

HyperBench: A Benchmark and Tool for Hypergraphs and Empirical Findings

WOLFGANG FISCHL, TU Wien

GEORG GOTTLÖB, University of Oxford

DAVIDE MARIO LONGO and REINHARD PICHLER, TU Wien

To cope with the intractability of answering Conjunctive Queries (CQs) and solving Constraint Satisfaction Problems (CSPs), several notions of hypergraph decompositions have been proposed—giving rise to different notions of width, noticeably, plain, generalized, and fractional hypertree width (hw, ghw, and fhw). Given the increasing interest in using such decomposition methods in practice, a publicly accessible repository of decomposition software, as well as a large set of benchmarks, and a web-accessible workbench for inserting, analyzing, and retrieving hypergraphs are called for.

We address this need by providing (i) concrete implementations of hypergraph decompositions (including new practical algorithms), (ii) a new, comprehensive benchmark of hypergraphs stemming from disparate CQ and CSP collections, and (iii) HyperBench, our new web-interface for accessing the benchmark and the results of our analyses. In addition, we describe a number of actual experiments we carried out with this new infrastructure.

CCS Concepts: • **Information systems** → **Database query processing**; • **Theory of computation** → **Design and analysis of algorithms**;

Additional Key Words and Phrases: Hypergraph decomposition methods, query answering, constraint satisfaction

ACM Reference format:

Wolfgang Fischl, Georg Gottlob, Davide Mario Longo, and Reinhard Pichler. 2020. HyperBench: A Benchmark and Tool for Hypergraphs and Empirical Findings. *J. Exp. Algorithmics* 26, 1, Article 1.6 (July 2021), 40 pages. <https://doi.org/10.1145/3440015>

1 INTRODUCTION

In this work, we study computational problems on hypergraph decompositions that are designed to speed up the evaluation of **Conjunctive Queries (CQs)** and the solution of **Constraint**

This is an extended and enhanced version of the paper “HyperBench: A Benchmark and Tool for Hypergraphs and Empirical Findings,” presented at PODS 2019 [18].

This work was supported by the Austrian Science Fund (FWF):P30930-N35 in the context of the project “HyperTrac.” Georg Gottlob is a Royal Society Research Professor and acknowledges support by the Royal Society for the present work in the context of the project “RAISON DATA” (Project reference: RP\R1\201074). Davide Mario Longo’s work was also supported by the FWF project W1255-N23.

Authors’ addresses: W. Fischl, D. Mario Longo, and R. Pichler, TU Wien; emails: {wfischl, dlongo, pichler}@dbai.tuwien.ac.at; G. Gottlob, University of Oxford; email: georg.gottlob@cs.ox.ac.uk.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Association for Computing Machinery.

1084-6654/2020/07-ART1.6 \$15.00

<https://doi.org/10.1145/3440015>

Satisfaction Problems (CSPs). Hypergraph decompositions have meanwhile found their way into commercial database systems such as LogicBlox [5, 8, 36, 37, 46] and advanced research prototypes such as EmptyHeaded [1, 2, 47, 55]. Hypergraph decompositions have also been successfully used in the CSP area [4, 33, 39]. In theory, the pros and cons of various notions of decompositions and widths are well understood (see Reference [24] for a survey). However, from a practical point of view, many questions have remained open.

We want to collect hypergraphs from different application contexts, analyze their structural properties and, in particular, their (generalized) hypertree width (ghw and hw , respectively), and make this hypergraph collection together with the results of our analyses publicly available. The investigation of millions of CQs [12, 48] posed at various SPARQL endpoints suggests that these real-world CQs with atoms of arity ≤ 3 have very low hw : the overwhelming majority is acyclic; almost all of the rest has $hw = 2$. It is, however, not clear if CQs with arbitrary arity and CSPs also have low hypertree width, say, $hw \leq 5$. Ghionna et al. [22] gave a positive answer to this question for a small set of TPC-H benchmark queries. We significantly extend their collection of CQs. This also motivates our main goal.

Main Goal: Create a comprehensive, easily extensible benchmark of hypergraphs corresponding to CQs or CSPs for the analysis of hypergraph structural properties and decomposition algorithms.

Answering CQs and solving CSPs are fundamental tasks in Computer Science. Formally, they are the same problem, since both correspond to the evaluation of first-order formulae over a finite structure, such that the formulae only use $\{\exists, \wedge\}$ as connectives but not $\{\forall, \vee, \neg\}$. Both problems, answering CQs and solving CSPs, are NP-complete [15]. Consequently, the search for tractable fragments of these problems has been an active research area in the database and artificial intelligence communities for several decades.

The most powerful methods known to date for defining tractable fragments are based on various decompositions of the hypergraph structure underlying a given CQ or CSP. The most important forms of decompositions are *hypertree decompositions (HDs)* [27], *generalized hypertree decompositions (GHDs)* [27], and *fractional hypertree decompositions (FHDs)* [31]. These decomposition methods give rise to three notions of width of a hypergraph H : the *hypertree width* $hw(H)$, *generalized hypertree width* $ghw(H)$, and *fractional hypertree width* $fhw(H)$, where, $fhw(H) \leq ghw(H) \leq hw(H)$ holds for every hypergraph H . For definitions, see Section 3.

Both, answering CQs and solving CSPs, become tractable if the underlying hypergraphs have bounded hw , ghw , or, fhw and an appropriate decomposition is given. This gives rise to the problem of recognizing if a given CQ or CSP has hw , ghw , or, fhw bounded by some constant k . Formally, for *decomposition* $\in \{\text{HD, GHD, FHD}\}$ and $k \geq 1$, we consider the following family of problems:

CHECK(*decomposition*, k)

Input hypergraph $H = (V, E)$;

Output *decomposition* of H of width $\leq k$ if it exists and answer “no” otherwise.

Clearly, bounded fhw defines the largest tractable class while bounded hw defines the smallest one. However, the problem CHECK(HD, k) is feasible in polynomial time [27] while the CHECK(GHD, k) [28] and CHECK(FHD, k) [19] problems are NP-complete even for $k = 2$.

Systems to solve the CHECK(HD, k) problem exist [30, 50]. In contrast, for CHECK(GHD, k) and CHECK(FHD, k), apart from exhaustive search over possible decomposition trees (which only

works for small hypergraphs), no implementations have been reported yet [1]—with one exception: very recently, an interesting approach is presented in Reference [17], where SMT-solving is applied to the $\text{CHECK}(\text{FHD}, k)$ problem. The same approach has been later extended to solve the $\text{CHECK}(\text{HD}, k)$ problem [52]. In Reference [30], tests of the $\text{CHECK}(\text{HD}, k)$ system are presented. However, a benchmark for systematically evaluating systems for the $\text{CHECK}(\text{decomposition}, k)$ problem with $\text{decomposition} \in \{\text{HD}, \text{GHD}, \text{FHD}\}$ and $k \geq 1$ were missing so far. This motivates our first research subgoal.

Subgoal 1: Use the benchmark from the Main Goal to find out if the hypertree width is, in general, small enough (say, ≤ 5) to allow for efficient evaluation of CQs of arbitrary arity and of CSPs.

Recently, in Reference [19], the authors have identified classes of CQs for which the $\text{CHECK}(\text{GHD}, k)$ and $\text{CHECK}(\text{FHD}, k)$ problems become tractable (from now on, we only speak about CQs; of course, all results apply equally to CSPs). To this end, the **Bounded Intersection Property (BIP)** and, more generally, the **Bounded Multi-Intersection Property (BMIP)** have been introduced. The maximum number d of attributes shared by two (respectively, c) atoms is referred to as the intersection size (respectively, c -multi-intersection size) of the CQ, which is similar to the notion of cutset width from the CSP literature [16]. We say that a class of CQs satisfies the BIP (respectively, BMIP) if the number of attributes shared by two (respectively, by a constant number c of) query atoms is bounded by some constant d .

A related property is that of bounded degree, i.e., each attribute only occurs in a constant number of query atoms. Clearly, the BMIP is an immediate consequence of bounded degree. It has been shown in Reference [19] that $\text{CHECK}(\text{GHD}, k)$ is solvable in polynomial time for CQs whose underlying hypergraphs satisfy the BMIP. For CQs, the BMIP and bounded degree seem natural restrictions. For CSPs, the situation is not so clear. This yields the following research subgoals.

Subgoal 2: Use the hypergraph benchmark from the Main Goal to analyze how realistic the restrictions to low (multi-)intersection size, or low degree of CQs and CSPs are.

Subgoal 3: Verify that for hypergraphs of low intersection size, the $\text{CHECK}(\text{GHD}, k)$ problem indeed allows for efficient algorithms that work well in practice.

The tractability results for $\text{CHECK}(\text{FHD}, k)$ [19] are significantly weaker than for $\text{CHECK}(\text{GHD}, k)$: They involve a factor that is at least double-exponential in some “constant” (namely, k , the bound δ on the degree and/or the bound d on the intersection size). Hence, we want to investigate if (generalized) hypertree decompositions could be “fractionally improved” by taking the integral edge cover at each node in the HD or GHD and replacing it by a fractional edge cover. We will thus introduce the notion of *fractionally improved* HD that checks if there exists an HD of width $\leq k$, such that replacing each integral cover by a fractional cover yields an FHD of width $\leq k'$ for given bounds k, k' with $0 < k' < k$.

Subgoal 4: Explore the potential of fractionally improved HDs, i.e., investigate if the improvements achieved are significant.

In cases where $\text{CHECK}(\text{GHD}, k)$ and $\text{CHECK}(\text{FHD}, k)$ are intractable, we may have to settle for good approximations of ghw and fhw . For GHDs, we may thus use the inequality $ghw(H) \leq 3 \cdot hw(H) + 1$, which holds for every hypergraph H [3]. In contrast, for FHDs, the best-known general, polynomial-time approximation is cubic. More precisely, in Reference [43], a polynomial-time algorithm is presented that, given a hypergraph H with $fhw(H) = k$, computes an FHD of

width $O(k^3)$. In Reference [19], it is shown that a polynomial-time approximation up to a logarithmic factor is possible for any class of hypergraphs with bounded Vapnik–Chervonenkis dimension (VC-dimension; see Section 3 for a precise definition). The problem of efficiently approximating the ghw and/or fhw leads us to the following subgoals.

Subgoal 5: Use the benchmark from the Main Goal to analyze if, in practice, hw and ghw indeed differ by factor 3 or, if hw is typically much closer to ghw than this worst-case bound.

Subgoal 6: Use the benchmark from the Main Goal to analyze how realistic the restriction to small VC-dimension of CQs and CSPs is.

Results. Our main results are as follows:

- We provide *HyperBench*, a comprehensive hypergraph benchmark of initially over 3,000 hypergraphs (see Section 5). This benchmark is exposed by a web interface, which allows the user to retrieve the hypergraphs or groups of hypergraphs together with a broad spectrum of properties of these hypergraphs, such as lower/upper bounds on hw and ghw , (multi-)intersection size, degree, and so on. The system is accessible at <http://hyperbench.dbai.tuwien.ac.at/>. We also release a tool for extracting hypergraphs from SQL queries that was used in the construction of HyperBench. The tool is available at <https://github.com/dmlongo/hgtools>.

- We extend the software for HD computation from Reference [30] to also solve the $\text{CHECK}(\text{GHD}, k)$ problem. For a given hypergraph H , our system first computes the intersection size of H and then applies the ghw -algorithm from Reference [19], which is parameterized by the intersection size. We implement several improvements and we further extend the system to compute also “fractionally improved” HDs. All implementations are available at <https://github.com/dmlongo/newdetkdecomp>.

- We carry out an empirical analysis of the hypergraphs in the HyperBench benchmark. This analysis demonstrates, especially for real-world instances, that the restrictions to BIP, BMIP, bounded degree, and bounded VC-dimension are astonishingly realistic. Moreover, on all hypergraphs in the HyperBench benchmark, we run our hw - and ghw -programs to identify (or at least bound) their hw and ghw . An interesting observation of our empirical study is that apart from the CQs also a significant portion of CSPs in our benchmark has small hypertree width (all non-random CQs have $hw \leq 3$ and over 60% of CSPs stemming from applications have $hw \leq 5$). Moreover, for $hw \leq 5$, in all of the cases where the ghw -computation terminates, hw and ghw have identical values.

- In our study of the ghw of the hypergraphs in the HyperBench benchmark, we observe that a straightforward implementation of the algorithm from Reference [19] for hypergraphs of low intersection size is too slow in many cases. We therefore present a new approach (based on so-called “balanced separators”) with promising experimental results. It is interesting to note that the new approach works particularly well in those situations that are particularly hard for the straightforward implementation, namely, hypergraphs H where the test if $ghw \leq k$ for given k gives a “no”-answer. Hence, combining the different approaches is very effective. Also, this implementation is provided at <https://github.com/dmlongo/newdetkdecomp>.

Structure. This article is structured as follows: In Section 2, we summarize related work. In Section 3, we recall some basic notions. In Section 4, we describe our algorithms for solving the $\text{CHECK}(\text{GHD}, k)$ problem. In Section 5, we present our system and test environment as well as our HyperBench benchmark. This section also contains the first results of our empirical study of the hypergraphs in this benchmark. In Section 6, we report on the performance of the implementations of our GHD-algorithms as well as a further extension of the system to allow for the computation of

fractionally improved HDs. We conclude in Section 7 by highlighting the most important lessons learned from our empirical study and by identifying some appealing directions for future work.

2 RELATED WORK

We distinguish several types of works that are highly relevant to ours. The works most closely related are the descriptions of HD, GHD, and FHD algorithms in References [19, 27] and the implementation of HD computation by the DetKDecomp program reported in Reference [30]. We have extended these works in several ways. Above all, we have incorporated our analysis tool (reported in Sections 5 and 6) and the GHD and FHD computations (reported in Sections 4 and 6.5) into the DetKDecomp program—resulting in our NewDetKDecomp library, which is openly available on GitHub at <https://github.com/dmlongo/newdetkdecomp>. For the GHD computation, we have added heuristics to speed up the basic algorithm from Reference [19]. Moreover, we have proposed a novel approach via balanced separators, which allowed us to significantly extend the range of instances for which the GHD computation terminates in reasonable time. We have also introduced a new form of decomposition method: the fractionally improved decompositions (see Section 6.5), which allow for a practical, lightweight form of FHDs.

The second important input to our work comes from the various sources [6, 9–11, 21, 30, 34, 40, 54] which we took our CQs and CSPs from. Note that our main goal was not to add further CQs and/or CSPs to these benchmarks. Instead, we have aimed at taking and combining existing, openly accessible benchmarks of CQs and CSPs, and convert them into hypergraphs, which are then thoroughly analyzed. Finally, the hypergraphs and the analysis results are made openly accessible again.

The third kind of works highly relevant to ours are previous analyses of CQs and CSPs. To the best of our knowledge, Ghionna et al. [22] presented the first systematic study of HDs of benchmark CQs from TPC-H. However, Ghionna et al. pursued a research goal different from ours in that they primarily wanted to find out to what extent HDs can actually speed up query evaluation. They achieved very positive results in this respect, which have recently been confirmed by the work of Perelman et al. [47], Tu et al. [55], and Aberger et al. [1] on query evaluation using FHDs. As a side result, Ghionna et al. also detected that CQs tend to have low hypertree width (a finding that was later confirmed in References [12, 13, 48] and also in our study). In a pioneering effort, Bonifati, Martens, and Timm [12] have recently analyzed an unprecedented, massive amount of queries: They investigated 180,653,910 queries from (not openly available) query logs of several popular SPARQL endpoints. After elimination of duplicate queries, there were still 56,164,661 queries left, out of which 26,157,880 queries were in fact CQs. The authors thus significantly extend previous work by Picalausa and Vansummeren [48], who analyzed 3,130,177 SPARQL queries posed by humans and software robots at the DBPedia SPARQL endpoint. The focus in Reference [48] is on structural properties of SPARQL queries such as keywords used and variable structure in optional patterns. There is one paragraph devoted to CQs, where it is noted that 99.99% of circa 2M CQs considered in Reference [48] are acyclic.

Many of the CQs (over 15M) analyzed in Reference [12] have arity 2 (here, we consider the maximum arity of all atoms in a CQ as the arity of the query), which means that all triples in such a SPARQL query have a constant at the predicate-position. Bonifati et al. made several interesting observations concerning the shape of these graph-like queries. For instance, they detected that exactly one of these queries has $tw = 3$, while all others have $tw \leq 2$ (and hence $hw \leq 2$). As far as the CQs of arity 3 are concerned (for CQs expressed as SPARQL queries, this is the maximum arity achievable), among many characteristics, also the hypertree width was computed by using the original DetKDecomp program from Reference [30]. Out of 6,959,510 CQs of arity 3, only 86 (i.e., 0.01%) turned out to have $hw = 2$ and 8 queries had $hw = 3$, while all other CQs of arity 3

are acyclic. Our analysis confirms that, also for non-random CQs of arity > 3 , the hypertree width indeed tends to be low, with the majority of queries being even acyclic.

Bonifati et al. continued on this line of work and analyzed in Reference [13] a yet bigger collection of SPARQL queries coming from Wikidata. This repository of 208,215,209 queries was first made available by Malyshev et al. for the work in Reference [42]. Bonifati et al. divided this dataset into four disjoint sets: queries for which the HTTP request was successful, further partitioned into organic and robotic queries; and timeout queries, again further partitioned into organic and robotic queries. Organic queries are the ones classified by Malyshev et al. as posed by humans, while robotic queries have been classified as produced by synthetic algorithms. While Malyshev et al. analyzed only the sets of successful queries, Bonifati et al. extended their study to timeout queries. The latter turned out to be the most interesting w.r.t. structural analysis. As in Reference [12], the focus of Reference [13] is on examining property paths of SPARQL queries and, in particular, having a clear picture of the structural characteristics of recursive properties.

In Reference [13], structural analysis builds upon conjunctive queries and variants thereof. In total, Bonifati et al. identify 176,679,495 robotic and 342,576 organic queries for the largest fragment of CQs (namely, C2RPQ+), which constitutes circa 85% of the whole dataset. Structural analysis showed that these queries are *mildly cyclic*, i.e., their treewidth is bounded by a small constant. In particular, for all these queries $tw \leq 4$ holds. For a different fragment of CQs (namely, CQ_{OF+}), a structural analysis based on hypergraphs is more suitable, thus they computed hw for a total of 1,915,550 CQ_{OF+} queries. It turns out that 590,005 queries have $hw = 2$, while the rest has $hw = 1$, i.e., they are acyclic.

For the analysis of CSPs, much less work has been done. Although it has been shown that exploiting (hyper-)tree decompositions may significantly improve the performance of CSP solving References [4, 33, 35, 39], a systematic study on the (generalized) hypertree width of CSP instances has only been carried out by few works [30, 39, 51]. To the best of our knowledge, we are the first to analyze the hw , ghw , and fhw of circa 2,000 CSP instances, where most of these instances have not been studied in this respect before.

It should be noted that the focus of our work is different from the above-mentioned previous works: Above all, we wanted to test the practical feasibility of various algorithms for HD, GHD, and FHD computation (including both, previously presented algorithms and new ones developed as part of this work). As far as our repository of hypergraphs (obtained from CQs and CSPs) is concerned, we emphasize open accessibility. Thus, users can analyze their CQs and CSPs (with our implementations of HD, GHD, and FHD algorithms) or they can analyze new decomposition algorithms (with our hypergraphs, which cover quite a broad range of characteristics). In fact, in the recent works on HD and FHD computation via SMT solving [17, 52], the HyperBench benchmark has already been used for the experimental evaluation. In Reference [17] a novel approach to fhw computation via an efficient encoding of the check-problem for FHDs to SMT (SAT modulo Theory) is presented. The tests were carried out with 2,191 hypergraphs from the initial version of the HyperBench. For all of these hypergraphs, we have established at least some upper bound on the fhw either by our hw -computation or by one of our new algorithms presented in Sections 4 and 6.5. In contrast, the exact algorithm in Reference [17] found FHDs only for 1.449 instances (66%). In 852 cases, both our algorithms and the algorithm in Reference [17] found FHDs of the same width; in 560 cases, an FHD of lower width was found in Reference [17]. By using the same benchmark for the tests, the results in Reference [17] and ours are comparable and have thus provided valuable input for future improvements of the algorithms by combining the different strengths and weaknesses of the two approaches.

The use of the same benchmark has also allowed us to provide feedback to the authors of Reference [17] for debugging their system: In 9 out of 2,191 cases, the “optimal” value for the fhw

computed in Reference [17] was apparently erroneous, since it was higher than the hw found out by our analysis; note that upper bounds on the width are, in general, more reliable than lower bounds, since it is easy to verify if a given decomposition indeed has the desired properties, whereas ruling out the existence of a decomposition of a certain width is a complex and error-prone task.

The work in Reference [29] represents a follow-up of this article and the conference version [18]. The main goal of Reference [29] consists in providing major improvements in computing GHDs. To this aim, the authors present a parallel algorithm for computing GHDs based on the balanced separator method described in this article and a hybrid approach that combines the parallel and sequential decomposition algorithms. More specifically, the hybrid approach uses a parallel version of the balanced separator algorithm to split a large hypergraph into smaller components and then uses the sequential version of DetKDecomp from Reference [30] to quickly decompose them. Moreover, the authors of Reference [29] propose new methods to simplify the input hypergraph and apply new heuristics to reduce the search space. Our HyperBench benchmark was fundamental for the experimental evaluation of all these techniques. Indeed, the version of HyperBench presented in Reference [18] was used as a baseline to show that the newly proposed algorithms can be used to efficiently compute GHDs on modern machines for a wide range of CSP instances.

Similarly to our analysis of properties that make $\text{CHECK}(\text{decomposition}, k)$ tractable or easy to approximate, HyperBench has been used to empirically test the validity of theoretical hypotheses. In Reference [38], the edge clique cover size of a graph is identified as a parameter allowing fixed-parameter-tractable algorithms for enumerating potential maximal cliques. The latter can be used to compute exact ghw and fhw . An edge clique cover of a graph is a set of cliques of the graph that covers all of its edges. In case of a CSP with n variables and m constraints, the set of constraints is an edge clique cover of the underlying (hyper)graph. Thus, this property can be exploited for CSPs having $n > m$ and HyperBench has been used to verify that it happens in circa 23% of the instances.

3 PRELIMINARIES

3.1 CQs, CSPs, and Hypergraphs

We treat **conjunctive queries(CQs)** and **constraint satisfaction problems(CSPs)** as first-order formulae using only connectives in $\{\exists, \wedge\}$ and disallowing $\{\forall, \vee, \neg\}$.

A *hypergraph* $H = (V(H), E(H))$ is a pair consisting of a set of vertices $V(H)$ and a set of non-empty (hyper)edges $E(H) \subseteq 2^{V(H)}$. We assume w.l.o.g. that there are no isolated vertices, i.e., for each $v \in V(H)$, there is at least one edge $e \in E(H)$ such that $v \in e$. We can thus identify a hypergraph H with its set of edges $E(H)$ with the understanding that $V(H) = \{v \in e \mid e \in E(H)\}$. A *subhypergraph* H' of H is then simply a subset of (the edges of) H .

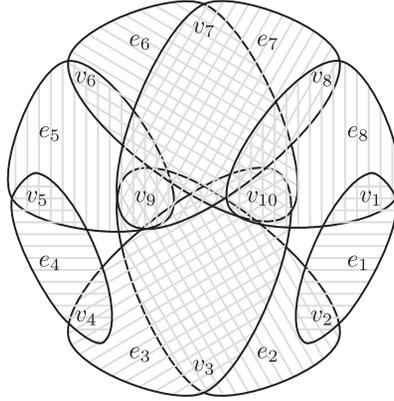
Given a formula ϕ corresponding to either a CQ or a CSP, the hypergraph H_ϕ corresponding to ϕ has $V(H_\phi) = \text{vars}(\phi)$ and, for each atom a of ϕ , $\text{vars}(a) \in E(H_\phi)$.

We are frequently dealing with sets of sets of vertices (e.g., sets of edges). For $S \subseteq 2^{V(H)}$, we write $\bigcup S$ and $\bigcap S$ as a short-hand for taking the union or intersection, respectively, of this set of sets of vertices, i.e., for $S = \{s_1, \dots, s_\ell\}$, we have $\bigcup S = \bigcup_{i=1}^\ell s_i$ and $\bigcap S = \bigcap_{i=1}^\ell s_i$.

Example 1. Consider the following Boolean conjunctive query Q_0 :

$$Q_0 \leftarrow e_1(v_1, v_2), e_2(v_2, v_3, v_9), e_3(v_3, v_4, v_{10}), e_4(v_4, v_5), e_5(v_5, v_6, v_9), \\ e_6(v_6, v_7, v_{10}), e_7(v_7, v_8, v_9), e_8(v_1, v_8, v_{10}).$$

The hypergraph H_0 underlying this query is shown in Figure 1.

Fig. 1. Hypergraph H_0 of Example 3.1.

3.2 Hypergraph Decompositions and Width Measures

We consider here three notions of hypergraph decompositions with their associated notions of width. To this end, we first introduce the notion of (fractional) edge covers.

Let $H = (V(H), E(H))$ be a hypergraph and consider an *edge weight function* $\gamma : E(H) \rightarrow [0, 1]$. We define the set $B(\gamma)$ of all vertices covered by γ and the weight of γ as

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

$$\text{weight}(\gamma) = \sum_{e \in E(H)} \gamma(e).$$

We call γ a *fractional edge cover* of a set $X \subseteq V(H)$ by edges in $E(H)$, if $X \subseteq B(\gamma)$. We also consider an *integral edge cover* as a function $\lambda : E(H) \rightarrow \{0, 1\}$, i.e., a fractional edge cover whose values are restricted to $\{0, 1\}$ values. Following Reference [27], we can also treat λ as a set with $\lambda \subseteq E(H)$ (namely, the set of edges e with $\lambda(e) = 1$) and the weight as the cardinality of such a set of edges. In the following, to emphasize the nature of the function we are dealing with, we will use γ for fractional edge covers and λ for integral edge covers.

We now introduce some relevant notions of decompositions.

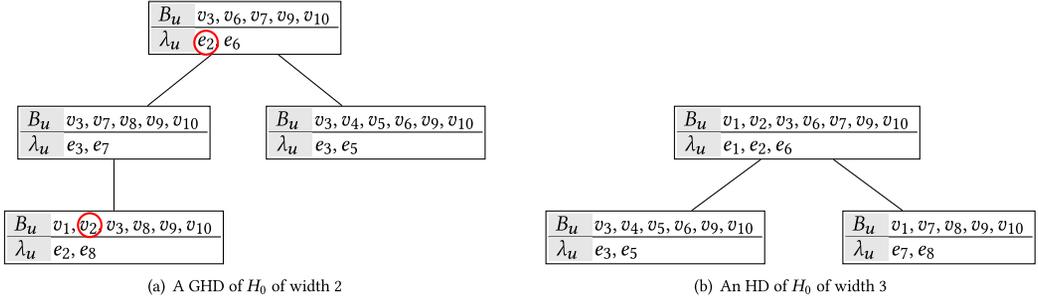
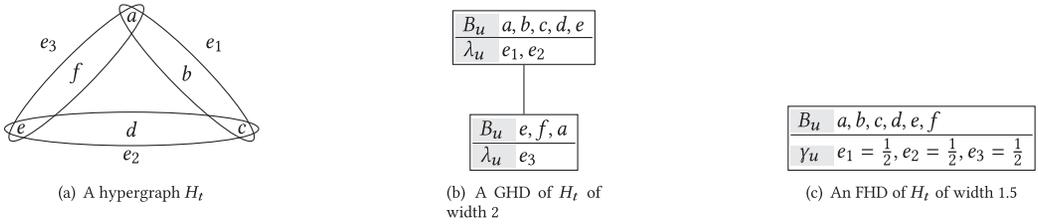
A tuple $\langle T, (B_u)_{u \in T} \rangle$ is a **tree decomposition (TD)** of hypergraph $H = (V(H), E(H))$, if $T = (N(T), E(T))$ is a tree, every B_u is a subset of $V(H)$, and the following conditions are satisfied:

- (1) For every edge $e \in E(H)$, there is a node u in T , such that $e \subseteq B_u$, and
- (2) for every vertex $v \in V(H)$, $\{u \in T \mid v \in B_u\}$ is connected in T .

The vertex sets B_u are usually referred to as the *bags* of the TD. By slight abuse of notation, we write $u \in T$ to express that u is a node in $N(T)$. Condition (2) is also called the “connectedness condition.”

We use the following notational conventions throughout this article: To avoid confusion, we will consequently refer to the elements in $V(H)$ as *vertices* of the hypergraph and to the elements in $N(T)$ as the *nodes* of the decomposition. For a node $u \in T$, we write T_u to denote the subtree of T rooted at u . By slight abuse of notation, we will often write $u' \in T_u$ to denote that u' is a node in the subtree T_u of T . Finally, we define $V(T_u) = \bigcup_{u' \in T_u} B_{u'}$.

A **fractional hypertree decomposition (FHD)** of a hypergraph $H = (V(H), E(H))$ is a tuple $\langle T, (B_u)_{u \in T}, (\gamma_u)_{u \in T} \rangle$, such that $\langle T, (B_u)_{u \in T} \rangle$ is a TD of H and the following condition holds:

Fig. 2. Comparison between a GHD and an HD of H_0 .Fig. 3. Comparison between a GHD and an FHD of H_t .

(3) For each $u \in T$, $B_u \subseteq B(\gamma_u)$ holds, i.e., γ_u is a fractional edge cover of B_u .

Following our notational convention, a **generalized hypertree decomposition (GHD)** is an FHD, where λ_u is an integral edge weight function for every $u \in T$. Hence, by condition (3), λ_u is an integral edge cover of B_u . A **hypertree decomposition (HD)** of H is a GHD with the following additional condition (referred to as the “special condition” in Reference [27]):

(4) For each $u \in T$, $V(T_u) \cap B(\gamma_u) \subseteq B_u$, where $V(T_u)$ denotes the union of all bags in the subtree of T rooted at u .

Because of condition (4), it is important to consider T as a *rooted tree* in case of HDs. For TDs, FHDs, and GHDs, the root of T can be arbitrarily chosen or simply ignored. The *width* of an FHD, GHD, or HD is defined as the maximum weight of the functions γ_u over all nodes $u \in T$. The fractional hypertree width, generalized hypertree width, and hypertree width of H (denoted $fhw(H)$, $ghw(H)$, and $hw(H)$) is the minimum width over all FHDs, GHDs, and HDs of H .

Example 2. Figure 2 shows two decompositions of the hypergraph H_0 of Figure 1. In particular, Figure 2(a) is a GHD of H_0 , while Figure 2(b) is an HD of H_0 . It is easy to check that the two decompositions satisfy Conditions (1)–(3), but only Figure 2(b) satisfies Condition (4). In fact, while Figure 2(b) is also a GHD (by definition), Figure 2(a) is not an HD. Indeed, the λ -label of the root of Figure 2(a) contains the edge e_2 , which in turn contains the vertex v_2 . Nevertheless, this vertex does not appear in the bag of the root. While this is allowed in a GHD, it violates Condition (4) of HDs, as v_2 could have been “covered” already in the root, but it appears only in a leaf of the decomposition. Moreover, the two decompositions have different widths. In fact, it can be proved that $ghw(H_0) = 2$ and $hw(H_0) = 3$.

Example 3. Figure 3 shows the difference between a GHD and an FHD of the simple triangle hypergraph H_t in Figure 3(a). Figure 3(b) is a GHD of H_t of width 2 and, since H_t is cyclic, no GHD of H_t of width 1 exists. However, Figure 3(c) is an FHD of H_t of width 1.5. This is possible because

in an FHD, we consider fractional edge covers of the bags instead of (integral) edge covers. Thus, for the hypergraph H_t , $fhw(H_t) < ghw(H_t)$ holds.

3.3 Components and Separators

For a set $U \subseteq V(H)$ of vertices, we define $[U]$ -components of a hypergraph H as follows:

- We define $[U]$ -adjacency as a binary relation on $E(H)$ as follows: two edges $e_1, e_2 \in E(H)$ are $[U]$ -adjacent, if $(e_1 \cap e_2) \setminus U \neq \emptyset$ holds.
- We define $[U]$ -connectedness as the transitive closure of the $[U]$ -adjacency relation.
- A $[U]$ -component of H is a maximally $[U]$ -connected subset $C \subseteq E(H)$.

For a set of edges $S \subseteq E(H)$, we say that C is “ $[S]$ -connected” or an “ $[S]$ -component” as a shortcut for C is “ $[W]$ -connected” or a “ $[W]$ -component,” respectively, with $W = \bigcup_{e \in S} e$. We also call S a *separator* in this context. The *size of an $[S]$ -component C* is defined as the number of edges in the component. For a hypergraph H and a set of edges $S \subseteq E(H)$, we say that S is a *balanced separator* if all $[S]$ -components of H have a size $\leq \frac{|E(H)|}{2}$. It was shown in Reference [3] that, for every GHD $\langle T, (B_u)_{u \in T}, (\lambda_u)_{u \in T} \rangle$ of a hypergraph H , there exists a node $u \in T$ such that λ_u is a balanced separator of H . This property can be made use of when searching for a GHD of size k of a hypergraph H , as we shall show in Section 4.4 below.

3.4 Computing Hypertree Decompositions (HDs)

We briefly recall the basic principles of the DetKDecomp program from Reference [30] for computing HDs. It is relevant for our work in that it is the first implementation of the original top-down nondeterministic HD algorithm presented in Reference [27]. Even though the first implementation of a deterministic HD algorithm is OptKDecomp [26], it is based on a different characterization of hw , which is not suitable for our purposes.

For fixed $k \geq 1$, DetKDecomp tries to construct an HD of a hypergraph H in a top-down manner. It thus maintains a set C of edges, which is initialized to $C = E(H)$. For a node u in the HD (initially, this is the root of the HD), it “guesses” an edge cover λ_u , i.e., $\lambda_u \subseteq E(H)$ and $|\lambda_u| \leq k$. For fixed k , there are only polynomially many possible values λ_u . DetKDecomp then proceeds by determining all $[\lambda_u]$ -components C_i with $C_i \subseteq C$. The special condition of HDs restricts the possible choices for B_u and thus guarantees that the $[\lambda_u]$ -components inside C and the $[B_u]$ -components inside C coincide. This is the crucial property for ensuring polynomial time complexity of HD-computation—at the price of possibly missing GHDs with a lower width.

Now let C_1, \dots, C_ℓ denote the $[\lambda_u]$ -components with $C_i \subseteq C$. By the maximality of components, these sets C_i are pairwise disjoint. Moreover, it was shown in Reference [27] that if H has an HD of width $\leq k$, then it also has an HD of width $\leq k$ such that the edges in each C_i are “covered” in different subtrees below u . More precisely, this means that u has ℓ child nodes u_1, \dots, u_ℓ , such that for every i and every $e \in C_i$, there exists a node u_e in the subtree rooted at u_i with $e \subseteq B_{u_e}$. Hence, DetKDecomp recursively searches for an HD of the hypergraphs H_i with $E(H_i) = C_i$ and $V(H_i) = \bigcup C_i$ with the slight extra feature that also edges from $E(H) \setminus C_i$ are allowed to be used in the λ -labels of these HDs.

Example 4. Consider again the hypergraph H_0 of Figure 1 and its HD in Figure 2(b). The HD tells us exactly what are the choices made by the algorithm while constructing the tree. Let us thus simulate the run of DetKDecomp, which produced the HD in Figure 2(b), for $k = 3$. The input of DetKDecomp is the whole set $C = E(H_0)$ of edges of H_0 . At the beginning, the algorithm guesses the edge cover $\lambda_u = \{e_1, e_2, e_6\}$ and determines the respective bag $B_u = \bigcup \lambda_u$. Then, it uses λ_u to separate C in two components $C_1 = \{e_3, e_4, e_5\}$, $C_2 = \{e_7, e_8\}$, which we show in Figure 4. Since each

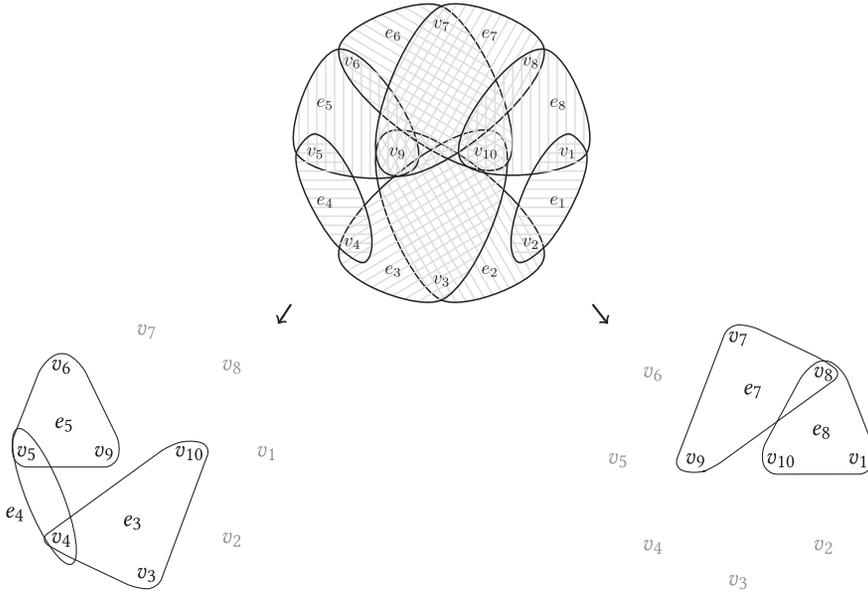


Fig. 4. The two components C_1, C_2 of Example 3.4 obtained by removing $\{e_1, e_2, e_6\}$ from H_0 .

of these components can be “covered” by at most two edges, DetKDecomp creates the two leaves of Figure 2(b), attaches them to the root node previously created, and outputs the resulting HD.

3.5 Favorable Properties of Hypergraphs

We are interested in certain structural properties of hypergraphs that make the CHECK(GHD, k) and CHECK(FHD, k) problems tractable or efficient to approximate for large classes of hypergraphs. We refer to the terminology of Reference [25], as it makes uniform the one originally introduced in Reference [19].

Definition 1. For $c \geq 1$, $d \geq 0$, a hypergraph $H = (V(H), E(H))$ is a (c, d) -hypergraph if the intersection of any c edges in $E(H)$ has at most d elements, i.e., for every subset $E' \subseteq E(H)$ with $|E'| = c$, we have $|\bigcap E'| \leq d$.

Definition 2. A hypergraph $H = (V(H), E(H))$ has c -multi-intersection size d if H is a (c, d) -hypergraph. In the special case of $c = 2$, we speak of intersection size of H , while if we do not have a particular c in mind, we simply speak of multi-intersection size of H .

Definition 3. A class C of hypergraphs satisfies the bounded multi-intersection property (BMIP), if there exist $c \geq 1$ and $d \geq 0$, such that every H in C is a (c, d) -hypergraph. As a special case, C satisfies the bounded intersection property (BIP), if there exists $d \geq 0$, such that every H in C is a $(2, d)$ -hypergraph.

There are further relevant properties of (classes of) hypergraphs: bounded degree and bounded **Vapnik–Chervonenkis dimension (VC-dimension)**.

Definition 4. The *degree* $\deg(H)$ of a hypergraph H is defined as the maximum number δ of hyperedges in which a vertex occurs, i.e., $\delta = \max_{v \in V(H)} |\{e \in E(H) \mid v \in e\}|$. A class C of hypergraphs has bounded degree, if there exists $\delta \geq 1$, such that every hypergraph H in C has degree $\leq \delta$.

Definition 5 ([56]). Let $H = (V(H), E(H))$ be a hypergraph, and $X \subseteq V(H)$ a set of vertices. Denote by $E(H)|_X = \{X \cap e \mid e \in E(H)\}$. X is called *shattered* if $E(H)|_X = 2^X$. The *Vapnik-Chervonenkis dimension* (*VC dimension*) of H is the maximum cardinality of a shattered subset of V . We say that a class \mathcal{C} of hypergraphs has bounded VC-dimension, if there exists $v \geq 1$, such that every hypergraph $H \in \mathcal{C}$ has VC-dimension $\leq v$.

Note that a hypergraph H with degree bounded by δ is a $(\delta + 1, 0)$ -hypergraph. Thus, bounded degree implies the BMIP, which in turn implies bounded VC-dimension [19].

The aforementioned properties help to solve or approximate the $\text{CHECK}(\text{GHD}, k)$ and $\text{CHECK}(\text{FHD}, k)$ problems:

THEOREM 1 ([19, 25]). *Let \mathcal{C} be a class of hypergraphs.*

- *If \mathcal{C} has the BMIP, then the $\text{CHECK}(\text{GHD}, k)$ problem is solvable in polynomial time for arbitrary $k \geq 1$. Consequently, this tractability holds if \mathcal{C} has bounded degree or the BIP (which each imply the BMIP) [19].*
- *If \mathcal{C} has bounded degree or the BIP, then the $\text{CHECK}(\text{FHD}, k)$ problem is solvable in polynomial time for arbitrary $k \geq 1$ [19, 25].*
- *If \mathcal{C} has bounded VC-dimension, then the fhw can be approximated in polynomial time up to a log-factor [19].*

4 GHD ALGORITHMS

In this section, we present two different ways to implement the tractable algorithm for computing GHDs in case of BIP [19]. We shall refer to these two implementations as `GlobalBIP` and `LocalBIP`. Moreover, we shall present a completely new approach to computing GHDs, which will be referred to as `BalSep`. These variants exploit low intersection size to compute a GHD of width $\leq k$ of a hypergraph in polynomial time. While `GlobalBIP` and `LocalBIP` differ in the time of the computation of certain sets of subedges, the `BalSep` algorithm is a novel approach in computing decompositions by means of *balanced separators*.

4.1 Theoretical Background

In Reference [19], the $\text{CHECK}(\text{GHD}, k)$ problem was proved to be intractable for $k \geq 2$, but the authors also identified tractable fragments corresponding to classes of hypergraphs having bounded (multi-)intersection size. In the following, we focus on the intersection size, i.e., the maximum intersection size of any two edges in a hypergraph, and explain how this property can be used for tractable ghw computation. Since the three variants presented later rely on the same central idea, we refer to the theoretical algorithm as the *ghw-algorithm*.

For a given hypergraph $H = (V(H), E(H))$, the ghw -algorithm adds a polynomial-time computable set $f(H, k)$ of subedges of edges in $E(H)$ to the hypergraph H . The set $f(H, k)$ has the property that $\text{ghw}(H) = k$ if and only if $\text{hw}(H') = k$, where $H' = (V(H), E(H) \cup f(H, k))$. Thus, it is possible to use $\text{CHECK}(\text{HD}, k)$ to solve $\text{CHECK}(\text{GHD}, k)$ and achieve tractability. The particular set

$$f(H, k) = \bigcup_{e \in E(H)} \left(\bigcup_{e_1, \dots, e_j \in (E(H) \setminus \{e\}), j \leq k} 2^{(e \cap (e_1 \cup \dots \cup e_j))} \right) \quad (1)$$

contains, for each $e \in E(H)$, all subsets of intersections of e with up to k edges of H different from e . Although $f(H, k)$ could in general contain an exponential number of elements, for fixed k and intersection size of H bounded by d , the set $e \cap (e_1 \cup \dots \cup e_j)$ contains at most $d \cdot k$ elements and, therefore, $|f(H, k)|$ is polynomially bounded.

Since the *ghw*-algorithm relies on the computation of HDs, we implemented a new version of the *hw*-algorithm in Reference [30] and called it NewDetKDecomp. Our program contains also the implementations of all the algorithms in this section, but we defer a detailed discussion of the complete library to Section 6.

4.2 The GlobalBIP Algorithm

Given H and $k \geq 1$, a straightforward implementation of the *ghw*-algorithm consists in computing the set $f(H, k)$, creating the hypergraph $H' = (V(H), E(H) \cup f(H, k))$, and finally computing an HD of H of width $\leq k$, if it exists. Since here we compute the set $f(H, k)$ *a priori globally* for the whole hypergraph, we call this algorithm GlobalBIP.

Algorithm 1 is a detailed description of GlobalBIP. The input consists of a hypergraph H and a constant k , while the output is a GHD of H of width $\leq k$, if it exists, and NULL otherwise. In line 2, we first compute $f(H, k)$ as in Equation (1) and then, in line 3, we create the hypergraph H' , which is obtained by adding the subedges in $f(H, k)$ to H . In line 4, we call NewDetKDecomp on H' and k as a black box and store its output in the variable HD . If HD is NULL, then a GHD of width $\leq k$ of H does not exist, therefore, we return NULL. Otherwise, we need to “fix” the decomposition as described in lines 6–10. In particular, for each node u of HD , and for each edge of $f(H, k)$ in λ_u , i.e., $e' \in (\lambda_u \cap f(H, k))$ that is not an edge in the original hypergraph H , we substitute e with an edge $e' \in E(H)$ such that $e \subseteq e'$. In this way, we obtain a new edge cover λ'_u such that $B_u \subseteq B(\lambda_u) \subseteq B(\lambda'_u)$, but also $|\lambda_u| = |\lambda'_u|$. Thus, the new decomposition still satisfies all the properties of a GHD and its width is still $\leq k$. Eventually, in line 11, we return HD .

ALGORITHM 1: GlobalBIP

Input: A hypergraph H .
Parameter: An integer $k \geq 1$.
Output: A GHD of H of width $\leq k$ if it exists, NULL otherwise.

```

1 begin
2    $f(H, k) \leftarrow$  compute as in Equation (1);
3    $H' \leftarrow (V(H), E(H) \cup f(H, k))$ ;
4    $HD \leftarrow$  NewDetKDecomp( $H', k$ );
5   if  $HD \neq \text{NULL}$  then
6     foreach  $u \in HD$  do
7       foreach  $e \in (\lambda_u \cap f(H, k))$  do
8          $e' \leftarrow e' \in E(H)$  such that  $e \subseteq e'$ ;
9          $\lambda_u(e) \leftarrow 0$ ;
10         $\lambda_u(e') \leftarrow 1$ ;
11  return  $HD$ 

```

4.3 The LocalBIP Algorithm

The main drawback of GlobalBIP is that the size of $f(H, k)$, though polynomial, could be huge for practical purposes. Therefore, we looked at ways to reduce the number of edges to add to H by restricting the computation only to those edges that might be actually necessary. The approach we used follows from an observation about the role played by $f(H, k)$ in the tractability proof in Reference [19].

Recall that the proof uses CHECK(HD, k) on the hypergraph H' to answer CHECK(GHD, k) for the hypergraph H . To do this in a sound way, the set $f(H, k)$ has to contain all the edges that could be used to cover possible bags of H' in an HD without changing the width. Consider a top-down construction of a GHD of H . At some point, we might want to choose, for some node u , a

bag B_u such that $x \notin B_u$ for some variable $x \in B(\lambda_u) \cap V(T_u)$. This choice would violate condition (4) of HDs and would not be allowed for the computation of an HD. In particular, there is an edge e with $x \in e$ and $\lambda_u(e) = 1$. For this reason, the set $f(H, k)$ contains an edge e' such that $e' \subset e$ and $x \notin e'$. Hence, we can substitute e with e' in the cover λ_u (i.e., $\lambda_u(e) = 0$, $\lambda_u(e') = 1$) to eliminate the violation of condition (4). Moreover, because of the connectedness condition, there is no need to look at the intersection of e with arbitrary edges in $E(H)$; instead, we consider only the intersections of e with unions of edges that may possibly occur in bags of T_u . In other words, for each node u of the decomposition, we consider only an appropriate subset $f_u(H, k) \subseteq f(H, k)$. More specifically, for the current node u , let $H_u \subseteq H$ be the component we want to decompose. Then, we define $f_u(H, k)$ as follows:

$$f_u(H, k) = \bigcup_{e \in E(H)} \left(\bigcup_{e_1, \dots, e_j \in (E(H_u) \setminus \{e\}), j \leq k} 2^{(e \cap (e_1 \cup \dots \cup e_j))} \right). \quad (2)$$

We call the resulting algorithm `LocalBIP`, because the set of edges $f_u(H, k)$ is computed *locally* for each node u during the construction of the decomposition. It follows `NewDetKDecomp` closely, but it differs in the search of the separators. In particular, while decomposing H , the algorithm first tries all possible ℓ -combinations of edges in $E(H)$ (with $\ell \leq k$) and only if the search does not succeed, it tries ℓ -combinations of subedges in $f_u(H, k)$.

4.4 The BalSep Algorithm

So far, we have presented two adaptations of the theoretical *ghw*-algorithm from Reference [19]. On the one hand, they extend `NewDetKDecomp` to compute GHDs and exploit bounded intersection size for tractability. On the other hand, they do not introduce any significant algorithmic innovation. In the following, we describe a novel approach to compute GHDs that makes use of sets of edges called *balanced separators*. We first extend the terminology of Section 3, then give a detailed description of the algorithm, and finally prove that our algorithm is sound and complete.

4.4.1 Balanced Separators and Special Edges. Recall that a *hypergraph* is a pair $H = (V(H), E(H))$, consisting of a set $V(H)$ of *vertices* and a set $E(H)$ of hyperedges (or, simply *edges*), which are non-empty subsets of $V(H)$. Since we assume that hypergraphs do not have isolated vertices, we can identify a hypergraph H with its set of edges $E(H)$. Then, a subhypergraph H' of H is a subset of (the edges of) H .

Starting off with a hypergraph H , the `BalSep` algorithm has to deal with subhypergraphs $H' \subseteq H$ augmented by a set S_p of special edges. A special edge is simply a set of vertices from H . Intuitively, special edges correspond to bags B_u in a GHD of H . Thus, an *extended subhypergraph* of H is of the form $H' \cup S_p$, where $H' \subseteq H$ is a subhypergraph and S_p is a set of special edges.

We now extend three crucial definitions from hypergraphs to extended subhypergraphs, namely, components, balanced separators, and GHDs. We recall that, even though a separator is a set of vertices, it can be defined as a set of edges. Then, for $S \subseteq E(H)$, an $[S]$ -component is a $[W]$ -component with $W = \bigcup_{e \in S} e$. In our algorithm, we use the fact that, because of the edges in $f(H, K)$, it is always possible to choose a separator λ_u such that $B(\lambda_u) = B_u$ [19]. Hence, there will not be any need to distinguish between vertex and edge separators. We start with components.

Definition 6 (Components of Extended Subhypergraphs). For a set $U \subseteq V(H)$ of vertices, we define $[U]$ -components of an extended subhypergraph $H' \cup S_p$ of H as follows:

- We define $[U]$ -adjacency as a binary relation on $H' \cup S_p$ as follows: two (possibly special) edges $f_1, f_2 \in H' \cup S_p$ are $[U]$ -adjacent, if $(f_1 \cap f_2) \setminus U \neq \emptyset$ holds.
- We define $[U]$ -connectedness as the transitive closure of the $[U]$ -adjacency relation.
- A $[U]$ -component of $H' \cup S_p$ is a maximally $[U]$ -connected subset $C \subseteq H' \cup S_p$.

Hence, if C_1, \dots, C_ℓ are the $[U]$ -components of $H' \cup S_p$, then $H' \cup S_p$ is partitioned into $C_0 \cup C_1 \cup \dots \cup C_\ell$, such that $C_0 = \{f \in H' \cup S_p \mid f \subseteq U\}$.

We next define balanced separators. While we give a definition w.r.t. sets of vertices $U \subseteq V(H)$, they can be alternatively defined in terms of sets of edges $S \subseteq E(H)$.

Definition 7 (Balanced Separators). Let $H' \cup S_p$ be an extended subhypergraph of a hypergraph H and let $U \subseteq V(H)$ be a set of vertices of H . The set U is a *balanced separator* of $H' \cup S_p$ if for each $[U]$ -component C_i of $H' \cup S_p$, $|C_i| \leq \frac{|H' \cup S_p|}{2}$ holds. In other words, no $[U]$ -component must contain more than half the edges of $H' \cup S_p$.

Finally, we extend GHDs to extended subhypergraphs.

Definition 8 (GHDs of Extended Subhypergraphs). Let H be a hypergraph and $H' \cup S_p$ an extended subhypergraph of H . A GHD of $H' \cup S_p$ is a tuple $\langle T, (B_u)_{u \in T}, (\lambda_u)_{u \in T} \rangle$, where $T = (N(T), E(T))$ is a tree, and B_u and λ_u are labeling functions, which map to each node $u \in T$ two sets, $B_u \subseteq V(H)$ and $\lambda_u \subseteq E(H) \cup S_p$. For a node u , we call B_u the *bag* and λ_u the *edge cover* of u . The set $B(\lambda_u)$ of vertices “covered” by λ_u is defined as $B(\lambda_u) = \{v \in V(H) \mid v \in f, f \in \lambda_u\}$. The functions λ_u and B_u have to satisfy the following conditions:

- (1) For each node $u \in T$, either
 - (a) $\lambda_u \subseteq E(H)$ and $B_u \subseteq B(\lambda_u)$, or
 - (b) $\lambda_u = \{s\}$ for some $s \in S_p$ and $B_u = s$.
- (2) If, for some $u \in T$, $\lambda_u = \{s\}$ for some $s \in S_p$, then u is a leaf node.
- (3) For each $e \in H' \cup S_p$, there is a node $u \in T$ s.t. $e \subseteq B_u$.
- (4) For each vertex $v \in V(H)$, $\{u \in T \mid v \in B_u\}$ is a connected subtree of T .

The *width* of a GHD is defined as $\max\{|\lambda_u| : u \in T\}$.

Clearly, also H itself is an extended subhypergraph of H with $H' = H$ and $S_p = \emptyset$. It is readily verified that the above definition of GHD of an extended subhypergraph $H' \cup S_p$ and the definition of GHD of a hypergraph H coincide for the special case of taking H as an extended subhypergraph of itself.

In Reference [27], a normal form of hypertree decompositions was introduced. We will carry the notion of normal form over to GHDs of extended subhypergraph. To this end, it is convenient to first define the set of edges *exclusively* covered by some subtree of a GHD:

Definition 9. Let $H' \cup S_p$ be an extended subhypergraph of some hypergraph H and $\mathcal{D} = \langle T, (B_u)_{u \in T}, (\lambda_u)_{u \in T} \rangle$ a GHD for $H' \cup S_p$. For a node $u \in T$, we write T_u to denote the subtree of T rooted at u . Moreover, we define the set of edges exclusively covered by T_u as $exCov(T_u) = \{f \in H' \cup S_p \mid \exists v \in T_u : f \subseteq B_v \wedge f \not\subseteq B_u\}$.

Our normal form of GHDs is then defined as follows:

Definition 10 (GHD of Extended Subhypergraphs Normal form). We say that a GHD $\langle T, (B_u)_{u \in T}, (\lambda_u)_{u \in T} \rangle$ of an extended subhypergraph $H' \cup S_p$ is in normal form, if for the root node r of T , the following property holds: let u_1, \dots, u_ℓ be the child nodes of r in T and let $T_{u_1}, \dots, T_{u_\ell}$ denote the subtrees in T rooted at u_1, \dots, u_ℓ , respectively. Then $exCov(T_{u_1}), \dots, exCov(T_{u_\ell})$ are precisely the $[V(\lambda_r)]$ -components of $H' \cup S_p$. Intuitively, each subtree T_i below the root “covers” the edges of *precisely one* $[B_r]$ -component of $H' \cup S_p$.

The following lemma is an immediate extensions of previous results for hypergraphs to extended subhypergraphs:

LEMMA 1. *Let $H' \cup S_p$ be an extended subhypergraph of some hypergraph H and suppose that there exists a GHD \mathcal{D} of width $\leq k$ for $H' \cup S_p$. Then there also exists a GHD \mathcal{D}' in normal form of width $\leq k$ for $H' \cup S_p$, such that B_r is a balanced separator of $H' \cup S_p$ for the root node r of \mathcal{D}' .*

PROOF. The lemma combines two results from References [27] and [3], respectively.

Our normal form relaxes the normal form of HDs introduced in Definition 5.1 in Reference [27]. The transformation into normal form can be taken over almost literally from the proof of Theorem 5.4 in Reference [27] for establishing the normal form of HDs.

The existence of a balanced separator as the root of a GHD is implicit in the definition of “hyperlinkedness” and Theorem 19 in Reference [3]. Again, it can be easily taken over to our case of an extended subhypergraph (i.e., to take also special edges into account). Actually, it can also be easily proved directly by starting off at the root r of an arbitrary GHD of $H' \cup S_p$ in normal form and checking if the components covered by the subtrees below all have size at most $\frac{|H' \cup S_p|}{2}$. If so, then we already have the desired form. If not, then there must be one subtree $T_{r'}$ rooted at a child r' of r , such that $\text{exCov}(T_{r'})$ is greater than $\frac{|H' \cup S_p|}{2}$. We then apply the normal form transformation also to $T_{r'}$ and check recursively if all the components covered by the subtrees below all have size at most $\frac{|H' \cup S_p|}{2}$. By repeating this recursive step, we will eventually reach a node u , such that B_u is a balanced separator. Then, we simply take this node as the root and again apply the normal form transformation of the proof of Theorem 5.4 in Reference [27] to this new root node and the subtrees immediately below it. \square

4.4.2 *Algorithm Description.* Here, we describe Algorithm 2, which we call `BalSep`. For a fixed integer $k \geq 1$, it takes as input a hypergraph H and computes a GHD of H of width $\leq k$ if it exists, or returns `NULL` otherwise. The main of Algorithm 2 consists of a call to the Function `Decompose` with parameters H and an empty set of special edges.

The recursive Function `Decompose` constitutes the core of the algorithm and, given as input a hypergraph H' and a set of special edges S_p , computes a GHD of $H' \cup S_p$ of width $\leq k$ if it exists. Lines 5–12 deal with the two base cases of the algorithm. If $H' \cup S_p$ has only one edge, then we create a decomposition made of a single node u whose label λ_u contains the only edge of $H' \cup S_p$ and the bag $B_u = V(H' \cup S_p)$, i.e., all the vertices of the extended subhypergraph, which are also the ones covered by λ_u . In a similar way, we deal with the case $|E(H' \cup S_p)| = 2$. We simply create two nodes u, v , one for each edge of $H' \cup S_p$, and then attach v as a child of u and return the decomposition.

If the extended subhypergraph has at least three edges, then we have to decompose it until we reach one of the two base cases. In line 13, we initialize the object `BalSepIt`, which is an iterator over the balanced separators of size $\leq k$ of $H' \cup S_p$ with edges in H . The iterator `BalSepIt` produces, one-by-one, all the ℓ -combinations of edges in H , for each $\ell \leq k$, to find a balanced separator for $H' \cup S_p$. Moreover, if all the combinations of full edges fail, then the function uses subedges of H to generate separators corresponding to elements of the set $f(H, k)$ of Equation (1).

In the while loop in lines 14–27, we recursively decompose $H' \cup S_p$. We are now creating the current node u of the GHD, and we have to compute λ_u, B_u and the children of u . In line 15, we set λ_u as the next balanced separator and, in line 16, we fix the bag $B_u = B(\lambda_u)$, as discussed above. We want to compute a GHD for each $[B_u]$ -component of $H' \cup S_p$ in such a way that it will be possible to attach each of them to the current node u without violating any condition of Definition 8.

Function `ComputeSubhypergraphs` computes the set of extended subhypergraphs corresponding to $[B_u]$ -components of $H' \cup S_p$ and introduces, in each of them, a new special edge B_u for connectedness. We assume here the existence of a function `ConnectedComponents`, which computes the connected components of a hypergraph in a standard way. First, Function

ComputeSubhypergraphs computes the hypergraph $H_u = (V_u, E_u)$ resulting from the removal of all vertices in B_u from $H' \cup S_p$. Then, in lines 5–11, it creates a new subhypergraph of $H' \cup S_p$ for each connected component of H_u . The new subhypergraph is stored in a variable c , which is a pair consisting of a hypergraph H and a set of special edges S_p . For a single subhypergraph, the new set of special edges $c.S_p$ is composed of the edges of S_p intersecting the current component $comp$ plus a new special edge $s = B_u$ corresponding to the separator B_u . We can then compute the hypergraph $c.H$. Its set of edges E contains the edges of H' intersecting the current component $comp$ and its set of vertices is the union of E and $c.S_p$. We then add c to the set of results res , which we finally return in line 12.

Back in Function Decompose, in lines 18–24, we recursively compute a GHD for each extended subhypergraph returned by Function ComputeSubhypergraphs. If the decomposition \mathcal{D} returned in line 19 is not NULL, then we add it to the set *subDecomps* of the children of the current node u ; otherwise, we set *subDecomps* to NULL and break the loop. At the end of the loop, we check whether *subDecomps* is NULL. If this is the case, it means that one of the recursive calls of Function Decompose was unsuccessful. We then have to continue the while loop of lines 14–27 and try the next balanced separator. In case all the recursive calls of Function Decompose were successful, Function BuildGHD builds the resulting GHD and returns it in line 27. If the algorithm exhausts all the choices of balanced separators for $H' \cup S_p$ and exits the while loop of lines 14–27, then it means that it is impossible to create a GHD of width $\leq k$ for $H' \cup S_p$ and we return NULL in line 28.

Finally, we describe the Function BuildGHD that, given a bag B_u , an edge cover λ_u , and a set *children* of GHDs, returns a GHD with root u and children *children*. We start off creating the node u with labels B_u and λ_u . Then, for each child $\mathcal{D} \in \textit{children}$, we find the node r in T (the tree of \mathcal{D}) having $B_r = B_u$ and reroot T to r . Now, for each child c_r of r , we attach c_r as a child of u . In other words, we attach every subtree rooted at a child node of r to u . Finally, we return the resulting decomposition.

4.4.3 Soundness and Completeness. Here, we prove that Algorithm 2 is sound and complete.

THEOREM 2. *Let H be a hypergraph and $k \geq 1$ an integer. Algorithm 2 called on H with parameter k returns a GHD of H of width $\leq k$ if and only if $ghw(H) \leq k$.*

We prove the soundness and the completeness of Algorithm 2 separately. Nevertheless, we want to point out that the main procedure of the algorithm consists of a call to Function Decompose with input H and an empty set of special edges. Thus, in the next proofs, we will actually prove that Function Decompose called on (H', S_p) with parameter k returns a GHD of $H' \cup S_p$ of width $\leq k$ if and only if $ghw(H' \cup S_p) \leq k$ w.r.t. Definition 8. Note that in case of a hypergraph H and a set of special edges $S_p = \emptyset$, Definition 8 coincides with the usual definition of GHD.

PROOF (SOUNDNESS). We show that if Function Decompose called on (H', S_p) with parameter k returns a GHD of $H' \cup S_p$ of width $\leq k$, then such a decomposition actually exists and $ghw(H' \cup S_p) \leq k$. We proceed by induction over the size of $H' \cup S_p$, i.e., $|E(H' \cup S_p)|$. For the base case, we assume $|H' \cup S_p| \leq 2$. In case $|H' \cup S_p| = 1$, we return a GHD made of a single node whose λ -label consists of the only edge $H' \cup S_p$, which also cover all of the vertices of the hypergraph. Such a decomposition has width 1 and clearly satisfies all the conditions of Definition 8. In case $|H' \cup S_p| = 2$, we create two nodes u, v , each one corresponding to an edge of $H' \cup S_p$ and we attach v as a child of u . Note that both u and v are leaf nodes. It is easy to verify that also, in this case, we return a valid GHD of $H' \cup S_p$.

For the induction step, suppose that the recursive function Decompose correctly returns a GHD of $H' \cup S_p$ of width $\leq k$ for each $H' \cup S_p$ such that $|H' \cup S_p| \leq j$, for some $j \geq 2$. Now suppose that

ALGORITHM 2: BalSep

Input: A hypergraph H .
Parameter: An integer $k \geq 1$.
Output: A GHD of H of width $\leq k$ if it exists, NULL otherwise.

```

1 Main
2   Make  $H$  globally visible
3   return Decompose( $H, \emptyset$ )
4 FUNCTION Decompose( $H'$ : hypergraph,  $S_p$ : set of special edges)
5   if  $|E(H' \cup S_p)| == 1$  then
6     return node  $u$  with  $B_u \leftarrow V(H' \cup S_p)$  and  $\lambda_u \leftarrow E(H' \cup S_p)$ 
7   if  $|E(H' \cup S_p)| == 2$  then
8     Let  $e_1, e_2$  be the two edges of  $H' \cup S_p$ 
9     Create node  $u$  with  $B_u \leftarrow e_1$  and  $\lambda_u \leftarrow \{e_1\}$ 
10    Create node  $v$  with  $B_v \leftarrow e_2$  and  $\lambda_v \leftarrow \{e_2\}$ 
11    AttachChild( $u, v$ )
12    return  $u$ ;
13  $BalSepIt \leftarrow$  InitBalSepIterator( $H, H', S_p, k$ )
14 while HasNext( $BalSepIt$ ) do
15    $\lambda_u \leftarrow$  Next( $BalSepIt$ )
16    $B_u \leftarrow B(\lambda_u)$ 
17    $subDecomps \leftarrow \{\}$ 
18   foreach  $c \in$  ComputeSubhypergraphs( $H', S_p, B_u$ ) do
19      $\mathcal{D} \leftarrow$  Decompose( $c.H, c.S_p$ )
20     if  $\mathcal{D} \neq$  NULL then
21        $subDecomps \leftarrow subDecomps \cup \{\mathcal{D}\}$ 
22     else
23        $subDecomps \leftarrow$  NULL
24     break
25   if  $subDecomps ==$  NULL then
26     continue
27   return BuildGHD( $B_u, \lambda_u, subDecomps$ )
28 return NULL

```

Function BuildGHD($B_u, \lambda_u, children$)

Input: A set of vertices B_u , a set of λ_u , a set of GHDs $children$.
Output: A GHD with root u and children $children$.

```

1 begin
2   Create node  $u$  with  $B_u$  and  $\lambda_u$ 
3   foreach  $\mathcal{D} \in children$  do
4     Let  $T$  be the tree structure of  $\mathcal{D}$ 
5      $\hat{r} \leftarrow$  Reroot( $T, B_u$ )
6     foreach  $c_{\hat{r}} \in$  Children( $\hat{r}$ ) do
7       AttachChild( $u, c_{\hat{r}}$ )
8   return  $u$ 

```

Function ComputeSubhypergraphs(H', S_p, B_u)**Input:** A hypergraph H' , a set of special edges S_p , a set of vertices B_u .**Output:** The set of subhypergraphs of $H' \cup S_p$ w.r.t. B_u .

```

1 begin
2    $V_u \leftarrow V(H') \setminus B_u$ 
3    $E_u \leftarrow \{e \cap V_u \mid e \in E(H' \cup S_p)\}$ 
4    $res \leftarrow \{\}$ 
5   foreach  $comp \in \text{ConnectedComponents}(V_u, E_u)$  do
6      $c \leftarrow \text{initialize pair } (H = \text{NULL}, S_p = \text{NULL})$ 
7      $c.S_p \leftarrow \{s \in S_p \mid s \cap comp \neq \emptyset\} \cup \{B_u\}$ 
8      $E \leftarrow \{e \mid e \in E(H') \wedge e \cap comp \neq \emptyset\}$ 
9      $V \leftarrow V(E) \cup V(c.S_p)$ 
10     $c.H \leftarrow (V, E)$ 
11     $res \leftarrow res \cup \{c\}$ 
12  return  $res$ 

```

$|H' \cup S_p| = j + 1 \geq 3$ and that Function Decompose(H', S_p) returns a GHD of $H' \cup S_p$ of width $\leq k$. We have to show that then there indeed exists such a GHD.

Algorithm 2 only returns a GHD in line 27. The program successfully reaches this line only if the following happens:

- In line 15, a balanced separator λ_u of $H' \cup S_p$ is chosen.
- In line 18, the extended subhypergraphs of $H' \cup S_p$ w.r.t. $B_u = B(\lambda_u)$ are computed; in particular, each extended subhypergraph corresponds to a $[B_u]$ -component of $H' \cup S_p$ plus a new special edge $s = B_u$.
- For each extended subhypergraph $c.H \cup c.S_p$, the call to Function Decompose in line 19 is successful, i.e., it returns a GHD of $c.H \cup c.S_p$ of width $\leq k$.

We are assuming $|H' \cup S_p| \geq 3$ and λ_u is a balanced separator of $H' \cup S_p$. Let C_1, \dots, C_ℓ be the ℓ $[B_u]$ -components of $H' \cup S_p$. All extended subhypergraphs $C_i \cup \{B_u\}$ are strictly smaller than j . Hence, by the induction hypothesis, for each $i \in \{1, \dots, \ell\}$, there indeed exists a GHD $\mathcal{D}_i = \langle T_i, (B_{i,u})_{u \in T}, (\lambda_{i,u})_{u \in T} \rangle$ of width $\leq k$ for the extended subhypergraph $C_i \cup \{B_u\}$.

It is left to show that Function BuildGHD correctly constructs a GHD of width $\leq k$ of $H' \cup S_p$. For each \mathcal{D}_i , let \hat{r}_i be the node of T_i with $B_{i,\hat{r}_i} = B_u$ and $\lambda_{i,\hat{r}_i} = \{B_u\}$. By construction, the node \hat{r}_i exists in every \mathcal{D}_i ; it is always a leaf and it has the same bag and λ -label everywhere. Let $T_i = (N(T_i), E(T_i))$ be the tree structure of \mathcal{D}_i and, w.l.o.g., assume that the node sets $N(T_i)$ are pairwise disjoint. We define the tree structure $T = (N(T), E(T))$ and the functions B_u and λ_u of \mathcal{D} as follows:

- $N = (N(T_1) \setminus \{\hat{r}_1\}) \cup \dots \cup (N(T_\ell) \setminus \{\hat{r}_\ell\}) \cup \{r\}$, where r is a new (root) node.
- For the definition of $E(T)$ recall that each \hat{r}_i is a leaf node in its decomposition. Let e_i denote the edge between \hat{r}_i and its parent. Then, we define $E(T)$ as $E(T) = (E(T_1) \setminus \{e_1\}) \cup \dots \cup (E(T_\ell) \setminus \{e_\ell\}) \cup R$ with $R = \{[r, \hat{r}_1], \dots, [r, \hat{r}_\ell]\}$.
- $\lambda_r = \lambda_u$ and $B_r = B_u$.
- For every $v \in N \setminus \{r\}$, there exists exactly one i , such that $v \in N(T_i)$. We set $\lambda_v = \lambda_{i,v}$ and $B_v = B_{i,v}$.

Intuitively, the GHD \mathcal{D} is obtained by taking in each GHD \mathcal{D}_i the node \hat{r}_i as the root node and combining all GHDs \mathcal{D}_i to a single GHD by merging their root nodes to a single node r . This is possible, since all nodes \hat{r}_i have the same λ -labels and bags. It is easy to verify that the resulting GHD is indeed a GHD of width $\leq k$ of the extended subhypergraph $H' \cup S_p$. \square

PROOF (COMPLETENESS). For any hypergraph H' and set of special edges S_p , we prove that if $ghw(H' \cup S_p) \leq k$, then Function Decompose on input (H', S_p) returns a GHD of width $\leq k$ of $H' \cup S_p$. Again, we proceed by induction on $|H' \cup S_p|$.

The base case of $|H' \cup S_p| \leq 2$ is dealt with in lines 5–12 of Algorithm 2. In this case, we simply construct a GHD of width 1 for $|H' \cup S_p|$ and return it.

For the induction step, suppose $|H' \cup S_p| \leq j$, for some $j \geq 2$, and $ghw(H' \cup S_p) \leq k$. Then, Function Decompose on input (H', S_p) returns a GHD of width $\leq k$ of $H' \cup S_p$. Now assume that $|H' \cup S_p| = j + 1 \geq 3$ and that $ghw(H' \cup S_p) \leq k$. We have to show that Function Decompose on input (H', S_p) returns a GHD of width $\leq k$ of $H' \cup S_p$.

By Lemma 1, we may assume w.l.o.g., that GHD $\mathcal{D} = \langle T, (B_u)_{u \in T}, (\lambda_u)_{u \in T} \rangle$ is in normal form and that B_r is a balanced separator of $H' \cup S_p$ for the root node r of \mathcal{D} . Let $S \subseteq E(H)$ denote the λ -label of r , i.e., $\lambda_r = S$.

When Function Decompose is called on input (H', S_p) , the while loop in lines 14–27 eventually generates all possible balanced separators of size $\leq k$ of $H' \cup S_p$, unless it returns on line 27 before the end of the loop. Remember that the object *BalSepIt* not only generates edge separators with edges in $E(H)$, i.e., the original hypergraph on which Algorithm 2 is called, but it also uses edges in $f(H, k)$, i.e., subedges of edges in $E(H)$. Thus, at some point, in line 15, we will choose the separator $\lambda_u = S$, equivalently, $B_u = B_r$.

Let C_1, \dots, C_ℓ denote the $[B_u]$ -components of $H' \cup S$. Since \mathcal{D} is in normal form, we know that the root node r has ℓ child nodes such that $C_i = \text{exCov}(T_i)$, where T_i is the subtree in T rooted at n_i for $i \in \{1, \dots, \ell\}$. Recall from Definition 9 that we write $\text{exCov}(T_i)$ to denote the set of edges exclusively covered by T_i .

Now consider the extended subhypergraph $C_i \cup \{B_u\}$ for arbitrary $i \in \{1, \dots, \ell\}$. Since B_r is a balanced separator of $H' \cup S_p$, we have $|C_i \cup \{B_u\}| \leq j$. Moreover, there exists a GHD \mathcal{D}_i of $C_i \cup \{B_u\}$, namely, the subtree of \mathcal{D} induced by the nodes in T_i plus r . Hence, by the induction hypothesis, calling Function Decompose with the input corresponding to the extended subhypergraph $C_i \cup \{B_u\}$ returns a valid GHD. This means that, in our call of Function Decompose with input (H', S_p) , we have the following behavior:

- On line 19, Function Decompose is called recursively for all extended subhypergraphs $C_i \cup \{B_u\}$.
- Each call of the Function Decompose returns a GHD for the respective extended subhypergraph.
- The results of these recursive calls are collected in line 21 in the variable *subDecomps*.
- Hence, after exiting the loop in lines 18–26, the return statement in line 27 is executed.

The call to Function BuildGHD correctly produces the desired GHD, as discussed in the soundness proof. Finally, Function Decompose indeed returns a GHD of width $\leq k$ of $H' \cup S_p$. \square

5 HYPERBENCH BENCHMARK AND TOOL

In this section, we describe *HyperBench*—our new benchmark and web tool. We first introduce our system and test environment used for the experiments, then we present a new method to extract simple conjunctive queries from complex SQL queries. Finally, we describe the CQs and CSPs we have collected.

5.1 System and Test Environment

Our system is composed of two libraries: a C++ library with implementations of the algorithms and a Java library for processing of SQL queries and hypergraphs. We have extended the algorithm DetKDecomp from Reference [30] and based our program on a new implementation, which we call NewDetKDecomp. The resulting C++ library comprising around 8,500 lines of code has been improved in many ways. While the underlying *hw* algorithm is still the one from Reference [30], the implementation makes use of modern C++ constructs such as smart pointers for better memory management and maintainability. The code itself has also been optimized in several parts, thus improving overall performance w.r.t. the previous release. The software now also goes beyond the computation of HDs and it has been expanded to compute GHDs through the algorithms presented in Section 4. Moreover, NewDetKDecomp can also compute an approximated form of FHDs, which will be described in Section 6.5. We designed our Java library *hg-tools* for preprocessing of SQL queries and collecting hypergraph statistics. It uses the open source libraries *JSqlParser* [57] for SQL processing and *JGraphT* [44] to deal with graph data structures. Our two libraries are available at <https://github.com/dmlongo/newdetkdecomp> and <https://github.com/dmlongo/hgtools>.

All the experiments reported in this article were performed on a cluster of 10 workstations each running Ubuntu 16.04. Every workstation has the same specification and is equipped with two Intel Xeon E5-2650 (v4) processors each having 12 cores and 256 GB main memory. Since all algorithms are single-threaded, we could run several experiments in parallel. For all upcoming runs of our algorithms, we set a timeout of 3,600 s.

5.2 Translation of Complex SQL Queries into Hypergraphs

Since the applicability of structural decomposition methods is limited to conjunctive queries, we have devised a strategy to transform also more complex queries into (a collection of) hypergraphs. Our program *hg-tools* takes as input a complex SQL query and produces a collection of simpler SQL queries. When facing a complex query, e.g., one containing nested queries, the idea is to identify simple queries, extract them, and transform them separately.

A conjunctive query can be expressed in SQL as a SELECT-FROM-WHERE statement in which the WHERE clause is only allowed to be a conjunction of equality conditions. Moreover, such a query must not contain negation, disjunction, and subqueries, i.e., nested SELECT statements. For our purposes, we neglect the SELECT clause, because in our experiments, we focus on computing decompositions instead of answering the queries. Hence, only the hypergraph structure determined by the FROM and WHERE clauses is important.

Query 1 is an example of an SQL query. The FROM clause uses two instances of the relation $tab(a, b, c)$ and in the WHERE clause there are some conditions on the two relation instances. The condition on line 3 is a join and falls in the scope of conjunctive queries. Though the condition on line 4 is not conjunctive, it is just a comparison with a constant value and it does not influence the query structure. Nevertheless, the condition on line 5 involves a negation, thus Query 1 is not a conjunctive query. This kind of query is not allowed by our framework.

```

1  SELECT *
2  FROM tab t1 , tab t2
3  WHERE t1.a = t2.a
4     AND t1.b > 5
5     AND t1.c <> t2.c ;

```

Listing 1. A simple SQL query.

Nevertheless, what really matters in defining the query structure are the relationships between the different variables involved in the query, in particular, join conditions. For this reason, when faced with a query that is not conjunctive, we consider a simplified version that contains the conjunctive core of the original query. In case of Query 1, we drop lines 4–5 and perform our experiments on the rest of the query.

Real-world SQL queries can be rather complicated and, in particular, they can contain nested SELECT statements, which we simply call *subqueries*. The presence of subqueries automatically makes the query non conjunctive, but we have decided to extract the single queries and analyze their conjunctive cores separately. Query 2 is a query with nested SELECT statements. The subquery on lines 4–6 can be examined separately. The subquery on lines 7–9 cannot be examined separately, because it contains a reference to a table defined in an external query and, more precisely, it requires the evaluation of the subquery for the value of the current row of that table.

```

1  SELECT *
2  FROM tab t1, tab t2
3  WHERE t1.a = t2.a
4     AND t1.b IN (SELECT tab.b
5                   FROM tab
6                   WHERE tab.c == 'ok')
7  AND EXISTS (SELECT *
8              FROM differentTable dt
9              WHERE dt.a = t1.a);

```

Listing 2. A complex SQL query.

In SQL, a subquery can appear in different places in a query. Depending in which statement or condition it appears, we treat it differently:

- If a query is of the form $q_1 \circ \dots \circ q_n$, where each q_i is a query and $\circ \in \{\cup, \cap, \setminus\}$, then we extract the single queries q_i and process them separately.
- If a subquery appears in the FROM clause, then we convert it into a view (see Section 5.4).
- If a subquery contains a reference to an external query, as in lines 7–9 of Query 2, then it must be discarded.

5.3 Extracting Simple Queries

In general, an SQL query can contain subqueries. We want to extract them and either integrate them into the *main* query, when possible, or analyze them separately. To do that in such a way that the end result closely resembles the original query as much as possible, we build a graph representing the dependencies between the subqueries. At the end of the process, we extract queries that are independent and eliminate those that are mutually dependent.

We say that a subquery s_1 *depends* on a subquery s_2 if the result of q_1 can be computed only after computing the result of q_2 . The dependency graph of a query Q is a graph $G = (S, D)$ where the set S contains the nodes corresponding to subqueries of Q and $(s_1, s_2) \in D$ is an arrow, for $s_1, s_2 \in S$, if s_1 depends on s_2 . Given a query Q , we create its dependency graph G as follows:

- (1) Create a node $q \in S$ representing the outer query.
- (2) For each nested query s_i of q , create a new node $s_i \in S$ and an edge $(q, s_i) \in D$.
- (3) If s_i contains a reference to a table defined in any ancestor s_j , then create an edge (s_i, s_j) .
- (4) Recursively examine s_i for nested queries.

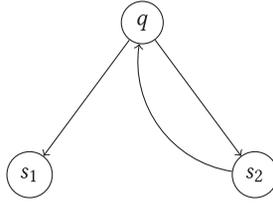


Fig. 5. Dependency graph of Query 2.

Once we have built the graph, we identify the nodes that are involved in cycles and eliminate them. In particular, we consider q as a root and navigate the graph. Whenever we find a node having an edge pointing to an ancestor, we eliminate it together with all of its incoming and outgoing edges. Eventually, we end up with a forest in which we extract a query from each node.

The dependency graph G of Query 2 is shown in Figure 5. The node q corresponds to the outer query and it is the root of the graph. The nodes s_1 and s_2 represent the subqueries in lines 4–6 and lines 7–9 of Query 2, respectively. As the result of Query 2 can be computed only after computing the subqueries, the two edges (q, s_1) and (q, s_2) are present in the graph. Since s_2 refers to the table t_1 defined in q (line 9), G contains also the edge (s_2, q) . After the creation, we look for cycles in the graph. In this case, we see that there is no way to evaluate s_2 independently from q . Then, we remove s_2 and all of its incident edges. Finally, we extract a simple query from each node of the remaining graph.

5.4 Converting Simple Queries into Hypergraphs

Once we have extracted and simplified subqueries, we are left with simple SQL queries of the type

$$\text{SELECT } r_{i_1}.A_{j_1}, \dots, r_{i_z}.A_{j_z} \text{ FROM } r_1, \dots, r_m \text{ WHERE } \textit{cond} \quad (3)$$

such that \textit{cond} is a conjunction of conditions of the form $r_i.A = r_j.B$ or $r_i.A = c$, where c is a constant. Such queries are equivalent to conjunctive queries, thus it is easy to draw a connection to a CQ and transform it into a hypergraph. Nevertheless, in our case it makes more sense to go directly from an SQL query to the hypergraph.

Let Q be an SQL query of the form ((3)), then the hypergraph $H_Q = (V(H_Q), E(H_Q))$ corresponding to Q is obtained as follows: We first build the hypergraph induced by the FROM clause. Consider a relation $r(A_1, \dots, A_\ell)$ in the FROM clause. For each attribute A_i of r , we create a vertex $v_{A_i} \in V(H_Q)$. Then, we create the edge $r = \{v_{A_1}, \dots, v_{A_\ell}\} \in E(H_Q)$. Now, we modify the hypergraph according to the conditions in the WHERE clause. Let \textit{cond} be such a condition. It can be of two forms:

- If \textit{cond} is of the form $r_i.A = r_j.B$, then we merge vertices v_A and v_B and modify their incident edges. W.l.o.g. assume v_A itself becomes the merged vertex. For each edge $r \in \{e \in E(H_Q) \mid v_B \in e\}$, we remove r from $E(H_Q)$ and add a new edge $r' = (r \setminus \{v_B\}) \cup \{v_A\}$.
- If \textit{cond} is of the form $r_i.A = c$, with c constant, then we remove v_A from $V(H_Q)$ and, for each edge $r \in \{e \in E(H_Q) \mid v_A \in e\}$, we remove r from $E(H_Q)$ and add a new edge $r' = r \setminus \{v_A\}$.

At the end of this procedure, we eliminate empty edges and multiple edges. Also, in our setting SELECT clauses do not contribute to the query structure, thus we simply ignore them.

We also take into account logical views. SQL views are virtual tables that are recreated every time the view is called; thus, we have decided to expand the main query by adding the view inside it. Consider Query 3. The view *crossView* has a cross structure composed by the two relations t_1, t_2 intersecting in the node b in the middle (see Figure 6(a)). The main query adds two relations t_1, t_2

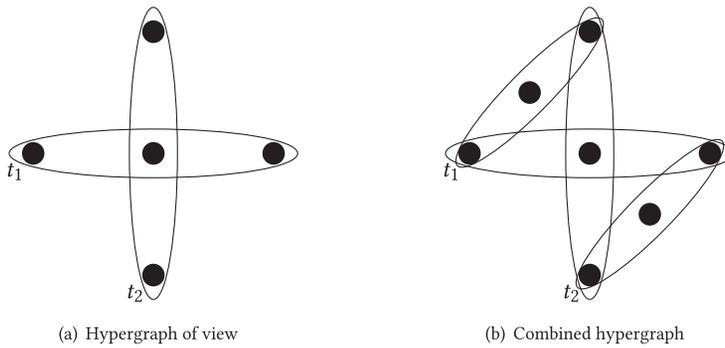


Fig. 6. Hypergraphs for Query 3.

that are distinct from the ones in the view *crossView* and intersect it in four points, thus creating two cycles. The end result is depicted in Figure 6(b). The resulting query can finally be converted into a hypergraph with the algorithm described above.

```

1  WITH crossView AS (
2      SELECT t1.a a1, t1.c c1, t2.a a2, t2.c c2
3      FROM tab t1, tab t2
4      WHERE t1.b = t2.b
5  )
6  SELECT *
7  FROM tab t1, tab t2, crossView cr
8  WHERE t1.a = cr.a1
9      AND t1.c = cr.a2
10     AND t2.a = cr.c1
11     AND t2.c = cr.c2;

```

Listing 3. An SQL query with a view.

5.5 From CSPs to Hypergraphs

An important part of our benchmark consists of instances of Constraint Satisfaction Problems. The set we have collected presents different characteristics w.r.t. the ones found in CQs, thus their analysis offers a more varied picture of the hypergraphs encountered in applications. For this reason, we have retrieved CSP instances that have a significant practical aspect.

The source of most of our CSPs is the website XCSP [7]. XCSP3 is an XML-based format used to represent constraint satisfaction problems. The language offers a wide variety of options to represent the most common constraints and frameworks, making it a solid intermediate format between different solvers. They also organize solver competitions for which they use their instances as a benchmark.

From XCSP, we have selected a total of 1,953 instances with less than 100 extensional constraints such that all constraints are extensional. The choice on the number of constraints allows us to have an adequate number of instances of increasing difficulty. Moreover, algorithms that use GHDs to solve CSPs typically need the constraints in a relational form. Therefore, we picked only CSPs such that all constraints are extensional, i.e., the constraints are provided in the form of a relation.

These instances are divided into CSPs from concrete applications, called *CSP Application* in the sequel (1,090 instances), and randomly generated CSPs, called *CSP Random* below (863 instances).

The instances we have fetched from the website are written in well-structured XML files in which variables and constraints are explicitly defined through the use of specific XML tags. The transformation of these instances into hypergraphs did not require a specific methodology, since the authors of the XCSP3 format provide an extensive library for parsing the instances where most of the process is already automatized. Obviously, we still had to convert the object in memory into a hypergraph. To this end, we have reimplemented the behavior of some callback methods in such a way that, whenever the program reads a variable, it adds a vertex to the hypergraph, and, whenever it reads a constraint, it adds an edge containing the vertices corresponding to the variables affected by the constraint.

Our collection of CSPs also includes a third class, which we call *CSP Other*. These instances have been used in previous hypertree width analyses available at <https://www.dbai.tuwien.ac.at/proj/hypertree/>. This set contains interesting examples coming from industry and a variety of different test examples [20]. In particular, a part of the hypergraphs is obtained from Daimler Chrysler and represents circuits and systems. A second part is a hypergraph translation of the circuits belonging to the well-known benchmark library of the **IEEE International Symposium on Circuits and Systems (ISCAS)**. Finally, some hypergraphs correspond to grids extracted from pebbling problems. Since the instances are provided already as hypergraphs, no additional processing was necessary to incorporate them.

5.6 Hypergraph Benchmark

Our benchmark contains 3,648 hypergraphs, which have been converted from CQs and CSPs collected from various sources. Out of these 3,648 hypergraphs, 3,142 hypergraphs have never been used in a hypertree width analysis before. The hypertree width of 424 CQs and of 82 CSPs has been analyzed in References [30], [11], and/or [12, 13]. In particular, the hypergraphs of the SPARQL queries from References [12, 13] have been kindly shared with us by the authors, and their hw has already been analyzed in those papers. An overview of all instances of CQs and CSPs is given in Table 1. They have been collected from various publicly available benchmarks and repositories of CQs and CSPs. In the first column, the names of each collection of CQs and CSPs are given together with references where they were first published. In the second column, we display the number of hypergraphs extracted from each collection. The hw of the CQs and CSPs in our benchmark will be discussed in detail in Section 6.1. To get a first feeling of the hw of the various sources, we mention the number of cyclic hypergraphs (i.e., those with $hw \geq 2$) in the last column. When gathering the CQs, we proceeded as follows: Of the huge benchmark reported in Reference [12], we have only included CQs, which were detected as having $hw \geq 2$ in Reference [12]. Of the other huge repository reported in Reference [13], we included the hypergraphs corresponding to the 273,974 unique SPARQL queries with $hw \geq 2$. Even though the queries are unique, most of them share the same hypergraph structure. Thus, after removing duplicates on the hypergraph level, we ended up with 354 unique hypergraphs with $hw \geq 2$. Of the big repository reported in Reference [34], we have included those CQs, which are not trivially acyclic (i.e., they have at least three atoms). Of all the small collections of queries, we have included all. It follows a detailed description of the different benchmarks.

Our benchmark contains 1,113 CQs from five main sources [9, 10, 12, 13, 34] and a set of 500 randomly generated queries using the query generator of Reference [49]. In the sequel, we shall refer to the former queries as *CQ Application* and to the latter as *CQ Random*. The CQs analyzed in Reference [12] constitute a big repository of CQs—namely, 26,157,880 CQs stemming from SPARQL

Table 1. Overview of Benchmark Instances

	Benchmark	No. instances	$hw \geq 2$
CQs	SPARQL [12]	70 (out of 26,157,880)	70
	WIKIDATA [13]	354 (out of 273,947)	354
	LUBM [10, 32]	14	2
	iBENCH [6, 10]	40	0
	DOCTORS [10, 21]	14	0
	DEEP [10]	41	0
	JOB (IMDB) [41]	33	7
	TPC-H [9, 54]	29	1
	TPC-DS [53]	228	5
	SQLSHARE [34]	290 (out of 15,170)	1
	RANDOM [49]	500	464
CSPs	APPLICATION [7]	1,090	1,090
	RANDOM [7]	863	863
	OTHER [11, 30]	82	82
	<i>Total:</i>	<i>3,648</i>	<i>2,939</i>

queries. The queries come from real users of SPARQL endpoints and their hypertree width was already determined in Reference [12]. Almost all of these CQs were shown to be acyclic. Our analysis comprises 70 CQs from Reference [12], which (apart from few exceptions) are essentially the ones in Reference [12] with $hw \geq 2$. In particular, we have analyzed all eight CQs with highest hw among the CQs analyzed in Reference [12] (namely, $hw = 3$). Bonifati et al. carried on this line of work and examined a bigger repository of SPARQL queries coming from Wikidata in Reference [13]. This repository of 208,215,209 SPARQL queries was originally released by Malyshev et al. with the study in Reference [42]. Bonifati et al. kindly sent us the unique 273,947 SPARQL queries with $hw \geq 2$ examined in Reference [42]. We extracted 354 different hypergraphs and all of them have $hw = 2$.

The LUBM [32], iBENCH [6], DOCTORS [21], and DEEP scenarios have been recently used to evaluate the performance of chase-based systems [10]. Their queries were especially tailored towards the evaluation of query answering tasks of such systems. Note that the LUBM benchmark [32] is a widely used standard benchmark for the evaluation of Semantic Web repositories. Its queries are designed to measure the performance of those repositories over large datasets. Strictly speaking, the iBENCH is a tool for generating schemas, constraints, and mappings for data integration tasks. However, in Reference [10], 40 queries were created for tests with the iBENCH. We therefore refer to these queries as iBENCH-CQs here. In summary, we have incorporated all queries that were either contained in the original benchmarks or created/adapted for the tests in Reference [10].

The goal of the **Join Order Benchmark (JOB)** [41] was to evaluate the impact of a good join order on the performance of query evaluation in standard RDBMSs. Those queries were formulated over the real-world dataset **Internet Movie Database (IMDB)**. All of the queries have between 3 and 16 joins. Clearly, as the goal was to measure the impact of a good join order, those 33 queries are of higher complexity, hence 7 out of the 33 queries have $hw \geq 2$.

The **Transaction Processing Performance Council (TPC)** is a well-known non-profit organization that develops benchmarks for the evaluation of DBMSs. Given their broad industry-wide relevance and since they reflect common workloads in decision support systems, we included the

TPC-H [54] and the TPC-DS [53] benchmarks. In Reference [18], we analyzed the TPC-H queries from the GitHub repository originally provided by Michael Benedikt and Efthymia Tsamoura [9] for the work on Reference [10]. Nevertheless, for this article, we downloaded the original dataset from Reference [54] and extracted the queries according to the methodology introduced in Sections 5.2–5.4. From the original set of 22 complex queries, we extracted 29 simple queries. The TPC-DS benchmark is more complex than TPC-H and it contains more queries. Indeed, from the original set of 113 complex queries, we extracted 228 simple queries.

From SQLShare [34], a multi-year SQL-as-a-service experiment with a large set of real-world queries, we extracted 15,170 queries by considering all queries in the log files. These queries are divided into two sets: materialized views and usual queries, which could possibly make use of the materialized views. Also, the whole dataset gathers data from different databases and the link between queries and databases is not explicitly defined. To execute the experiments, we had to clean the queries from trivial errors impeding the parsing, link the queries to the right database schema, incorporate the materialized views, and resolve ambiguities in the query semantics. After removing queries with complex syntactical errors, we obtained 12,483 queries. As a next step, we used the algorithm from Section 5.3 to obtain a set of CQs and, after removing duplicates, we got a collection of 6,086 simple SQL queries. From this set, we eliminated 5,796 queries with ≤ 2 atoms (whose acyclicity is immediate) and ended up with 290 queries.

The random queries were generated with a tool that stems from the work on query answering using views in Reference [49]. The query generator allows three options: chain/star/random queries. Since the former two types are trivially acyclic, we only used the third option. Here, it is possible to supply several parameters for the size of the generated queries. In terms of the resulting hypergraphs, one can thus fix the number of vertices, number of edges, and arity. We have generated 500 CQs with 5–100 vertices, 3–50 edges, and arities from 3 to 20. These values correspond to the values observed for the *CQ Application* hypergraphs. However, even though these size values have been chosen similarly, the structural properties of the hypergraphs in the two groups *CQ Application* and *CQ Random* differ significantly, as will become clear from our analysis in Section 6.1.

As was detailed in Section 5.5, our benchmark currently contains 2,035 hypergraphs from CSP instances, out of which 1,953 instances were obtained from xcs.org (see also Reference [7]). These instances, in turn, are divided into CSPs from concrete applications, called *CSP Application* in the sequel (1,090 instances), and randomly generated CSPs, called *CSP Random* below (863 instances). In addition, we have included 82 CSP instances, which were already used in previous *hw* experiments [20, 30]. These instances, which we refer to as *CSP Other*, are provided at <https://www.dbai.tuwien.ac.at/proj/hypertree/>.

Our HyperBench benchmark consists of the CQ and CSP instances converted to hypergraphs. In Figure 7, we show the number of vertices, the number of edges, and the arity (i.e., the maximum size of the edges) as three important metrics of the size of each hypergraph. The smallest are those coming from *CQ Application* (most of them have up to 10 edges), while the hypergraphs coming from CSPs can be significantly larger (up to 2,993 edges). Although some hypergraphs are very big, more than 50% of all hypergraphs have maximum arity less than 5. In Figure 7, we can easily compare the different types of hypergraphs, e.g., hypergraphs of arity greater than 20 only exist in the application classes; the *CSP Other* class contains the highest portion of hypergraphs with a big number of vertices and edges, and so on.

The hypergraphs and the results of our analysis can be accessed through our web tool, available at <http://hyperbench.dbai.tuwien.ac.at>.

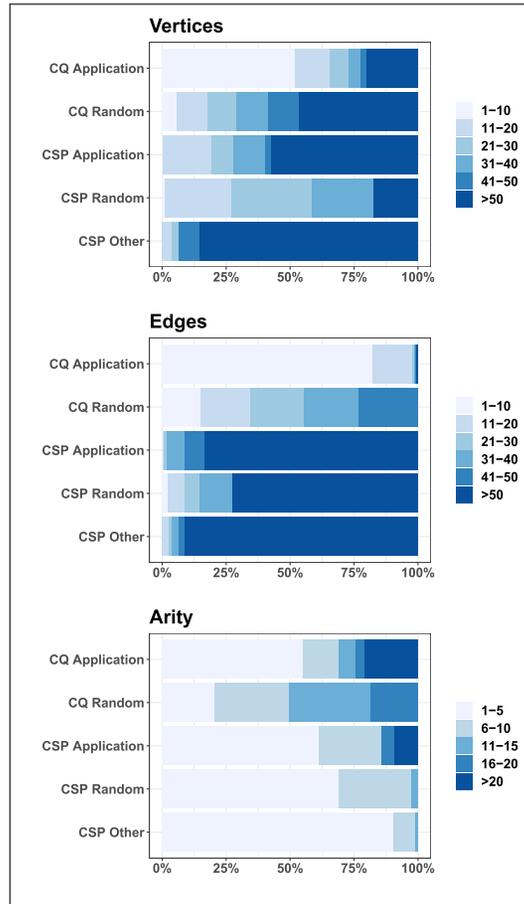


Fig. 7. Hypergraph sizes.

6 EXPERIMENTS

In this section, we present the empirical results obtained with the HyperBench benchmark. On the one hand, we want to get an overview of the hypertree width of the various types of hypergraphs in our benchmark (cf. Subgoal 1 in Section 1). On the other hand, we want to find out how realistic the restriction to low values for certain hypergraph invariants is (cf. Subgoal 2 stated in Section 1). After this first analysis of the structural properties of the hypergraphs we collected, we perform an evaluation of the different *ghw* algorithms presented in Section 4. Finally, we propose and evaluate two algorithms for computing approximated FHDs.

6.1 Hypergraph Properties

In References [19, 25], several invariants of hypergraphs were used to make the $\text{CHECK}(\text{GHD}, k)$ and $\text{CHECK}(\text{FHD}, k)$ problems tractable or, at least, easier to approximate. We thus investigate the following properties (cf. Definitions 3–5):

- *Deg*: the degree of the underlying hypergraph
- *BIP*: the intersection size
- *c-BMIP*: the *c*-multi-intersection size for $c \in \{3, 4\}$
- *VC-dim*: the VC-dimension

Table 2. Properties of All Benchmark Instances

<i>CQ Application</i>						<i>CQ Random</i>					
<i>i</i>	Deg	BIP	3-BMIP	4-BMIP	VC-dim	<i>i</i>	Deg	BIP	3-BMIP	4-BMIP	VC-dim
0	0	74	394	647	72	0	0	1	16	49	0
1	74	721	673	456	484	1	1	17	77	125	20
2	320	286	45	10	557	2	15	53	90	120	133
3	253	17	1	0	0	3	38	62	103	74	240
4	181	6	0	0	0	4	31	63	62	42	106
5	73	9	0	0	0	5	33	71	47	28	1
>5	212	0	0	0	0	>5	382	233	105	62	0

<i>CSP Application</i>						<i>CSP Random</i>					
<i>i</i>	Deg	BIP	3-BMIP	4-BMIP	VC-dim	<i>i</i>	Deg	BIP	3-BMIP	4-BMIP	VC-dim
0	0	0	596	597	0	0	0	0	0	0	0
1	0	1030	459	486	0	1	0	200	200	238	0
2	596	59	34	7	1064	2	0	224	312	407	220
3	1	0	1	0	26	3	0	76	147	95	515
4	1	0	0	0	0	4	12	181	161	97	57
5	2	0	0	0	0	5	8	99	14	1	71
>5	490	1	0	0	0	>5	843	83	29	25	0

<i>CSP Other</i>					
<i>i</i>	Deg	BIP	3-BMIP	4-BMIP	VC-dim
0	0	0	1	6	0
1	0	7	36	39	0
2	1	36	23	16	50
3	5	29	20	21	25
4	19	10	2	0	0
5	4	0	0	0	0
>5	53	0	0	0	0

The results obtained from computing *Deg*, *BIP*, *3-BMIP*, *4-BMIP*, and *VC-dim* for the hypergraphs in the HyperBench benchmark are shown in Table 2.

Table 2 has to be read as follows: In the first column, we distinguish different values of the various hypergraph metrics. In the columns labeled “Deg,” “BIP,” and so on, we indicate for how many instances each metric has a particular value. For instance, by the last row in the second column, only 212 non-random CQs have degree > 5 . Actually, for most CQs, the degree is less than 10. Moreover, for the BMIP, already with intersections of three edges, we get 3-multi-intersection size ≤ 2 for almost all non-random CQs. Also the VC-dimension is ≤ 2 .

For CSPs, all properties may have higher values. However, we note a significant difference between randomly generated CSPs and the rest: For hypergraphs in the groups *CSP Application* and *CSP Other*, 543 (46%) hypergraphs have a high degree (>5), but nearly all instances have BIP or BMIP of less than 3. And most instances have a VC-dimension of at most 2. In contrast, nearly all random instances have a significantly higher degree (843 out of 863 instances with a degree > 5). Nevertheless, many instances have small BIP and BMIP. For nearly all hypergraphs (838 out of 863), we have 4-multi-intersection size ≤ 4 . For seven instances the computation of the VC-dimension timed out. For all others, the VC-dimension is ≤ 5 for random CSPs. Clearly, as seen in Table 2, the random CQs resemble the random CSPs a lot more than the CQ and CSP Application instances. For example, random CQs have similar to random CSPs high degree (382, corresponding to 76%,

with degree > 5), higher BIP and BMIP. Nevertheless, similar to random CSPs, the values for BIP and BMIP are still small for many random CQ instances.

To conclude, for the proposed properties—in particular, BIP/BMIP and VC-dimension—most of the hypergraphs in our benchmark indeed have low values.

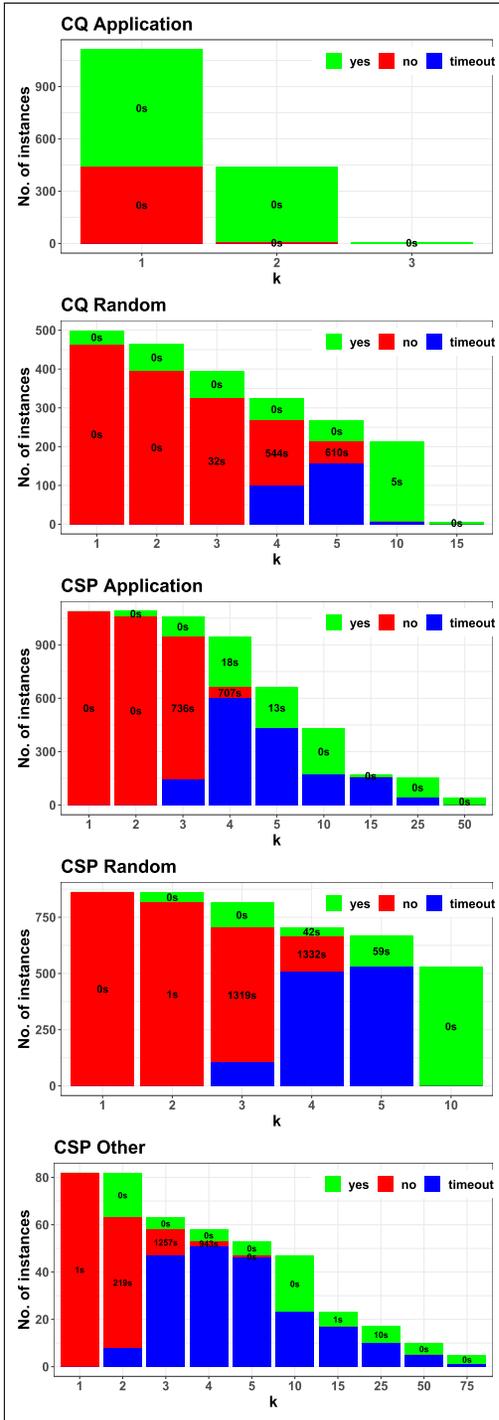
6.2 Hypertree Width

We have systematically applied the *hw*-computation from Reference [30] to all hypergraphs in the benchmark. The results are summarized in Figure 8. In our experiments, we proceeded as follows: We used the same classification of instances we used in the previous experiments, i.e., we distinguish the following classes: *CQ Application*, *CQ Random*, *CSP Application*, *CSP Random*, and *CSP Other*. For every hypergraph H , we first tried to solve the $\text{CHECK}(\text{HD}, k)$ problem for $k = 1$. In case of *CQ Application*, we thus got 673 yes-answers and 440 no-answers. The number in each bar indicates the average runtime to find these yes- and no-instances, respectively. Here, the average runtime was “0” (i.e., less than 1 second). For *CQ Random*, we got 36 yes- and 464 no-instances with an average runtime below 1 second. For all CSP-instances, we only got no-answers.

In the second round, we tried to solve the $\text{CHECK}(\text{HD}, k)$ problem for $k = 2$ for all hypergraphs that yielded a no-answer for $k = 1$. Now the picture is a bit more diverse: 432 of the remaining 440 CQs from *CQ Application* yielded a yes-answer in less than 1 second. For the hypergraphs stemming from *CQ Random*, only 68 instances yielded a yes-answer (in less than 1 second, on average), while 396 instances yielded a no-answer in less than 7 seconds, on average. The hypergraphs relative to CSP offer a different picture. The classes *CSP Application*, *CSP Random*, and *CSP Other* have 29, 47, and 19 yes-instances, respectively. Only eight instances from *CSP Other* give rise to a timeout (i.e., the program did not terminate within 3,600 seconds), while all the other instances give a no-answer within the timeout. Interestingly, our *hw*-algorithm gives a no-answer for 1,877 instances of *CSP Application* and *CSP Random* in less than 1 second, while it took the algorithm 219 seconds on average to answer “no” for 55 instances of *CSP Other*. This shows that the class *CSP Other* contains instances that are difficult to decompose.

This procedure is iterated by incrementing k and running the *hw*-computation for all instances, that either yielded a no-answer or a timeout in the previous round. For instance, for queries from *CQ Application*, one further round is needed after the second round. In other words, we confirm the observation of low *hw*, which was already made for CQs of arity ≤ 3 in References [12, 13, 48]. For the hypergraphs stemming from *CQ Random* (respectively, CSPs), 396 (respectively, 1,940) instances are left in the third round, of which 70 (respectively, 232) yield a yes-answer in less than 1 second, on average; 326 (respectively, 1,415) instances yield a no-answer in 32 (respectively, 988) seconds, on average; and no (respectively, 293) instances yield a timeout. Note that, as we increase k , the average runtime and the percentage of timeouts first increase up to a certain point and then they decrease. This is due to the fact that, as we increase k , the number of combinations of edges to be considered in each λ -label (i.e., the function λ_u at each node u of the decomposition) increases. In principle, we have to test $O(n^k)$ combinations, where n is the number of edges. However, if k increases beyond a certain point, then it gets easier to “guess” a λ -label, since an increasing portion of the $O(n^k)$ possible combinations leads to a solution (i.e., an HD of desired width).

To answer the question in *Subgoal 1*, it is indeed the case that for a big number of instances, the hypertree width is small enough to allow for efficient evaluation of CQs or CSPs: All instances of non-random CQs have $hw \leq 3$ no matter whether their arity is bounded by 3 (as in case of SPARQL queries) or not; and a large portion (at least 1,027, i.e., circa 50%) of all 2,035 CSP instances have $hw \leq 5$. In total, including random CQs, 2,427 (66.5%) out of 3,648 instances have $hw \leq 5$. And, out of these, we could determine the exact hypertree width for 2,356 instances; the others may even have lower *hw*.



CQ Application

k	Yes	No	Timeout
1	673	440	0
2	432	8	0
3	8	0	0

CQ Random

k	Yes	No	Timeout
1	36	464	0
2	68	396	0
3	70	326	0
4	59	167	100
5	54	55	158
10	206	0	7
15	7	0	0

CSP Application

k	Yes	No	Timeout
1	0	1090	0
2	29	1061	0
3	116	802	143
4	283	62	600
5	231	0	431
10	261	0	170
15	12	0	158
25	118	0	40
50	40	0	0

CSP Random

k	Yes	No	Timeout
1	0	863	0
2	47	816	0
3	111	602	103
4	39	160	506
5	136	0	530
10	530	0	0

CSP Other

k	Yes	No	Timeout
1	0	82	0
2	19	55	8
3	5	11	47
4	5	2	51
5	6	1	46
10	24	0	23
15	6	0	17
25	7	0	10
50	5	0	5
75	4	0	1

Fig. 8. HW analysis (labels are avg. runtimes in s).

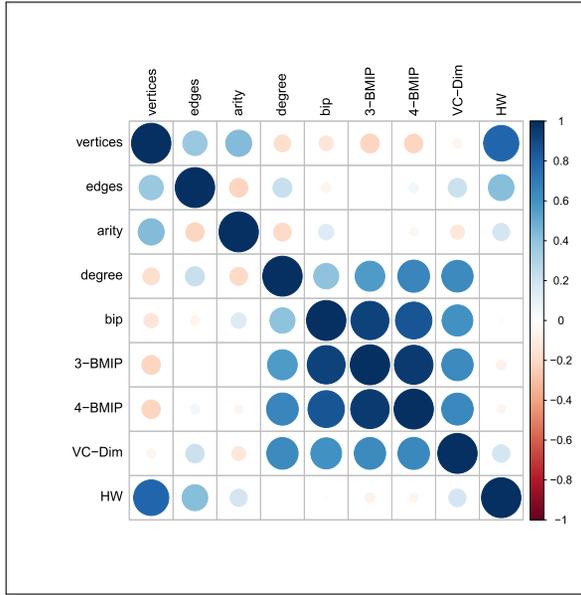


Fig. 9. Correlation analysis.

6.3 Correlation Analysis

Finally, we have analyzed the pairwise correlation between all properties. Of course, the different intersection sizes (BIP, 3-BMIP, 4-BMIP) are highly correlated. Other than that, we observe quite a strong correlation of the arity with the number of vertices and the hypertree width. Moreover, there is a significant correlation between number of vertices and arity and between number of vertices and hypertree width. Clearly, the correlation between arity and hypertree width is mainly due to the CSP instances and the random CQs, since, for non-random CQs, the hw never increases beyond 3, independently of the arity.

A graphical presentation of all pairwise correlations is given in Figure 9. Here, large, dark circles indicate a high correlation, while small, light circles stand for low correlation. Blue circles indicate a positive correlation, while red circles stand for a negative correlation. In Reference [19], it has been argued that Deg, BIP, 3-BMIP, 4-BMIP, and VC-dim are non-trivial restrictions to achieve tractability. It is interesting to note that, according to the correlations shown in Figure 9, these properties have almost no impact on the hypertree width of our hypergraphs. This underlines the usefulness of these restrictions in the sense that (a) they make the GHD computation and FHD approximation easier [19] but (b) low values of degree, (multi-)intersection-size, or VC-dimension do not pre-determine low values of the widths.

6.4 Comparison of ghw Algorithms

Here, we report on empirical results for the three ghw -algorithms described in Section 4. We have run the programs on each hypergraph from the HyperBench up to hypertree width 6, trying to get a smaller ghw than hw . We have thus run the ghw -algorithms with the following parameters: For all hypergraphs H with $hw(H) = k$ (or $hw \leq k$ and, due to timeouts, we do not know if $hw \leq k - 1$ holds), where $k \in \{3, 4, 5, 6\}$, try to solve the $\text{CHECK}(\text{GHD}, k - 1)$ problem. In other words, we just tried to improve the width by 1. Clearly, for $hw(H) \in \{1, 2\}$, no improvement is possible, since, in this case, $hw(H) = ghw(H)$ holds.

Table 3. GHW Algorithms with Avg. Runtimes in s

$hw \rightarrow ghw$	Total	GlobalBIP		LocalBIP		BalSep	
		yes (s)	no (s)	yes (s)	no (s)	yes (s)	no (s)
3 \rightarrow 2	310	-	128 (537)	-	195 (162)	-	307 (12)
4 \rightarrow 3	386	-	137 (2,809)	-	54 (2606)	-	249 (54)
5 \rightarrow 4	427	-	-	-	-	-	148 (13)
6 \rightarrow 5	459	13 (162)	-	13 (60)	-	-	180 (288)

In Table 3, for each algorithm, we report on the number of “successful” attempts to solve the $\text{CHECK}(\text{GHD}, k - 1)$ problem for hypergraphs with $hw = k$. Here, “successful” means that the program terminated within one hour. For instance, for the 310 hypergraphs with $hw = 3$ in the HyperBench, GlobalBIP terminated in 128 cases (i.e., 41%) when trying to solve $\text{CHECK}(\text{GHD}, 2)$. The average runtime of these “successful” runs was 537 seconds. For the 386 hypergraphs with $hw = 4$, GlobalBIP terminated in 137 cases (i.e., 35%) with average runtime 2,809 when trying to solve the $\text{CHECK}(\text{GHD}, 3)$ problem. For the 886 hypergraphs with $hw \in \{5, 6\}$, GlobalBIP only terminated in 13 cases (i.e., 1.4%). Overall, it turns out that the set $f(H, k)$ may be very big (even though it is polynomial if k and i are constants). Hence, H' can become considerably bigger than H . This explains the frequent timeouts in the GlobalBIP column in Table 3.

The results obtained with LocalBIP are shown in the corresponding column. Interestingly, for the hypergraphs with $hw = 3$, the “local” computation performs significantly better (namely, 63% solved with average runtime 162 seconds rather than 41% with average runtime 537 seconds). In contrast, for the hypergraphs with $hw = 4$, the “global” computation is significantly more successful. For $hw \in \{5, 6\}$, the “global” and “local” computations are equally bad. A possible explanation for the reverse behavior of “global” and “local” computation in case of $hw = 3$ as opposed to $hw = 4$ is that the restriction of the “global” set $f(H, k)$ of subedges to the “local” set $f_u(H, k)$ at each node u seems to be quite effective for the hypergraphs with $hw = 3$. In contrast, the additional cost of having to compute $f_u(H, k)$ at each node u becomes counter-productive, when the set of subedges thus eliminated is not significant. It is interesting to note that the sets of solved instances of the global computation and the local computation are incomparable, i.e., in some cases one method is better, while in other cases the other method is better.

If we look at the number of solved instances in Table 3, then we see that the recursive algorithm via balanced separators (reported in the last column labeled BalSep) has the least number of timeouts due to the fast identification of negative instances (i.e., those with no-answer), where it often detects quite fast that a given hypergraph does not have a balanced separator of desired width. As k increases, the performance of the balanced separators approach deteriorates. This is due to k in the exponent of the running time of our algorithm, i.e., we need to check for each of the possible $O(n^{k+1})$ combinations of $\leq k$ edges if it constitutes a balanced separator. Note that the balanced separators approach only terminated in case of no-answers.

We now look at Table 4, where we report for all hypergraphs with $hw \leq k$ and $k \in \{3, 4, 5, 6\}$, whether $ghw \leq k - 1$ could be verified. To this end, we run our three algorithms (GlobalBIP, LocalBIP, and BalSep) in parallel and stop the computation as soon as one terminates (with answer “yes” or “no”). The number in parentheses refers to the average runtime needed by the fastest of the three algorithms in each case. A timeout occurs if none of the three algorithms terminates within 3,600 seconds. It is interesting to note that in the vast majority of cases, no improvement of the width is possible when we switch from hw to ghw : In 97% of the solved cases with $hw \leq 6$, which form 65% of all instances, hw and ghw have identical values. Actually, we think that the

Table 4. GHW of Instances with Average Runtime in s

$hw \rightarrow ghw$	yes	no	timeout
$3 \rightarrow 2$	0	309 (10)	1
$4 \rightarrow 3$	0	262 (57)	124
$5 \rightarrow 4$	0	148 (13)	279
$6 \rightarrow 5$	18 (129)	180 (288)	261

high percentage of the *solved cases* gives a more realistic picture than the percentage of *all cases* for the following reason: Our algorithms (in particular, the “global” and “local” computations) need a particularly long time for negative instances. This is due to the fact that in a negative case, “all” possible choices of λ -labels for a node u in the GHD have to be tested before we can be sure that no GHD of H (or, equivalently, no HD of H') of desired width exists. Hence, it seems plausible that the timeouts are mainly due to negative instances. This also explains why our new BalSep algorithm, which is particularly well suited for negative instances, has the least number of timeouts.

A closer comparison of Table 3 and Table 4 makes clear that BalSep is superior to GlobalBIP and LocalBIP in solving negative instances. Indeed, the combined approach summarized in Table 4 relies almost completely on the runs of BalSep. The algorithms GlobalBIP and LocalBIP managed to solve only 15 out of 571 negative instances for $k \in \{3, 4\}$, while they could not give any negative answer for $k \in \{5, 6\}$. However, GlobalBIP and LocalBIP solved 18 positive instances, while BalSep did not terminate at all. Nevertheless, this does not significantly diminish the strength of BalSep as a powerful tool for negative instances.

We conclude this section with a final observation: In Figure 8, we had many cases for which only some upper bound k on the hw could be determined, namely, those cases where the attempt to solve $\text{CHECK}(\text{HD}, k)$ yields a yes-answer and the attempt to solve $\text{CHECK}(\text{HD}, k - 1)$ gives a timeout. In several such cases, we could get (with the balanced separator approach) a no-answer for the $\text{CHECK}(\text{GHD}, k - 1)$ problem, which implicitly gives a no-answer for the problem $\text{CHECK}(\text{HD}, k - 1)$. In this way, our new ghw -algorithm is also profitable for the hw -computation: For 827 instances with $hw \leq 6$, we were not able to determine the exact hypertree width. Using our new ghw -algorithm, we closed this gap for 297 instances; for these instances $hw = ghw$ holds.

To sum up, we now have a total of 2,356 (64.5%) instances for which we determined the exact hw and a total of 1,984 instances (54.4%) for which we determined the exact ghw . Out of these, 1,968 instances had identical values for hw and ghw . In 16 cases, we found an improvement of the width by 1 when moving from hw to ghw , namely, from $hw = 6$ to $ghw = 5$. In 2 further cases, we could show $hw \leq 6$ and $ghw \leq 5$, but the attempt to check $hw = 5$ or $ghw = 4$ led to a timeout. Hence, in response to *Subgoal 5*, hw is equal to ghw in 54.4% of the cases if we consider all instances and in 68.2% of the cases (1,968 of 2,886) with small width ($hw \leq 6$). However, if we consider the fully solved cases (i.e., where we have the precise value of hw and ghw), then hw and ghw coincide in 99.2% of the cases (1,968 of 1,984).

6.5 Fractionally Improved Decompositions

The algorithms for computing FHDs in the literature are very expensive and even the tractability result presented in Reference [19] involves a double exponential “constant.” Here, we propose two algorithms for computing an FHD of a hypergraph when we already have a GHD: `ImproveHD`, and `FracImproveHD`. They differ in the compromise between computational cost and quality of the approximation.

Table 5. Instances Solved with ImproveHD

hw	≥ 1	$[0.5, 1)$	$[0.1, 0.5)$	no	timeout
2	0	136	40	419	0
3	12	104	25	169	0
4	9	55	11	311	0
5	20	14	11	382	0
6	12	60	80	307	0

Table 6. Instances Solved with FracImproveHD

hw	≥ 1	$[0.5, 1)$	$[0.1, 0.5)$	no	timeout
2	0	194	46	353	2
3	14	116	21	135	24
4	11	81	2	8	284
5	18	126	59	2	222
6	28	149	95	4	183

The first algorithm we present is based on a simple observation: Given an (G)HD, we could substitute its integral edge covers with fractional edge covers and obtain an FHD. Formally, let $\mathcal{D} = \langle T, (B_u)_{u \in T}, (\lambda_u)_{u \in T} \rangle$ be either a GHD or an HD. Our algorithm ImproveHD computes an FHD $\mathcal{D}' = \langle T', (B'_u)_{u \in T'}, (\gamma_u)_{u \in T'} \rangle$ where:

- The tree T' is the same as T .
- For each node $u \in T'$, the bag $B'_u = B_u$.
- For each node $u \in T'$, γ_u is a minimum-weight fractional edge cover of B'_u .

To obtain the FHD \mathcal{D}' , we iterate over the nodes of \mathcal{D} and, for each B_u , we compute a minimum-weight fractional edge cover of B_u . Since computing such a fractional edge cover is polynomial and we assume to have already computed an HD to start with, the whole algorithm is very efficient. Nevertheless, it is clear that there is a strong dependence on the starting HD. This is unsatisfactory, and so we devised a more sophisticated algorithm.

The algorithm we describe here gets rid of the dependence on a particular HD and computes a fractionally improved (G)HD with a fixed improvement threshold. We call this algorithm FracImproveHD. It searches for an FHD \mathcal{D}' with $\mathcal{D}' = \text{SimpleImproveHD}(\mathcal{D})$ for some HD \mathcal{D} of H with $\text{width}(\mathcal{D}) \leq k$ and $\text{width}(\mathcal{D}') \leq k'$. Here, k is an upper bound on the hw and k' the desired fractionally improved hw . In other words, this algorithm searches for the best fractionally improved HD over all HDs of width $\leq k$. Hence, the result is independent of any concrete HD.

The algorithm FracImproveHD is built on top of the HD construction described in Section 3.4. Recall that, given a hypergraph H , this algorithm maintains a set of edges $C \subseteq E(H)$, which represents the current component to decompose. While searching for a separator λ_u of C , we do not only want that $|\lambda_u| \leq k$, but we also require that, among all possible choices of λ_u , we choose one such that $\text{weight}(\gamma_u) \leq k'$, where γ_u is a fractional edge cover of $B(\lambda_u)$. We thus guarantee that the output of the algorithm is the desired FHD.

The experimental results with these algorithms for computing fractionally improved HDs are summarized in Tables 5 and 6. We have applied these algorithms to all hypergraphs for which $hw \leq k$ with $k \in \{2, 3, 4, 5, 6\}$ is known from Figure 8. The various columns of the Tables 5 and 6

are as follows: The first column (labeled hw) refers to the (upper bound on the) hw according to Figure 8. The next three columns, labeled ≥ 1 , $[0.5, 1)$, and $[0.1, 0.5)$ tell us, by how much the width can be improved (if at all), if we compute an FHD by one of the two algorithms. We thus distinguish the three cases if, for a hypergraph of $hw \leq k$, we manage to construct an FHD of width $k - c$ for $c \geq 1$, $c \in [0.5, 1)$, or $c \in [0.1, 0.5)$. The column with label “no” refers to the cases where no improvement at all or at least no improvement by $c \geq 0.1$ was possible. The last column counts the number of timeouts.

For instance, in the first row of Table 5, we see that (with the ImproveHD algorithm and starting from the HD obtained by the hw -computation of Figure 8) out of 595 hypergraphs with $hw = 2$, no improvement was possible in 419 cases. In the remaining 176 cases, an improvement to a width of at most $2 - 0.5$ was possible in 40 cases and an improvement to $k - c$ with $c \in [0.1, 0.5)$ was possible in 136 cases. For the hypergraphs with $hw = 3$ in Figure 8, almost half of the hypergraphs (141 out of 310) allowed at least some improvement, in particular, 104 by $c \in [0.5, 1)$ and 12 even by at least 1. The improvements achieved for the hypergraphs with $hw \leq 4$ and $hw \leq 5$ are less significant.

The results obtained with our FracImproveHD implementation are displayed in Table 6. We see that the number of hypergraphs that allow for a fractional improvement of the width by at least 0.5 or even by 1 is often bigger than with ImproveHD—in particular, in the cases where $k' \leq k$ with $k \in \{4, 5\}$ holds. In the other cases, the results obtained with the naive ImproveHD algorithm are not much worse than with the more sophisticated FracImproveHD algorithm.

7 CONCLUSION

In this work, we have presented HyperBench, a new and comprehensive benchmark of hypergraphs derived from CQs and CSPs from various areas, together with the results of extensive empirical analyses with this benchmark.

Lessons learned. The empirical study has brought many insights. Below, we summarize the most important lessons learned from our studies.

- The finding of References [12, 13, 48] that non-random CQs have low hypertree width has been confirmed by our analysis, even if (in contrast to SPARQL queries) the arity of the CQs is not bounded by 3. For random CQs and CSPs, we have detected a correlation between the arity and the hypertree width, although also in this case, the increase of the hw with increased arity is not dramatic.

- In Reference [19], several hypergraph invariants were identified, which make the computation of GHDs and the approximation of FHDs tractable. We have seen that, at least for non-random instances, these invariants indeed tend to have low values.

- The reduction of the ghw -computation problem to the hw -computation problem in case of low intersection size turned out to be more problematical than the theoretical tractability results from Reference [19] had suggested. Even the improvement by “local” computation of the additional subedges did not help much. However, we were able to improve this significantly by presenting a new algorithm based on “balanced separators.” In particular, for negative instances (i.e., those with a no-answer), this approach proved very effective.

- An additional benefit of the new ghw -algorithm based on “balanced separators” is that it allowed us to also fill gaps in the hw -computation. Indeed, in several cases, we managed to verify $hw \leq k$ for some k but we could not show $hw \leq k - 1$, due to a timeout for CHECK(HD, $k - 1$). By establishing $ghw \leq k - 1$ with our new GHD-algorithm, we have implicitly shown $hw \leq k - 1$. This allowed us to compute the exact hw of many further hypergraphs.

- Most surprisingly, the discrepancy between hw and ghw is much lower than expected. Theoretically, only the upper bound $hw \leq 3 \cdot ghw + 1$ is known. However, in practice, when considering

hypergraphs of $hw \leq 6$, we could show that in circa 54% of all cases, hw and ghw are simply identical. Moreover, in *all* cases when one of our implementations of ghw -computation terminated on instances with $hw \leq 5$, we got identical values for hw and ghw .

Future work. Our empirical study has also given us many hints for future directions of research. We find the following tasks particularly urgent and/or rewarding:

- So far, we have only implemented the ghw -computation in case of low intersection size. In Reference [19], tractability of the $\text{CHECK}(\text{GHD}, k)$ problem was also proved for the more relaxed bounded multi-intersection size. Our empirical results in Table 2 show that, apart from the random CQs and random CSPs, the 3-multi-intersection size is ≤ 2 in almost all cases. It seems therefore worthwhile to implement and test also the BMIP-algorithm from Reference [19].

- The three approaches for ghw -computation presented here turned out to have complementary strengths and weaknesses. This was profitable when running all three algorithms in parallel and taking the result of the first one that terminates (see Table 4). In the future, we also want to implement a more sophisticated combination of the various approaches: For instance, one could try to apply our new “balanced separator” algorithm recursively only down to a certain recursion depth (say, depth 2 or 3) to split a big given hypergraph into smaller subhypergraphs and then continue with the “global” or “local” computation from Section 4. First promising results in this direction have recently been obtained in Reference [29].

- Our new approach to ghw -computation via “balanced separators” proved quite effective in our experiments. However, further theoretical underpinning of this approach is missing. The empirical results obtained for our new GHD algorithm via balanced separators suggest that the number of balanced separators is often drastically smaller than the number of arbitrary separators. We want to determine a realistic upper bound on the number of balanced separators in terms of n (the number of edges) and k (an upper bound on the width). This will then allow us to compute also a realistic upper bound on the runtime of this new algorithm.

- We want to further extend the HyperBench benchmark and tool in several directions. We will thus incorporate further implementations of decomposition algorithms from the literature such as the GHD- and FHD computation in Reference [45] or the polynomial-time FHD computation for hypergraphs of bounded multi-intersection size in Reference [25]. Moreover, we will continue to fill in hypergraphs from further sources of CSPs and CQs. For instance, in References [1, 14, 22, 23], a collection of CQs for the experimental evaluations in those papers is mentioned. We will invite the authors to disclose these CQs and incorporate them into the HyperBench benchmark.

- We also want to include hypergraphs coming from diverse domains and different applications. As a first step toward this aim, we plan to release an extended version of our software *hg-tools* for extracting hypergraphs from SPARQL queries. In this way, we hope to acquire more queries from applications using graph databases.

- Finally, we want to make use of HyperBench to test the practical feasibility of using decompositions to evaluate CQs and solving CSPs. To this end, we will extend our collection of hypergraphs with the data of these problems. In other words, we want to include in our benchmark the relations corresponding to the database for CQs and constraints for CSPs. Such a study would have as a primary goal to assess the usefulness of decompositions in solving related problems, but it would also help identify which characteristics should decompositions have to serve this purpose.

ACKNOWLEDGMENTS

We would like to thank Angela Bonifati, Wim Martens, and Thomas Timm for sharing most of the hypergraphs with $hw \geq 2$ from their work cited in References [12, 13] and for their effort in anonymizing these hypergraphs, which was required by the license restrictions.

REFERENCES

- [1] Christopher R. Aberger, Andrew Lamb, Susan Tu, Andres Nötzli, Kunle Olukotun, and Christopher Ré. 2017. Empty-Headed: A relational engine for graph processing. *ACM Trans. Datab. Syst.* 42, 4 (2017), 20:1–20:44. DOI: <https://doi.org/10.1145/3129246>.
- [2] Christopher R. Aberger, Susan Tu, Kunle Olukotun, and Christopher Ré. 2016. Old techniques for new join algorithms: A case study in RDF processing. *CoRR* abs/1602.03557 (2016).
- [3] Isolde Adler, Georg Gottlob, and Martin Grohe. 2007. Hypertree width and related hypergraph invariants. *Eur. J. Comb.* 28, 8 (2007), 2167–2181. DOI: <https://doi.org/10.1016/j.ejc.2007.04.013>.
- [4] Kamal Amroun, Zineb Habbas, and Wassila Aggoune-Mtalaa. 2016. A compressed generalized hypertree decomposition-based technique for non-binary constraint satisfaction problems. *AI Commun.* 29, 2 (2016), 371–392. DOI: <https://doi.org/10.3233/AIC-150694>.
- [5] Molham Aref, Balder ten Cate, Todd J. Green, Benny Kimelfeld, Dan Olteanu, Emir Pasalic, Todd L. Veldhuizen, and Geoffrey Washburn. 2015. Design and implementation of the LogicBlox system. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, Timos K. Sellis, Susan B. Davidson, and Zachary G. Ives (Eds.). ACM, 1371–1382. DOI: <https://doi.org/10.1145/2723372.2742796>.
- [6] Patricia C. Arocena, Boris Glavic, Radu Ciucanu, and Renée J. Miller. 2015. The iBench integration metadata generator. *PVLDB* 9, 3 (2015), 108–119. DOI: <https://doi.org/10.14778/2850583.2850586>.
- [7] Gilles Audemard, Frédéric Boussemart, Christophe Lecoutre, and Cédric Piette. 2016. XCSP3: an XML-based format designed to represent combinatorial constrained problems. Retrieved from <http://www.xcsp.org/>.
- [8] Nurzhan Bakibayev, Tomáš Kociský, Dan Olteanu, and Jakub Zavodny. 2013. Aggregation and ordering in factorised databases. *PVLDB* 6, 14 (2013), 1990–2001. DOI: <https://doi.org/10.14778/2556549.2556579>.
- [9] Michael Benedikt. 2017. CQ benchmarks. Personal Communication.
- [10] Michael Benedikt, George Constantinidis, Giansalvatore Mecca, Boris Motik, Paolo Papotti, Donatello Santoro, and Efthymia Tsamoura. 2017. Benchmarking the chase. In *Proceedings of the 36th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS'17)*, Emanuel Sallinger, Jan Van den Bussche, and Floris Geerts (Eds.). ACM, 37–52. DOI: <https://doi.org/10.1145/3034786.3034796>.
- [11] Jeremias Berg, Neha Lodha, Matti Järvisalo, and Stefan Szeider. 2017. MaxSAT benchmarks based on determining generalized hypertree-width. *MaxSAT Evaluation 2017: Solver and Benchmark Descriptions B-2017-2* (2017), 22.
- [12] Angela Bonifati, Wim Martens, and Thomas Timm. 2017. An analytical study of large SPARQL query logs. *PVLDB* 11, 2 (2017), 149–161. DOI: <https://doi.org/10.14778/3149193.3149196>.
- [13] Angela Bonifati, Wim Martens, and Thomas Timm. 2019. Navigating the maze of wikidata query logs. In *Proceedings of the World Wide Web Conference (WWW'19)*, Ling Liu, Ryen W. White, Amin Mantrach, Fabrizio Silvestri, Julian J. McAuley, Ricardo Baeza-Yates, and Leila Zia (Eds.). ACM, 127–138. DOI: <https://doi.org/10.1145/3308558.3313472>.
- [14] Nofar Carmeli, Batya Kenig, and Benny Kimelfeld. 2017. Efficiently enumerating minimal triangulations. In *Proceedings of the 36th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS'17)*, Emanuel Sallinger, Jan Van den Bussche, and Floris Geerts (Eds.). ACM, 273–287. DOI: <https://doi.org/10.1145/3034786.3056109>.
- [15] Ashok K. Chandra and Philip M. Merlin. 1977. Optimal implementation of conjunctive queries in relational data bases. In *Proceedings of the 9th ACM Symposium on Theory of Computing*, John E. Hopcroft, Emily P. Friedman, and Michael A. Harrison (Eds.). ACM, 77–90. DOI: <https://doi.org/10.1145/800105.803397>.
- [16] Rina Dechter. 2003. *Constraint Processing*. Elsevier. DOI: <https://doi.org/10.1016/b978-1-55860-890-0.x5000-2>.
- [17] Johannes Klaus Fichte, Markus Hecher, Neha Lodha, and Stefan Szeider. 2018. An SMT approach to fractional hypertree width. In *Proceedings of the 24th International Conference on Principles and Practice of Constraint Programming (Lecture Notes in Computer Science, Vol. 11008)*, John N. Hooker (Ed.). Springer, 109–127. DOI: https://doi.org/10.1007/978-3-319-98334-9_8.
- [18] Wolfgang Fischl, Georg Gottlob, Davide Mario Longo, and Reinhard Pichler. 2019. HyperBench: A benchmark and tool for hypergraphs and empirical findings. In *Proceedings of the 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS'19)*, Dan Suciu, Sebastian Skritek, and Christoph Koch (Eds.). ACM, 464–480. DOI: <https://doi.org/10.1145/3294052.3319683>.
- [19] Wolfgang Fischl, Georg Gottlob, and Reinhard Pichler. 2018. General and fractional hypertree decompositions: Hard and easy cases. In *Proceedings of the 37th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, Jan Van den Bussche and Marcelo Arenas (Eds.). ACM, 17–32. DOI: <https://doi.org/10.1145/3196959.3196962>.
- [20] Tobias Ganzow, Georg Gottlob, Nysret Musliu, and Marko Samer. 2005. *A CSP Hypergraph Library*. Technical Report. Institute of Information Systems (DBAI). <https://dbai.tuwien.ac.at/research/project/hypertree/cspghl.pdf>.
- [21] Floris Geerts, Giansalvatore Mecca, Paolo Papotti, and Donatello Santoro. 2014. Mapping and cleaning. In *Proceedings of the IEEE 30th International Conference on Data Engineering (ICDE'14)*, Isabel F. Cruz, Elena Ferrari, Yufei Tao, Elisa Bertino, and Goce Trajcevski (Eds.). IEEE Computer Society, 232–243. DOI: <https://doi.org/10.1109/ICDE.2014.6816654>.

- [22] Lucantonio Ghionna, Luigi Granata, Gianluigi Greco, and Francesco Scarcello. 2007. Hypertree decompositions for query optimization. In *Proceedings of the 23rd International Conference on Data Engineering (ICDE'07)*, Rada Chirkova, Asuman Dogac, M. Tamer Özsu, and Timos K. Sellis (Eds.). IEEE Computer Society, 36–45. DOI: <https://doi.org/10.1109/ICDE.2007.367849>.
- [23] Lucantonio Ghionna, Gianluigi Greco, and Francesco Scarcello. 2011. H-DB: A hybrid quantitative-structural sql optimizer. In *Proceedings of the 20th ACM Conference on Information and Knowledge Management (CIKM'11)*, Craig Macdonald, Iadh Ounis, and Ian Ruthven (Eds.). ACM, 2573–2576. DOI: <https://doi.org/10.1145/2063576.2064023>.
- [24] Georg Gottlob, Gianluigi Greco, Nicola Leone, and Francesco Scarcello. 2016. Hypertree decompositions: Questions and answers. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, (PODS'16)*, Tova Milo and Wang-Chiew Tan (Eds.). ACM, 57–74. DOI: <https://doi.org/10.1145/2902251.2902309>.
- [25] Georg Gottlob, Matthias Lanzinger, Reinhard Pichler, and Igor Razgon. 2020. Complexity analysis of generalized and fractional hypertree decompositions. *CoRR* abs/2002.05239 (2020).
- [26] Georg Gottlob, Nicola Leone, and Francesco Scarcello. 1999. On tractable queries and constraints. In *Proceedings of the 10th International Conference on Database and Expert Systems Applications (DEXA'99)*(Lecture Notes in Computer Science, Vol. 1677), Trevor J. M. Bench-Capon, Giovanni Soda, and A. Min Tjoa (Eds.). Springer, 1–15. DOI: https://doi.org/10.1007/3-540-48309-8_1.
- [27] Georg Gottlob, Nicola Leone, and Francesco Scarcello. 2002. Hypertree decompositions and tractable queries. *J. Comput. Syst. Sci.* 64, 3 (2002), 579–627. DOI: <https://doi.org/10.1006/jcss.2001.1809>.
- [28] Georg Gottlob, Zoltán Miklós, and Thomas Schwentick. 2009. Generalized hypertree decompositions: NP-hardness and tractable variants. *J. ACM* 56, 6 (2009), 30:1–30:32. DOI: <https://doi.org/10.1145/1568318.1568320>.
- [29] Georg Gottlob, Cem Okulmus, and Reinhard Pichler. 2020. Fast and parallel decomposition of constraint satisfaction problems. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI'20)*, Christian Bessiere (Ed.). ijcai.org, 1155–1162. DOI: <https://doi.org/10.24963/ijcai.2020/161>.
- [30] Georg Gottlob and Marko Samer. 2008. A backtracking-based algorithm for hypertree decomposition. *ACM J. Exp. Algor.* 13 (2008), 1:1.1–1:1.19. DOI: <https://doi.org/10.1145/1412228.1412229>.
- [31] Martin Grohe and Dániel Marx. 2014. Constraint solving via fractional edge covers. *ACM Trans. Algor.* 11, 1 (2014), 4:1–4:20. DOI: <https://doi.org/10.1145/2636918>.
- [32] Yuanbo Guo, Zhengxiang Pan, and Jeff Heflin. 2005. LUBM: A benchmark for OWL knowledge base systems. *J. Web Semant.* 3, 2-3 (2005), 158–182. DOI: <https://doi.org/10.1016/j.websem.2005.06.005>.
- [33] Zineb Habbas, Kamal Amroun, and Daniel Singer. 2015. A forward-checking algorithm based on a generalised hypertree decomposition for solving non-binary constraint satisfaction problems. *J. Exp. Theor. Artif. Intell.* 27, 5 (2015), 649–671. DOI: <https://doi.org/10.1080/0952813X.2014.993507>.
- [34] Shrainik Jain, Dominik Moritz, Daniel Halperin, Bill Howe, and Ed Lazowska. 2016. SQLShare: Results from a multi-year SQL-as-a-service experiment. In *Proceedings of the International Conference on Management of Data (SIGMOD'16)*, Fatma Özcan, Georgia Koutrika, and Sam Madden (Eds.). ACM, 281–293. DOI: <https://doi.org/10.1145/2882903.2882957>.
- [35] Shant Karakashian, Robert J. Woodward, and Berthe Y. Choueiry. 2011. Reformulating $R^*(m)C$ with tree decomposition. In *Proceedings of the 9th Symposium on Abstraction, Reformulation, and Approximation (SARA'11)*, Michael R. Genesereth and Peter Z. Revesz (Eds.). AAAI, 62–69. Retrieved from <http://www.aaai.org/ocs/index.php/SARA/SARA11/paper/view/4234>.
- [36] Mahmoud Abo Khamis, Hung Q. Ngo, Christopher Ré, and Atri Rudra. 2016. Joins via geometric resolutions: Worst case and beyond. *ACM Trans. Datab. Syst.* 41, 4 (2016), 22:1–22:45. DOI: <https://doi.org/10.1145/2967101>.
- [37] Mahmoud Abo Khamis, Hung Q. Ngo, and Atri Rudra. 2016. FAQ: Questions asked frequently. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS'16)*, Tova Milo and Wang-Chiew Tan (Eds.). ACM, 13–28. DOI: <https://doi.org/10.1145/2902251.2902280>.
- [38] Tuukka Korhonen. 2019. Potential Maximal Cliques Parameterized by Edge Clique Cover. arxiv:1912.10989 [cs.DS].
- [39] Mohammed Lalou, Zineb Habbas, and Kamal Amroun. 2009. Solving hypertree structured CSP: Sequential and parallel approaches. In *Proceedings of the 16th RCRA Workshop on Experimental Evaluation of Algorithms for Solving Problems with Combinatorial Explosion (RCRA@AI'IA'09)*(CEUR Workshop Proceedings, Vol. 589), Marco Gavanelli and Toni Mancini (Eds.). CEUR-WS.org. Retrieved from <http://ceur-ws.org/Vol-589/paper11.pdf>.
- [40] Viktor Leis, Andrey Gubichev, Atanas Mirchev, Peter A. Boncz, Alfons Kemper, and Thomas Neumann. 2015. How good are query optimizers, really? *PVLDB* 9, 3 (2015), 204–215. DOI: <https://doi.org/10.14778/2850583.2850594>.
- [41] Viktor Leis, Bernhard Radke, Andrey Gubichev, Atanas Mirchev, Peter A. Boncz, Alfons Kemper, and Thomas Neumann. 2018. Query optimization through the looking glass, and what we found running the join order benchmark. *VLDB J.* 27, 5 (2018), 643–668. DOI: <https://doi.org/10.1007/s00778-017-0480-7>.
- [42] Stanislav Malyshev, Markus Krötzsch, Larry González, Julius Gonsior, and Adrian Bielefeldt. 2018. Getting the Most Out of Wikidata: Semantic Technology Usage in Wikipedia's Knowledge Graph. In *Proceedings of the 17th International Semantic Web Conference (ISWC'18)*(Lecture Notes in Computer Science, Vol. 11137), Denny Vrandečić, Kalina

- Bontcheva, Mari Carmen Suárez-Figueroa, Valentina Presutti, Irene Celino, Marta Sabou, Lucie-Aimée Kaffee, and Elena Simperl (Eds.). Springer, 376–394. DOI : https://doi.org/10.1007/978-3-030-00668-6_23.
- [43] Dániel Marx. 2010. Approximating fractional hypertree width. *ACM Trans. Algor.* 6, 2 (2010), 29:1–29:17. DOI : <https://doi.org/10.1145/1721837.1721845>.
- [44] Dimitrios Michail, Joris Kinable, Barak Naveh, and John V. Sichi. 2020. JGraphT—A Java library for graph data structures and algorithms. *ACM Trans. Math. Softw.* 46, 2 (May 2020).
- [45] Lukas Moll, Siamak Tazari, and Marc Thurley. 2012. Computing hypergraph width measures exactly. *Inf. Proc. Lett.* 112, 6 (2012), 238–242. DOI : <https://doi.org/10.1016/j.ipl.2011.12.002>.
- [46] Dan Olteanu and Jakub Závodný. 2015. Size bounds for factorised representations of query results. *ACM Trans. Datab. Syst.* 40, 1 (2015), 2:1–2:44. DOI : <https://doi.org/10.1145/2656335>.
- [47] Adam Perelman and Christopher Ré. 2015. DunceCap: Compiling worst-case optimal query plans. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, Timos K. Sellis, Susan B. Davidson, and Zachary G. Ives (Eds.). ACM, 2075–2076. DOI : <https://doi.org/10.1145/2723372.2764945>.
- [48] François Picalausa and Stijn Vansummeren. 2011. What are real SPARQL queries like? In *Proceedings of the International Workshop on Semantic Web Information Management (SWIM'11)*, Roberto De Virgilio, Fausto Giunchiglia, and Letizia Tanca (Eds.). ACM, 7. DOI : <https://doi.org/10.1145/1999299.1999306>.
- [49] Rachel Pottinger and Alon Y. Halevy. 2001. MiniCon: A scalable algorithm for answering queries using views. *VLDB J.* 10, 2-3 (2001), 182–198. DOI : <https://doi.org/10.1007/s007780100048>.
- [50] Francesco Scarcello, Gianluigi Greco, and Nicola Leone. 2007. Weighted hypertree decompositions and optimal query plans. *J. Comput. Syst. Sci.* 73, 3 (2007), 475–506. DOI : <https://doi.org/10.1016/j.jcss.2006.10.010>.
- [51] Werner Schafhauser. 2006. *New heuristic methods for tree decompositions and generalized hypertree decompositions*. Master's thesis. Technische Universität Wien.
- [52] André Schidler and Stefan Szeider. 2020. Computing optimal hypertree decompositions. In *Proceedings of the Symposium on Algorithm Engineering and Experiments, (ALENEX'20)*, Guy E. Blelloch and Irene Finocchi (Eds.). SIAM, 1–11.
- [53] Transaction Processing Performance Council (TPC). 2006. TPC-DS decision support benchmark. Retrieved from <http://www.tpc.org/tpcds/>.
- [54] Transaction Processing Performance Council (TPC). 2014. TPC-H decision support benchmark. Retrieved from <http://www.tpc.org/tpch/default.asp>.
- [55] Susan Tu and Christopher Ré. 2015. DunceCap: Query plans using generalized hypertree decompositions. In *Proceedings of the ACM SIGMOD International Conference on Management of Data, Melbourne*, Timos K. Sellis, Susan B. Davidson, and Zachary G. Ives (Eds.). ACM, 2077–2078. DOI : <https://doi.org/10.1145/2723372.2764946>.
- [56] V. N. Vapnik and A. Ya. Chervonenkis. 1971. On the uniform convergence of relative frequencies of events to their probabilities. *Theor. Probab. Its Applic.* 16, 2 (Jan. 1971), 264–280. DOI : <https://doi.org/10.1137/1116025>.
- [57] Tobias Warneke. 2019. JSqlParser. Retrieved from <https://github.com/JSqlParser/JSqlParser>.

Received July 2020; accepted November 2020