

APPROXIMATING ACYCLICITY PARAMETERS OF SPARSE HYPERGRAPHS

FEDOR V. FOMIN¹ AND PETR A. GOLOVACH¹ AND DIMITRIOS M. THILIKOS²

¹ Department of Informatics, University of Bergen, N-5020 Bergen, Norway. Supported by the Norwegian Research Council.

² Department of Mathematics, University of Athens, Panepistimioupolis, GR15784 Athens, Greece. Supported by the project “Kapodistrias” (ΑΠ 02839/28.07.2008) of the National and Kapodistrian University of Athens.

ABSTRACT. The notions of hypertree width and generalized hypertree width were introduced by Gottlob, Leone, and Scarcello (PODS’99, PODS’01) in order to extend the concept of hypergraph acyclicity. These notions were further generalized by Grohe and Marx in SODA’06, who introduced the fractional hypertree width of a hypergraph. All these width parameters on hypergraphs are useful for extending tractability of many problems in database theory and artificial intelligence. Computing each of these width parameters is known to be an NP-hard problem. Moreover, the (generalized) hypertree width of an n -vertex hypergraph cannot be approximated within a logarithmic factor unless $P=NP$. In this paper, we study the approximability of (generalized, fractional) hypertree width of sparse hypergraphs where the criterion of sparsity reflects the sparsity of their incidence graphs. Our first step is to prove that the (generalized, fractional) hypertree width of a hypergraph is constant-factor sandwiched by the treewidth of its incidence graph, when the incidence graph belongs to some apex-minor-free graph class (the family of apex-minor-free graph classes includes planar graphs and graphs of bounded genus). This determines the combinatorial borderline above which the notion of (generalized, fractional) hypertree width becomes essentially more general than treewidth, justifying that way its functionality as a hypergraph acyclicity measure. While for more general sparse families of hypergraphs treewidth of incidence graphs and all hypertree width parameters may differ arbitrarily, there are sparse families where a constant factor approximation algorithm is possible. In particular, we give a constant factor approximation polynomial time algorithm for (generalized, fractional) hypertree width on hypergraphs whose incidence graphs belong to some H-minor-free graph class. This extends the results of Feige, Hajiaghayi, and Lee from STOC’05 on approximating treewidth of H-minor-free graphs.

1. Introduction

Many important theoretical and “real-world” problems can be expressed as constrained satisfaction problems (CSP). Among examples one can mention numerous problems from

1998 ACM Subject Classification: F.2.2,G.2.2.

Key words and phrases: Graph, hypergraph, hypertree width, treewidth.



© F. V. Fomin, P. A. Golovach, and D. M. Thilikos
Creative Commons Attribution-NoDerivs License

different domains like Boolean satisfiability, temporal reasoning, graph coloring, belief maintenance, machine vision, and scheduling. Another example is the conjunctive-query containment problem, which is a fundamental problem in database query evaluation. In fact, as it was shown by Kolaitis and Vardi [19], CSP, conjunctive-query containment, and finding homomorphism for relational structures are essentially the same problem. The problem is known to be **NP**-hard in general [2] and polynomial time solvable for restricted class of acyclic queries [25]. Recently, in the database and constraint satisfaction communities various extensions of query (or hypergraph) acyclicity were studied. The main motivation for the quest for a suitable measure of acyclicity of a hypergraph (query, or relational structure) is the extension of polynomial time solvable cases (like acyclic hypergraphs) to more general instances. In this direction, Chekuri and Rajaraman in [3] introduced the notion of query width. Gottlob, Leone, and Scarcello [13, 14, 16] defined hypertree width and generalized hypertree width. Furthermore, Grohe and Marx [18] have introduced the most general parameter known so far, fractional hypertree width, and proved that CSP, restricted to instances of bounded fractional hypertree width, is polynomial time solvable.

Unfortunately, all known variants of hypertree width are **NP**-complete [12, 17]. Moreover, generalized hypertree width is **NP**-complete even when checking whether its value is at most 3 (see [17]). In the case of hypertree width, the problem is $W[2]$ -hard when parameterized by k [12]. Both hypertree width and the generalized hypertree are hard to approximate. For example, the reduction of Gottlob et al. in [12] can be used to show that the generalized hypertree width of an n -vertex hypergraph cannot be approximated within a factor $c \log n$ for some constant $c > 0$ unless $P = NP$.

All these parameters for hypergraphs can be seen as generalizations of the treewidth of a graph. The treewidth is a fundamental graph parameter from Graph Minors Theory by Robertson and Seymour [22] and it has numerous algorithmic applications. It is an old open question whether the treewidth can be approximated within a constant factor and the best known approximation algorithm for treewidth is $\sqrt{\log OPT}$ -approximation due to Feige et al. [9]. However, as it was shown by Feige et al. [9], the treewidth of an H -minor-free graph is constant factor approximable.

Our results. Our first result is combinatorial. We show that for a wide family of hypergraphs (those where the incidence graph excludes an apex graph as a minor – that is a graph that can become planar after removing a vertex) the fractional and generalized hypertree width of a hypergraph is bounded by a linear function of treewidth of its incidence graph. Apex-minor-free graph classes include planar and bounded genus graphs.

For hypergraphs whose incidence graphs are apex graphs the two parameters may differ arbitrarily, and this result determines the boundary where fractional hypertree width starts being essentially different from treewidth of the incidence graph. This indicates that hypertree width parameters are more useful as the adequate version of acyclicity for non-sparse instances.

Our proof is based on theorems from bidimensionality theory and a min-max (in terms of fractional hyperbrambles) characterization of fractional hypertree width. The proof essentially identifies what is the obstruction analogue of fractional hypertree width for incidence graphs.

Our second result applies further for sparse classes where the difference between (generalized, fractional) hypertree width of a hypergraph and treewidth of its incidence graph can be arbitrarily large. In particular, we give a constant factor approximation algorithm

for generalized and fractional hypertree width of hypergraphs with H -minor-free incidence graphs extending the results of Feige et al. [9] from treewidth to (generalized, fractional) hypertree width. The algorithm is based on a series of theorems based on the main decomposition theorem of the Robertson-Seymour's Graph Minor project. As a combinatorial corollary of our results, it follows that generalized hypertree width and fractional hypertree width differ within constant multiplicative factor if the incidence graph of the hypergraph does not contain a fixed graph as a minor.

Due the space restrictions some proofs are omitted here. They will appear in the journal paper, but also they can be found in our technical report[10].

2. Definitions and preliminaries

Basic definitions and properties. We consider finite undirected graphs without loops or multiple edges. The vertex set of a graph G is denoted by $V(G)$ and its edge set by $E(G)$ (or simply by V and E if it does not create confusion).

Let G be a graph. For a vertex v , we denote by $N_G(v)$ its (*open*) *neighborhood*, i.e. the set of vertices which are adjacent to v . The *closed neighborhood* of v , i.e. the set $N_G(v) \cup \{v\}$, is denoted by $N_G[v]$. For $U \subseteq V(G)$, we define $N_G[U] = \bigcup_{v \in U} N_G[v]$ (we may omit index if the graph under consideration is clear from the context). If $U \subseteq V(G)$ (or $u \in V(G)$) then $G - U$ (or $G - u$) is the graph obtained from G by the removal of vertices of U (vertex u correspondingly). Given an edge $e = \{x, y\}$ of a graph G , the graph G/e is obtained from G by contracting e ; which is, to get G/e we identify the vertices x and y and remove all loops and replace all multiple edges by simple edges. A graph H obtained by a sequence of edge-contractions is said to be a *contraction* of G . A graph H is a *minor* of G if H is a subgraph of a contraction of G . We say that a graph G is *H -minor-free* when it does not contain H as a minor. We also say that a graph class \mathcal{G} is *H -minor-free* (or, excludes H as a minor) when all its members are H -minor-free. An *apex graph* is a graph obtained from a planar graph G by adding a vertex and making it adjacent to some of the vertices of G . A graph class \mathcal{G} is *apex-minor-free* if \mathcal{G} excludes a fixed apex graph H as a minor. The $(k \times k)$ -*grid* is the Cartesian product of two paths of lengths $k - 1$. A *surface* Σ is a compact 2-manifold (we always consider connected surfaces). Whenever we refer to a Σ -*embedded graph* G we consider a 2-cell embedding of G in Σ . To simplify notations, we do not distinguish between a vertex of G and the point of Σ used in the drawing to represent the vertex or between an edge and the line representing it. We also consider a graph G embedded in Σ as the union of the points corresponding to its vertices and edges. That way, a subgraph H of G can be seen as a graph H , where $H \subseteq G$. Recall that $\Delta \subseteq \Sigma$ is a (closed) disc if it is homeomorphic to $\{(x, y) : x^2 + y^2 \leq 1\}$. The *Euler genus* of a nonorientable surface Σ is equal to the nonorientable genus $\tilde{g}(\Sigma)$ (or the crosscap number). The *Euler genus* of an orientable surface Σ is $2g(\Sigma)$, where $g(\Sigma)$ is the orientable genus of Σ .

If $X \subseteq 2^A$ for some set A , then by $\bigcup X$ we denote the union of all elements of X . Recall that a *hypergraph* \mathcal{H} is a pair $\mathcal{H} = (V(\mathcal{H}), E(\mathcal{H}))$ where $V(\mathcal{H})$ is a finite nonempty set of vertices, and $E(\mathcal{H})$ is a set of nonempty subsets of $V(\mathcal{H})$ called *hyperedges*, $\bigcup E(\mathcal{H}) = V(\mathcal{H})$. We consider here only hypergraphs without isolated vertices (i.e. every vertex is in some hyperedge). For vertex $v \in V(\mathcal{H})$, we denote by $E_{\mathcal{H}}(v)$ the set of its incident hyperedges. The *incidence graph* of the hypergraph \mathcal{H} is the bipartite graph $I(\mathcal{H})$ with vertex set $V(\mathcal{H}) \cup E(\mathcal{H})$ such that $v \in V(\mathcal{H})$ and $e \in E(\mathcal{H})$ are adjacent in $I(\mathcal{H})$ if and only if $v \in e$.

Treewidth of graphs and hypergraphs. A *tree decomposition* of a hypergraph \mathcal{H} is a pair (T, χ) , where T is a tree and $\chi: V(T) \rightarrow 2^{V(\mathcal{H})}$ is a function associating a set of vertices $\chi(t) \subseteq V(\mathcal{H})$ (called a *bag*) to each node t of the decomposition tree T such that i) $V(\mathcal{H}) = \bigcup_{t \in V(T)} \chi(t)$, ii) for each $e \in E(\mathcal{H})$, there is a node $t \in V(T)$ such that $e \subseteq \chi(t)$, and iii) for each $v \in V(\mathcal{H})$, the set $\{t \in V(T) : v \in \chi(t)\}$ forms a subtree of T .

The *width* of a tree decomposition equals $\max\{|\chi(t)| - 1 : t \in V(T)\}$. The *treewidth* of a hypergraph \mathcal{H} is the minimum width over all tree decompositions of \mathcal{H} . We use notation $\mathbf{tw}(\mathcal{H})$ for the treewidth of a hypergraph \mathcal{H} .

It is easy to verify that for any hypergraph \mathcal{H} , $\mathbf{tw}(\mathcal{H}) + 1 \geq \mathbf{tw}(I(\mathcal{H}))$. However, these parameters can differ considerably on hypergraphs. For example, for the n -vertex hypergraph \mathcal{H} with one hyperedge which contains all vertices, $\mathbf{tw}(\mathcal{H}) = n - 1$ and $\mathbf{tw}(I(\mathcal{H})) = 1$.

Since $\mathbf{tw}(\mathcal{H}) \geq |e|$ for every $e \in E(\mathcal{H})$, we have that the presence of a large hyperedge results in a large treewidth of the hypergraph. The paradigm shift in the transition from treewidth to hypertree width consists in counting the covering hyperedges rather than counting the number of vertices in a bag. This parameter seems to be more appropriate, especially with respect to constraint satisfaction problems. We start with the introduction of even more general parameter of fractional hypertree width.

Hypertree width and its generalizations. In general, given a finite set A , we use the term *labelling of A* for any function $\gamma : A \rightarrow [0, 1]$. We also use the notation $\mathcal{G}(A)$ for the collection of all labellings of a set A . The *size* of a labelling of A is defined as $|\gamma| = \sum_{x \in A} \gamma(x)$. If the values of a labelling γ are restricted to be 0 or 1, then we say that γ is a *binary* labelling of A . Clearly, the size of a binary labelling is equal to the number of the elements of A that are labelled by 1. Given a hyperedge labelling γ of a hypergraph \mathcal{H} , we define the set of vertices of \mathcal{H} that are *blocked* by γ as

$$B(\gamma) = \{v \in V(\mathcal{H}) \mid \sum_{e \in E_{\mathcal{H}}(v)} \gamma(e) \geq 1\},$$

i.e. the set of vertices that are incident to hyperedges whose total labelling sums up to 1 or more.

A *fractional hypertree decomposition* [18] of \mathcal{H} is a triple (T, χ, λ) , where (T, χ) is a tree decomposition of \mathcal{H} and $\lambda: V(T) \rightarrow \mathcal{G}(E(\mathcal{H}))$ is a function, assigning a hyperedge labeling to each node of T , such that for every $t \in V(T)$, $\chi(t) \subseteq B(\lambda(t))$, i.e. all vertices of the bag $\chi(t)$ are blocked by the labelling $\lambda(t)$. The *width* of a fractional hypertree decomposition (T, χ, λ) is $\min\{|\lambda(t)| : t \in V(T)\}$, and the *fractional hypertree width* $\mathbf{fhw}(\mathcal{H})$ of \mathcal{H} is the minimum of the widths of all fractional hypertree decompositions of \mathcal{H} . If λ assigns a binary hyperedge labeling to each node of T , then (T, χ, λ) is a *generalized hypertree decomposition* [15]. Correspondingly, the *generalized hypertree width* $\mathbf{ghw}(\mathcal{H})$ of \mathcal{H} is the minimum of the widths of all generalized hypertree decompositions of \mathcal{H} . Clearly, $\mathbf{fhw}(\mathcal{H}) \leq \mathbf{ghw}(\mathcal{H})$ but, as it was shown in [18], there are families of hypergraphs of bounded fractional hypertree width but unbounded generalized hypertree width. Notice that computing the fractional hypertree width is an **NP**-complete problem even for sparse graphs. To see this, take a connected graph G that is not a tree and construct a new graph H by replacing every edge of G by $|V(G)| + 1$ paths of length 2. It is easy to check that $\mathbf{tw}(G) + 1 = \mathbf{fhw}(H)$.

The proof of the next lemma follows from results of [3] about query width. For completeness, a direct proof is given in [10].

Lemma 2.1. *For any hypergraph \mathcal{H} , $\mathbf{fhw}(\mathcal{H}) \leq \mathbf{ghw}(\mathcal{H}) \leq \mathbf{tw}(I(\mathcal{H})) + 1$. [Proof in [10]]*

It is necessary to remark here that the fractional hypertree width of a hypergraph can be arbitrarily smaller than the treewidth of its incidence graph. Suppose that a hypergraph \mathcal{H}' is obtained from the hypergraph \mathcal{H} by adding a hyperedge which includes all vertices. Then $\mathbf{fhw}(\mathcal{H}') = 1$ and $\mathbf{tw}(I(\mathcal{H}')) + 1 \geq \mathbf{tw}(I(\mathcal{H})) + 1 \geq \mathbf{fhw}(\mathcal{H})$.

Hyperbrambles. Let \mathcal{H} be a hypergraph. Two sets $X, Y \subseteq V(\mathcal{H})$ *touch* if $X \cap Y \neq \emptyset$ or there exists $e \in E(\mathcal{H})$ such that $e \cap X \neq \emptyset$ and $e \cap Y \neq \emptyset$. A *hyperbramble* of \mathcal{H} is a set \mathcal{B} of pairwise touching connected subsets of $V(\mathcal{H})$ [1]. We say that a labelling γ of $E(\mathcal{H})$ *covers* a vertex set $S \subseteq V(\mathcal{H})$ if some of its vertices are blocked by γ . The *fractional order* of a hyperbramble is the minimum k for which there is a labeling γ of size at most k covering all elements in \mathcal{B} . The *fractional hyperbramble number*, $\mathbf{fhn}(\mathcal{H})$, of \mathcal{H} is the maximum of the fractional orders of all hyperbrambles of \mathcal{H} . Using [18, Theorem 11], we can prove the following lemma.

Lemma 2.2. *For any hypergraph \mathcal{H} , $\mathbf{fhn}(\mathcal{H}) \leq \mathbf{fhw}(\mathcal{H})$.* [Proof in [10]]

i-brambles. An *i-labeled graph* G is a triple (G, N, M) where $N, M \subseteq V(G)$, $N \cup M = V(G)$, $M - N$ and $N - M$ are independent sets of G , and for any $v \in V(G)$ its closed neighborhood $N_G[v]$ is intersecting both N and M . Notice that $\{N, M\}$ is not necessarily a partition of $V(G)$. The incidence graph $I(\mathcal{H})$ of a hypergraph \mathcal{H} can be seen as an *i-labeled graph* $(I(\mathcal{H}), N, M)$ where $N = V(\mathcal{H})$, $M = E(\mathcal{H})$.

The result of the contraction of an edge $e = \{x, y\}$ of an *i-labeled graph* (G, N, M) to a vertex v_e is the *i-labeled graph* (G', N', M') where i) $G' = G/e$ ii) N' contains all vertices of $N - \{x, y\}$ and also the vertex v_e , in case $\{x, y\} \cap N \neq \emptyset$ and iii) M' contains all vertices of $M - \{x, y\}$ and also the vertex v_e , in case $\{x, y\} \cap M \neq \emptyset$. An *i-labeled graph* (G', N', M') is a *contraction* of an *i-labeled graph* (G, N, M) if (G', N', M') can be obtained after applying a (possibly empty) sequence of contractions to (G, N, M) . The following lemma is a direct consequence of the definitions.

Lemma 2.3. *Let (G, N, M) be an i-labeled graph and let G' be a contraction of G . Then there are $N', M' \subseteq V(G')$ such that the i-labeled graph (G', N', M') is a contraction of (G, N, M) .*

Let (G, N, M) be an *i-labeled graph*. We say that a set $S \subseteq N$ is *i-connected* if any pair $x, y \in S$ is connected by a path in $G[S \cup M]$. We say that two subsets $S, R \subseteq N$ *i-touch* either if i) $S \cap R \neq \emptyset$, or ii) there is an edge $\{x, y\}$ with $x \in S$ and $y \in R$, or iii) there is a vertex $z \in M$ such that $N_G[z]$ intersects both S and R .

Given an *i-labeled graph* (G, N, M) we define an *i-bramble* of (G, N, M) as any collection \mathcal{B} of *i-touching i-connected* sets of vertices in N . We say that a labeling γ of M *blocks* a vertex $x \in N$ if $\sum_{y \in N_G[x] \cap M} \gamma(y) \geq 1$. We say that γ *fractionally covers* a vertex set $S \subseteq N$ if some of its vertices is blocked by γ . The *order* of an *i-bramble* is the minimum k for which there is a labeling γ of M of size at most k that fractionally covers all sets of \mathcal{B} .

The *fractional i-bramble number* $\mathbf{fhn}(G, N, M)$ of an *i-labeled graph* (G, N, M) is the maximum order of all *i-brambles* of it.

The following statement follows immediately from the definitions of hyperbrambles and *i-brambles*.

Lemma 2.4. *For any hypergraph \mathcal{H} , $\mathbf{fhn}(I(\mathcal{H}), V(\mathcal{H}), E(\mathcal{H})) = \mathbf{fhn}(\mathcal{H})$.*

Also it can be easily seen that the fractional *i-bramble number* is a contraction-closed parameter.

Lemma 2.5. *If an i -labeled graph (G', N', M') is the contraction of an i -labeled graph (G, N, M) then $\mathbf{fibn}(G', N', M') \leq \mathbf{fibn}(G, N, M)$.*

Obviously, i -bramble number is not a subgraph-closed parameter (not even for induced subgraphs), but we can note the following useful claim.

Lemma 2.6. *Let (G, N, M) be an i -labeled graph and $X \subseteq V(G)$ such that $G - X$ has no isolated vertices, and for every $v \in X \cap M$, $N_G[v] \subseteq X$. Then $(G - X, N - X, M - X)$ is an i -labeled graph and $\mathbf{fibn}(G - X, N - X, M - X) \leq \mathbf{fibn}(G, N, M)$. [Proof in [10]]*

3. When hypertree width is sandwiched by treewidth

Influence and valency of i -brambles. Let (G, N, M) be an i -labelled graph and \mathcal{B} an i -bramble of it. We define the *influence* of \mathcal{B} , as $\mathbf{ifl}(\mathcal{B}) = \max_{v \in \cup \mathcal{B}} |\{x \in \cup \mathcal{B} \mid \mathbf{dist}_G(v, x) \leq 2\}|$. We also define the *valency* of \mathcal{B} as the quantity $\mathbf{val}(\mathcal{B}) = \max_{v \in \cup \mathcal{B}} |\{S \in \mathcal{B} \mid v \in S\}|$.

Lemma 3.1. *If \mathcal{B} is an i -bramble of an i -labeled graph (G, N, M) , then the order of \mathcal{B} is at least $\frac{|\mathcal{B}|}{\mathbf{ifl}(\mathcal{B}) \cdot \mathbf{val}(\mathcal{B})}$. [Proof in [10]]*

Triangulated grids. A *partially triangulated $(k \times k)$ -grid* is a graph G that is obtained from a $(k \times k)$ -grid (we refer to it as its *underlying grid*) after adding some edges without harming the planarity of the resulting graph. Each vertex of G will be denoted by a pair (i, j) corresponding to its coordinates in the underlying grid. We will also denote as $U(G)$ the vertices, we call them *non-marginal*, of G that in the underlying grid have degree 4 and we call the vertices in $V(G) - U(G)$ *marginal*.

Lemma 3.2. *Let (G, N, M) be an i -labeled graph, where G is a partially triangulated $(k \times k)$ -grid for $k \geq 4$. Then $\mathbf{fibn}(G, N, M) \geq k/50 - c$, for some constant $c \geq 0$.*

Proof. We use notation $C_{i,j}$ for the set vertices of $N \cap U(G)$ that belong to the i -th row or the j -th column of the underlying grid of G . We claim that $\mathcal{B} = \{C_{i,j} \mid 2 \leq i, j \leq k - 1\}$ is an i -bramble of G of order $\geq k/50 - c$, for some constant $c \geq 0$. Since $k \geq 4$, we have that each set $C_{i,j}$ is non-empty and i -connected. Notice also that the intersection of the i -th row and the j' -th column of the underlying grid of G is either a vertex in N and $C_{i,j} \cap C_{i',j'} \neq \emptyset$, or a vertex in $M - N$, but then all neighbors of it in G belong to N . Therefore, all $C_{i,j}$ and $C_{i',j'}$ should i -touch, and \mathcal{B} is an i -bramble. Each vertex $v = (i, j)$ in $N(\cup \mathcal{B})$ is contained in exactly $2k - 5$ sets of \mathcal{B} (that is $k - 2$ sets $C_{i',j'}$ that agree on the first coordinate plus $k - 2$ sets $C_{i',j'}$ that agree on the second, minus one set $C_{i,j}$ that agrees on both), therefore $\mathbf{val}(\mathcal{B}) = 2k - 5$. For each non-marginal vertex x in G , there are at most 25 non-marginal vertices within distance ≤ 2 in G (in the worst case, consider a triangulated (5×5) -grid subgraph of G that is centered at x) and thus $\mathbf{ifl}(\mathcal{B}) \leq 25$. As $|\mathcal{B}| = (k - 2)^2$, Lemma 3.1 implies that there is a constant c such that the order of \mathcal{B} is at least $k/50 - c$ and the lemma follows. ■

Theorem 3.3. *If \mathcal{H} is a hypergraph with a planar incidence graph $I(\mathcal{H})$, then $\mathbf{fhw}(\mathcal{H}) - 1 \leq \mathbf{ghw}(\mathcal{H}) - 1 \leq \mathbf{tw}(I(\mathcal{H})) \leq 300 \cdot \mathbf{fhw}(\mathcal{H}) + c$ for some constant $c \geq 0$.*

Proof. The left hand inequality follows directly from Lemma 2.1. Suppose now that \mathcal{H} is a hypergraph where $\mathbf{fhw}(\mathcal{H}) \leq k$. By Lemmata 2.2 and 2.4, $\mathbf{fbn}(I(\mathcal{H}), V(\mathcal{H}), E(\mathcal{H})) = \mathbf{fbn}(\mathcal{H}) \leq \mathbf{fhw}(\mathcal{H}) \leq k$. By Lemmata 2.5 and 3.2, $(I(\mathcal{H}), V(\mathcal{H}), E(\mathcal{H}))$ cannot be i -contracted to an i -labeled graph (G, N, M) where G is a partially triangulated $(l \times l)$ -grid, where $l = 50 \cdot k + O(1)$. By Lemma 2.3, $\mathcal{I}(\mathcal{H})$ cannot be contracted to a partially triangulated $(l \times l)$ -grid and thus $I(\mathcal{H})$ excludes an $(l \times l)$ -grid as a minor. From [21, (6.2)], $\mathbf{tw}(I(\mathcal{H})) \leq 6 \cdot l \leq 300 \cdot k + c$ and the result follows. ■

Brambles in Gridoids. We call a graph G by a (k, g) -gridoid if it is possible to obtain a partially triangulated $(k \times k)$ -grid after removing at most g edges from it (we call these edges *additional*).

Lemma 3.4. *Let (G, N, M) be an i -labeled graph where G is a (k, g) -gridoid. Then $\mathbf{fbn}(G, N, M) \geq k/50 - c \cdot g$ for some constant $c \geq 0$. [Proof in [10]]*

The proof of the next theorem is similar to the one of Theorem 3.3 (use Lemma 3.4 instead of Lemma 3.2 and [6, Theorem 4.12] instead of [21, (6.2)]).

Theorem 3.5. *If \mathcal{H} is a hypergraph with an incidence graph $I(\mathcal{H})$ of Euler genus at most g , then $\mathbf{fhw}(\mathcal{H}) - 1 \leq \mathbf{ghw}(\mathcal{H}) - 1 \leq \mathbf{tw}(I(\mathcal{H})) \leq 300 \cdot g \cdot \mathbf{fhw}(\mathcal{H}) + c \cdot g$, for some constant $c \geq 0$.*

Brambles in augmented grids. An *augmented $(r \times r)$ -grid of span s* is an $r \times r$ grid with some extra edges such that each vertex of the resulting graph is attached to at most s non-marginal vertices of the grid.

Lemma 3.6. *If (G, N, M) is an i -labeled graph where G is an augmented $(k \times k)$ -grid with span s , then $\mathbf{fbn}(G, N, M) \geq \frac{k}{2 \cdot s^2} - c$, for some constant $c \geq 0$. [Proof in [10]]*

As it was shown by Demaine et al. [5], every apex-minor-free graph with treewidth at least k can be contracted to a $(f(k) \times f(k))$ -augmented grid of span $O(1)$ (the hidden constants in the “ O ”-notation depend only on the excluded apex). Because, $f(k) = \Omega(k)$ (due to the results of Demaine and Hajiaghayi in [7]), we have the following proposition.

Proposition 3.7. *Let G be an H -apex-minor-free graph of treewidth at least $c_H \cdot k$. Then G contains as a contraction an augmented $(k \times k)$ -grid of span s_H , where constants c_H, s_H depend only on the size of apex graph H that is excluded.*

The proof of the next theorem is similar to the one of Theorem 3.3 (use Lemma 3.6 instead of Lemma 3.2 and Proposition 3.7 instead of [21, (6.2)]).

Theorem 3.8. *If \mathcal{H} is a hypergraph with an incidence graph $I(\mathcal{H})$ that is H -apex-minor-free, then $\mathbf{fhw}(\mathcal{H}) - 1 \leq \mathbf{ghw}(\mathcal{H}) - 1 \leq \mathbf{tw}(I(\mathcal{H})) \leq c_H \cdot \mathbf{fhw}(\mathcal{H})$ for some constant c_H that depends only on H .*

4. Hypergraphs with H -minor-free incidence graphs

The results of Theorem 3.8 cannot be extended to hypergraphs which incidence graph excludes an arbitrary fixed graph H as a minor. For example, for every integer k , it is possible to construct a hypergraph \mathcal{H} with the planar incidence graph such that $\mathbf{tw}(I(\mathcal{H})) \geq k$. By adding to \mathcal{H} an universal hyperedge containing all vertices of \mathcal{H} , we obtain a hypergraph \mathcal{H}' of generalized hypertree width one. Its incidence graph $I(\mathcal{H}')$ does not contain the complete graph K_6 as a minor, however its treewidth is at least k . Despite of that, in this section we prove that if a hypergraph has H -minor-free incidence graph, then its generalized hypertree width and fractional hypertree width can be approximated by the treewidth of a graph that can be constructed from its incidence graph in polynomial time. By making use of this result we show that in this case generalized hypertree width and fractional hypertree width are up to a constant multiplicative factor from each other. Another consequence of the combinatorial result is that there is a constant factor polynomial time approximation algorithm for both parameters on this class of hypergraphs. Our proof is based on the Excluded Minor Theorem by Robertson and Seymour [23].

Graph minor theorem. Before describing the Excluded Minor Theorem we need some definitions.

Definition 4.1 (CLIQUE-SUMS). Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be two disjoint graphs, and $k \geq 0$ an integer. For $i = 1, 2$, let $W_i \subseteq V_i$, form a clique of size h and let G'_i be the graph obtained from G_i by removing a set of edges (possibly empty) from the clique $G_i[W_i]$. Let $F : W_1 \rightarrow W_2$ be a bijection between W_1 and W_2 . We define the h -clique-sum of G_1 and G_2 , denoted by $G_1 \oplus_{h,F} G_2$, or simply $G_1 \oplus G_2$ if there is no confusion, as the graph obtained by taking the union of G'_1 and G'_2 by identifying $w \in W_1$ with $F(w) \in W_2$, and by removing all the multiple edges. The image of the vertices of W_1 and W_2 in $G_1 \oplus G_2$ is called the *join* of the sum.

Note that some edges of G_1 and G_2 are not edges of G , since it is possible that these graphs had edges which were removed by clique-sum operation. Such edges are called *virtual* edges of G . We remark that \oplus is not well defined; different choices of G'_i and the bijection F could give different clique-sums. A sequence of h -clique-sums, not necessarily unique, which result in a graph G , is called a *clique-sum decomposition* of G .

Definition 4.2 (h -nearly embeddable graphs). Let Σ be a surface with boundary cycles C_1, \dots, C_h , i.e. each cycle C_i is the border of a disc in Σ . A graph G is h -nearly embeddable in Σ , if G has a subset X of size at most h , called *apices*, such that there are (possibly empty) subgraphs G_0, \dots, G_h of $G - X$ such that i) $G - X = G_0 \cup \dots \cup G_h$, ii) G_0 is embeddable in Σ , we fix an embedding of G_0 , iii) graphs G_1, \dots, G_h (called *vortices*) are pairwise disjoint, iv) for $1 \leq \dots \leq h$, let $U_i := \{u_{i_1}, \dots, u_{i_{m_i}}\} = V(G_0) \cap V(G_i)$, G_i has a path decomposition (B_{ij}) , $1 \leq j \leq m_i$, of width at most h such that a) for $1 \leq i \leq h$ and for $1 \leq j \leq m_i$ we have $u_j \in B_{ij}$, b) for $1 \leq i \leq h$, we have $V(G_0) \cap C_i = \{u_{i_1}, \dots, u_{i_{m_i}}\}$ and the points $u_{i_1}, \dots, u_{i_{m_i}}$ appear on C_i in this order (either if we walk clockwise or anti-clockwise).

The following proposition is known as the Excluded Minor Theorem [23] and is the cornerstone of Robertson and Seymour's Graph Minors theory.

Theorem 4.3 ([23]). *For every non-planar graph H , there exists an integer h , depending only on the size of H , such that every graph excluding H as a minor can be obtained by*

h-clique-sums from graphs that can be h-nearly embedded in a surface Σ in which H cannot be embedded.

Let us remark that by the result of Demaine et al. [8] such a clique-sum decomposition can be obtained in time $O(n^c)$ for some constant c which depends only from H (see also [4]).

Approximation. Let \mathcal{H} be a hypergraph such that its incidence graph $G = I(\mathcal{H})$ excludes a fixed graph H as a minor. Every graph excluding a planar graph H as a minor has a constant treewidth [21]. Thus if H is planar, by Theorem 3.8, the generalized hypertree width does not exceed some constant. In what follows, we always assume that H is not planar.

By Theorem 4.3, there is an h -clique-sum decomposition of $G = G_1 \oplus G_2 \oplus \dots \oplus G_m$ such that for every $i \in \{1, 2, \dots, m\}$, the summand G_i can be h -nearly embedded in a surface Σ in which H can not be embedded. We assume that this clique-sum decomposition is *minimal*, in the sense that for every virtual edge $\{x, y\} \in E(G_i)$ there is an x, y -path in G with all inner vertices in $V(G) - V(G_i)$ (otherwise it is always possible to remove such edges and modify clique-sum operations correspondingly). Let A_i be the set of apices of G_i . We define $E_i = A_i \cap E(\mathcal{H})$ and $G'_i = G_i - (N_G[E_i] \cup A_i)$. For every virtual edge $\{x, y\}$ of G'_i we perform the following operation: if there is no x, y -path in $G - (N[E_i] \cup A_i)$ with all inner vertices in $G - V(G'_i)$, then $\{x, y\}$ is removed from G'_i . We denote the resulted graph by F_i .

In what remains we show that the maximal value of $\mathbf{tw}(F_i)$, where maximum is taken over all $i \in \{1, 2, \dots, m\}$, is a constant factor approximation of generalized and fractional hypertree widths of \mathcal{H} . The upper bound is given by the following lemma (the proof uses results from [1]).

Lemma 4.4. $\mathbf{ghw}(\mathcal{H}) \leq 3 \cdot \max\{\mathbf{tw}(F_i) : i \in \{1, 2, \dots, m\}\} + 6h + 4.$ [Proof in [10]]

To prove the lower bound we need the following property of the clique-sum decomposition which was observed by Demaine and Hajiaghayi [7] (with the reference on the personal communication by Seymour).

Proposition 4.5. *Let $G = G_1 \oplus G_2 \oplus \dots \oplus G_m$. Then every clique sum in this expression involves at most three vertices other than apices and vertices in vortices of the corresponding summand (i.e. at most three such vertices are identified by the operation).*

We also need a result roughly stating that if a graph G with a big grid as a surface minor is embedded on a surface Σ of small genus, then there is a disc in Σ containing a big part of the grid of G . This result is implicit in the work of Robertson and Seymour and there are simpler alternative proofs by Mohar and Thomassen [20, 24] (see also [6, Lemma 3.3]). We use the following variant of this result from Geelen et al. [11].

Proposition 4.6 ([11]). *Let g, l, r be positive integers such that $r \geq g(l + 1)$ and let G be an (r, r) -grid. If G is embedded in a surface Σ of Euler genus at most $g^2 - 1$, then some (l, l) -subgrid of G is embedded in a closed disc Δ in Σ such that the boundary cycle of the (l, l) -grid is the boundary of the disc.*

Now we are ready to prove the following lower bound.

Lemma 4.7. $\mathbf{fbn}(\mathcal{H}) \geq \varepsilon_H \cdot \max\{\mathbf{tw}(F_i) : i \in \{1, 2, \dots, m\}\}$ for some constant ε_H depending only on H . [Proof in [10]]

Proof. Let $i \in \{1, 2, \dots, m\}$. We assume that $G - (N[E_i] \cup A_i)$ is a connected graph which has at least one edge. (Otherwise one can consider the components of this graph separately and remove isolated vertices.) The main idea of the proof is to contract it to a planar graph with approximately the same treewidth as F_i and then apply same techniques that were used in the previous section for the planar case.

Structure of $G - (N[E_i] \cup A_i)$. Let us note that an h -clique-sum decomposition $G = G_1 \oplus G_2 \oplus \dots \oplus G_m$ induces an h -clique-sum decomposition of $G' = G - (N[E_i] \cup A_i)$ with the summand G_i replaced by F_i . Let G'_1, G'_2, \dots, G'_l be the connected components of $G' - V(F_i)$. Every such component G'_j is attached via clique-sum to F_i by some clique Q_j of F_i . Note that cliques Q_j contain all virtual edges of F_i . We assume that each clique Q_j does not separate vertices of F_i . Otherwise, it is possible to decompose F_i into the clique-sum of graphs $F_i^{(1)} \oplus F_i^{(2)}$ with the join Q_j and prove the bound for summands and, since $\mathbf{tw}(F_i) = \max\{F_i^{(1)}, F_i^{(2)}\}$, that will prove the lemma. To simplify the structure of the graph, for every component G'_j , we contract all its edges and denote by S_j the star whose central vertex is the result of the contraction and leaves are the vertices of Q_j .

Contracting vortices. The h -nearly embedding of the graph G_i induces the h -nearly embedding of $F_i = X_0 \cup X_1 \cup \dots \cup X_h$ without apices. Here we assume that X_0 is embedded in a surface Σ of genus depending on H and X_1, X_2, \dots, X_h are the vortices. For every vortex X_j , the vertices $V(X_0) \cap V(X_j)$ are on the boundary C_j of some face of X_0 . If for a star S_k some of its leaves Q_k are in X_j or C_j , we do the following operation: if $Q_k \cap (V(X_j) - V(C_j)) \neq \emptyset$ then all edges of S_k are contracted, and if $Q_k \cap (V(X_j) - V(C_j)) = \emptyset$ but $|Q_k \cap V(C_j)| \geq 2$, then we contract all edges of S_k that are incident to the vertices of $Q_k \cap V(C_j)$. These contractions results in the contraction of some edges of F_i . Particularly, all virtual edges of X_j and C_j are contracted. Additionally, we contract all remaining edges of X_j and C_j . We perform theses contractions for all vortices of F_i and denote the result by F'_i . It follows immediately from the definition of the h -clique-sum and Proposition 4.5, that F'_i coincides with the graph obtained from F_i by contractions of all vortices X_j and boundaries of faces C_j . It can be easily seen that F'_i is embedded in Σ . It is known (see e.g. [6, 7]) that there is a positive constant a_H which depends only on H such that $\mathbf{tw}(F'_i) \geq a_H \cdot \mathbf{tw}(F_i)$.

Contracting the part that lies outside of some planar disc. Since F'_i is embedded in Σ , we have that the graph F'_i contains some $(k \times k)$ -grid as a surface minor, where $k \geq b_H \cdot \mathbf{tw}(F'_i)$ for some constant b_H [6]. Combining this result with Proposition 4.6, we receive the following claim. There is a disc $\Delta \subseteq \Sigma$ such that i) the subgraph R of F'_i induced by vertices of $F'_i \cap \Delta$ is a connected graph; ii) the subgraph R' of F'_i induced by $N_{F'_i}[V(R)]$ is completely in some disc Δ' ; iii) vertices of $V(R') - V(R)$ induce a cycle C which is the border of Δ' , and iv) $\mathbf{tw}(R) \geq c_H \cdot \mathbf{tw}(F'_i)$ for some constant c_H . Now we treat the part of F'_i which is outside Δ exactly the same way we have treated vortices. For stars S_k intersecting $V(F'_i) - V(R')$ or C , we do the following: if $Q_k \cap (V(F'_i) - V(R')) \neq \emptyset$, then all edges of S_k are contracted, and if $Q_k \cap (V(F'_i) - V(R')) = \emptyset$ but $|Q_k \cap V(C)| \geq 2$, then all edges of S_k incident to the vertices of $Q_k \cap V(C)$ are contracted. These contractions result in the contraction of some edges of F'_i with endpoints on C or outside Δ' . Particularly, all such virtual edges are contracted. Additionally, we contract all remaining edges of $F'_i - V(R)$ and C . Thus this part of the graph is contracted to a single vertex. Denote the obtained graph X . This graph is planar, and since R is a subgraph of X , we have that $\mathbf{tw}(X) \geq \mathbf{tw}(R)$.

Embedding the stars. Some edges of X are virtual, and all such edges are in cliques Q_j . By Proposition 4.5, $|Q_j| \leq 3$. For every clique $Q = V(X) \cap Q_j$, we do the following. If $Q = \{x, y\}$, then the edge of the star S_j incident to x is contracted. If $Q = \{x, y, z\}$, then if two vertices of Q , say x and y , are joined by an edge in G , then the edge of S_j incident to z is contracted, and if there are no such edges and the triangle induced by $\{x, y, z\}$ is the boundary of some face of X , then we add a new vertex on this face, join it with x, y and z (it can be seen as S_j embedded in this face, and since our graph is i -labeled, it is assumed that this new vertex has same labels as the central vertex of S_j), and then remove virtual edges. Note that if the triangle is not a boundary of some face, then Q is a separator of our graph, but we assumed that there are no such separators. Denote by Y the obtained graph. Similar construction was used in the proof of the main theorem in [7], and by the same arguments as were used by Demaine et al. we immediately conclude that there is a positive constant d_H such that $\mathbf{tw}(X) \geq d_H \cdot \mathbf{tw}(Y)$.

Now all contractions are finished. Note that the graph Y is a planar graph which is a contraction of $G' = G - (N[E_i] \cup A_i)$. Also there is some positive constant e_H which depends only on H such that $\mathbf{tw}(Y) \geq e_H \cdot \mathbf{tw}(F_i)$. Recall that we consider the i -labeled graph $(G, V(\mathcal{H}), E(\mathcal{H}))$. By Lemma 2.4, $\mathbf{fbn}(\mathcal{H}) = \mathbf{fbn}(G, V(\mathcal{H}), E(\mathcal{H}))$. Because the sets $V(\mathcal{H})$ and $E(\mathcal{H})$ are independent, by Lemma 2.6, we have that $\mathbf{fbn}(G, V(\mathcal{H}), E(\mathcal{H})) \geq \mathbf{fbn}(G', N, M)$, where $N = V(\mathcal{H}) - (N[E_i] \cup A_i)$ and $M = E(\mathcal{H}) - (N[E_i] \cup A_i)$. By Lemma 2.5, $\mathbf{fbn}(G', N, M) \geq \mathbf{fbn}(Y, N', M')$, where N' and M' are sets which were obtained as the result of contractions of N and M . Finally, as in Theorem 3.3, one can show that $\mathbf{fbn}(Y, N', M') \geq f_H \cdot \mathbf{tw}(Y)$ for some constant f_H . By putting all these bounds together, we prove that there is a positive constant ε_H which depends only on H , such that $\mathbf{fbn}(\mathcal{H}) \geq \varepsilon_H \cdot \mathbf{tw}(F_i)$. ■

Combining Lemmata 2.1, 2.2, 4.4, and 4.7, we obtain the following theorem.

Theorem 4.8. $(1/c_H) \cdot w \leq \mathbf{fhw}(\mathcal{H}) \leq \mathbf{ghw}(\mathcal{H}) \leq c_H \cdot w$, where $w = \max\{\mathbf{tw}(F_i) : i \in \{1, 2, \dots, m\}\}$, and c_H is a constant depending only on H .

Remark. Notice that, by Theorem 4.8, the generalized hypertree width and the fractional hypertree width of a hypergraph with H -minor-free incidence graph may differ within a multiplicative constant factor. We stress that, as observed in [18], this is not the case for general hypergraphs.

Demaine et al. [8] (see also [4, 9, 23]) described an algorithm which constructs a clique-sum decomposition of an H -minor-free graph G on n vertices with the running time $n^{O(1)}$ (the hidden constant in the running time depends only on H). As far as we constructed summands G_i , the construction of graphs F_i can be done in polynomial time. Moreover, since the algorithm of Demaine et al. provides h -nearly embeddings of these graphs, it is possible to use it to construct a polynomial constant factor approximation algorithm for the computation of $\mathbf{tw}(F_i)$. This provides us with the main algorithmic result of this section.

Theorem 4.9. *For any fixed graph H , there is a polynomial time c_H -approximation algorithm computing the generalized hypertree width and the fractional hypertree width for hypergraphs with H -minor-free incidence graphs, where the constant c_H depends only on H .*

We finally remark that by making use of the results from [16], our results can be used not only to compute but to construct, up to constant multiplicative-factor, the corresponding decompositions.

References

- [1] I. ADLER, G. GOTTLÖB, AND M. GROHE, *Hypertree width and related hypergraph invariants*, European J. Combin., 28 (2007), pp. 2167–2181.
- [2] A. K. CHANDRA AND P. M. MERLIN, *Optimal implementation of conjunctive queries in relational data bases*, in STOC'77, ACM, 1977, pp. 77–90.
- [3] C. CHEKURI AND A. RAJARAMAN, *Conjunctive query containment revisited*, Theoret. Comput. Sci., 239 (2000), pp. 211–229.
- [4] A. DAWAR, M. GROHE, AND S. KREUTZER, *Locally excluding a minor*, in LICS'07, IEEE Computer Society, 2007, pp. 270–279.
- [5] E. D. DEMAINE, F. V. FOMIN, M. HAJIAGHAYI, AND D. M. THILIKOS, *Bidimensional parameters and local treewidth*, SIAM J. Discrete Math., 18 (2004/05), pp. 501–511.
- [6] ———, *Subexponential parameterized algorithms on graphs of bounded genus and H -minor-free graphs*, J. ACM, 52 (2005), pp. 866–893.
- [7] E. D. DEMAINE AND M. HAJIAGHAYI, *Graphs excluding a fixed minor have grids as large as treewidth, with combinatorial and algorithmic applications through bidimensionality*, in SODA'05, ACM, 2005, pp. 682–689.
- [8] E. D. DEMAINE, M. T. HAJIAGHAYI, AND K. ICHI KAWARABAYASHI, *Algorithmic graph minor theory: Decomposition, approximation, and coloring*, in FOCS'05, IEEE Computer Society, 2005, pp. 637–646.
- [9] U. FEIGE, M. HAJIAGHAYI, AND J. R. LEE, *Improved approximation algorithms for minimum weight vertex separators*, SIAM J. Computing, 38 (2008), pp. 629–657.
- [10] F. V. FOMIN, P. A. GOLOVACH, AND D. M. THILIKOS, *Approximating acyclicity parameters of sparse hypergraphs*, CoRR, abs/0809.3646 (2008).
- [11] J. F. GEELLEN, R. B. RICHTER, AND G. SALAZAR, *Embedding grids in surfaces*, European J. Combin., 25 (2004), pp. 785–792.
- [12] G. GOTTLÖB, M. GROHE, N. MUSLIU, M. SAMER, AND F. SCARCELLO, *Hypertree decompositions: Structure, algorithms, and applications*, in WG'05, vol. 3787 of Lecture Notes in Computer Science, Springer, 2005, pp. 1–15.
- [13] G. GOTTLÖB, N. LEONE, AND F. SCARCELLO, *A comparison of structural CSP decomposition methods*, Artificial Intelligence, 124 (2000), pp. 243–282.
- [14] ———, *The complexity of acyclic conjunctive queries*, J. ACM, 48 (2001), pp. 431–498.
- [15] ———, *Hypertree decompositions and tractable queries*, J. Comput. System Sci., 64 (2002), pp. 579–627.
- [16] ———, *Robbers, marshals, and guards: game theoretic and logical characterizations of hypertree width*, J. Comput. System Sci., 66 (2003), pp. 775–808.
- [17] G. GOTTLÖB, Z. MIKLÓS, AND T. SCHWENTICK, *Generalized hypertree decompositions: NP-hardness and tractable variants*, in PODS'07, ACM, 2007, pp. 13–22.
- [18] M. GROHE AND D. MARX, *Constraint solving via fractional edge covers*, in SODA'06, ACM, 2006, pp. 289–298.
- [19] P. G. KOLAITIS AND M. Y. VARDI, *Conjunctive-query containment and constraint satisfaction*, J. Comput. System Sci., 61 (2000), pp. 302–332.
- [20] B. MOHAR, *Combinatorial local planarity and the width of graph embeddings*, Canad. J. Math., 44 (1992), pp. 1272–1288.
- [21] N. ROBERTSON, P. SEYMOUR, AND R. THOMAS, *Quickly excluding a planar graph*, J. Combin. Theory Ser. B, 62 (1994), pp. 323–348.
- [22] N. ROBERTSON AND P. D. SEYMOUR, *Graph minors. II. Algorithmic aspects of tree-width*, J. Algorithms, 7 (1986), pp. 309–322.
- [23] ———, *Graph minors. XVI. Excluding a non-planar graph*, J. Combin. Theory Ser. B, 89 (2003), pp. 43–76.
- [24] C. THOMASSEN, *A simpler proof of the excluded minor theorem for higher surfaces*, J. Combin. Theory Ser. B, 70 (1997), pp. 306–311.
- [25] M. YANNAKAKIS, *Algorithms for acyclic database schemes*, in VLDB'81, IEEE Computer Society, 1981, pp. 82–94.