Duality for Geometric Set Cover and Geometric Hitting Set Problems on Pseudodisks*

Stephane Durocher† Robert Fraser†

Abstract

Given an instance of a geometric set cover problem on a set of points X and a set of objects $\mathcal{R}$, the dual is a geometric hitting set problem on a set of points P and a set of objects $\mathcal{Q}$, where there exists a one-to-one mapping from each $x_j \in X$ to a dual object $Q_j \in \mathcal{Q}$ and for each $R_i \in \mathcal{R}$ to a dual point in $p_i \in P$, so that a dual point p_i is contained in a dual object Q_j if and only if the corresponding primal point x_j is covered by the object R_i. In this work, we explore the setting of geometric duality for geometric set cover problems on pseudodisks. We first show that there does not always exist a geometric dual on pseudodisks. We initiate the search for a characterization of the class of objects that may be dualized by identifying a sufficient (but not necessary) property for a dual to exist on distinct pseudodisks, called the *pair-cover and crossing-quad free* property. We show that such problems may be dualized into hitting set instances on pseudodisks by building a planar support for the dual instance, and then constructing an orthogonal drawing of the support which we transform into a dual set of pseudodisks. A corollary of these results is a PTAS for dualizable set cover problems using the PTAS for hitting set on pseudodisks.

1 Introduction

Geometric duality is a beautiful and useful tool for computational geometers, as some problems are conceptually simpler to solve in the dual setting. The classic example is point-line duality in the plane, see e.g. [8, §8.2]. Duality has been the catalyst for breakthroughs such as the first optimal algorithm for the half-plane range query problem [6]. Our interest lies in geometric set cover and hitting set problems, motivated by the distinction that there exists a PTAS for the hitting set problem on pseudodisks [14],[1] while none is known for the set cover problem on pseudodisks. Therefore, we endeavoured to prove that any set cover problem on pseudodisks could be dualized to a hitting set problem on pseudodisks in polynomial time (and vice versa), which would permit the use of the PTAS to obtain an approximate solution.

Unfortunately, it turns out that the dualization we desired is not always possible, and we open our study by proving this fact. We show that dualization is straightforward for problems where all regions in the instance are translations of a geometric object. We proceed by studying dualization on pseudodisks having the *pair-cover and crossing-quad free* property, and we prove that such objects may always be dualized.

1.1 Definitions and Nomenclature

We begin by reviewing concepts required for our discussion.

Definition 1 Pseudodisk: *A pseudodisk is a region of the plane bounded by a closed Jordan curve, with the restriction that the boundaries of any two pseudodisks in a given instance may intersect only transversely and at most twice.*

Definition 1 could be generalized to allow pseudodisks to intersect exactly once (tangentially at a single point), but we ignore this detail for clarity of exposition.

Definition 2 Geometric Set Cover Problem: *Given a set of points X and a set of geometric objects $\mathcal{R}$, the* geometric set cover problem *is to find a subset $\mathcal{R}^\star \subseteq \mathcal{R}$ of minimum cardinality so that all elements of X are covered by $\mathcal{R}^\star$, i.e. $X \subseteq \cup_{R \in \mathcal{R}^\star} R \cap X$.*

Definition 3 Geometric Hitting Set Problem: *Given a set of points P and a set of geometric objects $\mathcal{Q}$, the* geometric hitting set problem *is to find a subset $P^\star \subseteq P$ of minimum cardinality so that all objects in $\mathcal{Q}$ contain at least one element of $P^\star$, i.e. $\forall Q \in \mathcal{Q}, Q \cap P^\star \neq \emptyset$.*

Definition 4 Geometric Dual: *Given an instance of the set cover problem $\mathcal{S} = (X, \mathcal{R})$ (the* primal *setting), an instance of the hitting set problem $\mathcal{H} = (P, \mathcal{Q})$ is a* geometric dual *of $\mathcal{S}$ (the* dual *setting) if there are bijections between X and $\mathcal{Q}$ as well as $\mathcal{R}$ and P and any point $p_i \in P$ is contained in an object $Q_j \in \mathcal{Q}$ if and only if the corresponding point $x_j \in X$ is covered by the object $R_i \in \mathcal{R}$ in the primal setting. An optimal solution $P^\star$ for the dual setting corresponds exactly to*

*This work was supported in part by the Natural Sciences and Engineering Council of Canada (NSERC).

†Department of Computer Science, University of Manitoba, Winnipeg, Canada {durocher,fraser}@cs.umanitoba.ca

[1] A set of pseudodisks is equivalent to a set of the 2-admissible regions used in [14].

an optimal solution $\mathcal{R}^\star$ for the primal setting. A set cover instance dualizing a hitting set instance is defined analogously.

In this discussion, we usually omit the prefix "geometric" from the concepts defined here. We consistently use X and $\mathcal{R}$ for set cover and P and $\mathcal{Q}$ for hitting set to differentiate the two settings. We write ∂R_i to denote the boundary of a pseudodisk R_i. We assume that all points and pseudodisks are distinct.

1.2 Related Work

The conventional application of duality is to points and lines (or pseudolines, which are curves that intersect pairwise at most once, and such intersections must not be tangential) [2, 10]. A set of pseudolines can define the boundaries of a set of pseudo-halfplanes so that the setting becomes a set cover or hitting set problem, and the preservation of the "above-below" property in the dual means that duality exists and is well defined for our purposes. Our work seeks to extend these results to pseudodisks.

A corollary of our result is the derivation of a PTAS for some set cover problems on pseudodisks by showing that they may be dualized to hitting set problems on pseudodisks, allowing the hitting set PTAS of Mustafa and Ray [14] to be applied. This is an active area of research; a QPTAS for broad classes of set cover problems was recently found [13]. Our algorithm makes use of the arrangement of the boundaries of a set of pseudodisks (this set of boundaries is also known as a set of pseudo-circles); the combinatorial properties of such arrangements are well studied [1].

A *hypergraph* is a generalization of a standard graph, where a *hyperedge* may contain any number of points. Therefore, a hypergraph $H = (V, F)$ may be used as an abstract representation of a set cover problem $\mathcal{S} = (X, \mathcal{R})$ (or hitting set problem) by creating a vertex v_j in V for each point x_j in X, and mapping each edge $f_i \in F$ to an object $R_i \in \mathcal{R}$ so that the vertices in $f_i \cap V$ correspond exactly to the points in $R_i \cap X$. A *planar support* of a hypergraph $H = (V, F)$ is a planar graph $G = (V, E)$ where the subgraph of G induced by any hyperedge $f_i \in F$ is connected. Much of the research with respect to planar supports has been to determine whether a planar support exists for a given hypergraph [5, 12]. Our use of planar supports is reversed, since finding the hypergraph corresponding to the hitting set instance dualizing a given set cover instance is straightforward. We build a planar support for the hypergraph, and the support is used to create the dual hitting set instance on pseudodisks. To our knowledge, supports have not been used in a similar manner before.

2 A Counterexample for Duality on Pseudodisks

Theorem 1 *There exists a family of set cover problems on pseudodisks for which there is no dual hitting set formulation on pseudodisks (and equivalently, such hitting set problems on pseudodisks cannot be dualized into set cover instances on pseudodisks).*

See Appendix A for the proof of Theorem 1.

3 Geometric Dual of Set Cover

Our aim is to identify classes of set cover problems on pseudodisks that may always be dualized. We begin with a more general result that is fairly trivial to establish.

Theorem 2 *Problems defined on translates of any single object can always be dualized.*

See Appendix A for the proof of Theorem 2.

For the remainder of the paper, we outline a method for reducing an instance of another class of geometric set cover problems on pseudodisks to an instance of a geometric hitting set problem on pseudodisks (or vice versa, of course).

Definition 5 Pair-Cover Free Property: *The* Pair-Cover Free *(PF) property holds for a set of geometric objects $\mathcal{R}$ if $R_i \not\subseteq R_j \cup R_k$ for all $R_i, R_j, R_k \in \mathcal{R}$.*

Definition 6 Crossing-Quad Free Property: *If $R_k \cap R_\ell \not\subseteq R_i \cup R_j$ for any four pseudodisks $\{R_i, R_j, R_k, R_\ell\} \subseteq \mathcal{R}$ where $R_i \cap R_j \subseteq R_k \cup R_\ell$, then the* Crossing-Quad Free *(CF) property holds.*[2]

We define the *pair-cover and crossing-quad free set cover* (PCF-SC) problem on pseudodisks as the set cover problem $\mathcal{S} = (X, \mathcal{R})$ where $\mathcal{R}$ is a set of pseudodisks with the PF and CF properties.

Theorem 3 *Any instance $\mathcal{S} = \{X, \mathcal{R}\}$ of PCF-SC may be reduced to an instance of a hitting set problem $\mathcal{H} = \{P, \mathcal{Q}\}$ in polynomial time, where P is a set of points, $\mathcal{Q}$ is a set of pseudodisks (both in $\mathbb{R}^2$), and $\mathcal{H}$ is a geometric dual of $\mathcal{S}$.*

The reduction progresses in two stages. First, the set cover instance is converted to a special graph known as a *planar support* $G = (V, E)$, where each vertex $v_i \in V$ corresponds to a pseudodisk $R_i \in \mathcal{R}$ and each point $x_j \in X$ maps to a connected induced subgraph S_j of G. Finally, we show how to fatten each of the subgraphs in the plane to form a pseudodisk, creating an instance of the hitting set problem on pseudodisks that is a geometric dual of the original set cover instance. We give an overview of the proof below, which is then proved formally in the remainder of Section 3:

[2]This property is only used in the proof of Lemma 9.

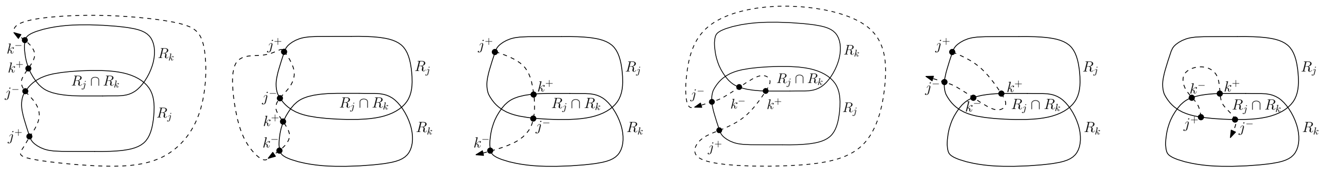

(a) Case 1a. R_i covers all of $R_j \cap R_k$. (b) Case 1b. R_i covers none of $R_j \cap R_k$. (c) Case 2 (and Case 5). R_i covers one point of $\partial R_j \cap \partial R_k$. (d) Case 3 violates the PF property (R_i encloses $\partial R_j \cap \partial R_k$). (e) Case 3 (and Case 4). R_i covers neither point in $\partial R_j \cap \partial R_k$. (f) Case 6. This case is forbidden by the PF property.

Figure 1: Cases illustrating the different ways in which one pseudodisk may intersect two others. The figures are oriented so the walk on R_i is clockwise.

1. Convert the set cover instance to a planar support:
   (a) Build the support by iterating over the regions of the plane in the arrangement of the pseudodisks in order of increasing depths (Algorithm 1).
   (b) Show that the subgraph of the support imposed by any point is connected (Lemma 7).
   (c) Show that the support is planar (Lemma 8).
2. Convert the planar support to a hitting set instance:
   (a) Show that the support may be embedded orthogonally (Theorem 10).
   (b) Show that each edge may be fattened to define a bounded region of the plane, and that these objects may be arranged and manipulated to become pseudodisks (Lemma 11).

3.1 Properties of Pair-Cover Free Set Cover

Various arguments in this discussion consider the possible ways in which three pseudodisks may interact, and so we enumerate all possible cases. Consider two pseudodisks R_j and R_k that each intersect a third pseudodisk R_i. Let j^+, j^-, k^+, k^- denote the set of events that occur during a clockwise walk around ∂R_i, where j^+ indicates the point of entry into R_j and j^- indicates the point of exit. We may arbitrarily begin our walk at j^+, and so there are 3! possible walks (see Figure 1), which we divide into cases. Note that for Cases 1–3, we begin outside R_k, while for Cases 4–6 we begin inside.

Case 1. j^+, j^-, k^+, k^-: A pseudodisk R_i may either cover $R_j \cap R_k$ completely or not at all depending on how the path is closed to create a pseudodisk. Call these Cases 1a (Figure 1a) and 1b (Figure 1b), respectively.

Case 2. j^+, k^+, j^-, k^-: R_i covers exactly one point in $\partial R_j \cap \partial R_k$, as shown in Figure 1c.

Case 3. j^+, k^+, k^-, j^-: R_i covers either both points or neither point in $\partial R_j \cap \partial R_k$. However, covering both points entails violating the PF property, since $R_k \subset R_i \cup R_j$ in this scenario (see Figure 1d). Therefore, the only valid scenario for Case 3 is that where R_i covers neither point in $\partial R_j \cap \partial R_k$, shown in Figure 1e.

Case 4. j^+, j^-, k^-, k^+: This is symmetric with Case 3 (swap the labels j and k).

Case 5. j^+, k^-, j^-, k^+: This is symmetric with Case 2.

Case 6. j^+, k^-, k^+, j^-: The beginning of the tour is in $R_j \cap R_k$. The next event is k^-, so the path is now in $R_j \setminus R_k$, and the following event is k^+, so both intersection points of $\partial R_i \cap \partial R_k$ are in R_j. The last event j^- is in R_k, and so both points of $\partial R_i \cap \partial R_j$ are in R_k, implying that $R_i \subseteq R_j \cup R_k$ which violates the PF property.

Therefore, the distinct cases are 1a, 1b, 2, and 3. Lemma 4 is an immediate consequence of Case 2:

Lemma 4 *Given any instance of PCF-SC, if a pseudodisk R_i intersects the boundaries of two other pseudodisks (say ∂R_j and ∂R_k) in the closed region $R_j \cap R_k$, then exactly one point of $\partial R_j \cap \partial R_k$ is covered by R_i.*

3.2 Building the Support

To build the support we describe how to construct an *adjacency list* that is a supergraph of the support and a subgraph of the intersection graph of $\mathcal{R}$; as we explain later, the support may be derived from the adjacency list. The support G has the property that if $(v_i, v_j) \in G$ then $R_i \cap R_j \neq \emptyset$, but the reverse is not necessarily true. Rule 1 holds for the adjacency list and the support G:

Rule 1 *For any three pseudodisks $\{R_i, R_j, R_k\}$, if $R_i \cap R_j \subseteq R_k$ then there is no edge (v_i, v_j) in the support G.*

The adjacency list is stored as a map from each vertex v_i to the set of neighbouring vertices, where a neighbour vertex corresponds to a pseudodisk that intersects R_i in the primal without violating Rule 1. A non-empty intersection between pseudodisks R_1 and R_2 does not necessarily imply the existence of the edge (v_1, v_2) in the adjacency list; it does, however, imply the following lemma:

Lemma 5 *For every $i \geq 1$ and $R' = \{R_{a_1}, \ldots, R_{a_i}\} \subseteq \mathcal{R}$, where R' is the set of all pseudodisks covering some cell of the arrangement, then G' is connected, where G' is the subgraph of the adjacency list induced by the vertices $v_{a_1}, v_{a_2}, \ldots, v_{a_i}$ associated with R'.*

Proof. Suppose otherwise. That is, there exists $R' = \{R_{a_1}, \ldots, R_{a_i}\} \subseteq \mathcal{R}$ such that $\bigcap_{1 \le j \le i} R_{a_j} \neq \emptyset$ and the corresponding induced subgraph of the adjacency list is disconnected. Suppose R_{a_a} and R_{a_b} in R' lie in separate components in the induced subgraph. Since there is no edge (v_{a_a}, v_{a_b}) in the subgraph, Rule 1 must have been applied such that $R_{a_a} \cap R_{a_b} \subseteq R_{a_c}$ for some R_{a_c}. Consequently, $R_{a_c} \in R'$. Without loss of generality, suppose R_{a_a} and R_{a_c} lie in separate components. Again, Rule 1 must have been applied such that $R_{a_a} \cap R_{a_c} \subseteq R_{a_d}$ for some $R_{a_d} \in R'$. Observe that $R_{a_a} \cap R_{a_b} \subseteq R_{a_a} \cap R_{a_c} \subseteq R_{a_d}$. Since R' is finite, this argument cannot be applied indefinitely, leading to a contradiction. □

The support is built iteratively using the *depths* of regions of the plane in the primal, where the depth is defined as the number of pseudodisks entirely covering that region of the plane (a *region* in this case is a cell in the arrangement of the set of pseudodisks), and so a region with depth k is covered by k pseudodisks. Let $\mathcal{A}(\mathcal{R})$ denote the arrangement defined by the pseudodisks in $\mathcal{R}$. Note that although some cells of the arrangement do not necessarily contain a point of X in the primal, we create a subgraph in the support for each cell in the arrangement. We show that a pseudodisk may be created in the dual for each cell, and those not needed in the dual may be discarded later. The algorithm for building the support is sketched in Algorithm 1.

Algorithm 1 BUILD-SUPPORT($\mathcal{S} = \{X, \mathcal{R}\}$)

Input: An instance of the PCF-SC problem $\mathcal{S}$.
Output: $G = (V, E)$, a planar support for the dual of $\mathcal{S}$.
Insert a vertex v_i in V for each R_i in $\mathcal{R}$.
Consider the arrangement of the plane imposed by the pseudodisks $\mathcal{R}$, call it $\mathcal{A}(\mathcal{R})$. Sort the cells of $\mathcal{A}(\mathcal{R})$ in order of increasing depth. An element $Z \in \mathcal{A}(\mathcal{R})$ is defined by a subset of $\mathcal{R}$.
For each vertex $v_i \in V$, compute the adjacency list.
For each region of depth 1, add a self-loop to the corresponding vertex in G (Figure 2).
for each region $Z \in \mathcal{A}(\mathcal{R})$ (in order of depth $= 2 \rightarrow |\mathcal{R}|$) **do**
    If the subgraph of G induced by the vertices corresponding to the set of pseudodisks $\mathcal{R}'(\subseteq \mathcal{R})$ that cover Z has two or more connected components, then iteratively add edges to join pairs of components using edges selected from the adjacency lists until the induced subgraph is connected.
end for
return G

Each region of depth 1 corresponds to a pseudodisk that uniquely covers some region of the plane in the primal (one pseudodisk may cover many such regions). Each vertex in G corresponding to such a pseudodisk is given a self-loop in E so that each cell in the arrangement corresponds to a non-empty set of edges in E.

The iterative procedure continues by considering regions of increasing depth, although edges are only added to G if the subgraph induced by the set of vertices corresponding to the region is not already connected. If the subgraph is not connected, then a pair of vertices (v_i, v_j) is selected so that v_i and v_j are in separate components of the subgraph and v_j is in the adjacency list for v_i. As shown in Lemma 5, some such pair must always exist, and so edges are added until the induced subgraph is connected. The algorithm can be made consistent by imposing a total ordering on the edges using their labels as keys, and always choosing the first edge in the ordering that connects components of the graph. The algorithm for connecting induced subgraphs operates somewhat analogously to Kruskal's minimum spanning tree algorithm, where edges already in G have zero weight, edges permitted by the adjacency lists are given unit weight, and those not permitted have infinite weight. Our approach has additional complexities, however, because we are operating on subgraphs of the support, and we must take care when adding edges to maintain the planarity of the support.

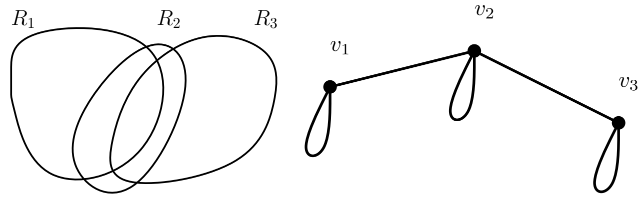

(a) (b)

Figure 2: Building the neighbourhood graph. (a) Three pseudodisks R_1, R_2, R_3, where $R_1 \cap R_3 \subseteq R_2$. (b) Each pseudodisk has a self-edge in G, and there are edges (v_1, v_2) and (v_2, v_3). The region of depth 3 in (a) is covered by R_1, R_2, R_3, but in the graph v_1, v_2, v_3 form a connected subgraph of G, so no further edges are required.

3.3 Planarity of the Support

The relative order of the edges around a vertex may be defined unambiguously by the objects in the primal. Call the pseudodisk R_2 a *neighbour* of pseudodisk R_1 if an edge between v_1 and v_2 exists in G. Let R_i be a region of the plane which is intersected by both R_j and R_k, as we examined in Figure 1. If there is an unambiguous sense that $R_j \cap \partial R_i$ is clockwise or counterclockwise of $R_k \cap \partial R_i$ w.r.t. any $R_\ell \cap \partial R_i$ on ∂R_i, then the edges (v_i, v_j) and (v_i, v_k) must have the same relative ordering w.r.t. (v_i, v_ℓ) around v_i in G, for each such R_ℓ where the edge (v_i, v_ℓ) exists. For brevity going forward, we simply refer to the relative order of R_j and R_k on ∂R_i.

The boundaries ∂R_j and ∂R_k may intersect 0, 1, or 2 times in R_i. If 0 times, then we are in either Case 1b or Case 3 (refer to §3.1). If 1b, then $R_i \cap R_j \cap R_k$ is empty and the relative order of R_j and R_k on ∂R_i is unambiguous. Case 3 does not arise, as it implies that either $R_i \cap R_j \subseteq R_k$ or $R_i \cap R_k \subseteq R_j$, which means $(v_i, v_j) \notin G$ or $(v_i, v_k) \notin G$ respectively by Rule 1. If $|\partial R_j \cap \partial R_k \cap R_i| = 1$, i.e. Case 2, then necessarily $|\partial R_i \cap \partial R_k \cap R_j| = 1$ and $|\partial R_i \cap \partial R_j \cap R_k| = 1$ as well, by Lemma 4. In this case, there is an unambiguous sense where one of R_j or R_k is clockwise of the other on R_i. If $|\partial R_j \cap \partial R_k \cap R_i| = 2$, then Case 1a applies and so $R_j \cap R_k \subseteq R_i$, which means the relative ordering of the pseudodisks R_j and R_k on R_i is again unambiguous. Therefore, relative orderings may always be consistently applied to the edges of G in the embedding.

A cycle in G corresponds to a set of pseudodisks that partitions the plane into an unbounded region and a (possibly empty) set of bounded regions.

Lemma 6 *There is a deterministic method for creating a cycle C in the embedding of G so that the clockwise ordering of the vertices in C is defined by corresponding pseudodisks in the primal.*

See Appendix A for the proof of Lemma 6.

The support now contains subgraphs corresponding to cells of the arrangement in the primal, where each vertex of the support corresponds to a pseudodisk in the primal that covers the point corresponding to the subgraph. This is immediate from the construction, since the point must exist in one of the regions of the plane used to build the graph, and a subgraph is built for each region. G adheres to the definition of a support for the dual, since each necessary subgraph is a connected induced subgraph of G, which gives the following lemma:

Lemma 7 *For any point x_j in the primal, and all vertices $v_i \in V$ in the support corresponding to pseudodisks $R_i \in \mathcal{R}$ in the primal, there exists a connected subgraph $S_j \subseteq G$ where $v_i \in S_j$ if and only if $x_j \in R_i$.*

Lemma 8 *The support G is planar.*[3]

See Appendix A for the proof of Lemma 8.

3.4 Dual Properties of the Support

The support G encapsulates some of the combinatorial structure of the dual; to complete the dual we must construct a set of pseudodisks defined by the connected subgraphs on G that correspond to the points in the primal. The subgraphs have several characteristics that allow the creation of the dual hitting set instance.

Lemma 9 *Let C denote a cycle in an induced subgraph S of G in the embedding, where S corresponds to a point x in the primal and C corresponds to pseudodisks $\mathcal{R}_C$. Any vertices on the interior of the bounded region defined by C must correspond to pseudodisks that cover x in the primal.*

See Appendix A for the proof of Lemma 9.

3.5 Building the Hitting Set Instance

We now describe how to embed the support G in the plane and transform it into the dual hitting set instance. To begin, remove subgraphs from the support that do not correspond to points in X in the primal. We construct a planar orthogonal box drawing[4] of the support G using the following result:

Theorem 10 (Biedl and Kaufmann (1997) [4]) *Given a planar triconnected graph $G = (V, E)$, a planar orthogonal box drawing of G can be drawn in $O(m + n)$ time on a $(m - n + 1) \times \min\{m - n + 1, m/2\}$ grid with $m - n$ edge bends, where $n = |V|$ and $m = |E|$.*

The drawing of the support remains planar, and while G is not necessarily triconnected, one may add dummy edges to make it triconnected and then remove the dummy edges once the drawing is computed [9]. The placement of the point set P for the dual hitting set instance is simple: place a point inside each of the vertices (boxes) of the orthogonal drawing of the support, and these will dualize the corresponding pseudodisks of $\mathcal{R}$ in the primal (and so $|P| = |V|$). Now we describe how to create the set of pseudodisks $\mathcal{Q}$ for the dual hitting set instance.

A finer grid is imposed upon the orthogonal drawing with a resolution of $1/(2m+1)$, where m is the number of pseudodisks needed in the dual, i.e., the cardinality of X in the primal. The pseudodisks that are created for the dual are orthogonal polygons with edges incident upon the lines in this finer grid. An edge e is made *k-fat* by taking the Minkowski sum of e with $[-k, k] \times [-k, k]$. We fatten parts of the edges as necessary using points on the refined grid, so that parts of the edges may be $(k/(2m + 1))$-fat, for $k \in [0, \ldots, m]$, (i.e., all edges of the drawing are less than $(1/2)$-fat).

If all edges of a subgraph are grown to be k-fat (for possibly varying values of k), the subgraph defines a polygon. Any introduced holes are removed (Lemma 9

[3] Note that if our goal was to simply derive a PTAS for this class of set cover problems, we could stop here by showing that the PTAS of [14] applies given that the support has the requisite locality property. However our primary goal is to demonstrate the existence of duality, so we proceed nonetheless.

[4] A box drawing is a graph with orthogonal non-overlapping edges, where vertices may be drawn as rectangles in order to accommodate all incident edges.

established the validity of this action, since removing a hole never causes the simplified region to cover any new vertices of G).

The planar regions are constructed iteratively. For subgraph $S_1 = (V_1, E_1)$, create a polygon so that the subgraph is $(1/(2m+1))$-fat. Now S_2 is added, and any overlap must be resolved. Overlapping edges may be resolved by making one of the edges $(2/(2m+1))$-fat in the area of overlap. Repeating this operation of inserting and fattening existing edges for each subgraph creates a hitting set instance that dualizes the set cover instance, although the resulting objects are not necessarily pseudodisks. However, the removal of extraneous intersections is always possible so that all objects are pseudodisks. See Appendix B for details.

This establishes the following lemma:

Lemma 11 *All subgraphs of the support induced by a set of vertices corresponding to a point in the primal may be enclosed with a region of the plane so that all such regions are pseudodisks.*

Finally, any cell of the arrangement that contained k points in the primal requires $k-1$ additional pseudodisks in the dual. Since we are not concerned with the PF property in the dual, we nest the missing dual pseudodisks just inside the existing dual pseudodisk so that they all have the same combinatorial structure with respect to all other points and pseudodisks in the hitting set instance. These pseudodisks form the objects $\mathcal{Q}$ for the dual hitting set instance. If the dualization may be completed in polynomial time, then Theorem 3 follows.

Theorem 12 *Dualization of an instance of PCF-SC on pseudodisks to an instance of the hitting set problem on pseudodisks can be completed in* $O(Im^5 \log m + mn)$ *time, where* $m = |\mathcal{R}| = |V| = |P|$, $n = |X| = |\mathcal{Q}|$, *and* I *denotes the time required to compute the intersection points of a pair of pseudodisks.*

See Appendix A for the proof of Theorem 12.

4 Conclusions

Our examination of the geometric duality of set cover and hitting set problems on pseudodisks has revealed positive and negative results. Perhaps surprising is the fact that not all instances are dualizable. The construction of a geometric dual is possible on translates of an object, or when we restrict instances on pseudodisks to those that have what we call the pair-cover and quad-crossing free property. A corollary of the dualization is that there exists a PTAS for set cover problem on pseudodisks with this property. Our algorithm for the construction of the dual applies interesting techniques, as we make use of graph drawing techniques to build the dual from a planar support.

There remain several open questions. Our dualization technique requires that the primal setting have the PF property, while the dual instance that is created does not necessarily have this property. It would be preferable if the dual instance also had the PF property, but we conjecture that there exists a counter-example to show that such duality does not always exist. Finally, the dualization also requires that the quad-crossing free property applies, but this property does not seem tremendously important to the dualization. We conjecture that duality is possible on instances without the CF restriction.

References

[1] P.K. Agarwal, E. Nevo, J. Pach, R. Pinchasi, M. Sharir, and S. Smorodinsky. Lenses in arrangements of pseudo-circles and their applications. *J. ACM*, 51(2):139–186, 2004.

[2] P.K. Agarwal and M. Sharir. Pseudo-line arrangements: Duality, algorithms, and applications. *SIAM J. Comp.*, 34(3):526–552, 2005.

[3] I.J. Balaban. An optimal algorithm for finding segments intersections. In *Proc. SoCG*, 211–219. ACM, 1995.

[4] T. Biedl and M. Kaufmann. Area-efficient static and incremental graph drawings. In *Proc. ESA*, vol. 1284 *LNCS*, 37–52. Springer, 1997.

[5] K. Buchin, M. van Kreveld, H. Meijer, B. Speckmann, and K. Verbeek. On planar supports for hypergraphs. *J. Graph Alg. & Appl.*, 15(4):533–549, 2011.

[6] B. Chazelle, L.J. Guibas, and D.-T. Lee. The power of geometric duality. *BIT Num. Math.*, 25(1):76–90, 1985.

[7] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms, 2nd ed.*. McGraw-Hill, 2001.

[8] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars. *Computational Geometry.* Springer, 2010.

[9] H. Fraysseix, J. Pach, and R. Pollack. How to draw a planar graph on a grid. *Combinatorica*, 10(1):41–51, 1990.

[10] J.E. Goodman. Proof of a conjecture of Burr, Grünbaum, and Sloane. *Disc. Math.*, 32(1):27–35, 1980.

[11] D. Bevan. Number of point subsets that can be covered by a disk. Math. Stack Exchange. http://math.stackexchange.com/q/75487.

[12] D.S. Johnson and H.O. Pollak. Hypergraph planarity and the complexity of drawing venn diagrams. *J. Graph Theory*, 11(3):309–325, 1987.

[13] N.H. Mustafa, R. Raman, and S. Ray. QPTAS for geometric set-cover problems via optimal separators. *CoRR*, abs/1403.0835, 2014.

[14] N.H. Mustafa and S. Ray. Improved results on geometric hitting set problems. *Disc. & Comp. Geom.*, 44(4):883–895, 2010.

[15] J. Steiner. *Gesammelte Werke*, vol. 1. Berlin, 1881.

A Appendix: Proofs

Proofs omitted from the main text due to space constraints appear in full in this section.

Theorem 1 *There exists a family of set cover problems on pseudodisks for which there is no dual hitting set formulation on pseudodisks (and equivalently, such hitting set problems on pseudodisks cannot be dualized into set cover instances on pseudodisks).*

Proof. Consider a set X of n points in the plane in general position. Let $\mathcal{R}$ be a maximal set of circular disks so that $R_i \cap X \neq R_j \cap X$ for any $R_i, R_j \in \mathcal{R}$; there are $\omega(n^2)$ such disks [11]. Since the arrangement of a set of n pseudodisks has at most $n^2 - n + 2$ cells [15], this is the maximum number of cells in the arrangement dualizing the n points of X. However, the $\omega(n^2)$ disks in the primal require $\omega(n^2)$ distinct cells in the arrangement, and so such an instance cannot be dualized. □

Theorem 2 *Problems defined on translates of any single object (including pseudodisks) may always be dualized.*

Proof. Given a canonical object C, choose a reference point r so that $r \in C$, and let $-C$ be the reflection of C through r. To create a dualization, replace every translated object with a similarly translated instance of r, and replace every point with a translation of $-C$ so that r is incident upon the point. Due to the reflection through r, every point (of the plane) in C maps to a unique point in $-C$, and this point maps back to the original point again. Therefore, a set R in the primal contains a point x if and only if the dual of R is in the dual of x. □

Lemma 6 *There is a deterministic method for creating a cycle C in the embedding of G so that the clockwise ordering of the vertices in C is defined by corresponding pseudodisks in the primal.*

Proof. Suppose we wish to add the edge (v_1, v_k) to the graph G, which will result in a new cycle C. Let (v_1, v_2) be the other edge incident upon v_1 in the cycle. We know that part of ∂R_1 lies outside of $R_2 \cup R_k$ by the PF property. Consider a very small pseudodisk R' that covers some point on $\partial R_1 \setminus R_2 \cup R_k$ as a reference, and assume that the edge (v_1, v') is required in G.

Any edge (v_i, v_{i+1}) in the cycle represents two vertices whose corresponding objects R_i and R_{i+1} in the primal have a non-empty area of intersection. By Rule 1, $R_i \cap R_{i+1}$ is not covered by any other object, and so there exists a point in $R_i \cap R_{i+1}$ outside of R_1. Therefore, we may place a point in $R_i \cap R_{i+1}$ in the primal for every edge (v_i, v_{i+1}) in the cycle. Given two consecutive edges of the cycle (v_i, v_{i+1}) and (v_{i+1}, v_{i+2}), there exists a path in the primal inside $R_{i+1} \setminus R_1$ from the point in $R_i \cap R_{i+1}$ to that in $R_{i+1} \cap R_{i+2}$, since the objects are pseudodisks (to prevent such a path, $R_{i+1} \setminus R_1$ would have to be disjoint). Let H be the path defined by joining all of the points defined by the cycle in this manner.

The points in $R_1 \cap R_2$ and $R_1 \cap R_k$ are in R_1, so the union of the path H with R_1 defines one unbounded region of the plane and at least one bounded region outside of R_1 (H need not be simple). The reference pseudodisk R' will either be on the boundary of the unbounded region or a bounded region. The relative order of the edges (v_1, v_2), (v_1, v_k), and (v_1, v') around v_1 is uniquely defined, as discussed earlier. Now the rule for closing the cycle is as follows: the cycle encloses v' in G if and only if R' is on the boundary of a bounded region of the plane defined by $H \cup R_i$ in the primal. □

Lemma 8 *The support G is planar.*

Proof. We establish that G is planar by demonstrating that edge crossings are never necessary under this scheme. Suppose the graph G has been built so that the support is planar so far, but the next edge to be added would violate planarity. Let one such edge be (v_1, v_2). Therefore, v_2 is in the adjacency list for v_1, and $R_1 \cap R_2$ contains the cell of the arrangement for which we are building an induced subgraph in the support.

If v_1 cannot be connected to v_2 with an edge while preserving planarity, then v_1 and v_2 belong to a connected component of G (otherwise we could place them in the same face), but they are separated by at least one cycle of edges C whose vertices are not part of the subgraph we are constructing (otherwise the vertex in the cycle would be part of some connected component that we could add an edge to). At minimum, the cycle C corresponds to a sequence of pseudodisks that are pairwise intersecting, and whose union is a bounded region of the plane and which may separate the plane into several regions. The cell of the arrangement must lie outside of all pseudodisks in C; any pseudodisk covering the cell has a vertex in the subgraph. One of the vertices in $\{v_1, v_2\}$ is on the interior of C and the other is not, and because the combinatorial structure of the primal is preserved in the support G, one of $\{R_1, R_2\}$ lies partially in the bounded region of the plane defined by the boundaries of the regions defining the cycle, and the other lies partially on the unbounded region. Therefore, either R_1 or R_2 must cross the pseudodisks in C to cover the cell in the arrangement. However, no pseudodisk may cross C. To do so would require either intersecting the boundary of a pseudodisk in C in at least four places (which violates the definition of a pseudodisk), or covering the area of intersection of at least two pseudodisks in C (which means that C is not a cycle, by Rule 1). □

Lemma 9 *Let C denote a cycle in an induced subgraph S of G in the embedding, where S corresponds to a point x in the primal and C corresponds to pseudodisks $\mathcal{R}_C$. Any vertices on the interior of the bounded region defined by C must correspond to pseudodisks that cover x in the primal.*

Proof. Consider a cycle C in S with a vertex v (corresponding to a pseudodisk R in the primal) where v has a neighbour

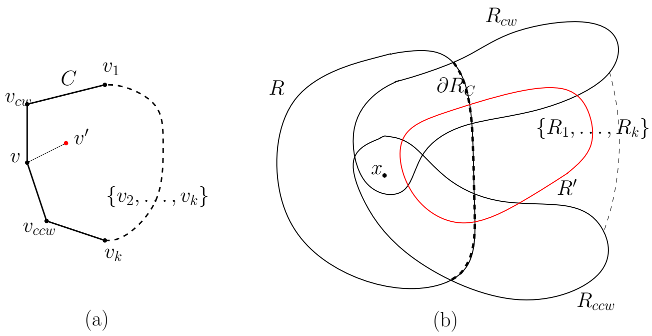

Figure 3: (a) Let C be a cycle in a subgraph S where a vertex $v' \notin S$ is in the interior of C, and the edge (v, v') exists in the support. v_{cw} (resp. v_{ccw}) is the clockwise (resp. counterclockwise) neighbour of v on the cycle. (b) All vertices in S correspond to pseudodisks covering a point x in the primal, and R' (the dual of v') does not cover x. The union of the pseudodisks dualizing the vertices in C (minus v) form a contiguous subset of the boundary of R which we call ∂R_C (drawn as the thick dashed line). R' covers a contiguous subset of ∂R_C.

v' on the interior of the cycle and $v' \notin S$. The pseudodisks in $\mathcal{R}_C \setminus R$ (in the primal) have a common area of intersection in R (since they all cover x and $x \in R$), and each edge of the cycle corresponds to a pair of pseudodisks that must intersect outside of R by Rule 1. It follows that the union of the pseudodisks in $\mathcal{R}_C$ covers a contiguous portion of the boundary of R, otherwise the boundaries of some pair of pseudodisks would have to intersect each other at four points. Let ∂R_C denote this portion of the boundary of R, as illustrated in Figure 3.

Now consider the two neighbours of v in C; call v_{cw} (resp. v_{ccw}) the clockwise (resp. counterclockwise) neighbour of v on C, and let R_{cw} (resp. R_{ccw}) denote the corresponding pseudodisk in the primal. Since v' lies on the interior of cycle C, by our construction algorithm, the entry point of R' lies between that of R_{cw} and R_{ccw} on ∂R_C. Consider a pseudodisk $R_i \in \mathcal{R}_C$ that contains the entry point of R'. The exit point of R' must be outside of R_i, otherwise either $R' \subset R_i \cup R$, which would violate the PF property, or $R \cap R' \subseteq R_i$, which would mean there is no edge (v, v') by Rule 1. Furthermore, the entry and exit points of R' cannot span those of any pseudodisk $R_j \in \mathcal{R}_C$. To do so would require that $|\partial R_j \cap \partial R' \cap R| = 2$, since R' cannot cover the point x in $R \cap R_j$, which would imply that $R_j \subseteq R \cup R'$, violating the PF property. Therefore, if the entry point of R' is in R_i, we can choose a pseudodisk R_j containing the exit point so that R_i and R_j are dualized by the vertices v_i and v_j and the edge (v_i, v_j) is in C.

Since the entry point of R' is in $R_i \setminus R_j$ and the exit point is in $R_j \setminus R_i$, we may conclude that $R_i \cap R_j \cap \partial R_C \subset R'$. Therefore, by Lemma 4, R' must cover either $R \cap R_i \cap R_j$ or $R_i \cap R_j \setminus R$. Of course, it cannot cover $R \cap R_i \cap R_j$ since this includes the point x, and we assumed that $x \notin R'$. However, R' cannot cover $R_i \cap R_j \setminus R$, since this implies that $\{R_i, R_j, R, R'\}$ is a crossing quad. Since either R' must cover x or v' must be outside the cycle C, we have a contradiction. □

Lemma 12 *Dualization of an instance of PCF-SC on pseudodisks to an instance of the hitting set problem on pseudodisks can be completed in $O(Im^5 \log m + mn)$ time, where $m = |\mathcal{R}| = |V| = |P|$, $n = |X| = |\mathcal{Q}|$, and I denotes the time required to compute the intersection points of a pair of pseudodisks.*

Proof. We show that the dualization can be completed in $O(Im^5 \log m + mn)$ time, where $m = |\mathcal{R}| = |V| = |P|$, $n = |X| = |\mathcal{Q}|$, and I denotes the time required to compute the intersection points of a pair of pseudodisks in the input set cover instance. Better analysis might establish a lower worst-case running time, but our goal is simply to establish polynomial running time.

The construction of the support takes $O(Im^5 \log m)$ time. The arrangement of a set of m pseudodisks has at most $m^2 - m + 2$ cells [15], and so the support has $O(m^2)$ subgraphs. In Algorithm 1, line 3 takes $O(m)$ time, line 4 may be completed in $O(Im^2 + m^2 \log m)$ time (sort the cells), line 5 may be done naïvely in $O(m^4)$ time, and line 6 may be completed in $O(m^2)$ time by traversing the arrangement. The loop iterates $O(m^2)$ times, and inside the loop we find the edges of the induced subgraph in G and then run Kruskal's minimum spanning tree algorithm to create a connected subgraph. Since we may have $O(m^2)$ edges in the graph, the running time is in $O(m^2 \log m)$ [7, p.570]. Checking whether a new edge violates planarity may be done in $O(m)$ time, but the bottleneck on the running time is when an edge closes a cycle. We can find the path in $O(Im)$ time, and the determination of whether the pseudodisk R' is on a closed face may be done in $O(Im)$ time by using one point in R' as a test for containment. This may be performed any time that an edge is added, so the work inside the loop takes $O(Im^3 \log m)$ time, giving the $O(Im^5 \log m)$ bound on the running time for building the support.

The algorithm for the construction of the hitting set instance from the support runs in $O(m^5 + mn)$ time. The orthogonal box drawing of the support runs in $O(m)$ time, and the determination of the positions of the points in P may be done at the same time. The insertion of a pseudodisk by fattening the edges of the corresponding subgraph may be done in $O(m^2)$ time, as each vertex of the subgraph may require that $O(m)$ other edges be fattened locally. There are $O(m^2)$ pseudodisks to be inserted, and so this may be done in $O(m^4)$ time. Any two objects have at most $O(m)$ intersections at this point, and each object is composed of $O(m)$ line segments, so the set of intersection points for the pair may be found in $O(m \log m)$ time [3], and these may placed in order around the boundary of one of the objects. Reducing the number of intersection points for the pair requires at most $O(m)$ iterations of the algorithm to remove pairs of intersection points. The algorithm may require determining which of two differences of the objects contains no points of P, and this may be done in $O(m^2)$ time by checking each point for containment, since each of the difference regions has $O(m)$ edges. Therefore, each pair may be

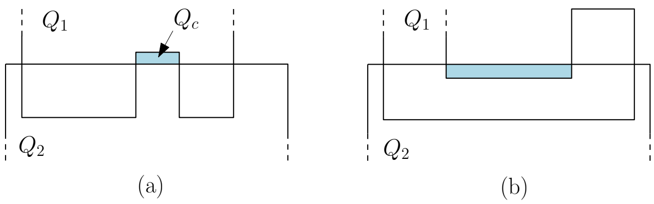

Figure 4: If ∂Q_1 and ∂Q_2 intersect in at least four places, each set of four consecutive intersections falls into one of two cases: (a) $Q_1 \setminus Q_2$ and $Q_2 \setminus Q_1$ are connected, or (b) $Q_1 \setminus Q_2$ and $Q_2 \setminus Q_1$ are not connected. The regions that may be resolved to reduce the number of intersections are shaded.

repaired into pseudodisks in $O(m^3)$ time, for a total upper bound of $O(m^5)$ time to convert all objects into pseudodisks. Finally, pseudodisks are added for each point in X that has no corresponding pseudodisk in $\mathcal{Q}$, i.e., from those cells in the arrangement containing more than one point from X. There are $O(n)$ such disks, and each may be placed in $O(m)$ time. □

B Appendix: Reducing the Number of Intersections

Consider a pair of objects, call them Q_1 and Q_2, enclosing S_1 and S_2 respectively, that are not pseudodisks. It has already been established that the objects are simple and enclosed by closed Jordan curves, so the only remaining possible violation is that the boundaries of Q_1 and Q_2 intersect more than twice. The sequence of events of a walk on ∂Q_1 (w.l.o.g.) must contain $2^+, 2^-, 2^+, 2^-$, and this gives rise to two cases to consider: local to these events, $Q_1 \setminus Q_2$ and $Q_2 \setminus Q_1$ are both either connected regions or not, as illustrated in Figure 4. The action taken to reduce the number of intersections while preserving the dual property depends on the case.

Case 1. If $Q_1 \setminus Q_2$ and $Q_2 \setminus Q_1$ are connected, then there must exist a bounded region of the plane Q_c outside of $Q_1 \cup Q_2$. We claim that we may remove the two points of intersection between Q_1 and Q_2 on the boundary of this region by moving the boundary of Q_1 (w.l.o.g.) to be just inside that of Q_2. This will not cause Q_1 to cover any additional points in $Q_2 \cap P$, so the only way that this move affects the hitting set combinatorially is if there exists any vertex v_c of G in Q_c, i.e. the bounded region of the plane that was formerly not covered by Q_1 or Q_2, but is now covered by Q_1.

Consider the cycle of $S_1 \cup S_2$ that encloses Q_c in the planar embedding of the support. As with the proof of Lemma 9, we note that if any vertex exists in Q_c, then at least one vertex v' exists that is a neighbour of a vertex v in $S_1 \cup S_2$, and say w.l.o.g. that $v \in S_1$. We define R_{cw} and R_{ccw} as before, and let R^i_{cw} be the first vertex in $S_1 \cap S_2$ on the cycle in a clockwise direction from v and R^i_{ccw} is defined analogously for the counterclockwise direction (R_{cw} and R^i_{cw} and also R_{ccw} and R^i_{ccw} are not necessarily distinct). There is a region of ∂R covered by the pseudodisks in the primal corresponding to the vertices of the cycle moving clockwise from R_{cw} to R^i_{cw}, and also analogously for R_{ccw} and R^i_{ccw}, call them ∂R_{cw} and ∂R_{ccw} respectively. In fact, since $R^i_{cw} \cap R^i_{ccw}$ contains points dualizing both S_1 and S_2, their pairwise area of intersection must extend across the boundary of R to cover points $x_1 \in R$ and $x_2 \notin R$ (S_i dualizes x_i), and so $\partial R_{cw} \cup \partial R_{ccw}$ defines a contiguous interval of ∂R. Now the same argument may be applied as in Lemma 9 to conclude that either R' covers x_1 (the dual of S_1), or else R' must be outside of the cycle. Therefore, removing intersections of the dual in this case may be done without covering additional vertices.

Case 2. If $Q_1 \setminus Q_2$ and $Q_2 \setminus Q_1$ are not connected, then one of the two regions of $Q_2 \setminus Q_1$ may be moved inside Q_1 unless both regions contain vertices of S_2. Suppose this is the case, and so an edge flip causes Q_1 to cover additional vertices of S_2. This implies that there is some path that is a subset of Q_1 for which Q_2 has vertices on both sides (i.e., the endpoints of the path could not be joined to create a cycle without enclosing vertices of Q_2). However, it was demonstrated that one subgraph cannot cross another in the proof of Lemma 9, so one of the regions of $Q_2 \setminus Q_1$ cannot contain vertices of S_2. Therefore, removing points of intersection may again be done without covering additional vertices. In both cases, there are no vertices in the regions where edges are moved to resolve conflicts. Therefore, these resolutions may be done without creating additional points of intersection with other objects by similarly translating other edges if necessary.