

Multilevel Algorithms for Acyclic Partitioning of Directed Acyclic Graphs

Julien Herrmann, Yusuf Özkaya, Bora Uçar, Kamer Kaya, Umit Catalyurek

► **To cite this version:**

Julien Herrmann, Yusuf Özkaya, Bora Uçar, Kamer Kaya, Umit Catalyurek. Multilevel Algorithms for Acyclic Partitioning of Directed Acyclic Graphs. *SIAM Journal on Scientific Computing, Society for Industrial and Applied Mathematics*, 2019, 41 (4), pp.A2117-A2145. 10.1137/18M1176865 . hal-02306566

HAL Id: hal-02306566

<https://hal.inria.fr/hal-02306566>

Submitted on 11 Oct 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

1 **MULTILEVEL ALGORITHMS FOR ACYCLIC PARTITIONING OF**
2 **DIRECTED ACYCLIC GRAPHS***

3 JULIEN HERRMANN[†], M. YUSUF ÖZKAYA[†], BORA UÇAR[‡],
4 KAMER KAYA[§], AND ÜMIT V. ÇATALYÜREK[†]

5 **Abstract.** We investigate the problem of partitioning the vertices of a directed acyclic graph
6 into a given number of parts. The objective function is to minimize the number or the total weight
7 of the edges having end points in different parts, which is also known as edge cut. The standard load
8 balancing constraint of having an equitable partition of the vertices among the parts should be met.
9 Furthermore, the partition is required to be *acyclic*, i.e., the inter-part edges between the vertices
10 from different parts should preserve an acyclic dependency structure among the parts. In this work,
11 we adopt the multilevel approach with coarsening, initial partitioning, and refinement phases for
12 acyclic partitioning of directed acyclic graphs. We focus on two-way partitioning (sometimes called
13 bisection), as this scheme can be used in a recursive way for multi-way partitioning. To ensure
14 the acyclicity of the partition at all times, we propose novel and efficient coarsening and refinement
15 heuristics. The quality of the computed acyclic partitions is assessed by computing the edge cut.
16 We also propose effective ways to use the standard undirected graph partitioning methods in our
17 multilevel scheme. We perform a large set of experiments on a dataset consisting of (i) graphs
18 coming from an application and (ii) some others corresponding to matrices from a public collection.
19 We report significant improvements compared to the current state of the art.

20 **Key words.** directed graph, acyclic partitioning, multilevel partitioning

21 **AMS subject classifications.** 05C70, 05C85, 68R10, 68W05

22 **1. Introduction.** The standard graph partitioning (GP) problem asks for a
23 partition of the vertices of an undirected graph into a number of parts. The objective
24 and the constraint of this well-known problem are to minimize the number of edges
25 having vertices in two different parts and to equitably partition the vertices among
26 the parts. The GP problem is NP-complete [13, ND14]. We investigate a variant of
27 this problem, called *acyclic partitioning*, for directed acyclic graphs. In this variant,
28 we have one more constraint: the partition should be acyclic. In other words, for a
29 suitable numbering of the parts, all edges should be directed from a vertex in a part
30 p to another vertex in a part q where $p \leq q$.

31 The directed acyclic graph partitioning (DAGP) problem arises in many appli-
32 cations. The stated variant of the DAGP problem arises in exposing parallelism in
33 automatic differentiation [6, Ch.9], and particularly in the computation of the Newton
34 step for solving nonlinear systems [4, 5]. The DAGP problem with some additional
35 constraints is used to reason about the parallel data movement complexity and to dy-
36 namically analyze the data locality potential [10, 11]. Other important applications
37 of the DAGP problem include (i) fusing loops for improving temporal locality, and en-
38 abling streaming and array contractions in runtime systems [19], such as Bohrium [20];
39 (ii) analysis of cache efficient execution of streaming applications on uniprocessors [1];
40 (iii) a number of circuit design applications in which the signal directions impose
41 acyclic partitioning requirement [7, 29].

42 Let us consider a toy example shown in Figure 1.1(a). A partition of the vertices

*A preliminary version appeared in CCGRID'17 [15].

[†]School of Computational Science and Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332-0250, julien.herrmann@cc.gatech.edu, myozka@gatech.edu, umit@gatech.edu.

[‡]CNRS and LIP (UMR5668 Université de Lyon - CNRS - ENS Lyon - Inria - UCBL 1), 46, allée d'Italie, ENS Lyon, 69364, France, bora.ucar@ens-lyon.fr.

[§]Sabancı University, Istanbul, Turkey, kaya@sabanciuniv.edu.




Fig. 1.1: **a)** A toy example with six tasks and six dependencies, **b)** a non-acyclic partitioning when edges are oriented, **c)** an acyclic partitioning of the same directed graph.

43 of this graph is shown in [Figure 1.1\(b\)](#) with a dashed curve. Since there is a cut edge
 44 from s to u and another from u to t , the partition is cyclic, and is not acceptable. An
 45 acyclic partition is shown in [Figure 1.1\(c\)](#), where all the cut edges are from one part
 46 to the other.

47 We adopt the multilevel partitioning approach [2, 14] with the coarsening, initial
 48 partitioning, and refinement phases for acyclic partitioning of DAGs. We propose
 49 heuristics for these three phases ([Subsections 4.1, 4.2 and 4.3](#), respectively) which
 50 guarantee acyclicity of the partitions at all phases and maintain a DAG at every
 51 level. We strived to have fast heuristics at the core. With these characterizations,
 52 the coarsening phase requires new algorithmic/theoretical reasoning, while the initial
 53 partitioning and refinement heuristics are direct adaptations of the standard methods
 54 used in undirected graph partitioning, with some differences worth mentioning. We
 55 discuss only the bisection case, as we were able to improve the direct k -way algorithms
 56 we proposed before [15] by using the bisection heuristics recursively—we give a brief
 57 comparison in [Subsection 5.4](#).

58 The acyclicity constraint on the partitions precludes the use of the state of the
 59 art undirected graph partitioning tools. This has been recognized before, and those
 60 tools were put aside [15, 21]. While this is sensible, one can still try to make use of the
 61 existing undirected graph partitioning tools [14, 16, 25, 27], as they have been very
 62 well engineered. Let us assume that we have partitioned a DAG with an undirected
 63 graph partitioning tool into two parts by ignoring the directions. It is easy to detect
 64 if the partition is cyclic since all the edges need to go from part one to part two.
 65 Furthermore, we can easily fix it as follows. Let v be a vertex in the second part;
 66 we can move all u vertices for which there is a path from v to u into the second
 67 part. This procedure breaks any cycle containing v and hence, the partition becomes
 68 acyclic. However, the edge cut may increase, and the partitions can be unbalanced.
 69 To solve the balance problem and reduce the cut, we can apply a restricted version
 70 of the move-based refinement algorithms in the literature. After this step, this final
 71 partition meets the acyclicity and balance conditions. Depending on the structure
 72 of the input graph, it could also be a good initial partition for reducing the edge
 73 cut. Indeed, one of our most effective schemes uses an undirected graph partitioning
 74 algorithm to create a (potentially cyclic) partition, fixes the cycles in the partition,
 75 and refines the resulting acyclic partition with a novel heuristic to obtain an initial
 76 partition. We then integrate this partition within the proposed coarsening approaches
 77 to refine it at different granularities. We elaborate on this scheme in [Subsection 4.4](#).

78 The rest of the paper is organized as follows: [Section 2](#) introduces the notation

79 and background on directed acyclic graph partitioning and Section 3 briefly surveys
 80 the existing literature. We propose multilevel partitioning heuristics for acyclic par-
 81 titioning of directed acyclic graphs in Section 4. Section 5 presents the experimental
 82 results, and Section 6 concludes the paper.

83 **2. Preliminaries and notation.** A *directed graph* $G = (V, E)$ contains a set of
 84 vertices V and a set of directed edges E of the form $e = (u, v)$, where e is directed
 85 from u to v . A *path* is a sequence of edges $(u_1, v_1) \cdot (u_2, v_2), \dots$ with $v_i = u_{i+1}$. A path
 86 $((u_1, v_1) \cdot (u_2, v_2) \cdot (u_3, v_3) \cdots (u_\ell, v_\ell))$ is of length ℓ , where it connects a sequence of
 87 $\ell + 1$ vertices $(u_1, v_1 = u_2, \dots, v_{\ell-1} = u_\ell, v_\ell)$. A path is called *simple* if the connected
 88 vertices are distinct. Let $u \rightsquigarrow v$ denote a simple path that starts from u and ends at
 89 v . A path $((u_1, v_1) \cdot (u_2, v_2) \cdots (u_\ell, v_\ell))$ forms a (simple) *cycle* if all v_i for $1 \leq i \leq \ell$
 90 are distinct and $u_1 = v_\ell$. A *directed acyclic graph*, DAG in short, is a directed graph
 91 with no cycles.

92 The path $u \rightsquigarrow v$ represents a dependency of v to u . We say that the edge (u, v)
 93 is *redundant* if there exists another $u \rightsquigarrow v$ path in the graph. That is, when we
 94 remove a redundant (u, v) edge, u remains to be connected to v , and hence, the
 95 dependency information is preserved. We use $\text{Pred}[v] = \{u \mid (u, v) \in E\}$ to represent
 96 the (immediate) predecessors of a vertex v , and $\text{Succ}[v] = \{u \mid (v, u) \in E\}$ to represent
 97 the (immediate) successors of v . We call the neighbors of a vertex v , its immediate
 98 predecessors and immediate successors: $\text{Neigh}[u] = \text{Pred}[v] \cup \text{Succ}[v]$. For a vertex u ,
 99 the set of vertices v such that $u \rightsquigarrow v$ are called the *descendants* of u . Similarly, the
 100 set of vertices v such that $v \rightsquigarrow u$ are called the *ancestors* of the vertex u . We will
 101 call vertices without any predecessors (and hence ancestors) as the *sources* of G , and
 102 vertices without any successors (and hence descendants) as the *targets* of G . Every
 103 vertex u has a weight denoted by w_u and every edge $(u, v) \in E$ has a cost denoted by
 104 $c_{u,v}$.

105 A k -way partitioning of a graph $G = (V, E)$ divides V into k disjoint subsets
 106 $\{V_1, \dots, V_k\}$. The weight of a part V_i denoted by $w(V_i)$ is equal to $\sum_{u \in V_i} w_u$, which
 107 is the total vertex weight in V_i . Given a partition, an edge is called a *cut edge* if its
 108 endpoints are in different parts. The *edge cut* of a partition is defined as the sum of
 109 the costs of the cut edges. Usually, a constraint on the part weights accompanies the
 110 problem. We are interested in acyclic partitions, which are defined below.

111 **DEFINITION 2.1 (Acyclic k -way partition).** A partition $\{V_1, \dots, V_k\}$ of $G =$
 112 (V, E) is called an *acyclic k -way partition* if two paths $u \rightsquigarrow v$ and $v' \rightsquigarrow u'$ do not
 113 co-exist for $u, u' \in V_i$, $v, v' \in V_j$, and $1 \leq i \neq j \leq k$.

114 There is a related definition in the literature [11], which is called a *convex par-*
 115 *tion*. A partition is *convex* if for all vertex pairs u, v in the same part, the vertices
 116 in any $u \rightsquigarrow v$ path are also in the same part. Hence, if a partition is acyclic it is also
 117 convex. On the other hand, convexity does not imply acyclicity. Figure 2.1 shows
 118 that the definitions of an acyclic partition and a convex partition are not equivalent.
 119 For the toy graph in Figure 2.1(a), there are three possible balanced partitions shown
 120 in Figure 2.1(b), Figure 2.1(c), and Figure 2.1(d). They are all convex, but only the
 121 one in Figure 2.1(d) is acyclic.

122 Deciding on the existence of a k -way acyclic partition respecting an upper bound
 123 on the part weights and an upper bound on the cost of cut edges is NP-complete [13].
 124 The formal problem treated in this paper is defined as follows.

125 **DEFINITION 2.2 (DAG partitioning problem).** *Given a DAG $G = (V, E)$ an im-*
 126 *balance parameter ε , find an acyclic k -way partition $P = \{V_1, \dots, V_k\}$ of V such that*




Fig. 2.1: A toy graph (left), two cyclic and convex partitions (middle two), and an acyclic and convex partition (right).

127 *the balance constraints*

$$128 \quad (2.1) \quad w(V_i) \leq (1 + \varepsilon) \frac{\sum_{v \in V} w_v}{k}$$

129 *are satisfied for $1 \leq i \leq k$, and the edge cut is minimized.*

130 **3. Related work.** Fauzia et al. [11] propose a heuristic for the acyclic partitioning
 131 problem to optimize data locality when analyzing DAGs. To create partitions,
 132 the heuristic categorizes a vertex as ready to be assigned to a partition when all of
 133 the vertices it depends on have already been assigned. Vertices are assigned to the
 134 current partition set until the maximum number of vertices that would be “active”
 135 during the computation of the part reaches a specified limit, which is the cache size
 136 in their application. This implies that part sizes are not limited by the sum of the
 137 total vertex weights but is a complex function that depends on an external schedule
 138 (order) of the vertices. This differs from our problem as we limit the size of each part
 139 by the total sum of the weights of the vertices on that part.

140 Kernighan [17] proposes an algorithm to find a minimum edge-cut partition of
 141 the vertices of a graph into subsets of size greater than a lower bound and inferior
 142 to an upper bound. The partition needs to use a fixed vertex sequence that cannot
 143 be changed. Indeed, Kernighan’s algorithm takes a topological order of the vertices
 144 of the graph as an input and partitions the vertices such that all vertices in a subset
 145 constitute a continuous block in the given topological order. This procedure is optimal
 146 for a given, fixed topological order and has a run time proportional to the number
 147 of edges in the graph, if the part weights are taken as constant. We used a modified
 148 version of this algorithm as a heuristic in the earlier version of our work [15].

149 Cong et al. [7] describe two approaches for obtaining acyclic partitions of di-
 150 rected Boolean networks, modeling circuits. The first one is a single-level Fiduccia-
 151 Mattheyses (FM)-based approach. In this approach, Cong et al. generate an initial
 152 acyclic partition by splitting the list of the vertices (in a topological order) from left
 153 to right into k parts such that the weight of each part does not violate the bound.
 154 The quality of the results is then improved with a k -way variant of the FM heuristic
 155 [12] taking the acyclicity constraint into account. Our previous work [15] employs
 156 a similar refinement heuristic. The second approach of Cong et al. is a two-level
 157 heuristic; the initial graph is first clustered with a special decomposition, and then it
 158 is partitioned using the first heuristic.

159 In a recent paper [21], Moreira et al. focus on an imaging and computer vision
 160 application on embedded systems and discuss acyclic partitioning heuristics. They

161 propose a single level approach in which an initial acyclic partitioning is obtained
 162 using a topological order. To refine the partitioning, they proposed four local search
 163 heuristics which respect the balance constraint and maintain the acyclicity of the
 164 partition. Three heuristics pick a vertex and move it to an eligible part if and only if
 165 the move improves the cut. These three heuristics differ in choosing the set of eligible
 166 parts for each vertex; some are very restrictive, and some allow arbitrary target parts
 167 as long as acyclicity is maintained. The fourth heuristic tentatively realizes the moves
 168 that increase the cut in order to escape from a possible local minima. It has been
 169 reported that this heuristic delivers better results than the others. In a follow-up
 170 paper, Moreira et al. [22] discuss a multilevel graph partitioner and an evolutionary
 171 algorithm based on this multilevel scheme. Their multilevel scheme starts with a
 172 given acyclic partition. Then, the coarsening phase contracts edges that are in the
 173 same part until there is no edge to contract. Here, matching-based heuristics from
 174 undirected graph partitioning tools are used without taking the directions of the
 175 edges into account. Therefore, the coarsening phase can create cycles in the graph;
 176 however the induced partitions are never cyclic. Then, an initial partition is obtained,
 177 which is refined during the uncoarsening phase with move-based heuristics. In order
 178 to guarantee acyclic partitions, the vertices that lie in cycles are not moved. In a
 179 systematic evaluation of the proposed methods, Moreira et al. note that there are
 180 many local minima and suggest using relaxed constraints in the multilevel setting.
 181 The proposed methods have high run time, as the evolutionary method of Moreira
 182 et al. is not concerned with this issue. Improvements with respect to the earlier
 183 work [21] are reported.

184 Previously, we had developed a multilevel partitioner [15]. In this paper, we
 185 propose methods to use an undirected graph partitioner to guide the multilevel par-
 186 titioner. We focus on partitioning the graph in two parts since we can handle the
 187 general case with a recursive bisection scheme. We also propose new coarsening, ini-
 188 tial partitioning, and refinement methods specifically designed for the 2-partitioning
 189 problem. Our multilevel scheme maintains acyclic partitions and graphs through all
 190 the levels.

191 Other related work on acyclic partitioning of directed graphs include an exact,
 192 branch-and-bound algorithm by Nossack and Pesch [23] which works on the integer
 193 programming formulation of the acyclic partitioning problem. This solution is, of
 194 course, too costly to be used in practice. Wong et al. [29] present a modification of
 195 the decomposition of Cong et al. [7] for clustering, and use this in a two-level scheme.

196 **4. Directed multilevel graph partitioning.** We propose a new multilevel tool
 197 for obtaining acyclic partitions of directed acyclic graphs. Multilevel schemes [2, 14]
 198 form the de-facto standard for solving graph and hypergraph partitioning problems
 199 efficiently, and used by almost all current state-of-the-art partitioning tools [3, 14, 16,
 200 25, 27]. Similar to other multilevel schemes, our tool has three phases: the coarsening
 201 phase, which reduces the number of vertices by clustering them; the initial partitioning
 202 phase, which finds a partition of the coarsest graph; and the uncoarsening phase, in
 203 which the initial partition is projected to the finer graphs and refined along the way,
 204 until a solution for the original graph is obtained.

205 **4.1. Coarsening.** In this phase, we obtain smaller DAGs by coalescing the ver-
 206 tices, level by level. This phase continues until the number of vertices becomes smaller
 207 than a specified bound or the reduction on the number of vertices from one level to the
 208 next one is lower than a threshold. At each level ℓ , we start with a finer acyclic graph
 209 G_ℓ , compute a valid clustering \mathcal{C}_ℓ ensuring the acyclicity, and obtain a coarser acyclic

210 graph $G_{\ell+1}$. While our previous work [15] discussed matching based algorithms for
 211 coarsening, we present agglomerative clustering based variants here. The new vari-
 212 ants supersede the matching based ones. Unlike the standard undirected graph case,
 213 in DAG partitioning, not all vertices can be safely combined. Consider a DAG with
 214 three vertices a, b, c and three edges $(a, b), (b, c), (a, c)$. Here, the vertices a and c
 215 cannot be combined, since that would create a cycle. We say that a set of vertices is
 216 contractible (all its vertices are matchable), if unifying them does not create a cycle.
 217 We now present a general theory about finding clusters without forming cycles, after
 218 giving some definitions.

219 **DEFINITION 4.1 (Clustering).** *A clustering of a DAG is a set of disjoint subsets*
 220 *of vertices. Note that we do not make any assumptions on whether the subsets are*
 221 *connected or not.*

222 **DEFINITION 4.2 (Coarse graph).** *Given a DAG G and a clustering C of G , we*
 223 *let $G_{|C}$ denote the coarse graph created by contracting all sets of vertices of C .*

224 The vertices of the coarse graph are the clusters in C . If $(u, v) \in G$ for two
 225 vertices u and v that are located in different clusters of C then $G_{|C}$ has an (directed)
 226 edge from the vertex corresponding to u 's cluster, to the vertex corresponding to v 's
 227 cluster.

228 **DEFINITION 4.3 (Feasible clustering).** *A feasible clustering C of a DAG G is*
 229 *a clustering such that $G_{|C}$ is acyclic.*

230 **THEOREM 4.1.** *Let $G = (V, E)$ be a DAG. For $u, v \in V$ and $(u, v) \in E$, the coarse*
 231 *graph $G_{|\{(u,v)\}}$ is acyclic if and only if there is no path from u to v in G avoiding the*
 232 *edge (u, v) .*

233 *Proof.* Let $G' = (V', E') = G_{|\{(u,v)\}}$ be the coarse graph, and w be the merged,
 234 coarser vertex of G' corresponding to $\{u, v\}$.

235 If there is a path from u to v in G avoiding the edge (u, v) , then all the edges of
 236 this path are also in G' , and the corresponding path in G' goes from w to w , creating
 237 a cycle.

238 Assume that there is a cycle in the coarse graph G' . This cycle has to pass through
 239 w ; otherwise, it must be in G which is impossible by the definition of G . Thus, there
 240 is a cycle from w to w in the coarse graph G' . Let $a \in V'$ be the first vertex visited
 241 by this cycle after w and $b \in V'$ be the last one, just before completing the cycle. Let
 242 \mathbf{p} be an $a \rightsquigarrow b$ path in G' such that $(w, a) \cdot \mathbf{p} \cdot (b, w)$ is the said $w \rightsquigarrow w$ cycle in G' .
 243 Note that a can be equal to b and in this case $\mathbf{p} = \emptyset$. By the definition of the coarse
 244 graph G' , $a, b \in V$ and all edges in the path \mathbf{p} are in $E \setminus \{(u, v)\}$. Since we have a
 245 cycle in G' , the following two items must hold:

- 246 • (i) either $(u, a) \in E$ or $(v, a) \in E$, or both; and
- 247 • (ii) either $(b, u) \in E$ or $(b, v) \in E$, or both.

248 Hence, overall we have nine (3×3) cases. Here, we investigate only four of them, as the
 249 “both” conditions in (i) and (ii) can be eliminated easily by the following statements.

- 250 • $(u, a) \in E$ and $(b, u) \in E$ is impossible because otherwise, $(u, a) \cdot \mathbf{p} \cdot (b, u)$
 251 would be a $u \rightsquigarrow u$ cycle in the original graph G .
- 252 • $(v, a) \in E$ and $(b, v) \in E$ is impossible because otherwise, $(v, a) \cdot \mathbf{p} \cdot (b, v)$
 253 would be a $v \rightsquigarrow v$ cycle in the original graph G .
- 254 • $(v, a) \in E$ and $(b, u) \in E$ is impossible because otherwise, $(u, v) \cdot (v, a) \cdot \mathbf{p} \cdot (b, u)$
 255 would be a $u \rightsquigarrow u$ cycle in the original graph G .

256 Thus $(u, a) \in E$ and $(b, v) \in E$. Therefore, $(u, a) \cdot \mathbf{p} \cdot (b, v)$ is a $u \rightsquigarrow v$ path in G
 257 avoiding the edge (u, v) , which concludes the proof. \square

258 **Theorem 4.1** can be extended to a set of vertices by noting that this time all
 259 paths connecting two vertices of the set should contain only the vertices of the set.
 260 The theorem (nor its extension) does not imply an efficient algorithm, as it requires
 261 at least one transitive reduction. Furthermore, it does not describe a condition about
 262 two clusters forming a cycle, even if both are individually contractible. In order to
 263 address both of these issues, we put a constraint on the vertices that can form a
 264 cluster, based on the following definition.

265 **DEFINITION 4.4 (Top level value).** *For a DAG $G = (V, E)$, the top level value*
 266 *of a vertex $u \in V$ is the length of the longest path from a source of G to that vertex.*
 267 *The top level values of all vertices can be computed in a single traversal of the graph*
 268 *with a complexity $O(|V| + |E|)$. We use $\text{top}[u]$ to denote the top level of the vertex u .*

269 The top level value of a vertex is independent of the topological order used for
 270 computation. By restricting the set of edges considered in the clustering to the edges
 271 $(u, v) \in E$ such that $\text{top}[u] + 1 = \text{top}[v]$, we ensure that no cycles are formed by
 272 contracting a unique cluster (the condition identified in **Theorem 4.1** is satisfied). Let
 273 C be a clustering of the vertices. Every edge in a cluster of C being contractible is a
 274 necessary condition for C to be feasible, but not a sufficient one. More restrictions on
 275 the edges of vertices inside the clusters should be found to ensure that C is feasible.
 276 We propose three coarsening heuristics based on clustering sets of more than two
 277 vertices, whose pair-wise top level differences are always zero or one.

278 **4.1.1. Acyclic clustering with forbidden edges.** To have an efficient heuris-
 279 tic, we rely only on static information computable in linear time while searching for
 280 a feasible clustering. As stated in the introduction of this section, we rely on the
 281 top level difference of one (or less) for all vertices in the same cluster, and an addi-
 282 tional condition to ensure that there will be no cycles when a number of clusters are
 283 contracted simultaneously. In **Theorem 4.2**, we give two sufficient conditions for a
 284 clustering to be feasible (that is, the graphs at all levels are DAGs) and prove their
 285 correctness.

286 **THEOREM 4.2 (Correctness of the proposed clustering).** *Let $G = (V, E)$ be a*
 287 *DAG and $C = \{C_1, \dots, C_k\}$ be a clustering. If C is such that:*

- 288 • *for any cluster C_i , for all $u, v \in C_i$, $|\text{top}[u] - \text{top}[v]| \leq 1$,*
- 289 • *for two different clusters C_i and C_j and for all $u \in C_i$ and $v \in C_j$ either*
 290 *$(u, v) \notin E$, or $\text{top}[u] \neq \text{top}[v] - 1$,*

291 *then, the coarse graph $G|_C$ is acyclic.*

292 *Proof.* Let us assume (for the sake of contradiction) that there is a clustering
 293 with the same properties above, but the coarsened graph has a cycle. We pick one
 294 such clustering $C = \{C_1, \dots, C_k\}$ with the minimum number of clusters. Let $t_i =$
 295 $\min\{\text{top}[u], u \in C_i\}$ be the smallest top level value of a vertex of C_i . According to the
 296 properties of C , for every vertex $u \in C_i$, either $\text{top}[u] = t_i$, or $\text{top}[u] = t_i + 1$. Let w_i
 297 be the coarse vertex in $G|_C$ obtained by contracting all vertices in C_i , for $i = 1, \dots, k$.
 298 By the assumption, there is a cycle in $G|_C$, and let \mathbf{c} be one with the minimum length.
 299 This cycle passes through all the w_i vertices. Otherwise, there would be a smaller
 300 cardinality clustering with the properties above and creating a cycle in the coarsened
 301 graph, contradicting the minimal cardinality of C . Let us renumber, without loss of
 302 generality, the w_i vertices such that \mathbf{c} is a $w_1 \rightsquigarrow w_1$ cycle which passes through all
 303 the w_i vertices in the non-decreasing order of the indices. This also renumbers the
 304 clusters accordingly.

305 After renumbering the w_i vertices, for every $i \in \{1, \dots, k\}$, there is a path in $G|_C$

306 from w_i to w_{i+1} . Given the definition of the coarsened graph, for every $i \in \{1, \dots, k\}$
 307 there exists a vertex $u_i \in C_i$, and a vertex $u_{i+1} \in C_{i+1}$ such that there exists a
 308 path $u_i \rightsquigarrow u_{i+1}$ in G . Thus, $\text{top}[u_i] + 1 \leq \text{top}[u_{i+1}]$. According to the second
 309 property, either there is at least one intermediate vertex between u_i and u_{i+1} and
 310 then $\text{top}[u_i] + 1 < \text{top}[u_{i+1}]$; or $\text{top}[u_i] + 1 \neq \text{top}[u_{i+1}]$ and then $\text{top}[u_i] + 1 <$
 311 $\text{top}[u_{i+1}]$. Thus, in any case, $\text{top}[u_i] + 1 < \text{top}[u_{i+1}]$ which can be rewritten as
 312 $\text{top}[u_i] < \text{top}[u_{i+1}] - 1$.

313 By definition, we know that $t_i \leq \text{top}[u_i]$ and $\text{top}[u_{i+1}] - 1 \leq t_{i+1}$. Thus for every
 314 $i \in \{1, \dots, k\}$, we have $t_i < t_{i+1}$, which leads to the self-contradicting statement
 315 $t_1 < t_{k+1} = t_1$ and concludes the proof. \square

316 The main heuristic based on [Theorem 4.2](#) is described in [Algorithm 1](#). This
 317 heuristic visits all vertices in an order, and adds the visited vertex to a cluster, if
 318 certain criteria are met; if not, the vertex stays as a singleton. When visiting a
 319 singleton vertex, the clusters of its in-neighbors and out-neighbors are investigated,
 320 and the best (according to an objective value) among those meeting the criterion
 321 described in [Theorem 4.2](#) is selected.

322 [Algorithm 1](#) returns the `leader` array of each vertex for the current coarsening
 323 step. Vertices with the same leader form a cluster (and will form a single vertex in
 324 the coarsened graph). For each vertex $u \in V$, `leader[u]` is the id of the representative
 325 vertex for the cluster that will contain u after [Algorithm 1](#). The `leader` table will
 326 be used to build the coarse graph. Any arbitrary vertex in a given cluster can be
 327 used as the leader of this cluster without impacting the rest of the algorithm. At the
 328 beginning, each vertex belongs to a singleton cluster, and `leader[u] = u`. To keep
 329 the track of trivial clusters (singleton vertices), we use an auxiliary `mark` array. The
 330 value `mark[u]` is *false* if u still belongs to a singleton cluster. Otherwise, the value is
 331 set to *true*.

332 For each singleton vertex u , we maintain an auxiliary array `nbbadneighbors` to
 333 keep the number of non-trivial *bad neighbor* clusters. That is to say, the number
 334 of clusters containing a neighbor of u that would violate the second condition of
 335 [Theorem 4.2](#) in case u was put in another cluster. Hence, if u has only one *bad*
 336 *neighbor* cluster, it can only be put into this cluster. For instance in [Figure 4.1\(a\)](#),
 337 at this point of the coarsening, vertex B can only be put in Cluster 1. Otherwise, if
 338 vertex B was matched with one of its other neighbors, the second condition of the
 339 theorem would be violated. Thus, if a vertex has more than one *bad neighbor* in
 340 different clusters, it has to stay as a singleton. For instance in [Figure 4.1\(b\)](#), vertex
 341 B has two bad neighbor clusters and cannot be put in any cluster without violating
 342 the second condition of [Theorem 4.2](#). To check if there exists another bad neighbor
 343 cluster previously formed, we maintain an array `leaderbadneighbor` that keeps the
 344 representative/leader of the first bad neighbor cluster for each vertex. Initially, this
 345 value is set to minus one.

346 In [Algorithm 1](#), the function *ValidNeighbors* selects the *compatible neighbors* of
 347 vertex u , that is the neighbors in clusters that vertex u can join. This selection is
 348 based on the top level difference (to respect the first condition of [Theorem 4.2](#)), the
 349 number of *bad neighbors* of u , and u 's neighbors (to respect the second condition of
 350 [Theorem 4.2](#)), and the size limitation (we do not want a cluster to be bigger than
 351 10% of the total weight of the graph). Then, a best neighbor, *BestNeigh*, according
 352 to an objective value, such as the edge cost, is selected. After setting the leader of
 353 vertex u to the same value as the leader of *BestNeigh*, some bookkeeping is done
 354 for the arrays related to the second condition of [Theorem 4.2](#). More precisely, at

355 Lines 16–22 of Algorithm 1, the neighbors of u are informed about u joining a new
 356 cluster, potentially becoming a bad neighbor. While doing that, the algorithm skips
 357 the vertices v such that $|\text{top}[u] - \text{top}[v]| > 1$, since u cannot form a bad neighbor
 358 cluster for such v . Similarly, if the best neighbor chosen for u was not in a cluster
 359 previously, i.e., was a singleton vertex, the number of *bad neighbors* of its neighbors
 360 are updated (Lines 24–30).




Fig. 4.1: Two examples of acyclic clustering.

361 In our framework, we also implemented the version in the preliminary study [15]
 362 where the size of cluster is limited to two, meaning that it computes a matching of
 363 the vertices.

364 It can be easily seen that Algorithm 1 has a worst case time complexity of $O(|V| +$
 365 $|E|)$. The array top is constructed in $O(|V| + |E|)$ time, and the best, valid neighbor
 366 of a vertex u is found in $O(|\text{Neigh}[u]|)$ time. The neighbors of a vertex are visited at
 367 most once to keep the arrays related to the second condition of Theorem 4.2 up to
 368 date at Lines 16 and 24.

369 **4.1.2. Acyclic clustering with cycle detection.** We now propose a less re-
 370 strictive clustering algorithm to ensure that the acyclicity of the coarse graph is
 371 maintained. As in the previous section, we rely on the top level difference of one
 372 (or less) for all vertices in the same cluster, i.e., for any cluster C_i , for all $u, v \in C_i$,
 373 $|\text{top}[u] - \text{top}[v]| \leq 1$. Knowing this invariant, when a new vertex is added to a cluster,
 374 a cycle-detection algorithm checks that no cycles are formed when all the clusters are
 375 contracted simultaneously. This algorithm does not traverse the entire graph by also
 376 using the fact that the top level difference within a cluster is at most one.

377 From the proof of Theorem 4.2, we know that with a feasible clustering, if adding
 378 a vertex to a cluster whose vertices' top level values are t and $t + 1$ creates a cycle
 379 in the contracted graph, then this cycle goes through only the vertices with top level
 380 values t or $t + 1$. Thus, when considering the addition of a vertex u to a cluster C
 381 containing v , we check potential cycle formations by traversing the graph starting
 382 from u in a breadth-first manner in the *DetectCycle* function used in Algorithm 2.
 383 Let t denote the minimum top level in C . When at a vertex w , we normally add a
 384 successor y of w into the queue, if $|\text{top}(y) - t| \leq 1$; if w is in the same cluster as one
 385 of its predecessors x , we also add x to the queue if $|\text{top}(x) - t| \leq 1$. This function
 386 uses markers to not to visit the same vertex multiple times, returns *true* if at some
 387 point in the traversal a vertex from cluster C is reached, and returns *false*, otherwise.
 388 In the worst-case, this cycle detection algorithm completes a full graph traversal but
 389 in practice, it stops quickly and does not introduce a significant overhead.

390 Here, we propose different clustering strategies. These algorithms consider all
 391 the vertices in the graph, one by one, and put them in a cluster if their top level

Algorithm 1: Clustering with forbidden edges

Data: Directed graph $G = (V, E)$, a traversal order of the vertices in V , a priority on edges

Result: The leader array for the coarsening

```

1 top  $\leftarrow$  CompTopLevels( $G$ )
/* Initialize all the auxiliary data to be used */
2 for  $u \in V$  do
3   mark[ $u$ ]  $\leftarrow$  false // all vertices are marked as singleton
4   leader[ $u$ ]  $\leftarrow$   $u$ 
5   weight[ $u$ ]  $\leftarrow$   $w_u$  // keeps the total weight for each cluster
/* nbbadneighbors[ $u$ ] stores the number of bad clusters for a vertex  $u$ . If
   it exceeds one,  $u$  is left alone (the second condition of Theorem 4.2.). */
6   nbbadneighbors[ $u$ ]  $\leftarrow$  0
7   leaderbadneighbors[ $u$ ]  $\leftarrow$  -1
8 for  $u \in V$  following the traversal order in input do
9   if mark[ $u$ ] then continue
/* The function ValidNeighbors returns the set of valid match candidates
   for  $u$  based on Theorem 4.2. It also checks the threshold for the
   maximum cluster size, and the number of bad neighbor clusters for  $u$ . */
10   $N \leftarrow$  ValidNeighbors( $u, G, nbbadneighbors, leaderbadneighbors, weight$ )
11  if  $N = \emptyset$  then continue
12  BestNeigh  $\leftarrow$  BestNeighbour( $N$ )
13   $\ell \leftarrow$  leader[BestNeigh]
14  leader[ $u$ ]  $\leftarrow$   $\ell$  // assign  $u$  to BestNeigh's cluster
15  weight[ $\ell$ ]  $\leftarrow$  weight[ $\ell$ ] +  $w_u$ 
/* Let the neighbors of  $u$  know that it is not a singleton anymore */
16  for  $v \in$  Neigh[ $u$ ] do
17    if  $|\text{top}[u] - \text{top}[v]| > 1$  then continue //  $u$  cannot form a bad cluster
18    if nbbadneighbors[ $v$ ] = 0 then
19      nbbadneighbors[ $v$ ]  $\leftarrow$  1
20      leaderbadneighbors[ $v$ ]  $\leftarrow$   $\ell$ 
21    else if nbbadneighbors[ $v$ ] = 1 and leaderbadneighbors[ $v$ ]  $\neq$   $\ell$  then
22      nbbadneighbors[ $v$ ]  $\leftarrow$  2 // mark  $v$  as unmatchable
/* If BestNeigh was forming a singleton cluster before  $u$ 's assignment */
23  if mark[BestNeigh] = false then
/* Let BestNeigh's neighbors know that it is not a singleton anymore */
24    for  $v \in$  Neigh[BestNeigh] do
25      if  $|\text{top}[BestNeigh] - \text{top}[v]| > 1$  then continue
26      if nbbadneighbors[ $v$ ] = 0 then
27        nbbadneighbors[ $v$ ]  $\leftarrow$  1 // The first bad neighbor cluster for  $v$ 
28        leaderbadneighbors[ $v$ ]  $\leftarrow$   $\ell$ 
29      else if nbbadneighbors[ $v$ ] = 1 and leaderbadneighbors[ $v$ ]  $\neq$   $\ell$  then
30        nbbadneighbors[ $v$ ]  $\leftarrow$  2 // mark  $v$  as unmatchable
31    mark[BestNeigh]  $\leftarrow$  true // BestNeigh is not a singleton anymore
32  mark[ $u$ ]  $\leftarrow$  true //  $u$  is not a singleton anymore
33 return leader

```

392 differences are at most one and if no cycles are introduced. The clustering algorithms
393 depending on different vertex traversal orders and priority definitions on the adjacent
394 edges are described in Algorithm 2. As Algorithm 1, this algorithm also returns the
395 leader array of each vertex for the current coarsening step. When a vertex is put in a
396 cluster with top level values t and $t + 1$, its *markup* (respectively *markdown*) value is
397 set to *true* if its top level value is t (respectively $t + 1$). Since the worst case complexity
398 of the cycle detection is $O(|V| + |E|)$, the worst case complexity of Algorithm 2 is
399 $O(|V|(|V| + |E|))$. However, the cycle detection stops quickly in practice and the

400 behavior of Algorithm 2 is closer to $O(|V| + |E|)$ as described in Subsection 5.6.

Algorithm 2: Clustering with cycle detection

Data: Directed graph $G = (V, E)$, a traversal order of the vertices in V , a priority on edges
Result: A feasible clustering C of G

```

1 top  $\leftarrow$  CompTopLevels( $G$ )
2 for  $u \in V$  do
3   markup[ $u$ ]  $\leftarrow$  false // if  $u$ 's cluster has a  $v$  with  $\text{top}[v] = \text{top}[u] + 1$ 
4   markdown[ $u$ ]  $\leftarrow$  false // if  $u$ 's cluster has a  $v$  with  $\text{top}[v] = \text{top}[u] - 1$ 
5   leader[ $u$ ]  $\leftarrow$   $u$  // the leader vertex id for  $u$ 's cluster
6 for  $u \in V$  following the traversal order in input do
7   if markup[ $u$ ] or markdown[ $u$ ] then continue
8   for  $v \in \text{Neigh}[u]$  following given priority on edges do
9     if  $(|\text{top}[u] - \text{top}[v]| > 1)$  then continue // we use  $|\text{top}[u] - \text{top}[v]| = 1$ 
10    /* If this is a  $(u, v)$  edge */
11    if  $v \in \text{Succ}[u]$  then
12      if markup[ $v$ ] then continue
13      if DetectCycle( $u, v, G, \text{leader}$ ) then continue
14      leader[ $u$ ]  $\leftarrow$  leader[ $v$ ]
15      markup[ $u$ ]  $\leftarrow$  markdown[ $v$ ]  $\leftarrow$  true
16    /* If this is a  $(v, u)$  edge */
17    if  $v \in \text{Pred}[u]$  then
18      if markdown[ $v$ ] then continue
19      if DetectCycle( $u, v, G, \text{leader}$ ) then continue
20      leader[ $u$ ]  $\leftarrow$  leader[ $v$ ]
21      markdown[ $u$ ]  $\leftarrow$  markup[ $v$ ]  $\leftarrow$  true
22 return leader

```

401 **4.1.3. Hybrid acyclic clustering.** The cycle detection based algorithm can
 402 suffer from quadratic run time for vertices with large in-degrees or out-degrees. To
 403 avoid this, we design a hybrid acyclic clustering which uses the clustering strategy
 404 described in Algorithm 2 by default and switches to the clustering strategy in Al-
 405 gorithm 1 for *large degree* vertices. We define a limit on the degree of a vertex
 406 (typically $\sqrt{|V|}/10$) for calling it *large degree*. When considering an edge (u, v) where
 407 $\text{top}[u] + 1 = \text{top}[v]$, if the degrees of u and v do not exceed the limit, we use the cycle
 408 detection algorithm to determine if we can contract the edge. Otherwise, if the out-
 409 degree of u or the indegree of v is too large, the edge will be contracted if Algorithm 1
 410 allows so. The complexity of this algorithm is in between those of Algorithm 1 and
 411 Algorithm 2 and will likely avoid the quadratic behavior in practice (if not, the degree
 412 parameter can be adapted).

413 **4.2. Initial partitioning.** After the coarsening phase, we compute an initial
 414 acyclic partitioning of the coarsest graph. We present two heuristics. One of them
 415 is akin to the greedy graph growing method used in the standard graph/hypergraph
 416 partitioning methods. The second one uses an undirected partitioning and then fixes
 417 the acyclicity of the partitions. Throughout this section, we use (V_0, V_1) to denote
 418 the bisection of the vertices of the coarsest graph G . The acyclic bisection (V_0, V_1) is
 419 such that there is no edge from the vertices in V_1 to those in V_0 .

420 **4.2.1. Greedy directed graph growing.** One approach to compute a bisection of a directed graph is to design a greedy algorithm that moves vertices from one part to another using local information. Greedy algorithms have shown to be effective for initial partitioning in multilevel schemes in the undirected case. We start with all vertices in V_1 and replace vertices towards V_0 by using heaps. At any time, the vertices that can be moved to V_0 are in the heap. These vertices are those whose all in-neighbors are in V_0 . Initially only the sources are in the heap, and when all the in-neighbors of a vertex v are moved to the first part, v is inserted into the heap. We separate this process into two phases. In the first phase, the key-values of the vertices in the heap are set to the weighted sum of their incoming edges, and the ties are broken in favor of the vertex closer to the first vertex moved. The first phase continues until the first part has more than 0.9 of the maximum allowed weight (modulo the maximum weight of a vertex). In the second phase, the actual gain of a vertex is used. This gain is equal to the sum of the weights of the incoming edges minus the sum of the weights of the outgoing edges. In this phase, the ties are broken in favor of the heavier vertices. The second phase stops as soon as the required balance is obtained. The reason that we separated this heuristic into two phases is that at the beginning, the gains are of no importance, and the more vertices become movable the more flexibility the heuristic has. Yet, towards the end, parts are fairly balanced, and using actual gains can help keeping the cut small.

440 Since the order of the parts is important, we also reverse the roles of the parts, and the directions of the edges. That is, we put all vertices in V_0 , and move the vertices one by one to V_1 , when all out-neighbors of a vertex have been moved to V_1 . The proposed greedy directed graph growing heuristic returns the best of the these two alternatives.

445 **4.2.2. Undirected bisection and fixing acyclicity.** In this heuristic, we partition the coarsest graph as if it were undirected and then move the vertices from one part to another in case the partition was not acyclic. Let (P_0, P_1) denote the (not necessarily acyclic) bisection of the coarsest graph treated as if it were undirected.

449 The proposed approach designates arbitrarily P_0 as V_0 and P_1 as V_1 . One way to fix the cycle is to move all ancestors of the vertices in V_0 to V_0 , thereby guaranteeing that there is no edge from vertices in V_1 to vertices in V_0 , making the bisection (V_0, V_1) acyclic. We do these moves in a reverse topological order, as shown in Algorithm 3. Another way to fix the acyclicity is to move all descendants of the vertices in V_1 to V_1 , again guaranteeing an acyclic partition. We do these moves in a topological order, as shown in Algorithm 4. We then fix the possible unbalance with a refinement algorithm.

457 Note that we can also initially designate P_1 as V_0 and P_0 as V_1 , and again use Algorithms 3 and 4 to fix a potential cycle in two different ways. We try all four of these choices, and return the best partition (essentially returning the best of the four choices to fix the acyclicity of (P_0, P_1)).

461 **4.3. Refinement.** This phase projects the partition obtained for a coarse graph to the next, finer one and refines the partition by vertex moves. As in the standard refinement methods, the proposed heuristic is applied in a number of passes. Within a pass, we repeatedly select the vertex with the maximum move gain among those that can be moved. We tentatively realize this move if the move maintains or improves the balance. Then, the most profitable prefix of vertex moves are realized at the end of the pass. As usual, we allow the vertices move only once in a pass; therefore once a vertex is moved, it is not eligible to move again during the same pass. We use heaps

Algorithm 3: fixAcyclicityUp**Data:** Directed graph $G = (V, E)$ and a bisection $part$ **Result:** An acyclic bisection of G

```

1 for  $u \in G$  (in reverse topological order) do
2   if  $part[u] = 0$  then
3     for  $v \in \text{Pred}[u]$  do
4        $part[v] \leftarrow 0$ 
5 return  $part$ 

```

Algorithm 4: fixAcyclicityDown**Data:** Directed graph $G = (V, E)$ and a bisection $part$ **Result:** An acyclic bisection of G

```

1 for  $u \in G$  (in topological order) do
2   if  $part[u] = 1$  then
3     for  $v \in \text{Succ}[u]$  do
4        $part[v] \leftarrow 1$ 
5 return  $part$ 

```

469 with the gain of moves as the key value, where we keep only movable vertices. We
470 call a vertex *movable*, if moving it to the other part does not create a cyclic partition.
471 As previously done, we use the notation (V_0, V_1) to designate the acyclic bisection
472 with no edge from vertices in V_1 to vertices in V_0 . This means that for a vertex to
473 move from part V_0 to part V_1 , one of the two conditions should be met (i) either all its
474 out-neighbors should be in V_1 ; (ii) or the vertex has no out-neighbors at all. Similarly,
475 for a vertex to move from part V_1 to part V_0 , one of the two conditions should be met
476 (i) either all its in-neighbors should be in V_0 ; (ii) or the vertex has no in-neighbors
477 at all. This is in a sense the adaptation of boundary Fiduccia-Mattheyses [12] (FM)
478 to directed graphs, where the boundary corresponds to the movable vertices. The
479 notion of movability being more restrictive results in an important simplification with
480 respect to the undirected case. The gain of moving a vertex v from V_0 to V_1 is

$$481 \quad (4.1) \quad \sum_{u \in \text{Succ}[v]} w(v, u) - \sum_{u \in \text{Pred}[v]} w(u, v),$$

482 and the negative of this value when moving it from V_1 to V_0 . This means that the gain
483 of vertices are static: once a vertex is inserted in the heap with the key value (4.1),
484 it is never updated. A move could render some vertices unmovable; if they were in
485 the heap, then they should be deleted. Therefore, the heap data structure needs to
486 support insert, delete, and extract max operations only.

487 We have also implemented a swapping based refinement heuristic akin to the
488 boundary Kernighan-Lin [18] (KL), and another one moving vertices only from the
489 maximum loaded part. For graphs with unit weight vertices, we suggest using the
490 boundary FM, and for others we suggest using one pass of boundary KL followed by
491 one pass of the boundary FM that moves vertices only from the maximum loaded
492 part.

493 **4.4. Constraint coarsening and initial partitioning.** There are a number
494 of highly successful, undirected graph partitioning libraries [16, 25, 27]. They are




Fig. 4.2: 8×8 grid graph whose vertices are ordered in a spiral way; a few of the vertices are labeled with their number. All edges are oriented from a lower numbered vertex to a higher numbered one. There is a unique bipartition with 32 vertices in each side. The edges defining the total order are shown in red and blue, except the one from 32 to 33; the cut edges are shown in gray; other internal edges are not shown.

495 not directly usable for our purposes, as the partitions can be cyclic. Fixing such
 496 partitions, by moving vertices to break the cyclic dependencies among the parts, can
 497 increase the edge cut dramatically (with respect to the undirected cut). Consider for
 498 example, the $n \times n$ grid graph, where the vertices are integer positions for $i = 1, \dots, n$
 499 and $j = 1, \dots, n$ and a vertex at (i, j) is connected to (i', j') when $|i - i'| = 1$ or
 500 $|j - j'| = 1$, but not both. There is an acyclic orientation of this graph, called spiral
 501 ordering, as described in Figure 4.2 for $n = 8$. This spiral ordering defines a total
 502 order. When the directions of the edges are ignored, we can have a bisection with
 503 perfect balance by cutting only $n = 8$ edges with a vertical line. This partition is
 504 cyclic; and it can be made acyclic by putting all vertices numbered greater than 32
 505 to the second part. This partition, which puts the vertices 1–32 to the first part and
 506 the rest to the second part, is the unique acyclic bisection with perfect balance for
 507 the associated directed acyclic graph. The edge cut in the directed version is 35 as
 508 seen in the figure (gray edges). In general one has to cut $n^2 - 4n + 3$ edges for $n \geq 8$:
 509 the blue vertices in the border (excluding the corners) have one edge directed to a red
 510 vertex; the interior blue vertices have two such edges; finally, the blue vertex labeled
 511 $n^2/2$ has three such edges.

512 Let us also investigate the quality of the partitions from a more practical stand-
 513 point. We used MeTiS [16] as the undirected graph partitioner on a dataset of 94
 514 matrices (their details are in Section 5). The results are given in Figure 4.3. For
 515 this preliminary experiment, we partitioned the graphs into two with the maximum
 516 allowed load imbalance $\varepsilon = 3\%$. In the experiment, for only two graphs, the output
 517 of MeTiS is acyclic, and the geometric mean of the normalized edge cut is 0.0012.
 518 Figure 4.3(a) shows the normalized edge cut and the load imbalance after fixing the
 519 cycles, while Figure 4.3(b) shows the two measurements after meeting the balance
 520 criteria. A normalized edge cut value is computed by normalizing the edge cut with
 521 respect to the number of edges.

522 In both figures, the horizontal lines mark the geometric mean of the normalized
 523 edge cuts, and the vertical lines mark the 3% imbalance ratio. In Figure 4.3(a), there
 524 are 37 instances in which the load balance after fixing the cycles is feasible. The




Fig. 4.3: Normalized edge cut (normalized with respect to the number of edges), and the balance obtained after using an undirected graph partitioner and fixing the cycles (left), and after ensuring balance with refinement (right).

525 geometric mean of the normalized edge cuts in this subfigure is 0.0045, while in the
 526 other subfigure, it is 0.0049. Fixing the cycles increases the edge cut with respect to
 527 an undirected partitioning, but not catastrophically (only by $0.0045/0.0012 = 3.75$
 528 times in these experiments), and achieving balance after this step increases the cut
 529 only a little (goes to 0.0049 from 0.0045). That is why we suggest using an undirected
 530 graph partitioner, fixing the cycles among the parts, and performing a refinement
 531 based method for load balancing as a good (initial) partitioner.

532 In order to refine the initial partition in a multilevel setting, we propose a scheme
 533 similar to the *iterated multilevel algorithm* used in the existing partitioners [3, 28]. In
 534 this scheme, first a partition P is obtained. Then, the coarsening phase is employed
 535 to match (or to agglomerate) the vertices that were in the same part in P . After
 536 the coarsening, an initial partitioning is freely available by using the partition P
 537 on the coarsest graph. The refinement phase then can work as before. Moreira
 538 et al. [22] use this approach for the directed graph partitioning problem. To be
 539 more concrete, we first use an undirected graph partitioner, then fix the cycles as
 540 discussed in Section 4.2.2, and then refine this acyclic partition for balance with the
 541 proposed refinement heuristics in Subsection 4.3. We then use this acyclic partition for
 542 constraint coarsening and initial partitioning. We expect this scheme to be successful
 543 in graphs with many sources and targets where the sources and targets can lie in any
 544 of the parts while the overall partition is acyclic. On the other hand, if a graph is such
 545 that its balanced acyclic partitions need to put sources in one part and the targets in
 546 another part, then fixing acyclicity may result in moving many vertices. This in turn
 547 will harm the edge cut found by the undirected graph partitioner.

548 **5. Experimental evaluation.** The partitioning tool presented (`dagP`) is imple-
 549 mented in C/C++ programming languages. The experiments are conducted on a
 550 computer equipped with dual 2.1 GHz, Xeon E5-2683 processors and 512GB memory.
 551 The source code and more information is available at [http://tda.gatech.edu/software/
 552 dagP/](http://tda.gatech.edu/software/dagP/).

553 We have performed an extensive evaluation of the proposed multilevel directed

Graph	Parameters	<i>#vertex</i>	<i>#edge</i>	max. deg.	avg. deg.	<i>#source</i>	<i>#target</i>
2mm	P=10, Q=20, R=30, S=40	36,500	62,200	40	1.704	2100	400
3mm	P=10, Q=20, R=30, S=40, T=50	111,900	214,600	40	1.918	3900	400
adi	T=20, N=30	596,695	1,059,590	109,760	1.776	843	28
atax	M=210, N=230	241,730	385,960	230	1.597	48530	230
covariance	M=50, N=70	191,600	368,775	70	1.925	4775	1275
doitgen	P=10, Q=15, R=20	123,400	237,000	150	1.921	3400	3000
durbin	N=250	126,246	250,993	252	1.988	250	249
fdtd-2d	T=20, X=30, Y=40	256,479	436,580	60	1.702	3579	1199
gemm	P=60, Q=70, R=80	1,026,800	1,684,200	70	1.640	14600	4200
gemver	N=120	159,480	259,440	120	1.627	15360	120
gesummv	N=250	376,000	500,500	500	1.331	125250	250
heat-3d	T=40, N=20	308,480	491,520	20	1.593	1280	512
jacobi-1d	T=100, N=400	239,202	398,000	100	1.664	402	398
jacobi-2d	T=20, N=30	157,808	282,240	20	1.789	1008	784
lu	N=80	344,520	676,240	79	1.963	6400	1
ludcmp	N=80	357,320	701,680	80	1.964	6480	1
mvt	N=200	200,800	320,000	200	1.594	40800	400
seidel-2d	M=20, N=40	261,520	490,960	60	1.877	1600	1
symm	M=40, N=60	254,020	440,400	120	1.734	5680	2400
syr2k	M=20, N=30	111,000	180,900	60	1.630	2100	900
syrk	M=60, N=80	594,480	975,240	81	1.640	8040	3240
trisolv	N=400	240,600	320,000	399	1.330	80600	1
trmm	M=60, N=80	294,570	571,200	80	1.939	6570	4800

Table 5.1: Instances from the Polyhedral Benchmark suite (PolyBench).

554 acyclic graph partitioning method on DAG instances coming from two sources. The
555 first set of instances is from the Polyhedral Benchmark suite (PolyBench) [26], whose
556 parameters are listed in Table 5.1. The graphs in the Polyhedral Benchmark suite
557 arise from various linear computation kernels. The parameters in the second column
558 of Table 5.1 represent the size of these computation kernels. For more details, we re-
559 fer the reader to the description of the Polyhedral Benchmark suite (PolyBench) [26].
560 The second set of instances is obtained from the matrices available in the SuiteS-
561 parse Matrix Collection (formerly known as the University of Florida Sparse Matrix
562 Collection) [8]. From this collection, we pick all the matrices satisfying the following
563 properties: listed as binary, square, and has at least 100000 rows and at most 2^{26}
564 nonzeros. There were a total of 95 matrices at the time of experimentation, where
565 two matrices (ids 1514 and 2294) having the same pattern. We discard the duplicate
566 and use the remaining 94 matrices for experiments. For each such matrix, we take
567 the strict upper triangular part as the associated DAG instance, whenever this part
568 has more nonzeros than the lower triangular part; otherwise we take the strict lower
569 triangular part. All edges have unit cost, and all vertices have unit weight.

570 Since the proposed heuristics have a randomized behavior (the traversals used
571 in the coarsening and refinement heuristics are randomized), we run them 10 times
572 for each DAG instance, and report the averages of these runs. We use performance
573 profiles [9] to present the edge-cut results. A performance profile plot shows the
574 probability that a specific method gives results within a factor θ of the best edge cut
575 obtained by any of the methods compared in the plot. Hence, the higher and closer
576 a plot to the y -axis, the better the method is.

577 We set the load imbalance parameter $\varepsilon = 0.03$ in (2.1) for all experiments. The
578 vertices are unit weighted, therefore, the imbalance is rarely an issue for a move-based
579 partitioner.




Fig. 5.1: Performance profiles of the proposed multilevel algorithm variants using three different coarsening heuristics in terms of edge cut.

580 **5.1. Coarsening evaluation.** We first evaluate the proposed coarsening heuristics.
 581 The aim is to find an effective one to set as a default coarsening heuristic.

582 The performance profile chart given in Figure 5.1 shows the effect of the coarsening
 583 heuristics on the final edge cut for the whole dataset. The variants of the proposed
 584 multilevel algorithm which use different coarsening schemes are named as CoTop (Sec-
 585 tion 4.1.1), CoCyc (Section 4.1.2), and CoHyb (Section 4.1.3). Here, and in the rest of
 586 the paper, we used a randomized Depth-First topological order for the node traversal
 587 in the coarsening heuristics, since it performed better in practice. In Figure 5.1, we
 588 see that CoCyc and CoHyb behave similarly; this is expected as not all graphs have
 589 vertices with large degrees. From this figure, we conclude that in general, the coarsening
 590 heuristics CoHyb and CoCyc are more helpful than CoTop in reducing the edge
 591 cut.

592 Another important characteristic to assess for a coarsening heuristic is its con-
 593 traction efficiency. It is important that the coarsening phase does not stop too early
 594 and that the coarsest graph is small enough to be partitioned efficiently. Table 5.2
 595 gives the maximum, the average, and the standard deviation of vertex and edge weight
 596 ratios, and the average, the minimum, and the maximum number of coarsening levels
 597 observed for the two datasets. An effective coarsening heuristic should have small
 598 vertex and edge weight ratios. We see that CoCyc and CoHyb behave similarly and
 599 provide slightly better results than CoTop on both datasets. The graphs from the two
 600 datasets have different characteristics. All coarsening heuristics perform better on the
 601 PolyBench instances compared to the UFL instances: they obtain smaller ratios in
 602 the number of remaining vertices, and yield smaller edge weights. Furthermore, the
 603 maximum vertex and edge weight ratios are smaller in PolyBench instances, again
 604 with all coarsening methods. To the best of our understanding, these happen due to
 605 two reasons; (i) the average degree in the UFL instances is larger than that of the
 606 PolyBench instances (3.63 vs. 1.72); (ii) the ratio of the total number of source and
 607 target vertices to the total number of vertices is again larger in the UFL instances
 608 (0.13 vs. 0.03). Based on Figure 5.1 and Table 5.2, we set CoHyb as the default
 609 coarsening heuristic, as it performs better than CoTop in terms of final edge cut, and
 610 is guaranteed to be more efficient than CoCyc in terms of run time.

611 **5.2. Constraint coarsening and initial partitioning.** We now investigate
 612 the effect of using undirected graph partitioners to obtain a more effective coarsen-

Algorithm	Vertex ratio (%)			Edge weight ratio (%)			Coarsening levels		
	avg	std. dev	max	avg	std. dev	max	avg	min	max
CoTop	1.29	6.34	46.72	26.07	24.95	87.00	12.45	2	17.0
CoCyc	1.06	6.31	47.29	25.97	24.86	87.90	12.74	2	17.6
CoHyb	1.08	6.27	46.70	26.00	24.80	87.00	12.69	2	17.7
CoTop	1.33	2.26	8.50	25.67	11.08	47.60	7.44	4	11.8
CoCyc	0.41	0.90	4.10	24.96	9.20	37.00	8.37	5	12.0
CoHyb	0.54	0.88	3.60	24.81	9.33	39.00	8.46	5	11.9

Table 5.2: The maximum, average, and standard deviation of vertex and edge weight ratios, and the average, the minimum, and the maximum number of coarsening levels for the UFL dataset on the upper half of the table, and for the PolyBench dataset on the lower half.

ing and better initial partitions as explained in [Subsection 4.4](#). We compare three variants of the proposed multilevel scheme. All of them use the refinement described in [Subsection 4.3](#) in the uncoarsening phase.

- **CoHyb**: this variant uses the hybrid coarsening heuristic described in [Section 4.1.3](#) and the greedy directed graph growing heuristic described in [Section 4.2.1](#) in the initial partitioning phase. This method does not use constraint coarsening.
- **CoHyb_C**: this variant uses an acyclic partition of the finest graph obtained as outlined in [Section 4.2.2](#) to guide the hybrid coarsening heuristic described in [Subsection 4.4](#), and uses the greedy directed graph growing heuristic in the initial partitioning phase.
- **CoHyb_CIP**: this variant uses the same constraint coarsening heuristic as the previous method, but inherits the fixed acyclic partition of the finest graph as the initial partitioning.

The comparison of these three variants are given in [Figure 5.2](#) for the whole dataset. From [Figure 5.2](#), we see that using the constraint coarsening is always helpful with respect to not using them. This clearly separates **CoHyb_C** and **CoHyb_CIP** from **CoHyb** after $\theta = 1.1$. Furthermore, applying the constraint initial partitioning (on top of the constraint coarsening) brings tangible improvements.

In the light of the experiments presented here, we suggest the variant **CoHyb_CIP** for general problem instances, as this has clear advantages over others in our dataset.

5.3. Evaluating CoHyb_CIP with respect to a single level algorithm. We compare **CoHyb_CIP** (the variant of the proposed approach with constraint coarsening and initial partitioning) with a single-level algorithm that uses an undirected graph partitioning, fixes the acyclicity, and refines the partitions. This last variant is denoted as **UndirFix**, and it is the algorithm described in [Section 4.2.2](#). Both variants use the same initial partitioning approach, which utilizes MeTiS [16] as the undirected partitioner. The difference between **UndirFix** and **CoHyb_CIP** is the latter’s ability to refine the initial partition at multiple levels. [Figure 5.3](#) presents this comparison. The plots show that the multilevel scheme **CoHyb_CIP** outperforms the single level scheme **UndirFix** at all appropriate ranges of θ , attesting to the importance of the multilevel scheme.

5.4. Comparison with existing work. Here we compare our approach with the evolutionary graph partitioning approach developed by Moreira et al. [21], and




Fig. 5.2: Performance profiles for the edge cut obtained by the proposed multilevel algorithm using the constraint coarsening and partitioning (CoHyb_CIP), using the constraint coarsening and the greedy directed graph growing (CoHyb_C), and the best identified approach without constraints (CoHyb).




Fig. 5.3: Performance profiles for the edge cut obtained by the proposed multilevel approach using the constraint coarsening and partitioning (CoHyb_CIP) and using the same approach without coarsening (UmdirFix).

647 briefly with our previous work [15].

648 Figure 5.4 shows how CoHyb_CIP and CoTop compare with the evolutionary ap-
 649 proach in terms of the edge cut on the 23 graphs of the PolyBench dataset, for the
 650 number of partitions $k \in \{2, 4, 8, 16, 32\}$. We use the average edge cut value of 10
 651 runs for CoTop and CoHyb_CIP and the average values presented in [21] for the evolu-
 652 tionary algorithm. As seen in the figure, the CoTop variant of the proposed multilevel
 653 approach obtains the best results on this specific dataset (all variants of the proposed
 654 approach outperform the evolutionary approach).

655 Tables A.1 and A.2 show the average and best edge cuts found by CoHyb_CIP and
 656 CoTop variants of our partitioner and the evolutionary approach on the PolyBench
 657 dataset. The two tables just after them (Tables A.3 and A.4) give the associated
 658 balance factors. The variants CoHyb_CIP and CoTop of the proposed algorithm obtain
 659 strictly better results than the evolutionary approach in 78 and 75 instances (out of
 660 115), respectively, when the average edge cuts are compared.




Fig. 5.4: Performance profiles for the edge cut obtained by CoHyb_CIP, CoTop, and Moreira et al.’s approach on the PolyBench dataset with $k \in \{2, 4, 8, 16, 32\}$.

661 As seen in the last row of Table A.2, CoHyb_CIP obtains 26% less edge cut than
 662 the evolutionary approach on average (geometric mean) when the average cuts are
 663 compared (0.74 vs. 1.00 in the table); when the best cuts are compared, CoHyb_CIP
 664 obtains 48% less edge cut (0.50 vs. 0.96). Moreover, CoTop obtains 37% less edge cut
 665 than the evolutionary approach when the average cuts are compared (0.63 vs. 1.00
 666 in the table); when the best cuts are compared, CoTop obtains 41% less cut (0.57
 667 vs. 0.96). In some instances (for example `covariance` and `gemm` in Table A.1 and
 668 `syrc` and `trmm` in Table A.2), we see large differences between the average and the
 669 best results of CoTop and CoHyb_CIP. Combined with the observation that CoHyb_CIP
 670 yields better results in general, this suggests that the neighborhood structure can be
 671 improved (see the notion of the strength of a neighborhood [24, Section 19.6]). All
 672 partitions attain 3% balance.

673 The proposed approach with all the reported variants take about 30 minutes to
 674 complete the whole set of experiments for this dataset, whereas the evolutionary ap-
 675 proach is much more compute-intensive, as it has to run the multilevel partitioning
 676 algorithm numerous times to create and update the population of partitions for the
 677 evolutionary algorithm. The multilevel approach of Moreira et al. [21] is more compa-
 678 rable in terms of characteristics with our multilevel scheme. When we compare CoTop
 679 with the results of the multilevel algorithm by Moreira et al., our approach provides
 680 results that are 37% better on average and CoHyb_CIP approach provides results that
 681 are 26% better on average, highlighting the fact that keeping the acyclicity of the
 682 directed graph through the multilevel process is useful.

683 Finally, CoTop and CoHyb_CIP also outperform the previous version of our mul-
 684 tilevel partitioner [15], which is based on a direct k -way partitioning scheme and
 685 matching heuristics for the coarsening phase, by 45% and 35% on average, respec-
 686 tively, on the same dataset.

687 **5.5. Single commodity flow-like problem instances.** In many of the in-
 688 stances of our dataset, graphs have many source and target vertices. We investigate
 689 how our algorithm performs on problems where all source vertices should be in a given
 690 part, and all target vertices should be in the other part, while also achieving balance.
 691 This is a problem close to the maximum flow problem, where we want to find the
 692 maximum flow (or minimum cut) from the sources to the targets with balance on
 693 part weights. Furthermore, addressing this problem also provides a setting for solving




Fig. 5.5: Performance profiles of CoHyb, CoHyb_CIP and UndirFix in terms of edge cut for single source, single target graph dataset. The average of 5 runs are reported for each approach.

694 partitioning problems with fixed vertices.

695 For these experiments, we used the UFL dataset. We discarded all isolated ver-
 696 tices, added to each graph a source vertex S (with an edge from S to all source vertices
 697 of the original graph with a cost equal to the number of edges) and target vertex T
 698 (with an edge from all target vertices of the original graph to T with a cost equal
 699 to the number of edges). A feasible partition should avoid cutting these edges, and
 700 separate all sources from the targets.

701 The performance profiles of CoHyb, CoHyb_CIP and UndirFix are given in Fig-
 702 ure 5.5 with the edge cut as the evaluation criterion. As seen in this figure, CoHyb
 703 is the best performing variant, and UndirFix is the worst performing variant. This is
 704 interesting as in the general setting, we saw a reverse relation. The variant CoHyb_CIP
 705 performs in the middle, as it combines the other two.

706 **5.6. Runtime performance.** We now assess the runtime performance of the
 707 proposed algorithms. Figure 5.6 shows the runtime comparison and distribution for
 708 13 graphs with the longest coarsening time for the CoTop variant. A description of
 709 these 13 graphs can be found in Table 5.3. In Figure 5.6, each graph has three bars
 710 representing the runtime for the multilevel algorithm using the coarsening heuristics
 711 described in Subsection 4.1: CoTop, CoCyc, and CoHyb. We can see that the run time
 712 performance of the three coarsening heuristics are similar. This means that, the cycle
 713 detection function in CoCyc does not introduce a large overhead, as stated in Sec-
 714 tion 4.1.2. Most of the time, CoCyc has a bit longer run time than CoTop, and CoHyb
 715 offers a good tradeoff. Note that in Figure 5.6, the computation time of the initial
 716 partitioning is negligible compared to that of the coarsening and uncoarsening phases,
 717 which means that the graphs have been efficiently contracted during the coarsening
 718 phase.

719 Figure 5.7 shows the comparison of the five variants of the proposed multilevel
 720 scheme and the single level scheme on the whole dataset. Each algorithm is run 10
 721 times on each graph. As expected, CoTop offers the best performance, and CoHyb
 722 offers a good trade-off between CoTop and CoCyc. An interesting remark is that these
 723 three algorithms have a better run time than the single level algorithm UndirFix. For
 724 example, on the average, CoTop is 1.44 times faster than UndirFix. This is mainly due
 725 to cost of fixing acyclicity. Undirected partitioning accounts for roughly 25% of the

Graph	#vertex	#edge	Max In	Max Out	Avg Deg	#source	#target
333SP	3,712,815	11,108,633	9	27	2.992	188,112	316,151
AS365	3,799,275	11,368,076	10	13	2.992	306,791	519,431
M6	3,501,776	10,501,936	10	10	2.999	280,784	472,230
cit-Patents	3,774,768	16,518,209	779	770	4.376	515,980	1,685,419
delaunay-n22	4,194,304	12,582,869	15	17	3	555,807	337,743
hugebubbles-00010	19,458,087	29,179,764	3	3	1.5	3,355,886	3,054,827
hugetrace-00020	16,002,413	23,998,813	3	3	1.5	2,514,461	2,407,017
hugetric-00010	6,592,765	9,885,854	3	3	1.5	1,085,866	1,006,163
italy-osm	6,686,493	7,013,978	5	8	1.049	155,509	458,561
rgg-n-2-22-s0	4,194,304	30,359,198	24	25	7.238	3,550	3,576
road-usa	23,947,347	28,854,312	8	8	1.205	6,392,288	8,010,032
wb-edu	9,845,725	29,494,732	17,489	3841	2.996	1,489,057	2,794,680
wikipedia-20060925	2,983,494	26,103,626	74,970	5,844	8.749	1,406,429	72,744

Table 5.3: 13 instances from the UFL dataset with the longest coarsening times for CoTop.




Fig. 5.6: Runtimes for CoTop, CoCyc, and CoHyb variants of the proposed multilevel scheme. For each bar group, the first, second, and the third bar present the detailed runtimes of CoTop, CoCyc, and CoHyb, respectively.

726 execution time of `UndirFix`, and fixing the acyclicity constitutes the remaining 75%.
 727 Finally, the variants of the multilevel algorithm using constraint coarsening heuristics
 728 provide satisfying run time performance with respect to the others.

729 **6. Conclusion.** We proposed a multilevel approach for acyclic partitioning of
 730 directed acyclic graphs. This problem is close to the standard graph partitioning in
 731 that the aim is to partition the vertices into a number of parts while minimizing the
 732 edge cut and meeting a balance criterion on the part weights. Unlike the standard
 733 graph partitioning problem, the directions of the edges are important and the resulting
 734 partitions should have acyclic dependencies.

735 We proposed coarsening, initial partitioning, and refinement heuristics for the
 736 target problem. The proposed heuristics take the directions of the edges into account
 737 and maintain the acyclicity through all the multilevel hierarchy. We also proposed
 738 efficient and effective approaches to use the standard undirected graph partitioning
 739 tools in the multilevel scheme for coarsening and initial partitioning. We performed
 740 a large set of experiments on a dataset with graphs having different characteristics
 741 and evaluated different combinations of the proposed heuristics. Our experiments
 742 suggested (i) the use of constraint coarsening and initial partitioning, where the main
 743 coarsening heuristic is a hybrid one which avoids the cycles, and in case it does not,
 744 performs a fast cycle detection (`CoHyb_CIP`) for the general case; (ii) a pure multilevel




Fig. 5.7: Runtime performance profile of CoCyc, CoHyb, CoTop, CoHyb_C, CoHyb_CIP and UndirFix on the whole dataset. The values are the averages of 10 runs.

745 scheme without constraint coarsening, using the hybrid coarsening heuristic (CoHyb)
 746 for the cases where a number of sources need to be separated from a number of targets;
 747 (iii) a pure multilevel scheme without constraint coarsening, using the fast coarsening
 748 algorithm (CoTop) for the cases where the degrees of the vertices are small. All three
 749 approaches are shown to be more effective and efficient than the current state of the
 750 art.

751 An avenue for the future work is applying the proposed multilevel scheme in real
 752 life applications that are based on task-graphs. This requires a scheduling step to be
 753 applied after the proposed partitioning scheme, which needs further investigations. A
 754 recent work uses a multilevel algorithm for recombination and mutation [22]. Plugging
 755 in our multilevel scheme to that framework can yield significant improvements.

756 **Acknowledgment.** We thank John Gilbert for his comments on an earlier ver-
 757 sion of this work presented at CSC'16. John suggested that we look at the spiral
 758 ordering of the grid graph.

759

REFERENCES

- 760 [1] K. AGRAWAL, J. T. FINEMAN, J. KRAGE, C. E. LEISERSON, AND S. TOLEDO, *Cache-conscious*
 761 *scheduling of streaming applications*, in Proc. Twenty-fourth Annual ACM Symposium on
 762 *Parallelism in Algorithms and Architectures*, SPAA '12, New York, NY, USA, 2012, ACM,
 763 pp. 236–245.
- 764 [2] T. N. BUI AND C. JONES, *A heuristic for reducing fill-in in sparse matrix factorization*, in Proc.
 765 *6th SIAM Conf. Parallel Processing for Scientific Computing*, SIAM, 1993, pp. 445–452.
- 766 [3] Ü. V. ÇATALYÜREK AND C. AYKANAT, *PaToH: A Multilevel Hypergraph Partitioning Tool,*
 767 *Version 3.0*, Bilkent University, Dept. Comp. Engineering, Ankara, 06533 Turkey. PaToH
 768 is available at <http://cc.gatech.edu/~umit/software.html>, 1999.
- 769 [4] T. F. COLEMAN AND W. XU, *Parallelism in structured Newton computations*, in *Parallel*
 770 *Computing: Architectures, Algorithms and Applications*, ParCo 2007, Forschungszentrum
 771 Jülich and RWTH Aachen University, Germany, 2007, pp. 295–302.
- 772 [5] T. F. COLEMAN AND W. XU, *Fast (structured) Newton computations*, *SIAM Journal on Scien-*
 773 *tific Computing*, 31 (2009), pp. 1175–1191.
- 774 [6] T. F. COLEMAN AND W. XU, *Automatic Differentiation in MATLAB using ADMAT with*
 775 *Applications*, SIAM, 2016.
- 776 [7] J. CONG, Z. LI, AND R. BAGRODIA, *Acyclic multi-way partitioning of Boolean networks*, in
 777 *Proceedings of the 31st Annual Design Automation Conference*, DAC'94, New York, NY,
 778 USA, 1994, ACM, pp. 670–675.
- 779 [8] T. A. DAVIS AND Y. HU, *The University of Florida sparse matrix collection*, *ACM Trans. Math.*

- 780 Softw., 38 (2011), pp. 1:1–1:25.
- 781 [9] E. D. DOLAN AND J. J. MORÉ, *Benchmarking optimization software with performance profiles*,
782 Mathematical programming, 91 (2002), pp. 201–213.
- 783 [10] V. ELANGO, F. RASTELLO, L.-N. POUCHET, J. RAMANUJAM, AND P. SADAYAPPAN, *On charac-*
784 *terizing the data access complexity of programs*, SIGPLAN Not., 50 (2015), pp. 567–580.
- 785 [11] N. FAUZIA, V. ELANGO, M. RAVISHANKAR, J. RAMANUJAM, F. RASTELLO, A. ROUNTEV, L.-N.
786 POUCHET, AND P. SADAYAPPAN, *Beyond reuse distance analysis: Dynamic analysis for*
787 *characterization of data locality potential*, ACM Trans. Archit. Code Optim., 10 (2013),
788 pp. 53:1–53:29.
- 789 [12] C. M. FIDUCCIA AND R. M. MATTHEYSES, *A linear-time heuristic for improving network par-*
790 *titions*, in Design Automation, 1982. 19th Conference on, IEEE, 1982, pp. 175–181.
- 791 [13] M. R. GAREY AND D. S. JOHNSON, *Computers and Intractability: A Guide to the Theory of*
792 *NP-Completeness*, W. H. Freeman & Co., New York, NY, USA, 1979.
- 793 [14] B. HENDRICKSON AND R. LELAND, *The Chaco user’s guide, version 1.0*, Tech. Report SAND93–
794 2339, Sandia National Laboratories, Albuquerque, NM, October 1993.
- 795 [15] J. HERRMANN, J. KHO, B. UÇAR, K. KAYA, AND Ü. V. ÇATALYÜREK, *Acyclic partitioning of*
796 *large directed acyclic graphs*, in Proceedings of the 17th IEEE/ACM International Sym-
797 posium on Cluster, Cloud and Grid Computing, CCGRID, Madrid, Spain, May 2017,
798 pp. 371–380.
- 799 [16] G. KARYPIS AND V. KUMAR, *MeTiS: A Software Package for Partitioning Unstructured*
800 *Graphs, Partitioning Meshes, and Computing Fill-Reducing Orderings of Sparse Matrices*
801 *Version 4.0*, University of Minnesota, Department of Comp. Sci. and Eng., Army HPC
802 Research Cent., Minneapolis, 1998.
- 803 [17] B. W. KERNIGHAN, *Optimal sequential partitions of graphs*, J. ACM, 18 (1971), pp. 34–40.
- 804 [18] B. W. KERNIGHAN AND S. LIN, *An efficient heuristic procedure for partitioning graphs*, The
805 Bell System Technical Journal, 49 (1970), pp. 291–307.
- 806 [19] M. R. B. KRISTENSEN, S. A. F. LUND, T. BLUM, AND J. AVERY, *Fusion of parallel array*
807 *operations*, in Proceedings of the 2016 International Conference on Parallel Architectures
808 and Compilation, New York, NY, USA, 2016, ACM, pp. 71–85.
- 809 [20] M. R. B. KRISTENSEN, S. A. F. LUND, T. BLUM, K. SKOVHEDE, AND B. VINTER, *Bohrium: A*
810 *virtual machine approach to portable parallelism*, in Proceedings of the 2014 IEEE Inter-
811 national Parallel & Distributed Processing Symposium Workshops, IPDPSW ’14, Wash-
812 ington, DC, USA, 2014, IEEE Computer Society, pp. 312–321.
- 813 [21] O. MOREIRA, M. POPP, AND C. SCHULZ, *Graph partitioning with acyclicity constraints*, in 16th
814 International Symposium on Experimental Algorithms, SEA, London, UK, 2017, Schloss
815 Dagstuhl–Leibniz-Zentrum fuer Informatik.
- 816 [22] O. MOREIRA, M. POPP, AND C. SCHULZ, *Evolutionary multi-level acyclic graph partitioning*, in
817 Proceedings of the Genetic and Evolutionary Computation Conference, GECCO, Kyoto,
818 Japan, 2018, ACM, pp. 332–339.
- 819 [23] J. NOSSACK AND E. PESCH, *A branch-and-bound algorithm for the acyclic partitioning problem*,
820 Computers & Operations Research, 41 (2014), pp. 174–184.
- 821 [24] C. H. PAPADIMITRIOU AND K. STEIGLITZ, *Combinatorial Optimization: Algorithms and Com-*
822 *plexity*, Dover Publications, (Corrected, unabridged reprint of Combinatorial Optimization:
823 Algorithms and Complexity originally published by Prentice-Hall Inc., New Jersey, 1982),
824 New York, 1998.
- 825 [25] F. PELLEGRINI, *SCOTCH 5.1 User’s Guide*, Laboratoire Bordelais de Recherche en Informa-
826 tique (LaBRI), 2008.
- 827 [26] L.-N. POUCHET, *Polybench: The polyhedral benchmark suite*, URL: [http://web.cse.ohio-](http://web.cse.ohio-state.edu/pouchet/software/polybench/)
828 state.edu/pouchet/software/polybench/, (2012).
- 829 [27] P. SANDERS AND C. SCHULZ, *Engineering multilevel graph partitioning algorithms*, in Algo-
830 rithms – ESA 2011: 19th Annual European Symposium, Saarbrücken, Germany, Septem-
831 ber 5-9, 2011. Proceedings, C. Demetrescu and M. M. Halldórsson, eds., Berlin, Heidelberg,
832 2011, Springer Berlin Heidelberg, pp. 469–480.
- 833 [28] C. WALSHAW, *Multilevel refinement for combinatorial optimisation problems*, Annals of Oper-
834 ations Research, 131 (2004), pp. 325–372.
- 835 [29] E. S. H. WONG, E. F. Y. YOUNG, AND W. K. MAK, *Clustering based acyclic multi-way parti-*
836 *tioning*, in Proceedings of the 13th ACM Great Lakes Symposium on VLSI, GLSVLSI ’03,
837 New York, NY, USA, 2003, ACM, pp. 203–206.

838 **Appendix A. Detailed results on the PolyBench instances.** We give in
839 Tables A.1 and A.2 the detailed edge cut results of the proposed CoTop, CoHyb_CIP and

840 of Moreira et al.'s evolutionary algorithm [21]. Tables A.3 and A.4 give the balance
841 attained in the partitions. In these two tables, the average balance of the ten runs
842 yielding the average edge cut of Tables A.1 and A.2 is reported per problem instance.
843 The balance of the partition yielding the best edge cut of the previous tables is also
844 given per problem instance.

Graph	k	Moreira et al. [21]		CoHyb_CIP		CoTop	
		Average	Best	Average	Best	Average	Best
2mm	2	200	200	200	200	200	200
	4	947	930	6134	2686	2160	1900
	8	7181	6604	8713	6300	5361	4027
	16	13330	13092	12135	9380	11196	10698
	32	14583	14321	15911	14829	15932	14838
3mm	2	1000	1000	7399	800	1000	1000
	4	38722	37899	16771	7653	9264	8634
	8	58129	49559	24330	9832	28121	24270
	16	64384	60127	37041	31036	39683	37194
	32	62279	58190	46437	43062	48567	43210
adi	2	134945	134675	142719	142174	143067	139672
	4	284666	283892	212938	211939	215399	214945
	8	290823	290672	271949	266349	256302	255522
	16	326963	326923	300755	292351	282485	281511
	32	370876	370413	324494	316241	306075	305411
atax	2	47826	47424	44942	38679	39876	39876
	4	82397	76245	60187	47184	48645	48645
	8	113410	111051	63353	51580	51243	50419
	16	127687	125146	70723	62697	59208	57085
	32	132092	130854	78264	67401	69556	63166
covariance	2	66520	66445	27269	4775	55195	17183
	4	84626	84213	82125	61793	61991	34307
	8	103710	102425	136946	122656	74325	50680
	16	125816	123276	142177	123221	119284	106422
	32	142214	137905	121155	103751	133522	117431
doitgen	2	43807	42208	5035	3000	5947	5947
	4	72115	71072	37767	22290	37051	31157
	8	76977	75114	51283	43572	53244	50795
	16	84203	77436	62296	56650	66483	64488
	32	94135	92739	68350	62576	74786	70168
durbin	2	12997	12997	12997	12997	12997	12997
	4	21641	21641	21572	21572	21566	21566
	8	27571	27571	27519	27518	27520	27520
	16	32865	32865	32852	32848	32912	32912
	32	39726	39725	39738	39732	39826	39826
fdtd-2d	2	5494	5494	6264	6003	6024	5896
	4	15100	15099	15294	13199	16965	16674
	8	33087	32355	23699	21886	35711	34361
	16	35714	35239	32917	30725	44643	43608
	32	43961	42507	42515	41258	53658	52420
gemm	2	383084	382433	4200	4200	44549	44549
	4	507250	500526	168962	12600	59854	46677
	8	578951	575004	183228	36273	116990	96059
	16	615342	613373	294777	241136	263050	238125
	32	626472	623271	330937	307225	332946	299774
gemver	2	29349	29270	26368	22824	20913	20913
	4	49361	49229	45689	38663	40299	40185
	8	68163	67094	56930	49776	55266	53759
	16	78115	75596	62143	57779	59072	56598
	32	85331	84865	75425	68673	73131	71349
gesummv	2	1666	500	24762	500	500	500
	4	98542	94493	24613	1783	10316	8710
	8	101533	98982	25342	13522	9618	9397
	16	112064	104866	37819	21155	35686	30954
	32	117752	114812	48775	42523	45050	40671
heat-3d	2	8695	8684	10165	9648	9378	9225
	4	14592	14592	17093	16321	16700	16424
	8	20608	20608	28388	25862	25883	25470
	16	31615	31500	47612	46825	42137	41261
	32	51963	50758	64614	62894	70462	69439

Table A.1: Comparing the edge cuts obtained by CoHyb_CIP and CoTop with those obtained by the evolutionary algorithm of Moreira et al. on the Polyhedral Benchmark Suite (first set of results).

Graph	k	Moreira et al. [21]		CoHyb_CIP		CoTop	
		Average	Best	Average	Best	Average	Best
jacobi-1d	2	596	596	646	472	682	660
	4	1493	1492	1617	1272	1789	1756
	8	3136	3136	2845	2560	3431	3216
	16	6340	6338	4519	3841	5089	4872
	32	8923	8750	6742	6026	6883	6634
jacobi-2d	2	2994	2991	4327	4002	3445	3342
	4	5701	5700	8405	7379	7370	7247
	8	9417	9416	14872	13802	13168	12895
	16	16274	16231	22626	21625	21565	21098
	32	22181	21758	30423	28911	29558	28979
lu	2	5210	5162	5351	4160	6085	6039
	4	13528	13510	21258	13141	22979	16959
	8	33307	33211	53643	44342	57437	49080
	16	74543	74006	105289	96617	108189	102868
	32	130674	129954	156187	147852	164737	158621
ludcmp	2	5380	5337	5731	5337	6942	5339
	4	14744	14744	25247	19339	22368	22065
	8	37228	37069	60298	50208	60255	50101
	16	78646	78467	106223	98324	109920	99798
	32	134758	134288	158619	151063	165018	155120
mvt	2	24528	23091	57216	33263	21281	19792
	4	74386	73035	55679	36564	38215	35788
	8	86525	82221	62453	47771	46776	43724
	16	99144	97941	71650	59399	54925	48385
	32	105066	104917	83635	79030	62584	60389
seidel-2d	2	4991	4969	4374	3401	4772	4638
	4	12197	12169	13177	12553	11784	11485
	8	21419	21400	24396	22452	21937	21619
	16	38222	38110	38065	35777	39747	38831
	32	52246	51531	58319	57012	59278	57885
symm	2	94357	94214	26374	24629	43597	43330
	4	127497	126207	59815	49450	85730	78379
	8	152984	151168	91892	75126	118259	111126
	16	167822	167512	105418	96322	135278	131127
	32	174938	174843	108950	99584	145903	141223
syr2k	2	11098	3894	4343	900	16124	14404
	4	49662	48021	12192	3121	22915	17959
	8	57584	57408	29194	24912	28787	27259
	16	59780	59594	29519	26327	31807	29132
	32	60502	60085	36111	34079	36689	35155
syrk	2	219263	218019	76767	3240	11740	9036
	4	289509	289088	72148	9995	56832	34893
	8	329466	327712	112236	66981	121664	109730
	16	354223	351824	179042	172076	184437	170781
	32	362016	359544	196173	186162	224330	213676
trisolv	2	6788	3549	367	280	336	336
	4	43927	43549	38148	1277	828	828
	8	66148	65662	20163	9364	2156	2156
	16	71838	71447	20421	12847	6240	5881
	32	79125	79071	25279	19949	13431	13172
trmm	2	138937	138725	50057	32720	13659	3440
	4	192752	191492	58477	16617	72276	35000
	8	225192	223529	92185	58957	134574	102693
	16	240788	238159	128838	122111	157277	145934
	32	246407	245173	153644	147551	171562	158113
Geomean		1.00	0.96	0.74	0.50	0.63	0.57

Table A.2: Comparing the edge cuts obtained by CoHyb_CIP and CoTop with those obtained by the evolutionary algorithm of Moreira et al. on the Polyhedral Benchmark Suite (second set of results). The last line (Geomean) is for the whole PolyBench dataset (i.e., computed by combining this table with the previous one), where the performance of the algorithms are normalized with respect to the average values shown under the column Moreira et al.

Graph	k	CoHyb_CIP		CoTop	
		Average	Best	Average	Best
2mm	2	1.001	1.001	1.001	1.001
	4	1.028	1.030	1.024	1.001
	8	1.030	1.030	1.030	1.030
	16	1.029	1.030	1.030	1.030
	32	1.030	1.030	1.030	1.030
3mm	2	1.021	1.009	1.017	1.017
	4	1.027	1.030	1.030	1.030
	8	1.030	1.030	1.030	1.030
	16	1.030	1.030	1.030	1.030
	32	1.030	1.030	1.030	1.030
adi	2	1.000	1.000	1.030	1.030
	4	1.030	1.030	1.030	1.029
	8	1.030	1.030	1.030	1.030
	16	1.030	1.030	1.030	1.030
	32	1.030	1.030	1.030	1.030
atax	2	1.010	1.011	1.030	1.030
	4	1.020	1.030	1.030	1.030
	8	1.027	1.016	1.029	1.030
	16	1.029	1.030	1.030	1.030
	32	1.030	1.030	1.030	1.030
covariance	2	1.022	1.023	1.030	1.030
	4	1.026	1.021	1.030	1.030
	8	1.028	1.030	1.030	1.030
	16	1.029	1.030	1.030	1.030
	32	1.030	1.030	1.030	1.030
doitgen	2	1.003	1.000	1.030	1.030
	4	1.030	1.030	1.030	1.030
	8	1.030	1.030	1.030	1.030
	16	1.030	1.030	1.030	1.030
	32	1.030	1.030	1.030	1.030
durbin	2	1.024	1.024	1.024	1.024
	4	1.018	1.018	1.023	1.023
	8	1.020	1.020	1.028	1.028
	16	1.028	1.028	1.030	1.030
	32	1.030	1.029	1.030	1.030
fdtd-2d	2	1.007	1.000	1.006	1.000
	4	1.023	1.026	1.021	1.025
	8	1.026	1.028	1.027	1.024
	16	1.027	1.027	1.029	1.028
	32	1.029	1.030	1.028	1.029
gemm	2	1.010	1.008	1.029	1.029
	4	1.024	1.025	1.030	1.030
	8	1.029	1.028	1.029	1.027
	16	1.030	1.030	1.027	1.030
	32	1.030	1.030	1.030	1.030
gemver	2	1.008	1.000	1.000	1.000
	4	1.030	1.030	1.029	1.030
	8	1.029	1.025	1.030	1.029
	16	1.029	1.029	1.030	1.030
	32	1.030	1.030	1.030	1.030
gesummv	2	1.014	1.010	1.022	1.022
	4	1.026	1.013	1.030	1.030
	8	1.028	1.027	1.027	1.030
	16	1.029	1.029	1.030	1.030
	32	1.030	1.030	1.030	1.030
heat-3d	2	1.008	1.030	1.030	1.030
	4	1.030	1.030	1.030	1.030
	8	1.020	1.016	1.030	1.030
	16	1.024	1.022	1.030	1.030
	32	1.030	1.028	1.030	1.030

Table A.3: The partition balances for the edge cuts given in table A.1

Graph	k	CoHyb_CIP		CoTop	
		Average	Best	Average	Best
jacobi-1d	2	1.009	1.010	1.016	1.006
	4	1.019	1.027	1.016	1.022
	8	1.016	1.006	1.024	1.028
	16	1.025	1.024	1.024	1.024
	32	1.027	1.027	1.028	1.028
jacobi-2d	2	1.027	1.030	1.028	1.030
	4	1.017	1.012	1.029	1.030
	8	1.027	1.027	1.030	1.030
	16	1.027	1.028	1.030	1.030
	32	1.029	1.028	1.030	1.030
lu	2	1.023	1.003	1.030	1.030
	4	1.027	1.030	1.029	1.027
	8	1.030	1.030	1.030	1.030
	16	1.030	1.030	1.030	1.030
	32	1.030	1.030	1.030	1.030
ludcmp	2	1.020	1.020	1.022	1.020
	4	1.027	1.030	1.030	1.030
	8	1.030	1.030	1.030	1.030
	16	1.030	1.030	1.030	1.030
	32	1.030	1.030	1.030	1.030
mvt	2	1.020	1.028	1.024	1.030
	4	1.021	1.015	1.028	1.021
	8	1.025	1.030	1.029	1.021
	16	1.028	1.030	1.029	1.030
	32	1.029	1.030	1.030	1.030
seidel-2d	2	1.012	1.011	1.016	1.008
	4	1.024	1.022	1.028	1.025
	8	1.026	1.030	1.030	1.030
	16	1.029	1.029	1.030	1.030
	32	1.029	1.028	1.030	1.030
symm	2	1.016	1.030	1.030	1.030
	4	1.021	1.019	1.030	1.030
	8	1.027	1.029	1.030	1.030
	16	1.030	1.030	1.030	1.030
	32	1.030	1.030	1.030	1.030
syr2k	2	1.018	1.016	1.026	1.000
	4	1.029	1.030	1.020	1.029
	8	1.030	1.027	1.029	1.030
	16	1.030	1.030	1.027	1.021
	32	1.030	1.030	1.030	1.030
syrk	2	1.021	1.022	1.024	1.026
	4	1.030	1.030	1.028	1.030
	8	1.029	1.027	1.030	1.030
	16	1.030	1.030	1.030	1.030
	32	1.030	1.030	1.030	1.030
trisolv	2	1.012	1.021	1.027	1.027
	4	1.026	1.028	1.020	1.020
	8	1.028	1.030	1.026	1.026
	16	1.030	1.030	1.030	1.030
	32	1.030	1.030	1.030	1.030
trmm	2	1.028	1.024	1.016	1.010
	4	1.027	1.021	1.030	1.030
	8	1.030	1.030	1.030	1.030
	16	1.030	1.030	1.030	1.030
	32	1.030	1.030	1.030	1.030
Min		1.000	1.000	1.000	1.000
Average		1.025	1.025	1.027	1.027
Max		1.030	1.030	1.030	1.030

Table A.4: The partition balances for the edge cuts given in table A.2. The last 3 lines (Min, Average, Max) are for the whole PolyBench dataset (i.e., computed by combining this table with the previous one).