

Improved Hardness for Cut, Interdiction, and Firefighter Problems^{*†}

Euiwoong Lee

Carnegie Mellon University, Pittsburgh, PA, USA
euiwoon1@cs.cmu.edu

Abstract

We study variants of the classic s - t cut problem and prove the following improved hardness results assuming the Unique Games Conjecture (UGC).

- For *Length-Bounded Cut* and *Shortest Path Interdiction*, we show that both problems are hard to approximate within any constant factor, even if we allow bicriteria approximation. If we want to cut vertices or the graph is directed, our hardness ratio for Length-Bounded Cut matches the best approximation ratio up to a constant. Previously, the best hardness ratio was 1.1377 for Length-Bounded Cut [4] and 2 for Shortest Path Interdiction [24].
- For any constant $k \geq 2$ and $\epsilon > 0$, we show that *Directed Multicut* with k source-sink pairs is hard to approximate within a factor $k - \epsilon$. This matches the trivial k -approximation algorithm. By a simple reduction, our result for $k = 2$ implies that *Directed Multiway Cut* with two terminals (also known as s - t *Bicut*) is hard to approximate within a factor $2 - \epsilon$, matching the trivial 2-approximation algorithm.
- Assuming a variant of the UGC (implied by another variant of Bansal and Khot [6]), we prove that it is hard to approximate *Resource Minimization Fire Containment* within any constant factor. Previously, the best hardness ratio was 2 [28]. For directed layered graphs with b layers, our hardness ratio $\Omega(\log b)$ matches the best approximation algorithm [3, 9].

Our results are based on a general method of *converting an integrality gap instance to a length-control dictatorship test* for variants of the s - t cut problem, which may be useful for other problems.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases length bounded cut, shortest path interdiction, multicut, firefighter, unique games conjecture

Digital Object Identifier 10.4230/LIPIcs.ICALP.2017.92

1 Introduction

One of the most important implications of the Unique Games Conjecture (UGC, [25]) is the results of Khot et al. [26] and Raghavendra [40], which say that for any maximum constraint satisfaction problem (Max-CSP), an integrality gap instance of the standard semidefinite programming (SDP) relaxation can be converted to the NP-hardness result with the same gap. These results initiated the study of beautiful connections between power of convex relaxations and hardness of approximation, from which surprising results for both subjects have been discovered.

* The full version of this paper is available as [30], <https://arxiv.org/abs/1607.05133>.

† Supported by the Samsung Scholarship, the Simons Award for Graduate Students in TCS, and Venkat Guruswami's NSF CCF-1115525.



© Euiwoong Lee;

licensed under Creative Commons License CC-BY

44th International Colloquium on Automata, Languages, and Programming (ICALP 2017).

Editors: Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl;

Article No. 92; pp. 92:1–92:14



Leibniz International Proceedings in Informatics

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



While their results hold for problems in Max-CSPs, the framework of *converting an integrality gap instance to hardness* has been successfully applied to covering and graph cut problems. For graph cut problems, Manokaran et al. [34] showed that for *Undirected Multiway Cut* and its generalizations, an integrality gap of the standard linear programming (LP) relaxation implies the hardness result assuming the UGC. Their result is further generalized by Ene et al. [16] by formulating them as *Min-CSPs*. In addition, Kumar et al. [29] studied *Strict CSPs* and showed the same phenomenon for the standard LP relaxation.

One of the limitations of the previous CSP-based transformations from LP gap instances to hard instances is based on the fact that they do not usually preserve the desired structure of the constraint hypergraph.¹ For example, consider the *Length-Bounded Edge Cut* problem where the input consists of a graph $G = (V, E)$, two vertices $s, t \in V$, and a constant $l \in \mathbb{N}$, and the goal is to remove the fewest edges to ensure there is no path from s to t of length less than l . This problem can be viewed as a special case of *Hypergraph Vertex Cover* (HVC) by viewing each edge as a vertex of a hypergraph and creating a hyperedge for every s - t path of length less than l . While HVC is in turn a Strict CSP, its integrality gap instance cannot be converted to hardness using Kumar et al. [29] as a black-box, since the set of hyperedges created in the resulting hard instance is not guaranteed to correspond to the set of short s - t paths of some graph.

For Undirected Multiway Cut, Manokaran et al. [34] bypassed this difficulty by using 2-ary constraints so that the resulting constraint hypergraph becomes a graph again. For *Undirected Node-weighted Multiway Cut*, Ene et al. [16] used the equivalence to *Hypergraph Multiway Cut* [38] so that the resulting hypergraph does not need to satisfy additional structure. These problems are then formulated as a Min-CSP by using many labels which are supposed to represent different connected components. However, these Min-CSP based techniques often require nontrivial problem-specific ideas and do not seem to be easily generalized to many other cut problems.

We study variants of the classical s - t cut problem in both directed and undirected graphs that have been actively studied. We prove the optimal hardness or the first super-constant hardness for them. See Section 1.1 for the definitions of the problems and our results. All our results are based on the general framework of converting an integrality gap instance to a *length-control dictatorship test*. The structure of our length-control dictatorship tests allows us to naturally convert an integrality gap instance for the basic LP for various cut problems to hardness based on the UGC. Section 1.2 provides more detailed intuition of this framework. We believe that our framework is general and will be useful to prove tight inapproximability of other cut problems.

1.1 Problems and Results

Length-Bounded Cut and Shortest Path Interdiction. The *Length-Bounded Cut* problem is a natural variant of s - t cut, where given a graph (directed or undirected), $s, t \in V$, and an integer l , we only want to cut s - t paths of length strictly less than l .² Its practical motivation is based on the fact that in most communication / transportation networks, short paths are preferred to be used to long paths [32].

Lovász et al. [31] gave an exact algorithm for Length-Bounded Vertex Cut ($l \leq 5$) in undirected graphs. Mahjoub and McCormick [32] proved that Length-Bounded Edge Cut

¹ One of notable exceptions we are aware is the result of Guruswami et al. [21], using Kumar et al. [29] to show that k -Uniform k -Partite Hypergraph Vertex Cover is hard to approximate within a factor $\frac{k}{2} - \epsilon$ for any $\epsilon > 0$.

² It is more conventional to cut s - t paths of length *at most* l . We use this slightly nonconventional way to be more consistent with Shortest Path Interdiction.

admits an exact polynomial time algorithm for $l \leq 4$ in undirected graphs. Baier et al. [4] showed that both Length-Bounded Vertex Cut ($l > 5$) and Length-Bounded Edge Cut ($l > 4$) are NP-hard to approximate within a factor 1.1377. They presented $O(\min(l, \frac{n}{l})) = O(\sqrt{n})$ -approximation algorithm for Length-Bounded Vertex Cut and $O(\min(l, \frac{n^2}{l^2}, \sqrt{m})) = O(n^{2/3})$ -approximation algorithm for Length-Bounded Edge Cut, with matching LP gaps. Length-Bounded Cut problems have been also actively studied in terms of their fixed parameter tractability [19, 15, 8, 17].

If we exchange the roles of the objective k and the length bound l , the problem becomes *Shortest Path Interdiction*, where we want to maximize the length of the shortest s - t path after removing at most k vertices or edges. It is also one of the central problems in a broader class of *interdiction* problems, where an *attacker* tries to remove some edges or vertices to destroy a desirable property (e.g., short s - t distance, large s - t flow, cheap MST) of a network (see the survey of [42]). The study of Shortest Path Interdiction started in 1980's when the problem was called as the *k-most-vital-arcs* problem [14, 33, 5] and proved to be NP-hard [5]. Khachiyan et al. [24] proved that it is NP-hard to approximate within a factor less than 2. While many heuristic algorithms were proposed [23, 7, 35] and hardness in planar graphs [39] was shown, whether the general version admits a constant factor approximation was still unknown.

Given a graph $G = (V, E)$ and $s, t \in V$, let $\text{dist}(G)$ be the length of the shortest s - t path. For $V' \subseteq V$, let $G \setminus V'$ be the subgraph induced by $V \setminus V'$. For $E' \subseteq E$, we use the same notation $G \setminus E'$ to denote the subgraph $(V, E \setminus E')$. We primarily study undirected graphs. We first present our results for the vertex version of both problems (collectively called as *Short Path Vertex Cut* onwards).

► **Theorem 1.** *Assuming the Unique Games Conjecture, for infinitely many values of constant $l \in \mathbb{N}$, the following three tasks are NP-hard: Given an undirected graph $G = (V, E)$ and $s, t \in V$ where there exists $C^* \subseteq V \setminus \{s, t\}$ such that $\text{dist}(G \setminus C^*) \geq l$,*

1. *Find $C \subseteq V \setminus \{s, t\}$ such that $|C| \leq \Omega(l) \cdot |C^*|$ and $\text{dist}(G \setminus C) \geq l$.*
2. *Find $C \subseteq V \setminus \{s, t\}$ such that $|C| \leq |C^*|$ and $\text{dist}(G \setminus C) \geq O(\sqrt{l})$.*
3. *Find $C \subseteq V \setminus \{s, t\}$ such that $|C| \leq \Omega(l^{\frac{\epsilon}{2}}) \cdot |C^*|$ and $\text{dist}(G \setminus C) \geq O(l^{\frac{1+\epsilon}{2}})$ for some $0 < \epsilon < 1$.*

The first result shows that Length Bounded Vertex Cut is hard to approximate within a factor $\Omega(l)$. This matches the best $O(l)$ -approximation [4] when l is a constant. The second result shows that Shortest Path Vertex Interdiction is hard to approximate with in a factor $\Omega(\sqrt{\text{OPT}})$, and the third result rules out *bicriteria approximation* – for any constant c , it is hard to approximate both l and $|C^*|$ within a factor of c .

The above results hold for directed graphs by definition. Our hard instances will have a natural *layered* structure, so it can be easily checked that the same results (up to a constant) hold for directed acyclic graphs. Since one vertex can be split as one directed edge, the same results hold for the edge version in directed acyclic graphs.

For Length-Bounded Edge Cut and Shortest Path Edge Interdiction in undirected graphs (collectively called *Short Path Edge Cut* onwards), we prove the following theorems.

► **Theorem 2.** *Assuming the Unique Games Conjecture, for infinitely many values of constant $l \in \mathbb{N}$, the following three tasks are NP-hard: Given an undirected graph $G = (V, E)$ and $s, t \in V$ where there exists $C^* \subseteq E$ such that $\text{dist}(G \setminus C^*) \geq l$,*

1. *Find $C \subseteq E$ such that $|C| \leq \Omega(\sqrt{l}) \cdot |C^*|$ and $\text{dist}(G \setminus C) \geq l$.*
2. *Find $C \subseteq E$ such that $|C| \leq |C^*|$ and $\text{dist}(G \setminus C) \geq l^{\frac{2}{3}}$.*
3. *Find $C \subseteq E$ such that $|C| \leq \Omega(l^{\frac{2\epsilon}{3}}) \cdot |C^*|$ and $\text{dist}(G \setminus C) \geq O(l^{\frac{2+2\epsilon}{3}})$ for some $0 < \epsilon < \frac{1}{2}$.*

Our hardness factors for the undirected edge versions, $\Omega(\sqrt{l})$ for Length-Bounded Edge Cut and $\Omega(\sqrt[3]{\text{OPT}})$ for Shortest Path Edge Interdiction, are slightly weaker than those for their vertex counterparts, but we are not aware of any approximation algorithm specialized for the undirected edge versions. It is an interesting open problem whether there exist better approximation algorithms for the undirected edge versions.

Directed Multicut and Directed Multiway Cut. Given a directed graph and two vertices s and t , one of the most natural variants of s - t cut is to remove the fewest edges to ensure that there is no directed path from s to t and no directed path from t to s . This problem is known as s - t *Bicut* and admits the trivial 2-approximation algorithm by computing the minimum s - t cut and t - s cut.

Directed Multiway Cut is a generalization of s - t Bicut that has been actively studied. Given a directed graph with k terminals s_1, \dots, s_k , the goal is to remove the fewest number of edges such that there is no path from s_i to s_j for any $i \neq j$. Directed Multiway Cut also admits 2-approximation [37, 11]. If k is allowed to increase polynomially with n , there is a simple reduction from Vertex Cover that shows $(2 - \epsilon)$ -approximation is hard under the UGC [18, 27].

Directed Multiway Cut can be further generalized to *Directed Multicut*. Given a directed graph with k source-sink pairs $(s_1, t_1), \dots, (s_k, t_k)$, the goal is to remove the fewest number of edges such that there is no path from s_i to t_i for any i . Computing the minimum s_i - t_i cut for all i separately gives the trivial k -approximation algorithm. Chuzhoy and Khanna [13] showed Directed Multicut is hard to approximate within a factor $2^{\Omega(\log^{1-\epsilon} n)} = 2^{\Omega(\log^{1-\epsilon} k)}$ when k is polynomially growing with n . Agarwal et al. [2] showed $\tilde{O}(n^{\frac{11}{23}})$ -approximation algorithm, which improves the trivial k -approximation when k is large.

Chekuri and Madan [11] showed simple approximation-preserving reductions from Directed Multicut with $k = 2$ to s - t Bicut (the other direction is trivially true), and (*Undirected*) *Node-weighted Multiway Cut* with $k = 4$ to s - t Bicut. Since Node-weighted Multiway Cut with $k = 4$ is hard to approximate within a factor $1.5 - \epsilon$ under the UGC [16] (matching the algorithm of Garg et al. [18]), the same hardness holds for s - t Bicut, Directed Multiway Cut, and Directed Multicut for constant k . To the best of our knowledge, $1.5 - \epsilon$ is the best hardness factor for constant k even assuming the UGC. In the same paper, Chekuri and Madan [11] asked whether a factor $2 - \epsilon$ hardness holds for s - t Bicut under the UGC.

We prove that for any constant $k \geq 2$, the trivial k -approximation for Directed Multicut might be optimal. Our result for $k = 2$ gives the optimal hardness result for s - t Bicut, answering the question of Chekuri and Madan.

► **Theorem 3.** *Assuming the Unique Games Conjecture, for every $k \geq 2$ and $\epsilon > 0$, Directed Multicut with k source-sink pairs is NP-hard to approximate within a factor $k - \epsilon$.*

► **Corollary 4.** *Assuming the Unique Games Conjecture, for any $\epsilon > 0$, s - t Bicut is hard to approximate within a factor $2 - \epsilon$.*

► **Remark.** Chekuri and Madan [12] obtained an independent and different proof of Theorem 3.

RMFC. *Resource Minimization for Fire Containment (RMFC)* is a problem closely related to Length-Bounded Cut with the additional notion of time. Given a graph G , a vertex s , and a subset T of vertices, consider the situation where fire starts at s on Day 0. For each Day i ($i \geq 1$), we can *save* at most k vertices, and the fire spreads from currently burning vertices to its unsaved neighbors. Once a vertex is burning or saved, it remains so from then onwards. The process is terminated when the fire cannot spread anymore. RMFC asks to

find a strategy to save k vertices each day with the minimum k so that no vertex in T is burnt. These problems model the spread of epidemics or ideas through a social network, and have been actively studied recently [9, 3, 1, 10].

RMFC, along with other variants, is first introduced by Hartnell [22]. Another well-studied variant is called the *Firefighter* problem, where we are only given $s \in V$ and want to maximize the number of vertices that are not burnt at the end. It is known to be NP-hard to approximate within a factor $n^{1-\epsilon}$ for any $\epsilon > 0$ [3]. King and MacGillivray [28] proved that RMFC is hard to approximate within a factor less than 2. Anshelevich et al. [3] presented an $O(\sqrt{n})$ -approximation algorithm for general graphs, and Chalermsook and Chuzhoy [9] showed that RMFC admits $O(\log^* n)$ -approximation in trees. Very recently, the approximation ratio in trees has been improved to $O(1)$ [1]. Both Anshelevich et al. [3] and Chalermsook and Chuzhoy [9] independently studied directed layer graphs with b layers, showing $O(\log b)$ -approximation.

Our final result on RMFC assumes a variant of the Unique Games Conjecture which is not known to be equivalent to the original UGC. Given a bipartite graph as an instance of Unique Games, it states that in the completeness case, all constraints incident on $(1 - \epsilon)$ fraction of vertices in one side are satisfied, and in the soundness case, in addition to having a low value, every $\frac{1}{10}$ fraction of vertices on one side have at least a $\frac{9}{10}$ fraction of vertices on the other side as neighbors. Our conjecture is implied by the conjecture of Bansal and Khot [6] that is used to prove the hardness of *Minimizing Weighted Completion Time with Precedence Constraints* and requires a more strict expansion condition. See [30] for the exact statement.

► **Theorem 5.** *Assuming Conjecture 7.5 of [30], it is NP-hard to approximate RMFC in undirected graphs within any constant factor.*

Again, our reduction has a natural layered structure and the result holds for directed layered graphs. With b layers, we prove that it is hard to approximate with in a factor $\Omega(\log b)$, matching the best approximation algorithms [9, 3].

1.2 Techniques

All our results are based on a general method of *converting an integrality gap instance to a dictatorship test*. This method has been successfully applied by Raghavendra [40] for Max-CSPs, Manokaran et al. [34] and Ene et al. [16] for Multiway Cut and Min CSPs, and Kumar et al. [29] for strict CSPs, and by Guruswami et al. [21] for k -uniform k -partite Hypergraph Vertex Cover, and Chekuri and Madan [12] for Directed Multicut. As mentioned in the introduction, the previous CSP-based results do not generally preserve the structure of constraint hypergraphs or use ingenious and specialized tricks to reduce the problem to a CSP.

We bypass this difficulty by constructing a special class of dictatorship tests that we call *length-control dictatorship tests*. Consider a meta-problem where given a directed graph $G = (V, E)$, some terminal vertices, and a set \mathcal{P} of *desired paths* between terminals, we want to remove the fewest number of non-terminal vertices to cut every path in \mathcal{P} . The integrality gap instances we use in this work [41, 4, 32, 9] share the common feature that every $p \in \mathcal{P}$ is of length at least r , and the fractional solution cuts $\frac{1}{r}$ fraction of each non-terminal vertex so that each path $p \in \mathcal{P}$ is cut. This gives a good LP value, and additional arguments are required to ensure that there is no efficient integral cut.

Given such an integrality gap instance, we construct our dictatorship test instance as follows. We replace every non-terminal vertex by a *hypercube* \mathbb{Z}_r^R and put edges such that

for two vertices (v, x) and (w, y) where $v, w \in V$ and $x, y \in \mathbb{Z}_r^R$, there is an edge from (v, x) to (w, y) if (1) $(v, w) \in E$ and (2) $y_j = x_j + 1$ for all $j \in [R]$. The set of desired paths \mathcal{P}' is defined to be $\{(s, (v_1, x_1), \dots, (v_l, x_l), t) : (s, v_1, \dots, v_l, t) \in \mathcal{P}\}$ (s, t denote some terminals). Note that each path in \mathcal{P}' is also of length at least r . We want to ensure that in the *completeness case* (i.e., every hypercube reveals the same *influential coordinate*), there is a very efficient cut, while in the *soundness case* (i.e., no hypercube reveals an influential coordinate), there is no such efficient cut.

In the completeness case, let $q \in [R]$ be an influential coordinate. For each vertex (v, x) where $v \in V, x \in \mathbb{Z}_r^R$, remove (v, x) if $x_q = 0$. Consider a desired path $p = (s, (v_1, x_1), \dots, (v_l, x_l), t) \in \mathcal{P}'$ for some terminals s, t and some $v_j \in V, x_j \in \mathbb{Z}_r^R$ ($1 \leq j \leq l$), and let $y_j = (x_j)_q$. By our construction, $y_{j+1} = y_j + 1$ for $0 \leq j < l$. Since p is desirable, $l \geq r$, so there exists j such that $y_j = (x_j)_q = 0$, but (v_j, x_j) is already removed by our previous definition. Therefore, every desired path is cut by this vertex cut. Note that this cut is integral and cuts exactly $\frac{1}{r}$ fraction of non-terminal vertices. This corresponds to the fractional solution to the gap instance that cuts $\frac{1}{r}$ fraction of every vertex.

For the soundness analysis, our final dictatorship test has additional *noise* vertices and edges to the test defined above. If no hypercube reveals an influential coordinate, the standard application of the invariance principle [36] proves that we can always take an edge between two hypercubes unless we almost completely cut one hypercube. We can then invoke the proof for the integrality gap instance to show that there is no efficient cut.

This idea is implicitly introduced by the work of Svensson [43] for Feedback Vertex Set (FVS) and DAG Vertex Deletion (DVD) by applying the *It ain't over till it's over theorem* to ingeniously constructed dictatorship tests with auxiliary vertices. Guruswami and Lee [20] gave a simpler construction and a new proof using the invariance principle instead of the *It ain't over till it's over theorem*. Our results are based on the observation that length-control dictatorship tests and LP gap instances *fool* algorithms in a similar way for various cut problems as mentioned above, so that the previous LP gap instances can be plugged into our framework to prove matching hardness results.

This method for the above meta-problem can be almost directly applied to Directed Multicut. For Length-Bounded Cut and RMFC in undirected graphs, we use the fact that the known integrality gap instances have a natural layered structure with s in the first layer and t in the last layer. Every edge is given a natural orientation, and the similar analysis can be applied. For Length-Bounded Cut, another set of edges called *long edges* are added to the dictatorship test. More technical work is required for edge cut versions in undirected graphs (Short Path Edge Cut), and the notion of time (RMFC).

Our framework seems general enough so that they can be applied to integrality gap instances to give strong hardness results. It would be interesting to further abstract this method of converting integrality gap instances to length-bounded dictatorship tests, as well as to apply it to other problems whose approximability is not well-understood.

2 Preliminaries

Graph Terminologies. Depending on whether we cut vertices or edges, we introduce weight $\text{wt}(v)$ for each vertex v , or weight $\text{wt}(e)$ for each edge e . Some weights can be ∞ , which means that some vertices or edges cannot be cut. For vertex-weighted graphs, we naturally have $\text{wt}(s) = \text{wt}(t) = \infty$. To reduce the vertex-weighted version to the unweighted version, we duplicate each vertex according to its weight and replace each edge by a complete bipartite graph between corresponding copies. To reduce the edge-weighted version to the unweighted

version, we replace a single edge with parallel edges according to its weight. To reduce to simple graphs, we split each parallel into two edges by introducing a new vertex.

For the Length-Bounded Cut problems, we also introduce length $\text{len}(e)$ for each edge e . It can be also dealt with serially splitting an edge according to its weight. We allow weights to be rational numbers, but as our hardness results are stated in terms of the length, all lengths in this work will be a positive integer.

For a path p , depending on the context, we abuse notation and interpret it as a set of edges or a set of vertices. The length of p is always defined to be the number of edges.

Gaussian Bounds for Correlated Spaces. We introduce the standard tools on correlated spaces from Mossel [36]. Given a probability space (Ω, μ) (we always consider finite probability spaces), let $\mathcal{L}(\Omega)$ be the set of functions $\{f : \Omega \rightarrow \mathbb{R}\}$ and for an interval $I \subseteq \mathbb{R}$, $\mathcal{L}_I(\Omega)$ be the set of functions $\{f : \Omega \rightarrow I\}$. For a subset $S \subseteq \Omega$, define *measure* of S to be $\mu(S) := \sum_{\omega \in S} \mu(\omega)$. A collection of probability spaces are said to be correlated if there is a joint probability distribution on them. We will denote k correlated spaces $\Omega_1, \dots, \Omega_k$ with a joint distribution μ as $(\Omega_1 \times \dots \times \Omega_k, \mu)$.

Given two correlated spaces $(\Omega_1 \times \Omega_2, \mu)$, we define the correlation between Ω_1 and Ω_2 by

$$\rho(\Omega_1, \Omega_2; \mu) := \sup \{ \text{Cov}[f, g] : f \in \mathcal{L}(\Omega_1), g \in \mathcal{L}(\Omega_2), \text{Var}[f] = \text{Var}[g] = 1 \}.$$

Given a probability space (Ω, μ) and a function $f \in \mathcal{L}(\Omega)$ and $p \in \mathbb{R}^+$, let $\|f\|_p := \mathbb{E}_{x \sim \mu} [|f(x)|^p]^{1/p}$.

Consider a product space $(\Omega^R, \mu^{\otimes R})$ and $f \in \mathcal{L}(\Omega^R)$. The *Efron-Stein decomposition* of f is given by

$$f(x_1, \dots, x_R) = \sum_{S \subseteq [R]} f_S(x_S)$$

where (1) f_S depends only on x_S and (2) for all $S \not\subseteq S'$ and all $x_{S'}, \mathbb{E}_{x' \sim \mu^{\otimes R}} [f_S(x') | x'_{S'} = x_{S'}] = 0$. The *influence* of the i th coordinate on f is defined by

$$\text{Inf}_i[f] := \mathbb{E}_{x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_R} [\text{Var}_{x_i} [f(x_1, \dots, x_R)]]$$

The influence has a convenient expression in terms of the Efron-Stein decomposition.

$$\text{Inf}_i[f] = \left\| \sum_{S:i \in S} f_S \right\|_2^2 = \sum_{S:i \in S} \|f_S\|_2^2.$$

We also define the *low-degree influence* of the i th coordinate.

$$\text{Inf}_i^{\leq d}[f] := \sum_{S:i \in S, |S| \leq d} \|f_S\|_2^2.$$

For $a, b \in [0, 1]$ and $\rho \in (0, 1)$, let

$$\Gamma_\rho(a, b) := \Pr[X \leq \Phi^{-1}(a), Y \geq \Phi^{-1}(1 - b)],$$

where X and Y are ρ -correlated standard Gaussian variables and Φ denotes the cumulative distribution function of a standard Gaussian. The following theorem bounds the product of two functions that do not share an influential coordinate in terms of their Gaussian counterparts.

► **Theorem 6** (Theorem 6.3 and Lemma 6.6 of [36]). *Let $(\Omega_1 \times \Omega_2, \mu)$ be correlated spaces such that the minimum nonzero probability of any atom in $\Omega_1 \times \Omega_2$ is at least α and such that $\rho(\Omega_1, \Omega_2; \mu) \leq \rho$. Then for every $\epsilon > 0$ there exist τ, d depending on ϵ and α such that if $f : \Omega_1^R \rightarrow [0, 1], g : \Omega_2^R \rightarrow [0, 1]$ satisfy $\min(\text{Inf}_i^{\leq d}[f], \text{Inf}_i^{\leq d}[g]) \leq \tau$ for all i , then $\mathbb{E}_{(x,y) \in \mu^{\otimes R}} [f(x)g(y)] \geq \Gamma_\rho(\mathbb{E}_x[f], \mathbb{E}_y[g]) - \epsilon$.*

Organization. The dictatorship tests for Short Path Edge Cut and Short Path Vertex Cut are presented in Section 3 and 4 respectively. Dictatorship tests for RMFC and Directed Multicut, as well as the reduction from Unique Games based on these tests, will appear in the full version of this paper [30].

3 Short Path Edge Cut

We propose our dictatorship test for Short Path Edge Cut that will be used for proving Unique Games hardness. It is parameterized by positive integers a, b, r, R . It is inspired by the integrality gap instances by Baier et al. [4] Mahjoub and McCormick [32], and made such that the edge cuts that correspond to *dictators* behave the same as the fractional solution that cuts $\frac{1}{r}$ fraction of every edge. All graphs in this section are undirected.

For positive integers a, b, r, R , we construct $\mathcal{D}_{a,b,r,R}^E = (V, E)$. Let $\Omega = \{0, \dots, r-1\}$, and $\mu : \Omega \rightarrow [0, 1]$ with $\mu(x) = \frac{1}{r}$ for each $x \in \Omega$. We also define a correlated probability space $(\Omega_1 \times \Omega_2, \nu)$ where both Ω_1, Ω_2 are copies of Ω . It is defined by the following process to sample $(x, y) \in \Omega^2$.

- Sample $x \in \{0, \dots, r-1\}$. Let $y = (x+1) \bmod r$.
- With probability $1 - \frac{1}{r}$, output (x, y) . Otherwise, resample $x, y \in \Omega$ independently and output (x, y) .

Note that the marginal distribution of both x and y is equal to μ . Given $x = (x_1, \dots, x_R) \in \Omega^R$ and $y = (y_1, \dots, y_R) \in \Omega^R$, let $\nu^{\otimes R}(x, y) = \prod_{i=1}^R \nu(x_i, y_i)$. We define $\mathcal{D}_{a,b,r,R}^E = (V, E)$ as follows.

- $V = \{s, t\} \cup \{v_x^i\}_{0 \leq i \leq b, x \in \Omega^R}$. Let v^i denote the set of vertices $\{v_x^i\}_{x \in \Omega^R}$.
- For any $x \in \Omega^R$, there is an edge from s to v_x^0 and an edge from v_x^b to t , both with weight ∞ and length 1.
- For $0 \leq i < b, x \in \Omega^R$, there is an edge (v_x^i, v_x^{i+1}) of length a and weight ∞ . Call it a *long edge*.
- For any $0 \leq i < b, x, y \in \Omega^R$, there is an edge (v_x^i, v_y^{i+1}) of length 1 and weight $\nu^{\otimes R}(x, y)$. Note that $\nu^{\otimes R}(x, y) > 0$ for any $x, y \in \Omega^R$. Call it a *short edge*. The sum of finite weights is b .

Completeness. We first prove that edge cuts that correspond to *dictators* behave the same as the fractional solution that gives $\frac{1}{r}$ to every short edge. Fix $q \in [R]$ and let E_q be the set of short edges defined by

$$E_q := \{(v_x^i, v_y^{i+1}) : 0 \leq i < b, y_q \neq x_q + 1 \bmod R \text{ or } (x_q, y_q) = (0, 1)\}.$$

When $(x, y) \in \Omega_1 \times \Omega_2$ is sampled according to ν , the probability that $y_q \neq x_q + 1 \bmod R$ or $(x_q, y_q) = (0, 1)$ is at most $\frac{2}{r}$. The total weight of E_q is $\frac{2b}{r}$.

► **Lemma 7.** *After removing edges in E_q , the length of the shortest path is at least $a(b-r+1)$.*

Proof. Let $p = (s, v_{x^1}^{i_1}, \dots, v_{x^z}^{i_z}, t)$ be a path from s to t where $i_j \in \{0, \dots, b\}$ and $x^j \in \Omega^R$ for each $1 \leq j \leq z$. Let $y_j := (x^j)_q \in \{0, \dots, r-1\}$ for each $1 \leq j \leq z$.

For each $1 \leq j < z$, the edge (p_j, p_{j+1}) is either a long edge or a short edge, and either taken forward (i.e., $i_j < i_{j+1}$) or backward (i.e., $i_j > i_{j+1}$). Let $z_{\text{LF}}, z_{\text{SF}}, z_{\text{LB}}, z_{\text{SB}}$ be the number of long edges taken forward, short edges taken forward, long edges taken backward, and short edges taken backward, respectively ($z_{\text{LF}} + z_{\text{SF}} + z_{\text{LB}} + z_{\text{SB}} = z - 1$). By considering how i_j changes,

$$z_{\text{LF}} + z_{\text{SF}} - z_{\text{LB}} - z_{\text{SB}} = b. \tag{1}$$

Consider how y_j changes. Taking a long edge does not change y_j . Taking a short edge forward increases y_j by 1 mod r , taking a short edge backward decreases y_j by 1 mod r . Since E_q is cut, y_j can never change from 0 to 1. This implies

$$z_{\text{SF}} - z_{\text{SB}} \leq r - 1. \quad (2)$$

(1) – (2) yields $z_{\text{LF}} - z_{\text{LB}} \geq b - r + 1$. The total length of p is at least $a \cdot z_{\text{LF}} \geq a(b - r + 1)$. ◀

Soundness. We first bound the correlation $\rho(\Omega_1, \Omega_2; \nu)$. The following lemma of Wenner [44] gives a convenient way to bound the correlation.

► **Lemma 8** (Corollary 2.18 of [44]). *Let $(\Omega_1 \times \Omega_2, \delta\mu + (1 - \delta)\mu')$ be two correlated spaces such that the marginal distribution of at least one of Ω_1 and Ω_2 is identical on μ and μ' . Then,*

$$\rho(\Omega_1, \Omega_2; \delta\mu + (1 - \delta)\mu') \leq \sqrt{\delta \cdot \rho(\Omega_1, \Omega_2; \mu)^2 + (1 - \delta) \cdot \rho(\Omega_1, \Omega_2; \mu')^2}.$$

When (x, y) is sampled from ν , they are completely independent with probability $\frac{1}{r}$. Therefore, we have $\rho := \rho(\Omega_1, \Omega_2; \nu) \leq \sqrt{1 - \frac{1}{r}}$. By Sheppard's Formula,

$$\Gamma_\rho\left(\frac{1}{2}, \frac{1}{2}\right) = \frac{1}{4} + \frac{1}{2\pi} \arcsin(-\rho) \geq \frac{1}{4} - \frac{1}{2\pi} \arccos\left(\frac{1}{\sqrt{r}}\right) = \sum_{n=0}^{\infty} \frac{(2n)!}{4^n (n!)^2 (2n+1)} \left(\frac{1}{\sqrt{r}}\right)^{2n+1} \geq \frac{1}{\sqrt{r}}.$$

Apply Theorem 6 ($\rho \leftarrow \rho, \alpha \leftarrow \frac{1}{r^3}, \epsilon \leftarrow \frac{\Gamma_\rho(\frac{1}{2}, \frac{1}{2})}{3}$) to get τ and d . We will later apply this theorem with the parameters obtained here.

Fix an arbitrary subset $C \subseteq E$ of short edges. For $0 \leq i < b$, let $C_i = C \cap (v^i \times v^{i+1})$. Call a pair $(i, i+1)$ as the i th layer, and say it is blocked when $\nu^{\otimes R}(C_i) \geq \frac{\Gamma_\rho(\frac{1}{2}, \frac{1}{2})}{2}$. Let b' be the number of blocked layers. For $0 \leq i \leq b$, let $S_i \subseteq v^i$ be such that $x \in S_i$ if there exists a path $(s, p_0, \dots, p_i = v_x^i)$ such that

- For $0 \leq i' \leq i$, $p_{i'} \in v^{i'}$.
- For $0 \leq i' < i$, $(p_{i'}, p_{i'+1})$ is short if and only if the i' th layer is unblocked.

Let $f_i : \Omega^R \mapsto [0, 1]$ be the indicator function of S_i . We prove that if none of f_i reveals any *influential coordinate*, S_b is nonempty, implying that there exists a path using b' long edges and $b - b'$ short edges. Therefore, even after removing edges in C , the length of the shortest path is at most $2 + ab' + (b - b')$.

► **Lemma 9.** *Suppose that for any $0 \leq i \leq b$ and $1 \leq j \leq R$, $\text{Inf}_j^{\leq d}[f_i] \leq \tau$. Then $S_b \neq \emptyset$.*

Proof. Assume towards contradiction that $S_b = \emptyset$. Since $S_0 = \Omega^R$ and $S_i = S_{i+1}$ if the i th layer is blocked (and we use long edges), there must exist i such that the i th layer is unblocked and $\mu^{\otimes R}(S_i) \geq \frac{1}{2}, \mu^{\otimes R}(S_{i+1}) < \frac{1}{2}$. All short edges between S_i and $v^{i+1} \setminus S_{i+1}$ are in C_i . Theorem 6 implies that $\nu^{\otimes R}(C_i) > \frac{2}{3} \Gamma_\rho(\frac{1}{2}, \frac{1}{2})$. This contradicts the fact that the i th layer is unblocked. ◀

In summary, in the completeness case, if we cut edges of total weight $k := k(a, b, r) = \frac{2b}{r}$, the length of the shortest path is at least $l := l(a, b, r) = a(b - r + 1)$. In the soundness case, even after cutting edges of total weight k' , at most $\frac{2k'}{\Gamma_\rho(\frac{1}{2}, \frac{1}{2})} \leq 2k' \sqrt{r}$ layers are blocked, the length of the shortest path is at most $l' = 2 + (b - 2k' \sqrt{r}) + 2ak' \sqrt{r}$.

- Let $a = 4, b = 2r - 1$ so that $k \leq 4, l = 4r$. Requiring $l' \geq l$ results in $k' = \Omega(\sqrt{r})$, giving a gap of $\Omega(\sqrt{r}) = \Omega(\sqrt{l})$ between the completeness case and the soundness case for Length-Bounded Edge Cut.
- Let $a = \sqrt{r}, b = 2r - 1$ so that $k \leq 4, l = r^{1.5}$. Requiring $k' \leq 4$ results in $l' = O(r)$, giving a gap of $\Omega(\sqrt{r}) = \Omega(l^{1/3})$ for Shortest Path Interdiction. Generally, $k' \leq O(r^\epsilon)$ results in $l' \leq O(r^{1+\epsilon})$, giving an $(O(r^\epsilon), O(r^{1/2-\epsilon}))$ -bicriteria gap for any $\epsilon \in (0, \frac{1}{2})$.

4 Short Path Vertex Cut

We propose our dictatorship test for Short Path Vertex Cut that will be used for proving Unique Games hardness. It is parameterized by positive integers a, b, r, R and small $\epsilon > 0$. It is inspired by the integrality gap instances by Baier et al. [4] Mahjoub and McCormick [32], and made such that the vertex cuts that correspond to *dictators* behave the same as the fractional solution that cuts $\frac{1}{r}$ fraction of every vertex. All graphs in this section are undirected.

For positive integers a, b, r, R , and $\epsilon > 0$, define $\mathcal{D}_{a,b,r,R,\epsilon}^V = (V, E)$ be the graph defined as follows. Consider the probability space (Ω, μ) where $\Omega := \{0, \dots, r-1, *\}$, and $\mu : \Omega \mapsto [0, 1]$ with $\mu(*) = \epsilon$ and $\mu(x) = \frac{1-\epsilon}{r}$ for $x \neq *$.

- $V = \{s, t\} \cup \{v_x^i\}_{0 \leq i \leq b, x \in \Omega^R}$. Let v_x^i denote the set of vertices $\{v_x^i\}_x$.
- For $0 \leq i \leq b$ and $x \in \Omega^R$, $\text{wt}(v_x^i) = \mu^{\otimes R}(x)$. Note that the sum of weights is $b + 1$.
- For any $0 \leq i \leq b$, there are edges from s to each vertex in v_i with length $ai + 1$ and edges from each vertex in v_i to t with length $(b - i)a + 1$.
- For $x, y \in \Omega^R$, we call that x and y are *compatible* if
 - For any $1 \leq j \leq R$: $[y_j = (x_j + 1) \bmod r]$ or $[y_j = *]$ or $[x_j = *]$.
- For any $0 \leq i < b$ and compatible $x, y \in \Omega^R$, we have an edge (v_x^i, v_y^{i+1}) of length 1 (called a *short edge*).
- For any i, j such that $0 \leq i < j - 1 < b$ and compatible $x, y \in \Omega^R$, we have an edge (v_x^i, v_y^j) of length $(j - i)a$ (called a *long edge*).

Completeness. We first prove that vertex cuts that correspond to *dictators* behave the same as the fractional solution that gives $\frac{1}{r}$ to every vertex. For any $q \in [R]$, let $V_q := \{v_x^i : 0 \leq i \leq b, x_q = * \text{ or } 0\}$. Note that the total weight of V_q is $(b + 1)(\epsilon + \frac{1-\epsilon}{r})$.

► **Lemma 10.** *After removing vertices in V_q , the length of the shortest path is at least $a(b - r + 2)$.*

Proof. Let $p = (s, v_{x^1}^{i_1}, \dots, v_{x^z}^{i_z}, t)$ be a path from s to t where $i_j \in \{0, \dots, b\}$ and $x^j \in \Omega^R$ for each $1 \leq j \leq z$. Let $y_j := (x^j)_q \in \{0, \dots, r-1\}$ for each $1 \leq j \leq z$.

For each $1 \leq j < z$, the edge $(v_{x^j}^{i_j}, v_{x^{j+1}}^{i_{j+1}})$ is either a long edge or a short edge, and either taken forward (i.e., $i_j < i_{j+1}$) or backward (i.e., $i_j > i_{j+1}$). Let $z_{\text{LF}}, z_{\text{SF}}, z_{\text{LB}}, z_{\text{SB}}$ be the number of long edges taken forward, short edges taken forward, long edges taken backward, and short edges taken backward, respectively ($z_{\text{LF}} + z_{\text{SF}} + z_{\text{LB}} + z_{\text{SB}} = z - 1$). For $1 \leq j \leq z_{\text{LF}}$ (resp. z_{LB}), consider the j th long edge taken forward (resp. backward) – it is $(v_{x^{j'}}^{i_{j'}}, v_{x^{j'+1}}^{i_{j'+1}})$ for some j' . Let s_j^{F} (resp. s_j^{B}) be $|i_{j'} - i_{j'+1}|$. The following equality holds by observing how i_j changes.

$$i_1 + \sum_{j=1}^{z_{\text{LF}}} s_j^{\text{F}} + z_{\text{SF}} - \sum_{j=1}^{z_{\text{LB}}} s_j^{\text{B}} - z_{\text{SB}} = i_z \quad \Rightarrow \quad i_1 + \sum_{j=1}^{z_{\text{LF}}} s_j^{\text{F}} + z_{\text{SF}} - z_{\text{LB}} - z_{\text{SB}} - i_z \geq 0. \quad (3)$$

Consider how y_j changes. Taking any edge forward increases y_j , and taking any edge backward decreases y_j . Since y_j can never be 0 or $*$, we can conclude that

$$z_{\text{LF}} + z_{\text{SF}} - z_{\text{LB}} - z_{\text{SB}} \leq r - 2. \quad (4)$$

(3) – (4) yields

$$i_1 - i_z + \sum_{j=1}^{z_{\text{LF}}} (s_j^{\text{F}} - 1) \geq 2 - r \quad \Rightarrow \quad i_1 - i_z + \sum_{j=1}^{z_{\text{LF}}} s_j^{\text{F}} \geq 2 - r. \quad (5)$$

The total length of p is

$$\begin{aligned} & 2 + a(i_1 + b - i_z + \sum_{j=1}^{z_{\text{LF}}} s_j^{\text{F}} + \sum_{j=1}^{z_{\text{LB}}} s_j^{\text{B}}) + z_{\text{SF}} + z_{\text{SB}} \\ & \geq a(i_1 + b - i_z + \sum_{j=1}^{z_{\text{LF}}} s_j^{\text{F}}) \\ & \geq a(b - r + 2). \end{aligned}$$

◀

Soundness. To analyze soundness, we define a correlated probability space $(\Omega_1 \times \Omega_2, \nu)$ where both Ω_1, Ω_2 are copies of $\Omega = \{0, \dots, r-1, *\}$. It is defined by the following process to sample $(x, y) \in \Omega^2$.

- Sample $x \in \{0, \dots, r-1\}$. Let $y = (x+1) \bmod r$.
- Change x to $*$ with probability ϵ . Do the same for y independently.

Note that the marginal distribution of both x and y is equal to μ . Assuming $\epsilon < \frac{1}{2r}$, the minimum probability of any atom in $\Omega_1 \times \Omega_2$ is ϵ^2 . Furthermore, in our correlated space, $\nu(x, *) > 0$ for all $x \in \Omega_1$ and $\nu(*, x) > 0$ for all $x \in \Omega_2$. We use the following lemma to bound the correlation.

► **Lemma 11** (Lemma 2.9 of [36]). *Let $(\Omega_1 \times \Omega_2, \mu)$ be two correlated spaces such that the probability of the smallest atom in $\Omega_1 \times \Omega_2$ is at least $\alpha > 0$. Define a bipartite graph $G = (\Omega_1 \cup \Omega_2, E)$ where $(a, b) \in \Omega_1 \times \Omega_2$ satisfies $(a, b) \in E$ if $\mu(a, b) > 0$. If G is connected, then $\rho(\Omega_1, \Omega_2; \mu) \leq 1 - \frac{\alpha^2}{2}$.*

Therefore, we can conclude that $\rho(\Omega_1, \Omega_2; \nu) \leq \rho := 1 - \frac{\epsilon^4}{2}$. Apply Theorem 6 ($\rho \leftarrow \rho, \alpha \leftarrow \epsilon^2, \epsilon \leftarrow \frac{\Gamma_{\rho}(\frac{\epsilon}{3}, \frac{\epsilon}{3})}{2}$) to get τ and d . We will later apply this theorem with the parameters obtained here. Fix an arbitrary subset $C \subseteq V$, and $C_i := C \cap v^i$. For $0 \leq i \leq b$, call v^i *blocked* if $\mu^{\otimes R}[C_i(x)] \geq 1 - \epsilon$. At most $\lfloor \frac{\text{wt}(C)}{1-\epsilon} \rfloor$ v^i 's can be blocked. Let k' be the number of blocked v^i 's, and $z = b + 1 - k'$ be the number of unblocked v^i 's. Let $\{v^{i_1}, \dots, v^{i_z}\}$ be the set of unblocked v^i 's with $i_1 < i_2 < \dots < i_z$.

For $1 \leq j \leq z$, let $S_j \subseteq v^{i_j}$ be such that $x \in S_j$ if there exists a path $(p_0 = s, p_1, \dots, p_{j-1}, v_x^{i_j})$ such that each $p_{j'} \in v^{i_{j'}} \setminus C$ ($1 \leq j' < j$). For $1 \leq j \leq z$, let $f_j : \Omega^R \mapsto [0, 1]$ be the indicator function of S_j .

We prove that if none of f_j reveals any *influential coordinate*, $\mu^{\otimes R}(S_z) > 0$. Since any path passing v^{i_1}, \dots, v^{i_z} (bypassing only blocked v^i 's) uses short edges at least $b - 2k'$ times, so the length of the shortest path after removing C is at most $2 + (b - 2k') + 2ak'$.

► **Lemma 12.** *Suppose that for any $1 \leq j \leq z$ and $1 \leq i \leq R$, $\text{Inf}_i^{\leq d}[f_j] \leq \tau$. Then $\mu^{\otimes R}(S_z) > 0$.*

Proof. We prove by induction that $\mu^{\otimes R}(S_j) \geq \frac{\epsilon}{3}$. It holds when $j = 1$ since v^{i_1} is unblocked. Assuming $\mu^{\otimes R}(S_j) \geq \frac{\epsilon}{3}$, since S_j does not reveal any influential coordinate, Theorem 6 shows that for any subset $T_{j+1} \subseteq v^{i_{j+1}}$ with $\mu^{\otimes R}(T_{j+1}) \geq \frac{\epsilon}{3}$, there exists an edge between S_j and T_{j+1} . If $S'_{j+1} \subseteq v^{i_{j+1}}$ is the set of neighbors of S_j , we have $\mu^{\otimes R}(S'_{j+1}) \geq 1 - \frac{\epsilon}{3}$. Since $v^{i_{j+1}}$ is unblocked, $\mu^{\otimes R}(S'_{j+1} \setminus C) \geq \frac{2\epsilon}{3}$, completing the induction. \blacktriangleleft

In summary, in the completeness case, if we cut vertices of total weight $k := k(a, b, r, \epsilon) = (b+1)(\epsilon + \frac{1-\epsilon}{r})$, the length of the shortest path is at least $l := l(a, b, r, \epsilon) = a(b-r+2)$. In the soundness case, even after cutting vertices of total weight k' , the length of the shortest path is at most $2 + (b - \frac{k'}{1-\epsilon}) + 2a(\frac{k'}{1-\epsilon})$.

- Let $a = 4, b = 2r - 2$ and ϵ small enough so that $k \leq 2, l = 4r$. Requiring $l' \geq l$ results in $k' = \Omega(r)$, giving a gap of $\Omega(r) = \Omega(l)$ for Length Bounded Cut.
- Let $a = r, b = 2r - 2$ and ϵ small enough so that $k \leq 2, l = r^2$. Requiring $k' \leq 2$ results in $l' = O(r)$, giving a gap of $\Omega(r) = \Omega(\sqrt{l})$ for Shortest Path Interdiction. Generally, $k' \leq O(r^\epsilon)$ results in $l' \leq O(r^{1+\epsilon})$, giving an $(O(r^\epsilon), O(r^{1-\epsilon}))$ -bicriteria gap for any $\epsilon \in (0, 1)$.

Acknowledgments. The author thanks Konstantin Makarychev for useful discussions on Directed Multicut, and Marek Elias for introducing Shortest Path Interdiction.

References

- 1 David Adjiashvili, Andrea Baggio, and Rico Zenklusen. Firefighting on trees beyond integrality gaps. *arXiv preprint arXiv:1601.00271*, 2016.
- 2 Amit Agarwal, Noga Alon, and Moses S. Charikar. Improved approximation for directed cut problems. In *Proceedings of the Thirty-ninth Annual ACM Symposium on Theory of Computing*, STOC'07, pages 671–680, New York, NY, USA, 2007. ACM. doi:10.1145/1250790.1250888.
- 3 Elliot Anshelevich, Deeparnab Chakrabarty, Ameya Hate, and Chaitanya Swamy. Approximability of the firefighter problem. *Algorithmica*, 62(1-2):520–536, 2012.
- 4 Georg Baier, Thomas Erlebach, Alexander Hall, Ekkehard Köhler, Petr Kolman, Ondřej Pangrác, Heiko Schilling, and Martin Skutella. Length-bounded cuts and flows. *ACM Transactions on Algorithms*, 7(1):4:1–4:27, December 2010. Preliminary version in ICALP'06. doi:10.1145/1868237.1868241.
- 5 Michael O Ball, Bruce L Golden, and Rakesh V Vohra. Finding the most vital arcs in a network. *Operations Research Letters*, 8(2):73–76, 1989.
- 6 N. Bansal and S. Khot. Optimal long code test with one free bit. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science*, FOCS'09, pages 453–462, Oct 2009. doi:10.1109/FOCS.2009.23.
- 7 Halil Bayrak and Matthew D Bailey. Shortest path network interdiction with asymmetric information. *Networks*, 52(3):133–140, 2008.
- 8 Cristina Bazgan, André Nichterlein, and Rolf Niedermeier. A refined complexity analysis of finding the most vital edges for undirected shortest paths. In Vangelis Th. Paschos and Peter Widmayer, editors, *Algorithms and Complexity*, volume 9079 of *Lecture Notes in Computer Science*, pages 47–60. Springer International Publishing, 2015. doi:10.1007/978-3-319-18173-8_3.
- 9 Parinya Chalermsook and Julia Chuzhoy. Resource minimization for fire containment. In *Proceedings of the Twenty-first Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA'10, pages 1334–1349, Philadelphia, PA, USA, 2010. Society for Industrial and Applied Mathematics. URL: <http://dl.acm.org/citation.cfm?id=1873601.1873709>.

- 10 Parinya Chalermsook and Daniel Vaz. New integrality gap results for the firefighters problem on trees. *arXiv preprint arXiv:1601.02388*, 2016.
- 11 Chandra Chekuri and Vivek Madan. Simple and fast rounding algorithms for directed and node-weighted multiway cut. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 797–807, 2016.
- 12 Chandra Chekuri and Vivek Madan. Approximating multicut and the demand graph. *arXiv preprint arXiv:1607.07200*, 2017. To appear in SODA’17.
- 13 Julia Chuzhoy and Sanjeev Khanna. Polynomial flow-cut gaps and hardness of directed cut problems. *Journal of the ACM*, 56(2), 2009.
- 14 HW Corley and Y Sha David. Most vital links and nodes in weighted networks. *Operations Research Letters*, 1(4):157–160, 1982.
- 15 Pavel Dvořák and Dušan Knop. Parametrized complexity of length-bounded cuts and multi-cuts. In Rahul Jain, Sanjay Jain, and Frank Stephan, editors, *Theory and Applications of Models of Computation*, volume 9076 of *Lecture Notes in Computer Science*, pages 441–452. Springer International Publishing, 2015. doi:10.1007/978-3-319-17142-5_37.
- 16 Alina Ene, Jan Vondrák, and Yi Wu. Local distribution and the symmetry gap: Approximability of multiway partitioning problems. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 306–325. SIAM, 2013.
- 17 Till Fluschnik, Danny Hermelin, André Nichterlein, and Rolf Niedermeier. Fractals for kernelization lower bounds, with an application to length-bounded cut problems. *arXiv preprint arXiv:1512.00333*, 2015.
- 18 Naveen Garg, Vijay V Vazirani, and Mihalis Yannakakis. Multiway cuts in directed and node weighted graphs. In *Automata, Languages and Programming*, pages 487–498. Springer, 1994.
- 19 Petr A. Golovach and Dimitrios M. Thilikos. Paths of bounded length and their cuts: Parameterized complexity and algorithms. *Discrete Optimization*, 8(1):72 – 86, 2011. Parameterized Complexity of Discrete Optimization. doi:http://dx.doi.org/10.1016/j.disopt.2010.09.009.
- 20 Venkatesan Guruswami and Euiwoong Lee. Simple proof of hardness of feedback vertex set. *Theory of Computing*, 12(6):1–11, 2016.
- 21 Venkatesan Guruswami, Sushant Sachdeva, and Rishi Saket. Inapproximability of minimum vertex cover on k-uniform k-partite hypergraphs. *SIAM Journal on Discrete Mathematics*, 29(1):36–58, 2015.
- 22 Bert Hartnell. Firefighter! an application of domination. presentation. In *25th Manitoba Conference on Combinatorial Mathematics and Computing, University of Manitoba in Winnipeg, Canada*, 1995.
- 23 Eitan Israeli and R Kevin Wood. Shortest-path network interdiction. *Networks*, 40(2):97–111, 2002.
- 24 Leonid Khachiyan, Endre Boros, Konrad Borys, Khaled Elbassioni, Vladimir Gurvich, Gabor Rudolf, and Jihui Zhao. On short paths interdiction problems: Total and node-wise limited interdiction. *Theory of Computing Systems*, 43(2):204–233, 2007. doi:10.1007/s00224-007-9025-6.
- 25 Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the 34th annual ACM Symposium on Theory of Computing*, STOC’02, pages 767–775, 2002.
- 26 Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O’Donnell. Optimal inapproximability results for Max-Cut and other 2-variable CSPs? *SIAM Journal on Computing*, 37(1):319–357, 2007.
- 27 Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within $2 - \epsilon$. *Journal of Computer and System Sciences*, 74(3):335–349, 2008.

- 28 Andrew King and Gary MacGillivray. The firefighter problem for cubic graphs. *Discrete Mathematics*, 310(3):614–621, 2010.
- 29 Amit Kumar, Rajsekar Manokaran, Madhur Tulsiani, and Nisheeth K. Vishnoi. On LP-based approximability for strict CSPs. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA’11, pages 1560–1573. SIAM, 2011. URL: <http://dl.acm.org/citation.cfm?id=2133036.2133157>.
- 30 Euiwoong Lee. Improved hardness for cut, interdiction, and firefighter problems. *arXiv preprint arXiv:1607.05133*, 2016. URL: <https://arxiv.org/abs/1607.05133>
- 31 László Lovász, V Neumann-Lara, and M Plummer. Mengerian theorems for paths of bounded length. *Periodica Mathematica Hungarica*, 9(4):269–276, 1978.
- 32 A Ridha Mahjoub and S Thomas McCormick. Max flow and min cut with bounded-length paths: complexity, algorithms, and approximation. *Mathematical programming*, 124(1-2):271–284, 2010.
- 33 K. Malik, A. K. Mittal, and S. K. Gupta. The k most vital arcs in the shortest path problem. *Oper. Res. Lett.*, 8(4):223–227, August 1989. doi:10.1016/0167-6377(89)90065-5.
- 34 Rajsekar Manokaran, Joseph Seffi 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*, pages 11–20. ACM, 2008.
- 35 David P Morton. Stochastic network interdiction. *Wiley Encyclopedia of Operations Research and Management Science*, 2011.
- 36 Elchanan Mossel. Gaussian bounds for noise correlation of functions. *Geometric and Functional Analysis*, 19(6):1713–1756, 2010. doi:10.1007/s00039-010-0047-x.
- 37 Joseph Naor and Leonid Zosin. A 2-approximation algorithm for the directed multiway cut problem. *SIAM Journal on Computing*, 31(2):477–482, 2001.
- 38 K. Okumoto, T. Fukunaga, and H. Nagamochi. Divide-and-conquer algorithms for partitioning hypergraphs and submodular systems. *Algorithmica*, 62(3):787–806, 2012.
- 39 Feng Pan and Aaron Schild. Interdiction problems on planar graphs. In Prasad Raghavendra, Sofya Raskhodnikova, Klaus Jansen, and JoséD.P. Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, volume 8096 of *Lecture Notes in Computer Science*, pages 317–331. Springer Berlin Heidelberg, 2013. doi:10.1007/978-3-642-40328-6_23.
- 40 Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the 40th annual ACM symposium on Theory of computing*, STOC’08, pages 245–254, 2008.
- 41 Michael Saks, Alex Samorodnitsky, and Leonid Zosin. A lower bound on the integrality gap for minimum multicut in directed networks. *Combinatorica*, 24(3):525–530, 2004.
- 42 J Cole Smith, Mike Prince, and Joseph Geunes. Modern network interdiction problems and algorithms. In *Handbook of Combinatorial Optimization*, pages 1949–1987. Springer, 2013.
- 43 Ola Svensson. Hardness of vertex deletion and project scheduling. *Theory of Computing*, 9(24):759–781, 2013. Preliminary version in APPROX’12. doi:10.4086/toc.2013.v009a024.
- 44 Cenny Wenner. Circumventing d -to-1 for approximation resistance of satisfiable predicates strictly containing parity of width four. *Theory of Computing*, 9(23):703–757, 2013.