

Being Even Slightly Shallow Makes Life Hard*

Irene Muzi¹, Michael P. O’Brien², Felix Reidl³, and Blair D. Sullivan⁴

- 1 University of Rome, “La Sapienza”, Rome, Italy
irene.muzi@gmail.com
- 2 North Carolina State University, Raleigh, USA
mpobrie3@ncsu.edu
- 3 North Carolina State University, Raleigh, USA
felix.reidl@gmail.com
- 4 North Carolina State University, Raleigh, USA
blair_sullivan@ncsu.edu

Abstract

We study the computational complexity of identifying dense substructures, namely $r/2$ -shallow topological minors and r -subdivisions. Of particular interest is the case $r = 1$, when these substructures correspond to very localized relaxations of subgraphs. Since DENSEST SUBGRAPH can be solved in polynomial time, we ask whether these slight relaxations also admit efficient algorithms.

In the following, we provide a negative answer: DENSE $r/2$ -SHALLOW TOPOLOGICAL MINOR and DENSE r -SUBDIVISION are already NP-hard for $r = 1$ in very sparse graphs. Further, they do not admit algorithms with running time $2^{o(\text{tw}^2)} n^{O(1)}$ when parameterized by the treewidth of the input graph for $r \geq 2$ unless ETH fails.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases Topological minors, NP Completeness, Treewidth, ETH, FPT algorithms

Digital Object Identifier 10.4230/LIPIcs.MFCS.2017.79


1 Introduction

Many structural graph theory results stem from the principle that excluding certain (dense) substructures affords algorithmic tractability [7, 8, 12, 13, 11, 16, 9, 20]. It is therefore natural to ask for each of these substructures whether computing its densest occurrence is efficiently possible. In this paper, we characterize the complexity of this problem for substructures formed by contracting short disjoint paths.

There has been significant recent work on a slightly more general substructure, r -shallow minors, which are formed by contracting disjoint connected *subgraphs* of radius at most r . Shallow minors form a gradient between the locality of subgraphs (0-shallow minors) and the global nature of (∞ -shallow) minor containment. Because finding r -shallow minors of density/degeneracy at least d is polynomial time solvable at $r = 0$ [15, 14] but NP-complete at $r = \infty$ [3], one might expect the problem to be fixed-parameter tractable with respect to r . However, Dvořák proved that for any fixed $r > 0$ both variations are NP-complete already in graphs of maximum degree four and $d \geq 4$ ($d \geq 2$ if degeneracy is the measure) [10]. Accordingly, a parameterization by d also cannot possibly yield an fpt-algorithm—a sharp

* This work supported in part by the DARPA GRAPHS Program and the Gordon & Betty Moore Foundation’s Data-Driven Discovery Initiative through Grants SPAWAR-N66001-14-1-4063 and GBMF4560 to Blair D. Sullivan.





■ **Figure 1** The graph H is a 1-shallow topological minor of G , as witnessed by the model marked with blue nails and golden paths.

contrast to unrestricted minors [3]. Dvořák further showed that the problem is in FPT if parameterized by the treewidth tw of the input graph and designed an $O(4^{\text{tw}^2} n)$ dynamic programming algorithm.

We consider whether more conservative relaxations of subgraph containment remain efficiently solvable, specifically $r/2$ -shallow *topological minors* and *r-subdivisions*. A graph H is an $r/2$ -shallow topological minor of a graph G if a $(\leq r)$ -subdivision of H is isomorphic to some subgraph of G . The case of $r = 1$ is of particular interest because it generalizes subgraphs to a proper subset of 1-shallow minors.

We show that DENSE $r/2$ -SHALLOW TOPOLOGICAL MINOR (DENSE $r/2$ -STM) and DENSE r -SUBDIVISION (DENSE r -SD) are NP-complete already in subcubic apex-graphs¹ for $r \geq 1$ via a reduction from POSITIVE 1-IN-3SAT. Accordingly, a parameterization by the target density d does not make these problems fixed-parameter tractable. The same reduction also implies that neither problem can be solved in time $O(2^{o(n)})$ unless the ETH fails. In other words, finding dense substructures which are just slightly ‘less local’ than subgraphs seems to be intrinsically difficult. Following Dvořák’s results, we then consider a parameterization by treewidth and ask whether an algorithm with running time better than $O(2^{\text{tw}^2} n)$ is possible. Surprisingly, we can rule out such an algorithm already for DENSE 1-SHALLOW TOPOLOGICAL MINOR: unless the ETH fails, no algorithm with running time $O(2^{o(\text{tw}^2)} n)$ can exist.

2 Preliminaries

For a graph G we use $|G|$ to denote the number of vertices and $\|G\|$ to denote the number of edges in G . A graph H appears as an *r-subdivision* in a graph G if the graph obtained from H by subdividing every edge r times is isomorphic to some subgraph of G . Similarly, H is a *$r/2$ -shallow topological minor*² of G if a graph obtained from H by subdividing every edge *up to* r times is isomorphic to a subgraph of G . In both cases, the subgraph witnessing the minor is the *model* and we call those vertices in it that correspond to (subdivided) edges *subdivision vertices* and all other vertices *nails*. If S_{uv} is the set of subdivision vertices on a subdivided uv edge, we say S_{uv} is *smoothed* into the uv edge.

The following two problems are the focus of this paper:

¹ That is, a graph in which the removal of a single vertex results in a subcubic planar graph.

² The somewhat cumbersome convention of letting an r -shallow topological minor contract paths of length $2r + 1$ is convenient in the broader context of sparse graph classes (cf. [18]).

DENSE $r/2$ -SHALLOW TOPOLOGICAL MINOR (DENSE $r/2$ -STM)

Input: A graph G and a rational number d .

Question: Is there an $r/2$ -shallow topological minor H of G with density $\|H\|/|H| \geq d$?

DENSE r -SUBDIVISION (DENSE r -SD)

Input: A graph G and a rational number d .

Question: Is there a graph H that is contained in G as an r -subdivision with density $\|H\|/|H| \geq d$?

The following variant, which we prove to be NP-complete in Section 4.1, might be of independent interest:

DENSE BIPARTITE SUBDIVISION

Input: A bipartite graph (X, Y, E) and a rational number d .

Question: Are there subsets $X' \subseteq X, Y' \subseteq Y$ such that all vertices in X' can be smoothed into unique edges in Y' and $|X'|/|Y'| \geq d$?

Our main tool will be linear reductions from the following SAT-variant:

POSITIVE 1-IN-3SAT

Input: A CNF boolean formula ψ with only positive literals.

Question: Does ϕ have a satisfying assignment such that each clause contains exactly one true variable?

Mulzer and Rote showed [17] that POSITIVE 1-IN-3SAT remains NP-hard when restricted to *planar* formulas. A formula ϕ is planar if the graph obtained from ϕ by creating one vertex for each clause and variable and connecting a variable-vertex to a clause-vertex if the clause contains said variable is planar.

Schaefer [21] provided a linear reduction from 3SAT to 1-IN-3SAT. We can further easily transform a formula ϕ with negative literals into one with only positive literals as follows: for each variable x , introduce the variables x^+, x^-, a_x, b_x, c_x . Replace every occurrence of x with x^+ and every occurrence of \bar{x} with x^- and add the clauses

$$\{x^+, x^-, a_x\}, \{x^+, x^-, b_x\}, \{a_x, b_x, c_x\},$$

to the formula. It is easy to verify that exactly one of x^+, x^- must be true in a 1-in-3 satisfying assignment and that the resulting formula ϕ' has size linear in $|\phi|$. In conclusion, there exists a linear reduction from 3SAT to POSITIVE 1-IN-3SAT which implies that under ETH, POSITIVE 1-IN-3SAT cannot be solved in time $2^{o(n)}(n+m)^{O(1)}$, where n is the number of variables and m is the number of clauses. Using sparsification one can further show that the ETH excludes algorithms for 3SAT with running time $2^{o(m)}(n+m)^{O(1)}$ (see *e.g.* the survey by Cygan *et al.* [5]). The above reduction implies the following lower bound:

► **Proposition 1.** Unless the ETH is false, POSITIVE 1-IN-3SAT cannot be solved in time $2^{o(m)}(n+m)^{O(1)}$.

3 Algorithmic considerations

We start with a basic observation about the problems in question with the smallest sensible depths of $r = 1$:

► **Lemma 1.** *The densest $\frac{1}{2}$ -shallow minor or 1-subdivision on a given set of nails can be computed in polynomial time.*

Proof. Assume we are to find the densest 1-subdivision with nail set X in a graph G . We construct an auxiliary bipartite graph \hat{G} with vertex set $V(G) \setminus X$ and $\binom{X}{2}$ where the vertex $v \in V(G) \setminus X$ is connected to $xy \in \binom{X}{2}$ iff $\{x, y\} \subseteq N(v)$ in G , that is, if v can be contracted into the edges x, y . Now simply note that a matching of cardinality ℓ in \hat{G} corresponds to a 1-subdivision in G with ℓ subdivisions. Finding a maximal matching in \hat{G} therefore provides us with the densest 1-subdivision in G with nail set X . The same proof works for $\frac{1}{2}$ -shallow minors if we subdivide all edges existing inside X and then construct \hat{G} . ◀

Consequently, DENSE 1-SD and DENSE $\frac{1}{2}$ -STM both admit a simple $2^n n^{O(1)}$ -algorithm: we guess the nail set X and apply the matching construction from Lemma 1. For the same reasons, both problems are in XP when parameterized by the number of nails. We cannot hope for much better since for $r = 0$ and $d \sim k^2$ we simply recover the problem of finding a k -clique. Besides being W[1]-hard and thus probably not in FPT, k -CLIQUE further does not admit algorithms with running time $f(k)n^{o(k)}$ unless the ETH fails [4].

The approach of guessing the nail sets also fails for larger depths: knowing the nails of a, say, 1-shallow minors leaves us with the problem of contracting paths of length two into X , which cannot be represented as a simple matching problem. The reduction presented in Section 4.2 proves as a corollary that DENSE 1-STM remains NP-hard when the nail set of the densest minor is known.

Finally, as we will see in Section 4.1, both problems are already NP-complete for very small densities d , making them paraNP-complete under this parameterization. Therefore none of the input variables will work well as a parameterization, and it is sensible to consider *structural* parameters, meaning parameters derived from the input graph. A good contender for such parameters are *width measures* like tree-, path-, or cliquewidth. Indeed, we can express the problem of finding a dense shallow minor or a dense subdivision in MSO₂ and apply variants of Courcelle's theorem to obtain the following:

► **Proposition 2.** DENSE $r/2$ -STM and DENSE r -SD are in FPT when parameterized by the treewidth of the input graph.

Proof. We can express a model for an r -shallow minor in MSO₂ as follows: it consists of a vertex-set W and an edge set F , where F induces a set of paths. We can further easily express that the paths formed by F are a) of length at most r , b) disjoint, and c) have endpoints in W . Lastly, we demand that for every pair $x, y \in W$ there exists at most one path in F that has x and y as endpoints.

From an optimization perspective, we can therefore express the *feasible* solutions to DENSE $r/2$ -STM (and DENSE r -SD with small modifications). In order to express our optimization goal, let us introduce one more set of vertices C with the property that every path induced by F contains at most one vertex from C —for example, we can express in MSO₂ that vertices of C are not pairwise reachable via the graph induced by F . With this auxiliary set, the density of the resulting minor is at least $|C|/|W|$ and exactly the density if C is maximal with respect to our choice of F . Accordingly, we find that there exists an r -shallow topological minor of density at least d if $|C| - d|W| \geq 0$. This constraint and the aforementioned MSO₂-description of a minor fall within the expressive power of the EMSO-framework introduced by Arnborg, Lagergren, and Seese [1] and we conclude that both DENSE $r/2$ -STM and DENSE r -SD are fpt when parameterized by treewidth. ◀

Furthermore, it is not difficult (albeit tedious) to design a dynamic programming algorithm that solves DENSE $r/2$ -STM and DENSE r -SD in time $2^{O(\text{tw}^2)}n$. The quadratic dependence on the treewidth stems from the fact that we have to keep track of which edges we have contracted so far and there is no obvious way to circumvent this. The important question was whether any of the known techniques to reduce the complexity of connectivity-problems [6, 2, 19] could be applied here. The answer is, to our surprise, negative as we will discuss in Section 4.2.

4 Hardness results

4.1 NP-hardness and ETH lower bounds

This section will be dedicated to the proof the following theorems which both follow directly via a linear reduction from POSITIVE 1-IN-3SAT.

► **Theorem 2.** DENSE $r/2$ -STM and DENSE r -SD are NP-hard for $r \geq 1$, even when restricted to graphs that can be turned into subcubic planar graphs by deleting a single vertex.

► **Theorem 3.** DENSE $r/2$ -STM and DENSE r -SD cannot be solved in time $2^{o(n)}n^{O(1)}$ on bipartite graphs unless the ETH fails.

A special case of our result might be of independent interest:

► **Theorem 4.** DENSE BIPARTITE SUBDIVISION is NP-hard even on instances $((X, Y, E), d)$ where vertices in X have degree at most 3 and $d \geq 3$.

In the following, we present two reduction from POSITIVE 1-IN-3SAT that depend on the parity of r . We describe the reduction for $r \in \{1, 2\}$ and then argue how to modify the construction for arbitrary values of r . Note that the resulting instances are such that the densest graph H that appears as a $r/2$ -shallow topological minor appears, in fact, as an r -subdivision and thus the reductions work for both problems.

Reduction for r odd


Let ψ be a POSITIVE 1-IN-3SAT instance with clauses C_1, \dots, C_m and variables x_1, \dots, x_p . We assume that every variable in ψ appears in at least 3 clauses; if not, we can duplicate clauses to achieve this without changing the satisfiability of ψ . We construct a graph G from ψ in the following manner (cf. Figure 2):

1. For each variable x_i , create a cycle D_i with as many vertices as the frequency of x_i in ψ .
2. Create an apex-vertex a that is connected to every vertex of the cycles D_1, \dots, D_p .
3. For each clause $\{x_i, x_j, x_k\}$, add a vertex u_{ijk} to the graph and connect it to one vertex in D_i, D_j, D_k each that has not yet been connected to any clause-vertex.
4. Subdivide every edge appearing in the cycles D_1, \dots, D_p and all edges incident to the apex a .

For easier presentation, let us color the vertices of G as follows: the vertices introduced in the first step are *white*, the vertices introduced in the third step *gray*, and the subdivision vertices created in the last step *black* (the apex vertex a remains uncolored). Note that the graph G is bipartite, where one side of the partition contains exactly the white vertices and a .

Note that if the input formula ψ is planar, then the constructed graph is planar and subcubic after removing the apex vertex a .

► **Lemma 5.** If ψ is satisfiable then G has a topological minor at depth $1/2$ of density $\frac{5m}{2m+1}$.



■ **Figure 2** Three variable gadgets D_i, D_j, D_k connected by the gray clause vertex u_{ijk} . Setting variable x_k to true and x_i, x_j to false corresponds to the contractions on the right-hand side. The apex vertex a is not shown.

Proof. We construct the minor H by first smoothing each black vertex. Then, for each variable set to true, we delete the corresponding cycle D_i . Since ψ is satisfiable and each clause has exactly one variable set to true, this step deletes exactly one neighbor from each gray vertex. We complete the construction of H by smoothing out each gray vertex.

$V(H)$ consists of exactly two vertices corresponding to each clause plus the apex a , for a total of $2m + 1$ vertices. Since all vertices of H were colored white in G , a has degree $2m$. Aside from the edges incident to a , there are m edges from smoothing gray vertices and $2m$ edges from smoothing black vertices, which yield a total of $5m$ edges. Thus, we have found a minor at depth $1/2$ of density $\frac{5m}{2m+1}$. ◀

► **Lemma 6.** *If G has a topological minor at depth $1/2$ of density at least $\frac{5m}{2m+1}$, then the formula ψ is satisfiable.*

Proof. Let H be the densest shallow topological minor at depth $1/2$ and fix some model of H in G . We first argue that the nails of H consist only of white vertex and potentially the apex vertex a .

► **Claim.** The nails of H consist of the apex vertex a and some subset of white vertices.

First, since the density of H is greater than two, its minimum degree is at least three (the removal of a degree-two vertex would increase the density). Since black vertices have degree two in G , the nails of the model forming H therefore cannot be black. Accordingly, every black vertex either does not participate in the formation of H or it is smoothed into an edge.

Let us define G_b to be the graph obtained from G by smoothing all black vertices. Since black vertices have degree exactly two, this operation is uniquely defined. By the previous observation, H can be obtained from G_b by only smoothing gray vertices and taking a subgraph. This of course implies that the nails of H are all either gray, white, or the apex vertex a . Let us now exclude the first of these three cases: assume y is a gray nail of H in G_b . Again, the degree of y in H must be at least three to ensure maximal density of H , and since y has degree three in G it must also have degree exactly three in H . Note that the three neighbors of y are necessarily white and independent in G_b , thus we can smooth y into an (arbitrary) edge between two of its neighbors. The newly obtained graph H' is again a

half-shallow topological minor of G and it contains one vertex and two edges less than H . Since the density of H is greater than two, this implies that the density of H' is greater than that of H , a contradiction. We conclude that the nails of H cannot be gray and therefore only consist of white vertices and, potentially, the apex vertex a . To see that a must be contained in H , simply note that otherwise the maximum degree of H would be three and as thus H 's density would lie strictly below the assumed $\frac{5m}{2m+1}$. In summary: H 's nails consist of the apex vertex a and some subset of white vertices of G , proving the claim.

Since the white vertices in G are independent, the above claim further implies that the construction of H can be accomplished without smoothing white vertices. We can therefore divide said construction into two steps: first we smooth all gray and black vertices to construct a graph G_{gb} from G and then we take the subgraph $H \subseteq G_{gb}$. In the following, we will refer to edges in G_{gb} or H as *gray* if they originated from smoothing a gray vertex and *black* if they originated from smoothing a black vertex. Note that the set of black and gray edges partition $E(G_{gb})$ and hence also $E(H)$.

We now denote by v_4 the number of degree-four vertices in H and by v_3 the number of degree-three vertices (as observed above, no vertex with degree lower than three can exist in H and a is the only vertex of degree greater than four). Since the number of gray edges is at most m and a degree-four vertex must be incident to a gray edge, we have that $v_4 \leq 2m$. Let $w = v_3 + v_4$ be the number of white vertices in H and $\alpha = v_4/w$ the ratio of degree-four vertices among them. Using these quantities, we can express H 's density as

$$\frac{2v_3 + 5/2v_4}{v_3 + v_4 + 1} = 2 \frac{w}{w+1} + \frac{\alpha}{2} \frac{w}{w+1} = \left(2 + \frac{\alpha}{2}\right) \frac{w}{w+1}$$

which we combine with the density-requirement on H to obtain

$$\left(2 + \frac{\alpha}{2}\right) \frac{w}{w+1} \geq \frac{5m}{2m+1} \iff \alpha \geq 2 \left(\frac{5m}{2m+1} \frac{w+1}{w} \right) - 4.$$

Note that the right-hand side is equal to one for $w = 2m$, smaller than one for $w > 2m$, and larger than one for $w < 2m$. This last regime would imply the impossible $\alpha > 1$ and we conclude $2m \leq w \leq 3m$, where the upper bound $3m$ is simply the total number of white vertices in G . Rewriting w as βm for $2 \leq \beta \leq 3$, we revisit the density-constraint on H :

$$\begin{aligned} \left(2 + \frac{\alpha}{2}\right) \frac{\beta m}{\beta m + 1} \geq \frac{5m}{2m + 1} &\iff \left(2\beta + \frac{\alpha\beta}{2}\right)(2m + 1) \geq 5(\beta m + 1) \\ &\implies (\alpha - 1)m + \frac{5}{2} \geq \frac{5}{\beta}. \end{aligned} \tag{*}$$

We will now show that α , the fraction of degree-four vertices among all w white vertices, needs to be one in order for (*) to hold. To that end, we distinguish the following two cases:

Case 1: $\beta = 2$. Assuming $\alpha \neq 1$, the largest possible value for α is achieved when $v_4 = 2m - 2$ (the case of exactly one gray edge missing from H), resulting in $\alpha = (2m - 2)/2m = 1 - \frac{1}{m}$. Plugging this value of α and $\beta = 2$ into (*), we obtain that

$$(1 - 1/m - 1)m + \frac{5}{2} = -1 + \frac{5}{2} \geq \frac{5}{2},$$

a contradiction. Smaller values of α lead to the same contradiction, and we conclude that necessarily $\alpha = 1$.

Case 2: $2 < \beta \leq 3$. Assuming $\alpha \neq 1$, the largest possible value for α is achieved when $v_4 = 2m$, resulting in $\alpha = 2m/\beta m = \frac{2}{\beta}$. Now (*) becomes

$$\left(\frac{2}{\beta} - 1\right)m + \frac{5}{2} \geq \frac{5}{\beta} \iff m \leq \frac{10 - 5\beta}{2\beta} \cdot \frac{\beta}{2 - \beta} = \frac{5}{2}.$$

Thus for formulas ψ with at least three clauses, we arrive at a contradiction and conclude that $\alpha = 1$.

We have now shown that a) H contains *only* vertices of degree four and b) that $|H| \geq 2m$. Since there cannot be more than $2m$ vertices of degree four, we conclude that H has exactly $2m$ vertices. Note that therefore H must consist of a collection of black-edge cycles with a total of $2m$ vertices, each of which is incident to exactly one of the m gray edges. Note that each black cycle B_i in H corresponds to a cycle D_i (associated with variable x_i) in G , where B_i was constructed from D_i by smoothing black vertices. Thus we can associate every black cycle B_i in H with a variable x_i in ψ . We claim that setting all such variables x_i that have a black cycle C_i in H to false and all other variables to true is a 1-in-3 satisfying assignment of ψ .

Consider any clause $\{x_i, x_j, x_k\}$ in ψ . The corresponding gray vertex u_{ijk} in G was smoothed into a gray edge e_{ijk} in H , since all m gray are present in H . Accordingly, exactly two of the three black cycles B_i, B_j, B_k are contained in H . Thus the assignment constructed above will set exactly two of the variables x_i, x_j, x_k to false and one variable to true. This argument holds for every clause in ψ and we conclude that the constructed assignment is 1-in-3 satisfying, proving the lemma. ◀

This concludes the reduction from POSITIVE 1-IN-3SAT. Note that an optimal solution in the reduction necessarily does not use any edges from the original graph, but only edges resulting from contractions. Therefore the reduction works for both DENSE $\frac{1}{2}$ -STM and DENSE $\frac{1}{2}$ -SD. As noted above, for $r = 1$ the constructed graph is bipartite. Since the latter set has degree at most three, Theorem 4 follows.

In order for the reduction to work for arbitrary odd r , we need to modify the construction in two places: first, we subdivide every edge in the clause gadget $(r - 1)/2$ times. Second, instead of subdividing all edges appearing in the cycles D_1, \dots, D_p and edges incident to the apex a once, we subdivide them r times. The correctness of this reduction follows from easy modifications to Lemma 5 and 6, concluding our proof of Theorem 2 for odd values of r . Finally, to see that the above reduction also proves Theorem 3 for odd r , simply note that the reduction results in a graph of size $\Theta(m)$ and the ETH lower bound follows from Proposition 1.

Reduction for r even

Let ψ be a POSITIVE 1-IN-3SAT instance as described above. Construct graph G in the following manner. We once again create a cycle D_i for each variable x_i , connect an apex vertex a to each vertex on the cycles, and color these vertices white. For this construction, however, we subdivide all edges between white vertices twice i.e. each white-white edge is replaced by a three-edge path. As with our previous construction, the subdivision vertices are all colored black. For each clause $C_i = \{x_j, x_k, x_\ell\}$, we add a triangle u_{ij}, u_{ik}, u_{il} and connect it to the vertices from D_j, D_k , and D_ℓ corresponding to C_i such that u_{ij} is incident to the vertex from D_j etc. We color each of these vertices gray.

► **Lemma 7.** *If ψ is satisfiable then G has a topological minor at depth 1 of density $\frac{5m}{2m+1}$.*

Proof. We construct the minor H by first smoothing each black vertex. Then, for each variable set to true, we delete the corresponding cycle D_i . Since ψ is satisfiable and each clause has exactly one variable set to true, each gray triangle has two vertices of degree three and one of degree two. The degree two gray vertices are deleted, leaving the remaining gray

vertices to lie on three-edge paths between white vertices. These paths are subsequently smoothed to create white-white edges.

$V(H)$ consists of exactly two vertices corresponding to each clause plus a , for a total of $2m + 1$ vertices. Since all vertices of H were colored white in G , a has degree $2m$. Aside from the edges incident to a , there are m edges from smoothing gray vertices and $2m$ edges from smoothing black vertices, which yield a total of $5m$ edges. Thus, we have found a minor at depth 1 of density $\frac{5m}{2m+1}$. ◀

► **Lemma 8.** *If G has a topological minor at depth 1 of density $\frac{5m}{2m+1}$ then ψ is satisfiable.*

Proof. Let H' be the densest topological minor at depth 1. For the same reasons presented in Lemma 6, H' has no black nails, and thus we can smooth all black vertices into white-white edges. This lack of black nails also implies that no white vertices can be smoothed to form a new edge incident to a gray vertex.

If H' contains all the gray and white vertices, it has $3m$ degree 4 white vertices, $3m$ degree 3 gray vertices, and a with degree $3m$ for a total of $12m$ edges and $6m + 1$ vertices. This implies a density below $\frac{5m}{2m+1}$, and thus not all white and gray vertices are nails.

Since the gray vertices induce triangles, there is no way to smooth gray vertices to create a new gray-gray edge. Consider one such triangle T_i . If we smooth two vertices in T_i to create a single gray-white edge, the gray nail has degree 2 and should be deleted instead to increase the density. On the other hand, smoothing exactly one gray vertex to create a gray-white edge cause the remaining gray vertex to have degree two. Thus, any gray nail in H' is adjacent to three white vertices. Note that instead of having a gray nail, we could delete one gray vertex and smooth the other two into a white-white edge. The proof in Lemma 6 already demonstrated that forming the white-white edges is necessary to yield a density of $\frac{5m}{2m+1}$, and thus H' has no gray nails.

Since the gray vertices must be smoothed and deleted to create two degree 4 vertices and one degree 3 vertex per clause, the arguments in Lemma 6 imply that for H' to have density $\frac{5m}{2m+1}$, ψ must be satisfiable. ◀

In order for the reduction to work for arbitrary even r , we again modify the construction in two places: first, we subdivide every edge of the triangle making up the clause gadget $r/2 - 1$ times. Second, instead of subdividing all edges appearing in the cycles D_1, \dots, D_p and edges incident to the apex a twice, we subdivide them r times. With both cases of r even or odd covered, we conclude that Theorem 2 and Theorem 3 hold true.


4.2 Excluding a $2^{o(\text{tw}^2)} n^{O(1)}$ -algorithm

We show in this section that the ETH implies that we cannot get a single-exponential algorithm parameterized by treewidth for DENSE $r/2$ -STM for $r \geq 2$.

► **Theorem 9.** *Unless the ETH fails, there is no algorithm that decides DENSE 1-SHALLOW TOPOLOGICAL MINOR on a graph with treewidth t in time $2^{o(t^2)} n^{O(1)}$.*

Our proof proceeds via a reduction from CNF-SAT. Assume that the CNF formula Φ with variables x_1, \dots, x_n and clauses C_1, \dots, C_m is such that \sqrt{n} is an even integer; if not, we pad Φ with dummy variables that appear in no clauses, which does not affect the answer to Φ . Figure 3 contains a sketch of the construction outlined in the following.

Decision gadget: The reduction will use sequences of vertices connected by decision gadgets. The decision gadget is a path of three vertices d_L, d_C, d_R which we will always connect to



■ **Figure 3** A sketch of the construction for Theorem 9, with an exemplary connection of the variable-path X_1 to the first clause gadget (here, x_1 appears negatively in C_1). Dashed edges denote parts that are actually connected via decision gadgets. 3-paths between the grid R and the clause gadgets (A_i, B_i) are not drawn.

a sequence of three vertices. For a sequence of vertices v_1, v_2, v_3 , connecting the decision gadget to the sequence involves adding the edges $\{d_L, v_1\}$, $\{d_C, v_2\}$, and $\{d_R, v_3\}$.

Variable gadgets: We construct a grid of vertices R with \sqrt{n} rows and m columns, denoting with $R[i, j]$ the vertex in the i th row and j th column. Each variable x_i will be represented by a sequence of m vertices X_i , one from each column. We will denote with $X_i[j]$ the j th vertex in the sequence X_i for any $1 \leq j \leq m$. In order to represent each of the n variables with m vertices using a $\sqrt{n} \times m$ grid, the sequences X_i must overlap and each vertex in R is part of the representation of \sqrt{n} variables. Specifically, let $X_i[j] = R[(i + j \lfloor i/\sqrt{n} \rfloor) \bmod \sqrt{n}, j]$. In other words, two vertices in successive columns will be in the same row in sequences $X_1, \dots, X_{\sqrt{n}}$; they will be one row apart (“wrapping around” to the top from the bottom) in sequences $X_{1+\sqrt{n}}, \dots, X_{2\sqrt{n}}$, two rows apart in $X_{1+2\sqrt{n}}, \dots, X_{3\sqrt{n}}$, and so on.

For each sequence X_i , we connect $X_i[j - 1]$, $X_i[j]$, and $X_i[j + 1]$ to a decision gadget. Denote such a decision gadget as $D_{i,j}$. We also “wrap around” X_i by connecting $X_i[m - 1]$, $X_i[m]$, $X_i[1]$ and $X_i[m]$, $X_i[1]$, $X_i[2]$ to their own decision gadgets.

Clause gadgets: Each clause C_i will be represented by a bipartition of vertices A_i, B_i where $|A_i| = |B_i| = \sqrt{n}$. Let $A_i[j]$ be the j th vertex in A_i and $B_i[j]$ likewise. Let σ be an ordering of the vertices in $A_i \cup B_i$ corresponding to an Eulerian tour of a biclique with bipartition A_i, B_i . Assume without loss of generality that $\sigma_i = A_i[\sqrt{n}], B_i[1], \dots, B_i[\sqrt{n}], A_i[\sqrt{n}]$ and note that every pair of vertices $a \in A_i$ and $b \in B_i$ appears consecutively exactly once in σ . For each consecutive triple of vertices in σ_i , attach a decision gadget (but do not “wrap around”).

Connecting variables and clauses: For each pair of vertices $(A_i[j], B_i[k])$ for $1 \leq j, k \leq \sqrt{n}$ assign the pair with a unique variable x_ℓ . Connect $X_\ell[i]$ to $A_i[j]$ and $B_i[k]$ via 3-edge paths. If x_ℓ appears in clause C_i positively, connect d_L of the decision gadget $D_{i,\ell}$ to $B_i[k]$ via an edge and to $A_i[j]$ via a 2-edge path. If it appears negatively, add the same connections to d_R of $D_{i,\ell}$ instead.

With the description of the reduction completed, let us now prove its correctness, *i.e.* We prove that G has a 1-STM of density $\rho = \frac{4\sqrt{n}}{3}$ if and only if Φ is satisfiable.

Forward direction: To prove the forward direction, we show how the satisfying assignment yields a topological minor of the desired density. We note that a cyclical sequence of vertices joined by decision gadgets can form a cycle in one of two ways: by smoothing each d_C and d_L or each d_C and d_R . Let the former be known as the *left configuration* of those sequenced gadgets and the latter the *right configuration*. Create a cycle on the vertices in X_i by choosing the right configuration if x_i is true and the left configuration if x_i is false.

For each clause, pick an arbitrary variable x_ℓ that satisfies it and let a and b be the pair of vertices from A_i and B_i assigned to that variable. If $a <_\sigma b$, set all decision gadgets preceding a in σ to the left configuration and all the decision gadgets succeeding b to the right configuration; do the reverse if $b <_\sigma a$. Thus A_i and B_i form a biclique missing the ab edge. Since there is a 3-edge path from a to b in G through a vertex in $D_{i,\ell}$ that has not been smoothed, we can use that path to form the ab edge. Smooth the remaining 3-edge induced paths in G and delete the vertices from the decision gadgets that were not contracted.

The nails of the resulting minor are exactly $\mathcal{A} \cup \mathcal{B} \cup R$. There are $m\sqrt{n}$ vertices in R and each one participates in \sqrt{n} variable gadgets. Since each variable gadget becomes a cycle there are mn edges within R . The m clause gadgets become bicliques on $2\sqrt{n}$ vertices each and thus contain mn edges in total. Each variable gadget ends up with two edges into each clause gadget, for a total of $2mn$ edges connecting them. In total, this makes $4mn$ edges and $3m\sqrt{n}$ vertices, exactly ρ .

Reverse direction: We now prove the reverse direction by assuming Φ is unsatisfiable. Let H' be a 1-STM with density ρ . Since ρ is $\Theta(\sqrt{n})$ and the vertices in the induced paths and decision gadgets have degree at most 4, we can assume that none of those vertices appears as a nail in H' . Thus, the nail set of H' is a subset of $\mathcal{A} \cup \mathcal{B} \cup R$. The only paths between these nail candidates that use at most three edges do not contain nail candidates as interior vertices, meaning nail candidates are never smoothed. Let $H(\Phi)$ be the minor constructed by the process described in the forward direction proof for an arbitrary satisfying assignment of Φ . Observe that for fixed n and m , the minor $H(\psi)$ is identical for every satisfiable formula ψ ; let $\mathcal{H}(n, m)$ be that minor. Moreover, every pair of nail candidates that has a 3-edge path between them in G is adjacent in $\mathcal{H}(n, m)$, meaning that H' is either a subgraph of or identical to $\mathcal{H}(n, m)$.

We now show that no proper induced subgraph of $\mathcal{H}(n, m)$ has density ρ . If a graph is d -regular and connected, it has edge density $d/2$; deleting part of a degree regular graph leaves vertices with degree less than d , so no proper subgraph reaches that density. Therefore R and each $A_i \cup B_i$ achieve their maximum densities of \sqrt{n} and $\sqrt{n}/2$ only when including the entire subgraph, implying a dense subgraph must contain portions of both vertex and clause gadgets. The only vehicle for increasing density is to use edges between vertex and clause gadgets, which means we should only include a vertex in R in a subgraph if it also contains its neighbors from \mathcal{A} and \mathcal{B} . Let G' be the subgraph of G induced on an $i \times j$ subgrid of R and all of its neighbors in $\mathcal{A} \cup \mathcal{B}$. The density within the subset of $\mathcal{A} \cup \mathcal{B}$ is greatest when those vertices induce $j \times j$ bicliques, so we assume they do. Under this assumption, the density of G' is $\frac{4i}{3}$ if $j = m$ and $\frac{(4j-1)i}{3j}$ if $j < m$ since the edges that wrap around R cannot be realized. In either case, the density is strictly less than ρ unless $j = m$ and $i = \sqrt{n}$ *i.e.* exactly $\mathcal{H}(n, m)$.

A decision gadget connected to sequential vertices v_1, v_2, v_3 can only create the edge v_1, v_2 or the edge v_2, v_3 , since d_C needs to be smoothed to construct either edge. Consequently, choosing to set some decision gadgets in the same variable gadget to opposite configurations

(or neither configuration) creates at least one fewer edge than if they were all set to the same decision. Thus, the configurations of the variable gadgets correspond to some truth assignment to Φ in the way intended in $H(\Phi)$. This indicates that there is a clause gadget that has no neighbors in a variable gadget that can be used to be smoothed into an edge in the clause gadget. However, because there is one fewer decision gadget than the number of biclique edges in the clause gadgets of $\mathcal{H}(n, m)$, H' cannot realize all possible edges in the biclique and thus there is no 1-STM of density ρ .

It stands to prove that the above reduction has the proper implications for a parameterization by treewidth. Using cops and robbers, we can show that G has treewidth $O(\sqrt{n})$ as follows: We permanently station \sqrt{n} cops on the first column of R . We use $2\sqrt{n}$ cops to walk through the columns of R sequentially; when a column is completely covered with cops we can explore its corresponding clause gadget with a separate unit of $2\sqrt{n}$ cops. In total, this requires $O(\sqrt{n})$ cops. Although the formula Φ may have been padded with additional variables, it would only have been enough to increase \sqrt{n} by 2. This means the number of variables in the unpadded instance is still $\Theta(n)$. Thus, if an algorithm parameterized by treewidth t could find a dense 1-STM in time $2^{o(t^2)}n^{O(1)}$, then we could use our reduction to solve CNF-SAT in time $2^{o(n)}n^{O(1)}$, violating ETH. This concludes the proof of Theorem 9.

An immediate consequence is that, unlike DENSE $\frac{1}{2}$ -STM, DENSE 1-STM is still NP-hard when the exact nail set is known.

5 Conclusion

We showed that finding dense substructures that are just slightly less local than subgraphs is computationally hard, and even a parameterization by treewidth cannot provide very efficient algorithms. While our first reduction excludes a subexponential exact algorithm assuming the ETH, we could not exclude an algorithm with a running time of $(2 - \varepsilon)^n n^{O(1)}$. Is such an algorithm possible for $r = 1$, or can one find a tighter reduction that provides a corresponding SETH lower bound? Our second reduction rules out a $2^{o(\text{tw}^2)}n^{O(1)}$ -algorithm for $r = 2$. Is a faster algorithm for $r = 1$ possible?

Finally, we ask whether there is a sensible notion of substructures that fit in between $\frac{1}{2}$ -shallow topological minors and subgraphs for which we can find the densest occurrence in polynomial time.

References

- 1 S. Arnborg, J. Lagergren, and D. Seese. Easy problems for tree-decomposable graphs. *Journal of Algorithms*, 12(2):308–340, 1991.
- 2 H. L. Bodlaender, M. Cygan, S. Kratsch, and J. Nederlof. Deterministic single exponential time algorithms for connectivity problems parameterized by treewidth. In *International Colloquium on Automata, Languages, and Programming*, pages 196–207. Springer, 2013.
- 3 H. L. Bodlaender, T. Wolle, and A. Koster. Contraction and treewidth lower bounds. *J. Graph Algorithms Appl.*, 10(1):5–49, 2006.
- 4 J. Chen, X. Huang, I. A. Kanj, and G. Xia. Strong computational lower bounds via parameterized complexity. *Journal of Computer and System Sciences*, 72(8):1346–1367, 2006.
- 5 M. Cygan, F.V. Fomin, L. Kowalik, D. Lokshantov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. Lower bounds based on the exponential-time hypothesis. In *Parameterized Algorithms*, pages 467–521. Springer, 2015.

- 6 M. Cygan, J. Nederlof, M. Pilipczuk, M. Pilipczuk, van J.M.M. Rooij, and J. O. Wojtaszczyk. Solving connectivity problems parameterized by treewidth in single exponential time. In *Proceedings of the 52nd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 150–159. IEEE Computer Society, 2011.
- 7 E. D. Demaine and M. Hajiaghayi. The bidimensionality theory and its algorithmic applications. *The Computer Journal*, 51(3):292–302, 2008.
- 8 F. Dorn, F. V. Fomin, and D. M. Thilikos. Subexponential parameterized algorithms. *Computer Science Review*, 2(1):29–39, 2008.
- 9 P. G. Drange, M. Dregi, F.V. Fomin, S. Kreutzer, D. Lokshtanov, M. Pilipczuk, M. Pilipczuk, F. Reidl, S. Saurabh, F. Sánchez Villaamil, S. Siebertz, and S. Sikdar. Kernelization and sparseness: the case of dominating set. In *33rd Symposium on Theoretical Aspects of Computer Science*, 2016.
- 10 Z. Dvořák. *Asymptotical Structure of Combinatorial Objects*. PhD thesis, Charles University, Faculty of Mathematics and Physics, 2007.
- 11 Z. Dvořák, D. Král, and R. Thomas. Deciding first-order properties for sparse graphs. In *Proceedings of the 51st IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 133–142. IEEE Computer Society, 2010.
- 12 F. V. Fomin, D. Lokshtanov, S. Saurabh, and D. M. Thilikos. Linear kernels for (connected) dominating set on H -minor-free graphs. In *Proceedings of the 23rd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 82–93. SIAM, 2012.
- 13 J. Gajarský, P. Hliněný, J. Obdržálek, S. Ordyniak, F. Reidl, P. Rossmanith, F. Sánchez Villaamil, and S. Sikdar. Kernelization using structural parameters on sparse graph classes. *To appear in Journal of Computer and System Sciences*, 2016.
- 14 G. Gallo, M. D. Grigoriadis, and R. E. Tarjan. A fast parametric maximum flow algorithm and applications. *SIAM Journal on Computing*, 18(1):30–55, 1989.
- 15 A. V. Goldberg. *Finding a maximum density subgraph*. University of California Berkeley, CA, 1984.
- 16 M. Grohe, S. Kreutzer, and S. Siebertz. Deciding first-order properties of nowhere dense graphs. In *Proceedings of the 46th ACM Symposium on Theory of Computing (STOC)*, pages 89–98, 2014.
- 17 W. Mulzer and G. Rote. Minimum-weight triangulation is NP-hard. *Journal of the ACM (JACM)*, 55(2):11, 2008.
- 18 J. Nešetřil and P. Ossona de Mendez. *Sparsity: Graphs, Structures, and Algorithms*, volume 28 of *Algorithms and Combinatorics*. Springer, 2012.
- 19 M. Pilipczuk. Problems parameterized by treewidth tractable in single exponential time: a logical approach. In *International Symposium on Mathematical Foundations of Computer Science*, pages 520–531. Springer, 2011.
- 20 F. Reidl. *Structural sparseness and complex networks*. Dr., Aachen, Techn. Hochsch., Aachen, 2016. Aachen, Techn. Hochsch., Diss., 2015. URL: <http://publications.rwth-aachen.de/record/565064>.
- 21 T. J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the tenth annual ACM symposium on Theory of computing*, pages 216–226. ACM, 1978.