Proc. 34th ICALP, 2007

Linear Problem Kernels for NP-Hard Problems on Planar Graphs

Jiong Guo* and Rolf Niedermeier

Institut für Informatik, Friedrich-Schiller-Universität Jena,
Ernst-Abbe-Platz 2, D-07743 Jena, Germany.
{guo,niedermr}@minet.uni-jena.de

Abstract. We develop a generic framework for deriving linear-size problem kernels for NP-hard problems on planar graphs. We demonstrate the usefulness of our framework in several concrete case studies, giving new kernelization results for Connected Vertex Cover, Minimum Edge Dominating Set, Maximum Triangle Packing, and Efficient Dominating Set on planar graphs. On the route to these results, we present effective, problem-specific data reduction rules that are useful in any approach attacking the computational intractability of these problems.

1 Introduction

Data reduction together with problem kernelization has been recognized as one of the primary contributions of parameterized complexity to practical algorithm design [9, 15, 21]. For instance, the NP-hard Vertex Cover problem, where one asks for a set of at most k vertices such that all edges of a given graph have at least one endpoint in this set, has a problem kernel of $2k$ vertices. That is, given a graph G and the parameter k, one can construct in polynomial time a graph G' consisting of only $2k$ vertices and with a new parameter $k' \leq k$ such that (G, k) is a yes-instance iff (G', k') is a yes-instance [20, 8]. In particular, this means that Vertex Cover can be efficiently preprocessed with a guaranteed quality of data reduction—the practical usefulness is confirmed by experimental work [1]. Note that a $2k$-vertex problem kernel is the best one may probably hope for because a $(2-\epsilon)k$-vertex kernel with $\epsilon > 0$ would imply a factor-$(2-\epsilon)$ polynomial-time approximation algorithm for Vertex Cover, solving a long standing open problem. Clearly, a k-vertex problem kernel for Vertex Cover would imply P=NP. That is why so-called linear-size problem kernels (a linear function in the parameter k) usually are considered as the "holy grail" in the field of kernelization and parameterized complexity analysis.

Unfortunately, so far there are not too many problems known with a problem kernel size as small as we have for Vertex Cover. Moreover, strictly speaking, the $2k$-vertex problem kernel for Vertex Cover is not really a linear-size problem kernel because the number of graph edges still may be $O(k^2)$. Apparently, the

* Supported by the Deutsche Forschungsgemeinschaft (DFG), Emmy Noether research group PIAF (fixed-parameter algorithms), NI 369/4.

situation changes when focussing attention on planar graphs where the number of vertices and the number of edges are linearly related. Although most NP-hard graph problems remain NP-hard when restricted to planar graphs, it has been observed that they behave much better in terms of approximability (see [5]) as well as in terms of fixed-parameter tractability (see [4]). In her seminal work, Baker [5] showed that a whole class of problems (including VERTEX COVER, INDEPENDENT SET, DOMINATING SET) possesses polynomial-time approximation schemes (PTAS), all derived from a general framework.

Concerning problem kernelization results on planar graphs where, in a sense, linear-size problem kernels can be seen as the parameterized counterpart of approximation schemes, so far only few isolated results are known [4, 7, 17, 19]. In particular, it has been shown that the DOMINATING SET problem—which is W[2]-complete on general graphs, meaning that there is no hope for a problem kernel at all [9, 21]—has a linear-size problem kernel when restricted to planar graphs [4]. This result goes along with the development of simple but effective data reduction rules whose practical usefulness has been empirically confirmed [2].[1] In this work, "in the spirit of Baker", we develop a general framework that allows for a systematic approach to derive linear-size problem kernels for planar graph problems. In particular, our methodology offers concrete starting-points for developing effective data reduction rules, the central part of any form of problem kernelization. Inspired by the work of Alber et al. [4], which focuses on the DOMINATING SET problem, we show what the common features are that lie at the heart of linear-size problem kernels for planar graph problems. In particular, we provide a concrete route of attack which serves as a tool for developing data reduction rules. Doing so, we provide a number of case studies together with new results, including the problems CONNECTED VERTEX COVER, EDGE DOMINATING SET, MAXIMUM TRIANGLE PACKING, and EFFICIENT DOMINATING SET, all of which are shown to have linear-size problem kernels with concrete upper bounds. Note that, although based on the general framework, all corresponding data reduction rules that had to be newly developed are—of course—problem-specific. The development of these rules still needs novel ideas in each specific case and is far from being routine. Still, our framework offers a guiding star to find them.

Most proofs are deferred to the full version of this paper.

2 Preliminaries

Parameterized algorithmics is a two-dimensional framework for studying the computational complexity of problems [9, 21]. A core tool in the development of fixed-parameter algorithms is polynomial-time preprocessing by *data reduction rules*, often yielding a *reduction to a problem kernel (kernelization)*. Herein, the goal is, given any problem instance x with parameter k, to transform it in polynomial time into a new instance x' with parameter k' such that the size of x' is

[1] Indeed, the data reduction rules can be applied to all sorts of graphs and not only to planar ones—the rules are particularly effective for sparse graphs.

bounded from above by some function only depending on k, $k' \leq k$, and (x, k) is a yes-instance iff (x', k') is a yes-instance. Then, the problem kernel is called to be *linear* if $|x'| = O(k)$. This transformation is accomplished by applying data reduction rules. A data reduction rule is *correct* if the new instance after an application of this rule is a yes-instance iff the original instance is a yes-instance. Throughout this paper, we call a problem instance *reduced* if the corresponding data reduction rules cannot be applied any more.

We only consider *undirected* graphs $G = (V, E)$, where V is the set of vertices and E is the set of edges. We use n to denote the number of vertices and m to denote the number of edges of a given graph. The *neighborhood* $N(v)$ of a vertex $v \in V$ is the set of vertices that are adjacent to v. The *degree* of a vertex v is the size of $N(v)$. We use $N[v]$ to denote the *closed* neighborhood of v, that is, $N[v] := N(v) \cup \{v\}$. For a set of vertices $V' \subseteq V$, the *induced subgraph* $G[V']$ is the graph with the vertex set V' and the edge set $\{\{v, w\} \in E \mid v, w \in V'\}$. A subset I of vertices is called an *independent set* if $G[I]$ has no edge. We implicitly assume that all paths that we deal with here are *simple*, that is, every vertex is contained at most once in a path. The length of a path is defined as the number of edges used by the path. The distance $d(u, v)$ between two vertices u, v is the length of a shortest path between u, v. The distance $d(e, w)$ between an edge $e = \{u, v\}$ and a vertex w is the minimum of $d(u, w)$ and $d(v, w)$. If a graph can be drawn in the plane without edge crossings then it is called a *planar graph*. A *plane graph* is a planar graph with a fixed embedding in the plane. Throughout this paper, we assume that we are working with an arbitrary but fixed embedding of G in the plane; whenever this embedding is of relevance, we refer to G as being *plane* instead of planar.

3 General Framework

In this section, we describe a general framework for systematically deriving linear problem kernels for NP-hard problems on planar graphs. Although in this (single) case not improving on previous results, for reason of simplicity, we use VERTEX COVER as a running example. The problem is, given a graph $G = (V, E)$ and $k \geq 0$, to find a subset $C \subseteq V$ of at most k vertices such that every edge has at least one endpoint in C. The remainder of this section is structured by exhibiting the four basic components of our methodology.

Component 1: Problem-specific distance property. The problems amenable to our framework have to admit a *distance property* defined as follows:

Definition 1. *A graph problem on input $G = (V, E)$ is said to admit a* distance property *with constants c_V and c_E if, for every solution set S with the vertex set $V(S)$, it holds that, for every vertex $u \in V$, there exists a vertex $v \in V(S)$ with $d(u, v) \leq c_V$, and, for every edge $e \in E$, there exists a vertex $v \in V(S)$ with $d(e, v) \leq c_E$.*

Note that $c_V - 1 \leq c_E \leq c_V$. The distance property is the only prerequisite for applying our framework to a specific graph problem.
Example: The distance property is valid for VERTEX COVER with $c_V = 1$ and $c_E = 0$, since every edge of E has to be incident to a covering vertex.

Component 2: Region decomposition. We divide the vertices not in $V(S)$ into two categories based on whether they lie in the vicinity of either at least two vertices of $V(S)$ or only one vertex of $V(S)$. The former vertices will build so-called *regions* leading to a decomposition of the planar graph.

Definition 2. *A* region $R(u, v)$ *between two distinct vertices* $u, v \in V(S)$ *is a closed subset of the plane with the following properties:*

1. *The* boundary *of* $R(u, v)$ *is formed by two length-at-most-*$(c_V + c_E + 1)$ *paths between* u *and* v*. (These two paths do not need to be disjoint or simple.)*
2. *All vertices which lie on the boundary or strictly inside of the region* $R(u, v)$ *have distance at most* c_V *to at least one of the vertices* u *and* v *and all edges whose both endpoints lie on the boundary or strictly inside of the region* $R(u, v)$ *have distance at most* c_E *to at least one of the vertices* u *and* v*.*
3. *With the exception of* u *and* v*, none of the vertices which lie inside of the region* $R(u, v)$ *are from* $V(S)$*.*

The vertices u *and* v *are called the* anchor vertices *of* $R(u, v)$*. A vertex is said to* lie inside of $R(u, v)$ *if it is either a boundary vertex of* $R(u, v)$ *or if it lies strictly inside of* $R(u, v)$*. We use* $V(R(u, v))$ *to denote the set of vertices that lie inside of a region* $R(u, v)$*.*

Using Definition 2, the graph can be partitioned by a so-called *region decomposition.*

Definition 3. *An* S-region decomposition *of a graph is a set* $\mathcal{R}$ *of regions such that there is no vertex that lies strictly inside of more than one region from* $\mathcal{R}$ *(the boundaries of regions may touch each other, however).*

For an S*-region decomposition* $\mathcal{R}$*, let* $V(\mathcal{R}) := \bigcup_{R \in \mathcal{R}} V(R)$*. An* S*-region decomposition* $\mathcal{R}$ *is called* maximal *if there is no region* $R \notin \mathcal{R}$ *such that* $\mathcal{R}' := \mathcal{R} \cup \{R\}$ *is an* S*-region decomposition with* $V(\mathcal{R}) \subsetneq V(\mathcal{R}')$*.*

As a basis for linear kernelization results, our framework makes use of the fact that the number of regions in a maximal region decomposition $\mathcal{R}$ for a given solution S can be upper-bounded by $c_V \cdot (3|V(S)| - 6)$. This generalizes a result of Alber et al. [4].

Lemma 1. *Let* P *be a graph problem admitting the distance property with* c_V *and* c_E *and let* S *be a solution of* P *on a plane graph* $G = (V, E)$*. Then, there is a maximal* S*-region decomposition* $\mathcal{R}$ *for the input graph* G *that consists of at most* $c_V \cdot (3|V(S)| - 6)$ *regions.*

Example: Since Vertex Cover admits the distance property with $c_V = 1$ and $c_E = 0$, the maximal region decomposition consists of regions with boundary paths of length at most two. By Lemma 1, we know that for a Vertex Cover solution C of size at most k we have at most $3k - 6$ regions in a maximal region decomposition.

Component 3: Local neighborhoods for data reduction rule design. The core algorithmic part of our methodology is based on the following two definitions that serve for developing data reduction rules yielding linear-size problem kernels.

Definition 4. *Given a problem admitting the distance property with constants c_V and c_E, the* private neighborhood $N_p(u)$ *of a vertex u consists of the vertices that have distance at most c_V to u, that are not adjacent to the vertices with distance at least $c_V + 1$ to u, and that are not incident to any edge with distance more than c_E to u.*

Example: For Vertex Cover, the private neighborhood $N_p(u)$ of a vertex u consists only of the degree-one vertices from $N(u)$. Therefore, the corresponding private neighborhood rule deals with degree-one vertices:
Private neighborhood rule for Vertex Cover**:** If $N_p(u) \neq \emptyset$, then add u to C and remove $N_p(u)$ from the graph. Decrease the parameter k by one.

Definition 5. *Given a problem admitting the distance property with constants c_V and c_E, the* joint private neighborhood $N_p(u, v)$ *of two vertices $u, v \in V$ consists of the vertices that have distance at most c_V to u or v, that are not adjacent to the vertices which have distance at least $c_V + 1$ to both u and v, and that are not incident to any edge with distance more than c_E to both u and v.*

Example: In Vertex Cover, the joint private neighborhood of u and v consists of their common neighbors and their degree-one neighbors. Since the private neighborhood rule deals with the degree-one neighbors, for the corresponding data reduction rule we only consider the common neighbors of u and v:
Joint private neighborhood rule for Vertex Cover**:** If two vertices u, v have at least two common degree-two neighbors, then add u and v to C and remove u, v and their common degree-two neighbors from the graph. Decrease the parameter k by two.

Generally speaking, if $N_p(v)$ of a vertex v (or $N_p(u, v)$ of vertices u and v) contains too many vertices, then any solution of a minimization problem has to include v (or at least one of u and v). For maximization problems, we can conclude that including only v from $N_p(v)$ (or only u and v from $N_p(u, v)$) cannot lead to a solution. This provides a useful argument for showing upper bounds on the region sizes in the problem kernel size analysis.

Component 4: Mathematical analysis of problem kernel size. Having derived problem-specific data reduction rules, the next step in our method is to prove that there are only constantly many vertices inside of a region. Together with

Lemma 1, this implies the upper bound $O(|V(S)|)$ for all vertices inside of all regions of the reduced graph.

Example: For Vertex Cover, the constant size for every region follows almost directly from the given two rules:

Lemma 2. *Given a vertex cover C of a reduced planar graph $G = (V, E)$, every region of a maximal region decomposition contains at most three vertices which are not from C.*

Proof. Consider a maximal region decomposition and let R denote an arbitrary region with boundary paths of length at most two. Let u, v be the two vertices in $V(R)$ that are from C. Clearly, at most two vertices on the boundary paths can be from $V \setminus C$. We claim that there is at most one vertex lying strictly inside of R. To show this, suppose that there are two vertices x, y strictly inside of R. Then, by the definition of regions, $x, y \in N(u) \cup N(v)$. Due to the private neighborhood rule, each of x and y has at least two neighbors in $V(R) \cap C$. Since u and v are the only C-vertices in R, the vertices x, y are common neighbors of u and v and have degree two. This implies that the joint private neighborhood rule can be applied, a contradiction to the fact that G is reduced. Therefore, at most one vertex lies strictly inside of R and the lemma follows. $\square$

Note that, if, in addition to the above two rules, the "folding" rule introduced by Chen et al. [8] is applied, one can show that every region contains at most two vertices from $V \setminus C$.

To complete the proof for a linear-size problem kernel, our method requires to upper-bound the number of vertices not contained in any region. To do so, the private neighborhood rule is crucial.

Example: In the case of Vertex Cover, the private neighborhood rule guarantees that there is no vertex lying outside of the regions of a maximal region decomposition:

Lemma 3. *Let $\mathcal{R}$ be a maximal region decomposition of a reduced planar graph for* Vertex Cover*. Then, there is no vertex lying outside of the regions in $\mathcal{R}$.*

Proof. Suppose that there is such a vertex x. It cannot be a degree-one vertex and it cannot be adjacent to a vertex not from C. Thus, $N(x) \subseteq C$. Then, we can arbitrarily pick two from x's neighbors and have a region R that is a path consisting of x and these two neighbors. By adding R to $\mathcal{R}$ we get a new region decomposition $\mathcal{R}'$ with $V(\mathcal{R}) \subsetneq V(\mathcal{R}')$, a contradiction to the fact that $\mathcal{R}$ is maximal. $\square$

Finally, to give an overall kernel size bound, we only need to add up the number of vertices inside of regions and the number of vertices outside of regions.

Example: With the two upper bounds given in Lemmas 2 and 3, we arrive at our linear kernelization result for Vertex Cover on planar graphs:

Proposition 1. Vertex Cover *on planar graphs admits a $10k$-vertex problem kernel.*

Proof. By Lemma 1, there are at most $3k - 6$ regions in a maximal region decomposition. Together with Lemma 2, there can be at most $9k - 18$ vertices from $V \setminus C$ lying inside of regions. By Lemma 3, no vertex can be outside of regions. Thus, altogether, we have $10k - 18$ vertices in the reduced graph. □

4 Case Studies

Now, we exhibit the versatility of our general methodology.

Connected Vertex Cover. Given a graph $G = (V, E)$ and a non-negative integer k, the CONNECTED VERTEX COVER problem asks for a set C of at most k vertices such that $G[C]$ is connected and C is a vertex cover of G. This problem is NP-complete on planar graphs [13]. Until now only an exponential-size kernel in general graphs is known [16].

Since a connected vertex cover is also a vertex cover, the distance property holds for this problem with $c_V = 1$ and $c_E = 0$. Thus, the regions in a maximal region decomposition for CONNECTED VERTEX COVER have also boundary paths of length at most two and we have at most $3k - 6$ regions in a maximal region decomposition. Moreover, the private neighborhood and the joint private neighborhood are defined in the same way as for VERTEX COVER.

The data reduction rules are similar to the ones for VERTEX COVER. However, to guarantee the resulting vertex cover being connected, we use gadgets:
Private neighborhood rule: If a vertex has more than one degree-one neighbor, then except for one remove all of these neighbors.
Joint private neighborhood rule: If two vertices have more than two common degree-two neighbors, then remove all of these neighbors except for two.

Theorem 1. CONNECTED VERTEX COVER *on planar graphs admits a $14k$-vertex problem kernel.*

Proof. First, consider a region R in a maximal C-region decomposition $\mathcal{R}$ for a connected vertex cover C with $|C| \leq k$. Note that strictly inside of a region there cannot be vertices of degree more than two because this would imply uncovered edges. Due to the joint private neighborhood rule, there can be at most two degree-two vertices lying strictly inside of R. Since there can be at most two vertices from $V \setminus C$ lying on the boundary of a region, each region can contain at most four vertices from $V \setminus C$. Since the graph is reduced with respect to the private neighborhood rule, each vertex in C can have at most one degree-one neighbor not lying in a region of $\mathcal{R}$. Therefore, we altogether have at most $4 \cdot (3k - 6)$ vertices from $V \setminus C$ which lie inside of regions and at most k vertices outside of regions. Together with $|C| \leq k$, the size bound follows. □

Edge Domination Set. Given a graph $G = (V, E)$ and a non-negative integer k, the EDGE DOMINATING SET problem asks for a set E' of at most k edges such that all edges in E share at least one endpoint with some edge in E'. This problem is NP-complete on planar graphs [14]. Based on its equivalence to the MINIMUM

Maximal Matching problem, a problem kernel with $O(k^2)$ vertices for general graphs has been derived [22]. It is easy to observe that Edge Dominating Set has the same distance parameters as Vertex Cover. Therefore, the same two data reduction rules for Connected Vertex Cover apply.

Theorem 2. Edge Dominating Set *on planar graphs admits a* $14k$*-vertex problem kernel.*

Maximum Triangle Packing. Given a graph $G = (V, E)$ and a non-negative integer k, the Maximum Triangle Packing problem asks for a set P of at least k vertex-disjoint triangles in G. The set P is called a *triangle packing* of G. This problem is NP-complete on planar graphs [14]. A problem kernel with $O(k^3)$ vertices is known for general graphs [10].

At first glance, Maximum Triangle Packing does not admit the required distance property. However, the following data reduction rule can be applied.
Cleaning rule: Remove all vertices and edges that are not in a triangle.

In an instance where the cleaning rule does not apply, every vertex and every edge has a distance at most $c_V = 1$ and $c_E = 1$, respectively, to some vertex occurring in a triangle packing that cannot be extended by a triangle. Then, the regions in a maximal region decomposition for Maximum Triangle Packing have boundary paths of length at most three.

Consider the private neighborhood $N_p(u)$ of a vertex u. According to Definition 4, all vertices $v \in N_p(u)$ have to satisfy $N[v] \subseteq N[u]$. We apply the following data reduction rule dealing with private neighborhoods.
Private neighborhood rule: If a vertex u has two neighbors v, w that form a triangle with u but are not involved in any other triangles that do not contain u, then remove u, v, w and decrease the parameter k by one.

Next, we consider the joint private neighborhood $N_p(u, v)$ of two vertices u, v. According to Definition 5, every vertex $x \in N_p(u, v)$ has to satisfy $N[x] \subseteq N[u] \cup N[v]$.
Joint private neighborhood rule: If two vertices u, v have more than two common neighbors, then consider the following cases.

- Case 1: If u and v have two common neighbors w_1 and w_2 such that w_1 has degree two and w_2 is only contained in triangles that also contain u or v, then remove w_1.
- Case 2: If u and v have three common neighbors w_1, w_2, w_3 such that $N(w_1) = \{u, v, w_2\}$ and $N(w_2) = \{u, v, w_1, w_3\}$, then remove edge $\{w_2, w_3\}$.
- Case 3: If there are four vertices w_1, w_2, w_3, w_4 such that u, w_1, w_2 form a triangle, v, w_3, w_4 form another one, and there is no other triangle that contains one of w_1, w_2, w_3, w_4 but none of u, v, then remove u, v, w_1, w_2, w_3, w_4 and decrease the parameter k by two.

Note that, after the application of the cleaning rule, every vertex has to be in a triangle. Therefore, in Case 1, there has to be an edge between u and v. The three cases of the joint private neighborhood rule are illustrated in Fig. 1.

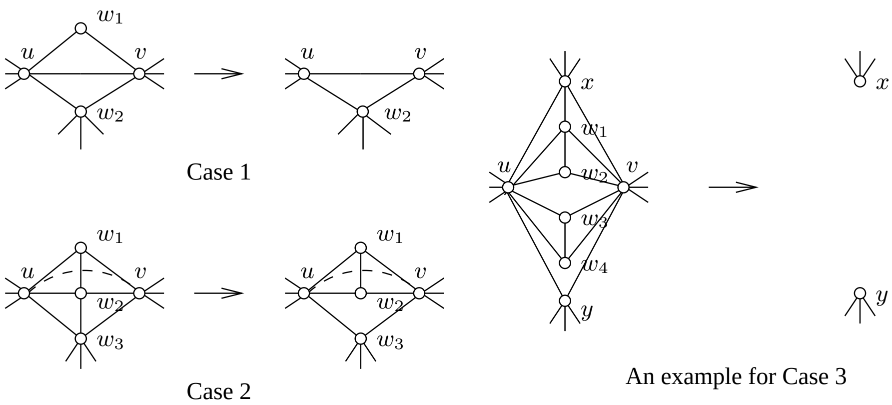

Fig. 1. Illustration of the three cases of the joint private neighborhood rule for MAXIMUM TRIANGLE PACKING. A dashed line means a possibly existing edge.

In order to give a linear-size problem kernel, we need only the first and the third case of the joint private neighborhood rule. However, including the second case allows us to give a better bound on the maximum size of regions in a maximal region decomposition of the reduced graphs as stated in the following, allowing for a smaller upper bound on the problem kernel size.

To prove that there is only a constant number of vertices inside of each region, graph structures that we call *diamonds*[2] are of great importance.

Definition 6. *Let u and v be two vertices in a plane graph G. A* diamond $D(u, v)$ *is a closed area of the plane that is bounded by two length-2 paths between u and v such that every vertex that lies inside this area is a neighbor of both u and v. If i vertices lie strictly inside a diamond, then it is said to have* $(i + 1)$ facets.

Lemma 4. *In a reduced planar graph, a diamond can have at most five facets.*

Now, we state upper bounds on the number of vertices in a region and on the number of vertices outside of all regions.

Lemma 5. *Consider a planar graph $G = (V, E)$ for which any triangle packing contains at most k triangles. If G is reduced, then, in a maximal region decomposition of G,*

1. *every region can contain at most* 71 *vertices, and*
2. *there are less than* $108k$ *vertices lying outside of regions.*

Theorem 3. MAXIMUM TRIANGLE PACKING *on planar graphs admits a* $732k$*-vertex problem kernel.*

[2] Note that standard graph theory uses the term "diamond" to denote a 4-cycle with exactly one chord. We abuse this term here for obvious reasons. We remark that diamonds also played a decisive role in proving a linear-size problem kernel for DOMINATING SET on planar graphs [4].

Efficient Dominating Set. Given a graph $G = (V, E)$ and a non-negative integer k, the Efficient Dominating Set problem is to decide whether there exists an independent set I such that every vertex in $V \setminus I$ has exactly one neighbor in I. A solution set of this problem is called an *efficient dominating set.* This problem is NP-complete on planar graphs of maximum degree three [11]. In the literature, Efficient Dominating Set also appears under the names Perfect Code, Independent Perfect Dominating set, and Perfect Dominating Set. Lu and Tang [18] provided an overview of complexity results for Efficient Dominating Set.

Bange et al. [6] showed that if a graph G has an efficient dominating set, then all efficient dominating sets of G have the same cardinality, and this is the same as the domination number of G, where the domination number is the cardinality of a minimum dominating set of G. Hence, the parameterized version of Efficient Dominating Set additionally has a non-negative integer k as input and asks for an efficient dominating set of size *exactly* k. Efficient Dominating Set is W[1]-hard in general graphs [9]. To our knowledge, there is no kernelization result known for this problem on planar graphs. Note that the linear-size problem kernel for Dominating Set on planar graphs does not imply a linear-size problem kernel for Efficient Dominating Set on planar graphs since the data reduction rules applied by Alber et al. [4] for deriving the problem kernel apparently do not work for Efficient Dominating Set.

Since every efficient dominating set also is a dominating set, the distance property holds for Efficient Dominating Set with $c_V = 1$ and $c_E = 1$. The boundary paths of the regions in a maximal region decomposition have length at most three. Note that Efficient Dominating Set has the same distance parameters c_V and c_E as Maximum Triangle Packing. Therefore, the private neighborhood and the joint private neighborhood are the same in both cases.

In the following we describe two data reduction rules for Efficient Dominating Set. Actually, the reduction rules apply to a more general setting where we are additionally given a subset of vertices $F \subseteq V$ which may not be added to the efficient dominating set. In the following rules, whenever we would be forced to add a vertex in F to a solution set I, we report that the given instance has no efficient dominating set.

Private neighborhood rule: Consider the following two cases for a vertex v with $N_p(u) \neq \emptyset$:

- Case 1. If there is no vertex $v \in N(u)$ such that v dominates all vertices in $N_p(u)$, then add u to the efficient dominating set I, remove all vertices in $N[u]$ from the graph, add to F all vertices which are not in $N[u]$ but adjacent to some vertex in $N(u)$, and decrease the parameter k by one.
- Case 2. If there is exactly one vertex $v \in N(u)$ dominating all vertices in $N_p(u)$, then remove all vertices in $N_p(u) \setminus \{v\}$ and add two new non-adjacent vertices x, y and connect them to both u and v.

Joint private neighborhood rule: Consider the following two cases for two vertices u, v with $N_p(u, v) \neq \emptyset$:

– Case 1. If u, v have two common neighbors x, y such that $\{x, y\} \notin E$, $N(x) \subsetneq N(u) \cap N(v)$, $N(y) \subsetneq N(u) \cap N(v)$, and $N(x) \cap N(y) = \{u, v\}$, then remove $(N(u) \cap N(v)) \setminus \{x, y\}$ and add those vertices to F that are not in $N(u) \cap N(v)$ but adjacent to some vertex in $N(u) \cap N(v)$.
– Case 2. Enumerate all subsets of $N[u, v] := N[u] \cup N[v]$ that induce independent sets of size at most two and whose vertices are adjacent to all vertices in $N_p(u, v)$. If there is a vertex w occurring in all of these sets, then add w to the efficient dominating set, remove $N[w]$ from the graph, add to F all vertices which are not in $N[w]$ but adjacent to some vertex in $N(w)$, and decrease the parameter k by one.

Lemma 6. *(1) In a reduced planar graph, a diamond can have at most four facets.*

(2) In a maximal region decomposition of a reduced planar graph, every region contains at most 28 *vertices.*

(3) In a maximal region decomposition of a reduced planar graph, there are at most $5k$ *vertices lying outside of regions.*

Theorem 4. EFFICIENT DOMINATING SET *on planar graphs admits a* $84k$*-vertex problem kernel.*

5 Outlook

There are numerous avenues for future research. First, it is promising to look into further improving the constant factors of our kernel bounds, similarly as Chen et al. [7] did for DOMINATING SET [4]. Second, again referring to Chen et al. [7] and the lower bounds on (linear) kernel sizes derived there for VERTEX COVER and DOMINATING SET on planar graphs, based on our framework, similar lower bound investigations may now be undertaken for other problems. Third, it appears natural to further extend the list of concrete problem kernel bounds for problems we did not touch here—further domination problems being obvious candidates. Observe that we can extend the list of linear problem kernel results using other further problems studied by Baker [5]. Fourth, as the linear-size problem kernel results for DOMINATING SET on planar graphs have been extended to graphs of bounded genus [12], it is tempting to generalize our whole framework in the same style. Fifth, for DOMINATING SET, a generic set of data reduction rules has been designed [3]—analogous studies now may be fruitful for all problems fitting into our framework.

References

1. F. N. Abu-Khzam, R. L. Collins, M. R. Fellows, M. A. Langston, W. H. Suters, and C. T. Symons. Kernelization algorithms for the Vertex Cover problem: theory and experiments. In *Proc. 6th ACM-SIAM ALENEX*, pages 62–69. ACM-SIAM, 2004.

2. J. Alber, N. Betzler, and R. Niedermeier. Experiments on data reduction for optimal domination in networks. *Annals of Operations Research*, 146(1):105–117, 2006.
3. J. Alber, B. Dorn, and R. Niedermeier. A general data reduction scheme for domination in graphs. In *Proc. 32nd SOFSEM*, volume 3831 of *LNCS*, pages 137–147. Springer, 2006.
4. J. Alber, M. R. Fellows, and R. Niedermeier. Polynomial time data reduction for Dominating Set. *Journal of the ACM*, 51(3):363–384, 2004.
5. B. S. Baker. Approximation algorithms for NP-complete problems on planar graphs. *Journal of the ACM*, 41(1):153–180, 1994.
6. D. W. Bange, A. E. Barkauskas, and P. J. Slater. Efficient dominating sets in graphs. In *Proc. 3rd Conference on Discrete Mathematics*, pages 189–199. SIAM, 1988.
7. J. Chen, H. Fernau, I. A. Kanj, and G. Xia. Parametric duality and kernelization: Lower bounds and upper bounds on kernel size. In *Proc. 22nd STACS*, volume 3404 of *LNCS*, pages 269–280. Springer, 2005.
8. J. Chen, I. A. Kanj, and W. Jia. Vertex Cover: Further observations and further improvements. *Journal of Algorithms*, 41:280–301, 2001.
9. R. G. Downey and M. R. Fellows. *Parameterized Complexity*. Springer, 1999.
10. M. R. Fellows, P. Heggernes, F. A. Rosamond, C. Sloper, and J. A. Telle. Finding k disjoint triangles in an arbitrary graph. In *Proc. 30th WG*, volume 3353 of *LNCS*, pages 235–244. Springer, 2004.
11. M. R. Fellows and M. Hoover. Perfect domination. *Australian Journal of Combinatorics*, 3:141–150, 1991.
12. F. V. Fomin and D. M. Thilikos. Fast parameterized algorithms for graphs on surfaces: Linear kernel and exponential speed-up. In *Proc. 31st ICALP*, volume 3142 of *LNCS*, pages 581–592. Springer, 2004.
13. M. R. Garey and D. S. Johnson. The rectilinear Steiner tree problem is NP-complete. *SIAM Journal on Applied Mathematics*, 32:826–834, 1977.
14. M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
15. J. Guo and R. Niedermeier. Invitation to data reduction and problem kernelization. *SIGACT News*, 38(1):31–45, 2007.
16. J. Guo, R. Niedermeier, and S. Wernicke. Parameterized complexity of generalized Vertex Cover problems. In *Proc. 9th WADS*, volume 3608 of *LNCS*, pages 36–48. Springer, 2005. Long version to appear under the title "Parameterized complexity of Vertex Cover variants" in *Theory of Computing Systems*.
17. J. Guo, R. Niedermeier, and S. Wernicke. Fixed-parameter tractability results for full-degree spanning tree and its dual. In *Proc. 2nd IWPEC*, volume 4196 of *LNCS*, pages 203–214. Springer, 2006.
18. C. L. Lu and C. Y. Tang. Weighted efficient domination problem on some perfect graphs. *Discrete Applied Mathematics*, 117:163–182, 2002.
19. H. Moser and S. Sikdar. The parameterized complexity of the induced matching problem in planar graphs. In *Proc. 1st International Frontiers of Algorithmics Workshop (FAW 2007)*, LNCS. Springer, 2007.
20. G. L. Nemhauser and L. E. Trotter. Vertex packing: structural properties and algorithms. *Mathematical Programming*, 8:232–248, 1975.
21. R. Niedermeier. *Invitation to Fixed-Parameter Algorithms*. Oxford University Press, 2006.
22. E. Prieto. *Systematic Kernelization in FPT Algorithm Design*. PhD thesis, Department of Computer Science, University of Newcastle, Australia, 2005.