

# Expander Hierarchies for Normalized Cuts on Graphs

Kathrin Hanauer

kathrin.hanauer@univie.ac.at  
University of Vienna  
Faculty of Computer Science  
Vienna, Austria

Monika Henzinger

monika.henzinger@ist.ac.at  
Institute of Science and Technology  
Austria (ISTA)  
Klosterneuburg, Austria

Robin Münk

robin.muenk@tum.de  
Technical University of Munich  
Munich, Germany

Harald Räcké

raecke@in.tum.de  
Technical University of Munich  
Munich, Germany

Maximilian Vötsch

maximilian.voetsch@univie.ac.at  
University of Vienna  
Faculty of Computer Science  
UniVie Doctoral School  
Computer Science DoCS  
Vienna, Austria

## ABSTRACT

Expander decompositions of graphs have significantly advanced the understanding of many classical graph problems and led to numerous fundamental theoretical results. However, their adoption in practice has been hindered due to their inherent intricacies and large hidden factors in their asymptotic running times. Here, we introduce the first practically efficient algorithm for computing expander decompositions and their hierarchies and demonstrate its effectiveness and utility by incorporating it as the core component in a novel solver for the normalized cut graph clustering objective.

Our extensive experiments on a variety of large graphs show that our expander-based algorithm outperforms state-of-the-art solvers for normalized cut with respect to solution quality by a large margin on a variety of graph classes such as citation, e-mail, and social networks or web graphs while remaining competitive in running time.

## CCS CONCEPTS

• Information systems → Clustering; • Theory of computation → Sparsification and spanners.

## KEYWORDS

normalized cut, expander decomposition, expander hierarchy, graph partitioning, graph clustering

## 1 INTRODUCTION

In recent years, expander decompositions and expander hierarchies have emerged as fundamental tools within the theory community for developing fast approximation algorithms and fast dynamic

algorithms for such diverse problems as, e.g., routing [12], connectivity [11], (approximate) maximum flow [6, 19], or triangle enumeration [5]. Informally, an expander is a well-connected graph. Its *expansion* or *conductance*  $\phi = \min_{S \subseteq V} \frac{\text{border}(S)}{\min(\text{vol}(S), \text{vol}(\bar{S}))}$  is the minimum ratio between the number of edges leaving any subset of vertices and the number of edges incident to vertices of the subset. A large value of  $\phi$  (close to 1) indicates a graph in which no subset of vertices can be easily disconnected from the rest of the graph.

An *expander decomposition* of a graph partitions the vertices such that the subgraph induced by each part is an expander and the number of edges between different components of the partition is low.<sup>1</sup> Many results over the years [24, 25, 34, 36, 39] have demonstrated that such a decomposition not only exists for every graph, but can be computed in near-linear time. At a high level, the success of expander decomposition-based algorithms is due to the fact that many problems are “easy” on expanders: One first identifies regions of a graph where a problem is easy to solve (the expanders), solves the problem (or a suitable sub-problem) within each region, and then combines the individual solutions to obtain a solution for the overall problem.

*Expander hierarchies* [11] apply expander decompositions in a hierarchical manner. This is done by computing an expander decomposition of the graph and then contracting the individual expanders into single vertices repeatedly. This process is repeated until the entire graph has been contracted to a single vertex. This recursive decomposition has a corresponding decomposition tree  $T$  (the so-called *(flow) sparsifier*) where the root corresponds to the entire graph, the leaves correspond to vertices of the original graph, and inner vertices correspond to the expanders found during the decomposition procedure. The sparsifier approximately preserves the cut-flow structure of the original graph in a rigorous sense. This relationship has been exploited in numerous applications in static and dynamic graph problems [12, 17, 21] to obtain fast algorithms with provable guarantees.

Full version of: Kathrin Hanauer, Monika Henzinger, Robin Münk, Harald Räcké, and Maximilian Vötsch. 2024. Expander Hierarchies for Normalized Cuts on Graphs. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24)*, August 25–29, 2024, Barcelona, Spain. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3637528.3671978> [13].



<sup>1</sup>This can be seen as a special form of a graph clustering.

While many algorithms for expander decomposition that offer strong asymptotic bounds on the running time have been suggested [24, 25, 34, 36, 39],<sup>2</sup> practical algorithms based on expander decomposition have not seen success thus far. This is due to these algorithms requiring one to solve many maximum flow problems, which means these algorithms are prohibitively slow in practice for graphs with many edges. Arvestad [1], e.g., reports that decomposing a graph with approximately  $10^5$  edges can take almost 5 min for various values of  $\phi$  in their implementation of [34].

In practice, the *multilevel graph partitioning* framework has been the de-facto approach for computing high quality solutions for cut problems on graphs. This framework consists of a *coarsening* step, in which a smaller representation of the graph is computed, a *solving* step, where a solution is computed on the coarse graph and a *refinement* step, during which the solution is improved using local methods during the graph uncoarsening process. Multilevel graph partitioning has been successfully applied for computing balanced cuts [2, 18, 29, 33] and normalized cuts in graphs [10].

We note that algorithms based on the expander hierarchy approach may also be regarded as a variant of multilevel graph partitioning. In this approach, a graph is first coarsened by recursive contractions to obtain a smaller graph that reflects the same basic cluster structure as the input graph. Then an initial partition is computed on the sparsifier and, afterwards, in a series of refinement steps the solution is mapped to the input graph while improving it locally as we undo the coarsening. Despite the theoretical guarantees brought forward by expander decompositions and expander hierarchies, the latter have not yet led to similar breakthroughs in the field of fast, high-quality experimental algorithms. The reason is that the expander decomposition required at each coarsening step becomes a significant performance bottleneck in practice. In contrast, multilevel graph partitioners like METIS [18] and KaHiP [33] use a fast matching approach here.

In this work we mitigate the computational bottleneck by introducing a novel, practically efficient random-walk-based algorithm for expander decomposition. Based on this, we give the first implementation of the expander hierarchy and, thus, an algorithm to compute tree flow sparsifiers, allowing us to solve various cut problems on graphs effectively. Exemplarily, we show that our approach is eminently suitable to compute normalized  $k$ -cuts on graphs, where the goal is to partition the vertex set into  $k$  clusters such that the sum over the number of edges leaving each cluster, normalized by the cluster’s volume, is minimized. Normalized cut is a popular graph clustering objective and particularly able to capture imbalanced clusterings. Its manifold applications include, e.g., community detection [22] and mining [4], topic reconstruction [3], story segmentation [42], bioinformatics [10, 40], tumor localization [32], and image segmentation [35, 38]. It is closely related to spectral clustering, which uses spectral properties of eigenvalues and eigenvectors of the graph’s Laplacian. However, the spectral approach suffers both from very large running times and memory requirements to compute and store the eigenvectors, and thus does not scale well [10]. This problem was addressed by Dhillon et al. [10] as well as Zhao et al. [43], who presented algorithms for normalized

cut that either use spectral methods only after coarsening [10] or apply them on a spectrally sparsified graph [43].

*Contributions.* We present the first practically efficient algorithm for expander decomposition and the first implementation to compute an expander hierarchy. Our approach is based on random walks and is justified by rigorous theoretic and empirical analysis. We report on a comprehensive experimental study on normalized cut solvers, comprising 50 medium-sized to very large graphs of various types, where we compare our expander-based algorithm, XCut, to Graclus by Dhillon et al. [10], the approach by Zhao et al. [43], as well as the state-of-the-art graph partitioners METIS and KaHiP, which do not specifically optimize towards the normalized cut objective, but are fast and in practice often used for this task. The experiments show that our algorithm produces superior normalized cuts on graph classes such as citation, e-mail, and social networks, web graphs, and generally scale-free graphs, and is only slightly worse on others. On average, it is still distinctly the best across all graphs and values of  $k$ .

If only a single value of  $k$  is desired, its running time is on average only 3 times slower than the runner-up, Graclus, and never exceeded 18 min. A notable advantage of XCut is that it can quickly compute solutions for multiple values of  $k$  once a sparsifier is computed, which can be faster than running Graclus multiple times.

## 2 RELATED WORK

Our work is motivated by recent theoretical results [5, 6, 11, 12, 19] building on expander decompositions [24, 25, 34, 36, 39] and the expander hierarchy [11] and inspired by non-spectral approaches [10] to tackle the normalized cut problem.

Mohar [23] showed that the computation of the so-called isoperimetric number or conductance of a graph (see section 3) is NP-hard, which implies the hardness of the normalized  $k$ -cut problem for  $k \geq 2$ . Normalized cut remains NP-hard on weighted trees [8].

A number of tools that have been used to solve normalized cut use spectral methods [7, 26–28, 30, 41]. This usually requires to compute  $k$  eigenvectors of the Laplacian matrix, which was shown to scale badly in practice [10, 43]. Afterwards, an additional discretization step is necessary to obtain the clustering.

Dhillon et al. [10] therefore suggest an algorithm called Graclus, which is based on the multilevel graph partitioning framework. It applies the same coarsening steps as METIS, but with a modified matching procedure. The coarsened graph can then be partitioned using different approaches, including a spectral one. In the refinement step, Graclus uses a kernel  $k$ -means-based local search algorithm for improving the normalized cut objective value. The authors evaluate their algorithm experimentally against METIS as well as a spectral clustering algorithm [41]. They show that it outperforms the spectral method w.r.t. normalized cut value, running time, and memory usage. It also produces better results than METIS and is comparable w.r.t. running time.

Zhao et al. [43] employ a joint spectral sparsification and coarsening scheme to produce a smaller representation of the graph that preserves the eigenvectors of the Laplacian in near-linear time  $\tilde{O}(m)$ . Afterwards, a normalized cut is computed on the sparsifier using spectral clustering. Their sparsification scheme obtains a significant reduction in the number of edges and nodes, which makes

<sup>2</sup>Including the near-linear-time algorithm by Saranurak and Wang [34], which finds a decomposition where each component has expansion at least  $\phi$  in time  $O(m\phi^{-1} \log^4 m)$ .

it feasible to apply the spectral method on the reduced graphs, without the quadratic term in the running time growing too large. The authors again perform an experimental comparison with METIS and observe that their algorithm overall outperforms METIS w.r.t. the normalized cut value, while being slightly slower.

METIS [18], KaHiP [33], and many other graph partitioning tools [15, 37] do not solve the normalized cut problem and instead aim to find good solutions to a balanced graph partitioning problem, where the vertex set is to be partitioned into  $k$  sets of (roughly) equal size, while minimizing the number of edges cut. CHACO [15] implements a spectral graph partitioning approach, but the number of clusters is limited to at most  $k = 8$ .

Nie et al. [28] recently presented a spectral normalized cut solver based on the coordinate descent method along with several speedup strategies. They evaluate their algorithm against two other spectral methods [7, 27] on a number of medium-sized data sets and show that it consistently computes the best solution and is the fastest. A notable difference is that  $k$  is not an input parameter.

### 3 PRELIMINARIES

For an undirected graph  $G = (V, E)$  we use  $d_v$  to denote the degree of vertex  $v \in V$ ,  $d$  to denote the *degree vector*, i.e., the vector of vertex degrees, and  $D = Id$  to denote the corresponding *degree matrix*.  $\Delta$  is used to denote the maximum degree of a vertex in  $G$ . For a subset  $S$  of vertices we define the *volume*  $\text{vol}(S)$  as the sum of vertex degrees of vertices in  $S$ . Its *border*  $\text{border}(S)$  (or *capacity*) is defined as  $|E(S, \bar{S})|$ , where  $E(X, Y) = \{\{x, y\} \in E \mid x \in X, y \in Y\}$  is the set of edges between subsets  $X, Y \subseteq V$ , and  $\bar{S} = V \setminus S$ .

A  $k$ -cut is a partition  $\mathcal{P}$  of the vertex set into  $k$  non-empty parts. Given such a partition  $\mathcal{P} = (S_1, \dots, S_k)$ , its *normalized cut value*  $\theta$  is defined as

$$\theta(\mathcal{P}) = \sum_{i=1}^k \frac{\text{border}(S_i)}{\text{vol}(S_i)}.$$

For  $k = 2$  we will usually specify a 2-cut (or just cut) by its smaller side, i.e., we will refer to the cut  $(S, \bar{S})$  by just  $S$ , where  $\text{vol}(S) \leq \text{vol}(\bar{S})$ . The *conductance* of a (2-)cut is defined as  $\Phi(S) = \text{border}(S)/\text{vol}(S)$  and the *conductance of a graph*  $G = (V, E)$  is  $\Phi(G) = \min_{\text{cuts } S} \Phi(S)$ .<sup>3</sup> The problem of finding a 2-cut with minimum conductance is closely related to the problem of finding a 2-cut with minimum normalized cut value, because for any 2-cut,  $\Phi(S) \leq \theta(S) \leq 2\Phi(S)$ . The normalized  $k$ -cut objective can be seen as a generalization of conductance to  $k$ -cuts.

To properly describe our random walks we need some further notation. For a cut  $S$  we call  $\text{vol}(S)/\text{vol}(V)$  the *balance* of the cut  $S$  and denote it with  $b(S)$ . For a subset  $A \subseteq V$  we use  $G\{A\}$  to denote the subgraph induced by  $A$  with self-loops added so that the vertex degrees do not change (a self-loop counts 1 to the degree of a vertex). This means the degree of a vertex  $a \in A$  is the same in  $G$  as in  $G\{A\}$ . Whenever the graph in question is not clear from the context we use subscripts to indicate the graph, i.e., we write  $\text{vol}_G(S)$ ,  $\Phi_G(S)$ ,  $E_G(X, Y)$ , etc. To avoid notational clutter we write  $\text{vol}_A(S)$ ,  $\Phi_A(S)$ ,  $E_A(X, Y)$  when we are referring to the graph  $G\{A\}$ .

We say a graph  $G = (V, E)$  is a  $\phi$ -*expander* if  $\Phi(G) \geq \phi$ , and we call a vertex partition  $(V_1, \dots, V_\ell)$  a  $\phi$ -*expander decomposition* if  $\Phi(G\{V_i\}) \geq \phi$  for all  $i$ .

<sup>3</sup>If there are no cuts in  $G$  (i.e.,  $|V| = 1$ ) we define its conductance to be 1.

## 4 EXPANDER DECOMPOSITION USING RANDOM WALKS

In this section we describe our random-walk-based approach for obtaining expander decompositions and present its theoretical guarantees. The complete proofs for these guarantees can be found in the appendix.

A natural approach for computing expander decompositions is to find a low conductance cut to split the graph into two parts and then to recurse on both sides. If no such cut exists, we have certified the (sub-)graph to be an expander. In the end, the whole procedure terminates if each subgraph is an expander, i.e., it terminates with an expander decomposition. Saranurak and Wang [34] used this general approach to obtain an expander decomposition that runs in time  $O(m \log^4 m / \phi)$  and only cuts  $O(\phi m \log^3 m)$  edges. While their flow-based techniques give very good theoretical guarantees, the hidden constants do not seem to allow for good practical performance (see e.g. [1]).

In this work we base the cut procedure of the expander decomposition on random walks. As a consequence, we can only guarantee that our decomposition cuts at most  $\tilde{O}(\sqrt{\phi} m)$  edges<sup>4</sup>, since we are limited by the intrinsic Cheeger barrier of spectral methods. However, random walks have a very simple structure, which leads to a simple algorithm with good practical performance. The weaker dependency on  $\phi$  ( $\sqrt{\phi}$  instead of  $\phi$ ) is not crucial for our graph partitioning application because when using the expander decomposition to build an expander hierarchy one chooses  $\phi$  as large as possible anyway.

**THEOREM 1 (EXPANDER DECOMPOSITION).** *Given a graph  $G$  with  $m$  edges and a parameter  $\phi$ , there is a random-walk-based algorithm that with high probability finds a  $\phi$ -expander decomposition of  $G$  and cuts at most  $O(\sqrt{\phi} m \log^{5/2} m)$  edges. The running time is  $O(\frac{m+n \log n}{\phi} \log^3 m)$ .*

The main part of our algorithm is the cut procedure described in Section 4.1. This procedure is then plugged into the framework of Saranurak and Wang to find an expander decomposition. Our cut procedure gives the following guarantees.

**THEOREM 2 (CUT PROCEDURE).** *Given a graph  $G = (V, E)$  with  $m$  edges and a parameter  $\phi$ , the cut procedure takes  $O((m + n \log n)/\phi)$  steps and terminates with one of these three cases:*

- (1) We certify that  $G$  has conductance  $\Phi(G) \geq \phi$ .
- (2) We find a cut  $(A, \bar{A})$  in  $G$  that has conductance at most  $\Phi_G(A) \in O(\sqrt{\phi} \log^{3/2} m)$ . Then one of the following holds:
  - (a) either  $\text{vol}(A), \text{vol}(\bar{A})$  are both  $\Omega(m/\log^2 m)$ , i.e.,  $(A, \bar{A})$  is a relatively balanced low conductance cut;
  - (b) or  $\text{vol}(\bar{A}) \in O(m/\log^2 m)$  and  $A$  is a near  $6\phi$ -expander.

Given the cut procedure, an expander decomposition is computed as follows. On the current subgraph  $G$ , execute the cut procedure, to either find a low conductance cut  $S$  or certify that none exists. If no such cut exists, then  $G$  is a certified  $\phi$ -expander and we terminate. Otherwise we check whether  $S$  is sufficiently balanced, i.e., the volume of the smaller side is at least  $\Omega(m/\log^2 m)$ . In that case we

<sup>4</sup> $\tilde{O}$  hides polylogarithmic factors.

cut the edges across  $(S, \bar{S})$  and recurse on both parts. As both parts are substantially smaller than  $G$  we obtain a low recursion depth.

Otherwise,  $S$  is very unbalanced but the larger side  $\bar{S}$  is a so-called *near expander* – a concept introduced by Saranurak and Wang [34]:

**DEFINITION 1 (NEAR  $\phi$ -EXPANDER).** *Given a graph  $G = (V, E)$ . A subset  $A \subseteq V$  is a near  $\phi$ -expander in  $G$  if for all sets  $X \subseteq A$  with  $\text{vol}(X) \leq \text{vol}(A)/2$ :  $|E(X, V \setminus X)| \geq \phi \text{vol}(X)$ .*

Note that if the LHS in the above equation were  $|E(X, A \setminus X)|$  then  $G\{A\}$  would be a  $\phi$ -expander.

If  $G\{\bar{S}\}$  (for  $\bar{S}$  returned by the cut-procedure) is indeed a  $\phi$ -expander we can just recurse on the smaller side  $S$  and would obtain a low recursion depth. Saranurak and Wang introduced a trimming procedure that, given a subset  $B$  that is a near  $6\phi$ -expander with volume  $\text{vol}(B) \geq 9 \text{vol}(V)/10$  and  $|E(B, \bar{B})| \leq \phi \text{vol}(B)/10$ , computes a subset  $B' \subseteq B$  that is a proper  $\phi$ -expander and has volume  $\text{vol}(B') \geq \frac{1}{2} \text{vol}(B)$ . By applying this trimming step to  $\bar{S}$  we can return  $G\{S'\}$  as a proper  $\phi$ -expander and recurse on the remaining graph – still with a small recursion depth. Overall, using our cut procedure within this framework gives Theorem 1.

## 4.1 Finding Low Conductance Cuts

We now give a detailed description of the cut procedure that forms the basis of Theorem 2. The goal is to either certify that  $G$  is a  $\phi$ -expander or to find a low conductance cut that is as balanced as possible. The idea is to exploit that random walks converge quickly on expanders and hence when they don't, we know there must be a low conductance cut. See Algorithm 1 for an outline.

We employ a concurrent random walk, where each node distributes its unique commodity in the graph. We are interested in the probability that after  $t$  steps a particle that started say at node  $i$  is at some other node  $j$ . The walk has converged if this distribution is essentially identical for every starting vertex. If we quickly reach this stationary distribution, there cannot be a low conductance cut and hence the graph must be an expander. Otherwise we can use information gathered from the walk to find a low conductance cut.

Such a cut may however be very unbalanced. The procedure therefore accumulates low conductance cuts until the combined cut is a balanced low conductance cut or the graph that remains does not have a low conductance cut anymore. Here we call a cut  $S$  *balanced* if its balance  $b(S) \geq \beta := 2/\log^2 m$ .

More precisely, the algorithm maintains a partition of  $V$  into two sets  $A, L$  for each iteration  $t$  with initial values  $A := V, L := \emptyset$ . We repeat the following for  $T = 1/(12\phi)$  steps: In iteration  $t$ , we generate a new random unit vector  $r$  and execute  $t - 1$  steps of the random walk, initialized according to  $r$ . This walk yields a vector  $u$ , on which we analyze the conductance of all sweep cuts, i.e., cuts of the form  $S_c := \{v \in A : u_v \leq c\}$  for some value  $c$ . Note that these conductance values can be calculated in linear time after sorting the entries of  $u$ .

We consider a cut to have low conductance if the value is below the threshold  $\gamma = O(\sqrt{\phi} \log^{3/2} m)$ . A lower threshold value  $\gamma$  would give better guarantees in case (2) of Theorem 2 but at the same time it would increase the number of rounds required to converge and hence worsen the guarantee of case (1) in Theorem 2. The value is chosen to ensure we can guarantee  $\Phi(G) \geq \phi$  in case (1) of the Theorem.

---

### Algorithm 1 Cut Procedure

---

**Input:** Graph  $G = (V, E)$ , Target expansion  $\phi$

**Output:** EXPANDER or BALANCED( $S, \bar{S}$ ) or UNBALANCED( $S, \bar{S}$ )

```

1:  $T \leftarrow 1/(12\phi)$ 
2:  $\gamma \leftarrow 343\sqrt{\phi} \log(32m^3) \log^2(n)$  ▶ cut threshold
3:  $\beta \leftarrow 2/\log^2 m$ , ▶ balance
4:  $W \leftarrow I$  ▶ random walk matrix
5:  $A \leftarrow V, L \leftarrow \emptyset$ 
6: for  $t = 1, \dots, T$  do
7:    $\vec{r} \leftarrow$  random unit vector in  $\mathbb{R}^n$ 
8:    $\vec{u} \leftarrow WD^{-1}r$  ▶ apply random walk
9:    $\vec{u} \leftarrow \vec{u} - (\vec{u}^\top \vec{d})/\text{vol}(V) \cdot \mathbb{1}$  ▶ ensure  $\vec{u} \perp \vec{d}$ 
10:   $\mathcal{A}, D_A \leftarrow$  adjacency and degree matrix of  $G\{A\}$ 
11:   $W \leftarrow (\frac{1}{2}I + \frac{1}{2}\mathcal{A}D_A^{-1})W$  ▶ extend walk matrix
12:   $S \leftarrow$  SweepCut( $\vec{u}, \gamma$ )
13:  if  $\text{vol}(S) \geq \beta m$  then return BALANCED( $S, \bar{S}$ )
14:  else if  $S \neq \emptyset$  then
15:     $L \leftarrow L \cup S, A \leftarrow A \setminus S$ 
16:    if  $\text{vol}(L) \geq \beta m$  then return BALANCED( $A, L$ )
17: if  $L = \emptyset$  then return EXPANDER
18: else return UNBALANCED( $A, L$ )
```

---

The analysis of the sweep cuts yields one of these three cases:

- (1) If all sweep cuts have conductance at least  $\gamma$ , we continue with the next iteration.
- (2) If there exists a sweep cut with conductance  $< \gamma$  and balance  $\geq \beta$  we return this low conductance cut.
- (3) Otherwise we consider the two-ended sweep cuts of the form  $S_{a,b} := \{v \in A : u_v \notin [a, b]\}$  for values  $a \leq b$  and find the one with largest volume among those with  $\Phi_A(S_{a,b}) < \gamma$ . If this cut has balance  $\geq \beta$  we return it, otherwise we move it from  $A$  to  $L$  and check whether  $L$  has become balanced.

The random walk can be interpreted as the projection of a much higher dimensional random walk onto the randomly chosen direction  $r$ . This projection step is crucial for the implementation to become computationally feasible and we show that the projections approximate the original structure sufficiently well.

In order to argue the correctness of the cut procedure we have to show that it is highly unlikely that the procedure does not find a cut on a graph that has expansion less than  $\phi$ . For this we argue that after  $T$  random walk steps (without finding a cut) the walk will have “converged” to its stationary distribution w.h.p. Because of the choice of parameters such a quick convergence is only possible if  $G$  is a near  $6\phi$ -expander. Whenever the cut procedure returns a cut, it is guaranteed to have conductance at most  $\gamma$ . From this it follows that the expander decomposition cuts at most  $O(\gamma m \log m)$  edges. The entire argument can be found in Section A in the appendix.

## 5 XCut – A NEW NORMALIZED CUT ALGORITHM

In this section, we introduce the algorithm XCut, which is based on the previous section’s novel random walk-based expander decomposition. We note the apparent similarity between multilevel graph partitioning and the expander hierarchy and use this as the basis of XCut. As an outline, we use the novel random walks to construct the expander hierarchy to obtain a coarse representation of the graph, the tree flow sparsifier. We then compute an initial solution on the tree. Finally, we use an iterative refinement step while descending the hierarchy to improve the solution we found. Compared to other contraction schemes, which lead to each vertex in the coarsest graph representing roughly the same number of nodes in the base graph, the subtrees on each level in the tree flow sparsifier can represent a vastly different number of vertices.

*Expander Decomposition.* While the expander decomposition outlined in section 4 is much simpler to implement than that of [34], as we do not rely on maximum flow computations at all, we made several choices in the implementation to speed up computation. We iterate a single random walk, and after each iteration, we check whether we can find a sparse cut. If we find a suitable cut, we disconnect the edges going across it, but contrary to the algorithm in section 4, we do not restart the random walk, as we find we can extract further information about the cut structure of the graph from the state of the random walk. For example, if the random walk has mixed very well on one of the new components, it is likely to be an expander, while if there is another sparse cut in the component, then the random walk will likely not have mixed well on the component. One may think that reducing the number of random walks might lead to a loss of guarantees and higher variance of the algorithm, but in experiments conducted while designing the algorithm, we found that on real graph instances running multiple concurrent random walks is not necessary. In fact, only a single graph in our 50 graph benchmark had noticeable variance. See also the discussion in subsection 6.3.

The main parameter of the expander decomposition is the cut value  $\gamma$ , which is the minimal sparsity of the cuts our algorithm makes. Additionally, we introduce a parameter  $\rho$ , to be used as a threshold for “certifying” that a component is an expander, as we found that choosing the threshold to be  $1/(4 \text{vol}(V)^2)$  does not offer any benefits over a much larger value. See also subsection 6.2 for details on choosing the value for this parameter. We make a final modification to the theoretical algorithm in that we omit the trimming step on unbalanced cuts, since it does not provide any further speedup of the expander decomposition routine in practice.

*Automatically Choosing  $\gamma$ .* The theoretical analysis in [11] suggests a choice for  $\gamma$  and  $\phi$  that is sufficient to prove the theoretical results. In preliminary experiments we found this choice to be too pessimistic and, in fact, by adapting  $\phi$  and  $\gamma$  to the graph we can obtain better results. However, there is a trade-off to be made: If  $\gamma$  is too small, the expander decomposition will not find many sparse cuts, and we obtain sparsifiers of low quality. On the other hand, if  $\gamma$  is too large, most cuts will be sparser than  $\gamma$ , which leads to many cuts being made and increases the running time. In the worst case, this can even prevent the algorithm from terminating.

Thus, our goal is to choose  $\gamma$  such that it offers a good quality vs. time trade-off. When choosing  $\gamma$ , we can only observe whether this was a good choice in a post hoc fashion. A naive strategy would be to start with a large  $\gamma$  and decrease it until the expander hierarchy terminates within a reasonable amount of time. However, this approach is wasteful, as we discard previously computed decompositions, even if they were good. Instead, we decrease  $\gamma$  by multiplying it with constant factor  $\epsilon < 1$  whenever the expander decomposition on a specific level cuts too many edges in  $G_i$ . We then use the new  $\gamma' = \epsilon\gamma$  for the remaining expander decompositions, decreasing it further as required.

*Solving on the Sparsifier.* To solve normalized  $k$ -cut on the tree sparsifier obtained from the hierarchy, we want to remove  $k - 1$  edges to decompose it into a forest of  $k$  trees. Given a solution, i.e., a tuple of edges  $(e_1, \dots, e_{k-1})$ , we assign vertex  $v \in V$  to the cluster  $C_j$  associated with  $e_j$  if  $e_j$  is the first edge we encounter on the path from  $v$  to the root. If no edge in the solution lies on the path to the root, we assign  $v$  to cluster  $C_k$ . As normalized cut is NP-hard also on trees (see section 2), we introduce two heuristic approaches that take time  $O(nk)$  each:

*Greedy:* The simple greedy heuristic picks the edge in the sparsifier which minimizes the increase in the normalized cut objective in every step. By simply computing the cost of cutting each remaining edge, it takes  $O(n)$  time to find this edge, assuming the number of vertices in the sparsifier is  $O(n)$ . To partition a graph into  $k$  clusters, we repeat this process  $k - 1$  times. For  $k = 2$ , this algorithm produces the optimal solution on the tree as we pick the edge that minimizes the cut objective.


*Dynamic Programming:* Each row of the dynamic program corresponds to a level, and we make one cell for each pair  $(v, i)$  of the row, where  $v$  is a vertex in the sparsifier and  $i \in [0, k]$ . The value of each cell is the normalized cut value  $\theta$  of decomposing the subtree rooted at  $v$  into  $i$  parts. We write  $DP(v, i)$  for this value. Additionally, we write  $\text{cut}(v, i)$  for the weight of cut edges in the solution incident to the subtree rooted at  $v$ , as well as  $\text{vol}(v, i)$  for the volume remaining in the subtree.

Without loss of generality, assume the tree is binary, as otherwise we can binarize the tree by inserting edges of infinite cost (see Henzinger et al. [16] for details). The value of a cell is computed according to the following rule:

$$DP(v, j) = \min(\text{cutParent}(DP(v, j - 1)), \min_{0 \leq i \leq j} (DP(v_l, i) + DP(v_r, j - i))),$$

where  $\text{cutParent}(DP(v, j - 1)) = DP(v, j - 1) + \frac{w(v, v') + \text{cut}(v, j - 1)}{\text{vol}(v, j - 1)}$  is the best solution where the edge going to the parent is cut. For each vertex  $v$  of  $G_0$ , we initialize the bottom row of the program with  $DP(v, 0) = 0$ ,  $DP(v, 1) = 1$  and  $DP(v, j) = \infty$  for  $j \in [2, k]$ . In the root vertex we use the special rule  $DP(v, j) = \min_i DP(v_l, i) + DP(v_r, j - i) + \frac{\text{cut}(v_l, i) + \text{cut}(v_r, j - i)}{\text{vol}(v_l, i) + \text{vol}(v_r, j - i)}$ , where the last term ensures we do not produce a solution with  $\text{vol}(v, i) = 0$ , leading to cluster  $j$  being empty.

*The Refinement Step.* Finally, we perform an iterative refinement as we descend the hierarchy. While descending, we introduce new clusters according to the edges in the solution. On each level we then perform vertex swaps that improve the normalized cut objective,



**Figure 1: Relative improvement over Graclus (y-axis) vs. the ratio of the maximum degree and the median degree (x-axis). The value on the x-axis is larger if the graphs exhibit a distribution with large outliers. Note the negative trend, except for the outlier towards the right corresponding to instance SN7.**

by either reducing the number of cut edges or making the partition more balanced. For details see the full version of the paper.

*Support for Variable Number of Partitions  $k$ .* A notable strength of our algorithm is that with both heuristics, it solves the problem for any value  $k' \leq k$  during their execution. Suppose we are still determining the number of clusters needed. In that case, we can compute the solution for the maximum  $k$  we are interested in and obtain solutions for all smaller numbers of partitions during exploratory data analysis. The only step that needs to be rerun is the refinement step.

## 6 EXPERIMENTAL EVALUATION

In the previous sections, we have shown that our approach provides provable guarantees on the approximation ratio for the value  $\theta$  of the normalized cut<sup>5</sup> (and its relatives, sparse cut, and low-conductance cut) if  $k = 2$ . We now turn to evaluate XCut experimentally in different configurations. We compare the objective value of normalized cuts produced by XCut and its running time against the normalized cut solver Graclus by Dhillon et al. [10] and the state-of-the-art graph partitioning packages METIS [18] and KaHiP (kaffpa) [33], all of which are available publicly. These algorithms are based on the multilevel graph partitioning framework and produce disjoint partitions of the vertices into  $k$  clusters, where  $k$  is a freely choosable parameter. We note that METIS and KaHiP solve the balanced  $k$ -partitioning problem rather than normalized cut. Nevertheless, we include these solvers in our comparison, as they are used for this task in practice and we found that they can outperform Graclus on some of the graphs in our benchmark dataset. By this, we follow the methodology of [10] and [43].

In addition, we compare our results to the values reported for the normalized cut algorithm by Zhao et al. [43]<sup>6</sup>. We omit comparisons to solvers employing spectral methods such as the recent works by Chen et al. [7] and Nie et al. [28], as this approach does not scale

<sup>5</sup>See Section 3 for the precise definition.

<sup>6</sup>Unfortunately, the code is not available publicly and also could not be provided by the authors upon request before the submission deadline.

well to large datasets of millions of nodes [10]. In preliminary experiments we found that the solver of Nie et al. [28] uses over 330GB of memory on instance CN3, whereas our solver used less than 400MB of memory. Furthermore, the algorithm presented in [28] does not necessarily produce  $k$  clusters, thus making direct comparisons difficult. Lastly, the space complexity of spectral methods becomes prohibitive for larger values of  $k$ , as noted by [10].

### 6.1 Experimental Setup

*Instances.* Our setup includes 50 graphs from various applications. See Table 1 for an overview and the full version of the paper for details. To facilitate comparability, our collection contains the eight instances used by Dhillon et al. [10] (BP1, CF1, CS1–3, DM1, OP1, and OP2) as well as the 21 instances used by Zhao et al. [43]. In addition, we selected 21 real-world networks that cover various application areas, including some of larger sizes. All instances are available publicly in the Network Repository [31] or the SuiteSparse Matrix collection [9].

Note that a graph with  $k$  or more connected components always has a (normalized)  $k$ -cut of size 0. Surprisingly, we found that XCut is the only solver tested here that finds the trivial optimal solution if  $k$  is less than the number of connected components. However, testing connectivity before starting a solver remedies this problem, which is why we decided to exclude graphs with more than 128 connected components except for one (graph ID 0), which we keep for consistency reasons as it was used in previous comparisons [10].

*Methodology.* As XCut, KaHiP, and METIS are randomized, we ran each of them  $\ell = 10$  times per instance with different seeds. Graclus is deterministic, and we ran it three times to obtain a stable value for the running time. We use the arithmetic mean over the  $\ell$  runs for each instance to approximate the expected value of  $\theta$  and the running time. When reporting values, we write  $\text{XCut}_{\text{mean}}$  for the mean value across these 10 runs, and  $\text{XCut}_{\text{min}}$  for the minimum. As KaHiP and METIS behaved almost identically over all  $\ell$  runs, we only report mean values for them. For each algorithm and each graph, we compute a partitioning consisting of  $k \in \{2, 4, 8, 16, 32, 64, 128\}$  clusters as well as  $\theta$  (smaller is better). As a second criterion, we compare the algorithms' running times.

All experiments were conducted on a server with an Intel Xeon 16 Core Processor and 1.5 TB of RAM running Ubuntu 22.04 with Linux kernel 5.15. XCut is implemented in C++ and compiled using gcc 11.4 with full optimization<sup>7</sup>. For all other solvers, we followed the build instructions shipped with their code. As Graclus is single-threaded, we ran the single-threaded version of every algorithm. METIS was run in its default configuration. We used KaHiP with the `fsocial` flag, which is tailored to quickly partitioning social network-like graphs<sup>8</sup> and Graclus with options `-l 20 -b` to enable the local search step and only consider boundary vertices during local search, as suggested by the authors for larger graphs [10].

### 6.2 Configuring XCut


*Greedy vs. Dynamic Programming (DP).* Section 5 describes two heuristics for computing normalized cuts on the sparsifier. We

<sup>7</sup>-O3 -march=native -mtune=native

<sup>8</sup>We chose this setting as we are especially interested in social network-like graphs.

**Table 1: Types and number of graphs in our benchmark dataset.  $\Delta$  is the maximum degree,  $k$  and  $M$  are shorthand notations for  $10^3$  and  $10^6$ , respectively. See the full version of the paper for a detailed list.**


Type (Abbreviation)	#	$ V $	$ E $	$\Delta$
Bipartite (BP)	1	1.4M	4.3M	1.7k
Computational Fluids (CF)	1	17k	1.4M	269
Clustering (CL)	2	4.8k-100k	6.8k-500k	3-17
Citation Network (CN)	9	226k-1.1M	814k-56M	238-1.1k
Circuit Simulation (CS)	3	5k-30k	9.4k-54k	31-573
Duplicate Materials (DM)	1	14k	477k	80
Email Network (EM)	2	33k-34k	54k-181k	623-1383
Finite Elements (FE)	2	78k-100k	453k-662k	39-125
Infrastructure Network (IF)	2	2.9k-49k	6.5k-16k	19-242
Numerical Simulation (NS)	2	11k-449k	75k-3.3M	28-37
Optimization (OP)	2	37k-62k	131k-2.1M	54-8.4k
Random Graph (RD)	1	14k	919k	293
Road Network (RN)	4	114k-6.7M	120k-7M	6-12
Social Network (SN)	7	404k-4M	713k-28M	626-106k
Triangle Mixture (TM)	7	10k-77k	54k-2M	22-18k
US Census Redistricting (US)	1	330k	789k	58
Web Graph (WB)	3	1.3k-1.9M	2.8k-4.5M	59-2.6k



**Figure 2: Running time vs. normalized cut for different choices of  $\rho$  on two different graphs for  $k = 16$ . Colors denote different levels of  $\rho$ , while shapes indicate the graph.**

compare for both heuristics the running time and  $\theta$  across ten precomputed sparsifiers each for 19 representative graphs. The quality returned by DP was never better than that of Greedy. While the running times for both heuristics are linear in the size of the sparsifier, the DP approach scaled worse in  $k$ . For example, for  $k = 32$ , DP was three times slower than Greedy, and seven times slower for  $k = 128$ . See the full version of the paper for a plot with the results for  $\rho = 10^{-4}$ . Greedy was faster and produced no worse quality than DP for all values of  $\rho$ . Thus, we only report results for Greedy for all further experiments.

*Parameter Choice for the Expander Decomposition ( $\rho, \gamma$ ).* In Figure 2, we examine the effect of the threshold parameter  $\rho$  on the quality of solutions and running time. If the parameter is chosen too large, especially greater than one, the entire graph will likely be certified as an expander before any cuts are made, leading to very large  $\theta$ . Furthermore, we no longer obtain any significant



**Figure 3: Geometric mean of the cut value  $\theta$  across all graphs for each  $k$  for  $XCut_{mean}$ , Graclus, METIS, and KaHiP.**

improvement in  $\theta$  for values of  $\rho$  below  $10^{-4}$ . At the same time, the running time increases inversely with  $\rho$  as extra iterations are needed to converge, which is shown by the vertical arrangement of the dots in Figure 2. For graph instance TM6, we also find that non-Pareto-optimal choices of  $\rho$  exist, namely 0.1 and 0.01, where both the running time and the returned value are worse than  $10^{-3}$  and  $10^{-4}$ . This suggests that the choice of  $\rho$  is an important design decision. On all our instances,  $\rho = 10^{-4}$  produced Pareto-optimal results, so we conclude that we can configure this parameter to be a constant independently of the graph’s structure.

For the automatic tuning of  $\gamma$ , we chose a starting threshold of 0.3, as this is around the largest value for which the expander decomposition finds structurally interesting cuts. Whenever  $\gamma$  is too large, i.e., the node reduction  $|G_{i+1}|/|G_i| > 0.95$ , we multiply  $\gamma$  by a factor of  $\epsilon = 0.8$  and restart the expander decomposition.

### 6.3 Comparison to Graclus, METIS, and KaHiP

Figure 5 depicts the solution quality of every solver relative to that produced by Graclus on all instances for  $k = 32$ , showing both the mean and the minimum of the ten runs for XCut. For other values of  $k$ , the overall picture remains the same, see Figure 3 and Appendix D. We group the graphs by type, with Figure 5 containing two plots, the upper showing the disconnected IMDB graph and email, citation, and social network graphs as well as infrastructure networks, as these are the graphs on which XCut is particularly strong.

Looking at absolute values of  $\theta$ , we find that across all instances, the geometric mean is at least 70% lower than our competitors, see Table 2. Interestingly, when we only consider the seven graphs from [10], the geometric means become 0.84 for  $XCut_{min}$  and 0.90 for  $XCut_{mean}$ , while they are 1.11, 1.19 and 1.03 for Graclus, METIS, and KaHiP respectively, which implies that KaHiP slightly outperforms Graclus in terms of  $\theta$  on their benchmark (but not XCut).

One point of note is that XCut does not always find small normalized cuts on the social network SN7 representing user-user interactions in the Foursquare social network. This appears to be due to a very high-degree node that connects to approximately 15% of all vertices, leading to fast convergence of the random walk. Due to the averaging effect of such a vertex, many nodes  $v$  have almost identical values  $u_v$  that vary only slightly between runs. Thus, when sorting by  $u$ -value, their order can vary greatly depending on the initial values, leading to very different candidate cuts. This is the only graph where the minimum and mean of the ten

**Table 2: Cut value ( $\theta$ ) of different algorithms. The geometric mean is taken across all graphs and values of  $k \in \{2, 4, 8, 16, 32, 64, 128\}$  for each graph type. \*For the overall geometric mean, instances (graph +  $k$ ) with  $\theta = 0$  were omitted. Only XCut detected such cases.**


Type	XCut <sub>mean</sub>	XCut <sub>min</sub>	Graclus	METIS	KaHiP
BP	<b>0.00</b>	<b>0.00</b>	1.18	1.38	1.58
CF	4.61	4.18	3.44	3.51	<b>3.15</b>
CL	0.80	<b>0.72</b>	1.04	1.13	1.02
CN	<b>0.07</b>	<b>0.07</b>	1.42	1.72	1.58
CS	0.44	<b>0.41</b>	0.51	0.59	0.57
DM	2.58	2.42	<b>2.11</b>	2.12	2.23
EM	0.44	<b>0.43</b>	3.44	3.78	3.80
FE	0.56	<b>0.53</b>	0.54	0.54	0.55
IF	<b>0.00</b>	<b>0.00</b>	0.93	1.11	1.12
NS	0.85	0.78	0.77	<b>0.76</b>	0.80
OP	0.71	<b>0.66</b>	1.47	1.48	0.99
RD	13.37	13.37	12.08	12.01	<b>11.94</b>
RN	<b>0.01</b>	<b>0.01</b>	<b>0.01</b>	<b>0.01</b>	<b>0.01</b>
SN	0.25	<b>0.19</b>	3.43	3.98	4.00
TM	2.04	<b>1.93</b>	2.87	3.29	3.04
US	0.04	<b>0.03</b>	0.06	0.06	0.05
WB	<b>0.01</b>	<b>0.01</b>	0.10	0.14	0.14
<b>All*</b>	0.25	<b>0.22</b>	0.89	1.0	0.94

runs of XCut differ significantly ( $-33\%$  for XCut<sub>min</sub> and  $+39\%$  for XCut<sub>mean</sub> relative to Graclus). This indicates that the theoretical algorithm sketched in section 4, which uses multiple concurrent random walks (corresponding to more attempts to find a good cut), would likely have reduced the variance here. This is also the only graph where the fact that we only use a single random walk impairs the quality of the result.

The graph classes on which XCut does not perform as well have fairly homogeneous degree distributions and often appear grid-like when drawn. We conjecture that in these grid-like graphs, there are no good expanders (which XCut is trying to find). Instead, sparse cuts arise mainly from the fact that the cut is balanced, i.e., the components we disconnect are all large enough, rather than being sparsely interconnected, which is exploited by the other solvers.

#### 6.4 Comparison to Zhao et al.

In Table D.5 we compare  $\theta$  for XCut and Graclus to the values reported by Zhao et al. [43] for  $k = 30$ . We observe a similar but weaker pattern as in the previous section. On graphs arising from numerical simulation and finite element problems, XCut performs worse than Zhao et al., but never by more than 36%. On the triangle mixture instances, XCut achieves better  $\theta$  on four instances, while their solver outperforms XCut on two instances. On citation networks, clustering instances, and those based on maps (RN1 and US1), XCut outperforms the algorithm by Zhao et al., with  $\theta$  being up to 2.5 times lower on US1 and CN7. Altogether, XCut is better than Zhao et al. on roughly 2/3 of the instances, while Graclus does best on one instance. The geometric mean across all instances is 1.46 for XCut<sub>mean</sub>, 1.39 for XCut<sub>min</sub>, 1.64 for Zhao et al., and 3.06 for Graclus. We note that ours is 11% lower for XCut<sub>mean</sub> and 15% lower for XCut<sub>min</sub>, while Graclus’s value is almost twice that of



**Figure 4: Relative improvement over Graclus (y-axis) vs. the average degree divided by the median degree (x-axis). Larger x-values signify that the graph exhibits a skewed, power-law-like degree distribution. Note the negative trend, except for the outlier towards the right corresponding to instance SN7.**

the other solvers in the comparison due to it producing 5–30 times greater  $\theta$  on the citation network (CN) instances.

While it is difficult to draw good conclusions for the running times from the values reported in [43] as no source code is available, we note that on some instances XCut takes less than 10% of the reported time while producing higher-quality solutions, which might be indicative of a running time advantage of our solver.

#### 6.5 Running Time

In our experiments, we found that on many graphs, the running time is spent mainly on computing the expander hierarchy, where on some instances, this step accounts for over 80% of the running time, even for  $k = 128$ . See Figure D.9 for some examples. However, even on the largest instances in our benchmark, the absolute running time never exceeded 18 minutes.

Overall, XCut is on average three times slower than Graclus, 20.8 times slower than METIS and 6.7 times slower than KaHiP for the mean execution time across all choices of  $k$  and all instances. Interestingly, we find that XCut’s running time is lower than Graclus’s on several social network graphs and the triangle mixture instances while it produces a better solution. See the full version of the paper for detailed running times on some exemplary instances.

Finally, recall that once our algorithm has computed a sparsifier, we can obtain a solution on the sparsifier for different values of  $k$  without recomputing the sparsifier, which is a unique feature among the solvers tested here. If we are interested in all seven values of  $k$ , e.g., and only count the time to compute the sparsifier once for XCut, our experiments take 0.56, 3.82, and 1.24 times the running time to compute partitions across all graphs and values of  $k$  using Graclus, METIS, and KaHiP, respectively. In particular, XCut then is 44% faster than Graclus. This demonstrates the utility of XCut as a tool for exploratory data analysis. We could achieve further speedups by only computing a solution for  $k = 128$  and then choosing the initial subsets of edges for the other values of  $k$ , only performing the refinement.





**Figure 5: Percentage deviation of the returned normalized cut value relative to Graclus for  $k = 32$ . This means that a value of  $-75\%$  indicates that the normalized cut value is  $75\%$  lower (i.e., better). The thin black bars indicate the standard error across our runs. The top graph shows the disconnected IMDB graph (BP1), citation network instances (CN), email networks (EM), infrastructure graphs (IF), social networks (SN), and web graphs (WB), while the bottom shows the remaining instances. See the full version of the paper for details.**

## 6.6 Discussion

XCut outperforms other software when computing normalized cuts on social, citation, email, and infrastructure networks and web graphs. It performs slightly worse on graphs arising from specific computational tasks, such as finite elements, circuit, or numerical simulations. The graphs on which this behavior occurs tend to have degree distributions concentrated around the average degree, suggesting that they are not graphs with a scale-free structure.

In Figure 4 we plot the relationship between our improvement over Graclus and the value of  $\frac{\frac{1}{n} \sum_{v \in V} d_v}{\text{median}_{v \in V} d_v}$  for the non-synthetic graphs of our benchmark. This value measures how much the mean and median diverge due to outlier nodes with very high degrees. The graphs with power-law distributions tend to have a higher value on this measure, and we find that there appears to be a negative correlation, with one outlier due to instance SN7.

## 7 CONCLUSION

In this work we introduced XCut, a new algorithm for solving the normalized cut problem. It is based on a novel expander decomposition algorithm and to the best of our knowledge, it is the first practical application of the expander hierarchy. XCut clearly outperforms other solvers in the experimental study on social, citation, email, and infrastructure networks and web graphs, and also in the geometric mean over all instances. It scales to instances with

tens of millions of edges and can produce solutions for multiple numbers of clusters  $k$  with little overhead by comparison. We are confident that with further optimization and the use of parallelism it will be possible to scale our algorithm to even larger graphs, while further improving the solution quality, especially since computing the expander decomposition appears to be highly parallelizable.

We also believe that the expander hierarchy and our expander decomposition can be applied to other graph cut problems in the future, as the tree flow sparsifiers approximate *all* cuts in the graph, and there are theoretical results that suggest this might be the case. XCut is open source software and its code is freely available on GitLab [14].

## ACKNOWLEDGMENTS

Monika Henzinger: This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant agreement No. 101019564) and the Austrian Science Fund (FWF) grant DOI 10.55776/Z422, grant DOI 10.55776/I5982, and grant DOI 10.55776/P33775 with additional funding from the netidee SCIENCE Stiftung, 2020–2024.

Harald Räcke, Robin Münk: This project has received funding from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 498605858 and 470029389.



## REFERENCES

- [1] Isaac Arvestad. 2022. Near-linear time expander decomposition in practice.
- [2] Bas Fagginger Auer and Rob H Bisseling. 2012. Graph coarsening and clustering on the GPU. *Graph Partitioning and Graph Clustering* 588, 223 (2012), 2.
- [3] Charles-Edmond Bichot. 2010. Co-clustering Documents and Words by Minimizing the Normalized Cut Objective Function. *J. Math. Model. Algorithms* 9, 2 (2010), 131–147. <https://doi.org/10.1007/S10852-010-9126-0>
- [4] Deng Cai, Zheng Shao, Xiaofei He, Xifeng Yan, and Jiawei Han. 2005. Mining hidden community in heterogeneous social networks. In *Proceedings of the 3rd international workshop on Link discovery, LinkKDD 2005, Chicago, Illinois, USA, August 21-25, 2005*, Jafar Adibi, Marko Grobelnik, Dunja Mladenic, and Patrick Pantel (Eds.). ACM, 58–65. <https://doi.org/10.1145/1134271.1134280>
- [5] Yi-Jun Chang, Seth Pettie, Thatchaphol Saranurak, and Hengjie Zhang. 2021. Near-optimal Distributed Triangle Enumeration via Expander Decompositions. *J. ACM* 68, 3 (2021), 21:1–21:36. <https://doi.org/10.1145/3446330>
- [6] Li Chen, Rasmus Kyng, Yang P. Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. 2022. Maximum Flow and Minimum-Cost Flow in Almost-Linear Time. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 - November 3, 2022*. IEEE, 612–623. <https://doi.org/10.1109/FOCS54457.2022.00064>
- [7] Xiaojun Chen, Feiping Nie, Joshua Zhexue Huang, and Min Yang. 2017. Scalable Normalized Cut with Improved Spectral Rotation. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, Carles Sierra (Ed.). ijcai.org, 1518–1524. <https://doi.org/10.24963/IJCAI.2017/210>
- [8] Amir Daneshgar and Ramin Javadi. 2012. On the complexity of isoperimetric problems on trees. *Discret. Appl. Math.* 160, 1-2 (2012), 116–131. <https://doi.org/10.1016/J.DAM.2011.08.015>
- [9] Timothy A Davis and Yifan Hu. 2011. The University of Florida sparse matrix collection. *ACM Transactions on Mathematical Software (TOMS)* 38, 1 (2011), 1–25.
- [10] Inderjit S Dhillon, Yuqiang Guan, and Brian Kulis. 2007. Weighted graph cuts without eigenvectors a multilevel approach. *IEEE transactions on pattern analysis and machine intelligence* 29, 11 (2007), 1944–1957.
- [11] Gramoz Goranci, Harald Räcke, Thatchaphol Saranurak, and Zihan Tan. 2021. The expander hierarchy and its applications to dynamic graph algorithms. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2212–2228.
- [12] Bernhard Haeupler, Harald Räcke, and Mohsen Ghaffari. 2022. Hop-constrained expander decompositions, oblivious routing, and distributed universal optimality. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. 1325–1338.
- [13] Kathrin Hanauer, Monika Henzinger, Robin Münk, Harald Räcke, and Maximilian Vötsch. 2024. Expander Hierarchies for Normalized Cuts on Graphs. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2023, August 25–29, 2024, Barcelona, Spain*. ACM. <https://doi.org/10.1145/3637528.3671978> To appear.
- [14] Kathrin Hanauer, Monika Henzinger, Robin Münk, Harald Räcke, and Maximilian Vötsch. 2024. XCut/XCut v1.0.0. <https://doi.org/10.5281/zenodo.12108189>
- [15] Bruce Hendrickson and Robert W. Leland. 1995. A Multi-Level Algorithm For Partitioning Graphs. In *Proceedings Supercomputing '95, San Diego, CA, USA, December 4-8, 1995*, Sidney Karin (Ed.). ACM, 28. <https://doi.org/10.1145/224170.224228>
- [16] Monika Henzinger, Stefan Neumann, Harald Räcke, and Stefan Schmid. 2023. Dynamic Maintenance of Monotone Dynamic Programs and Applications. In *40th International Symposium on Theoretical Aspects of Computer Science, STACS 2023, March 7-9, 2023, Hamburg, Germany (LIPIcs, Vol. 254)*, Petra Berenbrink, Patricia Bouyer, Anuj Dawar, and Mamadou Moustapha Kanté (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 36:1–36:16. <https://doi.org/10.4230/LIPICS.STACS.2023.36>
- [17] Yiding Hua, Rasmus Kyng, Maximilian Probst Gutenberg, and Zihang Wu. 2023. Maintaining expander decompositions via sparse cuts. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 48–69.
- [18] George Karypis and Vipin Kumar. 1998. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on scientific Computing* 20, 1 (1998), 359–392.
- [19] Jonathan A. Kelner, Yin Tat Lee, Lorenzo Orecchia, and Aaron Sidford. 2014. An Almost-Linear-Time Algorithm for Approximate Max Flow in Undirected Graphs, and its Multicommodity Generalizations. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, Chandra Chekuri (Ed.). SIAM, 217–226. <https://doi.org/10.1137/1.9781611973402.16>
- [20] B. Laurent and P. Massart. 2000. Adaptive estimation of a quadratic functional by model selection. *The Annals of Statistics* 28, 5 (2000), 1302 – 1338. <https://doi.org/10.1214/aos/1015957395>
- [21] Jason Li. 2021. Deterministic mincut in almost-linear time. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 384–395.
- [22] Jia Li, Honglei Zhang, Zhichao Han, Yu Rong, Hong Cheng, and Junzhou Huang. 2020. Adversarial Attack on Community Detection by Hiding Individuals. In *WWW '20: The Web Conference 2020, Taipei, Taiwan, April 20-24, 2020*, Yennun Huang, Irwin King, Tie-Yan Liu, and Maarten van Steen (Eds.). ACM / IW3C2, 917–927. <https://doi.org/10.1145/3366423.3380171>
- [23] Bojan Mohar. 1989. Isoperimetric numbers of graphs. *J. Comb. Theory, Ser. B* 47, 3 (1989), 274–291. [https://doi.org/10.1016/0095-8956\(89\)90029-4](https://doi.org/10.1016/0095-8956(89)90029-4)
- [24] Danupon Nanongkai and Thatchaphol Saranurak. 2017. Dynamic spanning forest with worst-case update time: adaptive, las vegas, and  $o(n^{1/2-\epsilon})$ -time. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. 1122–1129.
- [25] Danupon Nanongkai, Thatchaphol Saranurak, and Christian Wulff-Nilsen. 2017. Dynamic minimum spanning forest with subpolynomial worst-case update time. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 950–961.
- [26] Maxim Naumov and Timothy Moon. 2016. Parallel spectral graph partitioning. *NVIDIA, Santa Clara, CA, USA, Tech. Rep., NVR-2016-001* (2016).
- [27] Andrew Y. Ng, Michael I. Jordan, and Yair Weiss. 2001. On Spectral Clustering: Analysis and an algorithm. In *Advances in Neural Information Processing Systems 14 [Neural Information Processing Systems: Natural and Synthetic, NIPS 2001, December 3-8, 2001, Vancouver, British Columbia, Canada]*, Thomas G. Dietterich, Suzanna Becker, and Zoubin Ghahramani (Eds.). MIT Press, 849–856. <https://proceedings.neurips.cc/paper/2001/hash/801272ee79cfd7fa5960571fee369b-Abstract.html>
- [28] Feiping Nie, Jitao Lu, Danyang Wu, Rong Wang, and Xuelong Li. 2024. A Novel Normalized-Cut Solver With Nearest Neighbor Hierarchical Initialization. *IEEE Trans. Pattern Anal. Mach. Intell.* 46, 1 (2024), 659–666. <https://doi.org/10.1109/TPAMI.2023.3279394>
- [29] Vitaly Osipov and Peter Sanders. 2010. n-Level graph partitioning. In *Algorithms–ESA 2010: 18th Annual European Symposium, Liverpool, UK, September 6-8, 2010. Proceedings, Part 1* 18. Springer, 278–289.
- [30] Richard Peng, He Sun, and Luca Zanetti. 2017. Partitioning Well-Clustered Graphs: Spectral Clustering Works! *SIAM J. Comput.* 46, 2 (2017), 710–743. <https://doi.org/10.1137/15M1047209>
- [31] Ryan A. Rossi and Nesreen K. Ahmed. 2015. The Network Data Repository with Interactive Graph Analytics and Visualization. In *AAAI*. <https://networkrepository.com>
- [32] Tapasmini Sahoo and Kunal Kumar Das. 2023. Brain Tumor Localization Using N-Cut. *International Journal of Online & Biomedical Engineering* 19, 15 (2023).
- [33] Peter Sanders and Christian Schulz. 2012. High quality graph partitioning. *Graph Partitioning and Graph Clustering* 588, 1 (2012), 1–17.
- [34] Thatchaphol Saranurak and Di Wang. 2019. Expander Decomposition and Pruning: Faster, Stronger, and Simpler. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, Timothy M. Chan (Ed.). SIAM, 2616–2635. <https://doi.org/10.1137/1.9781611975482.162>
- [35] Jianbo Shi and Jitendra Malik. 2000. Normalized Cuts and Image Segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* 22, 8 (2000), 888–905. <https://doi.org/10.1109/34.868688>
- [36] Daniel A Spielman and Shang-Hua Teng. 2004. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*. 81–90.
- [37] Chris Walshaw, Mark Cross, and Martin G. Everett. 1995. A Localized Algorithm for Optimizing Unstructured Mesh Partitions. *Int. J. High Perform. Comput. Appl.* 9, 4 (1995), 280–295. <https://doi.org/10.1177/109434209500900403>
- [38] Yangtao Wang, Xi Shen, Yuan Yuan, Yuming Du, Maomao Li, Shell Xu Hu, James L. Crowley, and Dominique Vaufreydaz. 2023. TokenCut: Segmenting Objects in Images and Videos With Self-Supervised Transformer and Normalized Cut. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45, 12 (2023), 15790–15801. <https://doi.org/10.1109/TPAMI.2023.3305122>
- [39] Christian Wulff-Nilsen. 2017. Fully-dynamic minimum spanning forest with improved worst-case update time. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. 1130–1143.
- [40] Eric P. Xing and Richard M. Karp. 2001. CLIFF: clustering of high-dimensional microarray data via iterative feature filtering using normalized cuts. In *Proceedings of the Ninth International Conference on Intelligent Systems for Molecular Biology, July 21-25, 2001, Copenhagen, Denmark*. 306–315.
- [41] Stella X. Yu and Jianbo Shi. 2003. Multiclass Spectral Clustering. In *9th IEEE International Conference on Computer Vision (ICCV 2003), 14-17 October 2003, Nice, France*. IEEE Computer Society, 313–319. <https://doi.org/10.1109/ICCV.2003.1238361>
- [42] Jin Zhang, Lei Xie, Wei Feng, and Yanning Zhang. 2009. A Subword Normalized Cut Approach to Automatic Story Segmentation of Chinese Broadcast News. In *Information Retrieval Technology, 5th Asia Information Retrieval Symposium, AIRS 2009, Sapporo, Japan, October 21-23, 2009. Proceedings (Lecture Notes in Computer Science, Vol. 5839)*, Gary Geunbae Lee, Dawei Song, Chin-Yew Lin, Akiko N. Aizawa, Kazuko Kuriyama, Masaharu Yoshioka, and Tetsuya Sakai

(Eds.). Springer, 136–148. [https://doi.org/10.1007/978-3-642-04769-5\\_12](https://doi.org/10.1007/978-3-642-04769-5_12)

[43] Zhiqiang Zhao, Yongyu Wang, and Zhuo Feng. 2018. Nearly-linear time spectral graph reduction for scalable graph partitioning and data visualization. *arXiv preprint arXiv:1812.08942* (2018).

## A THEORETICAL ANALYSIS OF THE EXPANDER DECOMPOSITION

The key ingredient for the analysis is to show that it is very unlikely that the cut procedure does not return a cut within  $T$  iterations when started on a graph with conductance less than  $\phi$ . Consequently, if the algorithm does not return a cut for  $T$  iterations we can declare  $G$  to be a  $\phi$ -expander (or actually  $G\{A\}$  to be a near  $6\phi$ -expander), with a small probability of error.

For this we analyze the random walks in terms of flows. Each node  $v$  injects  $d_v$  units of flow of a unique commodity. This flow is distributed according to the random walk. Let  $F_{ij}(t)$  denote the amount of flow from node  $j$  that has reached node  $i$  after  $t$  steps. We define  $P_{ij}(t) := \frac{F_{ij}(t)}{d_i d_j}$  for all  $t$ . Whenever clear from context, we may omit the explicit indication of the round  $t$ .

Let  $A(t) \subseteq V$  be the set of nodes in the subgraph in round  $t$ . We define a natural average vector  $\mu(t) := \frac{1}{\text{vol}(A(t))} \sum_{i \in A(t)} d_i P_i(t)$  and track the convergence of our random walk to this stationary distribution with a potential function:

$$\varphi(t) := \sum_{i \in A(t)} d_i \cdot \|P_i(t) - \mu(t)\|^2.$$

Intuitively, a low potential indicates that the  $P(t)$  vectors have mixed well in the set  $A(t)$ . The following lemma shows that a low potential implies that  $A(t)$  is a near  $6\phi$ -expander in the graph  $G(G\{A(t)\})$  may not be a proper expander because the random walk may have used edges outside of  $G\{A(t)\}$ .

LEMMA 1. *If  $\varphi(t) \leq \frac{1}{4 \text{vol}(V)^2}$  in any step  $t \leq T$ , then  $A(t)$  is a near  $6\phi$ -expander in  $G$ .*

PROOF. We omit the time step  $t$  to avoid notational clutter. The random walk can be viewed as establishing a multicommodity flow within the network.  $d_i P_{ij}$  is the flow of commodity  $j$  that reached node  $i$ . The flow has congestion at most  $t$ —the number of steps of the random walk. We take at most  $T = 1/(12\phi)$  steps.

Now consider a cut  $S \subseteq A$ , with  $\text{vol}(S) \leq \text{vol}(A)/2$ . We have to show that  $|E(S, V \setminus S)| \geq \text{vol}(S)/(2T)$ . For the potential to be less than  $1/(4 \text{vol}(V)^2)$ , we must have  $P_{ij} \leq \mu_j + \frac{1}{2 \text{vol}(V)}$  for every node  $i \in A$  and any  $j$ . The flow that originates in  $S$  and stays in  $S$  is

$$\begin{aligned} \sum_{j \in S} \sum_{i \in S} d_i d_j P_{ij} &\leq \sum_{j \in S} \sum_{i \in S} d_i d_j \left( \mu_j + \frac{1}{2 \text{vol}(V)} \right) \\ &= \text{vol}(S) \sum_{j \in S} d_j \mu_j + \frac{\text{vol}(S)^2}{2 \text{vol}(V)} \\ &\leq \frac{\text{vol}(S)^2}{\text{vol}(A)} + \frac{\text{vol}(S)^2}{2 \text{vol}(V)} \leq \frac{3}{4} \text{vol}(S). \end{aligned}$$

Here the third step follows because  $d_j \mu_j \leq d_j / \text{vol}(A)$  as  $\mu_j$  is the fraction of the flow of commodity  $j$  that stays in  $A$  (normalized by  $1/\text{vol}(A)$ ). This means that at least  $1/4$  of the flow that starts in  $S$  has to leave  $S$ . At the same time an equivalent amount of flow has to enter  $S$ , which means that the traffic across the edges in  $E(S, V \setminus S)$  is at least  $\text{vol}(S)/2$ . As the congestion is only  $T$ , there must be at least  $\frac{1}{2T} \text{vol}(S) \geq 6\phi \text{vol}(S)$  edges across the cut.  $\square$

In the following we show that with constant probability during one iteration we either return a balanced low conductance cut or the potential decreases significantly. For this we first have to analyze by how much the potential decreases during a random walk step.

*Potential Decrease by Random Walk.* Fix a round  $t$ . Note that the random walk for round  $t$  only takes the  $t-1$  random walk steps for the graphs  $G\{A(1)\}, \dots, G\{A(t-1)\}$ , the step for  $G\{A(t)\}$  follows in the next round. In the following we develop an expression for how much the potential will decrease due to the random walk step in the current iteration  $t$ . For this we need a technical claim, which is proven in Section A.1.

CLAIM 1. *Let  $a_1, \dots, a_d, \mu$  be vectors of the same dimension. Then,*

$$d \left\| \frac{1}{d} \sum_i a_i - \mu \right\|^2 - \sum_i \|a_i - \mu\|^2 = \frac{1}{d} \left\| \sum_i a_i \right\|^2 - \sum_i \|a_i\|^2 \leq 0.$$

*Specifically, for  $d = 2$  the value is equal to  $-\frac{1}{2} \|a_1 - a_2\|^2$ .*

Let  $\delta(t) := (\varphi(t) - \varphi(t+1))/\varphi(t)$  be the relative factor by which the potential of round  $t$  decreases after the random walk step in graph  $G\{A(t)\}$ .

LEMMA 2. *The relative potential decrease due to the random walk step in iteration  $t$  is at least*

$$\delta(t) := \frac{1}{2} \frac{\sum_{\{i,j\} \in G\{A(t)\}} \|P_i(t) - P_j(t)\|_2^2}{\sum_{i \in A(t)} d_i \|P_i(t) - \mu(t)\|_2^2}.$$

PROOF. To simplify the analysis of a random walk step we view each vertex  $v$  of  $G\{A(t)\}$  as consisting of  $d_v$  sub-vertices as follows. For each edge  $\{x, y\}$  we introduce two sub-vertices – one at  $x$  and one at  $y$ , that are connected by an edge. In this way, every super-vertex  $i$  from  $A(t)$  receives  $d_i$  sub-vertices and every sub-vertex has a unique neighbour.

During a random walk step, each sub-vertex  $x$  first takes a copy of the flow vector of its super-vertex. Then the vectors are averaged along every edge, and finally a super-vertex computes the average of the vector of its sub-vertices.

We can express the potential by either summing over super-vertices or sub-vertices:

$$\varphi(t) = \sum_{i \in A(t)} d_i \|P_i - \mu\|_2^2 = \sum_x \|P_{\text{sup}(x)} - \mu\|_2^2,$$

where  $\text{sup}(x) \in A(t)$  is the super-vertex for sub-vertex  $x$ . Averaging the vectors between two sub-vertices  $x$  and  $y$  changes the potential by

$$\begin{aligned} &2 \left\| \frac{1}{2} (P_{\text{sup}(x)} + P_{\text{sup}(y)}) - \mu \right\|_2^2 \\ &- \|P_{\text{sup}(x)} - \mu\|_2^2 - \|P_{\text{sup}(y)} - \mu\|_2^2, \end{aligned} \quad (1)$$

which is equal to  $-\frac{1}{2} \|P_{\text{sup}(x)} - P_{\text{sup}(y)}\|_2^2$  by Claim 1. Let  $\psi$  denote the potential after averaging along edges but before averaging among sub-vertices. Similarly, let for a sub-node  $x$ ,  $Q_x$  denote the vector of  $x$  at this point. Summing Equation 1 over all edges gives

$$\psi - \varphi(t) \leq \sum_{\{x,y\} \in E(A(t))} -\frac{1}{2} \|P_{\text{sup}(x)} - P_{\text{sup}(y)}\|_2^2$$

and

$$\begin{aligned} \psi &= \sum_x \|Q_x - \mu\|_2^2 = \sum_{i \in A(t)} \sum_{x \in \text{sub}(i)} \|Q_x - \mu\|_2^2 \\ &\geq \sum_{i \in A(t)} d_i \left\| \frac{1}{d_i} \sum_{x \in \text{sub}(i)} Q_x - \mu \right\|_2^2, \end{aligned}$$

where the inequality also follows from Claim 1. This means averaging among edges decreased the potential by  $\frac{1}{2} \sum_{\{i,j\}} \|P_i - P_j\|_2^2$  and averaging among sub-nodes does not increase it again. Hence, the lemma follows.  $\square$

*Projected Potential Decrease.* The algorithm does not maintain the  $P_i$  vectors or the potential explicitly, as that would be computationally infeasible. Rather it operates on their projections onto some random vector  $r$ . The entries in the vector  $u$  of the algorithm are actually  $u_i = (P_i - \mu)^\top r$ . This follows since performing the random walk and then projecting the resulting  $P_i$  vectors onto  $r$  is equivalent to first projecting onto  $r$  and then running the random walk on the vector of projections. Ultimately, subtracting the weighted average of the  $u$  values corresponds to subtracting  $\mu$  from the  $P_i$  vectors before the projection.

Consider the quotient

$$R(u) := \frac{\sum_{\{i,j\} \in E(A(t))} (u_i - u_j)^2}{\sum_{i \in A(t)} d_i u_i^2},$$

and compare it to the bound  $\delta(t)$  for the relative potential decrease in Lemma 2. Up to a factor of 1/2,  $R(u)$  is obtained by individually performing a random projection on each vector in the expression for  $\delta(t)$ . Using properties of random projections, we will show that  $R(u)$  is a good indicator for the potential decrease  $\delta(t)$ .

LEMMA 3 (PROJECTION LEMMA). *Let  $v_1, \dots, v_n$  be a collection of  $d$ -dimensional vectors. Let  $r$  denote a random  $d$ -dimensional vector with each coordinate sampled independently from a Gaussian distribution  $\mathcal{N}(0, 1/d)$ . Then*

- (1)  $\Pr[\exists i, j : (v_i - v_j)^\top r \geq 11 \log n \cdot \|v_i - v_j\|^2/d] \leq 1/n$
- (2)  $\Pr[\sum_i (v_i^\top r)^2 \geq \frac{1}{20 \log n} \sum_i \|v_i\|^2/d] \geq 1/2$  for  $n \geq 8$ .

We call an iteration *good* if the vector of projections  $u$  retains sufficient information of the higher dimensional  $P_i$  vectors. Formally, we require the following.

DEFINITION 2. *An iteration  $t$  of the algorithm is good if*

- $\sum_{i \in A(t)} d_i u_i^2 \geq \frac{1}{20 n \log n} \cdot \sum_i d_i \|P_i - \mu\|^2$  and
- $(u_i - u_j)^2 \leq \frac{11 \log n}{n} \|P_i - P_j\|^2 \quad \forall \{i, j\} \in E(A(t))$ .

From Property 1 and Property 2 of the Projection Lemma we can directly infer the following.

CLAIM 2. *An iteration is good with probability at least 1/4 for  $n \geq 4$ .*

The next lemma shows that in a good round a large  $R(u)$  value implies a large  $\delta(t)$  value, i.e., a large potential decrease.

LEMMA 4. *In a good round  $t$  we have  $R(u) \leq 440 \log^2(n) \cdot \delta(t)$ .*

PROOF.

$$\begin{aligned} \delta(t) &= \frac{1}{2} \frac{\sum_{\{i,j\} \in G\{A(t)\}} \|P_i - P_j\|_2^2}{\sum_{i \in A(t)} d_i \|P_i - \mu\|_2^2} \\ &\geq \frac{1}{2} \frac{\frac{n}{11 \log n} \sum_{\{i,j\} \in G\{A(t)\}} (u_i - u_j)^2}{20 n \log n \sum_{i \in A(t)} d_i u_i^2} \end{aligned}$$

by the definition of  $\delta(t)$  and the definition of a good round.  $\square$

**For the further analysis we will always assume that we are in a good round.**

*Finding a Cut.* With the above lemma we have established, that a large  $R(u)$  value ensures the next step of the random walk reduces the potential a lot. Now we show that a small  $R(u)$  value ensures that the algorithm finds a low conductance cut.

Indeed the quotient  $R(u)$  (the Rayleigh quotient of  $u$  using the graph Laplacian), and its analysis lies at the heart of the proof for the celebrated Cheeger inequality. It states that a small  $R(u)$  value implies the existence of a low conductance sweep cut.

LEMMA 5 (CHEEGER). *For any  $x \in \mathbb{R}^n$ , with  $x \perp \vec{d}$  there is a value  $c$ , s.t. there is a non-trivial set  $S_c = \{i \in V : x_i \leq c\}$  with conductance  $\Phi_G(S_c) \leq \sqrt{2R(x)}$ .*

Note that  $u \perp \vec{d}$  holds by design of the algorithm due to Line 10 of the algorithm. With this Cheeger Lemma we thus get the guarantee that there has to be a sweep cut through  $u$  whose conductance  $\Phi$  satisfies  $\Phi \leq \sqrt{2R(u)}$ . Together with Lemma 4 this now implies the following: If we will only make little progress in the potential with the next random walk step, i.e.,  $\delta(t) \leq \alpha$  for some small  $\alpha$ , then we find a sweep cut through the vector of projections  $u$  with conductance  $\Phi \leq O(\sqrt{\alpha} \log n)$ . By contraposition, if all sweep cuts have conductance at least  $\Phi$ , then  $\delta(t) \geq \Omega(\Phi^2 / \log^2 n)$ .

This motivates our approach of analyzing the sweep cuts on  $u$  and differentiating three cases: 1) No low conductance sweep cut exists, 2) A balanced low conductance cut exists, or 3) An unbalanced low conductance cut exists.

If we find a balanced sweep cut, we return immediately as this ensures the recursion depth of the surrounding expander decomposition remains small. If no sweep cut has conductance below  $\gamma$ , the above reasoning ensures we make sufficient progress with the next random walk step. Next we prove that even if we find an unbalanced cut, the potential still decreases sufficiently.

*Handling Unbalanced Cuts.* Let  $(S, B)$  be an unbalanced cut, where  $\text{vol}(S) \leq 2 \log^2(\text{vol}(A)) / \text{vol}(A)$  and there is no  $\gamma$ -low conductance sweep cut through  $B$  in  $u$ . We show that if  $R(u)$  is small and we thus cannot ensure sufficient progress from the random walk, then the set  $S$  must contain a large fraction of the potential. Hence we instead make progress by removing  $S$  when we set  $A(t+1) = B$ .

Consider the following lemma, which gives an upper bound on how much the  $R(u)$  value can at most increase if we remove a cut and recenter the remaining projected values. Recentering the vector after removing the cut allows us to use Lemma 5 to argue about the existence of low conductance cuts in the remaining subgraph. Let  $z_B \in \mathbb{R}^{|B|}$  be the shortened and recentered vector of the projections  $u$ , i.e.,  $z_i := u_i - \bar{u}_B$  for all  $i \in B$ <sup>9</sup>. By design, it then holds  $z_B \perp \vec{d}$ .

LEMMA 6. *For any cut  $(S, B)$  in  $A$  and  $\lambda < \text{vol}(B) / \text{vol}(A)$  such that  $\sum_{i \in S} d_i u_i^2 \leq \lambda \cdot \sum_{i \in A} d_i u_i^2$ ,*

$$R(z_B) \leq \frac{\text{vol}(B)}{\text{vol}(B) - \lambda \text{vol}(A)} \cdot R(u).$$

The proof is deferred to Section A.1. We can put this in the context of our unbalanced cut  $(S, B)$  with the following lemma. For this, let  $\varphi_S$  and  $\varphi_B$  denote the potential generated by summing only over the nodes in  $S$  and  $B$  respectively.

<sup>9</sup>For any set  $X \subseteq V$  and vector  $x \in \mathbb{R}^{|X|}$  we denote by  $\bar{x}_X := \frac{1}{\text{vol}(X)} \sum_{i \in X} d_i x_i$  the degree-weighted average of the  $x$  values in  $X$ .

LEMMA 7. *If  $R(u) < \gamma^2/4$  and we find an unbalanced cut  $(S, B)$  with  $\text{vol}(S) \leq 2m/\log^2 m$ , then  $\varphi_B(t) \leq (1 - 1/(880 \log^2 n)) \cdot \varphi(t)$ .*

PROOF. Since  $(S, B)$  partitions  $A$  we get  $\varphi_S(t) + \varphi_B(t) = \varphi(t)$ . Assume for contradiction that  $\varphi_S(t) < 1/(880 \log^2 n) \cdot \varphi(t)$ , i.e., with  $S$  we only remove very little potential. It follows

$$\sum_{i \in S} d_i u_i^2 \leq \frac{11 \log n}{n} \varphi_S(t) < \frac{1}{80n \log n} \varphi(t) \leq \frac{1}{4} \sum_{i \in A} d_i u_i^2 \quad (2)$$

since round  $t$  is good. Note that  $\frac{\text{vol}(B)}{\text{vol}(A)} \geq 1 - \frac{2 \log^2(\text{vol}(A))}{\text{vol}(A)} \geq \frac{3}{4}$  and therefore

$$\frac{\text{vol}(B)}{\text{vol}(B) - 1/4 \cdot \text{vol}(A)} \leq \frac{\text{vol}(A)}{(3/4 - 1/4) \text{vol}(A)} \leq 2. \quad (3)$$

Summing over  $B$  gives  $\sum_i d_i z_i = \sum_i d_i u_i - \text{vol}(B) \bar{u}_B = 0$ . This allows us to use Lemma 5 in the subgraph  $G\{B\}$ . By design of the algorithm, since we did not find a balanced low conductance cut,  $(S, B)$  is a two-ended sweep cut where  $S$  has maximal volume among low conductance cuts. Hence there is no low conductance sweep cut through  $u$  that crosses the set  $B$ . We get

$$\begin{aligned} \gamma^2 &\leq 2R(z_B) && \text{by Lemma 5} \\ &\leq 2 \frac{\text{vol}(B)}{\text{vol}(B) - 1/4 \cdot \text{vol}(A)} \cdot R(u) && \text{by Lemma 6 using (2)} \\ &\leq 4R(u) && \text{by (3)}. \end{aligned}$$

This is a contradiction to our assumption that  $R(u) < \gamma^2/4$ .  $\square$

The above results show that a good round ensures sufficient progress in decreasing the potential or produces a balanced cut.

CLAIM 3. *In a good round we either find a balanced cut or the potential reduces by a factor at least  $1 - \gamma^2/(1760 \log^2 n)$ .*

PROOF. Let  $\Phi$  be the minimum value of any sweep cut on  $u$ . If  $\Phi \geq \gamma$ , then Lemmas 4 and 5 give  $\delta \geq \gamma^2/(880 \log^2 n)$  since

$$\gamma^2 \leq \Phi^2 \leq 2R(u) \leq 880 \log^2(n) \cdot \delta.$$

If  $\Phi < \gamma$ , we will find some low conductance cut  $S$ . If  $S$  is balanced, we can return it, otherwise we move  $S$  from  $A$  to  $R$ . Then either

- $R(u) \geq \gamma^2/4$  and Lemmas 4 and 5 give  $\delta \geq \gamma^2/(1760 \log^2 n)$ , or
- $R(u) < \gamma^2/4$  and Lemma 7 gives a factor of  $1 - 1/(880 \log^2 n)$ .

In any case, the decrease is at least  $1 - \gamma^2/(1760 \log^2 n)$ .  $\square$

*Iterations.* An iteration takes  $O(m)$  time for the random walk and  $O(n \log n)$  for sorting the  $u$  values. It only remains to prove the upper bound on the number of iterations. The threshold for the potential is chosen to guarantee the existence of a near  $6\phi$ -expander in Lemma 1.

LEMMA 8. *After  $T = 1/12\phi$ , iterations of the cut procedure, with high probability  $\varphi(T) \leq \frac{1}{4 \text{vol}(V)^2}$ .*

PROOF. Let  $\text{vol} = \text{vol}(V)$ . The initial potential is

$$\begin{aligned} \varphi(0) &= \sum_i d_i \cdot \left( (n-1) \cdot \frac{1}{\text{vol}^2} + \left(1 - \frac{1}{\text{vol}}\right)^2 \right) \\ &= \text{vol} \cdot \left( \frac{n}{\text{vol}^2} - \frac{2}{\text{vol}} + 1 \right) \leq \text{vol}. \end{aligned}$$

If we reach round  $T$ , then none of the previous iterations produced a balanced cut. By Claim 3, a good round decreases the potential by a factor of  $1 - \gamma^2/(1760 \log^2 n)$ . We first show that reducing the

initial potential by this factor  $\frac{1}{96\phi}$  times brings it below  $\frac{1}{4 \text{vol}^2}$ , then we show that with high probability there are sufficient good rounds. After  $k$  good rounds, the potential has decreased by a factor of

$$\left(1 - \frac{\gamma^2}{1760 \log^2 n}\right)^k < \exp\left(-\frac{\gamma^2}{1760 \log n}\right)^{\frac{k}{\log n}} = \exp\left(-\frac{\gamma^2}{1760 \log^2 n} \cdot k\right)$$

where the first inequality follows from  $(1+x/a)^a < e^x, \forall x \in \mathbb{R}, a > 0$  with  $a = \log n$ . Combining this with the above bound on  $\varphi(0)$ , we can take logarithms on both sides to see that  $\varphi(T) \leq 1/(4 \text{vol}^2)$  if  $k \cdot \gamma^2/(1760 \log^2 n) \geq \ln(4 \text{vol}^3)$  and thus we need at least

$$1760 \frac{1}{\gamma^2} \ln(4 \text{vol}^3) \log^2 n \leq 1220 \frac{1}{\gamma^2} \log(4 \text{vol}^3) \log^2(n)$$

good rounds. Setting

$$\gamma = 343 \sqrt{\phi \log(4 \text{vol}^3) \log^2(n)}$$

gives that after  $\frac{1}{96\phi}$  good rounds, the potential is below the threshold.

By Claim 2, a round is good with probability  $1/4$ . Let  $X_i$  be the indicator random variable representing whether round  $i$  was good and  $K = \sum_i^T X_i$  their sum. Clearly,  $K$  is the number of good rounds and its expectation is  $\mathbb{E}[K] = T/4$ . Since  $\phi \leq 1/\log^2 n$  we get  $T \geq \log^2(n)/12$  and thus  $\mathbb{E}[K] \geq \log^2(n)/48$ . Using a Chernoff bound then gives

$$\Pr\left[K \leq \frac{1}{2} \mathbb{E}[K]\right] \leq \exp\left(-\frac{1}{8} \mathbb{E}[K]\right) \leq n^{-\frac{\log n}{48}}.$$

It follows that with high probability the number of good rounds is

$$K \geq \frac{1}{2} \mathbb{E}[K] = \frac{T}{8} \geq \frac{1}{96\phi}.$$

$\square$

Altogether, the above arguments imply the correctness of our Theorem 2. Analogous to the reasoning by Saranurak and Wang [34], we can conclude that our modified guarantees yield Theorem 1 when using our cut step procedure in their expander decomposition framework.

## A.1 Deferred Proofs

LEMMA 9 (PROJECTION LEMMA). *Let  $v_1, \dots, v_n$  be a collection of  $d$ -dimensional vectors. Let  $r$  denote a random  $d$ -dimensional vector with each coordinate sampled independently from a Gaussian distribution  $\mathcal{N}(0, 1/d)$ . Then*


- (1)  $\mathbb{E}[(v_i^\top r)^2] = \|v_i\|^2/d$  for all  $i$
- (2)  $\Pr[(v_i^\top r)^2 \geq \alpha \cdot \|v_i\|^2/d] \leq e^{-\alpha/5}$  for  $\alpha \geq 1$
- (3)  $\Pr[\exists i, j : (v_i - v_j)^\top r \geq 11 \log n \cdot \|v_i - v_j\|^2/d] \leq 1/n$
- (4)  $\Pr[\sum_i (v_i^\top r)^2 \geq \frac{1}{20 \log n} \sum_i \|v_i\|^2/d] \geq 1/2$  for  $n \geq 8$ .

PROOF OF LEMMA 3.

1. Due to the rotation symmetry of the Gaussian distribution we can assume wlog. that  $v$  is an arbitrary vector of length  $\ell = \|v\|$ . In particular we can assume that  $v_1 = \sqrt{\ell}$  and all other entries are 0. Then  $\mathbb{E}[(v^\top r)^2] = \mathbb{E}[\ell \cdot \mathcal{N}(0, 1/d)^2] = \ell/d \cdot \mathbb{E}[\mathcal{N}(0, 1)^2] = \ell/d$ .

2. Using Lemma 1 from Laurent and Massart [20] one gets that for a  $\chi$ -squared distributed variable  $X \sim \mathcal{N}(0, 1)^2$  fulfills  $\Pr[X \geq 1+2\sqrt{x+2x}] \leq e^{-x}$ , which gives  $\Pr[X \geq 5x] \leq e^{-x}$  for  $x \geq 1$ . Since the distribution of  $(v^\top r)^2$  is  $\frac{\ell}{d} \cdot \mathcal{N}(0, 1)^2$  the statement follows.

3. Using Property 2 with  $\alpha = 15 \ln n$  gives that the probability that a vector is overstretched by a factor more than  $\alpha$  is at most



**Figure D.6: XCut with the greedy vs. dynamic programming strategy for a subset of instances for  $k = 32$  and  $\rho = 0.0001$ .**

$n^{-3}$ . Applying a union bound over the at most  $n^2$  vectors gives the statement. Finally, note that  $15 \ln x < 11 \log_2 x$  for all  $x > 1$ .

4. Let  $x_j$  denote the vector where the  $i$ -th entry is  $v_{ij}$ . Then  $Z := \sum_i (v_i^T r)^2 = \sum_i \sum_j (v_{ij} r_j)^2 = \sum_j \sum_i (x_{ij} r_j)^2 = \sum_j \|x_j\|_2^2 \cdot (r_j)^2$ . Let  $\ell_{\max}$  denote the largest length of an  $x_j$ -vector. We classify the  $x_j$ -vectors whose length fall in the range  $(\ell_{\max}/n, \ell_{\max}]$  into  $\log_2 n$  classes so that the length of vectors in the same class differs by at most a factor of 2. Formally, the  $i$ -th class for  $i \in 1, \dots, \log_2 n$  contains vectors with length in the range  $(\ell_{\max}/2^i, \ell_{\max}/2^{i-1}]$ . Scale the length for each classified vector  $x_j$  down to the lowest length in its class and scale the unclassified vectors down to length 0. Let  $\tilde{x}_j$  denote scaled length of a vector  $x_j$ . We have

$$\begin{aligned} \sum_j \|\tilde{x}_j\|_2^2 &\geq \sum_j \text{classified} \|x_j\|^2 / 4 \geq \sum_j \|x_j\|^2 / 4 - n(\ell_{\max}/n)^2 \\ &\geq \left(\frac{1}{4} - \frac{1}{n}\right) \sum_j \|x_j\|^2 \geq \frac{1}{8} \sum_j \|x_j\|^2. \end{aligned}$$

for  $n \geq 4$ . The second inequality holds because an unclassified vector has at most length  $\ell_{\max}/n$  and the third inequality because there is one vector  $x_j$  with length  $\ell_{\max}$ . Now, we choose the class where the total scaled length of vectors is largest. Let  $\mathcal{R}$  denote the index set of vectors in this class, let  $\ell_{\mathcal{R}}$  denote their scaled length and let  $n_{\mathcal{R}}$  denote their number.

Then  $n_{\mathcal{R}} \ell_{\mathcal{R}}^2 \geq \frac{1}{8 \log_2 n} \sum_j \|x_j\|^2$ . Now, we consider the sum of squared projections of the scaled vectors in  $\mathcal{R}$ . This is distributed according to a  $\chi^2$ -distribution with  $n_{\mathcal{R}}$  degrees of freedom scaled by  $\ell_{\mathcal{R}}^2/d$ . The median for this distribution is  $n_{\mathcal{R}}(1 - 2/(9n_{\mathcal{R}})) \geq 0.4n_{\mathcal{R}}$ . Hence, with probability  $1/2$

$$Z \geq Z_{\mathcal{R}} \geq 0.4n_{\mathcal{R}} \ell_{\mathcal{R}}^2 / d \geq \frac{1}{20 \log_2 n} \sum_j \|x_j\|^2.$$

□

**PROOF OF LEMMA 6.** For the numerator of  $R(z_B)$ , we get

$$\begin{aligned} \sum_{\{i,j\} \in E(B)} (z_i - z_j)^2 &= \sum_{\{i,j\} \in E(B)} (u_i - u_j)^2 \\ &\leq \sum_{\{i,j\} \in E(A)} (u_i - u_j)^2. \end{aligned}$$

For the denominator, recall that  $S$  and  $B$  form a partition of  $A$  and further  $u \perp \vec{d}$ . First, we get  $0 = \sum_{i \in A} d_i u_i = \text{vol}(S) \bar{u}_S + \text{vol}(B) \bar{u}_B$ , which can be rearranged to give

$$\bar{u}_B^2 = \frac{\text{vol}(S)^2}{\text{vol}(B)^2} \cdot \bar{u}_S^2. \quad (4)$$

Next, we observe that

$$\sum_{i \in S} d_i u_i^2 = \sum_{i \in S} d_i (u_i - \bar{u}_S)^2 + \text{vol}(S) \bar{u}_S^2 \geq \text{vol}(S) \bar{u}_S^2. \quad (5)$$

Putting everything together, we can bound the denominator of  $R(z_B)$  with

$$\begin{aligned} \sum_{i \in B} d_i z_i^2 &= \sum_{i \in B} d_i u_i^2 - \text{vol}(B) \cdot \bar{u}_B^2 \\ &\geq \sum_{i \in B} d_i u_i^2 - \frac{\text{vol}(S)}{\text{vol}(B)} \sum_{i \in S} d_i u_i^2 && \text{by (4), (5)} \\ &\geq (1 - \lambda) \sum_{i \in A} d_i u_i^2 - \lambda \frac{\text{vol}(S)}{\text{vol}(B)} \sum_{i \in A} d_i u_i^2 && \text{assumption} \\ &= \frac{\text{vol}(B) - \lambda \text{vol}(A)}{\text{vol}(B)} \cdot \sum_{i \in A} d_i u_i^2 && A = B \cup S \end{aligned}$$

and therefore  $R(z_B) \leq \frac{\text{vol}(B)}{\text{vol}(B) - \lambda \text{vol}(A)} \cdot R(u)$ . □

**PROOF OF CLAIM 1.**

$$\begin{aligned} d \left\| \frac{\sum_i a_i}{d} - \mu \right\|^2 - \sum_i \|a_i - \mu\|^2 &= d \left( \left\| \frac{1}{d} \sum_i a_i \right\|^2 - \frac{2}{d} \sum_i a_i^T \mu + \|\mu\|^2 \right) - \sum_i \|a_i - \mu\|^2 \\ &= \frac{1}{d} \left\| \sum_i a_i \right\|^2 - 2 \sum_i a_i^T \mu + d \|\mu\|^2 - \sum_i (\|a_i\|^2 - 2a_i^T \mu + \|\mu\|^2) \\ &= \frac{1}{d} \underbrace{\left\| \sum_i a_i \right\|^2}_{\leq \sum_i \|a_i\|^2} - \sum_i \|a_i\|^2 \leq 0. \end{aligned}$$

For  $d = 2$  we get

$$\begin{aligned} \frac{1}{2} \|a_1 + a_2\|^2 - \|a_1\|^2 - \|a_2\|^2 &= \frac{1}{2} \|a_1\|^2 + a_1^T a_2 + \frac{1}{2} \|a_2\|^2 - \|a_1\|^2 - \|a_2\|^2 \\ &= -\frac{1}{2} \|a_1 - a_2\|^2. \end{aligned}$$

□

## B ADDITIONAL RELATED WORK

*Multilevel Graph Partitioning.* This framework has been employed by many successful graph partitioning tools, e.g., METIS [18], Graclus [10], or KaHiP [33]. The general idea of the paradigm is to compute a small summary (“coarsening”) of the graph, on which it is easier to solve the problem we are interested in, and then mapping this coarse solution onto the original graph. In more technical terms, a solver following this paradigm has three phases: (i) *Coarsening*: Compute a series of successively smaller graphs  $G_0, G_1, \dots, G_\ell$  such that  $|V_0| > |V_1| > \dots > |V_\ell|$  with the aim to create a coarse summary of the original graph. (ii) *Solving*: Obtain a solution to the initial problem on the graph  $G_\ell$ . Although  $G_\ell$  is small, heuristics are usually employed here. (iii) *Refinement*: Successively map the solution from graph  $G_i$  to graph  $G_{i-1}$ , while applying heuristics that increase the solution quality.

## C INSTANCE LIST

See Table C.3 for a full list of instances and information on their degree distribution.

## D ADDITIONAL FIGURES

In this section we include a number of figures omitted from the main section of the paper.

Received 8 February 2024; accepted 17 May 2024

**Table C.3: Statistics on the number of vertices, edges and degree distribution of the graph instances in our benchmark dataset.  $\Delta$  denotes the maximum,  $d_{\text{mean}}$  the average vertex degree. The last four columns are percentiles of the degree distribution.**


#	Name	Type	$ V $	$ E $	$\Delta$	$d_{\text{mean}}$	25th	50th	75th	90th
BP1	imdb	Bipartite	1 403 278	4 303 383	1652	6.13	1	2	6	15
CF1	ramage02	Computational Fluids	16 830	1 424 761	269	169.31	131	170	170	269
CL1	uk	Clustering	4824	6837	3	2.83	3	3	3	3
CL2	smallworld	Clustering	100 000	499 998	17	10.00	9	10	11	12
CN1	citationCiteseer	Citation Network	268 495	1 156 647	1318	8.62	2	5	10	18
CN2	ca-hollywood-2009	Citation Network	1 069 126	56 306 653	11 467	105.33	13	31	75	212
CN3	coAuthorsDBLP	Citation Network	299 067	977 676	336	6.54	2	4	7	14
CN4	ca-MathSciNet	Citation Network	332 689	820 644	496	4.93	1	3	5	11
CN5	ca-coauthors-dblp	Citation Network	540 486	15 245 729	3299	56.41	13	34	74	135
CN6	ca-dblp-2012	Citation Network	317 080	1 049 866	343	6.62	2	4	7	14
CN7	ca-citeseer	Citation Network	227 320	814 134	1372	7.16	2	4	8	15
CN8	coPapersCiteseer	Citation Network	434 102	16 036 720	1188	73.88	15	39	92	177
CN9	ca-dblp-2010	Citation Network	226 413	716 460	238	6.33	2	4	7	13
CS1	add32	Circuit Simulation	4960	9462	31	3.82	2	3	4	9
CS2	rajat10	Circuit Simulation	30 202	50 101	101	3.32	2	4	4	4
CS3	memplus	Circuit Simulation	17 758	54 196	573	6.10	2	3	4	11
DM1	pcrystk02	Duplicate Materials	13 965	477 309	80	68.36	53	80	80	80
EM1	email-enron-large	Email Network	33 696	180 811	1383	10.73	1	3	7	19
EM2	email-EU	Email Network	32 430	54 397	623	3.35	1	1	1	3
FE1	fe_tooth	Finite Elements	78 136	452 591	39	11.58	6	12	15	18
FE2	fe_rotor	Finite Elements	99 617	662 431	125	13.30	11	13	14	17
IF1	inf-openflights	Infrastructure Network	2939	15 677	242	10.67	2	3	8	28
IF2	inf-power	Infrastructure Network	4941	6594	19	2.67	2	2	3	5
NS1	wing_nodal	Numerical Simulation	10 937	75 488	28	13.80	12	14	16	17
NS2	auto	Numerical Simulation	448 695	3 314 611	37	14.77	13	15	16	18
OP1	gupta2	Optimization	62 064	2 093 111	8412	67.45	25	38	47	49
OP2	finance256	Optimization	37 376	130 560	54	6.99	4	6	8	12
RD1	appu	Random Graph	14 000	919 552	293	131.36	109	131	154	176
RN1	inf-roadNet-CA	Road Network	1 957 027	2 760 388	12	2.82	2	3	3	4
RN2	inf-roadNet-PA	Road Network	1 087 562	1 541 514	9	2.83	2	3	4	4
RN3	inf-italy-osm	Road Network	6 686 493	7 013 978	9	2.10	2	2	2	3
RN4	luxembourg_osm	Road Network	114 599	119 666	6	2.09	2	2	2	3
SN1	soc-youtube-snap	Social Network	1 134 890	2 987 624	28 754	5.27	1	1	3	8
SN2	soc-flickr	Social Network	513 969	3 190 452	4369	12.41	1	1	5	17
SN3	soc-lastfm	Social Network	1 191 805	4 519 330	5150	7.58	1	2	4	11
SN4	soc-twitter-follows	Social Network	404 719	713 319	626	3.53	1	1	1	3
SN5	soc-pokec	Social Network	1 632 803	22 301 964	14 854	27.32	4	13	35	70
SN6	soc-livejournal	Social Network	4 033 137	27 933 062	2651	13.85	2	5	15	35
SN7	soc-FourSquare	Social Network	639 014	3 214 986	106 218	10.06	1	1	4	19
TM1	vsp_vibrobox_scagr7-2c_rlfddd	Triangle Mixture	77 328	435 586	669	11.27	3	5	8	26
TM2	vsp_bump2_e18_aa01_model1_crew1	Triangle Mixture	56 438	300 801	604	10.66	5	7	11	18
TM3	vsp_barth5_1Ksep_50in_5Kout	Triangle Mixture	32 212	101 805	22	6.32	6	6	7	7
TM4	vsp_model1_crew1_cr42_south31	Triangle Mixture	45 101	189 976	17 663	8.42	3	5	7	9
TM5	vsp_p0291_seymourl_iiasa	Triangle Mixture	10 498	53 868	229	10.26	3	7	11	16
TM6	vsp_bcsstk30_500sep_10in_1Kout	Triangle Mixture	58 348	2 016 578	219	69.12	45	64	92	106
TM7	vsp_befref_fxm_2_4_air02	Triangle Mixture	14 109	98 224	1531	13.92	6	9	11	15
US1	mi2010	US Census Redistricting	329 885	789 045	58	4.78	3	4	5	8
WB1	web-it-2004	Web Graph	509 338	7 178 413	469	28.19	14	16	19	39
WB2	web-google	Web Graph	1299	2773	59	4.27	1	2	5	12
WB3	web-wikipedia2009	Web Graph	1 864 433	4 507 315	2624	4.84	1	2	5	10

**Table D.4: Geometric mean of normalized cut value across graph types over  $k$  for the multilevel graph partitioning algorithms from subsection 6.3.**

Type	$k$ Algorithm Name	2	4	8	16	32	64	128
Bipartite	Graclus	0.05	0.17	0.45	1.09	3.55	9.18	24.41
	KaHiP	0.04	0.27	0.67	1.92	4.84	13.03	33.65
	METIS	0.03	0.17	0.59	1.66	4.66	13.07	34.55
	XCut	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>
Circuit Simulation	Graclus	<b>0.01</b>	0.06	<b>0.18</b>	<b>0.54</b>	1.57	4.76	14.04
	KaHiP	<b>0.01</b>	0.07	0.21	0.62	1.78	5.65	16.99
	METIS	0.02	0.07	0.22	0.59	1.72	5.49	16.90
	XCut	<b>0.01</b>	<b>0.05</b>	<b>0.18</b>	<b>0.54</b>	<b>1.41</b>	<b>4.21</b>	<b>11.90</b>
Citation Network	Graclus	0.08	0.26	0.67	1.57	3.75	8.22	17.76
	KaHiP	0.07	0.27	0.73	1.81	4.27	9.57	21.47
	METIS	0.08	0.29	0.81	2.01	4.80	11.04	24.32
	XCut	<b>0.00</b>	<b>0.01</b>	<b>0.03</b>	<b>0.08</b>	<b>0.22</b>	<b>0.61</b>	<b>1.69</b>
Clustering	Graclus	0.04	0.14	0.43	1.20	3.14	8.19	20.69
	KaHiP	0.04	0.13	0.43	1.15	3.11	7.91	19.95
	METIS	0.04	0.15	0.46	1.25	3.37	8.58	21.38
	XCut	<b>0.02</b>	<b>0.09</b>	<b>0.29</b>	<b>0.93</b>	<b>2.52</b>	<b>7.07</b>	<b>19.49</b>
Computational Fluids	Graclus	<b>0.12</b>	0.46	<b>1.32</b>	<b>3.86</b>	<b>10.49</b>	<b>27.76</b>	<b>70.28</b>
	KaHiP	<b>0.12</b>	0.46	1.47	4.18	10.98	29.64	74.51
	METIS	<b>0.12</b>	<b>0.44</b>	1.35	4.10	10.95	28.65	72.55
	XCut	0.15	0.55	1.63	4.75	14.30	44.39	108.44
Duplicate Materials	Graclus	<b>0.04</b>	<b>0.24</b>	<b>0.80</b>	<b>2.58</b>	<b>7.71</b>	<b>20.42</b>	<b>53.60</b>
	KaHiP	0.05	0.32	0.85	2.97	8.20	21.45	56.65
	METIS	<b>0.04</b>	<b>0.24</b>	<b>0.80</b>	2.60	7.75	20.67	55.85
	XCut	0.05	0.26	0.90	3.27	9.54	26.81	81.90
Email Network	Graclus	0.18	0.60	1.65	4.07	9.11	20.13	42.43
	KaHiP	0.19	0.66	1.70	4.36	10.11	23.19	51.73
	METIS	0.19	0.64	1.80	4.40	10.21	23.06	51.24
	XCut	<b>0.01</b>	<b>0.04</b>	<b>0.17</b>	<b>0.51</b>	<b>1.46</b>	<b>4.07</b>	<b>11.76</b>
Finite Elements	Graclus	<b>0.01</b>	<b>0.06</b>	0.21	<b>0.63</b>	<b>1.82</b>	<b>5.16</b>	<b>13.97</b>
	KaHiP	<b>0.01</b>	<b>0.06</b>	0.21	0.66	2.00	5.63	15.40
	METIS	<b>0.01</b>	<b>0.06</b>	<b>0.20</b>	0.65	1.88	5.23	14.42
	XCut	<b>0.01</b>	<b>0.06</b>	0.21	0.67	1.99	5.67	15.64
Infrastructure Network	Graclus	0.02	0.09	0.39	1.09	3.19	9.97	27.44
	KaHiP	0.02	0.09	0.44	1.51	4.41	12.24	32.27
	METIS	0.02	0.11	0.49	1.36	4.16	11.60	31.29
	XCut	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>	<b>0.23</b>	<b>1.38</b>	<b>5.30</b>	<b>17.73</b>
Numerical Simulation	Graclus	0.02	<b>0.08</b>	0.28	0.94	2.70	<b>7.24</b>	<b>19.40</b>
	KaHiP	0.02	0.09	0.31	0.93	2.76	7.77	20.76
	METIS	0.02	0.09	<b>0.26</b>	<b>0.89</b>	<b>2.63</b>	7.25	19.43
	XCut	<b>0.01</b>	0.09	0.33	1.10	3.18	8.88	23.89
Optimization	Graclus	0.04	0.14	0.40	1.85	5.65	16.68	42.32
	KaHiP	<b>0.00</b>	0.07	0.32	1.86	5.84	18.65	48.48
	METIS	0.04	0.09	0.40	1.83	5.55	16.60	47.46
	XCut	0.01	<b>0.03</b>	<b>0.22</b>	<b>1.38</b>	<b>4.63</b>	<b>13.60</b>	<b>33.06</b>
Random Graph	Graclus	0.91	<b>2.59</b>	<b>6.19</b>	<b>13.52</b>	28.37	<b>57.79</b>	<b>116.97</b>
	KaHiP	0.88	2.67	6.37	13.83	28.62	58.17	118.02
	METIS	<b>0.85</b>	2.61	<b>6.19</b>	13.53	<b>28.36</b>	58.18	117.98
	XCut	0.99	2.99	6.98	14.99	30.99	62.99	126.99
Road Network	Graclus	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>	<b>0.01</b>	<b>0.03</b>	0.10	<b>0.30</b>
	KaHiP	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>	<b>0.01</b>	<b>0.03</b>	<b>0.09</b>	<b>0.30</b>
	METIS	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>	<b>0.01</b>	<b>0.03</b>	0.10	<b>0.30</b>
	XCut	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>	<b>0.01</b>	<b>0.03</b>	0.11	0.34
Social Network	Graclus	0.16	0.66	1.64	4.01	9.03	19.96	44.26
	KaHiP	0.20	0.73	1.92	4.66	10.59	24.11	51.14
	METIS	0.19	0.68	1.87	4.67	10.91	24.54	53.15
	XCut	<b>0.01</b>	<b>0.03</b>	<b>0.11</b>	<b>0.31</b>	<b>0.83</b>	<b>2.19</b>	<b>6.02</b>
Triangle Mixture	Graclus	0.08	0.43	1.43	3.64	8.94	20.71	48.90
	KaHiP	0.08	0.50	1.54	3.94	9.62	22.73	52.85
	METIS	0.09	0.56	1.60	4.14	10.13	23.10	51.51
	XCut	<b>0.03</b>	<b>0.25</b>	<b>1.08</b>	<b>3.08</b>	<b>7.70</b>	<b>18.44</b>	<b>44.26</b>
US Census Redistricting	Graclus	<b>0.00</b>	0.01	<b>0.02</b>	0.07	0.21	0.64	1.92
	KaHiP	<b>0.00</b>	<b>0.00</b>	<b>0.02</b>	<b>0.06</b>	0.20	<b>0.59</b>	<b>1.80</b>
	METIS	<b>0.00</b>	0.01	<b>0.02</b>	0.07	0.21	0.68	1.99
	XCut	<b>0.00</b>	<b>0.00</b>	<b>0.02</b>	0.07	<b>0.19</b>	<b>0.59</b>	1.92
Web Graph	Graclus	<b>0.00</b>	0.02	0.04	0.10	0.26	0.77	2.02
	KaHiP	<b>0.00</b>	0.01	0.04	0.11	0.47	2.09	7.36
	METIS	0.01	0.02	0.03	0.08	0.42	2.46	9.01
	XCut	<b>0.00</b>	<b>0.00</b>	<b>0.00</b>	<b>0.01</b>	<b>0.02</b>	<b>0.10</b>	<b>0.36</b>




Expander Hierarchies for Normalized Cuts on Graphs



**Figure D.7: Percentage deviation of the returned normalized cut value relative to Graclus for  $k = 2$ . This means that a value of  $-75\%$  means that the normalized cut value is  $75\%$  lower (i.e., better). The thin black bars indicate the standard error across our runs. The top graph shows the disconnected IMDB graph (BP1), citation network instances (CN), email networks (EM), infrastructure graphs (IF), social networks (SN) and web graphs (WB), while the bottom shows the remaining instances. See [Table C.3](#) for details.**



**Figure D.8:** Percentage deviation of the returned normalized cut value relative to Graclus for  $k = 128$ . This means that a value of  $-75\%$  means that the normalized cut value is  $75\%$  lower (i.e., better). The thin black bars indicate the standard error across our runs. The top graph shows the disconnected IMDB graph (BP1), citation network instances (CN), email networks (EM), infrastructure graphs (IF), social networks (SN) and web graphs (WB), while the bottom shows the remaining instances. See [Table C.3](#) for details.



**Figure D.9:** Plot showcasing the time taken to compute the expander hierarchy in orange and the total time to compute a normalized cut in blue for five graph instances of different sizes.

Expander Hierarchies for Normalized Cuts on Graphs




Figure D.10: Plots of running time versus  $k$ . The first row shows the usual behavior of the solvers in our comparison, with our solver starting at a much higher level than the competition. This is due to the fact that we need to first compute the expander hierarchy, which gives a hard lower bound on the running time of our algorithm for any choice of  $k$ . The running time then grows linearly with  $k$ , due to Greedy taking time  $O(nk)$ . The bottom row contains some instances on which the running time results are unusual. Often Graclus running time grows very quickly, but we are unsure what causes this on these instances in particular.

Table D.5: Cut values and running time of our Algorithm vs the values reported by Zhao et al [43]. All values are for  $k = 30$ . The Cut-columns contain the normalized cut value (lower is better). The minimum value in each row is marked bold. The bottom row contains the geometric mean of all values.

GID	XCut <sub>mean</sub>	XCut <sub>min</sub>	Zhao	Graclus
CL1	0.87	<b>0.77</b>	1.05	1.15
CL2	6.00	<b>5.87</b>	7.05	7.61
CN1	0.31	<b>0.29</b>	0.52	4.06
CN3	<b>0.27</b>	<b>0.27</b>	0.49	3.66
CN5	<b>0.13</b>	<b>0.13</b>	0.14	3.24
CN7	0.17	<b>0.16</b>	0.41	2.21
CN8	<b>0.04</b>	<b>0.04</b>	0.06	1.88
FE1	2.06	1.96	<b>1.68</b>	1.74
FE2	1.58	1.55	1.50	<b>1.45</b>
NS1	5.73	5.35	<b>4.71</b>	<b>4.71</b>
NS2	1.45	1.39	<b>1.08</b>	1.17
RD1	28.98	28.98	<b>23.80</b>	26.48
RN4	<b>0.06</b>	<b>0.06</b>	0.07	0.07
TM1	9.53	8.91	<b>6.85</b>	8.48
TM2	12.34	<b>12.06</b>	13.55	16.63
TM3	2.78	<b>2.65</b>	2.72	2.78
TM4	10.09	<b>9.58</b>	10.48	14.45
TM5	7.77	<b>7.48</b>	7.88	13.25
TM6	2.11	<b>1.94</b>	2.09	2.1
TM7	17.51	16.90	<b>12.83</b>	15.6
US1	0.17	<b>0.16</b>	0.41	0.19
All	1.46	<b>1.39</b>	1.64	3.06