Polynomial Kernels for Hard Problems on Disk Graphs

Bart Jansen*

Utrecht University, PO Box 80.089, 3508 TB Utrecht, The Netherlands
bart@cs.uu.nl

Abstract. Kernelization is a powerful tool to obtain fixed-parameter tractable algorithms. Recent breakthroughs show that many graph problems admit small polynomial kernels when restricted to sparse graph classes such as planar graphs, bounded-genus graphs or H-minor-free graphs. We consider the intersection graphs of (unit) disks in the plane, which can be arbitrarily dense but do exhibit some geometric structure. We give the first kernelization results on these dense graph classes. Connected Vertex Cover has a kernel with $12k$ vertices on unit-disk graphs and with $3k^2 + 6k$ vertices on disk graphs with arbitrary radii. Red-Blue Dominating Set parameterized by the size of the smallest color class has a linear-vertex kernel on planar graphs, a quadratic-vertex kernel on unit-disk graphs and a quartic-vertex kernel on disk graphs. Finally we prove that H-Matching on unit-disk graphs has a linear-vertex kernel for every fixed graph H.

1 Introduction

Motivation A parameterized problem is a language $L \subseteq \Sigma^* \times \mathbb{N}$. An instance $(x, k) \in \Sigma^* \times \mathbb{N}$ is composed of x, the classical problem instance, and the parameter k that expresses some structural property of x. A kernelization algorithm (or *kernel*) is a mapping that transforms an instance (x, k) into an equivalent instance (x', k') in time $p(|x| + k)$ for some polynomial p, such that $(x, k) \in L \Leftrightarrow (x', k') \in L$ and $|x'|, k' \leq f(k)$ for a computable function f. The function f is the *size* of the kernel. Kernelization can be seen as a form of preprocessing with a performance guarantee on the data reduction that is obtained with respect to the parameter value k. Consult the survey by Guo and Niedermeier [1] for more background on kernelization, and the book by Downey and Fellows [2] for the theory of fixed parameter complexity. Kernelization has been an area of intensive study during the past few years. A rich theoretical framework is in development, yielding both positive and negative results with regard to the limits of kernelization. A celebrated result by Bodlaender et al. [3] states that there are problems in FPT that cannot have kernels of polynomial size, unless some unlikely complexity-theoretic collapse occurs. Recent results also

* This work was supported by the Netherlands Organisation for Scientific Research (NWO), project "KERNELS: Combinatorial Analysis of Data Reduction".

show that all problems that satisfy certain compactness and distance properties admit polynomial kernels on restricted graph classes such as planar graphs [4], graphs of bounded genus [5] and H-minor-free graphs [6]. The common theme in these frameworks is that they yield kernels for problems on *sparse graphs*, in which the number of edges is bounded linearly in the number of vertices.

Our work focuses on graph classes that exhibit some geometric structure but can be arbitrarily dense: the intersection graphs of disks in the plane ("disk graphs"). Each vertex is represented by a disk, and two vertices are adjacent if and only if their disks intersect. Disks that touch eachother but do not overlap are also said to be intersecting. It is a well-known consequence of the Koebe-Andreev-Thurston theorem that planar graphs are a strict subclass of disk graphs [7, Section 3.2]. If all the disks have the same radius then by a scaling argument we may assume that this radius is 1: the intersection graphs of such systems are therefore called *unit-disk* graphs. It is easy to see that (unit)disk graphs may contain arbitrarily large cliques. Many classical graph problems remain NP-complete when restricted to unit-disk graphs [8], and several important parameterized problems such as k-Independent Set and k-Dominating Set remain $W[1]$-hard [9,10] for unit-disk graphs. In this paper we will show how the structure of disk graphs graphs may be exploited to obtain polynomial kernels.

Previous Work Hardly any work has been done on kernelization for disk graphs: only a single result in this direction is known to us. Alber and Fiala [11] use the concept of a *geometric problem kernel* as in their work on k-Independent Set: they obtain a reduced instance in which the *area of the union of all disks* is bounded in k. This geometric kernel leads to subexponential exact algorithms, and to an FPT algorithm by applying restrictions that bound the maximum degree of the graph. We believe we are the first to present polynomial problem kernels on dense disk graph classes.

Our Results We show that the Connected Vertex Cover problem has a trivial $12k$-vertex kernel on unit-disk graphs, and obtain a more complex kernel with $3k^2 + 6k$ vertices for disk graphs with arbitrary radii. We prove that Red-Blue Dominating Set parameterized by the size of the smallest color class admits a kernel of linear size on planar graphs, of quadratic size on unit-disk graphs and of quartic size on disk graphs with arbitrary radii. Note that neither of these two problems admit polynomial kernels on general graphs unless the polynomial hierarchy collapses [12]. We also present a linear kernel for the H-Matching problem on unit-disk graphs, which asks whether a unit-disk graph G contains at least k vertex-disjoint copies of a fixed connected graph H.

2 Preliminaries

All graphs are undirected, finite and simple unless explicitly stated otherwise. Let $G = (V, E)$ be a graph. For $v \in V$ we denote the *open* (resp. closed) neighborhoods of v by $N_G(v)$ and $N_G[v]$. The degree of a vertex v in graph G is denoted by $\deg_G(v)$. We write $G' \subseteq G$ if G' is a subgraph of G. For $X \subseteq V$ we denote by $G[X]$ the subgraph of G that is induced by the vertices in X. We denote the set of connected components of G by $\textsc{Comp}(G)$. We treat geometric objects as closed sets of points in the Euclidean plane. If v is a vertex of a geometric intersection graph, then we use $\mathcal{D}(v)$ to denote the geometric representation of v; usually this is just a disk. We will write $\mathcal{D}(V')$ for $V' \subseteq V$ to denote the union of the geometric objects representing the vertices V'. Some proofs had to be omitted due to space restrictions.

Lemma 1. *Let $G = (X \cup Y, E)$ be a planar bipartite graph. If for all distinct $v, v' \in Y$ it holds that $N_G(v) \not\subseteq N_G(v')$ then $|Y| \leq 5|X|$.* □

Lemma 2 (Compare to Lemma 3.4 of [13]). *If v is a vertex in a unit-disk graph $G = (V, E)$ then there must be a clique of size at least $\lceil \deg_G(v)/6 \rceil$ in $G[N_G(v)]$.* □

3 Connected Vertex Cover

The Connected Vertex Cover problem asks for a given connected graph $G = (V, E)$ and integer k whether there is a subset $S \subseteq V$ with $|S| \leq k$ such that every edge in G has at least one endpoint in S, and such that $G[S]$ is connected. Guo and Niedermeier gave a kernel with $14k$ vertices for this problem [4] restricted to planar graphs. More recently it was shown that Connected Vertex Cover on general graphs does not admit a polynomial kernel unless the polynomial hierarchy collapses [12]; this situation is in sharp contrast with the regular Vertex Cover problem, which has a kernel with $2k$ vertices [1]. It is not hard to prove that a vertex cover for a connected unit-disk graph on n vertices must have size at least $n/12$ [14, Theorem 2] which yields a trivial linear kernel.

Observation 1 Connected Vertex Cover has a kernel with $12k$ vertices on unit-disk graphs.

The situation becomes more interesting when we consider disk graphs with arbitrary radii, where we show the existence of a quadratic kernel. To simplify the exposition we start by giving a kernel for an annotated version of the problem. Afterwards we argue that the annotation can be undone to yield a kernel for the original problem. An instance of the problem Annotated Connected Vertex Cover is a tuple $\langle G, k, C \rangle$ where G is the intersection graph of a set of disks in the plane, k is a positive integer and $C \subseteq V(G)$ is a subset of vertices. The question is whether there is a connected vertex cover $S \subseteq V(G)$ of cardinality at most k such that $C \subseteq S$. We introduce some notation to facilitate the description of the reduction rules.

Definition 1. *Let $\langle G, k, C \rangle$ be an instance of* Annotated Connected Vertex Cover. *A vertex $v \in C$ is* marked*; vertices in $V(G) \setminus C$ are* unmarked*. For unmarked vertices we distinguish between two types: an unmarked vertex is* dead *if all its neighbors are marked, and it is* live *if it has an unmarked neighbor. We call an edge* covered *if it is incident on a marked vertex, and* uncovered *otherwise.*

Observe that all edges incident on dead unmarked vertices will be covered by the marked vertices in any solution. Therefore the dead vertices can only be useful in a solution to ensure connectivity. The live unmarked vertices may be needed to cover additional edges.

Reduction Rule 1 If there is an unmarked vertex v with more than k neighbors, then add v to C (i.e. mark v).

This is an annotated analogue of Buss' rule for the Vertex Cover problem. It is easy to see that v must be part of any solution of size at most k; for if v is not taken in a vertex cover then all its $> k$ neighbors must be taken. We give a new definition that will simplify the exposition of the next reduction rule.

Definition 2. *We define the* component graph *that corresponds to the instance* $\langle G, k, C\rangle$ *with dead vertices* D *to be the bipartite graph* $G_C := (\textsc{Comp}(G[C]) \cup D, E)$ *where there is an edge between a connected component* $X \in \textsc{Comp}(G[C])$ *and a dead vertex* $d \in D$ *if and only if* $N_G(d) \cap V(X) \neq \emptyset$*, i.e. if* d *is adjacent in* G *to a vertex in the connected component* X*.*

Reduction Rule 2 If there are dead vertices u and v such that $N_{G_C}(u) \subseteq N_{G_C}(v)$ then delete u.

For the correctness of this rule observe that if S is a solution containing u, then $(S \setminus \{u\}) \cup \{v\}$ is also a solution: since all neighbors of u are marked, the removal of u does not cause edges to become uncovered. Removal of u cannot cause the vertex cover to become disconnected because all components of $G[C]$ that were connected through u remain connected through v.

Lemma 3. *Let* $\langle G, k, C\rangle$ *be a reduced instance of* Annotated Connected Vertex Cover *for a disk graph* G*, and let* D *denote the set of dead vertices. There is a vertex set* R *and a bipartite graph* $G^* = (R \cup D, E')$ *with the following properties:*

(i) The graph G^* *is planar.*
(ii) If $N_{G^*}(d) \subseteq N_{G^*}(d')$ *for* $d, d' \in D$ *then* $N_{G_C}(d) \subseteq N_{G_C}(d')$*.*
(iii) The number of vertices in R *is at most* $|C|$*.*

Proof. Assume the conditions stated in the lemma hold. Fix some realization of G by intersecting disks in the plane. We will create disjoint regions R that are subsets of the maximally connected regions in $\mathcal{D}(C)$. We define the regions R in two parts by setting $R := R_1 \cup R_2$.

- For every pair consisting of a component $X \in \textsc{Comp}(G[C])$ and dead vertex $d \in D$ such that $\mathcal{D}(V(X)) \subseteq \mathcal{D}(d)$ we let $\mathcal{D}(V(X))$ be a region in R_1.
- The set R_2 consists of the maximally connected regions of the plane that are obtained by taking the union of all disks of vertices in C, and subtracting the interior of all disks of vertices in D. Observe that a region $\mathcal{D}(V(X))$ for $X \in \textsc{Comp}(G[C])$ may be split into multiple regions by subtracting the interiors of $\mathcal{D}(D)$; see Figure 1(b).

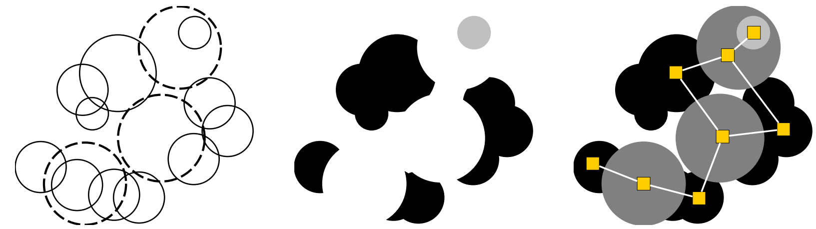

(a) Subgraph induced by D and C. Disks of dead vertices are drawn with broken lines.

(b) Regions R_1 colored grey, R_2 colored black.

(c) Embedding of bipartite graph G^*.

Fig. 1. Construction of the planar bipartite graph G^* from an instance $\langle G, k, C\rangle$ with dead vertices D.

Since the constructed regions R are subsets of the maximally connected regions induced by $\mathcal{D}(C)$, we know that for each $r \in R$ there is a unique $X \in \textsc{Comp}(G[C])$ such that $r \subseteq \mathcal{D}(V(X))$. Define the parent $\pi(r)$ of r to be the component X for which this holds, and define the parent of a set of regions $R' \subseteq R$ to be the union $\cup_{r\in R'}\{\pi(r)\}$. Let $d \in D$ be a dead vertex and consider the regions of R intersected by $\mathcal{D}(d)$. It is not difficult to verify that by construction of R, $\mathcal{D}(d)$ intersects at least one region in R for every component $X \in \textsc{Comp}(G[C])$ that d is adjacent to in G. Since the regions of R are subsets of $\mathcal{D}(C)$, the disk $\mathcal{D}(d)$ can *only* intersect a region $r \in R$ if d is adjacent in G_C to $\pi(r)$. Hence we establish that if R' is the set of regions in R intersected by $\mathcal{D}(d)$, then $\pi(R') = N_{G_C}(d)$. We will refer to this as the *parent property.*

We define the bipartite graph G^* with partite sets D and R as the intersection graph of D and R: there is an edge between a vertex $d \in D$ and a vertex corresponding to a region $r \in R$ if and only if $\mathcal{D}(d)$ intersects region r. We will show that G^* is planar by embedding it in the geometric model of the disk graph G. The planarity of G^* can be derived from the following observations:

(a) If some region intersects only one region in the other partite set, then it will become a degree-1 vertex in G^* and it will never violate planarity. So we only need to consider regions that intersect at least two regions in the other partite set.

(b) Every region $r \in R_1$ is completely contained within some $\mathcal{D}(d)$ for $d \in D$ and hence r only intersects one region in D since disks of dead vertices are disjoint; therefore we can ignore the regions in R_1 by the previous observation.

(c) For every disk $\mathcal{D}(d)$ for $d \in D$ the interior of the disk is not intersected by any other region in $D \cup R_2$, by construction of R_2. Similarly, for every region $r \in R_2$ the interior of that region is not intersected by any other region in $D \cup R_2$.

We show how to create an embedding in the plane of the intersection graph of $D \cup R_2$, which sufficies to prove that G^* is planar by observation (b). For every region in $D \cup R_2$ we select a location for the corresponding vertex in the interior of that region. Since the interior of every region is connected, we can draw edges from these points to all intersected neighboring regions. Since no three regions have a common intersection we can draw these edges in the plane without crossings - see Figure 1(c). This proves that G^* is planar and establishes (i). The parent property of the regions R establishes (ii), since the neighborhood of a dead vertex in G^* corresponds to the set of regions of R it intersects. Hence if two vertices $d, d' \in D$ satisfy $N_{G^*}(d) \subseteq N_{G^*}(d')$ then $\pi(N_{G^*}(d)) \subseteq \pi(N_{G^*}(d'))$ and therefore $N_{G_C}(d) \subseteq N_{G_C}(d')$.

As the last part of the proof we establish that $|R| \leq |C|$ by showing that we can charge every region in r to a vertex of C such that no vertex is charged twice. A region $r \in R_1$ corresponds to a component $X \in \text{Comp}(G[C])$ such that $\mathcal{D}(V(X))$ is contained in $\mathcal{D}(d)$ for some $d \in D$; for every such connected component X there is exactly one region in R_1, and we can charge the cost of vertex r to one vertex in the connected component X. For the regions in R_2 the situation is slightly more involved. Consider some vertex $v \in C$. It is not hard to see that if we start with the region $\mathcal{D}(v)$ and subtract the interiors of mutually *disjoint* disks from that region, then the result must be either an empty region or a connected region. This implies that for every vertex $v \in C$, there is at most one region $r \in R_2$ that has a non-empty intersection with $\mathcal{D}(v)$. Since every region of R_2 is a subset of $\mathcal{D}(C)$, every region $r \in R_2$ intersects $\mathcal{D}(v)$ for at least one $v \in C$, and by the previous argument r is the only region of R_2 for which this holds. Therefore we can charge the region r to such a vertex v and we can be sure that we never charge to a vertex of C twice; this proves that $|R| = |R_1| + |R_2| \leq |C|$, which establishes (iii) and completes the proof. □

Theorem 1. Annotated Connected Vertex Cover *has a kernel with* $2k^2 + 6k$ *vertices on disk graphs with arbitrary radii.*

Proof. Given an instance $\langle G, k, C \rangle$ we first exhaustively apply the reduction rules; it is not hard to verify that this can be done in polynomial time. Let $\langle G', k', C' \rangle$ be the resulting reduced instance. Since the reduction rules do not change the value of k we have $k' = k$. If the reduced instance contains more than k^2 uncovered edges then the answer to the decision problem must be no and we can output a trivial no-instance: all uncovered edges need to be covered by an unmarked vertex, and any unmarked vertex may cover at most k edges since it has degree at most k by Rule [1] - therefore any vertex cover must consist of more than k unmarked vertices if there are more than k^2 uncovered edges. If the number of uncovered edges is at most k^2 then the number of live vertices is at most $2k^2$ since every live vertex is incident on at least one uncovered edge.

If the number of marked vertices C exceeds k then clearly there is no solution set containing C of size at most k, and we output a no instance. Otherwise the number of marked vertices is at most k. Consider the dead vertices D in the reduced instance, and the bipartite planar graph $G^* = (R \cup D, E')$ whose existence is guaranteed by Lemma 3. By Rule [2] we know that $N_{G_C}(d) \not\subseteq N_{G_C}(d')$ for distinct vertices $d, d' \in D$, and by (ii) of the Lemma this implies that $N_{G^*}(d) \not\subseteq N_{G^*}(d')$. Therefore we may apply Lemma 1 to graph G^* where R plays the role of X, and D plays the role of Y, to conclude that $|D| \leq 5|R|$. Using (iii) of the Lemma gives $|D| \leq 5|C| \leq 5k$, which is the final ingredient of the size bound.

Now we can bound the total size of a reduced instance. The vertices are partitioned into marked vertices, live vertices and dead vertices. There are at most k marked vertices, at most $2k^2$ live vertices and at most $5k$ dead vertices which shows that $|V(G')| \leq 2k^2 + 6k$. Therefore we can output the reduced instance $\langle G', k', C' \rangle$ as the result of the kernelization; by the safety of the reduction rules we know that $\langle G, k, C \rangle$ is a YES-instance if and only if $\langle G', k', C' \rangle$ is. □

Theorem 2. CONNECTED VERTEX COVER *has a kernel with* $3k^2 + 6k$ *vertices on disk graphs with arbitrary radii.*

Proof. We show how to undo the annotation. Consider an instance $\langle G, k \rangle$ of CONNECTED VERTEX COVER, and derive an instance $\langle G, k, C \rangle$ for the annotated problem by setting $C := \emptyset$. If we apply the kernel to the annotated instance, we get as output either a trivial NO instance or a reduced instance $\langle G', k', C' \rangle$. The case of a trivial NO instance is easy to handle so we do not treat it further. If the output is a reduced instance then by the definitions of the reduction rules we know that G' is a vertex-induced subgraph of G that was obtained by deleting dead vertices from the annotated instance. By definition of dead vertices we know that these deleted vertices form an independent set in G. To undo the annotation we augment the graph G' to ensure that all vertices in C need to be used in any connected vertex cover of size at most k. If a vertex $v \in C$ has more than k neighbors in G' then this is satisfied automatically. Now suppose that v has k or fewer neighbors in G'; since the vertex was annotated by Rule [1] it must have had at least $k + 1$ neighbors in G, some of which were later deleted by Rule [2]. To undo the annotation of vertex v we just add vertices from $N_G(v) \setminus N_{G'}(v)$ that were deleted in G' back to the graph, until the degree of v is more than k. If we do this for every marked vertex then the resulting instance is equivalent to $\langle G, k, C \rangle$ since the added vertices are independent and only adjacent to marked vertices, and we force all marked vertices to be part of a solution by making their degree large. Hence we ensure that all marked vertices occur in any connected vertex cover of size k, and all the edges that we introduce are covered by the marked vertices. Since there are at most k marked vertices and we re-add at most k dead vertices for each $v \in C$ to undo the annotation, the total number of added vertices is at most k^2. Because the size of G' is at most $2k^2 + 6k$ the claim follows. □

The kernel for CONNECTED VERTEX COVER can be lifted to more general geometric graph classes. For ease of presentation we have presented the kernel for the intersection graphs of disks, but the stated results should easily generalize to intersection graphs of connected geometric objects in pseudo-disk relation [15].

4 *H*-matching

The H-Matching problem asks whether a given graph G contains at least k vertex-disjoint subgraphs that are isomorphic to some fixed connected graph H. For ease of notation we define $|H|$ as short-hand for $|V(H)|$. Subgraph matching problems have received considerable attention from the FPT community, resulting in $O(k)$ vertex kernels for H-Matching on sparse graphs [5,6] and a kernel with $O(k^{|H|-1})$ vertices on general graphs [16]. The restriction to unit-disk graphs has been considered in the context of approximation algorithms [17]. For every fixed H we give a linear-vertex kernel for H-Matching in the case that G is required to be a unit-disk graph.

Reduction Rule 3 Delete all vertices that are not contained in a H-subgraph of G.

This rule is clearly correct. We can obtain a reduced graph by trying all $\binom{|V(G)|}{|H|}$ possible vertex sets of size $|H|$ and marking the vertices for which the guessed set forms a H-subgraph. Afterwards we delete all vertices that were not marked. Since we assume H to be fixed, this can be done in polynomial time.

Theorem 3. *Let $G = (V, E)$ be a reduced unit-disk instance of H-Matching. If there is a* maximal *H-matching in G that consists of k^* copies of H, then $|V| \in O(k^*)$.*

Proof. Let $G = (V, E)$ be a reduced graph and consider a maximal H-matching in G consisting of k^* copies. Let S be the vertices that occur in a matched copy of H. Since the selected copies are vertex-disjoint we have $|S| = k^*|H|$. Let $O := V \setminus S$ be the vertices not used in a copy of H in the matching. We will prove that the size of O is bounded in $|S|$.

Define $\delta_G(u, v)$ as the length of a shortest path between u and v in G, or $+\infty$ if u and v are not connected in G. Let d_H be the diameter of the graph H, i.e. $d_H := \max_{u \in V(H)} \max_{v \in V(H)} \delta_H(u, v)$. Since we assume H to be connected we have $d_H \leq |H|$. We measure the distance of a vertex v to the closest vertex in S by $\delta_G(v, S) := \min_{u \in S} \delta(u, v)$. Now suppose that there is some $o \in O$ with $\delta(o, S) > d_H$. By Rule [3] the vertex o must be contained in some subgraph $G' \subseteq G$ that is isomorphic to H. All vertices $v \in V(G')$ involved in this subgraph have a distance to o in G' of at most d_H by the definition of diameter. But since the distance between vertices $u, v \in V(G')$ in graph G is at most the distance between these vertices in G' (since a path in G' is also a path in G), this would imply that all vertices $v \in V(G')$ in the H-subgraph have a distance to o in G of at most d_H, and since $\delta(o, S) > d_H$ none of the vertices in $V(G')$ are in S. But that means that we can add G' as an extra copy of H to the matching, contradicting the assumption that we started from a *maximal* H-matching. Therefore we may conclude that $\delta(o, S) \leq d_H$ for all $o \in O$.

For the next step in the analysis we show that any vertex has a bounded number of neighbors in O. If $v \in V$ has more than $6(|H|-1)$ neighbors in O, then by Lemma 2 there is a clique of size $|H|$ in $G[N_G(v) \cap O]$. This clique must contain a subgraph isomorphic to H and hence it can be added to the H-matching, again contradicting the assumption that we started with a maximal matching. Therefore every vertex in G has at most $6(|H|-1)$ neighbors in O. Since there are exactly $k^*|H|$ vertices in S this shows that there are at most $k^*|H| \cdot 6(|H|-1)$ vertices $o \in O$ for which $\delta(o, S) = 1$. Now observe that a vertex v with $\delta(v, S) = 2$ must be adjacent to some vertex u with $\delta(u, S) = 1$. Since all such vertices u have a bounded number of neighbors outside O, we find that there are at most $k^*|H| \cdot (6(|H|-1))^2$ vertices $o \in O$ for which $\delta(o, S) = 2$. By generalizing this step we obtain a bound of $k^*|H| \cdot (6(|H|-1))^r$ on the number of vertices v that have $\delta(v, S) = r$. Since we established $\delta(o, S) \leq d_H$ for all $o \in O$ we can conclude that $|O| \leq \sum_{i=1}^{d_H} k^*|H| \cdot (6(|H|-1))^i$ which implies that $|O| \in O(k^*(6|H|)^{d_H})$. Now we can bound the size of the instance G by noting that $|V| = |S| + |O|$ and therefore $|V| \in O(k^*(6|H|)^{d_H})$. Since the diameter d_H of H is at most $|H|$ this completes the proof: for every fixed H the term $(6|H|)^{d_H} \leq (6|H|)^{|H|}$ is constant and hence $|V| \in O(k^*)$. □

Corollary 1. *H-Matching in unit-disk graphs has a kernel with $O(k)$ vertices for every fixed H.* □

5 Red-Blue Dominating Set

In the Red-Blue Dominating Set problem we are given a graph $G = (V, E)$, a positive integer k and a partition of the vertices into red and blue color classes R, B such that $V = R \cup B$ and $R \cap B = \emptyset$; the goal is to determine whether there is a set $S \subseteq R$ consisting of at most k red vertices such that every blue vertex in B has at least one neighbor in S. In the literature it is often assumed that G is bipartite with the red and blue classes as the partite sets; we explicitly do not require this here since bipartite disk graphs are planar [18, Lemma 3.4]. Under these assumptions the Red-Blue Dominating Set problem on disk graphs graphs is not easier than k-Dominating Set on those graphs when parameterized by the size of the solution set, since we can reduce from the regular k-Dominating Set problem by making two copies of every vertex, marking one of them as red and the other as blue: hence Red-Blue Dominating Set is $W[1]$-hard on unit-disk graphs when parameterized by k [10]. The situation changes when we parameterize by $|R|$ or by $|B|$, which causes the problem to become fixed parameter tractable on general graphs. Dom et al. [12] have shown that the Red-Blue Dominating Set problem parameterized by $|R| + k$ or $|B| + k$ does not have a polynomial kernel on general graphs, unless the polynomial hierarchy collapses. We show that the situation is different for disk graphs by proving that Red-Blue Dominating Set when parameterized by the size of the smallest color class has polynomial kernels on planar graphs and (unit)disk graphs. We use the same reduction rules for all three graph classes.

Reduction Rule 4 If there are distinct red vertices $u, v \in R$ such that $N_G(u) \cap B \subseteq N_G(v) \cap B$ then delete u.

It is easy to see that this rule is correct: the red vertex v can dominate all blue vertices that can be dominated by u, and hence there is always a smallest dominating set that does not contain u. The following rule is similar in spirit, but works on the other vertex set.

Reduction Rule 5 If there are distinct blue vertices $u, v \in B$ such that $N_G(u) \cap R \subseteq N_G(v) \cap R$ then delete v.

In this case we may delete v because whenever u is dominated by some red $x \in N_G(u)$, the vertex v must be dominated as well since $x \in N_G(v)$. These reduction rules immediately lead to a kernel with $O(\min(|R|, |B|))$ vertices on planar graphs by invoking Lemma 1. For unit-disk graphs and disk graphs with arbitrary radii we need the following structural results.

Theorem 4. *Let $G = (V, E)$ be a unit-disk graph whose vertex set is partitioned into red and blue color classes by $V = R \cup B$ with $R \cap B = \emptyset$. If for all distinct $u, v \in R$ it holds that $N_G(u) \cap B \not\subseteq N_G(v) \cap B$ then $|R| \in O(|B|^2)$.* □

A well-known construction from the area of computational geometry (see [19, Section 5.2]) shows that the bound in Theorem 4 is asymptotically tight.

Theorem 5. *Let $G = (V, E)$ be a disk graph with arbitrary radii whose vertex set is partitioned into red and blue color classes by $V = R \cup B$ with $R \cap B = \emptyset$. If for all distinct $u, v \in R$ it holds that $N_G(u) \cap B \not\subseteq N_G(v) \cap B$ then $|R| \in O(|B|^4)$.*

Proof. Assume the conditions stated in the theorem hold, and fix a realization of the disk graph G that specifies a disk for every vertex. For a pair of blue disks $\mathcal{D}(b_1), \mathcal{D}(b_2)$ we consider all points for which the distances to the boundaries of $\mathcal{D}(b_1)$ and $\mathcal{D}(b_2)$ are the same, i.e. the points that are equidistant to the boundaries of $\mathcal{D}(b_1)$ and $\mathcal{D}(b_2)$. If the disks do not completely coincide then this set of points forms a curve in the plane. The type of curve depends on the relative orientation of the two disks. If the two disks have equal radius then the curve is a line. If the radii differ then the curve is an ellipse if one disk completely contains the other, and otherwise the curve is a branch of a hyperbola. Consider the arrangement in the plane that is induced by the set of all possible $\binom{|B|}{2}$ bisector curves for pairs of blue disks, and let F be a face of this arrangement (see Figure 2).

Suppose we choose a point p in the face F and grow a disk around p by gradually increasing its radius. If the disk grows large enough then it will intersect some of the blue disks. Once it intersects a blue disk $\mathcal{D}(b)$ at a certain radius, then naturally it will keep intersecting $\mathcal{D}(b)$ as its radius increases. Observe that the *order* in which the blue disks are encountered as the radius increases does not depend on where we place p within face F: the relative order in which two blue disks $\mathcal{D}(b_1)$ and $\mathcal{D}(b_2)$ are encountered is determined by the relative position of point p to the bisector curve of $\mathcal{D}(b_1)$ and $\mathcal{D}(b_2)$. Since F is a face in the arrangement of all bisector curves, this relative position will be the same for all points $p \in F$, and therefore all choices of a point p in F yield the same order in which blue vertices are encountered. Now observe that for every possible order

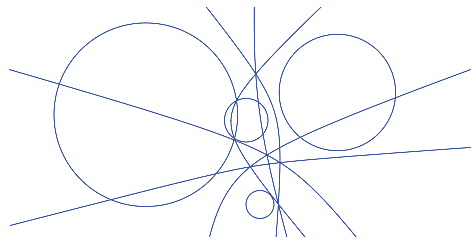

Fig. 2. Four circles and the bisector curves for all pairs.

of encountering the blue disks there can be at most one red vertex: if two red vertices are placed on positions that yield the same order, then the neighborhood of one must be a subset of the neighborhood of the other - which is not possible by the assumption in the statement of the theorem. Therefore F can contain at most one center of a red disk. The same argument shows that every edge or vertex in the arrangement induced by the bisector curves can contain at most one red vertex. Hence we can bound the number of red vertices by bounding the total number of vertices, edges and faces of the arrangement (i.e. the complexity of the arrangement).

The arrangement of curves consists of lines, branches of hyperbolas and ellipses. It is not hard to verify that each pair of these objects can intersect at most a constant number of times. In the area of computational geometry it is a well-known fact [20] that the complexity of an arrangement of n curves is $O(n^2)$ if the number of intersections between each pair of curves is bounded by a constant. In our case we find that there are $|B|^2$ objects and hence the complexity is $O(|B|^4)$. Since we have at most one red vertex per element in the arrangement, the claim follows. □

Corollary 2. *The* Red-Blue Dominating Set *problem admits a kernel with* $O(\min(|R|, |B|))$ *vertices on planar graphs, with* $O(\min(|R|^2, |B|^2))$ *vertices on unit-disk graphs and with* $O(\min(|R|^4, |B|^4))$ *vertices on disk graphs with arbitrary radii.*

6 Conclusion

We have shown that the geometric structure of dense (unit)disk graphs can be exploited to obtain polynomial size problem kernels for Connected Vertex Cover with parameter k and Red-Blue Dominating Set parameterized by $\min(|R|, |B|)$, while these problems do not admit polynomial kernels on general graphs. For fixed graphs H the H-Matching problem has a quadratic kernel in arbitrary graphs; we showed how to obtain a linear kernel in the restricted setting of unit-disk graphs. It would be interesting to see whether there are graph problems which are $W[1]$-hard in general, but have small kernels on (unit)disk graphs. Due to the structure of cliques in disk graphs it might be worthwhile to investigate whether Edge Clique Cover admits a polynomial kernel on unit-disk graphs, and whether Partition Into Cliques on unit-disk graphs is fixed-parameter tractable.

Acknowledgements I would like to thank Manu Basavaraju, Hans Bodlaender, Marc van Kreveld, Erik Jan van Leeuwen and Jan van Leeuwen for insightful discussions.

References

1. Guo, J., Niedermeier, R.: Invitation to data reduction and problem kernelization. SIGACT News **38** (2007) 31–45
2. Downey, R., Fellows, M.R.: Parameterized complexity. Monographs in computer science. Springer, New York (1999)
3. Bodlaender, H.L., Downey, R.G., Fellows, M.R., Hermelin, D.: On problems without polynomial kernels (extended abstract). In: Proc. ICALP. (2008) 563–574
4. Guo, J., Niedermeier, R.: Linear problem kernels for NP-hard problems on planar graphs. In: Proc. 34th ICALP. Number 4596 in LNCS, Springer (2007) 375–386
5. Bodlaender, H.L., Fomin, F.V., Lokshtanov, D., Penninkx, E., Saurabh, S., Thilikos, D.M.: (Meta) Kernelization. CoRR **abs/0904.0727** (2009)
6. Fomin, F., Lokshtanov, D., Saurabh, S., Thilikos, D.M.: Bidimensionality and kernels. In: Proceedings of SODA 2010. (2010)
7. van Leeuwen, E.: Optimization and Approximation on Systems of Geometric Objects. PhD thesis, CWI Amsterdam (2009)
8. Clark, B.N., Colbourn, C.J., Johnson, D.S.: Unit disk graphs. Discrete Mathematics **86** (1990) 165 – 177
9. Marx, D.: Efficient approximation schemes for geometric problems? In: ESA. (2005) 448–459
10. Marx, D.: Parameterized complexity of independence and domination on geometric graphs. In: IWPEC. (2006) 154–165
11. Alber, J., Fiala, J.: Geometric separation and exact solutions for the parameterized independent set problem on disk graphs. J. Algorithms **52** (2004) 134–151
12. Dom, M., Lokshtanov, D., Saurabh, S.: Incompressibility through colors and ids. In: Proceedings 36th ICALP, Berlin, Heidelberg, Springer-Verlag (2009) 378–389
13. Marathe, M., Breu, H., Iii, H.B.H., Ravi, S.S., Rosenkrantz, D.J.: Simple heuristics for unit disk graphs. Networks **25** (1995) 59–68
14. Wiese, A., Kranakis, E.: Local PTAS for independent set and vertex cover in location aware unit disk graphs. In: DCOSS. (2008) 415–431
15. Kratochvl, J., Pergel, M.: Intersection graphs of homothetic polygons. Electronic Notes in Discrete Mathematics **31** (2008) 277 – 280
16. Moser, H.: A problem kernelization for graph packing. In: SOFSEM '09, Berlin, Heidelberg, Springer-Verlag (2009) 401–412
17. Hunt, H.B., Marathe, M.V., Radhakrishnan, V., Ravi, S.S., Rosenkrantz, D.J., Stearns, R.E.: NC-approximation schemes for NP- and PSPACE-hard problems for geometric graphs. Journal of Algorithms **26** (1998) 238 – 274
18. Chan, T.M., Har-Peled, S.: Approximation algorithms for maximum independent set of pseudo-disks. In: Proceedings 25th SCG. (2009) 333–340
19. Efrat, A., Sharir, M., Ziv, A.: Computing the smallest k-enclosing circle and related problems. Comput. Geom. **4** (1994) 119–136
20. de Berg, M., Speckmann, B.: Computational geometry: Fundamental structures. In: Handbook Of Data Structures And Applications. (2004) 62–2

A Proofs

A.1 Proof of Lemma 1

Lemma *Let $G = (X \cup Y, E)$ be a planar bipartite graph. If for all distinct $v, v' \in Y$ it holds that $N_G(v) \not\subseteq N_G(v')$ then $|Y| \leq 5|X|$.*

Proof. Assume the conditions stated in the lemma hold. If there is an isolated vertex in Y then that vertex must be the only member of Y, because its neighborhood is a subset of any other neighborhood. So without loss of generality we may assume that the minimum degree of vertices in Y is at least 1. Now we deal with vertices of degree 1. If there is a vertex $v \in Y$ such that $N_G(v) = \{u\}$, then the graph G' obtained from G by deleting u and v satisfies the same conditions as the original: since any two neighborhoods are incomparable, there can be no other vertex $v' \in Y$ that is a neighbor of u. Therefore we may apply induction to find that $|Y \setminus \{v\}| \leq 5|X \setminus \{u\}|$ and hence that $|Y| \leq 5|X|$.

In the remainder we may assume that all vertices in Y have degree at least two. We obtain a simple bound on the size of Y in two steps: we first consider the set Y_2 of vertices of degree exactly two, and then consider the set Y_3 with vertices of degree at least 3.

Consider the planar graph G_2 obtained from G by taking the subgraph that is induced by $X \cup Y_2$, and replacing every vertex in Y_2 by a direct edge between its two neighbors. It is easy to verify that G_2 is planar. Since no two vertices in Y_2 have the same neighborhood there are no parallel edges in G_2 and therefore G_2 is a simple planar graph such that $|E(G_2)| = |Y_2|$. By Euler's formula the number of edges in a simple planar graph is at most three times the number of vertices, and we find that $|Y_2| = |E(G_2)| \leq 3|X|$.

To bound the size of Y_3 we consider the bipartite graph G_3 that is the subgraph of G induced by $X \cup Y_3$. Another application of Euler's formula shows that $|E| \leq 2n$ in a bipartite simple planar graph on n vertices. Since every vertex in Y_3 has degree at least 3 in G_3 we find $|E(G_3)| \geq 3|Y_3|$. Combining this with the upper bound on the edge count we find $3|Y_3| \leq |E(G_3)| \leq 2(|X| + |Y_3|)$, which implies that $|Y_3| \leq 2|X|$.

Since $|Y| = |Y_2| + |Y_3|$ we find that $|Y| \leq 5|X|$, which concludes the proof. □

A.2 Proof of Lemma 2

Lemma *If v is a vertex in a unit-disk graph $G = (V, E)$ then there must be a clique of size at least $\lceil \deg_G(v)/6 \rceil$ in $G[N_G(v)]$.*

Proof. Up to this point we have considered unit-disk graphs to be intersection graphs of disks of unit radius in the plane, where two vertices are connected by an edge if and only if their disks overlap - this is also known as the *intersection model* for unit-disk graphs. For this proof it is convienient to switch to a different model; by doubling the radius of every disk to 2 we obtain a new model where two vertices are connected by an edge if and only if their disks contain each other's center points - the containment model. Note that in this model two vertices are adjacent if and only if the distance between the center points of the disks is at most two.

Now consider a realization of G by a containment model that specifies a disk center $c_u \in \mathbb{R}^2$ for every vertex $u \in V$. Let v be the vertex for which we wish to prove the theorem and consider the disk $\mathcal{D}(v)$ of radius 2 that is placed at c_v. By definition the set of neighbors of v coincides with the set of vertices whose center point is contained in $\mathcal{D}(v)$. Now partition the disk $\mathcal{D}(v)$ into six equal sectors, such that one sector covers an arc of exactly 60 degrees. Suppose that the centers c_w and c_z of neighbors w and z fall into the same sector. Since w and z are neighbors of u we know that $|c_w - c_u|, |c_z - c_u| \leq 2$. Observe that for any two centers c_w, c_z in the same sector the angle $\angle c_w c_u c_z$ is at most 60 degrees, which implies that $|c_w - c_z| \leq \max\{|c_w - c_u|, |c_z - c_u|\} \leq 2$ by the law of cosines. Therefore there must be an edge between w and z if their centers fall in the same sector; all the centers that fall in the same sector must form a clique. By the pigeonhole principle there is at least one sector that contains $\lceil \deg_G(v)/6 \rceil$ vertices, and the claim follows. □

A.3 Proof of Corollary 1

Corollary *H-*Matching *in unit-disk graphs has a kernel with $O(k)$ vertices for every fixed H.*

Proof. Given an instance $\langle G, k \rangle$ we exhaustively apply the reduction rule to obtain a reduced graph $\langle G', k' \rangle$ with $k' = k$. We use a greedy algorithm to find a maximal H-matching in the reduced instance, which can be done in polynomial time for fixed H. Let k^* be the size of the resulting matching. If $k \leq k^*$ then the answer to the decision problem must be yes and we output a trivial yes-instance. Otherwise we know from Theorem 3 that $|V(G')| \in O(k^*)$ and since $k > k^*$ we find $|V(G')| \in O(k)$; we output the reduced instance $\langle G', k' \rangle$ as the result of the kernelization. □

A.4 Proof of Theorem 4

Theorem *Let $G = (V, E)$ be a unit-disk graph whose vertex set is partitioned into red and blue color classes by $V = R \cup B$ with $R \cap B = \emptyset$. If for all distinct $u, v \in R$ it holds that $N_G(u) \cap B \not\subseteq N_G(v) \cap B$ then $|R| \in O(|B|^2)$.*

Proof. Recall the containment model for unit-disk graphs described in the proof of Lemma 2, and fix a realization of G by a containment model. For each blue vertex vertex $v \in B$ of G we draw a circle of radius 2 around its center. Any vertex whose center lies inside that circle must have an edge to v. If we draw the circles for all blue vertices of G we obtain a subdivision of the plane into regions, also known as an *arrangement* in computational geometry. Consider a face F of the arrangement. If we take two points $p, p' \in F$ then these points are contained in exactly the same disks of $\mathcal{D}(B)$, since they fall within the same circles. Therefore if p and p' would be the centers of red vertices, then these red vertices would have exactly the same set of blue neighbors. But by the assumption that all red vertices have a different set of blue neighbors, we find that any face of the arrangement can contain at most one red vertex center in its interior. For every vertex or edge *of the arrangement* there is at most one red vertex, for the same reason: two red centers on the same vertex or edge of the arrangement implies that they have the same set of blue neighbors. It is well-known in the context of computational geometry that an arrangement of n circles has at most $O(n^2)$ faces, edges and vertices [20]. In this case we have $|B|$ circles and we conclude that the number of red vertices is $O(|B|^2)$. □