Subexponential Parameterized Algorithms

Frederic Dorn[1,⋆], Fedor V. Fomin[1,⋆], and Dimitrios M. Thilikos[2,⋆⋆]

[1] Department of Informatics, University of Bergen, N-5020 Bergen, Norway
{frederic.dorn,fedor.fomin}@ii.uib.no
[2] Department of Mathematics, National & Capodistrian University of Athens, Panepistimioupolis, GR-15784, Athens, Greece
sedthilk@math.uoa.gr

Abstract. We present a series of techniques for the design of subexponential parameterized algorithms for graph problems. The design of such algorithms usually consists of two main steps: first find a branch- (or tree-) decomposition of the input graph whose width is bounded by a sublinear function of the parameter and, second, use this decomposition to solve the problem in time that is single exponential to this bound. The main tool for the first step is Bidimensionality Theory. Here we present the potential, but also the boundaries, of this theory. For the second step, we describe recent techniques, associating the analysis of sub-exponential algorithms to combinatorial bounds related to Catalan numbers. As a result, we have $2^{O(\sqrt{k})} \cdot n^{O(1)}$ time algorithms for a wide variety of parameterized problems on graphs, where n is the size of the graph and k is the parameter.

1 Introduction

The theory of fixed-parameter algorithms and parameterized complexity has been thoroughly developed during the last two decades; see e.g. the books [23,27,35]. Usually, parameterizing a problem on graphs is to consider its input as a pair consisting of the graph G itself and a parameter k. Typical examples of such parameters are the size of a vertex cover, the length of a path or the size of a dominating set. Roughly speaking, a parameterized problem in graphs with parameter k is *fixed parameter tractable* if there is an algorithm solving the problem in $f(k) \cdot n^{O(1)}$ steps for some function f that depends only on the parameter.

While there is strong evidence that most of fixed-parameter algorithms cannot have running times $2^{o(k)} \cdot n^{O(1)}$ (see [33,7,27]), for planar graphs it is possible to design subexponential parameterized algorithms with running times of the type $2^{O(\sqrt{k})} \cdot n^{O(1)}$ (see [9,7] for further lower bounds on planar graphs). For example, Planar k-Vertex Cover can be solved in $O(2^{3.57\sqrt{k}}) + O(n)$ steps,

---

⋆ Supported by the Norwegian Research Council.

⋆⋆ Supported by the Project "Capodistrias" (ΑΠ 736/24.3.2006) of the National and Capodistrian University of Athens (project code: 70/4/8757).

L. Arge et al. (Eds.): ICALP 2007, LNCS 4596, pp. 15–27, 2007.

Planar k-Dominating Set can be solved in $O(2^{11.98\cdot\sqrt{k}}) + O(n^3)$ steps, and Planar k-Longest Path can be solved in $O(2^{10.52\cdot\sqrt{k}}\cdot n)+O(n^3)$ steps. Similar algorithms are now known for a wide class of parameterized problems, not only for planar graphs, but also for several other sparse graph classes.

Since the first paper in this area appeared [2], the study of fast subexponential algorithms attracted a lot of attention. In fact, it not only offered a good ground for the development of parameterized algorithms, but it also prompted combinatorial results, of independent interest, on the structure of several parameters in sparse graph classes such as planar graphs [1,3,5,8,11,26,29,32,34] bounded genus graphs [12,28], graphs excluding some single-crossing graph as a minor [17], apex-minor-free graphs [10] and H-minor-free graphs [12,13,14].

We here present general approaches for obtaining subexponential parameterized algorithms (Section 2) and we reveal their relation with combinatorial results related to the Graph Minors project of Robertson and Seymour. All these algorithms exploit the structure of graph classes that exclude some graph as a minor. This was used to develop techniques such as Bidimensionality Theory (Section 3) and the use of Catalan numbers for better bounding the steps of dynamic programming when applied to minor closed graph classes (Sections 4 and 5).

2 Preliminaries

Given an edge $e = \{x, y\}$ of a graph G, the graph G/e is obtained from G by contracting the edge e; that is, to get G/e we identify the vertices x and y and remove all loops and replace all multiple edges by simple edges. A graph H obtained by a sequence of edge-contractions is said to be a *contraction* of G. H is a *minor* of G if H is a subgraph of a contraction of G. We use the notation $H \preceq G$ (resp. $H \preceq_c G$) when H is a minor (a contraction) of G. It is well known that $H \preceq G$ or $H \preceq_c G$ implies $\mathbf{bw}(H) \leq \mathbf{bw}(G)$. We say that a graph G is *H-minor-free* when it does not contain H as a minor. We also say that a graph class $\mathcal{G}$ is *H-minor-free* (or, excludes H as a minor) when all its members are H-minor-free. E.g., the class of planar graphs is a K_5-minor-free graph class.

Let G be a graph on n vertices. A *branch decomposition* (T, μ) of a graph G consists of an unrooted ternary tree T (i.e. all internal vertices of degree three) and a bijection $\mu : L \to E(G)$ from the set L of leaves of T to the edge set of G. We define for every edge e of T the *middle set* $\mathbf{mid}(e) \subseteq V(G)$ as follows: Let T_1 and T_2 be the two connected components of $T \setminus \{e\}$. Then let G_i be the graph induced by the edge set $\{\mu(f) : f \in L \cap V(T_i)\}$ for $i \in \{1, 2\}$. The *middle set* is the intersection of the vertex sets of G_1 and G_2, i.e., $\mathbf{mid}(e) := V(G_1) \cap V(G_2)$. The *width* $\mathbf{bw}$ of (T, μ) is the maximum order of the middle sets over all edges of T, i.e., $\mathbf{bw}(T, \mu) := \max\{|\mathbf{mid}(e)| : e \in T\}$. An optimal branch decomposition of G is defined by the tree T and the bijection μ which give the minimum width, the *branchwidth*, denoted by $\mathbf{bw}(G)$.

A *parameter* P is any function mapping graphs to nonnegative integers. The *parameterized problem associated with* P asks, for some fixed k, whether $P(G) = k$ for a given graph G. We say that a parameter P is *closed under taking of* *minors/contractions* (or, briefly, *minor/contraction closed*) if for every graph H, $H \preceq G$ / $H \preceq_c G$ implies that $P(H) \leq P(G)$.

Many subexponential parameterized graph algorithms [1,17,28,29,32,34] are associated with parameters P on graph classes $\mathcal{G}$ satisfying the following two conditions for some constants α and β:

(A) For every graph $G \in \mathcal{G}$, $\mathbf{bw}(G) \leq \alpha \cdot \sqrt{P(G)} + O(1)$
(B) For every graph $G \in \mathcal{G}$ and given a branch decomposition (T, μ) of G, the value of $P(G)$ can be computed in $2^{\beta \cdot \mathbf{bw}(T,\mu)} n^{O(1)}$ steps.

Conditions (A) and (B) are essential due to the following generic result.

Theorem 1. *Let P be a parameter and let $\mathcal{G}$ be a class of graphs such that (A) and (B) hold for some constants α and β respectively. Then, given a branch decomposition (T, μ) where $\mathbf{bw}(T, \mu) \leq \lambda \cdot \mathbf{bw}(G)$ for a constant λ, the parameterized problem associated with P can be solved in $2^{O(\sqrt{k})} n^{O(1)}$ steps.*

Proof. Given a branch decomposition (T, μ) as above, one can solve the parameterized problem associated with P as follows. If $\mathbf{bw}(T, \mu) > \lambda \cdot \alpha \cdot \sqrt{k}$, then the answer to the associated parameterized problem with parameter k is "NO" if it is a minimization and "YES" if it is a maximization problem. Else, by (B), $P(G)$ can be computed in $2^{\lambda \cdot \alpha \cdot \beta \cdot \sqrt{k}} n^{O(1)}$ steps.

To apply Theorem 1, we need an algorithm that computes, in time $t(n)$, a branch decomposition (T, μ) of any n-vertex graph $G \in \mathcal{G}$ such that $\mathbf{bw}(T, \mu) \leq \lambda \cdot \mathbf{bw}(G) + O(1)$. Because of [38], $t(n) = n^{O(1)}$ and $\lambda = 1$ for planar graphs. For H-minor-free graphs (and thus, for all graph classes considered here), $t(n) = f(|H|) \cdot n^{O(1)}$ and $\lambda \leq f(|H|)$ for some function f depending only on the size of H (see [16,21,24]).

In this survey we discuss how

- to obtain a general scheme of proving bounds required by (A) and to extend parameterized algorithms to more general classes of graphs like graphs of bounded genus and graphs excluding a minor (Section 3);
- to improve the running times of such algorithms (Section 4), and
- to prove that the running time of many dynamic programming algorithms on planar graphs (and more general classes as well) satisfies (B) (Section 5).

The following three sample problems capture the most important properties of the investigated parameterized problems.

k-Vertex Cover. A *vertex cover* C of a graph is a set of vertices such that every edge of G has at least one endpoint in C. The k-Vertex Cover problem is to decide, given a graph G and a positive integer k, whether G has a vertex cover of size k. Let us note that vertex cover is closed under taking minors, i.e. if a graph G has a vertex cover of size k, then each of its minors has a vertex cover of size at most k.

k-DOMINATING SET. A *dominating* set D of a graph G is a set of vertices such that every vertex outside D is adjacent to a vertex of D. The k-DOMINATING SET problem is to decide, given a graph G and a positive integer k, whether G has a dominating set of size k. Let us note that the dominating set is not closed under taking minors. However, it is closed under contraction of edges.

Given a branch decomposition of G of width $\leq \ell$ both problems k-VERTEX COVER and k-DOMINATING SET can be solved in time $2^{O(\ell)}n^{O(1)}$. For the next problem, no such an algorithm is known.
k-LONGEST PATH. The k-LONGEST PATH problem is to decide, given a graph G and a positive integer k, whether G contains a path of length k. This problem is closed under the operation of taking minor but the best known algorithm solving this problem on a graph of branchwidth $\leq \ell$ runs in time $2^{O(\ell \log \ell)}n^{O(1)}$.

3 Property (A) and Bidimensionality

In this section we show how to obtain subexponential parameterized algorithms in the case when condition (B) holds for general graphs. The main tool for this is Bidimensionality Theory developed in [10,12,13,15,18]. For a survey on Bidimensionality Theory see [14].

Planar graphs. While the results of this subsection can be extended to wider graph classes, we start from planar graphs where the general ideas are easier to explain. The following theorem is the main ingredient for proving condition (A).

Theorem 2 ([37]). *Let $\ell \geq 1$ be an integer. Every planar graph of branchwidth $\geq \ell$ contains an $(\ell/4 \times \ell/4)$-grid as a minor.*

We start with PLANAR k-VERTEX COVER as an example. Let G be a planar graph of branchwidth $\geq \ell$. Observe that given a $(r \times r)$-grid H, the size of a vertex cover in H is at least $\lfloor r/2 \rfloor \cdot r$ (because of the existence of a matching of size $\lfloor r/2 \rfloor \cdot r$ in H). By Theorem 2, we have that G contains an $(\ell/4 \times \ell/4)$-grid as a minor. The size of any vertex cover of this grid is at least $\ell^2/32$. As such a grid is a minor of G, property (A) holds for $\alpha = 4\sqrt{2}$.

For the PLANAR k-DOMINATING SET problem, the arguments used above to prove (A) for PLANAR k-VERTEX COVER do not work. Since the problem is not minor-closed, we cannot use Theorem 2 as above. However, since the parameter is closed under edge contractions, we can use a *partially triangulated $(r \times r)$-grid* which is any planar graph obtained from the $(r \times r)$-grid by adding some edges. For every partially triangulated $(r \times r)$-grid H, the size of a dominating set in H is at least $\frac{(r-2)^2}{9}$ (every "inner" vertex of H has a closed neighborhood of at most 9 vertices). Theorem 2 implies that a planar graph G of branchwidth $\geq \ell$ can be contracted to a partially triangulated $(\ell/4 \times \ell/4)$-grid which yields that PLANAR k-DOMINATING SET also satisfies (A) for $\alpha = 12$.

These two examples induce the following idea: if the graph parameter is closed under taking minors or contractions, the only thing needed for the proof of (A) is to understand how this parameter behaves on a (partially triangulated) grid. This brings us to the following definition.

Definition 1 ([12]). *A parameter P is* minor bidimensional with density δ *if*

1. *P is closed under taking of minors, and*
2. *for the $(r \times r)$-grid R, $P(R) = (\delta r)^2 + o((\delta r)^2)$.*

A parameter P is called contraction bidimensional with density δ *if*

1. *P is closed under contractions,*
2. *for any partially triangulated $(r \times r)$-grid R, $P(R) = (\delta_R r)^2 + o((\delta_R r)^2)$, and*
3. *δ is the smallest δ_R among all paritally triangulated $(r \times r)$-grids.*

In either case, P is called bidimensional. *The* density *δ of P is the minimum of the two possible densities (when both definitions are applicable), $0 < \delta \leq 1$.*

Intuitively, a parameter is bidimensional if its value depends on the area of a grid and not on its width or height. By Theorem 2, we have the following.

Lemma 1. *If P is a bidimensional parameter with density δ then P satisfies property (A) for $\alpha = 4/\delta$, on planar graphs.*

Many parameters are bidimensional. Some of them, like the number of vertices or the number of edges, are not so much interesting from an algorithmic point of view. Of course the already mentioned parameter *vertex cover* (*dominating set*) is minor (contraction) bidimensional (with densities $1/\sqrt{2}$ for *vertex cover* and $1/9$ for *dominating set*). Other examples of bidimensional parameters are *feedback vertex set* with density $\delta \in [1/2, 1/\sqrt{2}]$, *minimum maximal matching* with density $\delta \in [1/\sqrt{8}, 1/\sqrt{2}]$ and *longest path* with density 1.

By Lemma 1, Theorem 1 holds for every bidimensional parameter satisfying (B). Also, Theorem 1 can be applied not only to bidimensional parameters but to parameters that are bounded by bidimensional parameters. For example, the *clique-transversal number* of a graph G is the minimum number of vertices intersecting every maximal clique of G. This parameter is not contraction-closed because an edge contraction may create a new maximal clique and cause the clique-transversal number to increase. On the other hand, it is easy to see that this graph parameter always exceeds the size of a minimum dominating set which yields (A) for this parameter.

Non-planar extensions and limitations. One of the natural approaches of extending Lemma 1 from planar graphs to more general classes of graphs is via extending of Theorem 2. To do this we have to treat separately minor closed and contraction closed parameters.

The following extension of Theorem 2 holds for bounded genus graphs:

Theorem 3 ([12]). *If G is a graph of Euler genus at most γ with branchwidth more than r, then G contains a $(r/4(\gamma+1) \times r/4(\gamma+1))$-grid as a minor.*

Working analogously to the planar case, Theorem 3 implies the following.

Lemma 2. *Let P be a minor bidimensional parameter with density δ. Then for any graph G of Euler genus at most γ, property (A) holds for $\alpha = 4(\gamma+1)/\delta$.*

Next step is to consider graphs excluding a fixed graph H as a minor. The proof extends Theorem 3 by making (nontrivial) use of the structural characterization of H-minor-free graphs by Robertson and Seymour in [36].

Theorem 4 ([13]). *If G is an H-minor-free graph with branchwidth more than r, then G has the $(\Omega(r) \times \Omega(r))$-grid as a minor (the hidden constants in the Ω notation depend only on the size of H).*

As before, Theorem 3 implies property (A) for all minor bidimensional parameters for some α depending only on the excluded minor H.

For contraction-closed parameters, the landscape is different. In fact, each possible extension of Lemma 2, requires a stronger version of bidimensionality. For this, we can use the notion of a (r, q)-*gridoid* that is obtained from a partially triangulated $(r \times r)$-grid by adding at most q edges. (Note that every (r, q)-gridoid has genus $\leq q$.) The following extends Theorem 2 for graphs of bounded genus.

Theorem 5 ([12]). *If a graph G of Euler genus at most γ excludes all $(k - 12\gamma, \gamma)$-gridoids as contractions, for some $k \geq 12\gamma$, then G has branchwidth at most $4k(\gamma + 1)$.*

A parameter is *genus-contraction bidimensional* if *a*) it is contraction closed and *b*) its value on every $(r, O(1))$-gridoid is $\Omega(r^2)$ (here the hidden constants in the big-O and the big-Ω notations depend only on the Euler genus). Then Theorem 5 implies property (A) for all genus-contraction bidimensional parameters for some constant that depends only on the Euler genus.

An *apex graph* is a graph obtained from a planar graph G by adding a vertex and making it adjacent to some vertices of G. A graph class is *apex-minor-free* if it does not contain a graph with some fixed apex graph as a minor. An (r, s)-*augmented grid* is an $(r \times r)$-grid with some additional edges such that each vertex is attached to at most s vertices that in the original grid had degree 4. We say that a contraction closed parameter P is *apex-contraction bidimensional* if *a*) it is closed under taking of contractions and *b*) its value on every $(r, O(1))$-augmented grid is $\Omega(r^2)$ (here the hidden constants in the big-O and the big-Ω notations depend only on the excluded apex graph). According to [10] and [13], every apex-contraction bidimensional parameter satisfies property (A) for some constant that depends only on the excluded apex graph.

A natural question appears: until what point property (A) can be satisfied for contraction-closed parameters (assuming a suitable concept of bidimensionality)? As it was observed in [10], for some contraction-closed parameters, like dominating set, the branchwidth of an apex graph cannot be bounded by any function of their values. Consequently, apex-free graph classes draw a natural combinatorial limit on the the above framework of obtaining subexponential parameterized algorithms for contraction-closed parameters. (On the other side, this is not the case for minor-closed parameters as indicated by Theorem 4.) However, it is still possible to cross the frontier of apex-minor-free graphs for the dominating set problem and some of its variants where subexponential parameterized algorithms exist, even for H-minor-free graphs, as shown in [12].

These algorithms are based on a combination of dynamic programming and the the structural characterization of H-minor-free graphs from [36].

4 Further Optimizations

In this section, we present several techniques for accelerating the algorithms emerging by the framework of Theorem 1.

Making algorithms faster. While proving properties (A) and (B), it is natural to ask for the best possible constants α and β, as this directly implies an exponential speed-up of the corresponding algorithms. While, Bidimensionality Theory provides some general estimation of α, in some cases, deep understanding of the parameter behavior can lead to much better constants in (A). For example, it holds that for Planar k-Vertex Cover, $\alpha \leq 3$ (see [30]) and for Planar k-Dominating Set, $\alpha \leq 6.364$ (see [29]). (Both bounds are based on the fact that planar graphs with n vertices have branchwidth at most $\sqrt{4.5}\sqrt{n}$, see [30].) Similar results hold also for bounded genus graphs [28].

On the other hand, there are several ways to obtain faster dynamic programming algorithms and to obtain better bounds for β in (B). A typical approach to compute a solution of size k works as follows:

– Root the branch decomposition (T, μ) of graph G picking any of the vertices of its tree and apply dynamic programming on the middle sets, bottom up, from the leaves towards the root.
– Each middle set $\mathbf{mid}(e)$ of (T, μ) represents the subgraph G_e of G induced by the leaves below. Recall that the vertices of $\mathbf{mid}(e)$ are separators of G.
– In each step of the dynamic programming, all optimal solutions for a subproblem in G_e are computed, subject to all possibilities of how $\mathbf{mid}(e)$ contributes to an overall solution for G. E.g., for Vertex Cover, there are up to $2^{\mathbf{bw}(T,\mu)}$ subsets of $\mathbf{mid}(e)$ that may constitute a vertex cover of G. Each subset is associated with an optimal solution in G_e with respect to this subset.
– The partial solutions of a middle set are computed using those of the already processed middle sets of the children and stored in an appropriate data structure.
– An optimal solution is computed at the root of T.

Encoding the middle sets in a refined way, may speed up the processing time significantly. Though, the same time is needed to scan all solutions assigned to a $\mathbf{mid}(e)$ after introducing vertex states, there are some methods to accelerate the update of the solutions of two middle sets to a parent middle set:

Using the right data structure: storing the solutions in a sorted list compensates the time consuming search for compatible solutions and allows a fast computing of the new solution. E.g., for k-Vertex Cover, the time to process two middle sets is reduced from $O(2^{3\cdot\mathbf{bw}(T,\mu)})$ (for each subset of the parent middle set, all pairs of solutions of the two children are computed) to $O(2^{1.5\cdot\mathbf{bw}(T,\mu)})$. In [19] matrices are used as a data structure for dynamic programming that allows an updating even in time $O(2^{\frac{\omega}{2}\mathbf{bw}(T,\mu)})$ for k-Vertex Cover (where ω is the fast matrix multiplication constant, actually $\omega < 2.376$).

A compact encoding: assign as few as possible vertex states to the vertices and reduce the number of processed solutions. Alber et al. [1], using the so-called "monotonicity technique", show that 3 vertex states are sufficient in order to encode a solution of k-Dominating Set. A similar approach was used in [29] to obtain, for the same problem, a $O(3^{1.5 \cdot \mathbf{bw}(T,\mu)})$-step updating process, that has been improved by [19] to $O(2^{2 \cdot \mathbf{bw}(T,\mu)})$.

Employing graph structures: as we will see in the Section 5, one can improve the runtime further for dynamic programming on branch decompositions whose middle sets inherit some structure of the graph. Using such techniques, the update process for Planar k-Dominating Set is done in time $O(3^{\frac{\omega}{2}\mathbf{bw}(T,\mu)})$ [19].

The above techniques can be used to prove the following result.

Theorem 6 ([19]). Planar k-Vertex Cover *can be solved in* $O(2^{3.56\sqrt{k}}) \cdot n^{O(1)}$ *runtime and* Planar k-Dominating Set *in* $O(2^{11.98\sqrt{k}}) \cdot n^{O(1)}$ *runtime.*

Kernels. Many of the parameterized algorithms discussed in this section can be further accelerated to time $O(n^{\theta}) + 2^{O(\sqrt{k})}$ for θ being a small integer (usually ranging from 1 to 3). This can be done using the technique of kernelization that is a prolynomial step preprocessing of the initial input of the problem towards creating an equivalent one, whose size depends exclusively on the parameter. Examples of such problems are Planar k-Dominating Set [4,8,28], k-Feedback Vertex Set [6], k-Vertex Cover and others [25]. See the books of [23,27,35] for a further reference.

5 Property (B) and Catalan Structures

All results of the previous sections provide subexponential parameterized algorithms when property (B) holds. However, there are many bidimensional parameters for which there is no known algorithm providing property (B) in general. The typical running times of dynamic programming algorithms for these problems are $O(\mathbf{bw}(G)!) \cdot n^{O(1)}$, $O(\mathbf{bw}(G)^{\mathbf{bw}(G)}) \cdot n^{O(1)}$, or even $O(2^{\mathbf{bw}(G)^2}) \cdot n^{O(1)}$. Examples of such problems are parameterized versions of k-Longest Path, k-Feedback Vertex Set, k-Connected Dominating Set, and k-Graph TSP. Usually, these are problems in NP whose certificate verifications involves some connectivity question. In this section, we show that for such problems one can prove that (B) actually holds for the graph class that we are interested in. To do this, one has to make further use of the structural properties of the class (again from the Graph Minors Theory) that can vary from planar graphs to H-minor-free graphs. In other words, we use the structure of the graph class not only for proving (A) but also for proving (B).

Planar graphs. The following type of decomposition for planar graphs follows from a proof by Seymour and Thomas in [38] and is extremely useful for making dynamic programming on graphs of bounded branchwidth faster (see [22,19]).

Let G be a planar graph embedded in a sphere $\mathbb{S}$. An O-arc is a subset of $\mathbb{S}$ homeomorphic to a circle. An O-arc in $\mathbb{S}$ is called a *noose* of the embedding of G if it meets G only in vertices. The length of a noose O is the number of vertices of G it meets. Every noose O bounds two open discs Δ_1, Δ_2 in $\mathbb{S}$, i.e., $\Delta_1 \cap \Delta_2 = \emptyset$ and $\Delta_1 \cup \Delta_2 \cup O = \mathbb{S}$.

We define a *sphere cut decomposition* or *sc-decomposition* (T, μ, π) as a branch decomposition with the following property: for every edge e of T, there exists a noose O_e meeting every face at most once and bounding the two open discs Δ_1 and Δ_2 such that $G_i \subseteq \Delta_i \cup O_e$, $1 \leq i \leq 2$. Thus O_e meets G only in $\mathbf{mid}(e)$ and its length is $|\mathbf{mid}(e)|$. A *clockwise traversal* of O_e in the embedding of G defines the cyclic ordering π of $\mathbf{mid}(e)$. We always assume that the vertices of every middle set $\mathbf{mid}(e) = V(G_1) \cap V(G_2)$ are enumerated according to π.

Theorem 7. *Let G be a planar graph of branchwidth at most ℓ without vertices of degree one embedded on a sphere. Then there exists an sc-decomposition of G of width at most ℓ that can be constructed in time $O(n^3)$.*

In what follows, we sketch the main idea of a $2^{O(\mathbf{bw}(T,\mu,\pi))}n^{O(1)}$ algorithm for the k-Planar Longest Path. One may use k-Longest path as an exemplar for other problems of the same nature.

Let G be a graph and let $E \subseteq E(G)$ and $S \subseteq V(G)$. To count the number of states at each step of the dynamic programming, we should estimate the number of collections of internally vertex disjoint paths using edges from E and having their (different) endpoints in S. We use the notation $\mathbf{P}$ to denote such a path collection and we define $\mathbf{paths}_G(E, S)$ as the set of all such path collections. Define equivalence relation $\sim$ on $\mathbf{paths}_G(E, S)$: for $\mathbf{P}_1, \mathbf{P}_2 \in \mathbf{paths}_G(E, S)$, $\mathbf{P}_1 \sim \mathbf{P}_2$ if there is a bijection between $\mathbf{P}_1$ and $\mathbf{P}_2$ such that bijected paths in $\mathbf{P}_1$ and $\mathbf{P}_2$ have the same endpoints. Denote by $\texttt{q-paths}_G(E, S) = |\mathbf{paths}_G(E, S)/\sim|$ the cardinality of the quotient set of $\mathbf{paths}_G(E, S)$ by $\sim$.

Recall that we define $\texttt{q-paths}_G(E, S)$ because, while applying dynamic programming on some middle set $\mathbf{mid}(e)$ of the branch decomposition (T, μ), the number of states for $e \in E(T)$ is bounded by $O(\texttt{q-paths}_{G_i}(E(G_i), \mathbf{mid}(e)))$.

Given a graph G and a branch decomposition (T, μ) of G, we say that (T, μ) has *Catalan structure* if for every edge $e \in E(T)$ and any $i \in \{1, 2\}$,

$$\texttt{q-paths}_{G_i}(E(G_i), \mathbf{mid}(e)) = 2^{O(\mathbf{bw}(T,\mu))} \tag{1}$$

Now, (B) holds for planar graphs because of the following combinatorial result.

Theorem 8 ([22]). *Every planar graph has an optimal branch decomposition with the Catalan structure that can be constructed in polynomial time.*

The proof of Theorem 8 uses an sc-decomposition (T, μ, π) (constructed using the polynomial algorithm in [38]). Let O_e be a noose meeting some middle set $\mathbf{mid}(e)$ of (T, μ, π). Let us count in how many ways this noose can cut paths of G. Observe that each path is cut into at most $\mathbf{bw}(T, \mu, \pi)$ parts. Each such part is itself a path whose endpoints are pairs of vertices in O_e. Notice also that, because of planarity, no two such pairs can cross. Therefore, counting the ways O_e can intersect paths of G is equivalent to counting non-crossing pairs of vertices in a cycle (the noose) of length $\mathbf{bw}(T, \mu, \pi)$ which, in turn, is bounded by the Catalan number of $\mathbf{bw}(T, \mu, \pi)$ that is $2^{O(\mathbf{bw}(T,\mu,\pi))}$.

We just concluded that the application of dynamic programming on an sc-decomposition (T, μ, π) is the $2^{O(\mathbf{bw}(T,\mu,\pi))} n^{O(1)}$ algorithm for proving property (B) for planar graphs. By further improving the way the members of $\texttt{q-paths}_{G_i}(E(G_i), \mathbf{mid}(e))$ are encoded during this procedure, one can bound the hidden constants in the big-O notation on the exponent of this algorithm (see [22]). For example, for Planar k-Longest Path $\beta \leq 2.63$. With analogous structures and arguments it follows that for Planar k-Graph TSP $\beta \leq 3.84$, for Planar k-Connected Dominating Set $\beta \leq 3.82$, for Planar k-Feedback Vertex Set $\beta \leq 3.56$ [19].

In [20], all above results were generalized for graphs with bounded genus (now constants for each problem depend also on the genus). This generalization requires a suitable "bounded genus"-extension of Theorem 8 and its analogues for other problems.

Excluding a minor. The final step is to prove property (B) for H-minor-free graphs. For the proof of this, we need the following analogue of Theorem 8.

Theorem 9 ([21]). *Let $\mathcal{G}$ be a graph class excluding some fixed graph H as a minor. Then every graph $G \in \mathcal{G}$ with $\mathbf{bw}(G) \leq \ell$ has an branch decomposition of width $O(\ell)$ with the Catalan structure (here the hidden constants in the big-O notations in $O(\ell)$ and the upper bound certifying the Catalan structure in Equation* (1) *depend only on H). Moreover, such a decomposition can be constructed in $f(|H|) \cdot n^{O(1)}$ steps, where f is a function depending only on H.*

The proof of Theorem 9 is based on an algorithm constructing the claimed branch decomposition using a structural characterization of H-minor-free graphs, given in [36]. Briefly, any H-minor-free graph can be seen as the result of gluing together (identifying constant size cliques and, possibly, removing some of their edges) graphs that, after the removal of some constant number of vertices (called *apices*) can be "almost" embedded in a surface of constant genus. Here, by "almost" we mean that we permit a constant number of non-embedded parts (called *vortices*) that are "attached" around empty disks of the embedded part and have a path-like structure of constant width. The algorithm of Theorem 9, as well as the proof of its correctness, has several phases, each dealing with some level of this characterisation, where an analogue of sc-decomposition for planar graphs is used. The core of the proof is based on the fact that the structure of the embeddible parts of this characterisation (along with vortices) is "close enough" to be plane, so to roughly maintain the Catalan structure property.

Theorem 9 implies (B) for k-Longest Path on H-minor-free graphs. Similar results can be obtained for all problems examined in this section on H-minor-free graphs. Since property (A) holds for minor/apex-contraction bidimensional parameters on H-minor-free/apex-minor-free graphs, we have that one can design $2^{O(\sqrt{k})} \cdot n^{O(1)}$ step parameterized algorithms for all problems examined in this section for H-minor-free/ apex-minor-free graphs (here the hidden constant in the big-O notation in the exponent depend on the size on the excluded minor).

6 Conclusion

In Section 3, we have seen that bidimensionality can serve as a general combinatorial criterion implying property (A). Moreover, no such a characterization is known, so far, for proving property (B). In Section 5, we have presented several problems where an analogue of Theorem 9 can be proven, indicating the existence of Catalan structures in H-minor-free graphs. It would be challenging to find a classification criterion (logical or combinatorial) for the problems that are amenable to this approach.

Acknowledgements. We take the opportunity to warmly thank Hans Bodlaender, Erik Demaine, MohammadTaghi Hajiaghayi, Naomi Nishimura, Prabhakar Ragde and Eelko Penninkx for collaborating with us on this project.

References

1. Alber, J., Bodlaender, H.L., Fernau, H., Kloks, T., Niedermeier, R.: Fixed parameter algorithms for dominating set and related problems on planar graphs. Algorithmica 33, 461–493 (2002)
2. Alber, J., Bodlaender, H.L., Fernau, H., Niedermeier, R.: Fixed parameter algorithms for planar dominating set and related problems. In: Halldórsson, M.M. (ed.) SWAT 2000. LNCS, vol. 1851, pp. 97–110. Springer, Heidelberg (2000)
3. Alber, J., Fan, H., Fellows, M.R., Fernau, H., Niedermeier, R., Rosamond, F.A., Stege, U.: Refined search tree technique for dominating set on planar graphs. In: Sgall, J., Pultr, A., Kolman, P. (eds.) MFCS 2001. LNCS, vol. 2136, pp. 111–122. Springer, Heidelberg (2001)
4. Alber, J., Fellows, M.R., Niedermeier, R.: Polynomial-time data reduction for dominating set. Journal of the ACM 51, 363–384 (2004)
5. Alber, J., Fernau, H., Niedermeier, R.: Parameterized complexity: exponential speed-up for planar graph problems. J. Algorithms 52, 26–56 (2004)
6. Bodlaender, H.: A cubic kernel for feedback vertex set. In: Thomas, W., Weil, P. (eds.) STACS 2007. LNCS, vol. 4393, Springer, Heidelberg (2007) (to appear)
7. Cai, L., Juedes, D.: On the existence of subexponential parameterized algorithms. J. Comput. System Sci. 67, 789–807 (2003)
8. Chen, J., Fernau, H., Kanj, I.A., Xia, G.: Parametric duality and kernelization: Lower bounds and upper bounds on kernel size. In: Ganter, B., Godin, R. (eds.) ICFCA 2005. LNCS (LNAI), vol. 3403, pp. 269–280. Springer, Heidelberg (2005)
9. [ERROR while converting LaTeX/Unicode], V.G., Klinz, B., Woeginger, G.J.: Exact algorithms for the Hamiltonian cycle problem in planar graphs. Oper. Res. Lett. 34, 269–274 (2006)
10. Demaine, E.D., Fomin, F.V., Hajiaghayi, M., Thilikos, D.M.: Bidimensional parameters and local treewidth. SIAM J. Discrete Math. 18, 501–511 (2004)
11. Demaine, E.D., Fomin, F.V., Hajiaghayi, M., Thilikos, D.M.: Fixed-parameter algorithms for (k,r)-center in planar graphs and map graphs. ACM Trans. Algorithms 1, 33–47 (2005)
12. Demaine, E.D., Fomin, F.V., Hajiaghayi, M., Thilikos, D.M.: Subexponential parameterized algorithms on graphs of bounded genus and H-minor-free graphs. Journal of the ACM 52, 866–893 (2005)

13. Demaine, E.D., Hajiaghayi, M.: Linearity of grid minors in treewidth with applications through bidimensionality. Combinatorica (to appear)
14. Demaine, E.D., Hajiaghayi, M.: The bidimensionality theory and its algorithmic applications. Computer Journal (to appear)
15. Demaine, E.D., Hajiaghayi, M.: Bidimensionality: new connections between FPT algorithms and PTASs. In: SODA'05, pp. 590–601. ACM-SIAM, New York (2005)
16. Demaine, E.D., Hajiaghayi, M.T., Nishimura, N., Ragde, P., Thilikos, D.M.: Approximation algorithms for classes of graphs excluding single-crossing graphs as minors. J. Comput. System Sci. 69, 166–195 (2004)
17. Demaine, E.D., Hajiaghayi, M.T., Thilikos, D.M.: Exponential speedup of fixed-parameter algorithms for classes of graphs excluding single-crossing graphs as minors. Algorithmica 41, 245–267 (2005)
18. Demaine, E.D., Hajiaghayi, M.T., Thilikos, D.M.: The bidimensional theory of bounded-genus graphs. SIAM J. Discrete Math. 20, 357–371 (2006)
19. Dorn, F.: Dynamic programming and fast matrix multiplication. In: Azar, Y., Erlebach, T. (eds.) ESA 2006. LNCS, vol. 4168, pp. 280–291. Springer, Heidelberg (2006)
20. Dorn, F., Fomin, F.V., Thilikos, D.M.: Fast subexponential algorithm for non-local problems on graphs of bounded genus. In: Arge, L., Freivalds, R. (eds.) SWAT 2006. LNCS, vol. 4059, pp. 172–183. Springer, Heidelberg (2006)
21. Dorn, F., Fomin, F.V., Thilikos, D.M.: Catalan structures and dynamic programming on H-minor-free graphs, manuscript (2007)
22. Dorn, F., Penninkx, E., Bodlaender, H., Fomin, F.V.: Efficient exact algorithms on planar graphs: Exploiting sphere cut branch decompositions. In: Brodal, G.S., Leonardi, S. (eds.) ESA 2005. LNCS, vol. 3669, pp. 95–106. Springer, Heidelberg (2005)
23. Downey, R.G., Fellows, M.R.: Parameterized complexity. Springer, New York (1999)
24. Feige, U., Hajiaghayi, M., Lee, J.R.: Improved approximation algorithms for minimum-weight vertex separators. In: STOC'05, pp. 563–572. ACM Press, New York (2005)
25. Fellows, M.R.: Blow-ups, win/win's, and crown rules: Some new directions in FPT. In: Bodlaender, H.L. (ed.) WG 2003. LNCS, vol. 2880, pp. 1–12. Springer, Heidelberg (2003)
26. Fernau, H., Juedes, D.W.: A geometric approach to parameterized algorithms for domination problems on planar graphs. In: Fiala, J., Koubek, V., Kratochvíl, J. (eds.) MFCS 2004. LNCS, vol. 3153, pp. 488–499. Springer, Heidelberg (2004)
27. Flum, J., Grohe, M.: Parameterized Complexity Theory. Springer, Berlin (2006)
28. Fomin, F.V., Thilikos, D.M.: Fast parameterized algorithms for graphs on surfaces: Linear kernel and exponential speed-up. In: Díaz, J., Karhumäki, J., Lepistö, A., Sannella, D. (eds.) ICALP 2004. LNCS, vol. 3142, pp. 581–592. Springer, Heidelberg (2004)
29. Fomin, F.V., Thilikos, D.M.: Dominating sets in planar graphs: Branch-width and exponential speed-up. SIAM J. Comput. 36, 281–309 (2006)
30. Fomin, F.V., Thilikos, D.M.: New upper bounds on the decomposability of planar graphs. Journal of Graph Theory 51, 53–81 (2006)
31. Grohe, M.: Local tree-width, excluded minors, and approximation algorithms. Combinatorica 23, 613–632 (2003)
32. Gutin, G., Kloks, T., Lee, C.M., Yeo, A.: Kernels in planar digraphs. J. Comput. System Sci. 71, 174–184 (2005)

33. Impagliazzo, R., Paturi, R., Zane, F.: Which problems have strongly exponential complexity. Journal of Computer and System Sciences 63, 512–530 (2001)
34. Kanj, I., Perković, L.: Improved parameterized algorithms for planar dominating set. In: Diks, K., Rytter, W. (eds.) MFCS 2002. LNCS, vol. 2420, pp. 399–410. Springer, Heidelberg (2002)
35. Niedermeier, R.: Invitation to fixed-parameter algorithms. Oxford University Press, Oxford (2006)
36. Robertson, N.: Graph minors. XVI. Excluding a non-planar graph. J. Combin. Theory Ser. B 89, 43–76 (2003)
37. Robertson, N., Seymour, P.D., Thomas, R.: Quickly excluding a planar graph. J. Combin. Theory Ser. B 62, 323–348 (1994)
38. Seymour, P.D., Thomas, R.: Call routing and the ratcatcher. Combinatorica 14, 217–241 (1994)