

Fast Parallel Hypertree Decompositions in Logarithmic Recursion Depth

Georg Gottlob
Matthias Lanzinger

University of Oxford
Oxford, UK
georg.gottlob@cs.ox.ac.uk
matthias.lanzinger@cs.ox.ac.uk

Cem Okulmus
Reinhard Pichler

TU Wien
Vienna, Austria
cokulmus@dbai.tuwien.ac.at
pichler@dbai.tuwien.ac.at

ABSTRACT

Various classic reasoning problems with natural hypergraph representations are known to be tractable when a hypertree decomposition (HD) of low width exists. The resulting algorithms are attractive for practical use in fields like databases and constraint satisfaction. However, algorithmic use of HDs relies on the difficult task of first computing a decomposition of the hypergraph underlying a given problem instance, which is then used to guide the algorithm for this particular instance. The performance of purely sequential methods for computing HDs is inherently limited, yet the problem is, theoretically, amenable to parallelisation.

In this paper we propose the first algorithm for computing hypertree decompositions that is well-suited for parallelisation. The newly proposed algorithm $\log\text{-}k\text{-decomp}$ requires only a logarithmic number of recursion levels and additionally allows for highly parallelised pruning of the search space by restriction to so-called balanced separators. We provide a detailed experimental evaluation over the HyperBench benchmark and demonstrate that $\log\text{-}k\text{-decomp}$ outperforms the current state-of-the-art significantly.

CCS CONCEPTS

• **Information systems** \rightarrow *Relational database query languages*;
• **Mathematics of computing** \rightarrow *Hypergraphs*; • **Computing methodologies** \rightarrow *Parallel algorithms*.

KEYWORDS

hypergraph decomposition, hypertree width, parallel algorithms

ACM Reference Format:

Georg Gottlob, Matthias Lanzinger, Cem Okulmus, and Reinhard Pichler. 2022. Fast Parallel Hypertree Decompositions in Logarithmic Recursion Depth. In *Proceedings of the 41st ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS '22)*, June 12–17, 2022, Philadelphia, PA, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3517804.3524153>

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 the author(s) 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.

PODS '22, June 12–17, 2022, Philadelphia, PA, USA.

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9260-0/22/06...\$15.00
<https://doi.org/10.1145/3517804.3524153>

1 INTRODUCTION

Hypertree decompositions (HDs) [20] have been demonstrated to be a valuable tool in a wide field of algorithmic applications. By way of structural decomposition of the hypergraph representation of problem instances, they induce tractable fragments for fundamental reasoning problems such as conjunctive query evaluation [20], constraint satisfaction problems [17], and related counting problems [24]. Other applications can be found in game theory, where problems such as determining Nash Equilibria [13] and combinatorial auctions [12] also become tractable in cases where HDs of bounded width exist.

In many of the listed cases, we do not only have theoretical tractability results but in fact know of algorithms which are suitable for practical applications. For example, in conjunctive query evaluation, HDs can be used for efficient reduction to an acyclic instance, which allows for linear-time solving using Yannakakis' algorithm [27]. Beyond practical algorithms, many of the listed problems are in fact known to be contained in the complexity class NC^2 [6] if a bounded width HD exists [3, 18, 19]. Importantly, problems in NC^2 are considered to be highly parallelisable [6] and thus the use of HDs in these areas can be even more attractive in parallelised and distributed scenarios. The promising theoretical properties of hypertree decompositions have also been experimentally verified. Implementations in specialised database systems have demonstrated the applicability of HDs in query evaluation by using them (and closely related variants), especially on difficult instances where current heuristic-based systems struggle [1, 10, 11].

Despite these desirable properties and a demand for worst-case guarantees in various potential fields of application, the adoption of hypertree decompositions in practice has been slow. One crucial challenge that is limiting their more widespread use is the computational difficulty of constructing good HDs. In general, finding an HD for a given hypergraph H and width at most k is NP-hard and $\text{W}[1]$ -hard when parametrised by k [20], but is tractable when k is fixed, i.e., the problem is in XP in the terminology of parametrised complexity. In fact, a significantly stronger upper bound can be given. Finding an HD of fixed width is in the complexity class LogCFL [20] (contained in NC^2) and therefore, in theory, highly parallelisable [6]. However, the theoretical parallelisability of the problem is demonstrated by construction of an appropriate Alternating Turing Machine [5] and no practical algorithm that allows for effective parallel computation of HDs is known. Here, we propose the first such algorithm.

Related Work. HD computation has received significant attention recently. This is witnessed, for instance, by the development of the large benchmark data set HyperBench [9], novel algorithmic approaches [8, 9, 22], and being the subject of a recent PACE competition [7]. Moreover, a number of new theoretical results [15, 16] have been presented, which have deepened our understanding of the problem. Still, the development of a parallel algorithm for hypertree decomposition remains a critical open question.

The two state-of-the-art approaches for computing HDs, *det-k-decomp* [23] and *HtdLEO* [25], both rely on techniques that are inherently unsuitable for parallelisation. *det-k-decomp* is heavily reliant on extensive caching and would therefore require excessive coordination between threads. In *HtdLEO*, the problem is encoded as an SMT instance and is therefore limited by the lack of parallelisation strategies for SMT solvers. While both algorithms perform well on current benchmarks, their lack of parallelisation ultimately limits them when it comes to solving large instances, i.e., finding HDs of large hypergraphs. This situation is especially disappointing as, on the one hand, single-core performance apparently does not suffice to solve larger instances and, on the other hand, the problem is in fact highly parallelisable in theory.

Interestingly, in [22], a parallel algorithm *BalancedGo* is proposed for a slightly more general problem of computing *generalised hypertree decompositions* (GHDs) [21]. In HDs, the so-called *special condition* enforces certain constraints on the parent/child nodes in the decomposition tree and the tree must therefore be treated as rooted. Crucially, this constraint is no longer enforced in GHDs and it is therefore also no longer necessary to consider the decomposition tree to be rooted. This additional degree of freedom is a key factor in the design of *BalancedGo*, where it ultimately allows for simple reassembly of individual decompositions of subproblems into a GHD of the full hypergraph. However, this freedom comes at a significant additional computational cost as the corresponding decision problem for computing GHDs is NP-hard even for constant width 2 [16, 21] (i.e., it is not even in XP in the parametrised setting). In practice, this leads to an additional exponential factor in the algorithms' complexity in contrast to the complexity of algorithms for computing HDs. In summary, current approaches either are not amenable to effective parallelisation or compute GHDs and therefore potentially cause exponential additional cost. The goal of this paper is to bridge this gap and develop a parallel algorithm for computing hypertree decompositions.

Our Contributions. As argued above, this goal is not achievable by a straightforward extension of current approaches. The two principal algorithms for HDs are inherently unsuited for parallelisation while the parallel algorithm for GHDs fundamentally relies on the fact that GHDs are unrooted. We therefore develop a new theoretical machinery which will allow us to construct HDs in an arbitrary order instead of being limited to a strict top-down or bottom-up construction of the HD. This machinery then allows us to build on some of the ideas of *BalancedGo* while avoiding the complexity of GHDs. Experimental evaluation demonstrates that the resulting algorithm combines the best of both worlds by scaling effectively with an increase of parallel threads while avoiding the exponential overhead of GHD computation. Our main contributions are as follows:

- We develop a new theoretical framework of extended hypergraphs and their balanced separation, and we show that extended hypergraphs always have a balanced separator. To actually find such balanced separators, it is crucial to apply a novel approach that determines pairs of parent and child nodes of an HD (rather than a single node) at a time.
- Based on these new results we propose a novel algorithm *log-k-decomp*, which searches for balanced separators at arbitrary positions in a potential HD. We argue that our algorithm is well-suited for parallelisation, in particular we prove a logarithmic upper bound on the recursion depth.
- We identify a number of further optimisations of our basic algorithm and we incorporate them into a parallelised reference implementation of *log-k-decomp*.
- We compare the performance of *log-k-decomp* to *det-k-decomp* and *HtdLEO* through experiments over the HyperBench benchmark [9]. We observe that *log-k-decomp* outperforms the state-of-the-art significantly. Furthermore, we experimentally verify the parallel scaling behaviour of *log-k-decomp*.

Structure. We formally introduce important concepts and notation in Section 2. The theoretical framework of extended hypergraphs and their balanced separation is established in Section 3. Building on this framework, we introduce the core ideas of the *log-k-decomp* algorithm and establish a logarithmic bound on its recursion depth in Section 4. The results of our empirical evaluation are presented in Section 5. We conclude with Section 6.

Full details for the proposed optimisations are presented in Appendix A. Further details of the algorithm and empirical evaluation are available in the full version of this paper [14].

2 PRELIMINARIES

CQs, CSPs, and hypergraphs. 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 may 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 $e \in E(H)$ with $v \in e$. 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 . By slight abuse of notation we may thus write $H' \subseteq H$ with the understanding that $E(H') \subseteq E(H)$ and, hence, implicitly also $V(H') \subseteq V(H)$. We are frequently dealing with sets of sets of vertices (e.g., sets of edges). For $S \subseteq 2^{V(H)}$, we write $\bigcup S$ as a short-hand for the union of such a set of sets, i.e., for $S = \{s_1, \dots, s_\ell\}$, we have $\bigcup S = \bigcup_{i=1}^\ell s_i$.

Conjunctive Queries (CQs), are arguably one of the most fundamental types of queries in the database world. Similarly, *Constraint Satisfaction Problems* (CSPs) are among the most fundamental formalisms in Artificial Intelligence and for modelling combinatorial problems. Formally, both are given by a first-order formula ϕ using only the connectives in $\{\exists, \wedge\}$ and disallowing $\{\forall, \vee, \neg\}$. Given such a formula ϕ , the hypergraph H_ϕ corresponding to ϕ is defined as follows: $V(H_\phi) = \text{vars}(\phi)$, i.e., the variables occurring in ϕ ; and $E(H_\phi) = \{\text{vars}(a) \mid a \text{ is an atom in } \phi\}$. In the sequel, we will only concentrate on hypergraphs, with the understanding that all results ultimately apply to CQs and CSPs.

Hypertree decompositions and hypertree width. We introduce the used notation first: given a rooted tree $T = \langle N(T), E(T) \rangle$ with node set $N(T)$ and edge set $E(T)$, we write T_u to denote the subtree of T rooted at u , where u is a node in $N(T)$. Analogously, we write T_u^\uparrow to denote the subtree of T induced by $N(T) \setminus N(T_u)$. Intuitively, T_u is the subtree of T “below” u and including u , while T_u^\uparrow is the subtree of T “above” u . By slight abuse of notation, we sometimes write $u \in T$ instead of $u \in N(T)$ to denote that u is a node in T . Below, we shall introduce node-labelling functions χ and λ , which assign to each node $u \in T$ a set of vertices or edges, respectively, from some hypergraph H , i.e., $\chi(u) \subseteq V(H)$ and $\lambda(u) \subseteq E(H)$. For a node-labelling function f with $f \in \{\chi, \lambda\}$ and a subtree T' of T , we define $f(T')$ as $f(T') = \bigcup_{u \in T'} f(u)$.

We are now ready to recall the definitions of hypertree decompositions and hypertree width from [20]: A *hypertree decomposition* (HD) \mathcal{D} of a hypergraph $H = (V(H), E(H))$ is a tuple $\mathcal{D} = \langle T, \chi, \lambda \rangle$, such that $T = \langle N(T), E(T) \rangle$ is a rooted tree, χ and λ are node-labelling functions with $\chi: N(T) \rightarrow 2^{V(H)}$ and $\lambda: N(T) \rightarrow 2^{E(H)}$ and the following conditions hold:

- (1) for each $e \in E(H)$, there exists a node $u \in N(T)$ with $e \subseteq \chi(u)$;
- (2) for each $v \in V(H)$, the set $\{u \in N(T) \mid v \in \chi(u)\}$ is connected in T ;
- (3) for each $u \in N(T)$, $\chi(u) \subseteq \bigcup \lambda(u)$;
- (4) for each $u \in N(T)$, $\chi(T_u) \cap (\bigcup \lambda(u)) \subseteq \chi(u)$.

The *width* of an HD $\mathcal{D} = \langle T, \chi, \lambda \rangle$ is the maximum size of the λ -labels over all nodes $u \in T$, i.e., $\text{width}(\mathcal{D}) = \max_{u \in T} |\lambda(u)|$. Moreover, the *hypertree width* of a hypergraph H , denoted $hw(H)$, is the minimum width over all HDs of H . Condition (2) is called the “connectedness condition” and condition (4) is referred to as the “special condition” in [20]. The set $\chi(u)$ is often referred to as the “bag” at node u and we will also call it the “ χ -label” of node u . Analogously, the set $\lambda(u)$ will be referred to as the “ λ -label” of u .

If we drop the special condition from the above definition then we get so-called generalized hypertree decompositions (GHD). The width of a GHD is again defined as the maximum size of the λ -labels over all nodes $u \in T$ and the *generalized hypertree width* of a hypergraph H , denoted $ghw(H)$, is the minimum width over all GHDs of H . The problem of checking if an HD of width $\leq k$ exists and, if so, computing a concrete HD of width $\leq k$ is known to be feasible in polynomial time for arbitrarily chosen but fixed k [20]. In contrast, for GHDs, this problem has been shown to be NP-complete even if we fix $k = 2$ [16, 21].

Throughout this paper, we will be dealing with a hypergraph H and a tree T of an HD of H . 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 T (of the decomposition).

3 EXTENDED HYPERGRAPHS AND THEIR BALANCED SEPARATION

The key idea of our algorithm is to split the task of constructing an HD into subtasks of constructing *parts* of the HD, which will be referred to as “HD-fragments” in the sequel. These HD-fragments can later be stitched together to form an HD of a given hypergraph. This splitting into HD-fragments is realised by choosing a node

u of the HD and splitting the HD into one subtree above node u and possibly several subtrees below u . In order to keep track of how to combine these subtrees later on, we introduce the notion of *special edges*. Intuitively, a special edge is the set $\chi(u)$ of vertices for some node u in the HD, and it is used to keep track of the interface between the HD-fragment “above” node u (we will denote this part of the HD as T_u^\uparrow) and the HD-fragments at subtrees below node u . Conversely, for each of the subtrees T_{u_i} rooted at the child nodes u_i of u , we have to keep track of the interface to $\chi(u)$ in the form of a set *Conn* of vertices, which is the intersection $\chi(T_{u_i}) \cap \chi(u)$.

At the heart of our decomposition algorithm in Section 4 will be a recursive function *Decomp*, which takes as input a subset E' of the edges $E(H)$, a set of special edges Sp , and a set of vertices *Conn*. The goal of *Decomp* is to construct a fragment of an HD, such that every edge $e \in E'$ is covered by some node u' in the HD-fragment (i.e., $e \subseteq \chi(u')$), all special edges are covered by some leaf node of this HD-fragment (hence, these are the interfaces to the HD-fragments “below”) and *Conn* must be fully contained in $\chi(r)$ of the root r of this HD-fragment (hence, this is the interface to the HD-fragment “above”). Formally, function *Decomp* deals with *extended subhypergraphs* of H in the following sense.

Definition 3.1 (extended subhypergraph). Let H be a hypergraph. An *extended subhypergraph* of H is a triple $\langle E', Sp, Conn \rangle$ with the following properties:

- E' is a subset of the edge set $E(H)$ of H ;
- Sp is a set of special edges, i.e., $Sp \subseteq 2^{V(H)}$;
- $Conn$ is a set of vertices, i.e., $Conn \subseteq V(H)$.

We now extend several crucial definitions introduced in [20] for hypergraphs to extended subhypergraphs.

Definition 3.2 (connectedness, components). Let H be a hypergraph, let $U \subseteq V(H)$ be a set of vertices, and let $H' = \langle E', Sp, Conn \rangle$ be an extended subhypergraph of H .

- We define $[U]$ -*adjacency* as a binary relation on $E' \cup Sp$ such that two (possibly special) edges $f_1, f_2 \in E' \cup Sp$ 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' is a maximally $[U]$ -connected subset $C \subseteq E' \cup Sp$.

Let S be a set of edges and special edges with $U = \bigcup S$. Then we will also use the terms $[S]$ -connectedness and $[S]$ -components as a short-hand for $[U]$ -connectedness and $[U]$ -components, respectively. Observe that the set *Conn* plays no role in the above definition of connectedness and components. This is in contrast to our definition of hypertree decompositions (HDs) of extended subhypergraphs, which we give next.

Definition 3.3 (hypertree decomposition). Let H be a hypergraph and let $H' = \langle E', Sp, Conn \rangle$ be an extended subhypergraph of H . A *hypertree decomposition* (HD) of H' is a tuple $\langle T, \chi, \lambda \rangle$, such that $T = \langle N(T), E(T) \rangle$ is a rooted tree, χ and λ are node-labelling functions and the following conditions hold:

- (1) for each $u \in N(T)$, either
 - a) $\lambda(u) \subseteq E(H)$ and $\chi(u) \subseteq \bigcup \lambda(u)$ or
 - b) $\lambda(u) = \{s\}$ for some $s \in Sp$ and $\chi(u) = s$;
- (2) each $f \in E' \cup Sp$ is “covered” by some $u \in N(T)$, i.e.:
 - a) if $f \in E'$, then $f \subseteq \chi(u)$;
 - b) if $f \in Sp$, then $\lambda(u) = \{f\}$ and, hence, $\chi(u) = f$;
- (3) for each $v \in (\bigcup E') \cup (\bigcup Sp)$, the set $\{u \in N(T) \mid v \in \chi(u)\}$ is connected in T ;
- (4) for each $u \in N(T)$, $\chi(T_u) \cap (\bigcup \lambda(u)) \subseteq \chi(u)$;
- (5) if $\lambda(u) = \{s\}$ for some $s \in Sp$, then u is a leaf of T ;
- (6) the root r of T satisfies $Conn \subseteq \chi(r)$.

Clearly, H can also be considered as an extended subhypergraph of itself by taking the triple $\langle E(H), \emptyset, \emptyset \rangle$. Then the HDs of the extended subhypergraph $\langle E(H), \emptyset, \emptyset \rangle$ and the HDs of hypergraph H coincide.

In [20], Definition 5.1, a normal form of HDs was introduced. Below, in Definition 3.5, we will carry the notion of normal form over to HDs of extended subhypergraphs. To this end, it is convenient to first define the set of (possibly special) edges *covered for the first time* by some node or by some subtree of an HD.

Definition 3.4. Let $H' = \langle E', Sp, Conn \rangle$ be an extended subhypergraph of some hypergraph H and let $\mathcal{D} = \langle T, \chi, \lambda \rangle$ be an HD of H' . For a node $u \in T$, we write $cov(u)$ to denote the set of edges and special edges *covered for the first time* at u , i.e.: $cov(u) = \{f \in E' \cup Sp \mid f \subseteq \chi(u) \text{ and for all ancestor nodes } u' \text{ of } u, f \not\subseteq \chi(u')\}$. For a subtree T' of T , we define $cov(T') = \bigcup_{u \in T'} cov(u)$.

Definition 3.5 (normal form). Let $H' = \langle E', Sp, Conn \rangle$ be an extended subhypergraph of some hypergraph H and let $\mathcal{D} = \langle T, \chi, \lambda \rangle$ be an HD of H' . We say that \mathcal{D} is in *normal form*, if for every node p in T and every child node c of p , the following properties hold:

- (1) There is exactly one $[\chi(p)]$ -component C_p of H' such that $C_p = cov(T_c)$;
- (2) there exists $f \in C_p$ with $f \subseteq \chi(c)$, where C_p is the $[\chi(p)]$ -component satisfying Condition 1;
- (3) $\chi(c) = (\bigcup \lambda(c)) \cap (\bigcup C_p)$, where again C_p is the $[\chi(p)]$ -component satisfying Condition 1.

By the connectedness condition, the following property holds in any HD: if C' is a $[\chi(p)]$ -component of H' with $C' \cap cov(T_c) \neq \emptyset$, then $C' \subseteq cov(T_c)$ must hold. That is, $cov(T_c)$ is the union of *one or several* $[\chi(p)]$ -components. Condition 1 of the normal form requires there to be *exactly one* $[\chi(p)]$ -component C_p of H' satisfying $C_p \subseteq cov(T_c)$.

Condition 2 intuitively requires that some “progress” must be made by the labelling of node c . Hence, in the first place, at least one vertex from $\bigcup C_p$ not already present in $\chi(p)$ must occur in $\chi(c)$. By the connectedness condition, this is only possible if one edge f from C_p occurs in $\lambda(c)$. Hence, by the special condition (i.e., condition (4) of the definition of HDs), $f \subseteq \chi(c)$ must hold.

Condition 3 is the only place where we deviate from the normal form in [20]. The purpose of Condition 3 in [20] is to make sure that $\chi(c)$ is uniquely determined whenever $\lambda(c)$, $\chi(p)$, and the $[\chi(p)]$ -component C_p from Condition 1 are known. However, there also would have been other choices to achieve this goal. Our Condition 3 chooses $\chi(c)$ *minimally*. That is, to ensure the special condition, $\chi(c)$ must contain all vertices from $\bigcup \lambda(c)$ that occur in $\chi(T_c)$. Since all edges in C_p are covered at some node in T_c , all vertices from

$(\bigcup \lambda(c)) \cap (\bigcup C_p)$ must occur in $\chi(c)$. On the other hand, there is no need to add further vertices to $\chi(c)$, since vertices not occurring in $\bigcup cov(T_c)$ can never violate the connectedness condition at node c as long as we stick to our strategy of choosing $\chi(u)$ minimally also for all nodes $u \in T_c$. In contrast, Condition 3 in [20] chooses $\chi(c)$ *maximally*. That is, also all vertices in $(\bigcup \lambda(c))$ that occur in $\chi(p)$ are added to $\chi(c)$. This deviation from the normal form in [20] is crucial since, in our construction of an HD, we will be able to derive the possible sets C_p as soon as we have $\lambda(p)$ and $\lambda(c)$ but we will “know” $\chi(p)$ only much later in the algorithm.

We now carry over two key results from [20], whose proofs can be easily adapted to our setting of extended subhypergraphs and are therefore omitted here.

Theorem 3.6 (cf. [20], Theorem 5.4). Let H' be an extended subhypergraph of some hypergraph H and let \mathcal{D} be an HD of H' of width k . Then there exists an HD \mathcal{D}' of H' in normal form, such that \mathcal{D}' also has width k .

Lemma 3.7 (cf. [20], Lemma 5.8). Let H' be an extended subhypergraph of some hypergraph H and let $\mathcal{D} = \langle T, \chi, \lambda \rangle$ be an HD in normal form of H' . Moreover, let p, c be nodes in T such that p is the parent of c and let $C_c \subseteq C_p$ for some $[\chi(p)]$ -component C_p of H' . Then the following equivalence holds: C_c is a $[\chi(c)]$ -component of H' if and only if C_c is a $[\lambda(c)]$ -component of H' .

Note that our deviation from [20] in the definition of the χ -label of nodes in a normal-form HD is inessential, since the “downward” components in an HD are not affected by adding or removing vertices from the parent node to the χ -label of the child node. However, for our purposes, we need a slightly stronger version of the above lemma: recall that the HD construction in [20] proceeds in a strict top-down fashion. Hence, when dealing with $\lambda(c)$, the bag $\chi(p)$ is already known. This is due to the fact that, initially at the root r , we have $\chi(r) = \bigcup \lambda(r)$ by the special condition. And then, whenever $\lambda(c)$ is determined and $\chi(p)$ plus a $[\chi(p)]$ -component are already known, also $\chi(c)$ can be computed. However, in our HD algorithm, which “jumps into the middle” of the HD to be constructed, we only have $\lambda(p)$ (but not $\chi(p)$) available when determining $\lambda(c)$. Hence, we need to slightly extend the above lemma to the following corollary, which follows from Lemma 3.7 by an easy induction argument over the distance from the root of the HD:

Corollary 3.8. Let H' be an extended subhypergraph of some hypergraph H and let $\mathcal{D} = \langle T, \chi, \lambda \rangle$ be an HD in normal form of H' . Moreover, let p, c be nodes in T such that p is the parent of c and let $C_c \subseteq C_p$ for some $[\lambda(p)]$ -component C_p of H' . Then the following equivalence holds: C_c is a $[\chi(c)]$ -component of H' if and only if C_c is a $[\lambda(c)]$ -component of H' .

As mentioned before, in our HD construction, we “jump into the middle” of the HD to be constructed. The motivation for this deviation from a strict top-down construction is that we want to split the work of recursively constructing fragments of the HD into pieces with a guaranteed upper bound on the size. Formally, we are thus aiming at a *balanced separator* of the HD. This concept was already studied in [2], and it was shown that HDs always have a balanced separator. In [9], balanced separators were used to design an algorithm for GHD computation. Below, we formally define

balanced separators for our notion of extended subhypergraphs and we show that in an HD, a balanced separator always exists.

Definition 3.9 (balanced separators). Let H' be an extended subhypergraph of some hypergraph H and let $\mathcal{D} = \langle T, \chi, \lambda \rangle$ be an HD of H' . A node u of T is a *balanced separator*, if the following holds:

- for every subtree T_{u_i} rooted at a child node u_i of u , we have $|cov(T_{u_i})| \leq \frac{|E'| + |Sp|}{2}$ and
- $|cov(T_u^\uparrow)| < \frac{|E'| + |Sp|}{2}$.

Intuitively, this means that none of the subtrees “below” u covers more than half of the edges of $E' \cup Sp$ and the subtree “above” u even covers less than half of the edges of $E' \cup Sp$.

Lemma 3.10. Let H' be an extended subhypergraph of some hypergraph H and let $\mathcal{D} = \langle T, \chi, \lambda \rangle$ be an HD of H' . Then there exists a balanced separator in \mathcal{D} .

PROOF OF LEMMA 3.10. We show that, given an arbitrary HD, we can always find a balanced separator as follows: Initially, we set $u = r$ for the root node r of T and distinguish two cases: if $|cov(T_{u_i})| \leq \frac{|E'| + |Sp|}{2}$ holds for every subtree T_{u_i} rooted at a child node u_i of u , then u is a balanced separator and we are done. Otherwise, there exists a child node u_i of u such that $|cov(T_{u_i})| > \frac{|E'| + |Sp|}{2}$ holds for the subtree T_{u_i} rooted at u_i . Of course, there can exist only one such child node u_i . Moreover, by $cov(T_{u_i}^\uparrow) \cap cov(T_{u_i}) = \emptyset$, we have $|cov(T_{u_i}^\uparrow)| < \frac{|E'| + |Sp|}{2}$.

Now set $u = u_i$ and repeat the case distinction: if $|cov(T_{u_i})| \leq \frac{|E'| + |Sp|}{2}$ holds for every subtree T_{u_i} rooted at a child node u_i of u , then u is a balanced separator and we are done. Otherwise, there exists a child node u_i of u such that $|cov(T_{u_i})| > \frac{|E'| + |Sp|}{2}$ holds for the subtree T_{u_i} rooted at u_i . Again, there can only be one such u_i . So we set $u = u_i$ and iterate the same considerations. This process is guaranteed to terminate since, eventually, we will reach a leaf node of T . \square

4 THE LOG-K-DECOMP ALGORITHM

We now describe the main ideas of algorithm log- k -decomp. A pseudo-code description of log- k -decomp is shown in Algorithm 1.

Algorithm log- k -decomp aims at constructing an HD in normal form according to Definition 3.5 of width $\leq k$ for a given hypergraph H and integer $k \geq 1$. The task of constructing an HD is split into subtasks that can then be processed in parallel. At the heart of log- k -decomp is the recursive function `Decomp`: it takes as input an extended subhypergraph H' of H in the form of parameter H' of H (with two fields $H'.E$ and $H'.Sp$ for the sets of edges and special edges of H' , respectively) plus parameter $Conn$ for the interface of the HD-fragment to be constructed with the parts “above” in the final HD. It returns “true” if an HD-fragment of width $\leq k$ of H' exists and “false” otherwise. The top-level calls to function `Decomp` (line 7) are from the main program of log- k -decomp which, in a loop (lines 3 – 9), searches for the λ -label of the root node r of the desired HD of H . By the special condition, we have $\chi(r) = \bigcup \lambda(r)$. Hence, the $[\lambda(r)]$ -components (computed at line 4) coincide with the $[\chi(r)]$ -components. Function `Decomp` is called (on line 7) for each of the extended subhypergraphs of H corresponding to the $[\lambda(r)]$ -components.

Algorithm 1: log- k -decomp

Type: $\text{Comp} = (E: \text{Edge set}, Sp: \text{Special Edge set})$
Input: H : Hypergraph
Parameter: k : width parameter
Output: true if hw of $H \leq k$, else false

```

1 begin
2    $H_{comp} := \text{Comp}(E: H, Sp: \emptyset)$ 
3   foreach  $\lambda_r \subseteq H$  s.t.  $1 \leq |\lambda_r| \leq k$  do ▷ RootLoop
4      $comps_r := [\lambda_r]$ -components of  $H_{comp}$ 
5     foreach  $y \in comps_r$  do
6        $Conn_y := V(y) \cap \bigcup \lambda_r$ 
7       if not(Decomp( $y, Conn_y$ )) then
8         continue RootLoop ▷ reject this root
9     return true
10  return false ▷ exhausted search space
11 function Decomp( $H': \text{Comp}, Conn: \text{Vertex set}$ )
12   if  $|H'.E| \leq k$  and  $|H'.Sp| = 0$  then ▷ Base Cases
13     return true
14   else if  $|H'.E| = 0$  and  $|H'.Sp| = 1$  then
15     return true
16   foreach  $\lambda_p \subseteq H$  s.t.  $1 \leq |\lambda_p| \leq k$  do ▷ ParentLoop
17      $comps_p := [\lambda_p]$ -components of  $H'$ 
18     if  $\exists i$  s.t.  $|comps_p[i]| > \frac{|H'|}{2}$  then
19        $comp_{down} := comps_p[i]$  ▷ found child comp.
20     else
21       continue ParentLoop
22   if  $V(comp_{down}) \cap Conn \not\subseteq \bigcup \lambda_p$  then
23     continue ParentLoop ▷ connect. check
24   foreach  $\lambda_c \subseteq H$  s.t.  $1 \leq |\lambda_c| \leq k$  do ▷ ChildLoop
25      $\chi_c := \bigcup \lambda_c \cap V(comp_{down})$ 
26     if  $V(comp_{down}) \cap \bigcup \lambda_p \not\subseteq \chi_c$  then
27       continue ChildLoop ▷ connect. check
28      $comps_c := [\chi_c]$ -components of  $comp_{down}$ 
29     if  $\exists i$  s.t.  $|comps_c[i]| > \frac{|H'|}{2}$  then
30       continue ChildLoop
31     foreach  $x \in comps_c$  do
32        $Conn_x := V(x) \cap \chi_c$ 
33       if not(Decomp( $x, Conn_x$ )) then
34         continue ChildLoop ▷ reject child
35      $comp_{up} := H' \setminus comp_{down}$  ▷ pointwise diff.
36      $comp_{up}.Sp = comp_{up}.Sp \cup \{\chi_c\}$ 
37     if not(Decomp( $comp_{up}, Conn$ )) then
38       continue ChildLoop ▷ reject child
39     return true ▷ hw of  $H' \leq k$ 
40  return false ▷ exhausted search space

```

The base case of function `Decomp` is reached (lines 12 – 15) when the existence of such an HD-fragment is trivial, i.e.: either there are at most k edges and no special edges left; or there is no edge and only one special edge left. In these cases, the desired HD-fragment simply consists of a single node whose λ -label either consists of the $\leq k$ edges or of the single special edge, respectively.

Function `Decomp` is controlled by two nested loops (lines 16 – 39 for the outer loop and lines 24 – 39 for the inner loop), which search for the λ -labels of two adjacent nodes p and c of the desired HD-fragment, such that p is the parent and c is the child. The idea of determining two nodes p and c is that, in an HD, we can determine $\chi(c)$ from $\lambda(c)$ if we know $\lambda(p)$ and the $[\lambda(p)]$ -component covered by the subtree T_c rooted at c , see Corollary 3.8 and Definition 3.5.

We want node c to be a balanced separator of the extended subhypergraph H' . By Lemma 3.10, a balanced separator is guaranteed to exist. To find a balanced separator c , we have to make sure that node c satisfies the two conditions of Definition 3.9, i.e.: (1) all of the subtrees rooted at a child of c cover at most half of the edges and special edges in H' and (2) the subtree T_c^\uparrow “above” c covers strictly fewer than half of the edges and special edges in H' . For the second condition, observe that $comp_{down}$ (chosen at line 19) is meant to be covered precisely by T_c . Note that, w.l.o.g., we are searching for an HD in normal form. This is why we may assume that T_c covers exactly one $[\lambda(p)]$ -component, namely $comp_{down}$. Further observe that the edges and special edges covered by T_c^\uparrow and the set $comp_{down}$ partition the edges and special edges in H' . Hence, checking if $comp_{down}$ contains more than half of H' (on line 18) is equivalent to checking condition (2), i.e., T_c^\uparrow covers strictly fewer than half of the edges and special edges in H' . In order to check that c also satisfies the first condition of Definition 3.9, we have to compute all $[\lambda(c)]$ -components inside $comp_{down}$ (line 28) and check that the size of each of them is at most half of the size of H' (line 29). Again, since we are only interested in HDs in normal form, we may assume here that each subtree rooted at a child of c covers exactly one of these $[\lambda(c)]$ -components.

If such a balanced separator $\lambda(c)$ together with the λ -label $\lambda(p)$ at its parent node has been found, several checks have to be performed to make sure that the HD-fragment under construction satisfies the connectedness condition. For instance, all vertices in the intersection of $Conn$ (i.e., the interface of the HD-fragment currently being constructed with the remaining HD “above” this HD-fragment) with component C_p (i.e., a component “below” node p) also have to occur in $\bigcup \lambda(p)$ (line 22).

Suppose that all these checks succeed. From $\lambda(p)$ and $\lambda(c)$, we can compute $\chi(c)$ according to Condition 3 of the normal form introduced in Definition 3.5 (line 25). In the HD \mathcal{D}' to be constructed for the extended subhypergraph H' , the edges and special edges of H' can be split into 3 disjoint categories:

- (1) the edges and special edges covered by $\chi(c)$,
- (2) the edges and special edges covered by a subtree rooted at some child node of c , and
- (3) the edges and special edges covered in the HD “above” c .

The edges and special edges covered by $\chi(c)$ are done and need no further consideration. The edges and special edges in the second

and third category are taken care of by recursive calls to the function `Decomp` (lines 33 and 37). To this end, we compute all $[\chi(c)]$ -components C_1, \dots, C_m (line 28). Now suppose that C_1, \dots, C_ℓ with $1 \leq \ell \leq m$ are the $[\chi(c)]$ -components inside the $[\lambda(p)]$ -component C_p . Then the function `Decomp` is called recursively for each of the $[\chi(c)]$ -components C_1, \dots, C_ℓ (line 33). In the call for component C_i , the interface $Conn_i$ is obtained simply as the intersection of the vertices in C_i and in $\chi(c)$ (line 32). All of the remaining $[\chi(c)]$ -components are taken care of by the HD-fragment “above” c , which we try to construct in another recursive call of function `Decomp` (line 37). In this recursive call, $\chi(c)$ is added as yet another special edge – in addition to the edges and special edges in the $[\chi(c)]$ -components outside C_p . The additional special edge in the recursive call for the HD-part “above” node c and the interfaces $Conn$ defined for each of the components as the intersection against $\chi(c)$, in the recursive calls for the HD-parts “below” node c ensure that we can (provided that all recursive calls of function `Decomp` are successful) stitch together the HD-fragments of these recursive calls to an HD-fragment of the extended subhypergraph H' of H .

To sum up, if all recursive calls return “true” then the overall result of this call to function `Decomp` is successful and returns “true” (line 39). If at least one of the recursive calls returns “false”, then we have to search for a different label $\lambda(c)$ (in the next iteration of the “ChildLoop”). If eventually all candidates for $\lambda(c)$ have been tried out and none of them was successful, then we have to search for a different label $\lambda(p)$ of the parent node p (in the next iteration of the “ParentLoop”) and restart the search for $\lambda(c)$ from scratch. Only when also all candidates for $\lambda(p)$ have been tried out and none of them was successful, then function `Decomp` returns the overall result “false” (line 40).

Below, we state the crucial property of $\log-k$ -decomp, which makes this approach particularly well-suited for a parallel implementation.

Theorem 4.1. Algorithm $\log-k$ -decomp correctly checks for given hypergraph H and integer $k \geq 1$, if $hw(H) \leq k$ holds. The algorithm is realised by a main program and the recursive function `Decomp`, whose recursion depth is bounded logarithmically in the number of edges of H , i.e., $O(\log(|H|))$.

Note that we have formulated algorithm $\log-k$ -decomp as a decision procedure that decides if $hw(H) \leq k$ holds for given H and k . In case of a successful computation (i.e. return-value true) it is easy to assemble a concrete HD of width $\leq k$ of H from the HD-fragments corresponding to the various calls of procedure `Decomp`.

We emphasize two further important properties of algorithm $\log-k$ -decomp: First, it should be noted that the logarithmic bound on the recursion depth does not restrict the form of the HD in any way. In particular, it does not imply a logarithmic bound on the depth of the HD. The bound on the recursion depth is achieved by our novel approach of constructing the HD by recursively “jumping” to a balanced separator of the HD-fragment to be constructed rather than constructing the HD in a strict top-down manner as proposed in previous approaches [20, 23].

Second, we stress that it is crucial in our approach that we search for appropriate λ -labels for a pair (p, c) of nodes, where p is the parent of c . The rationale is that we need the λ -label of the parent in order to determine $\chi(c)$ from $\lambda(c)$. And only when we know $\chi(c)$,

we can be sure, which edges are indeed covered by $\chi(c)$. This knowledge is crucial to guarantee that all of the recursive calls of function `Decomp` have to deal with an extended subhypergraph whose size is halved, which in turn guarantees the logarithmic upper bound on the recursion depth. This strategy is significantly different from all previous approaches of decomposition algorithms. In [9, 22], a parallel algorithm for generalised hypertree decompositions is presented. There, the problem of determining the χ -label of the balanced separator is solved by adding a big number of subedges to the hypergraph so that one may assume that $\chi(u) = \bigcup \lambda(u)$ holds for every node u . Clearly, this addition of subedges, in general, leads to a substantial increase of the hypergraph. In [4], a preliminary attempt to parallelise the computation of HDs was made without handling pairs of nodes. However, in the absence of $\lambda(p)$, we cannot determine $\chi(c)$ from $\lambda(c)$. Consequently, we do not know which edges covered by $\bigcup \lambda(c)$ are ultimately covered by $\chi(c)$. Hence, all the edges covered by $\bigcup \lambda(c)$ would have to be added to the recursive call of `Decomp` for the HD-part “above” c , thus destroying the balancedness and the logarithmic upper bound on the recursion depth.

By Theorem 4.1, Algorithm `log- k -decomp` guarantees a logarithmic bound on the recursion depth and thus provides a good basis for a parallel implementation. Nevertheless it still leaves room for several improvements. For instance, we can define also negative base cases to detect the overall answer “false” faster, we can restrict the edges that may possibly be used in the λ -labels of an extended subhypergraph (and provide them as an additional parameter of the function `Decomp`), etc. These ideas and several further improvements – together with the pseudo-code of the resulting improved algorithm – are presented in Appendix A.

5 IMPLEMENTATION AND EVALUATION

We report now on the empirical results obtained for our implementation of the `log- k -decomp` algorithm. Our experiments are based on the HyperBench benchmark from [9], which was already used for the evaluation of previous decomposition algorithms, notably `NewDetKDecomp` [9] (an enhanced re-implementation of `det- k -decomp` [23]) and `HtdLEO` [25].

Our goal was to determine the exact hypertree width of as many instances as possible. We compare here the performance of three different decomposition methods, namely `NewDetKDecomp` [9], `HtdLEO` [25], and our implementation of `log- k -decomp`. Note that while the tested implementations include the capability to compute GHDs or FHDs, we only consider the computation of HDs in our experiments here. Our new implementation of `log- k -decomp` is based on the open-source code of `BalancedGo` [22], a parallel algorithm for computing GHDs.

The full raw data of our experiments¹ as well as the source code of our implementation² of `log- k -decomp` are provided at the URLs below.

¹<https://zenodo.org/record/6389816>

²<https://github.com/cem-okulmus/log-k-decomp>

5.1 Benchmark Instances and Setting

For the evaluation, we use the benchmark library HyperBench [9]. It contains 3648 hypergraphs underlying CQs and CSPs from various sources in industry and the literature and is commonly used to evaluate decomposition algorithms. The instances are available at <http://hyperbench.dbai.tuwien.ac.at> and in the raw data accompanying this manuscript.

Hardware and Software. Our implementation is written in the programming language Go using version 1.14 and we will refer to it as `log- k -decomp`. We will give more details below on how it was configured for the experiments reported in Section 5.2. The hardware used for the evaluation was a cluster of 12 nodes, using Ubuntu 16.04.1 LTS, with Linux kernel 4.4.0-184-generic, GCC version 5.4.0. Each node has a 12 core Intel Xeon CPU E5-2650 v4, clocked at 2.20 GHz and using 264 GB of RAM.

Setup of Experiments. To ensure comparability of our experiments with results published in the literature, we employ the following test setup and restrictions: a timeout of one hour was used and available RAM was limited to 1 GB. We note that this corresponds to limits also used in previous experiments in this area [9, 22]. For `log- k -decomp`, each run needs two inputs: a hypergraph H and the width parameter $k \geq 1$. For these tests, we used width parameters in the range [1, 10]. When running tests for `HtdLEO`, we used different memory limits. Namely, we allowed `HtdLEO` to use up to 24 GB of RAM since SMT solving is significantly more memory intensive than the other two algorithms. Note that the other two algorithms have very low memory requirements and are not constrained in any way by the 1 GB limit and the respective experiments are therefore still comparable with `HtdLEO`. Furthermore, `HtdLEO` needs no width parameter since it directly tries to find an optimal solution.

We used the HTCondor system [26] to facilitate the tests, limits to memory and number of cores accessed by running test instances.

Throughout this section we will be interested in two key metrics. First, the number of *solved* instances, by which we mean instances for which an optimal (i.e., minimal width) hypertree decomposition was found and proven optimal. Second, the computation time that was necessary to compute the optimal width decomposition, which we will refer to as the *running time* or simply *runtime*. Importantly, this means that average running times are taken only over the instances that the respective algorithm is able to solve, while timed out instances are not considered in the running time calculation.

5.2 Empirical Evaluation

We report here on the main results of our experiments. We note that a number of additional experiments can be found in the full version of this paper [14], providing a variety of further details and insights. Our implementation of `log- k -decomp` also employs the following *hybridisation strategy*: as will be seen below, `NewDetKDecomp` performs very well on small hypergraphs but has difficulties with even slightly larger instances. In contrast, a particular strength of our new `log- k -decomp` algorithm is to quickly split a big hypergraph into significantly smaller extended subhypergraphs. To combine the best of both worlds, we use `log- k -decomp` to split the original HD computation problem until the subproblems become small, at which point we apply our own implementation of `det- k -decomp` (extended to handle extended subhypergraphs correctly) to the

Table 1: Comparison of prior methods and log- k -decomp: number of cases solved and runtimes (sec.) to find optimal-width HDs.

			Hypertree Decomposition Methods											
Origin of Instances	Size of Instances	Instances in Group	NewDetKDecomp [9]				HtdLEO [25]				log- k -decomp Hybrid			
			#solved	avg	max	stdev	#solved	avg	max	stdev	#solved	avg	max	stdev
Application	$75 < E \leq 100$	405	97	21.4	3296.0	192.8	65	809.5	3156.6	735.2	261	86.5	3555.8	332.4
	$50 < E \leq 75$	514	276	10.6	1906.0	104.7	448	250.0	3281.5	409.3	469	0.5	78.5	3.6
	$10 < E \leq 50$	369	253	0.0	0.0	0.0	237	60.1	1017.9	150.3	253	0.0	0.1	0.0
	$ E \leq 10$	915	906	0.0	0.0	0.0	876	56.6	1427.1	155.0	915	0.0	0.0	0.0
Synthetic	$ E > 100$	66	18	0.2	7.0	1.0	13	734.0	2507.1	711.7	34	46.9	2528.2	209.6
	$75 < E \leq 100$	422	87	77.2	3467.0	379.3	312	1045.2	3591.1	1287.0	235	48.9	2495.6	210.9
	$50 < E \leq 75$	215	38	18.8	1593.0	141.9	212	101.7	2560.1	246.1	215	4.1	476.3	32.7
	$10 < E \leq 50$	647	290	56.0	3240.0	336.3	303	412.2	3597.4	850.2	625	18.8	3526.3	174.7
	$ E \leq 10$	95	95	0.0	0.0	0.0	78	28.8	218.5	41.5	95	0.0	0.0	0.0
Total	-	3648	2060	20.6	3467.0	194.2	2544	280.2	3597.4	676.7	3102	30.5	3555.8	197.8

small subproblems. For details and an experimental evaluation of different parametrisations for our hybridisation strategy, we refer to the full version of this paper [14].

We compare the aforementioned hybrid version of the log- k -decomp algorithm with the two state-of-the-art implementations for finding HDs: NewDetKDecomp [9] and HtdLEO [25].

Our results are summarised in Table 1, distinguishing the hypergraphs in the HyperBench benchmark by size and origin³. We distinguish between two main categories, hypergraphs that are derived from applications and hypergraphs that were synthetically generated. In each group we report our results split by the number of edges $|E|$ in the instance. Note that the group $|E| > 100$ of instances with more than 100 edges is empty for the Application case and thus omitted from the table. *Instances in Group* reports the number of instances in each such group. For each algorithm and each group of instances, we list the number of solved instances (*#solved*) and statistics over the running times (*avg*, *max*, *stdev*). Times are all in seconds and rounded to a single digit after the comma. Results over all groups are given in the last row titled “Total”.

As mentioned above, some care is required when comparing times between algorithms. While NewDetKDecomp has low average time overall, this is partly due to solving fewer instances. The data therefore demonstrates that, in general, NewDetKDecomp either solves an instance quickly or fails to find an optimal width decomposition before timing out. Overall, we see that despite solving significantly more instances than its competitors, running times for log- k -decomp overall are comparable with NewDetKDecomp and noticeably lower than for HtdLEO.

It may be of further interest how these numbers compare to the performance of state of the art algorithms for finding generalised hypertree decompositions. The results reported for BalancedGo [22] (on a comparable system) show that the best method there solves

only 1730 instances optimally without timeout. In contrast log- k -decomp manages to solve 2491 of the instances tested there optimally⁴. Furthermore, in none of the cases where BalancedGo finds the optimal ghw is it lower than the optimal hw . In other words, in practice, the additional complexity of GHDs compared with HDs is not compensated by achieving lower width (even if, in theory, no better upper bound on the hw than $hw \leq 3 \cdot ghw + 1$ is known [2]).

In our experiments, we also observe that for low widths – i.e., cases where using HDs is most promising in practice – log- k -decomp is very close to solving all instances. In particular, of the 3224 instances with width at most 6, log- k -decomp solves 2930 (92%) instances. In contrast, NewDetKDecomp and HtdLEO time out on 1206 and 766, respectively, of those instances. This suggests that log- k -decomp can be a solid foundation for the integration of HDs in practice going forward. If we look at instances of $hw \leq 5$ the situation improves even further, with log- k -decomp solving 2450 out of 2482 (98.7%) instances; compared to 80% and 86% solved by NewDetKDecomp and HtdLEO, respectively.

The experiments reported in Table 1 were performed over the full set of HyperBench instances. However, for the additional experiments reported in this section, it is more meaningful to restrict our experiments to exclude hypergraphs that are, roughly speaking, too small or have high width. Small instances benefit only marginally from algorithmic improvements or parallelism, while very high width is of less algorithmic interest as it exponentially effects algorithms that make use of decompositions. Hence, we propose to exclude such instances to make more relevant observations. We therefore focus on instances with more than 50 edges and vertices that are known to have hypertree width at most 6. There are 465 instances in HyperBench which satisfy these conditions; we will refer to them as HB_{large} .

We performed a second set of experiments over the instances in HB_{large} to verify our claims that log- k -decomp is well-suited for parallelisation. For $1 \leq n \leq 5$, we observe the time taken to find and verify the optimal width of an instance using n CPU cores. We report on the times to find these optimum widths averaged

³HyperBench instances are often categorised more fine-grained in terms of their origin (cf., [9]). For our experiments we have found the direct effect of hypergraph size to be more informative and therefore report our results in this way instead.

⁴The evaluation in [22] considers only a subset of HyperBench with 3071 instances




Figure 1: Study of \log - k -decomp scaling behaviour w.r.t. the number of processing cores used.

over all instances in HB_{large} in Figure 1. To avoid a decreasing number of timeouts from skewing the data we report the average only over instances that do not timeout for any n for a given algorithm. For reference, we also report the (single core) performance of NewDetKDecomp for the same setting.

We observe approximately linear speedups up to 4 cores, from about 189 seconds on 1 core to 50 seconds for 4 cores for \log - k -decomp. This behaviour is expected since our parallelisation strategy relies on dividing up the search space for bounded separators uniformly over the available cores. Since this requires no communication between threads or other overhead that depends on the degree of parallelisation, the key task of searching for balanced separators scales linearly in the number of cores. In instances where the search for separators dominates the running time, such as negative instances where the full search space is explored, analysis of our algorithm therefore predicts effectively linear scaling of performance. In the data from Figure 1, we observe diminishing returns in the rate of improvement of average running time starting from 5 cores. However, preliminary experiments on additional different systems do not confirm this behaviour and there \log - k -decomp exhibits linear scaling up to much higher core counts. Further in-depth experimentation is therefore required to obtain a clearer picture for the scaling behaviour for a high number of cores.

Very similar scaling can be observed for our Hybrid version. Note that the reported times for the Hybrid algorithm are slightly higher only due to solving more (harder) instances.

6 CONCLUSION

In this paper we introduced a novel algorithm \log - k -decomp for computing hypertree decompositions. Based on new theoretical insights and results on HDs, we were able to propose an algorithm that constructs decompositions in arbitrary order (rather than, e.g., in a strict top-down manner) while achieving a *balanced* separation into subproblems. In this way, we have obtained a logarithmic bound on the recursion depth of our algorithm, making it particularly well suited for parallelisation. We evaluated an implementation of \log - k -decomp through experimental comparison with the state of the art. On the standard benchmark for hypertree decomposition,

we are able to achieve clear improvements both in the number of solved instances and in the time required to solve them.

In combination, our theoretical results and experiments demonstrate that \log - k -decomp achieves our goal of effective parallel HD computation. We believe that the performance improvements, especially on large hypergraphs lay a strong foundation for more widespread adoption of hypertree decompositions in practice, e.g., for complex query execution in high-performance database applications.

With HD computation for large and complex hypergraphs becoming practically feasible, one of the key challenges that block the use of HDs is quickly becoming less problematic. We therefore consider full integration of hypertree decompositions into existing database systems and constraint solvers to be a natural next step in this line of research.

Experiments suggest that there is significant potential in the study of metrics for hybrid approaches. In particular, how can we decide effectively when to switch from the balanced separation of \log - k -decomp to the greedy heuristic guided method underlying \det - k -decomp. This motivates a more in-depth study of hybridisation metrics in the future.

ACKNOWLEDGEMENTS

This work was supported by the Austrian Science Fund (FWF) project P30930-N35. 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). Matthias Lanzinger acknowledges support by the Royal Society project "RAISON DATA" (Project reference: RP\R1\201074).

REFERENCES

- [1] C. R. Aberger, A. Lamb, S. Tu, A. Nötzli, K. Olukotun, and C. Ré. Emptyheaded: A relational engine for graph processing. *ACM Trans. Database Syst.*, 42(4):20:1–20:44, 2017.
- [2] I. Adler, G. Gottlob, and M. Grohe. Hypertree width and related hypergraph invariants. *Eur. J. Comb.*, 28(8):2167–2181, 2007.
- [3] F. N. Afrati, M. R. Joglekar, C. Ré, S. Salihoglu, and J. D. Ullman. GYM: A multi-round distributed join algorithm. In M. Benedikt and G. Orsi, editors, *20th International Conference on Database Theory, ICDT 2017, March 21–24, 2017, Venice, Italy*, volume 68 of *LIPICs*, pages 4:1–4:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.
- [4] D. Akatov. *Exploiting parallelism in decomposition methods for constraint satisfaction*. PhD thesis, University of Oxford, UK, 2010.
- [5] A. K. Chandra and L. J. Stockmeyer. Alternation. In *17th Annual Symposium on Foundations of Computer Science, Houston, Texas, USA, 25–27 October 1976*, pages 98–108. IEEE Computer Society, 1976.
- [6] S. A. Cook. A taxonomy of problems with fast parallel algorithms. *Inf. Control.*, 64(1-3):2–21, 1985.
- [7] M. A. Dzulfikar, J. K. Fichte, and M. Hecher. The PACE 2019 parameterized algorithms and computational experiments challenge: The fourth iteration (invited paper). In *IPEC*, volume 148 of *LIPICs*, pages 25:1–25:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [8] J. K. Fichte, M. Hecher, N. Lodha, and S. Szeider. An SMT approach to fractional hypertree width. In J. N. Hooker, editor, *Principles and Practice of Constraint Programming - 24th International Conference, CP 2018, Lille, France, August 27–31, 2018, Proceedings*, volume 11008 of *Lecture Notes in Computer Science*, pages 109–127. Springer, 2018.
- [9] W. Fischl, G. Gottlob, D. M. Longo, and R. Pichler. HyperBench: A benchmark and tool for hypergraphs and empirical findings. *ACM Journal of Experimental Algorithmics*, 26:1.6:1–1.6:40, 2021.
- [10] L. Ghionna, L. Granata, G. Greco, and F. Scarcello. Hypertree decompositions for query optimization. In *ICDE*, pages 36–45. IEEE Computer Society, 2007.
- [11] L. Ghionna, G. Greco, and F. Scarcello. H-DB: a hybrid quantitative-structural sql optimizer. In C. Macdonald, I. Ounis, and I. Ruthven, editors, *Proceedings of*

- the 20th ACM Conference on Information and Knowledge Management, CIKM 2011, Glasgow, United Kingdom, October 24-28, 2011, pages 2573–2576. ACM, 2011.
- [12] G. Gottlob and G. Greco. Decomposing combinatorial auctions and set packing problems. *J. ACM*, 60(4):24:1–24:39, 2013.
 - [13] G. Gottlob, G. Greco, and F. Scarcello. Pure nash equilibria: Hard and easy games. *J. Artif. Intell. Res.*, 24:357–406, 2005.
 - [14] G. Gottlob, M. Lanzinger, C. Okulmus, and R. Pichler. Fast parallel hypertree decompositions in logarithmic recursion depth. *CoRR*, abs/2104.13793, 2021.
 - [15] G. Gottlob, M. Lanzinger, R. Pichler, and I. Razgon. Fractional covers of hypergraphs with bounded multi-intersection. In *MFCS*, volume 170 of *LIPIcs*, pages 41:1–41:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
 - [16] G. Gottlob, M. Lanzinger, R. Pichler, and I. Razgon. Complexity analysis of generalized and fractional hypertree decompositions (forthcoming). *J. ACM*, 68(1), 2021.
 - [17] G. Gottlob, N. Leone, and F. Scarcello. A comparison of structural CSP decomposition methods. *Artif. Intell.*, 124(2):243–282, 2000.
 - [18] G. Gottlob, N. Leone, and F. Scarcello. The complexity of acyclic conjunctive queries. *J. ACM*, 48(3):431–498, 2001.
 - [19] G. Gottlob, N. Leone, and F. Scarcello. Computing LOGCFL certificates. *Theor. Comput. Sci.*, 270(1-2):761–777, 2002.
 - [20] G. Gottlob, N. Leone, and F. Scarcello. Hypertree decompositions and tractable queries. *J. Comput. Syst. Sci.*, 64(3):579–627, 2002.
 - [21] G. Gottlob, Z. Miklós, and T. Schwentick. Generalized hypertree decompositions: Np-hardness and tractable variants. *J. ACM*, 56(6):30:1–30:32, 2009.
 - [22] G. Gottlob, C. Okulmus, and R. Pichler. Fast and parallel decomposition of constraint satisfaction problems. In C. Bessière, editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020 [scheduled for July 2020, Yokohama, Japan, postponed due to the Corona pandemic]*, pages 1155–1162. ijcai.org, 2020.
 - [23] G. Gottlob and M. Samer. A backtracking-based algorithm for hypertree decomposition. *ACM Journal of Experimental Algorithmics*, 13:1:1–1:19, 2008.
 - [24] R. Pichler and S. Skritek. Tractable counting of the answers to conjunctive queries. *J. Comput. Syst. Sci.*, 79(6):984–1001, 2013.
 - [25] A. Schidler and S. Szeider. Computing optimal hypertree decompositions with SAT. In Z. Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 1418–1424. ijcai.org, 2021.
 - [26] D. Thain, T. Tannenbaum, and M. Livny. Distributed computing in practice: the condor experience. *Concurrency - Practice and Experience*, 17(2-4):323–356, 2005.
 - [27] M. Yannakakis. Algorithms for acyclic database schemes. In *Proceedings of the 7th International Conference on Very Large Databases, VLDB 1981, Cannes*, pages 82–94. VLDB, 1981.

A FURTHER COMBINATORIAL OBSERVATIONS AND OPTIMISATIONS

As was shown in Theorem 4.1, algorithm $\log\text{-}k\text{-decomp}$ introduced in Section 4 reaches the primary goal of splitting the HD-construction into subtasks with guaranteed upper bound on their size. In theory, this is enough to support parallelism. However, this basic algorithm still leaves a lot of room for further improvements. In this section, we present several optimisations, which are crucial to achieve good performance in practice. The line numbers below refer to Algorithm 1. However, in Algorithm 2, we will ultimately also give the pseudo-code for the enhanced algorithm where all the optimisations mentioned below are included.

Extension of the base case. The recursive function `Decomp` starts (on lines 12 – 15) with some simple checks that immediately give a “true” answer. In contrast, a “false” answer is only obtained in case of unsuccessful execution of the entire procedure. We could add the following negative case to the top of the procedure: if $H'.E = \emptyset$, then $|H'.Sp| \leq 1$ must hold. The rationale of this condition is that, if there are no more edges in $H'.E$, then we would have to use only “old” edges (i.e., edges covered already at some node further up in the HD) in the λ -label to separate the remaining special edges. However, a λ -label consisting of “old” edges only is not allowed, since this would violate the second condition of the normal form in Definition 3.5 (i.e., “some progress has to be made”).

Root of the HD-fragment. In the current form of procedure `Decomp`, we always “guess” a pair (p, c) of nodes, such that p is the parent of c . This also covers the case that c is the root node of the HD-fragment for the current extended subhypergraph. In this case, the parent node p would actually be the node immediately above this HD-fragment (in other words, p was the node from which the current call of `Decomp` happened). However, it would be more efficient to consider the case of “guessing” the root node of this HD-fragment explicitly. More precisely, we would thus first check for the label λ_p guessed in `Decomp` on line 16 (which, in the current version of the algorithm, is automatically treated as the “parent”) if all $[\lambda_p]$ -components have at most half the size of the current extended subhypergraph.

If this is the case, then we may use this node as the root of the HD-fragment to cover the current extended subhypergraph. This makes sense since it corresponds precisely to the “search” for a balanced separator in the proof of Lemma 3.10. That is, if the root of the HD gives rise to components which are all at most half the size, then the root is the desired balanced separator.

If this is not the case, then we simply proceed with procedure `Decomp` in its present form, i.e.: there exists exactly one $[\lambda_p]$ -component whose size is bigger than half. So we take the guessed node as the parent and search for a balanced separator as a child of p in the direction of this oversized $[\lambda_p]$ -component.

Allowed edges. The main task of procedure `Decomp` is to compute labels (i.e., edge sets) λ_p and λ_c of nodes p, c , which will ultimately be in a parent-child relationship in the HD. For these labels, Algorithm 1 imposes no restriction. That is, in principle, we would try all possible sets of $\leq k$ edges for these labels. However, not all edges actually make sense. We should thus add one more parameter to procedure `Decomp` indicating the edges that are allowed in a λ -label of the HD-fragment for this extended subhypergraph.

More specifically, in our search for the λ -label of some node u , we may exclude from the HD of the extended subhypergraph comp_{up} (i.e., in the recursive call of function `Decomp` on line 37) all edges which are part of some component “below” u . The rationale of this restriction is that, by the special condition, using a “new” edge in a λ -label forces us to add all its vertices to the χ -label, i.e.: it is fully covered in such a node. But then it cannot be part of a component whose edges are covered for the first time further down in the tree.

Note that we can yet further restrict the search for the label λ_c by requiring that at least one edge must be from $H'.E$, since choosing only “old” edges would violate the second condition of the normal form. As far as the label λ_p is concerned, the same kind of restriction can be applied if we first implement the previous optimisation of handling the root node of the current HD-fragment separately. If we indeed have to guess the labels λ_p and λ_c of two nodes p and c (i.e., the label λ_p guessed first was not a balanced separator), then both nodes p and c are *inside* the current HD-fragment. Hence, also the label λ_p must contain at least one “new” edge.

No special treatment of the root of the HD. In the current form of the algorithm, we start in the main program by “guessing” the label λ_r of the root of the HD on line 3 and then branch into calls of procedure `Decomp` for each $[\lambda_r]$ -component. Of course, there is no guarantee that this λ -label is a balanced separator. Consequently,

there is no guarantee that the size of *all* HD-fragments to be constructed in these calls of procedure `Decomp` is significantly smaller than the entire HD.

In order to start with a balanced separator right from the beginning, we may instead call procedure `Decomp` straight away with parameters $H'.E = H$, $H'.Sp = \emptyset$, and $Conn = \emptyset$. We thus treat the search for the very first λ -label in the desired HD in exactly the same way as for any other HD-fragment. As far as the above mentioned optimisation of restricting the allowed edges is concerned, of course all edges of H would be initially allowed.

Searching for child nodes first. In Algorithm 1, we first look for λ -labels of potential parent nodes, and consider afterwards the λ -labels of potential child nodes. Only then do we check if the χ -label of the child is a balanced separator of the current subcomponent. We have observed that in many hypergraphs of HyperBench, balanced separators are rare, in the sense that only a small part of the search space will ever fulfil the properties required. Therefore we should first look for a potential child s.t. its λ -label is a balanced separator, and only afterwards try to find a fitting parent. While this may seem slightly unintuitive, it allows us to quickly detect cases where no balanced separator can be found at all.

Note that we can determine the precise bag χ_c for a child c only when we know the λ -label of its parent. Nevertheless, even if we only have λ_c , we can over-approximate the χ_c -label as $\bigcup \lambda_c$. Hence, if $\bigcup \lambda_c$ is not a balanced separator, then we may clearly conclude that neither is χ_c .

Finally, note that by searching for the child node first, we get the above described optimisation of treating the “Root of the HD-fragment” separately almost for free. Indeed, when computing λ_c , we can immediately check if $Conn \subseteq \bigcup \lambda_c$ holds. Recall that $Conn$ constitutes the interface to the HD-fragment *above* the current one. Hence, if $\bigcup \lambda_c$ fully covers this interface, c is in fact the root node of the current HD-fragment.

Speeding up the search for parent λ -labels. The previous optimisation means that, after having found a λ -label λ_c for the child which is a balanced separator of the current subcomponent, we need to find a *suitable* λ -label of the parent. By “suitable” we mean that we may limit ourselves to edges which have a non-empty intersection with $\bigcup \lambda_c$. A very high-level explanation why we may exclude edges e with $e \cap \bigcup \lambda_c = \emptyset$ from the search space of λ_p is that the control flow of function `Decomp` is mainly determined by the edges and special edges covered by $\bigcup \lambda_c$ and the $[\lambda_c]$ -components *below* c . By the connectedness condition, if e is covered *above* c and has empty intersection with $\bigcup \lambda_c$, then excluding or including e in λ_p has no effect on the $[\lambda_c]$ -components *below* c . In our experimental evaluation, we found that this restriction indeed significantly reduces the time it takes to either find a *suitable* λ_p , or detect that no such λ -label exists. Of course, this restriction of the search space cannot destroy soundness. We will show below that also the completeness of the algorithm is preserved.

Theorem A.1. The optimised \log - k -decomp algorithm for checking if a hypergraph H has $hw(H) \leq k$ given in Algorithm 2 is sound and complete. More specifically, for given hypergraph H and integer $k \geq 1$, the algorithm returns “true” if and only if there exists an HD of H of width $\leq k$. Moreover, by materialising the decompositions implicitly constructed in the recursive calls of the `Decomp` function,

an HD of H of width $\leq k$ can be constructed in polynomial time in case of a successful computation (i.e., return-value “true”).

PROOF. The soundness and completeness of Algorithm 2 follow almost immediately from the soundness and completeness of Algorithm 1 together with the above explanations of the various optimisations. Likewise, the polynomial-time upper bound on the time needed to construct an HD in case of a successful computation can again be easily shown as part of the soundness proof. The only non-trivial part is that the last optimisation (i.e., the restriction of the search space for $\lambda(p)$) does not destroy the completeness of the algorithm. The remainder of the proof will concentrate on this aspect.

Assume that hypergraph H has an HD of width $\leq k$. Then the optimised \log - k -decomp algorithm without the restriction on the search space for label λ_p (on line 22) returns the overall result *true*. This is due to the fact that, as was argued in Section A, the other optimisations mentioned there do not affect the completeness of the algorithm. Now consider a recursive call of function `Decomp` and suppose that it returns *true* if the restriction on the search space for label λ_p is dropped. Of course, if the value *true* is returned in one of the base cases (lines 6 or 8) or if λ_c turns out to be the λ -label of the root node of the current HD-fragment (and *true* is returned on line 21), then the restriction of the search space for λ_p has no effect at all. Hence, the only interesting case to consider is that the parent loop (lines 22 – 43) is indeed executed.

Let λ_p be the λ -label chosen on line 22 if no restriction is imposed on the search space. We claim that we may remove from λ_p all edges that have an empty intersection with $\bigcup \lambda_c$ without altering the control flow of this particular execution of function `Decomp`. Actually, it suffices to show that we may remove *one* edge e with an empty intersection with $\bigcup \lambda_c$ from λ_p without altering the control flow of this particular execution of function `Decomp`. Then the claim follows by an easy induction argument.

So suppose that λ_p contains at least one edge e such that $e \cap \bigcup \lambda_c = \emptyset$ and let $\lambda'_p = \lambda_p \setminus \{e\}$. An inspection of the code of the parent loop reveals that it suffices to show that this elimination of edge e from λ_p leaves $comp_{down}$ unchanged. Indeed, if $comp_{down}$ is still a $[\lambda'_p]$ -component, say the i -th $[\lambda'_p]$ -component, then the if-condition on line 24 is true. Of course, there can be only one $[\lambda'_p]$ -component satisfying the condition $comps_p[i] > \frac{|H|}{2}$. Hence, on line 25, for this particular i , exactly the same value is assigned to $comp_{down}$ for λ'_p as for λ_p . But then also χ_c on line 28 gets the same value as without the restriction on the search space of λ_p . Consequently, also the $[\chi_c]$ -components computed on line 33 and the parameters supplied to the recursive calls of function `Decomp` (on lines 36 and 41) remain the same as without the restriction on the search space. Hence, function `Decomp` will ultimately return the value *true* also if we choose λ'_p on line 22.

It remains to show that λ_p and λ'_p indeed give rise to the same component $comp_{down}$. To avoid confusion, let us write $comp_{down}$ to denote a $[\lambda_p]$ -component and $comp'_{down}$ to denote a $[\lambda'_p]$ -component. Let $comp_{down}$ be the unique $[\lambda_p]$ -component that satisfies the condition $comps_p[i] > \frac{|H|}{2}$ on line 24. We have $\lambda'_p \subseteq \lambda_p$. Decreasing a set can only increase the corresponding components. Hence, there exists a $[\lambda'_p]$ -component, call it $comp'_{down}$

Algorithm 2: Optimised log- k -decomp

Type: Comp=(E : Edge set, Sp : Special Edge set)**Input:** H : Hypergraph**Parameter:** k : width parameter**Output:** true if hw of $H \leq k$, else false

```
1 begin
2    $H_{comp} := \text{Comp}(E: H, Sp: \emptyset)$ 
3   return Decom( $H_{comp}, \emptyset, H$ ) ▷ initial call
4 function Decom( $H'$ : Comp, Conn: Vertex set,  $A$ : Edge set)
5   if  $|H'E| \leq k$  and  $|H'Sp| = 0$  then ▷ Base Cases
6     return true
7   else if  $|H'E| = 0$  and  $|H'Sp| = 1$  then
8     return true
9   else if  $|H'E| = 0$  and  $|H'Sp| > 1$  then
10    return false
11  foreach  $\lambda_c \subseteq A$  s.t.  $\lambda_c \cap H'E \neq \emptyset$ 
12    and  $1 \leq |\lambda_c| \leq k$  do ▷ ChildLoop
13     $comps_c := [\lambda_c]$ -components of  $H'$ 
14    if  $\exists i$  s.t.  $|comps_c[i]| > \frac{|H'|}{2}$  then
15      continue ChildLoop
16    else if  $Conn \subseteq \bigcup \lambda_c$  then ▷ check if  $\lambda_c$  is root
17       $\chi_c := \bigcup \lambda_c \cap V(H')$ 
18      foreach  $y \in comps_c$  do
19         $Conn_y := V(y) \cap \chi_c$ 
20        if not(Decomp( $y, Conn_y, A$ )) then
21          continue ChildLoop
22    return true ▷  $c$  is root of  $H'$ 
23  foreach  $\lambda_p \subseteq A$  s.t.  $\lambda_p \cap H'E \neq \emptyset$ 
24    and  $1 \leq |\lambda_p| \leq k$  do ▷ ParentLoop
25     $comps_p := [\lambda_p]$ -components of  $H'$ 
26    if  $\exists i$  s.t.  $|comps_p[i]| > \frac{|H'|}{2}$  then
27       $comp_{down} := comps_p[i]$  ▷ found child comp.
28    else
29      continue ParentLoop
30     $\chi_c := \bigcup \lambda_c \cap V(comp_{down})$ 
31    if  $V(comp_{down}) \cap Conn \not\subseteq \bigcup \lambda_p$  then
32      continue ParentLoop ▷ connect. check
33    if  $V(comp_{down}) \cap \bigcup \lambda_p \not\subseteq \chi_c$  then
34      continue ParentLoop ▷ connect. check
35     $new\_comps_c := [\chi_c]$ -components of  $comp_{down}$ 
36    foreach  $x \in new\_comps_c$  do
37       $Conn_x := V(x) \cap \chi_c$ 
38      if not(Decomp( $x, Conn_x, A$ )) then
39        continue ParentLoop ▷ reject parent
40     $comp_{up} := H' \setminus comp_{down}$  ▷ pointwise diff.
41     $comp_{up}.Sp = comp_{up}.Sp \cup \{\chi_c\}$ 
42     $A_{up} := A \setminus comp_{down}.E$  ▷ reducing  $A$ 
43    if not(Decomp( $comp_{up}, Conn, A_{up}$ )) then
44      continue ParentLoop ▷ reject parent
45    return true ▷ hw of  $H' \leq k$ 
46 return false ▷ exhausted search space
```

with $comp_{down} \subseteq comp'_{down}$. We have to show that $comp_{down} = comp'_{down}$ holds.

The set $comp_{down}$ consists of the edges and special edges of the $[\chi_c]$ -components contained in $comp_{down}$ (denoted as $comps_c$ in the algorithm), and the edges and special edges covered by χ_c . Let us refer to these $[\chi_c]$ -components as C_1, \dots, C_ℓ . By Corollary 3.8, these $[\chi(c)]$ -components are at the same time the $[\lambda_c]$ components contained in $comp_{down}$. And the edges and special edges covered by χ_c are of course also covered by $\bigcup \lambda_c$. Likewise, $comp'_{down}$ consists of the (special) edges of the $[\lambda_c]$ components contained in $comp'_{down}$ plus the (special) edges covered by λ_c .

By $comp_{down} \subseteq comp'_{down}$, all $[\lambda_c]$ -components C_1, \dots, C_ℓ contained in $comp_{down}$ are of course also contained in $comp'_{down}$. We have to show that there is no further $[\lambda_c]$ -component contained in $comp'_{down}$. Assume to the contrary that there exists a $[\lambda_c]$ -component C' in $comp'_{down}$ such that C' is not in $comp_{down}$. By definition, $comp_{down}$ is $[\lambda_p]$ connected while $comp'_{down}$ is $[\lambda'_p]$ connected. Hence, there exist (possibly special) edges $f' \in C'$ and $f \in C_i$ for some $i \in \{1, \dots, \ell\}$, such that there is a path π (represented as a sequence of edges) with $\pi = (f_0, f_1, \dots, f_m)$, such that $f = f_0$, $f' = f_m$, and $(f_\alpha \cap f_{\alpha+1}) \setminus \bigcup \lambda'_p \neq \emptyset$ for every $\alpha \in \{0, \dots, m-1\}$. W.l.o.g., choose f , f' , and π such that m is minimal. Since f and f' are not $[\lambda_p]$ -connected, there exists α with $f_\alpha \cap f_{\alpha+1} \cap e \neq \emptyset$ while $(f_\alpha \cap f_{\alpha+1}) \setminus \bigcup \lambda_p = \emptyset$.

Since all (special) edges in $comp'_{down}$ are either in some $[\lambda_c]$ -component contained in $comp'_{down}$ or covered by $\bigcup \lambda_c$, and since we are assuming that π is of minimal length, the path π starts with f in some $[\lambda_c]$ -component C_i , possibly goes through $\bigcup \lambda_c$ and ends with f' in component C' . Recall that e was chosen such that $e \cap \bigcup \lambda_c = \emptyset$. Hence, the edges f_α and $f_{\alpha+1}$ cannot be covered by $\bigcup \lambda_c$. By our assumption that π has minimal length, we can also exclude the case that both f_α and $f_{\alpha+1}$ are in C' . Hence, at least one of f_α and $f_{\alpha+1}$ must be in C_i . In other words, $e \cap C_i \neq \emptyset$. Hence, also $e \cap comp_{down} \neq \emptyset$. However, by the check on line 31 in Algorithm 2, we know that $V(comp_{down}) \cap \bigcup \lambda_p \subseteq \bigcup \lambda_c$. This contradicts the assumption that $e \in \lambda_p$ and $e \cap \bigcup \lambda_c = \emptyset$. \square