2014 IEEE Annual Symposium on Foundations of Computer Science

Network Sparsification for Steiner Problems on Planar and Bounded-Genus Graphs

Marcin Pilipczuk*, Michał Pilipczuk†, Piotr Sankowski‡ and Erik Jan van Leeuwen§
*Department of Informatics, University of Bergen, Norway, malcin@mimuw.edu.pl
†Department of Informatics, University of Bergen, Norway, michal.pilipczuk@ii.uib.no
‡Institute of Informatics, University of Warsaw, Poland, sank@mimuw.edu.pl
§Max-Planck Institut für Informatik, Saarbrücken, Germany, erikjan@mpi-inf.mpg.de

***Abstract*—We propose polynomial-time algorithms that sparsify planar and bounded-genus graphs while preserving optimal or near-optimal solutions to Steiner problems. Our main contribution is a polynomial-time algorithm that, given an unweighted graph G embedded on a surface of genus g and a designated face f bounded by a simple cycle of length k, uncovers a set $F \subseteq E(G)$ of size polynomial in g and k that contains an optimal Steiner tree for *any* set of terminals that is a subset of the vertices of f.**

We apply this general theorem to prove that:

- **given an unweighted graph G embedded on a surface of genus g and a terminal set $S \subseteq V(G)$, one can in polynomial time find a set $F \subseteq E(G)$ that contains an optimal Steiner tree T for S and that has size polynomial in g and $|E(T)|$;**
- **an analogous result holds for an optimal Steiner forest for a set $\mathcal{S}$ of terminal pairs;**
- **given an unweighted planar graph G and a terminal set $S \subseteq V(G)$, one can in polynomial time find a set $F \subseteq E(G)$ that contains an optimal (edge) multiway cut C separating S (i.e., a cutset that intersects any path with endpoints in different terminals from S) and has size polynomial in $|C|$.**

In the language of parameterized complexity, these results imply the first polynomial kernels for STEINER TREE and STEINER FOREST on planar and bounded-genus graphs (parameterized by the size of the tree and forest, respectively) and for (EDGE) MULTIWAY CUT on planar graphs (parameterized by the size of the cutset). STEINER TREE and similar "subset" problems were identified in [1] as important to the quest to widen the reach of the theory of bidimensionality [2, 3] . Therefore, our results can be seen as a leap forward to achieve this broader goal.

Additionally, we obtain a weighted variant of our main contribution: a polynomial-time algorithm that, given an edge-weighted planar graph G, a designated face f bounded by a simple cycle of weight $w(f)$, and an accuracy parameter $\varepsilon > 0$, uncovers a set $F \subseteq E(G)$ of total weight at most $\mathrm{poly}(\varepsilon^{-1})w(f)$ that, for any set of terminal pairs that lie on f, contains a Steiner forest within additive error $\varepsilon w(f)$ from the optimal Steiner forest. This result deepens the understanding of the recent framework of approximation schemes for network design problems on planar graphs ([4, 5] , and later works) by explaining the structure of the solution space within a brick of the so-called *mortar graph* – the central notion of this framework.

***Keywords*-Planar graphs, Steiner Tree, Steiner Forest, (Edge) Multiway Cut, polynomial kernel, parameterized algorithms.**

The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 267959 (Marcin Pilipczuk, Michał Pilipczuk), ERC Grant Agreement n. 259515 (Marcin Pilipczuk, Piotr Sankowski) and Foundation for Polish Science (Marcin Pilipczuk, Piotr Sankowski).

I. INTRODUCTION

Preprocessing algorithms seek out and remove chunks of instances of hard problems that are irrelevant or easy to resolve. The strongest preprocessing algorithms reduce instances to the point that even an exponential-time brute-force algorithm can solve the remaining instance within limited time. The power of many preprocessing algorithms can be explained through the relatively recent framework of kernelization [6, 7]. In this framework, each problem instance I has an associated parameter $k(I)$, often the desired or optimal size of a solution to the instance. Then a *kernel* is a polynomial-time algorithm that preprocesses the instance so that its size shrinks to at most $g(k(I))$, for some computable function g. If g is a polynomial, then we call it a *polynomial kernel*.

The ability to measure the strength of a kernel through the function g has led to a concerted research effort to determine, for each problem, the function g of smallest order that can be attained by a kernel for it. Initial insight into this function, in particular a proof of its existence, is usually given by a *parameterized algorithm*: an algorithm that solves an instance I in time $g(k(I)) \cdot |I|^{O(1)}$. Such an algorithm implies a kernel with the same function g, while, if the considered problem is decidable, then any kernel immediately gives a parameterized algorithm as well [6, 7]. However, if the problem is NP-hard, then this approach can only yield a kernel of superpolynomial size, unless P=NP. Therefore, different insights are needed to find the function g of smallest order, and in particular to find a polynomial kernel. This fact, combined with the discovery that for many problems the existence of a polynomial kernel implies a collapse in the polynomial hierarchy [8]–[10], has recently led to a spike in research on polynomial kernels.

A focal point of research into polynomial kernels are problems on planar graphs. Many problems that on general graphs have no polynomial kernel or even no kernel at all, possess a polynomial kernel on planar graphs. The existence of almost all of these polynomial kernels can be explained from the theory of bidimensionality [2, 3, 11]. The core assumption behind this theory is that the considered problem is *bidimensional*: informally speaking, the solution to an instance must be dense in the input graph. However, this assumption clearly fails for a lot of problems, which has led to gaps in our understanding of the power of preprocessing algorithms for planar graphs. In their survey, Demaine and Hajiaghayi [1, 12] pointed out 'subset' problems, in particular STEINER TREE, as an important research goal in the quest to generalize the theory of bidimensionality.

In this paper, we pick up this line of research and positively resolve the question to the existence of a polynomial kernel on planar graphs for three well-known 'subset' problems: STEINER TREE, STEINER FOREST, and MULTIWAY CUT. We remark that the theory of bidimensionality does not apply to any of these three problems, and that for the first two problems a polynomial kernel on general graphs is unlikely to exist [13] and for the third the existence of a polynomial kernel on general graphs is a major open problem [14, 15]. All kernelization results in this paper are a consequence of a single, generic sparsification algorithm for Steiner trees in planar graphs, which is of independent interest. This sparsification algorithm extends to edge-weighted planar graphs, and we demonstrate its impact on approximation algorithms for problems on planar graphs, in particular on the EPTAS for STEINER TREE on planar graphs [5].

A. Results

We present an overview of the three major results that make up this paper. First, we describe the generic sparsification algorithm for Steiner trees in planar graphs. Second, we show how this sparsification algorithm powers the kernelization results in this paper. Third, we exhibit the extension of the sparsification algorithm to edge-weighted planar graphs, and its implications for approximation algorithms on planar graphs.

The Main Theorem. In our main contribution, we characterize the behavior of Steiner trees in bricks. In our work, a *brick* is simply a connected plane graph B with one designated face formed by a simple cycle ∂B, which w.l.o.g. is the outer (infinite) face of the plane drawing of B, and called the *perimeter of* B. Recall that a *Steiner tree* of a graph G is a tree in G that contains a given set $S \subseteq V(G)$ (called *terminals*). We also say that the Steiner tree *connects* S. In the unweighted setting, a Steiner tree T that connects S is *optimal* if every Steiner tree that connects S has at least as many edges as T. We apply our characterization of Steiner trees in bricks to obtain the following sparsification algorithm:

Theorem I.1 (Main Theorem)**.** *Let B be a brick. Then one can find in $\mathcal{O}(|\partial B|^{142} \cdot |V(B)|)$ time a subgraph H of B such that*

(i) $\partial B \subseteq H$,

(ii) $|E(H)| = \mathcal{O}(|\partial B|^{142})$, *and*

(iii) for every set $S \subseteq V(\partial B)$, H contains some optimal Steiner tree in B that connects S.

The result of Theorem I.1 is stronger than just a polynomial kernel, because the graph H contains an optimal Steiner tree for *any* terminal set that is a subset of the brick's perimeter. We also emphasise that the purely combinatorial (non-algorithmic) statement of Theorem I.1, which asserts the existence of a subgraph H that has property (iii) and polynomial size, is nontrivial and, in our opinion, interesting on its own. A naive construction of a subgraph H that has property (iii) would mark an optimal Steiner tree for each set $S \subseteq V(\partial B)$. Combined with the observation that any optimal Steiner tree of a set $S \subseteq V(\partial B)$ has size at most $|\partial B|$ (as ∂B is a Steiner tree that connects S), we obtain a bound on the size of H of $|\partial B| \cdot 2^{|\partial B|}$. The polynomial bound of Theorem I.1 presents a significant improvement over this naive bound.

We give a detailed overview of the proof of Theorem I.1 in Section II. A full version is available on the arXiv [16]. There we also prove an analogue of Theorem I.1 for graphs of bounded genus, with a polynomial dependence on the genus in the size bound.

The approach that we take in this paper is very different from previous approaches to tackle problems on planar graphs or on bricks. In particular, our ideas are disjoint from those developed in both an EPTAS [5] and a subexponential-time parameterized algorithm [17] for PLANAR STEINER TREE. In those works, a brick was cut into so-called strips and then each strip was cut with a 'perpendicular column'. Here, we take a novel approach and start our construction by using an optimal Steiner tree to recursively decompose the brick. If no optimal Steiner tree decomposes the brick in a manner suitable for recursion, then we show that the brick has a very particular structure, which we can exploit in a different decomposition argument. The analysis of this case in fact constitutes the main part of our paper (we refer to Section II for details). We also stress that we do not employ any techniques used in the theory of bidimensionality. In particular, we do not use any tools from Graph Minors theory, such as the Excluded Grid Theorem [18, 19] — the engine of the theory of bidimensionality.

Applications of Theorem I.1. We give three applications of Theorem I.1. For each application, we state the result and its significance. Proof sketches can be found in Section III.

For the first application of Theorem I.1, we consider STEINER TREE. For this problem, a polynomial kernel on general graphs implies a collapse of the polynomial hierarchy [13]. At the same time, the core assumption of bidimensionality theory fails, and whether a polynomial kernel exists for STEINER TREE on planar graphs was hitherto unknown. Using Theorem I.1, we can resolve the existence of a polynomial kernel for STEINER TREE on planar graphs.

Theorem I.2. *Given a* PLANAR STEINER TREE *instance (G, S), one can in $\mathcal{O}(k_{OPT}^{142}|G|)$ time find a set $F \subseteq E(G)$ of $\mathcal{O}(k_{OPT}^{142})$ edges that contains an optimal Steiner tree connecting S in G, where k_{OPT} is the size of an optimal Steiner tree.*

We emphasise two aspects of Theorem I.2. First, the proposed algorithm *does not* need to be given an optimal solution nor its size, even though the running time and output size of the algorithm are polynomial in the size of an optimal solution. Second, the running time of the algorithm can be bounded by $\mathcal{O}(|G|^2)$: if $|G|$ is smaller than the promised kernel bound, then the algorithm may simply return the input graph without any modification. Similar remarks hold also for the second and third applications of Theorem I.1 that we present later.

For the second application of Theorem I.1, we modify the approach of Theorem I.2 for the closely related STEINER FOREST problem on planar graphs. Recall that a *Steiner forest* that *connects* a family $\mathcal{S} \subseteq V(G) \times V(G)$ of terminal pairs in a graph G is a forest in G such that both vertices of each pair in $\mathcal{S}$ are contained in the same connected component of the forest.

Theorem I.3. *Given a* PLANAR STEINER FOREST *instance $(G, \mathcal{S})$, one can in $\mathcal{O}(k_{OPT}^{710}|G|)$ time find a set $F \subseteq E(G)$ of $\mathcal{O}(k_{OPT}^{710})$ edges that contains an optimal Steiner forest connecting $\mathcal{S}$ in G, where k_{OPT} is the size of an optimal Steiner forest.*

Using the analogue of Theorem I.1 for bounded-genus graphs, we extend Theorems I.2 and I.3 to obtain a polynomial kernel for STEINER TREE and even STEINER FOREST on such graphs. Here, we assume that we are given an embedding of the input graph into a surface of genus g such that the interior of each face is homeomorphic to an open disc.

For the third application of Theorem I.1, we consider EDGE MULTIWAY CUT on planar graphs. Recall that an *edge multiway cut*[1] in a graph G is a set $X \subseteq E(G)$ such that no two vertices of a given set $S \subseteq V(G)$ are in the same component of $G \setminus X$. A recent breakthrough in the application of matroid theory to kernelization problems [15, 20] led to the discovery of a polynomial kernel for MULTIWAY CUT on general graphs with a constant number of terminals. It is a major open question whether this problem has a polynomial kernel for an arbitrary number of terminals [14, 15]. Here, we show that such a polynomial kernel does exist for EDGE MULTIWAY CUT on planar graphs.

Theorem I.4. *Given a* PLANAR EDGE MULTIWAY CUT *instance (G, S), one can in polynomial time find a set $F \subseteq E(G)$ of $\mathcal{O}(k_{OPT}^{568})$ edges that contains an optimal solution to (G, S), where k_{OPT} is the size of this optimal solution.*

We note that in contrast to the work on polynomial kernels for MULTIWAY CUT mentioned before [15, 20], we do not rely on matroid theory.

As an immediate consequence of Theorem I.2 and Theorem I.4, we observe that by plugging the kernels promised by these theorems into the algorithms of Tazari [21] for PLANAR STEINER TREE or its modification for PLANAR EDGE MULTIWAY CUT, or the algorithm of Klein and Marx [22] for PLANAR EDGE MULTIWAY CUT, respectively, we obtain faster parameterized algorithms for both problems.

[1]In the approximation algorithms literature, the term *multiway cut* usually refers to an edge cut, i.e., a subset of edges of the graph, and the node-deletion variants of the problem are often much harder. However, from the point of view of parameterized complexity, there is usually little or no difference between edge- and node-deletion variants of cut problems, and hence one often considers the (more general) node-deletion variant as the 'default one'. To avoid confusion, in this work we always explicitly state that we consider the edge-deletion variant.

Corollary I.5. *Given a planar graph G, a terminal set $S \subseteq V(G)$, and an integer k, one can*

1) *in $2^{\mathcal{O}(\sqrt{k \log k})} + \mathcal{O}(k^{142}|V(G)|)$ time decide whether the* PLANAR STEINER TREE *instance (G, S) has a solution with at most k edges;*
2) *in $2^{\mathcal{O}(\sqrt{k} \log k)} + \mathrm{poly}(|V(G)|)$ time decide whether the* PLANAR EDGE MULTIWAY CUT *instance (G, S) has a solution with at most k edges;*
3) *in $2^{\mathcal{O}(|S| + \sqrt{|S|} \log k)} + \mathrm{poly}(|V(G)|)$ time decide whether the* PLANAR EDGE MULTIWAY CUT *instance (G, S) has a solution with at most k edges.*

This corollary improves on the subexponential-time algorithm for PLANAR STEINER TREE previously proposed by the authors [17], and on the algorithm for PLANAR EDGE MULTIWAY CUT by Klein and Marx [22] if $k = o(\log |V(G)|)$. As Tazari's algorithm extends to graphs of bounded genus, combining it with our kernelization algorithm, we obtain the first subexponential-time algorithm for STEINER TREE on graphs of bounded genus. The running time is a computable function of the genus times the running time of the planar case.

We also remark that a similar corollary is unlikely to exist for the case of PLANAR STEINER FOREST. We observe that the lower bound for STEINER FOREST on graphs of bounded treewidth of Bateni et al. [23], with minor modifications, shows also that PLANAR STEINER FOREST does not admit a subexponential-time algorithm unless the Exponential Time Hypothesis of Impagliazzo, Paturi, and Zane [24] fails.

Theorem I.6. *Unless the Exponential Time Hypothesis fails, no algorithm can decide in $2^{o(k)}\mathrm{poly}(|G|)$ time whether* PLANAR STEINER FOREST *instances $(G, \mathcal{S})$ have a solution with at most k edges.*

Edge-Weighted Planar Graphs. Although the decomposition methods in the proof of Theorem I.1 were developed with applications in unweighted graphs in mind, they can be modified for graphs with positive edge weights (henceforth called *edge-weighted graphs*). That is, we show the following weighted and approximate variant of Theorem I.1:

Theorem I.7. *Let $\varepsilon > 0$ be a fixed accuracy parameter, and let B be an edge-weighted brick. Then one can find in $\mathrm{poly}(\varepsilon^{-1})\,|B| \log |B|$ time a subgraph H of B such that*[2]

(i) $\partial B \subseteq H$,
(ii) $w(H) \leq \mathrm{poly}(\varepsilon^{-1}) w(\partial B)$, *and*
(iii) for every set $\mathcal{S} \subseteq V(\partial B) \times V(\partial B)$ there exists a Steiner forest F_H that connects $\mathcal{S}$ in H such that $w(F_H) \leq w(F_B) + \varepsilon w(\partial B)$ for any Steiner forest F_B that connects $\mathcal{S}$ in B.

Notice that, contrary to Theorem I.1, we state Theorem I.7 in the language of Steiner *forest*, not Steiner tree. The reason is that the allowed error in Theorem I.7 is additive, and therefore the forest statement seems significantly stronger than

[2]We denote by $w(H)$ the total weight of all the edges of a graph H.

the tree one. Observe that for Theorem I.1, it would be of no consequence to state it in the language of Steiner forest instead of in the language of Steiner tree. We provide a sketch of the proof of Theorem I.7 in Section IV.

Theorem I.7 influences the known polynomial-time approximation schemes for network design as follows. The mortar graph framework of Borradaile, Klein, and Mathieu [5] may be understood as a method to decompose a brick into cells, such that each cell is equipped with θ evenly-spaced portal vertices, and there is an approximate Steiner tree that for each cell uses a subset of the portal vertices to enter and leave the cell. Then it suffices to preserve an approximate or optimal Steiner tree for any subset of portal vertices. Previously, only a bound that is exponential in θ on the preserved subgraph of each cell was known [5]. The impact of our work is that the dependency on θ can be reduced to a polynomial. This observation is not only used to prove Theorem I.7, but also leads to deeper understanding of the mortar graph framework.

Observe that one can directly derive an EPTAS for PLANAR STEINER TREE from Theorem I.7: cut the input graph G open along a 2-approximate Steiner tree (as in the kernel), apply Theorem I.7 to the resulting brick B, and project the obtained graph H back onto the original graph. An optimal Steiner tree in G becomes an optimal Steiner forest in B, and thus the projection of H preserves an approximate Steiner tree for the input instance. Since the total weight of H is within a multiplicative factor $\mathrm{poly}(\varepsilon^{-1})$ of the weight of the optimal solution for the input instance, an application of Baker's shifting technique [25] can find an approximate solution in H in $2^{\mathrm{poly}(\varepsilon^{-1})}|H|\log|H|$ time. However, we note that the polynomial dependency on ε in the exponent is worse than the one obtained by the currently known EPTAS [5], despite our substantially improved reduction of the cells. This is because that EPTAS utilizes Baker's technique in a more clever way that is aware of the properties of the mortar graph, and is indifferent to the actual replacement within each cell.

B. Discussion

A drawback of our methods is that the exponents in the kernel bounds and the polynomial dependency on ε^{-1} in the weighted variant are currently large, making the results theoretical. However, we see the strength of our results in that we prove that a polynomial kernel actually exists — thus proving that PLANAR STEINER TREE, PLANAR STEINER FOREST, and PLANAR EDGE MULTIWAY CUT belong to the class of problems that have a polynomial kernel — rather than in the actual size bound. We believe that, using the road paved by our work, it is possible to decrease the exponent in the bound of the kernel. In fact, we conjecture that the correct dependency in Theorem I.1 is quadratic, with a grid being the worst-case scenario.

Another limitation of our methods is that we need to parameterize by the *number of edges* of the Steiner tree. Although we believe that we can extend our results to kernelize PLANAR STEINER TREE with respect to the parameter *number of non-terminal vertices of the tree*, extending to the parameter *number of terminals* seems challenging. On general graphs, this parameter has already been studied, and it is known that the problem has a (tight) fixed-parameter algorithm [26, 27] and no polynomial kernel unless part of the polynomial hierarchy collapses [13]. However, our methods seem far from resolving whether PLANAR STEINER TREE has a polynomial kernel with respect to this smaller parameter. Similarly, one may consider graph-separation problems with vertex-based parameters, such as ODD CYCLE TRANSVERSAL or the node-deletion variant of MULTIWAY CUT. On planar graphs, both of these problems are some sort of Steiner problem on the dual graph. It would be interesting to show polynomial kernels for these problems (without using the matroid framework [15]).

To generalize our methods, it would be interesting to lift our results to more general graph classes, such as graphs with a fixed excluded minor. For EDGE MULTIWAY CUT, even the bounded-genus case remains open. Further work is also needed to improve the allowed error in Theorem I.7. Currently, this error is an additive error of $\varepsilon w(\partial B)$. In other words, a near-optimal Steiner forest is preserved only for "large" optimal forests, that is, for ones of size comparable to the perimeter of B. Is it possible to improve Theorem I.7 to ensure a $(1+\varepsilon)$ multiplicative error? That is, to obtain a variant of Theorem I.7 where the graph H satisfies $w(F_H) \leq (1+\varepsilon)w(F_B)$, and thus to preserve near-optimal Steiner forests at *all scales*? Finally, since our methods handle problems that are beyond the reach of the theory of bidimensionality, our contribution might open the door to a more general framework that is capable of addressing a broader range of problems.

C. Related work

The three problems considered in this paper (STEINER TREE, STEINER FOREST, and EDGE MULTIWAY CUT) are all NP-hard [28, 29] and unlikely to have a PTAS [29, 30] on general graphs. However, they do admit constant-factor approximation algorithms [31]–[33].

STEINER TREE has a $2^{|S|}\cdot\mathrm{poly}(|G|)$-time, polynomial-space algorithm on general graphs [26]; the exponential factor is believed to be optimal [27], but an improvement has not yet been ruled out under the Strong Exponential Time Hypothesis. The algorithm for STEINER TREE implies a $(2|\mathcal{S}|)^{|\mathcal{S}|}\cdot\mathrm{poly}(|G|)$-time, polynomial-space algorithm for STEINER FOREST. On the other hand, EDGE MULTIWAY CUT remains NP-hard on general graphs even when $|S| = 3$ [29], while for the parameterization by the size of the cut k, a $1.84^k\cdot\mathrm{poly}(|G|)$-time algorithm is known [34].

Neither STEINER TREE nor STEINER FOREST admits a polynomial kernel on general graphs [13], unless the polynomial hierarchy collapses. Recently, a polynomial kernel was given for EDGE and NODE MULTIWAY CUT for a constant number of terminals or deletable terminals [15]; nevertheless, the question for a polynomial kernel in the general case remains open.

STEINER TREE, STEINER FOREST, and EDGE MULTIWAY CUT all remain NP-hard on planar graphs [29, 35], even in restricted cases. All three problems do admit an EPTAS on planar graphs [5, 36, 37], and STEINER TREE admits an EPTAS on bounded-genus graphs [38].

We are not aware of any previous kernelization results for STEINER TREE, STEINER FOREST, or EDGE MULTIWAY CUT on planar graphs. The question of the existence of a subexponential-time algorithm for PLANAR STEINER TREE was first explicitly pursued by Tazari [21]. He showed that such a result would be implied by a subexponential or polynomial kernel. The current authors adapted the main ideas of the EPTAS for PLANAR STEINER TREE [5] to show a subexponential-time algorithm [17], without actually giving a kernel beforehand. The algorithm of [17] in fact finds subexponentially many subgraphs of subexponential size, one of which is a subexponential kernel if the instance is a YES-instance. Finally, for EDGE MULTIWAY CUT on planar graphs, a $2^{\mathcal{O}(|S|)} \cdot |G|^{\mathcal{O}(\sqrt{|S|})}$-time algorithm is known [22] and believed to be optimal [39].

II. OVERVIEW OF THE PROOF OF THEOREM I.1

Before we start, we set up some notation. For a subgraph H of G, we silently identify H with the edge set of H; that is, all our subgraphs are edge-induced. For a brick B, $\partial B[a,b]$ denotes the subpath of ∂B obtained by traversing ∂B in counter-clockwise direction from a to b. By Π we denote the standard Euclidean plane. For a closed curve γ on Π, we say that γ *strictly encloses* $c \in \Pi$ if $c \notin \gamma$ and γ is not continuously retractable in $\Pi \setminus \{c\}$ to a single point, and γ *encloses* c if it strictly encloses c or $c \in \gamma$. This notion naturally translates to cycles and walks in a plane graph G (strictly) enclosing vertices, edges, and faces of G.

The idea behind the proof of Theorem I.1 is to apply it recursively on subbricks (subgraphs enclosed by a simple cycle) of the given brick B. The main challenge is to devise an appropriate way to decompose B into subbricks, so that their "measure" decreases. Here we use the *perimeter* of a brick as a potential that measures the progress of the algorithm.

Intuitively, we would want to do the following. Let T be a tree in B that connects a subset of the vertices on the perimeter of B. Then T splits B into a number of smaller bricks $B_1, \ldots, B_r$, formed by the finite faces of $\partial B \cup T$ (see Figure 1a). We recurse on bricks B_i, obtaining graphs $H_i \subseteq B_i$, and return $H := \bigcup_{i=1}^{r} H_i$. We can prove that this decomposition yields a polynomial bound on $|H|$ if (i) all bricks B_i have multiplicatively smaller perimeter than B, and (ii) the sum of the perimeters of the subbricks is linear in the perimeter of B.

In this approach, there are two clear issues that need to be solved. The first issue is that we need an algorithm to decide whether there is a tree T for which the induced set of subbricks satisfies conditions (i) and (ii). We design a dynamic programming algorithm that either correctly decides that no such tree exists, or finds a set of subbricks of B that satisfies condition (i) and (ii). In the latter case, we can recurse on each of those subbricks.

The second issue is that there might be no trees T for which the induced set of subbricks satisfies conditions (i) and (ii). In this case, optimal Steiner trees, which are a natural candidate for such partitioning trees T, behave in a specific way. For example, consider the tree of Figure 1b, which consists of two small trees T_1, T_2 that lie on opposite sides of the brick B and that are connected through a shortest path P (of length slightly less than $|\partial B|/2$). Then both faces of $\partial B \cup T$ that neighbour P may have perimeter almost equal to $|\partial B|$, thus blocking our default decomposition approach.

To address this second issue, we propose a completely different decomposition. Intuitively, we find a cycle C of length linear in $|\partial B|$ that lies close to ∂B, such that all vertices of degree three or more of any optimal Steiner tree are hidden in the ring between C and ∂B (see Figure 1c). We then decompose the ring between ∂B and C into a number of smaller bricks. We recursively apply Theorem I.1 to these bricks, and return the result of these recursive calls together with a set of shortest paths inside C between any pair of vertices on C.

In Section II-A below, we formalise the above notions and give the algorithm that addresses the first issue. Then, Section II-B describes the default decomposition, whereas Section II-C describes the alternative decomposition that addresses the second issue.

A. Deciding on the Decomposition

In this section, we present some of the basic notions of our paper and describe the algorithm that decides which of the two possible decompositions is used.

Definition II.1. *For a brick B, a* brick covering *of B is a family $\mathcal{B} = \{B_1, \ldots, B_p\}$ of bricks, such that* (i) *each B_i, $1 \le i \le p$, is a subbrick of B, and* (ii) *each face of B is contained in at least one brick B_i, $1 \le i \le p$. A brick covering is called a* brick partition *if each face of B is contained in* exactly one *brick B_i.*

We note that if $\mathcal{B} = \{B_1, \ldots, B_p\}$ is a brick partition of B, then every edge of ∂B belongs to the perimeter of exactly one brick B_i, while every edge strictly enclosed by ∂B either is in the interior of exactly one brick B_i, or lies on the perimeters of exactly two bricks B_i, B_j for $i \neq j$.

Any connected set $F \subseteq B$ will be called a *connector*. Let S be the set of vertices of ∂B adjacent to at least one edge of F; the elements of S are the *anchors* of the connector F. We then say that F *connects* S. For a connector F, we say that F is *optimal* if there is no connector F' with $|F'| < |F|$ that connects a superset of the anchors of F. Clearly, each optimal connector F induces a tree, whose every leaf is an anchor of F. We say that a connector $F \subseteq B$ is *brickable* if the boundary of every inner face of $\partial B \cup F$ is a simple cycle, i.e., these boundaries form subbricks of B. Let $\mathcal{B}$ be the corresponding brick partition of B. Observe that $\sum_{B' \in \mathcal{B}} |\partial B'| \le |\partial B| + 2|F|$.

Next, we define the crucial notions for partitions and coverings that are used for the default decomposition.

Definition II.2. *The* total perimeter *of a brick covering $\mathcal{B} = \{B_1, \ldots, B_p\}$ is defined as $\sum_{i=1}^{p} |\partial B_i|$. For a constant $c > 0$,*

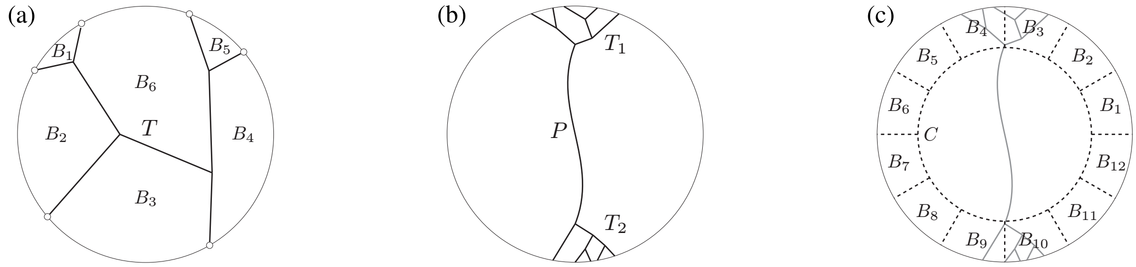

Fig. 1: (a) shows an optimal Steiner tree T and how it partitions the brick B into smaller bricks $B_1, \ldots, B_r$. (b) shows an optimal Steiner tree that connects a set of vertices on the perimeter of B and that consists of two small trees T_1, T_2 that are connected by a long path P; note that both bricks neighbouring P may have perimeter very close to $|\partial B|$. (c) shows a cycle C that (in particular) hides the small trees T_1, T_2 in the ring between C and ∂B, and a subsequent decomposition of B into smaller bricks.

*$\mathcal{B}$ is c-*short *if the total perimeter of $\mathcal{B}$ is at most $c \cdot |\partial B|$. For a constant $\tau > 0$, $\mathcal{B}$ is τ-*nice *if $|\partial B_i| \leq (1-\tau) \cdot |\partial B|$ for each $1 \leq i \leq p$.*

*Similarly, a brickable connector $F \subseteq B$, with $\mathcal{B} = \{B_1, \ldots, B_p\}$ being the corresponding brick partition, is c-*short *if $\mathcal{B}$ is c-short, is simply* short *if it is 3-short, and is* τ-nice *if $\mathcal{B}$ is τ-nice.*

Observe that if $F \subseteq B$ is a brickable connector, then F is c-short if $|F| \leq |\partial B| \cdot (c-1)/2$, and F is short if $|F| \leq |\partial B|$. Moreover, if F is an optimal connector, then F is a short brickable connector, as F must be a tree of length at most $|\partial B|$. Now we are ready to give the algorithm that decides what decomposition to use.

Theorem II.3. *Let $\tau > 0$ be a fixed constant. Given a brick B, in $\mathcal{O}(|\partial B|^8 |B|)$ time one can either correctly conclude that no short τ-nice tree exists in B or find a 3-short τ-nice brick covering of B.*

The proof of Theorem II.3 is a technical modification of the classical algorithm of Erickson et al. [40]. That algorithm computes an optimal Steiner tree in a planar graph assuming that all the terminals lie on the boundary of the infinite face. It uses the Dreyfus-Wagner dynamic-programming approach, where a state consists of a subset of already connected terminals, and the current "interface" vertex; the main observation is that only states with consecutive terminals on the boundary are relevant, yielding a polynomial bound on the number of them. In our case, we can proceed similarly: our state consists of the leftmost and rightmost chosen terminal, the "interface" vertex inside the brick, the total length of the tree, and the length of the leftmost and rightmost path in the constructed tree. Consequently, the terminals are chosen on-the-fly.

In case some short τ-nice tree exists, for technical reasons we cannot ensure that the output of the algorithm of Theorem II.3 will actually be a brick partition corresponding to some short τ-nice tree. Instead, the algorithm may output a brick covering, but one that is guaranteed to be 3-short and τ-nice. This is sufficient for our purposes.

We can now formally describe the main line of reasoning of our sparsification algorithm. Let $\tau > 0$ be some constant chosen later. If $|\partial B| \leq 2/\tau$, then for each $S \subseteq V(\partial B)$ we compute an optimal Steiner tree that connects S using the algorithm of Erickson et al. [40], and take the union of all such trees. If $|\partial B| > 2/\tau$, then we run the algorithm of Theorem II.3 for B and τ. If the algorithm returns a 3-short τ-nice brick covering, then we proceed to the default decomposition, formalized in Section II-B below. Otherwise, if the algorithm of Theorem II.3 concluded that B does not contain any short τ-nice tree, then we proceed to the arguments in Section II-C. We show that in all cases we obtain a subgraph of B that satisfies conditions (i)-(iii) of Theorem I.1.

B. The Default Decomposition

Suppose that the algorithm of Theorem II.3 returns a 3-short τ-nice brick covering $\mathcal{B} = \{B_1, \ldots, B_p\}$ of B. We can then use this brick covering as a decomposition and recurse on each brick individually. This is formalized in the following lemma.

Lemma II.4. *Let $c, \tau > 0$ be constants. Let B be a brick and let $\mathcal{B} = \{B_1, \ldots, B_p\}$ be a c-short τ-nice brick covering of B. Assume that the algorithm of Theorem I.1 was applied recursively to bricks $B_1, \ldots, B_p$, and let $H_1, \ldots, H_p$ be the subgraphs output by this algorithm for $B_1, \ldots, B_p$, respectively, where $|H_i| \leq C \cdot |\partial B_i|^\alpha$ for some constants $C > 0$ and $\alpha \geq 1$ such that $(1-\tau)^{\alpha-1} \leq \frac{1}{c}$. Let $H = \bigcup_{i=1}^{p} H_i$. Then H satisfies conditions (i)-(iii) of Theorem I.1, with $|H| \leq C \cdot |\partial B|^\alpha$.*

Proof: To see that H satisfies condition (i), note that every edge of ∂B is in the perimeter of some brick B_i, and that $\partial B_i \subseteq H_i$ for every $i = 1, 2, \ldots, p$. Therefore, $\partial B \subseteq H$.

To see that H satisfies condition (ii), recall that $\mathcal{B}$ is c-short and that $|\partial B_i| \leq (1-\tau) \cdot |\partial B|$ for each $i = 1, 2, \ldots, p$. Therefore, $|\partial B_i|^\alpha \leq |\partial B_i| \cdot (1-\tau)^{\alpha-1} |\partial B|^{\alpha-1}$, and $|H| \leq C \cdot |\partial B|^\alpha$.

Finally, to see that H satisfies condition (iii), let $S \subseteq V(\partial B)$ be a set of terminals lying on the perimeter of B, and let T be an optimal Steiner tree connecting S in B that contains a minimum number of edges that are not in H. We claim that $T \subseteq H$. Assume the contrary, and let

$e \in T \setminus H$. Since each face of B is contained in some brick of $\mathcal{B}$, there exists a brick B_i such that ∂B_i encloses e. As $\partial B_i \subseteq H_i \subseteq H$, we infer $e \notin \partial B_i$. Consider the subgraph of T strictly enclosed by ∂B_i, and let X be the connected component of this subgraph that contains e. Clearly, X is a connector inside B_i. Since H_i is obtained by a recursive application of Theorem I.1, there exists a connected subgraph $D \subseteq H_i$ that connects the anchors of X and that satisfies $|D| \leq |X|$. Let $T' = (T \setminus X) \cup D$. Observe that $|T'| \leq |T|$ and that T' contains strictly less edges that are not in H than T does. Since D connects the anchors of X in H_i, T' still connects the anchors of T in B, that is, T' connects S. However, T is an optimal Steiner tree that connects S, and thus T' is also an optimal Steiner tree that connects S. Since T' contains strictly less edges that are not in H than T, this contradicts the choice of T. Hence, $T \subseteq H$. ∎

C. The Alternative Decomposition — Mountain Ranges and the Core

Suppose that the algorithm of Theorem II.3 decides that no short τ-nice tree exists in B. As mentioned before, we want to find a cycle C of length linear in $|\partial B|$ that is close to ∂B, such that all vertices of degree three or more of any optimal Steiner tree are hidden in the ring between C and ∂B (see Figure 1c). In the following, we use a constant $\delta \in (0, \frac{1}{2})$, which depends on τ and is chosen later.

Definition II.5. *A* δ-carve L *from a brick* B *is a pair* (P, I)*, where* P*, called the* carvemark*, is a path in* B *between two distinct vertices* $a, b \in V(\partial B)$ *of length at most* $(\frac{1}{2} - \delta) \cdot |\partial B|$*, and* I*, called the* carvebase*, is a shortest of the two paths* $\partial B[a, b], \partial B[b, a]$*. The subgraph enclosed by the closed walk* $P \cup I$ *is called the* interior *of a* δ*-carve.*

Of particular interest will be the following special type of δ-carves.

Definition II.6. *For fixed* $l, r \in V(\partial B)$*, a* δ-mountain *of* B *for* l, r *is a* δ*-carve* M *in* B *such that*

1) *l and r are the endpoints of the carvemark and carvebase of M;*
2) *the edges of the carvemark can be partitioned into two paths P_L, P_R, where P_L is a shortest l–P_R path in the interior of M and P_R is a shortest r–P_L path in the interior of M.*

We write $M = (P_L \wedge P_R)$ to exhibit the partition of the carvemark into paths P_L and P_R. We use v_M to denote the unique vertex of $V(P_L) \cap V(P_R)$. We also say that a δ-mountain M *connects* the vertices l and r.

The following lemma motivates why we are interested in δ-mountains. For a tree T, $T[a, b]$ denotes the unique path in T between vertices a and b.

Lemma II.7. *Let B be a brick and let T be an optimal Steiner tree connecting $V(T) \cap V(\partial B)$ in B. Let $uv \in T$ be an edge of T, where v is of degree at least 3 in T, and let T_v be the connected component of $T \setminus \{uv\}$ containing v, rooted at v. Let l and r be the leftmost and rightmost elements of $V(T_v) \cap V(\partial B)$, that is, $V(T_v) \cap V(\partial B) \subseteq V(\partial B[l, r])$ and $T[l, r] \cup \partial B[r, l]$ encloses uv. Assume furthermore that $|\partial B[l, r]| < |\partial B|/2$. Then $M := (T[l, v] \wedge T[r, v])$ is a δ-mountain connecting l and r, for any $\delta < 1/2 - |T[l, r]|/|\partial B|$.*

Proof: As v is of degree at least 3 in T, v has degree at least 2 in T_v, and $T[l, v] \cap T[r, v] = \{v\}$. Therefore, $T[l, v] \cup T[v, r] = T[l, r]$, and $T[l, r]$ induces a δ-carve M with carvebase $\partial B[l, r]$.

Suppose that M is not a δ-mountain if we take $P_L = T[l, v]$ and $P_R = T[r, v]$. Without loss of generality, there exists a path P enclosed by M that connects l with $w \in V(P_R)$, $V(P_R) \cap V(P) = \{w\}$, and $|P| < |T[l, v]|$. Let D be the subgraph of M enclosed by the closed walk $T[l, v] \cup P \cup T[v, w]$. Define $T' := (T \setminus D) \cup T[v, w] \cup P$. As $T[v, w] \cup T[l, v] \subseteq D$, $|T'| < |T|$. By the definition of l and r, $T[l, v] \setminus P$ does not contain any vertex of ∂B. Therefore, T' is a connected subgraph of B connecting $V(T) \cap V(\partial B)$, a contradiction to the optimality of T. ∎

The above lemma shows that small subtrees of optimal Steiner trees in B are hidden in δ-mountains. Here, 'small' means that the leftmost and rightmost path in the subtree have total length at most $(1/2 - \delta) \cdot |\partial B|$. Note that an optimal Steiner tree in B has total size smaller than $|\partial B|$, as ∂B without an arbitrary edge connects any subset of $V(\partial B)$. Therefore, if we choose δ appropriately, then we can 'hide' almost an entire optimal non-τ-nice Steiner tree in at most two δ-mountains. To hide most of *all* optimal Steiner trees, we consider unions of δ-mountains. For fixed vertices $l, r \in V(\partial B)$, the δ*-mountain range* is the closed walk $W_{l,r}$ in B such that a face f of B is enclosed by $W_{l,r}$ if and only if f belongs to some δ-mountain that connects l and r.

Theorem II.8 (Mountain Range Theorem)**.** *Fix $\tau \in [0, 1/4)$ and $\delta \in [2\tau, 1/2)$, and assume that B does not admit any short τ-nice tree. Then for any fixed $l, r \in V(\partial B)$ with $|\partial B[l, r]| < |\partial B|/2$, $W_{l,r}$ has length at most $3 \cdot |\partial B[l, r]|$. Moreover, the set of the faces enclosed by $W_{l,r}$ can be computed in $\mathcal{O}(|B|)$ time.*

Proof sketch: By case analysis, we deduce that the set of all inclusion-wise maximal δ-mountains essentially looks as in Figure 2a, i.e., for any two maximal mountains there exists exactly one region of the plane that is in one of them but not in the other one.

Let $\{M^i = (P_L^i, P_R^i)\}_{i=1}^s$ be the set of all these maximal δ-mountains, ordered from left to right. By induction, we show that the perimeter of the union of the first i δ-mountains, denoted p^i, is at most $|\partial B[l, r]| + |P_R^1| + |P_L^i|$. This statement clearly holds for $i = 1$, and for $i = s$ it proves the bound on the perimeter of the δ-mountain range promised by Theorem II.8.

For the inductive step, define $b = |P_R^{i+1}|$ and $e = |P_L^i|$. Let v be the first point on P_L^{i+1} that lies on P_R^i. We denote the distance (along P_L^{i+1}) from l to v as d and the distance from v to $v_{M^{i+1}}$ as a. Finally, we denote by c the distance (along P_R^i) from r to v. These definitions are illustrated in Figure 2a.

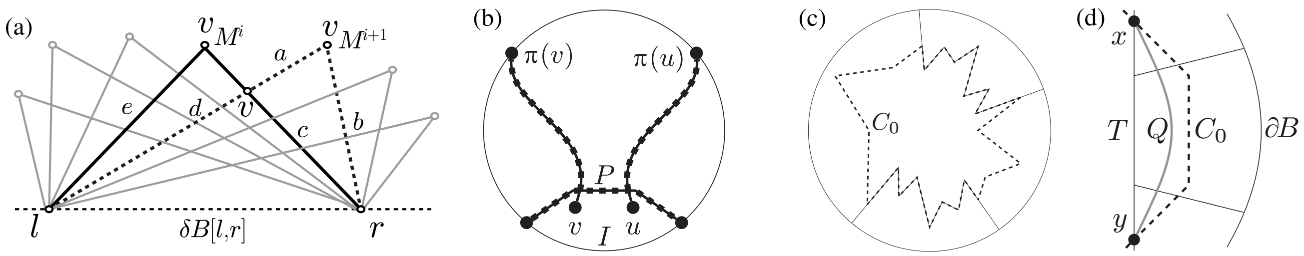

Fig. 2: (a) shows a mountain range. (b) shows a short τ-nice tree occurring if $\pi(v)$ and $\pi(u)$ are far from I. (c) shows the cycle C_0 formed by the union of the perimeters of the mountain ranges; example mountain ranges are drawn solid. (d) shows how to shortcut the tree T (solid) with a shortest xy-path Q (gray).

Observe that $d \geq e$, because M^i is a δ-mountain. Similarly, observe that $c \geq b$, because M^{i+1} is a δ-mountain. Hence, $p^{i+1} - p^i = a + b - c \leq a \leq a + d - e = |P_L^{i+1}| - |P_L^i|$. This concludes the inductive step. We omit the description of the algorithm that finds the mountain range. ■

We now designate $\mathcal{O}(\tau^{-1})$ vertices on ∂B, and construct the union $\mathcal{M}$ of all δ-mountain ranges for each pair of designated vertices. Using the following deep theorem, we can show that $\mathcal{M}$ is not the entire brick.

Theorem II.9 (Core Theorem)**.** *For any $\tau \in (0, \frac{1}{4})$ and any $\delta \in [2\tau, \frac{1}{2})$, if B has no short τ-nice tree, then there exists a face of B that is not enclosed by any δ-carve. Moreover, such a face can be found in $O(|B|)$ time.*

Proof sketch: Suppose, for sake of contradiction, that all faces of B are enclosed by some δ-carve. We first observe that, for any brickable short tree T with diameter not more than $(\frac{1}{2} - \delta) \cdot |\partial B|$, there exists an interval I_T of ∂B of length at most $(\frac{1}{2} - \frac{\delta}{2}) \cdot |\partial B|$ such that all anchors of T are in I_T. If no such interval exists, then every brick induced by T has perimeter less than $(\frac{1}{2} - \delta) \cdot |\partial B| + (\frac{1}{2} + \frac{\delta}{2}) \cdot |\partial B| \leq (1 - \tau) \cdot |\partial B|$. Hence, T would be τ-nice, a contradiction.

Define a map $v \to \pi(v)$ for $v \in V(B)$ such that $\pi(v)$ is a vertex of ∂B closest to v. The main observation is that if v and u belong to the interior of some δ-carve (P, I), then the distance between $\pi(v)$ and $\pi(u)$ along ∂B is at most $(\frac{1}{2} - \frac{\delta}{2}) \cdot |\partial B|$. To see this, consider the shortest paths P_v from v to $\pi(v)$. These paths can be used to form a tree T, consisting of P, the subpath of P_v to $\pi(v)$ from the last point of P_v on P, and the subpath of P_u to $\pi(u)$ from the last point of P_u on P (see Figure 2b). We observe that the diameter of T is bounded by $|P| \leq (\frac{1}{2} - \delta) \cdot |\partial B|$, because the paths that make up T always have length at most the corresponding part of P. Moreover, as T has only four leaves, $|T|$ is bounded by twice the diameter of T, so T is short. Hence, $\pi(v)$, $\pi(u)$, and $V(P) \cap V(\partial B)$ lie on the interval I_T, as observed above. We extend π to the edges of B by mapping uv onto the shorter subpath between $\pi(u)$ and $\pi(v)$ on ∂B. Now consider a face f that is enclosed by (P, I). We note that no point of any edge of f is mapped to a point lying exactly opposite on ∂B to any point in $V(P) \cap V(\partial B)$, as such points cannot belong to I_T. Hence, all edges of f are mapped to an interval of ∂B. Since an interval is a simply connected metric space, we can extend π from the boundary of face f to its interior in a continuous manner such that the whole face f is mapped into it. Consequently, since every face of B can be enclosed by a δ-carve, we have constructed a retraction of a closed disc onto its boundary. This contradicts Borsuk's non-retraction theorem [41]. ■

As each δ-mountain is actually a δ-carve, $\mathcal{M}$ does not contain an arbitrarily chosen core face $f_{\texttt{core}}$ promised by Theorem II.9. Hence, the union of the perimeters of the δ-mountain ranges that make up $\mathcal{M}$ contains a cycle C_0 that separates $f_{\texttt{core}}$ from the mountain ranges. Moreover, as we construct only $\mathcal{O}(\tau^{-2})$ mountain ranges, each of perimeter $\mathcal{O}(|\partial B|)$ by Theorem II.8, we have that $|C_0| = \mathcal{O}(|\partial B|)$; see Figure 2c.

We observe that certain optimal Steiner trees in B may behave nontrivially in the subgraph enclosed by C_0, and in particular, may still have a vertex of degree three or more that is enclosed by C_0. However, this behavior is easily dealt with as follows. Consider the situation in Figure 2d. If Q is a shortest path between x and y, then we may replace the part of the tree to the left of Q by Q. Hence, we shortcut C_0 whenever possible while keeping $f_{\texttt{core}}$ enclosed by C_0. By choosing $\delta = 4\tau$, we then obtain the following result.

Theorem II.10. *Let $\tau \in (0, 1/36]$. Assume that B does not admit any short τ-nice tree and that $|\partial B| > 2/\tau$. Then one can in $\mathcal{O}(|B|)$ time compute a simple cycle C in B with the following properties:*

1) *the length of C is at most $\frac{16}{\tau^2}|\partial B|$;*
2) *for each vertex $x \in V(C)$, there exists a path from x to $V(\partial B)$ of length at most $(\frac{1}{4} - 2\tau) \cdot |\partial B|$ such that no edge of the path is strictly enclosed by C;*
3) *C encloses $f_{\texttt{core}}$, where $f_{\texttt{core}}$ is a face of B, promised by Theorem II.9, that is not enclosed by any 2τ-carve;*
4) *for any $S \subseteq V(\partial B)$, there exists an optimal Steiner tree T_S that connects S in B such that no vertex of degree at least 3 in T_S is strictly enclosed by C.*

Finally, we are ready to describe the decomposition. Apply the algorithm of Theorem II.10 to B, and let C denote the resulting cycle. We can then decompose the brick as in Figure 1c, meaning that the area between C and ∂B is partitioned into a number of small subbricks of total perimeter $\mathcal{O}(|\partial B|)$. Here we use the second property of C that is promised by Theorem II.10 to build the sides of the subbricks. We recursively apply Theorem I.1 to these subbricks, and let H denote the union of the resulting subgraphs. Then we add to H for each pair of vertices of C a shortest path in the area enclosed by C between the two vertices if that shortest path has length at most $|\partial B|$. The linear bound on the total perimeter of the subbricks enables a similar analysis as in the proof of Lemma II.4. We then choose $\tau = 1/36$. This concludes the proof of Theorem I.1.

III. Applications of Theorem I.1

In this section, we briefly sketch how to prove Theorems I.2, I.3, and I.4.

Proof sketch of Theorem I.2: We manipulate the graph such that all terminals lie on the outer face. We first find a 2-approximate Steiner tree T_{apx} for S in G. We then cut the plane open along T_{apx}, cf. [5]. That is, we create an Euler tour of T_{apx} that traverses each edge twice in different directions and respects the plane embedding of T_{apx}. Then we duplicate every edge of T_{apx}, replace each vertex v of T_{apx} with $d-1$ copies of v, where d is the degree of v in T_{apx}, and distribute the copies in the plane embedding so that we obtain a new face F with boundary corresponding to the aforementioned Euler tour. Fix the embedding of the resulting graph $\hat{G}$ such that F is its outer face. Note that the terminals S lie only on the outer face of $\hat{G}$, and that $|\partial\hat{G}| \leq 4k_{OPT}$. Apply Theorem I.1 to $\hat{G}$ to obtain $\hat{H}$, which is of size $\mathcal{O}(|\partial\hat{G}|^{142}) = \mathcal{O}(k_{OPT}^{142})$. As an optimal Steiner tree T in G splits into a family of trees in $\hat{G}$ that each connect subsets of $V(T) \cap V(\partial\hat{G})$, the projection of $\hat{H}$ onto G yields the desired set $F \subseteq E(G)$. ■

To prove Theorem I.3, we compute a simple approximate solution and remove all edges that are farther from a terminal than the size of this approximate solution. We then apply the same idea as in Theorem I.2 to each of the resulting connected components.

The idea behind the proof of Theorem I.4 is that the Edge Multiway Cut problem becomes a Steiner Forest-like problem in the dual graph. Hence, we cut open the dual of G similarly as we cut open G in Theorem I.2: for each terminal t of G, we take the cycle C_t in the dual of G that consists of all edges incident to t, and cut the dual along a short connected subgraph containing all cycles C_t for all terminals of G. We show that to preserve an optimal solution for Edge Multiway Cut in G it suffices to preserve an optimal Steiner tree for any choice of the terminals on the perimeter of the obtained brick. Hence, to apply Theorem I.1, we need to bound the length the perimeter, that is, the length of the subgraph of the dual of G that we cut along. By standard reductions, the total length of the cycles C_t (i.e., the total number of edges incident to terminals) is bounded by $2k_{OPT}$, where k_{OPT} is the optimal solution size. Hence, it suffices to bound the diameter of the dual of G.

To this end, we fix a terminal t and choose an inclusion-wise maximal laminar family of minimal separators that separate t from the remaining terminals and that are maximally "pushed away" from t (that is, they are important separators in the sense of [42]). By the "pushed away" property of the chosen family, each chosen separator is of different size, and as there are at most $2k_{OPT}$ edges incident to the terminals, the largest chosen separator is of size at most $2k_{OPT}$. Hence, there are $\mathcal{O}(k_{OPT}^2)$ edges in this chosen laminar family of minimal separators.

The essence of the proof is to show that an edge that is "far" from the chosen family of separators is irrelevant for the problem, and may be safely contracted. Here, "far" means ck_{OPT} for some universal constant c. Intuitively, if such an edge e is chosen in an optimal solution X, then the connected component of X of the dual of G that contains e lives between two separators from the chosen family, and we can show that it can be replaced by (a part of) one of these two separators.

Hence, after this reduction is performed exhaustively, the diameter of the dual of G is bounded by $\mathcal{O}(k_{OPT}^3)$. Consequently, cutting the graph open and applying Theorem I.1 leads to a polynomial kernel.

Using the extension of Theorem I.1 to graphs of bounded genus, we can extend Theorem I.2 and I.3, and part 1 of Corollary I.5 to such graphs.

IV. Overview of the proof of Theorem I.7

We now focus on the weighted variant, and sketch the proof of Theorem I.7.

We start by considering a base case, where $\mathcal{S}$ consists of a single terminal pair and H must contain a Steiner forest F_H that connects $\mathcal{S}$ such that $w(F_H) \leq (1+\varepsilon)w(F_B)$ for any Steiner forest F_B in B that connects $\mathcal{S}$. To this end, we first partition the input brick into *strips* [43]. Informally speaking, a *strip* is a brick whose perimeter can be partitioned into a shortest path (called the *south boundary*) and an "almost" shortest path (called the *north boundary*). Note that any Steiner tree in a strip that cannot be replaced by a part of the perimeter, even with a small loss in weight, needs to contain terminals both on the south and north boundary. We use this observation to provide an explicit construction of the graph H in a single strip, using so-called columns (similar to the columns introduced in [5]).

With the base case of a single terminal pair in mind, we move to the θ-variant of Theorem I.7, where $\mathcal{S}$ is allowed to contain only θ terminal pairs and the obtained bound for $w(H)$ depends polynomially both on ε^{-1} and θ. In this proof, we use the entire power of the structural results and decomposition methods developed for the proof of Theorem I.1, adjusted to the edge-weighted case. In short, we show that if we decompose each brick recursively into smaller bricks, stopping when the perimeter of the brick drops below some threshold $\mathrm{poly}(\varepsilon/\theta)w(\partial B)$, then we can take the single-pair graph H developed previously in each such small brick, and the union of all such graphs has the desired properties. The crux of the analysis is that the bound θ ensures that we can "buy" the entire perimeter of each small brick in which some vertex of degree at least three of an optimal Steiner forest of B is present.

Finally, we use the partitioning methods from the EPTAS [5], the so-called mortar graph framework, to derive Theorem I.7 from the θ-variant. The mortar graph constructed by [5] is essentially a brickable connector. We call the bricks induced by this connector *cells*. The mortar graph has the property that there exists a near-optimal Steiner forest in B that crosses each cell at most $\alpha(\varepsilon) = o(\varepsilon^{-5.5})$ times. Therefore, we construct the mortar graph of the input brick and then apply θ-variant to each cell independently, for an appropriate choice of $\theta = \text{poly}(\varepsilon^{-1})$. This then yields the desired graph H.

ACKNOWLEDGMENT

We thank Daniel Lokshtanov and Saket Saurabh for showing us the application of Baker's approach to PLANAR MULTIWAY CUT and for allowing us to include the proof in this paper. Moreover, we acknowledge the discussions with Daniel Lokshtanov that lead to the discovery that the NP-hardness proof for STEINER FOREST on planar graphs of treewidth 3 of Bateni et al. [23] can be strengthened to also refute a subexponential-time algorithm.

We would like also to acknowledge the support and extremely productive atmosphere at Dagstuhl Seminars 13121, 13421 and 14071.

REFERENCES

[1] E. D. Demaine and M. Hajiaghayi, "The bidimensionality theory and its algorithmic applications," *Comput. J.*, vol. 51, no. 3, pp. 292–302, 2008.

[2] E. D. Demaine, F. V. Fomin, M. Hajiaghayi, and D. M. Thilikos, "Subexponential parameterized algorithms on bounded-genus graphs and H-minor-free graphs," *J. ACM*, vol. 52, no. 6, pp. 866–893, 2005.

[3] F. V. Fomin, D. Lokshtanov, S. Saurabh, and D. M. Thilikos, "Bidimensionality and kernels," in *SODA*, 2010, pp. 503–510.

[4] P. N. Klein, "A linear-time approximation scheme for tsp in undirected planar graphs with edge-weights," *SIAM J. Comput.*, vol. 37, no. 6, pp. 1926–1952, 2008.

[5] G. Borradaile, P. N. Klein, and C. Mathieu, "An $O(n \log n)$ approximation scheme for Steiner tree in planar graphs," *ACM T. on Algorithms*, vol. 5, no. 3, 2009.

[6] R. G. Downey and M. R. Fellows, *Parameterized Complexity*. Springer, 1999.

[7] R. Niedermeier, *Invitation to Fixed-Parameter Algorithms*. Oxford University Press, 2006, vol. 31.

[8] H. L. Bodlaender, R. G. Downey, M. R. Fellows, and D. Hermelin, "On problems without polynomial kernels," *J. Comput. Syst. Sci.*, vol. 75, no. 8, pp. 423–434, 2009.

[9] L. Fortnow and R. Santhanam, "Infeasibility of instance compression and succinct PCPs for NP," *J. Comput. Syst. Sci.*, vol. 77, no. 1, pp. 91–106, 2011.

[10] A. Drucker, "New limits to classical and quantum instance compression," in *FOCS*. IEEE Computer Society, 2012, pp. 609–618.

[11] H. L. Bodlaender, F. V. Fomin, D. Lokshtanov, E. Penninkx, S. Saurabh, and D. M. Thilikos, "(Meta) kernelization," in *FOCS*, 2009, pp. 629–638.

[12] E. D. Demaine and M. Hajiaghayi, "Bidimensionality," in *Encyclopedia of Algorithms*, M.-Y. Kao, Ed. Springer, 2008.

[13] M. Dom, D. Lokshtanov, and S. Saurabh, "Incompressibility through colors and IDs," in *ICALP (1)*, 2009, pp. 378–389.

[14] M. R. Fellows, J. Guo, D. Marx, and S. Saurabh, "Data reduction and problem kernels (dagstuhl seminar 12241)," *Dagstuhl Reports*, vol. 2, no. 6, pp. 26–50, 2012.

[15] S. Kratsch and M. Wahlström, "Representative sets and irrelevant vertices: New tools for kernelization," in *FOCS*, 2012, pp. 450–459.

[16] M. Pilipczuk, M. Pilipczuk, P. Sankowski, and E. J. van Leeuwen, "Network sparsification for steiner problems on planar and bounded-genus graphs," *CoRR*, vol. abs/1306.6593, 2013.

[17] ——, "Subexponential-time parameterized algorithm for steiner tree on planar graphs," in *STACS*, 2013, pp. 353–364.

[18] E. D. Demaine and M. Hajiaghayi, "Linearity of grid minors in treewidth with applications through bidimensionality," *Combinatorica*, vol. 28, no. 1, pp. 19–36, 2008.

[19] N. Robertson, P. D. Seymour, and R. Thomas, "Quickly excluding a planar graph," *J. Comb. Theory, Ser. B*, vol. 62, no. 2, pp. 323–348, 1994.

[20] S. Kratsch and M. Wahlström, "Compression via matroids: a randomized polynomial kernel for odd cycle transversal," in *SODA*, 2012, pp. 94–103.

[21] S. Tazari, "Faster approximation schemes and parameterized algorithms on H-minor-free and odd-minor-free graphs," in *MFCS*, 2010, pp. 641–652.

[22] P. N. Klein and D. Marx, "Solving planar k-terminal cut in $O(n^{c\sqrt{k}})$ time," in *ICALP*, 2012, pp. 569–580.

[23] M. Bateni, M. Hajiaghayi, and D. Marx, "Approximation schemes for Steiner forest on planar graphs and graphs of bounded treewidth," *J. ACM*, vol. 58, no. 5, p. 21, 2011.

[24] R. Impagliazzo, R. Paturi, and F. Zane, "Which problems have strongly exponential complexity?" *J. Comput. Syst. Sci.*, vol. 63, no. 4, pp. 512–530, 2001.

[25] B. S. Baker, "Approximation algorithms for NP-complete problems on planar graphs," *J. ACM*, vol. 41, no. 1, pp. 153–180, 1994.

[26] J. Nederlof, "Fast polynomial-space algorithms using Möbius inversion: Improving on Steiner Tree and related problems," in *ICALP (1)*, 2009, pp. 713–725.

[27] M. Cygan, H. Dell, D. Lokshtanov, D. Marx, J. Nederlof, Y. Okamoto, R. Paturi, S. Saurabh, and M. Wahlström, "On problems as hard as CNF-SAT," in *IEEE CCC*, 2012, pp. 74–84.

[28] R. Karp, "On the computational complexity of combinatorial problems," *Networks*, vol. 5, pp. 45–68, 1975.

[29] E. Dahlhaus, D. S. Johnson, C. H. Papadimitriou, P. D. Seymour, and M. Yannakakis, "The complexity of multiterminal cuts," *SIAM J. Comput.*, vol. 23, no. 4, pp. 864–894, 1994.

[30] M. W. Bern and P. E. Plassmann, "The Steiner problem with edge lengths 1 and 2," *Inf. Process. Lett.*, vol. 32, no. 4, pp. 171–176, 1989.

[31] J. Byrka, F. Grandoni, T. Rothvoß, and L. Sanità, "An improved LP-based approximation for Steiner tree," in *STOC*, L. J. Schulman, Ed. ACM, 2010, pp. 583–592.

[32] A. Agrawal, P. N. Klein, and R. Ravi, "When trees collide: An approximation algorithm for the generalized Steiner problem on networks," *SIAM J. Comput.*, vol. 24, no. 3, pp. 440–456, 1995.

[33] D. R. Karger, P. N. Klein, C. Stein, M. Thorup, and N. E. Young, "Rounding algorithms for a geometric embedding of minimum multiway cut," *Math. Oper. Res.*, vol. 29, no. 3, pp. 436–461, 2004.

[34] Y. Cao, J. Chen, and J.-H. Fan, "An $O^*(1.84^k)$ parameterized algorithm for the multiterminal cut problem," in *FCT*, 2013, pp. 84–94.

[35] M. R. Garey and D. S. Johnson, "The Rectilinear Steiner Tree problem is NP complete," *SIAM J. Appl. Math.*, vol. 32, pp. 826–834, 1977.

[36] D. Eisenstat, P. N. Klein, and C. Mathieu, "An efficient polynomial-time approximation scheme for steiner forest in planar graphs," in *SODA*, 2012, pp. 626–638.

[37] M. Bateni, M. Hajiaghayi, P. N. Klein, and C. Mathieu, "A polynomial-time approximation scheme for planar multiway cut," in *SODA*, 2012, pp. 639–655.

[38] G. Borradaile, E. D. Demaine, and S. Tazari, "Polynomial-time approximation schemes for subset-connectivity problems in bounded-genus graphs," *Algorithmica*, vol. 68, no. 2, pp. 287–311, 2014.

[39] D. Marx, "A tight lower bound for planar multiway cut with fixed number of terminals," in *ICALP*, 2012, pp. 677–688.

[40] R. E. Erickson, C. L. Monma, and A. F. J. Veinott, "Send-and-split method for minimum-concave-cost network flows," *Mathematics of Operations Research*, vol. 12, no. 4, pp. 634–664, 1987.

[41] K. Borsuk, "Sur les rétractes," *Fundamenta Mathematicae*, vol. 17, pp. 2–20, 1931.

[42] D. Marx, "Parameterized graph separation problems," *Theor. Comput. Sci.*, vol. 351, no. 3, pp. 394–406, 2006.

[43] P. N. Klein, "A subset spanner for planar graphs, with application to Subset TSP," in *STOC*, J. M. Kleinberg, Ed. ACM, 2006, pp. 749–756.