

Finding Detours is Fixed-Parameter Tractable^{*†}

Ivona Bezáková¹, Radu Curticapean², Holger Dell³, and
Fedor V. Fomin⁴

- 1 Department of Computer Science, Rochester Institute of Technology,
Rochester, NY, USA
ib@cs.rit.edu
- 2 Institute for Computer Science and Control, Hungarian Academy of Sciences
(MTA SZTAKI), Budapest, Hungary
radu.curticapean@gmail.com
- 3 Saarland University and Cluster of Excellence, MMCI, Saarbrücken, Germany
hdell@mmci.uni-saarland.de
- 4 University of Bergen, Bergen, Norway
fomin@ii.uib.no

Abstract

We consider the following natural “above guarantee” parameterization of the classical LONGEST PATH problem: For given vertices s and t of a graph G , and an integer k , the problem LONGEST DETOUR asks for an (s, t) -path in G that is at least k longer than a shortest (s, t) -path. Using insights into structural graph theory, we prove that LONGEST DETOUR is fixed-parameter tractable (FPT) on undirected graphs and actually even admits a single-exponential algorithm, that is, one of running time $\exp(O(k)) \cdot \text{poly}(n)$. This matches (up to the base of the exponential) the best algorithms for finding a path of length at least k .

Furthermore, we study the related problem EXACT DETOUR that asks whether a graph G contains an (s, t) -path that is exactly k longer than a shortest (s, t) -path. For this problem, we obtain a randomized algorithm with running time about $2.746^k \cdot \text{poly}(n)$, and a deterministic algorithm with running time about $6.745^k \cdot \text{poly}(n)$, showing that this problem is FPT as well. Our algorithms for EXACT DETOUR apply to both undirected and directed graphs.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems, G.2.2 Graph Theory

Keywords and phrases longest path, fixed-parameter tractable algorithms, above-guarantee parameterization, graph minors

Digital Object Identifier 10.4230/LIPIcs.ICALP.2017.54

1 Introduction

The LONGEST PATH problem asks, given an undirected n -vertex graph G and an integer k , to decide whether G contains a path of length at least k , that is, a self-avoiding walk with at least k edges. This problem is a natural generalization of the classical NP-complete HAMILTONIAN PATH problem, and the parameterized complexity community has paid exceptional attention to it. For instance, Monien [28] and Bodlaender [4] showed *avant*

* Full version at <https://arxiv.org/abs/1607.07737>.

† Most of this work was done while the authors were visiting the Simons Institute for the Theory of Computing. IB is supported by NSF grant CCF-1319987. RC is supported by ERC grant PARAMTIGHT (No. 280152).



© Ivona Bezáková, Radu Curticapean, Holger Dell, and Fedor V. Fomin;
licensed under Creative Commons License CC-BY

44th International Colloquium on Automata, Languages, and Programming (ICALP 2017).

Editors: Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl;

Article No. 54; pp. 54:1–54:14



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



la lettre that LONGEST PATH is fixed-parameter tractable with parameter k and admits algorithms with running time $2^{O(k \log k)} n^{O(1)}$. This led Papadimitriou and Yannakakis [29] to conjecture that LONGEST PATH is solvable in polynomial time for $k = \log n$, and indeed, this conjecture was resolved in a seminal paper of Alon, Yuster, and Zwick [2], who introduced the method of color coding and derived from it the first algorithm with running time $2^{O(k)} n$. Since this breakthrough of Alon et al. [2], the problem LONGEST PATH occupied a central place in parameterized algorithmics, and several novel approaches were developed in order to reduce the base of the exponent in the running time [19, 22, 9, 8, 23, 33, 15, 15, 3]. We refer to two review articles in Communications of ACM [14, 24] as well as to the textbook [12, Chapter 10] for an extensive overview of parameterized algorithms for LONGEST PATH. Let us however note that the fastest known randomized algorithm for LONGEST PATH is due to Björklund et al. [3] and runs in time $1.657^k \cdot n^{O(1)}$, whereas the fastest known deterministic algorithm is due to Zehavi [34] and runs in time $2.597^k \cdot n^{O(1)}$.

In the present paper, we study the problem LONGEST PATH from the perspective of an “above guarantee” parameterization that can attain small values even for long paths: For a pair of vertices $s, t \in V(G)$, we use $d_G(s, t)$ to denote the distance, that is, the length of a shortest path from s to t . We then ask for an (s, t) -path of length at least $d_G(s, t) + k$, and we parameterize by this offset k rather than the actual length of the path to obtain the problem LONGEST DETOUR. In other words, the first $d_G(s, t)$ steps on a path sought by LONGEST DETOUR are complimentary and will not be counted towards the parameter value. This reflects the fact that shortest paths can be found in polynomial time and could (somewhat embarrassingly) be much better solutions for LONGEST PATH than the paths of logarithmic length found by algorithms that parameterize by the path length.

We study two variants of the detour problem, one asking for a detour of length at least k , and another asking for a detour of length exactly k .

LONGEST DETOUR

Parameter: k

Input: Graph G , vertices $s, t \in V(G)$, and integer k .

Task: Decide whether there is an (s, t) -path in G of length at least $d_G(s, t) + k$.

EXACT DETOUR

Parameter: k

Input: Graph G , vertices $s, t \in V(G)$, and integer k .

Task: Decide whether there is an (s, t) -path in G of length exactly $d_G(s, t) + k$.

Our parameterization above the length of a shortest path is a new example in the general paradigm of “above guarantee” parameterizations, which was introduced by Mahajan and Raman [26]. Their approach was successfully applied to various problems, such as finding independent sets in planar graphs (where an independent set of size at least $\frac{n}{4}$ is guaranteed to exist by the Four Color Theorem), or the maximum cut problem, see e.g. [1, 11, 17, 16, 27].

Our results

We show the following tractability results for LONGEST DETOUR and EXACT DETOUR:

- LONGEST DETOUR is fixed-parameter tractable (FPT) on undirected graphs. The running time of our algorithm is single-exponential, i.e., it is of the type $2^{O(k)} \cdot n^{O(1)}$ and thus asymptotically matches the running time of algorithms for LONGEST PATH. Our approach requires a non-trivial argument in graph structure theory to obtain the single-exponential algorithm; a mere FPT-algorithm could be achieved with somewhat less effort. It should also be noted that a straightforward reduction rules out a running time of $2^{o(k)} \cdot n^{O(1)}$ unless the exponential-time hypothesis of Impagliazzo and Paturi [20] fails.

- EXACT DETOUR is FPT on directed and undirected graphs. Actually, we give a polynomial-time Turing reduction from EXACT DETOUR to the standard parameterization of LONGEST PATH, in which we ask on input u, v and $k \in \mathbf{N}$ whether there is a (u, v) -path of length k . This reduction only makes queries to instances with parameter at most $2k + 1$. Pipelined with the fastest known algorithms for LONGEST PATH mentioned above, this implies that EXACT DETOUR admits a bounded-error randomized algorithm with running time $2.746^k n^{O(1)}$, and a deterministic algorithm with running time $6.745^k n^{O(1)}$.

By a self-reducibility argument, we also show how to construct the required paths rather than just detect their existence. This reduction incurs only polynomial overhead.

Techniques

The main idea behind the algorithm for LONGEST DETOUR is the following combinatorial theorem, which shows the existence of specific large planar minors in large-treewidth graphs while circumventing the full machinery used in the Excluded Grid Theorem [31]. Although the Excluded Grid Theorem already shows that graphs of sufficiently large treewidth contain arbitrary fixed planar graphs, resorting to more basic techniques allows us to show that linear treewidth suffices for our specific cases. More specifically, we show that there exists a global constant $c \in \mathbf{N}$ such that every graph of treewidth at least $c \cdot k$ contains as a subgraph a copy of a graph $K_4^{\geq k}$, which is any graph obtained from the complete graph K_4 by replacing every edge by a path with at least k edges. The proof of this result is based on the structural theorems of Leaf and Seymour [25] and Raymond and Thilikos [30].

With the combinatorial theorem at hand, we implement the following win/win approach: If the treewidth of the input graph is less than $c \cdot k$, we use known algorithms [5, 15] to solve the problem in single-exponential time. Otherwise the treewidth of the input graph is at least $c \cdot k$ and there must be a $K_4^{\geq k}$, which we use to argue that any path visiting the same two-connected component as $K_4^{\geq k}$ can be prolonged by rerouting it through $K_4^{\geq k}$. To this end, we set up a fixed system of linear inequalities corresponding to the possible paths in $K_4^{\geq k}$ such that rerouting is possible if and only if the system is unsatisfiable. We then verify the unsatisfiability of this fixed system by means of a computer-aided proof (more specifically, a linear programming solver). From LP duality, we also obtain a short certificate for the unsatisfiability, which we include in the full version of this extended abstract.

The algorithm for EXACT DETOUR is based on the following idea. We run breadth-first search (BFS) from vertex v to vertex u . Then, for every (u, v) -path P of length $d_G(u, v) + k$, all but at most k levels of the BFS-tree contain exactly one vertex of P . Using this property, we are able to devise a dynamic programming algorithm for EXACT DETOUR, provided it is given access to an oracle for LONGEST PATH.

The remaining part of the paper is organized as follows. Section 2 contains definitions and preliminary results used in the technical part of the paper. In Section 3, we give an algorithm for LONGEST DETOUR while Section 4 is devoted to EXACT DETOUR. Due to space constraints, we defer some proofs, some figures, a search-to-decision reduction for LONGEST DETOUR and EXACT DETOUR to the full version of this extended abstract. Statements whose proofs are omitted here are marked with \star .

2 Preliminaries

We consider graphs G to be undirected, and we denote by uv an undirected edge joining vertices $u, v \in V(G)$. A *path* is a self-avoiding walk in G ; the *length* of the path is its number of edges. An (s, t) -path for $s, t \in V(G)$ is a path that starts at s and ends at t . We allow

paths to have length 0, in which case $s = t$ holds. For a vertex set $X \subseteq V(G)$, denote by $G[X]$ the subgraph induced by X .

Tree decompositions. A *tree decomposition* \mathcal{T} of a graph G is a pair $(T, \{X_t\}_{t \in V(T)})$, where T is a tree in which every node t is assigned a vertex subset $X_t \subseteq V(G)$, called a bag, such that the following three conditions hold:

- (T1) Every vertex of G is in at least one bag, that is, $V(G) = \bigcup_{t \in V(T)} X_t$.
- (T2) For every $uv \in E(G)$, there exists a node $t \in V(T)$ such that X_t contains both u and v .
- (T3) For every $u \in V(G)$, the set T_u of all nodes of T whose corresponding bags contain u , induces a connected subtree of T .

The *width* of the tree decomposition \mathcal{T} is the integer $\max_{t \in V(T)} |X_t| - 1$, that is, the size of its largest bag minus 1. The *treewidth* of a graph G , denoted by $\text{tw}(G)$, is the smallest possible width that a tree decomposition of G can have.

We will need the following algorithmic results about treewidth.

► **Proposition 1** ([6]). *There is a $2^{O(k)} \cdot n$ time algorithm that, given a graph G and an integer k , either outputs a tree decomposition of width at most $5k + 4$, or correctly decides that $\text{tw}(G) > k$.*

► **Proposition 2** ([5, 15]). *There is an algorithm with running time $2^{O(\text{tw}(G))} \cdot n^{O(1)}$ that computes a longest path between two given vertices of a given graph.*

Let us note that the running time of Proposition 2 can be improved to $2^{O(\text{tw}(G))} \cdot n$ by making use of the matroid-based approach from [15].

Our main theorem is based on graph minors, and we introduce some notation here.

► **Definition 3.** A *topological minor model* of H in G is a pair of functions (f, p) with $f: V(H) \rightarrow V(G)$ and $p: E(H) \rightarrow 2^{E(G)}$ such that

1. f is injective, and
2. for every edge $uv \in E(H)$, the graph $G[p(uv)]$ is a path from $f(u)$ to $f(v)$ in G , and
3. for edges $e, g \in E(H)$ with $e \neq g$, the paths $G[p(e)]$ and $G[p(g)]$ intersect only in endpoints or not at all.

The graph T induced by the topological minor model (f, p) is the subgraph of G that consists of the union of all paths $G[p(uv)]$ over all $uv \in E(H)$. The vertices in $f(V(H))$ are the *branch vertices* of T , and $G[p(e)]$ realizes the edge e in T .

3 Win/Win algorithm for Longest Detour

Throughout this section, let G be an undirected graph with n vertices and m edges, let $s, t \in V(G)$ and $k \in \mathbb{N}$. We wish to decide in time $2^{O(k)} \cdot n^{O(1)}$ whether G contains an (s, t) -path of length at least $d_G(s, t) + k$. To avoid trivialities, we assume without loss of generality that G is connected and $s \neq t$ holds. Moreover, we can safely remove vertices v that are not part of any (s, t) -path.

► **Definition 4.** Let G be a graph and let $s, t \in V(G)$. The (s, t) -*relevant part* of G is the graph induced by all vertices contained in some (s, t) -path. We denote it by $G_{s,t}$.

The graph $G_{s,t}$ can be computed efficiently from the block-cut tree of G . Recall that the *block-cut tree* of a connected graph G is a tree where each vertex corresponds to a *block*, that is, a maximal biconnected component $B \subseteq V(G)$, or to a *cut vertex*, that is, a vertex whose removal disconnects the graph. A block B and a cut vertex v are adjacent in the block-cut tree if and only if there is a block B' such that $B \cap B' = \{v\}$.

► **Lemma 5** (\star). *Let B_s and B_t denote the blocks of G that contain s and t , respectively. Furthermore, let P be the unique (B_s, B_t) -path in the block-cut tree of G . Then $G_{s,t}$ is the graph induced by the union of all blocks visited by P .*

We formulate an immediate implication of Lemma 5 that will be useful later.

► **Corollary 6.** *The block-cut tree of $G_{s,t}$ is a (B_s, B_t) -path.*

Hopcroft and Tarjan [18] proved that the block-cut tree of a graph can be computed in linear time using DFS. Hence we obtain an algorithm for computing $G_{s,t}$ from G .

► **Corollary 7.** *There is a linear-time algorithm that computes $G_{s,t}$ from G .*

3.1 The algorithm

By definition, the graph $G_{s,t}$ contains the same set of (s,t) -paths as G . Our algorithm for LONGEST DETOUR establishes a “win/win” situation as follows: We prove that, if the treewidth of $G_{s,t}$ is “sufficiently large”, then (G, s, t, k) is a YES-instance of LONGEST DETOUR. Otherwise the treewidth is small, and we use a known treewidth-based dynamic programming algorithm for computing the longest (s,t) -path. Hence the algorithm builds upon the following subroutines:

1. The algorithm from Corollary 7, computing the relevant part $G_{s,t}$ of G in time $O(n+m)$.
2. COMPUTE TREewidth(G, w) from Proposition 1, which is given G and $w \in \mathbb{N}$ as input, and either constructs a tree-decomposition T of G whose width is bounded by $5w+4$, or outputs LARGE. If the algorithm outputs LARGE, then $\text{tw}(G) > w$ holds. The running time is $2^{O(w)} \cdot n$.
3. LONGEST PATH(G, T, s, t) from Proposition 2, which is given G, s, t and additionally a tree-decomposition T of G , and outputs a longest (s,t) -path in G . The running time is $2^{O(w)} \cdot n^{O(1)}$, where w denotes the width of T .

We now formalize what we mean by “sufficiently large” treewidth.

► **Definition 8.** A function $f : \mathbb{N} \rightarrow \mathbb{N}$ is *detour-enforcing* if, for all $k \in \mathbb{N}$ and all graphs G with vertices s and t , the following implication holds: If $\text{tw}(G_{s,t}) > f(k)$, then G contains an (s,t) -path of length at least $d_G(s,t) + k$.

► **Theorem 9.** *The function $f : k \mapsto 32k + 2$ is detour-enforcing.*

We defer the proof of this theorem to the next section, and instead state Algorithm D, which uses f to solve LONGEST DETOUR. Algorithm D turns out to be an FPT-algorithm already when any detour-enforcing function f is known (as long as it is polynomial-time computable), and it becomes faster when detour-enforcing f of slower growth are used.

Algorithm D (LONGEST DETOUR) *Given (G, s, t, k) , this algorithm decides whether the graph G contains an (s,t) -path of length at least $d_G(s,t) + k$.*

D1 (Restrict to relevant part) Compute $G_{s,t}$ using Corollary 7.

D2 (Compute shortest path) Compute the distance d between s and t in $G_{s,t}$.

D3 (Compute tree-decomposition) Call COMPUTE TREewidth($G_{s,t}, f(k)$).

D3a (Small treewidth) If the subroutine found a tree-decomposition T of width at most $f(k)$, call LONGEST PATH($G_{s,t}, T, s, t$). Output YES if there is an (s,t) -path of length at least $d+k$, otherwise output NO.

D3b (Large treewidth) If the subroutine returned LARGE, output YES.

We prove the running time and correctness of Algorithm D.

► **Lemma 10** (\star). *For every polynomial-time computable detour-enforcing function $f : \mathbb{N} \rightarrow \mathbb{N}$, Algorithm D solves LONGEST DETOUR in time $2^{O(f(k))} \cdot n^{O(1)}$.*

Theorem 9 and Lemma 10 imply a $2^{O(k)} \cdot n^{O(1)}$ time algorithm for LONGEST DETOUR.

3.2 Overview of the proof of Theorem 9

In our proof of Theorem 9, large subdivisions of K_4 play an important role. Intuitively speaking, a sufficiently large subdivision of K_4 in $G_{s,t}$ allows us to route some (s,t) -path through it and then exhibit a long detour within that subdivision.

► **Definition 11.** For $k \in \mathbb{N}$, a graph F is a $K_4^{\geq k}$ if it can be obtained by subdividing each edge of K_4 at least k times. Please note that the numbers of subdivisions do not need to agree for different edges.

We show in Section 3.3 that graphs G containing $K_4^{\geq k}$ subgraphs in $G_{s,t}$ have k -detours.

► **Lemma 12.** *Let G be a graph and $k \in \mathbb{N}$. If $G_{s,t}$ contains a $K_4^{\geq k}$ subgraph, then G contains an (s,t) -path of length at least $d_G(s,t) + k$.*

Since the graph obtained by subdividing each edge of K_4 exactly k times is a planar graph on $O(k)$ vertices, the Excluded Grid Theorem yields a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that every graph of treewidth at least $f(k)$ contains some $K_4^{\geq k}$ minor. Furthermore, since every $K_4^{\geq k}$ has maximum degree 3, this actually shows that G contains some $K_4^{\geq k}$ as a subgraph. Thus, Lemma 12 implies that f is detour-enforcing, and a proof of this lemma immediately implies a weak version of Theorem 9.

By recent improvements on the Excluded Grid Theorem [7, 10], the function f above is at most a polynomial. However, even equipped with this deep result we cannot obtain a single-exponential algorithm for LONGEST DETOUR using the approach of Lemma 10: It would require f to be linear. In fact, excluding grids is too strong a requirement for us, since every function f obtained as a corollary of the full Excluded Grid Theorem must be super-linear [32]. We circumvent the use of the Excluded Grid Theorem and prove the following lemma from more basic principles.

► **Lemma 13.** *For graphs G and $k \in \mathbb{N}$, if $\text{tw}(G) \geq 32k+2$, then G contains a $K_4^{\geq k}$ subgraph.*

Together, Lemmas 13 and 12 imply Theorem 9.

Proof of Theorem 9. Let G and $s, t \in V(G)$ and $k \in \mathbb{N}$ be such that $\text{tw}(G_{s,t}) > f(k)$. By Lemma 13, the graph $G_{s,t}$ contains a $K_4^{\geq k}$ subgraph, so Lemma 12 implies that G contains an (s,t) -path of length $d_G(s,t) + k$. This shows that f is indeed detour-enforcing. ◀

3.3 Proof of Lemma 12: Rerouting in subdivided tetrahedra

Let (G, s, t, k) be an instance for LONGEST DETOUR such that $G_{s,t}$ contains a $K_4^{\geq k}$ subgraph M . We want to prove that $G_{s,t}$ has a path of length at least $d_G(s,t) + k$; in fact, we construct the desired detour entirely in the subgraph M , for which reason we first need to route some (s,t) -path through M .

► **Lemma 14** (\star). *There are two distinct vertices $u, v \in V(M)$ and two vertex-disjoint paths P_s and P_t in G such that P_s is an (s,u) -path, P_t is a (v,t) -path, and they only intersect with $V(M)$ at u and v .*




Figure 1 *Left*: One of the three possible cases for the relative positions of vertices u and v (red squares) in a subdivided tetrahedron $K_4^{\geq k}$ with degree-3 vertices b_1, \dots, b_4 (gray dots) and at least $k = 5$ subdivision vertices (small gray dots). Here, u and v lie in the same subdivided edge. See the full version of the paper for the remaining cases. *Right*: An exhaustive list of all (u, v) -paths (thick red); Lemma 15 implies that the longest among them is at least k longer than the shortest one.

The proof of this lemma uses the fact that every block in the block-cut tree is biconnected. Next we show that every $K_4^{\geq k}$ -graph M contains long detours.

► **Lemma 15.** *Let M be a $K_4^{\geq k}$ -graph. For every two distinct vertices $u, v \in V(M)$, there is a (u, v) -path of length at least $d_M(u, v) + k$ in M .*

The proof idea is to distinguish cases depending on where u, v lie in M relative to each other. For each case, we can exhaustively list all (u, v) -paths (see Figure 1). We do not quite know the lengths of these paths, but we do know that each has length at least $d_M(u, v)$; moreover, each (b_i, b_j) -path in M for two distinct degree-3 vertices b_i and b_j has length at least k , since we subdivided K_4 at least k times. The claim of Lemma 15 is that one of the (u, v) -paths must have length at least $d_M(u, v)$. To prove this, we set up a linear program where the variables are $d_M(u, v)$, k , and the various path segment lengths; its infeasibility informs us that indeed a path that is longer by k must exist.

Proof Sketch for Lemma 15. Let M be a $K_4^{\geq k}$ -graph, let $u, v \in V(M)$, and let b_1, \dots, b_4 denote the four degree-3 vertices of M . Let P_u be a path in M that realizes an edge of K_4 and satisfies $u \in V(P_u)$, and let P_v be such a path with $v \in V(P_v)$. We distinguish three cases, one of which is depicted in Figure 1:

1. The two paths are the same, that is, $P_u = P_v$.
2. The two paths share a degree-3 vertex, that is, $|V(P_u) \cap V(P_v)| = 1$.
3. The two paths are disjoint, that is, $V(P_u) \cap V(P_v) = \emptyset$.

By the symmetries of K_4 , this case distinction is exhaustive. Since K_4 has automorphisms that map any edge to any other edge, we can further assume that P_u is the path implementing the edge b_1b_2 such that P_u visits the vertices b_1, u, v , and b_2 in this order, see Figure 1.

We exhaustively list the set \mathcal{P} of (u, v) -paths of M in Figure 1. Each path is uniquely specified by the sequence of the degree-3 vertices it visits. For example, consider the path $ub_1b_4b_2v$: This path consists of the four edge-disjoint segments ub_1, b_1b_4, b_4b_2 , and b_2v ; in the example figure, these segments have length 3, 6, 6, and 4, respectively. Given a path $P \in \mathcal{P}$, let $S(P)$ be the set of its segments between u, v , and the degree-3 vertices. For a path or a path segment s , we denote its length by $\ell(s)$.

Since M is a $K_4^{\geq k}$, every edge of K_4 is realized by a path of length at least k in M . Hence, $\ell(b_ib_j) \geq k$ holds for all i, j with $i \neq j$. Moreover, we have $\ell(b_1b_2) = \ell(b_1u) + \ell(uv) + \ell(vb_2)$ in case 1. Let $d = d_M(u, v)$; clearly $\ell(P) \geq d$ holds for all $P \in \mathcal{P}$. Our goal is to show that M has a (u, v) -path P with $\ell(P) \geq d + k$. To this end, we treat d, k , and all path

segment lengths $\ell(b_i b_j)$ for $i \neq j$ and $\ell(b_1 u), \ell(uv), \ell(vb_2)$ as variables in a system of linear inequalities and establish that the claim holds if this system is unsatisfiable:

$$\ell(b_i b_j) \geq k, \quad \text{for all } i, j \text{ with } i \neq j, \quad (1)$$

$$\ell(b_1 u) + \ell(uv) + \ell(vb_2) = \ell(b_1 b_2), \quad (2)$$

$$\sum_{s \in S(P)} \ell(s) \geq d, \quad \text{for all } P \in \mathcal{P}, \quad (3)$$

$$\sum_{s \in S(P)} \ell(s) \leq d + k - 1, \quad \text{for all } P \in \mathcal{P}. \quad (4)$$

This system has eleven variables. Please note that d and k are also considered as variables in our formulation. The constraints in (1) express that M realizes each edge of K_4 by a path of length at least k . The constraints in (2) express that u and v lie on the path $b_1 b_2$ and break it up into segments. The constraints in (3) express that no (u, v) -path is shorter than d in length, and the constraints in (4) express that every (u, v) -path has length strictly less than $d + k$. In the full version of this extended abstract, we prove that the linear program is infeasible, and so every setting for the variables that satisfies (1)–(3) must violate an inequality from (4); this means that M must contain a (u, v) -path of length at least $d + k$ in case 1. The proof is analogous when u and v are on different subdivided edges of the subdivided tetrahedron. We conclude that, no matter how u and v lie relative to each other in M , there is always a (u, v) -path that is at least k longer than a shortest one. ◀

This allows us to conclude Lemma 12 rather easily.

Proof of Lemma 12. Let $d = d_G(s, t)$ be the length of a shortest (s, t) -path in G . Let M be a $K_4^{\geq k}$ in $G_{s,t}$, and let P_s, P_t, u , and v be the objects guaranteed by Lemma 14. Let P_{uv} be a shortest (u, v) -path that only uses edges of M ; its length is $d_M(u, v)$. Since the combined path P_s, P_{uv}, P_t is an (s, t) -path, its length is at least d .

Finally, Lemma 15 guarantees that there is a (u, v) -path Q_{uv} in M whose length is at least $d_M(u, v) + k$. Therefore, the length of the (s, t) -path P_s, Q_{uv}, P_t satisfies

$$\begin{aligned} \ell(P_s) + \ell(Q_{uv}) + \ell(P_t) &\geq \ell(P_s) + (d_M(u, v) + k) + \ell(P_t) \\ &= \ell(P_s) + \ell(P_{uv}) + \ell(P_t) + k \geq d + k. \end{aligned}$$

We constructed a path of at least length $d + k$ as required. ◀

3.4 Proof of Lemma 13: Large treewidth entails subdivided tetrahedra

To prove Lemma 13, we require some preliminaries from graph minors theory, among them a term for vertex sets that enjoy very favorable connectivity properties.

► **Definition 16** ([13]). Let G be a graph and $A, B \subseteq V(G)$. The pair (A, B) is a *separation* in G if the sets $A \setminus B$ and $B \setminus A$ are non-empty and no edge runs between them. The *order* of (A, B) is the cardinality of $A \cap B$.

For $S \subseteq V(G)$, we say that S is *linked* in G if, for every $X, Y \subseteq S$ with $|X| = |Y|$, there are $|X|$ vertex-disjoint paths between X and Y that intersect S exactly at its endpoints.

With these definitions at hand, we can adapt a result by Leaf and Seymour [25] to prove the following lemma on topological minor containment in graphs of sufficiently large treewidth. For any forest F on k vertices, with maximum degree 3, it asserts that graphs G of treewidth $\Omega(k)$ admit a separation such that one side contains F as a topological minor,

with the branch vertices of this topological minor being contained in $A \cap B$ and linked in G . We will use this lemma to complete the topological F -minor in $G[A]$ to a larger graph by using disjoint paths between vertices in $A \cap B$.

► **Lemma 17** (\star). *Let F be a forest on $k > 0$ vertices with maximum degree 3 and let G be a graph. If $\text{tw}(G) \geq \frac{3}{2}k - 1$, then G has a separation (A, B) of order $|V(F)|$ such that:*

1. *There is a topological minor model (f, p) of F in $G[A]$.*
2. *For every vertex $v \in V(F)$ of degree ≤ 2 , we have $f(v) \in A \cap B$.*
3. *$A \cap B$ is linked in $G[B]$.*

Building upon Lemma 17, we then prove Lemma 13 by adapting work of Raymond and Thilikos [30], who used a variant of Lemma 17 to prove the existence of k -wheel minors in graphs of treewidth $\Omega(k)$. To this end, let T and P be obtained by k -subdividing the full binary tree with 8 leaves, and the path with 8 vertices, respectively. We invoke Lemma 17 with F instantiated to the disjoint union $T \cup P$. Since F has $21k + 2$ vertices, we obtain from Lemma 17 that any graph G with $\text{tw}(G) \geq 32k + 2 \geq \frac{3}{2} \cdot (21k + 2)$ has a separation (A, B) of order $|V(F)|$ that contains F in $G[A]$ and has $A \cap B$ linked in $G[B]$.

Let X_F denote the eight leaves of T , and let Y_F denote the eight non-subdivision vertices of P . Furthermore, let $X_G, Y_G \subseteq A \cap B$ denote the images of X_F and Y_F in $G[A]$ under a topological minor model guaranteed by Lemma 17. Since $A \cap B$ is linked, we can find eight disjoint paths connecting X_G and Y_G in $G[B]$. We then prove that, regardless of how these paths connect X_G and Y_G , they always complete the topological minor model of F to one of $K_4^{\geq k}$ in G . (The full version illustrates this in a figure.) Lemma 13 then follows.

Proof of Lemma 13. Let $k \in \mathbb{N}$ and let G be a graph with $\text{tw}(G) \geq 32k + 2$. As before, let T denote the full binary tree with 8 leaves, with root r , after each edge was subdivided k times. Let P denote the path on 8 vertices after subdividing each edge k times.

We write $X_F = \{x_1, \dots, x_8\}$ for the leaves of T , and we write $Y_F = \{y_1, \dots, y_8\}$ for the vertices in P that were not obtained as subdivision vertices. Finally, we write F for the disjoint union $T \cup P$ and consider $X_F, Y_F \subseteq V(F)$. Note that $|V(F)| = 21k + 2$ and that the degree of all vertices in $X_F \cup Y_F$ is bounded by 2.

By Lemma 17, there is a separation (A, B) in G of order $|V(F)|$ such that $A \cap B$ is linked, and there is a topological minor model (f, p) of F in $G[A]$ with $f(X_F \cup Y_F) \subseteq A \cap B$. We write $X_G = \{f(v) \mid v \in X_F\}$ and $Y_G = \{f(v) \mid v \in Y_F\}$. In the following, we aim at completing the subgraph induced by (f, p) in G to a $K_4^{\geq k}$ subgraph.

Since $A \cap B$ is linked in $G[B]$, there are vertex-disjoint paths L_1, \dots, L_8 between X_G and Y_G in $G[B]$ that avoid $A \cap B$ except at their endpoints. For $i \in [8]$, denote the endpoints of L_i in X_G and Y_G by s_i and t_i , respectively. Assume without limitation of generality (by reordering paths) that $t_i = f(y_i)$ holds for all $i \in [8]$. Furthermore, for $x \in X_G$, write $\sigma(x)$ for the vertex of Y_G that x is connected to via its path among L_1, \dots, L_8 .

Let S denote the image of T under (f, p) , which is a tree; let $\text{root}(S) = f(r)$. Write S_1, S_2 for the two subtrees of S rooted at the children of $\text{root}(S)$. Let $\text{lca}(s_1, s_8)$ denote the lowest common ancestor of s_1 and s_8 in S . We distinguish two cases:

Case 1: We have $\text{lca}(s_1, s_8) \neq \text{root}(S)$. That is, s_1 and s_8 are both in S_1 or both in S_2 .

Assume without limitation of generality that $s_1, s_8 \in V(S_1)$, as the argument proceeds symmetrically otherwise. Let x and x' be two distinct leaves of S_2 . Then we find a $K_4^{\geq k}$ in G by defining branch vertices $w = \text{lca}(s_1, s_8)$, $p = \text{lca}(x, x')$, $a = \sigma(x)$, and $b = \sigma(x')$. Note that $p \notin \{x, x'\}$ and that the four vertices are distinct.

We realize the edge pw along the (p, w) -path present in S , and ab along the (a, b) -path present in P . We realize pa by concatenating the (p, x) -path in S and the (x, a) -path in

$G[B]$, and we realize pb likewise. To realize wa , we proceed as follows: If a precedes b in the order on P , then concatenate the (w, s_1) -path in S with L_1 and the (y_1, a) -path in P . If b precedes a , then concatenate the (w, s_8) -path in S with L_8 and the (y_8, a) -path in P . Realize wb symmetrically. Then every edge between pairs in $\{w, p, a, b\}$ is realized, and it is so by a path of length at least k . This gives a topological minor model of $K_4^{\geq k}$ in G .

Case 2: We have $\text{lca}(s_1, s_8) = \text{root}(S)$. That is, s_1 and s_8 are in different subtrees S_1 and S_2 . Let R be a subtree of height 2 in S that is disjoint from the (s_1, s_8) -path in S . It is easy to verify that such a subtree indeed exists; denote its root by p , its leaves by x, x' , and its parent in S by w . Furthermore, define $a = \sigma(x)$ and $b = \sigma(x')$. We declare $\{w, p, a, b\}$ as branch vertices and connect them as in the previous case.

In both cases, the constructed topological minor model shows that G contains a $K_4^{\geq k}$ subgraph. This proves the lemma. \blacktriangleleft

4 Dynamic programming algorithm for exact detour

We devise an algorithm for EXACT DETOUR using a reduction to EXACT PATH, the problem that is given (G, s, t, k) to determine whether there is an (s, t) -path of length exactly k .

► **Theorem 18** (\star). EXACT DETOUR is fixed-parameter tractable. In particular, it has a bounded-error randomized algorithm with running time $2.746^k \text{poly}(n)$, and a deterministic algorithm with running time $6.745^k \text{poly}(n)$.

Before we state the algorithm, let us introduce some notation. Let $s, t \in V(G)$. For any $x \in V(G)$, we abbreviate $d_G(s, x)$, that is, the distance from s to x in G , with $d(x)$, and we let the i -th layer of G be the set of vertices x with $d(x) = i$. For $u, v \in V(G)$ with $d(u) < d(v)$, we write $G_{[u, v]}$ for the graph $G[X]$ induced by the vertex set X that contains u, v , and all vertices x with $d(u) < d(x) < d(v)$. We also write $G_{[u, \infty]}$ for the graph $G[X]$ induced by the vertex set X that contains u and all vertices x with $d(u) < d(x)$. These graphs can be computed in linear time using breadth-first search starting at s . We now describe an algorithm for EXACT DETOUR that makes queries to an oracle for EXACT PATH.

The general idea is as follows. Let G be an undirected graph, and consider an (s, t) -path P of length $d + k$ where $d = d_G(s, t)$, and let x be a token that travels along this path from s to t . As the token advances one step in the path, the number $d(x)$ can be incremented, decremented, or stay the same. When x moves from s to t , we must increment $d(x)$ at least d times, can decrement it at most $k/2$ times, and keep it unchanged at most k times; the reason is that the path must reach t but must use exactly k edges more than a shortest path. The crucial observation is that there are at most k different layers whose intersection with the path P contains more than one vertex. The idea for the algorithm is to guess the layers with more than one vertex and run an algorithm for EXACT PATH on them.

Algorithm A (EXACT DETOUR) Given (G, s, t, k) , this algorithm decides whether the graph G contains an (s, t) -path of length exactly $d_G(s, t) + k$.

A1 (Initialize table) For each $x \in V(G)$ with $d(x) \leq d(t)$, set $T[x] = \emptyset$.

When the algorithm halts, every entry $T[x]$ of the table is meant to satisfy the following property Q_x : For each integer ℓ with $d(t) - d(x) \leq \ell \leq d(t) - d(x) + k$, the set $T[x]$ contains ℓ if and only if $G_{[x, \infty]}$ contains an (x, t) -path of length ℓ .

A2 (Compute entries for the last $k+1$ layers) For each $x \in V(G)$ with $d(t) - k \leq d(x) \leq d(t)$, let $T[x]$ be the set of all integers ℓ with $\ell \in \{0, \dots, 2k\}$ such that there is an (x, t) -path of length ℓ in $G_{[x, \infty]}$ (that is, call EXACT PATH $(G_{[x, \infty]}, x, t, \ell)$).

When this step finishes, all vertices x in the last $k+1$ layers satisfy property Q_x .

A3 (Inductively fill in earlier layers) For each d from $d(t) - k - 1$ down to 0, for each x with $d(x) = d$, and for each y with $d(x) < d(y) \leq d(x) + k + 1$, we do the following:

A3a Compute the set L of all $\ell' \in \{0, \dots, 2k + 1\}$ such that there is an (x, y) -path of length ℓ' in $G_{[x, y]}$ (that is, call EXACT PATH $(G_{[x, y]}, x, y, \ell')$).

A3b Set $T[x] := T[x] \cup (L + T[y])$.

We will show that, when all vertices of a layer d have been considered, all vertices x in the layers d and higher satisfy property Q_x .

A4 Accept if and only if $(d_G(s, t) + k) \in T[s]$ holds.

► **Lemma 19.** *Algorithm A is a polynomial-time Turing reduction from EXACT DETOUR to EXACT PATH; on instances with parameter k , all queries have parameter at most $2k + 1$.*

Proof. The running time of A is polynomially bounded since breadth-first search can be used to discover all partial graphs $G_{[x, y]}$ and $G_{[x, \infty)}$, and we loop at most over every pair of vertices in A2 and A3. For the parameter bound, note that the queries in A2 and A3 are for paths of length at most $2k$ and $2k + 1$, respectively. It remains to prove the correctness.

We execute algorithm A on an instance (G, s, t, k) . For the correctness, it suffices to prove that property Q_s holds at the end of the execution: Note that ℓ with $\ell = d_G(s, t) + k$ lies in the interval $[d(t) - d(s), d(t) - d(s) + k]$ since $d(s) = 0$ and $d(t) = d_G(s, t)$ holds. Moreover, we have $G_{[s, \infty)} = G$. Thus Q_s guarantees that $\ell \in T[s]$ holds if and only if G contains an (s, t) -path of length ℓ , which by step A4 implies that A accepts if and only if (G, s, t, k) is a yes-instance of EXACT DETOUR. Therefore it remains to prove that Q_s holds at the end of the execution of A. We do so using the following claim.

Claim: For all x with $0 \leq d(x) \leq d(t)$, property Q_x holds forever after the entry $T[x]$ is written to for the last time.

We prove this claim by induction on $d(x)$. For the base case, let x be a vertex with $d(x) \geq d(t) - k$. The entry $T[x]$ is only written to in step A2. To prove that Q_x holds after A2, let ℓ be an integer with $d(t) - d(x) \leq \ell \leq d(t) - d(x) + k$. Note that $d(t) - d(x) \geq 0$ and $d(t) - d(x) + k \leq d(t) - (d(t) - k) + k \leq 2k$ holds, and so step A2 adds ℓ to $T[x]$ if and only if the graph $G_{[x, \infty)}$ contains an (x, t) -path of length ℓ . Therefore, Q_x holds forever after A2 has been executed.

For the induction step, let x be a vertex with $d(x) < d(t) - k$. By the induction hypothesis, Q_y holds for all y with $d(y) > d(x)$. The entry $T[x]$ is only written to in step A3b, and when it is first written to, the outer d -loop in A3 has fully processed all layers larger than $d(x)$. Thus already when $T[x]$ is written to for the first time, Q_y holds for all y with $d(y) > d(x)$. Let T be the table right after A3b writes to $T[x]$ for the last time. It remains to prove that $T[x]$ satisfies Q_x . Let ℓ be an integer with $d(t) - d(x) \leq \ell \leq d(t) - d(x) + k$.

Claim: There is an (x, t) -path of length ℓ if and only if $T[x]$ contains ℓ .

For the forward direction, let P be an (x, t) -path in $G_{[x, \infty)}$ of length exactly ℓ . There are exactly ℓ vertices $u \in V(P) \setminus \{x\}$. Moreover, since every edge $uv \in E(P)$ satisfies $|d(u) - d(v)| \leq 1$, every $d \in \{d(x) + 1, \dots, d(t)\}$ must have some vertex $u \in V(P)$ with $d(u) = d$. Since $\ell \leq d(t) - d(x) + k$, there are at most k distinct d where more than one vertex $u \in V(P)$ satisfies $d(u) = d$. By the pigeon hole principle, there exists an integer d in the $(k + 1)$ -element set $\{d(x) + 1, \dots, d(x) + k + 1\}$ such that there is exactly one vertex $y \in V(P)$ with $d(y) = d$.

Let $P_{[x, y]}$ be the subpath of P between x and y , and let ℓ' be its length. By construction, $P_{[x, y]}$ is an (x, y) -path in $G_{[x, y]}$. Moreover, we have $\ell' \leq \ell - (d(t) - d(y))$ since $V(P) \setminus \{x\}$ contains ℓ vertices u , at least $d(t) - d(y)$ of which satisfy $d(u) > d(y)$. By choice of ℓ and y , we obtain $\ell' \leq d(y) - d(x) + k \leq 2k + 1$. For this setting of y and ℓ' , step A3a detects the

path $P_{[x,y]}$ and ℓ' is added to the set L . The second piece $P_{[y,\infty]}$ of the path P is a (y, t) -path in $G_{[y,\infty]}$ of some length ℓ'' between $d(t) - d(y)$ and $d(t) - d(y) + k$; since Q_y holds when A3b is executed for x and y , the set $T[y]$ contains ℓ'' , and so $\ell = \ell' + \ell''$ gets added to $T[x]$. Since elements never get removed from $T[x]$, the forward direction of the claim holds.

For the backward direction of the claim, assume that $T[x]$ contains ℓ . This means that ℓ is added in step A3b during the execution of the algorithm; in particular, consider the variables $y \in V(G)$, $\ell' \in L$, and $\ell'' \in T[y]$ when $\ell = \ell' + \ell''$ is added to $T[x]$. By the induction hypothesis, $\ell'' \in T[y]$ implies that there is a (y, t) -path in $G_{[y,\infty]}$ of length ℓ'' . Moreover, ℓ' was set in A3a in such a way that there is an (x, y) -path of length ℓ' in the graph $G_{[x,y]}$. Combined, these two paths yield a single (x, t) -path in $G_{[x,\infty]}$ of length ℓ . The backward direction of the claim follows. ◀

The randomized algorithm of Björklund et al. [3] is for a variant of EXACT PATH where the terminal vertices s and t are not given, that is, any path of length exactly k yields a YES-instance. Their algorithm applies to our problem as well, with the same running time. We sketch an argument for this observation here. Recall that the idea is to reduce the problem to checking whether a certain polynomial is identically zero – this polynomial is defined by summing over all possible labelled walks of length k (see [12, Sec. 10.4.3]). We modify the polynomial by adding two leaf-edges, one incident to s and one to t , and restricting our attention only to $(k + 2)$ -walks that contain these two edges. The required information for such walks can still be computed efficiently as before. The crux of the proof is that walks that are not paths cancel out over a field of characteristic two; this argument works by a local re-orientation of segments of the walk – an operation that does not change the vertices of the walk and must therefore keep s and t fixed. The graph G contains a k -path if and only if the polynomial is not identically zero; this property remains true in our case. The rest of the argument goes through as before, so the algorithm of Björklund et al. applies to EXACT PATH with no significant loss in the running time.

The deterministic algorithm of Zehavi [34] also does not expect the terminal vertices to be given, but this algorithm works for the weighted version of the problem. In the weighted k -path problem, we are given a graph G , weights w_e on each edge, a number k , and a number W , and the question is whether there is a path of length exactly k such that the sum of all edge weights along the path is at most W . We observe the following simple reduction from EXACT PATH (with terminal vertices s and t) to the weighted k -path problem (without terminal vertices): Every edge gets assigned the same edge weight 2, except for the new leaf-edges at s and t , which get edge weight 1. Now every path with exactly $k + 2$ edges has weight at most $W = 2k + 2$ if and only if the first and the last edges of the path are the leaf-edges we added. Due to this reduction, Zehavi's algorithm applies to EXACT PATH with no significant loss in the running time.

Theorem 18 follows from Algorithm A by using either the algorithm of Björklund et al. [3] or Zehavi [34] as the oracle. We remark that Theorem 18 and Algorithm A apply to directed graphs as well, in which case an algorithm for EXACT PATH in directed graphs needs to be used (color coding yields the fastest randomized algorithm [2], while Zehavi's deterministic algorithm also applies to directed graphs).

Open problem. We conclude with an open problem: what is the complexity of LONGEST DETOUR in directed graphs? Neither directed treewidth nor cylindrical grid minors [21] seem to help. Can one even find an (s, t) -path of length $\geq d_G(s, t) + 1$ in polynomial time?

Acknowledgments. We thank Daniel Lokshtanov, Meirav Zehavi, Petr Golovach, Saket Saurabh, Stephan Kreutzer, and Tobias Mömke for helpful discussions and answers.

References

- 1 Noga Alon, Gregory Gutin, Eun Jung Kim, Stefan Szeider, and Anders Yeo. Solving MAX- r -SAT above a tight lower bound. *Algorithmica*, 61(3):638–655, 2011. doi:10.1007/s00453-010-9428-7.
- 2 Noga Alon, Raphael Yuster, and Uri Zwick. Color-coding. *Journal of the ACM*, 42(4):844–856, 1995. doi:10.1145/210332.210337.
- 3 Andreas Björklund, Thore Husfeldt, Petteri Kaski, and Mikko Koivisto. Narrow sieves for parameterized paths and packings. *Journal of Computer and System Sciences*, 87:119–139, 2017. doi:10.1016/j.jcss.2017.03.003.
- 4 Hans L. Bodlaender. On linear time minor tests with depth-first search. *Journal of Algorithms*, 14(1):1–23, 1993. doi:10.1006/jagm.1993.1001.
- 5 Hans L. Bodlaender, Marek Cygan, Stefan Kratsch, and Jesper Nederlof. Deterministic single exponential time algorithms for connectivity problems parameterized by treewidth. *Information and Computation*, 243:86–111, 2015. doi:10.1016/j.ic.2014.12.008.
- 6 Hans L. Bodlaender, Pål Grønås Drange, Markus S. Dregi, Fedor V. Fomin, Daniel Lokshtanov, and Michał Pilipczuk. A $c^k n$ 5-approximation algorithm for treewidth. *SIAM Journal on Computing*, 45(2):317–378, 2016. doi:10.1137/130947374.
- 7 Chandra Chekuri and Julia Chuzhoy. Polynomial bounds for the grid-minor theorem. *Journal of the ACM*, 63(5):40:1–40:65, 2016. doi:10.1145/2820609.
- 8 Jianer Chen, Joachim Kneis, Songjian Lu, Daniel Mölle, Stefan Richter, Peter Rossmanith, Sing-Hoi Sze, and Fenghui Zhang. Randomized divide-and-conquer: Improved path, matching, and packing algorithms. *SIAM Journal on Computing*, 38(6):2526–2547, 2009. doi:10.1137/080716475.
- 9 Jianer Chen, Songjian Lu, Sing-Hoi Sze, and Fenghui Zhang. Improved algorithms for path, matching, and packing problems. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 298–307. SIAM, 2007.
- 10 Julia Chuzhoy. Improved bounds for the excluded grid theorem. *CoRR*, abs/1602.02629, 2016. URL: <http://arxiv.org/abs/1602.02629>.
- 11 Robert Crowston, Mark Jones, Gabriele Muciaccia, Geevarghese Philip, Ashutosh Rai, and Saket Saurabh. Polynomial kernels for lambda-extendible properties parameterized above the Poljak-Turzik bound. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, pages 43–54, 2013. doi:10.4230/LIPIcs.FSTTCS.2013.43.
- 12 Marek Cygan, Fedor V. Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015. doi:10.1007/978-3-319-21275-3.
- 13 Reinhard Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, third edition, 2005.
- 14 Fedor V. Fomin and Petteri Kaski. Exact exponential algorithms. *Communications of the ACM*, 56(3):80–88, 2013. doi:10.1145/2428556.2428575.
- 15 Fedor V. Fomin, Daniel Lokshtanov, and Saket Saurabh. Efficient computation of representative sets with applications in parameterized and exact algorithms. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 142–151, 2014. doi:10.1137/1.9781611973402.10.
- 16 Gregory Gutin, Eun Jung Kim, Michael Lampis, and Valia Mitsou. Vertex cover problem parameterized above and below tight bounds. *Theory of Computing Systems*, 48(2):402–410, 2011. doi:10.1007/s00224-010-9262-y.

- 17 Gregory Gutin, Leo van Iersel, Matthias Mnich, and Anders Yeo. Every ternary permutation constraint satisfaction problem parameterized above average has a kernel with a quadratic number of variables. *Journal of Computer and System Sciences*, 78(1):151–163, 2012. doi:10.1016/j.jcss.2011.01.004.
- 18 John Hopcroft and Robert Tarjan. Algorithm 447: Efficient algorithms for graph manipulation. *Communications of the ACM*, 16(6):372–378, June 1973. doi:10.1145/362248.362272.
- 19 Falk Hüffner, Sebastian Wernicke, and Thomas Zichner. Algorithm engineering for color-coding with applications to signaling pathway detection. *Algorithmica*, 52(2):114–132, 2008. doi:10.1007/s00453-007-9008-7.
- 20 Russell Impagliazzo and Ramamohan Paturi. On the complexity of k -SAT. *Journal of Computer and System Sciences*, 62(2):367–375, 2001. doi:10.1006/jcss.2000.1727.
- 21 Ken-ichi Kawarabayashi and Stephan Kreutzer. The directed grid theorem. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing (STOC)*, pages 655–664. ACM, 2015. doi:10.1145/2746539.2746586.
- 22 Joachim Kneis, Daniel Mölle, Stefan Richter, and Peter Rossmanith. Divide-and-color. In *Proceedings of the 32nd International Workshop on Graph-Theoretic Concepts in Computer Science (WG)*, pages 58–67, 2006. doi:10.1007/11917496_6.
- 23 Ioannis Koutis. Faster algebraic algorithms for path and packing problems. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming (ICALP)*, volume 5125, pages 575–586. Springer, 2008. doi:10.1007/978-3-540-70575-8_47.
- 24 Ioannis Koutis and Ryan Williams. Algebraic fingerprints for faster algorithms. *Communications of the ACM*, 59(1):98–105, 2016. doi:10.1145/2742544.
- 25 Alexander Leaf and Paul D. Seymour. Tree-width and planar minors. *Journal of Combinatorial Theory, Series B*, 111:38–53, 2015. doi:10.1016/j.jctb.2014.09.003.
- 26 Meena Mahajan and Venkatesh Raman. Parameterizing above guaranteed values: MaxSat and MaxCut. *Journal of Algorithms*, 31(2):335–354, 1999. doi:10.1006/jagm.1998.0996.
- 27 Meena Mahajan, Venkatesh Raman, and Somnath Sikdar. Parameterizing above or below guaranteed values. *Journal of Computer and System Sciences*, 75(2):137–153, 2009. doi:10.1016/j.jcss.2008.08.004.
- 28 Burkhard Monien. How to find long paths efficiently. In *Analysis and design of algorithms for combinatorial problems (Udine, 1982)*, volume 109 of *North-Holland Math. Stud.*, pages 239–254. North-Holland, Amsterdam, 1985. doi:10.1016/S0304-0208(08)73110-4.
- 29 Christos H. Papadimitriou and Mihalis Yannakakis. On limited nondeterminism and the complexity of the V-C dimension. *Journal of Computer and System Sciences*, 53(2):161–170, 1996. doi:10.1006/jcss.1996.0058.
- 30 Jean-Florent Raymond and Dimitrios M. Thilikos. Low polynomial exclusion of planar graph patterns. *Journal of Graph Theory*, 84(1):26–44, 2017. doi:10.1002/jgt.22009.
- 31 Neil Robertson and Paul D. Seymour. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, 1986. doi:10.1016/0095-8956(86)90030-4.
- 32 Neil Robertson, Paul D. Seymour, and Robin Thomas. Quickly excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 62(2):323–348, 1994. doi:10.1006/jctb.1994.1073.
- 33 Ryan Williams. Finding paths of length k in $O^*(2^k)$ time. *Information Processing Letters*, 109(6):315–318, 2009. doi:10.1016/j.ipl.2008.11.004.
- 34 Meirav Zehavi. Mixing color coding-related techniques. In *Proceedings of the 23rd Annual European Symposium on Algorithms (ESA)*, volume 9294, pages 1037–1049. Springer, 2015. doi:10.1007/978-3-662-48350-3_86.