

Multilevel Hypergraph Partitioning with Vertex Weights Revisited

Tobias Heuer  

Karlsruhe Institute of Technology, Germany

Nikolai Maas 

Karlsruhe Institute of Technology, Germany

Sebastian Schlag  

Karlsruhe Institute of Technology, Germany

Abstract

The balanced hypergraph partitioning problem (HGP) is to partition the vertex set of a hypergraph into k disjoint blocks of bounded weight, while minimizing an objective function defined on the hyperedges. Whereas real-world applications often use vertex and edge weights to accurately model the underlying problem, the HGP research community commonly works with unweighted instances.

In this paper, we argue that, in the presence of vertex weights, current balance constraint definitions either yield infeasible partitioning problems or allow unnecessarily large imbalances and propose a new definition that overcomes these problems. We show that state-of-the-art hypergraph partitioners often struggle considerably with weighted instances and tight balance constraints (even with our new balance definition). Thus, we present a recursive-bipartitioning technique that is able to reliably compute balanced (and hence feasible) solutions. The proposed method balances the partition by pre-assigning a small subset of the heaviest vertices to the two blocks of each bipartition (using an algorithm originally developed for the job scheduling problem) and optimizes the actual partitioning objective on the remaining vertices. We integrate our algorithm into the multilevel hypergraph partitioner **KaHyPar** and show that our approach is able to compute balanced partitions of high quality on a diverse set of benchmark instances.

2012 ACM Subject Classification Mathematics of computing → Hypergraphs; Mathematics of computing → Graph algorithms

Keywords and phrases multilevel hypergraph partitioning, balanced partitioning, vertex weights

Digital Object Identifier 10.4230/LIPIcs.SEA.2021.8

Supplementary Material *Software (Source Code)*: <https://github.com/kahypar/kahypar>
archived at `swh:1:dir:c66f21225656426f1f902c5dc86404a622964a04`

Dataset (Benchmark Set & Experimental Results): <http://algo2.iti.kit.edu/heuer/sea21/>

1 Introduction

Hypergraphs are a generalization of graphs where each hyperedge can connect more than two vertices. The k -way hypergraph partitioning problem (HGP) asks for a partition of the vertex set into k disjoint blocks, while minimizing an objective function defined on the hyperedges. Additionally, a balance constraint requires that the weight of each block is smaller than or equal to a predefined upper bound (most often $L_k := (1 + \varepsilon) \lceil \frac{c(V)}{k} \rceil$ for some parameter ε , where $c(V)$ is the sum of all vertex weights). The hypergraph partitioning problem is NP-hard [32] and it is even NP-hard to find good approximations [8]. The most commonly used heuristic to solve HGP in practice is the multilevel paradigm [1, 11, 29] which consists of three phases: First, the hypergraph is *coarsened* to obtain a hierarchy of smaller hypergraphs. After an *initial partitioning* algorithm is applied to the smallest hypergraph, *coarsening* is undone, and, at each level, *refinement* algorithms are used to improve the quality of the solution.



© Tobias Heuer, Nikolai Maas, and Sebastian Schlag;
licensed under Creative Commons License CC-BY 4.0

19th International Symposium on Experimental Algorithms (SEA 2021).

Editors: David Coudert and Emanuele Natale; Article No. 8; pp. 8:1–8:20

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The two most prominent application areas of HGP are very large scale integration (VLSI) design [3, 29] and parallel computation of the sparse matrix-vector product [11]. In the former, HGP is used to divide a circuit into two or more blocks such that the number of external wires interconnecting circuit elements in different blocks is minimized. In this setting, each vertex is associated with a weight equal to the area of the respective circuit element [2] and tightly-balanced partitions minimize the total area required by the physical circuit [18]. In the latter, HGP is used to optimize the communication volume for parallel computations of sparse matrix-vector products [11]. In the simplest hypergraph model, vertices correspond to rows and hyperedges to columns of the matrix (or vice versa) and a partition of the hypergraphs yields an assignment of matrix entries to processors [11]. The work of a processor (which can be measured in terms of the number of non-zero entries [7]) is integrated into the model by assigning each vertex a weight equal to its degree [11]. Tightly-balanced partitions hence ensure that the work is distributed evenly among the processors.

Despite the importance of weighted instances for real-world applications, the HGP research community mainly uses unweighted hypergraphs in experimental evaluations [38]. The main rationale hereby being that even unweighted instances become weighted implicitly due to vertex contractions during the coarsening phase. Many partitioners therefore incorporate techniques that prevent the formation of heavy vertices [13, 24, 27] during coarsening to facilitate finding a feasible solution during the initial partitioning phase [38]. However, in practice, many weighted hypergraphs derived from real-world applications already contain heavy vertices – rendering the mitigation strategies of today’s multilevel hypergraph partitioners ineffective. The popular ISPD98 VLSI benchmark set [2], for example, includes instances in which vertices can weigh up to 10% of the total weight of the hypergraph.

Contributions and Outline. After introducing basic notation in Section 2 and presenting related work in Section 3, we first formulate an alternative balance constraint definition in Section 4 that overcomes some drawbacks of existing definitions in presence of vertex weights. In Section 5, we then present an algorithm that enables partitioners based on the recursive bipartitioning (RB) paradigm to reliably compute balanced partitions for weighted hypergraphs. Our approach is based on the observation that usually only a small subset of the heaviest vertices is *critical* to satisfy the balance constraint. We show that pre-assigning these vertices to the two blocks of each bipartition (i.e., treating them as *fixed vertices*) and optimizing the actual objective function on the remaining vertices yields provable balance guarantees for the resulting k -way partition. We implemented our algorithms in the open source HGP framework **KaHyPar** [38]. The experimental evaluation presented in Section 6 shows that our new approach (called **KaHyPar-BP**) is able to compute balanced partitions for all instances of a large real-world benchmark set (without increasing the running time or decreasing the solution quality), while other partitioners such as the latest versions of **KaHyPar**, **hMetis**, and **PaToH** produced imbalanced partitions on 4.9% up to 42% of the instances for $\varepsilon = 0.01$ (4.3% up to 23.1% for $\varepsilon = 0.03$). Section 7 concludes the paper.

2 Preliminaries

A *weighted hypergraph* $H = (V, E, c, \omega)$ is defined as a set of vertices V and a set of hyperedges/nets E with vertex weights $c : V \rightarrow \mathbb{R}_{>0}$ and net weights $\omega : E \rightarrow \mathbb{R}_{>0}$, where each net e is a subset of the vertex set V (i.e., $e \subseteq V$). We extend c and ω to sets in the natural way, i.e., $c(U) := \sum_{v \in U} c(v)$ and $\omega(F) := \sum_{e \in F} \omega(e)$. Given a subset $V' \subseteq V$, the *subhypergraph* $H_{V'}$ is defined as $H_{V'} := (V', \{e \cap V' \mid e \in E : e \cap V' \neq \emptyset\}, c, \omega)$.

A k -way partition of a hypergraph H is a partition of the vertex set V into k non-empty disjoint subsets $\Pi_k = \{V_1, \dots, V_k\}$. We refer to a k -way partition $\Psi_k = \{P_1, \dots, P_k\}$ of a subset $P \subseteq V$ as a k -way prepacking. We call a vertex $v \in P$ a *fixed* vertex and a vertex $v \in V \setminus P$ an *ordinary* vertex. During partitioning, fixed vertices are not allowed to be moved to a different block of the partition. A k -way partition Π_k is ε -balanced if each block V_i satisfies the *balance constraint*: $c(V_i) \leq L_k := (1 + \varepsilon) \lceil \frac{c(V)}{k} \rceil$ for some parameter ε . The k -way hypergraph partitioning problem initialized with a k -way prepacking $\Psi_k = \{P_1, \dots, P_k\}$ is to find an ε -balanced k -way partition $\Pi_k = \{V_1, \dots, V_k\}$ of a hypergraph H that minimizes an objective function and satisfies that $\forall i \in \{1, \dots, k\} : P_i \subseteq V_i$. In this paper, we optimize the *connectivity* metric $(\lambda - 1)(\Pi) := \sum_{e \in E} (\lambda(e) - 1) \omega(e)$, where $\lambda(e) := |\{V_i \in \Pi \mid V_i \cap e \neq \emptyset\}|$.

The *most balanced partition problem* is to find a k -way partition Π_k of a weighted hypergraph $H = (V, E, c, \omega)$ such that $\max(\Pi_k) := \max_{V' \in \Pi_k} c(V')$ is minimized. For an optimal solution Π_{OPT} it holds that there exists no other k -way partition Π'_k with $\max(\Pi'_k) < \max(\Pi_{\text{OPT}})$. We use $\text{OPT}(H, k) := \max(\Pi_{\text{OPT}})$ to denote the weight of the heaviest block of an optimal solution. Note that the problem is equivalent to the most common version of the *job scheduling* problem: Given a sequence $J = \langle j_1, \dots, j_n \rangle$ of n computing jobs each associated with a *processing time* p_i for $i \in [1, n]$, the task is to find an assignment of the n jobs to k identical machines (each job j_i runs exclusively on a machine for exactly p_i time units) such that the latest completion time of a job is minimized.

3 Related Work

In the following, we will focus on work closely related to our main contributions. For an extensive overview on hypergraph partitioning we refer the reader to existing literature [3, 5, 35, 38]. Well-known multilevel HGP software packages with certain distinguishing characteristics include PaToH [4, 11] (originating from scientific computing), hMetis [29, 30] (originating from VLSI design), KaHyPar [26, 27] (general purpose, n -level), Moondrian [41] (sparse matrix partitioning), UMPa [14] (multi-objective) and Zoltan [16] (distributed partitioner).

Partitioning with Vertex Weights. The most widely used techniques to improve the quality of a k -way partition are move-based local search heuristics [19, 31] that greedily move vertices according to a *gain* value (i.e., the improvement in the objective function). Vertex moves violating the balance constraint are usually rejected, which can significantly deteriorate solution quality in presence of varying vertex weights [10]. This issue is addressed using techniques that allow intermediate balance violations [18] or use temporary relaxations of the balance constraint [9, 10]. Caldwell et al. [10] proposed to preassign each vertex with a weight greater than the average block weight L_k to a separate block before partitioning (treated as fixed vertices) and build the actual k -way partition around them. All of these techniques were developed and evaluated for flat (i.e., non-multilevel) partitioning algorithms. In the multilevel setting, even unweighted instances become implicitly weighted due to vertex contractions in the coarsening phase, which is why the formation of heavy vertices is prevented by penalizing the contraction of vertices with large weights [13, 24, 40] or enforcing a strict upper bound for vertex weights throughout the coarsening process [1, 27]. If the input hypergraph is unweighted, the aforementioned techniques often suffice to find a feasible solution [38]. PaToH [12] additionally uses bin packing techniques during initial partitioning.

Job Scheduling Problem. The job scheduling problem is NP-hard [20] and we refer the reader to existing literature [23, 36] for a comprehensive overview of the research topic. In this work, we make use of the *longest processing time* (LPT) algorithm proposed by Graham [22].

We will explain the algorithm in the context of the most balanced partition problem defined in Section 2: For a weighted hypergraph $H = (V, E, c, \omega)$, the algorithm iterates over the vertices of V sorted in decreasing vertex-weight order and assigns each vertex to the block of the k -way partition with the lowest weight. The algorithm can be implemented to run in $\mathcal{O}(|V| \log |V|)$ time, and for a k -way partition Π_k produced by the algorithm it holds that $\max(\Pi_k) \leq (\frac{4}{3} - \frac{1}{3k}) \text{OPT}(H, k)$.

KaHyPar. The **K**arlsruhe **H**ypergraph **P**artitioning framework takes the multilevel paradigm to its extreme by only contracting a single vertex in every level of the hierarchy. **KaHyPar** provides recursive bipartitioning [37] as well as direct k -way partitioning algorithms [1] (direct k -way uses RB in the initial partitioning phase). It uses a community detection algorithm as preprocessing step to restrict contractions to densely connected regions of the hypergraph during coarsening [27]. Furthermore, it employs a portfolio of bipartitioning algorithms for initial partitioning of the coarsest hypergraph [25, 37], and, during the refinement phase, improves the partition with a highly engineered variant of the classical FM local search [1] and a refinement technique based on network flows [21, 26].

During RB-based partitioning, **KaHyPar** ensures that the solution is balanced by adapting the imbalance ratio for each bipartition individually. Let $H_{V'}$ be the subhypergraph of the current bipartition that should be partitioned recursively into $k' \leq k$ blocks. Then,

$$\varepsilon' := \left((1 + \varepsilon) \frac{c(V)}{k} \cdot \frac{k'}{c(V')} \right)^{\frac{1}{\lceil \log_2(k') \rceil}} - 1 \quad (1)$$

is the imbalance ratio used for the bipartition of $H_{V'}$. The equation is based on the observation that the worst-case block weight of the resulting k' -way partition of $H_{V'}$ obtained via RB is smaller than $(1 + \varepsilon')^{\lceil \log_2(k') \rceil} \frac{c(V')}{k'}$, if ε' is used for all further bipartitions. Requiring that this weight must be smaller or equal to $L_k = (1 + \varepsilon) \lceil \frac{c(V)}{k} \rceil$ leads to the formula defined in Equation 1.

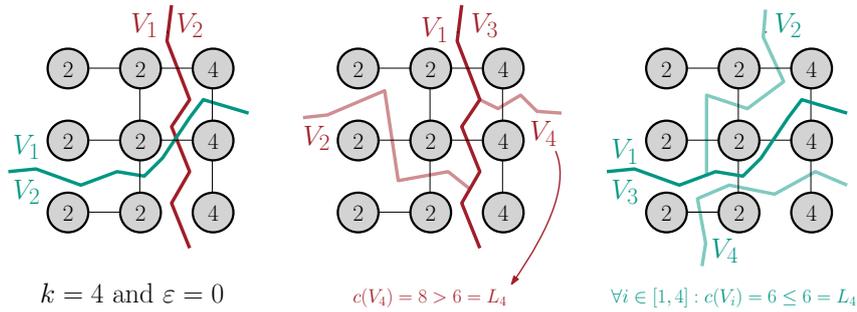
4 A New Balance Constraint For Weighted Hypergraphs

A k -way partition of a weighted hypergraph $H = (V, E, c, \omega)$ is balanced, if the weight of each block is below some predefined upper bound. In the literature, the most commonly used bounds are $L_k := (1 + \varepsilon) \lceil \frac{c(V)}{k} \rceil$ (standard definition) and $L_k^{\max} := L_k + \max_{v \in V} c(v)$ [19, 38, 39]. The latter was initially proposed by Fiduccia and Mattheyses [19] for bipartitioning to ensure that the highest-gain vertex can always be moved to the opposite block.

Both definitions exhibit shortcomings in the presence of heavy vertices: As soon as the hypergraph contains even a single vertex with $c(v) > L_k$, no feasible solution exists when the block weights are constrained by L_k , while for L_k^{\max} it follows that $L_k^{\max} > 2L_k$ – allowing large variations in block weights even if ε is small. In the following, we therefore propose a new balance constraint that (i) guarantees the existence of an ε -balanced k -way partition and (ii) avoids unnecessarily large imbalances.

While the optimal solution of the most balanced partition problem would yield a partition with the best possible balance, it is not feasible in practice to use $L_k^{\text{OPT}} := (1 + \varepsilon) \text{OPT}(H, k)$ as balance constraint, because finding such a k -way partition is NP-hard [20]. Hence, we propose to use the bound provided by the LPT algorithm instead:

$$L_k^{\text{LPT}} := (1 + \varepsilon) \text{LPT}(H, k) \leq \left(\frac{4}{3} - \frac{1}{3k} \right) L_k^{\text{OPT}}. \quad (2)$$



■ **Figure 1** Illustration of a deeply (left, green line) and a non-deeply balanced bipartition (left, red line). The numbers in each circle denotes the vertex weights. In both cases, the hypergraph is partitioned into $k = 4$ blocks with $\varepsilon = 0$ via recursive bipartitioning. Thus, the weight of heaviest block must be smaller or equal to $L_4 = 6$ and for the first bipartition, we use $L_2 = 12$ as an upper bound.

Note that if the hypergraph is unweighted, the LPT algorithm will always find an optimal solution with $\text{OPT}(H, k) = \lceil \frac{|V|}{k} \rceil$ and thus, L_k^{LPT} is equal to L_k . Since all of today's partitioning algorithms bound the maximum block weight by L_k , Section 6 gives more details on how we employ this new balance constraint definition in our experimental evaluation.

5 Multilevel Recursive Bipartitioning with Vertex Weights Revisited

Most multilevel hypergraph partitioners either employ recursive bipartitioning directly [11, 16, 29, 37, 41] or use RB-based algorithms in the initial partitioning phase to compute an initial k -way partition of the coarsest hypergraph [1, 4, 14, 30]. In both settings, a k -way partition is derived by first computing a bipartition $\Pi_2 = \{V_1, V_2\}$ of the (input/coarse) hypergraph H and then recursing on the subhypergraphs H_{V_1} and H_{V_2} by partitioning V_1 into $\lceil \frac{k}{2} \rceil$ and V_2 into $\lfloor \frac{k}{2} \rfloor$ blocks. Although KaHyPar adaptively adjusts the allowed imbalance at each bipartitioning step (using the imbalance factor ε' as defined in Equation 1), an *unfortunate* distribution of the vertices in some bipartitions Π_2 can easily lead to instances for which it is impossible to find a balanced solution during the recursive calls – even though the current bipartition Π_2 satisfies the adjusted balance constraint. An example is shown in Figure 1 (left): Although the current bipartition (indicated by the red line) is perfectly balanced, it will not be possible to recursively partition the subhypergraph induced by the vertices of V_2 into two blocks of equal weight, because each of the three vertices has a weight of four.

To capture this problem, we introduce the notion of *deep balance*:

► **Definition 1.** (Deep Balance). *Let $H = (V, E, c, \omega)$ be a weighted hypergraph for which we want to compute an ε -balanced k -way partition, and let $H_{V'}$ be a subhypergraph of H which should be partitioned into $k' \leq k$ blocks via recursive bipartitioning. A subhypergraph $H_{V'}$ is deeply balanced w.r.t. k' , if there exists a k' -way partition $\Pi_{k'}$ of $H_{V'}$ such that $\max(\Pi_{k'}) \leq L_k := (1 + \varepsilon) \lceil \frac{c(V)}{k} \rceil$. A bipartition $\Pi_2 = \{V_1, V_2\}$ of $H_{V'}$ is deeply balanced w.r.t. k' , if the subhypergraphs H_{V_1} and H_{V_2} are deeply balanced with respect to $\lceil \frac{k'}{2} \rceil$ resp. $\lfloor \frac{k'}{2} \rfloor$.*

If a subhypergraph $H_{V'}$ is deeply balanced with respect to k' , there always exists a k' -way partition $\Pi_{k'}$ of $H_{V'}$ such that weight of the heaviest block satisfies the original balance constraint L_k imposed on the partition of the input hypergraph H . Moreover, there also always exists a deeply balanced bipartition $\Pi_2 := \{V_1, V_2\}$ (V_1 is the union of the first $\lceil \frac{k'}{2} \rceil$

8:6 Multilevel Hypergraph Partitioning with Vertex Weights Revisited

and V_2 of the last $\lfloor \frac{k'}{2} \rfloor$ blocks of $\Pi_{k'}$). Hence, a RB-based partitioning algorithm that is able to compute deeply balanced bipartitions on deeply balanced subhypergraphs will always compute ε -balanced k -way partitions (assuming the input hypergraph is deeply balanced).

Deep Balance and Adaptive Imbalance Adjustments. Computing deeply balanced bipartitions in the RB setting guarantees that the resulting k -way partition is ε -balanced. Thus, the concept of deep balance could replace the adaptive imbalance factor ε' employed in KaHyPar [37] (see Equation 1). However, as we will see in the following example, combining both approaches gives the partitioner more flexibility (in terms of feasible vertex moves during refinement). Assume that we want to compute a 4-way partition via recursive bipartitioning and that the first bipartition $\Pi_2 := \{V_1, V_2\}$ is deeply balanced with $c(V_1) = (1 + \varepsilon) \lceil \frac{c(V)}{2} \rceil$. The deep-balance property ensures that we can further partition V_1 into two blocks such that the weight of the heavier block is smaller than L_4 . However, this bipartition has to be perfectly balanced:

$$L_2 = (1 + \bar{\varepsilon}) \lceil \frac{c(V_1)}{2} \rceil = (1 + \bar{\varepsilon}) \lceil \frac{(1 + \varepsilon) \lceil \frac{c(V)}{2} \rceil}{2} \rceil \leq (1 + \varepsilon) \lceil \frac{c(V)}{4} \rceil = L_4 \Rightarrow \bar{\varepsilon} \approx 0. \quad (3)$$

If we would have computed the first bipartition with an adjusted imbalance factor ε' , then $\max(\Pi_2) \leq (1 + \varepsilon') \lceil \frac{c(V)}{2} \rceil = \sqrt{1 + \varepsilon'} \lceil \frac{c(V)}{2} \rceil$ – providing more flexibility for subsequent bipartitions. In the following, we therefore focus on computing deeply ε' -balanced bipartitions.

Deep Balance and Multilevel Recursive Bipartitioning. In general, computing a deeply balanced bipartition $\Pi_2 := \{V_1, V_2\}$ w.r.t. k is NP-hard, as we must show that there exists a k -way partition Π_k of H with $\max(\Pi_k) \leq L_k$, which can be reduced to the most balanced partition problem presented in Section 2. However, we can first compute a k -way partition $\Pi_k := \{V'_1, \dots, V'_k\}$ using the LPT algorithm, thereby approximating an optimal solution. If $\max(\Pi_k) \leq L_k$, we can then construct a deeply balanced bipartition $\Pi_2 = \{V_1, V_2\}$ by choosing $V_1 := V'_1 \cup \dots \cup V'_{\lfloor \frac{k}{2} \rfloor}$ and $V_2 := V'_{\lfloor \frac{k}{2} \rfloor + 1} \cup \dots \cup V'_k$. Unfortunately, this approach completely ignores the optimization of the objective function – yielding balanced partitions of low quality. If such a bipartition were to be used as initial solution in the multilevel setting, the objective could still be optimized during the refinement phase. However, this would necessitate that refinement algorithms are aware of the concept of deep balance and that they only perform vertex moves that don't destroy the deep-balance property of the starting solution. Since this is infeasible in practice, we propose a different approach that involves fixed vertices.

The key idea of our approach is to compute a prepacking $\Psi = \{P_1, P_2\}$ of the $m = |P_1| + |P_2|$ heaviest vertices of the hypergraph and to show that this prepacking suffices to ensure that each ε' -balanced bipartition $\Pi_2 = \{V_1, V_2\}$ with $P_1 \subseteq V_1$ and $P_2 \subseteq V_2$ is deeply balanced. Note that the upcoming definitions and theorems are formulated from the perspective of the first bipartition of the input hypergraph H to simplify notation. They can be generalized to subhypergraphs $H_{V'}$ in a similar fashion as was done in Definition 1. Furthermore, we say that the bipartition $\Pi_2 = \{V_1, V_2\}$ respects a prepacking $\Psi = \{P_1, P_2\}$, if $P_1 \subseteq V_1$ and $P_2 \subseteq V_2$, and that the bipartition is balanced, if $\max(\Pi_2) \leq L_2 := (1 + \varepsilon') \lceil \frac{c(V_1 \cup V_2)}{2} \rceil$ (with ε' as defined in Equation 1). The following definition formalizes our idea.

► **Definition 2.** (Sufficiently Balanced Prepacking). *Let $H = (V, E, c, \omega)$ be a hypergraph for which we want to compute an ε -balanced k -way partition via recursive bipartitioning. We call a prepacking Ψ of H sufficiently balanced if every balanced bipartition Π_2 respecting Ψ is deeply balanced with respect to k .*

Our approach to compute ε -balanced k -way partitions is outlined in Algorithm 1. We first compute a bipartition Π_2 . Before recursing on each of the two induced subhypergraphs, we check if Π_2 is deeply balanced using the LPT algorithm in a similar fashion as described in the beginning of this paragraph. If it is not deeply balanced, we compute a sufficiently balanced prepacking Ψ and re-compute Π_2 – treating the vertices of the prepacking as fixed vertices. If this second bipartitioning call was able to compute a balanced bipartition, we found a deeply balanced partition and proceed to partition the subhypergraphs recursively.

Note that, in general, we may not detect that Π_2 is deeply balanced or fail to find a sufficiently balanced prepacking Ψ or a balanced bipartition Π_2 , since all involved problems are NP-hard. However, as we will see in Section 6, Algorithm 1 computes balanced partitions for all instances of our large real-world benchmark set. This seems to indicate that the above-mentioned problems only happen rarely in practice.

■ **Algorithm 1** Recursive Bipartitioning Algorithm.

Data: Hypergraph H for which we seek an ε -balanced k -way partition and subhypergraph $H_{V'}$ of H which is to be bipartitioned recursively into $k' \leq k$ blocks.

```

1 Function recursiveBipartitioning( $H, k, \varepsilon, H_{V'}, k'$ ):
2    $L_2 \leftarrow (1 + \varepsilon') \lceil \frac{c(V')}{2} \rceil$  // with  $\varepsilon'$  as defined in Equation 1
3    $\Pi_2 := \{V_1, V_2\} \leftarrow \text{multilevelBipartitioning}(H_{V'}, L_2, \emptyset)$  //  $\emptyset = \text{empty prepacking}$ 
4   if  $k' = 2$  then return  $\Pi_2$ 
5   else if  $\Pi_2$  is not deeply balanced w.r.t.  $k'$  then
6      $\Psi \leftarrow \text{sufficientlyBalancedPrepacking}(H, k, \varepsilon, H_{V'}, k')$  // see Algorithm 2
7      $\Pi_2 \leftarrow \text{multilevelBipartitioning}(H_{V'}, L_2, \Psi)$  // treating  $\Psi$  as fixed vertices
8    $\Pi_{k_1} \leftarrow \text{recursiveBipartitioning}(H, k, \varepsilon, H_{V_1}, k_1)$  with  $k_1 := \lceil \frac{k'}{2} \rceil$ 
9    $\Pi_{k_2} \leftarrow \text{recursiveBipartitioning}(H, k, \varepsilon, H_{V_2}, k_2)$  with  $k_2 := \lfloor \frac{k'}{2} \rfloor$ 
10  return  $\Pi_{k_1} \cup \Pi_{k_2}$ 

```

Computing a Sufficiently Balanced Prepacking. The prepacking Ψ is constructed by incrementally assigning vertices to Ψ in decreasing order of weight and checking a property \mathcal{P} after each assignment that, if satisfied, implies that the current prepacking is sufficiently balanced. In the proof of property \mathcal{P} , we will extend a k -way prepacking Ψ_k to an ε -balanced k -way partition Π_k using the LPT algorithm and use the following upper bound on the weight of the heaviest block of Π_k .

► **Lemma 3.** (LPT Bound). *Let $H = (V, E, c, \omega)$ be a weighted hypergraph, Ψ_k be a k -way prepacking for a set of fixed vertices $P \subseteq V$, and let $O := \langle v_1, \dots, v_m \mid v_i \in V \setminus P \rangle$ be the sequence of all ordinary vertices of $V \setminus P$ sorted in decreasing order of weight. If we assign the remaining vertices O to the blocks of Ψ_k by using the LPT algorithm, we can extend Ψ_k to a k -way partition Π_k of H such that the weight of the heaviest block is bound by:*

$$\max(\Pi_k) \leq \max\left\{\frac{1}{k}c(P) + h_k(O), \max(\Psi_k)\right\}, \text{ with } h_k(O) := \max_{i \in \{1, \dots, m\}} c(v_i) + \frac{1}{k} \sum_{j=1}^{i-1} c(v_j).$$

The proof of Lemma 3 can be found in Appendix A. O is sorted in decreasing order of weight because for any permutation O' of O , it holds that $h_k(O) \leq h_k(O')$ – resulting in the tightest bound for $\max(\Pi_k)$.

Assuming that the number k of blocks is even (i.e., $k_1 = k_2 = k/2$) to simplify notation, the balance property \mathcal{P} is defined as follows (the generalized version can be found in Appendix B):

► **Definition 4.** (Balance Property \mathcal{P}). Let $H = (V, E, c, \omega)$ be a hypergraph for which we want to compute an ε -balanced k -way partition and let Ψ be a prepacking of H for a set of fixed vertices $P \subseteq V$. Furthermore, let $O_t := \langle v_1, \dots, v_t \rangle$ be the sequence of the t heaviest ordinary vertices of $V \setminus P$ sorted in decreasing order of weight such that t is the smallest number that satisfies $\max(\Psi) + c(O_t) \geq L_2$ (see Line 2, Algorithm 1). We say that a prepacking Ψ satisfies the balance property \mathcal{P} if the following two conditions hold:

- (i) the prepacking Ψ is deeply balanced
- (ii) $\frac{1}{k/2} \max(\Psi) + h_{k/2}(O_t) \leq L_k$.

In the following, we will show that the LPT algorithm can be used to construct a $k/2$ -way partition $\Pi_{k/2}$ for both blocks of any balanced bipartition $\Pi_2 = \{V_1, V_2\}$ that respects Ψ , such that the weight of the heaviest block can be bound by the left term of Condition (ii). This implies that $\max(\Pi_{k/2}) \leq L_k$ (right term of Condition (ii)) and thus proves that any balanced bipartition Π_2 respecting Ψ is deeply balanced. Note that choosing t as the smallest number that satisfies $\max(\Psi) + c(O_t) \geq L_2$ minimizes the left term of Condition (ii) (since $h_k(O_t) \leq h_k(O_{t+1})$).

► **Theorem 5.** A prepacking Ψ of a hypergraph $H = (V, E, c, \omega)$ that satisfies the balance property \mathcal{P} is sufficiently balanced with respect to k .

Proof. For convenience, we use $k' := k/2$. Let $\Pi_2 = \{V_1, V_2\}$ be an arbitrary balanced bipartition that respects the prepacking $\Psi = \{P_1, P_2\}$ with $\max(\Pi_2) \leq L_2$. Since Ψ is deeply balanced (see Definition 4(i)), there exists a k' -way prepacking $\Psi_{k'}$ of P_1 such that $\max(\Psi_{k'}) \leq L_k$. We define the sequence of the ordinary vertices of block V_1 sorted in decreasing weight order with $O_1 := \langle v_1, \dots, v_m \mid v_i \in V_1 \setminus P_1 \rangle$. We can extend $\Psi_{k'}$ to a k' -way partition $\Pi_{k'}$ of V_1 by assigning the vertices of O_1 to the blocks in $\Psi_{k'}$ using the LPT algorithm. Lemma 3 then establishes an upper bound on the weight of the heaviest block.

$$\begin{aligned} \max(\Pi_{k'}) &\stackrel{\text{Lemma 3}}{\leq} \max\left\{\frac{1}{k'}c(P_1) + h_{k'}(O_1), \max(\Psi_{k'})\right\} \\ &\leq^{\max(\Psi_{k'}) \leq L_k} \max\left\{\frac{1}{k'}c(P_1) + h_{k'}(O_1), L_k\right\} \end{aligned}$$

Let O_t be the sequence of the t heaviest ordinary vertices of $V \setminus P$ with $P := P_1 \cup P_2$ as defined in Definition 4.

▷ **Claim 6.** It holds that: $\frac{1}{k'}c(P_1) + h_{k'}(O_1) \leq \frac{1}{k'} \max(\Psi) + h_{k'}(O_t)$.

For a proof of Claim 6 see Appendix C. We can conclude that

$$\frac{1}{k'}c(P_1) + h_{k'}(O_1) \stackrel{\text{Claim 6}}{\leq} \frac{1}{k'} \max(\Psi) + h_{k'}(O_t) \stackrel{\text{Definition 4(ii)}}{\leq} L_k.$$

This proves that the subhypergraph H_{V_1} is deeply balanced. The proof for block V_2 can be done analogously, which then implies that Π_2 is deeply balanced. Since Π_2 is an arbitrary balanced bipartition respecting Ψ , it follows that Ψ is sufficiently balanced. ◀

Algorithm 2 outlines our approach to efficiently compute a sufficiently balanced prepacking Ψ . In Line 6, we compute a k' -way prepacking $\Psi_{k'}$ of the i heaviest vertices with the LPT algorithm and if $\Psi_{k'}$ satisfies $\max(\Psi_{k'}) \leq L_k$, then Line 7 constructs a deeply balanced prepacking Ψ (which fulfills Condition (i) of Definition 4). We store the blocks P'_j of $\Psi_{k'}$ together with their weights $c(P'_j)$ as key in an addressable priority queue such that we can

determine and update the block with the smallest weight in time $\mathcal{O}(\log k')$ (Line 6). In Line 9, we compute the smallest t that satisfies $\max(\Psi) + c(O_t) \geq L_2$ via a binary search in logarithmic time over an array containing the vertex weight prefix sums of the sequence O , which can be precomputed in linear time. Furthermore, we construct a range maximum query data structure over the array $H_{k'/2} = \langle c(v_1), c(v_2) + \frac{1}{k'/2}c(v_1), \dots, c(v_n) + \frac{1}{k'/2} \sum_{j=1}^{n-1} c(v_j) \rangle$. Calculating $h_{k'/2}(O_t)$ (Line 10) then corresponds to a range maximum query in the interval $[i+1, i+t]$ in $H_{k'/2}$, which can be answered in constant time after $H_{k'/2}$ has been precomputed in time $\mathcal{O}(n)$ [6]. In total, the running time of the algorithm is $\mathcal{O}(n(\log k' + \log n))$. Note that if the algorithm reaches Line 12, we could not prove that any of the intermediate constructed prepackings were sufficiently balanced, in which case Ψ represents a bipartition of $H_{V'}$ computed by the LPT algorithm.

■ **Algorithm 2** Prepacking Algorithm.

Data: Hypergraph $H = (V, E, c, \omega)$ for which we seek an ε -balanced k -way partition and subhypergraph $H_{V'} = (V', E', c, \omega)$ of H which is to be bipartitioned recursively into $k' \leq k$ blocks.

```

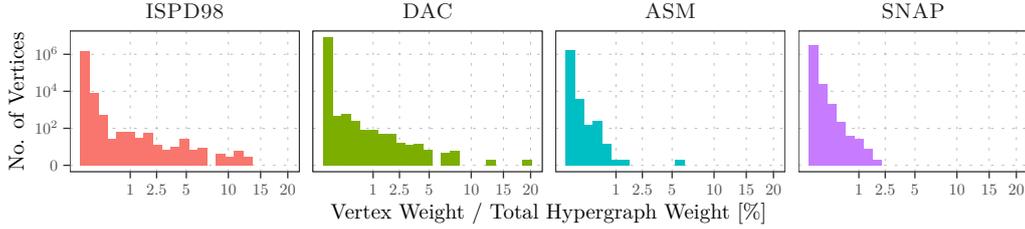
1 Function sufficientlyBalancedPrepacking( $H, k, \varepsilon H_{V'}, k'$ ):
2    $\Psi = \langle P_1, P_2 \rangle \leftarrow \langle \emptyset, \emptyset \rangle$  and  $\Psi_{k'} = \langle P'_1, \dots, P'_{k'} \rangle \leftarrow \langle \emptyset, \dots, \emptyset \rangle$  // Initialization
3    $L_2 \leftarrow (1 + \varepsilon') \lceil \frac{c(V')}{2} \rceil$  and  $L_k \leftarrow (1 + \varepsilon) \lceil \frac{c(V)}{k} \rceil$  // with  $\varepsilon'$  as defined in Equation 1
4    $O \leftarrow \langle v_1, \dots, v_n \mid v_i \in V' \rangle$  //  $V'$  sorted in decreasing order of weight  $\Rightarrow \mathcal{O}(n \log n)$ 
5   for  $i = 1$  to  $n$  do
6     Add  $v_i \in O$  to bin  $P'_j \in \Psi_{k'}$  with smallest weight // LPT algorithm
7      $\Psi \leftarrow \{P'_1 \cup \dots \cup P'_x, P'_{x+1} \cup \dots \cup P'_{k'}\}$  with  $x := \lceil \frac{k'}{2} \rceil$ 
8     if  $\max(\Psi) \leq L_2$  and  $\max(\Psi_{k'}) \leq L_k$  then //  $\Rightarrow \Psi$  is deeply ( $\varepsilon'$ -)balanced
9        $t \leftarrow \min(\{t \mid \max(\Psi) + c(O_t) \geq L_2\})$  //  $O_t := \langle v_{i+1}, \dots, v_{i+t} \rangle$ 
10      if  $\frac{2}{k'} \max(\Psi) + h_{k'/2}(O_t) \leq L_k$  then // Condition (ii) of Definition 4
11        return  $\Psi$  //  $\Rightarrow \Psi$  is sufficiently balanced (Theorem 5)
12 return  $\Psi$  // No sufficiently balanced prepacking found  $\Rightarrow$  treat all vertices as fixed vertices

```

6 Experimental Evaluation

We integrated the prepacking technique (see Algorithms 1 and 2) into the recursive bipartitioning algorithm of **KaHyPar**. Our implementation is available from <http://www.kahypar.org>. The code is written in C++17 and compiled using g++9.2 with the flags `-mtune=native -O3 -march=native`. Since **KaHyPar** offers both a recursive bipartitioning and direct k -way partitioning algorithm (which uses the RB algorithm in the initial partitioning phase), we refer to the RB-version using our improvements as **KaHyPar-BP-R** and to the direct k -way version as **KaHyPar-BP-K** (BP = **B**alanced **P**artitioning).

Instances. The following experimental evaluation is based on two benchmark sets. The **REALWORLD** benchmark set consists of 50 hypergraphs originating from the VLSI design and scientific computing domain. It contains instances from the ISPD98 VLSI Circuit Benchmark Suite [2] (18 instances), the DAC 2012 Routability-Driven Placement Benchmark Suite [42] (9 instances), 16 instances from the Stanford Network Analysis (SNAP) Platform [33], and 7 highly asymmetric matrices of Davis et al. [15] (referred to as ASM). For VLSI instances



■ **Figure 2** Vertex weight distributions for each instance type. Each bucket of a histogram shows the number of vertices (y-axis) that contribute $x\%$ to the total weight of the corresponding hypergraph.

(ISPD98 and DAC), we use the area of a circuit element as the weight of its corresponding vertex. We translate sparse matrices (SNAP and ASM instances) to hypergraphs using the row-net model [13] and use the degree of a vertex as its weight. The vertex weight distributions of the individual instance types are depicted in Figure 2.¹

Additionally, we generate ten ARTIFICIAL instances that use the net structure of the ten largest ISPD98 instances. Instead of using the area as weight, we assign new vertex weights that yield instances for which it is difficult to satisfy the balance constraint: Each vertex is assigned either unit weight or a weight chosen randomly from an uniform distribution in $[1, W] \subseteq \mathbb{N}_+$. Both the probability that a vertex has non-unit weight and the parameter W are determined (depending on the total number of vertices) such that the expected number of vertices with non-unit weight is 120 and the expected total weight of these vertices is half the expected total weight of the resulting hypergraph.

System and Methodology. All experiments are performed on a single core of a cluster with Intel Xeon Gold 6230 processors running at 2.1 GHz with 96GB RAM. We compare KaHyPar-BP-R and KaHyPar-BP-K with the latest recursive bipartitioning (KaHyPar-R) and direct k -way version (KaHyPar-K) of KaHyPar [21], the default (PaToH-D) and quality preset (PaToH-Q) of PaToH 3.3 [11], as well as with the recursive bipartitioning (hMetis-R) and direct k -way version (hMetis-K) of hMetis 2.0 [29, 30]. Details about the choices of config parameters that influence partitioning quality or imbalance can be found in Appendix D.

We perform experiments using $k \in \{2, 4, 8, 16, 32, 64, 128\}$, $\varepsilon \in \{0.01, 0.03, 0.1\}$, ten repetitions using different seeds for each combination of k and ε , and a time limit of eight hours. We call a combination of a hypergraph $H = (V, E, c, \omega)$, k , and ε an *instance*. Before partitioning an instance, we remove all vertices $v \in V$ from H with a weight greater than $L_k = (1 + \varepsilon) \lceil \frac{c(V)}{k} \rceil$ as proposed by Caldwell et al. [10] and adapt k to $k' := k - |V_R|$, where V_R represents the set of removed vertices. We repeat that step recursively until there is no vertex with a weight greater than $L_{k'} := (1 + \varepsilon) \lceil \frac{c(V \setminus V_R)}{k'} \rceil$. The input for each partitioner is the subhypergraph $H_{V \setminus V_R}$ of H for which we compute a k' -way partition with $L_{k'}^{\text{LPT}}$ as maximum allowed block weight. Note that since all evaluated partitioners internally employ $L_{k'}$ as balance constraint, we initialize each partitioner with a modified imbalance factor $\hat{\varepsilon}$

¹ The benchmark sets and detailed statistics of their properties are publicly available from <http://algo2.iti.kit.edu/heuer/sea21/>.

■ **Table 1** Percentage of imbalanced instances produced by each partitioner for each combination of instance type and ε .

ε	ISPD98			DAC			ASM			SNAP			ARTIFICIAL		
	0.01	0.03	0.1	0.01	0.03	0.1	0.01	0.03	0.1	0.01	0.03	0.1	0.01	0.03	0.1
KaHyPar-BP-K	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
KaHyPar-BP-R	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
KaHyPar-K	6.3	5.6	0.8	9.5	7.9	6.3	4.1	4.1	2.0	0.9	0.9	0.0	27.1	22.9	11.4
KaHyPar-R	10.3	8.7	7.1	19.0	19.0	14.3	6.1	4.1	4.1	6.2	2.7	0.9	28.6	24.3	12.9
hMetis-K	43.7	22.2	9.5	33.3	22.2	11.1	67.3	32.7	4.1	33.9	20.5	3.6	51.4	38.6	24.3
hMetis-R	17.5	15.1	7.1	20.6	15.9	12.7	8.2	6.1	4.1	15.2	10.7	4.5	58.6	54.3	34.3
PaToH-Q	15.9	11.1	5.6	23.8	17.5	9.5	24.5	6.1	4.1	33.9	8.0	1.8	31.4	24.3	14.3
PaToH-D	9.5	7.9	3.2	20.6	17.5	9.5	28.6	6.1	4.1	22.3	11.6	2.7	20.0	15.7	8.6

instead of ε which is calculated as follows:

$$L_{k'} = (1 + \hat{\varepsilon}) \left\lceil \frac{c(V \setminus V_R)}{k'} \right\rceil = (1 + \varepsilon) \text{LPT}(H_{V \setminus V_R}, k') = L_{k'}^{\text{LPT}} \Rightarrow \hat{\varepsilon} = \frac{L_{k'}^{\text{LPT}}}{\lceil \frac{c(V \setminus V_R)}{k'} \rceil} - 1.$$

We consider the resulting k' -way partition $\Pi_{k'}$ to be imbalanced, if it is not $\hat{\varepsilon}$ -balanced. Each partitioner optimizes the connectivity metric, which we also refer to as the quality of a partition. Partition $\Pi_{k'}$ can be extended to a k -way partition Π_k by adding each of the removed vertices $v \in V_R$ to Π_k as a separate block. Note that adding the removed vertices increases the connectivity metric of a k' -way partition only by a constant value $\alpha \geq 0$. Thus, we report the quality of $\Pi_{k'}$, since $(\lambda - 1)(\Pi_k)$ will be always equal to $(\lambda - 1)(\Pi_{k'}) + \alpha$.

For each instance, we average quality and running times using the arithmetic mean (over all seeds). To further average over multiple instances, we use the geometric mean for absolute running times to give each instance a comparable influence. Runs with imbalanced partitions are not excluded from averaged running times. If *all ten runs* of a partitioner produced imbalanced partitions on an instance, we consider the instance as *imbalanced* and mark it with \mathbf{X} in the plots.

To compare the solution quality of different algorithms, we use *performance profiles* [17]. Let \mathcal{A} be the set of all algorithms we want to compare, \mathcal{I} the set of instances, and $q_A(I)$ the quality of algorithm $A \in \mathcal{A}$ on instance $I \in \mathcal{I}$. For each algorithm A , we plot the fraction of instances (y -axis) for which $q_A(I) \leq \tau \cdot \min_{A' \in \mathcal{A}} q_{A'}(I)$, where τ is on the x -axis. For $\tau = 1$, the y -value indicates the percentage of instances for which an algorithm $A \in \mathcal{A}$ performs best. Note that these plots relate the quality of an algorithm to the best solution and thus do not permit a full ranking of three or more algorithms.

Balanced Partitioning. In Table 1, we report the percentage of imbalanced instances produced by each partitioner for each instance type and ε . Both KaHyPar-BP-K and KaHyPar-BP-R compute balanced partitions for all tested benchmark sets and parameters. For the remaining partitioners, the number of imbalanced solutions increases as the balance constraint becomes tighter. For the previous KaHyPar versions, the number of imbalanced partitions is most pronounced on VLSI instances: For $\varepsilon = 0.01$, KaHyPar-K and KaHyPar-R compute infeasible solutions for 6.3% (10.3%) of the ISPD98 and for 9.5% (19.0%) of the DAC instances. Comparing the distribution of vertex weights reveals that these instances tend to have a larger proportion of *heavier* vertices compared to the ASM and SNAP instances (see Figure 2). The largest benefit of using our approach can be observed on the artificially generated instances, where KaHyPar-K and KaHyPar-R only computed balanced partitions for 72.9% (71.4%) of the instances (for $\varepsilon = 0.01$).

■ **Table 2** Percentage of imbalanced instances produced by each partitioner on our REALWORLD benchmark set for each combination of k and ε .

ε	$k \in \{2, 4, 8\}$			$k \in \{16, 32\}$			$k \in \{64, 128\}$			REALWORLD		
	0.01	0.03	0.1	0.01	0.03	0.1	0.01	0.03	0.1	0.01	0.03	0.1
KaHyPar-BP-K	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
KaHyPar-BP-R	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
KaHyPar-K	0.7	0.0	0.0	5.0	6.0	2.0	11.1	9.1	4.0	4.9	4.3	1.7
KaHyPar-R	2.0	0.7	0.7	11.0	9.0	7.0	21.0	18.0	13.0	10.0	8.0	6.0
hMetis-K	12.0	2.0	0.0	53.0	21.0	11.0	76.0	57.0	14.0	42.0	23.1	7.1
hMetis-R	2.7	2.0	0.0	18.0	14.0	7.0	34.0	27.0	17.0	16.0	12.6	6.9
PaToH-Q	15.3	2.7	0.7	28.0	11.0	5.0	34.0	22.0	11.0	24.3	10.6	4.9
PaToH-D	9.3	2.7	0.7	18.0	11.0	4.0	32.0	22.0	10.0	18.3	10.6	4.3

■ **Table 3** Occurrence of prepacked vertices (i.e., vertices that are fixed to a specific block during partitioning) for each combination of k and ε when using KaHyPar-BP-R on REALWORLD instances: Minimum/average/maximum percentage of prepacked vertices (left), and percentage of instances for which the prepacking is executed at least once (right).

k	$\varepsilon = 0.01$			$\varepsilon = 0.03$			$\varepsilon = 0.1$			Prepacking Triggered		
	Min	Avg	Max	Min	Avg	Max	Min	Avg	Max	$\varepsilon = 0.01$	$\varepsilon = 0.03$	$\varepsilon = 0.1$
2	-	-	-	-	-	-	-	-	-	-	-	-
4	-	-	-	-	-	-	-	-	-	-	-	-
8	≤ 0.1	≤ 0.1	0.2	≤ 0.1	≤ 0.1	≤ 0.1	-	-	-	5.0	1.7	-
16	≤ 0.1	≤ 0.1	≤ 0.1	≤ 0.1	≤ 0.1	≤ 0.1	≤ 0.1	≤ 0.1	≤ 0.1	8.3	5.0	3.3
32	≤ 0.1	8.1	59.0	≤ 0.1	6.2	68.4	≤ 0.1	1.9	14.7	20.0	18.3	10.0
64	≤ 0.1	23.2	87.7	≤ 0.1	17.3	90.9	≤ 0.1	2.7	35.9	18.3	13.3	10.0
128	≤ 0.1	67.9	100.0	≤ 0.1	42.0	96.3	≤ 0.1	15.4	97.0	26.7	20.0	15.0

With some notable exceptions, the number of imbalanced partitions of both variants of PaToH and hMetis-R is comparable to that of KaHyPar-R: PaToH computes significantly fewer feasible solutions on sparse matrix instances (ASM and SNAP) for $\varepsilon = 0.01$, while hMetis-R performs considerably worse on the ARTIFICIAL benchmark set. Out of all partitioners, hMetis-K yields the most imbalanced instances across all benchmark sets. As can be seen in Table 2, the number of imbalanced partitions produced by each competing partitioner increases with decreasing ε and increasing k .

Table 3 shows (i) how often our prepacking algorithm is triggered at least once in KaHyPar-BP-R (see Line 5 in Algorithm 1) and (ii) the percentage of vertices that are treated as fixed vertices (see Table 4 in Appendix E for the results of KaHyPar-BP-K). Except for $k = 128$, on average less than 25% of the vertices are treated as fixed vertices (even less than 10% for $k < 64$), which provides sufficient flexibility to optimize the connectivity objective on the remaining ordinary vertices. However, in a few cases there are also runs where almost all vertices are added to the prepacking. As expected, the triggering frequency and the percentage of fixed vertices increases for larger values of k and smaller ε .

Quality and Running Times. Comparing the different KaHyPar configurations in Figure 3 (left), we can see that our new configurations provide the same solution quality as their non-prepacking counterparts. Furthermore, we see that, in general, the direct k -way algorithm still performs better than its RB counterpart [38]. Figure 3 (middle) therefore compares the

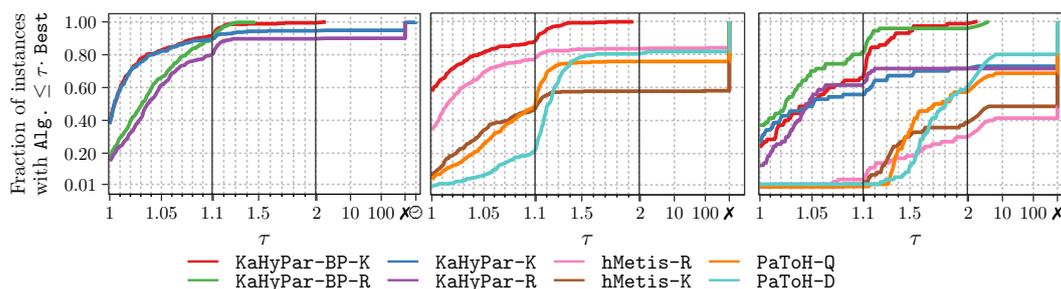


Figure 3 Performance profiles comparing the solution quality of KaHyPar-BP-K and KaHyPar-BP-R with KaHyPar-K (left), KaHyPar-R (left), PaToH (middle), and hMetis (middle) on our REALWORLD benchmark set, and with all systems on our ARTIFICIAL benchmark set ($\varepsilon = 0.01$).

strongest configuration KaHyPar-BP-K with PaToH and hMetis. We see that KaHyPar-BP-K performs considerably better than the competitors. If we compare KaHyPar-BP-K with each partitioner individually on the REALWORLD benchmark set, KaHyPar-BP-K produces partitions with higher quality than those of KaHyPar-K, KaHyPar-BP-R, KaHyPar-R, hMetis-R, hMetis-K, PaToH-Q and PaToH-D on 48.9%, 70.2%, 73.2%, 76.4%, 84.3%, 92.9% and 97.9% of the instances, respectively. KaHyPar-BP-K outperforms KaHyPar-BP-R on the REALWORLD benchmark set. On artificial instances, both algorithms produce partitions with comparable quality for $\varepsilon = \{0.01, 0.03\}$, while the results are less clear for $\varepsilon = 0.1$ (see Figure 3 (right), as well as Figures 4 in Appendix F).

The running time plots (see Figure 5 and 6 in Appendix G) show that our new approach does not impose any additional overheads in KaHyPar. On average, KaHyPar-BP-K is slightly faster than KaHyPar-K as our new algorithm has replaced the previous balancing strategy in KaHyPar (restarting the bipartition with a tighter bound on the weight of the heaviest block if the bipartition is imbalanced). The running time difference is less pronounced for KaHyPar-BP-R and KaHyPar-R. This can be explained by the fact that, in KaHyPar-BP-R, our prepacking algorithm is executed on the input hypergraph, whereas it is executed on the coarsest hypergraph in KaHyPar-BP-K.

7 Conclusion and Future Work

In this work, we revisited the problem of computing balanced partitions for weighted hypergraphs in the multilevel setting and showed that many state-of-the-art hypergraph partitioners struggle to find balanced solutions on hypergraphs with weighted vertices – especially for tight balance constraints. We therefore developed an algorithm that enables partitioners based on the recursive bipartitioning scheme to reliably compute balanced partitions. The method is based on the concept of *deeply balanced* bipartitions and is implemented by pre-assigning a small subset of the heaviest vertices to the two blocks of each bipartition. For this pre-assignment, we established a property that can be verified in polynomial time and, if fulfilled, leads to provable balance guarantees for the resulting k -way partition. We integrated the approach into the recursive bipartitioning algorithm of KaHyPar. Our new algorithms KaHyPar-BP-K and KaHyPar-BP-R are capable of computing balanced solutions on all instances of a diverse benchmark set, without negatively affecting the solution quality or running time of KaHyPar.

Interesting opportunities for future research include replacing the LPT algorithm with an algorithm that additionally optimizes the partitioning objective to construct sufficiently balanced prepackings with improved solution quality [34], and integrating rebalancing strategies similar to the techniques proposed for non-multilevel partitioners [9, 10, 18] into multilevel refinement algorithms.

References

- 1 Y. Akhremtsev, T. Heuer, P. Sanders, and S. Schlag. Engineering a Direct k -way Hypergraph Partitioning Algorithm. In *19th Workshop on Algorithm Engineering and Experiments (ALENEX)*, pages 28–42. SIAM, January 2017.
- 2 C. J. Alpert. The ISPD98 Circuit Benchmark Suite. In *International Symposium on Physical Design (ISPD)*, pages 80–85, April 1998.
- 3 C. J. Alpert and A. B. Kahng. Recent Directions in Netlist Partitioning: A Survey. *Integration: The VLSI Journal*, 19(1-2):1–81, 1995.
- 4 C. Aykanat, B. B. Cambazoglu, and B. Uçar. Multi-Level Direct k -Way Hypergraph Partitioning with Multiple Constraints and Fixed Vertices. *Journal of Parallel and Distributed Computing*, 68(5):609–625, 2008.
- 5 D. A. Bader, H. Meyerhenke, P. Sanders, and D. Wagner, editors. *Graph Partitioning and Graph Clustering, 10th DIMACS Implementation Challenge Workshop*, volume 588 of *Contemporary Mathematics*. American Mathematical Society, February 2013.
- 6 M. A. Bender and M. Farach-Colton. The LCA Problem Revisited. In *Latin American Symposium on Theoretical Informatics*, pages 88–94. Springer, 2000.
- 7 R. H. Bisseling, B. O. Auer Fagginger, A. N. Yzelman, T. van Leeuwen, and Ü. V. Çatalyürek. Two-Dimensional Approaches to Sparse Matrix Partitioning. *Combinatorial Scientific Computing*, pages 321–349, 2012.
- 8 T. N. Bui and C. Jones. Finding Good Approximate Vertex and Edge Partitions is NP-Hard. *Information Processing Letters*, 42(3):153–159, May 1992.
- 9 A. E. Caldwell, A. B. Kahng, and I. L. Markov. Improved Algorithms for Hypergraph Bipartitioning. In *Asia South Pacific Design Automation Conference (ASP-DAC)*, pages 661–666, 2000.
- 10 A. E. Caldwell, A. B. Kahng, and I. L. Markov. Iterative Partitioning with Varying Node Weights. *VLSI Design*, 11(3):249–258, 2000.
- 11 Ü. V. Çatalyürek and C. Aykanat. Decomposing Irregularly Sparse Matrices for Parallel Matrix-Vector Multiplication. In *International Workshop on Parallel Algorithms for Irregularly Structured Problems*, pages 75–86. Springer, 1996.
- 12 Ü. V. Çatalyürek and C. Aykanat. PaToH: Partitioning Tool for Hypergraphs. <https://www.cc.gatech.edu/~umit/PaToH/manual.pdf>, 2011.
- 13 Ü. V. Çatalyürek and Cevdet Aykanat. Hypergraph-Partitioning-Based Decomposition for Parallel Sparse-Matrix Vector Multiplication. *IEEE Transactions on Parallel and Distributed Systems*, 10(7):673–693, 1999.
- 14 Ü. V. Çatalyürek, M. Deveci, K. Kaya, and B. Uçar. UMPa: A Multi-Objective, Multi-Level Partitioner for Communication Minimization. In *Graph Partitioning and Graph Clustering, 10th DIMACS Implementation Challenge Workshop*, pages 53–66, February 2012.
- 15 T. Davis, I. S. Duff, and S. Nakov. Design and Implementation of a Parallel Markowitz Threshold Algorithm. *SIAM Journal on Matrix Analysis and Applications*, 41(2):573–590, April 2020.
- 16 K. D. Devine, E. G. Boman, R. T. Heaphy, R. H. Bisseling, and Ü. V. Çatalyürek. Parallel Hypergraph Partitioning for Scientific Computing. In *20th International Parallel and Distributed Processing Symposium (IPDPS)*, April 2006.
- 17 E. D. Dolan and J. J. Moré. Benchmarking Optimization Software with Performance Profiles. *Mathematical Programming*, 91(2):201–213, 2002.

- 18 S. Dutt and H. Theny. Partitioning Around Roadblocks: Tackling Constraints with Intermediate Relaxations. In *International Conference on Computer-Aided Design (ICCAD)*, pages 350–355, November 1997.
- 19 C. M. Fiduccia and R. M. Mattheyses. A Linear-Time Heuristic for Improving Network Partitions. In *19th Conference on Design Automation (DAC)*, pages 175–181, 1982.
- 20 M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*, volume 174. W.H. Freeman, San Francisco, 1979.
- 21 Lars Gottesbüren, Michael Hamann, Sebastian Schlag, and Dorothea Wagner. Advanced Flow-Based Multilevel Hypergraph Partitioning. In *18th International Symposium on Experimental Algorithms (SEA 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
- 22 R. L. Graham. Bounds on Multiprocessing Timing Anomalies. *SIAM Journal on Applied Mathematics*, 17(2):416–429, 1969.
- 23 R. L. Graham, E. L. Lawler, J. K. Lenstra, and R. Kan. Optimization and Approximation in Deterministic Sequencing and Scheduling: A Survey. In *Annals of Discrete Mathematics*, volume 5, pages 287–326. Elsevier, 1979.
- 24 S. A. Hauck. *Multi-FPGA Systems*. PhD thesis, University of Washington, 1995.
- 25 T. Heuer. Engineering Initial Partitioning Algorithms for direct k -way Hypergraph Partitioning. Bachelor thesis, Karlsruhe Institute of Technology, August 2015.
- 26 T. Heuer, P. Sanders, and S. Schlag. Network Flow-Based Refinement for Multilevel Hypergraph Partitioning. *ACM Journal of Experimental Algorithmics (JEA)*, 24(1):2.3:1–2.3:36, September 2019.
- 27 T. Heuer and S. Schlag. Improving Coarsening Schemes for Hypergraph Partitioning by Exploiting Community Structure. In *16th International Symposium on Experimental Algorithms (SEA)*, Leibniz International Proceedings in Informatics (LIPIcs), pages 21:1–21:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, June 2017.
- 28 G. Karypis. A Software Package for Partitioning Unstructured Graphs, Partitioning Meshes, and Computing Fill-Reducing Orderings of Sparse Matrices, Version 5.1.0. Technical report, University of Minnesota, 2013.
- 29 G. Karypis, R. Aggarwal, V. Kumar, and S. Shekhar. Multilevel Hypergraph Partitioning: Application in VLSI Domain. In *34th Conference on Design Automation (DAC)*, pages 526–529, June 1997.
- 30 G. Karypis and V. Kumar. Multilevel k -way Hypergraph Partitioning. *VLSI Design*, 11(3):285–300, 2000.
- 31 B. W. Kernighan and S. Lin. An Efficient Heuristic Procedure for Partitioning Graphs. *The Bell System Technical Journal*, 49(2):291–307, February 1970.
- 32 T. Lengauer. *Combinatorial Algorithms for Integrated Circuit Layout*. John Wiley & Sons, Inc., 1990.
- 33 J. Leskovec and A. Krevl. SNAP Datasets: Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data>, 2014.
- 34 N. Maas. Multilevel Hypergraph Partitioning with Vertex Weights Revisited. Bachelor thesis, Karlsruhe Institute of Technology, May 2020.
- 35 D. A. Papa and I. L. Markov. Hypergraph Partitioning and Clustering. In *Handbook of Approximation Algorithms and Metaheuristics*. Citeseer, 2007.
- 36 M. Pinedo. *Scheduling*, volume 29. Springer, 2012.
- 37 S. Schlag, V. Henne, T. Heuer, H. Meyerhenke, P. Sanders, and C. Schulz. k -way Hypergraph Partitioning via n -Level Recursive Bisection. In *18th Workshop on Algorithm Engineering and Experiments (ALENEX)*, pages 53–67. SIAM, January 2016.
- 38 Sebastian Schlag. *High-Quality Hypergraph Partitioning*. PhD thesis, Karlsruhe Institute of Technology, 2020.
- 39 C. Schulz. *High Quality Graph Partitioning*. PhD thesis, Karlsruhe Institute of Technology, 2013.

- 40 H. Shin and C. Kim. A Simple Yet Effective Technique for Partitioning. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 1(3):380–386, 1993.
- 41 B. Vastenhouw and R. H. Bisseling. A Two-Dimensional Data Distribution Method for Parallel Sparse Matrix-Vector Multiplication. *SIAM Review*, 47(1):67–95, 2005.
- 42 N. Viswanathan, C. J. Alpert, C. C. N. Sze, Z. Li, and Y. Wei. The DAC 2012 Routability-Driven Placement Contest and Benchmark Suite. In *49th Conference on Design Automation (DAC)*, pages 774–782. ACM, June 2012.

A Proof of Lemma 3

► **Lemma 3.** (LPT Bound). *Let $H = (V, E, c, \omega)$ be a weighted hypergraph, Ψ_k be a k -way prepacking for a set of fixed vertices $P \subseteq V$, and let $O := \langle v_1, \dots, v_m \mid v_i \in V \setminus P \rangle$ be the sequence of all ordinary vertices of $V \setminus P$ sorted in decreasing order of weight. If we assign the remaining vertices O to the blocks of Ψ_k by using the LPT algorithm, we can extend Ψ_k to a k -way partition Π_k of H such that the weight of the heaviest block is bound by:*

$$\max(\Pi_k) \leq \max\left\{\frac{1}{k}c(P) + h_k(O), \max(\Psi_k)\right\}, \text{ with } h_k(O) := \max_{i \in \{1, \dots, m\}} c(v_i) + \frac{1}{k} \sum_{j=1}^{i-1} c(v_j).$$

Proof. We define $\Psi_k := \{P_1, \dots, P_k\}$ and $\Pi_k := \{V_1, \dots, V_k\}$. Let assume that the LPT algorithm assigned the i -th vertex v_i of O to block $V_j \in \Pi_k$. We define $V_j^{(i)}$ as a subset of block V_j that only contains vertices of $\langle v_1, \dots, v_i \rangle \subseteq O$ and P . Since the LPT algorithm always assigns an vertex to a block with the smallest weight (see Section 3), the weight of $V_j^{(i-1)}$ must be smaller or equal to $\frac{1}{k}(c(P) + \sum_{j=1}^{i-1} c(v_j))$ (average weight of all previously assigned vertices), otherwise $V_j^{(i-1)}$ would be not the block with the smallest weight.

$$\Rightarrow c(V_j^{(i)}) = c(V_j^{(i-1)}) + c(v_i) \leq \frac{1}{k}(c(P) + \sum_{j=1}^{i-1} c(v_j)) + c(v_i) \leq \frac{1}{k}c(P) + h_k(O)$$

We can establish an upper bound on the weight of all blocks to which the LPT algorithm assigns an vertex to with $\frac{1}{k}c(P) + h_k(O)$. If the LPT algorithm does not assign any vertex to a block $V_j \in \Pi_k$, its weight is equal to $c(P_j) \leq \max(\Psi_k)$.

$$\Rightarrow \max(\Pi_k) \leq \max\left\{\frac{1}{k}c(P) + h_k(O), \max(\Psi_k)\right\} \quad \blacktriangleleft$$

B Generalized Balance Property

► **Definition 7.** (Generalized Balance Property). *Let $H = (V, E, c, \omega)$ be a hypergraph for which we want to compute an ε -balanced k -way partition and $\Psi := \{P_1, P_2\}$ be a prepacking of H for a set of fixed vertices $P \subseteq V$. Furthermore, let O_{t_1} resp. O_{t_2} be the sequence of the t_1 resp. t_2 heaviest ordinary vertices of $V \setminus P$ sorted in decreasing vertex weight order such that t_1 resp. t_2 is the smallest number that satisfies $c(P_1) + c(O_{t_1}) \geq L_2$ resp. $c(P_2) + c(O_{t_2}) \geq L_2$ (see Line 2, Algorithm 1). We say that a prepacking Ψ satisfies the balance property with respect to k if the following conditions hold:*

- (i) Ψ is deeply balanced
- (ii) $\frac{1}{k_1}c(P_1) + h_{k_1}(O_{t_1}) \leq L_k$ with $k_1 := \lceil \frac{k}{2} \rceil$
- (iii) $\frac{1}{k_2}c(P_2) + h_{k_2}(O_{t_2}) \leq L_k$ with $k_2 := \lfloor \frac{k}{2} \rfloor$

The proof of Theorem 5 can be adapted such that we show that there exist a k_1 - resp. k_2 -way partition Π_{k_1} resp. Π_{k_2} for V_1 resp. V_2 of any balanced bipartition $\Pi_2 := \{V_1, V_2\}$ that respects the prepacking Ψ with $\max(\Pi_{k_1}) \leq \frac{1}{k_1}c(P_1) + h_{k_1}(O_{t_1}) \leq L_k$ (Defintion (ii)) and $\max(\Pi_{k_2}) \leq \frac{1}{k_2}c(P_2) + h_{k_2}(O_{t_2}) \leq L_k$ (Defintion (iii)).

C Proof of Claim 6

► **Lemma 8.** *Let $L = \langle a_1, \dots, a_n \rangle$ be a sequence of elements sorted in decreasing weight order with respect to a weight function $c : L \rightarrow \mathbb{R}_{\geq 0}$ (for a subsequence $A := \langle a_1, \dots, a_l \rangle$ of L , we define $c(A) := \sum_{i=1}^l c(a_i)$), L' be an arbitrary subsequence of L sorted in decreasing weight order and $L_m = \langle a_1, \dots, a_m \rangle$ the subsequence of the $m \leq n$ heaviest elements in L . Then the following conditions hold:*

- (i) *If $c(L') \leq c(L_m)$, then $h_k(L') \leq h_k(L_m)$*
- (ii) *If $c(L') > c(L_m)$, then $h_k(L') - \frac{1}{k}c(L') \leq h_k(L_m) - \frac{1}{k}c(L_m)$*

Proof. For convenience, we define $L' := \langle b_1, \dots, b_l \rangle$. Note that $\forall i \in \{1, \dots, \min(m, l)\} : c(a_i) \geq c(b_i)$, since L_m contains the m heaviest elements in decreasing order. We define $i := \arg \max_{i \in \{1, \dots, l\}} c(b_i) + \frac{1}{k} \sum_{j=1}^{i-1} c(b_j)$ (index that maximizes $h_k(L')$).

(i) + (ii): If $i \leq m$, then

$$h_k(L') = c(b_i) + \frac{1}{k} \sum_{j=1}^{i-1} c(b_j) \stackrel{\forall j \in [1, i]: c(b_j) \leq c(a_j)}{\leq} c(a_i) + \frac{1}{k} \sum_{j=1}^{i-1} c(a_j) \leq h_k(L_m)$$

(i): If $m < i \leq l$, then

$$\begin{aligned} h_k(L') &= c(b_i) + \frac{1}{k} \sum_{j=1}^{i-1} c(b_j) = c(b_i) - \frac{1}{k} \sum_{j=i}^n c(b_j) + \frac{1}{k} c(L') \leq \left(1 - \frac{1}{k}\right) c(b_i) + \frac{1}{k} c(L') \\ &\stackrel{\substack{c(b_i) \leq c(a_m) \\ c(L') \leq c(L_m)}}{\leq} \left(1 - \frac{1}{k}\right) c(a_m) + \frac{1}{k} c(L_m) = c(a_m) + \frac{1}{k} \sum_{j=1}^{m-1} c(a_j) \leq h_k(L_m) \end{aligned}$$

(ii): If $m < i \leq l$, then

$$\begin{aligned} h_k(L') - \frac{1}{k}c(L') &= c(b_i) + \frac{1}{k} \sum_{j=1}^{i-1} c(b_j) - \frac{1}{k}c(L') = c(b_i) - \frac{1}{k} \sum_{l=i}^n c(b_l) \leq \left(1 - \frac{1}{k}\right) c(b_i) \\ &\stackrel{c(b_i) \leq c(a_m)}{\leq} \left(1 - \frac{1}{k}\right) c(a_m) = c(a_m) + \frac{1}{k} \sum_{j=1}^{m-1} c(a_j) - \frac{1}{k}c(L_m) \leq h_k(L_m) - \frac{1}{k}c(L_m) \blacktriangleleft \end{aligned}$$

▷ **Claim 6.** It holds that: $\frac{1}{k'}c(P_1) + h_{k'}(O_1) \leq \frac{1}{k'} \max(\Psi) + h_{k'}(O_t)$.

Proof. Remember, $\Psi = \{P_1, P_2\}$, $\Pi_2 = \{V_1, V_2\}$ with $P_1 \subseteq V_1$ and $P_2 \subseteq V_2$, O_1 is equal to $V_1 \setminus P_1$ and O_t represents the t heaviest vertices of $(V_1 \cup V_2) \setminus (P_1 \cup P_2)$ with $\max(\Psi) + c(O_t) \geq L_2$ as defined in Definition 4. The following proof distinguish two cases based on Lemma 8.

If $c(O_1) \leq c(O_t)$, then

$$\frac{1}{k'}c(P_1) + h_{k'}(O_1) \stackrel{\text{Lemma 8(i)}}{\leq} \frac{1}{k'}c(P_1) + h_{k'}(O_t) \stackrel{c(P_1) \leq \max(\Psi)}{\leq} \frac{1}{k'} \max(\Psi) + h_{k'}(O_t)$$

If $c(O_1) > c(O_t)$, then

$$\begin{aligned}
 \frac{1}{k'}c(P_1) + h_{k'}(O_1) &= \frac{1}{k'}c(P_1) + h_{k'}(O_1) - \frac{1}{k'}c(O_1) + \frac{1}{k'}c(O_1) \\
 &\stackrel{\text{Lemma 8(ii)}}{\leq} \frac{1}{k'}(c(P_1) + c(O_1)) + h_{k'}(O_t) - \frac{1}{k'}c(O_t) \\
 &\quad \stackrel{c(P_1)+c(O_1)=c(V_1)}{=} \frac{1}{k'}(c(V_1) - c(O_t)) + h_{k'}(O_t) \\
 \stackrel{c(V_1) \leq L_2}{\leq} \frac{1}{k'}(L_2 - c(O_t)) + h_{k'}(O_t) &\stackrel{\max(\Psi)+c(O_t) \geq L_2}{\leq} \frac{1}{k'}\max(\Psi) + h_{k'}(O_t) \quad \triangleleft
 \end{aligned}$$

D Configuration of Evaluated Partitioners

hMetis does not directly optimize the $(\lambda - 1)$ -metric. Instead it optimizes the *sum-of-external-degrees* (SOED), which is closely related to the connectivity metric: $(\lambda - 1)(\Pi) = \text{SOED}(\Pi) - \text{cut}(\Pi)$. We therefore configure **hMetis** to optimize SOED and calculate the $(\lambda - 1)$ -metric accordingly. The same approach is also used by the authors of **hMetis** [30]. Additionally, **hMetis-R** defines the maximum allowed imbalance of a partition differently [28]. For example, an imbalance value of 5 means that a block weight between $0.45 \cdot c(V)$ and $0.55 \cdot c(V)$ is allowed at each bisection step. We therefore translate the imbalance parameter ε to a modified parameter ε' such that the correct allowed block weight is matched after $\log_2(k)$ bisections:

$$\varepsilon' := 100 \cdot \left(\left((1 + \varepsilon) \frac{\lceil \frac{c(V)}{k} \rceil}{c(V)} \right)^{\frac{1}{\log_2(k)}} - 0.5 \right)$$

PaToH is evaluated with both the default (**PaToH-D**) and the quality preset (**PaToH-Q**). However, there are also more fine-grained parameters available for **PaToH** as described in [12]. In our case, the **balance** parameter is of special interest as it might affect the ability of **PaToH** to find a balanced partition. Therefore, we evaluated the performance of **PaToH** on our benchmark set with each of the possible options **Strict**, **Adaptive** and **Relaxed**. The configuration using the **Strict** option (which is also the default) consistently produced fewest imbalanced partitions and had similar quality to the other configurations. Consequently, we only report the results of this configuration.

E Prepacking Algorithm Statistics for KaHyPar-BP-K

Table 4 Occurrence of prepacked vertices (i.e., vertices that are fixed to a specific block during partitioning) for each combination of k and ϵ when using KaHyPar-BP-K on REALWORLD instances: Minimum/average/maximum percentage of prepacked vertices (left), and percentage of instances for which the prepacking is executed at least once (right).

k	$\epsilon = 0.01$			$\epsilon = 0.03$			$\epsilon = 0.1$			Prepacking Triggered [%]		
	Min	Avg	Max	Min	Avg	Max	Min	Avg	Max	$\epsilon = 0.01$	$\epsilon = 0.03$	$\epsilon = 0.1$
2	-	-	-	-	-	-	-	-	-	-	-	-
4	-	-	-	-	-	-	-	-	-	-	-	-
8	6.7	17.1	41.6	0.4	0.5	0.6	2.3	2.3	2.3	5.0	3.3	1.7
16	3.1	15.6	34.0	0.2	2.0	7.2	1.9	2.1	2.3	8.3	6.7	3.3
32	0.3	29.9	56.0	0.1	11.7	42.3	0.2	3.4	26.3	13.3	15.0	6.7
64	0.2	54.4	94.3	0.3	23.0	69.3	0.4	6.6	94.7	21.7	10.0	8.3
128	0.5	76.5	100.0	0.4	42.4	91.0	0.3	15.7	59.8	28.3	21.7	11.7

F Quality Comparison for $\epsilon = 0.03$ and $\epsilon = 0.1$

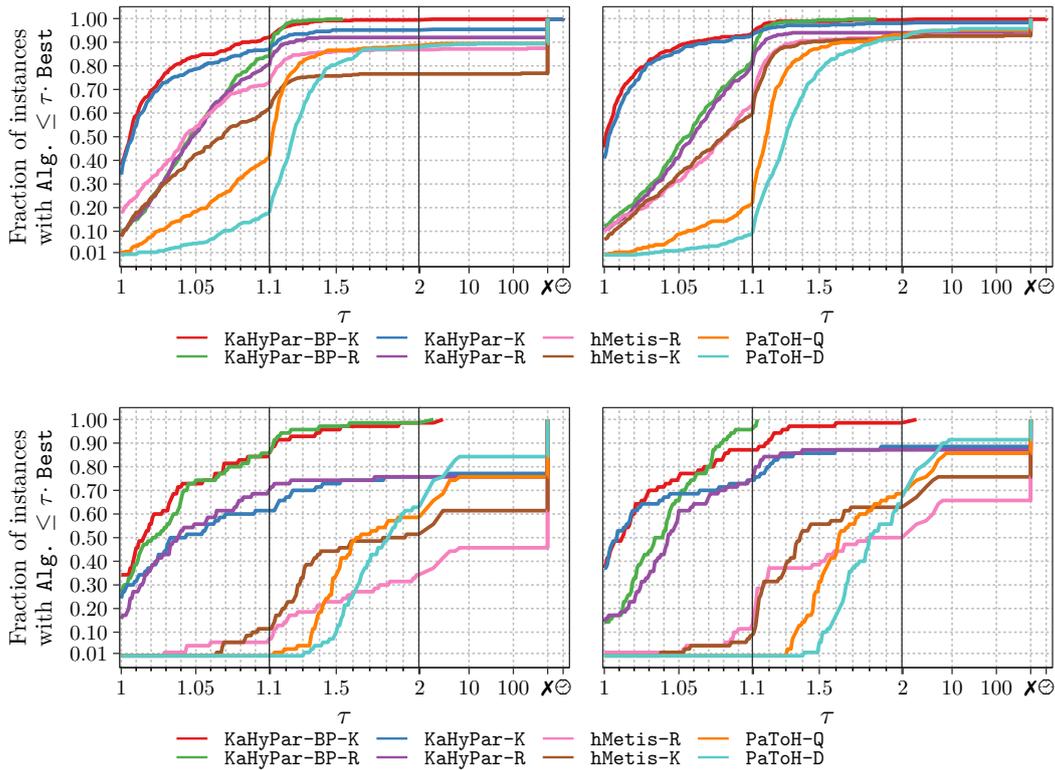


Figure 4 Comparing the solution quality of each evaluated partitioner for $\epsilon = 0.03$ (left) and $\epsilon = 0.1$ (right) on our REALWORLD (top) and ARTIFICIAL (bottom) benchmark set. Note, \odot marks instances that exceeded the time limit.

G Absolute Running Times

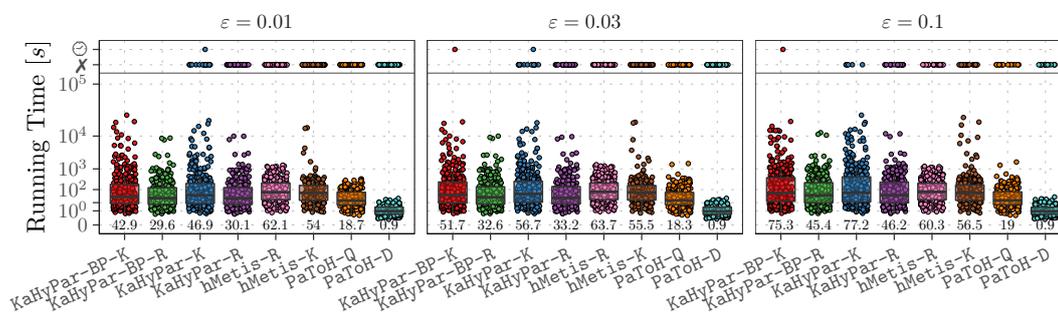


Figure 5 Comparing the running time of each evaluated partitioner for different values of ϵ on our REALWORLD benchmark set. The number under each boxplot denotes the average running time of the corresponding partitioner. Note, \otimes marks instances that exceeded the time limit.

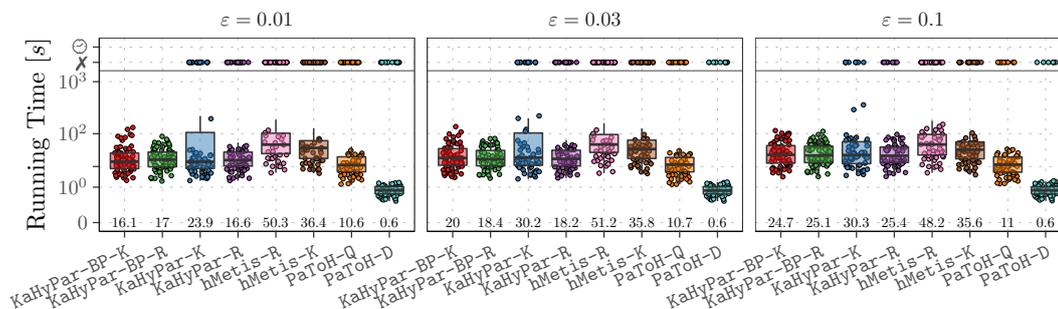


Figure 6 Comparing the running time of each evaluated partitioner for different values of ϵ on our ARTIFICIAL benchmark set. The number under each boxplot denotes the average running time of the corresponding partitioner. Note, \otimes marks instances that exceeded the time limit.