

Global and fixed-terminal cuts in digraphs

Kristóf Bérczi* Karthekeyan Chandrasekaran† Tamás Király* Euiwoong Lee‡
Chao Xu†

Abstract

The computational complexity of multicut-like problems may vary significantly depending on whether the terminals are fixed or not. In this work we present a comprehensive study of this phenomenon in two types of cut problems in directed graphs: double cut and bicut.

1. The fixed-terminal edge-weighted double cut is known to be solvable efficiently. We show a tight approximability factor of 2 for the fixed-terminal node-weighted double cut. We show that the global node-weighted double cut cannot be approximated to a factor smaller than $3/2$ under the Unique Games Conjecture (UGC).
2. The fixed-terminal edge-weighted bicut is known to have a tight approximability factor of 2. We show that the global edge-weighted bicut is approximable to a factor strictly better than 2, and that the global node-weighted bicut cannot be approximated to a factor smaller than $3/2$ under UGC.
3. In relation to these investigations, we also prove two results on undirected graphs which are of independent interest. First, we show NP-completeness and a tight inapproximability bound of $4/3$ for the node-weighted 3-cut problem. Second, we show that for constant k , there exists an efficient algorithm to solve the minimum $\{s, t\}$ -separating k -cut problem.

Our techniques for the algorithms are combinatorial, based on LPs and based on enumeration of approximate min-cuts. Our hardness results are based on combinatorial reductions and integrality gap instances.

1 Introduction

The minimum two-terminal cut problem (min $s - t$ cut) and its global version (min cut) are classic interdiction problems with fast algorithms. Generalizations of the fixed-terminal versions, including the multi-cut and the multi-way cut, as well as generalizations of the global versions, including the k -cut, have been well-studied in the algorithmic literature [9, 13]. In this work, we study two generalizations of global cut problems in directed graphs, namely double cut and bicut (that we describe below). We study the power and limitations of fixed terminal versions of these cut problems in order to solve the global versions. In the process, we examine “intermediate” multicut problems where only a subset of the terminals are fixed, and obtain some results of independent interest. In particular, we show that the undirected $\{s, t\}$ -separating k -cut problem, where two of the k terminals are fixed, is polynomial-time solvable for constant k . In what follows, we describe the problems along with the results.

The starting point of this work is node-weighted double cut, that we describe below. We recall that an arborescence in a directed graph $D = (V, E)$ is a minimal subset $F \subseteq E$ of arcs such that there exists a node $r \in V$ with every node $u \in V$ having a unique path from r to u in the subgraph (V, F) (e.g., see [26]).

*MTA-ELTE Egerváry Research Group, Budapest {berkri,tkiraly}@cs.elte.hu

†University of Illinois, Urbana-Champaign {karthe,chaoxu3}@illinois.edu

‡Carnegie Mellon University, Pittsburgh euiwoonl@cs.cmu.edu

Double Cut. The input to the NODEDOUBLECUT problem is a directed graph and the goal is to find the smallest number of nodes whose deletion ensures that the remaining graph has no arborescence. This problem is key to understanding fault tolerant consensus in networks. We briefly describe this connection.

Significance of double cut. In a recent work, Tseng and Vaidya [28] showed that *consensus* in a directed graph can be achieved in the *synchronous model* subject to the failure of f nodes *if and only if* the removal of any subset of f nodes still leaves an arborescence in the remaining graph. Thus, the number of nodes whose failure can be tolerated for the purposes of achieving consensus in a network is *exactly* one less than the smallest number of nodes whose removal ensures that there is no arborescence in the network. So, it is imperative for the network authority to be able to compute this number.

A directed graph $D = (V, E)$ has no arborescence if and only if¹ there exist two distinct nodes $s, t \in V$ such that every node $u \in V$ can reach at most one node in $\{s, t\}$ [2]. By this characterization, every directed graph that is not a tournament has a feasible solution to NODEDOUBLECUT. This characterization motivates the following fixed-terminal version, denoted $\{s, t\}$ -NODEDOUBLECUT: Given a directed graph with two specified nodes s and t , find the smallest number of nodes whose deletion ensures that every remaining node u can reach at most one node in $\{s, t\}$ in the resulting graph. An instance of $\{s, t\}$ -NODEDOUBLECUT has a feasible solution provided that the instance has no edge between s and t . An efficient algorithm to solve/approximate $\{s, t\}$ -NODEDOUBLECUT immediately gives an efficient algorithm to solve/approximate NODEDOUBLECUT.

Previous work. In the edge-weighted variation of two-terminal double cut, namely $\{s, t\}$ -EDGEDOUBLECUT, the goal is to delete the smallest number of edges to ensure that every node in the graph can reach at most one node in $\{s, t\}$. Similarly, in the global variant, denoted EDGEDOUBLECUT, the goal is to delete the smallest number of edges to ensure that there exist nodes s, t such that every node u can reach at most one node in $\{s, t\}$. Thus, EDGEDOUBLECUT is equivalent to deleting the smallest number of edges to ensure that the graph has no arborescence. The fixed-terminal variant $\{s, t\}$ -EDGEDOUBLECUT is solvable in polynomial time using maximum flow and, consequently, EDGEDOUBLECUT is also solvable in polynomial time [2].

Results for double cut. We show the following inapproximability results for $\{s, t\}$ -NODEDOUBLECUT.

Theorem 1.1. $\{s, t\}$ -NODEDOUBLECUT is NP-hard, has no efficient $(100/99 - \epsilon)$ -approximation for any $\epsilon > 0$ assuming $P \neq NP$, and has no efficient $(2 - \epsilon)$ -approximation for any $\epsilon > 0$ assuming the Unique Games Conjecture.

We show a 2-approximation for $\{s, t\}$ -NODEDOUBLECUT matching the inapproximability factor. This also leads to a 2-approximation for NODEDOUBLECUT.

Theorem 1.2. There exists an efficient 2-approximation algorithm for $\{s, t\}$ -NODEDOUBLECUT and NODEDOUBLECUT.

Next we focus on the inapproximability of NODEDOUBLECUT.

Theorem 1.3. NODEDOUBLECUT is NP-hard, has no efficient $(100/99 - \epsilon)$ -approximation for any $\epsilon > 0$ assuming $P \neq NP$, and has no efficient $(3/2 - \epsilon)$ -approximation for any $\epsilon > 0$ assuming the Unique Games Conjecture.

¹We believe that this characterization led earlier authors [2] to coin the term *double cut* to refer to the edge deletion variant of the problem and we are following this naming convention.

We show the NP-hardness of NODEDOUBLECUT in Theorem 1.3 by a reduction from the node-weighted 3-cut problem in undirected graphs. In the node weighted 3-cut problem, denoted NODE-3-CUT, the input is an undirected graph and the goal is to find the smallest subset of nodes whose deletion leads to at least 3 connected components in the remaining graph. A classic result of Goldschmidt and Hochbaum [13] showed that the edge-weighted variant, denoted EDGE-3-CUT (more commonly known as 3-cut)—namely find a smallest subset of edges of a given undirected graph whose deletion leads to at least 3 connected components—is solvable in polynomial time. Surprisingly, the complexity of NODE-3-CUT is open. We present the first results on the complexity of NODE-3-CUT.

Theorem 1.4. *NODE-3-CUT is NP-hard, has no efficient $(100/99 - \epsilon)$ -approximation for any $\epsilon > 0$ assuming $P \neq NP$, and has no efficient $(4/3 - \epsilon)$ -approximation for any $\epsilon > 0$ assuming the Unique Games Conjecture.*

The inapproximability factor of $4/3$ mentioned in the above theorem is tight: the $4/3$ -approximation factor can be achieved by guessing 3 terminals that are separated and then using well-known approximation algorithms to solve the resulting node-weighted 3-terminal cut instance. The best known approximation factor achievable for the node-weighted 3-terminal cut problem is $4/3$ [12].

Returning to NODEDOUBLECUT, we note that our results (Theorems 1.2 and 1.3) still have a gap in the approximability of NODEDOUBLECUT. A similar phenomenon can be observed for bicut problems, where the approximability of the fixed-terminal variants is already well-understood, but the complexity of the global problems is unknown. In the following we describe these problems and report on some progress concerning their complexity.

Bicuts. The edge-weighted version of two-terminal bicut, denoted $\{s, t\}$ -EDGEBICUT, is the following: Given a directed graph with two specified nodes s and t , find the smallest number of edges whose deletion ensures that s cannot reach t and t cannot reach s in the resulting graph. It is well-known that $\{s, t\}$ -EDGEBICUT is equivalent to the 2-terminal multiway-cut and is a special case of 2-pairs multi-cut in directed graphs. This problem has a rich history and has seen renewed interest in the last few months culminating in inapproximability results matching the best-known approximability factor: $\{s, t\}$ -EDGEBICUT admits a 2-factor approximation [23, 3] and has no efficient $(2 - \epsilon)$ -approximation assuming the Unique Games Conjecture [19, 4]. In the global version, denoted EDGEBICUT, the goal is to find the smallest number of edges whose deletion ensures that there exist two distinct nodes s and t such that s cannot reach t and t cannot reach s in the resulting digraph.

The dichotomy between global cut problems and fixed-terminal cut problems in undirected graphs is well-known. For concreteness, we recall EDGE-3-CUT and EDGE-3-WAY-CUT. In EDGE-3-CUT, the input is an undirected graph and the goal is to find the smallest number of edges to delete so that the resulting graph has at least 3 connected components. In EDGE-3-WAY-CUT, the input is an undirected graph with 3 specified nodes and the goal is to find the smallest number of edges to delete so that the resulting graph has at least 3 connected components with at most one of the 3 specified nodes in each component. While EDGE-3-WAY-CUT is NP-hard [9], EDGE-3-CUT is solvable efficiently [13]. However, such a dichotomy is unknown for directed graphs. In particular, it is unknown whether EDGEBICUT is solvable efficiently. Our next result shows evidence of such a dichotomic behaviour. While $\{s, t\}$ -EDGEBICUT is inapproximable to a factor better than 2 assuming UGC, we show that EDGEBICUT is approximable to a factor strictly better than 2.

Theorem 1.5. *There exists an efficient $(2 - \frac{1}{448})$ -approximation algorithm for EDGEBICUT.*

We mention one main ingredient of our algorithm to prove Theorem 1.5 that has been of independent interest. As a sub-problem in the algorithm for Theorem 1.5, we solve the following,

denoted $(s, *, t)$ -EDGE_{LIN}3CUT (abbreviating edge-weighted linear 3-cut): Given a directed graph $D = (V, E)$ and two specified nodes $s, t \in V$, find the smallest number of edges to delete so that there exists a node r with the property that s cannot reach r and t , and r cannot reach t in the resulting graph. This problem is a global variant of (s, r, t) -EDGE_{LIN}3CUT, introduced in [10], where the input specifies three terminals s, r, t and the goal is to find the smallest number of edges whose removal achieves the property above. A simple reduction from EDGE-3-WAY-CUT shows that (s, r, t) -EDGE_{LIN}3CUT is NP-hard. The approximability of (s, r, t) -EDGE_{LIN}3CUT was studied by Chekuri and Madan [4]. They showed that the inapproximability factor coincides with the flow-cut gap of the associated path-blocking linear program.

There exists a simple combinatorial 2-approximation algorithm for (s, r, t) -EDGE_{LIN}3CUT. An approach to approximate $(s, *, t)$ -EDGE_{LIN}3CUT is to guess the terminal r and use the above-mentioned combinatorial approximation to obtain an efficient 2-approximation. We obtain the following improved approximation factor for $(s, *, t)$ -EDGE_{LIN}3CUT using combinatorial properties that hold for the global version of the problem:

Theorem 1.6. *There exists an efficient $3/2$ -approximation algorithm for $(s, *, t)$ -EDGE_{LIN}3CUT.*

A problem closely related to linear 3-cut is the problem of simultaneously blocking in- and out- r -arborescences. We recall that an *out- r -arborescence* (*in- r -arborescence*) in a directed graph $D = (V, E)$ for a node $r \in V$ is a minimal subset $F \subseteq E$ of arcs such that every node $u \in V$ has a unique path from r to u (from u to r) in the subgraph (V, F) . In the blocking problem for in- and out- r -arborescences, denoted r -EDGE_{INOUT}BLOCKER, the input is a directed graph with a specified node r and the goal is to find the smallest number of edges to delete so that the resulting graph has no in- r -arborescence and no out- r -arborescence. We show that the approximability factor of this problem coincides with that of (s, r, t) -EDGE_{LIN}3CUT.

Theorem 1.7. *There exists an efficient α -approximation algorithm for r -EDGE_{INOUT}BLOCKER if and only if there exists an efficient α -approximation for (s, r, t) -EDGE_{LIN}3CUT.*

In contrast to (s, r, t) -EDGE_{LIN}3CUT, we do not have a hardness result for $(s, *, t)$ -EDGE_{LIN}3CUT. Our next result shows that the analogous problem in undirected graphs is in fact solvable in polynomial time. The analogous problem in undirected graphs is the following: given an undirected graph $G = (V, E)$ with two specified nodes $s, t \in V$, find the smallest subset of edges to delete so that there exists a node r with the property that s cannot reach r and t , and r cannot reach t in the resulting graph. We note that this is equivalent to min $\{s, t\}$ -separating 3-cut, denoted $\{s, t\}$ -SEP_{EDGE}3CUT: given an undirected graph with two specified nodes s, t , remove the smallest subset of edges so that the resulting graph has at least 3 connected components with s and t being in different components. This is a generalization of two efficiently solvable problems in undirected graphs, namely min $s - t$ cut and EDGE-3-CUT. The complexity of $\{s, t\}$ -SEP_{EDGE}3CUT was posed as an open problem by Queyranne [25]. More generally, we consider $\{s, t\}$ -SEP_{EDGE} k CUT, where the goal is to delete the smallest subset of edges from a given undirected graph so that the resulting graph has at least k connected components with s and t being in different components. We show that $\{s, t\}$ -SEP_{EDGE} k CUT is solvable in polynomial-time for every constant k .

Theorem 1.8. *For every constant k , there exists an efficient algorithm to solve $\{s, t\}$ -SEP_{EDGE} k CUT.*

While Theorem 1.5 shows that EDGE_{BIC}CUT is approximable to a factor strictly smaller than 2, we do not have a hardness result. One approach to solving EDGE_{BIC}CUT is by solving the single fixed terminal variant, denoted $\{s, *\}$ -EDGE_{BIC}CUT: Given a directed graph with a specified node s , find the smallest number of edges to delete so that there exists a node t such that s cannot reach t and t cannot reach s in the resulting graph. $\{s, *\}$ -EDGE_{BIC}CUT

admits a 2-approximation by guessing the terminal t and then using the 2-approximation for $\{s, t\}$ -EDGEBiCUT. We show the following inapproximability results for $\{s, *\}$ -EDGEBiCUT:

Theorem 1.9. $\{s, *\}$ -EDGEBiCUT is NP-hard, has no efficient $(100/99 - \epsilon)$ -approximation for any $\epsilon > 0$ assuming $P \neq NP$, and has no efficient $(4/3 - \epsilon)$ -approximation for any $\epsilon > 0$ assuming the Unique Games Conjecture.

Furthermore, motivated towards closing the approximability gap of NODEDOUBLECUT, we consider the node-weighted variant of bicut, denoted NODEBiCUT: Given a directed graph, find the smallest number of nodes whose deletion ensures that there exist nodes s and t such that s cannot reach t and t cannot reach s in the resulting graph. Every directed graph that is not a tournament has a feasible solution to NODEBiCUT. NODEBiCUT admits a 2-approximation by a simple reduction to $\{s, t\}$ -EDGEBiCUT. We show the following inapproximability results.

Theorem 1.10. NODEBiCUT is NP-hard, has no efficient $(53/52 - \epsilon)$ -approximation for any $\epsilon > 0$ assuming $P \neq NP$, and has no efficient $(3/2 - \epsilon)$ -approximation for any $\epsilon > 0$ assuming the Unique Games Conjecture.

We observe that our approximability and inapproximability results for NODEBiCUT and NODEDOUBLECUT coincide—2 and $3/2$ respectively (Theorems 1.2, 1.3 and 1.10).

We mention that all our algorithmic/approximation results hold for the weighted variant while the inapproximability results hold for the cardinality variant. For ease of presentation, we do not make this distinction.

1.1 Related Work

In recent work, Bernáth and Pap [2] studied the problem of deleting the smallest number of arcs to block all minimum cost arborescences of a given directed graph. They gave an efficient algorithm to solve this problem through combinatorial techniques. As a special case, their algorithm also solves $\{s, t\}$ -EDGEDOUBLECUT. There are two known techniques to solve $\{s, t\}$ -EDGEDOUBLECUT: a flow-based algorithm as well as a combinatorial algorithm based on the tree structure of *in-solid* sets. However, both techniques fail to extend to the node weighted variant $\{s, t\}$ -NODEDOUBLECUT.

The node-weighted 3-cut problem—NODE-3-CUT—is a generalization of the classic EDGE-3-CUT. Various other generalizations of EDGE-3-CUT have been studied in the literature showing the existence of efficient algorithms. These include the edge-weighted 3-cut in hypergraphs [30, 11] and the more general submodular 3-way partitioning [31, 24]. However, none of these known generalizations address NODE-3-CUT as a special case. Feasible solutions to NODE-3-CUT are also known as shredders in the node-connectivity literature. In the unit-weight case, shredders whose cardinality is equal to the node connectivity of the graph play a crucial role in the problem of min edge deletion to augment node connectivity by one [5, 14, 20, 29]. There are at most linear number of such shredders and all of them can be found efficiently [5, 14]. The complexity of finding a min cardinality shredder was open until our results (Theorem 1.4).

In the edge-weighted multiway cut in undirected graphs, the input is an undirected graph with k terminal nodes and the goal is to find the smallest cardinality subset of edges whose deletion ensures that there is no path between any pair of terminal nodes. For $k = 3$, a $12/11$ -approximation is known [6, 15], while for constant k , the current-best approximation factor is 1.2975 due to Sharma and Vondrák [27]. These results are based on an LP-relaxation proposed by Călinescu, Karloff and Rabani [8], known as the CKR relaxation. Manokaran, Naor, Raghavendra and Shwartz [21] showed that the inapproximability factor coincides with the integrality gap of the CKR relaxation. Recently, Angelidakis, Makarychev and Manurangsi [1] exhibited instances with integrality gap at least $6/(5 + (1/k - 1)) - \epsilon$ for every $k \geq 3$ and every $\epsilon > 0$.

The node-weighted multiway cut in undirected graphs exhibits very different structure in comparison to the edge-weighted multiway cut. It reduces to edge-weighted multiway cut in hypergraphs. Garg, Vazirani and Yannakakis [12] gave a $(2 - 2/k)$ -approximation by exploiting the extreme point structure of a natural LP-relaxation.

The edge-weighted multiway cut in directed graphs has a 2-approximation, due to Naor and Zosin [23], as well as Chekuri and Madan [3]. Matching inapproximability results were shown recently for $k = 2$ [19, 4]. The node-weighted multiway cut in directed graphs reduces to the edge-weighted multiway cut by exploiting the fact that the terminals are fixed. Such a reduction is unknown for the global version.

Preliminaries. Let $D = (V, E)$ be a directed graph. For two disjoint sets $X, Y \subset V$, we denote $\delta(X, Y)$ to be the set of edges (u, v) with $u \in X$ and $v \in Y$ and $d(X, Y)$ to be the cut value $|\delta(X, Y)|$. We use $\delta^{in}(X) := \delta(V \setminus X, X)$, $\delta^{out}(X) := \delta(X, V \setminus X)$, $d^{in}(X) := |\delta^{in}(X)|$ and $d^{out}(X) := |\delta^{out}(X)|$. We use a similar notation for undirected graphs by dropping the superscripts. For two nodes $s, t \in V$, a subset $X \subset V$ is called an $\bar{s}t$ -set if $t \in X \subseteq V - s$. The *cut value* of an $\bar{s}t$ -set X is $d^{in}(X)$.

We frequently use the following characterization of directed graphs with no arborescence for the purposes of double cut.

Theorem 1.11. (e.g., see [2]) *Let $D = (V, E)$ be a directed graph. The following are equivalent:*

1. D has no arborescence.
2. There exist two distinct nodes $s, t \in V$ such that every node u can reach at most one node in $\{s, t\}$ in D .
3. There exist two disjoint non-empty sets $S, T \subset V$ with $\delta^{in}(S) \cup \delta^{in}(T) = \emptyset$.

2 Approximation for NodeDoubleCut

In this section, we present an efficient 2-approximation algorithm for $\{s, t\}$ -NODEDOUBLECUT which also leads to a 2-approximation for NODEDOUBLECUT by guessing the pair of nodes s, t . Our algorithm is LP-based. At the end of this section, we give an example showing that the integrality gap of the LP nearly matches the approximation factor achieved by our rounding algorithm.

Proof of Theorem 1.2. We recall the problem: Given a directed graph $D = (V, E)$ with two specified nodes $s, t \in V$ and node costs $c : V \setminus \{s, t\} \rightarrow \mathbb{R}_+$, the goal is to find a least cost subset $U \subseteq V \setminus \{s, t\}$ of nodes such that every node $u \in V \setminus U$ can reach at most one node in $\{s, t\}$ in the subgraph $D - U$. We will denote a path P by the set of nodes in the path and the collection of paths from node u to node v by $\mathcal{P}^{u \rightarrow v}$. For a fixed function $d : V \rightarrow \mathbb{R}_+$, the d -distance of a path P is defined to be $\sum_{u \in P} d_u$ and the shortest d -distance from node u to node v is the minimum d -distance among all paths from node u to node v . We use the following LP-relaxation, where we have a variable d_u for every node $u \in V$:

$$\begin{aligned}
 & \min \sum_{v \in V \setminus \{s, t\}} c_v d_v && \text{(Path-Blocking-LP)} \\
 & \sum_{v \in P} d_v + \sum_{v \in Q} d_v - d_u \geq 1 \quad \forall P \in \mathcal{P}^{u \rightarrow s}, Q \in \mathcal{P}^{u \rightarrow t}, \forall u \in V \\
 & d_s, d_t = 0 \\
 & d_v \geq 0 \quad \forall v \in V
 \end{aligned}$$

We first observe that Path-Blocking-LP can be solved efficiently. The separation problem is the following: given $d : V \rightarrow \mathbb{R}_+$, verify if there exists a node $u \in V$ such that the sum of the shortest d -distance path from u to s and the shortest d -distance path from u to t is at most $1 + d_u$. Thus, the separation problem can be solved efficiently by solving the shortest path problem in directed graphs.

Let $d : V \rightarrow \mathbb{R}_+$ be a feasible solution to Path-Blocking-LP. We now present a rounding algorithm that achieves a 2-factor approximation. We note that our algorithm rounds an arbitrary feasible solution d to obtain an integral solution whose cost is at most twice the LP-cost of the solution d . For a subset U of nodes, let $\Delta^{in}(U)$ be the set of nodes $v \in V \setminus U$ that have an edge to a node $u \in U$.

Rounding Algorithm for $\{s, t\}$ -NodeDoubleCut

1. Pick θ uniformly from the interval $(0, 1/2)$.
2. Let $\mathbb{B}^{in}(s, \theta)$ and $\mathbb{B}^{in}(t, \theta)$ be the set of nodes whose shortest d -distance to s and t respectively, is at most θ .
3. Return $U := \Delta^{in}(\mathbb{B}^{in}(s, \theta)) \cup \Delta^{in}(\mathbb{B}^{in}(t, \theta))$.

The rounding algorithm can be implemented to run in polynomial-time. We first show the feasibility of the solution returned by the rounding algorithm. We use the following claim.

Claim 2.1. *For every $\theta \in (0, 1/2)$, we have $\mathbb{B}^{in}(s, \theta) \cap \mathbb{B}^{in}(t, \theta) = \emptyset$.*

Proof. Say $u \in \mathbb{B}^{in}(s, \theta) \cap \mathbb{B}^{in}(t, \theta)$. Then there exists a path $P \in \mathcal{P}^{u \rightarrow s}$ and a path $Q \in \mathcal{P}^{u \rightarrow t}$ such that $\sum_{v \in P} d_v + \sum_{v \in Q} d_v \leq 2\theta < 1$, a contradiction to the fact that d is feasible for Path-Blocking-LP. \square

Claim 2.2. *The solution U returned by the algorithm is such that every node $u \in V \setminus U$ can reach at most one node in $\{s, t\}$ in the subgraph $D - U$.*

Proof. Suppose not. Then there exists $u \in V \setminus U$ that can reach both s and t in $D - U$. If $u \notin \mathbb{B}^{in}(s, \theta)$, then u cannot reach s in $D - U$ since $\mathbb{B}^{in}(s, \theta)$ has no entering edges in $D - U$. Thus, $u \in \mathbb{B}^{in}(s, \theta)$. Similarly, $u \in \mathbb{B}^{in}(t, \theta)$. However, this contradicts the above claim that $\mathbb{B}^{in}(s, \theta) \cap \mathbb{B}^{in}(t, \theta) = \emptyset$. \square

We next bound the expected cost of the solution returned by the rounding algorithm. Let $\bar{d}(v, a)$ denote the shortest d -distance from node v to node a in D . We use the following claim.

Claim 2.3. *Let $\theta \in (0, 1/2)$. If $v \in \Delta^{in}(\mathbb{B}^{in}(s, \theta))$ then $\theta < \bar{d}(v, s) \leq \theta + d_v$ and $d_v \neq 0$.*

Proof. If $\bar{d}(v, s) \leq \theta$, then $v \in \mathbb{B}^{in}(s, \theta)$, a contradiction to $v \in \Delta^{in}(\mathbb{B}^{in}(s, \theta))$. If $\bar{d}(v, s) > \theta + d_v$, then $v \notin \Delta^{in}(\mathbb{B}^{in}(s, \theta))$, a contradiction. If $d_v = 0$, then $\theta < \bar{d}(v, s) \leq \theta + d_v = \theta$, a contradiction. \square

Claim 2.4. *For every $v \in V$, the probability that v is chosen in U is at most $2d_v$.*

Proof. The claim holds if $v \in \{s, t\}$. Let us fix $v \in V \setminus \{s, t\}$. By the claim above, if $v \in \Delta^{in}(\mathbb{B}^{in}(s, \theta))$ then $\theta < \bar{d}(v, s) \leq \theta + d_v$ and $d_v \neq 0$. Similarly, if $v \in \Delta^{in}(\mathbb{B}^{in}(t, \theta))$, then $\theta < \bar{d}(v, t) \leq \theta + d_v$ and $d_v \neq 0$. Now, the probability that v is in U is at most

$$\Pr(\theta \in (\bar{d}(v, s) - d_v, \min\{\bar{d}(v, s), 1/2\}) \cup (\bar{d}(v, t) - d_v, \min\{\bar{d}(v, t), 1/2\})).$$

Without loss of generality, let $\bar{d}(v, s) \leq \bar{d}(v, t)$. We may assume that $d_v > 0$ and $\bar{d}(v, s) - d_v < 1/2$, since otherwise, the probability that v is in U is 0 and the claim is proved. Now, by the




Figure 1: $D_{a,b}$ in the proof of Lemma 2.1 and $(2-\epsilon)$ -inapproximability of $\{s,t\}$ -NODEDOUBLE-CUT.

feasibility of the solution d to Path-Blocking-LP, we have that $\bar{d}(v,s) + \bar{d}(v,t) - d_v \geq 1$ and hence $\bar{d}(v,t) \geq 1/2$. Therefore,

$$\begin{aligned} \Pr(v \in U) &\leq \Pr(\theta \in (\bar{d}(v,s) - d_v, \min(\bar{d}(v,s), 1/2))) + \Pr(\theta \in (\bar{d}(v,t) - d_v, 1/2)) \\ &= \frac{1}{(1/2)} (1/2 - \bar{d}(v,s) + d_v + 1/2 - \bar{d}(v,t) + d_v) \\ &= 2(1 - (\bar{d}(v,s) + \bar{d}(v,t) - d_v) + d_v) \\ &\leq 2d_v. \end{aligned}$$

The first equality in the above is because θ is chosen uniformly from the interval $(0, 1/2)$ while the last inequality is because of the feasibility of the solution d to Path-Blocking-LP. \square

By the above claim, the expected cost of the returned solution is

$$\mathbb{E} \left(\sum_{v \in U} c_v \right) = \sum_{v \in V} \Pr(v \in U) c_v \leq 2 \sum_{v \in V} c_v d_v.$$

Although our rounding algorithm is a randomized algorithm, it can be derandomized using standard techniques. \square

Our next lemma shows a lower bound on the integrality gap that nearly matches the approximation factor achieved by our rounding algorithm.

Lemma 2.1. *The integrality gap of the Path-Blocking-LP for directed graphs containing n nodes is at least $2 - 7/n^{1/3}$.*

Our integrality gap instance is also helpful in the hardness of approximation proofs. So, we define the instance below and prove certain properties of the instance which will be used later.

For two integers $a, b \in \mathbb{N}$, consider the directed graph $D_{a,b} = (V_D, A_D)$ obtained as follows (see Figure 1): Let $V_D := \{s, t\} \cup ([a] \times [b])$. There are $ab + 2$ nodes. Let $I_D := [a] \times [b]$ and call them as the *internal nodes*. The set of arcs A_D are as follows:

1. For each $1 \leq i \leq a$, there is a bidirected arc between s and $(i, 1)$, and a bidirected arc between (i, b) and t .

2. For each $1 \leq i \leq a$ and $1 \leq j < b$, there is a bidirected arc between (i, j) and $(i, j + 1)$.
3. For each $1 \leq i < a$ and $2 \leq j \leq b - 1$, there is an arc from (i, j) to $(i + 1, j - 2)$, and an arc from (i, j) to $(i + 1, j + 2)$ (let $(i, 0) := s$ and $(i, b + 1) := t$ for every i). Call them *jumping arcs*.

Lemma 2.2. $D_{a,b}$ has the following properties:

1. For each internal node $\alpha = (\alpha_1, \alpha_2) \in I_D$, each $\alpha \rightarrow s$ path has at least $\alpha_2 - a$ internal nodes other than α . Similarly, each $\alpha \rightarrow t$ path has at least $b - \alpha_2 - a + 1$ internal nodes other than α .
2. If $S \subseteq I_D$ is such that the subgraph induced by $V_D \setminus S$ has no node v that has paths to both s and t , then $|S| \geq 2a - 1$.

Proof. 1. Jumping arcs are the only arcs that change α_2 by 2 while all other arcs change α_2 by 1. However, a path to s can use at most $a - 1$ jumping arcs because they strictly increase α_1 . The first property follows from these observations.

2. Suppose that $S \subseteq I_D$ is such that the subgraph induced by $V_D \setminus S$ has no node v that has paths to both s and t . For $i = 1, \dots, a$, let $s_i := |S \cap \{\{i\} \times [b]\}|$. We note that $s_i \geq 1$ for each i , otherwise s can reach t and t can reach s .

Suppose $s_i = 1$ for some $1 < i \leq a$ and let j be such that $S \cap \{\{i\} \times [b]\} = (i, j)$. If $j = 1$, then $(i, 2) \in V_D \setminus S$ and $(i, 2)$ can reach both s and t . If $j = b$, then $(i, b - 1) \in V_D \setminus S$ and $(i, b - 1)$ can reach both s and t . Therefore, we have $1 < j < b$. Then $s_{i-1} \geq 3$ because $(i - 1, j - 1), (i - 1, j), (i - 1, j + 1)$ can reach both s and t using one jumping arc followed by regular arcs in the i th row.

Therefore, $|S| = \sum_{i=1}^a s_i \geq 1 + 2(a - 1) = 2a - 1$. □

Proof of Lemma 2.1. The integer optimum of Path-Blocking-LP on $D_{a,b}$ is at least $2a - 1$ by the second property of Lemma 2.2. Let $r := b - 2a + 1$. We set $d_v := 1/r$ for every internal node v . The resulting solution is feasible to Path-Blocking-LP: Indeed, consider $\alpha = (\alpha_1, \alpha_2)$. By the first property of Lemma 2.2, any $\alpha \rightarrow s$ path and $\alpha \rightarrow t$ path have to together traverse at least $\alpha_2 - a + (b - \alpha_2 - a + 1) = r$ internal nodes.

Setting $b = a^2$, the integrality gap is at least $(2a - 1)/(a^3/r) = 2 - 1/a^3 + 4/a^2 - 5/a \geq 2 - 6/a$ for $a \geq 2$. Using the fact that $a = (|V(D_{a,b})| - 2)^{1/3}$, we get the desired bound on the integrality gap. □

3 Hardness of Approximation

In this section, we prove the hardness results, namely Theorem 1.1 for $\{s, t\}$ -NODEDOUBLE-CUT, Theorem 1.3 for NODEDOUBLECUT, Theorem 1.4 for NODE-3-CUT, Theorem 1.9 for $\{s, *\}$ -EDGEBICUT, and Theorem 1.10 for NODEBICUT. All our reductions begin from VERTEXCOVER ON k -REGULAR GRAPHS, where the input is an undirected k -regular graph, and the goal is to find the smallest subset S of nodes such that every edge in the graph has at least one end-vertex in S .

We use VERTEXCOVER ON k -PARTITE GRAPHS as an intermediate problem, where the input is an undirected k -partite graph $G = (V_1 \cup \dots \cup V_k, E)$ (we emphasize that the partitioning V_1, \dots, V_k is specified explicitly in the input) and the goal is to find the smallest subset $S \subset V_1 \cup \dots \cup V_k$ such that every edge in E has at least one end-vertex in S . Our hardness results are structured as follows.

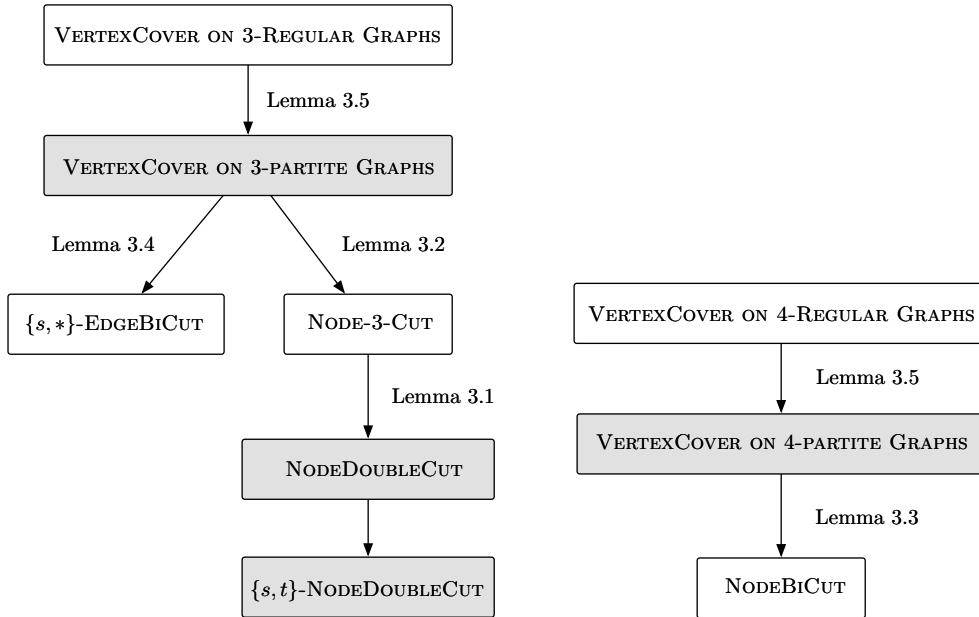


Figure 2: Approximation-preserving reductions in Section 3.1. Our NP-hardness results follow from these reductions and the hardness of VERTEXCOVER in 3-regular and 4-regular graphs [7]. For the problems in shaded cells, we show direct reductions from UNIQUEGAMES to prove stronger hardness results assuming the Unique Games Conjecture.

1. We first show approximation-preserving (combinatorial) reductions from VERTEXCOVER ON k -REGULAR GRAPHS (for $k = 3$ or 4) to the above-mentioned problems in Section 3.1 (see Fig. 2). These reductions prove all the inapproximability results under the assumption that $P \neq NP$.
2. For improved hardness of approximation results, we show that VERTEXCOVER ON k -PARTITE GRAPHS is hard to approximate within a factor of $2(k-1)/k - \epsilon$ for any $\epsilon > 0$ assuming the Unique Games Conjecture (Section 3.5). Considering $k = 3$ and $k = 4$, this result in conjunction with the combinatorial reductions show $(4/3 - \epsilon)$ -inapproximability for NODEDOUBLECUT and $\{s, *\}$ -EDGEBiCUT and $(3/2 - \epsilon)$ -inapproximability for NODEBiCUT assuming the Unique Games Conjecture.
3. We further improve the hardness of approximation for NODEDOUBLECUT and $\{s, t\}$ -NODEDOUBLECUT by directly reducing from UNIQUEGAMES via the *length-control dictatorship tests* introduced in [19]. We obtain $(3/2 - \epsilon)$ -inapproximability for NODEDOUBLECUT in Section 3.3 and $(2 - \epsilon)$ -inapproximability for $\{s, t\}$ -NODEDOUBLECUT in Section 3.4.

We show the combinatorial reductions in Section 3.1. We summarize the preliminaries on discrete Fourier analysis that we need for hardness based on UNIQUEGAMES in Section 3.2. Similar to most hardness results based on the UNIQUEGAMES, the main technical contribution of our work is the construction of *dictatorship tests*. In order to avoid repetition, we present the dictatorship tests for the respective problems in Sections 3.3, 3.4, and 3.5, and show the full reduction from UNIQUEGAMES to all four problems in Section 3.6.

3.1 Combinatorial Reductions

Lemma 3.1. *For every $\alpha \geq 1$, an α -approximation algorithm for NODEDOUBLECUT implies an α -approximation algorithm for NODE-3-CUT.*

Proof. Given an instance $G = (V, E)$ of NODE-3-CUT, we use the following algorithm:

1. For each $s \in V$,
 - (a) Construct an instance of NODEDOUBLECUT $D = (V, A)$ as follows:
 - i. For each edge $\{u, v\} \in E$, add a bidirected arc between u and v .
 - ii. For each vertex $v \in V \setminus \{s\}$, add an arc from v to s .
 - iii. Set the weight of s to ∞ and the weight of all other vertices to 1.
 - (b) Run the α -approximation algorithm for NODEDOUBLECUT on D and obtain the solution U_s .
2. Output the solution U_s with the smallest size.

Let $S \subseteq V$ be an optimal solution of NODE-3-CUT, so that $V \setminus S$ is partitioned into V_1, V_2, V_3 such that there are no edges between V_i and V_j in $G - S$.

Consider $s \in V_1$. Let D be the instance of NODEDOUBLECUT generated for s . We first show that S is a feasible solution for NODEDOUBLECUT in D . Indeed, consider $u \in V_2$ and $v \in V_3$. In the subgraph of D induced by $V \setminus S$, no vertex can reach both u and v , since vertices in V_i have outgoing arcs only to vertices in V_i or to s , and s can only reach vertices in V_1 .

Therefore, the α -approximation algorithm for NODEDOUBLECUT will yield $T \subseteq V \setminus \{s\}$ such that $T \leq \alpha \cdot |S|$ and there exist two vertices $u, v \in V \setminus T$ such that no vertex in $V \setminus T$ can reach both u and v . Let U_2 be the set of vertices strongly connected to u , and U_3 be the set of vertices strongly connected to v . They must be disjoint since otherwise u and v can reach each other. Furthermore, neither U_2 nor U_3 can contain s , since otherwise v can reach u (if $s \in U_2$) or u can reach v (if $s \in U_3$).

Since every edge of G becomes a bidirected arc in D , if some vertices x and y are connected in the subgraph of G induced by $V \setminus T$, then they are strongly connected in the subgraph of D induced by $V \setminus T$. This implies that in the subgraph of G induced by $V \setminus T$, both U_2 and U_3 are disconnected from the rest of the graph. Therefore, U_2 and U_3 form a (union of) disjoint connected components, but their union is not $V \setminus T$ (since s is not contained). This implies that T is a feasible solution to NODE-3-CUT in G , finishing the proof. □

Lemma 3.2. *There is an approximation-preserving reduction from VERTEXCOVER ON 3-PARTITE GRAPHS to NODE-3-CUT.*

Proof. Given a 3-partite graph $G = (V_1 \cup V_2 \cup V_3, E)$, we construct an instance G' for NODE-3-CUT by adding three vertices s_1, s_2, s_3 of infinite weight (all other vertices have weight 1), and for all $i \in [3]$ and $v \in V_i$, adding an edge between s_i and v . Then a subset $S \subseteq V_1 \cup V_2 \cup V_3$ is a vertex cover in G if and only if $G' - S$ has at least three connected components. □

Lemma 3.3. *There is an approximation-preserving reduction from VERTEXCOVER ON 4-PARTITE GRAPHS to NODEBICUT.*

Proof. Given a 4-partite graph $G = (V_1 \cup V_2 \cup V_3 \cup V_4, E)$, we construct an instance $D = (V_D, A_D)$ for NODEBICUT as follows: Let $V_D := V_1 \cup V_2 \cup V_3 \cup V_4 \cup \{s, t\}$. The set of arcs A_D are obtained as follows:

1. For every $u, v \in V_i$ for some $i \in [4]$, we add a bidirected arc between u and v .
2. For every $(u, v) \in E$, we add a bidirected arc between u and v .
3. For every $u \in V_1$, we add a bidirected arc between s and u .
4. For every $u \in V_2$, we add an arc from s to u and an arc from t to u .

5. For every $u \in V_3$, we add an arc from u to s and an arc from u to t .
6. For every $u \in V_4$, we add a bidirected arc between t and u .

We now show the completeness of the reduction. Suppose $R \subseteq V_1 \cup V_2 \cup V_3 \cup V_4$ is a vertex cover in G . Then $D - R$ has no $s \rightarrow t$ path, since s can only reach vertices in V_1 and V_2 , only vertices in V_3 and V_4 can reach t , and there is no arc between V_i and V_j for any $i \neq j$. Similarly, there is no $t \rightarrow s$ path. Therefore, R is a feasible solution to NODEBICUT in D .

Next we show soundness of the reduction. Suppose $R \subseteq V_1 \cup V_2 \cup V_3 \cup V_4$ is a feasible solution to NODEBICUT in D . There exists two vertices $u, v \in V_D \setminus R$ such that there is no $u \rightarrow v$ path and no $v \rightarrow u$ path in the subgraph of D induced by $V_D \setminus R$. We note that v and u cannot be in the same V_i since V_i is a clique in V_D . We also rule out the following cases:

1. If $v \in V_1, u \in V_2$, then (v, s, u) is a path from v to u , a contradiction.
2. If $v \in V_1, u \in V_3$, then (u, s, v) is a path from u to v , a contradiction.
3. If $v \in V_2, u \in V_3$, then (u, s, v) is a path from u to v , a contradiction.
4. If $v \in V_2, u \in V_4$, then (u, t, v) is a path from u to v , a contradiction.
5. If $v \in V_3, u \in V_4$, then (v, t, u) is a path from v to u , a contradiction.

Thus, $v \in V_1$ and $u \in V_4$. We will show that if R is not a vertex cover, then there is a $v \rightarrow u$ path or $u \rightarrow v$ path, a contradiction. Suppose there exists $\{a, b\} \in E$ such that $a, b \notin R$.

1. If $a \in V_1, b \in V_2$, then (u, t, b, a, v) is a path from u to v , a contradiction.
2. If $a \in V_1, b \in V_3$, then (v, a, b, t, u) is a path from v to u , a contradiction.
3. If $a \in V_1, b \in V_4$, then (v, a, b, u) is a path from v to u , a contradiction.
4. If $a \in V_2, b \in V_3$, then (v, s, a, b, t, u) is a path from v to u , a contradiction.
5. If $a \in V_2, b \in V_4$, then (v, s, a, b, u) is a path from v to u , a contradiction.
6. If $a \in V_3, b \in V_4$, then (u, b, a, s, v) is a path from u to v , a contradiction.

Therefore, R must be a vertex cover. This establishes the soundness of the reduction and completes the proof. \square

Lemma 3.4. *There is an approximation-preserving reduction from VERTEXCOVER ON 3-PARTITE GRAPHS to $\{s, *\}$ -EDGEBICUT.*

Proof. Given a 3-partite graph $G = (A \cup B \cup C, E)$, we construct an instance $D = (V_D, A_D)$ for $\{s, *\}$ -EDGEBICUT as follows: Let $V_D := A_1 \cup A_2 \cup B_1 \cup B_2 \cup C_1 \cup C_2 \cup \{s, t\}$. For a vertex $v \in A \cup B \cup C$ and $i \in \{1, 2\}$, let v_i denote the corresponding vertex in V (e.g., if $v \in A$, then $v_1 \in A_1$ and $v_2 \in A_2$). We introduce three types of arcs in A_D .

1. Vertex arcs: For every $v \in A \cup B \cup C$, create an arc (v_1, v_2) with weight 1.
2. Forward arcs: Create arcs with weight ∞
 - (a) (s, a_1) for all $a \in A$, (s, b_1) for all $b \in B$, (b_2, s) for all $b \in B$, (c_2, s) for all $c \in C$.
 - (b) (t, a_1) for all $a \in A$, (c_2, t) for all $c \in C$.
 - (c) (a_2, b_1) for every $\{a, b\} \in E$, $a \in A, b \in B$ (call them AB arcs), (a_2, c_1) for every $\{a, c\} \in E$, $a \in A, c \in C$ (call them AC arcs), (b_2, c_1) for every $\{b, c\} \in E$, $b \in B, c \in C$ (call them BC arcs).

3. Backward arcs: Create arcs with weight ∞

- (a) (v_2, u_1) for all $u, v \in A$ (call them AA arcs), (v_2, u_1) for all $u, v \in C$ (call them CC arcs), (c_1, a_1) for all $a \in A, c \in C$ (call them CA_1 arcs), (c_2, a_2) for all $a \in A, c \in C$ (call them CA_2 arcs).

We first show completeness of the reduction. Suppose $R \subseteq A \cup B \cup C$ is a vertex cover in G . Let $F = \{(v_1, v_2) : v \in R\}$. We will show that there is no $s \rightarrow t$ path and no $t \rightarrow s$ path in $D - F$.

1. Suppose there is a $t \rightarrow s$ path in $D - F$. Fix the shortest such $t \rightarrow s$ path P . Then, the path P has the following properties:

- (a) Path P does not contain AA arcs or CA_1 arcs, since t has direct arcs to vertices in A_1 . Similarly, P does not contain CC arcs or CA_2 arcs, since vertices in C_2 have direct arcs to s . So, P does not contain any backward arcs.
- (b) Path P does not contain BC arcs, since vertices in B_2 have direct arcs to s .

Thus, the only possibility for the path P is $P = (t, a_1, a_2, v_1, v_2, s)$ for $a \in A, v \in B \cup C$, and $\{a, v\} \in E$. This contradicts that R is a vertex cover.

2. Suppose there is a $s \rightarrow t$ path in $D - F$. Fix the shortest such $s \rightarrow t$ path P . Then, the path P has the following properties:

- (a) Path P does not contain AA arcs or CA_1 arcs, since s has direct arcs to vertices in A_1 . Similarly, P does not contain CC arcs or CA_2 arcs since vertices in C_2 have direct arcs to t . So, P does not contain any backward arcs.
- (b) Path P does not contain AB arcs, since s has direct arcs to vertices in B_1 .

Thus, the only possibility for the path P is $P = (t, v_1, v_2, c_1, c_2, s)$ for $v \in A \cup B, c \in C$, and $\{v, c\} \in E$. This contradicts that R is a vertex cover.

Therefore, s and t cannot reach each other in $D - F$. Consequently, the existence of a vertex cover R in G implies the existence of a feasible solution to $\{s, *\}$ -EDGEbicUT in D of the same size.

Next we show soundness of the reduction. Suppose $F \subseteq E_D$ is a feasible solution to $\{s, *\}$ -EDGEbicUT in D . Let $R \subseteq A \cup B \cup C$ be the set of vertices whose vertex arcs are in F . We will show that if R is not a vertex cover in G , then every vertex $v \in V_D$ has either a path to s or a path from s . Since vertices in A_1, B_1, B_2, C_2 have a direct arc either from or to s , we only need to check vertices in A_2, C_1 and t . We verify these cases below:

1. Suppose there exist $a \in A \setminus R, b \in B \setminus R$ such that $\{a, b\} \in E$.

- (i) Considering t , we have $(t, a_1, a_2, b_1, b_2, s)$ as a path from t to s .
- (ii) For every $a' \in A_2$, we have $(a', a_1, a_2, b_1, b_2, s)$ as a path from a' to s .
- (iii) For every $c' \in C_1$, we have $(c', a_1, a_2, b_1, b_2, s)$ as a path from c' to s .

2. Suppose there exist $a \in A \setminus R, c \in C \setminus R$ such that $\{a, c\} \in E$.

- (i) Considering t , we have $(t, a_1, a_2, c_1, c_2, s)$ as a path from t to s .
- (ii) For every $a' \in A_2$, we have $(a', a_1, a_2, c_1, c_2, s)$ as a path from a' to s .
- (iii) For every $c' \in C_1$, we have $(c', a_1, a_2, c_1, c_2, s)$ as a path from c' to s .

3. Suppose there exist $b \in B \setminus R, c \in C \setminus R$ such that $\{b, c\} \in E$.

- (i) Considering t , we have $(s, b_1, b_2, c_1, c_2, t)$ as a path from s to t .
- (ii) For every $a' \in A_2$, we have $(s, b_1, b_2, c_1, c_2, a')$ as a path from s to a' .
- (iii) For every $c' \in C_1$, we have $(s, b_1, b_2, c_1, c_2, c')$ as a path from s to c' .

Therefore, the existence of a feasible solution to $\{s, *\}$ -EDGEBICUT in D implies the existence of a vertex cover in G of the same size. This establishes the soundness of the reduction, and proves the lemma. □

The following lemma proves hardness of approximation for VERTEXCOVER ON k -PARTITE GRAPHS for $k = 3, 4$. We use the result of Chlebík and Chlebíková [7] on the hardness of VERTEXCOVER on 3-regular and 4-regular graphs.

Lemma 3.5. *For every $\epsilon > 0$, VERTEXCOVER ON 3-PARTITE GRAPHS and VERTEXCOVER ON 4-PARTITE GRAPHS are NP-hard to approximate within a factor $100/99 - \epsilon$ and $53/52 - \epsilon$ respectively.*

Proof. Chlebík and Chlebíková [7] proved that VERTEXCOVER on 3-regular and 4-regular graphs are NP-hard to approximate within a factor of $100/99 - \epsilon$ and $53/52 - \epsilon$ respectively for any $\epsilon > 0$.

For any k -regular graph ($k = 3, 4$ in our case), we recall that Brooks' theorem gives an efficient algorithm to find a k -coloring (except when the graph is K_k for which it gives a $(k + 1)$ -coloring), which gives a k -partition of the graph. Therefore, any α -approximation algorithm for VERTEXCOVER ON k -PARTITE GRAPHS can be used to get an α -approximation for VERTEXCOVER on k -regular graphs. The lemma follows. □

3.2 Preliminaries for Unique Games Hardness

Gaussian Bounds for Correlated Spaces. We introduce the standard tools on correlated spaces from Mossel [22]. 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 $I \subseteq \mathbb{R}$, let $\mathcal{L}_I(\Omega)$ be the set of functions $\{f : \Omega \rightarrow I\}$. For two functions $f, g : \Omega \rightarrow \mathbb{R}$, we define

$$\begin{aligned}\mathbb{E}[f] &:= \mathbb{E}_x[f(x)], \\ \text{Var}[f] &:= \mathbb{E}_x[(f(x) - \mathbb{E}[f])^2], \\ \text{Cov}[f, g] &:= \mathbb{E}_x[(f(x) - \mathbb{E}[f])(g(x) - \mathbb{E}[g])],\end{aligned}$$

where x is sampled from the probability space (Ω, μ) .

For a subset $S \subseteq \Omega$, we define the *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 $\|f\|_p := \mathbb{E}_{x \sim \mu}[|f(x)|^p]^{1/p}$.

Given (Ω, μ) , let $(\Omega^R, \mu^{\otimes R})$ be the *product space* where for $(x_1, \dots, x_R) \in \Omega^R$, $\mu^{\otimes R}(x_1, \dots, x_R) = \prod_{i=1}^R \mu(x_i)$. Consider $f \in \mathcal{L}(\Omega^R)$. The *Efron-Stein decomposition* of f is given by

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




Figure 3: D in the proof of $(3/2 - \epsilon)$ -inapproximability of NODEDOUBLECUT.

where (1) $f_S : \Omega^R \rightarrow \mathbb{R}$ depends only on $\{x_i\}_{i \in S}$ 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 3.1 (Theorem 6.3 and Lemma 6.6 of [22]). *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 $\mathbb{E}_{(x,y) \in \mu^{\otimes R}}[f(x)g(y)] \geq \Gamma_\rho(\mathbb{E}_x[f], \mathbb{E}_y[g]) - \epsilon$.*

3.3 $(3/2 - \epsilon)$ -Inapproximability for NodeDoubleCut

Consider the directed graph $D = (V_D, A_D)$ (see Figure 3) defined by

$$\begin{aligned} V_D &:= \{s, t, a, b, c, d\}, \\ A_D &:= \{(a, s), (s, a), (s, c), (c, a), (a, b), (b, c), (c, b), (d, c), (b, d), (d, t), (t, d), (t, b)\}. \end{aligned}$$

Let $I_D := \{a, b, c, d\}$ be the set of internal vertices.

We summarize the properties of D that can be verified easily.

Proposition 3.1. *D has the following three properties.*

- (i) *For any vertex $v \in V$, there exists a vertex $u \in \{s, t\}$ such that every $v \rightarrow u$ path has at least three internal vertices.*
- (ii) *Every $v \in I_D$ has an incoming arc from either s or t .*

(iii) Even after deleting one vertex from I_D , there exists an $s \rightarrow t$ path or a $t \rightarrow s$ path with exactly three remaining internal vertices.

Based on D , we define the *dictatorship test* graph $\mathcal{D}_{R,\epsilon}^{\text{global}} = (V, A)$ as follows, for a positive integer R and $\epsilon > 0$. It will be used to show hardness results under the Unique Games Conjecture in Section 3.6. Let $r = 3$. Consider the probability space (Ω, μ) where $\Omega := \{0, \dots, r-1, *\}$, and $\mu : \Omega \mapsto [0, 1]$ with $\mu(*) = \epsilon$ and $\mu(x) = (1 - \epsilon)/r$ for $x \neq *$.

1. We take $V := \{s, t\} \cup \{v_x^\alpha\}_{\alpha \in I_D, x \in \Omega^R}$. Let v^α denote the set of vertices $\{v_x^\alpha\}_{x \in \Omega^R}$.
2. For $\alpha \in I_D$ and $x \in \Omega^R$, we define the weight as $c(v_x^\alpha) := \mu^{\otimes R}(x)$. We note that the sum of weights is 4. The terminals s and t have infinite weight.
3. There are arcs from s to all vertices in v^c , from v^a to s , s to v^a , from v^d to t , from t to v^d , from t to v^b .
4. For each $(\alpha, \beta) \in \{(c, a), (a, b), (b, c), (c, b), (d, c), (b, d)\}$ and $x, y \in \Omega^R$, we have an arc from v_x^α to v_y^β if there exists $1 \leq j \leq R$ such that $y_j = (x_j + 1) \bmod r$ or $y_j = *$ or $x_j = *$.

Completeness. We first prove that removing a set of vertices that correspond to *dictators* behaves the same as the fractional solution that gives $1/r$ to every internal vertex. For any $q \in [R]$, let $V_q := \{v_x^\alpha : \alpha \in I_D, x_q = * \text{ or } 0\}$. We note that the total weight of V_q is $4(\epsilon + (1 - \epsilon)/r) = 4(1 + 2\epsilon)/3$.

Lemma 3.6. *After removing vertices in V_q , no vertex in V can reach both s and t .*

Proof. Suppose towards contradiction that there exists a vertex that can reach both s and t . First, assume that this vertex is $v_{x_0}^{\alpha_0}$ for some $\alpha_0 \in I_D$ and $x_0 \in \Omega^R$. By Property (i) of Proposition 3.1, there exists $u \in \{s, t\}$ such that every $\alpha_0 \rightarrow u$ path has at least three internal vertices in D . Let $(v_{x_0}^{\alpha_0}, v_{x_1}^{\alpha_1}, \dots, v_{x_k}^{\alpha_k}, u)$ be a path from $v_{x_0}^{\alpha_0}$ to u in $\mathcal{D}_{R,\epsilon}^{\text{global}} - V_q$. Note that $k \geq 2$.

Consider the sequence $((x_0)_q, (x_1)_q, \dots, (x_k)_q)$. Recall that v_x^α has an arc to v_y^β for some α, β, x, y only if $y_q = (x_q + 1) \bmod r$ or $y_q = *$ or $x_q = *$. Since we removed V_q , $(x_i)_q \notin \{0, *\}$, $(x_i)_q = (x_{i-1})_q + 1$. This forces $k \leq 1$, leading to contradiction.

Finally, assume that s can reach t , and let $(s, v_{x_0}^{\alpha_0}, v_{x_1}^{\alpha_1}, \dots, v_{x_k}^{\alpha_k}, t)$ be a $s \rightarrow t$ path for some $\alpha_i \in I_D$, $x_i \in \Omega^R$. Every $s \rightarrow t$ path in D has to have at least three internal vertices, which forces $k \geq 2$, but considering the sequence $((x_0)_q, (x_1)_q, \dots, (x_k)_q)$ forces $k \leq 1$, which leads to contradiction. Paths from t to s can be ruled out in the same way. \square

Soundness. Suppose that we removed some vertices C such that there exist two vertices $u, v \in V \setminus C$ where no vertex $w \in V \setminus C$ can reach both u and v . This implies that no vertex $w \in V \setminus C$ can reach both s and t , since both u and v have an incoming arc from either s or t . Therefore, it suffices to show that unless C reveals an influential coordinate or $c(C) \geq 2(1 - \epsilon)$, either s can reach t or t can reach s .

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 $(x, y) \in \Omega^2$.

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

We note that the marginal distribution of both x and y is equal to μ . Assuming $\epsilon < 1/2r$, the minimum probability of any atom in $\Omega_1 \times \Omega_2$ is ϵ^2 . We use the following lemma to bound the correlation $\rho(\Omega_1, \Omega_2; \nu)$.

Lemma 3.7 (Lemma 2.9 of [22]). *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 $\rho(\Omega_1, \Omega_2; \mu) \leq 1 - \alpha^2/2$.*

In our correlated space, the bipartite graph on $\Omega_1 \cup \Omega_2$ is connected since every $x \in \Omega_1$ is connected to $*$ in Ω_2 and vice versa. Therefore, we can conclude that $\rho(\Omega_1, \Omega_2; \nu) \leq \rho := 1 - \epsilon^4/2$.

Apply Theorem 3.1 by setting $\rho \leftarrow \rho, \alpha \leftarrow \epsilon^2, \epsilon \leftarrow \Gamma_\rho(\epsilon/3, \epsilon/3)/2$ to get τ and d . We will later apply this theorem with the parameters obtained here. Fix an arbitrary subset $C \subseteq V$, and let $C_\alpha := C \cap v^\alpha$. For $\alpha \in I_D$, call v^α *blocked* if $\mu^{\otimes R}(C_\alpha) \geq 1 - \epsilon$. The number of blocked v^α 's is at most $c(C)/(1 - \epsilon)$.

By Property (iii) of Proposition 3.1, unless $c(C) \geq 2(1 - \epsilon)$ (i.e., unless two vertices are blocked), there exists a path $(s, v^{\alpha_1}, v^{\alpha_2}, v^{\alpha_3}, t)$ or $(t, v^{\alpha_1}, v^{\alpha_2}, v^{\alpha_3}, s)$ where each v^{α_i} is unblocked. Without loss of generality, suppose we have a path $(s, v^{\alpha_1}, v^{\alpha_2}, v^{\alpha_3}, t)$.

For $1 \leq j \leq 3$, let $S_j \subseteq v^{\alpha_j}$ be such that $x \in S_j$ if there exists a path $(v_x^{\alpha_j}, v_{x^{j+1}}^{\alpha_{j+1}}, \dots, v_{x^3}^{\alpha_3}, t)$ for some x^{j+1}, \dots, x^3 . For $1 \leq j \leq 3$, 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*, then $\mu^{\otimes R}(S_1) > 0$, which shows that there exists a $s \rightarrow t$ path even after removing vertices in C .

Lemma 3.8. *Suppose that for any $1 \leq j \leq 3$ and $1 \leq i \leq R$, $\text{Inf}_i^{\leq d}[f_j] \leq \tau$. Then $\mu^{\otimes R}(S_1) > 0$.*

Proof. We prove by induction that $\mu^{\otimes R}(S_j) \geq \epsilon/3$ for $j = 3, 2, 1$. It holds when $j = 3$ since v^{α_3} is unblocked. Assuming $\mu^{\otimes R}(S_j) \geq \epsilon/3$, since S_j does not reveal any influential coordinate, Theorem 3.1 shows that for any subset $T_{j-1} \subseteq v^{\alpha_{j-1}}$ with $\mu^{\otimes R}(T_{j-1}) \geq \epsilon/3$, there exists an arc from S_j and T_{j-1} . If $S'_{j-1} \subseteq v^{\alpha_{j-1}}$ is the set of in-neighbors of S_j , we have $\mu^{\otimes R}(S'_{j-1}) \geq 1 - \epsilon/3$. Since $v^{\alpha_{j-1}}$ is unblocked, $\mu^{\otimes R}(S'_{j-1} \setminus C) \geq 2\epsilon/3$, completing the induction. \square

In summary, in the completeness case, if we remove vertices of total weight $\leq 4(1 + 2\epsilon)/3$, no vertex can reach both s and t . In the soundness case, unless we reveal an influential coordinate or we remove vertices of total weight at least $2(1 - \epsilon)$, there is a $s \rightarrow t$ path or $t \rightarrow s$ path, which means that either s or t can reach every vertex. The gap between the two cases is at least

$$\frac{2(1 - \epsilon)}{4(1 + 2\epsilon)/3},$$

which approaches to $3/2$ as $\epsilon \rightarrow 0$.

3.4 $(2 - \epsilon)$ -Inapproximability for $\{s, t\}$ -NodeDoubleCut

Consider the digraph $D_{a,b}$ introduced in Section 2. Let $r = b - 2a + 1$. Based on $D_{a,b}$, we define the *dictatorship test* graph $\mathcal{D}_{a,b,R,\epsilon}^{st} = (V, A)$ as follows, for a positive integer R and $\epsilon > 0$. It will be used to show hardness results under the Unique Games Conjecture in Section 3.6. Consider the probability space (Ω, μ) where $\Omega := \{0, \dots, r - 1, *\}$, and $\mu : \Omega \mapsto [0, 1]$ with $\mu(*) = \epsilon$ and $\mu(x) = (1 - \epsilon)/r$ for $x \neq *$.

1. $V = \{s, t\} \cup \{v_x^\alpha\}_{\alpha \in I_D, x \in \Omega^R}$. Let v^α denote the set of vertices $\{v_x^\alpha\}_{x \in \Omega^R}$.
2. For $\alpha \in I_D$ and $x \in \Omega^R$, define the weight as $c(v_x^\alpha) = \mu^{\otimes R}(x)$. We note that the sum of weights is ab . The terminals s and t have infinite weight.
3. For each arc between s and $\alpha \in I_D$, for each $x \in \Omega^R$, add an arc with the same direction between s and v_x^α . Do the same for each arc between t and $\alpha \in I_D$.
4. For each arc $(\alpha, \beta) \in A_D$ with $\alpha = (\alpha_1, \alpha_2), \beta = (\beta_1, \beta_2) \in I_D$ and $x, y \in \Omega^R$, we have an arc from v_x^α to v_y^β according to the following rule (note that $\alpha_2 \neq \beta_2$).

- (a) $\alpha_2 < \beta_2$: add an arc if for any $1 \leq j \leq R$: $[y_j = (x_j + 1) \bmod r]$ or $[y_j = *]$ or $[x_j = *]$. Call them *forward* arcs.
- (b) $\alpha_2 > \beta_2$: add an arc if for any $1 \leq j \leq R$: $[y_j = (x_j - 1) \bmod r]$ or $[y_j = *]$ or $[x_j = *]$. Call them *backward* arcs.
- (c) If $(\alpha, \beta) \in A_D$ is a jumping arc, call (v_x^α, v_y^β) also a jumping arc.

Completeness. We first prove that removing a set of vertices that correspond to *dictators* behaves the same as the fractional solution that gives $1/r$ to every vertex. For any $q \in [R]$, let $V_q := \{v_x^\alpha : \alpha \in I_D, x_q = * \text{ or } 0\}$. We note that the total weight of V_q is

$$ab \left(\epsilon + \frac{1 - \epsilon}{r} \right) \leq ab\epsilon + \frac{ab}{b - 2a}.$$

Lemma 3.9. *After removing vertices in V_q , no vertex in V can reach both s and t .*

Proof. Suppose towards contradiction that there exists a vertex that can reach both s and t . First, assume that this vertex is $v_{x_0}^{\alpha_0}$ for some $\alpha_0 = ((\alpha_0)_1, (\alpha_0)_2) \in I_D$ and $x_0 \in \Omega^R$. Let $p_1 = (v_{x_0}^{\alpha_0}, v_{y_1}^{\beta_1}, \dots, v_{y_l}^{\beta_l}, s)$ be a $v_{x_0}^{\alpha_0} \rightarrow s$ path and $p_2 = (v_{x_0}^{\alpha_0}, v_{x_1}^{\alpha_1}, \dots, v_{x_k}^{\alpha_k}, t)$ be a $v_{x_0}^{\alpha_0} \rightarrow t$ path in $\mathcal{D}_{R,\epsilon}^{\text{st}} - V_q$ for some $k, l \in \mathbb{N}$, and $\alpha_1, \dots, \alpha_k, \beta_1, \dots, \beta_l \in I_D$, and $x_1, \dots, x_k, y_1, \dots, y_l \in \Omega^R$.

Proposition 3.2. $(x_k)_q \geq (x_0)_q + b - (\alpha_0)_2 - a + 1$.

Proof. Consider the two sequences $(\alpha_0)_q, \dots, (\alpha_k)_q$ and $(x_0)_q, \dots, (x_k)_q$. Since we removed V_q , $(\alpha_{i+1})_2 > (\alpha_i)_2$ if and only if $(x_{i+1})_q > (x_i)_q$. Let $n_{\text{jf}}, n_{\text{jb}}, n_{\text{rf}}, n_{\text{rb}}$ be the number forward jumping arcs, backward jumping arcs, forward non-jumping arcs, backward non-jumping arcs in p_2 respectively. Jumping forward arcs, jumping backward arcs, non-jumping forward arcs, and non-jumping backward arcs change $(\alpha_i)_2$ by $+2, -2, +1$, and -1 respectively. By considering $(\alpha_0)_q, \dots, (\alpha_k)_q$,

$$2n_{\text{jf}} + n_{\text{rf}} - 2n_{\text{jb}} - n_{\text{rb}} = (b + 1) - (\alpha_0)_2.$$

Since using a jumping arc increases $(\alpha_i)_1$ by 1,

$$n_{\text{jf}} + n_{\text{jb}} \leq a - 1.$$

Forward arcs (whether they are jumping or not) increase $(x_i)_q$ by 1 and backward arc decrease it by 1. Consider $(x_0)_q, \dots, (x_k)_q$,

$$\begin{aligned} (x_k)_q - (x_0)_q &\geq n_{\text{jf}} + n_{\text{rf}} - n_{\text{jb}} - n_{\text{rb}} - 1 \\ &\geq (2n_{\text{jf}} + n_{\text{rf}} - 2n_{\text{jb}} - 2n_{\text{rb}}) - (n_{\text{jf}} - n_{\text{jb}}) - 1 \\ &\geq b - (\alpha_0)_2 - a + 1, \end{aligned}$$

as claimed. □

The same proof for p_1 shows that $(x_0)_q \geq (y_l)_q + (\alpha_0)_2 - a$. Therefore, $(x_k)_q \geq (y_l)_q + b - 2a + 1$ and $(y_l)_q \geq 1$. This implies $(x_k)_q > b - 2a + 1 = r$, leading to contradiction. □

Soundness. Suppose that we removed some vertices C such that no vertex $w \in V \setminus C$ can reach both s and t . We show this happens only if C reveals an influential coordinate or $c(C) \geq 2a(1 - \epsilon)$.

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 $(x, y) \in \Omega^2$.

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

We note that the marginal distribution of both x and y is equal to μ . Assuming $\epsilon < 1/2r$, the minimum probability of any atom in $\Omega_1 \times \Omega_2$ is ϵ^2 . By the same arguments as in Section 3.3, $\rho(\Omega_1, \Omega_2; \nu) \leq \rho := 1 - \epsilon^4/2$.

Apply Theorem 3.1 $\rho \leftarrow \rho, \alpha \leftarrow \epsilon^2, \epsilon \leftarrow \Gamma_\rho(\epsilon/3, \epsilon/3)/2$ to get τ and d . We will later apply this theorem with the parameters obtained here. Fix an arbitrary subset $C \subseteq V$, and let $C_\alpha := C \cap v^\alpha$. For $\alpha \in I_D$, call v^α *blocked* if $\mu^{\otimes R}(C_\alpha) \geq 1 - \epsilon$. The number of blocked v^α 's is at most $c(C)/(1 - \epsilon)$.

By Property 2. of Lemma 2.2, unless $c(C) \geq (2a-1)(1-\epsilon)$ (i.e., unless $2a-1$ vertices are blocked), there exists $\alpha_0 \in I_D$ and a path $(v^{\alpha_0}, v^{\alpha_0-1}, \dots, v^{\alpha_0-k}, s)$ and $(v^{\alpha_0}, v^{\alpha_0+1}, \dots, v^{\alpha_0+l}, t)$ where each v^{α_i} is unblocked for $-k \leq i \leq l$.

For $-k \leq j \leq -1$, let $S_j \subseteq v^{\alpha_j}$ be such that $x \in S_j$ if there exists a path $(v_x^{\alpha_j}, v_{x^{j-1}}^{\alpha_{j-1}}, \dots, v_{x^{-k}}^{\alpha_{-k}}, s)$ for some x^{j-1}, \dots, x^{-k} . Similarly, For $1 \leq j \leq l$, let $S_j \subseteq v^{\alpha_j}$ be such that $x \in S_j$ if there exists a path $(v_x^{\alpha_j}, v_{x^{j+1}}^{\alpha_{j+1}}, \dots, v_{x^l}^{\alpha_l}, t)$ for some x^{j+1}, \dots, x^l . 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*, that there exists a $x^0 \in \Omega^R$ such that $v_{x^0}^{\alpha_0}$ can reach both s and t even after removing vertices in C .

Lemma 3.10. *Suppose that for any $j \in \{-k, \dots, -1\} \cup \{1, \dots, l\}$ and $1 \leq i \leq R$, $\text{Inf}_i^{\leq d}[f_j] \leq \tau$. Then there exists a $x^0 \in \Omega^R$ such that $v_{x^0}^{\alpha_0}$ can reach both s and t*

Proof. We prove that $\mu^{\otimes R}(S_1) \geq \epsilon/3$ by induction on $j = l, \dots, 1$. It holds when $j = l$ since v^{α_l} is unblocked. Assuming $\mu^{\otimes R}(S_j) \geq \epsilon/3$, since S_j does not reveal any influential coordinate, Theorem 3.1 shows that for any subset $T_{j-1} \subseteq v^{\alpha_{j-1}}$ with $\mu^{\otimes R}(T_{j-1}) \geq \epsilon/3$, there exists an arc from S_j and T_{j-1} . If $S'_{j-1} \subseteq v^{\alpha_{j-1}}$ is the set of in-neighbors of S_j , we have $\mu^{\otimes R}(S'_{j-1}) \geq 1 - \epsilon/3$. Since $v^{\alpha_{j-1}}$ is unblocked, $\mu^{\otimes R}(S'_{j-1} \setminus C) \geq 2\epsilon/3$, completing the induction.

The same argument also proves that $\mu^{\otimes R}(S_{-1}) \geq \epsilon/3$ by induction on $j = -k, \dots, -1$. The total weight of the in-neighbors of S_{-1} in v^{α_0} is at least $1 - \epsilon/3$, and the total weight of the in-neighbors of S_1 in v^{α_0} is at least $1 - \epsilon/3$. Therefore, the total weight of vertices in v^{α_0} that has outgoing arcs to both S_{-1} and S_1 is at least $1 - 2\epsilon/3$. Since α_0 is not blocked, there exists a vertex $v_{x^0}^{\alpha_0}$ that has outgoing arcs to both S_1 and S_{-1} , and is not contained C . This vertex can reach both s and t . \square

In summary, in the completeness case, if we remove vertices of total weight at most $abe + ab/(b-2a)$, no vertex can reach both s and t . In the soundness case, unless we reveal an influential coordinate or we remove vertices of total weight at least $(2a-1)(1-\epsilon)$, there exists a vertex that can reach both s and t . The gap between the two cases is at least

$$\frac{(2a-1)(1-\epsilon)}{abe + ab/(b-2a)},$$

which approaches to 2 as a increases, $b = a^2$ and $\epsilon = 1/a^4$.

3.5 Hardness of VertexCover on k -partite Graphs

Fix $k \geq 3$ and $\epsilon > 0$. Let $\Omega := \{*, 0, 1\}$. Let $R \in \mathbb{N}$ be another parameter. Our dictatorship test $\mathcal{D}_{k,R,\epsilon}^{\text{vc}}$ ($([k] \times \Omega^R, E)$) is defined as follows. Each vertex is represented by v_x^i where $i \in [k]$ and $x \in \Omega^R$ is a R -dimensional vector. Let $v^i := \{v_x^i\}_{x \in \Omega^R}$. There will be no edge within each v^i , so $\mathcal{D}_{k,R,\epsilon}^{\text{vc}}$ will be k -partite. Consider the probability space (Ω, μ) where $\Omega := \{0, 1, *\}$, and $\mu : \Omega \mapsto [0, 1]$ with $\mu(*) = \epsilon$ and $\mu(x) = (1-\epsilon)/2$ for $x \neq *$. We define the weight $c(v_x^i) := \mu^{\otimes R}(x) = \prod_{i=1}^R \mu(x_i)$. The sum of weights is k . The edges are constructed as follows.

1. There is an edge between v_x^i with $x = (x_1, \dots, x_R)$ and v_y^j with $y = (y_1, \dots, y_R)$ if and only if
 - (a) $i \neq j$.
 - (b) For any $1 \leq l \leq R$: $[x_l \neq y_l]$ or $[y_l = *]$ or $[x_l = *]$.

Completeness. Fix $j \in [R]$ and let $U_q := \{v_x^i : x_j = 0 \text{ or } *\}$. The weight of U_q is $c(U_q) = k(1 + \epsilon)/2$.

Lemma 3.11. U_q is a vertex cover.

Proof. Let $\{v_x^i, v_y^j\}$ be an edge of $\mathcal{D}_{k,R,\epsilon}^{\text{vc}}$. If both endpoints do not belong to U_q , it implies $x_q = y_q = 1$. It contradicts our construction. \square

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 Ω . It is defined by the following process to sample $(x, y) \in \Omega^2$.

1. Sample $x \in \{0, 1\}$ uniformly at random. Let $y = 1 - x$.
2. Change x to $*$ with probability ϵ . Do the same for y independently.

We note that the marginal distribution of both x and y is equal to μ . Assuming $\epsilon < 1/3$, the minimum probability of any atom in $\Omega_1 \times \Omega_2$ is ϵ^2 . By the same arguments as in Section 3.3, $\rho(\Omega_1, \Omega_2; \nu) \leq \rho := 1 - \epsilon^4/2$. Apply Theorem 3.1 ($\rho \leftarrow \rho, \alpha \leftarrow \epsilon^2, \epsilon \leftarrow \Gamma_\rho(\epsilon, \epsilon)/2$) to get τ and d . We will later apply this theorem with the parameters obtained here.

Fix an arbitrary vertex cover $U \subseteq V$, and let $U_i := U \cap v^i$ for $i \in [k]$. Let $f_i : \Omega^R \mapsto \{0, 1\}$ be the indicator function of U_i . Call v^i *blocked* if $\mathbb{E}[f_i] = \mu^{\otimes R}(U_i) \geq 1 - \epsilon$. The number of blocked v^i 's is at most $c(U)/(1 - \epsilon)$. We prove that if none of f_i reveals any *influential coordinate*, all but one v^i 's must be blocked.

Lemma 3.12. Suppose that for any $1 \leq i \leq k$ and $1 \leq j \leq R$, $\text{Inf}_j^{\leq d}[f_i] \leq \tau$. Then at least $k - 1$ v^i 's must be blocked.

Proof. Assume towards contradiction that there exist $i_1 \neq i_2 \in [k]$ such that v^{i_1} and v^{i_2} are unblocked. Since both f_{i_1} and f_{i_2} do not reveal influential coordinates and $\mathbb{E}[f_{i_1}], \mathbb{E}[f_{i_2}] \leq 1 - \epsilon$, Theorem 3.1 ($f \leftarrow 1 - f_1, g \leftarrow 1 - f_2$) shows that $\mathbb{E}_{(x,y) \sim \nu^{\otimes R}}[(1 - f_1)(x) \cdot (1 - f_2)(y)]$ is strictly greater than 0. This implies that there exists x, y such that there is an edge between $v_x^{i_1}$ and $v_y^{i_2}$ but neither $v_x^{i_1}$ nor $v_y^{i_2}$ is contained in U . This contradicts that U is a vertex cover. \square

Therefore, if U does not reveal any influential coordinate, then $c(U) \geq (k - 1)(1 - \epsilon)$. In summary, in the completeness case, there exists a vertex cover of weight $k(1 + \epsilon)/2$. In the soundness case, unless we reveal an influential coordinate, every vertex cover has weight at least $(k - 1)(1 - \epsilon)$. The gap between the two cases is at least

$$\frac{2(k - 1)(1 - \epsilon)}{k(1 + \epsilon)},$$

which approaches to $2(k - 1)/k$ as $\epsilon \rightarrow 0$.

3.6 Reduction from UniqueGames

UGC. We introduce the Unique Games Conjecture and its equivalent variant.

Definition An instance $\mathcal{L} = (B(U_B \cup W_B, E_B), [R], \{\pi(u, w)\}_{(u, w) \in E_B})$ of UNIQUEGAMES consists of a biregular bipartite graph $B(U_B \cup W_B, E_B)$ and a set $[R]$ of labels. For each edge $(u, w) \in E_B$ there is a constraint specified by a permutation $\pi(u, w) : [R] \rightarrow [R]$. The goal is to find a labeling $l : U_B \cup W_B \rightarrow [R]$ of the vertices such that as many edges as possible are satisfied, where an edge $e = (u, w)$ is said to be satisfied if $l(u) = \pi(u, w)(l(w))$.

Definition Given a UNIQUEGAMES instance $\mathcal{L} = (B(U_B \cup W_B, E_B), [R], \{\pi(u, w)\}_{(u, w) \in E_B})$, let $\text{OPT}(\mathcal{L})$ denote the maximum fraction of simultaneously-satisfied edges of \mathcal{L} by any labeling, i.e.

$$\text{OPT}(\mathcal{L}) := \frac{1}{|E_B|} \max_{l: U_B \cup W_B \rightarrow [R]} |\{e \in E : l \text{ satisfies } e\}|.$$

Conjecture 3.1 (The Unique Games Conjecture [17]). *For any constants $\eta > 0$, there is $R = R(\eta)$ such that, for a UNIQUEGAMES instance \mathcal{L} with label set $[R]$, it is NP-hard to distinguish between*

1. $\text{OPT}(\mathcal{L}) \geq 1 - \eta$.
2. $\text{OPT}(\mathcal{L}) \leq \eta$.

To show the optimal hardness result for VERTEXCOVER, Khot and Regev [18] introduced the following seemingly stronger conjecture, and proved that it is in fact equivalent to the original Unique Games Conjecture.

Conjecture 3.2 (Khot and Regev [18]). *For any constants $\eta > 0$, there is $R = R(\eta)$ such that, for a UNIQUEGAMES instance \mathcal{L} with label set $[R]$, it is NP-hard to distinguish between*

1. *There is a set $W' \subseteq W_B$ such that $|W'| \geq (1 - \eta)|W_B|$ and a labeling $l : U_B \cup W_B \rightarrow [R]$ that satisfies every edge (u, w) for $v \in U_B$ and $w \in W'$.*
2. $\text{OPT}(\mathcal{L}) \leq \eta$.

General Reduction. We now introduce our reduction from UNIQUEGAMES to our problems NODEDOUBLECUT, $\{s, t\}$ -NODEDOUBLECUT, and VERTEXCOVER ON k -PARTITE GRAPHS. We constructed three dictatorship tests for $\mathcal{D}_{R, \epsilon}^{\text{global}}$, $\mathcal{D}_{a, b, R, \epsilon}^{\text{st}}$, $\mathcal{D}_{k, R, \epsilon}^{\text{vc}}$. The first two are directed and $\mathcal{D}_{k, R, \epsilon}^{\text{vc}}$ is undirected, but they are all vertex-weighted. Fix a problem and the parameters, and let $\mathcal{D} = (V_{\mathcal{D}}, E_{\mathcal{D}})$ be the dictatorship test with the weight function $c : V_{\mathcal{D}} \rightarrow \mathbb{R}$.

Given an instance \mathcal{L} of UNIQUEGAMES, we describe how to reduce it to a graph $G = (V_G, E_G)$. G will be directed or undirected depending on the problem we reduce to. We assign to each vertex $w \in W_B$ a copy of $V_{\mathcal{D}}$ and for each *terminal* of $V_{\mathcal{D}}$, merge all $|W_B|$ copies into one. The merged terminals are $\{s, t\}$ for NODEDOUBLECUT and $\{s, t\}$ -NODEDOUBLECUT, and VERTEXCOVER has no terminal. For any $w \in W_B, v \in V_{\mathcal{D}}$, the vertex weight of (w, v) is $c(v)/|W_B|$, so that the sum of vertex weights (except terminals) is 4 for NODEDOUBLECUT, ab for $\{s, t\}$ -NODEDOUBLECUT, and k for VERTEXCOVER ON k -PARTITE GRAPHS. Let $I_{\mathcal{D}} := V_{\mathcal{D}} \setminus \{s, t\}$ for NODEDOUBLECUT and $\{s, t\}$ -NODEDOUBLECUT, and $I_{\mathcal{D}} := V_{\mathcal{D}}$ for VERTEXCOVER ON k -PARTITE GRAPHS.

For a permutation $\sigma : [R] \rightarrow [R]$, let $x \circ \sigma := (x_{\sigma(1)}, \dots, x_{\sigma(R)})$. To describe the set of edges, consider the random process where $u \in U_B$ is sampled uniformly at random, and its two neighbors w^1, w^2 are independently sampled. For each edge $(v_{x_1}^{i_1}, v_{x_2}^{i_2}) \in E_{\mathcal{D}}$, we create an edge $((w_1, v_{x_1}^{i_1 \circ \pi(u, w^1)}), (w_2, v_{x_2}^{i_2 \circ \pi(u, w^2)}))$. Call this edge is *created by* u . For each edge incident on a terminal (i.e., (X, v_x^i) or (v_x^i, X) where $X \in \{s, t\}$), we add the corresponding edge $(X, (w, v_x^i))$ or $((w, v_x^i), X)$ for each $w \in W_B$.

Completeness. Suppose there exists a labeling l and a subset $W' \subseteq W_B$ with $|W'| \geq (1 - \eta)|W_B|$ such that l satisfy every edge incident on W' .

NodeDoubleCut. Let $D = (V_D, A_D)$ be the graph constructed in Section 3.3 and I_D be $V_D \setminus \{s, t\}$. For every $w \in W'$, we remove the following vertices.

$$\{(w, v_x^\alpha) : \alpha \in I_D, x_{l(w)} = * \text{ or } 0\}.$$

For $w \notin W'$, we remove every vertex in $\{w\} \times I_D$. The total weight is at most $4(1 + 2\epsilon)/3 + 4\eta$. The completeness analysis for the dictatorship test ensures that no vertex in V_G can reach both s and t . The proof of Lemma 3.6 works verbatim — for each vertex $(w_j, v_{x_j}^{\alpha_j})$ with $x_j \in \Omega^R$, consider $(x_j)_{l(w_j)}$ in place of $(x_j)_q$.

$\{s, t\}$ -NodeDoubleCut. Let $D = (V_D, A_D)$ be the graph constructed in Section 3.4 and I_D be $V_D \setminus \{s, t\}$. For every $w \in W'$, we remove the following vertices.

$$\{(w, v_x^\alpha) : \alpha \in I_D, x_{l(w)} = * \text{ or } 0\}.$$

For $w \notin W'$, we remove every vertex in $\{w\} \times I_D$. The total weight is at most $ab/(b - 2a) + ab\epsilon + ab\eta$. The completeness analysis for the dictatorship test ensures that no vertex in V_G can reach both s and t . The proof of Lemma 3.9 works verbatim — for each vertex $(w_j, v_{x_j}^{\alpha_j})$ with $x_j \in \Omega^R$, consider $(x_j)_{l(w_j)}$ in place of $(x_j)_q$.

VertexCover on k -partite Graphs. For every $w \in W'$, we remove the following vertices.

$$\{(w, v_x^\alpha) : \alpha \in [k], x_{l(w)} = * \text{ or } 0\}.$$

For $w \notin W'$, we remove every vertex in $\{w\} \times V_D$. The total weight is at most $k(1 + \epsilon)/2 + k\eta$. The completeness analysis for the dictatorship test, Lemma 3.11, ensures that every edge of G is covered — for each edge $\{(w, v_x^i), (w', v_y^j)\}$, consider $x_{l(w)}$ and $y_{l(w')}$ in place of x_q and y_q .

Soundness. We present the soundness analysis. We first discuss how to extract an influential coordinate for each $u \in U_B$.

NodeDoubleCut and $\{s, t\}$ -NodeDoubleCut. Fix an arbitrary $C \subseteq V_G \setminus \{s, t\}$, and consider the graph after removing vertices in C . We will show that if $c(C)$ is small and no vertex can reach both s and t , we can decode influential coordinates for many vertices of the UNIQUE-GAMES instance. For NODEDOUBLECUT, since every vertex in V_G has an incoming arc from either s or t , it implies that any solution to NODEDOUBLECUT must reveal influential coordinates or $c(C)$ must be large. Recall the graph $D = (V_D, A_D)$ constructed in Section 3.3 (for NODEDOUBLECUT) or Section 3.4 (for $\{s, t\}$ -NODEDOUBLECUT), and $I_D = V_D \setminus \{s, t\}$.

For each $w \in W_B$, $r \in \{s, t\}$, $1 \leq j \leq |I_D|$, and a sequence $\bar{\alpha} = (\alpha_1, \dots, \alpha_j) \in (I_D)^j$, let $g_{w,r,j,\bar{\alpha}} : \Omega^R \mapsto \{0, 1\}$ such that $g_{w,i,j,\bar{\alpha}}(x) = 1$ if and only if there exists a path $p = ((w, v_x^{\alpha_1}), (w_2, v_{x_2}^{\alpha_2}), \dots, (w_j, v_{x_j}^{\alpha_j}), r)$ for some $w_2, \dots, w_j \in W_B$ and $x^2, \dots, x^j \in \Omega^R$.

For $u \in U_B$, $1 \leq j \leq |I_D|$, and $\bar{\alpha} \in (I_D)^j$, let $f_{u,r,j,\bar{\alpha}} : \Omega_R \mapsto [0, 1]$ be such that

$$f_{u,r,j,\bar{\alpha}}(x) = \mathbb{E}_{w \in N(u)} [g_{w,r,j,\bar{\alpha}}(x \circ \pi^{-1}(u, w))],$$

where $N(u)$ is the set of neighbors of u in the UNIQUEGAMES instance.

Let $S := 2(1 - \epsilon)$ (for NODEDOUBLECUT) or $S := (2a - 1)(1 - \epsilon)$ (for $\{s, t\}$ -NODEDOUBLECUT) be the lower bound on the weight in the soundness analysis of the respective dictatorship tests. Let $S' := (1 - \beta)S$ for some $\beta > 0$ that will be determined later. Let $\gamma(u)$ be the expected weight of $C \cap (\{w\} \times I_D)$, where w is a random neighbor of u . Since $\mathbb{E}_u[\gamma(u)] = S' = (1 - \beta)S$, at least β fraction of u 's have $\mathbb{E}_u[\gamma(u)] \leq S$. For such u , since no vertex can reach both s and t , the soundness analysis for the dictatorship test shows that there exists $q \in [R]$, $r \in \{s, t\}$, $1 \leq j \leq |I_D|$, $\bar{\alpha}$ such that $\text{Inf}_q^{\leq d}[f_{u,r,j,\bar{\alpha}}] \geq \tau$ (d and τ do not depend on u).

VertexCover on k -partite Graphs. Fix an arbitrary $C \subseteq V_G$, and consider the graph after removing vertices in C . We will show that if $c(C)$ is small and every edge is removed, we can decode influential coordinates for many vertices of the UNIQUEGAMES instance.

For each $w \in W_B$ and $j \in [k]$, let $g_{w,j} : \Omega^R \mapsto \{0, 1\}$ such that $g_{w,j}(x) = 1$ if and only if $(w, v_x^j) \notin C$. For $u \in U_B$ and $1 \leq j \leq [k]$, let $f_{u,j} : \Omega_R \mapsto [0, 1]$ be such that

$$f_{u,j}(x) = \mathbb{E}_{w \in N(u)} [g_{w,j}(x \circ \pi^{-1}(u, w))],$$

where $N(u)$ is the set of neighbors of u in the UNIQUEGAMES instance.

Let $S := (1 - \epsilon)(k - 1)$ be the lower bound on the weight in the soundness analysis of the dictatorship test. Let $S' := (1 - \beta)S$ for some $\beta > 0$ that will be determined later. Let $\gamma(u)$ be the expected weight of $C \cap (\{w\} \times V_D)$, where w is a random neighbor of u . Since $\mathbb{E}_u[\gamma(u)] = S' = (1 - \beta)S$, at least β fraction of u 's have $\mathbb{E}_u[\gamma(u)] \leq S$. For such u , since every edge is removed, the soundness analysis for the dictatorship test shows that there exists $q \in [R], 1 \leq j \leq [k]$ such that $\text{Inf}_q^{\leq d}[f_{u,j}] \geq \tau$ (d and τ do not depend on u).

Finishing Up. The above analyses NODEDOUBLECUT, $\{s, t\}$ -NODEDOUBLECUT, VERTEXCOVER ON k -PARTITE GRAPHS can be abstracted as follows. Each vertex $u \in U_B$ is associated with functions $\{f_{u,h} : \Omega^R \mapsto [0, 1]\}_{h \in I}$ for some index set I ($|I|$ is upper bounded by some constant for NODEDOUBLECUT, some function of a and b for $\{s, t\}$ -NODEDOUBLECUT, some function on k on VERTEXCOVER ON k -PARTITE GRAPHS). For at least β fraction of $u \in U_B$, there exist $i \in I$ and $q \in [R]$ such that $\text{Inf}_q^{\leq d}[f_{u,i}] \geq \tau$. Set $l(u) = q$ for those vertices. Since

$$\begin{aligned} \text{Inf}_q^{\leq d}(f_{u,i}) &= \sum_{\alpha_q \neq 0, |\alpha| \leq d} \widehat{f_{u,i}}(\alpha)^2 = \sum_{\alpha_q \neq 0, |\alpha| \leq d} (\mathbb{E}_w[\widehat{f_{u,i}}(\pi(u, w)^{-1}(\alpha))])^2 \\ &\leq \sum_{\alpha_q \neq 0, |\alpha| \leq d} \mathbb{E}_w[\widehat{f_{u,i}}(\pi(u, w)^{-1}(\alpha))^2] = \mathbb{E}_w[\text{Inf}_{\pi(u, w)^{-1}(q)}^{\leq d}(f_{u,i})], \end{aligned}$$

at least $\tau/2$ fraction of u 's neighbors satisfy $\text{Inf}_{\pi(u, w)^{-1}(q)}^{\leq d}(f_{u,i}) \geq \tau/2$. There are at most $2d/\tau$ coordinates with degree- d influence at least $\tau/2$ for a fixed h , so their union over $i \in I$ yields at most $2d \cdot |I|/\tau$ coordinates. Choose $l(w)$ uniformly at random among those coordinates (if there is none, set it arbitrarily). The above probabilistic strategy satisfies at least $\beta(\tau/2)(\tau/(2d \cdot |I|))$ fraction of all edges. Taking η smaller than this quantity proves the soundness of the reductions.

Final Results. Combining our completeness and soundness analyses and taking ϵ and η small enough, we prove our main results.

NodeDoubleCut. It is hard to distinguish the following cases.

1. Completeness: There is a $\{s, t\}$ -double cut of weight at most $4(1 + 2\epsilon)/3 + 4\eta$.
2. Soundness: There is no global double cut of weight less than $2(1 - \epsilon)(1 - \beta)$.

The gap is

$$\frac{2(1 - \epsilon)(1 - \beta)}{\frac{4(1 + 2\epsilon)}{3} + 4\eta},$$

which approaches to 1.5 by taking ϵ, η, β small. This proves Theorem 1.3.

$\{s, t\}$ -NodeDoubleCut. It is hard to distinguish the following cases.

1. Completeness: There is a $\{s, t\}$ -double cut of weight at most $ab/(b - 2a) + ab\epsilon + ab\eta$.
2. Soundness: There is no $\{s, t\}$ -double cut of weight less than $(2a - 1)(1 - \epsilon)(1 - \beta)$.

The gap is

$$\frac{(2a - 1)(1 - \epsilon)(1 - \beta)}{\frac{ab}{b - 2a} + ab\epsilon + ab\eta},$$

which approaches to 2 by taking a large, b larger, and ϵ, η, β small. This proves Theorem 1.1.

VertexCover on k -partite Graphs. It is hard to distinguish the following cases.

1. Completeness: There is a vertex cover of weight at most $k(1 + \epsilon)/2 + k\eta$.
2. Soundness: There is no vertex cover of weight less than $(k - 1)(1 - \epsilon)(1 - \beta)$.

The gap is

$$\frac{(k - 1)(1 - \epsilon)(1 - \beta)}{\frac{k(1 + \epsilon)}{2} + k\eta},$$

which approaches to $2(k - 1)/k$ by taking ϵ, η, β small. In particular, it approaches to $4/3$ for $k = 3$ and $3/2$ for $k = 4$. Take large r and small ϵ, β, η . With Lemma 3.2, this implies Theorem 1.4. With Lemma 3.4, this implies Theorem 1.9. With Lemma 3.3, this implies Theorem 1.10.

4 EdgeLin3Cut problems

Given a directed graph $D = (V, E)$, a feasible solution to (s, r, t) -EDGE_{LIN3}CUT in D is a subset F of arcs whose deletion from the graph eliminates all directed $s \rightarrow r$, $r \rightarrow t$ and $s \rightarrow t$ paths. One of our main tools used in the approximation algorithm for EDGE_{BIC}CUT is a $3/2$ -approximation algorithm for $(s, *, t)$ -EDGE_{LIN3}CUT. We present this algorithm now. For two sets $A, B \subseteq V$, let $\beta(A, B) := |\delta^{in}(A) \cup \delta^{in}(B)|$.

Proof of Theorem 1.6. We first rephrase the problem in a more convenient way.

Claim 4.1. $(s, *, t)$ -EDGE_{LIN3}CUT in a directed graph $D = (V, E)$ is equivalent to

$$\min \{ \beta(A, B) : t \in A \subset B \subseteq V - \{s\} \}.$$

Proof. Let $F \subseteq E$ be an optimal solution for $(s, *, t)$ -EDGE_{LIN3}CUT in D and let $(A, B) := \operatorname{argmin} \{ \beta(A, B) : t \in A \subset B \subseteq V - s \}$. Fix an arbitrary node $r \in B - A$. Since the deletion of $\delta^{in}(A) \cup \delta^{in}(B)$ results in a graph with no directed path from s to r , from r to t and from s to t , the edge set $\delta^{in}(A) \cup \delta^{in}(B)$ is a feasible solution to (s, r, t) -EDGE_{LIN3}CUT in D , thus implying that $|F| \leq \beta(A, B)$.

On the other hand, F is a feasible solution for (s, r, t) -EDGE_{LIN3}CUT in D for some $r \in V - \{s, t\}$. Let A be the set of nodes that can reach t in $D - F$, and R be the set of nodes that can reach r in $D - F$. Then, $F \supseteq \delta^{in}(A)$. Moreover, $F \supseteq \delta^{in}(R \cup A)$ since $R \cup A$ has in-degree 0 in $D - F$, and s is not in $R \cup A$ because it cannot reach r and t in $D - F$. Therefore, taking $B = R \cup A$ we get $F \supseteq \delta^{in}(A) \cup \delta^{in}(B)$. \square

Our algorithm for determining an optimal pair $(A, B) := \operatorname{argmin} \{ \beta(A, B) : t \in A \subset B \subseteq V - s \}$ proceeds as follows: We build a chain \mathcal{C} of $\bar{s}t$ -sets with the property that, for some value $k \in \mathbb{Z}_+$,

- (i) \mathcal{C} contains only cuts of value at most k , and
- (ii) every $\bar{s}t$ -set of cut value strictly less than k is in \mathcal{C} .

We start with k being the minimum $\bar{s}t$ -cut value and \mathcal{C} consisting of a single minimum $\bar{s}t$ -cut. In a general step, we find two $\bar{s}t$ -sets: a minimum $\bar{s}t$ -cut Y compatible with the current chain \mathcal{C} , i.e. $\mathcal{C} \cup \{Y\}$ forming a chain, and a minimum $\bar{s}t$ -cut Z not compatible with the current chain \mathcal{C} , i.e. crossing at least one member of \mathcal{C} . These two sets can be found in polynomial time. Indeed, let $t \in C_1 \subset \dots \subset C_q \subseteq V - s$ denote the members of \mathcal{C} . Find a minimum cut Y_i with $C_i \subseteq Y_i \subseteq V \setminus C_{i+1}$ for $i = 1, \dots, q$, and choose Y to be a minimum one among these cuts. Concerning Z , for each pair x, y of nodes with $y \in C_i \subseteq V - x$ for some $i \in \{1, \dots, q\}$, find a minimum cut Z_{xy} with $\{t, x\} \subseteq Z_{xy} \subseteq V - \{s, y\}$, and choose Z to be a minimum one among these cuts. If $d^{in}(Y) \leq d^{in}(Z)$, then we add Y to \mathcal{C} , and set k to $d^{in}(Y)$; otherwise we set k to $d^{in}(Z)$, and stop.

Let \mathcal{C} denote the chain constructed by the algorithm, and let Y be an arbitrary set crossing some of its members.

Claim 4.2. $d^{in}(Y) \geq d^{in}(C)$ for all $C \in \mathcal{C}$.

Proof. Suppose indirectly that $d^{in}(Y) < d^{in}(C)$ for some $C \in \mathcal{C}$. Let \mathcal{C}' denote the chain consisting of those members of \mathcal{C} that were added before C . As C is a set of minimum cut value compatible with \mathcal{C}' , Y crosses at least one member of \mathcal{C}' . Hence, by $d^{in}(Y) < d^{in}(C)$, the algorithm stops before adding C , a contradiction. \square

The claim implies that \mathcal{C} satisfies (i) and (ii) with the k obtained at the end of the algorithm. Indeed, (i) is obvious from the construction, while (ii) follows from the claim and the fact that \mathcal{C} contains all cuts of value strictly less than k that are compatible with \mathcal{C} .

By the above, the procedure stops with a chain \mathcal{C} containing all $\bar{s}t$ -sets of cut value less than k , and an $\bar{s}t$ -set Z of cut value exactly k which crosses some member X of \mathcal{C} . If the optimum value of our problem is less than k , then both members of the optimal pair (A, B) belong to the chain \mathcal{C} , and we can find them by taking the minimum of $\beta(A', B')$ where $A' \subset B'$ with $A', B' \in \mathcal{C}$.

We can thus assume that the optimum is at least k . As $d^{in}(Z) = k$ and $d^{in}(X) \leq k$, the submodularity of the in-degree function implies $d^{in}(X \cap Z) + d^{in}(X \cup Z) \leq d^{in}(Z) + d^{in}(X) \leq 2k$. Hence at least one of $d^{in}(X \cap Z) \leq k$ and $d^{in}(X \cup Z) \leq k$ holds. As $d(X \setminus Z, X \cap Z) + d(Z \setminus X, X \cap Z) \leq d^{in}(X \cap Z)$ and $d(V \setminus (X \cup Z), X \setminus Z) + d(V \setminus (X \cup Z), Z \setminus X) \leq d^{in}(X \cup Z)$, at least one of the following four possibilities is true:

1. $d^{in}(X \cap Z) \leq k$ and $d(X \setminus Z, X \cap Z) \leq \frac{1}{2}k$. Choose $A = X \cap Z$, $B = X$. Then $\beta(A, B) = d(X \setminus Z, X \cap Z) + d^{in}(X) \leq \frac{1}{2}k + k = \frac{3}{2}k$.
2. $d^{in}(X \cap Z) \leq k$ and $d(Z \setminus X, X \cap Z) \leq \frac{1}{2}k$. Choose $A = X \cap Z$, $B = Z$. Then $\beta(A, B) = d(Z \setminus X, X \cap Z) + d^{in}(Z) \leq \frac{1}{2}k + k = \frac{3}{2}k$.
3. $d^{in}(X \cup Z) \leq k$ and $d(V \setminus (X \cup Z), X \setminus Z) \leq \frac{1}{2}k$. Choose $A = Z$, $B = X \cup Z$. Then $\beta(A, B) = d^{in}(Z) + d(V \setminus (X \cup Z), X \setminus Z) \leq k + \frac{1}{2}k = \frac{3}{2}k$.
4. $d^{in}(X \cup Z) \leq k$ and $d(V \setminus (X \cup Z), Z \setminus X) \leq \frac{1}{2}k$. Choose $A = X$, $B = X \cup Z$. Then $\beta(A, B) = d^{in}(X) + d(V \setminus (X \cup Z), Z \setminus X) \leq k + \frac{1}{2}k = \frac{3}{2}k$.

Thus a pair (A, B) can be obtained by taking the minimum among the four possibilities above and $\beta(A', B')$ where $A' \subset B'$ with $A', B' \in \mathcal{C}$, concluding the proof of the theorem. \square

Before proving Theorem 1.7, we introduce the STRONGMINCUT problem, in which the input is a directed graph $D = (V, E)$, and the goal is to find a minimum number of edges whose deletion results in at least two disjoint weakly connected components. We observe that STRONGMINCUT is equivalent to min cut in the underlying undirected graph, hence it can be solved in polynomial-time.

Proof of Theorem 1.7. We first show that r -EDGEINOUTBLOCKER is a combination of (s, r, t) -EDGELIN3CUT and STRONGMINCUT.

Claim 4.3. *For every directed graph $D = (V, E)$ with $r \in V$, the optimal solution to r -EDGEINOUTBLOCKER has value equal to*

$$\min \left\{ \min_{s, t \in V - r} \{ (s, r, t)\text{-EDGELIN3CUT} \}, \text{STRONGMINCUT} \right\}$$

Proof. Let F be an optimal solution of r -EDGEINOUTBLOCKER in $D = (V, E)$ with $r \in V$. Both (s, r, t) -EDGELIN3CUT and STRONGMINCUT are upper bounds for the size of F . If $|F|$ is strictly smaller than STRONGMINCUT, then $D - F$ is weakly connected. By the definition of F , we have that $D - F$ does not contain an in- r -arborescence and hence it has a strongly connected component C_1 not containing r with $\delta_{D-F}^{\text{out}}(C_1) = \emptyset$. Similarly, since $D - F$ does not contain an out- r -arborescence, it has a strongly connected component C_2 not containing r with $\delta_{D-F}^{\text{in}}(C_2) = \emptyset$. Since $D - F$ is weakly connected, $C_1 \neq C_2$ and since C_1 and C_2 are strongly connected components, they should be disjoint. For arbitrary nodes $s \in C_1$ and $t \in C_2$, there are no directed $s \rightarrow r$, $r \rightarrow t$ and $s \rightarrow t$ paths in $D - F$, thus F contains an (s, r, t) -EDGELIN3CUT. \square

Now we turn to the proof of the theorem. The ‘if’ part follows from the claim. To see the other direction, consider an instance $D = (V, E)$ of (s, r, t) -EDGELIN3CUT. For each node $v \in V$, add $|E|$ parallel arcs from t to v and $|E|$ parallel arcs from v to s . This step does not affect the optimum value of (s, r, t) -EDGELIN3CUT. Let F be an α -approximate solution to r -EDGEINOUTBLOCKER in the resulting digraph. Clearly, F does not contain any of the (v, s) or (t, v) arcs. If $D - F$ contains a directed path from s to r or from r to t or from s to t , then the additional arcs can be used to obtain either an in- r -arborescence or an out- r -arborescence. Thus, F gives an α -approximate solution to (s, r, t) -EDGELIN3CUT in D . \square

In [4], Chekuri and Madan showed that under the UGC, the hardness of approximation for (s, r, t) -EDGELIN3CUT in directed graphs coincides with the so-called flow-cut gap of the associated path-blocking linear program. Theorem 1.7 implies the same hardness result for r -EDGEINOUTBLOCKER.

In the rest of the section, we show that $\{s, t\}$ -SEPEDEGE k CUT is solvable in polynomial time if k is a fixed constant.

Let $G = (V, E)$ be an undirected graph. Let the minimum size of an $\{s, t\}$ -cut in G be denoted by $\lambda_G(s, t)$. For two subsets of nodes X, Y , let $d(X, Y)$ denote the number of edges between X and Y and let $d(X) := d(X, V \setminus X)$. The cut value of a partition $\{V_1, \dots, V_q\}$ of V is defined to be the total number of crossing edges, that is, $(1/2) \sum_{i=1}^q d(V_i)$, and is denoted by $\gamma(V_1, \dots, V_q)$. Let $\gamma^q(G)$ denote the value of an optimum EDGE- q -CUT in G , i.e.,

$$\min \{ \gamma(V_1, \dots, V_q) : V_i \neq \emptyset \forall i \in [q], V_i \cap V_j = \emptyset \forall i, j \in [q], \cup_{i=1}^q V_i = V \}.$$

Proof of Theorem 1.8. Let γ^* denote the optimum value of $\{s, t\}$ -SEPEDEGE k CUT in $G = (V, E)$ and let H denote the graph obtained from G by adding an edge of infinite capacity between s and t . The algorithm is based on the following observation (we recommend the reader to consider $k = 3$ for ease of understanding):

Proposition 4.1. *Let $\{V_1, \dots, V_k\}$ be a partition of V corresponding to an optimal solution of $\{s, t\}$ -SEPEDGE k CUT, where s is in V_{k-1} and t is in V_k . Then $\gamma(V_1, \dots, V_{k-2}, V_{k-1} \cup V_k) \leq 2\gamma^{k-1}(H)$.*

Proof. Let W_1, \dots, W_{k-1} be a minimum $(k-1)$ -cut in H . Clearly, s and t are in the same part, so we may assume that they are in W_{k-1} . Let U_1, U_2 be a minimum $\{s, t\}$ -cut in $G[W_{k-1}]$. Then $\{W_1, \dots, W_{k-2}, U_1, U_2\}$ gives an $\{s, t\}$ -separating k -cut, showing that

$$\gamma^* \leq \gamma(W_1, \dots, W_{k-2}, U_1, U_2) = \gamma^{k-1}(H) + \lambda_{G[W_{k-1}]}(s, t). \quad (1)$$

By Menger's theorem, we have $\lambda_G(s, t)$ pairwise edge-disjoint paths $P_1, \dots, P_{\lambda_G(s, t)}$ between s and t in G . Consider one of these paths, say P_i . If all nodes of P_i are from $V_{k-1} \cup V_k$, then P_i has to use at least one edge from $\delta(V_{k-1}, V_k)$. Otherwise, P_i uses at least two edges from $\delta(V_1 \cup \dots \cup V_{k-2}) \cup \bigcup_{\substack{i, j \leq k-2 \\ i \neq j}} \delta(V_i, V_j)$. Hence the maximum number of pairwise edge-disjoint paths between s and t is

$$\lambda_G(s, t) \leq d(V_{k-1}, V_k) + \frac{1}{2} \left(d(V_1 \cup \dots \cup V_{k-2}) + \sum_{\substack{i, j \leq k-2 \\ i \neq j}} d(V_i, V_j) \right).$$

Thus, we have

$$\begin{aligned} \gamma^* &= d(V_{k-1}, V_k) + d(V_1 \cup \dots \cup V_{k-2}) + \sum_{\substack{i, j \leq k-2 \\ i \neq j}} d(V_i, V_j) \\ &\geq \lambda_G(s, t) + \frac{1}{2} \left(d(V_1 \cup \dots \cup V_{k-2}) + \sum_{\substack{i, j \leq k-2 \\ i \neq j}} d(V_i, V_j) \right) \\ &= \lambda_G(s, t) + \frac{1}{2} \gamma(V_1, \dots, V_{k-2}, V_{k-1} \cup V_k) \\ &\geq \lambda_{G[W_{k-1}]}(s, t) + \frac{1}{2} \gamma(V_1, \dots, V_{k-2}, V_{k-1} \cup V_k) \end{aligned}$$

that is,

$$\gamma^* \geq \lambda_{G[W_{k-1}]}(s, t) + \frac{1}{2} \gamma(V_1, \dots, V_{k-2}, V_{k-1} \cup V_k). \quad (2)$$

By combining (1) and (2), we get $\gamma(V_1, \dots, V_{k-2}, V_{k-1} \cup V_k) \leq 2\gamma^{k-1}(H)$, proving the proposition. \square

Karger and Stein [16] showed that the number of feasible solutions to EDGE- k -CUT in G with value at most $2\gamma^k(G)$ is $O(n^{4k})$, and all these solutions can be enumerated in (randomized) polynomial-time for fixed k . This observation together with Proposition 4.1 gives the following algorithm for finding an optimal solution to $\{s, t\}$ -SEPEDGE k CUT:

Step 1. Let H be the graph obtained from G by adding an edge of infinite capacity between s and t . In H , enumerate all feasible solutions to EDGE- $(k-1)$ -CUT—namely the vertex partitions $\{W_1, \dots, W_{k-1}\}$ —whose cut value $\gamma_H(W_1, \dots, W_{k-1})$ is at most $2\gamma^{k-1}(H)$. Without loss of generality, assume $s, t \in W_{k-1}$.

Step 2. For each feasible solution to EDGE- $(k-1)$ -CUT in H listed in Step 1, find a minimum $\{s, t\}$ -cut in $G[W_{k-1}]$, say U_1, U_2 .

Step 3. Among all feasible solutions $\{W_1, \dots, W_{k-1}\}$ to EDGE- $(k-1)$ -CUT listed in Step 1 and the corresponding U_1, U_2 found in Step 2, return the k -cut $\{W_1, \dots, W_{k-2}, U_1, U_2\}$ with minimum $\gamma(W_1, \dots, W_{k-2}, U_1, U_2)$.

□

5 Approximation for EdgeBiCut

In this section we describe an efficient $(2 - 1/448)$ -approximation algorithm for EDGEBiCUT (Theorem 1.5). We recall that in EDGEBiCUT, the goal is to find the smallest number of edges in a directed graph whose deletion ensures that there exist two distinct nodes s and t such that s cannot reach t and t cannot reach s . An equivalent formulation of the problem that is convenient for our purposes is as follows: Two sets A and B are called *uncomparable* if $A \setminus B \neq \emptyset$ and $B \setminus A \neq \emptyset$. Given a directed graph $D = (V, E)$, EDGEBiCUT is equivalent to finding an uncomparable pair $A, B \subseteq V$ with minimum $|\delta^{in}(A) \cup \delta^{in}(B)|$. Indeed, if A and B are uncomparable and we remove $\delta^{in}(A) \cup \delta^{in}(B)$ from the directed graph, then nodes in $A \setminus B$ cannot reach nodes in $B \setminus A$ and vice versa. On the other hand, if s cannot reach t and t cannot reach s , then the set of nodes that can reach s and the set of nodes that can reach t are uncomparable, and have in-degree 0.

We introduce some definitions and notation in order to describe the $(2 - \epsilon)$ -approximation algorithm, where the value of ϵ is computed at the end of the proof.

Definition For $A, B \subseteq V$, let $\beta(A, B) := |\delta^{in}(A) \cup \delta^{in}(B)|$ and let $\sigma(A, B) := d^{in}(A) + d^{in}(B)$. Furthermore, let

$$\begin{aligned}\beta &:= \min\{\beta(A, B) \mid A \text{ and } B \text{ are uncomparable}\}, \\ \sigma &:= \min\{\sigma(A, B) \mid A \text{ and } B \text{ are uncomparable}\}.\end{aligned}$$

The pair where the latter value is attained is called the *minimum uncomparable cut-pair*.

Definition If c is a capacity function on a directed graph D , then $d_c^{in}(U) = \sum_{e \in \delta^{in}(U)} c(e)$ is the sum of the capacities of incoming edges of U . Similarly, $d_c^{out}(U) = \sum_{e \in \delta^{out}(U)} c(e)$. For two disjoint set of vertices A and B , the number of edges from A to B is defined as $d(A, B) = |\delta^{out}(A) \cap \delta^{in}(B)|$.

Clearly, $\sigma(A, B) \geq \beta(A, B)$ for any A, B . The following lemma shows that σ can be computed efficiently. This means that we immediately have a $(2 - \epsilon)$ -approximation if $\sigma \leq (2 - \epsilon)\beta$.

Lemma 5.1. *For a directed graph $D = (V, E)$, there exists a polynomial time algorithm to find a minimum uncomparable cut-pair.*

Proof. For fixed vertices a and b , there is an efficient algorithm to find A and B such that $a \in A \setminus B$ and $b \in B \setminus A$ and $\sigma(A, B)$ is minimized. Indeed, this is precisely finding the sink side of an $a - b$ min-cut and that of a $b - a$ min-cut. Trying all possible a and b and taking the minimum gives the desired result. □

We also need the following lemma showing that we can minimize $\beta(A, B)$ among pairs whose intersection is fixed.

Lemma 5.2. *Given a directed graph $D = (V, E)$ and $Z \subseteq V$, there exists a polynomial time algorithm to find an uncomparable pair A, B satisfying $A \cap B = Z$ that minimizes $\beta(A, B)$ among pairs with this property.*

Proof. Let $D' = D[V \setminus Z]$ be the directed graph induced on $V \setminus Z$. The EDGEDOUBLECUT problem can be solved in polynomial time in D' [2]; let X' and Y' be the disjoint sets whose incoming edges give the optimal double cut. We claim that the pair $X' \cup Z, Y' \cup Z$ forms a minimum bicut among all bicuts with intersection Z . Indeed, assume the optimal solution is $\beta(A, B)$. Let $X = A \setminus B, Y = B \setminus A$ and $W = V - (A \cup B)$. Then

$$\begin{aligned} \beta(X' \cup Z, Y' \cup Z) &= d_{D'}^{in}(X') + d_{D'}^{in}(Y') + d^{in}(Z) \\ &\leq d_{D'}^{in}(X) + d_{D'}^{in}(Y) + d^{in}(Z) \\ &= d^{in}(Z) + d(W, X) + d(W, Y) + d(X, Y) + d(Y, X) \\ &= \beta(A, B). \end{aligned}$$

□

A similar argument proves the following.

Lemma 5.3. *Given a directed graph $D = (V, E)$ and $W \subseteq V$, there exists a polynomial time algorithm to find an uncomparable pair A, B satisfying $V \setminus (A \cup B) = W$ that minimizes $\beta(A, B)$ among pairs with this property.*

Proof of Theorem 1.5. If the minimum bicut (A, B) has the property that either $|A \cap B| \leq 2$ or $|V \setminus (A \cup B)| \leq 2$, then the optimal bicut can be found by applying the above algorithms by setting Z or W to every possible choice of subsets of nodes of size at most 2 and considering the minimum. The algorithm that we present below makes use of this observation.

APPROXIMATEGLOBALBICUT for directed graph $D = (V, E)$

Find minimum bicut if $|Z| \leq 2$ or $|W| \leq 2$ using Lemmas 5.2 and 5.3

Compute the minimum uncomparable cut-pair

For each tuple of nodes $(x, y, w_1, w_2, z_1, z_2)$

$X' \leftarrow$ sink-side of the minimum $\{w_1, w_2, y\} \rightarrow \{x, z_1, z_2\}$ -cut

$Y' \leftarrow$ sink-side of the minimum $\{w_1, w_2, x\} \rightarrow \{y, z_1, z_2\}$ -cut

$E_1 \leftarrow E[X'] \cup E[Y']$

$E_2 \leftarrow E[V \setminus X'] \cup E[V \setminus Y']$

$D_1 \leftarrow D$ with the arcs in E_1 duplicated

$D_2 \leftarrow D$ with the arcs in E_2 duplicated

$Z' \leftarrow$ sink-side of minimum $\{w_1, w_2, x, y\} \rightarrow \{z_1, z_2\}$ -cut in D_1

$W' \leftarrow$ source-side of minimum $\{w_1, w_2\} \rightarrow \{x, y, z_1, z_2\}$ -cut in D_2

$D' \leftarrow$ contract $X' \cap Y'$ to z' , contract $V \setminus X'$ to w' , remove all $w'z'$ arcs

In D' , find $\overline{w'z'}$ -sets $A' \subsetneq B'$ such that $\beta(A', B')$ is

at most $3/2$ times minimum, using Theorem 1.6

Find all bicuts that can be generated using set operations on

X', Y', Z', W', A', B' .

Output the minimum bicut over all the bicuts found.

To show that the algorithm is correct, first we fix a global min-bicut (A, B) , i.e. the optimum is $\beta = \beta(A, B)$. Let $X = A \setminus B, Y = B \setminus A, Z = A \cap B$, and $W = V \setminus (A \cup B)$. We assume that both Z and W are of size at least 3, otherwise it is clear that the algorithm finds the optimum. Let $\epsilon > 0$ be a constant whose value will be determined later.

Lemma 5.4. *If one of the following is true, then the minimum uncomparable cut-pair is a $(2 - \epsilon)$ -approximation:*

(i) $d(Z, W) \leq (1 - \epsilon)\beta$,

- (ii) For every $z_1, z_2 \in Z$, there exists a subset U of nodes containing z_1, z_2 but not Z with $d^{in}(U) < (1 - \epsilon)\beta$.
- (iii) For every $w_1, w_2 \in W$, there exists a subset U of nodes not containing w_1, w_2 but intersecting W with $d^{in}(U) < (1 - \epsilon)\beta$.

Proof.

- (i) The pair (A, B) is uncomparable, and $\sigma(A, B) \leq (2 - \epsilon)\beta$ if $d(Z, W) \leq (1 - \epsilon)\beta$. Therefore the minimum uncomparable cut-pair is a $(2 - \epsilon)$ -approximation if (i) holds.
- (ii) Among the sets with in-degree less than $(1 - \epsilon)\beta$ which do not contain every node of Z , let T be the one with inclusionwise maximal intersection with Z . Let $z_1 \in Z \setminus T$ and $z_2 \in Z \cap T$. There exists a set U such that $d^{in}(U) < (1 - \epsilon)\beta$ and $z_1, z_2 \in U$ that contains z_1 and z_2 but not the whole Z . Because of the maximal intersection of T with Z , we have that $T \not\subseteq U$. Hence T and U are uncomparable and $\sigma(T, U) \leq (2 - 2\epsilon)\beta$.
- (iii) Argument similar to the proof of (ii) shows that the minimum uncomparable cut-pair is a $(2 - 2\epsilon)$ -approximation if (iii) holds.

□

By Lemma 5.4, we only have to consider the case where $d(W, Z) \geq (1 - \epsilon)\beta$. We will also assume that z_1, z_2, w_1, w_2 are chosen such that $d^{in}(U) \geq (1 - \epsilon)\beta$ for all subsets U of nodes containing z_1, z_2 but not Z , and $d^{in}(U) \geq (1 - \epsilon)\beta$ for all subsets U of nodes not containing w_1, w_2 but intersecting W .

We may assume that $\beta(X', Y') \geq (2 - \epsilon)\beta$, otherwise X' and Y' forms a $(2 - \epsilon)$ -approximation. Also, $d^{in}(X') \leq d^{in}(X \cup Z) \leq \beta$ because X' is the sink-side of a $\min \{w_1, w_2, y\} \rightarrow \{x, z_1, z_2\}$ cut. Similarly, $d^{in}(Y') \leq d^{in}(Y \cup Z) \leq \beta$. Let c be the capacity function obtained by increasing the capacity of each edge in E_1 to 2, and let \bar{c} be the capacity function obtained by increasing the capacity of each edge in E_2 to 2. We consider four cases depending on the relations between W and $X' \cup Y'$, and between Z and $X' \cap Y'$.

Case 0. Suppose $W \cap (X' \cup Y') = \emptyset$, $Z \subseteq X' \cap Y'$. In this case $\delta^{in}(X')$ and $\delta^{in}(Y')$ both contain all edges counted in $d(W, Z)$. Hence $\beta(X', Y') \leq \sigma(X', Y') - d(W, Z) \leq (1 + \epsilon)\beta$. This shows that (X', Y') is a $(1 + \epsilon)$ -approximation.

For the remaining three cases, we will use the following proposition.

Proposition 5.1. *If $d^{in}(X' \cap Z') \geq (1 - \epsilon)\beta$ and $d^{in}(Y' \cap Z') \geq (1 - \epsilon)\beta$, then $\beta(X' \cup Z', Y' \cup Z') \leq 2\epsilon\beta + d_c^{in}(Z)$.*

Proof. If $d^{in}(X' \cap Z') \geq (1 - \epsilon)\beta$, then $d^{in}(X') - d^{in}(X' \cap Z') \leq \epsilon\beta$, so

$$\begin{aligned} d^{in}(X' \cup Z') &\leq d^{in}(Z') + \epsilon\beta - d(X' \setminus Z', Z' \setminus X') - d(Z' \setminus X', X' \setminus Z') \\ &\leq d^{in}(Z') + \epsilon\beta - d(X' \setminus Z', Z' \setminus X'). \end{aligned} \quad (3)$$

Similarly,

$$d^{in}(Y' \cup Z') \leq d^{in}(Z') + \epsilon\beta - d(Y' \setminus Z', Z' \setminus Y'). \quad (4)$$

We need the following claim.

Claim 5.1.

$$\begin{aligned} \beta(X' \cup Z', Y' \cup Z') &\leq \sigma(X' \cup Z', Y' \cup Z') + d_c^{in}(Z') - 2d^{in}(Z') \\ &\quad + d(X' \setminus Z', Z' \setminus X') + d(Y' \setminus Z', Z' \setminus Y'). \end{aligned} \quad (5)$$

Proof. By counting the edges entering Z' , we have

1. $d_c^{in}(Z') = d^{in}(Z') + |\delta^{in}(Z') \cap E_1|$.
2. $d^{in}(Z') = d(V \setminus (X' \cup Y' \cup Z'), Z') + |\delta^{in}(Z') \cap E_1| + d(X' \setminus Z', Z' \setminus X') + d(Y' \setminus Z', Z' \setminus Y') - d((X' \cap Y') \setminus Z', Z' \setminus (X' \cup Y'))$.

The first equation can be rewritten as $d_c^{in}(Z') - 2d^{in}(Z') = -d^{in}(Z') + |\delta^{in}(Z') \cap E_1|$. Using this and the second equation, we get $d_c^{in}(Z') - 2d^{in}(Z') + d(X' \setminus Z', Z' \setminus X') + d(Y' \setminus Z', Z' \setminus Y') = -d(V \setminus (X' \cup Y' \cup Z'), Z') + d((X' \cap Y') \setminus Z', Z' \setminus (X' \cup Y'))$. Thus the desired inequality (5) simplifies to

$$\begin{aligned} \beta(X' \cup Z', Y' \cup Z') &\leq \sigma(X' \cup Z', Y' \cup Z') \\ &\quad - d(V \setminus (X' \cup Y' \cup Z'), Z') + d((X' \cap Y') \setminus Z', Z' \setminus (X' \cup Y')). \end{aligned}$$

The elements counted by $d(V \setminus (X' \cup Y' \cup Z'), Z')$ are counted twice in $\sigma(X' \cup Z', Y' \cup Z')$. Hence we have the desired relation (5). \square

Using (3),(4) and (5) we get

$$\begin{aligned} &\beta(X' \cup Z', Y' \cup Z') \\ &\leq d^{in}(X' \cup Z') + d^{in}(Y' \cup Z') \\ &\quad + d_c^{in}(Z') - 2d^{in}(Z') + d(X' \setminus Z', Z' \setminus X') + d(Y' \setminus Z', Z' \setminus Y') \\ &\leq d^{in}(Z') + \epsilon\beta + d^{in}(Z') + \epsilon\beta + d_c^{in}(Z') - 2d^{in}(Z') \\ &\leq 2\epsilon\beta + d_c^{in}(Z') \\ &\leq 2\epsilon\beta + d_c^{in}(Z). \end{aligned}$$

\square

Case 1. Suppose $W \cap (X' \cup Y') = \emptyset$ and $Z \not\subseteq X' \cap Y'$. Without loss of generality, let $Z \not\subseteq X'$. The set $X' \cap Z'$ contains z_1, z_2 but not the whole Z , hence $d^{in}(X' \cap Z') \geq (1 - \epsilon)\beta$.

We first consider the subcase when $d^{in}(Y' \cap Z') \geq (1 - \epsilon)\beta$ also holds. Then, by Proposition 5.1, we get

$$\beta(X' \cup Z', Y' \cup Z') \leq 2\epsilon\beta + d_c^{in}(Z).$$

We are in the case where $(X' \cup Y') \cap W = \emptyset$, so $d_c^{in}(Z) \leq d^{in}(Z) + d(X, Z) + d(Y, Z) \leq (1 + \epsilon)\beta$ since $d(W, Z) \geq (1 - \epsilon)\beta$. Hence we have $\beta(X' \cup Z', Y' \cup Z') \leq (1 + 3\epsilon)\beta$.

Next we consider the other subcase where $d^{in}(Y' \cap Z') < (1 - \epsilon)\beta$. By the choice of z_1, z_2 , this means that $Z \subseteq Y' \cap Z'$. In this case $Y' \cap Z'$ crosses X' , because X' does not contain the whole Z , and $Y' \cap Z'$ does not contain x and thus X' and $Y' \cap Z'$ are incomparable. Since $\sigma(X', Y' \cap Z') \leq (2 - \epsilon)\beta$, the minimum incomparable cut-pair is a $(2 - \epsilon)$ -approximation.

Case 2. Suppose $W \cap (X' \cup Y') \neq \emptyset$ and $Z \subseteq X' \cap Y'$. This is similar to Case 1 by symmetry.

Case 3. Suppose $W \cap (X' \cup Y') \neq \emptyset$ and $Z \not\subseteq X' \cap Y'$.

We may assume that $Z \not\subseteq X'$ without loss of generality. Because of the choice of z_1, z_2 , we have $d^{in}(X' \cap Z') \geq (1 - \epsilon)\beta$. By the same argument as in Case 1 (last paragraph), we may assume that $d^{in}(Y' \cap Z') \geq (1 - \epsilon)\beta$ as well. The inequality $\beta(X' \cup Z', Y' \cup Z') \leq 2\epsilon\beta + d_c^{in}(Z)$ holds using Proposition 5.1. If $d_c^{in}(Z) \leq (2 - 3\epsilon)\beta$, then these imply $\beta(X' \cup Z', Y' \cup Z') \leq (2 - \epsilon)\beta$. Similarly, if $d_c^{out}(W) \leq (2 - 3\epsilon)\beta$, then we obtain $\beta(X' \setminus W', Y' \setminus W') \leq (2 - \epsilon)\beta$. Thus, we may assume that both $d_c^{in}(Z)$ and $d_c^{out}(W)$ are at least $(2 - 3\epsilon)\beta$.

Let us define the following quantities (see Figure 4).




Figure 4: The quantities $\alpha_1, \dots, \alpha_6$.

1. $\alpha_1 = d(W \setminus (X' \cup Y'), W \cap (X' \setminus Y'))$,
2. $\alpha_2 = d(W \setminus (X' \cup Y'), W \cap (Y' \setminus X'))$,
3. $\alpha_3 = d(W \cap (X' \setminus Y'), Z \cap (X' \setminus Y'))$,
4. $\alpha_4 = d(W \cap (Y' \setminus X'), Z \cap (Y' \setminus X'))$,
5. $\alpha_5 = d(Z \cap (X' \setminus Y'), X' \cap Y' \cap Z')$, and
6. $\alpha_6 = d(Z \cap (Y' \setminus X'), X' \cap Y' \cap Z')$.

In propositions 5.2, 5.3, 5.4, 5.5, 5.6 and 5.7, we show a sequence of inequalities involving these quantities.

Proposition 5.2. $(1 - \epsilon)\beta \leq d^{in}(X' \cap Y'), d^{in}(X' \cup Y'), d^{in}(X' \cap Z), d^{in}(X' \cup Z) \leq (1 + \epsilon)\beta$.

Proof. By submodularity, $d^{in}(X' \cap Y') + d^{in}(X' \cup Y') \leq d^{in}(X') + d^{in}(Y') \leq 2\beta$. We note that $d^{in}(X' \cap Y') \geq (1 - \epsilon)\beta$ by the choice of z_1, z_2 . This shows $d^{in}(X' \cup Y') \leq (1 + \epsilon)\beta$. Similarly, $d^{in}(X' \cup Y') \geq (1 - \epsilon)\beta$ by the choice of w_1, w_2 , and hence $d^{in}(X' \cap Y') \leq (1 + \epsilon)\beta$.

By our assumption, X' and Z are uncomparable, hence $X' \cap Z$ contains both z_1, z_2 but not all of Z . By the choice of z_1, z_2 , $d^{in}(X' \cap Z) \geq (1 - \epsilon)\beta$. By submodularity, $d^{in}(X' \cup Z) \leq d^{in}(X') + d^{in}(Z) - d^{in}(X' \cap Z) \leq (1 + \epsilon)\beta$. For the remaining inequalities, we notice that $X' \cup Z$ and Y' are uncomparable, so $\sigma(X' \cup Z, Y') \geq (2 - \epsilon)\beta$, and therefore $d^{in}(X' \cup Z) \geq (1 - \epsilon)\beta$. Submodularity gives $d^{in}(X' \cap Z) \leq (1 + \epsilon)\beta$. \square

Proposition 5.3. $(1 - 6\epsilon)\beta \leq \alpha_3 + \alpha_4 \leq \beta$.

Proof. We recall that $d_c^{in}(Z) = d^{in}(Z) + |\delta^{in}(Z) \cap E_1| \geq (2 - 3\epsilon)\beta$ and $d_c^{out}(W) = d^{out}(W) + |\delta^{out}(W) \cap E_2| \geq (2 - 3\epsilon)\beta$. Let C be the set of edges from W to Z , i.e. those counted by $d(W, Z)$. Let $a = |\delta^{in}(Z) \setminus C|$ and $b = |\delta^{out}(W) \setminus C|$. We note that $\alpha_3 + \alpha_4 = |C \cap E_1 \cap E_2|$ and $|C| + a + b \leq \beta$.

We have $|C \cap E_1| \geq |\delta^{in}(Z) \cap E_1| - a$ and $|C \cap E_2| \geq |\delta^{out}(W) \cap E_2| - b$. From all the above, we get the following sequence of inequalities.

$$\begin{aligned}
& |C \cap E_1 \cap E_2| \\
& \geq |C| - |C \setminus E_1| - |C \setminus E_2| \\
& \geq |C| - (|C| - (|\delta^{in}(Z) \cap E_1| - a)) - (|C| - (|\delta^{out}(W) \cap E_2| - b)) \\
& = |\delta^{in}(Z) \cap E_1| + |\delta^{out}(W) \cap E_2| - |C| - a - b \\
& \geq (2 - 3\epsilon)\beta - d^{in}(Z) + (2 - 3\epsilon)\beta - d^{out}(W) - |C| - a - b \\
& \geq (4 - 6\epsilon)\beta - 3\beta \\
& = (1 - 6\epsilon)\beta.
\end{aligned}$$

□

Proposition 5.4. $(1 - 8\epsilon)\beta \leq \alpha_1 + \alpha_2 \leq (1 + \epsilon)\beta$ and $(1 - 8\epsilon)\beta \leq \alpha_5 + \alpha_6 \leq (1 + \epsilon)\beta$.

Proof. The upper bounds follow from $\alpha_1 + \alpha_2 \leq d^{in}(X' \cup Y') \leq (1 + \epsilon)\beta$ and $\alpha_5 + \alpha_6 \leq d^{in}(X' \cap Y') \leq (1 + \epsilon)\beta$. For the lower bound, we observe that $d^{in}(X' \cap Y' \cap Z) \geq (1 - \epsilon)\beta$ but $|\delta^{in}(X') \cap \delta^{in}(Y')| \leq \epsilon\beta$. This shows that at most $\epsilon\beta$ edges enter $X' \cap Y' \cap Z$ from outside of $X' \cup Y'$. By Proposition 5.3, $|\delta^{in}(Z) \cap \delta^{in}(X' \cap Y' \cap Z)| \leq \delta^{in}(Z) - \alpha_3 - \alpha_4 \leq 6\epsilon\beta$. These imply $\alpha_5 + \alpha_6 \geq (1 - \epsilon)\beta - 6\epsilon\beta - \epsilon\beta = (1 - 8\epsilon)\beta$.

A similar argument shows the lower bound for $\alpha_1 + \alpha_2$. The choice of w_1, w_2 implies $d^{out}(W \setminus (X' \cup Y')) \geq (1 - \epsilon)\beta$, and at most $\epsilon\beta$ edges enter $X' \cap Y'$ from $W \setminus (X' \cup Y')$. By Proposition 5.3, $|\delta^{out}(W) \cap \delta^{out}(W \setminus (X' \cup Y'))| \leq \delta^{out}(W) - \alpha_3 - \alpha_4 \leq 6\epsilon\beta$. These imply $\alpha_1 + \alpha_2 \geq (1 - \epsilon)\beta - 6\epsilon\beta - \epsilon\beta = (1 - 8\epsilon)\beta$. □

Proposition 5.5. $(1 - 16\epsilon)\beta \leq \alpha_1 + \alpha_6 \leq \beta$ and $(1 - 16\epsilon)\beta \leq \alpha_2 + \alpha_5 \leq \beta$.

Proof. The upper bounds follow by $\alpha_1 + \alpha_6 \leq d^{in}(X') \leq \beta$ and $\alpha_2 + \alpha_5 \leq d^{in}(Y') \leq \beta$. On the other hand, combining the two inequalities in Proposition 5.4 gives $(2 - 16\epsilon)\beta \leq \alpha_1 + \alpha_2 + \alpha_5 + \alpha_6$, so we also get the lower bounds. □

Proposition 5.6. $(1 - 23\epsilon)\beta \leq \alpha_3 + \alpha_6 \leq (1 + \epsilon)\beta$.

Proof. Consider the set $T = X' \cap Z$. By Proposition 5.2, we have $\alpha_3 + \alpha_6 \leq d^{in}(T) \leq (1 + \epsilon)\beta$, which gives the upper bound. Also, $\alpha_1 + \alpha_6 + d(Z \setminus (X' \cup Y'), T) \leq d^{in}(X') \leq \beta$, hence $d(Z \setminus (X' \cup Y'), T) \leq 16\epsilon\beta$. Since $\alpha_3 + \alpha_4 \geq (1 - 6\epsilon)\beta$, the remaining contribution to $d^{in}(Z)$ from elsewhere is at most $6\epsilon\beta$. We obtain

$$\begin{aligned}
(1 - \epsilon)\beta & \leq d^{in}(T) \\
& \leq \alpha_6 + d(Z \setminus (X' \cup Y'), T) + d(V \setminus Z, T') \\
& \leq \alpha_6 + 16\epsilon\beta + \alpha_3 + 6\epsilon\beta \\
& \leq \alpha_3 + \alpha_6 + 22\epsilon\beta.
\end{aligned}$$

Hence, $(1 - 23\epsilon)\beta \leq \alpha_3 + \alpha_6$. □

Proposition 5.7. $\alpha_1 + \alpha_5 \geq 2\alpha_3 - 51\epsilon\beta$.

Proof. The above claims give us a chain of relations:

$$\begin{aligned}
(1 - 16\epsilon)\beta - \alpha_6 &\leq \alpha_1 \leq \beta - \alpha_6, \\
(1 - 8\epsilon)\beta - \alpha_1 &\leq \alpha_2 \leq (1 + \epsilon)\beta - \alpha_1, \\
(1 - 16\epsilon)\beta - \alpha_2 &\leq \alpha_5 \leq \beta - \alpha_2, \\
(1 - 23\epsilon)\beta - \alpha_3 &\leq \alpha_6 \leq (1 + \epsilon)\beta - \alpha_3.
\end{aligned}$$

By substitution, we get

$$\begin{aligned}
\alpha_3 - 17\epsilon\beta &\leq \alpha_1 \leq \alpha_3 + 23\epsilon\beta, \\
\alpha_1 - 17\epsilon\beta &\leq \alpha_5 \leq \alpha_1 + 8\epsilon\beta, \\
\alpha_3 - 34\epsilon\beta &\leq \alpha_5 \leq \alpha_1 + 31\epsilon\beta.
\end{aligned}$$

Therefore $\alpha_1 + \alpha_5 \geq 2\alpha_3 - 51\epsilon\beta$. □

Without loss of generality, let $\alpha_3 \geq (\alpha_3 + \alpha_4)/2$, since if not, there is another iteration of the algorithm where x and y are switched. Therefore, $\alpha_3 \geq (1/2 - 3\epsilon)\beta$.

Let $A = (X' \cap Z) \cup (X' \cap Y')$ and $B = (X' \setminus W) \cup (X' \cap Y')$. Let H be the directed graph obtained by contracting $X' \cap Y'$ to a node z' , contracting $V \setminus X'$ to a node w' , and removing all $w'z'$ arcs.

Claim 5.2. $|\delta_H^{in}(A) \cup \delta_H^{in}(B)| \leq \alpha_3 + 39\epsilon\beta$.

Proof. The left hand side is precisely

$$\begin{aligned}
d(V \setminus A, (X' \cap Z) \setminus Y') + d(X' \setminus (Y' \cup Z), X' \cap Y') \\
+ d((V \setminus X') \cup (X' \cap W \setminus Y'), X' \setminus (Y' \cup W \cup Z)).
\end{aligned}$$

We consider each term individually.

1. The term $d(V \setminus A, (X' \cap Z) \setminus Y')$ counts a subset of the edges going into Z . All but $\epsilon\beta$ edges are from W . Hence $d((V \setminus A) \cap W, (X' \cap Z) \setminus Y') \geq d(V \setminus A, (X' \cap Z) \setminus Y') - \epsilon\beta$. All but at most $6\epsilon\beta$ edges counted by $d((V \setminus A) \cap W, (X' \cap Z) \setminus Y')$ are counted by α_3 . This means $d((V \setminus A) \cap W, (X' \cap Z) \setminus Y') \leq \alpha_3 + 6\epsilon\beta$. In total, it shows $d(V \setminus A, (X' \cap Z) \setminus Y') \leq \alpha_3 + 7\epsilon\beta$.
2. The term $d(X' \setminus (Y' \cup Z), X' \cap Y')$ counts a subset of the edges going into Y' . We have $d^{in}(Y') \geq d(X' \setminus (Y' \cup Z), X' \cap Y') + \alpha_2 + \alpha_5$. Using Proposition 5.5, we obtain $d(X' \setminus (Y' \cup Z), X' \cap Y') \leq 16\epsilon\beta$.
3. Let $t = d((V \setminus X') \cup (X' \cap W \setminus Y'), X' \setminus (Y' \cup W \cup Z))$. So t counts a subset of the edges going into X' . Using Proposition 5.5 and $d^{in}(X') \geq t + \alpha_1 + \alpha_6$, we obtain $t \leq 16\epsilon\beta$.

Thus, the total contribution is at most $\alpha_3 + 39\epsilon\beta$. □

Using the 3/2-approximation algorithm of Theorem 1.6, we can find $\overline{w'}z'$ -sets $A' \subsetneq B'$ such that $|\delta_H^{in}(A') \cup \delta_H^{in}(B')| \leq 3(\alpha_3 + 39\epsilon\beta)/2$. Now consider $\beta(X' \cap B', Y' \cup A')$ in the original directed graph. We have the following inequality by counting the edges on the left hand side (see Figure 5 for a proof).

$$\beta(X' \cap B, Y' \cup A') + \alpha_5 + \alpha_1 \leq \sigma(X', Y') + |\delta_H^{in}(A') \cup \delta_H^{in}(B')|. \tag{6}$$




Figure 5: Every arrow indicates the edges counted in the left hand side of (6). \circ indicates that the edge is in $\delta^{in}(X')$, \diamond indicates that the edge is in $\delta^{in}(Y')$ and \square indicates that the edge is in $\delta_H^{in}(A') \cup \delta_H^{in}(B')$. Irrelevant edges are not shown.

Using Proposition 5.7 and the assumption that $\alpha_3 \geq (1/2 - 3\epsilon)\beta$, we get

$$\begin{aligned}
\beta(X' \cap B, Y' \cup A') &\leq \sigma(X', Y') + |\delta_H^{in}(A') \cup \delta_H^{in}(B')| - \alpha_5 - \alpha_1 \\
&\leq 2\beta + \frac{3}{2}\alpha_3 + 117\epsilon\beta/2 - \alpha_5 - \alpha_1 \\
&\leq 2\beta + \frac{3}{2}\alpha_3 + 117\epsilon\beta/2 - (2\alpha_3 - 51\epsilon\beta) \\
&\leq 2\beta - \frac{1}{2}\alpha_3 + 117\epsilon\beta/2 + 51\epsilon\beta \\
&\leq (2 + (219/2)\epsilon)\beta - (\frac{1}{2} - 3\epsilon)\beta/2 \\
&= (7/4 + 111\epsilon)\beta.
\end{aligned}$$

Based on all the cases analyzed above, the approximation factor is $\max\{1 + \epsilon, 2 - \epsilon, 7/4 + 111\epsilon\}$. In order to minimize the factor, we set $\epsilon = 1/448$ to get the desired approximation factor. \square

Acknowledgements. Karthik would like to thank Chandra Chekuri, Neil Olver and Chaitanya Swamy for helpful discussions at various stages of this work. Kristóf and Tamás are supported by the Hungarian National Research, Development and Innovation Office – NKFIH grants K109240 and K120254.

References

- [1] H. Angelidakis, Y. Makarychev, and P. Manurangsi, *An Improved Integrality Gap for the Călinescu-Karloff-Rabani Relaxation for Multiway Cut*, Preprint arXiv:1611.05530, 2016.
- [2] A. Bernáth and G. Pap, *Blocking optimal arborescences*, Proceedings of the 16th International Conference on Integer Programming and Combinatorial Optimization (IPCO), 2013, pp. 74–85.

- [3] C. Chekuri and V. Madan, *Simple and fast rounding algorithms for directed and node-weighted multiway cut*, Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '16, 2016, pp. 797–807.
- [4] ———, *Simple and fast rounding algorithms for directed and node-weighted multiway cut*, To appear in Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '17, 2017.
- [5] J. Cheriyan and R. Thurimella, *Fast algorithms for k -shredders and k -node connectivity augmentation*, Journal of Algorithms **33** (1999), no. 1, 15–50.
- [6] K. Cheung, W. Cunningham, and L. Tang, *Optimal 3-terminal cuts and linear programming*, Mathematical Programming **106** (2006), no. 1, 1–23.
- [7] Miroslav Chlebík and Janka Chlebíková, *Complexity of approximating bounded variants of optimization problems*, Theoretical Computer Science **354** (2006), no. 3, 320–338.
- [8] G. Călinescu, H. Karloff, and Y. Rabani, *An improved approximation algorithm for multiway cut*, Journal of Computer and System Sciences **60** (2000), no. 3, 564–574.
- [9] E. Dahlhaus, D. Johnson, C. Papadimitriou, P. Seymour, and M. Yannakakis, *The complexity of multiterminal cuts*, SIAM Journal on Computing **23** (1994), no. 4, 864–894.
- [10] R. Erbacher, T. Jaeger, N. Talele, and J. Teutsch, *Directed multicut with linearly ordered terminals*, Preprint arXiv:1407.7498, 2014.
- [11] T. Fukunaga, *Computing minimum multiway cuts in hypergraphs from hypertree packings*, Proceedings of the 14th International Conference on Integer Programming and Combinatorial Optimization (IPCO), 2010, pp. 15–28.
- [12] N. Garg, V. Vazirani, and M. Yannakakis, *Multiway cuts in node weighted graphs*, Journal of Algorithms **50** (2004), no. 1, 49–61.
- [13] O. Goldschmidt and D. Hochbaum, *A polynomial algorithm for the k -cut problem for fixed k* , Math. Oper. Res. **19** (1994), no. 1, 24–37.
- [14] T. Jordán, *On the number of shredders*, Journal of Graph Theory **31** (1999), no. 3, 195–200.
- [15] D. Karger, P. Klein, C. Stein, M. Thorup, and N. Young, *Rounding algorithms for a geometric embedding of minimum multiway cut*, Mathematics of Operations Research **29** (2004), no. 3, 436–461.
- [16] D. Karger and C. Stein, *A new approach to the minimum cut problem*, Journal of ACM **43** (1996), no. 4, 601–640.
- [17] S. Khot, *On the power of unique 2-prover 1-round games*, Proceedings of the 34th annual ACM Symposium on Theory of Computing, STOC '02, 2002, pp. 767–775.
- [18] S. Khot and O. Regev, *Vertex cover might be hard to approximate to within $2 - \epsilon$* , Journal of Computer and System Sciences **74** (2008), no. 3, 335–349.
- [19] E. Lee, *Improved Hardness for Cut, Interdiction, and Firefighter Problems*, Preprint arXiv:1607.05133, 2016.
- [20] G. Liberman and Z. Nutov, *On shredders and vertex connectivity augmentation*, Journal of Discrete Algorithms **5** (2007), no. 1, 91–101.

- [21] R. Manokaran, J. Naor, P. Raghavendra, and R. Schwartz, *SDP Gaps and UGC Hardness for Multiway Cut, 0-extension, and Metric Labeling*, Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing, STOC '08, 2008, pp. 11–20.
- [22] E. Mossel, *Gaussian bounds for noise correlation of functions*, Geometric and Functional Analysis **19** (2010), no. 6, 1713–1756.
- [23] J. Naor and L. Zosin, *A 2-approximation algorithm for the directed multiway cut problem*, SIAM Journal on Computing **31** (2001), no. 2, 477–482.
- [24] K. Okumoto, T. Fukunaga, and H. Nagamochi, *Divide-and-conquer algorithms for partitioning hypergraphs and submodular systems*, Algorithmica **62** (2012), no. 3, 787–806.
- [25] M. Queyranne, *On Optimum k -way Partitions with Submodular Costs and Minimum Part-Size Constraints*, Talk url: <https://smartech.gatech.edu/bitstream/handle/1853/43309/Queyranne.pdf>, 2012.
- [26] A. Schrijver, *Combinatorial optimization: Polyhedra and efficiency*, Algorithms and Combinatorics, Springer, 2003.
- [27] A. Sharma and J. Vondrák, *Multiway cut, pairwise realizable distributions, and descending thresholds*, Proceedings of the 46th Annual ACM Symposium on Theory of Computing, STOC '14, 2014, pp. 724–733.
- [28] L. Tseng and N. Vaidya, *Fault-Tolerant Consensus in Directed Graphs*, Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing (PODC 2015), 2015, pp. 451–460.
- [29] L. Végh, *Augmenting undirected node-connectivity by one*, SIAM J. Discrete Math. **25** (2011), no. 2, 695–718.
- [30] M. Xiao, *Finding minimum 3-way cuts in hypergraphs*, Information Processing Letters **110** (2010), no. 14, 554–558.
- [31] L. Zhao, H. Nagamochi, and T. Ibaraki, *Greedy splitting algorithms for approximating multiway partition problems*, Mathematical Programming **102** (2005), no. 1, 167–183.