

An Improved Integrality Gap for the Călinescu-Karloff-Rabani Relaxation for Multiway Cut

Haris Angelidakis*
TTIC

Yury Makarychev†
TTIC

Pasin Manurangsi‡
UC Berkeley

November 18, 2016

Abstract

We construct an improved integrality gap instance for the Călinescu-Karloff-Rabani LP relaxation of the Multiway Cut problem. In particular, for $k \geq 3$ terminals, our instance has an integrality ratio of $6/(5 + \frac{1}{k-1}) - \varepsilon$, for every constant $\varepsilon > 0$. For every $k \geq 4$, our result improves upon a long-standing lower bound of $8/(7 + \frac{1}{k-1})$ by Freund and Karloff [FK00]. Due to Manokaran et al.'s result [MNRS08], our integrality gap also implies Unique Games hardness of approximating Multiway Cut of the same ratio.

1 Introduction

In the Multiway Cut problem, we are given a weighted undirected graph and a set of k terminal vertices, and we are asked to find a set of edges with minimum total weight whose removal disconnects all the terminals. Equivalently, the objective can also be stated as partitioning the graph into k clusters, each containing exactly one terminal, such that the total weight of edges across clusters is minimized. Such a partition is sometimes called a *k-way cut* and the total weight of edges across clusters is called the *cost* of the cut.

Since its introduction in the early 1980s by Dahlhaus, Johnson, Papadimitriou, Seymour, and Yannakakis [DJP+83], the problem has been extensively studied in the approximation algorithms community [Cun89, CR91, DJP+94, BTV99, CT99, CKR00, FK00, KKS+04, BNS13, SV14] and has, by now, become one of the classic graph partitioning problems taught in many graduate approximation algorithms courses. Despite this, the approximability of the problem when $k \geq 4$ still remains open.

When $k = 2$, Multiway Cut is simply Minimum s - t Cut, which is solvable in polynomial time. For $k \geq 3$, Dahlhaus et al. [DJP+94] showed that the problem becomes APX-hard and gave the first approximation algorithm for the problem, which achieves an approximation ratio of $(2 - 2/k)$. Due to the combinatorial nature of the algorithm, many linear programs were proposed following their work (see e.g. [Cun89, CR91, BTV99]). However, it was not until Călinescu, Karloff and Rabani's work [CKR00] that a significant improvement in the approximation ratio was made. Their linear programming relaxation, commonly referred to as the CKR relaxation, on a graph $G = (V, E, w)$ and terminals $\{t_1, \dots, t_k\} \subseteq V$, can be formulated as

$$\begin{aligned} & \text{minimize} && \sum_{e=(u,v) \in E} w(e) \cdot \frac{1}{2} \|x^u - x^v\|_1 \\ & \text{subject to} && \forall u \in V, x^u \in \Delta_k, \\ & && \forall i \in [k], x^{t_i} = e^i, \end{aligned}$$

*Email: hangel@ttic.edu.

†Email: yury@ttic.edu. Supported by NSF awards CAREER CCF-1150062 and IIS-1302662.

‡Email: pasin@berkeley.edu. Supported by NSF Grants No. CCF-1540685 and CCF-1655215. Part of this work was done while the author was visiting TTIC.

where $[k] = \{1, \dots, k\}$, $\Delta_k = \{(x_1, \dots, x_k) \in [0, 1]^k \mid x_1 + \dots + x_k = 1\}$ denotes the k -simplex and e^i is the i -th vertex of the simplex, i.e. $e_i^i = 1$. Călinescu, Karloff and Rabani gave a rounding scheme for this LP that yields a $(3/2 - 1/k)$ -approximation algorithm for Multiway Cut [CKR00]. Exploiting the geometric nature of the relaxation even further, Karger, Klein, Stein, Thorup and Young [KKS⁺04] gave improved rounding schemes for the relaxation; for general k , they gave a 1.3438-approximation algorithm for the problem. They also conducted computational experiments that led to improvements over small k 's. For $k = 3$, in particular, they gave a 12/11-approximation algorithm and proved that this is tight by constructing an integrality gap example of ratio $12/11 - \varepsilon$, for every $\varepsilon > 0$. The same result was also independently discovered by Cunningham and Tang [CT99]. More recently, Buchbinder, Naor and Schwartz [BNS13] gave a neat 4/3-approximation algorithm to the problem for general k and additionally showed how to push the ratio down to 1.3239. Their algorithm was later improved by Sharma and Vondrák [SV14] to get an approximation ratio of 1.2965. This remains the state-of-the-art approximation for sufficiently large k . We remark that the Sharma-Vondrák algorithm is quite complicated and requires a computer-assisted proof. To circumvent this, Buchbinder, Schwartz and Weizman [BSW17] recently came up with a simplified algorithm and analytically showed that it yielded roughly the same approximation ratio as Sharma and Vondrák's.

The CKR relaxation not only leads to improved algorithms for Multiway Cut but also has a remarkable consequence on the approximability of the problem. In a work by Manokaran, Naor, Raghavendra and Schwartz [MNR08], it was shown, assuming the Unique Games Conjecture (UGC), that, if there exists an instance of Multiway Cut with integrality gap τ for the CKR relaxation, then it is NP-hard to approximate Multiway Cut to within a factor of $\tau - \varepsilon$ of the optimum, for every constant $\varepsilon > 0$. Note here that the integrality gap of an instance of Multiway Cut is the ratio between the integral optimum and the value of the CKR relaxation of the instance.

Roughly speaking, Manokaran et al.'s result means that, if one believes in the UGC, the CKR relaxation achieves essentially the best approximation ratio one can hope to get in polynomial time for Multiway Cut. Despite this strong connection, few lower bounds for the CKR relaxation are known. Apart from the aforementioned $12/11 - \varepsilon$ integrality gap for $k = 3$ by Karger et al. [KKS⁺04] and Cunningham and Tang [CT99], the only other known lower bound is an $8/(7 + \frac{1}{k-1})$ -integrality gap which was constructed by Freund and Karloff [FK00] not long after the introduction of the CKR relaxation. Hence, it is a natural and intriguing task to construct improved integrality gap examples for the relaxation in an attempt to bridge the gap between approximation algorithms and hardness of approximation for the Multiway Cut.

1.1 Our Contributions

In this paper, we provide a new construction of integrality gap instances for Multiway Cut. Our construction achieves an integrality gap of $6/(5 + \frac{1}{k-1}) - \varepsilon$, for every $k \geq 3$, as stated formally below. For every $k \geq 4$, our integrality gap improves on the integrality gap of $8/(7 + \frac{1}{k-1})$ of Freund and Karloff [FK00]. When $k = 3$, our gap matches the known results of Karger et al. [KKS⁺04] and Cunningham and Tang [CT99].

Theorem 1 *For every $k \geq 3$ and every $\varepsilon > 0$, there exists an instance $\mathcal{I}_{k,\varepsilon}$ of Multiway Cut with k terminals such that the integrality gap of the CKR relaxation for $\mathcal{I}_{k,\varepsilon}$ is at least $6/(5 + \frac{1}{k-1}) - \varepsilon$.*

As mentioned earlier, Manokaran et al. showed that integrality gaps for the CKR relaxation can be translated directly to a Unique Games hardness of approximation [MNR08]. Hence, the following corollary holds as an immediate consequence of our integrality gap.

Corollary 2 *Assuming the Unique Games Conjecture, for every $k \geq 3$ and every $\varepsilon > 0$, it is NP-hard to approximate Multiway Cut with k terminals to within a factor of $6/(5 + \frac{1}{k-1}) - \varepsilon$ of the optimum.*

1.1.1 Techniques

To see the motivation behind our construction, it is best to first gain additional intuition for the geometry of the CKR relaxation. For brevity, we will not be thoroughly formal in this section.

Geometric Interpretation of the CKR Relaxation. A solution of the CKR relaxation embeds the graph into a simplex; each vertex u becomes a point $x^u \in \Delta_k$, while each edge (u, v) becomes a segment (x^u, x^v) . As a result, to construct a randomized rounding scheme for the relaxation, it is sufficient to define a randomized k -partitioning scheme of the simplex. More specifically, this is a distribution \mathcal{P} on k -way cuts of Δ_k (i.e. $\forall P \in \mathcal{P}, P : \Delta_k \rightarrow [k]$ such that $P(e^i) = i$ for every $i \in [k]$) and its performance is measured by its *maximum density*, which is defined as

$$\tau_k(\mathcal{P}) := \sup_{x \neq y \in \Delta_k} \frac{\Pr_{P \sim \mathcal{P}}[P(x) \neq P(y)]}{\frac{1}{2}\|x - y\|_1}.$$

In words, for any two different points $x, y \in \Delta_k$, a random k -way cut P sampled from \mathcal{P} assigns the two points to different clusters with probability at most $\tau_k(\mathcal{P}) \cdot \frac{1}{2}\|x - y\|_1$. This immediately yields the following randomized rounding scheme for the CKR relaxation: pick a random $P \sim \mathcal{P}$ and place each $u \in V$ in cluster $P(x^u)$. From the definition of maximum density, the expected contribution of each edge $e = (u, v)$ in the rounded solution is at most $\tau_k(\mathcal{P}) \cdot w(e) \cdot \frac{1}{2}\|x^u - x^v\|_1$, meaning that this rounding algorithm indeed gives a $\tau_k(\mathcal{P})$ -approximation for Multiway Cut with k terminals.

The above observation was implicit in [CKR00] and was first made explicit by Karger et al. [KKS⁺04], who also proved the opposite side of the inequality: for every $\varepsilon > 0$, there is an instance \mathcal{I} of Multiway Cut with k terminals whose integrality gap is at least $\tau_k^* - \varepsilon$, where τ_k^* is the minimum¹ of $\tau_k(\mathcal{P})$ among all \mathcal{P} 's. In other words, constructing an integrality gap for Multiway Cut is equivalent to proving a lower bound for τ_k^* .

Non-Opposite Cuts. Let us now consider a special class of partitions of the simplex, which we call *non-opposite cuts*. A non-opposite cut of Δ_k is a function $P : \Delta_k \rightarrow [k + 1]$ such that $P(e^i) = i$ for every $i \in [k]$ and, for every $x \in \Delta_k$, $P(x) \in \text{supp}(x) \cup \{k + 1\}$, where $\text{supp}(x) := \{i \in [k] \mid x_i \neq 0\}$ is the set of all non-zero coordinates of x . Note here that P now partitions Δ_k into $k + 1$ parts and the additional condition requires that a point that lies on any face of the simplex is either assigned to the $(k + 1)$ -th cluster or to a cluster corresponding to a vertex of the simplex that belongs to this face. The representative case to keep in mind is when $k = 3$; in this case, a non-opposite cut partitions the triangle Δ_3 into four parts and any point on the border of Δ_3 is not assigned to the cluster corresponding to the opposite vertex of the triangle. Figure 1 demonstrates the difference between non-opposite cuts and normal 3-way cuts of Δ_3 .




Figure 1: Illustrations of cuts of Δ_3 . The 3-way cut in Figure 1a is not non-opposite since some points on the line between e^2 and e^3 belong to the cluster corresponding to e^1 . On the other hand, the 3-way cut in Figure 1b is non-opposite. Finally, Figure 1c demonstrates an example of a non-opposite cut that is not a 3-way cut because the fourth cluster is not empty.

We again define the maximum density τ_k for a distribution on non-opposite cuts similarly. Additionally, let $\tilde{\tau}_k^*$ be the infimum of $\tau_k(\mathcal{P})$ among all distributions \mathcal{P} on non-opposite cuts of Δ_k . Observe first that $\tau_k^* \leq \tilde{\tau}_k^*$

¹In [KKS⁺04], τ_k^* is defined as the infimum of $\tau_k(\mathcal{P})$ among all \mathcal{P} 's but it was proved in the same work that there exists \mathcal{P} that achieves the infimum.

because each non-opposite cut P of Δ_k can be turned into a k -way cut without separating additional pairs of points by merging the $(k+1)$ -th cluster into the first cluster. This means that it will be easier for us to show a lower bound for $\tilde{\tau}_k^*$. Unfortunately, it is so far unclear whether such lower bound implies any lower bound for τ_k^* . However, we will see in a moment that we can define an approximate version of $\tilde{\tau}_k^*$ that will give us a lower bound of τ_K^* for all $K > k$ as well.

The approximation of $\tilde{\tau}_k^*$ we use is the one obtained by the discretization of the problem in which we restrict ourselves to the triangular grid $\Delta_{k,n} := \{x \in \Delta_k \mid \forall i \in [k], x_i \text{ is a multiple of } 1/n\}$ instead of Δ_k . For every distribution \mathcal{P} on non-opposite cuts of Δ_k , we can similarly define $\tau_{k,n}(P)$ as

$$\tau_{k,n}(\mathcal{P}) := \max_{x \neq y \in \Delta_{k,n}} \frac{\Pr_{P \sim \mathcal{P}}[P(x) \neq P(y)]}{\frac{1}{2}\|x - y\|_1}$$

and $\tilde{\tau}_{k,n}^*$ as the minimum² of $\tau_{k,n}(\mathcal{P})$ among all such distributions \mathcal{P} 's on non-opposite cuts. The simple observation that allows us to construct our integrality gap is the following relation between τ_K^* and $\tilde{\tau}_{k,n}^*$: for every $K > k$, $\tilde{\tau}_{k,n}^* - O(kn/(K-k)) \leq \tau_K^*$. In fact, we only use this inequality when $k = 3$ to construct our gap and, hence, we only sketch the proof of this case here. For completeness, we prove the general case of the inequality in Appendix B.

Before we sketch the proof of this inequality, let us note that Freund and Karloff's result [FK00] can be seen as a proof of $\tilde{\tau}_{3,2}^* \geq 8/7$; please refer to Appendix A for a more detailed explanation. Note that $\tilde{\tau}_{3,2}^* \geq 8/7$, together with the above fact, immediately implies a lower bound of $\tau_K^* \geq 8/7 - O(1/K)$. Barring the dependency on K , this is exactly the gap proven in [FK00]. With this interpretation in mind, the rest of our work can be mostly seen as proving that $\tilde{\tau}_{3,n}^* \geq 6/5 - O(1/n)$. By selecting $n = \Theta(\sqrt{K})$, this implies a lower bound of $6/5 - O(1/\sqrt{K})$ for τ_K^* ; more care can then be taken to get the right dependency on K .

We now sketch the proof of $\tilde{\tau}_{3,n}^* - O(n/K) \leq \tau_K^*$. Suppose that \mathcal{P} is the distribution on K -way cuts such that $\tau_K(\mathcal{P}) = \tau_K^*$. We sample a non-opposite cut \tilde{P} of $\Delta_{3,n}$ as follows. First, sample $P \sim \mathcal{P}$. Then, randomly select three different indices i_1, i_2, i_3 from $[K]$. Let $P_{\{i_1, i_2, i_3\}} : \Delta_3 \rightarrow [4]$ be the cut induced by P on the face induced by the vertices i_1, i_2, i_3 . Specifically, let $f(i_j) = j$ for every $j \in [3]$. We define $P_{\{i_1, i_2, i_3\}}$ by

$$P_{\{i_1, i_2, i_3\}}(x) = \begin{cases} f(P(x_1 e^{i_1} + x_2 e^{i_2} + x_3 e^{i_3})) & \text{if } P(x_1 e^{i_1} + x_2 e^{i_2} + x_3 e^{i_3}) \in \{i_1, i_2, i_3\}, \\ 4 & \text{otherwise.} \end{cases}$$

If $P_{\{i_1, i_2, i_3\}}$ is a non-opposite cut, then let $\tilde{P} = P_{\{i_1, i_2, i_3\}}$. Otherwise, let $\tilde{P}(x) = P_{\{i_1, i_2, i_3\}}(x)$ for every x such that $P_{\{i_1, i_2, i_3\}}(x) \in \text{supp}(x) \cup \{4\}$. For the other x 's, we simply set $\tilde{P}(x) = 4$; we will call these x 's "bad". It is obvious that \tilde{P} is a non-opposite cut. Moreover, if any two points $x, y \in \Delta_{3,n}$ are separated in \tilde{P} , then either $P_{\{i_1, i_2, i_3\}}(x) \neq P_{\{i_1, i_2, i_3\}}(y)$ or exactly one of x, y is bad. From the definition of $P_{\{i_1, i_2, i_3\}}$, the former implies $P(x_1 e^{i_1} + x_2 e^{i_2} + x_3 e^{i_3}) \neq P(y_1 e^{i_1} + y_2 e^{i_2} + y_3 e^{i_3})$, which happens with probability at most $\tau_K^* \cdot \frac{1}{2}\|x - y\|_1$. Hence, we are left to show that the probability that any point $x \in \Delta_{3,n}$ is bad is at most $O(1/K)$; this immediately yields the intended bound since $\|x - y\|_1 \geq 2/n$ for every $x \neq y \in \Delta_{3,n}$.

Recall that a point x is bad iff $P_{\{i_1, i_2, i_3\}}(x) \in [3] \setminus \text{supp}(x)$, which can only happen when $|\text{supp}(x)| = 2$. Assume without loss of generality that $\text{supp}(x) = \{1, 2\}$, i.e. $x_3 = 0$. We want to bound the probability that $P_{\{i_1, i_2, i_3\}}(x) = 3$. Fix the cut P and the choice of i_1, i_2 . Since $x_3 = 0$, this already determines $P(x_1 e^{i_1} + x_2 e^{i_2} + x_3 e^{i_3})$. From the definition of $P_{\{i_1, i_2, i_3\}}$, $P_{\{i_1, i_2, i_3\}}(x) = 3$ if and only if $i_3 = P(x_1 e^{i_1} + x_2 e^{i_2} + x_3 e^{i_3})$. Since i_3 is a random element of $[K] \setminus \{i_1, i_2\}$, this happens with probability at most $1/(K-2)$ and we have completed our proof sketch.

Lower Bound for $\tilde{\tau}_{3,n}^*$. The final component in our construction is a lower bound for $\tilde{\tau}_{3,n}^*$. To lower bound $\tilde{\tau}_{3,n}^*$, it is enough to construct a weighted undirected graph with vertices $\Delta_{3,n}$ such that the CKR relaxation value (with respect to the solution in which each vertex is assigned the vector representation of itself) is

²The minimum exists since there are only finite number of cuts when restricted to the discretization $\Delta_{k,n}$ of the simplex.

small and every non-opposite cut has large cost. In other words, we need to construct an integrality gap for Multiway Cut with 3 terminals but only against non-opposite cuts, instead of against all 3-way cuts.

We will set the edge set $E_{3,n}$ to be all two consecutive points on grid lines, i.e., $E_{3,n} = \{(x, y) \mid x, y \in \Delta_{3,n}, \|x - y\|_1 = 2/n\}$. Fortunately for us, Karger et al. [KKS⁺04] gave an elegant characterization of 3-way cuts of the graph $(\Delta_{3,n}, E_{3,n})$, which ultimately led to their $12/11 - \varepsilon$ ratio integrality gap for Multiway Cut with 3 terminals. By extending this, we arrive at a similar characterization for all non-opposite cuts of the graph, which also eventually enables us to construct the integrality gap against such cuts.

1.2 Organization of the paper

The rest of the paper is organized as follows. We start by formally defining our notation in Section 2. In Section 3, we prove the lower bound of $6/5 - O(1/n)$ on $\tilde{\tau}_{3,n}$. In Section 4, we use this to construct an integrality gap of ratio $6/5 - O(1/\sqrt{k})$, which essentially contains all of our main ideas. Then, in Section 5, we show how to nail down the dependency on k and achieve the integrality ratio of $6/(5 + \frac{1}{k-1}) - \varepsilon$, completing the proof of our main result. Finally, we provide some discussion regarding our result and open questions in Section 6.

2 Preliminaries and Notation

From now on, our graphs will have vertex sets being $\Delta_{k,n} \subseteq \Delta_k$ for some k, n , and the edge sets being $E_{k,n} := \{(x, y) \mid x, y \in \Delta_{k,n}, \|x - y\|_1 = 2/n\}$. The terminals are naturally associated with the k vertices of the simplex. To avoid confusion, we reserve the word *vertices* for the simplex vertices and we instead refer to the vertices of such graphs as *points*.

We refer to the LP solution in which each point is assigned to itself as the *canonical LP solution* and we use $\mathcal{L}(w)$ to denote its value with respect to a weight function $w : E_{k,n} \rightarrow \mathbb{R}_{\geq 0}$, i.e.,

$$\mathcal{L}(w) = \sum_{(x,y) \in E_{k,n}} w(x, y) \cdot \frac{1}{2} \|x - y\|_1 = \frac{1}{n} \sum_{(x,y) \in E_{k,n}} w(x, y),$$

where the latter equality comes from the fact that $\|x - y\|_1 = 2/n$ for all $(x, y) \in E_{k,n}$. Moreover, for every k -way cut P , we denote its cost with respect to w by $\mathcal{C}(P, w)$; in other words,

$$\mathcal{C}(P, w) = \sum_{(x,y) \in E_{k,n}} w(x, y) \cdot \mathbb{1}[P(x) \neq P(y)],$$

where $\mathbb{1}[P(x) \neq P(y)]$ is the indicator variable of the event $P(x) \neq P(y)$.

3 A Lower Bound on $\tilde{\tau}_{3,n}^*$

In this section, we will construct an integrality gap on the graph $(\Delta_{3,n}, E_{3,n})$ such that its CKR LP value is small but every non-opposite cut has large cost. As our graph is already defined, we only need to define the weights of the edges. The main result of this section can be stated as follows.

Lemma 3 *For every n divisible by 3, there exists $w : E_{3,n} \rightarrow \mathbb{R}_{\geq 0}$ such that*

- *The value of the canonical LP solution of $(\Delta_{3,n}, E_{3,n}, w)$ is at most $5/6 + O(1/n)$.*
- *The cost of every non-opposite cut of $(\Delta_{3,n}, E_{3,n}, w)$ is at least one.*

To prove Lemma 3, we will first characterize candidate optimal non-opposite cuts.

3.1 A Characterization of Non-Opposite Cuts of $\Delta_{3,n}$

To characterize the 3-way cuts, Karger et al. [KKS⁺04] consider the augmented version of $(\Delta_{3,n}, E_{3,n})$, in which each vertex of the simplex has an edge heading out infinitely. They then look at the dual graph of this graph. The augmentation creates three outer faces; the vertices in the dual graph corresponding to these faces are named O_1, O_2 and O_3 where O_1, O_2 and O_3 are opposite to e^1, e^2 and e^3 respectively. For convenience, we disregard the edges among O_1, O_2, O_3 in the dual. Figure 2 contains illustrations of the augmented graph and its planar dual. The figure is a recreation of Figure 1 and Figure 2 in [KKS⁺04].




Figure 2: The augmented $(\Delta_{3,6}, E_{3,6})$ and its dual.

A cut in the original graph can now be viewed as a collection of edges in the dual graph; an edge in the dual graph between two faces corresponding to the shared edge of the faces being cut. With this interpretation, Karger et al. give the following characterization for candidate optimal cuts of any weight function.

Observation 4 ([KKS⁺04]) *For any $w : E_{3,n} \rightarrow \mathbb{R}_{\geq 0}$, there exists a least-cost³ 3-way cut P that is of one of the following forms:*

- P contains three non-intersecting paths from a triangle to O_1, O_2 and O_3 . Such P is called a ball cut.
- P contains two non-intersecting paths among O_1, O_2 and O_3 . Such P is called a 2-corner cut.

Examples of a ball cut and a 2-corner cut can be found in Figure 3a and Figure 3b respectively. These figures are reproduced from Figure 3 and Figure 4 of [KKS⁺04] respectively.

The observation implies that, if one wants to prove that every cut incurs large cost with respect to w , then it is enough to show that all ball cuts and 2-corner cuts have large cost according to w . This is much more convenient to work with and it turns out that finding such w can be formulated as a linear program of size polynomial in n . Once this observation was made, Karger et al. then solved the LP for large n and used the solution as an inspiration to come up with the $12/11 - \varepsilon$ gap for Multiway Cut with three terminals.

We make the following analogous observation for non-opposite cuts. This observation allows us to modify Karger et al.'s LP to express the problem of finding an integrality gap against non-opposite cuts on $(\Delta_{3,n}, E_{3,n})$. We then solve the LP and use it as an inspiration for our explicit gap construction. Since the LP will not appear in our proof, we refer interested readers to Section 3.1 of [KKS⁺04] for more details on the LP.

³Note that it is possible that there are cuts not of these forms that also achieve the same (optimal) cost.




Figure 3: Examples of a ball cut, a 2-corner cut and a 3-corner cut of $(\Delta_{3,6}, E_{3,6})$.

Observation 5 For any $w : E_{3,n} \rightarrow \mathbb{R}_{\geq 0}$, there exists a least-cost non-opposite cut P that is of one of the following forms:

- P is a ball cut.
- P contains three non-intersecting paths among O_1, O_2 and O_3 . Such P is called a 3-corner cut.

An example of a 3-corner cut can be found in Figure 3c. For brevity, we will sometimes abbreviate 3-corner cuts as simply *corner cuts*. Note that we reserve this abbreviation only for 3-corner cuts and we will never shorten 2-corner cuts. The proof of our observation is quite simple and is given below.

Proof of Observation 5. Let P be any non-opposite cut of $(\Delta_{3,n}, E_{3,n})$ (not necessary a ball cut or a corner cut). To prove Observation 5, we will provide a sequence of transformations of P such that all uncut edges in P remain uncut and that, at the end of the transformation, P becomes either a ball cut or a corner cut.

For convenience, let G_P be the graph $(\Delta_{3,n}, E_{3,n})$ after all the edges cut by P are removed. Observe that every point in a connected component of G_P is assigned by P to the same cluster. The transformations can be described as follows:

- For each connected component C of G_P that is assigned to 4, if C intersects with all the three sides of the triangle, then do nothing. Otherwise, reassign the whole component to either 1, 2 or 3, whichever one still maintains the non-oppositeness of P ; note that since C does not intersect with all three sides of the triangle, this is always possible.
- As long as there exists a component C of G_P that is assigned to $i \in [3]$ but does not contain e^i , reassign C to the same cluster as one of its neighbors. Moreover, it is again easy to check that we can always pick the neighboring cluster such that P remains non-opposite.

Clearly, any edge that was initially uncut remains uncut in the final P . Moreover, since the second transformation reduces the number of connected components of G_P by at least one, the sequence of transformations always halts. Finally, note that, in the resulting P , all points assigned to any $i \in [4]$ form a connected component. Moreover, if any point is assigned to 4, then the corresponding connected component must contain at least one point from each of the three sides. It is now easy to see that, if no point is assigned to 4, then the final cut is a ball cut. On the other hand, if at least one point is assigned to 4, the final cut is a corner cut. \square

3.2 The Integrality Gap

With the characterization of non-opposite cuts in place, we are now ready to describe our integrality gap and prove Lemma 3. We remark here that, although we have so far focused only on Karger et al.'s result, our construction will be more similar to that of Cunningham and Tang [CT99], in that we will partition the simplex into “corner triangles” and “middle hexagon”. We would like to stress however that, even with these similarities, these two gaps are significantly different and nothing carries over from there as a blackbox.

Proof of Lemma 3. The gap can be described as follows. First, we divide the vertex set into *corner triangles* T_1, T_2, T_3 and a *middle hexagon* H as follows. The middle hexagon contains all points $x \in \Delta_{3,n}$ such that all the coordinates of x do not exceed $2/3$, whereas the corner triangle T_i contains all the points x 's such that $x_i \geq 2/3$. Note that this is not a partition since the middle hexagon H and each corner triangle T_i still share the line $x_i = 2/3$, but this notation will be more convenient for us.

Every edge in the middle hexagon including its border is assigned weight $\rho := 1/(2n)$. Moreover, for each i , every non-border edge in T_i that is not parallel to the opposite side of e^i is also assigned weight ρ , whereas the non-border edges parallel to the opposite side of e^i are assigned weight zero. Finally, for each of the two borders of T_i containing e^i , we define its weight as follows. The edge closest to e^i is assigned weight $(n/3)\rho$, the second closest is assigned $(n/3 - 1)\rho$, and so on. An illustration of the construction is shown in Figure 4.




Figure 4: The constructed integrality gap when $n = 9$. The middle hexagon H is the shaded area. The weights of the edges on the borders of the triangle are shown. For the non-border edges, the dashed edges have weight zero whereas the rest of the edges all have weight ρ .

Let us now compute the value of the canonical LP solution with respect to w . Recall that the value $\mathcal{L}(w)$ is simply $\frac{1}{n} \sum_{(x,y) \in E_{3,n}} w(x,y)$. Note that due to symmetry the sum of all the weights is three times the sum of the weights of all $(x,y) \in E_{3,n}$ that are parallel to (e^1, e^2) . The total weight on the line (e^1, e^2) is

$$((n/3)\rho + (n/3 - 1)\rho + \cdots + \rho) + (\rho + \cdots + \rho) + (\rho + \cdots + (n/3)\rho) = (n^2/9 + 2n/3)\rho$$

and the total weight on the non-border edges parallel to (e^1, e^2) is simply

$$(n-1)\rho + \dots + (n/3)\rho = (4n^2/9 - n/3)\rho.$$

Hence, the total weight of all edges parallel to (e^1, e^2) is $(5n^2/9 + n/3)\rho$. As a result, we have

$$\mathcal{L}(w) = \frac{1}{n} \cdot 3 \cdot \left(\frac{5n^2}{9} + \frac{n}{3} \right) \cdot \frac{1}{2n} = \frac{5}{6} + O(1/n),$$

as desired.

We will next prove that, for any non-opposite cut P , $\mathcal{C}(P, w) \geq 1$. From Observation 5, we can assume that P is either a ball cut or a corner cut. We will prove that $\mathcal{C}(P, w) \geq 1$ by a potential function argument. Recall that each node in the dual graph is either O_1, O_2, O_3 or a triangle. We represent each triangle by its middle point (median). For each i , we define the potential function Φ_i on O_i and all the triangles as follows.

$$\Phi_i(F) = \begin{cases} 0 & \text{if } F = O_i, \\ (4n/3)\rho & \text{if } x \in T_i, \\ (n/3 + n(x_i - x_w))\rho & \text{if } x \in T_j, \\ (n/3 + n(x_i - x_j))\rho & \text{if } x \in T_w, \\ \lceil 2nx_i \rceil \rho & \text{if } x \in H, \end{cases}$$

where $\{i, j, w\} = \{1, 2, 3\}$ and $x = (x_1, x_2, x_3)$ is the middle point of F . An example illustrating the potential function can be found in Figure 5.




Figure 5: The potential function Φ_1 when $n = 9$.

It is easy to check that, for any two triangles F_1, F_2 that share an edge, the difference $|\Phi_i(F_1) - \Phi_i(F_2)|$ is no more than the weight of the shared edge. The same remains true even when one of F_1, F_2 is O_i . This implies that, in the dual graph, the shortest path from O_i to any triangle F is at least $\Phi_i(F)$. With these

observations in mind, we are ready to show that $\mathcal{C}(P, w) \geq 1$. Let us consider the two cases based on whether P is a corner cut or a ball cut.

Case 1. Suppose that P is a corner cut. Observe that, for any $j \in [3] \setminus \{i\}$ and for any triangle F sharing an edge with the outer face O_j , we always have $\Phi_i(F)$ plus the weight of the shared edge being at least $(2n/3)\rho = 1/3$. Hence, the shortest path from O_i to O_j in the dual graph has weight at least $1/3$. Since P contains three paths among O_1, O_2, O_3 , the cost of P is at least 1.

Case 2. Suppose that P is a ball cut. Let F be the triangle which the three paths to O_1, O_2, O_3 in P originate from. To show that P has cost at least one, it is enough to show that the total length of the shortest paths from F to O_1, O_2 and O_3 is at least one. Since these shortest paths are lower bounded by the potential functions, we only need to prove that $\Phi_1(F) + \Phi_2(F) + \Phi_3(F) \geq 1$.

Let x be the middle point of F . If $x \in H$, then

$$\Phi_1(F) + \Phi_2(F) + \Phi_3(F) = \lceil 2nx_1 \rceil \rho + \lceil 2nx_2 \rceil \rho + \lceil 2nx_3 \rceil \rho \geq (2nx_1)\rho + (2nx_2)\rho + (2nx_3)\rho = 1.$$

On the other hand, if $x \notin H$, assume without loss of generality that $x \in T_1$. We have

$$\Phi_1(F) + \Phi_2(F) + \Phi_3(F) = (4n/3)\rho + (n/3 + n(x_2 - x_3))\rho + (n/3 + n(x_3 - x_2))\rho = (2n)\rho = 1.$$

In all cases, we have $\mathcal{C}(P, w) \geq 1$, thus completing the proof of Lemma 3. \square

4 An Integrality Gap of $6/5 - O(1/\sqrt{k})$

We will now use the integrality gap from Lemma 3 to construct an integrality gap for Multiway Cut with k terminals, for $k \geq 3$. The vertex set and the edge set of the graph will be $\Delta_{k,n} \subseteq \Delta_k$ and $E_{k,n} = \{(x, y) \in \Delta_{k,n} \times \Delta_{k,n} \mid \|x - y\|_1 = 2/n\}$ respectively. The goal of this section is to prove the following lemma.

Lemma 6 *For every n divisible by 3 and every $k \geq 3$, there exists $\bar{w} : E_{k,n} \rightarrow \mathbb{R}_{\geq 0}$ such that*

- *The value of the canonical LP solution of $(\Delta_{k,n}, E_{k,n}, \bar{w})$ is at most $5/6 + O(1/n)$.*
- *The cost of every k -way cut of the graph is at least $1 - O(n/k)$.*

By setting $n = \sqrt{k}$, we have an integrality gap of $6/5 - O(1/\sqrt{k})$. While the dependency on k here is still worse than that in Theorem 1, the construction and proof encapsulate the main ideas of this paper. The gap from this section will finally be slightly tweaked in the next section to yield the right dependency on k .

The core of the proof of Lemma 6 is the following proposition, which is essentially equivalent to $\tilde{\tau}_{3,n}^* - O(n/k) \leq \tau_k^*$, albeit stated in more convenient integrality gap terminologies.

Proposition 7 *Let P be any k -way cut of $\Delta_{k,n}$ for any $k \geq 3$. For any $\{i_1, i_2, i_3\} \subseteq [k]$, let $P_{\{i_1, i_2, i_3\}}$ be as defined in the introduction. If i_1, i_2, i_3 are 3 randomly selected distinct elements of $[k]$, then*

$$\Pr_{i_1, i_2, i_3} [P_{\{i_1, i_2, i_3\}} \text{ is non-opposite}] \geq 1 - O(n/k).$$

Before we prove the proposition, let us first use it to prove Lemma 6.

Proof of Lemma 6. Our gap is simply created by embedding the gap from Section 3 to every triangular face of $\Delta_{k,n}$. More specifically, for each $\{i_1, i_2, i_3\} \in \binom{[k]}{3}$, the weights $\bar{w}_{\{i_1, i_2, i_3\}} : E_{k,n} \rightarrow \mathbb{R}_{\geq 0}$ of the triangular integrality gap embedded to the face induced by $e^{i_1}, e^{i_2}, e^{i_3}$ can be defined as

$$\bar{w}_{\{i_1, i_2, i_3\}}(x, y) = \begin{cases} w((x_{i_1}, x_{i_2}, x_{i_3}), (y_{i_1}, y_{i_2}, y_{i_3})) & \text{if } \text{supp}(x), \text{supp}(y) \subseteq \{i_1, i_2, i_3\}, \\ 0 & \text{otherwise,} \end{cases}$$

where w is the weight function from Lemma 3. Finally, the overall weighting scheme is just the average over all $\bar{w}_{\{i_1, i_2, i_3\}}$'s, i.e.,

$$\bar{w} = \mathbb{E}_{\{i_1, i_2, i_3\} \in \binom{[k]}{3}} [\bar{w}_{\{i_1, i_2, i_3\}}].$$

Let us now compute the value of the canonical LP solution of $(\Delta_{k,n}, E_{k,n}, \bar{w})$. From the definition of \bar{w} , $\mathcal{L}(\bar{w})$ is simply the average of $\mathcal{L}(\bar{w}_{\{i_1, i_2, i_3\}})$ over all $\{i_1, i_2, i_3\} \in \binom{[k]}{3}$. Since $\mathcal{L}(\bar{w}_{\{i_1, i_2, i_3\}}) = \mathcal{L}(w) = 5/6 + O(1/n)$, we can conclude that $\mathcal{L}(\bar{w}) = 5/6 + O(1/n)$ as desired.

Next, we will argue that the cost of every k -way cut of the graph is at least $1 - O(n/k)$. Consider any k -way cut P . Before we formally prove that $\mathcal{C}(P, \bar{w}) \geq 1 - O(n/k)$, let us give a short intuition behind the proof. First, by Proposition 7, for at least $1 - O(n/k)$ fraction of $\{i_1, i_2, i_3\} \in \binom{[k]}{3}$, $P_{\{i_1, i_2, i_3\}}$ is non-opposite. From the second condition of Lemma 3, this means that the cost of P with respect to $\bar{w}_{\{i_1, i_2, i_3\}}$ is at least one. This implies that the total cost of P with respect to \bar{w} is at least $1 - O(n/k)$.

More formally, the cost of P can be written as

$$\begin{aligned} \mathcal{C}(P, \bar{w}) &= \mathbb{E}_{\{i_1, i_2, i_3\} \in \binom{[k]}{3}} [\mathcal{C}(P, \bar{w}_{\{i_1, i_2, i_3\}})] \\ &\text{(from definitions of } \bar{w}_{\{i_1, i_2, i_3\}}, P_{\{i_1, i_2, i_3\}}) \geq \mathbb{E}_{\{i_1, i_2, i_3\} \in \binom{[k]}{3}} [\mathcal{C}(P_{\{i_1, i_2, i_3\}}, w)]. \end{aligned}$$

Let $E_{\{i_1, i_2, i_3\}}$ be the event that $P_{\{i_1, i_2, i_3\}}$ is non-opposite. We can then lower bound the last expression by

$$\begin{aligned} &\mathbb{E}_{\{i_1, i_2, i_3\} \in \binom{[k]}{3}} [\mathcal{C}(P_{\{i_1, i_2, i_3\}}, w)] \geq \Pr[E_{\{i_1, i_2, i_3\}}] \cdot \mathbb{E}_{\{i_1, i_2, i_3\} \in \binom{[k]}{3}} [\mathcal{C}(P_{\{i_1, i_2, i_3\}}, w) \mid E_{\{i_1, i_2, i_3\}}] \\ &\text{(from Proposition 7)} \geq (1 - O(n/k)) \cdot \mathbb{E}_{\{i_1, i_2, i_3\} \in \binom{[k]}{3}} [\mathcal{C}(P_{\{i_1, i_2, i_3\}}, w) \mid E_{\{i_1, i_2, i_3\}}] \\ &\geq 1 - O(n/k), \end{aligned}$$

where the last inequality comes from the second property of w in Lemma 3. \square

We will next prove Proposition 7; the proof is essentially the same as the proof of $\tilde{\tau}_{3,n}^* - O(n/k) \leq \tau_k^*$ sketched earlier in the introduction.

Proof of Proposition 7. By definition of non-opposite cuts of $\Delta_{3,n}$, if $P_{\{i_1, i_2, i_3\}}$ is not non-opposite, then there exists a point x on one of the sides of $\Delta_{3,n}$ (i.e. $|\text{supp}(x)| = 2$) such that $P_{\{i_1, i_2, i_3\}}(x) \notin \text{supp}(x) \cup \{4\}$. Hence, by union bound, we have

$$\Pr_{i_1, i_2, i_3} [P_{\{i_1, i_2, i_3\}} \text{ is not non-opposite}] \leq \sum_{\substack{x \in \Delta_{3,n}: \\ |\text{supp}(x)|=2}} \Pr_{i_1, i_2, i_3} [P_{\{i_1, i_2, i_3\}}(x) \notin \text{supp}(x) \cup \{4\}].$$

Consider an $x \in \Delta_{3,n}$ with $|\text{supp}(x)| = 2$. Assume without loss of generality that $\text{supp}(x) = \{1, 2\}$. If $P_{\{i_1, i_2, i_3\}}(x) \notin \text{supp}(x) \cup \{4\}$, then $P_{\{i_1, i_2, i_3\}}(x) = 3$. The latter occurs if and only if $P(x_1 e^{i_1} + x_2 e^{i_2} + x_3 e^{i_3}) = i_3$. Fix i_1, i_2 ; this already determines $P(x_1 e^{i_1} + x_2 e^{i_2} + x_3 e^{i_3})$. Since i_3 is now a random element from $[k] \setminus \{i_1, i_2\}$, the probability that $P(x_1 e^{i_1} + x_2 e^{i_2} + x_3 e^{i_3}) = i_3$ is at most $1/(k-2)$. Hence, we have

$$\Pr_{i_1, i_2, i_3} [P_{\{i_1, i_2, i_3\}} \text{ is not non-opposite}] \leq \sum_{\substack{x \in \Delta_{3,n}: \\ |\text{supp}(x)|=2}} \frac{1}{k-2} = \frac{3(n-1)}{k-2} = O(n/k),$$

completing the proof of the proposition. \square

5 Getting the Right Dependency on k

We will modify our gap from the previous section to arrive at the following quantitatively better bound.

Theorem 8 *For every n divisible by 3 and every $k \geq 3$, there exists $\tilde{w} : E_{k,n} \rightarrow \mathbb{R}_{\geq 0}$ such that*

- *The value of the canonical LP solution of $(\Delta_{k,n}, E_{k,n}, \tilde{w})$ is at most $(5 + \frac{1}{k-1})/6 + O(1/n)$.*
- *The cost of every k -way cut of the graph is at least one.*

Clearly, by taking n sufficiently large (i.e. $n = \Omega(1/\varepsilon)$), Theorem 8 implies Theorem 1 since the value of the canonical LP solution is no more than the value of the optimal LP solution.

The rest of the section is arranged as follows. First, in the first subsection, we refine our analysis from Section 4 to get improved bounds for the costs of k -way cuts against the weight function. Unfortunately, this will not be enough to prove Theorem 8 and we will have to provide another simple gap to complement worst cases for the gap from Section 4; this is done in Subsection 5.2. Finally, we take an appropriate linear combination of the two gaps and prove Theorem 8 in the last subsection.

5.1 Refined Analysis for the Gap from Section 4

There are two main points of the analysis from the previous section that need to be tightened in order to get the bound in Theorem 8. First, notice that the bound in Proposition 7 can be improved if, for each two vertices e^i, e^j of the simplex, the points on the line between e^i and e^j are assigned to few clusters. For instance, in the extreme case where every point on the line from e^i to e^j is assigned to either cluster i or cluster j , $P_{\{i_1, i_2, i_3\}}$ is always non-opposite for every i_1, i_2, i_3 . The following proposition gives a more refined bound on the probability that $P_{\{i_1, i_2, i_3\}}$ is non-opposite based on the number of clusters that the points on the line from e^i to e^j are assigned to for every $i, j \in [k]$.

Proposition 9 *Let P be any k -way cut of $\Delta_{k,n}$ for any $k \geq 3$. For each $\{i, j\} \subseteq [k]$, let us define $D_{i,j}(P)$ to be the set of all clusters that the points on the line between e^i and e^j (inclusive) are assigned to, i.e.,*

$$D_{i,j}(P) = \{P(x) \mid x \in \Delta_{k,n}, \text{supp}(x) \subseteq \{i, j\}\}.$$

Finally, let $D(P)$ be the average of the size of $D_{i,j}(P)$ among all $\{i, j\} \subseteq [k]$, i.e., $D(P) = \mathbb{E}_{\{i,j\} \subseteq [k]} [|D_{i,j}(P)|]$. If i_1, i_2, i_3 are three randomly selected distinct elements of $[k]$, then

$$\Pr_{i_1, i_2, i_3} [P_{\{i_1, i_2, i_3\}} \text{ is non-opposite}] \geq 1 - 3(D(P) - 2)/(k - 2).$$

Proof. We want to bound the probability that $P_{\{i_1, i_2, i_3\}}$ is not non-opposite. Again, this happens only when there exists x such that $\text{supp}(x) \subseteq \{i_1, i_2, i_3\}$ and $P(x) \in \{i_1, i_2, i_3\} \setminus \text{supp}(x)$. Note that such x can only have $|\text{supp}(x)| = 2$. Thus, by union bound, we have

$$\begin{aligned} & \Pr_{i_1, i_2, i_3} [P_{\{i_1, i_2, i_3\}} \text{ is not non-opposite}] \\ & \leq \sum_{\{j_1, j_2\} \subseteq \{i_1, i_2, i_3\}} \Pr_{i_1, i_2, i_3} [\exists x, \text{supp}(x) = \{j_1, j_2\} \text{ and } P(x) \in \{i_1, i_2, i_3\} \setminus \{j_1, j_2\}] \\ & = 3 \Pr_{i_1, i_2, i_3} [\exists x, \text{supp}(x) = \{i_1, i_2\} \text{ and } P(x) = i_3], \end{aligned}$$

where the last equality comes from symmetry. Now, let us fix i_1, i_2 . There is an x with $\text{supp}(x) = \{i_1, i_2\}$ such that $P(x) = i_3$ if and only if $i_3 \in D_{i_1, i_2}(P)$. Observe that $i_1, i_2 \in D_{i_1, i_2}(P)$. Hence, the probability that i_3 , a uniformly random element from $[k] \setminus \{i_1, i_2\}$, lies in $D_{i_1, i_2}(P)$ is exactly $(|D_{i_1, i_2}(P)| - 2)/(k - 2)$. Thus,

$$\Pr_{i_1, i_2, i_3} [\exists x, \text{supp}(x) = \{i_1, i_2\} \text{ and } P(x) = i_3] = \mathbb{E}_{i_1, i_2} [\Pr_{i_3} [\exists x, \text{supp}(x) = \{i_1, i_2\} \text{ and } P(x) = i_3]]$$

$$\begin{aligned}
&= \mathbb{E}_{i_1, i_2} [|D_{i_1, i_2}(P)| - 2] / (k - 2) \\
&= (D(P) - 2) / (k - 2).
\end{aligned}$$

Hence, $\Pr_{i_1, i_2, i_3} [P_{\{i_1, i_2, i_3\}} \text{ is non-opposite}] \geq 1 - 3(D(P) - 2)/(k - 2)$ as desired. \square

Another point in our analysis that was not tight was that, in the proof of Lemma 6, we only used the fact that $\mathcal{C}(P, \bar{w}_{\{i_1, i_2, i_3\}})$ is at least one when $P_{\{i_1, i_2, i_3\}}$ is non-opposite. However, even when $P_{\{i_1, i_2, i_3\}}$ is not non-opposite, $\mathcal{C}(P, \bar{w}_{\{i_1, i_2, i_3\}})$ is non-zero and we should include this in our argument, too. To do so, we need the following observation.

Observation 10 *Let w be as defined in Lemma 3. The cost of any 3-way cut of $(\Delta_{3,n}, E_{3,n})$ with respect to w is at least $2/3$.*

To see that the observation is true, recall Observation 4; to prove that any 3-way cut of $(\Delta_{3,n}, E_{3,n})$ has cost at least $2/3$, we only need to consider ball cuts and 2-corner cuts. However, we have shown in Subsection 3.2 that the cost of any ball cut and 3-corner cut with respect to w is at least one. This immediately implies that the cost of any 2-corner cut is at least $2/3$, meaning that Observation 10 is indeed true.

With the above proposition and the observation, we can now prove a refined version of Lemma 6, as stated formally below. The proof is very similar to that of the original lemma.

Lemma 11 *For \bar{w} as defined in Lemma 6 and for any k -way cut P , we have $\mathcal{C}(P, \bar{w}) \geq 1 - (D(P) - 2)/(k - 2)$.*

Proof. Recall that we have

$$\mathcal{C}(P, \bar{w}) \geq \mathbb{E}_{\{i_1, i_2, i_3\} \subseteq [k]} [\mathcal{C}(P_{\{i_1, i_2, i_3\}}, w)].$$

Again, let $E_{\{i_1, i_2, i_3\}}$ be the event that $P_{\{i_1, i_2, i_3\}}$ is non-opposite. We can then lower bound $\mathcal{C}(P, \bar{w})$ as follows.

$$\begin{aligned}
\mathbb{E}_{\{i_1, i_2, i_3\} \subseteq [k]} [\mathcal{C}(P_{\{i_1, i_2, i_3\}}, w)] &= \Pr[E_{\{i_1, i_2, i_3\}}] \cdot \mathbb{E}_{\{i_1, i_2, i_3\} \subseteq [k]} [\mathcal{C}(P_{\{i_1, i_2, i_3\}}, w) \mid E_{\{i_1, i_2, i_3\}}] + \\
&\quad \Pr[\neg E_{\{i_1, i_2, i_3\}}] \cdot \mathbb{E}_{\{i_1, i_2, i_3\} \subseteq [k]} [\mathcal{C}(P_{\{i_1, i_2, i_3\}}, w) \mid \neg E_{\{i_1, i_2, i_3\}}] \\
&\stackrel{\text{(from Lemma 3 and Observation 10)}}{\geq} \Pr[E_{\{i_1, i_2, i_3\}}] + \frac{2}{3} \cdot \Pr[\neg E_{\{i_1, i_2, i_3\}}] \\
&= 1 - \frac{1}{3} \cdot \Pr[\neg E_{\{i_1, i_2, i_3\}}] \\
&\stackrel{\text{(from Proposition 9)}}{\geq} 1 - \frac{D(P) - 2}{k - 2}.
\end{aligned}$$

\square

5.2 An Integrality Gap Against Large $D(P)$

As seen in the previous subsection, the weight \bar{w} works well against cuts P with small $D(P)$ but it does not work well against P with large $D(P)$. As a result, we need to construct a different integrality gap instance that deals with this case. Formally, the new gap will have the following properties.

Lemma 12 *For every $n \geq 2$ and every $k \geq 2$, there exists $w' : E_{k,n} \rightarrow \mathbb{R}_{\geq 0}$ such that*

- *The value of the canonical LP solution of $(\Delta_{k,n}, E_{k,n}, w')$ is exactly one.*
- *The cost of every k -way cut P of the graph is at least $D(P) - 1$.*

Proof. The gap construction is intuitive; since $D(P)$ being large means that many edges on the lines between two vertices of the simplex are cut, we simply define the gap so that we place all the weights equally on such edges. Specifically, define $w'(x, y)$ for all $(x, y) \in E_{k,n}$ as follows.

$$w'(x, y) = \begin{cases} \frac{1}{\binom{k}{2}} & \text{if } |\text{supp}(x) \cup \text{supp}(y)| = 2, \\ 0 & \text{otherwise.} \end{cases}$$

Clearly, $\mathcal{L}(w') = 1$ since each line from e^i to e^j contributes $1/\binom{k}{2}$ to the value of the solution.

Now, consider any k -way cut P . The cost of the cut can be written as

$$\begin{aligned} \mathcal{C}(P, w') &= \sum_{(x, y) \in E_{k,n}} w'(x, y) \cdot \mathbb{1}[P(x) \neq P(y)] \\ &= \sum_{\substack{(x, y) \in E_{k,n}: \\ |\text{supp}(x) \cup \text{supp}(y)| = 2}} \frac{1}{\binom{k}{2}} \cdot \mathbb{1}[P(x) \neq P(y)] \\ &= \frac{1}{\binom{k}{2}} \sum_{\{i, j\} \subseteq [k]} \sum_{\substack{(x, y) \in E_{k,n}: \\ \text{supp}(x) \cup \text{supp}(y) = \{i, j\}}} \mathbb{1}[P(x) \neq P(y)] \\ &= \mathbb{E}_{\{i, j\} \subseteq [k]} \left[\sum_{\substack{(x, y) \in E_{k,n}: \\ \text{supp}(x) \cup \text{supp}(y) = \{i, j\}}} \mathbb{1}[P(x) \neq P(y)] \right]. \end{aligned}$$

The term in the expectation is simply the number of edges on the line between e^i and e^j cut by P . Since we know that the points on this line are assigned to $|D_{i,j}(P)|$ different clusters, we can lower bound the number of edges cut by $|D_{i,j}(P)| - 1$, which gives

$$\mathcal{C}(P, w') \geq \mathbb{E}_{\{i, j\} \subseteq [k]} [|D_{i,j}(P)| - 1] = D(P) - 1.$$

This completes the proof of Lemma 12. □

5.3 Putting Things Together

It is now easy to prove our main theorem by simply picking the weight \tilde{w} to be an appropriate linear combination of \bar{w} and w' .

Proof of Theorem 8. Let \bar{w} and w' be as defined in Lemma 6 and Lemma 12 respectively. We define \tilde{w} to be

$$\tilde{w} = \left(\frac{k-2}{k-1} \right) \bar{w} + \left(\frac{1}{k-1} \right) w'.$$

The value of the canonical LP solution of $(\Delta_{k,n}, E_{k,n}, \tilde{w})$ is simply

$$\begin{aligned} \mathcal{L}(\tilde{w}) &= \left(\frac{k-2}{k-1} \right) \mathcal{L}(\bar{w}) + \left(\frac{1}{k-1} \right) \mathcal{L}(w') \\ &= \left(\frac{k-2}{k-1} \right) (5/6 + O(1/n)) + \left(\frac{1}{k-1} \right) \\ &= \left(5 + \frac{1}{k-1} \right) / 6 + O(1/n). \end{aligned}$$

Moreover, for any k -way cut P , we have

$$\begin{aligned}\mathcal{C}(P, \tilde{w}) &= \left(\frac{k-2}{k-1}\right) \mathcal{C}(P, \overline{w}) + \left(\frac{1}{k-1}\right) \mathcal{C}(P, w') \\ (\text{from Lemma 11 and Lemma 12}) &\geq \left(\frac{k-2}{k-1}\right) \left(1 - \frac{D(P)-2}{k-2}\right) + \left(\frac{1}{k-1}\right) (D(P)-1) \\ &= 1.\end{aligned}$$

Hence, \overline{w} satisfies the two properties in Theorem 8. \square

6 Conclusion and Discussions

In this paper, we construct integrality gap instances of ratio $6/(5 + \frac{1}{k-1}) - \varepsilon$ for every $\varepsilon > 0$ for the Călinescu-Karloff-Rabani linear program relaxation of Multiway Cut with $k \geq 3$ terminals. Combined with the result of Manokaran et al. [MNRS08], our result implies that it is UG-hard to approximate Multiway Cut with k terminals to within $6/(5 + \frac{1}{k-1}) - \varepsilon$ ratio of the optimum. Our gap construction is based on the observation that $\tilde{\tau}_{3,n}^* - O(n/K)$ lower bounds τ_K^* ; with this observation, we extend Karger et al.’s characterization of 3-way cuts [KKS⁺04] to non-opposite cuts, which ultimately leads to a lower bound of $6/5 - O(1/n)$ for $\tilde{\tau}_{3,n}^*$. We also observe that Freund and Karloff’s integrality gap instance [FK00] is subsumed by our approach.

Even with our result, the approximability of Multiway Cut is far from resolved as the best known approximation algorithm by Sharma and Vondrák [SV14] only achieves approximation ratio of 1.2965 as k goes to infinity whereas our gap is only 1.2. It may be tempting to try to prove a stronger lower bound for $\tilde{\tau}_{3,n}^*$ in order to construct improved integrality gaps. Unfortunately, this is impossible since a slight modification of Karger et al.’s algorithm for Multiway Cut with 3 terminals shows that $\tilde{\tau}_3^* \leq 1.2$. For completeness, we include the modified algorithm in Appendix C. With this in mind, it is likely that, to construct better integrality gaps, we need to exploit geometric properties of higher-dimensional simplexes. This seems much harder to do as all the known gaps, including ours, only deal with the two-dimensional simplex Δ_3 , whose nice properties allow the cuts to be characterized in a convenient manner.

Finally, we would like to note that the proof in Section 5 can easily be converted to a proof of $\tilde{\tau}_3^* - O(1/K) \leq \tau_K^*$. Notice that $\tilde{\tau}_3^*$ appears in the inequality as opposed to its discretized approximation $\tilde{\tau}_{3,n}^*$ as sketched in the introduction. Unfortunately, we do not know of any way of extending this proof to show a direct relation between $\tilde{\tau}_k^*$ and τ_K^* for $k > 3$. Such relation may be crucial if one wants to construct an integrality gap instance by proving a lower bound for $\tilde{\tau}_k^*$.

References

- [BNS13] Niv Buchbinder, Joseph Naor, and Roy Schwartz. Simplex partitioning via exponential clocks and the multiway cut problem. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC ’13, pages 535–544, New York, NY, USA, 2013. ACM. 1, 2
- [BSW17] Niv Buchbinder, Roy Schwartz, and Baruch Weizman. Simplex transformations and the multiway cut problem. In *Proceedings of the Twenty-eight Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA ’17, 2017. To appear. 2
- [BTV99] Dimitris Bertsimas, Chung-Piaw Teo, and Rakesh Vohra. Analysis of lp relaxations for multiway and multicut problems. *Networks*, 34(2):102–114, 1999. 1
- [CKR00] Grigori Călinescu, Howard J. Karloff, and Yuval Rabani. An improved approximation algorithm for multiway cut. *Journal of Computer and System Sciences*, 60(3):564–574, 2000. 1, 2, 3
- [CR91] Sunil Chopra and M. R. Rao. On the multiway cut polyhedron. *Networks*, 21(1):51–89, 1991. 1

- [CT99] William H. Cunningham and Lawrence Tang. Optimal 3-terminal cuts and linear programming. In *Integer Programming and Combinatorial Optimization, 7th International IPCO Conference, Graz, Austria, June 9-11, 1999, Proceedings*, pages 114–125, 1999. [1](#), [2](#), [8](#)
- [Cun89] William H. Cunningham. The optimal multiterminal cut problem. In *Reliability Of Computer And Communication Networks, Proceedings of a DIMACS Workshop, New Brunswick, New Jersey, USA, December 2-4, 1989*, pages 105–120, 1989. [1](#)
- [DJP⁺83] Elias Dahlhaus, David S. Johnson, Christos H. Papadimitriou, Paul D. Seymour, and Mihalis Yannakakis. The complexity of multiway cuts (extended abstract). 1983. [1](#)
- [DJP⁺94] Elias Dahlhaus, David S. Johnson, Christos H. Papadimitriou, Paul D. Seymour, and Mihalis Yannakakis. The complexity of multiterminal cuts. *SIAM Journal on Computing*, 23(4):864–894, August 1994. [1](#)
- [FK00] Ari Freund and Howard Karloff. A lower bound of $8/(7 + \frac{1}{k-1})$ on the integrality ratio of the Călinescu-Karloff-Rabani relaxation for multiway cut. *Information Processing Letters*, 75(1-2):43–50, July 2000. [1](#), [2](#), [4](#), [15](#), [16](#)
- [KKS⁺04] David R. Karger, Philip N. Klein, Clifford Stein, Mikkell Thorup, and Neal E. Young. Rounding algorithms for a geometric embedding of minimum multiway cut. *Mathematics of Operations Research*, 29(3):436–461, 2004. [1](#), [2](#), [3](#), [5](#), [6](#), [15](#), [17](#), [18](#)
- [MNRS08] Rajsekar Manokaran, Joseph Naor, Prasad Raghavendra, and Roy Schwartz. SDP gaps and UGC hardness for multiway cut, 0-extension, and metric labeling. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing, STOC '08*, pages 11–20, New York, NY, USA, 2008. ACM. [1](#), [2](#), [15](#)
- [SV14] Ankit Sharma and Jan Vondrák. Multiway cut, pairwise realizable distributions, and descending thresholds. In *Proceedings of the Forty-sixth Annual ACM Symposium on Theory of Computing, STOC '14*, pages 724–733, New York, NY, USA, 2014. ACM. [1](#), [2](#), [15](#)

A The Freund-Karloff Gap, Revisited

In this section, we will rephrase the Freund-Karloff integrality gap [\[FK00\]](#) in terms of a lower bound for $\tilde{\tau}_{3,2}^*$. In $\Delta_{3,2}$, there are only six vertices: the three vertices of the simplex and the three middle points. The Freund-Karloff gap sets the weight of the edges between a simplex vertex and a middle point to be $1/6$ and sets the weight of the edges among the middle points to be $1/4$.




Figure 6: The Freund-Karloff Integrality Gap and Candidate Optimal Non-Opposite Cuts. Figure [6a](#) demonstrates the Freund-Karloff integrality gap whereas Figure [6b](#) shows the only three possible ball and corner cuts (up to rotations and reflections).

The illustrations of the gap can be found in Figure 6a. Clearly, the canonical LP solution has value $7/8$. Moreover, discounting symmetry, there are only three ball and corner cuts, as shown in Figure 6b. Obviously, every cut has value at least one. Hence, this is an $8/7$ lower bound on $\tilde{\tau}_{3,2}^*$ as intended.

Finally, we note that, by using the technique from Section 5, one can arrive at $8/(7 + \frac{1}{k-1})$ integrality gap for Multiway Cut with k terminals, which is exactly the result of Freund and Karloff.

B An Inequality Between $\tilde{\tau}_{k,n}^*$ and τ_K^*

In this section, we prove a relation between $\tilde{\tau}_{k,n}^*$ and τ_K^* for $K > k$. This generalizes the inequality between $\tilde{\tau}_{3,n}^*$ and τ_K^* , whose proof sketch was presented in the introduction.

Lemma 13 *For every $K > k$, $\tilde{\tau}_{k,n}^* - O(kn/(K - k)) \leq \tau_K^*$.*

Proof. Suppose that \mathcal{P} is the distribution on K -way cuts such that $\tau_K(\mathcal{P}) = \tau_K^*$. We define a distribution $\tilde{\mathcal{P}}$ on non-opposite cuts of Δ_k by a sampling process for $\tilde{P} \sim \tilde{\mathcal{P}}$ as follows. First, sample $P \sim \mathcal{P}$. Then, randomly select an injection $f : [k] \rightarrow [K]$. In short, we will set \tilde{P} to be the restriction of P on the face induced by $e^{f(1)}, \dots, e^{f(k)}$ and “fix” \tilde{P} so that it is non-opposite.

For convenience, let $x \circ f$ denote $\sum_{i \in [k]} x_i e^{f(i)} \in \Delta_K$ for every $x \in \Delta_k$. \tilde{P} can then be defined as follows.

$$\tilde{P}(x) = \begin{cases} f^{-1}(P(x \circ f)) & \text{if } P(x \circ f) \in \text{supp}(x \circ f), \\ k + 1 & \text{otherwise.} \end{cases}$$

Since $\text{supp}(x \circ f) = f(\text{supp}(x))$, it is clear from the definition above that \tilde{P} is indeed a non-opposite cut of Δ_k . We are left to only argue that $\tau_{k,n}(\tilde{\mathcal{P}}) \leq \tau_K(\mathcal{P}) + O(k/K)$. Consider any two different points $x, y \in \Delta_{k,n}$. Observe that, if $\tilde{P}(x) \neq \tilde{P}(y)$, then either $P(x \circ f) \neq P(y \circ f)$ or $P(x \circ f) = P(y \circ f) \in \text{supp}(x \circ f) \oplus \text{supp}(y \circ f)$ where \oplus denotes the symmetric difference between the two sets. The latter also implies that either $P(x \circ f) \in f([k] \setminus \text{supp}(x))$ or $P(y \circ f) \in f([k] \setminus \text{supp}(y))$. By union bound, we have

$$\begin{aligned} \Pr_{\tilde{P} \sim \tilde{\mathcal{P}}} [\tilde{P}(x) \neq \tilde{P}(y)] &\leq \Pr_{P \sim \mathcal{P}, f} [P(x \circ f) \neq P(y \circ f)] + \Pr_{P \sim \mathcal{P}, f} [P(x \circ f) = P(y \circ f) \in \text{supp}(x \circ f) \oplus \text{supp}(y \circ f)] \\ &\leq \tau_K^* \cdot \frac{1}{2} \|x - y\|_1 + \Pr_{P \sim \mathcal{P}, f} [P(x \circ f) \in f([k] \setminus \text{supp}(x))] + \Pr_{P \sim \mathcal{P}, f} [P(y \circ f) \in f([k] \setminus \text{supp}(y))]. \end{aligned}$$

Let us now bound $\Pr_{P \sim \mathcal{P}, f} [P(x \circ f) \in f([k] \setminus \text{supp}(x))]$. Fix P and fix $g = f|_{\text{supp}(x)}$. Note that fixing g already determines $x \circ f$ and hence $P(x \circ f)$. If $P(x \circ f)$ is in $g(\text{supp}(x)) = f(\text{supp}(x))$, then clearly $P(x \circ f) \notin f([k] \setminus \text{supp}(x))$. On the other hand, if $P(x \circ f) \notin f(\text{supp}(x))$, then $P(x \circ f) \in f([k] \setminus \text{supp}(x))$ if and only if $P(x \circ f) = f(i)$ for some $i \in [k] \setminus \text{supp}(x)$. For a fixed i , the probability that $P(x \circ f) = f(i)$ is only $1/(K - \text{supp}(x)) \leq 1/(K - k)$. Hence, by union bound, $\Pr_{P \sim \mathcal{P}, f} [P(x \circ f) \in f([k] \setminus \text{supp}(x))] \leq k/(K - k)$.

A similar bound can be derived for $\Pr_{P \sim \mathcal{P}, f} [P(y \circ f) \in f([k] \setminus \text{supp}(y))]$, which implies that

$$\Pr_{\tilde{P} \sim \tilde{\mathcal{P}}} [\tilde{P}(x) \neq \tilde{P}(y)] \leq \tau_K^* \cdot \frac{1}{2} \|x - y\|_1 + 2k/(K - k).$$

Finally, note that, for all $x \neq y \in \Delta_{k,n}$, the distance $\|x - y\|_1$ is at least $2/n$. Hence, we have $\Pr_{\tilde{P} \sim \tilde{\mathcal{P}}} [\tilde{P}(x) \neq \tilde{P}(y)] \leq (\tau_K^* + O(kn/(K - k))) \cdot \frac{1}{2} \|x - y\|_1$, meaning that $\tilde{\tau}_k^* - O(kn/(K - k)) \leq \tau_K^*$ as desired. \square

C A Rounding Scheme for Non-Opposite Cuts of Δ_3

In this section, we provide a probability distribution \mathcal{P} on non-opposite cuts of Δ_3 such that $\tau(\mathcal{P}) \leq 6/5$, thereby proving that $\tilde{\tau}_3^* \leq 6/5$. The probability distribution is a slight modification of the algorithm for Multiway Cut with 3 terminals by Karger et al. [KKS⁺04].

A cut $P \sim \mathcal{P}$ is drawn as follows:

- With probability $1/5$, we select $r \in [2/3, 1)$ uniformly at random. Let P be the 3-corner cut in which each $x \in \Delta_3$ is assigned to i if and only if $x_i > r$ and is assigned to 4 if none of its coordinates exceeds r .
- With probability $4/5$, we pick P to be a ball cut by the following process. With probability $1/2$, let r be a random point on the line between $(2/3, 1/3, 0)$ and $(0, 2/3, 1/3)$. Otherwise, let r be a random point on the line between $(2/3, 0, 1/3)$ and $(0, 1/3, 2/3)$. Consider segments starting from r which are parallel to the sides of Δ_3 . There are six such segments with two of them ending on each side of the triangle. For each side, randomly pick one of the two segments. Finally, let P be the cut induced by the three segments picked.

We note that the distribution is the same as that in Theorem 4.3 of [KKS⁺04] except that here the corner cut is chosen with probability $1/5$ instead of $3/11$ as in [KKS⁺04]; for illustrations of the cuts, please refer to Figure 5 in [KKS⁺04]. Since the proof remains essentially the same as that in the mentioned theorem, we do not provide it here.