Submodular Hypergraph Partitioning: Metric Relaxations and Fast Algorithms via an Improved Cut-Matching Game

Antares Chen
University of Chicago, IL, USA

Lorenzo Orecchia
University of Chicago, IL, USA

Erasmo Tani
University of Chicago, IL, USA
Sapienza University, Rome, Italy

Abstract

Despite there being significant work on developing spectral- [12, 36, 35], and metric-embedding-based [50] approximation algorithms for hypergraph conductance, little is known regarding the approximability of other hypergraph partitioning objectives.

This work proposes algorithms for a general model of hypergraph partitioning that unifies both undirected and directed versions of many well-studied partitioning objectives. The first contribution of this paper introduces *polymatroidal cut functions*, a large class of cut functions amenable to approximation algorithms via metric embeddings and routing multicommodity flows. We demonstrate a simple $O(\sqrt{\log n})$-approximation, where n is the number of vertices in the hypergraph, for these problems by rounding relaxations to metrics of negative-type.

The second contribution of this paper generalizes the cut-matching game framework of Khandekar et al. [31] to tackle polymatroidal cut functions. This yields an almost-linear time $O(\log n)$-approximation algorithm for standard versions of undirected and directed hypergraph partitioning [35]. A technical contribution of our construction is a novel cut-matching game, which greatly relaxes the set of allowed actions by the cut player and allows for the use of *approximate s-t maximum flows* by the matching player. We believe this to be of independent interest.

2012 ACM Subject Classification Theory of computation → Facility location and clustering

Keywords and phrases Hypergraph Partitioning, Cut Improvement, Cut-Matching Game

Digital Object Identifier 10.4230/LIPIcs.ICALP.2025.49

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: https://arxiv.org/abs/2301.08920 [2]

Funding *Antares Chen*: NSF DGE 2140001.
Lorenzo Orecchia: NSF CAREER 1943510.

Acknowledgements The authors would like to thank anonymous reviewers for suggesting related work on polymatroidal networks, Thatchaphol Saranurak for suggesting related work on the use of cut-matching games beyond partitioning graphs, and Jeffrey Negrea for assisting on topics related to online learning. AC would like to dedicate this submission to Luca Trevisan. A portion of this work was done while AC visited Bocconi; among many other fond memories, AC recalls long discussions with Luca on this project and other topics of research. AC would also like to thank Max Ovsiankin and Marco Pirazzini for their steadfast support in finding the SPARSEST of cuts.

52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025).
Editors: Keren Censor-Hillel, Fabrizio Grandoni, Joël Ouaknine, and Gabriele Puppis
Article No. 49; pp. 49:1–49:16
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Increasing complexity in real-world data has necessitated generalizing graph models to capture relations between multiple objects [1, 11]. Hypergraphs provide such a model, with tasks involving hypergraphs being featured prominently in recent practical [9, 67] and theoretical developments [12, 13, 49, 50].

Hypergraph partitioning is such a task. In machine learning, hypergraph partitioning enables detecting clusters in complex higher-order relations arising in social [67, 71] and biological [24, 33] networks. In combinatorial optimization, partitioning algorithms are a fundamental primitive, used in the development of divide-and-conquer methods for other hypergraph problems. Unlike in graphs, hypergraphs admit multiple manners of quantifying the cost of how a hyperedge is cut. This yields a wide spectrum partitioning objectives, and allows end-users to better design partitioning problems that suite their application domain. While there has been significant work on approximation algorithms for specific hypergraph partitioning objectives, such as undirected hypergraph conductance [12, 36], vertex expansion [51, 50, 35] and conductance in directed graphs [36], little is known about the approximability of hypergraph partitioning objectives *at broad*, with even less known regarding *fast algorithms* for these objectives.

The first part of this work defines *ratio-cut problems with polymatroidal cut functions*, a class of partitioning objects inspired by classical work on polymatroidal networks [27, 42]. This is a subset of submodular cut functions considered in [45, 46], but with structure well-suited for approximation methods via metric embeddings. It captures undirected and directed versions of many partitioning objectives found in current literature, and for all objectives in this class, we construct polynomial-time $O(\sqrt{\log n})$-approximation algorithms by rounding ℓ_2^2-metric relaxations. The second part of the paper constructs fast algorithms for polymatroidal cut function objectives. Our algorithms do not require solving super-quadratic-size semidefinite programs, and achieve $O(\log n)$-approximations. Instead, we *generalize the cut-matching framework* of [31], and produce the first almost-linear time polylogarithmic approximation for undirected and directed hypergraph conductance.

To describe partitioning task in more detail, recall that a *hypergraph* is a tuple $G = (V, E)$, where $E \subseteq 2^V$ is a set of *hyperedges*. The *rank* of a hyperedge $h \in E$ is its cardinality $|h|$. If every $h \in E$ satisfies $|h| = k$, the hypergraph G is called *k-uniform*. Hence, an undirected graph is a 2-uniform hypergraph. With each hyperedge $h \in E$, we associate a suitable *cut function* $\delta_h : 2^h \to [0, 1]$. This expresses the cost incurred when h is cut by $(A, \bar{A})$ into $A \cap h$ and $\bar{A} \cap h$. This work focuses on *submodular cut functions*, defined by [45, 46]:

▶ **Definition 1** (Submodular Cut Function). *A function $\delta_h : 2^h \to [0, 1]$ is a submodular cut function if it is submodular and satisfies $\delta_h(\varnothing) = \delta_h(h) = 0$.*

The space of cut functions for an edge $\{i, j\}$ is spanned by the undirected edge cut function $\delta_{\{i,j\}}(S) = |\mathbf{1}_i^S - \mathbf{1}_j^S|$ and the directed edge cut function $\delta_{(i,j)}(S) = (\mathbf{1}_i^S - \mathbf{1}_j^S)_+$, giving a complete classification of cut functions for rank-2 hyperedges. Higher-rank hyperedges support broader choices of cut functions. The most common example is the *standard hypergraph cut function* [12, 49]: $\delta_h^{\mathsf{cut}}(S) \stackrel{\text{def}}{=} \min\{1, |S|, |h \setminus S|\}$, which takes value 1 if h is cut by S and 0 otherwise. If the hypergraph $G = (V, E)$ is associated with positive hyperedge weights $\mathbf{w} \in \mathbb{Z}_{>0}^E$, we can succinctly write the cut function for the entire graph as $\delta_G(S) \stackrel{\text{def}}{=} \sum_{h \in E} w_h \cdot \delta_h(S)$.

With this, we can define our hypergraph partitioning objective of interest: *submodular hypergraph ratio-cut.*

▶ **Definition 2** (Submodular Hypergraph Ratio-Cut)**.** *Given a hypergraph $G = (V, E)$ with positive hyperedge weights $\mathbf{w} \in \mathbb{Z}^E_{>0}$, non-negative vertex weights $\boldsymbol{\mu} \in \mathbb{Z}^V_{\geq 0}$, and a collection of submodular cut functions $\{\delta_h\}_{h\in E}$, the* ratio-cut objective *on G is defined for all $S \subseteq V$ as:*

$$\Psi_G(S) \stackrel{\text{def}}{=} \frac{\delta_G(S)}{\min\{\mu(S), \mu(\bar{S})\}} = \frac{\sum_{h\in E} w_h \cdot \delta_h(S \cap h)}{\min\{\mu(S), \mu(\bar{S})\}}.$$

The ratio-cut *of G is $\Psi^*_G \stackrel{\text{def}}{=} \min_{S\subset V} \Psi_G(S)$.*

When specialized to graphs[1] ratio-cut objectives include both graph expansion ($\boldsymbol{\mu} = \mathbf{1}$), graph conductance ($\forall i \in V, \mu_i = \sum_{h\in E: i\in h} w_h$). It also captures their directed counterparts, by replacing the undirected cut function $\delta_{\{i,j\}}$ with its directed analogue $\delta_{(i,j)}$.

Previous Work on Polymatroidal Networks

Polymatroidal networks are a model of network flows first introduced independently by Hassin [27], and Lawler and Martel [42], with the latter drawing inspiration from scheduling problems [54, 55, 53]. This class of networks generalized the standard notion of *s*-*t* flows over graphs by replacing edge capacity constraints with submodular function constraints over subsets of edges [40] (in [7]). Among many others which studied submodular constrained flow networks [20, 23, 25, 41, 43], this body of work explored how to prove the existence of integral solutions to different flow and partitioning linear programs, and subsequently generalize combinatorial strong duality results such as the maxflow-mincut of Ford & Fulkerson, and Kőnig's matching theorem [63, 64]. Along this theme, Chekuri et al. [14] defined a notion of multicommodity flow for polymatroidal networks, and used metric embedding techniques to prove polylogarithmic upper bounds on multicommodity flow-cut gaps in this setting.

Previous Work on Submodular Hypergraph Partitioning

Hypergraph partitioning has been extensively studied from both a practical and theoretical perspective. Early work was driven by interest in optimal placement of circuit components [38, 52, 57], designing scheduling layouts [62, 22], and information management [39, 56]. Later interest was driven by applications in parallel numerical algorithms [11], and scientific computing [21].

In the realm of tractable algorithms with provable approximation guarantees, a number of metric embedding and spectral algorithms are known for approximating specific hypergraph partitioning objectives. Louis & Makarychev [50] gave the first randomized polynomial-time $O\big(\sqrt{\log|V|}\big)$-approximation algorithm for hypergraph expansion, matching the best known bound on graph expansion [6]. Louis [49] and Chan et al. [12] show Cheeger-like guarantees for hypergraph conductance. Their algorithms require solving a semidefinite program (SDP) to approximate a notion of hypergraph spectral gap, returning a partition with conductance at most $O\big(\sqrt{\Psi^*_G \log \max_{h\in E} |h|}\big)$. Their results are known to be tight under the small-set expansion conjecture. Recently, Lau et al. [36] introduced a different notion of hypergraph spectral gap, called reweighted eigenvalues, and use it to obtain Cheeger-like guarantees of the form $O\big(\sqrt{\Psi^*_G \log(1/\Psi^*_G)}\big)$ for undirected and directed hypergraph conductance.

[1] In the rest of the paper, for any undirected (resp. directed) graphs G, unless otherwise specified, we assume that δ_G and Ψ_G are formed using the undirected (resp. directed) cut functions.

Less is known for larger classes of ratio-cut objectives. Yoshida [72] showed a Cheeger-like bound on ratio-cut with submodular cut functions of $O\big(\sqrt{\Psi_G^* \cdot \log|V|}\big)$. Li & Milenkovic [46] improved this to $O\big(\sqrt{\Psi_G^* \cdot \max_{h\in E}|h|}\big)$ and conjectured that improving the dependence on rank is NP-hard. In broader generality, Svitkina & Fleischer [66] showed that an $\omega\big(\sqrt{|V|/\log|V|}\big)$-approximation for a submodular version of sparsest cut requires an exponential number of function-value queries. Thus, one must assume additional structure to construct non-trivial approximations to submodular hypergraph partitioning objectives.

Previous Work on Cut-Matching Games

Many of the spectral and metric embedding methods mentioned above solve SDP relaxations over $|V| \times |V|$ symmetric matrices, with at least $|E|$ constraints. No faster algorithms are currently known beyond generic SDP solvers [30], which run in $\Omega\big(|V|^2|E|\big)$ time. Even for the simplest form of hypergraph partitioning, the resulting SDP relaxation is a mixed packing-and-covering program for which no almost-linear time algorithm is currently known [29].

To circumvent this issue, one can apply primal-dual methods which solve the dual to the metric embedding SDPs. However, the dual is typically a multicommodity flow problem whose demand graph is dense, and solving this black-box may require quadratic time [5]. The cut-matching game [32] was designed to approximately reduce this multicommodity flow to a polylogarithmic number of *exact single-commodity flows*. Each single-commodity flow problem would route a perfect matching in the graph, until the union of matchings approximates the desired demand graph.

Using the cut-matching game, Khandekar et al. [32] obtained a $O(\log^2|V|)$-approximation to undirected graph expansion by exactly solving $O(\log^3|V|)$ maximum flows. Orecchia et al. [61] reduced both of these measures by $O(\log n)$. Subsequent work has vastly generalized the applicability of the cut-matching game to directed graph expansion [48], to vertex and hypergraph expansion [47]. Nanongkai & Saranurak [58] demonstrate how to replace exact single-commodity flows with approximate ones at the cost of an additional $O(\log n)$-factor in the approximation ratio. The cut-matching game has since become ubiquitous in designing fast deterministic algorithms for various static, and dynamic graph problems [3, 10, 17, 18, 19, 26].

Our Work

We provide the following technical contributions. Further detail is given in the parenthesized subsections.

- (1.1.1) *A metric approach to submodular hypergraph ratio-cut* – We identify a subset of submodular cut functions, called *polymatroidal cut functions*, that are amenable to metric and flow techniques. For this class, we develop a notion of hypergraph flow, and show how to certify lower bounds to the ratio-cut objective by flow-embedding graphs into hypergraphs. We give polynomial-time $O(\sqrt{\log|V|})$-approximation algorithms for minimum ratio-cut for submodular hypergraphs with polymatroidal cut functions by providing an ℓ_2^2-metric relaxation, generalizing the result of Arora, Rao & Vazirani [5] to submodular hypergraph partitioning.
- (1.1.2) *Local Formulations of the Submodular Hypergraph Ratio-Cut Problem* – Generalizing the cut-matching game requires an algorithm that can locally probe the hypergraph cut structure by finding a sparse hypergraph cut near a given input partition. The cut-improvement algorithm of Andersen & Lang [4] does this for graphs. We generalize this to submodular hypergraphs. Our cut improvement problem can be solved using decomposable submodular minimization. We provide a simple interpretation of this problem by characterizing it as a *surrogate estimate* for the submodular hypergraph ratio-cut objective that is both *local and convex.*

- (1.1.3) *An Improved Cut-Matching Game* – We extend the cut-matching game framework to approximate minimum ratio-cuts on submodular hypergraphs with polymatroidal cut functions. Our approach relies on an improved formulation of the cut-matching games that allows for a larger set of cut and matching strategies. As a result, we produce an $O(\log|V|)$-approximation algorithm for minimum ratio-cut that uses only *a single* $O(1)$*-approximate* decomposable submodular minimization solve per round. For the special case of graph partitioning, our algorithm improves on existing applications of the cut-matching game by enabling the use of approximate maxflow solutions *at no cost* to the approximation ratio. This is to be compared with previous work that require either a single $o_n(1)$-approximate, or a polylogarithmic number of $O(1)$-approximate maxflow computations per round [10, 47, 58], which results in a polylogarithmic loss in approximation ratio.

Concurrent and Subsequent Work

Subsequent to an initial version of this paper [2], Veldt [68] gave a $O(\log|V|)$-approximation for ratio-cut with cardinality-based symmetric submodular cut functions, a special case of submodular hypergraph partitioning where $\delta_h(S) = \delta_h(h \setminus S)$ and $\delta_h(S) = \delta_h(T)$ if $|S| = |T|$. This assumption on δ_h allows a certain combinatorial gadget, *hypergraph cut preservers*, to certify lower bounds to the ratio-cut objective [70, 69]. Their work gives an algorithm that generalizes the cut-matching framework to allow for a matching player to construct such a gadget using exact single commodity flows. [68] additionally provides compelling empirical evaluations.

Independently, Lau et al. [37] gave the first almost-linear time $O(\sqrt{\log n})$-approximation for undirected and directed hypergraph expansion. Their algorithm does not use the cut-matching game, and instead solves the dual to the relaxation of the minimum reweighted eigenvalue problem in [36] equipped with additional ℓ_2^2-metric constraints. This produces a multicommodity flow problem which they show to be approximable by a sequence of single-commodity flows via a generalization of Sherman's algorithmic chaining [65].

Finally in very recent work, Chekuri & Louis [15] extend our algorithm and analysis to produce a $O(\sqrt{\log|V|})$-approximation to the larger class of ratio-cut problems in polymatroidal networks (see section 1.2 of [15] for a comparison). However, their work explicitly does not address the design of fast algorithms.

1.1 Technical Summary of Contributions

In this section, we provide a technical overview of our work, including the key definitions and the statement of the main theorems and lemmata. Full proofs, together with more detailed discussions and examples, are available in the full version of the paper [16].

1.1.1 A Metric Approach to Submodular Hypergraph Ratio-Cut

The best polynomial-time approximation algorithms for graph ratio-cut problems solve for a metric relaxation, then round using a metric embedding result. Relaxing to general metrics results in an $O(\log n)$-approximation [44], while relaxing to ℓ_2^2-metrics achieve $O(\sqrt{\log n})$ [6]. Our goal is to understand the applicability of this approach to submodular hypergraph ratio-cut problems. First, one should not hope to obtain polylogarithmic approximations for general submodular cut functions without further assumptions, as Svitkina & Fleischer (Section 3 in [66]) exhibit a submodular cut function for which any $o(\sqrt{n})$-approximation requires exponentially many queries to a value oracle.

What kind of restrictions allow for metric relaxation methods to work? The cut function must exhibit some kind of monotonicity, so that its value is preserved under low-distortion metric embeddings. However, a cut function cannot be monotone without being equal to 0, so the monotone structure must appear in some other way. To solve this issue, we introduce the class of *polymatroidal cut functions.*

Polymatroidal Cut Functions and Metric Embeddings

Our first contribution is to define *polymatroidal cut functions*, a subclass of submodular cut functions given by the infimal symmetrization of non-decreasing submodular functions.

▶ **Definition 3** (Polymatroidal Cut Function). *A cut function $\delta_h : 2^h \to \mathbb{R}_{\geq 0}$ is polymatroidal if, for all $S \subseteq h$, it can be expressed as*

$$\delta_h(S) = \min\left\{F_h^-(S),\ F_h^+(h \setminus S)\right\},$$

where the associated functions $F_h^-, F_h^+ : 2^h \to \mathbb{R}$ are non-decreasing submodular functions such that $F_h^-(\varnothing) = F_h^+(\varnothing) = 0$. When the associated functions F_h^- and F_h^+ are identical, we refer to the cut function δ_h as symmetric polymatroidal.

Note, δ_h penalizes a cut $(A, \overline{A})$ intersecting h in a directed manner: F^- penalizes $A \cap h$ and F^+ penalizes $\overline{A} \cap h$. For a symmetric polymatroidal cut function δ_h, such penalty is the same, so that $\delta_h(A \cap h) = \delta_h(\overline{A} \cap h)$, generalizing the cut function of undirected graphs and hypergraphs. In the full version [16], we show that Definition 3 captures most of the cut functions proposed in previous work, including all directed and undirected standard hypergraph cut functions. Moreover, our setup also captures possible combinations of all these types of cut functions into a single framework. We highlight two examples relevant in applications:

- F_h is a cardinality-based non-decreasing submodular function [69], a non-decreasing concave function of the cardinality. This is useful when we wish to modify the star-expansion cut function by diminishing the penalty on balanced partitions of h, e.g., by taking $\delta_h(S) = \min\left\{|S|^p, |h \setminus S|^p\right\}$ for $p \in (0, 1)$.
- F_h is the entropy of a subset of random variables associated with vertices of h. When the hypergraph G is the factor graph of an undirected graphical model, this yields a variant of the minimum information partition [28, 59].

We investigate the polynomial-time approximability of ratio-cut problems over submodular hypergraphs with polymatroidal cut functions. In particular, we prove the following result.

▶ **Theorem 4.** *There exists a randomized polynomial-time $O\left(\sqrt{\log n}\right)$-approximation algorithm for solving the minimum ratio-cut problem on weighted submodular hypergraphs equipped with polymatroidal cut functions.*

This result extends the $O\left(\sqrt{\log n}\right)$-approximation of Louis and Makarychev [50] for standard hypergraph sparsest cut to all ratio-cut problems for the more general class of polymatroidal functions. Our algorithm for Theorem 4 solves a new ℓ_2^2-metric relaxation of the ratio-cut problem for weighted submodular hypergraphs This relaxation crucially exploits the form of the Lovász extension of polymatroidal cut functions. To the best of our knowledge, our results constitutes the broadest known generalization of the seminal result of Arora, Rao and Vazirani for sparsest cut, capturing a wealth of partitioning objectives, including directed and undirected graph ratio-cut, vertex-based ratio cut and hypergraph ratio-cut within a single simple analysis. It also improves on the best known approximation ratio for many other objectives in this class, such as cardinality-based cut functions [69]. For details, see the full version of this paper [16].

Polymatroidal Cut Functions and Hypergraph Flows

In the second part of the paper, we explore the efficient solution of the minimum submodular hypergraph ratio cut problem via flow methods based on the cut-matching game. We start by developing the notions of flow-cut duality and flow embeddings over hypergraphs. For this purpose, we define a novel notion of *hypergraph flows* for weighted submodular hypergraphs equipped with polymatroidal cut functions. Specifically, for a weighted submodular hypergraph $G = (V, E, \mathbf{w}, \boldsymbol{\mu})$, we show that the subgradients of polymatroidal cut functions behave like network flows over the *factor graph*, a graph analogue of the input submodular hypergraph that includes an auxiliary vertex for each hyperedge $h \in E$ and an edge $\{i, h\}$ for each $i \in h$. For each hyperedge $h \in E$, the monotone functions F_h^- and F_h^+ define capacity constraints on the flows entering and leaving the corresponding auxiliary node. The main algorithmic result in this part of the paper is a *flow decomposition result for hypergraph flows*, which allows us to construct flow embeddings of directed and undirected graphs into hypergraphs while lower bounding the size of the hypergraph cut with the size of the corresponding graph cut. This will become a fundamental building block for our cut-matching game construction.

Connection to Polymatroidal Networks and Flows

A polymatroidal network [42] is characterized by a directed (2-uniform) graph H representing a flow network where the edges are not individually capacitated, but rather, for every vertex $v \in V(H)$, there exist monotone submodular functions ρ^- and ρ^+ on the constraining the amount of flow entering and exiting a vertex respectively. When viewed as network flows over the factor graph, our hypergraph flows are exactly an instantiation of polymatroidal flows, where each auxiliary node $h \in E$ has associated submodular capacity functions F_h^- and F_h^+. In the same way, it is also possible to interpret the submodular hypergraph partitioning problem with polymatroidal cut functions, as a special case of the ratio-cut problem over polymatroidal networks [14]. This is achieved by casting the factor graph as a polymatroidal network in which the original vertices $v \in V$ have no restrictions on their capacities and the auxiliary vertices $h \in H$ do not contribute to the denominator in the ratio-cut problem. See the recent paper of Chekuri & Louis [15] for more details of this reduction.

1.1.2 Local Formulations of the Submodular Hypergraph Ratio-Cut Problem

To generalize the cut-matching game framework to the setting of submodular hypergraphs with polymatroidal cut functions, we introduce an optimization approach based on a continuous non-convex formulation (`RC-NonCvx`) of the submodular hypergraph ratio-cut problem. Given a weighted submodular hypergraph $G = (V, E, \mathbf{w}, \boldsymbol{\mu})$, and cut functions $\{\delta_h\}_{h \in E}$ with Lovász extensions $\{\bar{\delta}_h\}_{h \in E}$, we define (`RC-NonCvx`) as the following non-convex optimization problem over $\mathbb{R}^V$:

$$\bar{\Psi}_G(\mathbf{x}) \stackrel{\text{def}}{=} \frac{\sum_{h \in E} w_h \bar{\delta}_h(\mathbf{x})}{\min_u \|\mathbf{x} - u\mathbf{1}\|_{\boldsymbol{\mu},1}}$$

$$\bar{\Psi}_G^* \stackrel{\text{def}}{=} \min_{\mathbf{x} \in \mathbb{R}^V} \bar{\Psi}_G(\mathbf{x}) \qquad (\texttt{RC-NonCvx})$$

The fact that solving (`RC-NonCvx`) is equivalent to solving the minimum ratio-cut problem (expressed in the following lemma) is a direct consequence of the submodularity of the cut functions.

▶ **Lemma 5.** *Given a weighted hypergraph $G^* = (V, E, \boldsymbol{\mu}, \mathbf{w})$, we have $\Psi^*_G = \bar{\Psi}^*_G$. Furthermore, there exists an algorithm that, given any $\mathbf{x} \in \mathbb{R}^V$, recovers a cut $S \subseteq V$ satisfying*

$$\Psi_G(S) \leq \bar{\Psi}_G(\mathbf{x})$$

in time $O\big(|V| \log |V| + \mathsf{sp}(G)\big)$, where the sparsity of G is defined as $\mathsf{sp}(G) \stackrel{\text{def}}{=} \sum_{h \in E} |h|$.

The main idea in our approach is to address the non-convexity of the $\bar{\Psi}_G$ objective by replacing the non-concave denominator with a linear lower bound. In particular, we can exploit the dual characterization of the ℓ_1-norm in the lower bound to write:

$$\min_{u \in \mathbb{R}} \|\mathbf{x} - u\mathbf{1}\|_{\boldsymbol{\mu},1} = \min_{u \in \mathbb{R}} \max_{\|\mathbf{y}\|_\infty \leq 1} \langle \mathbf{y}, \mathbf{x} - u\mathbf{1} \rangle_{\boldsymbol{\mu}} = \max_{\substack{\|\mathbf{y}\|_\infty \leq 1 \\ \langle \mathbf{y}, \mathbf{1} \rangle_{\boldsymbol{\mu}} = 0}} \langle \mathbf{y}, \mathbf{x} \rangle_{\boldsymbol{\mu}} \tag{2}$$

This calculation directly leads us to define a novel *localized, convex formulation* of the (RC-NonCvx) problem for each vector $\mathbf{s} \in \mathbb{R}^V$ with $\|\mathbf{s}\|_\infty \leq 1$ and $\langle \mathbf{s}, \mathbf{1} \rangle_{\boldsymbol{\mu}} = 0$, which we call a *seed*:

$$\bar{\Psi}_{G,\mathbf{s}}(\mathbf{x}) \stackrel{\text{def}}{=} \frac{\sum_{h \in E} w_h \bar{\delta}_h(\mathbf{x})}{\max\{0, \langle \mathbf{s}, \mathbf{x} \rangle_{\boldsymbol{\mu}}\}}$$

$$\bar{\Psi}^*_{G,\mathbf{s}} \stackrel{\text{def}}{=} \min_{\mathbf{x} \in \mathbb{R}^V} \bar{\Psi}_{G,\mathbf{s}}(\mathbf{x}) \tag{Local-RC}$$

Intuitively, the program (Local-RC) seeks a distribution over cuts $\mathbf{x}$ with small expected cut size and high correlation with the seed $\mathbf{s}$. In the full version [16], we investigate the properties of this formulation, including its equivalence with an integral cut problem, the ratio-cut improvement problem, which generalizes the cut improvement problem of Andersen and Lang [4] to submodular hypergraphs. For polymatroidal cut functions, we show that the natural notion of a dual solution for the convex problem Local-RC is exactly a hypergraph flow with demand vector $\text{diag}(\boldsymbol{\mu})\mathbf{s}$. By applying our flow decomposition result, we can turn such a hypergraph flow into a *dual graph certificate*, a graph D such that $\Psi^*_{G,\mathbf{s}} \cdot \delta_D(S) \leq \delta_G(S)$ for all $S \subseteq V$. To construct these objects, we prove the following algorithmic result:

▶ **Theorem 6** (Informal. See full version)**.** *An approximate primal-dual solution to* (Local-RC) *can be computed by solving $O(\log |V|)$-many decomposable submodular minimization problems for general polymatroidal cut functions or $O(\log |V|)$-many $1/2$-approximate maximum flow problems over a graph of size $O(\mathsf{sp}(G))$.*

Finally, and crucially for our purposes, for all $\mathbf{x} \in \mathbb{R}^V$, we have $\bar{\Psi}_G(\mathbf{x}) \leq \bar{\Psi}_{G,\mathbf{s}}(\mathbf{x})$ by Equation 2. Hence, any (not necessarily optimal) solution $\mathbf{x}$ to (Local-RC) for a seed $\mathbf{s}$ yields a solution to (RC-NonCvx) of value at most $\bar{\Psi}_{G,\mathbf{s}}(\mathbf{x})$ and, via Lemma 5, a cut S with $\Psi_G(S) \leq \bar{\Psi}_{G,\mathbf{s}}(\mathbf{x})$. In other words, if we can somehow find a seed $\mathbf{s}$ such that $\bar{\Psi}^*_{G,\mathbf{s}} \leq \alpha \Psi^*_G$, then solving the (Local-RC) will yield an α-approximation algorithm for Ψ^*_G.

This reduction may not seem very useful, as finding the required seed may be as hard as solving the original problem. However, we can now exploit the dual feedback obtained when solving (Local-RC) for a sequence of seeds $\mathbf{S} = (\mathbf{s}_1, \mathbf{s}_2, \ldots, \mathbf{s}_T)$ to effectively guide our search for a good seed and boost our approximation ratio. In particular, suppose we can adaptively construct a sequence of T seeds $\mathbf{S}$ such that the sum $H = \sum_{t=1}^T D_t$ of the corresponding dual demand graphs D_t has large ratio-cut, i.e., $\Psi^*_H \geq \beta$. Then, it is easy to show that we must have $\min_{t=1}^T \Psi^*_{G,\mathbf{s}_t} \leq T/\beta \cdot \Psi^*_G$. Hence, solving (Local-RC) for one of the seeds in the sequence $\mathbf{S}$ must yield a T/β-approximation to the minimum ratio-cut problem. It turns out that the problem of constructing such sequence $\mathbf{S}$ is precisely the task addressed by the cut-matching framework of Khandekar et al. [32], which we generalize and strengthen in our next contribution.

1.1.3 An Improved Cut-Matching Game

We now describe our construction of the cut-matching game. Later, we use this to design an approximation algorithm for hypergraph partitioning. We first define what constitutes admissible actions played by the cut and matching players.

▶ **Definition 7.** *Let V be a set of vertices with vertex weights $\boldsymbol{\mu} \in \mathbb{Z}^V_{\geq 0}$. A cut action is a pair (A, B) of non-empty disjoint subsets $A, B \subseteq V$ such that $\boldsymbol{\mu}(A) \leq \boldsymbol{\mu}(B)$.*

The cut player is not restricted to returning a partition of the vertex set, let alone a bisection. The set of valid actions for a matching player is similarly relaxed, as the edges of the response do not need to form a perfect matching. It is only required that enough of the available degree is routed across the cut action.

▶ **Definition 8.** *An* approximate matching response *to a cut action (A, B) is a weighted bipartite graph $D = \big(A \cup B, E_D, \mathbf{w}^D, \boldsymbol{\mu}\big)$ satisfying the following conditions.*

1. *Bounded degree: for all $i \in V$*

$$\deg_D(i) \leq \begin{cases} \mu_i & i \in A \\ \frac{\boldsymbol{\mu}(A)}{\boldsymbol{\mu}(B)} \cdot \mu_i & i \in B \end{cases}.$$

2. *Largeness: $\mathbf{w}^D\big(E(A, B)\big) \geq \frac{\boldsymbol{\mu}(A)}{2}$.*

We may emphasize D is a directed approximate matching response *if edges of D are directed from A to B.*

This definition is motivated by how an approximate matching response naturally arises from the dual graph certificate obtained by approximately solving the (`Local-RC`) program above. The matching response will be undirected when we wish to approximate the minimum ratio-cut of a hypergraph with symmetric cut functions and directed when generically considering asymmetric cut functions. We now define the cut-matching game.

▶ **Definition 9.** *Let $n > 0$, and $\boldsymbol{\mu} \in \mathbb{Z}^V_{\geq 0}$ be a weighting over a set of vertices V where $|V| = n$. A* $(n, \boldsymbol{\mu})$-generalized cut-matching game *is a multi-round, two-player game between a cut player $\mathcal{C}$, and a matching player $\mathcal{M}$ that proceeds as follows. In the t-th round of the game, following definitions 7 and 8:*

1. *$\mathcal{C}$ plays a cut action (A_t, B_t),*
2. *$\mathcal{M}$ responds with an approximate matching response D_t to (A_t, B_t).*

The state graph *after t rounds is $H_t = \sum_{s=1}^t D_s$ the edge-wise sum of the matching response seen thus far.*

For comparison, the original cut-matching game of [32] is captured by the above definition when one specifically considers graph expansion $\boldsymbol{\mu} = \mathbf{1}$, restricts cut actions to be exact bisections $(A, \bar{A})$ where $|A| = |\bar{A}|$, and requires exact matching responses: $\mathbf{w}^D\big(E(A, \bar{A})\big) = \boldsymbol{\mu}(A)$.

To approximate minimum ratio-cut using our cut-matching games, we seek a cut player $\mathcal{C}$ that outputs a sequence of cuts which, in as few cuts as possible, forces the matching player $\mathcal{M}$ to place edges that make the minimum ratio-cut objective for the state graph H_T large. We call a cut player *good* if it can achieve this.

▶ **Definition 10.** *A* cut strategy *is a randomized algorithm $\mathcal{A}_{\mathsf{cut}}$ that takes as input the current state graph H, and outputs a cut action (A, B). A cut strategy $\mathcal{A}_{\mathsf{cut}}$ is $\big(f(n), g(n)\big)$*-good *if a cut player $\mathcal{C}$ using $\mathcal{A}_{\mathsf{cut}}$ at every iteration, achieves:*

$$\Psi_{H_{g(n)}} \geq f(n)$$

with constant probability, for any (potentially adaptive) sequence of approximate matching responses $D_1, \ldots, D_{g(n)}$.

Note that, when dealing with symmetric cut functions, H will be undirected, so that Ψ_H refers to the undirected ratio-cut objective. If the cut functions are asymmetric, then Ψ_H is naturally replaced with the directed ratio-cut objective.

These definitions reveal the mechanism by which the cut-matching framework adaptively constructs a sequence of seeds to (`Local-RC`): the seed we choose at every iteration of cut-matching game is just a $\boldsymbol{\mu}$-weighted indicator vector for the cut action of that round. A good strategy guarantees that, no matter what dual graph certificates we receive for each seed, at the end of the game at $T = g(n)$ iterations, we can guarantee H_T has ratio-cut objective at least $\beta = f(n)$. This shows that a good strategy yields an approximation ratio of $O(g(n)/f(n))$ for minimum ratio cut. Our main technical theorem regarding the new cut-matching game shows the existence of good strategies for the cut-player for both the cases of symmetric and general polymatroidal cut functions.

▶ **Theorem 11.** *Let $H = \big(V, E_H, \mathbf{w}^H, \boldsymbol{\mu}\big)$ be the state graph for an $(n, \boldsymbol{\mu})$-generalized cut-matching game. There exists a cut strategy $\mathcal{A}_{\mathsf{cut}}$ satisfying the following:*
1. *If H is undirected, then $\mathcal{A}_{\mathsf{cut}}$ is $\big(\Omega(\log n), O(\log^2 n)\big)$-good with probability $O(1)$.*
2. *If H is directed, then $\mathcal{A}_{\mathsf{cut}}$ is $\big(\Omega(\log^2 n), O(\log^3 n)\big)$-good with probability $O(1)$.*

Furthermore, $\mathcal{A}_{\mathsf{cut}}$ outputs a cut action in time $\tilde{O}\big(\mathsf{sp}(H)\big)$.

Together with Theorem 6, our results on good cut-player strategies directly imply the following results for the solution of the minimum ratio-cut problem.

▶ **Corollary 12.** *There exists an $O(\log n)$-approximation algorithm for the minimum ratio-cut problem over submodular hypergraphs with polymatroidal cut functions whose running time is dominated by the solution of a polylogarthmic number of decomposable submodular minimization problems over the hypergraph cut function.*

▶ **Corollary 13.** *There exists an $O(\log n)$-approximation algorithm for the minimum ratio-cut problem over submodular hypergraphs equipped with the directed or standard hypergraph cut function whose running time is almost linear in the hypergraph sparsity.*

In many cases of interest, the decomposable submodular minimization problem can be solved in almost linear time [8], so that Corollary 12 yields the first $O(\log V|)$-approximation algorithms for many submodular hypergraph ratio-cut problems with polymatroidal cut functions. Corollary 13 shows that the specialization of this algorithm to the directed and standard hypergraph cut functions also yields $O(\log |V|)$-approximation almost linear time algorithms. Concurrently to our work, this latter result was improved to an $O(\sqrt{\log |V|}$-approximation in almost linear time by Lau et al. [37]. However, our algorithm is still the fastest one to achieve an $O(\log |V|)$-approximation ratio for general polymatroidal cut functions.

Approximate Matching Responses

The concept of an approximate matching response in Definition 8 is carefully tailored to enable its implementation by a $O(1)$-approximate maximum flow algorithm (for directed and standard hypergraph problems) or a $O(1)$-approximate decomposable submodular minimization problem (for general polymatroidal cut functions). This contrasts with previous applications of the cut-matching game, which either require a single $o_n(1)$-approximate solve per round [32], at the cost of higher running time, or a logarithmic number of $O(1)$-approximate solves per round [58], at the cost of a logarithmic loss in approximation ratio. As a result, our novel cut-matching game can be plugged in to improve the analyses of down-stream applications of the cut-matching game [10, 17, 18, 19, 47, 58]. The low-precision required by our algorithm may also impact the deployment of cut-matching games in practice as high precision solves for maximum flow or submodular minimization are a running-time bottleneck of algorithms based on cut-matching games [60, 68].

Note that the additional freedom in the matching responses, which can now fail to match up to a constant fraction of vertices, makes it harder to establish the existence of good cut strategies. Showing Theorem 11 thus requires a new geometric result regarding well-separated sets of embedding vectors, discussed below.

Separated Sets

In our new cut-matching game, the cut player must pay considerable more care in choosing a cut action (A, B) than in the original version, as the approximate nature of the matching response means that the matching player can easily keep some small subset of vertices disconnected from the rest of the graph. Suppose for instance that the cut player restricted itself to play μ-bisections. Then, the matching player could select a small cut T, with $\mu(T) \leq 1/10 \cdot \mu(V)$ and never place any edge across the boundary of T. The only way the cut player can force progress on T is to depart from playing bisections and play a cut action (A, B) where A is close or equal to T.

How is this intuition formalized in our design of cut-player strategies? As in the analysis of Orecchia et al. [61], our cut strategy maintains at each iteration t a spectral embedding for the current state graph H_t. We make progress if we can force the matching player to connect vertices that are far away in this embedding. We show that we can achieve this, even against an approximate matching player, by choosing a cut action coming from the following notion of separated sets:

▶ **Definition 14** (σ-separated sets). *Given an embedding $\{\mathbf{v}_i\}_{i\in[n]}$ and a measure $\boldsymbol{\mu} \in \mathbb{R}^n_{\geq 0}$ we say two sets $S, T \subseteq [n]$ are* σ-separated *if:*

$$\boldsymbol{\mu}(S) \cdot \boldsymbol{\mu}(T) \cdot \min_{\substack{i\in S\\ j\in T}} \|\mathbf{v}_i - \mathbf{v}_j\|^2 \geq \sigma \cdot \sum_{i,j\in\binom{V}{2}} \mu_i\mu_j \cdot \|\mathbf{v}_i - \mathbf{v}_j\|^2.$$

This definition generalizes the idea of well-separated sets [6], which was applied only to large balanced sets, to possibly unbalanced sets, as long as S and T contribute a σ-fraction to the total variance of the embedding. In the full version [16], we show how to efficiently find separated sets and prove the following result.

▶ **Theorem 15.** *Given an embedding $\{\mathbf{v}_i\}_{i\in[n]} \subseteq \mathbb{R}^d$ and a measure $\boldsymbol{\mu} \in \mathbb{Z}^n_{\geq 0}$, with $\mu(V) =$ poly(n), there exists a randomized algorithm running in time $\tilde{O}(nd)$ that outputs σ-separated sets S and T for $\sigma = \Omega(1/\log n)$ with high probability.*

1.1.4 Open Problems

Chekuri and Louis [15] have recently shown that our metric analysis, originally devised for submodular hypergraphs, can be applied more generally to partitioning problems over polymatroidal networks. This naturally leads to the question of whether our second contribution, the design of fast algorithms based on the cut-matching game, can also be extended to the general setting of polymatroidal networks to yield $O(\log |V|)$ approximation algorithms.

An even more ambitious goal is to adapt the recent work of Kolmogorov [34] which provides a simpler construction of Sherman's $O(\sqrt{\log n})$-approximation for sparsest cut, to the minimum-ratio cut problem over general polymatroidal networks. The ultimate goal would be to obtain $O(\sqrt{\log |V|/\epsilon})$-approximation algorithms that only requires the solution of $O(|V|^\epsilon)$ decomposable sumbodular minimization problems, for $\epsilon \in \Theta(1/\log |V|, \Theta(1))$.

References

1 Sameer Agarwal, Kristin Branson, and Serge Belongie. Higher order learning with graphs. In *Proceedings of the 23rd international conference on machine learning*, pages 17–24. Association for Computing Machinery, 2006. doi:10.1145/1143844.1143847.

2 Konstantinos Ameranis, Antares Chen, Lorenzo Orecchia, and Erasmo Tani. Efficient flow-based approximation algorithms for submodular hypergraph partitioning via a generalized cut-matching game. *arXiv preprint arXiv:2301.08920*, 2023. doi:10.48550/arXiv.2301.08920.

3 Aditya Anand, Euiwoong Lee, Jason Li, and Thatchaphol Saranurak. Approximating small sparse cuts. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 319–330, 2024. doi:10.1145/3618260.3649747.

4 Reid Andersen and Kevin J. Lang. An algorithm for improving graph partitions. In Shang-Hua Teng, editor, *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2008, San Francisco, California, USA, January 20-22, 2008*, SODA '08, pages 651–660, USA, 2008. SIAM. URL: http://dl.acm.org/citation.cfm?id=1347082.1347154.

5 Sanjeev Arora, Elad Hazan, and Satyen Kale. O(sqrt(log(n)) approximation to SPARSEST CUT in õ(n^2) time. *SIAM J. Comput.*, 39(5):1748–1771, January 2010. doi:10.1137/080731049.

6 Sanjeev Arora, Satish Rao, and Umesh Vazirani. Expander flows, geometric embeddings and graph partitioning. *Journal of the ACM (JACM)*, 56(2):1–37, 2009. doi:10.1145/1502793.1502794.

7 Giorgio Ausiello, Mario Lucertini, et al. *Analysis and design of algorithms in combinatorial optimization*, volume 266. Springer, 1981.

8 Kyriakos Axiotis, Adam Karczmarz, Anish Mukherjee, Piotr Sankowski, and Adrian Vladu. Decomposable submodular function minimization via maximum flow. In *Proceedings of the 38th International Conference on Machine Learning*, pages 446–456. PMLR, 2021. URL: https://proceedings.mlr.press/v139/axiotis21a.html.

9 Austin R. Benson, David F. Gleich, and Jure Leskovec. Higher-order organization of complex networks. *CoRR*, abs/1612.08447(6295):163–166, 2016. doi:10.48550/arXiv.1612.08447.

10 Aaron Bernstein, Maximilian Probst Gutenberg, and Thatchaphol Saranurak. Deterministic decremental SSSP and approximate min-cost flow in almost-linear time. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 1000–1008. IEEE, IEEE, 2021. doi:10.1109/FOCS52979.2021.00100.

11 Umit V Catalyurek and Cevdet Aykanat. Hypergraph-partitioning-based decomposition for parallel sparse-matrix vector multiplication. *IEEE Transactions on parallel and distributed systems*, 10(7):673–693, 1999. doi:10.1109/71.780863.

12 T-H Hubert Chan, Anand Louis, Zhihao Gavin Tang, and Chenzi Zhang. Spectral properties of hypergraph laplacian and approximation algorithms. *Journal of the ACM (JACM)*, 65(3):1–48, 2018. doi:10.1145/3178123.

13 Yuk Hei Chan and Lap Chi Lau. On linear and semidefinite programming relaxations for hypergraph matching. *Mathematical programming*, 135(1-2):123–148, 2012. doi:10.1007/s10107-011-0451-5.

14 Chandra Chekuri, Sreeram Kannan, Adnan Raja, and Pramod Viswanath. Multicommodity flows and cuts in polymatroidal networks. *SIAM J. Comput.*, 44(4):912–943, January 2015. doi:10.1137/130906830.

15 Chandra Chekuri and Anand Louis. On sparsest cut and conductance in directed polymatroidal networks. *CoRR*, abs/2410.20525, 2024. doi:10.48550/arXiv.2410.20525.

16 Antares Chen, Lorenzo Orecchia, and Erasmo Tani. Submodular Hypergraph Partitioning: Metric Relaxations and Fast Algorithms via an Improved Cut-Matching Game, July 2023. doi:10.48550/arXiv.2301.08920.

17 Julia Chuzhoy. A Distanced Matching Game, Decremental APSP in Expanders, and Faster Deterministic Algorithms for Graph Cut Problems. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Proceedings, pages 2122–2213. Society for Industrial and Applied Mathematics, January 2023. doi:10.1137/1.9781611977554.ch82.

18 Julia Chuzhoy, Yu Gao, Jason Li, Danupon Nanongkai, Richard Peng, and Thatchaphol Saranurak. A deterministic algorithm for balanced cut with applications to dynamic connectivity, flows, and beyond. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1158–1167. IEEE, 2020. doi:10.1109/FOCS46700.2020.00111.

19 Julia Chuzhoy and Sanjeev Khanna. A new algorithm for decremental single-source shortest paths with applications to vertex-capacitated flow and cut problems. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 389–400, 2019. doi:10.1145/3313276.3316320.

20 William H Cunningham and András Frank. A primal-dual algorithm for submodular flows. *Mathematics of Operations Research*, 10(2):251–262, 1985. doi:10.1287/moor.10.2.251.

21 Karen D. Devine, Erik G. Boman, Robert T. Heaphy, Rob H. Bisseling, and Ümit V. Çatalyürek. Parallel hypergraph partitioning for scientific computing. In *20th International Parallel and Distributed Processing Symposium (IPDPS 2006), Proceedings, 25-29 April 2006, Rhodes Island, Greece*, pages 10–pp. IEEE, IEEE, 2006. doi:10.1109/IPDPS.2006.1639359.

22 Henry Harcis Driver. *Partitioning of weighted hypergraphs.* PhD thesis, University of Delaware, 1977.

23 Jack Edmonds and Rick Giles. A min-max relation for submodular functions on graphs. In P.L. Hammer, E.L. Johnson, B.H. Korte, and G.L. Nemhauser, editors, *Studies in Integer Programming*, volume 1 of *Annals of Discrete Mathematics*, pages 185–204. Elsevier, 1977. doi:10.1016/S0167-5060(08)70734-9.

24 Song Feng, Emily Heath, Brett Jefferson, Cliff Joslyn, Henry Kvinge, Hugh D Mitchell, Brenda Praggastis, Amie J Eisfeld, Amy C Sims, Larissa B Thackray, et al. Hypergraph models of biological networks to identify genes critical to pathogenic viral response. *BMC bioinformatics*, 22(1):1–21, 2021. doi:10.1186/s12859-021-04197-2.

25 Satoru Fujishige. Algorithms for solving the independent-flow problems. *Journal of the Operations Research Society of Japan*, 21(2):189–204, 1978.

26 Bernhard Haeupler, Jonas Hübotter, and Mohsen Ghaffari. A cut-matching game for constant-hop expanders. In Yossi Azar and Debmalya Panigrahi, editors, *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025, New Orleans, LA, USA, January 12-15, 2025*, pages 1651–1678. SIAM, SIAM, 2025. doi:10.1137/1.9781611978322.51.

27 Refael Hassin. *On network flows.* PhD thesis, Yale University, 1978.

28 Shohei Hidaka and Masafumi Oizumi. Fast and exact search for the partition with minimal information loss. *PLoS One*, 13(9):e0201126, 2018. doi:10.48550/arXiv.1708.01444.

29 Arun Jambulapati, Yin Tat Lee, Jerry Li, Swati Padmanabhan, and Kevin Tian. Positive semidefinite programming: mixed, parallel, and width-independent. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 789–802, 2020. doi:10.1145/3357713.3384338.

30 Haotian Jiang, Tarun Kathuria, Yin Tat Lee, Swati Padmanabhan, and Zhao Song. A faster interior point method for semidefinite programming. In *2020 IEEE 61st annual symposium on foundations of computer science (FOCS)*, pages 910–918. IEEE, 2020. doi:10.1109/FOCS46700.2020.00089.

31 Rohit Khandekar, Subhash Khot, Lorenzo Orecchia, and Nisheeth K Vishnoi. On a cut-matching game for the sparsest cut problem. *Univ. California, Berkeley, CA, USA, Tech. Rep. UCB/EECS-2007-177*, 6(7):12, 2007.

32 Rohit Khandekar, Satish Rao, and Umesh Vazirani. Graph partitioning using single commodity flows. *Journal of the ACM (JACM)*, 56(4):1–15, 2009. doi:10.1145/1538902.1538903.

33 Steffen Klamt, Utz-Uwe Haus, and Fabian Theis. Hypergraphs and cellular networks. *PLoS computational biology*, 5(5):e1000385, 2009. doi:10.1371/journal.pcbi.1000385.

34 Vladimir Kolmogorov. A simpler and parallelizable $O(\sqrt{\log n})$-approximation algorithm for sparsest cut, April 2024. doi:10.48550/arXiv.2307.00115.

35 Tsz Chiu Kwok, Lap Chi Lau, and Kam Chuen Tung. Cheeger inequalities for vertex expansion and reweighted eigenvalues. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 366–377. IEEE, 2022. doi:10.1109/FOCS54457.2022.00042.

36 Lap Chi Lau, Kam Chuen Tung, and Robert Wang. Cheeger inequalities for directed graphs and hypergraphs using reweighted eigenvalues. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1834–1847, 2023. doi:10.1145/3564246.3585139.

37 Lap Chi Lau, Kam Chuen Tung, and Robert Wang. Fast algorithms for directed graph partitioning using flows and reweighted eigenvalues. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 591–624. ACM-SIAM, 2024. doi:10.1137/1.9781611977912.22.

38 EL Lawer. Electrical assemblies with a minimum number of interconnections. *IRE transactions on Electronic Computers*, 1:86–88, 1962. doi:10.1109/TEC.1962.5219327.

39 Eugene L. Lawler. Cutsets and partitions of hypergraphs. *Networks*, 3(3):275–285, May 1973. doi:10.1002/net.3230030306.

40 Eugene L Lawler. An introduction to polymatroidal network flows. In *Analysis and Design of Algorithms in Combinatorial Optimization*, pages 129–145. Springer, 1981.

41 Eugene L Lawler. Generalizations of the polymatroidal network flow model, 1982.

42 Eugene L. Lawler and Charles U. Martel. Computing maximal "polymatroidal" network flows. *Math. Oper. Res.*, 7(3):334–347, 1982. doi:10.1287/moor.7.3.334.

43 Eugene L Lawler and Charles U Martel. Flow network formulations of polymatroid optimization problems. In *North-Holland Mathematics Studies*, volume 66, pages 189–200. Elsevier, 1982.

44 Tom Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM (JACM)*, 46(6):787–832, 1999. doi:10.1145/331524.331526.

45 Pan Li and Olgica Milenkovic. Inhomogeneous hypergraph clustering with applications. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, volume 30, pages 2308–2318, 2017. URL: https://proceedings.neurips.cc/paper/2017/hash/a50abba8132a77191791390c3eb19fe7-Abstract.html.

46 Pan Li and Olgica Milenkovic. Submodular hypergraphs: p-laplacians, cheeger inequalities and spectral clustering. In Jennifer G. Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 3020–3029. PMLR, PMLR, 2018. URL: http://proceedings.mlr.press/v80/li18e.html.

47 Yaowei Long and Thatchaphol Saranurak. Near-optimal deterministic vertex-failure connectivity oracles. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1002–1010. IEEE, 2022. doi:10.1109/FOCS54457.2022.00098.

48 Anand Louis. Cut-matching games on directed graphs. *CoRR*, abs/1010.1047, 2010. doi:10.48550/arXiv.1010.1047.

49 Anand Louis. Hypergraph markov operators, eigenvalues and approximation algorithms. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 713–722, 2015. doi:10.1145/2746539.2746555.

50 Anand Louis and Yury Makarychev. Approximation algorithms for hypergraph small-set expansion and small-set vertex expansion. *Theory of Computing*, 12(1):1–25, 2016. doi:10.4086/toc.2016.v012a017.

51 Anand Louis, Prasad Raghavendra, and Santosh Vempala. The complexity of approximating vertex expansion. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 360–369. IEEE, 2013. doi:10.1109/FOCS.2013.46.

52 F. Luccio and M. G. Sami. On the decomposition of networks into minimally interconnected subnetworks. *IEEE Transactions on Circuit Theory*, 2:184–188, 1972.

53 Charles Martel. Scheduling uniform machines with release times, deadlines and due times. In *Deterministic and Stochastic Scheduling: Proceedings of the NATO Advanced Study and Research Institute on Theoretical Approaches to Scheduling Problems held in Durham, England, July 6–17, 1981*, pages 89–99. Springer, 1982.

54 Charles U Martel. *Generalized network flows with an application to multiprocessor scheduling*. University of California, Berkeley, 1980.

55 Charles U. Martel. Preemptive scheduling with release times, deadlines, and due times. *J. ACM*, 29(3):812–829, July 1982. doi:10.1145/322326.322337.

56 G. Martella and M. G. Sami. A proposal for the organization of large files in a hierarchical structure. Technical report, Instituto di Elettotecnica ed Elettronica del Politecnico di Milano, Italy, 1972.

57 A. N. Melikhov, L. S. Bershtein, V. V. Selyankin, and M. I. Khil. Use of hypergraphs for arranging circuits in a unit. *Engineering Cybernetics*, 12(3):157–162, 1974.

58 Danupon Nanongkai and Thatchaphol Saranurak. Dynamic spanning forest with worst-case update time: adaptive, las vegas, and o($n^{1/2 - \epsilon}$)-time. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 1122–1129. ACM, 2017. doi:10.1145/3055399.3055447.

59 Mukund Narasimhan, Nebojsa Jojic, and Jeff A. Bilmes. Q-clustering. In *Advances in Neural Information Processing Systems 18 [Neural Information Processing Systems, NIPS 2005, December 5-8, 2005, Vancouver, British Columbia, Canada]*, volume 18, pages 979–986, 2005. URL: https://proceedings.neurips.cc/paper/2005/hash/b0bef4c9a6e50d43880191492d4fc827-Abstract.html.

60 Lorenzo Orecchia, Konstantinos Ameranis, Charalampos Tsourakakis, and Kunal Talwar. Practical almost-linear-time approximation algorithms for hybrid and overlapping graph clustering. In *International Conference on Machine Learning*, pages 17071–17093. PMLR, 2022. URL: https://proceedings.mlr.press/v162/orecchia22a.html.

61 Lorenzo Orecchia, Leonard J Schulman, Umesh V Vazirani, and Nisheeth K Vishnoi. On partitioning graphs via single commodity flows. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 461–470, 2008. doi:10.1145/1374376.1374442.

62 Shailendra C Parikh and William S Jewell. Decomposition of project networks. *Management Science*, 11(3):444–459, 1965.

63 Alexander Schrijver. Total dual integrality from directed graphs, crossing families, and sub-and supermodular functions. In *Progress in combinatorial optimization*, pages 315–361. Elsevier, 1984.

64 Alexander Schrijver et al. *Combinatorial optimization: polyhedra and efficiency*. Springer, 2003.

65 Jonah Sherman. Breaking the multicommodity flow barrier for o(vlog n)-approximations to sparsest cut. In *50th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2009, October 25-27, 2009, Atlanta, Georgia, USA*, pages 363–372. IEEE, IEEE Computer Society, 2009. doi:10.1109/FOCS.2009.66.

66 Zoya Svitkina and Lisa Fleischer. Submodular approximation: Sampling-based algorithms and lower bounds. *SIAM Journal on Computing*, 40(6):1715–1737, 2011. doi:10.1137/100783352.

67 Charalampos E Tsourakakis, Jakub Pachocki, and Michael Mitzenmacher. Scalable motif-aware graph clustering. In *Proceedings of the 26th International Conference on World Wide Web*, pages 1451–1460, 2017. doi:10.1145/3038912.3052653.

68 Nate Veldt. Cut-matching games for generalized hypergraph ratio cuts. In Ying Ding, Jie Tang, Juan F. Sequeda, Lora Aroyo, Carlos Castillo, and Geert-Jan Houben, editors, *Proceedings of the ACM Web Conference 2023, WWW 2023, Austin, TX, USA, 30 April 2023 - 4 May 2023*, WWW '23, pages 694–704, New York, NY, USA, 2023. ACM. doi:10.1145/3543507.3583539.

69 Nate Veldt, Austin R Benson, and Jon Kleinberg. Approximate decomposable submodular function minimization for cardinality-based components. *Advances in Neural Information Processing Systems*, 34:3744–3756, 2021. URL: `https://proceedings.neurips.cc/paper/2021/hash/1e8a19426224ca89e83cef47f1e7f53b-Abstract.html`.

70 Nate Veldt, Austin R Benson, and Jon Kleinberg. Hypergraph cuts with general splitting functions. *SIAM Review*, 64(3):650–685, 2022. `doi:10.1137/20m1321048`.

71 Wenyin Yang, Guojun Wang, Md Zakirul Alam Bhuiyan, and Kim-Kwang Raymond Choo. Hypergraph partitioning for social networks based on information entropy modularity. *Journal of Network and Computer Applications*, 86:59–71, 2017. `doi:10.1016/j.jnca.2016.10.002`.

72 Yuichi Yoshida. Cheeger inequalities for submodular transformations. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2582–2601. SIAM, 2019. `doi:10.1137/1.9781611975482.160`.