

Tractable hypergraph properties for constraint satisfaction and conjunctive queries¹

Dániel Marx, Computer and Automation Research Institute, Hungarian Academy of Sciences (MTA SZTAKI), Budapest, Hungary, dmarx@cs.bme.hu.²

An important question in the study of constraint satisfaction problems (CSP) is understanding how the graph or hypergraph describing the incidence structure of the constraints influences the complexity of the problem. For binary CSP instances (i.e., where each constraint involves only two variables), the situation is well understood: the complexity of the problem essentially depends on the treewidth of the graph of the constraints [Grohe 2007; Marx 2010b]. However, this is not the correct answer if constraints with unbounded number of variables are allowed, and in particular, for CSP instances arising from query evaluation problems in database theory. Formally, if \mathcal{H} is a class of hypergraphs, then let $\text{CSP}(\mathcal{H})$ be CSP restricted to instances whose hypergraph is in \mathcal{H} . Our goal is to characterize those classes of hypergraphs for which $\text{CSP}(\mathcal{H})$ is polynomial-time solvable or fixed-parameter tractable, parameterized by the number of variables. Note that in the applications related to database query evaluation, we usually assume that the number of variables is much smaller than the size of the instance, thus parameterization by the number of variables is a meaningful question.

The most general known property of \mathcal{H} that makes $\text{CSP}(\mathcal{H})$ polynomial-time solvable is bounded fractional hypertree width. Here we introduce a new hypergraph measure called *submodular width*, and show that bounded submodular width of \mathcal{H} (which is a strictly more general property than bounded fractional hypertree width) implies that $\text{CSP}(\mathcal{H})$ is fixed-parameter tractable. In a matching hardness result, we show that if \mathcal{H} has unbounded submodular width, then $\text{CSP}(\mathcal{H})$ is not fixed-parameter tractable (and hence not polynomial-time solvable), unless the Exponential Time Hypothesis (ETH) fails. The algorithmic result uses tree decompositions in a novel way: instead of using a single decomposition depending on the hypergraph, the instance is split into a set of instances (all on the same set of variables as the original instance), and then the new instances are solved by choosing a different tree decomposition for each of them. The reason why this strategy works is that the splitting can be done in such a way that the new instances are “uniform” with respect to the number extensions of partial solutions, and therefore the number of partial solutions can be described by a submodular function. For the hardness result, we prove via a series of combinatorial results that if a hypergraph H has large submodular width, then a 3SAT instance can be efficiently simulated by a CSP instance whose hypergraph is H . To prove these combinatorial results, we need to develop a theory of (multicommodity) flows on hypergraphs and vertex separators in the case when the function $b(S)$ defining the cost of separator S is submodular, which can be of independent interest.

Categories and Subject Descriptors: F.2 [Theory of Computing]: Analysis of Algorithms and Problem Complexity; G.2.2 [Mathematics of Computing]: Discrete Mathematics—*Graph Theory*

General Terms: Algorithms

Additional Key Words and Phrases: constraint satisfaction, hypergraphs, hypertree width, fractional edge covers

1. INTRODUCTION

There is a long line of research devoted to identifying hypergraph properties that make the evaluation of conjunctive queries tractable (see e.g. [Gottlob et al. 2002a; Scarcello et al. 2008; Grohe 2006; 2007]). Our main contribution is giving a complete theoretical answer to this question: in a very precise technical sense, we characterize those hypergraph properties that imply tractability for the evaluation of a query. Efficient evaluation of queries is originally a question of database theory; however, it has been noted that the problem can be treated as a constraint satisfaction problem (CSP) and this connection led to a fruitful

¹An extended abstract of the paper was presented at the 42nd ACM Symposium on Theory of Computing (STOC 2010) [Marx 2010c].

²Research supported by the European Research Council (ERC) grant “PARAMTIGHT: Parameterized complexity and the search for tight complexity results,” reference 280152.

interaction between the two communities [Kolaitis and Vardi 2000a; Gottlob and Szeider 2008; Scarcello et al. 2008]. Most of the literature relevant to the current paper use the language of constraint satisfaction. Therefore, after a brief explanation of the database-theoretic motivation, we switch to the language of CSPs.

Conjunctive queries. Evaluation of conjunctive queries (or equivalently, Select-Project-Join queries) is one of the most basic and most studied tasks in relational databases. A relational database consists of a fixed set of relations. A conjunctive query defines a new relation that can be obtained as first taking the join of some relations and then projecting it to a subset of the variables. As an example, consider a relational database that contains three relations: $\text{enrolled}(\text{Person}, \text{Course}, \text{Date})$, $\text{teaches}(\text{Person}, \text{Course}, \text{Year})$, $\text{parent}(\text{Person1}, \text{Person2})$. The following query Q defines a unary relation $\text{ans}(P)$ with the meaning that “ P is enrolled in a course taught by her parent.”

$$Q : \text{ans}(P) \leftarrow \text{enrolled}(P, C, D) \wedge \text{teaches}(P2, C, Y) \wedge \text{parent}(P2, P).$$

In the *Boolean Conjunctive Query* problem, the task is only to decide if the answer relation is empty or not, that is, if the join of the relations is empty or not. This is usually denoted as the relation “ans” not having any variables. Boolean Conjunctive Query contains most of the combinatorial difficulty of the general problem without complications such that the size of the output being exponentially large. Therefore, the current paper focuses on this decision problem.

In a natural way, we can define the *hypergraph* of a query: its vertices are the variables appearing in the query and for each relation there is a corresponding hyperedge containing the variables appearing in the relation. Intuitively, if the hypergraph has “simple structure,” then the query is easy to solve. For example, compare the following two queries:

$$\begin{aligned} Q_1 : \text{ans} &\leftarrow R_1(A, B, C) \wedge R_2(C, D) \wedge R_3(D, E, F) \wedge R_4(E, F, G, H) \wedge R_5(H, I) \\ Q_2 : \text{ans} &\leftarrow R_1(A, B) \wedge R_2(A, C) \wedge R_3(A, D) \wedge R_4(B, C) \wedge R_5(B, D) \wedge R_6(C, D) \end{aligned}$$

Even though more variables appear in Q_1 , evaluating it seems to be easier: its hypergraph is “path like,” thus the query can be answered efficiently by, say, dynamic programming techniques. On the other hand, the hypergraph of Q_2 is a clique on 4 vertices and no significant shortcut is apparent compared to trying all possible combinations of values for (A, B, C, D) .

What are those hypergraph properties that make Boolean Conjunctive Query tractable? In the early 80s, it has been noted that acyclicity is one such property [Beeri et al. 1983; Fagin 1983; Yannakakis 1981; Beeri et al. 1981]. Later, more general such properties were identified in the literature: for example, bounded query width [Chekuri and Rajaraman 2000], bounded hypertree width [Gottlob et al. 2002a], and bounded fractional hypertree width [Marx 2010a; Grohe and Marx 2012]. Our goal is to find the most general hypergraph property that guarantees an efficient solution for query evaluation.

Constraint satisfaction. Constraint satisfaction is a general framework that includes many standard algorithmic problems such as satisfiability, graph coloring, database queries, etc. [Grohe 2006; Feder and Vardi 1999]. A constraint satisfaction problem (CSP) consists of a set V of variables, a domain D , and a set C of constraints, where each constraint is a relation on a subset of the variables. The task is to assign a value from D to each variable in such a way that every constraint is satisfied (see Definition 2.1 in Section 2 for the formal definition). For example, 3SAT can be interpreted as a CSP problem where the domain is $D = \{0, 1\}$ and the constraints in C correspond to the clauses (thus the arity of each constraint is 3). As another example, let us observe that the k -Clique problem (Is there a k -clique in a given graph G ?) can be easily expressed as a CSP instance the following way. Let D be the set of vertices of G , let V contain k variables, and let C contain $\binom{k}{2}$ constraints, one constraint on each pair of variables. The binary relation of these constraints require

that the two vertices are distinct and adjacent. Therefore, the CSP instance has a solution if and only if G has a k -clique.

It is easy to see that Boolean Conjunctive Query can be formulated as the problem of deciding if a CSP instance has a solution: the variables of the CSP instance correspond to the variables appearing in the query and the constraints correspond to the database relations. A distinctive feature of CSP instances obtained this way is that the number of variables is small (as queries are typically small), while the domain of the variables are large (as the database relations usually contain a large number of entries). This has to be contrasted with typical CSP problems from AI, such as 3-colorability and satisfiability, where the domain is small, but the number of variables is large. As our motivation is database-theoretic, in the rest of the paper the reader should keep in mind that we are envisioning scenarios where the number of variables is small and the domain is large.

As the examples above show, solving constraint satisfaction problems is NP-hard in general if there are no additional restrictions on the input instances. The main goal of the research on CSP is to identify tractable special cases of the general problem. The theoretical literature on CSP investigates two main types of restrictions. The first type is to restrict the *constraint language*, that is, the type of constraints that are allowed. This direction includes the classical work of Schaefer [1978] and its many generalizations [Bulatov 2006; 2003; Bulatov et al. 2001; Feder and Vardi 1999; Jeavons et al. 1997]. The second type is to restrict the *structure* induced by the constraints on the variables. The *hypergraph* of a CSP instance is defined to be a hypergraph on the variables of the instance such that for each constraint $c \in C$ there is a hyperedge e_c containing exactly the variables that appear in c . If the hypergraph of the CSP instance has very simple structure, then the instance is easy to solve. For example, it is well-known that a CSP instance I with hypergraph H can be solved in time $\|I\|^{O(\text{tw}(H))}$ [Freuder 1990], where $\text{tw}(H)$ denotes the treewidth of H and $\|I\|$ is the size of the representation of I in the input.

Our goal is to characterize the “easy” and “hard” hypergraphs from the viewpoint of constraint satisfaction. However, formally speaking, CSP is polynomial-time solvable for every fixed hypergraph H : since H has a constant number k of vertices, every CSP instance with hypergraph H can be solved by trying all $\|I\|^k$ possible combinations on the k variables. It makes more sense to characterize those *classes* of hypergraphs where CSP is easy. Formally, for a class \mathcal{H} of hypergraphs, let $\text{CSP}(\mathcal{H})$ be the restriction of CSP where the hypergraph of the instance is assumed to be in \mathcal{H} . For example, as discussed above, we know that if \mathcal{H} is a class of hypergraphs with bounded treewidth (i.e., there is a constant w such that $\text{tw}(H) \leq w$ for every $H \in \mathcal{H}$), then $\text{CSP}(\mathcal{H})$ is polynomial-time solvable.

For the characterization of the complexity of $\text{CSP}(\mathcal{H})$, we can investigate two notions of tractability. $\text{CSP}(\mathcal{H})$ is *polynomial-time solvable* if there is an algorithm solving every instance of $\text{CSP}(\mathcal{H})$ in time $(\|I\|)^{O(1)}$, where $\|I\|$ is the length of the representation of I in the input. The following notion interprets tractability in a less restrictive way: $\text{CSP}(\mathcal{H})$ is *fixed-parameter tractable (FPT)* if there is an algorithm solving every instance I of $\text{CSP}(\mathcal{H})$ in time $f(H)(\|I\|)^{O(1)}$, where f is an arbitrary computable function of the hypergraph H of the instance. Equivalently, the factor $f(H)$ in the definition can be replaced by a factor $f(k)$ depending only on the number k of vertices of H : as the number of hypergraphs on k vertices (without parallel edges) is bounded by a function of k , the two definitions result in the same notion. The motivation behind this definition is that if the number of variables is assumed to be much smaller than the the domain size, then we can afford even exponential dependence on the number of variables, as long as the dependence on the size of the instance is polynomial. For a more background on fixed-parameter tractability, the reader is referred to the parameterized complexity literature [Downey and Fellows 1999; Flum and Grohe 2006; Niedermeier 2006].

The case of bounded arities. If the constraints have bounded arity (i.e., the edge size in \mathcal{H} is bounded by a constant r), then the complexity of $\text{CSP}(\mathcal{H})$ is well understood. In this case, bounded treewidth is the only polynomial-time solvable case:

THEOREM 1.1 (GROHE [2007]). *If \mathcal{H} is a recursively enumerable class of hypergraphs with bounded edge size, then (assuming $\text{FPT} \neq \text{W}[1]$) the following are equivalent:*

- (1) $\text{CSP}(\mathcal{H})$ is polynomial-time solvable.
- (2) $\text{CSP}(\mathcal{H})$ is fixed-parameter tractable.
- (3) \mathcal{H} has bounded treewidth.

The assumption $\text{FPT} \neq \text{W}[1]$ is a standard hypothesis of parameterized complexity. Thus in the bounded arity case bounded treewidth is the only property of the hypergraph that can make the problem polynomial-time solvable. By definition, polynomial-time solvability implies fixed-parameter tractability, but Theorem 1.1 proves the surprising result that whenever $\text{CSP}(\mathcal{H})$ is fixed-parameter tractable, it is polynomial-time solvable as well.

The following sharpening of Theorem 1.1 shows that there is no algorithm whose running time is significantly better than the $\|I\|^{O(\text{tw}(H))}$ bound of the treewidth based algorithm, and this is true if we restrict the problem to *any* class \mathcal{H} of hypergraphs. The result is proved under the Exponential Time Hypothesis (ETH) [Impagliazzo et al. 2001] stating that there is no $2^{o(n)}$ time algorithm for n -variable 3SAT, which is a somewhat stronger assumption than $\text{FPT} \neq \text{W}[1]$.

THEOREM 1.2 (MARX [2010B]). *If there is a computable function f and a recursively enumerable class \mathcal{H} of hypergraphs with bounded edge size and unbounded treewidth such that the problem $\text{CSP}(\mathcal{H})$ can be solved in time $f(H)\|I\|^{o(\text{tw}(H)/\log \text{tw}(H))}$ for instances I with hypergraph $H \in \mathcal{H}$, then ETH fails.*

This means that the treewidth-based algorithm is almost optimal on every class of hypergraphs: in the exponent only an $O(\log \text{tw}(H))$ factor improvement is possible. It is conjectured in [Marx 2010b] that Theorem 1.2 can be made tight, i.e., the lower bound holds even if the logarithmic factor is removed from the exponent.

CONJECTURE 1.3 (MARX [2010B]). *If \mathcal{H} is a class of hypergraphs with bounded edge size, then there is no algorithm that solves $\text{CSP}(\mathcal{H})$ in time $f(H)\|I\|^{o(\text{tw}(H))}$ for instances I with hypergraph $H \in \mathcal{H}$, where f is an arbitrary function.*

Unbounded arities. The situation is less understood in the unbounded arity case, i.e., when there is no bound on the maximum edge size in \mathcal{H} . First, the complexity in the unbounded-arity case depends on how the constraints are represented. In the bounded-arity case, if each constraint contains at most r variables (r being a fixed constant), then every reasonable representation of a constraint has size $|D|^{O(r)}$. Therefore, the size of the different representations can differ only by a polynomial factor. On the other hand, if there is no bound on the arity, then there can be exponential difference between the size of succinct representations (e.g., formulas [Chen and Grohe 2010]) and verbose representations (e.g., truth tables [Marx 2011]). The running time of an algorithm is expressed as a function of the input size, hence the complexity of the problem can depend on how the input is represented: longer representation means that it is potentially easier to obtain a polynomial-time algorithm.

The most well-studied representation of constraints is listing all the tuples that satisfy the constraint. This representation is perfectly compatible with our database-theoretic motivation: the constraints are relations of the database, and a relation is physically stored as a table containing all the tuples in the relation. For this representation, there are classes \mathcal{H} with unbounded treewidth such that CSP restricted to this class is polynomial-time solvable. A trivial example is the class \mathcal{H} of all hypergraphs having only a single hyperedge of

arbitrary size. The treewidth of such hypergraphs can be arbitrarily large (as the treewidth of a hypergraph consisting of a single edge e is exactly $|e| - 1$), but $\text{CSP}(\mathcal{H})$ is trivial to solve: we can pick any tuple from the constraint corresponding to the single edge. There are other, nontrivial, classes of hypergraphs with unbounded treewidth such that $\text{CSP}(\mathcal{H})$ is solvable in polynomial time: for example, classes with bounded *(generalized) hypertree width* [Gottlob et al. 2002b], bounded *fractional edge cover number* [Grohe and Marx 2012], and bounded *fractional hypertree width* [Grohe and Marx 2012; Marx 2010a]. Thus, unlike in the bounded-arity case, treewidth is not the right measure for characterizing the complexity of the problem.

Our results. We introduce a new hypergraph width measure that we call *submodular width*. Small submodular width means that for every monotone submodular function b on the vertices of the hypergraph H , there is a tree decomposition where $b(B)$ is small for every bag B of the decomposition. (This definition makes sense only if we normalize the considered functions: for this reason, we require that $b(e) \leq 1$ for every edge e of H .) The main result of the paper is showing that bounded submodular width is the property that precisely characterizes the complexity of $\text{CSP}(\mathcal{H})$:

THEOREM 1.4 (MAIN). *Let \mathcal{H} be a recursively enumerable class of hypergraphs. Assuming the Exponential Time Hypothesis, $\text{CSP}(\mathcal{H})$ parameterized by H is fixed-parameter tractable if and only if \mathcal{H} has bounded submodular width.*

Theorem 1.4 has an algorithmic side (algorithm for bounded submodular width) and a complexity side (hardness result for unbounded submodular width). Unlike previous width measures in the literature, where small value of the measure suggests a way of solving $\text{CSP}(\mathcal{H})$ it is not at all clear how bounded submodular width is of any help. In particular, it is not obvious what submodular functions have to do with CSP instances. The main idea of our algorithm is that a CSP instance can be “split” into a small number of “uniform” CSP instances; for this purpose, we use a partitioning procedure inspired by a result of Alon et al. [2007]. More precisely, splitting means that we partition the set of tuples appearing in the constraint relations in a certain way and each new instance inherits only one class of the partition (thus each new instance has the same set of variables as the original). Uniformity means that for any subset $B \subseteq A$ of variables, every solution for the problem restricted to B has roughly the same number of extensions to A . The property of uniformity allows us to bound the logarithm of the number of solutions on the different subsets by a submodular function. Therefore, bounded submodular width guarantees that each uniform instance has a tree decomposition where only a polynomially bounded number of solutions has to be considered in each bag.

Conceptually, our algorithm goes beyond previous decomposition techniques in two ways. First, the tree decomposition that we use depends not only on the hypergraph, but on the actual constraint relations in the instance (we remark that this idea first appeared in [Marx 2011] in a different context that does not directly apply to our problem). Second, we are not only decomposing the set of variables, but we also split the constraint relations. This way, we can apply different decompositions to different parts of the solution space.

The proof of the complexity side of Theorem 1.4 follows the same high-level strategy as the proof of Theorem 1.2 in [Marx 2010b]. In a nutshell, the argument of [Marx 2010b] is the following: if treewidth is large, then there is subset of vertices which is highly connected in the sense that the set does not have a small balanced separator; such a highly connected set implies that there is uniform concurrent flow (i.e., a compatible set of flows connecting every pair of vertices in the set); the paths in the flows can be used to embed the graph of a 3SAT formula; and finally this embedding can be used to reduce 3SAT to CSP. These arguments build heavily on well-known characterizations of treewidth and results from combinatorial optimization (such as the $O(\log k)$ integrality gap of sparsest cut). The proof of Theorem 1.4 follows this outline, but now no such well-known tools are available: we are




Fig. 1. Connections between different domains.

dealing with hypergraphs and submodular functions in a way that was not explored before in the literature. Thus we have to build from scratch all the necessary tools. One of the main difficulties of obtaining Theorem 1.4 is that we have to work in three different domains:

- **CSP instances.** As our goal is to investigate the existence of algorithms solving CSP, the most obvious domain is CSP instances. In light of previous results, we are especially interested in algorithms based on tree decompositions. For such algorithms, what matters is the existence of subsets of vertices such that restricting the instance to any of these subsets gives an instance with “small” number of solutions. In order to solve the instance, we would like to find a tree decomposition where every bag is such a small set.

- **Submodular functions.** Submodular width is defined in terms of submodular functions, thus submodular functions defined on hypergraphs is our second natural domain. We need to understand what large submodular width means, that is, what property of the submodular function and the hypergraph makes it impossible to obtain a tree decomposition where every bag has small value.

- **Flows and embeddings in hypergraphs.** In the hardness proof, our goal is to embed the graph of a 3SAT formula into a hypergraph. Thus we need to define an appropriate notion of embedding and study what guarantees the existence of embeddings with suitable properties. As in [Marx 2010b], we use the paths appearing in flows to construct embeddings. For our purposes, the right notion of flow is a collection of weighted paths where the total weight of the paths intersecting each hyperedge is at most 1. This notion of flows has not been studied in the literature before, thus we need to obtain basic results on such flows, such as exploring the duality between flows and separators.

A key question is how to find connections between these domains. As mentioned above and detailed in Section 4, we have a procedure that reduces a CSP instance into a set of uniform CSP instances, and the number of solutions on the different subsets of variables in a uniform CSP instance can be described by a submodular function. This method allows us to move from the domain of CSP instances to the domain of submodular functions.

Section 5 is devoted to showing that if submodular width of a hypergraph is large, then there is a certain “highly connected” set in the hypergraph. Highly connected set is defined as a property of the hypergraph (it requires the existence of certain flows) and has no longer anything to do with submodular functions. Thus this connection allows us to move from the domain of submodular functions to the study of hypergraphs. In Section 6, we show that a highly connected set in a hypergraph means that graphs can be efficiently embedded into the hypergraph. In particular, the graph of a 3SAT formula can be embedded into the hypergraph, which gives us (as shown in Section 7) a reduction from 3SAT to $\text{CSP}(\mathcal{H})$. This connection allows us to move from the domain of embeddings back to the domain of CSP instances. We remark that Sections 4–7 are written in a self-contained way: only the first theorem of each section is used outside the section.

A consequence of our characterization of submodular width implies the surprising fact that bounded submodular width equals bounded adaptive width (defined in [Marx 2011]):

THEOREM 1.5. *A class of hypergraphs has bounded submodular width if and only if it has bounded adaptive width.*

It is proved in [Marx 2011] that there are classes of hypergraphs having bounded adaptive width (and hence bounded submodular width), but unbounded fractional hypertree width. Previously, bounded fractional hypertree width was the most general property that was known to guarantee fixed-parameter tractability [Grohe and Marx 2012]. Thus Theorem 1.4 not only gives a complete characterization of the parameterized complexity of $\text{CSP}(\mathcal{H})$, but its algorithmic side proves fixed-parameter tractability in a strictly more general case than what was known before.

Why fixed-parameter tractability? We argue that investigating the fixed-parameter tractability of $\text{CSP}(\mathcal{H})$ is at least as interesting as investigating polynomial-time solvability. In problems coming from our database-theoretic motivation, the size of the hypergraph (that is, the size of the query) is assumed to be much smaller than the input size (which is usually dominated by the size of the database), hence a constant factor in the running time depending only on the number of variables (or on the hypergraph) is acceptable³. Even the STOC 1977 landmark paper of Chandra and Merlin [1977], which started the complexity research on conjunctive queries, suggests spending exponential time (in the size of the query) on finding the best possible evaluation order. Furthermore, the notion of fixed-parameter tractability formalizes the usual viewpoint of the literature on conjunctive queries: in the complexity analysis, we should analyze separately the contribution of the query size and the contribution of the database size.

By aiming for fixed-parameter tractability, we can focus more on the core algorithmic question: is there some method for decomposing the space of all solutions in a way that allows efficient evaluation of the query? Some of the progress in this area was made by introducing new decomposition techniques, without showing how to actually find such decompositions. For example, this was the case for the papers introducing query width [Chekuri and Rajaraman 2000] and fractional hypertree width [Grohe and Marx 2012]: it was shown that if a certain type of decomposition is given, then the problem can be solved in polynomial time. In our terminology, these results already show the fixed-parameter tractability of $\text{CSP}(\mathcal{H})$ for the classes \mathcal{H} where such decompositions exist (since the time required to find an appropriate decomposition can be bounded by a function of the hypergraph H only), but do not give polynomial-time algorithms. It took some more time and effort to come up with polynomial-time (approximation) algorithms for finding such decompositions [Gottlob et al. 2002a; Marx 2010a]. While investigating algorithms for finding decompositions

³This assumption is valid only for evaluation problems (where the problem instance includes a large database) and not for problems that involves only queries, such as the Conjunctive Query Containment problem.

give rise to interesting and important problems, they are purely combinatorial problems on graphs and hypergraphs, and no longer has anything to do with query evaluation, constraints, or databases. Thus fixed-parameter tractability gives us a formal way of ignoring these issues and focusing exclusively on the evaluation problem.

On the complexity side, fixed-parameter tractability of $\text{CSP}(\mathcal{H})$ seems to be a more robust question than polynomial-time solvability. For example, any polynomial-time reduction to $\text{CSP}(\mathcal{H})$ should be able to pick a member of \mathcal{H} , thus it seems that polynomial-time reduction to $\text{CSP}(\mathcal{H})$ is only possible if certain artificial technical conditions are imposed on \mathcal{H} (such as there is an algorithm efficiently generating appropriate members of \mathcal{H}). Furthermore, there are classes \mathcal{H} for which $\text{CSP}(\mathcal{H})$ is polynomial-time equivalent to LOG CLIQUE [Grohe 2007], thus we cannot hope to classify $\text{CSP}(\mathcal{H})$ into polynomial-time solvable and NP-hard cases. Another difficulty in understanding polynomial-time solvability is that it can depend on the “irrelevant” parts of the hypergraph. Suppose for example that there is class \mathcal{H} for which $\text{CSP}(\mathcal{H})$ is not polynomial-time solvable, but it is fixed-parameter tractable: it can be solved in time $f(H) \cdot (\|I\|)^{O(1)}$. Let \mathcal{H}' be constructed the following way: for every $H \in \mathcal{H}$, class \mathcal{H}' contains a hypergraph H' that is obtained from H by adding a new component that is a path of length $f(H)$. This new path is trivial with respect to the CSP problem, thus any algorithm for $\text{CSP}(\mathcal{H})$ can be used for $\text{CSP}(\mathcal{H}')$ as well. Consider an instance I of $\text{CSP}(\mathcal{H}')$ having hypergraph H' , which was obtained from hypergraph H . After taking care of the path, the assumed algorithm for $\text{CSP}(\mathcal{H})$ can solve this instance in time $f(H) \cdot (\|I\|)^{O(1)}$, which is polynomial in $\|I\|$: instance I contains a representation of H' , which has at least $f(H)$ vertices, thus $\|I\|$ is at least $f(H)$. Therefore, $\text{CSP}(\mathcal{H}')$ is polynomial-time solvable. This example shows that aiming for polynomial-time solvability instead of fixed-parameter tractability might require understanding such subtle, but mostly irrelevant phenomena.

In the hardness results obtained so far, evidence for the non-existence of polynomial-time algorithms is given not in the form of NP-hardness, but by giving evidence that the problem is not even fixed-parameter tractable. In Theorem 1.1, it is a remarkable coincidence that polynomial-time solvability and fixed-parameter tractability are equivalent. However, there is no reason to expect this to remain true in more general cases. Therefore, as discussed above, it makes sense to focus first on understanding the fixed-parameter tractability of the problem.

Organization. For convenience, Section 2 collects many of the definitions appearing in the paper. The reader might want to skim through this at first and refer to appropriate parts of it later. Submodular width and other width measures are defined in Section 3. Section 4 contains the algorithmic part of the paper: the algorithm for classes with bounded submodular width. Section 5 characterizes large submodular width with highly connected sets, while Section 6 uses highly connected sets to find good embeddings in hypergraph. The main hardness result of the paper is proved in Section 7.

2. PRELIMINARIES

Constraint satisfaction problems. We briefly recall the most important notions related to CSP. For more background, see e.g., [Grohe 2006; Feder and Vardi 1999].

DEFINITION 2.1. *An instance of a constraint satisfaction problem is a triple (V, D, C) , where:*

- V is a set of variables,
- D is a domain of values,
- C is a set of constraints, $\{c_1, c_2, \dots, c_q\}$. Each constraint $c_i \in C$ is a pair $\langle s_i, R_i \rangle$, where:
 - s_i is a tuple of variables of length m_i , called the constraint scope, and
 - R_i is an m_i -ary relation over D , called the constraint relation.

For each constraint $\langle s_i, R_i \rangle$ the tuples of R_i indicate the allowed combinations of simultaneous values for the variables in s_i . The length m_i of the tuple s_i is called the *arity* of the constraint. A *solution* to a constraint satisfaction problem instance is a function f from the set of variables V to the domain of values D such that for each constraint $\langle s_i, R_i \rangle$ with $s_i = \langle v_{i_1}, v_{i_2}, \dots, v_{i_m} \rangle$, the tuple $\langle f(v_{i_1}), f(v_{i_2}), \dots, f(v_{i_m}) \rangle$ is a member of R_i . We say that an instance is *binary* if each constraint relation is binary, i.e., $m_i = 2$ for each constraint.⁴ It can be assumed that the instance does not contain two constraints $\langle s_i, R_i \rangle, \langle s_j, R_j \rangle$ with $s_i = s_j$, since in this case the two constraints can be replaced by the constraint $\langle s_i, R_i \cap R_j \rangle$.

In the input, the relation appearing in a constraint is represented by listing all the tuples of the constraint. We denote by $\|I\|$ the size of the representation of the instance $I = (V, D, C)$. It can be assumed that $|D| \leq \|I\|$: elements of D that do not appear in any relation can be safely removed.

Let $I = (V, D, C)$ be a CSP instance and let $V' \subseteq V$ be a nonempty subset of variables. If f is a solution of I , then $\text{pr}_{V'} f$ is the *projection* of f to V' , which is simply the restriction of the function $f : V \rightarrow D$ to $V' \subseteq V$. If R is a set of solutions for I , then we let $\text{pr}_{V'} R = \{\text{pr}_{V'} f \mid f \in R\}$.

The *projection* $\text{pr}_{V'} I$ of I to V' is a CSP $I' = (V', D, C')$, where C' is defined the following way: For each constraint $c = \langle (v_1, \dots, v_k), R \rangle$ having at least one variable in V' , there is a corresponding constraint c' in C' . Suppose that $v_{i_1}, \dots, v_{i_\ell}$ are the variables among v_1, \dots, v_k that are in V' . Then the constraint c' is defined as $\langle (v_{i_1}, \dots, v_{i_\ell}), R' \rangle$, where the relation R' is the projection of R to the coordinates i_1, \dots, i_ℓ , that is, R' contains an ℓ -tuple $\langle d'_1, \dots, d'_\ell \rangle \in D^\ell$ if and only if there is a k -tuple $\langle d_1, \dots, d_k \rangle \in R$ such that $d'_j = d_{i_j}$ for $1 \leq j \leq \ell$. Clearly, if f is a solution of I , then $\text{pr}_{V'} f$ is a solution of $\text{pr}_{V'} I$ (but the converse is not true). For a subset $V' \subseteq V$, we denote by $\text{sol}_I(V')$ the set of all solutions of $\text{pr}_{V'} I$ (which can contain a solution which is not the projection of any solution of I). If the instance I is clear from the context, we drop the subscript.

The *primal graph* (or *Gaifman graph*) of a CSP instance $I = (V, D, C)$ is a graph with vertex set V such that $u, v \in V$ are adjacent if and only if there is a constraint whose scope contains both u and v . The *hypergraph* of a CSP instance $I = (V, D, C)$ is a hypergraph H with vertex set V , where $e \subseteq V$ is an edge of H if and only if there is a constraint whose scope is e (more precisely, where the scope is an $|e|$ -tuple s , whose coordinates form a permutation of the elements of e). For a class \mathcal{H} of graphs, we denote by $\text{CSP}(\mathcal{H})$ the problem restricted to instances whose hypergraph is in \mathcal{H} .

Graphs and hypergraphs. If G is a graph or hypergraph, then we denote by $V(G)$ and $E(G)$ the set of vertices and the set of edges of G , respectively. Vertices $u, v \in V(G)$ are *adjacent* if there is an edge $e \in E(G)$ with $u, v \in e$. A set $K \subseteq V(G)$ is a *clique* if the vertices in K are pairwise adjacent. If H is a hypergraph and $V' \subseteq V(H)$, then the *subhypergraph induced by V'* is a hypergraph H' with vertex set S and $\emptyset \subset e' \subseteq V'$ is an edge of H' if and only if there is an edge $e \in E(H)$ with $e \cap V' = e'$. We denote by $H \setminus S$ the subhypergraph of H induced by $V(H) \setminus S$.

Paths, separators, and flows in hypergraphs. A *path* P in hypergraph H is an ordered sequence v_0, v_1, \dots, v_r of distinct vertices such that v_i and v_{i-1} are adjacent for every $1 \leq i < r$. We distinguish the endpoints of a path: vertex v_0 is the *first endpoint* of P and v_r is the *second endpoint* of P . For a path of length zero, the first and second endpoints coincide. A path is an $X - Y$ *path* if its first endpoint is in X and its second endpoint is in Y . A path $P = v_1 v_2 \dots v_t$ is *minimal* if there are no shortcuts, i.e., v_i and v_j are not adjacent if $|i - j| > 1$. Note that a minimal path intersects each edge at most twice.

Let H be a hypergraph and $X, Y \subseteq V(H)$ be two (not necessarily disjoint) sets of vertices. An (X, Y) -*separator* is a set $S \subseteq V(H)$ of vertices such that there is no $(X \setminus S) - (Y \setminus S)$

⁴It is unfortunate that while some communities use the term “binary CSP” in the sense that each constraint is binary (as does this paper), others use it in the sense that the variables are 0-1, i.e., the domain size is 2.

path in $H \setminus S$, or in other words, every $X - Y$ path of H contains at least one vertex of S . In particular, this means that $X \cap Y \subseteq S$.

An assignment $s : E(H) \rightarrow \mathbb{R}^+$ is a *fractional (X, Y) -separator* if every $X - Y$ path P is covered by s , that is, $\sum_{e \in E(H), e \cap P \neq \emptyset} s(e) \geq 1$. The *weight* of the fractional separator s is $\sum_{e \in E(H)} s(e)$.

Let H be a hypergraph and let \mathcal{P} be the set of all paths in H . A *flow* of H is an assignment $f : \mathcal{P} \rightarrow \mathbb{R}^+$ such that $\sum_{P \in \mathcal{P}, P \cap e \neq \emptyset} f(P) \leq 1$ for every $e \in E(H)$. The *value* of the flow f is $\sum_{P \in \mathcal{P}} f(P)$. We say that a path P *appears* in flow f , or simply P is a *path of f* if $f(P) > 0$. For some $X, Y \subseteq V(H)$, an *(X, Y) -flow* is a flow f such that only $X - Y$ paths appear in f . A standard LP duality argument shows that the minimum weight of a fractional (X, Y) -separator is equal to the maximum value of an (X, Y) -flow.

If f, f' are flows such that $f'(P) \leq f(P)$ for every path P , then f' is a *subflow* of f . The *sum* of the flows f_1, \dots, f_r is a mapping that assigns weight $\sum_{i=1}^r f_i(P)$ to each path P . Note that the sum of flows is not necessarily a flow itself as the total weight of the paths intersecting a certain edge can be more than 1 in the sum. If the sum of f_1, \dots, f_r happens to be a flow, then we say that f_1, \dots, f_r are *compatible*.

Highly connected sets. An important step in understanding various width measures is showing that if the measure is large, then the (hyper)graph contains a highly connected set (in a certain sense). We define here the notion of highly connectedness that will be used in the paper. First, recall that a *fractional independent set* of a hypergraph H is a mapping $\mu : V(H) \rightarrow [0, 1]$ such that $\sum_{v \in e} \mu(v) \leq 1$ for every $e \in E(H)$. We extend functions on the vertices of H to subsets of vertices of H the natural way by setting $\mu(X) := \sum_{v \in X} \mu(v)$, thus we can equivalently say that $\mu : V(H) \rightarrow [0, 1]$ is a fractional independent set if and only if $\mu(e) \leq 1$ for every $e \in E(H)$.

Let μ be a fractional independent set of hypergraph H and let $\lambda > 0$ be a constant. We say that a set $W \subseteq V(H)$ is *(μ, λ) -connected* if for any two disjoint sets $A, B \subseteq W$, the minimum weight of a fractional (A, B) -separator is at least $\lambda \cdot \min\{\mu(A), \mu(B)\}$. Note that if W is (μ, λ) -connected, then W is (μ, λ') -connected for every $\lambda' < \lambda$ and every $W' \subseteq W$ is also (μ, λ) -connected. Informally, if W is (μ, λ) -lambda connected for some fractional independent set μ such that $\mu(W)$ is “large”, then we call W a highly connected set. For $\lambda > 0$, we denote by $\text{con}_\lambda(H)$ the maximum of $\mu(W)$, taken over every fractional independent set μ and (μ, λ) -connected set W of H . Note that if $\lambda' \leq \lambda$, then $\text{con}_{\lambda'}(H) \geq \text{con}_\lambda(H)$. Throughout the paper, λ can be thought of as a sufficiently small universal constant, say, 0.001.

Embeddings. The hardness result presented in the paper and earlier hardness results for CSP(\mathcal{H}) [Grohe 2007; Marx 2011; 2010b] are based on embedding some other problem (with a certain graph structure) in a CSP instance whose hypergraph is a member of \mathcal{H} . Thus we need appropriate notions of embedding a graph in a (hyper)graph. Let us first recall the definition of minors in graphs. A graph F is a *minor* of G if F can be obtained from G by a sequence of vertex deletions, edge deletions, and edge contractions. The following alternative definition is more relevant from the viewpoint of embeddings: a graph F is a minor of G if there is a mapping ψ that maps each vertex of F to a connected subset of $V(G)$ such that $\psi(u) \cap \psi(v) = \emptyset$ for $u \neq v$, and if $u, v \in V(F)$ are adjacent in F , then there is an edge in $E(G)$ connecting $\psi(u)$ and $\psi(v)$.

A crucial difference between the proof of Theorem 1.1 in [Grohe 2007] and the proof of Theorem 1.2 in [Marx 2010b] is that the former result is based on finding a minor embedding of a grid, while the latter result uses a more general notion of embedding where the images of distinct vertices are not necessarily disjoint, but can overlap in a controlled way. We define such embeddings the following way. We say that two sets of vertices $X, Y \subseteq V(H)$ *touch* if either $X \cap Y \neq \emptyset$, or there is an edge $e \in E(H)$ intersecting both X and Y . An *embedding* of graph G into hypergraph H is a mapping ψ that maps each vertex of G

to a connected subset of $V(H)$ such that if u and v are adjacent in G , then $\psi(u)$ and $\psi(v)$ touch. The *depth* of a vertex $v \in V(H)$ in embedding ψ is $d_\psi(v) := |\{u \in V(G) \mid v \in \psi(u)\}|$, the number of vertices of G whose images contain v . The *vertex depth* of the embedding is $\max_{v \in V(H)} d_\psi(v)$. Observe that ψ is a minor mapping if and only if it has vertex depth 1. Because in our case we want to control the size of the constraint relations, we need a notion of depth that is sensitive to “what the edges see.” We define the depth $d_\psi(e)$ of an edge $e \in E(H)$ as $d_\psi(e) = \sum_{v \in e} d_\psi(v)$ and the *edge depth* to be the maximum of e taken over all edges $e \in E(H)$. Equivalently, we can define the depth of an edge $e \in H$ as $d_\psi(e) = \sum_{v \in V(G)} |\psi(v) \cap e|$, that is, each vertex v contributes $|\psi(v) \cap e|$ to the depth. (A different, perhaps more natural, definition of edge depth would be to define it simply as a maximum number of sets $\psi(v)$ that intersect an edge. Somewhat unexpectedly, most results of the paper remain true with both notions; see Remarks 7.6–7.7.)

Trivially, for any graph G and hypergraph H , there is an embedding of G into H having vertex depth and edge depth at most $|V(G)|$. If G has m edges and no isolated vertices, then $|V(G)|$ is at most $2m$. We are interested in how much we can gain compared to this trivial solution of depth $O(m)$. We define the *embedding power* $\text{emb}(H)$ to be the maximum (supremum) value of α for which there is an integer m_α such that every graph G with $m \geq m_\alpha$ edges has an embedding into H with edge depth m/α . It might look unmotivated that we define embedding power in terms of the number of edges of G : defining it in terms of the number of vertices might look more natural. However, if we replace the number m of edges with the number n of vertices in the definition, then the worst case occurs if G is a clique on n vertices. Such a definition would describe how well cliques can be embedded, and would give us no information about how sparse graphs can be embedded.

3. WIDTH PARAMETERS

Treewidth and its various generalizations are defined in this section. We follow the framework of width functions introduced by Adler [2006]. A *tree decomposition* of a hypergraph H is a tuple $(T, (B_t)_{t \in V(T)})$, where T is a tree and $(B_t)_{t \in V(T)}$ is a family of subsets of $V(H)$ satisfying the following two conditions: (1) for each $e \in E(H)$ there is a node $t \in V(T)$ such that $e \subseteq B_t$, and (2) for each $v \in V(H)$ the set $\{t \in V(T) \mid v \in B_t\}$ is connected in T . The sets B_t are called the *bags* of the decomposition. Let $f : 2^{V(H)} \rightarrow \mathbb{R}^+$ be a function that assigns a nonnegative real number to each nonempty subset of vertices. The *f-width* of a tree-decomposition $(T, (B_t)_{t \in V(T)})$ is $\max\{f(B_t) \mid t \in V(T)\}$. The *f-width* of a hypergraph H is the minimum of the *f-widths* of all its tree decompositions. In other words, $f\text{-width}(H) \leq w$ if and only if there is a tree decomposition of H where $f(B) \leq w$ for every bag B .

The main idea of tree decomposition based algorithms is that if we have a tree decomposition for instance I such that at most C assignments on B_t have to be considered for each bag B_t , then the problem can be solved by dynamic programming in time polynomial in C and $\|I\|$. The various width notions try to guarantee the existence of such decompositions. The simplest such notion, treewidth, can be defined as follows:

DEFINITION 3.1. *Let $s(B) = |B| - 1$. The treewidth of H is $\text{tw}(H) := s\text{-width}(H)$.*

Further width notions defined in the literature can also be conveniently defined using this setup. A subset $E' \subseteq E(H)$ is an *edge cover* if $\bigcup_{e \in E'} e = V(H)$. The *edge cover number* $\rho(H)$ is the size of the smallest edge cover (here we assume that H has no isolated vertices). For $X \subseteq V(H)$, let $\rho_H(X)$ be the size of the smallest set of edges covering X .

DEFINITION 3.2. *The generalized hypertree width of H is $\text{hw}(H) := \rho_H\text{-width}(H)$.*

The original (nongeneralized) definition [Gottlob et al. 2002a] of hypertree width includes an additional requirement on the decomposition (we omit the details), thus it cannot be

less than generalized hypertree. However, it is known that hypertree width and generalized hypertree width can differ by at most a constant factor [Adler et al. 2007].

Grohe and Marx [2012] further generalized hypertree width by considering linear relaxations of edge covers. A function $\gamma : E(H) \rightarrow [0, 1]$ is a *fractional edge cover* of H if $\sum_{e \in E(H): v \in e} \gamma(e) \geq 1$ for every $v \in V(H)$. The *fractional cover number* $\rho^*(H)$ of H is the minimum of $\sum_{e \in E(H)} \gamma(e)$ taken over all fractional edge covers of H (it is well known that this minimum is achieved by some rational γ). We define $\rho_H^*(X)$ analogously to $\rho_H(X)$: the requirement $\sum_{e: v \in e} \gamma(e) \geq 1$ is restricted to vertices $v \in X$.

DEFINITION 3.3. *The fractional hypertree width of H is $\text{fhw}(H) := \rho_H^*\text{-width}(H)$.*

A crucial idea in [Marx 2011] is to make the choice of tree decomposition adaptive: instead of assigning a single decomposition to each hypergraph, we choose the best decomposition based on additional properties of the current instance. Motivated by this idea, we generalize the notion of f -width from a single function f to a class of functions \mathcal{F} . Let H be a hypergraph and let \mathcal{F} be an arbitrary (possibly infinite) class of functions that assign nonnegative real numbers to nonempty subsets of vertices of H . The \mathcal{F} -width of H is $\mathcal{F}\text{-width}(H) := \sup \{f\text{-width}(H) \mid f \in \mathcal{F}\}$. Thus if $\mathcal{F}\text{-width}(H) \leq k$, then for every $f \in \mathcal{F}$, hypergraph H has a tree decomposition with f -width at most k . Note that this tree decomposition can be different for the different functions f . For normalization purposes, we consider only functions f on $V(H)$ that satisfy $f(\emptyset) = 0$ and that are *edge-dominated*, that is, $f(e) \leq 1$ holds for every $e \in E(H)$.

Using these definitions, we can define adaptive width, introduced in [Marx 2011], as follows. Recall that in Section 2, we stated that if μ is a fractional independent set, then μ is extended to subsets of vertices by defining $\mu(X) := \sum_{v \in X} \mu(v)$ for every $X \subseteq V(H)$.

DEFINITION 3.4. *The adaptive width $\text{adw}(H)$ of a hypergraph H is $\mathcal{F}\text{-width}(H)$, where \mathcal{F} is the set of all fractional independent sets of H .*

A function $f : 2^{V(H)} \rightarrow \mathbb{R}$ is *modular* if $f(X) = \sum_{v \in X} c_v$ for some constants c_v ($v \in V(H)$). The function $\mu(X)$ arising from a fractional independent set is clearly a modular and edge dominated function, in fact, in Definition 3.4 we can equivalently define \mathcal{F} as the set of all nonnegative modular edge-dominated functions on $V(H)$. The main new definition of the paper is a new width measure, which is obtained by imposing a requirement weaker than modularity on the functions in \mathcal{F} (hence the considered set \mathcal{F} of functions is larger):

DEFINITION 3.5. *A function $b : 2^{V(H)} \rightarrow \mathbb{R}^+$ is submodular if $b(X) + b(Y) \geq b(X \cap Y) + b(X \cup Y)$ holds for every $X, Y \subseteq V(H)$. Given a hypergraph H , let \mathcal{F} contain every edge-dominated monotone submodular function b on $V(H)$ with $b(\emptyset) = 0$. The submodular width of hypergraph H is $\text{subw}(H) := \mathcal{F}\text{-width}(H)$.*

It is well-known that submodular functions can be equivalently characterized by the property that $b(X \cup v) - b(X)$, the *marginal value* of v with respect to X , is a nonincreasing function of X . That is, for every v and $X \subseteq Y$,

$$b(X \cup v) - b(X) \geq b(Y \cup v) - b(Y). \quad (1)$$

It is clear that $\text{subw}(H) \geq \text{adw}(H)$: Definition 3.5 considers a larger set of functions than Definition 3.4. Furthermore, we show that $\text{subw}(H)$ is at most the fractional hypertree width $\text{fhw}(H)$. This is a straightforward consequence of the fact that an edge-dominated submodular function is always bounded by the fractional cover number:

LEMMA 3.6. *Let H be a hypergraph and b be a monotone edge-dominated submodular function with $b(\emptyset) = 0$. Then $b(S) \leq \rho_H^*(S)$ for every $S \subseteq V(H)$.*

PROOF. The statement can be proved along the same lines as the proof of Shearer’s Lemma [Chung et al. 1986] attributed to Radhakrishnan goes. It is sufficient to prove the statement for the case $S = V(H)$: otherwise, we can consider the subhypergraph of H induced by S and the function b restricted to S . Let $\gamma : E(H) \rightarrow \mathbb{R}^+$ be a minimum fractional edge cover of S . Let v_1, \dots, v_n be an arbitrary ordering of $V(H)$ and let $V_i = \{v_1, \dots, v_i\}$, $V_0 = \emptyset$. For every $e \in E(H)$, we have

$$b(e) = \sum_{v_i \in e} (b(e \cap V_i) - b(e \cap V_{i-1})) \geq \sum_{v_i \in e} (b(V_i) - b(V_{i-1}))$$

(the equality is a simple telescopic sum; the inequality uses (1), i.e., the marginal value of v_i with respect to V_{i-1} is not greater than with respect to $e \cap V_{i-1}$).

$$\begin{aligned} \rho_H^*(V(H)) &= \sum_{e \in E(H)} \gamma(e) \geq \sum_{e \in E(H)} \gamma(e)b(e) \geq \sum_{e \in E(H)} \gamma(e) \sum_{v_i \in e} (b(V_i) - b(V_{i-1})) \\ &= \sum_{i=1}^n \left((b(V_i) - b(V_{i-1})) \sum_{e \in E(H), v_i \in e} \gamma(e) \right) \geq \sum_{i=1}^n (b(V_i) - b(V_{i-1})) = b(V(H)) \end{aligned}$$

(in the first inequality, we use that f is edge dominated; in the last inequality, we use that γ is a fractional edge cover). \square

PROPOSITION 3.7. *For every hypergraph H , $\text{subw}(H) \leq \text{fhw}(H)$.*

PROOF. Let $(T, B_{t \in V(T)})$ be a tree decomposition of H whose ρ_H^* -width is $\text{fhw}(H)$. If b is an edge-bounded monotone submodular function with $b(\emptyset) = 0$, then by Lemma 3.6, $b(B_t) \leq \rho_H^*(B_t) \leq \text{fhw}(H)$ for every bag B_t of the decomposition, i.e., b -width(H) \leq $\text{fhw}(H)$. This is true for every such function b , hence $\text{subw}(H) \leq \text{fhw}(H)$. \square

Since $\text{adw}(H) \leq \text{subw}(H) \leq \text{fhw}(H)$, if a class \mathcal{H} of hypergraphs has bounded fractional hypertree width, then it has bounded submodular width, and if a class \mathcal{H} has bounded submodular width, then it has bounded adaptive width. Surprisingly, it turns out that the latter implication is actually an equivalence: Corollary 6.10 shows that $\text{subw}(H)$ is at most $O(\text{adw}(H)^4)$, thus a class of hypergraphs has bounded submodular width if and only if it has bounded adaptive width. In other words, large submodular width can be certified already by modular functions: if submodular width is unbounded in \mathcal{H} and we want to choose an $H \in \mathcal{H}$ and a submodular function b such that the b -width of H is larger than some constant k , then we can choose H and b such that b is actually modular. There is no intuitive reason why this is true: submodular functions seem to be much more powerful than modular functions. Still, we obtain this result as a byproduct of our characterization of submodular width.

There is no such connection between adaptive width and fractional hypertree width: it is shown in [Marx 2011] that there is a class of hypergraphs with bounded adaptive width and unbounded fractional hypertree width. Thus the property bounded fractional hypertree width is a strictly weaker property than bounded adaptive/submodular width.

Figure 2 shows the relations of the hypergraph properties defined in this section (note that the elements of this Venn diagram are sets of hypergraphs; e.g., the set “bounded treewidth” contains every set \mathcal{H} of hypergraphs with bounded treewidth). As discussed above, all the inclusions in the figure are proper.

Finally, let us remark that there have been investigations of tree decompositions and branch decompositions of submodular functions and matroids in the literature [Hliněný and Oum 2008; Oum and Seymour 2007; Hliněný and Whittle 2006; Hliněný 2005; Amini et al. 2009]. However, in those results the submodular function is a connectivity function: $b(S)$ describes the boundary of S , that is, the cost of separating S from its complement. In




Fig. 2. Hypergraph properties that make CSP fixed-parameter tractable.

our case, $b(S)$ describes the cost of the separator S itself. Therefore, we are in a completely different setting and the previous results cannot be used.

4. FROM CSP INSTANCES TO SUBMODULAR FUNCTIONS

In this section, we prove the main algorithmic result of the paper: $\text{CSP}(\mathcal{H})$ is fixed-parameter tractable if \mathcal{H} has bounded submodular width.

THEOREM 4.1. *Let \mathcal{H} be a class of hypergraphs such that $\text{subw}(H) \leq c_0$ for every $H \in \mathcal{H}$. Then $\text{CSP}(\mathcal{H})$ can be solved in time $2^{c_0} \cdot 2^{O(|V(H)|)} \cdot \|I\|^{O(c_0)}$.*

The proof of Theorem 4.1 is based on two main ideas:

(1) A CSP instance I can be decomposed into a bounded number of “uniform” CSP instances I_1, \dots, I_t (Lemma 4.11). Here uniform means that if $B \subseteq A$ are two sets of variables, then every solution of $\text{pr}_B I_j$ has roughly the same number of extensions to $\text{pr}_A I_j$.

(2) If I is a uniform CSP instance, then (the logarithm of) the number of solutions on the different projections of I can be described by an edge-dominated monotone submodular function b (Lemma 4.12). Therefore, if the hypergraph H of I has bounded submodular width, then it follows that there is a tree decomposition where every bag has a polynomially bounded number of solutions. This means that the existence of a solution can be tested by standard techniques.

While our algorithm is based on these two ideas, the technical implementation is slightly different. First, we can achieve uniformity only on “small sets” of variables. For technical reasons, we have to ensure a certain consistency condition (for example, in order to ensure that the submodular function b is monotone). It follows from the consistency condition that when we find a tree decomposition for a uniform instance such that every bag has a small number of solutions, then this automatically implies that the instance has a solution; we do not even have to use the tree decomposition (see Lemma 4.7).

In Section 4.1 we define the notion of consistency that we use and discuss how it can be achieved. Section 4.2 describes how the instance can be partitioned into uniform instances.

Section 4.3 shows how a submodular function can be defined based on a uniform instance, connecting our algorithm to submodular width.

4.1. Consistency

Recall from Section 2 that $\text{pr}_A I$ is instance I projected to a set A of variables and $\text{sol}_I(A)$ is the set of all solutions of $\text{pr}_A I$. Note that, in general, it is possible that $|\text{sol}_I(S')| > |\text{sol}_I(S)|$ for some $S' \subset S$. In the implementation of the first idea (Lemma 4.11), we guarantee uniformity only to subsets of variables that are “small” in the following hereditary sense

DEFINITION 4.2. *Let I be a CSP instance and $M \geq 1$ an integer. We say that $S \subseteq V$ is M -small if $|\text{sol}_I(S')| \leq M$ for every $S' \subseteq S$.*

It is not difficult to find all the M -small sets, and every solution of the instances projected onto these sets:

LEMMA 4.3. *Let $I = (V, D, C)$ be a CSP instance and $M \geq 1$ an integer. There is an algorithm with running time $2^{O(|V|)} \cdot \text{poly}(\|I\|, M)$ that finds the set \mathcal{S} of all M -small sets $S \subseteq V$ and constructs $\text{sol}_I(S)$ for each such $S \in \mathcal{S}$.*

PROOF. For $i = 1, 2, \dots, |V|$, we find every M -small set S of size i and construct $\text{sol}_I(S)$. This is trivial to do for $i = 1$. Suppose that we have already found the collection \mathcal{S}_i of all M -small sets of size exactly i . By definition, every size i subset S of an M -small set S' of size $i + 1$ is an M -small set. Thus we can find every M -small set of size $i + 1$ by enumerating every $S \in \mathcal{S}_i$ and checking for every $v \in V \setminus S$ whether $S' := S \cup \{v\}$ is M -small. To check whether S' is M -small, we first check whether every subset of size i is M -small, which is easy to do using the set \mathcal{S}_i . Then we construct $\text{sol}_I(S')$: this can be done by enumerating every tuple $s \in \text{sol}_I(S)$ and every extension of s by a new value from D . Thus we need to consider at most $|\text{sol}_I(S)| \cdot |D| \leq M \cdot |D|$ tuples as possible members in $\text{sol}_I(S')$, which means that $\text{sol}_I(S')$ can be constructed in time polynomial in M and $\|I\|$. If $|\text{sol}_I(S')| \leq M$, then we put S' into \mathcal{S}_{i+1} . As the collection \mathcal{S}_i contains at most $2^{|V|}$ sets and every operation is polynomial in M and $\|I\|$, the total running time is $2^{O(|V|)} \cdot \text{poly}(\|I\|, M)$. \square

We want to avoid dealing with assignments $b \in \text{sol}(B)$ that cannot be extended to any member of $\text{sol}(A)$ for some $A \supseteq B$. Of course, there is no easy way to avoid this in general (or even to detect if there is such a b): for example, if A is the set of all variables, then we would need to check if b can be extended to a solution. Therefore, we require only that there is no such unextendable b if A and B are M -small:

DEFINITION 4.4. *A CSP instance I is M -consistent if $\text{sol}_I(B) = \text{pr}_B \text{sol}_I(A)$ for all M -small sets $B \subseteq A$.*

The notion of M -consistency is very similar to k -consistency, a standard notion in the constraint satisfaction literature [Atserias et al. 2007; Dalmau et al. 2002; Kolaitis and Vardi 2000b]. However, we restrict the considered subsets not by the number of variables, but by the number of solutions (more precisely, by considering only M -small sets). Thus the notion of M -consistency could be interpreted in the framework of Greco and Scarcello [2010], where consistency is defined with respect to an arbitrary set of views.

Similarly to usual k -consistency, we can achieve M -consistency by throwing away partial solutions that violate the requirements: if we use the algorithm of Lemma 4.3 to find all possible assignments of the M -small sets, then we can check if there is such an unextendable b for some M -small sets A and B . If there is such a b , then we can exclude it from consideration (without losing any solution of the instance) by introducing a new constraint on B . By repeatedly excluding the unextendable assignments, we can avoid all such problems. We say that $I' = (V, D, C')$ is a *refinement* of $I = (V, D, C)$ if for every constraint $\langle s, R \rangle \in C$, there is a constraint $\langle s, R' \rangle \in C'$ such that $R' \subseteq R$.

LEMMA 4.5. *Let $I = (V, D, C)$ be a CSP instance and $M \geq 1$ an integer. There is an algorithm with running time $2^{O(|V|)} \cdot \text{poly}(\|I\|, M)$ that produces an M -consistent CSP instance I' that is a refinement of I with $\text{sol}(I) = \text{sol}(I')$.*

PROOF. Using the algorithm of Lemma 4.3, we can find all the M -small sets and then we can easily check if there are two M -small sets $S \subseteq S'$ violating consistency, i.e., $\text{sol}(S) \not\subseteq \text{pr}_S \text{sol}(S')$. In this case, let s be an $|S|$ -tuple whose coordinates contain S in an arbitrary order and let us add the constraint $\langle s, \text{pr}_S \text{sol}(S') \rangle$; it is clear that $\text{sol}(V)$ does not change but $|\text{sol}(S)|$ strictly decreases. We repeat this step until the instance becomes M -consistent. Note that adding the new constraint can make a set M -small that was not M -small before, thus we need to rerun the algorithm of Lemma 4.3. To bound the number of iterations before M -consistency is reached, observe that adding a new constraint does not increase $|\text{sol}(A)|$ for any A and strictly decreases $|\text{sol}(S)|$ for some M -small set S . As there are at most $2^{|V|}$ sets S and $|\text{sol}(S)| \leq M$ for every M -small set S , it follows that this step can be repeated at most $2^{|V|} \cdot M$ times. The size of the instance increases in each step by adding a new constraint with at most M tuples, thus the size of the instance at the end of the process can be still bounded by $2^{O(|V|)} \cdot \text{poly}(\|I\|, M)$. Thus the total time required to ensure that instance I is M -consistent can be bounded by $2^{O(|V|)} \cdot \text{poly}(\|I\|, M)$. \square

We want to avoid degenerate cases where there is no solution even for some M -small sets. Consistency implies that it is sufficient to require this for sets of size 1. We say that a CSP instance is *nontrivial* if $\text{sol}(\{v\}) \neq \emptyset$ for every $v \in V$. The following is immediate:

PROPOSITION 4.6. *If I is an M -consistent nontrivial CSP instance, then $\text{sol}(S) \neq \emptyset$ for every M -small set S .*

It is well known that by achieving k -consistency, we can solve CSP instances with treewidth k : the key observation is that if an instance I with treewidth at most k has a k -consistent nontrivial refinement I' , then I has a solution. The following lemma adapts this statement to our setting.

LEMMA 4.7. *Let $I = (V, D, C)$ be a CSP instance and $M \geq 1$ an integer. Let I' be an M -consistent nontrivial refinement of I . If the hypergraph H of I has a tree decomposition where every bag B is M -small in I' , then I has a solution.*

PROOF. Suppose that there is such a tree decomposition $(T, (B_t)_{t \in V(T)})$. Assume that T is rooted and for every node $t \in V(T)$, let V_t be the union of the bags that are descendants of t (including B_t). We claim that every assignment in $\text{sol}_{I'}(B_t)$ can be extended to an assignment of V_t that satisfies every constraint of I whose scope is fully contained in V_t . Applying this statement to the root of T proves that there exists a solution for I . (Recall that every edge of the hypergraph H , and hence the scope of every constraint, is fully contained in one of the bags.)

We prove the claim for every node of T in a bottom up order. The statement is trivial for the leaves. Let t_1, \dots, t_ℓ be the children of t and suppose the claim is true for these nodes. Consider an assignment $g \in \text{sol}_{I'}(B_t)$. Since I' is M -consistent and B_{t_i} is M -small, assignment $g|_{B_t \cap B_{t_i}}$ can be extended to an assignment $g_i \in \text{sol}_{I'}(B_{t_i})$. As the claim is true for node t_i , assignment g_i can be extended to an assignment g'_i of V_{t_i} . The assignments g, g'_1, \dots, g'_ℓ can be combined to obtain an assignment g' on V_t (note that this is well defined: the intersection of V_{t_i} and V_{t_j} is in V_t , which means that a variable appearing in both V_{t_i} and V_{t_j} has the same value in g, g'_i , and g'_j). Furthermore, every edge e of H that is fully contained in V_t is fully contained in at least one of $B_t, V_{t_1}, \dots, V_{t_\ell}$, and the corresponding assignment among g, g'_1, \dots, g'_ℓ shows that g' satisfies the constraint corresponding to e . \square

Note the subtle detail that Lemma 4.7 does not claim that I' has a solution. Furthermore, when Lemma 4.5 creates an M -consistent instance, then it possibly adds many new constraints and the hypergraph of I' can be very dense even if the hypergraph of I has nice structure. However, this is not a problem, as Lemma 4.7 does not require any property on the hypergraph of I' . Note also that every M -small set in I is M -small in I' as well, thus I' has a potentially larger collection of M -small sets, which makes finding the required tree decomposition of the hypergraph H of I easier.

4.2. Decomposition into uniform CSP instances

Our algorithm for decomposing a CSP instance into uniform CSP instances is inspired by a combinatorial result of Alon et al. [2007], which shows that, for every fixed n , an n -dimensional point set S can be partitioned into $\text{polylog}(|S|)$ classes such that each class is $O(1)$ -uniform. We follow the same proof idea: the instance is split into two instances if uniformity is violated somewhere, and we analyze the change of an appropriately defined weight function to bound the number of splits performed. However, the parameter setting is different in our proof: we want to partition into $f(|V|)$ classes, but we are satisfied with somewhat weaker uniformity. Another minor technical difference is that we require uniformity only on the N^c -small sets.

The following definitions gives the precise notion of uniformity that we use:

DEFINITION 4.8. *Let $I = (V, D, C)$ be a CSP instance. For $B \subseteq A \subseteq V$ and an assignment $b : B \rightarrow D$, let $\text{sol}_I(A|B = b) := \{a \in \text{sol}_I(A) \mid \text{pr}_B a = \text{pr}_B b\}$, the set of all extensions of b to a solution of $\text{pr}_A I$. Let $\max_I(A|B) = \max_{b \in \text{sol}_I(B)} |\text{sol}_I(A|B = b)|$ (if $\text{sol}_I(B) = \emptyset$, then $\max_I(A|B) = 0$). We define $\max_I(A|\emptyset) = |\text{sol}_I(A)|$ and $\max_I(\emptyset|\emptyset) = 1$. We will drop I from the subscript of \max if it is clear from the context.*

Let us prove two straightforward properties of the function $\max(A|B)$:

PROPOSITION 4.9. *For every $B \subseteq A \subseteq V$ with $\text{sol}(A) \neq \emptyset$ and $C \subseteq V$, we have*

- (1) $\max(A|B) \geq |\text{sol}(A)|/|\text{sol}(B)|$,
- (2) $\max(A|B) \geq \max(A \cup C|B \cup C)$.

PROOF. If every $b \in \text{sol}(B)$ has at most $\max(A|B)$ extensions to A , then clearly $|\text{sol}(A)|$ is at most $|\text{sol}(B)| \cdot \max(A|B)$, proving the first statement. To show the second statement, consider an $x \in \text{sol}(B \cup C)$ with $\max(A \cup C|B \cup C)$ extensions to $A \cup C$. For any two $y_1, y_2 \in \text{sol}(A \cup C|B \cup C = x)$ with $y_1 \neq y_2$, we have $\text{pr}_C y_1 = \text{pr}_C y_2 = \text{pr}_C x$, hence y_1 and y_2 can be different only if $\text{pr}_A y_1 \neq \text{pr}_A y_2$. This means that $\text{pr}_A y_1$ and $\text{pr}_A y_2$ are two different extensions of $\text{pr}_B x$ to A . Therefore,

$$\max(A|B) \geq |\text{sol}(A|B = \text{pr}_B x)| \geq |\text{sol}(A \cup C|B \cup C = x)| = \max(A \cup C|B \cup C),$$

what we had to show. \square

Notice that (2) in Prop. 4.9 gives a hint that submodularity will be relevant: it is analogous to inequality (1) (Section 3) expressing that marginal value is larger with respect to a smaller set.

DEFINITION 4.10. *We say that $A \subseteq V$ is c -uniform (for some integer c) if $\text{sol}(A) \neq \emptyset$ and, for every $B \subseteq A$,*

$$\max_I(A|B) \leq c |\text{sol}_I(A)|/|\text{sol}_I(B)|.$$

A CSP instance is (N, c, ϵ) -uniform if every N^c -small set is N^ϵ -uniform.

That is, A is c -uniform if every solution of $\text{sol}_I(B)$ has at most c times as many extensions as the average number of extensions. The following lemma states the main combinatorial tool of our algorithm: splitting an instance into a constant number of uniform instances. Note

that instances arising from this splitting are more constrained than the original instance, hence it is possible that they contain N^c -small sets that are not N^c -small in the original instance. Therefore, it may happen that the hypergraph of the original instance has no tree decomposition using the N^c -small sets of the original instance, but has a tree decomposition using the N^c -small sets of one of the new instances (and then we can invoke Lemma 4.7 to show that the original instance has a solution).

LEMMA 4.11. *Let $I = (V, D, C)$ be a CSP instance, let N be an integer, and let $c \geq 1$, $\epsilon > 0$ be real numbers. There is an algorithm with running time $2^{2^{O(|V|)} \cdot c/\epsilon} \cdot \text{poly}(\|I\|, N^c)$ that produces a set of (N, c, ϵ) -uniform N^c -consistent nontrivial instances I_1, \dots, I_t with $0 \leq t \leq 2^{2^{O(|V|)} \cdot c/\epsilon}$, all on the set V of variables, such that*

- (1) every solution of I is a solution of exactly one instance I_i ,
- (2) for every $1 \leq i \leq t$, instance I_i is a refinement of I .

PROOF. The main step of the algorithm takes a CSP instance I and either makes it (N, c, ϵ) -uniform and N^c -consistent without losing any solutions, or splits it into two instances $I_{\text{small}}, I_{\text{large}}$. By applying the main step recursively on I_{small} and I_{large} , we eventually arrive to a set of (N, c, ϵ) -uniform N^c -consistent instances. We will argue that the number of constructed instances is $2^{2^{O(|V|)} \cdot c/\epsilon}$.

In the main step, we first check if the instance is trivial; in this case we can stop with $t = 0$. Otherwise, we invoke the algorithm of Lemma 4.5 to obtain an N^c -consistent refinement of the instance, without losing any solution. Next we check if this N^c -consistent instance I is (N, c, ϵ) -uniform. This can be tested in time $2^{O(|V|)} \cdot \text{poly}(\|I\|, N^c)$ if we use Lemma 4.3 to find all the N^c -small sets and the corresponding sets of solutions. Suppose that N^c -small sets $B \subseteq A$ violate uniformity, that is,

$$\max(A|B) > N^\epsilon |\text{sol}(A)|/|\text{sol}(B)|. \quad (2)$$

Let $\text{sol}_{\text{small}}(B)$ contain those tuples b for which $|\text{sol}(A|B = b)| \leq \sqrt{N^\epsilon} |\text{sol}(A)|/|\text{sol}(B)|$ and let $\text{sol}_{\text{large}}(B) = \text{sol}(B) \setminus \text{sol}_{\text{small}}(B)$. Note that $|\text{sol}(A)| \geq |\text{sol}_{\text{large}}(B)| \cdot (\sqrt{N^\epsilon} |\text{sol}(A)|/|\text{sol}(B)|)$ (as every tuple $b \in \text{sol}_{\text{large}}(B)$ has at least $\sqrt{N^\epsilon} |\text{sol}(A)|/|\text{sol}(B)|$ extensions to A), hence

$$|\text{sol}_{\text{large}}(B)| \leq |\text{sol}(B)|/\sqrt{N^\epsilon}. \quad (3)$$

Let instance I_{small} (resp., I_{large}) be obtained from I by adding the constraint $\langle B, \text{sol}_{\text{small}}(B) \rangle$ (resp., $\langle B, \text{sol}_{\text{large}}(B) \rangle$). Clearly, the set of solutions of I is the disjoint union of the sets of solutions of I_{small} and I_{large} . This completes the description of the main step.

It is clear that if the recursive procedure stops, then the instances at the leaves of the recursion satisfy the two requirements: the application of Lemma 4.5 does not lose any solution and each resulting instance is N^c -consistent and (N, c, ϵ) -uniform. We show that the height of the recursion tree can be bounded from above by a function $h(|V|, c, \epsilon) = 2^{O(|V|)} \cdot c/\epsilon$ depending only on $|V|$, c , and ϵ ; in particular, this shows that the recursive algorithm eventually stops and produces at most $t = 2^{h(|V|, c, \epsilon)} = 2^{2^{O(|V|)} \cdot c/\epsilon}$ instances.

Let us consider a path in the recursion tree starting at the root, and let I^1, I^2, \dots, I^p be the corresponding N^c -consistent instances. If a set S is N^c -small in I^j , then it is N^c -small in $I^{j'}$ for every $j' > j$: the main step cannot increase $|\text{sol}(S)|$ for any S . Thus, with the exception of at most $2^{|V|}$ values of j , instances I^j and I^{j+1} have the same N^c -small sets. Let us consider a subpath I^x, \dots, I^y such that all these instances have the same N^c -small sets. We show that the length of this subpath is $O(3^{|V|} \cdot c/\epsilon)$, hence $p = O(2^{|V|} \cdot 3^{|V|} \cdot c/\epsilon)$. As this holds for any path starting at the root, we obtain that the height of the recursion tree is $2^{O(|V|)} \cdot c/\epsilon$ and hence $t = 2^{2^{O(|V|)} \cdot c/\epsilon}$.

For the instance I^j , let us define the following weight:

$$W^j = \sum_{\substack{\emptyset \subseteq B \subseteq A \subseteq V \\ A, B \text{ are } N^c\text{-small in } I^j}} \log_2 \max_{I^j} (A|B).$$

We bound the length of the subpath I^x, \dots, I^y by analyzing how this weight changes in each step. Observe first that when invoking the algorithm of Lemma 4.5 to find an N^c -consistent refinement, then the weight does not increase: adding new constraints cannot increase $\max(A|B)$ for any $A, B \subseteq V$ and cannot create new N^c -small sets by the assumption on the subpath I^x and I^y . Thus it is sufficient to analyze how the weight decreases in I_{large} and I_{small} compared to I . Note that $0 \leq W^j \leq 3^{|V|} \log_2 N^c = 3^{|V|} \cdot c \log_2 N$: the sum consists of at most $3^{|V|}$ terms and (as A is N^c -small and the instance I^j is N^c -consistent and nontrivial) $\max_{I^j} (A|B)$ is between 1 and N^c . We show that $W^{j+1} \leq W^j - (\epsilon/2) \log_2 N$, which immediately implies that the length of the subpath is $O(3^{|V|} \cdot c/\epsilon)$. Let us inspect how W^{j+1} changes compared to W^j . Since I^j and I^{j+1} have the same N^c -small sets by assumption, no new term can appear in W^{j+1} . It is clear that $\max_{I^{j+1}} (A|B)$ cannot be greater than $\max_{I^j} (A|B)$ for any A, B . Moreover, we show that there is at least one term that strictly decreases. Suppose first that I^{j+1} was obtained from I^j by adding the constraint $\langle B, \text{sol}_{\text{small}}(B) \rangle$. Then

$$\log_2 \max_{I^{j+1}} (A|B) \leq \log_2 \sqrt{N^\epsilon} \frac{|\text{sol}_{I^j}(A)|}{|\text{sol}_{I^j}(B)|} \leq \log_2 \frac{\max_{I^j} (A|B)}{\sqrt{N^\epsilon}} = \log_2 \max_{I^j} (A|B) - (\epsilon/2) \log_2 N,$$

where we have used (2) in the second inequality. On the other hand, if I^{j+1} was obtained by adding the constraint $\langle B, \text{sol}_{\text{large}}(B) \rangle$, then

$$\log_2 \max_{I^{j+1}} (B|\emptyset) = \log_2 |\text{sol}_{I^{j+1}}(B)| \leq \log_2 (|\text{sol}_{I^j}(B)|/\sqrt{N^\epsilon}) = \log_2 \max_{I^j} (B|\emptyset) - (\epsilon/2) \log_2 N,$$

where the inequality follows from (3). In both cases, we get that at least one term decreases by at least $(\epsilon/2) \log_2 N$. \square

4.3. Uniform CSP instances and submodularity

Assume for a moment that we have a 1-uniform instance I with hypergraph H . Note that by Prop 4.9(1), this means that $\max(A|B) = |\text{sol}(A)|/|\text{sol}(B)|$. Suppose that every constraint contains at most N tuples and let us define the function $b(S) = \log_N |\text{sol}(S)|$. For every edge $e \in E(H)$, there is a corresponding constraint, which has at most N tuples by the definition of N . Thus $|\text{sol}(e)| \leq N$ and hence $b(e) \leq 1$ for every $e \in E(H)$, that is, b is edge dominated. The crucial observation of this section is that this function b is submodular:

$$\begin{aligned} b(X) + b(Y) &= \log_N |\text{sol}(X)| + \log_N \left(|\text{sol}(X \cap Y)| \frac{|\text{sol}(Y)|}{|\text{sol}(X \cap Y)|} \right) \\ &= \log_N |\text{sol}(X)| + \log_N (|\text{sol}(X \cap Y)| \cdot \max(Y|X \cap Y)) \\ &\geq \log_N |\text{sol}(X)| + \log_N (|\text{sol}(X \cap Y)| \cdot \max(X \cup Y|X)) \\ &= \log_N |\text{sol}(X)| + \log_N \left(|\text{sol}(X \cap Y)| \cdot \frac{|\text{sol}(X \cup Y)|}{|\text{sol}(X)|} \right) \\ &= \log_N |\text{sol}(X \cap Y)| + \log_N |\text{sol}(X \cup Y)| \\ &= b(X \cap Y) + b(X \cup Y) \end{aligned}$$

(the equalities follow from 1-uniformity; the inequality uses Prop. 4.9(2) with $A = Y$, $B = X \cap Y$, $C = X$). Therefore, if the submodular width of H is at most c , then H has a tree decomposition where $b(B) \leq c$ and hence $|\text{sol}(B)| \leq N^c$ for every bag B . Thus we can find a solution of the instance by dynamic programming in time polynomial in N^c .

Lemma 4.11 guarantees some uniformity for the created instances, but not perfect 1-uniformity and only for the N^c -small sets. Thus in Lemma 4.12, we need to define b in a slightly different and more technical way: we add some small terms to correct errors arising from the weaker uniformity and we truncate the function at large values (i.e., for sets that are not N^c -small). Also, we state Lemma 4.12 for an arbitrary hypergraph H , possibly different from the hypergraph of I ; then to guarantee that b is edge dominated, we need that $|\text{sol}(e)| \leq N$ for every edge e of H . The reason we need the statement in this form is that in an instance I_i produced by Lemma 4.11 from instance I , the maximum size of a constraint relation in I is an upper bound on $|\text{sol}_{I_i}(e)|$ only if e corresponds to a constraint of I , but we have no such bound if e corresponds to a constraint of I_i not appearing in I .

LEMMA 4.12. *Let $I = (V, D, C)$ be a CSP instance and let H be a hypergraph on V (possibly different from the hypergraph of I) such that $|\text{sol}(e)| \leq N$ holds for every $e \in E(H)$. If I is N^c -consistent and (N, c, ϵ^3) -uniform for some $c \geq 1$ and $\epsilon := 1/|V|$, then the following function b is an edge-dominated, monotone, submodular function on H with $b(\emptyset) = 0$:*

$$b(S) := \begin{cases} (1 - \epsilon) \log_N |\text{sol}(S)| + 2\epsilon^2 |S| - \epsilon^3 |S|^2 & \text{if } S \text{ is } N^c\text{-small,} \\ (1 - \epsilon)c + 2\epsilon^2 |S| - \epsilon^3 |S|^2 & \text{otherwise.} \end{cases}$$

PROOF. Let $h(S) := 2\epsilon^2 |S| - \epsilon^3 |S|^2$. The function $h(S)$ is a quadratic function of $|S|$; it is 0 when $|S| = 0$ or $|S| = 2/\epsilon$, hence its maximum is at $|S| = 1/\epsilon = |V(H)|$ with maximum value ϵ . Therefore, $h(S)$ is monotone in the range $0 \leq |S| \leq |V(H)|$. Furthermore, h is a submodular function:

$$\begin{aligned} & h(X) + h(Y) - h(X \cap Y) - h(X \cup Y) \\ &= 2\epsilon^2(|X| + |Y| - |X \cap Y| - |X \cup Y|) + \epsilon^3(-|X|^2 - |Y|^2 + |X \cap Y|^2 + |X \cup Y|^2) \\ &= \epsilon^3(-(|X \cap Y| + |X \setminus Y|)^2 - (|X \cap Y| + |Y \setminus X|)^2 \\ &+ |X \cap Y|^2 + (|X \cap Y| + |X \setminus Y| + |Y \setminus X|)^2) \\ &= 2\epsilon^3 |X \setminus Y| \cdot |Y \setminus X| \geq 0. \end{aligned}$$

This calculation shows that if $|X \setminus Y|, |Y \setminus X| \geq 1$, then we actually have $h(X) + h(Y) \geq h(X \cap Y) + h(X \cup Y) + 2\epsilon^3$. We will use this extra $2\epsilon^3$ term to dominate the error terms arising from assuming only (N, c, ϵ^3) -uniformity instead of perfect 1-uniformity.

Let us first verify the monotonicity of b . If Y is N^c -small, then every $X \subseteq Y$ is N^c -small, which implies $|\text{sol}(X)| \leq |\text{sol}(Y)|$ as I is N^c -consistent. Therefore, $b(X) \leq b(Y)$ follows from the monotonicity of h . If Y is not N^c small, then $b(Y) = (1 - \epsilon)c + h(Y)$ and $b(X) \leq b(Y)$ is clear for every $X \subseteq Y$, no matter whether X is N^c -small or not.

To see that b is edge-dominated, consider an edge $e \in E(H)$. By assumption, $\log_N |\text{sol}(e)| \leq 1$ for every $e \in E(H)$ and hence (using N^c -consistency and $c \geq 1$) e is N^c -small. Thus $b(e) \leq (1 - \epsilon) + h(S) \leq 1$, as required.

Finally, let us verify the submodularity of b for some $X, Y \subseteq V$. If $X \subseteq Y$ or $Y \subseteq X$, then there is nothing to show. Thus we can assume that $|X \setminus Y|, |Y \setminus X| \geq 1$. We consider three cases depending on which of X and Y are N^c -small. Suppose first that X and Y are both N^c -small. In this case,

$$\begin{aligned}
b(X) + b(Y) &= (1 - \epsilon) \log_N |\text{sol}(X)| + (1 - \epsilon) \log_N |\text{sol}(Y)| + h(X) + h(Y) \\
&= (1 - \epsilon) \log_N |\text{sol}(X)| + (1 - \epsilon) \log_N \left(|\text{sol}(X \cap Y)| \cdot \frac{|\text{sol}(Y)|}{|\text{sol}(X \cap Y)|} \right) \\
&\quad + h(X) + h(Y) \\
&\geq (1 - \epsilon) \log_N |\text{sol}(X)| + (1 - \epsilon) \log_N |\text{sol}(X \cap Y)| \\
&\quad + (1 - \epsilon) \log_N (\max(Y|X \cap Y)/N^{\epsilon^3}) + h(X) + h(Y) \\
&\geq (1 - \epsilon) \log_N |\text{sol}(X \cap Y)| + (1 - \epsilon) \log_N (|\text{sol}(X)| \max(X \cup Y|X)) \\
&\quad - (1 - \epsilon) \cdot \epsilon^3 + h(X \cap Y) + h(X \cup Y) + 2\epsilon^3 \\
&\geq (1 - \epsilon) \log_N |\text{sol}(X \cap Y)| + (1 - \epsilon) \log_N |\text{sol}(X \cup Y)| + h(X \cap Y) + h(X \cup Y) \\
&\geq b(X \cap Y) + b(X \cup Y)
\end{aligned}$$

(in the first inequality, we used the definition of (N, c, ϵ^3) -uniformity on $X \cap Y$ and Y ; in the second inequality, we used the submodularity of h and Prop. 4.9(2) for $A = Y$, $B = X \cap Y$, and $C = X$; in the third inequality, we used Prop. 4.9(1) for $A = X \cup Y$, $B = X$; the last inequality is strict only if $X \cup Y$ is not N^c -small).

For the second case, suppose that, say, X is N^c -small but Y is not. In this case, $X \cap Y$ is N^c -small but $X \cup Y$ is not. Thus

$$\begin{aligned}
b(X) + b(Y) &= (1 - \epsilon) \log_N |\text{sol}(X)| + (1 - \epsilon)c + h(X) + h(Y) \\
&\geq (1 - \epsilon) \log_N |\text{sol}(X \cap Y)| + (1 - \epsilon)c + h(X \cap Y) + h(X \cup Y) \\
&= b(X \cap Y) + b(X \cup Y)
\end{aligned}$$

(in the inequality, we used the N^c -consistency on $X \cap Y$ and Y , and the submodularity of h).

Finally, suppose that neither X nor Y is N^c -small. In this case, $X \cup Y$ is not N^c -small either. Now

$$b(X) + b(Y) = 2(1 - \epsilon)c + h(X) + h(Y) \geq 2(1 - \epsilon)c + h(X \cap Y) + h(X \cup Y) \geq b(X \cap Y) + b(X \cup Y).$$

□

Having constructed the submodular function b as in Lemma 4.12, we can use the argument described at the beginning of the section: if H has submodular width at most $(1 - \epsilon)c$, then there is a tree decomposition where every bag is N^c -small, and we can use this tree decomposition to find a solution. In fact, by Lemma 4.7, in this case N^c -consistency implies that every nontrivial instance has a solution.

PROOF (OF THEOREM 4.1). Let I be an instance of $\text{CSP}(\mathcal{H})$ having hypergraph $H \in \mathcal{H}$. We decide the solvability of I the following way. Let $N \leq \|I\|$ be the size of the largest constraint relation in I , i.e., every constraint has at most N satisfying assignments. Trivially, we have $|\text{sol}_I(e)| \leq N$ for every $e \in E(H)$. Set $\epsilon := 1/|V(H)|$ (we may assume that $|V(H)| \geq 2$), and let $c := c_0/(1 - \epsilon) \leq 2c_0$. Let us use the algorithm of Lemma 4.11 to produce the nontrivial N^c -consistent (N, c, ϵ^3) -uniform instances I_1, \dots, I_t . The running time of this step is $2^{2^{O(|V|)} \cdot c/\epsilon} \cdot \text{poly}(\|I\|, N^c)$, which is $2^{c_0 \cdot 2^{O(|V(H)|)}} \cdot \|I\|^{O(c_0)}$.

If $t = 0$, then we can conclude that I has no solution. Otherwise, we argue that I has a solution. Consider any I_i ; as I_i is a refinement of I , we have $|\text{sol}_{I_i}(e)| \leq |\text{sol}_I(e)| \leq N$ for any $e \in E(H)$. Let us use Lemma 4.12 with I_i and H (the hypergraph of I , not I_i !) to define the edge-dominated monotone submodular function b . By definition of submodular width, H has a tree decomposition $(T, (B_t)_{t \in V(T)})$ such that $b(B_t) \leq \text{subw}(H) \leq c_0 = (1 - \epsilon)c$ for every $t \in V(T)$. Since $b(S) \leq (1 - \epsilon)c$ implies $|\text{sol}_{I_i}(S)| \leq N^c$ and b is monotone, this means

that B_t is N^c -small in I_i for every $t \in V(T)$. Therefore, the conditions of Lemma 4.7 hold, and I has a solution. \square

5. FROM SUBMODULAR FUNCTIONS TO HIGHLY CONNECTED SETS

The aim of this section is to show that if a hypergraph H has large submodular width, then there is a large highly connected set in H . Recall that we say that a set W is (μ, λ) -connected for some fractional independent set μ and $\lambda > 0$, if for every disjoint $A, B \subseteq W$, every fractional (A, B) -separator has weight at least $\lambda \cdot \min\{\mu(A), \mu(B)\}$ (see Section 2). Equivalently, we can say that for every disjoint $A, B \subseteq W$, there is an (A, B) -flow of value $\lambda \cdot \min\{\mu(A), \mu(B)\}$. Recall also that $\text{con}_\lambda(H)$ denotes the maximum value of $\mu(W)$ taken over every fractional independent set μ and (μ, λ) -connected set W .

The main result of this section allows us to identify a highly connected set if submodular width is large:

THEOREM 5.1. *For every sufficiently small constant $\lambda > 0$, the following holds. Let b be an edge-dominated monotone submodular function of H with $b(\emptyset) = 0$. If the b -width of H is greater than $\frac{3}{2}(w + 1)$, then $\text{con}_\lambda(H) \geq w$.*

For the proof of Theorem 5.1, we need to show that if there is no tree decomposition where $b(B)$ is small for every bag B , then a highly connected set exists. There is a standard recursive procedure that either builds a tree decomposition or finds a highly connected set (see e.g., [Flum and Grohe 2006, Section 11.2]). Simplifying somewhat, the main idea is that if the graph can be decomposed into smaller graphs by splitting a certain set of vertices into two parts, then a tree decomposition for each part is constructed using the algorithm recursively, and the tree decompositions for the parts are joined in an appropriate way to obtain a tree decomposition for the original graph. On the other hand, if the set of vertices cannot be split, then we can conclude that it is highly connected. This high-level idea has been applied for various notions of tree decompositions [Oum and Seymour 2006; Oum 2005; Adler et al. 2007; Oum and Seymour 2007; Marx 2010a], and it turns out to be useful in our context as well. However, we need to overcome two major difficulties:

(1) Highly connected set in our context is defined as not having certain *fractional separators* (i.e., weight assignments). However, if we want to build a tree decomposition in a recursive manner, we need *integer separators* (i.e., subsets of vertices) that decompose the hypergraph into smaller parts.

(2) Measuring the sizes of sets with a submodular function b can lead to problems, since the size of the union of two sets can be much smaller than the sum of the sizes of the two sets. We need the property that, roughly speaking, removing a “large” part from a set makes it “much smaller.” For example, if A and B are components of $H \setminus S$, and both $b(A)$ and $b(B)$ are large, then we need the property that both of them are much smaller than $b(A \cup B)$. Adler [2006, Section 4.2] investigates the relation between some notion of highly connected sets and f -width, but assumes that f is additive: if A and B do not touch, then $f(A \cup B) = f(A) + f(B)$. However, for a submodular function b , there is no reason to assume that additivity holds: for example, it very well may be that $b(A) = b(B) = b(A \cup B)$.

To overcome the first difficulty, we have to understand what fractional separation really means. The first question is whether fractional separation is equivalent to some notion of integral separation, perhaps up to constant factors. The first, naive, question is whether a fractional (X, Y) -separator of weight w implies that there are $O(w)$ edges whose union is an (X, Y) -separator, i.e., there is an (X, Y) -separator S with $\rho_H(S) = O(w)$. There is a simple counterexample showing that this is not true. It is well-known that for every integer $k > 0$, there is a hypergraph H_k such that $\rho^*(H_k) = 2$ and $\rho(H_k) = k$. Let V be the set of vertices of H_k and let H'_k be obtained from H_k by extending it with two independent sets X, Y , each of size k , and connecting every vertex of $X \cup Y$ with every vertex of V . It is

clear that there is a fractional (X, Y) -separator of weight 2, but every (X, Y) -separator S has to fully contain at least one of X, Y , or V , implying $\rho_{H'}(S) \geq k$.

A less naive question is whether a fractional (X, Y) -separator with weight w in H implies that there exists an (X, Y) -separator S with $\rho_H^*(S) = O(w)$ (or at most $f(w)$ for some function f). It can be shown that this is not true either: using the hypergraph family presented in [Marx 2011, Section 5], one can construct counterexamples where the minimum weight of a fractional (X, Y) -separator is a constant, but $\rho_H^*(S)$ has to be arbitrarily large for every (X, Y) -separator S (we omit the details).

We will characterize fractional separation in a very different way. We show that if there is a fractional (A, B) -separator of weight w , then there is an (A, B) -separator S with $b(S) = O(w)$ for *every* edge-dominated monotone submodular function b . Note that this separator S can be different for different functions b , so we are not claiming that there is a single (A, B) -separator S that is small in every b . The converse is also true, thus this gives a novel characterization of fractional separation, tight up to a constant factor. This result is the key idea that allows us to move from the domain of submodular functions to the domain of pure hypergraph properties: if there is no (A, B) -separator such that $b(S)$ is small, then we know that there is no small fractional (A, B) -separator, which is a property of the hypergraph H only and has no longer anything to do with the submodular function b .

To overcome the second difficulty, we introduce a transformation that turns a monotone submodular function b on $V(H)$ into a function b^* that encodes somehow the neighborhood structure of H as well. The new function b^* is no longer monotone and submodular, but it has a number of remarkable properties, for example, b^* remains edge dominated and $b^*(S) \geq b(S)$ for every set $S \subseteq V(H)$, implying that b^* -width is not smaller than b -width. The main idea is to prove Theorem 5.1 for b^* -width instead of b -width (note that this makes the statement stronger). Because of the way b^* encodes the neighborhoods, the second difficulty will disappear: for example, it will be true that $b^*(A \cup B) = b^*(A) + b^*(B)$ if there are no edges between A and B , that is, b^* is additive on disjoint components. Lemma 5.6 formulates (in a somewhat technical way) the exact property of b^* that we will need. Furthermore, luckily it turns out that the result mentioned in the previous paragraph remains true with b replaced by b^* : if there is a fractional (A, B) -separator of weight w , then there is an (A, B) -separator S such that not only $b(S)$, but even $b^*(S)$ is $O(w)$.

5.1. The function b^*

We define the function b^* the following way. Let H be a hypergraph and let b be a monotone submodular function defined on $V(H)$. Let $S_{V(H)}$ be the set of all permutations of $V(H)$. For a permutation $\pi \in S_{V(H)}$, let $N_\pi^-(v)$ be the neighbors of v preceding v in the ordering π . For $\pi \in S_{V(H)}$ and $Z \subseteq V(H)$, we define

$$\partial b_{\pi, Z}(v) := b(v \cup (N_\pi^-(v) \cap Z)) - b(N_\pi^-(v) \cap Z).$$

In other words, $\partial b_{\pi, Z}(v)$ is the marginal value of v with respect to the set of its neighbors in Z preceding it. We abbreviate $\partial b_{\pi, V(H)}$ by ∂b_π . As usual, we extend the definition to subsets by letting $\partial b_{\pi, Z}(S) := \sum_{v \in S} \partial b_{\pi, Z}(v)$. Furthermore, we define

$$b_\pi(Z) := \partial b_{\pi, Z}(Z) = \sum_{v \in Z} \partial b_{\pi, Z}(v),$$

$$b^*(Z) := \min_{\pi \in S_{V(H)}} b_\pi(Z).$$

Thus $b_\pi(Z)$ is the sum of the marginal values with respect to a given ordering, while $b^*(Z)$ is the smallest possible sum taken over all possible orderings. Let us prove some simple properties of the function b^* . Properties (1)–(3) and their proofs show why b^* was defined

this way: $b^*(Z)$ is never smaller than $b(Z)$, but it is still edge dominated. Properties (4)–(5) are technical statements that we will need later.

PROPOSITION 5.2. *Let H be a hypergraph and let b be a monotone submodular function defined on $V(H)$ with $b(\emptyset) = 0$. For every $\pi \in S_{V(H)}$ and $Z \subseteq V(H)$ we have*

- (1) $b_\pi(Z) \geq b(Z)$,
- (2) $b^*(Z) \geq b(Z)$,
- (3) $b^*(Z) = b_\pi(Z) = b(Z)$ if Z is a clique,
- (4) $\partial b_{\pi, Z_1}(v) \leq \partial b_{\pi, Z_2}(v)$ if $Z_2 \subseteq Z_1$,
- (5) $\partial b_\pi(v) \leq \partial b_{\pi, Z}(v)$,
- (6) $b^*(X \cup Y) \leq b^*(X) + b^*(Y)$.

PROOF. (1) We prove the statement by induction on $|Z|$; for $Z = \emptyset$, the claim is true (as $b(\emptyset) = 0$). Otherwise, let v be the last element of Z according to the ordering π . As v is not preceding any element of Z , for every $u \in Z$ we have $N_\pi^-(u) \cap Z = N_\pi^-(u) \cap (Z \setminus v)$, and hence $\partial b_{\pi, Z}(u) = \partial b_{\pi, Z \setminus v}(u)$.

$$\begin{aligned} b_\pi(Z) &= \sum_{u \in Z \setminus v} \partial b_{\pi, Z}(u) + \partial b_{\pi, Z}(v) = \sum_{u \in Z \setminus v} \partial b_{\pi, Z \setminus v}(u) + \partial b_{\pi, Z}(v) \\ &= b_\pi(Z \setminus v) + \partial b_{\pi, Z}(v) \geq b(Z \setminus v) + b(v \cup (N_\pi^-(v) \cap Z)) - b(N_\pi^-(v) \cap Z) \geq b(Z). \end{aligned}$$

In the first inequality, we used the induction hypothesis and the definition of $\partial b_{\pi, Z}(v)$; in the second inequality, we used the submodularity of b : the marginal value of v with respect to $Z \setminus v$ is not greater than with respect to $N_\pi^-(v) \cap Z \subseteq Z \setminus v$.

(2) Follows immediately from (1) and from the definition of b^* .

(3) As $b_\pi(Z) \geq b^*(Z) \geq b(Z)$ (by property (1) and the definition of $b^*(Z)$), we need to prove $b_\pi(Z) = b(Z)$ only. We prove the statement by induction on $|Z|$. As in (1), let v be the last vertex of Z in π . Note that since Z is a clique, $N_\pi^-(v) \cap Z$ is exactly $Z \setminus v$.

$$\begin{aligned} b_\pi(Z) &= \sum_{u \in Z \setminus v} \partial b_{\pi, Z}(u) + \partial b_{\pi, Z}(v) = \sum_{u \in Z \setminus v} \partial b_{\pi, Z \setminus v}(u) + b(v \cup (N_\pi^-(v) \cap Z)) - b(N_\pi^-(v) \cap Z) \\ &= b_\pi(Z \setminus v) + b(v \cup (Z \setminus v)) - b(Z \setminus v) = b(Z \setminus v) + b(Z) - b(Z \setminus v) = b(Z). \end{aligned}$$

(4) Follows from the submodularity of b : $\partial b_{\pi, Z_1}(v)$ is the marginal value of v with respect to $N_\pi^-(v) \cap Z_1$, while $\partial b_{\pi, Z_2}(v)$ is the marginal value of v with respect to the subset $N_\pi^-(v) \cap Z_2$ of $N_\pi^-(v) \cap Z_1$.

(5) Immediate from (4).

(6) Let π_X and π_Y be the orderings such that $b_{\pi_X}(X) = b^*(X)$ and $b_{\pi_Y}(Y) = b^*(Y)$. Let us define ordering π such that it starts with the elements of X , in the order of π_X , followed by the elements of $Y \setminus X$, in the order of π_Y , and completed by an arbitrary ordering of $V(H) \setminus (X \cup Y)$. It is clear that for every $v \in X$, we have $\partial b_{\pi, X \cup Y}(v) = \partial b_{\pi_X, X}(v)$. Furthermore, for every $v \in Y \setminus X$, we have $N_\pi^-(v) \cap Y \subseteq N_\pi^-(v) \cap (X \cup Y)$: if u is a neighbor of v in Y that precedes it in π_Y , then u is either in X or in $Y \setminus X$; in both cases u precedes v in π . Thus, similarly to (4), we have $\partial b_{\pi, X \cup Y}(v) \leq \partial b_{\pi_Y, Y}(v)$ for every $v \in Y \setminus X$: $\partial b_{\pi, X \cup Y}(v)$ is the marginal value of v with respect to $N_\pi^-(v) \cap (X \cup Y)$, while $\partial b_{\pi_Y, Y}(v)$ is the marginal value of v with respect to the subset $N_{\pi_Y}^-(v) \cap Y$. Now we have

$$\begin{aligned} b^*(X \cup Y) &\leq b_\pi(X \cup Y) = \sum_{v \in X \cup Y} \partial b_{\pi, X \cup Y}(v) \leq \sum_{v \in X} \partial b_{\pi_X, X}(v) + \sum_{v \in Y \setminus X} \partial b_{\pi_Y, Y}(v) \\ &\leq \sum_{v \in X} \partial b_{\pi_X, X}(v) + \sum_{v \in Y} \partial b_{\pi_Y, Y}(v) = b^*(X) + b^*(Y). \quad \square \end{aligned}$$

Prop. 5.2(3) implies that $\partial b_{w,Z}$ can be used to define a fractional independent set:

LEMMA 5.3. *Let H be a hypergraph and let b be an edge-dominated monotone submodular function defined on $V(H)$ with $b(\emptyset) = 0$. Let $W \subseteq V(H)$ and let π be an ordering of W . Let us define $\mu(v) = \partial b_{\pi,W}(v)$ for $v \in W$ and $\mu(v) = 0$ otherwise. Then μ is a fractional independent set of H with $\mu(W) = b_\pi(W)$.*

PROOF. Let e be an edge of H and let $Z := e \cap W$. We have

$$\mu(e) = \mu(Z) = \partial b_{\pi,W}(Z) \leq \partial b_{\pi,Z}(Z) = b_\pi(Z) = b(Z) \leq 1,$$

where the first inequality follows from Prop. 5.2(4), the last equality follows from Prop. 5.2(3), and the second inequality follows from the fact that b is edge dominated. Furthermore, we have $\mu(W) = \partial b_{\pi,W}(W) = b_\pi(W)$. \square

We close this section by proving the main property of b^* that allows us to avoid the second difficulty described at the beginning of Section 5. First, although it is not used directly, let us state that b^* is additive on sets that are independent from each other:

LEMMA 5.4. *Let H be a hypergraph, let b be an edge-dominated monotone submodular function defined on $V(H)$ with $b(\emptyset) = 0$, and let $A, B \subseteq V(H)$ be disjoint sets such that there is no edge intersecting both A and B . Then $b^*(A \cup B) = b^*(A) + b^*(B)$.*

PROOF. By Prop. 5.2(6), we have to show only $b^*(A \cup B) \geq b^*(A) + b^*(B)$. Let π be an ordering of $V(H)$ such that $b_\pi(A \cup B) = b^*(A \cup B)$; we can assume that π starts with the vertices of $A \cup B$. Since there is no edge that intersects both A and B , and no vertex outside $A \cup B$ precedes a vertex $u \in A \cup B$, we have $N_{\pi^-}(u) \subseteq A$ for every $u \in A$ and $N_{\pi^-}(u) \subseteq B$ for every $u \in B$. Thus $\partial b_{\pi,A \cup B}(u) = \partial b_{\pi,A}(u)$ for every $u \in A$ and $\partial b_{\pi,A \cup B}(u) = \partial b_{\pi,B}(u)$ for every $u \in B$. Therefore, $b^*(A \cup B) = b_\pi(A \cup B) = b_\pi(A) + b_\pi(B) \geq b^*(A) + b^*(B)$, what we had to show. \square

The actual statement that we use is more complicated than Lemma 5.4: there can be edges between A and B , but we assume that there is a small (A, B) -separator. We want to generalize the following simple statement to our setting:

PROPOSITION 5.5. *Let G be a graph, $W \subseteq V(G)$ a set of vertices, $A, B \subseteq W$ two disjoint subsets, and an (A, B) -separator S . If $|S| < |A|, |B|$, then $|(C \cap W) \cup S| < |W|$ for every component C of $G \setminus S$.*

The proof of Prop. 5.5 is easy to see: every component C of $G \setminus S$ is disjoint from either A or B , thus $|C \cap W|$ is at most $|W| - \min\{|A|, |B|\} < |W| - |S|$, implying that $|(C \cap W) \cup S|$ is less than $|W|$. Statements of this form are useful when constructing tree decompositions in a recursive way. In our setting, we want to measure the size of the sets using the function b^* , not by the number of vertices. More precisely, we measure the size of S and $(C \cap W) \cup S$ using b^* , while the size of W , A , and B are measured using the fractional independent set μ defined by Lemma 5.3. The reason for this will be apparent in the proof of Lemma 5.10: we want to claim that if such a separator S does not exist for any $A, B \subseteq W$, then W is a (μ, λ) -connected set for this fractional independent set μ .

LEMMA 5.6. *Let H be a hypergraph, let b be an edge-dominated monotone submodular function defined on $V(H)$ with $b(\emptyset) = 0$ and let W be a set of vertices. Let π_W be an ordering of $V(H)$, and let $\mu(v) := \partial b_{\pi_W,W}(v)$ for $v \in W$ and $\mu(v) = 0$ otherwise. Let $A, B \subseteq W$ be two disjoint sets, and let S be an (A, B) -separator. If $b^*(S) < \mu(A), \mu(B)$, then $b^*((C \cap W) \cup S) < \mu(W)$ for every component C of $H \setminus S$.*

PROOF. Let C be a component of $H \setminus S$ and let $Z := (C \cap W) \cup S$. Let π_S be the ordering reaching the minimum in the definition of $b^*(S)$. Let us define the ordering π that starts with S in the order of π_S , followed by $C \cap W$ in the order of π_W , and finished by

an arbitrary ordering of the remaining vertices. It is clear that for every $v \in S$, we have $\partial b_{\pi,Z}(v) = \partial b_{\pi_S,S}(v)$. Let us consider a vertex $v \in C \cap W$ and let $u \in W$ be a neighbor of v that precedes it in π_W . Since $v \in C$ and C is a component of $H \setminus S$, either $u \in S$ or $u \in C \cap W$. In both cases, u precedes v in π . This means that $N_{\pi_W}^-(v) \cap W \subseteq N_{\pi}^-(v) \cap Z$, which implies that $\partial b_{\pi,Z}(v) \leq \partial b_{\pi_W,W}(v) = \mu(v)$ for every $v \in C \cap W$. As S separates A and B , component C intersects at most one of A and B ; suppose, without loss of generality, that C is disjoint from A . Thus

$$\begin{aligned} b^*(Z) \leq b_{\pi}(Z) &= \sum_{v \in S} \partial b_{\pi,Z}(v) + \sum_{v \in C \cap W} \partial b_{\pi,Z}(v) \leq b^*(S) + \mu(C \cap W) \\ &< \mu(A) + \mu(W \setminus A) = \mu(W). \quad \square \end{aligned}$$

5.2. Submodular separation

This section is devoted to understanding what fractional separation means: we show that having a small fractional (A, B) -separator is essentially equivalent to the property that for every edge-dominated submodular function b , there is an (A, B) -separator S such that $b(S)$ is small. The proof is based on a standard trick that is often used for rounding fractional solutions for separation problems: we define a distance function and show by an averaging argument that cutting at some distance t gives a small separator. However, in our setting, we need significant new ideas to make this trick work: the main difficulty is that the cost function b is defined on *subsets* of vertices and is not a modular function defined by the cost of vertices. To overcome this problem, we use the definitions in Section 5.1 (in particular, the function $\partial b_{\pi}(v)$) to assign a cost to every single vertex.

THEOREM 5.7. *Let H be a hypergraph, $X, Y \subseteq V(H)$ two sets of vertices, and $b : V(H) \rightarrow \mathbb{R}^+$ an edge-dominated monotone submodular function with $b(\emptyset) = 0$. Suppose that s is a fractional (X, Y) -separator of weight at most w . Then there is an (X, Y) -separator $S \subseteq V(H)$ with $b(S) \leq b^*(S) = O(w)$.*

PROOF. The total weight of the edges covering a vertex v is $\sum_{e \in E(H), v \in e} s(e)$; let us define $x(v) := \min\{1, \sum_{e \in E(H), v \in e} s(e)\}$. It is clear that if P is a path from X to Y , then $\sum_{v \in P} x(v) \geq 1$. We define the distance $d(v)$ to be the minimum of $\sum_{v' \in P} x(v')$, taken over all paths from X to v (this means that $d(v) = x(v)$ for every $v \in X$, that is, $d(v) > 0$ is possible for $v \in X$). It is clear that $d(v) \geq 1$ for every $v \in Y$. Let us associate the closed interval $\iota(v) = [d(v) - x(v), d(v)]$ to each vertex v . If v is in X , then the left endpoint of $\iota(v)$ is 0, while if v is in Y , then the right endpoint of $\iota(v)$ is at least 1.

Let u and v be two adjacent vertices in H such that $d(u) \leq d(v)$. It is easy to see that $d(v) \leq d(u) + x(u)$: there is a path P from X to u such that $\sum_{u' \in P} x(u') = d(u)$, thus the path P' obtained by appending v to P has $\sum_{v' \in P'} x(v') = \sum_{u' \in P} x(u') + x(v) = d(u) + x(v)$. Therefore, we have:

Claim 1. If u and v are adjacent, then $\iota(u) \cap \iota(v) \neq \emptyset$.

The *class* of a vertex $v \in V(H)$ is the largest integer $\kappa(v)$ such that $x(v) \leq 2^{-\kappa(v)}$, and we define $\kappa(v) := \infty$ if $x(v) = 0$. Recall that $x(v) \leq 1$, thus $\kappa(v)$ is nonnegative. The *offset* of a vertex v is the unique value $0 \leq \alpha < 2 \cdot 2^{-\kappa(v)}$ such that $d(v) = i(2 \cdot 2^{-\kappa(v)}) + \alpha$ for some integer i . In other words, if the class is 0, 1, 2, ..., the offset is $d(v)$ modulo 2, 1, 1/2, ..., respectively. Let us define an ordering $\pi = (v_1, \dots, v_n)$ of $V(H)$ such that

- $\kappa(v)$ is nondecreasing,
- among vertices having the same class, the offset is nondecreasing.

Let directed graph D be the orientation of the primal graph of H such that if v_i and v_j are adjacent and $i < j$, then there is a directed edge $\vec{v}_i v_j$ in D . Figure 3 shows a directed




Fig. 3. The intervals corresponding to a directed path v_1, \dots, v_8 . The shaded lines show the offsets of the vertices.




Fig. 4. Proof of Claim 2: Two examples of directed paths where every vertex has the same class κ (and $h := 2^{-\kappa}$). The shaded lines show the offsets of the vertices.

path in D . If P is a directed path in D , then the *width* of P is the length of the interval $\bigcup_{v \in P} \iota(v)$ (note that by Claim 1, this union is indeed an interval). The following claim bounds the maximum possible width of a directed path:

Claim 2. If P is a directed path D starting at v , then the width of P is at most $16x(v)$.

Proof. We first prove that if every vertex of P has the same class $\kappa(v)$, then the width of P is at most $4 \cdot 2^{-\kappa(v)}$. Since the class is nondecreasing along the path, we can partition the path into subpaths such that every vertex in a subpath has the same class and the classes are distinct on the different subpaths. The width of P is at most the sum of the widths of the subpaths, which is at most $\sum_{i \geq \kappa(v)} 4 \cdot 2^{-i} = 8 \cdot 2^{-\kappa(v)} \leq 16x(v)$.

Suppose now that every vertex of P has the same class $\kappa(v)$ as the first vertex v and let $h := 2^{-\kappa(v)}$. As the offset is nondecreasing, path P can be partitioned into two parts: a subpath P_1 containing vertices with offset less than h , followed by a subpath P_2 containing vertices with offset at least h (one of P_1 and P_2 can be empty). See Figure 4 for examples. We show that each of P_1 and P_2 has width at most $2h$, which implies that the width of P is at most $4h$. Observe that if $u \in P_1$ and $\iota(u)$ contains a point $i \cdot 2h - h$ for some integer i , then, considering $x(u) \leq h$ and the bounds on the offset of u , this is only possible if $\iota(u) = [i \cdot 2h - h, i \cdot 2h]$, i.e., $i \cdot 2h - h$ is the left endpoint of $\iota(u)$. Thus if $I_1 = \bigcup_{u \in P_1} \iota(u)$ contains $i \cdot 2h - h$, then it is the left endpoint of I_1 . Therefore, I_1 can contain $i \cdot 2h - h$ for at most one value of i , which immediately implies that the length of I_1 is at most $2h$.

We argue similarly for P_2 . If $u \in P_2$, then $\iota(u)$ can contain the point $i \cdot 2h$ only if $\iota(u) = [i \cdot 2h, i \cdot 2h + h]$. Thus if $I_2 = \bigcup_{u \in P_2} \iota(u)$ contains $i \cdot 2h$, then it is the left endpoint of I_2 . We get that I_2 can contain $i \cdot 2h$ for at most one value of i , which immediately implies that the width of I_2 is at most $2h$. This concludes the proof of Claim 2. \square

Let $c(v) := \partial b_\pi(v)$.

Claim 3. $\sum_{v \in V(H)} x(v)c(v) \leq w$.

Proof. Let us examine the contribution of an edge $e \in E(H)$ with value $s(e)$ to the sum. For every vertex $v \in e$, edge e increases the value $x(v)$ by at most $s(e)$ (the contribution may be less than $s(e)$, since we defined $x(v)$ to be at most 1). Thus the total contribution of edge e is at most

$$s(e) \cdot \sum_{v \in e} c(v) = s(e) \cdot \sum_{v \in e} \partial b_\pi(v) \leq s(e) \cdot \sum_{v \in e} \partial b_{\pi,e}(v) = s(e)b_\pi(e) = s(e)b(e) \leq s(e),$$

where the first inequality follows Prop. 5.2(5); the last equality follows from Prop. 5.2(3); the last inequality follows from the fact that b is edge dominated. Therefore, $\sum_{v \in V(H)} x(v)c(v) \leq \sum_{e \in E(H)} s(e) \leq w$, proving Claim 3. \square

Let S be a set of vertices. We define \widehat{S} to be the ‘‘inneighbor closure’’ of S , that is, the set of all vertices from which a vertex of S is reachable on a directed path in D (in particular, this means that $S \subseteq \widehat{S}$).

Claim 4. For every $S \subseteq V(H)$, $\sum_{v \in \widehat{S}} c(v) = b_\pi(\widehat{S})$.

Proof. Observe that for any $v \in \widehat{S}$, every inneighbor of v is also in \widehat{S} , hence $N_\pi^-(v) \subseteq \widehat{S}$. Therefore, $\partial b_{\pi, \widehat{S}}(v) = \partial b_\pi(v) = c(v)$ and Claim 4 follows. \square

Let $S(t)$ be the set of all vertices $v \in V(H)$ for which $t \in \iota(v)$. Observe that for every $0 \leq t \leq 1$, the set $S(t)$ (and hence $\widehat{S}(t)$) separates X from Y . We use an averaging argument to show that there is a $0 \leq t \leq 1$ for which $b_\pi(\widehat{S}(t))$ is $O(w)$. As $b^*(\widehat{S}(t)) \leq b_\pi(\widehat{S}(t))$ by definition, the set $\widehat{S}(t)$ satisfies the requirement of the lemma.

If we are able to show that $\int_0^1 b_\pi(\widehat{S}(t)) dt = O(w)$, then the existence of the required t clearly follows. Let $I_v(t) = 1$ if $v \in \widehat{S}(t)$ and let $I_v(t) = 0$ otherwise. If $I_v(t) = 1$, then there is a path P in D from v to a member of $S(t)$. By Claim 2, the width of this path is at most

$16x(v)$, thus $t \in [d(v) - 16x(v), d(v) + 15x(v)]$. Therefore, $\int_0^1 I_v(t)dt \leq 31x(v)$. Now we have

$$\begin{aligned} \int_0^1 b_\pi(\widehat{S}(t))dt &= \int_0^1 \sum_{v \in \widehat{S}(t)} c(v)dt = \int_0^1 \sum_{v \in V(H)} c(v)I_v(t)dt \\ &= \sum_{v \in V(H)} c(v) \int_0^1 I_v(t)dt \leq 31 \sum_{v \in V(H)} x(v)c(v) \leq 31w \end{aligned}$$

(we used Claim 4 in the first equality and Claim 3 in the last inequality). \square

Although it is not used in this paper, we can prove the converse of Theorem 5.7 in a very simple way.

THEOREM 5.8. *Let H be a hypergraph, and let $X, Y \subseteq V(H)$ be two sets of vertices. Suppose that for every edge-dominated monotone submodular function b on H with $b(\emptyset) = 0$, there is an (X, Y) -separator S with $b(S) \leq w$. Then there is a fractional (X, Y) -separator of weight at most w .*

PROOF. If there is no fractional (X, Y) -separator of weight at most w , then by LP duality, there is an (X, Y) -flow F of value greater than w . Let $b(Z)$ be defined as the total weight of the paths in F intersecting Z ; it is easy to see that f is a monotone submodular function, and since F is a flow, $b(e) \leq 1$ for every $e \in E(H)$. Thus by assumption, there is an (X, Y) -separator S with $b(S) \leq w$. However, every $X - Y$ path of F intersects (X, Y) -separator S , which implies $b(S) > w$, a contradiction. \square

The problem of finding a small separator in the sense of Theorem 5.7 might seem related to submodular function minimization at a first look. We close this section by pointing out that finding an (A, B) -separator S with $b(S)$ small for a given submodular function b is *not* an instance of submodular function minimization, and hence the well-known algorithms (see [Iwata 2008; Iwata et al. 2001; Schrijver 2000]) cannot be used for this problem. If a submodular function $g(X)$ describes the weight of the *boundary* of X , then finding a small (A, B) -separator is equivalent to minimizing $g(X)$ subject to $A \subseteq X$, $X \cap B = \emptyset$, which can be expressed as an instance of submodular function minimization (and hence solvable in polynomial time). In our case, however, $b(S)$ is the weight of S itself, which means that we have to minimize $g(S)$ subject to S being an (A, B) -separator and this latter constraint cannot be expressed in the framework of submodular function minimization. A possible workaround is to define $\delta(X)$ as the neighborhood of X (the set of vertices outside X adjacent to X) and $b'(X) := b(\delta(S))$; now minimizing $b'(X)$ subject to $A \subseteq X \cup \delta(X)$, $X \cap B = \emptyset$ is the same as finding an (X, Y) -separator S minimizing $b(S)$. However, the function b' is not necessarily a submodular function in general. Therefore, transforming b to b' this way does not lead to a polynomial-time algorithm using submodular function minimization. In fact, it is quite easy to show that finding an (A, B) -separator S with $b(S)$ minimum possible can be an NP-hard problem even if b is a submodular function of very simple form.

THEOREM 5.9. *Given a graph G , subsets of vertices X, Y , and collection \mathcal{S} of subsets of vertices, it is NP-hard to find an (X, Y) -separator that intersects the minimum number of members of \mathcal{S} .*

PROOF. The proof is by reduction from 3-COLORING. Let H be a graph with n vertices and m edges; we identify the vertices of H with the integers from 1 to n . We construct a graph G consisting of $3n + 2$ vertices, vertex sets X, Y , and a collection \mathcal{S} of $6m$ sets such that there is an (X, Y) -separator S in G intersecting at most $3m$ members of \mathcal{S} if and only if H is 3-colorable.

The graph G consists of two vertices x, y , and for every $1 \leq i \leq n$, a path $xv_{i,1}v_{i,2}v_{i,3}y$ of length 4 connecting x and y . The collection \mathcal{S} is constructed such that for every edge $ij \in E(H)$ and $1 \leq a, b \leq 3$, $a \neq b$, there is a corresponding set $\{v_{i,a}, v_{j,b}, x, y\}$. Let $X := \{x\}$ and $Y := \{y\}$. Observe that the set $\{v_{i,a}, v_{j,b}\}$ intersects exactly 3 sets of \mathcal{S} if $a \neq b$ and exactly 4 sets of \mathcal{S} if $a = b$.

Let $c : V(H) \rightarrow \{1, 2, 3\}$ be a 3-coloring of H . The set $S = \{v_{i,c(i)} \mid 1 \leq i \leq n\}$ is clearly an (X, Y) -separator. For every $ij \in E(H)$, separator S intersects only 3 of the 6 sets $\{v_{i,a}, v_{i,b}, x, y\}$. Therefore, S intersects exactly $3m$ members of \mathcal{S} .

Consider now an (X, Y) -separator S intersecting at most $3m$ members of \mathcal{S} . Since every member of \mathcal{S} contains both x and y , it follows that $x, y \notin S$. Thus S has to contain at least one internal vertex of every path $xv_{i,1}v_{i,2}v_{i,3}y$. For every $1 \leq i \leq n$, let us fix a vertex $v_{i,c(i)} \in S$. We claim that c is a 3-coloring of H . For every $ij \in E(H)$, S intersects at least 3 of the sets $\{v_{i,a}, v_{i,b}, x, y\}$, and intersects 4 of them if $c(i) = c(j)$. Thus the assumption that S intersects at most $3m$ members of \mathcal{S} immediately implies that c is a proper 3-coloring. \square

5.3. Obtaining a highly connected set

The following lemma is the same as the main result of Section 5 (Theorem 5.1) we are trying to prove, with the exception that b -width is replaced by b^* -width. By Prop 5.2(2), $b^*(S) \geq b(S)$ for every set $S \subseteq V(H)$, thus b^* -width is not less than b -width. Therefore, the following is actually a stronger statement and immediately implies Theorem 5.1.

LEMMA 5.10. *For every sufficiently small constant $\lambda > 0$, the following holds. Let b be an edge-dominated monotone submodular function of H with $b(\emptyset) = 0$. If the b^* -width of H is greater than $\frac{3}{2}(w + 1)$, then $\text{con}_\lambda(H) \geq w$.*

PROOF. Suppose that $\lambda < 1/c$, where c is the universal constant of Lemma 5.7 hidden by the big-O notation. Suppose that $\text{con}_\lambda(H) < w$, that is, there is no fractional independent set μ and (μ, λ) -connected set W with $\mu(W) \geq w$. We show that H has a tree decomposition of b^* -width at most $\frac{3}{2}(w + 1)$, or more precisely, we show the following stronger statement:

For every subhypergraph H' of H and every $W_0 \subseteq V(H')$ with $b^*(W_0) \leq w + 1$, there is a tree decomposition of H' having b^* -width at most $\frac{3}{2}(w + 1)$ such that W_0 is contained in one of the bags.

We prove this statement by induction on $|V(H')|$. If $b^*(V(H')) \leq \frac{3}{2}(w + 1)$, then a decomposition consisting of a single bag proves the statement. Otherwise, let W be a superset of W_0 such that $w \leq b^*(W) \leq w + 1$; let us choose a W that is inclusionwise maximal with respect to this property. Observe that there has to be at least one such set: from the fact that $b^*(v) \leq 1$ for every vertex v and from Prop. 5.2(6), we know that adding a vertex increases $b^*(W)$ by at most 1. Since $b^*(V(H')) \geq \frac{3}{2}(w + 1)$, by adding vertices to W_0 in an arbitrary order, we eventually find a set W with $b^*(W) \geq w$, and the first such set satisfies $b^*(W) \leq w + 1$ as well.

Let π be an ordering of $V(H')$ such that $b_\pi(W) = b^*(W)$. As in Lemma 5.3, let us define the fractional independent set μ by $\mu(v) := \partial b_{\pi, W}(v)$ if $v \in W$ and $\mu(v) = 0$ otherwise. Clearly, we have $\mu(W) = b_\pi(W) = b^*(W) \geq w$.

By assumption, W is not (μ, λ) -connected, hence there are disjoint sets $A, B \subseteq W$ and a fractional (A, B) -separator of weight less than $\lambda \cdot \min\{\mu(A), \mu(B)\}$. Thus by Theorem 5.7, there is an (A, B) -separator $S \subseteq V(H')$ with $b^*(S) < c \cdot \lambda \cdot \min\{\mu(A), \mu(B)\} < \min\{\mu(A), \mu(B)\} \leq \mu(W)/2 \leq (w + 1)/2$ (the second inequality follows from the fact that A and B are disjoint subsets of W). Let C_1, \dots, C_r be the connected components of $H' \setminus S$; by Lemma 5.6, $b^*((C_i \cap W) \cup S) < \mu(W) = b_\pi(W) = b^*(W) \leq w + 1$ for every $1 \leq i \leq r$. As $b^*(V(H')) \geq \frac{3}{2}(w + 1)$ and $b^*(S) \leq (w + 1)/2$, it is not possible that $S = V(H')$, hence $r > 0$. It is not possible that $r = 1$ either: $(C_1 \cap W) \cup S$ is a proper superset of W with

b^* -value strictly less than $b^*(W) \leq w + 1$, and (as $b^*(V(H')) \geq \frac{3}{2}(w + 1)$) we could find a set between $(C_1 \cap W) \cup S$ and $V(H')$ contradicting the maximality of the choice of W . Thus $r \geq 2$, which means that each hypergraph $H'_i := H'[C_i \cup S]$ has strictly fewer vertices than H' for every $1 \leq i \leq r$.

By the induction hypothesis, each H'_i has a tree decomposition \mathcal{T}_i having b^* -width at most $\frac{3}{2}(w + 1)$ such that $W_i := (C_i \cap W) \cup S$ is contained in one of the bags. Let B_i be the bag of \mathcal{T}_i containing W_i . We build a tree decomposition \mathcal{T} of H by joining together the tree decompositions $\mathcal{T}_1, \dots, \mathcal{T}_r$: let $B_0 := W_0 \cup S$ be a new bag that is adjacent to bags B_1, \dots, B_r . It can be easily verified that \mathcal{T} is indeed a tree decomposition of H . Furthermore, by Prop. 5.2(6), $b^*(B_0) \leq b^*(W_0) + b^*(S) < w + 1 + (w + 1)/2 = \frac{3}{2}(w + 1)$ and by the assumptions on $\mathcal{T}_1, \dots, \mathcal{T}_r$, every other bag has b^* value at most $\frac{3}{2}(w + 1)$. \square

6. FROM HIGHLY CONNECTED SETS TO EMBEDDINGS

The main result of this section is showing that the existence of highly connected sets implies that the hypergraph has large embedding power. Recall from Section 2 that W is a (μ, λ) -connected set for some $\lambda > 0$ and fractional independent set μ if for every disjoint $X, Y \subseteq W$, the minimum weight of a fractional (X, Y) -separator is at least $\lambda \cdot \{\mu(X), \mu(Y)\}$. We denote by $\text{con}_\lambda(H)$ the maximum value of $\mu(W)$ taken over every fractional independent set μ and (μ, λ) -connected set W . Recall also that the edge depth of an embedding ϕ of G into H is the maximum of $\sum_{v \in V(G)} |\phi(v) \cap e|$, taken over every $e \in E(H)$.

THEOREM 6.1. *For every sufficiently small $\lambda > 0$ and hypergraph H , there is a constant $m_{H,\lambda}$ such that every graph G with $m \geq m_{H,\lambda}$ edges has an embedding into H with edge depth $O(m/(\lambda^{\frac{3}{2}} \text{con}_\lambda^{\frac{1}{4}}(H)))$. Furthermore, there is an algorithm that, given G, H , and λ , produces such an embedding in time $f(H, \lambda)n^{O(1)}$.*

In other words, Theorem 6.1 gives a lower bound on the embedding power of H :

COROLLARY 6.2. *For every sufficiently small $\lambda > 0$ and hypergraph H , $\text{emb}(H) = \Omega(\lambda^{\frac{3}{2}} \text{con}_\lambda^{\frac{1}{4}}(H))$.*

Theorem 6.1 is stated in algorithmic form, since the reduction in the hardness result of Section 7 needs to find such embeddings. For the proof, our strategy is similar to the embedding result of [Marx 2010b]: we show that a highly connected set implies that a uniform concurrent flow exists, the paths appearing in the uniform concurrent flow can be used to embed (a blowup of) the line graph of a complete graph, and every graph has an appropriate embedding in the line graph of a complete graph. To make this strategy work, we need generalizations of concurrent flows, multicuts, and multicommodity flows in our hypergraph setting and we need to obtain results that connect these concepts to highly connected sets. Some of these results are similar in spirit to the $O(\sqrt{n})$ -approximation algorithms appearing in the combinatorial optimization literature [Gupta 2003; Hajiaghayi and Räcke 2006; Agarwal et al. 2007]. However, those approximation algorithms are mostly based on clever rounding of fractional solutions, while in our setting rounding is not an option: as discussed in Section 5, the existence of a fractional (X, Y) -separator of small weight does not imply the existence of a small integer separator. Thus we have to work directly with the fractional solution and use the properties of the highly connected set.

It turns out that the right notion of uniform concurrent flow for our purposes is a collection of flows that connect cliques: that is, a collection $F_{i,j}$ ($1 \leq i < j \leq k$) of compatible flows, each of value ϵ , such that $F_{i,j}$ is a (K_i, K_j) -flow, where K_1, \dots, K_k are disjoint cliques. Thus our first goal is to find a highly connected set that can be partitioned into k cliques in an appropriate way.

6.1. Highly connected sets with cliques

Let $(X_1, Y_1), \dots, (X_k, Y_k)$ be pairs of vertex sets such that the minimum weight of a fractional (X_i, Y_i) -separator is s_i . Analogously to multicut problems in combinatorial optimization, we investigate weight assignments that *simultaneously* separate all these pairs. Clearly, the minimum weight of such an assignment is at least the minimum of the s_i 's and at most the sum of the s_i 's. The following lemma shows that in a highly connected set, such a simultaneous separator cannot be very efficient: roughly speaking, its weight is at least the square root of the sum of the s_i 's.

LEMMA 6.3. *Let μ be a fractional independent set in hypergraph H and let W be a (μ, λ) -connected set for some $0 < \lambda \leq 1$. Let $(X_1, \dots, X_k, Y_1, \dots, Y_k)$ be a partition of W , let $w_i := \min\{\mu(X_i), \mu(Y_i)\} \geq 1/2$, and let $w := \sum_{i=1}^k w_i$. Let $s : E(H) \rightarrow \mathbb{R}^+$ be a weight assignment of total weight p such that s is a fractional (X_i, Y_i) -separator for every $1 \leq i \leq k$. Then $p \geq (\lambda/7) \cdot \sqrt{w}$.*

PROOF. Let us define the function s' by $s'(e) = 6s(e)$ and let $x(v) := \sum_{e \in E(H), v \in e} s'(e)$. We define the distance $d(u, v)$ to be the minimum of $\sum_{r \in P} x(r)$, taken over all paths P from u to v . It is clear that the triangle inequality holds, i.e., $d(u, v) \leq d(u, z) + d(z, v)$ for every $u, v, z \in V(H)$. If s covers every $u - v$ path, then $d(u, v) \geq 6$: every edge e intersecting a $u - v$ path P contributes at least $s'(e)$ to the sum $\sum_{r \in P} x(r)$ (as e can intersect P in more than one vertices, e can increase the sum by more than $s'(e)$). On the other hand, we claim that if $d(u, v) \geq 2$, then s' covers every $u - v$ path. Clearly, it is sufficient to verify this for minimal paths. Such a path P can intersect an edge e at most twice, hence e contributes at most $2s'(e)$ to the sum $\sum_{r \in P} x(r) \geq 2$, implying that the edges intersecting P have total weight at least 1 in s' .

Suppose for contradiction that $p < (\lambda/7) \cdot \sqrt{w}$, that is, $w > 49p^2/\lambda^2$. As s is an (X_i, Y_i) -separator, we have that $p \geq 1$. Let $A := \emptyset$ and $B := \bigcup_{i=1}^k (X_i \cup Y_i)$. Note that $\mu(B) \geq 2 \sum_{i=1}^k w_i = 2w$. We will increase A and decrease B while maintaining the invariant condition that the distance of A and B is at least 2 in d . Let T be the smallest integer such that $\sum_{i=1}^T w_i > 6p/\lambda$; if there is no such T , then $w \leq 6p/\lambda$, a contradiction. As $w_i \geq 1/2$ for every i , it follows that $T \leq 12p/\lambda + 1 \leq 13p/\lambda$ (since $p \geq 1$ and $\lambda \leq 1$).

For $i = 1, 2, \dots, T$, we perform the following step. Let X'_i (resp., Y'_i) be the set of all vertices of W that are at distance at most 2 from X_i (resp., Y_i). As the distance of X_i and Y_i is at least 6, by the triangle inequality the distance of X'_i and Y'_i is at least 2, hence s' is a fractional (X'_i, Y'_i) -separator. Since W is (μ, λ) -connected and s' is an assignment of weight $6p$, we have $\min\{\mu(X'_i), \mu(Y'_i)\} \leq 6p/\lambda$. If $\mu(X'_i) \leq 6p/\lambda$, then let us put X_i (note: *not* X'_i) into A and let us remove X'_i from B . The set X'_i , which we remove from B , contains all the vertices that are at distance at most 2 from any new vertex in A , hence it remains true that the distance of A and B is at least 2. Similarly, if $\mu(X'_i) > 6p/\lambda$ and $\mu(Y'_i) \leq 6p/\lambda$, then let us put Y_i into A and let us remove Y'_i from B . Note that we may put a vertex into A even if it was removed from B in an earlier step.

In the i -th step of the procedure, we increase $\mu(A)$ by at least w_i (as $\mu(X_i), \mu(Y_i) \geq w_i$ and these sets are disjoint from the sets already contained in A) and $\mu(B)$ is decreased by at most $6p/\lambda$. Thus at the end of the procedure, we have $\mu(A) \geq \sum_{i=1}^T w_i > 6p/\lambda$ and

$$\mu(B) \geq 2w - T \cdot 6p/\lambda > 98p^2/(\lambda^2) - (13p/\lambda)(6p/\lambda) > 6p/\lambda,$$

that is, $\min\{\mu(A), \mu(B)\} > 6p/\lambda$. By the invariant condition, the distance of A and B is at least 2, thus s' is a fractional (A, B) -separator of weight exactly $6p$, contradicting the assumption that W is (μ, λ) -connected. \square

In the rest of the section, we need a more constrained notion of flow, where the endpoints "respect" a particular fractional independent set. Let μ_1, μ_2 be fractional independent sets

of hypergraph H and let $X, Y \subseteq V(H)$ be two (not necessarily disjoint) sets of vertices. A (μ_1, μ_2) -demand (X, Y) -flow is an (X, Y) -flow F such that for each $x \in X$, the total weight of the paths in F having first endpoint x is at most $\mu_1(x)$, and similarly, the total weight of the paths in F having second endpoint $y \in Y$ is at most $\mu_2(y)$. Note that there is no bound on the weight of the paths going through an $x \in X$, we only bound the paths whose first/second endpoint is x . The definition is particularly delicate if X and Y are not disjoint, in this case, a vertex $z \in X \cap Y$ can be the first endpoint of some paths and the second endpoint of some other paths, or it can be even both the first and second endpoint of a path of length 0. We use the abbreviation μ -demand for (μ, μ) -demand.

The following lemma shows that if a flow connects a set U with a highly connected set W , then U is highly connected as well (“ W can be moved to U ”). This observation will be used in the proof of Lemma 6.5, where we locate cliques and show that their union is highly connected, since there is a flow that connects the cliques to a highly connected set.

LEMMA 6.4. *Let H be a hypergraph, μ_1, μ_2 fractional independent sets, and $W \subseteq V(H)$ a (μ_1, λ) -connected set for some $0 < \lambda \leq 1$. Suppose that $U \subseteq V(H)$ is a set of vertices and F is a (μ_1, μ_2) -demand (W, U) -flow of value $\mu_2(U)$. Then U is $(\mu_2, \lambda/6)$ -connected.*

PROOF. Suppose that there are disjoint sets $A, B \subseteq U$ and a fractional (A, B) -separator s of weight $w < (\lambda/6) \cdot \min\{\mu_2(A), \mu_2(B)\}$. (Note that this means $\mu_2(A), \mu_2(B) > 6w/\lambda \geq 6w$.) For a path P , let $s(P) = \sum_{e \in E(H), e \cap P \neq \emptyset} s(e)$ be the total weight of the edges intersecting P . Let $A' \subseteq W$ (resp., $B' \subseteq W$) contain a vertex $v \in W$ if there is a path P in F with first endpoint v and second endpoint in A (resp., B) and $s(P) \leq 1/3$. If $A' \cap B' \neq \emptyset$, then it is clear that there is a path P with $s(P) \leq 2/3$ connecting a vertex of A and a vertex of B via a vertex of $A' \cap B'$, a contradiction. Thus we can assume that A' and B' are disjoint.

Since F is a flow and s has weight w , the total weight of the paths in F with $s(P) \geq 1/3$ is at most $3w$. As the value of F is exactly $\mu_2(U)$, the total weight of the paths in F with second endpoint in A is exactly $\mu_2(A)$. If $s(P) \leq 1/3$ for such a path, then its first endpoint is in A' by definition. Therefore, the total weight of the paths in F with first endpoint in A' is at least $\mu_2(A) - 3w$, which means that $\mu_1(A') \geq \mu_2(A) - 3w \geq \mu_2(A)/2$. Similarly, we have $\mu_1(B') \geq \mu_2(B)/2$. Since W is (μ_1, λ) -connected and s is an assignment with weight less than $(\lambda/6) \cdot \min\{\mu_2(A), \mu_2(B)\} \leq (\lambda/3) \cdot \min\{\mu_1(A'), \mu_1(B')\}$, there is an $A' - B'$ path P with $s(P) < 1/3$. Now the concatenation of an $A' - A$ path P_A having $s(P_A) \leq 1/3$, the path P , and a $B' - B$ path P_B having $s(P_B) \leq 1/3$ forms an $A - B$ path that is not covered by s , a contradiction. \square

A μ -demand *multicommodity flow* between pairs $(A_1, B_1), \dots, (A_r, B_r)$ is a set F_1, \dots, F_r of compatible flows such that F_i is a μ -demand (A_i, B_i) -flow (recall that a set of flows is compatible if their sum is also a flow, that is, does not violate the edge constraints). The *value* of a multicommodity flow is the sum of the values of the r flows. Let $A = \bigcup_{i=1}^r A_i$, $B = \bigcup_{i=1}^r B_i$, and let us restrict our attention to the case when $(A_1, \dots, A_r, B_1, \dots, B_r)$ is a partition of $A \cup B$. In this case, the maximum value of a μ -demand multicommodity flow between pairs $(A_1, B_1), \dots, (A_r, B_r)$ can be expressed as the optimum values of the primal and dual linear programs in Figure 5.

The following lemma shows that if $\text{con}_\lambda(H)$ is sufficiently large, then there is a highly connected set that has the additional property that it is the union of k cliques K_1, \dots, K_k with $\mu(K_i) \geq 1/2$ for every clique. The high-level idea of the proof is the following. Take a (μ, λ) -connected set W with $\mu(W) = \text{con}_\lambda(H)$ and find a large multicommodity flow between some pairs $(A_1, B_1), \dots, (A_r, B_r)$ in W . Consider the dual solution y . By complementary slackness, every edge with nonzero value in y covers exactly 1 unit of the multicommodity flow. If most of the weight of the dual solution is on the edge variables, then we can choose k edges that cover at least $\Omega(k)$ units of flow. These edges are connected to

Primal LP	Dual LP
maximize $\sum_{i=1}^r \sum_{\substack{u \in A_i, v \in B_i \\ P \in \mathcal{P}_{uv}}} x(P)$	minimize $\sum_{e \in E(H)} y(e) + \sum_{u \in A} \mu(u)y(u) + \sum_{v \in B} \mu(v)y(v)$
s.t. $\sum_{i=1}^r \sum_{\substack{u \in A_i, v \in B_i \\ P \in \mathcal{P}_{uv}, P \cap e \neq \emptyset}} x(P) \leq 1 \quad \forall e \in E(H)$	s.t. $\sum_{\substack{e \in E(H), \\ e \cap P \neq \emptyset}} y(e) + y(u) + y(v) \geq 1$
$\sum_{v \in B_i, P \in \mathcal{P}_{uv}} x(P) \leq \mu(u) \quad \forall 1 \leq i \leq r, u \in A_i$	$\forall 1 \leq i \leq r, \\ u \in A_i, v \in B_i, \\ P \in \mathcal{P}_{uv}$
$\sum_{u \in A_i, P \in \mathcal{P}_{uv}} x(P) \leq \mu(v) \quad \forall 1 \leq i \leq r, v \in B_i$	$y(e) \geq 0 \quad \forall e \in E(H)$
$x(P) \geq 0 \quad \begin{array}{l} \forall 1 \leq i \leq r, \\ u \in A_i, v \in B_i, \\ P \in \mathcal{P}_{uv} \end{array}$	$y(u) \geq 0 \quad \forall u \in A$
	$y(v) \geq 0 \quad \forall v \in B$

Fig. 5. Primal and dual linear programs for μ -demand multicommodity flow between pairs $(A_1, B_1), \dots, (A_r, B_r)$. We denote by \mathcal{P}_{uv} the set of all $u - v$ paths.

W by a flow, and therefore by Lemma 6.4 the union of these edges is also highly connected and obviously can be partitioned into a small number cliques.

There are two things that can go wrong with this argument. First, it can happen that the dual solution assigns most of the weight to the vertex variables $y(u), y(v)$ ($u \in A, v \in B$). The cost of covering the $A_i - B_i$ paths using vertex variables only is $\min\{\mu(A_i), \mu(B_i)\}$, thus this case is only possible if the value of the dual (and hence the primal) solution is close to $\sum_{i=1}^r (\min\{\mu(A_i), \mu(B_i)\})$. To avoid this situation, we want to select the pairs (A_i, B_i) such that they are only “moderately connected”: there is a fractional (A_i, B_i) -separator of weight $2\lambda \min\{\mu(A_i), \mu(B_i)\}$, that is, at most twice the minimum possible. This means that the weight of the dual solution is at most $2\lambda \sum_{i=1}^r (\min\{\mu(A_i), \mu(B_i)\})$, which is much less than $\sum_{i=1}^r (\min\{\mu(A_i), \mu(B_i)\})$ (if λ is small). If we are not able to find sufficiently many such pairs, then we argue that a larger highly connected set can be obtained by scaling μ by a factor of 2. More precisely, we show that there is a large subset $W' \subseteq W$ that is $(2\mu, \lambda)$ -connected and $2\mu(W') > \text{con}_\lambda(H)$, a contradiction (a technical difficulty here that we have to make sure first that 2μ is also a fractional independent set).

The second problem we have to deal with is that the value of the dual solution can be so small that we find a very small set of edges that already cover a large fraction of the multicommodity flow. However, we can use Lemma 6.3 to argue that a weight assignment on the edges that covers a large multicommodity flow in a (μ, λ) -connected set cannot have very small weight.

LEMMA 6.5. *Let H be a hypergraph and let $0 < \lambda < 1/16$ be a constant. Then there is fractional independent set μ , a $(\mu, \lambda/6)$ -connected set W , and a partition (K_1, \dots, K_k) of W such that $k = \Omega(\lambda \sqrt{\text{con}_\lambda(H)})$, and for every $1 \leq i \leq k$, K_i is a clique with $\mu(K_i) \geq 1/2$.*

PROOF. Let k be the largest integer such that $\text{con}_\lambda(H) \geq 3T + 2k$ holds, where $T := (56/\lambda)^2 \cdot k^2$; it is clear that $k = \Omega(\lambda \sqrt{\text{con}_\lambda(H)})$. Let μ_0 be a fractional independent set and W_0 be a (μ_0, λ) -connected set with $\mu_0(W_0) = \text{con}_\lambda(H)$. We can assume that $\mu_0(v) > 0$ if and only if $v \in W_0$. This also implies that W_0 is in one connected component of H .

Highly loaded edges. First, we want to modify μ_0 such that there is no edge e with $\mu_0(e) \geq 1/2$. The following claim shows that we can achieve this by restricting μ_0 to an appropriate subset W of W_0 .

Claim 1. There is a subset $W \subseteq W_0$ such that $\mu_0(W) \geq \text{con}_\lambda(H) - k$ and $\mu_0(e \cap W) < 1/2$ for every edge e .

Proof.

Let us choose edges g_1, g_2, \dots as long as possible with the requirement $\mu_0(K_i) \geq 1/2$ for $K_i := (g_i \cap W_0) \setminus \bigcup_{j=1}^{i-1} K_j$. If we can select at least k such edges, then the cliques K_1, \dots, K_k satisfy the requirements of Lemma 6.5 and we are done. Indeed, $W' := \bigcup_{i=1}^k K_i \subseteq W_0$ is a (μ_0, λ) -connected set, $\mu_0(K_i) \geq 1/2$, and (K_1, \dots, K_k) is a partition of W' into cliques.

Thus we can assume that the selection of the edges stops at edge g_t for some $t < k$. Let $W := W_0 \setminus \bigcup_{i=1}^t K_i$. Observe that there is no edge $e \in E(H)$ with $\mu_0(e \cap W) \geq 1/2$, as in this case the selection of the edges could be continued with $g_{t+1} := e$. Furthermore, we have $\mu_0(W) = \mu_0(W_0 \setminus \bigcup_{i=1}^t K_i) > \mu_0(W_0) - k = \text{con}_\lambda(H) - k$, as required. \lrcorner

Moderately connected pairs. Let us define μ such that $\mu(v) = 2\mu_0(v)$ if $v \in W$ and $\mu(v) = 0$ otherwise. By Claim 1, μ is a fractional independent set. The set W is (μ_0, λ) -connected (recall that a subset of (μ_0, λ) -connected is also (μ_0, λ) -connected). However, W is not necessarily (μ, λ) -connected. In the next step, we find a large collection of pairs (A_i, B_i) that violate (μ, λ) -connectivity. Informally, we can say that these pairs (A_i, B_i) are “moderately connected”: denoting $w_i = \min\{\mu(A_i), \mu(B_i)\}$, the minimum value of a fractional (A_i, B_i) -separator for such a pair is less than λw_i (because the pair (A_i, B_i) violates (μ, λ) -connectivity), but at least $\lambda w_i/2 = \lambda \min\{\mu_0(A_i), \mu_0(B_i)\}$ (because W is (μ_0, λ) -connected).

Claim 2. There are disjoint sets $A_1, B_1, \dots, A_r, B_r \subseteq W$ such that for every $1 \leq i \leq r$ there is a fractional (A_i, B_i) -separator with weight less than λw_i for $w_i := \min\{\mu(A_i), \mu(B_i)\}$ and $w := \sum_{i=1}^r w_i \geq T$.

Proof. Let us greedily select a maximal collection of pairs $(A_1, B_1), \dots, (A_r, B_r)$ with the property that there is a fractional (A_i, B_i) -separator with weight less than λw_i for $w_i := \min\{\mu(A_i), \mu(B_i)\}$. Note that every fractional separator has value at least 1 (as W is in a single component of H), thus $\lambda w_i > 1$ holds, implying $w_i > 1/\lambda > 1$. We can assume that $\mu(A_i), \mu(B_i) \leq w_i + 1$: if, say, $\mu(A_i) > \mu(B_i) + 1$, then removing an arbitrary vertex of A_i decreases $\mu(A_i)$ by at most one (as μ is a fractional independent set) without changing $\min\{\mu(A_i), \mu(B_i)\}$, hence there would be a smaller pair of sets with the required properties. Therefore, we have $2w_i \leq \mu(A_i \cup B_i) \leq 2w_i + 1 \leq 3w_i$ for every $1 \leq i \leq r$.

Suppose for contradiction that $w := \sum_{i=1}^r w_i < T$. Let $W' := W \setminus \bigcup_{i=1}^r (A_i \cup B_i)$. As $\mu(\bigcup_{i=1}^r (A_i \cup B_i)) \leq \sum_{i=1}^r 3w_i = 3w < 3T$, we have $\mu(W') > \mu(W) - 3T = 2\mu_0(W) - 3T \geq 2\text{con}_\lambda(H) - 2k - 3T \geq \text{con}_\lambda(H)$. Since the greedy selection stopped, there is no fractional (A', B') -separator of value less than $\lambda \cdot \min\{\mu(A'), \mu(B')\}$ for any disjoint $A', B' \subseteq W'$, that is, W' is (μ, λ) -connected with $\mu(W') > \text{con}_\lambda(H)$, contradicting the definition of $\text{con}_\lambda(H)$. \lrcorner

Finding a multicommodity flow. Let $(A_1, B_1), \dots, (A_r, B_r)$ be as in Claim 2. Since there is a fractional (A_i, B_i) -separator of value less than λw_i , the maximum value of a μ -demand multicommodity flow between pairs $(A_1, B_1), \dots, (A_r, B_r)$ is less than λw . Let y be an optimum dual solution; we give a lower bound on the total weight of the edge variables.

Claim 3. $\sum_{e \in E(H)} y(e) \geq 2k$.

Proof. Let $A := \bigcup_{i=1}^r A_i$ and $B := \bigcup_{i=1}^r B_i$. Let $A^* := \{u \in A \mid y(u) \leq 1/4\}$, $B^* := \{v \in B \mid y(v) \leq 1/4\}$, $A_i^* = A_i \cap A^*$, $B_i^* = B_i \cap B^*$, and $w_i^* = \min\{\mu(A_i^*), \mu(B_i^*)\}$. For each i ,

the value of w_i^* is either at least $w_i/2$, or less than that. Assume without loss of generality that there is a $1 \leq r^* \leq r$ such that $w_i^* \geq w_i/2$ if and only if $i \leq r^*$. Let $w^* = \sum_{i=1}^{r^*} w_i^*$.

We claim that $w^* \geq w/4$. Note that $w_i^* < w_i/2$ means that either $\mu(A_i^*) < w_i/2$ or $\mu(B_i^*) < w_i/2$; as $\mu(A_i), \mu(B_i) \geq w_i$, this is only possible if $\mu(A_i \setminus A^*) + \mu(B_i \setminus B^*) > w_i/2$. Suppose first that $\sum_{i=r^*+1}^r w_i > w/2$. This would imply

$$\mu((A \setminus A^*) \cup (B \setminus B^*)) \geq \sum_{i=r^*+1}^r (\mu(A_i \setminus A^*) + \mu(B_i \setminus B^*)) > \sum_{i=r^*+1}^r w_i/2 > w/4.$$

However, $y(u) > 1/4$ for every $u \in (A \setminus A^*) \cup (B \setminus B^*)$, thus $\sum_{v \in A \cup B} \mu(v)y(v) \geq \mu((A \setminus A^*) \cup (B \setminus B^*)) / 4 \geq w/16 > \lambda w$ (since $\lambda < 1/16$), a contradiction with the assumption that the optimum is at most λw . Thus we can assume that $\sum_{i=r^*+1}^r w_i \leq w/2$ and hence $\sum_{i=1}^{r^*} w_i \geq w/2$. Together with $w_i^* \geq w_i/2$ for every $1 \leq i \leq r^*$, this implies $w^* \geq w/4$.

As $y(a), y(b) \leq 1/4$ for every $a \in A_i^*, b \in B_i^*$, it is clear that for every $A_i^* - B_i^*$ path P , the total weight of the edges intersecting P has to be at least $1/2$ in assignment y . Therefore, if we define $y^* : E(H) \rightarrow \mathbb{R}^+$ by $y^*(e) = 2y(e)$ for every $e \in E(H)$, then y^* covers every $A_i^* - B_i^*$ path. Let $W^* = \bigcup_{i=1}^{r^*} (A_i^* \cup B_i^*)$. We use Lemma 6.3 for the (μ, λ) -connected set W^* , the pairs $(A_1^*, B_1^*), \dots, (A_{r^*}^*, B_{r^*}^*)$, and for the weight assignment y^* . Note that $w_i^* \geq w_i/2 \geq 1/2$ for every i . It follows that the total weight of y^* on the edges is at least $(\lambda/7) \cdot \sqrt{w^*} \geq (\lambda/14) \cdot \sqrt{w}$, which means that $\sum_{e \in E(H)} y(e) \geq (\lambda/28) \cdot \sqrt{w} \geq (\lambda/28) \cdot \sqrt{T} \geq 2k$. \square

Locating the cliques. Let y be an optimum dual solution for the maximum multicommodity flow problem with pairs $(A_1, B_1), \dots, (A_r, B_r)$ and let flow F be the sum of the flows obtained from an optimum primal solution.

Claim 4. There are k pairwise-disjoint cliques K_1, \dots, K_k and a set of k subflows f_1, \dots, f_k of F , each of them having value at least $1/2$, such that every path appearing in f_i intersects K_i and is disjoint from K_j for every $j \neq i$.

Proof. Let $F^{(0)} = F$ and for $i = 1, 2, \dots$, let $F^{(i)}$ be the flow obtained from $F^{(0)}$ by removing f_1, \dots, f_i . Let $c(e, F^{(i)})$ be the total weight of the paths in $F^{(i)}$ intersecting edge e and let $C_i = \sum_{e \in E(H)} y(e)c(e, F^{(i)})$. By complementary slackness, $c(e, F^{(0)}) = 1$ for each $e \in E(H)$ with $y(e) > 0$ and hence $C_0 = \sum_{e \in E(H)} y(e) \geq 2k$.

Let us select e_i to be an edge such that $c(e_i, F^{(i-1)})$ is maximum possible and let $K_i := e_i \setminus \bigcup_{j=1}^{i-1} e_j$. Let the flow f_i contain all the paths of $F^{(i-1)}$ intersecting e_i . Observe that the paths appearing in f_i do not intersect e_1, \dots, e_{i-1} (otherwise they would be in one of f_1, \dots, f_{i-1} and hence they would no longer be in $F^{(i-1)}$), thus clique K_i intersects every path in f_i .

For every $u - v$ path P appearing in $F^{(0)}$, we get $\sum_{e \in E(H), e \cap P \neq \emptyset} y(e) + y(u) + y(v) = 1$ from complementary slackness: if the primal variable corresponding to P is nonzero, then the corresponding dual constraint is tight. In particular, this means that the total weight of the edges intersecting such a path P is at most 1 in y . As $F^{(i-1)}$ is a subflow of $F^{(0)}$, this is also true for every path P in $F^{(i-1)}$. This means that when we remove a path of weight γ from $F^{(i-1)}$ to obtain $F^{(i)}$, then the total weight of the edges e for which $c(e, F^{(i-1)})$ decreases by γ is at most 1, i.e., C_{i-1} decreases by at most γ . As only the paths intersecting e_i are removed from $F^{(i-1)}$ and the total weight of the paths intersecting e_i is at most 1, we get that $C_i \geq C_{i-1} - 1$ and hence $C_i \geq C_0 - k \geq C_0/2$ for $i \leq k$. Since $C_0 = \sum_{e \in E(H)} y(e)$ and $C_i = \sum_{e \in E(H)} y(e)c(e, F^{(i)}) \geq C_0/2$, it follows that there has to be at least one edge e with $c(e, F^{(i)}) \geq 1/2$. Thus in each step, we can select an edge e_i such that the total

Primal LP	Dual LP
<p style="text-align: center;">maximize ϵ</p> <p style="text-align: center;">s.t.</p> $\sum_{1 \leq i < j \leq k} \sum_{\substack{P \in \mathcal{P}_{i,j}, \\ P \cap e \neq \emptyset}} x(P) \leq 1 \quad \forall e \in E(H)$ $\sum_{P \in \mathcal{P}_{i,j}} x(P) \geq \epsilon \quad \forall 1 \leq i < j \leq k$ $x(P) \geq 0 \quad \forall 1 \leq i < j \leq k, \\ P \in \mathcal{P}_{i,j}$	<p style="text-align: center;">minimize $\sum_{e \in E(H)} y(e)$</p> <p style="text-align: center;">s.t.</p> $\sum_{e \in E(H), e \cap P \neq \emptyset} y(e) \geq \ell_{i,j} \quad \forall 1 \leq i < j \leq k, \\ P \in \mathcal{P}_{i,j}$ $\sum_{1 \leq i < j \leq k} \ell_{i,j} \geq 1$ $y(e) \geq 0 \quad \forall e \in E(H)$ $\ell_{i,j} \geq 0 \quad \forall 1 \leq i < j \leq k$

Fig. 6. Primal and dual linear programs for uniform concurrent flow on $W = (X_1, \dots, X_k)$. We denote by $\mathcal{P}_{i,j}$ the set of all $X_i - X_j$ paths.

weight of the paths in $F^{(i)}$ intersecting e_i is at least $1/2$, and hence the value of f_i is at least $1/2$ for every $1 \leq i \leq k$. \lrcorner

Moving the highly connected set. Let $U = \bigcup_{i=1}^k K_i$.

Claim 5. There is a fractional independent set μ' such that U is a $(\mu', \lambda/6)$ -connected set with $\mu'(K_i) \geq 1/2$ for every $1 \leq i \leq r$.

Proof. Each path P in f_i is a path with endpoints in W and intersecting K_i . Let us truncate each path P in f_i such that its first endpoint is still in W and its second endpoint is in K_i ; let f'_i be the (W, K_i) -flow obtained by truncating every path in f_i . Note that f'_i is still a flow and the sum F' of f'_1, \dots, f'_k is a (W, U) -flow. Let $\mu_1 = \mu$ and let $\mu_2(v)$ be the total weight of the paths in F' with second endpoint v . It is clear that μ_2 is a fractional independent set, $\mu_2(K_i) \geq 1/2$, and F is a (μ_1, μ_2) -demand (W, U) -flow with value $\mu_2(U)$. Thus by Lemma 6.4, U is a $(\mu_2, \lambda/6)$ -connected set with the required properties. \lrcorner

The set U , the partition (K_1, \dots, K_r) , and the fractional independent set μ' clearly satisfy the requirements of the lemma. \square

6.2. Concurrent flows and embedding

Let W be a set of vertices and let (X_1, \dots, X_k) be a partition of W . A *uniform concurrent flow of value ϵ on (X_1, \dots, X_k)* is a compatible set of $\binom{k}{2}$ flows $F_{i,j}$ ($1 \leq i < j \leq k$) where $F_{i,j}$ is an (X_i, X_j) -flow of value ϵ . The maximum value of a uniform concurrent flow on W can be expressed as the optimum values of the primal and dual linear programs in Figure 6. Intuitively, the dual linear program expresses that the “distance” of X_i and X_j is at least $\ell_{i,j}$ (where distance is measured as the minimum total weight of the edges intersected by an $X_i - X_j$ path) and the sum of these $\binom{k}{2}$ distances is at least 1.

If H is connected, then the maximum value of a uniform concurrent flow on (X_1, \dots, X_k) is at least $1/\binom{k}{2} = \Omega(k^{-2})$: if each of the $\binom{k}{2}$ flows has value $1/\binom{k}{2}$, then they are clearly compatible. The following lemma shows that in a (μ, λ) -connected set, if the sets X_1, \dots, X_k are cliques and $\mu(X_i) \geq 1/2$ for every i , then we can guarantee a better bound of $\Omega(k^{-\frac{3}{2}})$.

LEMMA 6.6. *Let H be a hypergraph, μ a fractional independent set of H , and $W \subseteq V(H)$ a (μ, λ) -connected set for some $0 < \lambda < 1$. Let (K_1, \dots, K_k) (for some $k \geq 1$) be a partition of W such that K_i is a clique and $\mu(K_i) \geq 1/2$ for every $1 \leq i \leq k$. Then there is a uniform concurrent flow of value $\Omega(\lambda/k^{\frac{3}{2}})$ on (K_1, \dots, K_k) .*

PROOF. Suppose that there is no uniform concurrent flow of value $\beta \cdot \lambda/k^{\frac{3}{2}}$, where $\beta > 0$ is a sufficiently small universal constant specified later. This means that the dual linear program has a solution having value less than that. Let us fix such a solution $(y, \ell_{i,j})$ of the dual linear program. In the following, for every path P , we denote by $y(P) := \sum_{e \in E(H), e \cap P \neq \emptyset} y(e)$ the total weight of the edges intersecting P . It is clear from the dual linear program that $y(P) \geq \ell_{i,j}$ for every $P \in \mathcal{P}_{i,j}$.

We construct two graphs G_1 and G_2 : the vertex set of both graphs is $\{1, \dots, k\}$ and for every $1 \leq i < j \leq k$, vertices i and j are adjacent in G_1 (resp., G_2) if and only if $\ell_{i,j} > 1/(3k^2)$ (resp., $\ell_{i,j} > 1/k^2$). Note that G_2 is a subgraph of G_1 . First we prove the following claim:

Claim 1. If the distance of u and v is at most 3 in the complement of G_1 , then u and v are not adjacent in G_2 .

Proof. Suppose that uw_1w_2v is a path of length 3 in the complement of G_1 (the same argument works for paths of length less than 3). By definition of G_1 , there is a $K_u - K_{w_1}$ path P_1 , a $K_{w_1} - K_{w_2}$ path P_2 , and a $K_{w_2} - K_v$ path P_3 such that $y(P_1), y(P_2), y(P_3) \leq 1/(3k^2)$. Since K_{w_1} and K_{w_2} are cliques, paths P_1 and P_2 touch, and paths P_2 and P_3 touch. Thus by concatenating the three paths, we can obtain a $K_u - K_v$ path P with $y(P) \leq y(P_1) + y(P_2) + y(P_3) \leq 1/k^2$, implying that u and v are not adjacent in G_2 , proving the claim. Note that the proof of this claim is the only point where we use that the K_i 's are cliques. \square

Let $y' : E(H) \rightarrow \mathbb{R}^+$ be defined by $y'(e) := 3k^2 \cdot y(e)$, thus y' has total weight less than $3\beta \cdot \lambda\sqrt{k}$. Suppose first that G_1 has a matching a_1b_1, \dots, a_mb_m of size $m = \lceil k/4 \rceil$. This means that y' covers every $K_{a_i} - K_{b_i}$ path for every $1 \leq i \leq \lceil k/4 \rceil$. Therefore, by Lemma 6.3, y' has weight at least $(\lambda/7) \cdot \sqrt{\lceil k/4 \rceil} \cdot (1/2) > 3\beta \cdot \lambda\sqrt{k}$, if β is sufficiently small, yielding a contradiction.

Thus the size of the maximum matching in G_1 is less than $\lceil k/4 \rceil$, which means that there is a vertex cover S_1 of size at most $k/2$. Let $S_2 \subseteq S_1$ contain those vertices of S_1 that are adjacent to every vertex outside S_1 in G_1 . We claim that S_2 is a vertex cover of G_2 . Suppose that there is an edge uv of G_2 for some $u, v \notin S_2$. By the definition of S_2 , either $u \notin S_1$, or there is a vertex $w_1 \notin S_1$ such that u and w_1 are not adjacent in G_1 . Similarly, either v is not in S_1 , or it is not adjacent in G_1 to some $w_2 \notin S_1$. Since vertices not in S_1 are not adjacent in G_1 (as S_1 is a vertex cover of G_1), we get that the distance of u and v is at most 3 in the complement of G_1 . Thus by the claim, u and v are not adjacent in G_2 .

Let us give an upper bound on $\sum_{1 \leq i < j \leq k} \ell_{i,j}$ by bounding $\ell_{i,j}$ separately for pairs that are adjacent in G_2 and for pairs that are not adjacent in G_2 . The number of edges in G_2 is at most $|S_2|k$ (as S_2 is vertex cover). The total weight of y , which is less than $\beta \cdot \lambda/k^{\frac{3}{2}}$, is an upper bound on any $\ell_{i,j}$. Furthermore, if i and j are not adjacent in G_2 , then we have $\ell_{i,j} \leq 1/k^2$. Therefore,

$$1 \leq \sum_{1 \leq i < j \leq k} \ell_{i,j} \leq |S_2|k \cdot \beta \cdot \lambda/k^{\frac{3}{2}} + \binom{k}{2}(1/k^2) \leq \beta \cdot \lambda|S_2|/\sqrt{k} + 1/2,$$

which implies that $|S_2| \geq \sqrt{k}/(2\beta\lambda)$. Let $A := \bigcup_{i \in S_2} K_i$ and $B := \bigcup_{i \notin S_1} K_i$; we have $\mu(A) \geq |S_2| \cdot (1/2) \geq \sqrt{k}/(4\beta\lambda)$ and $\mu(B) \geq (1/2) \cdot (k - |S_1|) \geq k/4$. As every vertex of S_2 is adjacent in G_1 with every vertex outside S_1 , assignment y' covers every $A - B$ path. However, y' has weight less than $3\beta \cdot \lambda\sqrt{k} < \min\{\sqrt{k}/(4\beta\lambda), k/4\}$ (using that $\lambda \leq 1$ and assuming that β is sufficiently small), contradicting the assumption that W is (μ, λ) -connected. \square

Intuitively, the intersection structure of the paths appearing in a uniform concurrent flow on cliques K_1, \dots, K_k is reminiscent of the edges of the complete graph on k vertices: if $\{i_1, j_1\} \cap \{i_2, j_2\} \neq \emptyset$, then every path of F_{i_1, j_1} touches every path of F_{i_2, j_2} . We use the following result from [Marx 2010b], which shows that the line graph of cliques have good embedding properties. If G is a graph and $q \geq 1$ is an integer, then the *blow up* $G^{(q)}$ is obtained from G by replacing every vertex v with a clique K_v of size q and for every edge uv of G , connecting every vertex of the clique K_u with every vertex of the clique K_v . Let L_k be the line graph of the complete graph on k vertices.

LEMMA 6.7 (MARX [2010B]). *For every $k > 1$ there is a constant $n_k > 0$ such that for every G with $|E(G)| > n_k$ and no isolated vertices, the graph G is a minor of $L_k^{(q)}$ for $q = \lceil 130|E(G)|/k^2 \rceil$. Furthermore, a minor mapping can be found in time polynomial in the size of G .*

Using the terminology of embeddings, a minor mapping of G into $L_k^{(q)}$ can be considered as an embedding from G to L_k where every vertex of L_k appears in the image of at most q vertices, i.e., the vertex depth of the embedding is at most q . Thus we can restate Lemma 6.7 the following way:

LEMMA 6.8. *For every $k > 1$ there is a constant $n_k > 0$ such that for every G with $|E(G)| > n_k$ and no isolated vertices, the graph G has an embedding into L_k with vertex depth $O(|E(G)|/k^2)$. Furthermore, such an embedding can be found in time polynomial in the size of G .*

Now we are ready to prove Theorem 6.1, the main result of the section:

PROOF (OF THEOREM 6.1). By Lemma 6.5 and Lemma 6.6, for some $k = \Omega(\lambda\sqrt{\text{con}_\lambda(H)})$, there are cliques K_1, \dots, K_k and a uniform concurrent flow $F_{i,j}$ ($1 \leq i < j \leq k$) of value $\epsilon = \Omega(\lambda/k^{\frac{3}{2}})$ on (K_1, \dots, K_k) . By trying all possibilities for the cliques and then solving the uniform concurrent flow linear program, we can find these flows (the time required for this step is a constant $f(H, \lambda)$ depending only on H and λ). Let w_0 be the smallest positive weight appearing in the flows.

Let $m = |E(G)|$ and suppose that $m \geq n_k$, for the constant n_k in Lemma 6.8. Thus the algorithm of Lemma 6.8 can be used to find an embedding ψ from G to L_k with vertex depth $q = O(m/k^2)$. Let us denote by $v_{\{i,j\}}$ ($1 \leq i < j \leq k$) the vertices of L_k with the meaning that distinct vertices $v_{\{i_1, j_1\}}$ and $v_{\{i_2, j_2\}}$ are adjacent if and only if $\{i_1, j_1\} \cap \{i_2, j_2\} \neq \emptyset$.

We construct an embedding ϕ from G to H the following way. The set $\phi(u) \subseteq V(H)$ is obtained from the set $\psi(u) \subseteq V(L_k)$ by replacing each vertex of $v_{\{i,j\}} \in \psi(u) \subseteq V(L_k)$ by a path from the flow $F_{i,j}$ (thus $\phi(u)$ is the union of $|\psi(u)|$ paths of H). We select the paths in such a way that the following requirement is satisfied: a path P of $F_{i,j}$ having weight w is selected into the image of at most $\lceil (q/\epsilon) \cdot w \rceil$ vertices of G . We set $m_{H,\lambda}$ sufficiently large that $(q/\epsilon) \cdot w_0 \geq 1$ (note that q depends on m , but ϵ and w_0 depends only on H and λ). Thus if $m \geq m_{H,\lambda}$, then $\lceil (q/\epsilon) \cdot w \rceil \leq 2(q/\epsilon) \cdot w$. Since the total weight of the paths in $F_{i,j}$ is ϵ , these paths can accommodate the image of at least $(q/\epsilon) \cdot \epsilon = q$ vertices. As each vertex $v_{\{i,j\}}$ of L_k appears in the image of at most q vertices of G in the mapping ψ , we can satisfy the requirement.

It is easy to see that if u_1 and u_2 are adjacent in G , then $\phi(u_1)$ and $\phi(u_2)$ touch: in this case, there are vertices $v_{\{i_1, j_1\}} \in \psi(u_1)$, $v_{\{i_2, j_2\}} \in \psi(u_2)$ that are adjacent or the same in L_k (that is, there is a $t \in \{i_1, j_1\} \cap \{i_2, j_2\}$), and the corresponding paths of F_{i_1, j_1} and F_{i_2, j_2} selected into $\phi(u_1)$ and $\phi(u_2)$ touch, as they both intersect the clique K_t . With a similar argument, we can show that $\phi(u)$ is connected.

To bound the edge depth of the embedding ϕ , consider an edge e . The total weight of the paths intersecting e is at most 1 and a path with weight w is used in the image of at most $2(q/\epsilon) \cdot w$ vertices. Each path intersects e in at most 2 vertices (as we can assume that the paths appearing in the flows are minimal), thus a path with weight w contributes at most $4(q/\epsilon) \cdot w$ to the depth of e . Thus the edge depth of ϕ is at most $4(q/\epsilon) = O(m/(\lambda\sqrt{k})) = O(m/(\lambda^{\frac{3}{2}} \text{con}_\lambda(H)^{\frac{1}{4}}))$. \square

6.3. Connection with adaptive width

As an easy consequence of the embedding result Corollary 6.2, we can show that large submodular width implies large adaptive width:

LEMMA 6.9. *For every hypergraph H , $\text{adw}(H) = \Omega(\text{emb}(H))$.*

PROOF. Suppose that $\text{emb}(H) > \alpha$. This means that there is an integer m_α such that every graph with $m \geq m_\alpha$ edges has an embedding into H with edge depth m/α . It is well-known that there are arbitrarily large sparse graphs whose treewidth is linear in the number of vertices (for example, bounded-degree expanders, see e.g., [Grohe and Marx 2009]): for some universal constant β , there is a graph G with $m \geq m_\alpha$ edges and treewidth at least βm . Thus there is an embedding ϕ from G to H with edge depth at most $q \leq m/\alpha$. Let $d(v)$ be the depth of vertex v in the embedding and let us define $\mu(v) := d(v)/q$. From the definition of edge depth, it is clear that μ is a fractional independent set. Suppose that there is a tree decomposition $(T, B_{v \in V(T)})$ of H having μ -width w . This tree decomposition can be turned into a tree decomposition $(T, B'_{v \in V(T)})$ of G : for every $B_t \subseteq V(H)$, let $B'_t := \{u \in V(G) \mid \phi(u) \cap B_t \neq \emptyset\}$ contain those vertices of G whose images intersect B_t . Now $\mu(B_t) \leq w$ means that $\sum_{v \in B_t} d(v) \leq qw$, which implies that $|B'_t| \leq qw$. Thus the width of $(T, B'_{v \in V(T)})$ is less than qw , which means that w has to be at least $\beta m/q = \Omega(\alpha)$, the required lower bound on the adaptive width of H . \square

Combining Theorem 5.1 and Lemma 6.9 gives:

COROLLARY 6.10. *For every hypergraph H , $\text{subw}(H) = O(\text{adw}(H)^4)$.*

7. FROM EMBEDDINGS TO HARDNESS OF CSP

We prove the main hardness result of the paper in this section:

THEOREM 7.1. *Let \mathcal{H} be a recursively enumerable class of hypergraphs with unbounded submodular width. If there is an algorithm \mathbb{A} and a function f such that \mathbb{A} solves every instance I of $\text{CSP}(\mathcal{H})$ with hypergraph $H \in \mathcal{H}$ in time $f(H) \cdot \|I\|^{o(\text{subw}(H)^{1/4})}$, then the Exponential Time Hypothesis fails.*

In particular, Theorem 7.1 implies that $\text{CSP}(\mathcal{H})$ for such a \mathcal{H} is not fixed-parameter tractable:

COROLLARY 7.2. *If \mathcal{H} is a recursively enumerable class of hypergraphs with unbounded submodular width, then $\text{CSP}(\mathcal{H})$ is not fixed-parameter tractable, unless the Exponential Time Hypothesis fails.*

The Exponential Time Hypothesis (ETH) states that there is no $2^{o(n)}$ time algorithm for n -variable 3SAT. The Sparsification Lemma of Impagliazzo et al. [2001] shows that ETH is equivalent to the assumption that there is no algorithm for 3SAT whose running time is subexponential *in the number of clauses*. This result will be crucial for our hardness proof, as our reduction from 3SAT is sensitive to the number of clauses.

THEOREM 7.3 (IMPAGLIAZZO ET AL. [2001]). *If there is a $2^{o(m)}$ time algorithm for m -clause 3SAT, then there is a $2^{o(n)}$ time algorithm for n -variable 3SAT.*

To prove Theorem 7.1, we show that a subexponential-time algorithm for 3SAT exists if $\text{CSP}(\mathcal{H})$ can be solved “too fast” for some \mathcal{H} with unbounded submodular width. We use the characterization of submodular width from Section 5 and the embedding results of Section 6 to reduce 3SAT to $\text{CSP}(\mathcal{H})$ by embedding the incidence graph of a 3SAT formula into a hypergraph $H \in \mathcal{H}$. The basic idea of the proof is that if the 3SAT formula has m clauses and the edge depth of the embedding is m/r , then we can gain a factor r in the exponent of the running time. If submodular width is unbounded in \mathcal{H} , then we can make this gap r between the number of clauses and the edge depth arbitrary large, and hence the exponent can be arbitrarily smaller than the number of clauses, i.e., the algorithm is subexponential in the number of clauses.

The following simple lemma from [Marx 2010b] gives a transformation that turns a 3SAT instance into a binary CSP instance. We include the proof for completeness.

LEMMA 7.4. *Given an instance of 3SAT with n variables and m clauses, it is possible to construct in polynomial time an equivalent CSP instance with $n+m$ variables, $3m$ binary constraints, and domain size 3.*

PROOF. Let ϕ be a 3SAT formula with n variables and m clauses. We construct an instance of CSP as follows. The CSP instance contains a variable x_i ($1 \leq i \leq n$) corresponding to the i -th variable of ϕ and a variable y_j ($1 \leq j \leq m$) corresponding to the j -th clause of ϕ . Let $D = \{1, 2, 3\}$ be the domain. We try to describe a satisfying assignment of ϕ with these $n+m$ variables. The intended meaning of the variables is the following. If the value of variable x_i is 1 (resp., 2), then this represents that the i -th variable of ϕ is true (resp., false). If the value of variable y_j is ℓ , then this represents that the j -th clause of ϕ is satisfied by its ℓ -th literal. To ensure consistency, we add $3m$ constraints. Let $1 \leq j \leq m$ and $1 \leq \ell \leq 3$, and assume that the ℓ -th literal of the j -th clause is a positive occurrence of the i -th variable. In this case, we add the binary constraint $(x_i = 1 \vee y_j \neq \ell)$: either x_i is true or some other literal satisfies the clause. Similarly, if the ℓ -th literal of the j -th clause is a negated occurrence of the i -th variable, then we add the binary constraint $(x_i = 2 \vee y_j \neq \ell)$. It is easy to verify that if ϕ is satisfiable, then we can assign values to the variables of the CSP instance such that every constraint is satisfied, and conversely, if the CSP instance has a solution, then ϕ is satisfiable. \square

Next we show that an embedding from graph G to hypergraph H can be used to simulate a binary CSP instance I_1 having primal graph G by a CSP instance I_2 whose hypergraph is H . The domain size and the size of the constraint relations of I_2 can grow very large in this transformation: the edge depth of the embedding determines how large this increase is.

LEMMA 7.5. *Let $I_1 = (V_1, D_1, C_1)$ be a binary CSP instance with primal graph G and let ϕ be an embedding of G into a hypergraph H with edge depth q . Given I_1 , H , and the embedding ϕ , it is possible to construct (in time polynomial in the size of the output) an equivalent CSP instance $I_2 = (V_2, D_2, C_2)$ with hypergraph H where the size of every constraint relation is at most $|D_1|^q$.*

PROOF. For every $v \in V(H)$, let $U_v := \{u \in V(G) \mid v \in \phi(u)\}$ be the set of vertices in G whose images contain v , and for every $e \in E(H)$, let $U_e := \bigcup_{v \in e} U_v$. Observe that for every $e \in E(H)$, we have $|U_e| \leq \sum_{v \in e} |U_v| \leq q$, since the edge depth of ϕ is q . Let D_2 be the set of integers between 1 and $|D_1|^q$. For every $v \in V(H)$, the number of assignments from U_v to D_1 is clearly $|D_1|^{|U_v|} \leq |D_1|^q$. Let us fix a bijection h_v between these assignments on U_v and the set $\{1, \dots, |D_1|^{|U_v|}\} \subseteq D_2$.

The set C_2 of constraints of I_2 are constructed as follows. For each $e \in E(H)$, there is a constraint $\langle s_e, R_e \rangle$ in C_2 , where s_e is an $|e|$ -tuple containing an arbitrary ordering of the elements of e . The relation R_e is defined the following way. Suppose that v_i is

the i -th coordinate of s_e and consider a tuple $t = (d_1, \dots, d_{|e|}) \in D_2^{|e|}$ of integers where $1 \leq d_i \leq |D_1|^{|U_{v_i}|}$ for every $1 \leq i \leq |e|$. This means that d_i is in the image of h_{v_i} and hence $f_i := h_{v_i}^{-1}(d_i)$ is an assignment from U_{v_i} to D_1 . We define relation R_e such that it contains tuple t if the following two conditions hold. First, we require that the assignments $f_1, \dots, f_{|e|}$ are *consistent* in the sense that $f_i(u) = f_j(u)$ for any i, j and $u \in U_{v_i} \cap U_{v_j}$. In this case, $f_1, \dots, f_{|e|}$ together define an assignment f on $\bigcup_{i=1}^{|e|} U_{v_i} = U_e$. The second requirement is that this assignment f satisfies every constraint of I_1 whose scope is contained in U_e , that is, for every constraint $\langle (u_1, u_2), R \rangle \in C_1$ with $\{u_1, u_2\} \subseteq U_e$, we have $(f(u_1), f(u_2)) \in R$. This completes the description of the instance I_2 .

Let us bound the maximum size of a relation of I_2 . Consider the relation R_e constructed in the previous paragraph. It contains tuples $(d_1, \dots, d_{|e|}) \in D_2^{|e|}$ where $1 \leq d_i \leq |D_1|^{|U_{v_i}|}$ for every $1 \leq i \leq |e|$. This means that

$$|R_e| \leq \prod_{i=1}^{|e|} |D_1|^{|U_{v_i}|} = |D_1|^{\sum_{i=1}^{|e|} |U_{v_i}|} \leq |D_1|^q, \quad (4)$$

where the last inequality follows from the fact that ϕ has edge depth at most q .

To prove that I_1 and I_2 are equivalent, assume first that I_1 has a solution $f_1 : V_1 \rightarrow D_1$. For every $v \in V_2$, let us define $f_2(v) := h_v(\text{pr}_{U_v} f_1)$, that is, the integer between 1 and $|D_1|^{|U_v|}$ corresponding to the projection of assignment f_1 to U_v . It is easy to see that f_2 is a solution of I_2 .

Assume now that I_2 has a solution $f_2 : V_2 \rightarrow D_2$. For every $v \in V(H)$, let $f_v := h_v^{-1}(f_2(v))$ be the assignment from U_v to D_1 that corresponds to $f_2(v)$ (note that by construction, $f_2(v)$ is at most $|D_1|^{|U_v|}$, hence $h_v^{-1}(f_2(v))$ is well-defined). We claim that these assignments are compatible: if $u \in U_{v'} \cap U_{v''}$ for some $u \in V(G)$ and $v', v'' \in V(H)$, then $f_{v'}(u) = f_{v''}(u)$. Recall that $\phi(u)$ is a connected set in H , hence there is a path between v' and v'' in $\phi(u)$. We prove the claim by induction on the distance between v' and v'' in $\phi(u)$. If the distance is 0, that is, $v' = v''$, then the statement is trivial. Suppose now that the distance of v' and v'' is $d > 0$. This means that v' has a neighbor $z \in \phi(u)$ such that the distance of z and v'' is $d - 1$. Therefore, $f_z(u) = f_{v''}(u)$ by the induction hypothesis. Since v' and z are adjacent in H , there is an edge $E \in E(H)$ containing both v' and z . From the way I_2 is defined, this means that $f_{v'}$ and f_z are compatible and $f_{v'}(u) = f_z(u) = f_{v''}(u)$ follows, proving the claim. Thus the assignments $\{f_v \mid v \in V(H)\}$ are compatible and these assignments together define an assignment $f_1 : V(G) \rightarrow D_1$. We claim that f_1 is a solution of I_1 . Let $c = \langle (u_1, u_2), R \rangle$ be an arbitrary constraint of I_1 . Since $u_1 u_2 \in E(G)$, sets $\phi(u_1)$ and $\phi(u_2)$ touch, thus there is an edge $e \in E(H)$ that contains a vertex $v_1 \in \phi(u_1)$ and a vertex $v_2 \in \phi(u_2)$ (or, in other words, $u_1 \in U_{v_1}$ and $u_2 \in U_{v_2}$). The definition of c_e in I_2 ensures that f_1 restricted to $U_{v_1} \cup U_{v_2}$ satisfies every constraint of I_1 whose scope is contained in $U_{v_1} \cup U_{v_2}$; in particular, f_1 satisfies constraint c . \square

Now we are ready to prove Theorem 7.1, the main result of the section. We show that if there is a class \mathcal{H} of hypergraphs with unbounded submodular width such that $\text{CSP}(\mathcal{H})$ is FPT, then this algorithm can be used to solve 3SAT in subexponential time. The main ingredients are the embedding result of Theorem 6.1, and Lemmas 7.4 and 7.5 above on reduction to CSP. Furthermore, we need a way of choosing an appropriate hypergraph from the set \mathcal{H} . As discussed earlier, the larger the submodular width of the hypergraph is, the more we gain in the running time. However, we should not spend too much time on constructing the hypergraph and on finding an embedding. Therefore, we use the same technique as in [Marx 2010b]: we enumerate a certain number of hypergraphs and we try all of them simultaneously. The number of hypergraphs enumerated depends on the size of the 3SAT instance. This will be done in such a way that guarantees that we do not spend

too much time on the enumeration, but eventually every hypergraph in \mathcal{H} is considered for sufficiently large input sizes.

PROOF (OF THEOREM 7.1). Let us fix a $\lambda > 0$ that is sufficiently small for Theorems 5.1 and 6.1. Suppose that there is an $f_1(H)n^{o(\text{subw}(H)^{1/4})}$ time algorithm \mathbb{A} for $\text{CSP}(\mathcal{H})$. We can express the running time as $f_1(H)n^{\text{subw}(H)^{1/4}/\iota(\text{subw}(H))}$ for some unbounded nondecreasing function ι with $\iota(1) > 0$. We construct an algorithm \mathbb{B} that solves 3SAT in subexponential time by using algorithm \mathbb{A} as subroutine.

Given an instance I of 3SAT with n variables and m clauses and a hypergraph $H \in \mathcal{H}$, we can solve I the following way. First we use Lemma 7.4 to transform I into a CSP instance $I_1 = (V_1, D_1, C_1)$ with $|V_1| = n + m$, $|D_1| = 3$, and $|C_1| = 3m$. Let G be the primal graph of I_1 , which is a graph having $3m$ edges. It can be assumed that m is greater than some constant $m_{H,\lambda}$ of Theorem 6.1, otherwise the instance can be solved in constant time. Therefore, the algorithm of Theorem 6.1 can be used to find an embedding ϕ of G into H with edge depth $q = O(m/(\lambda^{\frac{3}{2}} \text{con}_\lambda(H)^{1/4}))$; by Theorem 5.1, we have that $\text{con}_\lambda(H) = \Omega(\text{subw}(H))$ and hence $q \leq c_\lambda m / \text{subw}(H)^{1/4}$ for some constant c_λ depending only on λ . By Lemma 7.5, we can construct an equivalent instance $I_2 = (V_2, D_2, C_2)$ whose hypergraph is H . By solving I_2 using the assumed algorithm \mathbb{A} for $\text{CSP}(\mathcal{H})$, we can answer if I_1 has a solution, or equivalently, if the 3SAT instance I has a solution.

We will call “running algorithm $\mathbb{A}[I, H]$ ” this way of solving the 3SAT instance I . Let us determine the running time of $\mathbb{A}[I, H]$. The two dominating terms are the time required to find embedding ϕ using the $f(H, \lambda)m^{O(1)}$ time algorithm of Theorem 7.1 and the time required to run \mathbb{A} on I_2 . The size of every constraint relation in I_2 is at most $|D_1|^q = 3^q$, hence $\|I_2\| = O((|E(H)| + |V(H)|)3^q)$. Let $k = \text{subw}(H)$. The total running time of $\mathbb{A}[I, H]$ can be bounded by

$$\begin{aligned} f(H, \lambda)m^{O(1)} + f_1(H)\|I_2\|^{k^{1/4}/\iota(k)} &= f(H, \lambda)m^{O(1)} + f_1(H)(|E(H)| + |V(H)|)^{k^{1/4}/\iota(k)} \cdot 3^{q \cdot k^{1/4}/\iota(k)} \\ &= f_2(H, \lambda) \cdot m^{O(1)} \cdot 3^{c_\lambda m/\iota(k)} \end{aligned}$$

for an appropriate function $f_2(H, \lambda)$ depending only on H and λ .

Algorithm \mathbb{B} for 3SAT proceeds as follows. Let us fix an arbitrary computable enumeration H_1, H_2, \dots of the hypergraphs in \mathcal{H} . Given an m -clause 3SAT formula I , algorithm \mathbb{B} spends the first m steps on enumerating these hypergraphs; let H_ℓ be the last hypergraph produced by this enumeration (we assume that m is sufficiently large that $\ell \geq 1$). Next we start simulating the algorithms $\mathbb{A}[I, H_1], \mathbb{A}[I, H_2], \dots, \mathbb{A}[I, H_\ell]$ in *parallel*. When one of the simulations stops and returns an answer, then we stop all the simulations and return the answer. It is clear that algorithm \mathbb{B} will correctly decide the satisfiability of I .

We claim that there is a universal constant d such that for every s , there is an m_s such that for every $m > m_s$, the running time of \mathbb{B} is at most $(m \cdot 2^{m/s})^d$ on an m -clause formula. Clearly, this means that the running time of \mathbb{B} is $2^{o(m)}$.

For any positive integer s , let k_s be the smallest positive integer such that $\iota(k_s) \geq s$ (as ι is unbounded, this is well defined). Let i_s be the smallest positive integer such that $\text{subw}(H_{i_s}) \geq k_s$ (as \mathcal{H} has unbounded submodular width, this is also well defined). Set m_s sufficiently large that $m_s \geq f_2(H_{i_s}, \lambda)$ and the fixed enumeration of \mathcal{H} reaches H_{i_s} in less than m_s steps. This means that if we run \mathbb{B} on a 3SAT formula I with $m \geq m_s$ clauses, then $\ell \geq i_s$ and hence $\mathbb{A}[I, H_{i_s}]$ will be one of the ℓ simulations started by \mathbb{B} . The simulation of $\mathbb{A}[I, H_{i_s}]$ terminates in

$$f_2(H_{i_s}, \lambda)m^{O(1)} \cdot 3^{c_\lambda m/\iota(\text{subw}(H_{i_s}))} \leq m \cdot m^{O(1)} \cdot 3^{c_\lambda m/s}$$

steps. Taking into account that we simulate $\ell \leq m$ algorithms in parallel and all the simulations are stopped not later than the termination of $\mathbb{A}[I, H_{i_s}]$, the running time of \mathbb{B} can

be bounded polynomially by the running time of $\mathbb{A}[I, H_{i_s}]$. Therefore, there is a constant d such that the running time of \mathbb{B} is at most $(m \cdot 2^{m/s})^d$, as required. \square

Remark 7.6. Recall that if ϕ is an embedding of G into H , then the depth of an edge $e \in E(H)$ is $d_\phi(e) = \sum_{v \in V(G)} |\phi(v) \cap e|$. A variant of this definition would be to define the depth of e as $d'_\phi(e) = |\{v \in V(G) \mid \phi(v) \cap e \neq \emptyset\}|$, i.e., if $\phi(v)$ intersects e , then v contributes only 1 to the depth of e , not $|\phi(v) \cap e|$ as in the original definition. Let us call this variant *weak edge depth*, it is clear that the weak edge depth of an embedding is at most the edge depth of the embedding.

Lemma 7.5 can be made stronger by requiring only that the weak edge depth is at most q . Indeed, the only place where we use the bound on edge depth is in Inequality (4). However, the size of the relation R_e can be bounded by the number of possible assignments on U_e in instance I_1 . If weak edge depth is at most q , then $|U_e| \leq q$, and the $|D_1|^q$ bound on the size of R_e follows.

Remark 7.7. A different version of CSP was investigated in [Marx 2011], where each variable has a different domain, and each constraint relation is represented by a full truth table (see the exact definition in [Marx 2011]). Let us denote by $\text{CSP}_{\text{tt}}(\mathcal{H})$ this variant of the problem. It is easy to see that $\text{CSP}_{\text{tt}}(\mathcal{H})$ can be reduced to $\text{CSP}(\mathcal{H})$ in polynomial time, but a reduction in the other direction can possibly increase the representation of a constraint by an exponential factor. Nevertheless, the hardness results of this section apply to the “easier” problem $\text{CSP}_{\text{tt}}(\mathcal{H})$ as well. What we have to verify is that the proof of Lemma 7.5 works even if I_2 is an instance of CSP_{tt} , i.e., the constraint relations have to be represented by truth tables. Inspection of the proof shows that it indeed works: the product in Inequality (4) is exactly the size of the truth table describing the constraint corresponding to edge e , thus the $|D_1|^q$ upper bound remains valid even if constraints are represented by truth tables. Therefore, the hardness results of [Marx 2011] are subsumed by the following corollary:

COROLLARY 7.8. *If \mathcal{H} is a recursively enumerable class of hypergraphs with unbounded submodular width, then $\text{CSP}_{\text{tt}}(\mathcal{H})$ is not fixed-parameter tractable, unless the Exponential Time Hypothesis fails.*

8. CONCLUSIONS

The main result of the paper is introducing submodular width and proving that bounded submodular width is the property that determines the fixed-parameter tractability of $\text{CSP}(\mathcal{H})$. The hardness result is proved assuming the Exponential Time Hypothesis. This conjecture was formulated relatively recently [Impagliazzo et al. 2001], but it turned out to be very useful in proving lower bounds in a variety of settings [Marx 2010b; Andoni et al. 2006; Marx 2007; Pătraşcu and Williams 2010].

For the hardness proof, we had to understand what large submodular width means and we had to explore the connection between submodular width and other combinatorial properties. We have obtained several equivalent characterizations of bounded submodular width, in particular, we have showed that bounded submodular width is equivalent to bounded adaptive width:

COROLLARY 8.1. *The following are equivalent for every class \mathcal{H} of hypergraphs:*

- (1) *There is a constant c_1 such that $\mu\text{-width}(H) \leq c_1$ for every $H \in \mathcal{H}$ and fractional independent set μ .*
- (2) *There is a constant c_2 such that $b\text{-width}(H) \leq c_2$ for every $H \in \mathcal{H}$ and edge-dominated monotone submodular function b on $V(H)$ with $b(\emptyset) = 0$.*
- (3) *There is a constant c_3 such that $b^*\text{-width}(H) \leq c_3$ for every $H \in \mathcal{H}$ and edge-dominated monotone submodular function b on $V(H)$ with $b(\emptyset) = 0$.*

(4) There is a constant c_4 such that $\text{con}_\lambda(H) \leq c_4$ for every $H \in \mathcal{H}$, where $\lambda > 0$ is a universal constant.

(5) There is a constant c_5 such that $\text{emb}(H) \leq c_5$ for every $H \in \mathcal{H}$.

Implications (2) \Rightarrow (1) and (3) \Rightarrow (2) are trivial; (4) \Rightarrow (3) follows from Lemma 5.10; (5) \Rightarrow (4) follows from Corollary 6.2; (1) \Rightarrow (5) follows from Lemma 6.9.

Let us briefly review the main ideas that were necessary for proving the main result of the paper:

- Recognizing that submodular width is the right property characterizing the complexity of the problem.

- A CSP instance can be partitioned into a bounded number of uniform instances (Section 4.2).

- The number of solutions in a uniform CSP instance can be described by a submodular function (Section 4.3).

- There is a connection between fractional separation and finding a separator minimizing an edge-dominated submodular cost function (Section 5.2).

- The transformation that turns b into b^* , and the properties of b^* that are more suitable than b for recursively constructing a tree decomposition (Section 5.1).

- Our results on fractional separation and the standard framework of finding tree decompositions show that large submodular width implies that there is a highly connected set (Section 5.3).

- A highly connected set can be turned into a highly connected set that is partitioned into cliques in an appropriate way (Section 6.1).

- A highly connected set with appropriate cliques implies that there is a uniform concurrent flow of large value between the cliques (Section 6.2).

- Similarly to [Marx 2010b], we use the observation that a concurrent flow is analogous to a line graph of a clique, hence it has good embedding properties (Section 6.2).

- Similarly to [Marx 2010b], an embedding in a hypergraph gives a way of simulating 3SAT with $\text{CSP}(\mathcal{H})$ (Section 7).

It is possible that the main result can be proved in a simpler way by bypassing some of the ideas above. In particular, a surprising consequence of our results is that bounded submodular width and bounded adaptive width are the same, i.e., if a class \mathcal{H} has unbounded submodular width, then for every k there is a $H_k \in \mathcal{H}$ and a fractional independent set μ_k such that μ_k -width(H_k) $\geq k$, or in other words, large submodular width can be certified by the *modular* function μ_k . To prove this, we need all the results of Sections 5 and 6. Having a better understanding and an independent proof of this fact could simplify the proofs considerably. Another possible target for simplification is Section 6.1, where a lot of effort is spent on proving that if there is a large highly connected set, then there is a large highly connected set that is partitioned into cliques in an appropriate way. It might be possible to strengthen the results of Section 5 (perhaps by better understanding the role of cliques in separators) so that they give such a highly connected set directly.

An obvious question for further research is whether it is possible to prove a similar dichotomy result with respect to polynomial-time solvability. At this point, it is hard to see what the answer could be if we investigate the same question using the more restricted notion of polynomial time solvability. We know that bounded fractional hypertree width implies polynomial-time solvability [Marx 2010a] and Theorem 7.1 shows that unbounded submodular width implies that the problem is not polynomial-time solvable (as it is not even fixed-parameter tractable). So only those classes of hypergraphs are in the “gray zone” that have bounded submodular width but unbounded fractional hypertree width.

What could be the truth in this gray zone? A first possibility is that $\text{CSP}(\mathcal{H})$ is polynomial-time solvable for every such class, i.e., Theorem 4.1 can be improved from

fixed-parameter tractability to polynomial-time solvability. However, Theorem 4.1 uses the power of fixed-parameter tractability in an essential way (splitting into a double-exponential number of uniform instances), so it is not clear how such improvement is possible. A second possibility is that unbounded fractional hypertree width implies that $\text{CSP}(\mathcal{H})$ is not polynomial-time solvable. Substantially new techniques would be required for such a hardness proof. The hardness proofs of this paper and of [Grohe 2007; Marx 2010b] are based on showing that a large problem space can be efficiently embedded into an instance with a particular hypergraph. However, the fixed-parameter tractability results show that no such embedding is possible in case of classes with bounded submodular width. Therefore, a possible hardness proof should embed a problem space that is comparable (in some sense) with the size of the hypergraph and should create instances where the domain size is bounded by a function of the size of the hypergraph. A third possibility is that the boundary of polynomial-time solvability is somewhere between bounded fractional hypertree width and bounded submodular width. Currently, there is no natural candidate for a property that could correspond to this boundary and, again, the hardness part of the characterization should be substantially different than what was done before. Finally, there is a fourth possibility: the boundary of the polynomial-time cases cannot be elegantly characterized by a simple combinatorial property. In general, if we consider the restriction of a problem to all possible classes of (hyper)graphs, then there is no a priori reason why an elegant characterization should exist that describes the easy and hard classes. For example, it is highly unlikely that there is an elegant characterization of those classes of graphs where solving the MAXIMUM INDEPENDENT SET problem is polynomial-time solvable. As discussed earlier, the fixed-parameter tractability of $\text{CSP}(\mathcal{H})$ is a more robust question than its polynomial-time solvability, hence it is very well possible that only the former question has an elegant answer.

REFERENCES

- Isolde Adler. 2006. *Width functions for hypertree decompositions*. Ph.D. Dissertation. Albert-Ludwigs-Universität Freiburg.
- Isolde Adler, Georg Gottlob, and Martin Grohe. 2007. Hypertree width and related hypergraph invariants. *European J. Combin.* 28, 8 (2007), 2167–2181.
- Amit Agarwal, Noga Alon, and Moses Charikar. 2007. Improved approximation for directed cut problems. In *STOC'07—Proceedings of the 39th Annual ACM Symposium on Theory of Computing*. ACM, New York, 671–680.
- Noga Alon, Ilan Newman, Alexander Shen, Gábor Tardos, and Nikolai Vereshchagin. 2007. Partitioning multi-dimensional sets in a small number of “uniform” parts. *European J. Combin.* 28, 1 (2007), 134–144.
- Omid Amini, Frédéric Mazoit, Nicolas Nisse, and Stéphan Thomassé. 2009. Submodular partition functions. *Discrete Mathematics* 309, 20 (2009), 6000 – 6008. DOI: <http://dx.doi.org/DOI:10.1016/j.disc.2009.04.033>
- Alexandr Andoni, Piotr Indyk, and Mihai Patrascu. 2006. On the Optimality of the Dimensionality Reduction Method. In *FOCS '06: Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*. IEEE Computer Society, Washington, DC, USA, 449–458. DOI: <http://dx.doi.org/10.1109/FOCS.2006.56>
- Albert Atserias, Andrei A. Bulatov, and Víctor Dalmau. 2007. On the Power of k -Consistency. In *ICALP*. 279–290.
- Catriel Beeri, Ronald Fagin, David Maier, Alberto O. Mendelzon, Jeffrey D. Ullman, and Mihalis Yannakakis. 1981. Properties of Acyclic Database Schemes. In *STOC*. 355–362.
- Catriel Beeri, Ronald Fagin, David Maier, and Mihalis Yannakakis. 1983. On the desirability of acyclic database schemes. *J. Assoc. Comput. Mach.* 30, 3 (1983), 479–513.
- Andrei A. Bulatov. 2003. Tractable conservative Constraint Satisfaction Problems. In *18th Annual IEEE Symposium on Logic in Computer Science (LICS'03)*. IEEE Computer Society, Los Alamitos, CA, USA, 321.
- Andrei A. Bulatov. 2006. A dichotomy theorem for constraint satisfaction problems on a 3-element set. *J. ACM* 53, 1 (2006), 66–120.

- A. A. Bulatov, A. A. Krokhin, and P. Jeavons. 2001. The complexity of maximal constraint languages. In *Proceedings of the 33rd ACM Symposium on Theory of Computing*. 667–674.
- Ashok K. Chandra and Philip M. Merlin. 1977. Optimal Implementation of Conjunctive Queries in Relational Data Bases. In *STOC*. 77–90.
- Chandra Chekuri and Anand Rajaraman. 2000. Conjunctive query containment revisited. *Theoret. Comput. Sci.* 239, 2 (2000), 211–229.
- Hubie Chen and Martin Grohe. 2010. Constraint satisfaction with succinctly specified relations. *J. Comput. Syst. Sci.* 76, 8 (2010), 847–860.
- F. R. K. Chung, R. L. Graham, P. Frankl, and J. B. Shearer. 1986. Some intersection theorems for ordered sets and graphs. *J. Combin. Theory Ser. A* 43, 1 (1986), 23–37.
- Víctor Dalmau, Phokion G. Kolaitis, and Moshe Y. Vardi. 2002. Constraint Satisfaction, Bounded Treewidth, and Finite-Variable Logics. In *CP '02: Proceedings of the 8th International Conference on Principles and Practice of Constraint Programming*. Springer-Verlag, London, UK, 310–326.
- R. G. Downey and M. R. Fellows. 1999. *Parameterized Complexity*. Springer, New York. xvi+533 pages.
- Ronald Fagin. 1983. Degrees of acyclicity for hypergraphs and relational database schemes. *J. Assoc. Comput. Mach.* 30, 3 (1983), 514–550.
- Tomás Feder and Moshe Y. Vardi. 1999. The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory. *SIAM J. Comput.* 28, 1 (1999), 57–104.
- Jörg Flum and Martin Grohe. 2006. *Parameterized Complexity Theory*. Springer, Berlin.
- E. C. Freuder. 1990. Complexity of K-Tree Structured Constraint Satisfaction Problems. In *Proc. of AAAI-90*. Boston, MA, 4–9.
- G. Gottlob, N. Leone, and F. Scarcello. 2002a. Hypertree Decompositions and Tractable Queries. *J. Comput. System Sci.* 64 (2002), 579–627.
- Georg Gottlob, Francesco Scarcello, and Martha Sideri. 2002b. Fixed-parameter complexity in AI and nonmonotonic reasoning. *Artificial Intelligence* 138, 1-2 (2002), 55–86.
- G. Gottlob and S. Szeider. 2008. Fixed-Parameter Algorithms For Artificial Intelligence, Constraint Satisfaction and Database Problems. *Comput. J.* 51, 3 (2008), 303–325.
- Gianluigi Greco and Francesco Scarcello. 2010. The power of tree projections: local consistency, greedy algorithms, and larger islands of tractability. In *Proceedings of the twenty-ninth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems (PODS '10)*. ACM, New York, NY, USA, 327–338. DOI:<http://dx.doi.org/10.1145/1807085.1807127>
- Martin Grohe. 2006. The Structure of Tractable Constraint Satisfaction Problems.. In *MFCS 2006*. 58–72.
- Martin Grohe. 2007. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *J. ACM* 54, 1 (2007), 1. DOI:<http://dx.doi.org/10.1145/1206035.1206036>
- Martin Grohe and Dániel Marx. 2009. On tree width, bramble size, and expansion. *Journal of Combinatorial Theory Ser. B* 99, 1 (2009), 218–228.
- Martin Grohe and Dániel Marx. 2012+. Constraint solving via fractional edge covers. (2012+). To appear in *ACM Transactions on Algorithm*.
- Anupam Gupta. 2003. Improved results for directed multicut. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (Baltimore, MD, 2003)*. ACM, New York, 454–455.
- Mohammad Taghi Hajiaghayi and Harald Räcke. 2006. An $O(\sqrt{n})$ -approximation algorithm for directed sparsest cut. *Inform. Process. Lett.* 97, 4 (2006), 156–160.
- Petr Hliněný. 2005. A parametrized algorithm for matroid branch-width. *SIAM J. Comput.* 35, 2 (2005), 259–277.
- Petr Hliněný and Sang-il Oum. 2008. Finding branch-decompositions and rank-decompositions. *SIAM J. Comput.* 38, 3 (2008), 1012–1032.
- Petr Hliněný and Geoff Whittle. 2006. Matroid tree-width. *European J. Combin.* 27, 7 (2006), 1117–1128.
- Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. 2001. Which problems have strongly exponential complexity? *J. Comput. System Sci.* 63, 4 (2001), 512–530.
- Satoru Iwata. 2008. Submodular function minimization. *Math. Program.* 112, 1, Ser. B (2008), 45–64.
- Satoru Iwata, Lisa Fleischer, and Satoru Fujishige. 2001. A combinatorial strongly polynomial algorithm for minimizing submodular functions. *J. ACM* 48, 4 (2001), 761–777 (electronic).
- P. Jeavons, D. A. Cohen, and M. Gyssens. 1997. Closure properties of constraints. *J. ACM* 44, 4 (1997), 527–548.
- Phokion G. Kolaitis and Moshe Y. Vardi. 2000a. Conjunctive-query containment and constraint satisfaction. *J. Comput. Syst. Sci.* 61, 2 (2000), 302–332. DOI:<http://dx.doi.org/10.1006/jcss.2000.1713>

- Phokion G. Kolaitis and Moshe Y. Vardi. 2000b. A Game-Theoretic Approach to Constraint Satisfaction. In *AAAI/IAAI*. 175–181.
- Dániel Marx. 2007. On the optimality of planar and geometric approximation schemes. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*. 338–348.
- Dániel Marx. 2010a. Approximating fractional hypertree width. *ACM Trans. Algorithms* 6, 2 (2010), 1–17. DOI:<http://dx.doi.org/10.1145/1721837.1721845>
- Dániel Marx. 2010b. Can You Beat Treewidth? *Theory of Computing* 6, 1 (2010), 85–112. DOI:<http://dx.doi.org/10.4086/toc.2010.v006a005>
- Dániel Marx. 2010c. Tractable hypergraph properties for constraint satisfaction and conjunctive queries. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*. 735–744.
- Dániel Marx. 2011. Tractable Structures for Constraint Satisfaction with Truth Tables. *Theory of Computing Systems* 48 (2011), 444–464.
- Rolf Niedermeier. 2006. *Invitation to fixed-parameter algorithms*. Oxford Lecture Series in Mathematics and its Applications, Vol. 31. Oxford University Press, Oxford. xii+300 pages.
- Sang-il Oum. 2005. Approximating Rank-width and Clique-width Quickly. In *Proceedings of the 31st International Workshop on Graph-Theoretic Concepts in Computer Science*. 49–58.
- Sang-il Oum and Paul Seymour. 2007. Testing branch-width. *J. Combin. Theory Ser. B* 97, 3 (2007), 385–393.
- Sang-il Oum and Paul Seymour. 2006. Approximating clique-width and branch-width. *J. Combin. Theory Ser. B* 96, 4 (2006), 514–528.
- Mihai Pătraşcu and Ryan Williams. 2010. On the Possibility of Faster SAT Algorithms. In *Proc. 21st ACM/SIAM Symposium on Discrete Algorithms (SODA)*. 1065–1075.
- Francesco Scarcello, Georg Gottlob, and Gianluigi Greco. 2008. Uniform Constraint Satisfaction Problems and Database Theory. In *Complexity of Constraints*. 156–195.
- Thomas J. Schaefer. 1978. The complexity of satisfiability problems. In *Conference Record of the Tenth Annual ACM Symposium on Theory of Computing (San Diego, Calif., 1978)*. ACM, New York, 216–226.
- Alexander Schrijver. 2000. A combinatorial algorithm minimizing submodular functions in strongly polynomial time. *J. Combin. Theory Ser. B* 80, 2 (2000), 346–355.
- Mihalis Yannakakis. 1981. Algorithms for Acyclic Database Schemes. In *VLDB*. 82–94.