SIAM J. DISCRETE MATH.
Vol. 25, No. 3, pp. 1331–1348

APPROXIMATING WIDTH PARAMETERS OF HYPERGRAPHS WITH EXCLUDED MINORS*

FEDOR V. FOMIN†, PETR A. GOLOVACH‡, AND DIMITRIOS M. THILIKOS§

Abstract. The notions of hypertree width and generalized hypertree width were introduced by Gottlob, Leone, and Scarcello in order to extend the concept of hypergraph acyclicity. These notions were further generalized by Grohe and Marx, who introduced the fractional hypertree width of a hypergraph. All these width parameters on hypergraphs are useful for extending the tractability of many problems in database theory and artificial intelligence. 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 $\mathcal{H}$ 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 in 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.

Key words. hypergraphs, minor-free graphs, hypertree width, approximation

AMS subject classifications. 05C65, 05C83, 05C85, 68R10

DOI. 10.1137/080743226

1. Introduction.

Many important theoretical and "real-world" problems can be expressed as constrained satisfaction problems (CSPs). Among examples, one can mention numerous problems from 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 was shown by Kolaitis and Vardi [21], CSPs, conjunctive-query containment, and finding homomorphism for relational structures are essentially the same problem. The problem is known to be NP-hard in general [3] and polynomial time solvable for a restricted class of acyclic queries [29]. 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 [4] introduced the notion of query width. Gottlob,

*Received by the editors December 10, 2008; accepted for publication (in revised form) June 16, 2011; published electronically September 13, 2011. A preliminary version of this paper appeared in *Proceedings of the 26th Symposium on Theoretical Aspects of Computer Science*, Freiburg, Germany, 2009.

http://www.siam.org/journals/sidma/25-3/74322.html

†Department of Informatics, University of Bergen, N-5020 Bergen, Norway (fomin@ii.uib.no). This author's research was supported by the Norwegian Research Council.

‡School of Engineering and Computing Sciences, Durham University, DH1 3LE Durham, UK (petr.golovach@durham.ac.uk).

§Department of Mathematics, University of Athens, Panepistimioupolis, GR15784 Athens, Greece (sedthilk@math.uoa.gr). This author's research was supported by the Kapodistrias project (ΑΠ 02839/28.07.2008) of the National and Kapodistrian University of Athens.

Leone, and Scarcello [15], [16], [18] defined hypertree width and generalized hypertree width. Furthermore, Grohe and Marx [20] introduced the most general parameter known so far, fractional hypertree width, and proved that CSPs, restricted to instances of bounded fractional hypertree width, are polynomial time solvable.

Unfortunately, all known variants of hypertree width are NP-complete [14], [19], [22]. Moreover, generalized hypertree width is NP-complete even when checking whether its value is at most 3 (see [19]). In the case of hypertree width, the problem is $W[2]$-hard when parameterized by k [14]. Both the hypertree width and generalized hypertree width are hard to approximate. For example, the reduction of Gottlob et al. in [14] 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 $\mathsf{P} = \mathsf{NP}$.

All these parameters for hypergraphs can be seen as generalizations of the treewidth of a graph. The treewidth is a fundamental graph parameter from the graph minors theory by Robertson and Seymour [26], and it has numerous algorithmic applications (for a survey, see [2]). 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 the $\sqrt{\log OPT}$-approximation of Feige, Hajiaghayi, and Lee [10]. However, as was shown by Feige, Hajiaghayi, and Lee [10], 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 the treewidth of the incidence graph. This indicates that hypertree width parameters are more useful as the adequate version of acyclicity for nonsparse 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 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 the treewidth of its incidence graph can be arbitrarily large. In particular, we give a constant factor approximation algorithm for the generalized and the fractional hypertree width of hypergraphs with H-minor-free incidence graphs, extending the approximation results of Feige, Hajiaghayi, and Lee [10] from treewidth to (generalized, fractional) hypertree width. The algorithm uses a series of theorems based on the main decomposition theorem of the Robertson–Seymour graph minor project. As a combinatorial corollary of our results, it follows that generalized hypertree width and fractional hypertree width differ within a constant multiplicative factor if the incidence graph of the hypergraph does not contain a fixed graph as a minor.

2. Definitions and preliminaries.

2.1. Basic definitions. 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, respectively, 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, respectively).

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 notation, 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 an open (respectively, closed) disc if it is homeomorphic to $\{(x, y) : x^2 + y^2 < 1\}$ (respectively, $\{(x, y) : x^2 + y^2 \leq 1\}$). 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. We refer to the book of Mohar and Thomassen [24] for more details on graph embeddings.

Given an edge $e = \{x, y\}$ of a graph G, the graph G/e is obtained from G by contracting e; that is, to get G/e we identify the vertices x and y and remove all loops and replace all multiple edges by simple edges. A graph H obtained by a sequence of edge-contractions is said to be a *contraction* of G. If H is a contraction of G, then for a vertex $v \in V(H)$, the set of vertices of G which are contracted to v is called the *model* of v. A graph H is a *minor* of G if H is a subgraph of a contraction of G. Let G be a graph embedded in some surface Σ, and let H be a contraction of G. We say that H is a *surface contraction* of G if for each vertex $v \in V(H)$, the model of v is embedded in some open disk in Σ. It can be easily noted that if H is a surface contraction of a graph G embedded in Σ, then it can be assumed that H is embedded in a surface Σ' homeomorphic to Σ. For simplicity, we always assume in such cases that Σ' and Σ are the same surface. We say that H is a *surface minor* of a graph G if H is the surface contraction of some subgraph of G. Observe that H is a graph embedded in Σ.

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 length $k - 1$.

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 $v \in e$ (see Figure 2.1).

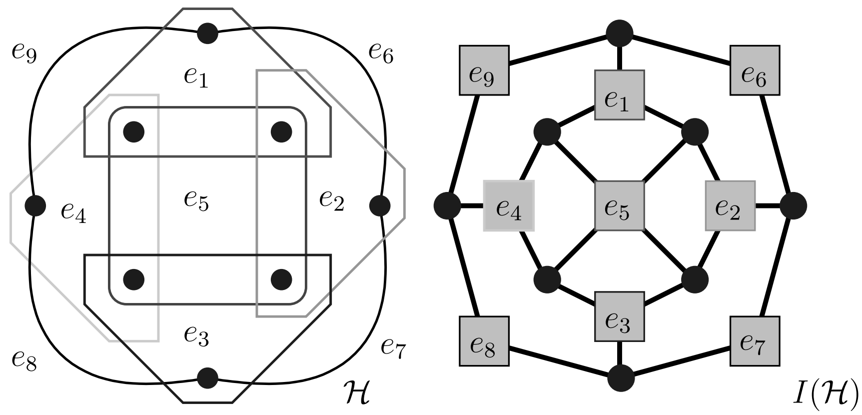

FIG. 2.1. *A hypergraph $\mathcal{H}$ and its incidence graph $I(\mathcal{H})$.*

2.2. Treewidth of graphs and hypergraphs. A *tree decomposition* of a hypergraph $\mathcal{H}$ is a pair (T, χ), where T is a tree and $\chi : V(T) \to 2^{V(\mathcal{H})}$ 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 $v \in V(G)$, 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 the notation $\mathbf{tw}(\mathcal{H})$ for the treewidth of a hypergraph $\mathcal{H}$. If, in the above definitions, we restrict the tree T to be a path, then we define the notions of *path decomposition* and *pathwidth*.

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 may 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 $\mathbf{tw}(\mathcal{H}) \geq |e| - 1$ 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 the even more general parameter of fractional hypertree width.

2.3. Hypertree width, its generalizations, and related notions. In general, given a finite set A, we use the term *labeling of* A for any function $\gamma : A \to [0, 1]$. We also use the notation $\mathcal{G}(A)$ for the collection of all labelings of a set $\mathcal{A}$.

The *size* of a labeling of A is defined as

$$|\gamma| = \sum_{x \in A} \gamma(x).$$

If the values of a labeling γ are restricted to be 0 or 1, then we say that γ is a *binary* labeling of A. Clearly, the size of a binary labeling is equal to the number of the elements of A that are labeled by 1. Given a hyperedge labeling γ of a hypergraph $\mathcal{H}$, we define the set of vertices of $\mathcal{H}$ that are *blocked* by γ as

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

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

A *fractional hypertree decomposition* [20] of $\mathcal{H}$ is a triple (T, χ, λ), where (T, χ) is a tree decomposition of $\mathcal{H}$ and $\lambda: V(T) \to \mathcal{G}(E(\mathcal{H}))$ 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 labeling $\lambda(t)$. The *width* of a fractional hypertree decomposition (T, χ, λ) is $\max\{|\lambda(t)| : t \in V(T)\}$, 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* [17]. Correspondingly, the *generalized hypertree width* $\mathbf{ghw}(\mathcal{H})$ of $\mathcal{H}$ is the minimum of the widths of all generalized hypertree decompositions of $\mathcal{H}$. There are several width parameters defined on hypergraphs, including hypertree width, which, up to a multiplicative constant factor, are the same as the generalized hypertree width [1].

Clearly, $\mathbf{fhw}(\mathcal{H}) \leq \mathbf{ghw}(\mathcal{H})$, but, as was shown in [20], 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 $\mathbf{tw}(G) + 1 = \mathbf{fhw}(H)$ (see also [22]).

The proof of the next lemma follows from the results of [4] about query width. For completeness, we provide a direct proof here.

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. Let (T, χ) be a tree decomposition of $I(\mathcal{H})$ of width $\leq k$. It is enough to describe a generalized hypertree decomposition (T, χ', λ) for $\mathcal{H}$ that has width at most $k+1$. For every $t \in V(T)$, let $\chi'(t) = (\chi(t) - E(\mathcal{H})) \cup (\bigcup(\chi(t) \cap E(\mathcal{H})))$. We include in $\lambda(t)$ all hyperedges $\chi(t) \cap E(\mathcal{H})$, and for every $v \in \chi(t) \cap V(\mathcal{H})$, a hyperedge e such that $v \in e$ is chosen arbitrarily and included in $\lambda(t)$. Clearly,

$$V(\mathcal{H}) = \bigcup_{t \in V(T)} \chi'(t)$$

for each $e \in E(\mathcal{H})$ there is a node $t \in V(T)$ such that $e \subseteq \chi'(t)$, and for every $t \in V(T)$, $\chi'(t) \subseteq \bigcup \lambda(t)$. We have to prove that for each $v \in V(\mathcal{H})$, the set $\{t \in V(T) : v \in \chi'(t)\}$ forms a subtree of T. Suppose that there are $s, t \in V(T)$ at a distance of at least two, $v \in \chi'(s) \cap \chi'(t)$, and $v \notin \chi'(x)$ for all inner vertices x of s, t-path in T. Since (T, χ) is a tree decomposition of $I(\mathcal{H})$, $v \in \chi'(t) - \chi(t)$ or $v \in \chi'(s) - \chi(s)$. Assume that $v \in \chi'(t) - \chi(t)$. It means that there is $e \in \chi(t)$ such that $v \in e$. Note that $e \notin \chi(x)$ for inner vertices x of s, t-path and $e \notin \chi(s)$ (otherwise $v \in \chi'(x)$ by the definition). If $v \in \chi(s)$, then there is no bag in (T, χ) that contains both endpoints of the edge $\{v, e\} \in E(I(\mathcal{H}))$. So $v \in \chi'(s) - \chi(s)$, and there is $e' \in \chi(s)$ such that $v \in e'$. As before, $e' \notin \chi(x)$ for inner vertices and $e' \notin \chi(t)$ (hence $e \neq e'$). But since v is adjacent to e and e' in $I(\mathcal{H})$, bags $\chi(x)$ should contain v, and we receive a contradiction. ∎

Let us remark here that the fractional hypertree width of a hypergraph can be arbitrarily smaller than the treewidth of its incidence graph. Adding a hyperedge to a hypergraph never decreases the treewidth of its incidence graph but can decrease dramatically the fractional and generalized hypertree widths. For example, 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 $\mathbf{tw}(I(\mathcal{H}')) + 1 \geq \mathbf{tw}(I(\mathcal{H})) + 1 \geq \mathbf{fhw}(\mathcal{H})$.

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 $V(\mathcal{H})$. We say that a labeling 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* $\mathbf{fbn}(\mathcal{H})$ of $\mathcal{H}$ is the maximum of the fractional orders of all hyperbrambles of $\mathcal{H}$.

Many graph and hypergraph width parameters can be expressed in terms of graph searching [12]. We need the following game-theoretical interpretation of fractional hypertree width. The *robber and army game* was introduced by Grohe and Marx in [20]. The game is played on a hypergraph $\mathcal{H}$ by two players: a robber and a general who commands the army. A position of the game is a pair (γ, v), where γ is a labeling of $E(\mathcal{H})$ and $v \in V(\mathcal{H})$. The choice of γ is a distribution of the army on the hyperedges of $\mathcal{H}$, chosen by the general, while v is the position of the robber. During the game, a vertex of the hypergraph is blocked only if the total amount of army on the hyperedges that contain this vertex adds up to the strength of at least one battalion. To start the game, the robber picks a position v_0, and the initial position is $(\mathcal{O}, v_0)$, where $\mathcal{O}$ denote the constant zero mapping. In each round, the players move from the current position (γ, v) to a new position (γ', v') as follows: The general selects γ', and then the robber selects v' such that there is a path from v to v' in the hypergraph $\mathcal{H}$ that avoids the vertices in $B(\gamma) \cap B(\gamma')$. Under these circumstances, the positions (γ, v) and (γ', v') are called *compatible*. A *game sequence* is a sequence of compatible positions, and its cost is the maximum size of a distribution γ in it. If, at some moment, the position of the game is (γ, v) where $v \in B(\gamma)$, then the general wins. If this never happens, then the robber wins. A *winning strategy of cost at most k* for the general is a program that provides a response on each possible position such that any game sequence generated by this program is finite and has cost at most k. The *army width* $\mathbf{aw}(\mathcal{H})$ of $\mathcal{H}$ is the least k for which there exists a winning strategy of cost at most k.

Using the fact that $\mathbf{aw}(\mathcal{H}) \leq \mathbf{fhw}(\mathcal{H})$ (see [20, Theorem 11]), we can prove the following lemma.

LEMMA 2.2. *For any hypergraph $\mathcal{H}$, $\mathbf{fbn}(\mathcal{H}) \leq \mathbf{aw}(\mathcal{H}) \leq \mathbf{fhw}(\mathcal{H})$.*

Proof. Let $\mathcal{B}$ be a hyperbramble of $\mathcal{H}$ of fractional order at least k. Our aim is to provide an escape strategy for the robber against any possible winning strategy of cost at most $<k$. In particular, the robber will always be on a vertex of some set $S \in \mathcal{B}$ such that S is not covered by γ and at any position (γ, v) of the game there will be a new unblocked vertex for the robber to move. Indeed, if the response of the general at position (γ, v) is γ', we have that $|\gamma| < k$, and therefore γ cannot cover all elements of $\mathcal{B}$. If $S' \in \mathcal{B}$ is such a set, the new position of the robber will be any vertex v' of S'. Clearly, the robber can move from v to v', as S and S' touch and all of their vertices are unblocked. This implies that $\mathbf{fbn}(\mathcal{H}) \leq \mathbf{aw}(\mathcal{H})$, and the result follows from the fact that $\mathbf{aw}(\mathcal{H}) \leq \mathbf{fhw}(\mathcal{H})$, which was proved in [20, Theorem 11]. ∎

The variant of the robber and army game where the labelings are restricted to be binary labelings is called the marshals and robbers game and was introduced by Gottlob, Leone, and Scarcello [18]. The corresponding parameter is called the *marshal width* and is denoted by $\mathbf{mw}$. Clearly, for any hypergraph $\mathcal{H}$, $\mathbf{aw}(\mathcal{H}) \leq \mathbf{mw}(G)$.

2.4. i-brambles. For a hypergraph $\mathcal{H}$, the graph $I(\mathcal{H})$ has vertices of two types: the vertices corresponding to the vertices of $\mathcal{H}$ and the vertices corresponding to the hyperedges. To capture this distinction, we introduce some additional notions here.

An *i-labeled graph* G is a triple (G, N, M) such that

(i) $N, M \subseteq V(G)$;
(ii) $N \cup M = V(G)$;
(iii) $M - N$ and $N - M$ are independent sets of G;
(iv) for every $v \in V(G)$, the closed neighborhood $N_G[v]$ is intersecting both N and M.

Let us remark that if G has no isolated vertices, then (iv) follows from (i)–(iii). Another remark is that we do not require in this definition that $N \cap M = \varnothing$. 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 $N = V(\mathcal{H})$ and $M = E(\mathcal{H})$ (see Figure 2.2).

By the definition of the i-labeled graph, for every edge $\{x, y\}$ of an i-labeled graph (G, N, M), $\{x, y\} \cap N \neq \varnothing$ and $\{x, y\} \cap M \neq \varnothing$. 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) $N' = (N - \{x, y\}) \cup \{v_e\}$;
(iii) $M' = (M - \{x, y\}) \cup \{v_e\}$.

Indeed, (G', N', M') is an i-labeled graph because $N' \cup M' = V(G')$; because v_e is in both sets N' and M', we have that $M' - N'$ and $N' - M'$ are independent sets in G' and for each $v \in V(G')$, $N_{G'}[v]$ is intersecting both N' and M' (see Figure 2.2 for an example). 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 definition.

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 $G[S \cup (M - N)]$. We say that two subsets $S, R \subseteq N$ *i-touch* if one of the following holds:

(i) $S \cap R \neq \varnothing$;
(ii) there is an edge $\{x, y\}$ with $x \in S$ and $y \in R$; or
(iii) there is a vertex $z \in M - N$ such that $N_G(z)$ intersects both S and R.

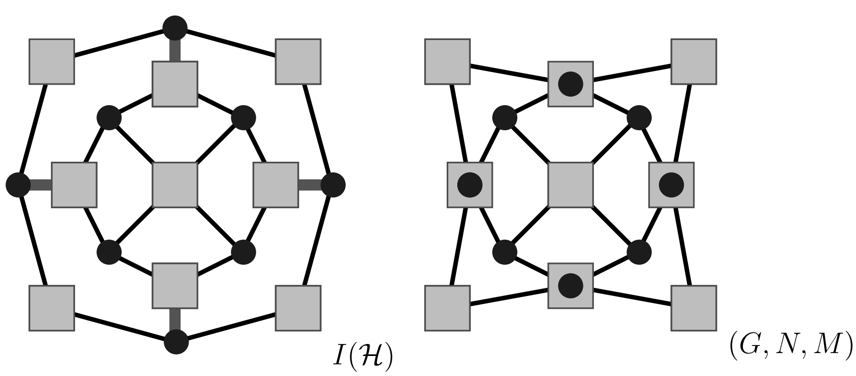

FIG. 2.2. *The i-labeled graph $I(\mathcal{H})$ and a contraction (G, N, M) of it.*

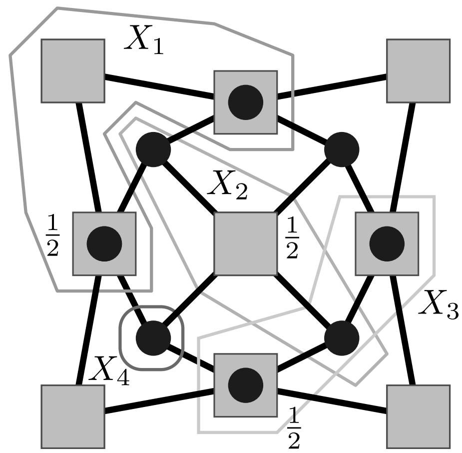

FIG. 2.3. *A bramble* $\mathcal{B} = \{X_1 \cap N, X_2 \cap N, X_3 \cap N, X_4 \cap N\}$ *in the i-labeled graph* (G, N, M).

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 (see Figure 2.3). 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 at least one vertex of S 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 $\mathcal{B}$. For example, in Figure 2.3, if γ labels by $1/2$ the three indicated vertices of M and labels by 0 all the others, then γ fractionally covers all sets of $\mathcal{B}$. It is not hard to verify that the order of the bramble in Figure 2.3 is $3/2$.

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

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

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

It can be 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* $\mathbf{fibn}(G', N', M') \leq \mathbf{fibn}(G, N, M)$.

Proof. For each vertex $x \in V(G')$, let $U_x \subseteq V(G)$ be the model of x. Let $\mathcal{B}'$ be an i-bramble in (G', N', M'). For each set B' in $\mathcal{B}'$, we define the set of vertices

$$B = N \cap \left(\bigcup_{x \in B'} U_x \right).$$

Let $\mathcal{B}$ be a collection of all such sets. Then $\mathcal{B}$ is an i-bramble in G. Now let γ be a labeling of M of size at most $\mathbf{fibn}(G, N, M)$ which fractionally covers all sets of $\mathcal{B}$. We define a labeling γ' of M' as follows: For each $x \in M'$,

$$\gamma'(x) = \sum_{v \in M \cap U_x} \gamma(v).$$

Thus $|\gamma'| = |\gamma|$. It remains to notice that all sets of $\mathcal{B}'$ are fractionally covered by γ'. ∎

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. Let $G' = G - X$, $N' = N - X$, and $M' = M - X$. Since G' has no isolated vertices, (G', N', M') is an i-labeled graph. Let $\mathcal{B}$ be an i-bramble of (G', N', M'). Obviously, $\mathcal{B}$ is an i-bramble of (G, N, M), and there is a labeling γ of M of size $k \leq \mathbf{fibn}(G, N, M)$ which fractionally covers all sets of $\mathcal{B}$. Because $N_G[v] \subseteq X$ for every vertex $v \in X \cap M$, we have that the restriction γ' of γ to M is the labeling of M' which covers all sets of $\mathcal{B}$, and $|\gamma'| \leq k$. ∎

3. When hypertree width is sandwiched by treewidth.

Vertex removal of an incidence graph corresponding to an edge removal of the corresponding hypergraph can significantly increase each of the hypertree width parameters. This is the main reason why in this section we have to develop tools based on edge-contractions.

3.1. Influence and valency of i-brambles.

Let (G, N, M) be an i-labeled graph and $\mathcal{B}$ be an i-bramble of it. By slightly abusing the notation, we use $\cup\,\mathcal{B}$ to denote the subset of N which is the union of all vertices contained in elements of $\mathcal{B}$. 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

$$\mathbf{val}(\mathcal{B}) = \max_{v \in \cup \mathcal{B}} |\{S \in \mathcal{B} \mid v \in S\}|.$$

See Figure 3.1 for an example.

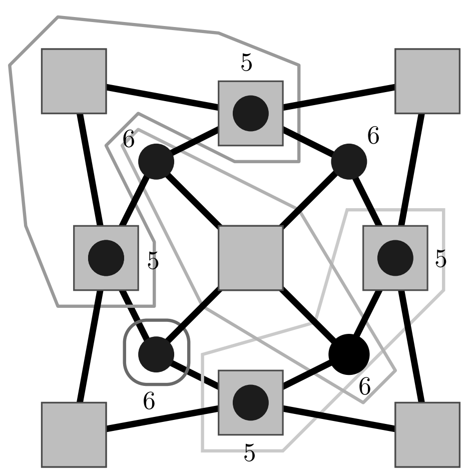

FIG. 3.1. *The values of $|\{x \in \cup\, \mathcal{B} \mid \mathbf{dist}_G(v, x) \leq 2\}|$ for each member for N in the bramble $\mathbf{B}$. It follows that $\mathbf{ifl}(\mathcal{B}) = 6$. The black vertex in N is the unique vertex belonging in two sets of $\mathcal{B}$, and therefore $\mathbf{val}(\mathcal{B}) = 2$.*

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. Let γ be a labeling of M that fractionally covers all sets of $\mathcal{B}$. We first prove the following claim.

Claim. Labeling γ blocks at most $\mathbf{ifl}(\mathcal{B}) \cdot |\gamma|$ vertices in $\cup\,\mathcal{B}$.

Proof. Let R be a subset of $\cup\,\mathcal{B}$ such that every vertex in R is blocked by γ. We define G_R as the graph whose vertex set is R and where two vertices $x, y \in R$ are adjacent if their distance in G is 1 or 2. By the definition of influence, we obtain that the maximum degree of G_R is at most $\mathbf{ifl}(\mathcal{B}) - 1$, and therefore G_R has an independent set I of size at least $|R|/\mathbf{ifl}(\mathcal{B})$. As $I \subseteq R$, all vertices of I are blocked by γ. This implies that for every $x \in I$,

$$\sum_{y\in N_G[x]\cap M} \gamma(x) \geq 1.$$

By definition, for each pair $x, x' \in I$, $x \neq x'$, $N_G[x] \cap N_G[x'] = \varnothing$. Therefore,

$$|\gamma| = \sum_{x\in M}\gamma(x) \geq \sum_{x\in N_G[R]\cap M}\gamma(x) \geq \sum_{x\in N_G[I]\cap M}\gamma(x) \geq \sum_{x\in I}\sum_{y\in N[x]\cap M}\gamma(y) \geq |I| \geq \frac{|R|}{\mathbf{ifl}(\mathcal{B})},$$

and the claim follows. ◻

The above claim, along with the definition of valency, implies that γ fractionally covers no more than $\mathbf{ifl}(\mathcal{B}) \cdot |\gamma| \cdot \mathbf{val}(\mathcal{B})$ sets of $\mathcal{B}$. We conclude that $|\mathcal{B}| \leq \mathbf{ifl}(\mathcal{B}) \cdot |\gamma| \cdot \mathbf{val}(\mathcal{B})$, and the lemma follows. ◻

3.2. 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 destroying the planarity of the resulting graph. Each vertex of G is denoted by a pair (i, j) corresponding to its coordinates in the underlying grid. The *nonmarginal* vertices of the partially triangulated grid G, denoted by $U(G)$, are the vertices that have degree 4 in the underlying grid. We refer to the remaining vertices $V(G) - U(G)$ as *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 the notation $C_{i,j}$ for the set vertices of $N \cap U(G)$ that belong to the ith row or the jth 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 nonempty and i-connected. Notice also that the intersection of the ith row and the j'th column of the underlying grid of G is either a vertex in N and $C_{i,j} \cap C_{i',j'} \neq \varnothing$, or a vertex in $M - N$, but then all of its neighbors in G belong to N. Therefore, all $C_{i,j}$'s and $C_{i',j'}$'s should i-touch, and $\mathcal{B}$ is an i-bramble. Each vertex $v = (i, j)$ in $\cup\,\mathcal{B}$ 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), and therefore $\mathbf{val}(\mathcal{B}) = 2k - 5$. For each nonmarginal vertex x in G, there are at most 25 nonmarginal vertices within a distance ≤ 2 in G (in the worst case, consider a triangulated (5×5)-grid subgraph of G that is centered at x), and thus $\mathbf{ifl}(\mathcal{B}) \leq 25$ (see Figure 3.2). 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. ◻

We require the following result.

PROPOSITION 3.3 (see [25, Theorem 6.2]). *Let k be a positive integer. Then every planar graph excluding $(k \times k)$-grid as a minor has treewidth at most $6k - 5$.*

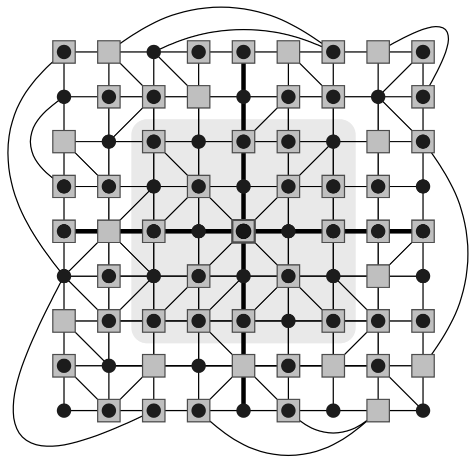

FIG. 3.2. *An i-labeled triangulated* (9×9)*-grid. The vertices meeting the two bold lines correspond to one of the* 49 *elements of the i-bramble* $\mathcal{B}$ *of the proof of Lemma* 3.2. *The vertex in the center meets* 13 *such elements, and there are* 25 *vertices of* N *in its closed* 2*-neighborhood.*

THEOREM 3.4. *If* $\mathcal{H}$ *is a hypergraph with a planar incidence graph* $I(H)$*, 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. Now suppose that $\mathcal{H}$ is a hypergraph where $\mathbf{fhw}(\mathcal{H}) \leq k$. By Lemmas 2.2 and 2.4, $\mathbf{fibn}(I(\mathcal{H}), V(\mathcal{H}), E(\mathcal{H})) = \mathbf{fbn}(\mathcal{H}) \leq \mathbf{fhw}(\mathcal{H}) \leq k$. By Lemmas 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 Proposition 3.3, $\mathbf{tw}(I(\mathcal{H})) \leq 6 \cdot l - 5 \leq 300 \cdot k + c$, and the result follows. ∎

3.3. Brambles in gridoids. We call a graph G 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 3.5. *Let* (G, N, M) *be an i-labeled graph, where* G *is a* (k, g)*-gridoid. Then* $\mathbf{fibn}(G, N, M) \geq k/50 - c \cdot (g + 1)$ *for some constant* $c \geq 0$.

Proof. The proof follows in the same way as the proof of Lemma 3.2. The only difference is that now we exclude from $\mathcal{B}$ all the $C_{i,j}$'s where either i or j is the coordinate of some endpoint of an additional edge. Notice that again $\mathbf{val}(\mathcal{B}) \leq 2k - 5$. Moreover, it also holds that $\mathbf{ifl}(\mathcal{B}) \leq 25$ as none of the endpoints belongs to elements of the bramble. Finally, $|\mathcal{B}| \geq (k - 2 - 2 \cdot g)^2$, and the result follows from Lemma 3.1. ∎

We need the following extension of Proposition 3.3 for graphs of bounded genus.

PROPOSITION 3.6 (see [7, Theorem 4.12]). *Let* k *be a positive integer. Then every graph of Euler genus* g *and excluding* $(k \times k)$*-grid as a minor has treewidth at most* $6k \cdot (g + 1)$.

The proof of the next theorem is very similar to the one of Theorem 3.4 (use Lemma 3.5 instead of Lemma 3.2, and Proposition 3.6 instead of Proposition 3.3).

THEOREM 3.7. *If* $\mathcal{H}$ *is a hypergraph with an incidence graph* $I(H)$ *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* $c \geq 0$.

3.4. 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 adjacent to at most s nonmarginal vertices of the grid.

LEMMA 3.8. *If (G, N, M) is an i-labeled graph where G is an augmented $(k \times k)$-grid with span s, then* $\mathbf{fibn}(G, N, M) \geq \frac{k}{2 \cdot s^2} - c$ *for some constant* $c \geq 0$.

Proof. We consider the i-bramble $\mathcal{B} = \{C_{i,j} \mid 2 \leq i, j \leq k-1\}$ of the proof of Lemma 3.2, and we directly observe that $\mathbf{val}(\mathcal{B}) \leq 2k - 5$ and $|\mathcal{B}| \geq (k-2)^2$. By the definition of the augmented $(k \times k)$-grid with span s, we obtain that $\mathbf{ifl}(\mathcal{B}) \leq s^2$, and the result follows by applying Lemma 3.1. ∎

As was shown by Demaine et al. [6], 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 graph). Because $f(k) = \Omega(k)$ (due to the results of Demaine and Hajiaghayi in [8]), we have the following proposition.

PROPOSITION 3.9. *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 3.4 (use Lemma 3.8 instead of Lemma 3.2, and Proposition 3.9 instead of Proposition 3.3).

THEOREM 3.10. *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 3.10 cannot be extended to hypergraphs whose 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 $\mathcal{H}$ a 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 this fact, in this section we prove that if a hypergraph has an 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 the generalized hypertree width and fractional hypertree width may differ 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 [27].

4.1. 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 $k \geq 0$ be 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 $F\colon W_1 \to W_2$ be a bijection between W_1 and W_2. We define the h-clique-sum of G_1 and G_2, denoted by $G_1 \bigoplus_{h,F} G_2$, or simply $G_1 \bigoplus G_2$ 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 $G_1 \bigoplus G_2$ 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 they had edges which were removed by clique-sum operation. Such edges are called *virtual* edges of G. We remark that $\bigoplus$ 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 and $h > 0$ be an integer. A graph G is h-nearly embeddable in Σ if there is a set of vertices $X \subseteq V(G)$ (called apexes) of size at most h such that graph $G - X$ is the union of subgraphs $G_0, \ldots, G_h$ with the following properties:*

(i) *There is a set of cycles $C_1, \ldots, C_h$ in Σ such that the cycles C_i are the borders of open pairwise disjoint discs Δ_i in Σ.*
(ii) *G_0 has an embedding in Σ in such a way that $G_0 \cap \bigcup_{i=1,\ldots,h} \Delta_i = \varnothing$.*
(iii) *Graphs $G_1, \ldots, G_h$ (called* vortices*) are pairwise disjoint, and for $1 \leq i \leq h$, $V(G_0) \cap V(G_i) \subset C_i$.*
(iv) *For $1 \leq i \leq h$, let $U_i := \{u_1^i, \ldots, u_{m_i}^i\}$ be the vertices of $V(G_0) \cap V(G_i) \subset C_i$ appearing in an order obtained by clockwise traversing of C_i; we call vertices of U_i* bases *of G_i. Then G_i has a path decomposition (P_i, χ_i) of width at most h such that P_i is the path on m_i vertices $1, \ldots, m_i$ and for $1 \leq j \leq m_i$, we have $u_j^i \in \chi_i(j)$.*

The following proposition is known as the excluded minor theorem [27] and is the cornerstone of Robertson and Seymour's graph minors theory.

THEOREM 4.3 (see [27]). *For every nonplanar 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. Moreover, while applying each of the clique-sums, at most three vertices from each summand other than apexes and vertices in vortices are identified.*

Let us remark that by the result of Demaine, Hajiaghayi, and Kawarabayashi [9] such a clique-sum decomposition can be obtained in time $O(n^c)$ for some constant c which depends only on H (see also [5]).

4.2. 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 [25]. Thus if H is planar, by the results of Theorem 3.10, 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 $G = G_1 \bigoplus G_2 \bigoplus \cdots \bigoplus G_m$ such that, for every $i \in \{1, 2, \ldots, m\}$, the summand G_i can be h-nearly embedded in a surface Σ in which H cannot 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 apexes 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 resulting graph by F_i.

In what remains, we show that the maximal value of $\mathbf{tw}(F_i)$, where the maximum is taken over all $i \in \{1, 2, \ldots, m\}$, is a constant factor approximation of generalized and fractional hypertree widths of $\mathcal{H}$. The upper bound is given by the following lemma. Its proof uses the fact that $\mathbf{ghw}(\mathcal{H}) \leq 3 \cdot \mathbf{mw}(\mathcal{H}) + 1$ (see [1]) and is based on the description of a winning strategy for $k = \max\{\mathbf{tw}(F_i) : i \in \{1, 2, \ldots, m\}\} + 2h + 1$ marshals on $\mathcal{H}$.

LEMMA 4.4. $\mathbf{ghw}(\mathcal{H}) \leq 3 \cdot \max\{\mathbf{tw}(F_i) : i \in \{1, 2, \ldots, m\}\} + 6h + 4$.

Proof. Let $w = \max\{\mathbf{tw}(F_i) : i \in \{1, 2, \ldots, m\}\}$ and $k = w + 2h + 1$. By the result of Adler, Gottlob, and Grohe [1], we have that $\mathbf{ghw}(\mathcal{H}) \leq 3 \cdot \mathbf{mw}(\mathcal{H}) + 1$, and it is enough to describe a winning strategy for k marshals on $\mathcal{H}$.

The clique-sum decomposition $G = G_1 \bigoplus G_2 \bigoplus \cdots \bigoplus G_m$ can be considered as a tree decomposition (T, χ) of G for some tree T with nodes $\{1, 2, \ldots, m\}$ such that $\chi(i) = V(G_i)$; i.e., the vertex sets of the summands are the bags of this decomposition. The idea behind the winning strategy for marshals is to "chase" the robber in the hypergraph along $m + 1$ decompositions for its incidence graph: One is induced by the clique-sum decomposition, and others are tree decompositions of F_i. We say that marshals *block* a set $X \subseteq V(G)$ if all hyperedges $X \cap E(\mathcal{H})$ are occupied by them, and for every $v \in X \cap V(\mathcal{H})$, there is a hyperedge $e \in E(\mathcal{H})$, occupied by a marshal, such that $v \in e$.

Let us note that the definition of F_i yields the following: If $x, y \in V(F_i)$ and there is an x, y-path in $G - (N[E_i] \cup A_i)$ with all inner vertices not in F_i, then $\{x, y\}$ is an edge of F_i. (Indeed, if $\{x, y\}$ is an edge of G, then it is also an edge of F_i. If $\{x, y\} \notin E(G)$ but such a path exists, then $\{x, y\}$ is a virtual edge in G_i and by the definition of F_i, such an edge is also an edge of F_i.)

For $i \in \{1, 2, \ldots, m\}$, let $(T^{(i)}, \chi_i)$ be a tree decomposition of F_i of width at most w. We assume that trees T and $T^{(1)}, T^{(2)}, \ldots, T^{(m)}$ are rooted trees with roots r and $r_1, r_2, \ldots, r_m$, respectively.

For a node $i \in V(T)$ and its parent j (in T), we define $S = V(G_i) \cap V(G_j)$. (If $i = r$, then we set $S = \varnothing$.) By the definition of the clique-sum, $|S| \leq h$. Assume that at most h marshals are already placed on the hypergraph in such a way that they block S. Assume also that the robber occupies some vertex of $\chi(T^{(i)})$. We put at most h marshals on hyperedges to block the set of apexes A_i. Then the set $N_G[E_i] \cup A_i$ is blocked by these marshals.

Now marshals start to "chase" the robber in the subhypergraph induced by the vertex set $V(F_i) \cap V(\mathcal{H})$ along $T^{(i)}$. We put at most $w + 1$ marshals to block the set $\chi_i(r_i)$. Now assume that some set $\chi_i(x)$ for $x \in V(T^{(i)})$ is blocked; i.e., for any hyperedge $e \in \chi_i(x) \cap E(\mathcal{H})$, e is occupied by a marshal, and for every $v \in \chi_i(x) \cap V(\mathcal{H})$, there is a hyperedge $e \in E(\mathcal{H})$, occupied by a marshal, such that $v \in e$. In the last case we say that this marshal is *assigned* to v. We remove marshals that do not occupy the hyperedges $\chi_i(x) \cap \chi_i(y) \cap E(\mathcal{H})$ that are not assigned to the vertices of $\chi_i(x) \cap \chi_i(y) \cap V(\mathcal{H})$. Observe that at most $|\chi_i(x) \cap \chi_i(y)|$ marshals are used now, and $\chi_i(x) \cap \chi_i(y)$ remains blocked. Then we place marshals on the hyperedges $(\chi_i(y) \setminus \chi_i(x)) \cap E(\mathcal{H})$, and for any $v \in (\chi_i(y) \setminus \chi_i(x)) \cap V(\mathcal{H})$, we choose a hyperedge e such that $v \in e$ and place a marshal on e. This maneuver is done by making use of at most $w + 1 - |\chi_i(x) \cap \chi_i(y)|$ marshals. We set $x = y$ and repeat this operation until there is a child y of x such that the robber can be in $\chi_i(T_y^{(i)})$. Thus by repeating at most $|V(T^{(i)})|$ times this operation, marshals "push" the robber out of $V(F_i) \cap V(\mathcal{H})$.

Let j be a child of i in T such that the robber now can occupy only the vertices of $\chi(T_j)$, where T_j is the subtree of T rooted at j. Let $S' = V(G_i) \cap V(G_j)$. Since $|S'| \leq h$, we have that h marshals can block this set and, after that, all other marshals can be removed from $\mathcal{H}$.

We apply the described strategy of marshals starting from $i = r$ until the robber is captured in some leaf-node of T. For every node of T we have used at most h marshals to occupy apexes, at most h marshals to block the vertices of the clique-sum, and at most $w + 1$ marshals to push the robber out of F_i. Thus in total at most $2h + w + 1$ marshals have a winning strategy on $\mathcal{H}$. ∎

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 [23] and Thomassen [28] (see also [7, Lemma 3.3]). We use the following variant of this result from Geelen, Richter, and Salazar [13].

Proposition 4.5 (see [13]). *Let g, l, and 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 4.6. $\mathbf{fbn}(\mathcal{H}) \geq \varepsilon_H \cdot \max\{\mathbf{tw}(F_i) : i \in \{1,2,\ldots,m\}\}$ *for some constant ε_H depending only on H.*

Proof. Let $i \in \{1,2,\ldots,m\}$. 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 the 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 $G = G_1 \bigoplus G_2 \bigoplus \cdots \bigoplus G_m$ 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 $G'_1, G'_2, \ldots, G'_l$ 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 $F_i^{(1)} \bigoplus F_i^{(2)}$ 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 whose leaves are the vertices of Q_j.

Contracting vortices. The h-nearly embedding of the graph G_i induces the h-nearly embedding of $F_i = X_0 \cup X_1 \cup \cdots \cup X_h$ without apexes. Here we assume that X_0 is embedded in a surface Σ of genus depending on H and $X_1, X_2, \ldots, X_h$ 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 $Q_k \cap (V(X_j) - V(C_j)) \neq \varnothing$, then all edges of S_k are contracted, and if $Q_k \cap (V(X_j) - V(C_j)) = \varnothing$ 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 result in the contraction of some edges of F_i. In particular, 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 these 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 the fact that at most three vertices that do not belong in vortices or apexes are identified, that F'_i coincides with the graph obtained from F_i by contracting all vortices X_j and all boundaries of faces C_j. It can be easily seen that F'_i is embedded in Σ. It is known (see, e.g., [7], [8]) 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 [7]. Combining this result with Proposition 4.5, we obtain the following claim.

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.

The claim above permits us to treat the part of F'_i which is outside Δ exactly in the same way we have treated vortices. For stars S_k intersecting $V(F'_i) - V(R')$ or C, we do the following: If $Q_k \cap (V(F'_i) - V(R')) \neq \varnothing$, then all edges of S_k are contracted, and if $Q_k \cap (V(F'_i) - V(R')) = \varnothing$ 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'. In particular, 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 by 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. Given the fact that while taking clique-sums, at most three vertices that do not belong in vortices or apexes are identified, we obtain that $|Q_j| \leq 3$. 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 the 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. A similar construction was used in the proof of the main theorem in [8, Theorem 1.2], and by the same arguments as were used by Demaine and Hajiaghayi we immediately conclude that there is a positive constant d_H such that $\mathbf{tw}(X) \leq d_H \cdot \mathbf{tw}(Y)$.

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, $\mathbf{fbn}(\mathcal{H}) = \mathbf{fibn}(G, V(\mathcal{H}), E(\mathcal{H}))$. Because the sets $V(\mathcal{H})$ and $E(\mathcal{H})$ are independent, by Lemma 2.6, we have that $\mathbf{fibn}(G, V(\mathcal{H}), E(\mathcal{H})) \geq \mathbf{fibn}(G', N, M)$, 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, $\mathbf{fibn}(G', N, M) \geq \mathbf{fibn}(Y, N', M')$, where N' and M' are sets which were obtained as the result of contractions of N and M. Finally, as in Theorem 3.4, one can show that $\mathbf{fibn}(Y, N', M') \geq f_H \cdot \mathbf{tw}(Y)$ 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 Lemmas 2.1, 2.2, 4.4, and 4.6, we obtain the following theorem.

THEOREM 4.7. *$(1/c_H) \cdot w \leq \mathbf{fhw}(\mathcal{H}) \leq \mathbf{ghw}(\mathcal{H}) \leq c_H \cdot w$, where $w = \max\{\mathbf{tw}(F_i) : i \in \{1, 2, \ldots, m\}\}$ and c_H is a constant depending only on H.*

Remark. Notice that, by Theorem 4.7, 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 [20], this is not the case for general hypergraphs.

Demaine, Hajiaghayi, and Kawarabayashi [9] (see also [5], [10], [27]) 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 long as we construct 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 4.8. *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.*

5. Conclusion. Let us remark that while the winning strategy for marshals used in the proof of Lemma 4.4 is not monotone (a strategy is *monotone* if the territory available for the robber decreases only in the game), it can be turned into monotone by choosing marshals' positions in a slightly more careful way. By making use of the results from [18], the monotone strategy can be used to construct a generalized hypertree decomposition (or fractional hypertree decomposition). Thus our results can be used not only to approximate but also to construct the corresponding decompositions as well.

A long-standing open question in graph algorithms is whether the treewidth of a planar graph is computable in polynomial time. It would be very interesting to see whether an NP-hardness proof can be obtained for at least one of the hypertree width parameters of planar hypergraphs. Moreover, a polynomial time algorithm for any of these hypertree width parameters on planar hypergraphs would be a significant breakthrough. While we do not know whether there exists a polynomial time algorithm for any of these problems, it would be challenging to ask if some of the variants of the problem are fixed parameter tractable on planar hypergraphs. On the other hand, the treewidth of a planar graph admits a constant factor approximation. Our results extend this algorithmic result to all the hypertree width parameters on planar hypergraphs.

Finally, the sparsity of hypergraphs studied in this paper is expressed in terms of their incidence graphs. It is an interesting question whether there are other sparsity measures where further algorithmic or complexity results can be obtained for hypertree width parameters.

Acknowledgment. We are grateful to the anonymous referees for their constructive suggestions and remarks.

REFERENCES

[1] I. Adler, G. Gottlob, and M. Grohe, *Hypertree width and related hypergraph invariants*, Eur. J. Combin., 28 (2007), pp. 2167–2181.

[2] H. L. Bodlaender, *A partial k-arboretum of graphs with bounded treewidth*, Theoret. Comput. Sci., 209 (1998), pp. 1–45.

[3] A. K. Chandra and P. M. Merlin, *Optimal implementation of conjunctive queries in relational data bases*, in Proceedings of the 9th Annual ACM Symposium on Theory of Computing, Boulder, CO, 1977, pp. 77–90.

[4] C. Chekuri and A. Rajaraman, *Conjunctive query containment revisited*, Theoret. Comput. Sci., 239 (2000), pp. 211–229.

[5] A. Dawar, M. Grohe, and S. Kreutzer, *Locally excluding a minor*, in Proceedings of the 22nd IEEE Symposium on Logic in Computer Science, Wroclaw, Poland, 2007, pp. 270–279.

[6] E. D. Demaine, F. V. Fomin, M. Hajiaghayi, and D. M. Thilikos, *Bidimensional parameters and local treewidth*, SIAM J. Discrete Math., 18 (2004), pp. 501–511.

[7] E. D. Demaine, F. V. Fomin, M. Hajiaghayi, and D. M. Thilikos, *Subexponential parameterized algorithms on graphs of bounded genus and H-minor-free graphs*, J. ACM, 52 (2005), pp. 866–893.

[8] 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 Proceedings of the 6th Annual ACM-SIAM Symposium on Discrete Algorithms, Vancouver, Canada, 2005, pp. 682–689.

[9] E. D. Demaine, M. T. Hajiaghayi, and K. Kawarabayashi, *Algorithmic graph minor theory: Decomposition, approximation, and coloring*, in Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science, Pittsburgh, PA, 2005, pp. 637–646.

[10] U. Feige, M. Hajiaghayi, and J. R. Lee, *Improved approximation algorithms for minimum weight vertex separators*, SIAM J. Comput., 38 (2008), pp. 629–657.

[11] F. V. Fomin, P. A. Golovach, and D. M. Thilikos, *Approximating acyclicity parameters of sparse hypergraphs*, in Proceedings of the 26th Symposium on Theoretical Aspects of Computer Science (STACS 2009), Freiburg, Germany, Schloss Dagstuhl - Leibniz Center for Informatics, 2009.

[12] F. V. Fomin and D. M. Thilikos, *An annotated bibliography on guaranteed graph searching*, Theoret. Comput. Sci., 399 (2008), pp. 236–245.

[13] J. F. Geelen, R. B. Richter, and G. Salazar, *Embedding grids in surfaces*, Eur. J. Combin., 25 (2004), pp. 785–792.

[14] G. Gottlob, M. Grohe, N. Musliu, M. Samer, and F. Scarcello, *Hypertree decompositions: Structure, algorithms, and applications*, in Graph-Theoretic Concepts in Computer Science, 31st International Workshop, Lecture Notes in Comput. Sci. 2787, Springer, Berlin, 2005, pp. 1–15.

[15] G. Gottlob, N. Leone, and F. Scarcello, *A comparison of structural CSP decomposition methods*, Artificial Intelligence, 124 (2000), pp. 243–282.

[16] G. Gottlob, N. Leone, and F. Scarcello, *The complexity of acyclic conjunctive queries*, J. ACM, 48 (2001), pp. 431–498.

[17] G. Gottlob, N. Leone, and F. Scarcello, *Hypertree decompositions and tractable queries*, J. Comput. System Sci., 64 (2002), pp. 579–627.

[18] G. Gottlob, N. Leone, and F. Scarcello, *Robbers, marshals, and guards: Game theoretic and logical characterizations of hypertree width*, J. Comput. System Sci., 66 (2003), pp. 775–808.

[19] G. Gottlob, Z. Miklós, and T. Schwentick, *Generalized hypertree decompositions: NP-hardness and tractable variants*, J. ACM, 56 (2009), paper 6.

[20] M. Grohe and D. Marx, *Constraint solving via fractional edge covers*, in Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms, Miami, FL, 2006, pp. 289–298.

[21] P. G. Kolaitis and M. Y. Vardi, *Conjunctive-query containment and constraint satisfaction*, J. Comput. System Sci., 61 (2000), pp. 302–332.

[22] D. Marx, *Approximating fractional hypertree width*, in Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms, New York, 2009, pp. 902–911.

[23] B. Mohar, *Combinatorial local planarity and the width of graph embeddings*, Canad. J. Math., 44 (1992), pp. 1272–1288.

[24] B. Mohar and C. Thomassen, *Graphs on Surfaces*, Johns Hopkins University Press, Baltimore, MD, 2001.

[25] N. Robertson, P. Seymour, and R. Thomas, *Quickly excluding a planar graph*, J. Combin. Theory Ser. B, 62 (1994), pp. 323–348.

[26] N. Robertson and P. D. Seymour, *Graph minors.* II. *Algorithmic aspects of tree-width*, J. Algorithms, 7 (1986), pp. 309–322.

[27] N. Robertson and P. D. Seymour, *Graph minors.* XVI. *Excluding a non-planar graph*, J. Combin. Theory Ser. B, 89 (2003), pp. 43–76.

[28] C. Thomassen, *A simpler proof of the excluded minor theorem for higher surfaces*, J. Combin. Theory Ser. B, 70 (1997), pp. 306–311.

[29] M. Yannakakis, *Algorithms for acyclic database schemes*, in Proceedings of the 7th International IEEE Very Large Databases Conference, Cannes, France, 1981, pp. 82–94.