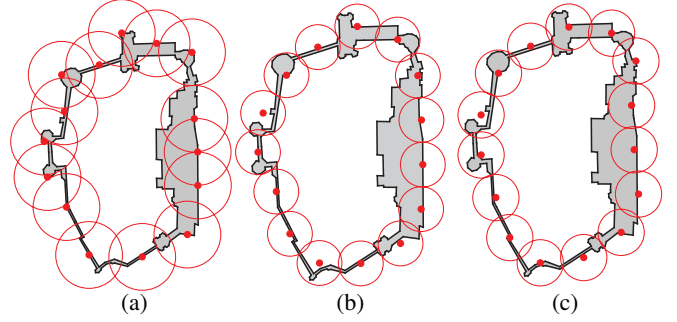


Fig. 12: Solutions for deploying 15 mobile sensors to guard the perimeter of the Warwick Castle. Methods: (a) $(2 + \varepsilon)$ -optimal. (b) AL_OPG_2D_CONT. (c) AL_OPG_2D_ILP.




Fig. 13: A near-optimal solution for deploying 40 mobile robots for monitoring the Brazil National Museum, which caught fire in 2019.

VI. CONCLUSION AND DISCUSSIONS

In this study, we examine OSG_{2D}, the problem of directly computing a deployment strategy for covering 1D or 2D critical sets using many mobile sensors while minimizing the maximum sensing radius. After showing that OSG_{2D} is computationally intractable to even approximate within 1.152, we describe several algorithmic solutions with optimality and/or computation time guarantees. Subsequent thorough evaluation demonstrates the effectiveness of these algorithmic solutions. Finally, we demonstrate the utility of our algorithmic solutions with two application scenarios. Due to space limit, guarding perimeters with gaps (see, e.g., [1]) is not discussed in this work. However, because our algorithms work with a grid-based discretization, the results directly apply to arbitrary bounded 1D and 2D sets.

Many intriguing questions follow; we mention two here concerning the sensing capabilities. First, OSG_{2D} works with circular regions which is perhaps the simplest one due to symmetry. What if the sensor region is not circular? Whereas such cases appear to be hard [35], effective scalable solutions may still be possible. Secondly, currently we assume that all parts of the critical set to be guarded have equal importance. What if certain subsets are more important?

REFERENCES

- [1] S. W. Feng, S. D. Han, K. Gao, and J. Yu, "Efficient algorithms for optimal perimeter guarding," in *Robotics: Sciences and Systems*, 2019.
- [2] M. A. Batalin and G. S. Sukhatme, "Spreading out: A local approach to multi-robot coverage," in *Distributed Autonomous Robotic Systems 5*. Springer, 2002, pp. 373–382.
- [3] J. Cortés, S. Martínez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Transactions on Robotics & Automation*, vol. 20, no. 2, pp. 243–255, 2004.
- [4] N. Correll, J. Bachrach, D. Vickery, and D. Rus, "Ad-hoc wireless network coverage with networked robots that cannot localize," in *Proceedings IEEE International Conference on Robotics and Automation*. IEEE, 2009, pp. 3878–3885.
- [5] S. Gil, D. Feldman, and D. Rus, "Communication coverage for independently moving robots," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 4865–4872.
- [6] S. Martínez, J. Cortés, and F. Bullo, "Motion coordination with distributed information," *IEEE Control Systems Magazine*, vol. 27, no. 4, pp. 75–88, 2007.
- [7] M. Schwager, B. J. Julian, and D. Rus, "Optimal coverage for multiple hovering robots with downward facing cameras," in *Proceedings IEEE International Conference on Robotics and Automation*. IEEE, 2009, pp. 3515–3522.
- [8] M. Pavone, A. Arsie, E. Frazzoli, and F. Bullo, "Equitable partitioning policies for robotic networks," in *2009 IEEE International Conference on Robotics and Automation*. IEEE, 2009, pp. 2356–2361.
- [9] M. Schwager, D. Rus, and J.-J. Slotine, "Decentralized, adaptive coverage control for networked robots," *The International Journal of Robotics Research*, vol. 28, no. 3, pp. 357–375, 2009.
- [10] A. Pierson, L. C. Figueiredo, L. C. Pimenta, and M. Schwager, "Adapting to sensing and actuation variations in multi-robot coverage," *The International Journal of Robotics Research*, vol. 36, no. 3, pp. 337–354, 2017.
- [11] S. W. Feng and J. Yu, "Optimal perimeter guarding with heterogeneous robot teams: complexity analysis and effective algorithms," *IEEE Robotics and Automation Letters*, 2020.
- [12] A. Weber, *Theory of the Location of Industries*. University of Chicago Press, 1929.
- [13] Z. Drezner, *Facility Location: A Survey of Applications and Methods*. Springer Verlag, 1995.
- [14] S. Har-Peled, *Geometric Approximation Algorithms*. American Mathematical Soc., 2011, no. 173.
- [15] T. Feder and D. Greene, "Optimal algorithms for approximate clustering," in *Proceedings ACM Symposium on Theory of Computing*. ACM, 1988, pp. 434–444.
- [16] D. S. Hochbaum and D. B. Shmoys, "A best possible heuristic for the k-center problem," *Mathematics of Operations Research*, vol. 10, no. 2, pp. 180–184, 1985.
- [17] T. F. Gonzalez, "Clustering to minimize the maximum intercluster distance," *Theoretical Computer Science*, vol. 38, pp. 293–306, 1985.
- [18] M. S. Daskin, "A new approach to solving the vertex p-center problem to optimality: Algorithm and computational results," *Communications of the Operations Research Society of Japan*, vol. 45, no. 9, pp. 428–436, 2000.
- [19] M. I. Shamos and D. Hoey, "Closest-point problems," in *16th Annual Symposium on Foundations of Computer Science (sfcs 1975)*. IEEE, 1975, pp. 151–162.
- [20] T. C. Hales, "A proof of the kepler conjecture," *Annals of Mathematics*, vol. 162, no. 3, pp. 1065–1185, 2005.
- [21] A. Thue, *Über die dichteste Zusammenstellung von kongruenten Kreisen in einer Ebene*. na, 1910.
- [22] J. O'Rourke, *Art Gallery Theorems and Algorithms*. Oxford University Press, 1987.
- [23] T. C. Shermer, "Recent results in art galleries (geometry)," *Proceedings of the IEEE*, vol. 80, no. 9, pp. 1384–1399, 1992.
- [24] M. R. Garey and D. S. Johnson, "The rectilinear steiner tree problem is np-complete," *SIAM Journal on Applied Mathematics*, vol. 32, no. 4, pp. 826–834, 1977.
- [25] B. Mohar, "Face covers and the genus problem for apex graphs," *Journal of Combinatorial Theory, Series B*, vol. 82, no. 1, pp. 102–117, 2001.
- [26] S. W. Feng and J. Yu, "Optimally guarding perimeters and regions with mobile range sensors," *arXiv preprint arXiv:2002.08477*, 2020.
- [27] J. Petersen, "Die theorie der regulären graphs," *Acta Math.*, vol. 15, pp. 193–220, 1891. [Online]. Available: <https://doi.org/10.1007/BF02392606>
- [28] J. Edmonds, "Paths, trees, and flowers," *Canadian Journal of Mathematics*, vol. 17, p. 449–467, 1965.
- [29] E. Welzl, "Smallest enclosing disks (balls and ellipsoids)," in *New Results and New Trends in Computer Science*. Springer, 1991, pp. 359–370.
- [30] M. Berg, de, O. Cheong, M. Kreveld, van, and M. Overmars, *Computational Geometry: Algorithms and Applications*. Springer, 2008.
- [31] N. Megiddo, "Linear-time algorithms for linear programming in r^3 and related problems," *SIAM Journal on Computing*, vol. 12, no. 4, pp. 759–776, 1983.
- [32] V. V. Vazirani, *Approximation Algorithms*. Springer Science & Business Media, 2013.
- [33] G. Optimization, "Gurobi optimizer 9.0," *Gurobi: <http://www.gurobi.com>*, 2019.
- [34] M. Haklay and P. Weber, "Openstreetmap: User-generated street maps," *IEEE Pervasive Computing*, vol. 7, no. 4, pp. 12–18, 2008.
- [35] J. C. Culberson and R. A. Reckhow, "Covering polygons is hard," in *29th Annual Symposium on Foundations of Computer Science*. IEEE, 1988, pp. 601–611.