

Evaluation of a Flow-Based Hypergraph Bipartitioning Algorithm

Lars Gottesbüren

Institute of Theoretical Informatics, Karlsruhe Institute of Technology, Germany
lars.gottesbueren@kit.edu

Michael Hamann

Institute of Theoretical Informatics, Karlsruhe Institute of Technology, Germany
michael.hamann@kit.edu

Dorothea Wagner

Institute of Theoretical Informatics, Karlsruhe Institute of Technology, Germany
dorothea.wagner@kit.edu

Abstract

In this paper, we propose HyperFlowCutter, an algorithm for balanced hypergraph bipartitioning that is based on minimum S - T hyperedge cuts and maximum flows. It computes a sequence of bipartitions that optimize cut size and balance in the Pareto sense, being able to trade one for the other. HyperFlowCutter builds on the FlowCutter algorithm for partitioning graphs. We propose additional features, such as handling disconnected hypergraphs, novel methods for obtaining starting S, T pairs as well as an approach to refine a given partition with HyperFlowCutter. Our main contribution is ReBaHFC, a new algorithm which obtains an initial partition with the fast multilevel hypergraph partitioner PaToH and then improves it using HyperFlowCutter as a refinement algorithm. ReBaHFC is able to significantly improve the solution quality of PaToH at little additional running time. The solution quality is only marginally worse than that of the best-performing hypergraph partitioners KaHyPar and hMETIS, while being one order of magnitude faster. Thus ReBaHFC offers a new time-quality trade-off in the current spectrum of hypergraph partitioners. For the special case of perfectly balanced bipartitioning, only the much slower plain HyperFlowCutter yields slightly better solutions than ReBaHFC, while only PaToH is faster than ReBaHFC.

2012 ACM Subject Classification Mathematics of computing → Graph algorithms

Keywords and phrases Hypergraph Partitioning, Maximum Flows, Algorithm Engineering

Digital Object Identifier 10.4230/LIPIcs.ESA.2019.52

Related Version A full version of the paper is available at <https://arxiv.org/abs/1907.02053>.

Supplement Material Source code at <https://github.com/kit-algo/HyperFlowCutter>

Funding This work was supported by the DFG under grants WA654/19-2 and WA654/22-2. The authors acknowledge support by the state of Baden-Württemberg through bwHPC.

Acknowledgements We thank Ben Strasser, Tim Zeitz and Lukas Barth for helpful discussions.

1 Introduction

Given a hypergraph $H = (V, E)$, a hyperedge cut $C \subset E$ is a set of hyperedges whose removal disconnects H . The *balanced hypergraph bipartitioning problem* is to find a hyperedge cut of minimum cardinality whose removal separates H into two blocks of roughly equal size – up to $(1 + \varepsilon)\frac{|V|}{2}$. Hypergraph partitioning has applications in VLSI design, database sharding, and high performance computing, in particular load balancing and reducing communication for highly parallel computations as well as accelerating sparse matrix vector multiplications. This problem is NP-hard [35] and it is hard to approximate, even for graphs [12]. Therefore, practical algorithms use heuristics. Most of them are based on the *multilevel* framework [28].



© Lars Gottesbüren, Michael Hamann, and Dorothea Wagner;
licensed under Creative Commons License CC-BY

27th Annual European Symposium on Algorithms (ESA 2019).

Editors: Michael A. Bender, Ola Svensson, and Grzegorz Herman; Article No. 52; pp. 52:1–52:17

Leibniz International Proceedings in Informatics



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

In this paper we consider a different approach based on the max-flow min-cut duality. The basic idea has already been used in the Flow-Balanced-Bipartition algorithm (FBB) of Yang and Wong [51]. So far it has not been of further consideration due to being too slow compared to multilevel algorithms and too slow to solve current instances in feasible time. More recently, FlowCutter [27] for graph bipartitioning has been introduced independently of FBB. It is designed for computing very small node separators in road networks with a rather loose balance constraint; $\varepsilon = 0.33$ is recommended for the application of accelerating shortest path computations [20]. Based on similar ideas as FBB but equipped with more engineering, it computes both unbalanced and highly balanced bipartitions of high quality on the Walshaw graph partitioning benchmark [47].

Contribution. We present HyperFlowCutter, an algorithm which computes a series of hypergraph bipartitions with increasing balance, up to perfectly balanced solutions. With HyperFlowCutter, we extend FlowCutter to hypergraphs and contribute additional features. We provide a method to handle disconnected hypergraphs, which FlowCutter and FBB cannot handle. Our main contribution is ReBaHFC, an algorithm to refine a given partition using HyperFlowCutter. It is a natural extension of the max-flow based refinement of the k -way hypergraph partitioner KaHyPar [29]. We provide a thoroughly engineered implementation as well as an extensive experimental evaluation on the benchmark set of Heuer and Schlag [30], comparing HyperFlowCutter and ReBaHFC against the state-of-the-art hypergraph partitioning tools KaHyPar [30, 29], PaToH [13] and hMETIS [31, 32].

In our experiments we use the fast algorithm PaToH to obtain initial partitions for ReBaHFC. When using the quality preset of PaToH, ReBaHFC computes solutions for $\varepsilon = 0.03$, which are only slightly worse than those of the best-performing partitioners KaHyPar and hMETIS and significantly better than those of PaToH. ReBaHFC is only marginally slower than PaToH and thus, like PaToH, it is one order of magnitude faster than KaHyPar and hMETIS, when using the quality preset, and two orders of magnitude faster, when using the default preset. Furthermore, ReBaHFC with the PaToH default preset computes better solutions than PaToH with its quality preset. Thus ReBaHFC offers new time-quality trade-offs. For the special case of perfectly balanced bipartitioning, only the much slower plain HyperFlowCutter yields marginally better solutions than ReBaHFC, while only PaToH is faster than ReBaHFC.

Outline. After discussing related work in Section 2 and presenting notation and preliminaries in Section 3, we introduce the core algorithm of HyperFlowCutter for S - T hyperedge cuts in Section 4.1. Then we show how to handle disconnected hypergraphs in Section 4.2, propose our refinement algorithm ReBaHFC in Section 4.3 and finally discuss the experimental evaluation in Section 5.

2 Related Work

For an overview of the field of hypergraph partitioning we refer to survey articles [7, 40, 5]. The most common approach among hypergraph partitioning tools is the multilevel framework. Multilevel algorithms repeatedly *contract* vertices to obtain a hierarchy of *coarser* hypergraphs while trying to preserve the cut structure. On the coarsest hypergraph an *initial partition* is computed in some way. Then the contractions are reversed step-by-step and after every uncontraction a *refinement* algorithm tries to improve the solution. Most multilevel algorithms use a variant of the Fiduccia-Mattheyses (FM) [23] or Kernighan-Lin (KL) [33] local vertex

moving heuristics. These algorithms move vertices between blocks, prioritized by cut improvement. The multilevel framework has been immensely successful because it provides a global view on the problem through local operations on the coarse levels. Furthermore, it allows a great deal of engineering and tuning, which have a drastic impact on solution quality. Even though this framework has been used since the 1990s, the implementations are still improving today. A selection of well-known multilevel hypergraph partitioners are PaToH [13] (scientific computing), hMETIS [31, 32] (VLSI design) KaHyPar [30, 29] (general purpose, n-level), Zoltan [19] (scientific computing, parallel), Zoltan-AlgD [46] (algebraic distances based coarsening, sequential), Mondriaan [49] (sparse matrices), MLPart [4] (circuit partitioning) and Parkway [48] (parallel).

Compared to graph partitioning, the performance of local vertex moving suffers from the presence of large hyperedges with vertices scattered over multiple blocks, since many moves have zero cut improvement. On coarse levels of the multilevel hierarchy, this problem is alleviated since hyperedges contain fewer vertices. A second remedy are flow-based refinement algorithms. For graphs, Sanders and Schulz [43] extract a size-constrained corridor around the cut and compute a minimum cut within this corridor. If the cut is balanced, an improved solution was found, otherwise the step is repeated with a smaller corridor. Heuer et al. [29] extend their approach to hypergraphs by using *Lawler networks* [34]. The Lawler network of a hypergraph is a flow network such that minimum S - T hyperedge cuts can be computed via max-flow.

In their Flow-Balanced-Bipartition algorithm (FBB), Yang and Wong [51] use incremental maximum flows on the Lawler network to compute ε -balanced hypergraph bipartitions. Liu and Wong [37] enhance FBB with a *most-desirable-minimum-cut* heuristic, which is inspired by the correspondence between S - T minimum cuts and closed node sets due to Picard and Queyranne [41]. It is similar to the *most-balanced-minimum-cut* heuristics used in the multilevel graph partitioning tool KaHiP [43] and KaHyPar-MF [29] as well as the *avoid-augmenting-paths* piercing heuristics of FlowCutter [27] and HyperFlowCutter. Li et al. [36] propose a push-relabel algorithm, which operates directly on the hypergraph. Furthermore they present heuristics rooted in VLSI design for choosing sets of initial seed vertices S and T as well as piercing vertices. The performance of their approach in other contexts than VLSI design remains unclear.

For perfectly balanced graph partitioning, diffusion-based methods have been successful [38]. Furthermore Sanders and Schulz [44] propose an algorithm based on detecting negative cycles, which is used on top of their evolutionary partitioner. Delling and Werneck [18] provide an efficient implementation of an optimal branch-and-bound algorithm. Additionally there are metaheuristic approaches such as PROBE [14], as well as multilevel memetic algorithms due to Benlic and Hao [9, 10, 11].

3 Preliminaries

A *hypergraph* $H = (V, E)$ consists of a set of n vertices V and a set of m hyperedges E , where a hyperedge e is a subset of the vertices V . A vertex $v \in V$ is *incident* to hyperedge $e \in E$ if $v \in e$. The vertices incident to e are called the *pins* of e . We denote the incident hyperedges of v by $I(v)$ and its degree by $\deg(v) := |I(v)|$. Furthermore let $p := \sum_{e \in E} |e|$ denote the total number of pins in H . All hypergraphs in this paper are unweighted. H can be represented as a bipartite graph $G = (V \cup E, \{(v, e) \in V \times E \mid v \in e\})$ with bipartite node set $V \cup E$ and an edge for every pin. This is also referred to as the *star expansion* of H . H is *connected* if its star expansion is connected. Let $V[E'] := \bigcup_{e' \in E'} e'$ denote the vertex set induced by the hyperedge set E' . To avoid confusion, we use the terms vertices, hyperedges and pins for hypergraphs, and we use the terms nodes and edges for graphs.

3.1 Hypergraph Partitioning

A *bipartition* of a hypergraph H is a partition (A, B) of the vertices V into two non-empty, disjoint sets (called blocks). The *cut* $\text{cut}(A, B) := \{e \in E \mid e \cap A \neq \emptyset \wedge e \cap B \neq \emptyset\}$ consists of all hyperedges with pins in both blocks. The *size* of a cut is the number of cut hyperedges $|\text{cut}(A, B)|$. Let $\varepsilon \in [0, 1)$. A bipartition (A, B) is ε -balanced if $\max(|A|, |B|) \leq \lceil (1 + \varepsilon) \frac{n}{2} \rceil$. The *balanced hypergraph bipartitioning problem* is to find an ε -balanced bipartition (A, B) which minimizes the cut. The special case $\varepsilon = 0$ is called *perfectly balanced bipartitioning*.


3.2 Maximum Flows

A flow network $\mathcal{N} = (\mathcal{V}, \mathcal{E}, S, T, c)$ is a simple symmetric directed graph $(\mathcal{V}, \mathcal{E})$ with two non-empty *terminal* node sets $S, T \subsetneq \mathcal{V}$, $S \cap T = \emptyset$, the source and target node set, as well as a capacity function $c : \mathcal{E} \mapsto \mathbb{R}_{\geq 0}$. Any node that is not a source node and not a target node is a *non-terminal* node. A flow in \mathcal{N} is a function $f : \mathcal{E} \mapsto \mathbb{R}$ subject to the *capacity constraint* $f(e) \leq c(e)$ for all edges e , *flow conservation* $\sum_{(u,v) \in \mathcal{E}} f((u,v)) = 0$ for all non-terminal nodes v and *skew symmetry* $f((u,v)) = -f((v,u))$ for all edges (u,v) . The *value* of a flow $|f| := \sum_{s \in S, (s,u) \in \mathcal{E}} f((s,u))$ is the amount of flow leaving S . The *residual capacity* $r_f(e) := c(e) - f(e)$ is the additional amount of flow that can pass through e without violating the capacity constraint. The residual network with respect to f is the directed graph $\mathcal{N}_f = (\mathcal{V}, \mathcal{E}_f)$ where $\mathcal{E}_f := \{e \in \mathcal{E} \mid r_f(e) > 0\}$. A node v is *source-reachable* if there is a path from S to v in \mathcal{N}_f , it is *target-reachable* if there is a path from v to T in \mathcal{N}_f . We denote the source-reachable and target-reachable nodes by S_r and T_r , respectively. An *augmenting path* is an S - T path in \mathcal{N}_f . The flow f is a *maximum flow* if $|f|$ is maximal of all possible flows in \mathcal{N} . This is the case iff there is no augmenting path in \mathcal{N}_f . An S - T edge cut is a set of edges whose removal disconnects S and T . The value of a maximum flow equals the weight of a minimum-weight S - T edge cut [24]. The *source-side cut* consists of the edges from S_r to $\mathcal{V} \setminus S_r$ and the *target-side cut* consists of the edges from T_r to $\mathcal{V} \setminus T_r$. The bipartition $(S_r, \mathcal{V} \setminus S_r)$ is induced by the source-side cut and $(\mathcal{V} \setminus T_r, T_r)$ is induced by the target-side cut. Note that $\mathcal{V} \setminus S_r \setminus T_r$ is not necessarily empty.

3.3 Hyperedge Cuts via Maximum Flows

Lawler [34] uses maximum flows to compute minimum S - T hyperedge cuts without balance constraints. On the star expansion, the standard construction to model node capacities as edge capacities [1] is applied to the hyperedge-nodes. A hyperedge e is expanded into an *in-node* e_i and an *out-node* e_o joined by a directed *bridge edge* (e_i, e_o) with unit capacity. For every pin $u \in e$ there are two directed *external edges* $(u, e_i), (e_o, u)$ with infinite capacity. The transformed graph is called the *Lawler network*. A minimum S - T edge cut in the Lawler network consists only of bridge edges, which directly correspond to S - T cut hyperedges in H .

Instead of using the Lawler network, we emulate max-flow algorithms directly on the hypergraph, using an approach first proposed by Pistorius and Minoux [42]. In the paper it is formulated for unit weight hyperedges and the Edmonds-Karp flow algorithm [22] but it can be easily extended to handle weighted hyperedges and emulate any flow algorithm. For every hyperedge $e \in E$, we store the pins sending and receiving flow via e . In this work, we consider only unit weight hyperedges and thus need to store only one pin $\text{flow-from}(e)$ sending flow into e , and one pin $\text{flow-to}(e)$ receiving flow from e . To keep the description simple, it relies on this assumption as well. Let u be a fixed vertex. The idea is to enumerate short paths of the form $u \rightarrow e \in I(u) \rightarrow v \in e$ that correspond to paths in the residual Lawler network. This allows us to emulate algorithms for traversing the residual Lawler network directly on



■ **Figure 1** Flow augmentation and computing S_r, T_r in Fig.1a; adding S_r to S and piercing the source-side cut in Fig.1b. S in blue, $S_r \setminus S$ in yellow, T in green, $T_r \setminus T$ in red, $V \setminus S_r, \setminus T_r$ in white.

the hypergraph, such as Breadth-First-Search or Depth-First-Search, as well as other types of local operations, e. g., the *discharge* and *push* operations in push-relabel algorithms [25]. For every hyperedge $e \in I(u)$ we do the following. If e has no flow, we enumerate all pins $v \in e$. These paths correspond to (u, e_i, e_o, v) in the Lawler network. If e has flow and $u = \text{flow-to}(e)$ we also enumerate all pins $v \in e$. However, these paths correspond to (u, e_o, e_i, v) in the Lawler network. If e has flow and $u = \text{flow-from}(e)$, there is no path in the residual Lawler network starting at u that uses e . If e has flow and $\text{flow-from}(e) \neq u \neq \text{flow-to}(e)$, we enumerate just one path $(u, e, \text{flow-from}(e))$, corresponding to $(u, e_i, \text{flow-from}(e))$ in the Lawler network. If we can push flow from the vertex $\text{flow-from}(e)$ to T , we can redirect the flow that the vertex $\text{flow-from}(e)$ sends into e , and instead route flow coming from u to $\text{flow-to}(e)$. Then u becomes the new $\text{flow-from}(e)$.

We use this approach in our implementation because it is significantly more efficient than the Lawler network in practice. In the last case we can avoid scanning all pins of e . In a preliminary experiment, computing flow directly on the hypergraph yielded a speedup of 15 over using the Lawler network, for a hypergraph with maximum hyperedge size of only 36. The speedup will be more extreme on hypergraphs with larger hyperedges.

Via the Lawler network and the above emulation approach, the notions of flow, source-reachable vertices and source-side cuts translate naturally from graphs to hypergraphs. We use the notation and terminology already known from max-flows in graphs.

4 HyperFlowCutter

In the following we outline the core HyperFlowCutter algorithm, which can only be used on connected hypergraphs. Then we discuss how to handle disconnected hypergraphs, and finally show how to improve an existing partition using HyperFlowCutter.

4.1 The Core Algorithm

The idea of the *core* HyperFlowCutter algorithm is to solve a sequence of incremental S - T max-flow min-cut problems with monotonically increasing cut size and balance, until an ε -balanced bipartition is found. They are incremental in the sense that the terminals S, T of the current flow problem are subsets of the terminals in the next flow problem, which allows us to reuse previously computed flows.

Given starting terminal sets $S_{\text{init}}, T_{\text{init}}$, we set $S := S_{\text{init}}, T := T_{\text{init}}$. First, we compute a maximum S - T flow. We terminate if the bipartition $(S_r, V \setminus S_r)$ induced by the source-side cut or $(V \setminus T_r, T_r)$ induced by the target-side cut is ε -balanced. Otherwise, we add the source-reachable vertices S_r to S , if $|S_r| \leq |T_r|$, or we add T_r to T if $|S_r| > |T_r|$. Assume

$|S_r| \leq |T_r|$ without loss of generality. Further, we add one or more vertices, called *piercing vertices*, to S . This step is called *piercing the cut*. It ensures that the next flow problem yields a different bipartition. Subsequently, we augment the previous flow to a max-flow that respects the new sources. We repeat these steps until an ε -balanced bipartition is found.

Note that after adding S_r to S , the flow is still a maximum S - T flow, even though the added vertices are now exempt from flow conservation. Using the smaller side allows it to catch up with the larger side. In particular, this ensures that ε -balance is always possible, as neither side grows beyond $\lceil n/2 \rceil$ vertices.

A hyperedge with pins in both S and T is *mixed*, all other hyperedges are *non-mixed*. We consider two options to find piercing vertices. The preferred option is to choose all pins $e \setminus S \setminus T$ of a non-mixed hyperedge e in the source-side cut. Adding multiple vertices is faster in practice. This small detail is a major difference to FBB [51] and is necessary to make the running time feasible on the used benchmark set. If the source-side cut contains only mixed hyperedges, we choose a single non-terminal vertex incident to the source-side cut. We prefer piercing vertices which are not reachable from T , as these avoid augmenting paths in the next iteration, and thus the cut size does not increase. Avoiding augmenting paths has the highest precedence, followed by choosing hyperedges over single vertices. Ties are broken randomly. If the piercing vertices are not reachable from T , we do not recompute T_r and we skip flow augmentation, but we do recompute S_r .

We experimented with other piercing methods, including trying to avoid mixed hyperedges, the distance-based piercing of FlowCutter, as well as piercing based on a machine learning technique named ensemble classification. We discuss ensemble classification again in the experimental section, although in a different context. None of these approaches yielded consistent or significant quality improvements over just avoiding augmenting paths and random tie-breaking, which is why we use only those.

Asymptotic Complexity. The asymptotic running time of Core HyperFlowCutter is $\mathcal{O}(cp)$ where c is the cut size of the ε -balanced partition and p is the number of pins in the hypergraph. Roughly speaking, the term p stems from performing up to one hypergraph traversal per flow unit. Here we use that $n \leq p, m \leq p$ holds for connected hypergraphs. Note that the running time is output-sensitive, however the factor c is rather pessimistic in practice, since the flow algorithm finds many augmenting paths in a single traversal. The original proof for FlowCutter [27] is applicable, but implementing the piercing step requires a little care. Selecting piercing vertices takes $\mathcal{O}(c)$ per iteration, and we have at most $n \leq p$ iterations. Selecting a non-mixed hyperedge for piercing takes $\mathcal{O}(c)$ time, by iterating over the cut hyperedges. Selecting single piercing vertices which avoid augmenting paths whenever possible, is slightly more involved, when restricted to $\mathcal{O}(c)$ time. For every cut hyperedge e we additionally track its pins that are not reachable from either side. This can be implemented with a linked list, from which we delete vertices once they get reachable from a side. An alternative implementation divides the memory storing the pins of e into three regions: the pins in S_r , in T_r or not reachable. Then we can check, whether e has pins that are not reachable from either side, and if so pick one. In practice, this adds significantly to the complexity of the implementation and the piercing step is never critical for running time, so our implementation simply scans all non-terminal boundary vertices.

Our implementation has $\mathcal{O}(n + m)$ memory footprint by transferring and re-engineering the implementation details from [27]. This is dominated by the $\mathcal{O}(n + m + p)$ memory for storing the hypergraph.

Isolated Vertices. We call a vertex $v \notin S \cup T$ *isolated* if every incident hyperedge $e \in I(v)$ is mixed. Figure 1c illustrates an isolated vertex. An isolated vertex cannot be reached from either S or T via hyperedges not in the cut. Mixed hyperedges remain in both the source-side cut and the target-side cut. Thus isolated vertices can be moved freely between the two blocks to increase balance. It never makes sense to permanently add them to a side, so we exclude them from the piercing step. Furthermore, this needs to be reflected when checking for ε -balance. For checking the bipartition $(S_r, V \setminus S_r)$ we assume up to $n/2 - |S_r|$ of the isolated vertices were part of S_r , analogously for T_r .

Maximum Flow Algorithm. In our implementation, we adapt Dinic maximum flow algorithm [21] to operate directly on the hypergraph, as described in Section 3.3. Dinic algorithm has two alternating phases that are repeated until a maximum flow is found: computing hop distance labels of nodes, using Breadth-First-Search, and finding augmenting paths using Depth-First-Search, such that the distance labels always increase by one along the path. We interleave the first phase of Dinic algorithm with the computation of reachable vertices, in order to avoid duplicate hypergraph traversal. This intrusive optimization is important for improving the running time of flow augmentation in practice, as the part of the running time of the flow algorithm that cannot be charged towards computing reachable vertices is dominated by the part that can be. This is not possible with push-relabel algorithms [25], which is why they were ruled out after short experimentation. We experimented with the Edmonds-Karp flow algorithm [22], modified to augment flow along multiple vertex-disjoint paths in one Breadth-First-Search by propagating vertex labels through the layers. It is slightly slower than Dinic for plain HyperFlowCutter but unfeasible for the refinement variant of HyperFlowCutter, since there are fewer vertices to route flow through and thus the amount of flow augmented per Breadth-First-Search is limited by a few bottleneck vertices.

4.2 Disconnected Hypergraphs

The HyperFlowCutter core algorithm is limited to connected hypergraphs since it computes S - T -cuts. An obvious approach for handling disconnected hypergraphs is connecting components artificially. We refrained from this because a component that intersects neither S nor T would be added to S or T in its entirety. Instead, we run the core algorithm up to $\varepsilon = 0$ on every connected component. The core algorithm computes multiple bipartitions with different balances. We systematically try all possible ways to combine the bipartitions of the components into a bipartition of H .

This can be stated as a generalization of the well-known SUBSETSUM problem. SUBSETSUM is a weakly NP-hard decision problem, which asks whether a subset of an input multiset of positive integers $A = \{a_1, \dots, a_z\}$, the *items*, sums to an input target sum W . Finding a bipartition with zero cut is equivalent to SUBSETSUM, where the items are the sizes of the components and W is the minimum size of the smaller block. We are interested in any subset summing to at least W . Let A be sorted in increasing order and let $Q(i, S)$ be a boolean variable, which is true iff a subset of the first i items sums to S . The standard pseudo-polynomial time dynamic program (DP) [15, Section 35.5] for SUBSETSUM computes solutions for all possible target sums. It fills the DP table Q by iterating through the items in increasing order and setting $Q(i, S)$ to true if $Q(i-1, S - a_i)$ or $Q(i-1, S)$ is true. For filling row i , only row $i-1$ is required, so the memory footprint is not quadratic.

We now turn to non-zero cut bipartitions by allowing to split items in different ways and associating costs with the splits. The core algorithm computes multiple bipartitions P_i on component C_i , at most one for every possible size of the smaller side. These correspond

directly to the different ways we can split items. The associated cost is the cut size. We modify the standard SUBSETSUM DP to minimize the added cuts instead of finding any subset, ensuring every component/item is split only one way in a solution, i. e., select a bipartition, and trying the smaller and larger side of the component bipartition for the smaller side of the bipartition of H . The worst case asymptotic running time of this DP is $\mathcal{O}(\sum_{i=1}^z \sum_{j=1}^{i-1} |P_i||P_j|)$.

We propose some optimizations to make the approach faster in practice. First we solve standard SUBSETSUM to check whether H has an ε -balanced bipartition with zero cut. For the *gap-filler* optimization, we find the largest $g \in \mathbb{N}$ such that for every $x \in [0, g]$ there are connected components, whose sizes sum to x . Computing g is possible in $\mathcal{O}(n)$ time. Let C_1, \dots, C_z be sorted by increasing size, which takes $\mathcal{O}(n)$ time using counting sort [15, Section 8.2]. Then $g = \sum_{j=1}^{k-1} |C_j|$ for the smallest k such that $|C_k| > 1 + \sum_{j=1}^{k-1} |C_j|$. It is never beneficial to split the components C_1, \dots, C_{k-1} . For most hypergraphs we consider in our experiments, we do not invoke the DP because, due to gap-filler, we split only the largest component. For the hypergraphs on which we do invoke the DP, its running time is negligible. Nonetheless, it is easy to construct a worst case instance, where the quadratic running time is prohibitive. For a robust algorithm, we propose to sample bipartitions from every P_i so that the worst case running time falls below some input threshold. The samples should include balanced bipartitions to guarantee that a balanced partition on H can be combined from those on $H[C_i]$.

4.3 HyperFlowCutter as a Refinement and Balancing Algorithm

Instead of partitioning from scratch, HyperFlowCutter can be used to refine an existing balanced bipartition $\pi = (V_1, V_2)$, or repair the balance of an unbalanced bipartition. We fix two blocks of vertices $V'_i \subset V_i$ such that $|V'_i| \leq \alpha \cdot n$ for a *relative block-size threshold* parameter $\alpha \in [0, 0.5]$. To obtain V'_i we run Breadth-First-Search from the boundary vertices on the side of V_i until $|V_i| - \alpha \cdot n$ vertices have been visited. The $\alpha \cdot n$ vertices not visited by the Breadth-First-Search are set as V'_i . Then we run HyperFlowCutter with $S = V'_1, T = V'_2$. We call this algorithm *ReBaHFC*. This idea is equivalent to the way KaHyPar and KaHiP extract *corridors* around the cut for their flow-based refinement. Only the semantics of the size constraint are different to our approach. However, KaHyPar and KaHiP only compute one flow. If the associated bipartition is not balanced, a smaller flow network is derived. This is repeated until the bipartition is balanced. ReBaHFC does not need to repeatedly re-scale flow networks.

In this work we only perform refinement as a post-processing step to a given partition, whereas KaHyPar and KaHiP employ flow-based refinement on the multilevel hierarchy. In a future work we hope to integrate HyperFlowCutter based refinement into KaHyPar. Using ReBaHFC could eliminate the significant overhead of repeated rescaling and improve solution quality – in particular when the minimum cut is just short of being balanced.

We use the fast multilevel partitioner PaToH [13] to obtain initial partitions. We briefly discuss properties of PaToH and differences between its presets. For coarsening, PaToH uses agglomerative clustering, based on the number of common hyperedges divided by cluster size. For initial partitioning, a portfolio of graph growing, bin packing and neighborhood expansion algorithms is used. For refinement PaToH uses a pass of FM [23] followed by a pass of KL [33], each initialized with boundary nodes.

In order to improve cut size, Walshaw [50] proposed to iterate the multilevel scheme by contracting only nodes from the same block, which maintains the cut size, thus allowing refinement on coarse levels starting from a relatively high quality partition. The existing

■ **Table 1** Average and quantile speedups of the hybrid and interleaved execution strategies.

	avg	min	0.1	0.25	median	0.75	0.9	max
hybrid	2.21	0.4	0.96	1.07	1.29	1.55	2.57	49.66
interleaved	3.88	0.55	1.0	1.14	1.38	1.74	2.66	175.47

partition serves as initial partition on the coarsest level. One iteration is called a *V-cycle*. Contraction can be stopped at different stages. The quality preset PaToH-Q, uses 3 full *V-cycles* and 3 shorter *V-cycles* as opposed to the single *V-cycle* of the default preset PaToH-D. To accelerate partitioning, both presets temporarily discard hyperedges which have more pins than some threshold. PaToH-D sets a lower threshold than PaToH-Q.

5 Experimental Evaluation

We implemented HyperFlowCutter and ReBaHFC in C++17 and compiled our code using g++8 with flags `-O3 -mtune=native -march=native`. The source code is available on GitHub¹. Experiments are performed sequentially on a cluster of Intel Xeon E5-2670 (Sandy Bridge) nodes with two Octa-Core processors clocked at 2.6 GHz with 64 GB RAM, 20 MB L3- and 8×256 KB L2-Cache, using only one core of a node.

We use the benchmark set² of Heuer and Schlag [30], which has been used to evaluate KaHyPar. It consists of 488 hypergraphs from four sources: the ISPD98 VLSI Circuit Benchmark Suite [3] (VLSI, 18 hypergraphs), the DAC 2012 Routability-Driven Placement Benchmark Suite [39] (DAC, 10), the SuiteSparse Matrix Collection [17] (SPM, 184) and the international SAT Competition 2014 [8] (Literal, Primal, Dual, 92 hypergraphs each). The set contains 173 disconnected hypergraphs, in particular all DAC instances are disconnected. Refer to [30] for more information on how the hypergraphs were derived. Unless mentioned otherwise, experiments are performed on the full benchmark set.

In the following we describe the configuration of ReBaHFC and plain HyperFlowCutter, before comparing them to competing algorithms in Section 5.4.


5.1 General HyperFlowCutter Configuration

To improve the solution quality of HyperFlowCutter, we run it $q \in \mathbb{N}$ times with different terminal pairs and take the minimum cut. To improve running time we run them simultaneously, in an *interleaved* fashion, as already described in [27], so that the output-sensitive running time depends on the smallest found ε -balanced cut, not the largest. We always schedule the terminal pair with the currently smallest cut to progress.

Table 1 shows the average and some quantile speedups when interleaving the execution of 20 random terminal vertex pairs, instead of running them one after another; repeated for 5 random seeds. Because consecutive execution exhibits more memory locality, we also tested a hybrid strategy where the instance with the currently smallest cut is allowed to make multiple progress iterations. Interleaving outperforms consecutive execution by a factor of 3.88 on average and also consistently beats hybrid execution. This shows that saving work is more important than memory locality. These numbers stem from a preliminary experiment

¹ Source code available at <https://github.com/kit-algo/HyperFlowCutter>

² Benchmark set and detailed statistics available at <https://algo2.iti.kit.edu/schlag/sea2017/>



■ **Figure 2** Improvement ratios $1 - \frac{\text{cut}(\text{ReBaHFC})}{\text{cut}(\text{PaToH})}$ of the two ReBaHFC variants and $\varepsilon = 0, 0.03$ compared to their initial partitions. The curves are independent from one another. Higher values are better. Ratios of zero mean ReBaHFC did not improve the initial partition. We count partitions with zero cut as an improvement ratio greater than 1, since ReBaHFC does not invoke PaToH in this case.

on a 139 hypergraph subset of the full benchmark set. The subset contains 78 sparse matrices, 17 Primal, 23 Dual, 6 Literal SAT instances, 15 VLSI and 0 DAC instances. It contains only connected hypergraphs (all DAC instances are disconnected) in order to measure the impact of interleaving, not the setup overhead for many small connected components.

5.2 ReBaHFC Configuration

We now discuss the configuration for ReBaHFC. The imbalance for the initial partition is set to the same value as the desired imbalance ε for the output partition, which proved superior to larger imbalances on initial partitions. The block-size threshold parameter α should depend on ε , so we settled on $\alpha = 0.4$ for $\varepsilon = 0.03$ and $\alpha = 0.46$ for $\varepsilon = 0$. In the TR [26] we conduct a thorough parameter study for these choices. We resize the blocks once and run HyperFlowCutter five times, interleaved as described in the previous section. This number seems to provide decent quality without increasing running time too much. We consider two variants: ReBaHFC-D, which uses PaToH with default preset and ReBaHFC-Q, which uses PaToH with quality preset.

Figure 2 shows how much ReBaHFC improves the initial partition. We run the two ReBaHFC variants for $\varepsilon = 0, 0.03$ on all hypergraphs of the benchmark set with five different random seeds and plot the ratio $1 - \frac{\text{cut}(\text{ReBaHFC})}{\text{cut}(\text{PaToH})}$ per run. Note that there is no comparison between the curves, and higher values are better for ReBaHFC. Table 2 reports how often ReBaHFC improves the initial partition, for different hypergraph classes. As expected, ReBaHFC-Q could improve fewer solutions than ReBaHFC-D since the PaToH baseline is already better. Furthermore, ReBaHFC has more opportunities for refinement with $\varepsilon = 0.03$, in particular on the DAC and VLSI instances, whereas it struggles with the Primal and Literal SAT instances for $\varepsilon = 0$. Note that other refinement algorithms do not always improve solutions either. In particular, local moving based refinement algorithms struggle with zero-gain moves in the presence of large hyperedges, and the flow-based refinement in KaHyPar can yield unbalanced solutions or reproduce the existing solution. These results show that HyperFlowCutter is a promising candidate for a refinement algorithm integrated in a multilevel partitioner, which is a direction we hope to investigate in future work.

■ **Table 2** Overview by hypergraph class, how often ReBaHFC improves the initial partition.

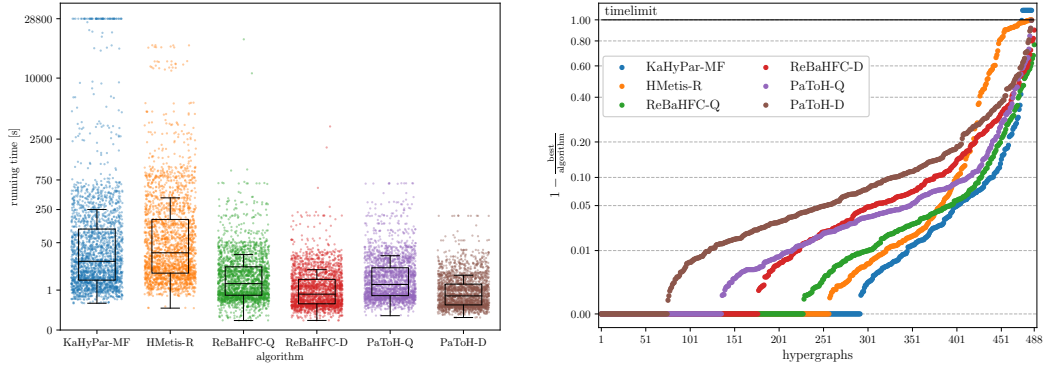
Algorithm	ε	all	SPM	Dual	Primal	Lit	DAC	VLSI
ReBaHFC-Q	0.00	37.3	39.8	47.4	23.9	33.0	66.0	34.4
ReBaHFC-D	0.00	49.2	58.4	59.8	27.4	41.3	58.0	47.8
ReBaHFC-Q	0.03	52.5	47.2	51.5	50.9	57.8	82.0	76.7
ReBaHFC-D	0.03	64.5	61.5	65.9	56.7	71.1	76.0	86.7

The PaToH runs in the experiments from Section 5.4 use other random seeds than those used internally in ReBaHFC. This makes sure that stand-alone PaToH can find smaller cuts than ReBaHFC.

5.3 Plain HyperFlowCutter Configuration

For the experiments on perfectly balanced partitioning we run plain HyperFlowCutter with up to $q = 100$ terminal pairs and take the minimum cut. This value was used already for FlowCutter [27]. With plain HyperFlowCutter we want to push the envelope on solution quality for $\varepsilon = 0$, regardless of running time – because, as the experiments show, ReBaHFC already provides a good time-quality trade-off. The most simple method for choosing starting terminals is to select random vertices. We unsuccessfully experimented with *pseudo-peripheral* terminals, i. e. two vertices that are intuitively far away from each other and at the boundary of the hypergraph. Instead we propose a selection method based on ensemble classification. Ensemble classification is a technique used in machine learning to build a strong classifier from multiple weak ones. We compute 10 bipartitions π_1, \dots, π_{10} with PaToH-D. Let $x \equiv y \Leftrightarrow \pi_i(x) = \pi_i(y)$ for all $i = 1, \dots, 10$ be the equivalence relation, in which two vertices are equivalent if they are in the same block for all ensemble bipartitions. An equivalence class is likely in the same block of a good bipartition and is thus suited as a terminal set. We order the equivalence classes by size in descending order and group two successive classes as one terminal pair. Generally speaking, the larger equivalence classes make for better terminal pairs. Based on experiments in the TR [26], we use 3 ensemble terminal pairs and 97 random vertex pairs. The reported running time for plain HyperFlowCutter always includes the running time for the 10 PaToH-D runs.

On 42 of the 488 hypergraphs, plain HyperFlowCutter with 100 terminal pairs exceeds the eight hour time limit. One downside of interleaving executions is that the solution is only available once all terminal pairs have been processed. Instead of interleaving all 100 executions, we run four waves of $\langle 1, 5, 14, 80 \rangle$ terminal pairs consecutively and interleave execution within waves. An improved bipartition is available after every wave, so that, even if the time limit is exceeded, a solution is available as long as the first wave has been completed. We chose wave sizes, so that completing waves four and three corresponds to 100 and 20 terminal pairs, respectively, as these values were used in [27]. The first wave consists of the first ensemble terminal pair, the second/third wave consist of 5/14 random terminal pairs and the fourth wave consists of 78 random as well as two additional ensemble terminal pairs. There are 438 hypergraphs for which the fourth wave finishes, 35 for which the third but not the fourth wave finishes, 6 for the second, 1 for the first and there are 8 hypergraphs which are partitioned with zero cut, using just the subset sum preprocessing.



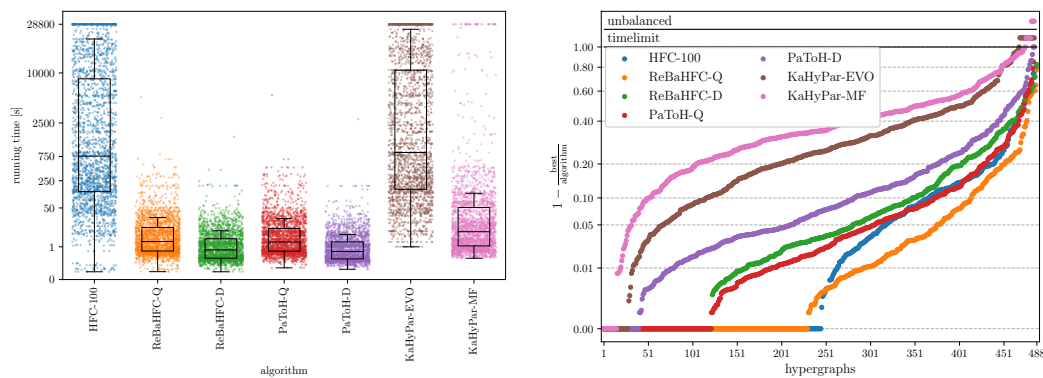
■ **Figure 3** Comparison between the algorithms for $\varepsilon = 0.03$. Left: Absolute running times for every hypergraph and random seed. Right: Performance plot relating the minimum cut per algorithm and hypergraph to the overall best cut for that hypergraph. Lower values are better.

5.4 Comparing ReBaHFC and HyperFlowCutter against State-of-the-Art Hypergraph Partitioners

In this section, we compare ReBaHFC and plain HyperFlowCutter against state-of-the-art hypergraph partitioners. After discussing our comparison methodology, we present results for two settings, namely $\varepsilon = 0.03$ and $\varepsilon = 0$.

Methodology. We run each partitioner five times with different random seeds and report the minimum cut. For every run we set a time limit of eight hours. We use the *performance plots* introduced in [45] to compare algorithms on a per-hypergraph basis regarding cut size. For each algorithm and hypergraph these plots contain a *performance ratio* $1 - \text{best}/\text{algorithm}$, which relates the minimum cut found by any algorithm to the minimum cut found by this algorithm. The ratios of each algorithm are sorted in increasing order. A ratio of 0 means that this algorithm found the smallest overall cut, the number of achieved ratios of 0 is the number of hypergraphs on which this algorithm is the best. Furthermore, algorithm A dominates algorithm B if the curve of A is strictly below that of B. We use values greater than 1 to indicate that algorithms exceeded the time limit or produced unbalanced solutions. This is clearly marked in the plots. To include partitions with zero cut, we set the performance ratio to 0, if the algorithm found the zero cut partition, and 1 otherwise. The performance plots use a cube root scaled y-axis in order to reduce right skewness [16] and to give a fine-grained view on the smaller improvements. For comparing algorithms regarding running time we use a combination of a scatter plot, which shows every measured running time, and a box plot (0.25, median, 0.75 quantiles, whiskers at most extreme points within distance $1.5 \cdot \text{IQR}$ from the upper/lower quartile). The running time plots use a fifth root scaled y-axis for a fine-grained view on areas with smaller running times, which contain more data points.

Comparison for 3% imbalance. For $\varepsilon = 0.03$ we compare ReBaHFC against the state-of-the-art hypergraph partitioning tools KaHyPar-MF (the latest version of KaHyPar with flow-based refinement) and hMETIS-R (the recursive bisection variant of hMETIS), as well as PaToH-D (default preset) and PaToH-Q (quality preset). We use the library interface of PaToH. According to the hMETIS manual, hMETIS-R is preferred over hMETIS-K (direct k-way) for bipartitions, so we exclude hMETIS-K. These tools were chosen because



■ **Figure 4** Comparison between the algorithms for $\varepsilon = 0$. Left: Absolute running times for every hypergraph and random seed. Right: performance plot relating the minimum cut per algorithm and hypergraph to the overall best cut for that hypergraph. Lower values are better.

they provide the best solution quality according to [2, 30]. We chose $\varepsilon = 0.03$ as this is a commonly used value in the literature. Plain HyperFlowCutter is excluded from this part of the experiments because it is not competitive.

Figure 3 shows the running times and a performance plot on the full benchmark set for $\varepsilon = 0.03$. In addition to the running time plot, we compare algorithms by the geometric mean of their running times. We use the geometric mean in order to give instances of different sizes a comparable influence. KaHyPar-MF finds the smallest cut on 292 hypergraphs, hMETIS-R on 257, ReBaHFC-Q on 228, ReBaHFC-D on 177, PaToH-Q on 136 and PaToH-D on 75 of the 488 hypergraphs. While KaHyPar-MF is the best algorithm regarding solution quality, it is also the slowest, exceeding the time limit on 11 hypergraphs. For the instances on which ReBaHFC-Q does not find the best solution it provides solution quality similar to hMETIS-R and only marginally worse than KaHyPar-MF. In particular its solution quality compared to the best cut deteriorates less than that of hMETIS-R. With 2.23s PaToH-Q is one order of magnitude faster than KaHyPar (34.1s) and hMETIS-R (20.1s), whereas ReBaHFC-Q (2.32s) is only slightly slower than PaToH-Q. Furthermore ReBaHFC-D (0.68s) finds more of the best solutions than PaToH-Q at a running time between PaToH-D (0.5s) and PaToH-Q. Thus ReBaHFC-Q and ReBaHFC-D provide new Pareto points in the time-quality trade-off. In the TR [26], we also report performance plots for the different hypergraph classes of the benchmark set. ReBaHFC is particularly good on the DAC and SPM instances. There are hypergraphs on which ReBaHFC is faster than PaToH. These are disconnected hypergraphs, for which ReBaHFC invokes PaToH on smaller sub-hypergraphs, due to the gap-filler optimization and the SUBSETSUM preprocessing described in Section 4.2.

Comparison for perfectly balanced partitioning. Even though the setting $\varepsilon = 0$ has received no attention in hypergraph partitioning and only some attention in graph partitioning [44, 38, 14, 9, 10, 11, 18], we consider it here. Previous studies on perfectly balanced partitioning for graphs have focused on high quality solutions through running time intensive metaheuristics such as evolutionary algorithms [44, 10, 9] or tabu search [11] and even an exact branch-and-bound algorithm [18]. Therefore, we include KaHyPar-EVO [6] (the evolutionary algorithm of KaHyPar) as well as plain HyperFlowCutter in addition to the already considered algorithms. We exclude hMETIS-R from this comparison since it rejects $\varepsilon < 0.002$ for bipartitions.

We include plain HyperFlowCutter with up to 100 terminal pairs as described in Section 5.1 and denote this configuration as HFC-100. The evolutionary algorithm KaHyPar-EVO generates, manages and improves a pool of solutions until a time limit is exceeded, and outputs the minimum cut out of all generated solutions. We set the instance-wise time limit to the maximum of the running times of HFC-100 and KaHyPar-MF to evaluate whether KaHyPar-EVO can yield better solution quality when given the same running time as HFC-100. As opposed to the original paper, we configure KaHyPar-EVO to use flow-based refinement, which further improves solution quality.

KaHyPar-MF is unable to find any balanced bipartition on 4 hypergraphs, whereas KaHyPar-EVO always finds one. Furthermore, KaHyPar-MF exceeds the time limit on 7 hypergraphs and KaHyPar-EVO on an additional 17, without reporting intermediate solutions. Figure 4 shows the running times and a performance plot of all tested algorithms. HFC-100 produces the best solutions on 245 hypergraphs, followed by ReBaHFC-Q (230), ReBaHFC-D (122), PaToH-Q (121), PaToH-D (40), KaHyPar-EVO (28) and finally KaHyPar-MF (15). This shows that with exorbitant running time, HFC-100 produces high quality solutions for $\varepsilon = 0$. However the time-quality trade-off is clearly in favor of ReBaHFC-Q, especially since the solution quality of the latter is closer to the best cut for the instances on which it does not find the best cut, as opposed to HFC-100. PaToH is better than KaHyPar for $\varepsilon = 0$ because it includes a KL [33] refinement pass as opposed to KaHyPar which only uses FM [23].

6 Conclusion

In this paper we propose and evaluate HyperFlowCutter, a hypergraph bipartitioning algorithm based on maximum flow computations. It enumerates partitions with increasing balance up to perfect balance. We also propose and evaluate ReBaHFC, a refinement algorithm based on HyperFlowCutter.

In our experimental evaluation on a large set of hypergraphs, we show that while ReBaHFC is unable to beat the state-of-the-art hypergraph partitioners in terms of quality, it is still close in terms of quality and at the same time an order of magnitude faster. Thus, it offers a new trade-off between quality and running time. For the special case of perfectly balanced bipartitioning, the plain HyperFlowCutter algorithm, while being slow, computes the highest-quality solutions. In this setting, ReBaHFC not only still beats all other partitioners but is also much faster.

In future work, it would be interesting to integrate the refinement step of ReBaHFC into multilevel partitioners to see if it can further improve their solution quality.

References

- 1 Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.
- 2 Yaroslav Akhremtsev, Tobias Heuer, Sebastian Schlag, and Peter Sanders. Engineering a direct k-way Hypergraph Partitioning Algorithm. In *Proceedings of the 19th Meeting on Algorithm Engineering and Experiments (ALENEX'17)*, pages 28–42. SIAM, 2017. doi:10.1137/1.9781611974768.3.
- 3 Charles J. Alpert. The ISPD98 Circuit Benchmark Suite. In *Proceedings of the 1998 International Symposium on Physical Design*, pages 80–85, 1998. doi:10.1145/274535.274546.
- 4 Charles J. Alpert, Jen-Hsin Huang, and Andrew B. Kahng. Multilevel Circuit Partitioning. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 17(8):655–667, August 1998. doi:10.1109/43.712098.

- 5 Charles J. Alpert and Andrew B. Kahng. Recent Directions in Netlist Partitioning: A Survey. *Integration: The VLSI Journal*, 19(1-2):1–81, 1995. doi:10.1016/0167-9260(95)00008-4.
- 6 Robin Andre, Sebastian Schlag, and Christian Schulz. Memetic Multilevel Hypergraph Partitioning. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 347–354. ACM Press, July 2018. doi:10.1145/3205455.3205475.
- 7 David A. Bader, Henning Meyerhenke, Peter Sanders, and Dorothea Wagner. *Graph Partitioning and Graph Clustering: 10th DIMACS Implementation Challenge*, volume 588. American Mathematical Society, 2013. doi:10.1090/conm/588.
- 8 Anton Belov, Daniel Diepold, Marijn JH Heule, and Matti Järvisalo. The SAT Competition 2014, 2014. URL: <http://www.satcompetition.org/2014/>.
- 9 Una Benlic and Jin-Kao Hao. An Effective Multilevel Memetic Algorithm for Balanced Graph Partitioning. In *IEEE International Conference on Tools with Artificial Intelligence*, pages 121–128. IEEE Computer Society, 2010. doi:10.1109/ICTAI.2010.25.
- 10 Una Benlic and Jin-Kao Hao. A Multilevel Memetic Approach for Improving Graph k-Partitions. *IEEE Transactions on Evolutionary Computation*, 15(5):624–642, October 2011. doi:10.1109/TEVC.2011.2136346.
- 11 Una Benlic and Jin-Kao Hao. An Effective Multilevel Tabu Search Approach for Balanced Graph Partitioning. *Computers & Operations Research*, 38(7):1066–1075, July 2011. doi:10.1016/j.cor.2010.10.007.
- 12 Thang Nguyen Bui and Curt Jones. Finding good approximate vertex and edge partitions is NP-hard. *Information Processing Letters*, 42(3):153–159, 1992. doi:10.1016/0020-0190(92)90140-Q.
- 13 Umit Catalyurek 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. doi:10.1109/71.780863.
- 14 Pierre Chardaire, Musbah Barake, and Geoff P. McKeown. A PROBE-Based Heuristic for Graph Partitioning. *IEEE Transactions on Computers*, 56(12):1707–1720, 2007. doi:10.1109/TC.2007.70760.
- 15 Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. MIT Press, second edition, 2001.
- 16 Nicholas J. Cox. Stata Tip 96: Cube Roots. *The Stata Journal*, 11(1):149–154, 2011. doi:10.1177/1536867X1101100112.
- 17 Timothy A. Davis and Yifan Hu. The University of Florida Sparse Matrix Collection. *ACM Transactions on Mathematical Software*, 38(1), 2011. doi:10.1145/2049662.2049663.
- 18 Daniel Delling and Renato F. Werneck. Better Bounds for Graph Bisection. In Leah Epstein and Paolo Ferragina, editors, *Proceedings of the 20th Annual European Symposium on Algorithms (ESA’12)*, volume 7501 of *Lecture Notes in Computer Science*, pages 407–418. Springer, 2012. doi:10.1007/978-3-642-33090-2_36.
- 19 Karen Devine, Erik Boman, Robert Heaphy, Rob Bisseling, and Umit Catalyurek. Parallel Hypergraph Partitioning for Scientific Computing. In *20th International Parallel and Distributed Processing Symposium (IPDPS’06)*. IEEE Computer Society, 2006. doi:10.1109/IPDPS.2006.1639359.
- 20 Julian Dibbelt, Ben Strasser, and Dorothea Wagner. Customizable Contraction Hierarchies. *ACM Journal of Experimental Algorithmics*, 21(1):1.5:1–1.5:49, April 2016. doi:10.1145/2886843.
- 21 Yefim Dinitz. Algorithm for Solution of a Problem of Maximum Flow in a Network with Power Estimation. *Soviet Mathematics-Doklady*, 11(5):1277–1280, September 1970.
- 22 Jack Edmonds and Richard M. Karp. Theoretical Improvements in Algorithmic Efficiency for Network Flow Problems. *Journal of the ACM*, 19(2):248–264, April 1972. doi:10.1145/321694.321699.

- 23 Charles M. Fiduccia and Robert M. Mattheyses. A linear-time heuristic for improving network partitions. In *Proceedings of the 19th ACM/IEEE Conference on Design Automation*, pages 175–181, 1982. doi:10.1145/800263.809204.
- 24 Lester R. Ford, Jr. and Delbert R. Fulkerson. Maximal flow through a network. *Canadian Journal of Mathematics*, 8:399–404, 1956.
- 25 Andrew V. Goldberg and Robert E. Tarjan. A new approach to the maximum-flow problem. *Journal of the ACM*, 35(4):921–940, 1988. doi:10.1145/48014.61051.
- 26 Lars Gottesbüren, Michael Hamann, and Dorothea Wagner. Evaluation of a Flow-Based Hypergraph Bipartitioning Algorithm. Technical report, Institute for Theoretical Informatics, Karlsruhe Institute of Technology, 2019. arXiv:1907.02053.
- 27 Michael Hamann and Ben Strasser. Graph Bisection with Pareto Optimization. *ACM Journal of Experimental Algorithmics*, 23(1):1.2:1–1.2:34, 2018. doi:10.1145/3173045.
- 28 Bruce Hendrickson and Robert Leland. A Multilevel Algorithm for Partitioning Graphs. In *Proceedings of the 1995 ACM/IEEE conference on Supercomputing (SC’05)*, page 28. ACM Press, 1995. doi:10.1145/224170.224228.
- 29 Tobias Heuer, Peter Sanders, and Sebastian Schlag. Network Flow-Based Refinement for Multilevel Hypergraph Partitioning. In *Proceedings of the 17th International Symposium on Experimental Algorithms (SEA’18)*, Leibniz International Proceedings in Informatics, pages 1–19, 2018. doi:10.4230/LIPIcs.SEA.2018.1.
- 30 Tobias Heuer and Sebastian Schlag. Improving Coarsening Schemes for Hypergraph Partitioning by Exploiting Community Structure. In Costas S. Iliopoulos, Solon P. Pissis, Simon J. Puglisi, and Rajeev Raman, editors, *Proceedings of the 16th International Symposium on Experimental Algorithms (SEA’17)*, volume 75 of *Leibniz International Proceedings in Informatics*, 2017. doi:10.4230/LIPIcs.SEA.2017.21.
- 31 George Karypis, Rajat Aggarwal, Vipin Kumar, and Shashi Shekhar. Multilevel Hypergraph Partitioning: Applications in VLSI Domain. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 7(1):69–79, 1999. doi:10.1109/92.748202.
- 32 George Karypis and Vipin Kumar. Multilevel K-Way Hypergraph Partitioning. In *Proceedings of the 36th Annual ACM/IEEE Design Automation Conference*, 1999. doi:10.1145/309847.309954.
- 33 Brian W. Kernighan and Shen Lin. An Efficient Heuristic Procedure for Partitioning Graphs. *Bell System Technical Journal*, 49(2):291–307, February 1970. doi:10.1002/j.1538-7305.1970.tb01770.x.
- 34 Eugene L. Lawler. Cutsets and Partitions of hypergraphs. *Networks*, 3:275–285, 1973. doi:10.1002/net.3230030306.
- 35 Thomas Lengauer. *Combinatorial Algorithms for Integrated Circuit Layout*. Wiley, 1990.
- 36 Jianmin Li, John Lillis, and C.K. Cheng. Linear decomposition algorithm for VLSI design applications. In *Proceedings of the 1995 IEEE/ACM International Conference on Computer-Aided Design*, pages 223–228. IEEE Computer Society, 1995. doi:10.1109/ICCAD.1995.480016.
- 37 Huiqun Liu and D.F. Wong. Network-Flow-Based Multiway Partitioning with Area and Pin Constraints. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 17(1):50–59, 1998. doi:10.1109/43.673632.
- 38 Henning Meyerhenke, Burkhard Monien, and Thomas Sauerwald. A new diffusion-based multilevel algorithm for computing graph partitions. *Journal of Parallel and Distributed Computing*, 69(9):750–761, 2009. doi:10.1016/j.jpdc.2009.04.005.
- 39 Viswanath Nagarajan, Charles J. Alpert, Cliff Sze, Zhuo Li, and Yaoguang Wei. The DAC 2012 Routability-driven Placement Contest and Benchmark Suite. In *Proceedings of the 49th Annual Design Automation Conference*, 2012. doi:10.1145/2228360.2228500.
- 40 David A. Papa and Igor L. Markov. Hypergraph Partitioning and Clustering. In Teofilo F. Gonzalez, editor, *Handbook of Approximation Algorithms and Metaheuristics*. Chapman and Hall/CRC, 2007. doi:10.1201/9781420010749.

- 41 Jean-Claude Picard and Maurice Queyranne. On the Structure of All Minimum Cuts in a Network and Applications. *Mathematical Programming, Series A*, 22(1):121, December 1982. doi:10.1007/BF01581031.
- 42 Joachim Pistorius and Michel Minoux. An Improved Direct Labeling Method for the Max-Flow Min-Cut Computation in Large Hypergraphs and Applications. *International Transactions in Operational Research*, 10(1):1–11, 2003. doi:10.1111/1475-3995.00389.
- 43 Peter Sanders and Christian Schulz. Engineering Multilevel Graph Partitioning Algorithms. In *Proceedings of the 19th Annual European Symposium on Algorithms (ESA'11)*, volume 6942 of *Lecture Notes in Computer Science*, pages 469–480. Springer, 2011. doi:10.1007/978-3-642-23719-5_40.
- 44 Peter Sanders and Christian Schulz. Think Locally, Act Globally: Highly Balanced Graph Partitioning. In *Proceedings of the 12th International Symposium on Experimental Algorithms (SEA'13)*, volume 7933 of *Lecture Notes in Computer Science*, pages 164–175. Springer, 2013. doi:10.1007/978-3-642-38527-8_16.
- 45 Sebastian Schlag, Vitali Henne, Tobias Heuer, Henning Meyerhenke, Peter Sanders, and Christian Schulz. k-way Hypergraph Partitioning via n-Level Recursive Bisection. In *Proceedings of the 18th Meeting on Algorithm Engineering and Experiments (ALENEX'16)*, pages 53–67. SIAM, 2016. doi:10.1137/1.9781611974317.5.
- 46 Ruslan Shaydulín, Jie Chen, and Ilya Safro. Relaxation-Based Coarsening for Multilevel Hypergraph Partitioning. *Multiscale Modeling and Simulation*, 17(1):482–506, 2019. doi:10.1137/17M1152735.
- 47 A. J. Soper, Chris Walshaw, and Mark Cross. The Graph Partitioning Archive, 2004. URL: <http://chriswalshaw.co.uk/partition/>.
- 48 Aleksandar Trifunovic and William J. Knottenbelt. Parallel Multilevel Algorithms for Hypergraph Partitioning. *Journal of Parallel and Distributed Computing*, 68(5):563–581, 2008. doi:10.1016/j.jpdc.2007.11.002.
- 49 Brendan Vastenhouw and Rob Bisseling. A Two-Dimensional Data Distribution Method for Parallel Sparse Matrix-Vector Multiplication. *SIAM Review*, 47(1):67–95, 2005. doi:10.1137/S0036144502409019.
- 50 Chris Walshaw. Multilevel Refinement for Combinatorial Optimisation Problems. *Annals of Operations Research*, 131(1):325–372, October 2004. doi:10.1023/B:ANOR.0000039525.80601.15.
- 51 Hannah Honghua Yang and D.F. Wong. Efficient Network Flow Based Min-Cut Balanced Partitioning. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 15(12):1533–1540, 1996. doi:10.1007/978-1-4615-0292-0_41.