

Approximating Multicut and the Demand Graph*

Chandra Chekuri

Vivek Madan

November 2, 2016

Abstract

In the minimum Multicut problem, the input is an edge-weighted supply graph $G = (V, E)$ and a demand graph $H = (V, F)$. Either G and H are directed (DIR-MULC) or both are undirected (UNDIR-MULC). The goal is to remove a minimum weight set of supply edges $E' \subseteq E$ such that in $G - E'$ there is no path from s to t for any demand edge $(s, t) \in F$. UNDIR-MULC admits $O(\log k)$ -approximation where k is the number of edges in H while the best known approximation for DIR-MULC is $\min\{k, \tilde{O}(|V|^{11/23})\}$. These approximations are obtained by proving corresponding results on the multicommodity flow-cut gap. In this paper we consider the role that the structure of the demand graph plays in determining the approximability of Multicut. We obtain several new positive and negative results.

In undirected graphs our main result is a 2-approximation in $n^{O(t)}$ time when the demand graph excludes an induced matching of size t . This gives a constant factor approximation for a specific demand graph that motivated this work, and is based on a reduction to uniform metric labeling and not via the flow-cut gap.

In contrast to the positive result for undirected graphs, we prove that in directed graphs such approximation algorithms can not exist. We prove that, assuming the Unique Games Conjecture (UGC), that for a large class of fixed demand graphs DIR-MULC cannot be approximated to a factor better than the worst-case flow-cut gap. As a consequence we prove that for any fixed k , assuming UGC, DIR-MULC with k demand pairs is hard to approximate to within a factor better than k . On the positive side, we obtain a k approximation when the demand graph excludes certain graphs as an induced subgraph. This generalizes the known 2 approximation for directed Multiway Cut to a larger class of demand graphs.

1 Introduction

The minimum MULTICUT problem is a generalization of the classical s - t cut problem to multiple pairs. The input to the MULTICUT problem is an edge-weighted graph $G = (V, E)$ and k source-sink pairs $(s_1, t_1), (s_2, t_2), \dots, (s_k, t_k)$. The goal is to find a minimum weight subset of edges $E' \subseteq E$ such that all the given pairs are disconnected in $G - E'$; that is, for $1 \leq i \leq k$, there is no path from s_i to t_i in $G - E'$. In this paper we consider an equivalent formulation that exposes, more directly, the structure that the source-sink pairs may have.

The input now consists of an edge-weighted *supply* graph $G = (V, E)$ and a *demand* graph $H = (V, F)$. The goal is to find a minimum weight set of edges $E' \subseteq E$ such that for each edge $f = (s, t) \in F$, there is no path from s to t in $G - E'$. In other words the source-sink pairs are encoded in the form of the demand graph H . Either both G and H are directed in which case we refer to the problem as DIR-MULC (directed MULTICUT) or both are undirected in which case we refer to the problem as UNDIR-MULC (undirected MULTICUT).

MULTICUT in both directed and undirected graphs has been extensively studied. It is a natural cut problem, has several applications, and important connections to several other well-known problems such as sparsest cut and multicommodity flows. UNDIR-MULC and DIR-MULC are NP-Hard even in very restricted settings. For instance UNDIR-MULC is NP-Hard even when H has 3 edges and it generalizes VERTEX COVER even when G is a tree. DIR-MULC is NP-Hard and APX-Hard even in the special case when H is a cycle of length 2 which is better understood as removing a minimum weight set of edges to disconnect s from t and t from s in a directed graph. Consequently there has been substantial effort towards developing approximation algorithm for these problems as well as understanding special cases. We briefly summarize some of the known results. We use k to denote the number of edges in the demand graph H . For UNDIR-MULC there is an $O(\log k)$ -approximation [13] which improves to an $O(r)$ -approximation if the supply graph G excludes K_r .

*Department of Computer Science, University of Illinois, Urbana, IL 61801. {chekuri, vmadan2}@illinois.edu. Work on this paper partly supported by NSF grant CCF-1319376.

$\min \sum_{e \in E} w_e x_e$ $\sum_{e \in p} x_e \geq 1 \quad p \in \mathcal{P}_{st}, st \in F$ $x_e \geq 0 \quad e \in E$
<hr/> $\min \sum_{e \in E} w_e x_e$ $d(s, t) \geq 1 \quad st \in F$ $d(u, v) + d(v, w) - d(u, w) \geq 0 \quad u, v, w \in V$ $d(u, v) - x_e = 0 \quad e = uv \in E$ $d(u, v) \geq 0 \quad u, v \in V$

Figure 1: DISTANCE-LP for MultiCut written using exponential number of constraints (top) and in a compact fashion using additional variables (bottom)

as a minor [1, 12, 14] (in particular this yields a constant factor approximation in planar graphs). In terms of inapproximability, UNDIR-MULC is at least as hard as VERTEX COVER even in trees and hence APX-Hard. Under the Unique Game Conjecture (UGC) it is known to be super-constant hard [4]. DIR-MULC is known to be a harder problem. Assuming $NP \neq ZPP$ it is hard to approximate DIR-MULC to within a factor of $\Omega(2^{\log^{1-\epsilon} n})$ [9]; evidence is also presented in [9] that it could be hard to approximate to within an $\Omega(n^\delta)$ factor for some fixed $\delta > 0$. The best known approximation is $\min\{k, \tilde{O}(n^{11/23})\}$ [2]; here $n = |V|$. Note that a k -approximation is trivial.

We note that all the preceding positive results for MULTICUT are based on bounding the integrality gap of a natural LP relaxation shown in Figure 1. This is the standard cut formulation with a variable x_e for each edge e and an exponential set of constraints; the constraints specify that for each demand edge $st \in F$ the length of any path p between s and t is at least 1. One can solve this LP in polynomial-time using the Ellipsoid method. One can also write a compact polynomial-sized formulation using distance variables and it is shown in the same figure. The dual is a maximum multicommodity flow LP. We henceforth refer to the integrality gap of this LP as the *flow-cut gap* and the LP as DISTANCE-LP. Most multicut approximation algorithms are based on bounding the flow-cut gap.

The role of the demand graph: Our preceding discussion has focused on the approximability of MULTICUT when H is arbitrary, with some improved results when G is restricted. Here we are interested in the set-

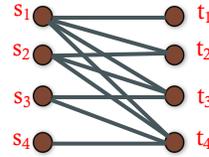


Figure 2: Demand graph for SKEW-MULTICUT

ting where H is restricted and G is arbitrary. Before we describe a concrete application that motivated us, we mention the well-known MULTIWAY CUT problem in undirected graphs (UNDIR-MULTIWAY-CUT) and directed graphs (DIR-MULTIWAY-CUT). In MULTIWAY CUT H is the complete (directed) graph on a set of k terminals. This problem has been extensively studied over the years. UNDIR-MULTIWAY-CUT admits a 1.29 approximation [21] and DIR-MULTIWAY-CUT admits a 2-approximation [19, 7]. In a recent work [5], motivated by connections to the problem of understanding the information capacity of networks with delay constraints, the following special case of MULTICUT was considered. The authors of [5] referred to it as the TRIANGLE-CAST problem since they were unaware of that the same demand graph had been studied previously in the fixed parameter tractability (FPT) literature as the SKEW-MULTICUT problem. For the sake of continuity we use the SKEW-MULTICUT nomenclature. The demand graph H is a bipartite graph with k terminals s_1, \dots, s_k on one side and k terminals t_1, \dots, t_k on the other side: (s_i, t_j) is an edge in H iff $i \leq j$. See Figure 2 for an example with $k = 4$. If the edges are directed from s 's to the t 's we obtain the directed version. It was shown in [5] that the flow-cut gap for DIR-SKEW-MULTICUT gives an upper bound on the capacity advantage of network coding with delays. Further, it was established that the flow-cut gap for DIR-SKEW-MULTICUT is $O(\log k)$ which is in contrast to the general setting where the gap can be as large as k . The following natural questions arose from this application.

QUESTION 1. *What is the approximability of UNDIR-SKEW-MULTICUT and DIR-SKEW-MULTICUT? Is the flow-cut gap $O(1)$ for UNDIR-SKEW-MULTICUT and even DIR-SKEW-MULTICUT?*

Answering the preceding questions has not been easy. In fact we do not yet know whether the flow-cut gap is $O(1)$ for UNDIR-SKEW-MULTICUT. One can show APX-hardness and a constant factor lower bound on the flow-cut gap for UNDIR-SKEW-MULTICUT (and hence also for DIR-SKEW-MULTICUT) via a reduction from UNDIR-MULTIWAY-CUT problem. In this paper we provide two results. First, we obtain a 2-approximation

for UNDIR-SKEW-MULTICUT; this result is not established via the DISTANCE-LP relaxation and requires a new LP relaxation. Second, we show that under UGC, for any fixed constant k , the hardness of approximation for DIR-SKEW-MULTICUT co-incides with its worst-case flow-cut gap. At this moment we only know that the flow-cut gap is $O(\log k)$ and at least some fixed constant $c > 1$. We mention that DIR-SKEW-MULTICUT is approximation equivalent to another problem called LIN-CUT [11]; here the demand graph H consists of k terminals s_1, \dots, s_k and there is a directed edge (s_i, s_j) for all $i < j$. As we mentioned, DIR-SKEW-MULTICUT has been considered earlier in [8] where it was shown to be in FPT parameterized by the solution size and is then used to prove that the well-known Directed Feedback Vertex set problem is in FPT. The demand pattern for SKEW-MULTICUT also arises in parameterized complexity of disjoint paths problem [18].

Our results for SKEW-MULTICUT are special cases of more general results that examine the role that the demand graph H plays in the approximability of MULTICUT. What structural aspects of H allow for better bounds than the worst-case results? For instance, can the constant factor approximation algorithms for MULTIWAY CUT be understood in a more general setting? Some previous work has also examined the role that demand graph plays in MULTICUT. Two examples are the original paper of Garg, Vazirani and Yannakakis [13] who showed that one can obtain an $O(\log h)$ -approximation for UNDIR-MULC where h is the vertex cover size of the demand graph. This was generalized by Steurer and Vishnoi [22] who showed that h can be chosen to be $\min_S \max_T |S \cap T|$ where S is a vertex cover in H and T is an independent set in H . Note that both these results are based on the flow-cut gap and yield only an $O(\log k)$ upper bound for UNDIR-SKEW-MULTICUT.

We now describe our results for both UNDIR-MULC and DIR-MULC which yield, as corollaries, the results that we already mentioned and several others.

1.1 Overview of Results We first discuss our result for UNDIR-MULC. We obtain a 2-approximation for a class of demand graphs. This class is inspired by the observation that the SKEW-MULTICUT demand graph does not contain a matching with two edges as an *induced subgraph*¹. More generally, a graph is said to be tK_2 -free for an integer $t > 1$ if it does not contain a matching of size t as an induced subgraph.

¹ G' is an induced subgraph of $G = (V, E)$ if $G' = G[V']$ for some $V' \subset V$.

THEOREM 1.1. *There is a 2-approximation algorithm with running time $\text{poly}(n, k^{O(t)})$ on instances of UNDIR-MULC with supply graph G and tK_2 -free demand graph H . Here $n = V(G)$ and $k = V(H)$.*

Since demand graph in UNDIR-SKEW-MULTICUT instances are $2K_2$ -free we obtain the following corollary.

COROLLARY 1.1. *UNDIR-SKEW-MULTICUT admits a polynomial-time 2-approximation.*

Note that a graph containing $t - 1$ parallel edges is tK_2 -free. Via known lower bounds [13], DISTANCE-LP has integrality gap $\Omega(\log(t-1))$ on tK_2 -free demand graphs. Thus, one cannot obtain a constant approximation for all fixed t using DISTANCE-LP. Our algorithm relies on a different relaxation and a reduction to uniform metric labeling [15]. Also, as we noted earlier, assuming UGC, UNDIR-MULC is NP-hard to approximate within a constant factor [4]. Hence, assuming UGC, it is unlikely that one can obtain a fixed constant factor approximation for UNDIR-MULC with tK_2 -free demand graphs in $\text{poly}(n, k, t)$ time.

We now turn our attention to DIR-MULC. As we mentioned earlier, DIR-MULC admits $\tilde{O}(n^{11/23})$ approximation [2] and is hard to approximate within a factor of $\Omega(2^{\log^{1-\epsilon} n})$ assuming $NP \neq ZPP$ [9]. A k approximation is trivial. It is natural to ask whether we can improve the bound of k when k is a fixed constant; however, it is known that the worst-case flow-cut gap in this case is k [20], and hence one needs a different relaxation. What about the case of a specific demand graph H such as the one for DIR-SKEW-MULTICUT? To formalize this, for a *fixed* demand graph H , we define the problem DIR-MULC- H as the special case of DIR-MULC where G is arbitrary but the demand graph is constrained to be H . To be formal we need to define H as a “pattern” since we need to specify the nodes of G to which the nodes of H are mapped. However we avoid further notation and assume that $V(H) \subset V(G)$. We define α_H to be the worst-case flow-cut gap over all instances with demand graph H . We conjecture the following general hardness of approximation result.

CONJECTURE 1. *For any fixed demand graph H and any fixed $\epsilon > 0$, unless $P = NP$, there is no polynomial-time $(\alpha_H - \epsilon)$ -approximation for DIR-MULC- H .*

In this paper we prove weaker forms of the conjecture, captured in the following two theorems:

THEOREM 1.2. *Assuming UGC, for any fixed directed bipartite graph H , and for any fixed $\varepsilon > 0$ there is no polynomial-time $(\alpha_H - \varepsilon)$ approximation for DIR-MULC-H.*

THEOREM 1.3. *Assuming UGC, for any fixed directed graph H on k vertices and for any fixed $\varepsilon > 0$, there is no polynomial-time $\frac{\alpha_H}{2^{\lceil \log k \rceil}} - \varepsilon$ approximation for DIR-MULC-H.*

Via known flow-cut gap results [20] and some standard reductions we obtain the following corollary.

COROLLARY 1.2. *Assuming UGC, for any fixed $\varepsilon > 0$ the following hold.*

- *For any fixed k , if H is a collection of k disjoint directed edges then DIR-MULC-H is hard to approximate within a factor of $k - \varepsilon$.*
- *Separating s from t and t from s in a directed graph (DIR-MULTIWAY-CUT with 2 terminals) is hard to approximate within a factor of $2 - \varepsilon$.*
- *For any fixed k , DIR-SKEW-MULTICUT's approximability coincides with its flow-cut gap.*

In a recent independent work, Lee [16] has also shown that for fixed k , assuming UGC, DIR-MULC-H is $k - \varepsilon$ hard when H is a collection of k disjoint edges.

Our last result is on upper bounds for DIR-MULC. Corollary 1.2 shows that if H contains, a matching of size k as an induced subgraph then we cannot obtain a better than k approximation. The following question arises naturally.

QUESTION 2. *Let H be a fixed demand graph that does not contain a matching of size k as an induced subgraph. Is there a k -approximation for DIR-MULC-H?*

A positive answer to above question would imply a 2-approximation for DIR-SKEW-MULTICUT. Since we are currently unable to improve the $O(\log k)$ -approximation for DIR-SKEW-MULTICUT, we consider a relaxed version of the preceding question and give a positive answer. We say that a directed demand graph $H = (V, F)$ contains an induced k -matching-extension if there are two subsets of V , $S = \{s_1, s_2, \dots, s_k\}$ and $T = \{t_1, t_2, \dots, t_k\}$ satisfying the following properties: (i) for $1 \leq i \leq k$, $(s_i, t_i) \in F$ and (ii) for $i > j$, $(s_i, t_j) \notin F$. Note that s_1, s_2, \dots, s_k are distinct since S is a set and similarly t_1, \dots, t_k are distinct but some s_i may be the same as a t_j for $i \neq j$. Intuitively speaking, k -matching extension is a graph which lies between a matching and DIR-SKEW-MULTICUT demand pattern.

We give two examples to illustrate the utility of considering this special case of instances.

Consider DIR-MULTIWAY-CUT which corresponds to the demand graph H being a complete directed graph. It can be verified that H does *not* contain an induced 3-matching-extension. Now consider $H = (V, F)$ being a complete graph on an even number of nodes and remove edges from H corresponding to a perfect matching M on V (if $uv \in M$ we remove (u, v) and (v, u) from F). We claim that the resulting graph H' does not contain an induced 4-matching-extension. What is the approximability of DIR-MULC-H'? It is fair to say that previous work would not have found it easy to answer this question since H' does not appear to have the same nice structure that the complete graph has. The theorem below shows that one can obtain a 3 approximation for DIR-MULC-H'.

THEOREM 1.4. *Consider DIR-MULC-H where H does not contain an induced k -matching-extension. Then the flow-cut gap is at most $k - 1$ and there is a polynomial-time rounding algorithm that achieves this upper bound.*

The rounding scheme that proves the preceding theorem is built upon our recent insight for DIR-MULTIWAY-CUT [7]. Interestingly the rounding scheme is itself *oblivious* to the demand graph H . It either provably obtains a $(k - 1)$ approximation via the LP solution or provides a certificate that H contains an induced k -matching-extension.

Techniques: At a high-level our results are based on a *labeling* view for MULTICUT. For undirected graphs we show that this yields Theorem 1.1. In directed graphs we show that a labeling based LP is no more powerful than DISTANCE-LP which is stark contrast to the undirected graph setting. The labeling LP allows us to relate the hardness of DIR-MULC-H to the hardness of a min constraint satisfaction problems (Min CSPs) via a standard labeling LP for CSPs called BASIC-LP. We crucially rely on a general hardness result for MIN- β -CSP due to Ene, Vondrak and Wu [10] that generalized prior work of Manokaran *et al.* [17]. Finally, Theorem 1.4 builds upon our recent work on a simpler rounding scheme for DIR-MULTIWAY-CUT [7].

Organization: Section 2 describes the factor 2-approximation for UNDIR-MULC with tK_2 -free demand graph. Section 3 describes the hardness of approximation results for DIR-MULC-H. Section 4 describes the $(k - 1)$ approximation for DIR-MULC-H when H does not contain an induced k -matching extension.

2 Approximating Undir-MulC with tK_2 -free demand graph

In this section we obtain 2-approximation for tK_2 -free demand graphs and prove Theorem 1.1.

THEOREM 1.1. *There is a 2-approximation algorithm with running time $\text{poly}(n, k^{O(t)})$ on instances of UNDIR-MULC with supply graph G and tK_2 -free demand graph H . Here $n = V(G)$ and $k = V(H)$.*

Before we prove the theorem, we consider the UNDIR-MULC problem where the demand graph has k vertices. Given supply graph $G = (V, E)$ let $S = \{s_1, \dots, s_k\} \subset V$ be the *terminals* participating in the demand edges specified by H . A feasible solution $E' \subset E_G$ of the UNDIR-MULC instance will induce a partition over S such that if $s_i s_j$ is an edge in the demand graph H , then s_i and s_j belong to different components in $G - E'$. Note that two terminals that are not connected by a demand edge may be in the same connected component of $G - E'$. If k is a fixed constant we can “guess” the partition of the terminals induced by an optimum solution. With the guess in place it is easy to see that the problem reduces to an instance of UNDIR-MULTIWAY-CUT which admits a constant factor approximation. Thus, one can obtain a constant factor approximation for UNDIR-MULC in $2^{O(k \log k)} \text{poly}(n)$ time by trying all possible partitions of the terminals.

To prove Theorem 1.1, we use this idea of enumerating feasible partitions. However, H is not necessarily of fixed size, and enumerating all possible partitions of the terminals is not feasible. Instead, we make use of the following theorem which bounds the number of *maximal* independent sets in a tK_2 -free graph.

THEOREM 2.1. *(Balas and Yu [3]) Any s -vertex tK_2 -free graph has at most $s^{O(t)}$ maximal independent sets and these can be found in $s^{O(t)}$ time.*

We prove Theorem 1.1 by using the preceding theorem and reducing the UNDIR-MULC problem to the UNIFORM-METRICLABELING problem. We now describe the general METRICLABELING problem.

METRICLABELING: The input consists of an undirected edge-weighted graph $G = (V, E)$, a set of labels $L = \{1, \dots, h\}$ and a metric $d(i, j), i, j \in L$ defined over the labels. In addition, for each vertex $u \in V$ and label $i \in L$ there is a non-negative assignment cost $c(u, i)$. Given an assignment $f : V \rightarrow L$ of vertices to labels we define its cost as $\sum_{u \in V} c(u, f(u)) + \sum_{uv \in E} w(uv) d(f(u), f(v))$. The goal is to find an assignment of minimum cost. The special case when the metric is uniform, that is

$d(i, j) = 1$ for $i \neq j$, is referred to as UNIFORM-METRICLABELING.

THEOREM 2.2. *(Kleinberg and Tardos [15]) There is a 2-approximation for UNIFORM-METRICLABELING.*

Proof of Theorem 1.1: Let the demand graph H of the UNDIR-MULC instance be tK_2 -free. Using Theorem 2.1, we can find all maximal independent sets in H . Let these independent sets be I_1, \dots, I_r where $r \leq |V_H|^{O(t)}$. Note that the independent sets are considered only in the demand graph. Note that instance also need to specify the mapping of vertices of H to vertices of G . However, for ease of notation we will simply assume that $V_H \subset V_G$.

Consider the following instance of UNIFORM-METRICLABELING: The supply graph $G = (V, E)$ of the UNDIR-MULC instance is the input graph to the UNIFORM-METRICLABELING instance. The label set $L = \{1, 2, \dots, r\}$, one for each maximal independent set in H . For each $u \in V_H$ let $c(u, i) = 0$ if $u \in I_i$ and $c(u, i) = \infty$ otherwise. And for each $u \in V \setminus V_H$, $c(u, i) = 0$ for all i .

We claim that the preceding reduction is approximation preserving. Assuming the claim, we can obtain the desired 2-approximation by solving the UNIFORM-METRICLABELING instance using Theorem 2.2. The size of the UNIFORM-METRICLABELING instance that is generated from the given UNDIR-MULC instance is $\text{poly}(n, |V_H|^{O(t)})$ which explains the running time. We now prove the claim.

Let $f : V \rightarrow L$ be an assignment of labels to the nodes whose cost is finite (such an assignment always exists since each terminal is in some independent set). Let $E' \subset E$ be the set of edges “cut” by this assignment; that is, $uv \in E'$ iff $f(u) \neq f(v)$. The cost of this assignment is equal to the weight of E' since the metric is uniform and the labeling costs are 0 or ∞ . We argue that E' is a feasible solution for the UNDIR-MULC instance. Suppose not. Then there are terminals u, v such that uv is an edge in the demand graph H and u, v belong to the same connected component of $G - E'$. The label $j = f(u)$ corresponds to a maximal independent set I_j in H which means that $v \notin I_j$. Thus $f(v) \neq j$ since $c(v, j) = \infty$. Therefore, u and v are assigned different labels and cannot be in the same connected component.

Conversely, let $E' \subset E$ be a feasible solution for UNDIR-MULC instance and let V_1, \dots, V_ℓ be vertex sets of the connected components of $G - E'$. Let T_j be the terminals in V_j . Since, all pairs of terminals connected by an edge in H are separated in $G - E'$, T_j must be an

independent set in H . For each T_j , consider a maximal independent set in H containing all the vertices of T_j ; pick arbitrary one if more than one exists. Let this independent set be I_{i_j} . We construct a labeling f by labeling all vertices of V_j by label i_j . It is easy to see that all terminals are assigned a label corresponding to an independent set in H containing that terminal. Hence, labeling cost is equal to zero. Also, all vertices corresponding to same connected component in $G - E'$ are assigned the same label. Hence, cost of the edges cut by the assignment f is at most the cost of the edges in E' . \square

3 UGC-based hardness of approximation results for Dir-MulC

In this section we prove hardness of approximation for DIR-MULC-H, in particular Theorem 1.2 relating the hardness of approximation to the flow-cut gap. Recall that α_H is the worst-case flow-cut gap (equivalently, the integrality gap of the DISTANCE-LP) for instances of DIR-MULC-H.

THEOREM 1.2. *Assuming UGC, for any fixed directed bipartite graph H , and for any fixed $\varepsilon > 0$ there is no polynomial-time $(\alpha_H - \varepsilon)$ approximation for DIR-MULC-H.*

We prove the theorem via a reduction to MIN- β -CSP and the hardness result of Ene, Vondrák and Wu [10]. We note that the result is technical and involves several steps. This is partly due to the fact that the theorem is establishing a meta-result. The theorem of [10] is in a similar vein. In particular [10] establishes that the hardness of MIN- β -CSP depends on the integrality gap of a specific LP formulation BASIC-LP (defined later). Our proof is based on establishing a correspondence between DIR-MULC-H and a specific constraint satisfaction problem MIN- β_H -CSP where β_H is constructed from H (this is the heart of the reduction) and proving the following properties:

- (I) Establish approximation equivalence between DIR-MULC-H and MIN- β_H -CSP. That is, prove that each of them reduces to the other in an approximation preserving fashion.
- (II) Prove that if the flow-cut gap for DIR-MULC-H (equivalently the integrality gap of DISTANCE-LP) is α_H then the integrality gap of BASIC-LP for MIN- β_H -CSP is also α_H .

From (I), we obtain that the hardness of approximation factor for DIR-MULC-H and MIN- β_H -CSP coincide. From (II), we can apply the result in [10] which shows that, assuming UGC, the hardness of approximation

for MIN- β_H -CSP is the same as the integrality gap of BASIC-LP. Putting together these two claims give us our desired result.

It is not straightforward to directly relate DISTANCE-LP for DIR-MULC-H and BASIC-LP for MIN- β_H -CSP. BASIC-LP appears to be stronger on first glance. In order to relate them we show that a seemingly strong LP for DIR-MULC that we call LABEL-LP is in fact no stronger than DISTANCE-LP. In fact this can be seen as the key technical fact underlying the entire proof and is independently interesting since it is quite different from the undirected graph setting. It is much easier to relate LABEL-LP and BASIC-LP. The rest of this section is organized as follows. In Section 3.1 we describe LABEL-LP and prove its equivalence with DISTANCE-LP. In Section 3.2 we describe MIN- β -CSP and BASIC-LP and formally state the theorem of [10] that we rely on. We then subsequently describe our reduction from DIR-MULC-H to MIN- β_H -CSP and complete the proof.

3.1 Label-LP and equivalence with Distance-LP for Dir-MulC

In Section 2, we saw that if demand graph H has size k , then there is a labeling LP for MULTICUT (the undirected problem) with size $\text{poly}(2^k, n)$ and integrality gap at most 2 which improves upon the integrality gap of DISTANCE-LP which can be $\Omega(\log k)$. Here we describe a natural labeling LP for DIR-MULC (LABEL-LP), but in contrast to the undirected case, we show that it is not stronger than DISTANCE-LP. We show this equivalence on an instance by instance basis. That is, for any DIR-MULC instance I , given a solution to DISTANCE-LP, we can find a solution to LABEL-LP with same cost and vice versa.

Let the demand graph be H with vertex set $V_H = \{s_1, \dots, s_k\}$, and the supply graph be $G = (V_G, E)$ with n vertices. We will assume here, for ease of notation, that $V_H \subset V_G$. Define a labeling set $L = \{0, 1\}^k$ which corresponds to all subsets of V_H . We interpret the labels in L as k -length bit-vectors; if $\sigma \in L$ we use $\sigma[i]$ to denote the i 'th bit of σ . For two labels $\sigma_1, \sigma_2 \in L$ we say $\sigma_1 \leq \sigma_2$ if $\forall i, \sigma_1[i] \leq \sigma_2[i]$. To motivate the formulation consider any set of edges $E' \subseteq E$ that can be cut. In $G' = G - E'$ we consider, for each $v \in V$, the reachability information from each of the terminals s_1, s_2, \dots, s_k . For each v this can be encoded by assigning a label $\sigma_v \in L$ where $\sigma_v[i] = 1$ iff v is reachable from s_i in G' . E' is a feasible solution if s_i cannot reach s_j whenever (s_i, s_j) is an edge of H . The goal of the formulation to assign labels to vertices and to ensure that demand pairs are separated. An edge $e = (u, v)$ is cut if there is some s_i such that s_i can reach

u but s_i cannot reach v . We add several constraints to ensure that the label assignment is consistent. The basic variables are $z_{v,\sigma}$ for each $v \in V_G$ and $\sigma \in L$ which indicate whether v is assigned the label σ . We also a variable x_e for each edge $e = (u, v) \in E_G$ that is derived from the label assignment variables. We start with the basic constraints involving these variables and then add additional variables that ensure consistency of the assignment.

- Each vertex is labelled by exactly one label. For $v \in V_G$, $\sum_{\sigma \in L} z_{v,\sigma} = 1$.
- Vertex s_i is reachable from s_i . For $s_i \in V_H$ and any $\sigma \in L$ such that $\sigma[i] = 0$, $z_{s_i,\sigma} = 0$
- Demand edges are separated. That is, if $(s_i, s_j) \in E_H$, then s_j is not reachable from s_i . That is, $z_{s_j,\sigma} = 0$ for any σ where $\sigma[i] = 1$ and $(s_i, s_j) \in E_H$.

For each edge $e = (u, v)$ we have variables of the form $z_{e,\sigma_1\sigma_2}$ where the intention is that u is labeled σ_1 and v is labeled σ_2 . To enforce consistency between edge assignment variables and vertex assignment variables we add the following set of constraints.

- For $e = (u, v) \in E_G$, $z_{u,\sigma_1} = \sum_{\sigma_2 \in L} z_{e,\sigma_1\sigma_2}$ and $z_{v,\sigma_2} = \sum_{\sigma_1 \in L} z_{e,\sigma_1\sigma_2}$.

Finally, the auxiliary variable x_e indicates whether e is cut.

- For $e = (u, v) \in E_G$, $x_e = 1$ if for some i , u is reachable from s_i and v is not reachable from s_i . Then, $x_e = 1$ if $z_{e,\sigma_1\sigma_2} = 1$ for $\sigma_1 \not\leq \sigma_2$. We thus set $x_e = \sum_{\sigma_1, \sigma_2 \in L: \sigma_1 \not\leq \sigma_2} z_{e,\sigma_1\sigma_2}$.

It is not hard to show that if one constraints all the variables to be binary then the resulting integer program is a valid formulation for DIR-MULC. Note that the number of variables is exponential in $k = |V_H|$. Relaxing the integrality constraint on the variables, we get LABEL-LP as shown in Fig 3.

THEOREM 3.1. *For any instance G, H of DIR-MULC-H, the optimum solution values for the formulations LABEL-LP and DISTANCE-LP are the same both in the fractional and integral settings.*

LABEL-LP has similarities to the earth-mover LP for metric labeling considered in [15, 6] except that the “distance” between labels is not a metric. Define a cost function $c : L \times L \rightarrow \{0, 1\}$ as follows: $c(\sigma, \sigma') = 0$ if $\sigma \leq \sigma'$ and 1 otherwise. In fact, given the basic labeling variables $z_{v,\sigma}$, other variables are decided in a min-cost solution. We explain this formally.

Interpreting Variables $z_{e,\sigma_1\sigma_2}$ and x_e as flow: Let $e = (u, v)$ be an edge in G . Consider a directed complete bipartite digraph B_{uv} with vertex set $\Gamma_u = \{u_\sigma \mid \sigma \in L\}$

Label-LP

$$\begin{aligned} \min \quad & \sum_{e \in E} w_e x_e \\ \sum_{\sigma \in L} z_{v,\sigma} &= 1 \quad v \in V_G, \sigma \in L \\ z_{s_i,\sigma} &= 0 \quad s_i \in V_H, \sigma \in L, \sigma[i] = 0 \\ z_{s_j,\sigma} &= 0 \quad \sigma \in L, \sigma[i] = 1, (s_i, s_j) \in E_H \\ \sum_{\sigma_2 \in L} z_{e,\sigma_1\sigma_2} &= z_{u,\sigma_1} \quad e = (u, v) \in E_G, \sigma_1 \in L \\ \sum_{\sigma_1 \in L} z_{e,\sigma_1\sigma_2} &= z_{v,\sigma_2} \quad e = (u, v) \in E_G, \sigma_2 \in L \\ \sum_{\sigma_1, \sigma_2 \in L: \sigma_1 \not\leq \sigma_2} z_{e,\sigma_1\sigma_2} &= x_e \quad e \in E_G \\ z_{v,\sigma}, z_{e,\sigma_1\sigma_2} &\leq 1 \quad v \in V_G, e \in E_G, \sigma, \sigma_1, \sigma_2 \in L \\ z_{v,\sigma}, z_{e,\sigma_1\sigma_2} &\geq 0 \quad v \in V_G, e \in E_G, \sigma, \sigma_1, \sigma_2 \in L \end{aligned}$$

Figure 3: LABEL-LP for DIR-MULC

and $\Gamma_v = \{v_\sigma \mid \sigma \in L\}$. We assign cost $c(\sigma, \sigma')$ on the edge $(u_\sigma, v_{\sigma'})$. We assign a supply of $z_{u,\sigma}$ on the vertex u_σ and a demand of $z_{v,\sigma}$ on the vertex v_σ . The values $z_{e,\sigma_1\sigma_2}$ can be thought of as flow from u_{σ_1} to v_{σ_2} satisfying the following properties: (i) total flow out of u_{σ_1} must be equal to the supply z_{u,σ_1} ($z_{u,\sigma_1} = \sum_{\sigma_2 \in L} z_{e,\sigma_1\sigma_2}$) (ii) total flow into v_{σ_2} must be equal to z_{v,σ_2} ($z_{v,\sigma_2} = \sum_{\sigma_1 \in L} z_{e,\sigma_1\sigma_2}$) (iii) flow is non-negative ($z_{e,\sigma_1\sigma_2} \geq 0$). The cost of the flow according to c is precisely x_e ($= \sum_{\sigma_1 \not\leq \sigma_2} z_{e,\sigma_1\sigma_2}$). In particular, given the values of the labeling variables $z_{u,\sigma}$, $\sigma \in L$ and $z_{v,\sigma'}$, $\sigma' \in L$ which can be thought of as two distributions on the labels, the smallest value of x_e that can be achieved is basically the min-cost flow in B_{uv} with supplies and demands defined by the two distributions. In other words the other variables are completely determined by the distributions if one wants a minimum cost solution.

In the sequel we use \bar{z}_u to denote the vector of assignment value $z_{u,\sigma}$, $\sigma \in L$ and refer to \bar{z}_u as the distribution corresponding to u . We present the high-level reduction between the solutions of the two LP’s and refer to Section A for the full proof.

From Label-LP to Distance-LP: Let (\mathbf{x}, \mathbf{z}) be a feasible solution to LABEL-LP for an instance (G, H) . This solution satisfies the following two conditions: (i) If $(s_i, s_j) \in E_H$, then $\sum_{\sigma \in \{0,1\}^k: \sigma[i]=1} z_{s_i,\sigma} = 1$ and $\sum_{\sigma \in \{0,1\}^k: \sigma[i]=1} z_{s_j,\sigma} = 0$. (ii) For an edge $e = (u, v) \in E_G$, and any terminal s_i , $x_e \geq \sum_{\sigma \in \{0,1\}^k: \sigma[i]=1} z_{v,\sigma} - \sum_{\sigma \in \{0,1\}^k: \sigma[i]=1} z_{u,\sigma}$.

Suppose $(s_i, s_j) \in E_H$ and $s_i, a_1, \dots, a_t, s_j$ is a path from s_i to s_j in G . Then, plugging in the above inequalities for the edges of the path, we get $\sum_{e \in p} x_e \geq$

1. Hence, \mathbf{x} is a feasible solution to DISTANCE-LP and has same cost as (\mathbf{x}, \mathbf{z}) .

From Distance-LP to Label-LP: Let \mathbf{x} be a feasible solution to DISTANCE-LP. We obtain a label assignment \mathbf{z}' as follows. For a vertex $u \in V_G$, let $d(s_1, u) \leq d(s_2, u) \leq \dots \leq d(s_k, u)$ (if not, rename the terminals accordingly). Here, $d(u, v)$ denotes the shortest path distance from u to v as per lengths x_e . For $i \in [0, k]$, let $\sigma_i = 1^i 0^{k-i}$. Then, set $z'_{u, \sigma_0} = d(s_1, u)$, $z'_{u, \sigma_k} = 1 - d(s_k, u)$ and for $i \in [1, k-1]$, $z'_{u, \sigma_i} = d(s_{i+1}, u) - d(s_i, u)$. For $\sigma \notin \{\sigma_0, \dots, \sigma_k\}$, $z'_{u, \sigma} = 0$. Once the label assignment for vertices is defined, we obtain the values of other variables by considering each edge $e = (u, v)$ and using the min-cost flow between \bar{z}'_u and \bar{z}'_v as described earlier. In Section A, we prove that this flow has cost ($:= x'_e$) at most $\max_{i \in [1, k]} d(s_i, v) - d(s_i, u)$ which is upper bounded by x_e . Hence, cost of solution $(\mathbf{x}', \mathbf{z}')$ to LABEL-LP is upper bounded by cost of \mathbf{x} .

3.2 Min-CSP and Basic-LP Min-CSP refers to a minimization version of constraint satisfaction problems. We set up the formalism borrowed from [10]. Let L denote a finite set of labels. A real-valued function $f : L^i \rightarrow \mathbb{R}$ has *arity* i . Let $\Gamma = \{\psi \mid \psi : L^i \rightarrow [0, 1] \cup \{\infty\}, i \leq k\}$ be the set of functions defined on L with arity at most k and range $[0, 1] \cup \{\infty\}$. Let $\beta \subset \Gamma$ be a finite subset of ψ . These functions are also referred to as *predicates*. k denotes the arity and L denotes the alphabet of β . Each β induces an optimization problem MIN- β -CSP.

DEFINITION 3.1. *An instance of Min- β -CSP consists of the following:*

- A vertex set V and a set of tuples $T \subset \cup_{i=1}^k V^i$.
- A predicate $\psi_t \in \beta$ for each tuple $t \in T$ where cardinality of t matches the arity of ψ_t .
- A non-negative weight function over the set of tuples, $w : T \rightarrow \mathbb{R}^+$.

The goal is to find a label assignment $\ell : V \rightarrow L$ to minimize $\sum_{t=(v_{i_1}, \dots, v_{i_j}) \in T} w_t \cdot \psi_t(\ell(v_{i_1}), \dots, \ell(v_{i_j}))$.

Consider an integer programming formulation with following variables: for each vertex $v \in V$ and label $\sigma \in L$, we have a variable $z_{v, \sigma}$ which is 1 if v is assigned label σ . Also, for each tuple $t = (v_{i_1}, \dots, v_{i_j}) \in T$ and $\alpha \in L^{|t|}$, we have a boolean variable $z_{t, \alpha}$ which is 1 if v_{i_p} is labelled $\alpha[p]$ for $p \in [1, j]$. These variables satisfy following constraints:

- Each vertex receives unique label: $\sum_{\sigma \in L} z_{v, \sigma} = 1$.
- Variables $z_{v, \sigma}$ and $z_{t, \alpha}$ are consistent. That is, if $v \in t$ is assigned label σ , then $z_{t, \alpha}$ must be zero if α does not assign label σ to v . For every

Basic-LP

$$\begin{aligned} \min \quad & \sum_{t \in T} w_t \cdot \sum_{\alpha \in L^{|t|}} z_{t, \alpha} \cdot \psi_t(\alpha) \\ \sum_{\sigma \in L} z_{v, \sigma} &= 1 \quad v \in V \\ \sum_{\alpha \in L^{|t|} : \alpha[i] = \sigma} z_{t, \alpha} &= z_{v, \sigma} \quad t \in T, v = t[i], \sigma \in L \\ z_{v, \sigma}, z_{t, \alpha} &\geq 0 \quad v \in V, \sigma \in L, t \in T, \alpha \in L^{|t|} \\ z_{v, \sigma}, z_{t, \alpha} &\leq 1 \quad v \in V, \sigma \in L, t \in T, \alpha \in L^{|t|} \end{aligned}$$

Figure 4: Basic LP for Min- β -CSP

tuple $t \in T, v = t[i], \sigma \in L$, we have: $z_{v, \sigma} = \sum_{\alpha \in L^{|t|} : \alpha[i] = \sigma} z_{t, \alpha}$.

The objective is minimize $\sum_{t \in T} w_t \cdot \sum_{\alpha \in L^{|t|}} z_{t, \alpha} \cdot \psi_t(\alpha)$.

BASIC-LP is the LP relaxation obtained by allowing the variables to take on values in $[0, 1]$ and is described in Fig 4. For instance \mathcal{I} , $LP(\mathcal{I})$ and $OPT(\mathcal{I})$ refer to the fractional and integral optimum values respectively.

A particular type of predicate termed NAE (for *not all equal*) is important in subsequent discussion.

DEFINITION 3.2. *For $i \geq 2$, $NAE_i : L^i \rightarrow \{0, 1\}$ be a predicate such that $NAE_i(\sigma_1, \dots, \sigma_i) = 0$ if $\sigma_1 = \sigma_2 = \dots = \sigma_i$ and 1 otherwise.*

The following theorem shows that the hardness of MIN- β_H -CSP coincides with the integrality gap of BASIC-LP if NAE_2 is in β .

THEOREM 3.2. *(Ene, Vonnak, Wu [10]) Suppose we have a MIN- β -CSP instance $\mathcal{I} = (V, T, \Psi, w)$ with fractional optimum (of Basic LP) $LP(\mathcal{I}) = c$, integral optimum $OPT(\mathcal{I}) = s$, and β contains the predicate NAE_2 . Then, assuming UGC, for any ϵ , for some $\lambda > 0$, it is NP-hard to distinguish between instances of MIN- β -CSP where the optimum value is at least $(s - \epsilon)\lambda$ and instances where the optimum value is less than $(c + \epsilon)\lambda$.*

3.3 Dir-MulC-H and an equivalent Min- β -CSP Problem

In this section, we show that given a bipartite directed graph $H = (S \cup T, E_H)$, we can construct a set of predicates β_H such that DIR-MULC-H is equivalent to MIN- β_H -CSP. The notion of equivalence is as follows. We give a reduction from instances of DIR-MULC-H to instances of MIN- β_H -CSP which preserves the cost of optimal integral solution and in addition also preserves the cost of optimum fractional solution to LABEL-LP and BASIC-LP. Similarly we give a reduction from MIN- β_H -CSP to DIR-MULC-H which

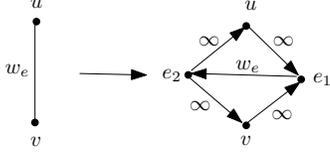


Figure 5: Gadget to convert undirected edge/ NAE_2 predicate to a directed graph

preserves the cost of both the integral and fractional solutions.

The basic idea behind the construction of β_H from H is to simulate the constraints of LABEL-LP via the predicates of β_H . In addition to setting up β_H correctly, we also need to preprocess the supply graph to prove the correctness of the reductions. Let the bipartite demand graph H be $(S \cup T, E_H)$ with $S = \{a_1, \dots, a_p\}$ and $T = \{b_1, \dots, b_q\}$ as the bipartition. For $u \in S$ let $N_H^+(u) = \{v \in T \mid (u, v) \in E_H\}$ be the neighbors of u in H . For $i \in [1, p]$, let $Y_i = \{j \in [1, p] \mid N_H^+(a_j) \subseteq N_H^+(a_i)\}$. That is, if $a_j \in Y_i$, the set of terminals that a_j needs to be separated from is a subset of the terminals that a_i needs to be separated from. For $j \in [1, q]$ let $Z_j = \{i \in [1, p] \mid a_i b_j \notin E_H\}$. That is, Z_j is the set of all terminals in S that do not need to be separated from b_j .

Assumptions on supply graph: We will assume that the supply graph G in the instances of DIR-MULC-H satisfy the following properties.

- Assumption I: G may contain *undirected* edges. The meaning of this is that a path may include this edge in either direction. A simple and well-known gadget shown in Fig 5 shows that this is without loss of generality.
- Assumption II: For $1 \leq j \leq q$ and $i \in Z_j$, there is an infinite weight edge from a_i to b_j in G . Moreover b_j has no outgoing edge.
- Assumption III: For $1 \leq i \leq p$, and $i' \in Y_i$, there is an infinite weight edge from $a_{i'}$ to a_i in G . Moreover a_i has no other incoming edges.

The preceding assumptions are to make the construction of β_H and the subsequent proof of equivalence with DIR-MULC-H somewhat more transparent and technically easier. Undirected edges allow us to use the NAE_2 predicate in β_H . Assumption II and III simplify the reachability information of terminals that needs to be kept track of and this allows for a simpler label set definition and easier proof of equivalence.

DISTANCE-LP easily generalizes to handle undirected edges; in examining paths from s_i to t_i for a demand pair we allow an undirected edge to be used in both directions. A more technical part is to generalize LABEL-LP to handle undirected edges in the

supply graph. For a directed edge e recall that $x_e = \sum_{\sigma_1, \sigma_2 \in L: \sigma_1 \not\leq \sigma_2} z_{e, \sigma_1 \sigma_2}$. For an undirected edge e we set $x_e = \sum_{\sigma_1, \sigma_2 \in L: \sigma_1 \neq \sigma_2} z_{e, \sigma_1 \sigma_2}$. See Section B for the justification of the assumptions.

Constructing β_H from H : Next, we formally define β_H for a bipartite graph $H = (S \cup T, E_H)$ where $S = \{a_1, \dots, a_p\}$ and $T = \{b_1, \dots, b_q\}$. Recall the definitions of Y_i for $1 \leq i \leq p$ and Z_j for $1 \leq j \leq q$ based on E_H . Observe that no vertex other than b_j is reachable from b_j . And, since labels encode the reachability from terminals, we can ignore the reachability from b_j and define β_H with respect to terminal set S . For $\sigma \in \{0, 1\}^p$, let $J_\sigma = \{i \in [1, p] \mid \sigma[i] = 1\}$

- Alphabet (Label Set) $L = \{0, 1\}^p$. Labels encode the list of a_i 's from which a vertex is reachable.
- For $i \in [1, p]$, a unary predicate ψ_{a_i} encode the correct label for a_i and is defined as follows: $\psi_{a_i}(\sigma) = 0$ if $J_\sigma = Y_i$, otherwise $\psi_{a_i}(\sigma) = \infty$.
- For $j \in [1, q]$, predicate ψ_{b_j} that encodes the correct label for b_j . $\psi_{b_j}(\sigma) = 0$ if $J_\sigma = Z_j$, otherwise $\psi_{b_j}(\sigma) = \infty$.
- A binary predicate \mathcal{C} that encodes whether a directed edge is cut or not. If $\sigma_1 \leq \sigma_2$ $\mathcal{C}(\sigma_1, \sigma_2) = 0$, otherwise $\mathcal{C}(\sigma_1, \sigma_2) = 1$.
- A binary predicate NAE_2 that encodes whether an undirected edge is cut or not. If $\sigma_1 = \sigma_2$ $NAE_2(\sigma_1, \sigma_2) = 0$, otherwise $NAE_2(\sigma_1, \sigma_2) = 1$.

Thus $\beta_H = \{\mathcal{C}, NAE_2\} \cup \{\psi_{a_i} \mid i \in [1, p]\} \cup \{\psi_{b_j} \mid j \in [1, q]\}$. MIN- β_H -CSP has label set L , predicate set β_H and arity 2.

The main technical theorem we prove is the following. We remark that when we refer to DIR-MULC-H we are referring to the problem where the supply graph satisfies the assumptions I, II, III that we outlined previously.

THEOREM 3.3. *Let H be a directed bipartite graph. There is a polynomial time reduction that given a DIR-MULC-H instance $I_M = (G = (V_G, E_G, w_G : E_G \rightarrow R^+), H = (S \cup T, E_H))$, outputs a MIN- β_H -CSP instance $I_C = (V_C, T_C, \psi_{T_C} : T_C \rightarrow \beta_H, w_{T_C} : T_C \rightarrow R^+)$ such that the following holds: given a solution \mathbf{z} of the Label LP for I_M , we can construct a solution \mathbf{z} of Basic LP for I_C with cost at most that of (\mathbf{x}, \mathbf{z}) and vice versa. More over, if (\mathbf{x}, \mathbf{z}) is an integral solution, then \mathbf{z} is also an integral solution and vice versa. A similar reduction exists from MIN- β_H -CSP to DIR-MULC-H.*

With the preceding theorem in place we can formally prove Theorem 1.2

Proof of Theorem 1.2: Let I_M be some fixed instance of DIR-MULC-H with flow-cut gap α_H . From

Theorem 3.1 the integrality gap of LABEL-LP on I_m is also α_H . Let I_C be the MIN- β_H -CSP instance obtained via the reduction guaranteed by Theorem 3.3. I_M and I_C have the same integral cost. Fractional cost of LABEL-LP for I_M and BASIC-LP for I_C are also the same. Therefore the integrality gap of BASIC-LP on I_C is also α_H . Via Theorem 3.2, assuming UGC, MIN- β_H -CSP is hard to approximate within a factor of $\alpha_H - \varepsilon$ for any fixed $\varepsilon > 0$.

Theorem 3.3 (the second part) implies that MIN- β_H -CSP reduces to DIR-MULC-H in an approximation preserving fashion. Thus, DIR-MULC-H is at least as hard to approximate as MIN- β_H -CSP which implies that assuming UGC, the hardness of DIR-MULC-H is at least $\alpha_H - \varepsilon$. \square

BASIC-LP and LABEL-LP are almost identical except for the fact that LABEL-LP is defined with label set $\{0, 1\}^k$ where $k = p + q$ is the total number of terminals whereas BASIC-LP is defined with label set $\{0, 1\}^p$. However, since b_i 's do not have any outgoing edge, reachability from b_i is trivial. The formal proof of equivalence is long and somewhat tedious. We need to consider a reduction from MIN- β_H -CSP to DIR-MULC-H and vice-versa. In each direction we need to establish the equivalence of the cost of LABEL-LP and BASIC-LP for both integral and fractional settings. We will briefly sketch the reduction here. Full proofs can be found in Section B.

Reduction from Min- β_H -CSP to Dir-MulC-H:

Given a MIN- β_H -CSP instance I_C , equivalent DIR-MULC-H instance I_M is constructed as follows: (i) Vertex set of I_M is same as that of I_C (ii) For $i \in [1, p]$, name one of the vertex $v \in V_C$ with constraint $\psi_{a_i}(v)$ as vertex a_i and for rest of the vertices $u \in V_C$ with $\psi_{a_i}(u)$, add an undirected infinite weight edge between a_i and u (iii) For constraint $\mathcal{C}(u, v)$, add a directed edge $e_t = (u, v)$ and for constraint $\text{NAE}_2(u, v)$, add an undirected edge $e_t = uv$ (iv) Add edges among a_i 's and b_j 's so as to satisfy Assumption II and III. Next, we show how to convert a solution for one LP to a solution to other LP while preserving cost.

From Label-LP to Basic-LP: Let (\mathbf{x}, \mathbf{z}) be a feasible solution to LABEL-LP for I_M . Then, a feasible solution \mathbf{z}' to BASIC-LP for I_C is simply a projection of \mathbf{z} from label space $\{0, 1\}^{p+q}$ to label space $\{0, 1\}^p$. Formally, \mathbf{z}' is defined as follows (i) For $\sigma \in \{0, 1\}^p$, $z'_{v, \sigma} = \sum_{\sigma' \in \{0, 1\}^q} z_{v, \sigma \cdot \sigma'}$ (ii) For $\sigma_1, \sigma_2 \in \{0, 1\}^p$, $z'_{t, \sigma_1 \sigma_2} = \sum_{\sigma', \sigma'' \in \{0, 1\}^q} z_{e_t, \sigma_1 \cdot \sigma' \sigma_2 \cdot \sigma''}$. We can argue that $\sum_{\sigma_1, \sigma_2 \in \{0, 1\}^p: \sigma_1 \not\leq \sigma_2} z'_{t, \sigma_1 \sigma_2} \leq \sum_{\sigma_3, \sigma_4 \in \{0, 1\}^{p+q}: \sigma_3 \not\leq \sigma_4} z_{e_t, \sigma_3 \sigma_4} = x_e$. Hence, cost of solution \mathbf{z}' is at most the cost of solution (\mathbf{x}, \mathbf{z}) .

From Basic-LP to Label-LP: Let \mathbf{z} be a feasible solution to BASIC-LP for I_C . Let $\sigma_0 = 1^q$, then a feasible solution $(\mathbf{x}', \mathbf{z}')$ to LABEL-LP can be defined as an extension of \mathbf{z} along σ_0 . Formally, \mathbf{z}' is defined as follows: For $\sigma \in \{0, 1\}^p$, $\sigma' \in \{0, 1\}^q$, $v \in V_C$, $z'_{v, \sigma \cdot \sigma'} = z_{v, \sigma}$ if $\sigma' = \sigma_0$ and 0 otherwise. Similarly, for $\sigma_1, \sigma_2 \in \{0, 1\}^p$, $\sigma', \sigma'' \in \{0, 1\}^q$, $z'_{t, \sigma_1 \cdot \sigma' \sigma_2 \cdot \sigma''} = z_{e_t, \sigma_1 \sigma_2}$ if $\sigma' = \sigma'' = \sigma_0$ and 0 otherwise. We prove that $x'_e = \sum_{\sigma_1, \sigma_2 \in \{0, 1\}^{p+q}: \sigma_1 \not\leq \sigma_2} z'_{t, \sigma_1 \sigma_2} = \sum_{\sigma_3, \sigma_4 \in \{0, 1\}^{p+q}: \sigma_3 \not\leq \sigma_4} z_{e_t, \sigma_3 \sigma_4}$. Hence, cost of solution $(\mathbf{x}', \mathbf{z}')$ is equal to cost of solution \mathbf{z} .

3.4 Hardness for Non-bipartite Demand graphs

Here we prove Theorem 1.3 on the hardness of approximation of DIR-MULC-H when H is fixed and may not be bipartite. Let γ_H denote the hardness of approximation for DIR-MULC-H. Recall that α_H is the worst-case flow-cut gap for DIR-MULC-H.

Let the demand graph be H with 2^p vertices, $V_H = \{s_\sigma \mid \sigma \in \{0, 1\}^p\}$. If number of vertices is not a power of two, then we can add dummy isolated vertices without changing the problem. We find $r = 2^p$ subgraphs H_1, \dots, H_r such that $H = H_1 \cup \dots \cup H_r$ and

- Each H_i is a directed bipartite graph.
- $\alpha_H \leq \sum_{i=1}^r \alpha_{H_i}$.
- For $1 \leq i \leq r$, there is an approximation preserving reduction from DIR-MULC- H_i to DIR-MULC-H. Hence, $\gamma_H \geq \gamma_{H_i}$.

Since, H_i is bipartite, Theorem 1.2 implies, under UGC, that $\gamma_{H_i} \geq \alpha_{H_i} - \varepsilon$. Since, $\gamma_H \geq \gamma_{H_i}$ for all $i \in [1, r]$, we have $\gamma_H \geq \frac{1}{r} \sum_{i=1}^r \gamma_{H_i}$. Therefore,

$$\gamma_H \geq \frac{1}{r} \sum_{i=1}^r \gamma_{H_i} \geq \frac{1}{r} \sum_{i=1}^r (\alpha_{H_i} - \varepsilon) \geq \frac{1}{r} \alpha_H - \varepsilon.$$

Since $r = 2^{\lceil \log k \rceil}$ where $k = |V_H|$, we obtain the proof of Theorem 1.3.

Next, we show how to construct H_i which satisfy the properties above. For each number $j \in [1, p]$, define $A_j = \{s_\sigma \mid \sigma \in \{0, 1\}^p, \sigma(j) = 0\}$, $B_j = \{s_\sigma \mid \sigma \in \{0, 1\}^p, \sigma(j) = 1\}$. Let H_{2j-1} be the subgraph of H with vertex set V_H and edge set containing edges of H with head in B_j and tail in A_j . H_{2j} be the subgraph of H with vertex set V_H and edge set containing edges of H with head in A_j and tail in B_j .

$$\begin{aligned} V_{H_{2j-1}} &= V_{H_{2j}} = V_H \\ E_{H_{2j-1}} &= \{(s_{\sigma_1}, s_{\sigma_2}) \in E_H \mid s_{\sigma_1} \in A_j, s_{\sigma_2} \in B_j\} \\ E_{H_{2j}} &= \{(s_{\sigma_1}, s_{\sigma_2}) \in E_H \mid s_{\sigma_1} \in B_j, s_{\sigma_2} \in A_j\} \end{aligned}$$

By construction, it is clear that H_{2j-1}, H_{2j} are bipartite.

LEMMA 3.3. H_i as defined above satisfy the following properties:

- $E_H = \cup_{i=1}^r E_{H_i}$.
- $\alpha_H \leq \sum_{i=1}^r \alpha_{H_i}$.
- For $i \in [1, r]$, $\gamma_H \geq \gamma_{H_i}$.

Proof. • Let $e = (s_{\sigma_1}, s_{\sigma_2}) \in E_H$. Since, there are no self-loops in H , there exists $j \in [1, p]$ such that either $\sigma_1[j] = 0, \sigma_2[j] = 1$ or $\sigma_1[j] = 1, \sigma_2[j] = 0$. In the first case, $e \in E_{H_{2j-1}}$ and in the second case $e \in E_{H_{2j}}$.

- Given a DIR-MULC-H instance (G, H) , idea is to solve (G, H_i) for $i \in [1, p]$. Let $I = (G, H)$ be a DIR-MULC-H instance. Let \mathbf{x} be the optimal solution to DISTANCE-LP on I . Let $I_i = (G, H_i)$ be the instance with the same supply graph G but demand graph H_i . It is easy to see that \mathbf{x} is a feasible fractional solution to I_i since H_i is a subgraph of H . Since the worst-case integrality gap for DIR-MULC- H_i is α_{H_i} , there is a set $E'_i \subseteq E_G$ such that $w(E'_i) \leq \alpha_{H_i} w(\mathbf{x})$ and $G - E'_i$ disconnects all demand pairs in H_i . Clearly $\cup_i E'_i$ is a feasible integral solution to (G, H) since $H = \cup_i H_i$. The cost of $\cup_i E'_i$ is at most $(\sum_i \alpha_{H_i}) w(\mathbf{x})$. Since (G, H) was an arbitrary instance of DIR-MULC-H, this proves that $\alpha_H \leq \sum_i \alpha_{H_i}$.
- We prove that there is an approximation preserving reduction from DIR-MULC- H_i to DIR-MULC-H which in turn proves that $\gamma_H \geq \gamma_{H_i}$. Assume that $i = 2j - 1$ (case when $i = 2j$ is similar). Let (G, H_i) be a DIR-MULC- H_i instance. G' is defined as follows:

- $V_{G'} = V_G \cup A'_j \cup B'_j$ where $A'_j = \{s'_\sigma \mid s_\sigma \in A_j\}, B'_j = \{s'_\sigma \mid s_\sigma \in B_j\}$.
- G' contains all the edges of G and an infinite edge from s'_σ to s_σ for every $s_\sigma \in A_j$ and infinite weight edge from s_σ to s'_σ for every $s_\sigma \in B_j$.

Let H' be a demand graph with vertex s_σ of H renamed as s'_σ . Then, (G', H') is a DIR-MULC-H instance. Note that for $s_\sigma \in A_j$, s'_σ in G' has no incoming edge and for $s_\sigma \in B_j$, s'_σ in G' has no outgoing edge. Hence, for DIR-MULC instance (G', H') , we only need to separate $(s'_{\sigma_1}, s'_{\sigma_2})$ if $s_{\sigma_1} \in A_j, s_{\sigma_2} \in B_j$. Hence, DIR-MULC instances (G, H_i) and (G', H') are equivalent. \square

4 Approximating Dir-MulC

In this section, we prove Theorem 1.4 which improves the approximation ratio for DIR-MULC with restricted

class of demand graphs. Recall that a directed demand graph $H = (V, F)$ contains an induced k -matching extension if there are two subsets of V , $S = \{s_1, \dots, s_k\}$ and $T = \{t_1, \dots, t_k\}$ satisfying the following properties: (i) for $1 \leq i \leq k$, $(s_i, t_i) \in F$ and (ii) for $i > j$, $(s_i, t_j) \notin F$.

THEOREM 1.4. Consider DIR-MULC-H where H does not contain an induced k -matching-extension. Then the flow-cut gap is at most $k - 1$ and there is a polynomial-time rounding algorithm that achieves this upper bound.

Let $G = (V, E)$ and $H = (V, F)$ be the supply and demand graph for a DIR-MULC instance. We provide a generic randomized rounding algorithm that given a fractional solution \mathbf{x} to LP 1 for an instance (G, H) of DIR-MULC returns a feasible solution; the rounding does not depend on H . We can prove that the returned solution is a $(k - 1)$ -approximation with respect to the fractional solution \mathbf{x} or show that H contains an induced k -matching extension. This algorithm is inspired by our recent rounding scheme for DIR-MULTIWAY-CUT [7].

Let \mathbf{x} be a feasible solution to LP 1. For $u, v \in V$, define $d(u, v)$ to be the shortest path length in G from vertex u to vertex v using lengths x_e . We also define another parameter $d_1(u, v)$ for each pair of vertices $u, v \in V$. $d_1(u, v)$ is the minimum non-negative number such that if we add an edge uv in G with $x_{uv} = d_1(u, v)$ then u is still separated from all the vertices it has to be separated from. Formally, for $u, v \in V$, $d_1(u, v) := \max(0, 1 - \min_{v' \in V, (u, v') \in F} d(v, v'))$. If for some vertex u , there is no demand edge leaving u in F then we define $d_1(u, v) = 0$ for all $v \in V$. The following properties of d_1 are easy to verify.

LEMMA 4.1. $d_1(u, v)$ satisfies the following properties:

- $\forall u \in V, d_1(u, u) = 0$
- $\forall (u, v) \in F, v' \in V, d_1(u, v') + d(v', v) \geq 1$. Hence, $\forall (u, v) \in F, d_1(u, v) \geq 1$.
- If $d_1(u, v') \neq 0$, then there exists $(u, v) \in F$ such that $d_1(u, v') + d(v', v) = 1$
- $\forall u \in V, (a, b) \in E, d_1(u, b) - d_1(u, a) \leq x_{ab}$

Algorithm is a simple ball cut rounding around all the vertices as per $d_1(u, v)$. We pick a number $\theta \in (0, 1)$ uniformly at random. For all $u \in V$, we consider θ radius ball around u for all $u \in V$; $B_u = \{v \in V \mid d_1(u, v) \leq \theta\}$. And then cut all the edges leaving the set B_u ; $\delta^+(B_u) = \{(v, v') \in E_G \mid v \in B_u, v' \notin B_u\}$. Note that it is crucial that the same θ is used for all u .

Proving that E' is a feasible solution is easy. However, to bound the expected cost of the solution, we need the following lemma which shows that for any vertex v ,

Algorithm 1 Rounding for DIR-MULC

- 1: Given a feasible solution \mathbf{x} to LP 1
 - 2: For all $u, v \in V$, compute $d(u, v) =$ shortest path length from u to v according to lengths x_e
 - 3: For all $u, v \in V$, compute $d_1(u, v) = \max(0, 1 - \min_{v' \in V, uv' \in F} d(v, v'))$
 - 4: Pick $\theta \in (0, 1)$ uniformly at random
 - 5: $B_u = \{v \in V \mid d_1(u, v) \leq \theta\}$
 - 6: $E' = \cup_{u \in V} \delta^+(B_u)$
 - 7: Return E'
-

number of u_i with different non-zero values of $d_1(u_i, v)$ is at most $k - 1$.

LEMMA 4.2. *If for some $v \in V$ there exists u_1, \dots, u_k such that $0 \neq d_1(u_i, v) \neq d_1(u_j, v)$ for all $i \neq j$, then the demand graph H contains an induced k -matching extension.*

Proof. Rename the vertices u_1, \dots, u_k such that $d_1(u_1, v) > \dots > d_1(u_k, v) > 0$. By Lemma 4.1, there exists v'_1, \dots, v'_k such that $u_i v'_i \in F$ and $d_1(u_i, v) + d(v, v'_i) = 1$. Consider the subgraph of H induced by the vertices $s_1, \dots, s_k, t_1, \dots, t_k$ where $s_i = u_i, t_i = v'_i$. Edge $(s_i, t_i) \in F$ as $(u_i, v'_i) \in F$. By construction s_1, \dots, s_k are distinct. We also argue that t_1, \dots, t_k are distinct. Suppose $t_i = t_j$, that is $v'_j = v'_i$ for $i < j$. Then, $d_1(u_j, v) + d(v, v'_j) < d_1(u_i, v) + d(v, v'_j) = d_1(u_i, v) + d(v, v'_i) = 1$. Since $u_j v'_j \in F$, by Lemma 4.1 $d_1(u_j, v) + d(v, v'_j) \geq 1$ which contradicts the inequality above.

For $i > j$, $d_1(s_i, v) + d(v, t_j) = d_1(u_i, v) + 1 - d_1(u_j, v) < 1$. By lemma 4.1, $(s_i, t_j) \notin F$. Thus we have shown that $(s_i, t_i) \in F$ for $i \in [1, k]$ and $(s_i, t_j) \notin F$ for $i \neq j$. Thus H contains an induced k -matching extension. \square

Proof of Theorem 1.4: We start by solving LP 1 and then perform the rounding scheme as per Algorithm 1. By Lemma 4.1, for all $(u, v) \in F$, $d_1(u, v) \geq 1$ and since $\theta < 1$, we have $u \in B_u, v \notin B_u$. We remove all the edges going out of the set B_u and hence, cut all the paths from u to v . As argued above, for all $uv \in E_H$, $u \in B_u, v \notin B_u$ and we cut the edges going out of B_u . Hence, there is no path from u to v in $G - E'$ and E' is a feasible DIR-MULC solution.

We claim that $\Pr[e \in E'] \leq (k-1)x_e$ for all $e \in E_G$. Once we have this property, by linearity of expectation, the expected cost of E' can be bounded by $(k-1)$ times the LP cost: $\mathbb{E}[\sum_{e \in E'} w_e] \leq (k-1) \sum_{e \in E_G} w_e x_e$.

Now we prove the preceding claim. Consider an edge $e = (a, b) \in E$. Edge $e \in E'$ only if for some $u \in V$,

$e \in \delta^+(B_u)$ and this holds only if $\theta \in [d_1(u, a), d_1(u, b))$. By Lemma 4.1, $d_1(u, b) \leq d_1(u, a) + x_e$. Hence, $e \in \delta^+(B_u)$, if $\theta \in [d_1(u, b) - x_e, d_1(u, b))$. Denote this interval by $I_u(e)$.

By Lemma 4.2, there are at most $k - 1$ distinct elements in the set $\{d_1(u, b) \mid u \in V\}$. This implies that there are at most $k - 1$ distinct intervals $I_u(e)$. In other words there exists $u_1, \dots, u_r, r \leq k - 1$ such that $\cup_{u \in V} I_u(e) = \cup_{i=1}^r I_{u_i}(e)$.

$$\begin{aligned} \Pr[(a, b) \in E'] &\leq \Pr[\theta \in \cup_{u \in V} I_u(e)] \\ &= \Pr[\theta \in \cup_{i=1}^r I_{u_i}(e)] \\ &\leq \sum_{i=1}^r \Pr[\theta \in I_{u_i}(e)] \\ &\leq r \cdot x_e \leq (k-1)x_e. \end{aligned}$$

Penultimate inequality follows from the fact that $I_{u_i}(e)$ has length x_e and θ is chosen uniformly at random from $[0, 1)$. \square

5 Concluding Remarks and Open Problems

Question 1 asks for the approximability and flow cut gap for UNDIR-SKEW-MULTICUT and DIR-SKEW-MULTICUT. In this paper, we proved that UNDIR-SKEW-MULTICUT admits a constant factor approximation. However, the question of determining the flow-cut gap is still open. Also, if this flow-cut gap turns out to be constant, we could ask the question about tK_2 -free demand graphs as well.

QUESTION 3. *Does there exist a function $f : N \rightarrow N$ such that the flow-cut gap for UNDIR-MULC with tK_2 -free demand graph is at most $f(t)$? If yes, what is the tightest possible bound on $f(t)$?*

For DIR-SKEW-MULTICUT, we proved that flow-cut gap is a lower bound on hardness of approximation. However, determining the exact flow-cut gap for DIR-SKEW-MULTICUT is still an open question. If this flow-cut gap turns out to be constant, we could also ask about flow-cut gap when demand graph does not contain matching of size k as an induced subgraph. This may also help answer Question 2 which asks for better approximation algorithms for such graphs.

Hardness results in our paper relies on UGC. Can we prove these results assuming $P \neq NP$; specifically when demand graph is a collection of k disjoint directed edges.

QUESTION 4. *Assuming $P \neq NP$, can we prove that DIR-MULC-H does not admit $k - \epsilon$ approximation algorithm if H is a collection of k disjoint edges?*

References

- [1] Ittai Abraham, Cyril Gavoille, Anupam Gupta, Ofer Neiman, and Kunal Talwar. Cops, robbers, and threatening skeletons: Padded decomposition for minor-free graphs. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 79–88. ACM, 2014.
- [2] Amit Agarwal, Noga Alon, and Moses S Charikar. Improved approximation for directed cut problems. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 671–680. ACM, 2007.
- [3] Egon Balas and Chang Sung Yu. On graphs with polynomially solvable maximum-weight clique problem. *Networks*, 19(2):247–253, 1989.
- [4] Shuchi Chawla, Robert Krauthgamer, Ravi Kumar, Yuval Rabani, and D Sivakumar. On the hardness of approximating multicut and sparsest-cut. *computational complexity*, 15(2):94–114, 2006.
- [5] Chandra Chekuri, Sudeep Kamath, Sreeram Kannan, and Pramod Viswanath. Delay-constrained unicast and the triangle-cast problem. In *Information Theory (ISIT), 2015 IEEE International Symposium on*, pages 804–808. IEEE, 2015.
- [6] Chandra Chekuri, Sanjeev Khanna, Joseph Naor, and Leonid Zosin. A linear programming formulation and approximation algorithms for the metric labeling problem. *SIAM Journal on Discrete Mathematics*, 18(3):608–625, 2004.
- [7] Chandra Chekuri and Vivek Madan. Simple and fast rounding algorithms for directed and node-weighted multiway cut. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10–12, 2016*, pages 797–807, 2016.
- [8] Jianer Chen, Liu Yang, Lu Songjian, O’sullivan Barry, and Igor Razgon. A fixed-parameter algorithm for the directed feedback vertex set problem. *Journal of the ACM (JACM)*, 55(5), 2008.
- [9] Julia Chuzhoy and Sanjeev Khanna. Polynomial flow-cut gaps and hardness of directed cut problems. *Journal of the ACM (JACM)*, 56(2):6, 2009.
- [10] Alina Ene, Jan Vondrák, and Yi Wu. Local distribution and the symmetry gap: Approximability of multiway partitioning problems. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 306–325. SIAM, 2013.
- [11] Robert F. Erbacher, Trent Jaeger, Nirupama Talele, and Jason Teutsch. Directed multicut with linearly ordered terminals. *CoRR*, abs/1407.7498, 2014.
- [12] Jittat Fakcharoenphol and Kunal Talwar. An improved decomposition theorem for graphs excluding a fixed minor. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 36–46. Springer, 2003.
- [13] Naveen Garg, Vijay V Vazirani, and Mihalis Yannakakis. Approximate max-flow min-(multi) cut theorems and their applications. *SIAM Journal on Computing*, 25(2):235–251, 1996.
- [14] Philip Klein, Serge A Plotkin, and Satish Rao. Excluded minors, network decomposition, and multicommodity flow. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 682–690. ACM, 1993.
- [15] Jon M. Kleinberg and Éva Tardos. Approximation algorithms for classification problems with pairwise relationships: Metric labeling and Markov random fields. *Journal of the ACM (JACM)*, 49(5):616–639, 2002. Preliminary version in FOCS 1999.
- [16] Euiwoong Lee. Improved hardness for cut, interdiction, and firefighter problems. *CoRR*, abs/1607.05133, 2016.
- [17] Rajsekar Manokaran, Joseph Seffi Naor, Prasad Raghavendra, and Roy Schwartz. Sdp gaps and ugc hardness for multiway cut, 0-extension, and metric labeling. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 11–20. ACM, 2008.
- [18] Dániel Marx and Paul Wollan. An exact characterization of tractable demand patterns for maximum disjoint path problems. *CoRR*, abs/1411.0871, 2014.
- [19] J Naor and L Zosin. A 2-approximation algorithm for the directed multiway cut problem. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 548–548. IEEE Computer Society, 1997.
- [20] Michael Saks, Alex Samorodnitsky, and Leonid Zosin. A lower bound on the integrality gap for minimum multicut in directed networks. *Combinatorica*, 24(3):525–530, 2004.
- [21] Ankit Sharma and Jan Vondrák. Multiway cut, pairwise realizable distributions, and descending thresholds. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 724–733. ACM, 2014.
- [22] David Steurer and Nisheeth Vishnoi. Connections between multi-cut and unique games. *Technical Report TR09-125*, 2009.

A Proof of Theorem 3.1

From Label-LP to Distance-LP: Let (\mathbf{x}, \mathbf{z}) be a feasible solution to LABEL-LP for the given instance of G, H . Consider a solution \mathbf{x}' to DISTANCE-LP where we set $x'_e = x_e$. We claim that \mathbf{x}' is a feasible solution to DISTANCE-LP for G, H . That is, for $(s_i, s_j) \in E_H$, and a path p from s_i to s_j , we have $\sum_{e \in p} x'_e \geq 1$.

LEMMA A.1. *For any edge $e = (u, v) \in E_G$ and $i \in \{1, \dots, k\}$, $x_e \geq \sum_{\sigma \in L, \sigma[i]=1} z_{u, \sigma} - \sum_{\sigma \in L, \sigma[i]=1} z_{v, \sigma}$.*

Proof. Recall the interpretation of variables $z_{e, \sigma_1 \sigma_2}$ as flow from set $\Gamma_u = \{u_\sigma \mid \sigma \in L\}$ to $\Gamma_v = \{v_\sigma \mid \sigma \in L\}$. Consider the following partition of Γ_u into $\Gamma_u^1 = \{u_\sigma \mid \sigma \in L, \sigma[i] = 1\}$ and $\Gamma_u^2 = \{u_\sigma \mid \sigma \in L, \sigma[i] = 0\}$. Similarly, consider the partition of Γ_v

into Γ_v^1 and Γ_v^2 . Amount of flow out of Γ_u^1 is equal to $\sum_{\sigma \in L, \sigma[i]=1} z_{u,\sigma}$ and amount of flow coming into Γ_v^1 is equal to $\sum_{\sigma \in L, \sigma[i]=1} z_{v,\sigma}$. Amount of flow from Γ_u^1 to Γ_v^1 is at most $\sum_{\sigma \in L, \sigma[i]=1} z_{v,\sigma}$. Hence, flow from Γ_u^1 to Γ_v^2 is at least $\sum_{\sigma \in L, \sigma[i]=1} z_{u,\sigma} - \sum_{\sigma \in L, \sigma[i]=1} z_{v,\sigma}$. For $u_{\sigma_1} \in \Gamma_u^1, v_{\sigma_2} \in \Gamma_v^2$, we have $\sigma_1 \not\leq \sigma_2$ and hence,

$$\begin{aligned} x'_e = x_e &= \sum_{\sigma_1, \sigma_2 \in L: \sigma_1 \not\leq \sigma_2} z_{e, \sigma_1 \sigma_2} \\ &\geq \sum_{\sigma \in L: \sigma[i]=1} z_{u, \sigma} - \sum_{\sigma \in L: \sigma[i]=1} z_{v, \sigma}. \end{aligned}$$

□

Let $(s_i, s_j) \in E_H$. We prove that for any path p from s_i to s_j in G has $\sum_{e \in p} x'_e \geq 1$. Let the path p be $s_i, a_1, \dots, a_\ell, s_j$. Then, by Lemma A.1

$$\begin{aligned} x_{(s_i, a_1)} + \sum_{t=1}^{\ell-1} x_{(a_t, a_{t+1})} + x_{(a_\ell, s_j)} &\geq \\ \sum_{\sigma \in L: \sigma[i]=1} \left[\left(z_{s_i, \sigma} - z_{a_1, \sigma} \right) + \sum_{t=1}^{\ell-1} \left(z_{a_t, \sigma} - z_{a_{t+1}, \sigma} \right) \right. \\ &\quad \left. + \left(z_{a_\ell, \sigma} - z_{s_j, \sigma} \right) \right] \\ &= \sum_{\sigma \in L: \sigma[i]=1} (z_{s_i, \sigma} - z_{s_j, \sigma}) \end{aligned}$$

LABEL-LP ensures that $z_{s_i, \sigma} = 0$ if $\sigma[i] = 0$ and $z_{s_j, \sigma} = 0$ if $\sigma[i] = 1$. Hence, $\sum_{\sigma \in L: \sigma[i]=1} z_{s_i, \sigma} = 1$ and $\sum_{\sigma \in L: \sigma[i]=1} z_{s_j, \sigma} = 0$. Hence the right hand side in the preceding inequality is 1.

From Distance-LP to Label-LP: Suppose \mathbf{x} is a feasible solution to DISTANCE-LP for the given instance G, H . We construct a solution $(\mathbf{x}', \mathbf{z})$ for LABEL-LP such that $x'_e \leq x_e$ for all $e \in E_G$. The edge lengths given by \mathbf{x} induce shortest path distances in G and we use $d(u, v)$ to denote this distance from u to v . By adding dummy edges with zero cost as needed we can assume that $d(u, v) \leq 1$ for each vertex pair (u, v) . With this assumption in place we have that for any edge $e = (u, v)$ and any terminal s_i , $d(s_i, v) \leq d(s_i, u) + x_e$; hence $x_e \geq \max_{1 \leq i \leq k} (d(s_i, v) - d(s_i, u))$. We will in fact prove that $x'_e \leq \max_{1 \leq i \leq k} (d(s_i, v) - d(s_i, u))$.

We start by describing how to assign values to the variables $z_{v, \sigma}$. Recall that these induce values to the other variables if one is interested in a minimum cost solution. Let $d(u, v)$ denote the shortest distance from u to v in G as per lengths x_e .

For a vertex u , consider the permutation $\pi^u : \{1, \dots, k\} \rightarrow \{1, \dots, k\}$ such that $d(s_{\pi^u(1)}, u) \leq \dots \leq$

$d(s_{\pi^u(k)}, u)$. In other words π^u is an ordering of the terminals based on distance to u (breaking ties arbitrarily). Define $\sigma_0^u, \dots, \sigma_k^u$ as follows:

$$\sigma_i^u[j] = \begin{cases} 1 & j \in \{\pi^u(1), \dots, \pi^u(i)\} \\ 0 & j \notin \{\pi^u(1), \dots, \pi^u(i)\} \end{cases}$$

In the assignment above it is useful to interpret σ_i^u as a set of indices of the terminals. Hence σ_0^u corresponds to \emptyset and σ_i^u to $\{\pi^u(1), \dots, \pi^u(i)\}$. Thus, these sets form a chain with.

The assignment of values to the variables $z_{u, \sigma}$, $\sigma \in L$ is done as follows:

$$z_{u, \sigma} = \begin{cases} d(s_{\pi^u(1)}, u) & \sigma = \sigma_0^u \\ d(s_{\pi^u(i+1)}, u) - d(s_{\pi^u(i)}, u) & \sigma = \sigma_i^u, i \in [1, k-1] \\ 1 - d(s_{\pi^u(k)}, u) & \sigma = \sigma_k^u \\ 0 & \text{otherwise} \end{cases}$$

LEMMA A.2. $z_{u, \sigma}$ as defined above satisfy the following properties:

- $\forall u \in V_G, \sigma \in L, z_{u, \sigma} \geq 0$.
- $\forall u \in V_G, \sum_{\sigma \in L} z_{u, \sigma} = 1$.
- For $A \subseteq \{1, \dots, k\}$, define $\sigma_A \in L$ as: $\sigma_A[i] = 1$ for $i \in A$ and 0 otherwise. Then,

$$\sum_{\sigma \geq \sigma_A} z_{u, \sigma} = 1 - \max_{i \in A} d(s_i, u)$$

- Terminals are labelled correctly. That is, for each s_j and $\sigma \in L$, $z_{s_j, \sigma} = 0$ if $\sigma[j] = 0$.
- If $(s_i, s_j) \in E_H$, then $z_{s_j, \sigma} = 0$ for $\sigma \in L$ such that $\sigma[i] = 1$.

Proof. For $u \in V_G$, consider $\sigma_0^u, \sigma_1^u, \dots, \sigma_k^u$ as defined above.

- $z_{u, \sigma} \geq 0$ is true by definition.
- By definition, $z_{u, \sigma} = 0$ if $\sigma \notin \{\sigma_0^u, \dots, \sigma_k^u\}$. Hence,

$$\begin{aligned} \sum_{\sigma \in L} z_{u, \sigma} &= \sum_{i=0}^k z_{u, \sigma_i^u} \\ &= d(s_{\pi^u(1)}, u) + 1 - d(s_{\pi^u(k)}, u) \\ &\quad + \sum_{i=1}^{k-1} d(s_{\pi^u(i+1)}, u) - d(s_{\pi^u(i)}, u) \\ &= 1 \end{aligned}$$

- Let $j = \arg \max_{i: \pi^u(i) \in A} d(s_{\pi^u(i)}, u)$. Then,

$\sigma_j^u, \dots, \sigma_k^u \geq \sigma_A$ and $\sigma_0^u, \dots, \sigma_{j-1}^u \not\geq \sigma_A$. Hence,

$$\begin{aligned} \sum_{\sigma \geq \sigma_A} z_{u,\sigma} &= \sum_{i=j}^k z_{u,\sigma_i^u} \\ &= \sum_{i=j}^{k-1} d(s_{\pi^u(i+1)}, u) - d(s_{\pi^u(i)}, u) \\ &\quad + 1 - d(s_{\pi^u(k)}, u) \\ &= 1 - d(s_{\pi^u(j)}, u) \\ &= 1 - \max_{i:\pi^u(i) \in A} d(s_{\pi^u(i)}, u) \\ &= 1 - \max_{i \in A} d(s_i, u) \end{aligned}$$

- By definition of distance, $d(s_j, s_j) = 0$. Consider $A = \{j\}$. Applying the result from previous part, we get $\sum_{\sigma \geq \sigma_A} z_{s_j, \sigma} = 1 - 0 = 1$. Hence, $z_{s_j, \sigma} = 0$ if $\sigma \not\geq \sigma_A$. Equivalently speaking, $z_{s_j, \sigma} = 0$ if $\sigma[j] = 0$.
- Let $(s_i, s_j) \in E_H$. Then, for the solution \mathbf{x} to be feasible, we must have $d(s_i, s_j) = 1$. Consider $A = \{i\}$. Then, using result from previous part, we get $\sum_{\sigma \geq \sigma_A} z_{s_j, \sigma} = 1 - 1 = 0$. Hence, $z_{s_j, \sigma} = 0$ if $\sigma \geq \sigma_A$. Equivalently speaking, $z_{s_j, \sigma} = 0$ if $\sigma[i] = 1$.

□

Consider an edge $e = (u, v)$. Recall that once the distributions of \bar{z}_u and \bar{z}_v are fixed, x'_e is simply the min-cost flow between these two distributions in the digraph B_{uv} with costs given by c . Our goal is to show that this cost is at most $\max\{0, \max_i(d(s_i, v) - d(s_i, u))\}$. Suppose we define a partial flow between \bar{z}_u and \bar{z}_v on zero-cost edges such that the total amount of this flow is γ where $\gamma \in [0, 1]$. Then it is easy to see that we can complete this flow to achieve a cost of $(1 - \gamma)$. This is because the graph is a complete bipartite graph and costs are either 0 or 1 and \bar{z}_u and \bar{z}_v are distributions that have a total of exactly one unit of mass on each side.

Next, we define a partial flow of zero cost between \bar{z}_u and \bar{z}_v by setting some variables $z_{e, \sigma_1 \sigma_2}$ in a greedy fashion as follows. Initially all flow values are zero. We consider the vertex $u_{\sigma_i^u}$ with supply z_{u, σ_i^u} in order from $i = 0$ to k . Our goal is to send as much flow as possible from this vertex on zero-cost edges to demand vertices $v_{\sigma_j^v}$ which requires that $\sigma_i^u \leq \sigma_j^v$. We maintain the invariants that we do not exceed supply or demand in this process. While trying to send flow out of $u_{\sigma_i^u}$ we again use a greedy process; if there are $j < j'$ such that σ_j^v and $\sigma_{j'}^v$ are both eligible to receive flow on zero-cost edges and have capacity left, we use j first; recall that σ_j^v corresponds to a subset of $\sigma_{j'}^v$. Let $z_{e, \sigma_1 \sigma_2}$ be the partial flow created by the algorithm.

LEMMA A.3. *The total flow sent by the greedy algorithm described is at least $1 - \max\{0, \max_h(d(s_h, v) - d(s_h, u))\}$.*

Assuming the lemma we are done because the zero-cost flow is at least $1 - x_e$ and hence total cost of the flow is at most x_e , thus proving $x'_e \leq x_e$ as desired. We now prove the lemma.

Consider the greedy flow. Let ℓ be the maximum integer such that $v_{\sigma_\ell^v}$ is not saturated by the flow. If no such ℓ exists then the greedy algorithm has sent a total flow of one unit on zero-cost edges and hence $x'_e = 0$. Thus, we can assume ℓ exists. Moreover, in this case we can also assume that $\ell < k$ for if $\ell = k$ the greedy algorithm can send more flow since $\sigma_i^u \leq \sigma_k^v$ for all i . Let ℓ' be the maximum integer such that $\sigma_{\ell'}^u \leq \sigma_\ell^v$. Such an ℓ' exists since $\ell' = 0$ is a candidate (corresponding to the empty set). Moreover, $\ell' < k$ since $\sigma_k^u \not\leq \sigma_\ell^v$ since $\ell < k$. Let ℓ'' be the minimum integer such that $\sigma_{\ell''+1}^u \leq \sigma_{\ell'}^v$. ℓ'' exists because k is a candidate for it.

CLAIM 1. $\pi_{\ell'+1}^u = \pi_{\ell''}^v$.

Proof. By choice of ℓ, ℓ', ℓ'' we have $\sigma_{\ell'}^u \leq \sigma_\ell^v$ and $\sigma_{\ell'+1}^u \not\leq \sigma_\ell^v$ while $\sigma_{\ell'+1}^u \leq \sigma_{\ell'}^v$. Thus $\ell'' \geq \ell + 1$ and $\sigma_{\ell'}^u \leq \sigma_\ell^v \leq \sigma_{\ell'+1}^v$. Moreover, since ℓ'' is chosen to smallest, $\sigma_{\ell'+1}^u \not\leq \sigma_{\ell'+1}^v$. These facts imply the desired claim. □

We now claim several properties of the partial flow and justify them.

- $\forall i \in [0, \ell'], j \in [\ell + 1, k], z_{e, \sigma_i^u \sigma_j^v} = 0$. This follows from the fact that the greedy algorithm did not saturate z_{v, σ_ℓ^v} .
- $\forall i \in [\ell' + 1, k], j \in [0, \ell'' - 1], z_{e, \sigma_i^u \sigma_j^v} = 0$. From the definition of ℓ', ℓ'' , this is not a zero cost edge.
- $\forall i \in [0, \ell'], \sum_{j=0}^{\ell'} z_{e, \sigma_i^u \sigma_j^v} = z_{u, \sigma_i^u}$. From definition of ℓ' , for each $i \leq \ell'$, there is a zero-cost edge from $u_{\sigma_i^u}$ to $v_{\sigma_\ell^v}$. Since the greedy algorithm did not saturate $v_{\sigma_\ell^v}$, it means that $u_{\sigma_i^u}$ is saturated and sends flow only to $v_{\sigma_1^v}, \dots, v_{\sigma_\ell^v}$.
- $\forall j \in [\ell'', k], \sum_{i=\ell'+1}^k z_{e, \sigma_i^u \sigma_j^v} = z_{v, \sigma_j^v}$. By definition of ℓ , for $j \geq \ell + 1$ we have the property that $v_{\sigma_j^v}$ is saturated. As we argued above, for $i \in [\ell' + 1, k], j \in [0, \ell'' - 1]$ we have $z_{e, \sigma_i^u \sigma_j^v} = 0$. Hence, for $j \geq \ell'' \geq \ell + 1$, we have $\sum_{i=\ell'+1}^k z_{e, \sigma_i^u \sigma_j^v} = \sum_{i=0}^k z_{e, \sigma_i^u \sigma_j^v} = z_{v, \sigma_j^v}$.

From the preceding claim we see that the total value of the partial flow can be summed up as

$$\sum_{\sigma_1, \sigma_2 \in L} z_{e, \sigma_1 \sigma_2} = \sum_{i=0}^{\ell'} z_{u, \sigma_i^u} + \sum_{j=\ell''+1}^k z_{v, \sigma_j^v}.$$

Moreover, by construction of \bar{z}_u and \bar{z}_v ,

$$\sum_{i=0}^{\ell'} z_{u, \sigma_i^u} = d(s_{\pi^u(\ell'+1)}, u)$$

and $\sum_{j=\ell''+1}^k z_{v, \sigma_j^v} = 1 - d(s_{\pi^v(\ell'')}, v).$

Letting $h = \pi_{\ell'+1}^u = \pi_{\ell''}^v$, we see that from the preceding equalities that the total flow routed on the zero-cost edges is

$$d(s_h, u) + 1 - d(s_h, v) = 1 - (d(s_h, v) - d(s_h, u)) \geq 1 - x_e.$$

This finishes the proof.

B Proof of Theorem 3.3

The first two lemmas help establish that we can safely assume that the supply graph satisfies the assumptions I, II, and III. We omit the proof of the first lemma which involves tedious reworking of some of the details on equivalence of LABEL-LP and DISTANCE-LP.

LEMMA B.1. *For any instance G, H of DIR-MULC-H where the supply graph has undirected edges, the optimum solution values for the formulations LABEL-LP and DISTANCE-LP are the same both in the fractional and integral settings.*

Assuming the preceding lemma, following lemma is easy to prove:

LEMMA B.2. *For bipartite H , DIR-MULC-H with a general supply graph and DIR-MULC-H restricted to supply graphs satisfying Assumptions I, II and III are equivalent in terms of approximability and in terms of the integrality gap of DISTANCE-LP (equal to integrality gap of LABEL-LP).*

Proof. We sketch the proof. Undirected edges can be handled by the gadget shown in Fig 5. It is easy to see that given any instance of DIR-MULC-H with supply graph G and bipartite demand graph H we can first add dummy terminals to G and assume that each terminal a_i has only one outgoing infinite weight edge (to the original terminal) and each b_j has only one incoming infinite weight edge. With this in place adding edges to satisfy Assumptions II and III can be seen to not affect the integral or fractional solutions to DISTANCE-LP. \square

We will assume for simplicity that all weights (for edges and constraints) are either 1 or ∞ . Generic weights can be easily simulated by copies and the proofs make no essential use of weights other than that some are finite and others are infinite.

B.1 Reduction from Min- β_H -CSP to Dir-MulC-H Let the MIN- β -CSP instance be $I_C = (V_C, T_C, \psi_{T_C} : T_C \rightarrow \beta_H, w_{T_C} : T_C \rightarrow R^+)$. We refer to tuple $t = (u)$ with $\psi_{T_C}(t) = \psi_{a_i}$ as constraint $\psi_{a_i}(u)$, $t = (u)$, $\psi_{T_C}(u) = \psi_{b_j}$ as constraint $\psi_{b_j}(u)$, $t = (u, v)$, $\psi_{T_C}(t) = \mathcal{C}$ as constraint $\mathcal{C}(u, v)$ and $t = (u, v)$, $\psi_{T_C}(u) = NAE_2$ as constraint $NAE_2(u, v)$. We assume that for every $i \in [1, p]$, there is a constraint $\psi_{a_i}(u_i)$ for some vertex $u_i \in V_C$, and similarly for every $j \in [1, q]$ there is a constraint $\psi_{b_j}(v_j)$ for some vertex $v_j \in V_C$; moreover we will assume that $u_1, \dots, u_p, v_1, \dots, v_q$ are distinct vertices. One can ensure that this assumption holds by adding dummy vertices and dummy constraints with zero weight. We create an instance $I_M = (G = (V_G, E_G, w_G : E_G \rightarrow R^+), (S \cup T, E_H))$ of DIR-MULC-H as follows.

- $V_G = V_C$, the vertex remains the same. Pick vertices u_1, \dots, u_p and v_1, \dots, v_q that are all distinct such that for $1 \leq i \leq p$ there is a constraint $\psi_{a_i}(u_i)$ in I_C and for $1 \leq j \leq q$ there is a constraint $\psi_{b_j}(v_j)$ in I_C . This holds by our assumption. For $i \in [1, p]$ associate the terminal $a_i \in V_H$ with u_i and for $j \in [1, q]$ associate the terminal $b_j \in V_H$ with v_j .
- E_G and w_G are defined as follows:
 - For each constraint $\psi_{a_i}(u)$ in I_C where $u \neq a_i$ add an undirected edge $e_t = a_i u$ to E_G with $w_G(e_t) = \infty$.
 - For each constraint $\psi_{b_j}(v)$ in I_C where $v \neq b_j$ add an undirected edge $e_t = b_j v$ to E_G with $w_G(e_t) = \infty$.
 - For each constraint $\mathcal{C}(u, v)$ in I_C add a directed edge $e_t = (u, v)$ in G with $w_G(e_t)$ equal to the weight of the constraint in I_C .
 - For each constraint $NAE_2(u, v)$, add an undirected edge $e_t = uv$ with $w_G(e_t)$ equal to the weight of the constraint in I_C .
 - For each $i \in [1, p]$ and for each $i' \in Y_i$, add a directed edge $e = (a_{i'}, a_i)$ with $w_G(e) = \infty$.
 - For each $j \in [1, q]$ and each $i \in Z_j$, add a directed edge $e = (a_i, b_j)$ with $w_G(e) = \infty$.

We now prove the equivalence of I_C and I_M from the point of view solutions to BASIC-LP and LABEL-LP respectively.

Given two labels σ and σ' which can be interpreted as binary strings, we use the notation $\sigma \cdot \sigma'$ to denote the label obtained by concatenating σ and σ' .

From Label-LP to Basic-LP: Suppose (\mathbf{x}, \mathbf{z}) is a feasible solution to LABEL-LP for I_M . We construct a solution \mathbf{z}' to BASIC-LP for I_C in the following way. \mathbf{z}' is simply a projection of \mathbf{z} from label set $\{0, 1\}^{p+q}$ onto label set $\{0, 1\}^p$. Recall that in the instance I_M the terminals b_1, \dots, b_q do not have any outgoing edges.

Hence, in the solution (\mathbf{x}, \mathbf{z}) with label space $\{0, 1\}^{p+q}$, which encodes reachability from both the a_i s and the b_j 's the information on reachability from the b_j s does not play any essential role. We formalize this below.

- For $v \in V_C, \sigma \in \{0, 1\}^p, z'_{v,\sigma} = \sum_{\sigma' \in \{0,1\}^q} z_{v,\sigma\sigma'}$.
- For unary constraint $t = (v) \in T_C$ and $\sigma \in \{0, 1\}^p, z'_{t,\sigma} = z'_{v,\sigma}$.
- For binary constraint $t = (u, v) \in T_C$, for $\sigma_1, \sigma_2 \in \{0, 1\}^p$,

$$z'_{t,\sigma_1\sigma_2} = \sum_{\sigma' \in \{0,1\}^q} \sum_{\sigma'' \in \{0,1\}^q} z_{e_t,\sigma_1\sigma'\sigma_2\sigma''}$$

Note that if (\mathbf{x}, \mathbf{z}) is an integral solution then \mathbf{z}' as defined above is also an integral solution.

Feasibility of \mathbf{z}' for BASIC-LP is an “easy” consequence of the projection operation but we prove it formally.

LEMMA B.3. \mathbf{z}' as defined above is a feasible solution to BASIC-LP for instance I_C .

Proof. From the definition of \mathbf{z}' , for each vertex v ,

$$\sum_{\sigma \in \{0,1\}^p} z'_{v,\sigma} = \sum_{\sigma \in \{0,1\}^p} \sum_{\sigma' \in \{0,1\}^q} z_{v,\sigma\sigma'} = 1$$

which proves that one set of constraints holds.

Next, we prove that for $t \in T_C, v = t[i], \sigma \in L = \{0, 1\}^p$, the constraint $z'_{v,\sigma} - \sum_{\alpha \in L^{|t|}: \alpha[i]=\sigma} z'_{t,\alpha} = 0$ holds. We consider unary and binary predicates separately.

- For $t = (v)$ s.t. $v = t[i], \sigma \in L = \{0, 1\}^p$,

$$z'_{v,\sigma} - \sum_{\alpha \in L^{|t|}: \alpha[i]=\sigma} z'_{t,\alpha} = z'_{v,\sigma} - z'_{t,\sigma} = z'_{v,\sigma} - z'_{v,\sigma} = 0.$$

- For $t = (u, v) \in T_C, \sigma \in \{0, 1\}^p$

$$\begin{aligned} & z'_{v,\sigma} - \sum_{\sigma_1 \in \{0,1\}^p} z'_{t,\sigma_1\sigma} \\ &= z'_{v,\sigma} - \sum_{\sigma_1 \in \{0,1\}^p} \sum_{\sigma', \sigma'' \in \{0,1\}^q} z_{e_t,\sigma_1\sigma'\sigma''\sigma} \\ &= z'_{v,\sigma} - \sum_{\sigma' \in \{0,1\}^q} \sum_{\sigma_1 \in \{0,1\}^p, \sigma'' \in \{0,1\}^q} z_{e_t,\sigma_1\sigma'\sigma''\sigma} \\ &= z'_{v,\sigma} - \sum_{\sigma' \in \{0,1\}^q} z_{v,\sigma\sigma'} = 0. \end{aligned}$$

Similar argument holds for u as well. \square

LEMMA B.4. The cost of \mathbf{z}' is at most $\sum_{e \in E_G} w_e x_e$ which is the cost of (\mathbf{x}, \mathbf{z}) to I_M .

Before we prove Lemma B.4 we establish some properties satisfied by (\mathbf{x}, \mathbf{z}) .

LEMMA B.5. If the solution (\mathbf{x}, \mathbf{z}) to LABEL-LP has finite cost, then the following conditions hold:

- For directed edge $e = (u, v)$, and for $i \in [1, p]$ $x_e \geq \sum_{\sigma \in \{0,1\}^{p+q}: \sigma[i]=1} z_{u,\sigma} - \sum_{\sigma \in \{0,1\}^{p+q}: \sigma[i]=1} z_{v,\sigma}$. Hence, if edge e has infinite weight ($w_G(e) = \infty$), then $\sum_{\sigma \in \{0,1\}^{p+q}: \sigma[i]=1} z_{u,\sigma} \leq \sum_{\sigma \in \{0,1\}^{p+q}: \sigma[i]=1} z_{v,\sigma}$.
- For $i \in [1, p], \sigma \in \{0, 1\}^p, \sigma' \in \{0, 1\}^q$ s.t. $J_\sigma \neq Y_i$, we have $z_{a_i,\sigma\sigma'} = 0$. Hence, for $\sigma \in \{0, 1\}^p, z'_{a_i,\sigma} = 1$ if $J_\sigma = Y_i$ and 0 otherwise.
- For $j \in [1, q], \sigma \in \{0, 1\}^p, \sigma' \in \{0, 1\}^q$ s.t. $J_\sigma \neq Z_j$ we have $z_{b_j,\sigma\sigma'} = 0$. Hence, for $\sigma \in \{0, 1\}^p, z'_{b_j,\sigma} = 1$ if $J_\sigma = Z_j$ and 0 otherwise.
- For an undirected edge $e = uv \in E_G$ with $w_G(e) = \infty$, and $\sigma_1, \sigma_2 \in \{0, 1\}^{p+q}$, $z_{e,\sigma_1\sigma_2} = 0$ if $\sigma_1 \neq \sigma_2$. For $\sigma \in \{0, 1\}^{p+q}, z_{u,\sigma} = z_{v,\sigma}$ and for $\sigma_1 \in \{0, 1\}^p, z'_{u,\sigma_1} = z'_{v,\sigma_1}$. Hence, for $t = (u) \in T_C$ s.t. $\psi_{T_C}(t) = \psi_{a_i}, z'_{u,\sigma} = 1$ if $J_\sigma = Y_i$ and 0 otherwise.

Proof. If (\mathbf{x}, \mathbf{z}) has finite cost, then for an edge e with infinite weight ($w_G(e) = \infty$), we must have $x_e = 0$.

- Let $e = (u, v)$ be a directed edge, and $i \in [1, p]$

$$\begin{aligned} x_e &= \sum_{\sigma_1, \sigma_2 \in \{0,1\}^{p+q}: \sigma_1 \not\leq \sigma_2} z_{e,\sigma_1\sigma_2} \\ &\geq \sum_{\sigma_1, \sigma_2 \in \{0,1\}^{p+q}: \sigma_1[i]=1, \sigma_2[i]=0} z_{e,\sigma_1\sigma_2} \\ &= \sum_{\sigma_1 \in \{0,1\}^{p+q}: \sigma_1[i]=1} z_{u,\sigma_1} \\ &\quad - \sum_{\sigma_1, \sigma_2 \in \{0,1\}^{p+q}: \sigma_1[i]=1, \sigma_2[i]=1} z_{e,\sigma_1\sigma_2} \\ &\geq \sum_{\sigma_1 \in \{0,1\}^{p+q}: \sigma_1[i]=1} z_{u,\sigma_1} \\ &\quad - \sum_{\sigma_2 \in \{0,1\}^{p+q}: \sigma_2[i]=1} z_{v,\sigma_2} \\ &= \sum_{\sigma \in \{0,1\}^{p+q}: \sigma[i]=1} z_{u,\sigma} - \sum_{\sigma \in \{0,1\}^{p+q}: \sigma[i]=1} z_{v,\sigma} \end{aligned}$$

If edge e has infinite weight, then $x_e = 0$ and $\sum_{\sigma \in \{0,1\}^{p+q}: \sigma[i]=1} z_{u,\sigma} \leq \sum_{\sigma \in \{0,1\}^{p+q}: \sigma[i]=1} z_{v,\sigma}$

- We prove the following two statements which in turn imply that for $\sigma \in \{0, 1\}^p, \sigma' \in \{0, 1\}^q$, if $J_\sigma \neq Y_i$, then $z_{a_i,\sigma\sigma'} = 0$.

$$\forall j \in Y_i, \sum_{\sigma \in \{0,1\}^p \sigma' \in \{0,1\}^q: \sigma[j]=1} z_{a_i,\sigma\sigma'} = 1$$

$$\forall j \in [1, p] \setminus Y_i, \sum_{\sigma \in \{0,1\}^p \sigma' \in \{0,1\}^q: \sigma[j]=1} z_{a_i,\sigma\sigma'} = 0$$

Let $j \in Y_i$, then by construction of G , there exists an infinite weight edge from a_j to a_i . Using the result from previous part we get $\sum_{\sigma \in \{0,1\}^p \sigma' \in \{0,1\}^q: \sigma[j]=1} z_{a_i, \sigma \cdot \sigma'} \geq \sum_{\sigma \in \{0,1\}^p \sigma' \in \{0,1\}^q: \sigma[j]=1} z_{a_j, \sigma \cdot \sigma'}$. LABEL-LP enforces that term on the right side is lower bounded by 1 (a_j reachable from itself). Hence, term on the left side is lower bounded by 1. Since, it is also upper bounded by 1, it must be equal to 1.

Let $j \in [1, p] \setminus Y_i$. By definition of Y_i , we have $N_H^+(a_j) \not\subseteq N_H^+(a_i)$. That is, there exists $j' \in [1, q]$ such that $a_j b_{j'} \in E_H$ and $a_i b_{j'} \notin E_H$. Since $a_j b_{j'} \in E_H$, LABEL-LP enforces that

$$\sum_{\sigma \in \{0,1\}^p \sigma' \in \{0,1\}^q: \sigma[j]=1} z_{b_{j'}, \sigma \cdot \sigma'} = 0$$

Also, we have $a_i b_{j'} \notin E_H$ and hence, there is an infinite weight edge from a_i to $b_{j'}$ in G . Applying the result from previous part, we get $\sum_{\sigma \in \{0,1\}^p \sigma' \in \{0,1\}^q: \sigma[j]=1} z_{a_i, \sigma \cdot \sigma'} \leq \sum_{\sigma \in \{0,1\}^p \sigma' \in \{0,1\}^q: \sigma[j]=1} z_{b_{j'}, \sigma \cdot \sigma'} = 0$. Next, to prove that $z'_{a_i, \sigma} = 1$ if $J_\sigma = Y_i$ and 0 otherwise, we argue as follows:

$$\begin{aligned} 1 &= \sum_{\sigma \in \{0,1\}^p, \sigma' \in \{0,1\}^q} z_{a_i, \sigma \cdot \sigma'} \\ &= \sum_{\sigma \in \{0,1\}^p: J_\sigma = Y_i} \sum_{\sigma' \in \{0,1\}^q} z_{a_i, \sigma \cdot \sigma'} \\ &\quad + \sum_{\sigma \in \{0,1\}^p: J_\sigma \neq Y_i} \sum_{\sigma' \in \{0,1\}^q} z_{a_i, \sigma \cdot \sigma'} \\ &= \sum_{\sigma \in \{0,1\}^p: J_\sigma = Y_i} z'_{a_i, \sigma} \end{aligned}$$

- Again, we prove the following two statements which in turn implies that for $\sigma \in \{0,1\}^p, \sigma' \in \{0,1\}^q$ if $J_\sigma \neq Z_j$, then $z_{b_j, \sigma \cdot \sigma'} = 0$:

$$\begin{aligned} \forall i \in Z_j, \quad \sum_{\sigma \in \{0,1\}^p \sigma' \in \{0,1\}^q: \sigma[i]=1} z_{b_j, \sigma \cdot \sigma'} &= 1 \\ \forall i \in [1, p] \setminus Z_j, \quad \sum_{\sigma \in \{0,1\}^p \sigma' \in \{0,1\}^q: \sigma[i]=1} z_{b_j, \sigma \cdot \sigma'} &= 0 \end{aligned}$$

Let $i \in Z_j$. Hence, there is an infinite weight directed edge from a_i to b_j in G . Applying the result from first part, we get $\sum_{\sigma \in \{0,1\}^p \sigma' \in \{0,1\}^q: \sigma[i]=1} z_{b_j, \sigma \cdot \sigma'} \geq \sum_{\sigma \in \{0,1\}^p \sigma' \in \{0,1\}^q: \sigma[i]=1} z_{a_i, \sigma \cdot \sigma'}$. LABEL-LP enforces that right side is lower bounded by 1 (a_i reachable from itself). Hence, left side is lower bounded by 1. It is also upper bounded by 1 and hence, it must be equal to 1.

Let $i \in [1, p] \setminus Z_j$. Then, $a_i b_j \in E_H$ and hence, from the constraint in LABEL-LP

$$\sum_{\sigma \in \{0,1\}^p \sigma' \in \{0,1\}^q: \sigma[i]=1} z_{b_j, \sigma \cdot \sigma'} = 0$$

Next, to prove that $z'_{b_j, \sigma} = 1$ if $J_\sigma = Z_j$ and 0 otherwise, we argue as follows:

$$\begin{aligned} 1 &= \sum_{\sigma \in \{0,1\}^p, \sigma' \in \{0,1\}^q} z_{b_j, \sigma \cdot \sigma'} \\ &= \sum_{\sigma \in \{0,1\}^p: J_\sigma = Z_j} \sum_{\sigma' \in \{0,1\}^q} z_{b_j, \sigma \cdot \sigma'} \\ &\quad + \sum_{\sigma \in \{0,1\}^p: J_\sigma \neq Z_j} \sum_{\sigma' \in \{0,1\}^q} z_{b_j, \sigma \cdot \sigma'} \\ &= \sum_{\sigma \in \{0,1\}^p: J_\sigma = Z_j} z'_{b_j, \sigma} \end{aligned}$$

- For an undirected edge $e = uv$, $x_e = \sum_{\sigma_1, \sigma_2 \in \{0,1\}^{p+q}: \sigma_1 \neq \sigma_2} z_{e, \sigma_1 \sigma_2}$. Since, weight of e is infinite, x_e must be 0. Hence, $z_{e, \sigma_1 \sigma_2} = 0$ if $\sigma_1 \neq \sigma_2$. For $\sigma_1 \in \{0,1\}^{p+q}$

$$\begin{aligned} z_{u, \sigma_1} &= \sum_{\sigma_2 \in \{0,1\}^{p+q}} z_{e, \sigma_1 \sigma_2} = z_{e, \sigma_1 \sigma_1} \\ &= \sum_{\sigma_2 \in \{0,1\}^{p+q}} z_{e, \sigma_2 \sigma_1} = z_{v, \sigma_1} \end{aligned}$$

Let $t = (u) \in T_C$ s.t. $\psi_{T_C}(t) = \psi_{a_i}$. If $u = a_i$, then we have already proved that $z'_{u, \sigma} = 1$ if $J_\sigma = Y_i$ and 0 otherwise. If $u \neq a_i$, then there is an infinite weight undirected edge between u and a_i in G . Hence, $z'_{u, \sigma} = z'_{a_i, \sigma}$ for all $\sigma \in \{0,1\}^p$ and the result follows. \square

Proof of Lemma B.4: Next, we argue about the cost of the solution \mathbf{z}' . We assume here that (\mathbf{x}, \mathbf{z}) has finite cost. For a constraint $t \in T_C$, the cost according to \mathbf{z}' is $w_{T_C}(t) \sum_{\alpha \in L|t|} z'_{t, \alpha} \cdot \psi_t(\alpha)$. We consider four cases based on the type of t .

- t corresponds to constraint of the form $\psi_{a_i}(v)$. As argued in Lemma B.5, then $z'_{t, \sigma} = z_{v, \sigma} = 0$ if $J_\sigma \neq Y_i$ and 1 if $J_\sigma = Y_i$. On the other hand, $\psi_{a_i}(\sigma) = 0$ if $J_\sigma = Y_i$ and ∞ if $J_\sigma \neq Y_i$. Hence, $z'_{t, \sigma} \psi_{a_i}(\sigma) = 0$ for all σ . Therefore this constraint contributes zero to the cost.
- t corresponds to constraint of the form $\psi_{b_j}(v)$. From Lemma B.5, $z'_{t, \sigma} = z'_{v, \sigma} = 0$ if $J_\sigma \neq Z_j$ and 1 if $J_\sigma = Z_j$. And $\psi_{b_j}(\sigma) = 0$ if $J_\sigma = Z_j$ and ∞ if $J_\sigma \neq Z_j$. Hence, $z'_{t, \sigma} \psi_{b_j}(\sigma) = 0$ for all σ . Therefore the contribution of this constraint is zero.

- t corresponds to constraint $\mathcal{C}(u, v)$. This corresponds to a directed edge $e_t = (u, v)$ in G and the cost paid by (\mathbf{x}, \mathbf{z}) is x_{e_t} . The cost for t in \mathbf{z}' is given by:

$$\begin{aligned}
& \sum_{\sigma_1, \sigma_2 \in \{0, 1\}^p} z'_{t, \sigma_1 \sigma_2} \cdot \mathcal{C}(\sigma_1, \sigma_2) \\
&= \sum_{\sigma_1, \sigma_2 \in \{0, 1\}^p: \sigma_1 \not\leq \sigma_2} z'_{t, \sigma_1 \sigma_2} \\
&= \sum_{\sigma_1, \sigma_2 \in \{0, 1\}^p: \sigma_1 \not\leq \sigma_2} \sum_{\sigma', \sigma'' \in \{0, 1\}^q} z_{e_t, \sigma_1 \cdot \sigma' \sigma_2 \cdot \sigma''} \\
&\leq \sum_{\sigma_1, \sigma_2 \in \{0, 1\}^p, \sigma', \sigma'' \in \{0, 1\}^q: \sigma_1 \cdot \sigma' \not\leq \sigma_2 \cdot \sigma''} z_{e_t, \sigma_1 \cdot \sigma' \sigma_2 \cdot \sigma''} \\
&= x_{e_t}.
\end{aligned}$$

First equality follows from the fact that $\mathcal{C}(\sigma_1, \sigma_2) = 0$ if $\sigma_1 \leq \sigma_2$ and 1 otherwise. Penultimate inequality follows because if $\sigma_1 \not\leq \sigma_2$, then $\sigma_1 \cdot \sigma' \not\leq \sigma_2 \cdot \sigma''$ for any $\sigma', \sigma'' \in \{0, 1\}^q$.

- t corresponds to constraint $\text{NAE}_2(u, v)$. This corresponds to an undirected edge $e_t = uv$ in G and the cost paid by (\mathbf{x}, \mathbf{z}) is x_{e_t} . The cost for t in \mathbf{z}' is given by:

$$\begin{aligned}
& \sum_{\sigma_1, \sigma_2 \in \{0, 1\}^p} z'_{t, \sigma_1 \sigma_2} \cdot \text{NAE}_2(\sigma_1, \sigma_2) \\
&= \sum_{\sigma_1, \sigma_2 \in \{0, 1\}^p: \sigma_1 \neq \sigma_2} z'_{t, \sigma_1 \sigma_2} \\
&= \sum_{\sigma_1, \sigma_2 \in \{0, 1\}^p: \sigma_1 \neq \sigma_2} \sum_{\sigma', \sigma'' \in \{0, 1\}^q} z_{e_t, \sigma_1 \cdot \sigma' \sigma_2 \cdot \sigma''} \\
&\leq \sum_{\sigma_1, \sigma_2 \in \{0, 1\}^p, \sigma', \sigma'' \in \{0, 1\}^q: \sigma_1 \cdot \sigma' \neq \sigma_2 \cdot \sigma''} z_{e_t, \sigma_1 \cdot \sigma' \sigma_2 \cdot \sigma''} \\
&= x_{e_t}.
\end{aligned}$$

Combining the four cases, the total cost of the solution \mathbf{z}' is equal to the cost of the binary constraints each of which corresponds to an edge in G with the same weight. From the above inequalities we see that the cost is at most $\sum_{e \in E_G} w_G(e) x_e$ which is the cost of (\mathbf{x}, \mathbf{z}) . \square

From Basic-LP to Label-LP: Let \mathbf{z} be a BASIC-LP solution to I_C . Let $\sigma_0 = 1^q$. We define a solution $(\mathbf{x}', \mathbf{z}')$ to LABEL-LP for I_M as follows:

- For $v \in V_C, \forall \sigma_1 \in \{0, 1\}^p, \sigma_2 \in \{0, 1\}^q$,

$$z'_{v, \sigma_1 \cdot \sigma_2} = \begin{cases} z_{v, \sigma_1} & \sigma_2 = \sigma_0 \\ 0 & \text{otherwise} \end{cases}$$

- For unary constraint $t = (u)$ s.t. $u \notin \{a_1, \dots, a_p, b_1, \dots, b_q\}$ and $\sigma_1, \sigma_2 \in \{0, 1\}^p, \sigma_3, \sigma_4 \in \{0, 1\}^q$,

$$z'_{e_t, \sigma_1 \cdot \sigma_3 \sigma_2 \cdot \sigma_4} = \begin{cases} z_{u, \sigma_1} & \sigma_1 = \sigma_2, \sigma_3 = \sigma_4 = \sigma_0 \\ 0 & \text{otherwise} \end{cases}$$

- For binary constraint $t = (u, v) \in T_C$ such that $\psi_{T_C}(t) = \mathcal{C}$ or NAE_2 , and $\sigma_1, \sigma_2 \in \{0, 1\}^p, \sigma_3, \sigma_4 \in \{0, 1\}^q$

$$z'_{e_t, \sigma_1 \cdot \sigma_3 \sigma_2 \cdot \sigma_4} = \begin{cases} z_{t, \sigma_1 \sigma_2} & \sigma_3 = \sigma_4 = \sigma_0 \\ 0 & \text{otherwise} \end{cases}$$

- The edge variables x'_e are induced by the z' variables. We explicitly write them down. For directed edge $e \in E_G$, $x'_e = \sum_{\sigma_1, \sigma_2 \in \{0, 1\}^p, \sigma_3, \sigma_4 \in \{0, 1\}^q: \sigma_1 \cdot \sigma_3 \leq \sigma_2 \cdot \sigma_4} z'_{e, \sigma_1 \cdot \sigma_3 \sigma_2 \cdot \sigma_4}$. For undirected edge $e \in E_G$, $x'_e = \sum_{\sigma_1, \sigma_2 \in \{0, 1\}^p, \sigma_3, \sigma_4 \in \{0, 1\}^q: \sigma_1 \cdot \sigma_3 \neq \sigma_2 \cdot \sigma_4} z'_{e, \sigma_1 \cdot \sigma_3 \sigma_2 \cdot \sigma_4}$

It is easy to check that $(\mathbf{x}', \mathbf{z}')$ is integral if \mathbf{z} is

integral.

LEMMA B.6. $(\mathbf{x}', \mathbf{z}')$ is a feasible solution to LABEL-LP for I_M .

Proof. It is easy to check that all the variables are non-negative and upper bounded by 1.

We show that the other constraints are satisfied one at a time. Recall that LABEL-LP considered here has a constraint for undirected edges in addition to the constraints showed in Fig 3. The label set for LABEL-LP is $\{0, 1\}^{p+q}$ which we can write as $\{\sigma_1 \cdot \sigma_2 \mid \sigma_1 \in \{0, 1\}^p, \sigma_2 \in \{0, 1\}^q\}$.

Constraint 1: For each v , $\sum_{\sigma \in \{0, 1\}^{p+q}} z'_{v, \sigma} = 1$

$$\begin{aligned}
\sum_{\sigma_1 \in \{0, 1\}^p} \sum_{\sigma_2 \in \{0, 1\}^q} z'_{v, \sigma_1 \cdot \sigma_2} &= \sum_{\sigma_1 \in \{0, 1\}^p} z'_{v, \sigma_1 \cdot \sigma_0} \\
&= \sum_{\sigma_1 \in \{0, 1\}^p} z_{v, \sigma_1} = 1
\end{aligned}$$

Constraint 2: For $\sigma_1 \in \{0, 1\}^p, \sigma_2 \in \{0, 1\}^q$, $z'_{a_i, \sigma_1 \sigma_2} = 0$ if $\sigma_1[i] = 0$. And $z'_{b_j, \sigma_1 \sigma_2} = 0$ if $\sigma_2[j] = 0$.

There is $t = (a_i) \in T_C$ such that $\psi_{T_C}(t) = \psi_{a_i}$. For \mathbf{z} to be a finite valued solution, we must have $z_{t, \sigma_1} = z_{a_i, \sigma_1} = 0$ if $J_{\sigma_1} \neq Y_i$. Since, $i \in Y_i$, we have that $z_{a_i, \sigma_1} = 0$ if $\sigma_1[i] = 0$. And hence, $z'_{a_i, \sigma_1 \sigma_2} = 0$ if $\sigma_1[i] = 0$.

For $v \in V_C, z'_{v, \sigma_1 \sigma_2} = 0$ if $\sigma_2 \neq \sigma_0 = 1^q$. Hence, $z'_{v, \sigma_1 \sigma_2} = 0$ if $\sigma_2[j] = 0$. In particular, $z'_{b_j, \sigma_1 \sigma_2} = 0$ if $\sigma_2[j] = 0$.

Constraint 3: For $e = (u, v) \in E_G, \sigma_1 \in \{0, 1\}^p, \sigma_3 \in \{0, 1\}^q, z'_{u, \sigma_1 \cdot \sigma_3} = \sum_{\sigma_2 \in \{0, 1\}^p, \sigma_4 \in \{0, 1\}^q} z_{e, \sigma_1 \cdot \sigma_3 \sigma_2 \cdot \sigma_4}$. If $\sigma_3 \neq \sigma_0$, then all the terms are zero and hence, the equality holds. Else, $\sigma_3 = \sigma_0$ and there are two types of edges:

– For $t = (u)$, $e = e_t$

$$\begin{aligned} & \sum_{\sigma_2 \in \{0,1\}^p, \sigma_4 \in \{0,1\}^q} z'_{e, \sigma_1 \cdot \sigma_0 \sigma_2 \cdot \sigma_4} \\ &= \sum_{\sigma_2 \in \{0,1\}^p} z'_{e, \sigma_1 \cdot \sigma_0 \sigma_2 \cdot \sigma_0} \\ &= z_{u, \sigma_1} = z'_{u, \sigma_1 \cdot \sigma_0} = z'_{u, \sigma_1 \cdot \sigma_3} \end{aligned}$$

– For $t = (u, v)$, $e = e_t$,

$$\begin{aligned} & \sum_{\sigma_2 \in \{0,1\}^p, \sigma_4 \in \{0,1\}^q} z'_{e, \sigma_1 \cdot \sigma_0 \sigma_2 \cdot \sigma_4} \\ &= \sum_{\sigma_2 \in \{0,1\}^p} z'_{e, \sigma_1 \cdot \sigma_0 \sigma_2 \cdot \sigma_4} \\ &= \sum_{\sigma_2 \in \{0,1\}^p} z_{e, \sigma_1 \sigma_2} \\ &= z_{u, \sigma_1} = z'_{u, \sigma_1 \cdot \sigma_0} = z'_{u, \sigma_1 \cdot \sigma_3} \end{aligned}$$

Constraint 4: For $e = (u, v) \in E_G$, $\sigma_2 \in \{0, 1\}^p$, $\sigma_4 \in \{0, 1\}^q$, $z'_{v, \sigma_2 \cdot \sigma_4} = \sum_{\sigma_1 \in \{0,1\}^p, \sigma_3 \in \{0,1\}^q} z_{e, \sigma_1 \cdot \sigma_3 \sigma_2 \cdot \sigma_4}$. Proof is similar to the previous part.

Constraint 5: For directed edge e , $x'_e - \sum_{\sigma_1, \sigma_3 \in \{0,1\}^p, \sigma_2, \sigma_4 \in \{0,1\}^q: \sigma_1 \cdot \sigma_3 \not\leq \sigma_2 \cdot \sigma_4} z'_{e, \sigma_1 \cdot \sigma_3 \sigma_2 \cdot \sigma_4} = 0$. This is true by definition of x'_e .

Constraint 6: For undirected edge e , $x'_e - \sum_{\sigma_1, \sigma_3 \in \{0,1\}^p, \sigma_2, \sigma_4 \in \{0,1\}^q: \sigma_1 \cdot \sigma_3 \neq \sigma_2 \cdot \sigma_4} z'_{e, \sigma_1 \cdot \sigma_3 \sigma_2 \cdot \sigma_4} = 0$. This is true as well from the definition of x'_e . \square

LEMMA B.7. *The cost $(\mathbf{x}', \mathbf{z}')$ is upper bounded by the cost of \mathbf{z} .*

Proof. Recall that $\sigma_0 = 1^q$. We consider three cases based on the type of edge e

• $e = e_t = (u, v)$ for constraint $\mathcal{C}(u, v)$. Then, $x'_{e_t} =$

$$\begin{aligned} & \sum_{\sigma_1, \sigma_2 \in \{0,1\}^p, \sigma_3, \sigma_4 \in \{0,1\}^q: \sigma_1 \cdot \sigma_3 \not\leq \sigma_2 \cdot \sigma_4} z'_{e_t, \sigma_1 \cdot \sigma_3 \sigma_2 \cdot \sigma_4} \\ &= \sum_{\sigma_1, \sigma_2 \in \{0,1\}^p: \sigma_1 \cdot \sigma_0 \not\leq \sigma_2 \cdot \sigma_0} z'_{e_t, \sigma_1 \cdot \sigma_0 \sigma_2 \cdot \sigma_0} \\ &= \sum_{\sigma_1, \sigma_2 \in \{0,1\}^p: \sigma_1 \not\leq \sigma_2} z_{t, \sigma_1 \sigma_2} \end{aligned}$$

• $e = e_t = (u, v)$ for constraint $NAE_2(u, v)$. Then,

$x'_{e_t} =$

$$\begin{aligned} & \sum_{\sigma_1, \sigma_2 \in \{0,1\}^p, \sigma_3, \sigma_4 \in \{0,1\}^q: \sigma_1 \cdot \sigma_3 \neq \sigma_2 \cdot \sigma_4} z'_{e_t, \sigma_1 \cdot \sigma_3 \sigma_2 \cdot \sigma_4} \\ &= \sum_{\sigma_1, \sigma_2 \in \{0,1\}^p: \sigma_1 \cdot \sigma_0 \neq \sigma_2 \cdot \sigma_0} z'_{e_t, \sigma_1 \cdot \sigma_0 \sigma_2 \cdot \sigma_0} \\ &= \sum_{\sigma_1, \sigma_2 \in \{0,1\}^p: \sigma_1 \neq \sigma_2} z_{t, \sigma_1 \sigma_2} \end{aligned}$$

• $e = e_t = (u, a_i)$ or (u, b_j) for constraint $\psi_{a_i}(u)$ or $\psi_{b_j}(u)$. In such a case $z'_{e_t, \sigma_1 \cdot \sigma_3 \sigma_2 \cdot \sigma_4}$ is non-zero only if $\sigma_1 = \sigma_2, \sigma_3 = \sigma_4 = \sigma_0$. Hence, $x'_{e_t} =$

$$\sum_{\sigma_1, \sigma_2 \in \{0,1\}^p, \sigma_3, \sigma_4 \in \{0,1\}^q: \sigma_1 \cdot \sigma_3 \neq \sigma_2 \cdot \sigma_4} z'_{e_t, \sigma_1 \cdot \sigma_3 \sigma_2 \cdot \sigma_4} = 0$$

Combining the above three facts we get the following. First, infinite any infinite weight edge e in G has $x'_e = 0$. For any finite weight edge x'_e is the same as the fractional cost paid by the corresponding finite weight binary constraint in I_C . Hence, cost of $(\mathbf{x}', \mathbf{z}')$ is upper bounded by cost of \mathbf{z} . \square

B.2 Reduction from Dir-MulC-H to Min- β -CSP Let the DIR-MULC-H instance be $I_M = (G = (V_G, E_G, w_G : E_G \rightarrow R^+), (S \cup T, E_H))$. Recall that the supply graph satisfies assumptions I, II, and III. We reduce it an equivalent MIN- β -CSP instance $I_C = (V_C, T_C, \psi_{T_C} : T_C \rightarrow \beta_H, w_{T_C} : T_C \rightarrow R^+)$ as follows.

- Vertex Set $V_C = V_G$.
- T_C, ψ_{T_C}, w_{T_C} are defined as follows:
 - For every $a_i \in S$, add a tuple $t = (a_i)$ in T_C with $\psi_{T_C}(t) = \psi_{a_i}$ and $w_{T_C}(t) = 1$.
 - For every $b_j \in T$, add a tuple $t = (b_j)$ in T_C with $\psi_{T_C}(t) = \psi_{b_j}$ and $w_{T_C}(t) = 1$.
 - For every directed edge $e = (u, v) \in E_G$, add a tuple $t = (u, v)$ in T_C with $\psi_{T_C}(t) = \mathcal{C}$ and $w_{T_C}(t) = w_G(e)$.
 - For every undirected edge $e = uv \in E_G$, add a tuple $t = (u, v)$ in T_C with $\psi_{T_C}(t) = NAE_2$ and $w_{T_C}(t) = w_G(e)$.

The proof of equivalence between LABEL-LP for I_M and BASIC-LP for I_C is essentially identical to the proof for the reduction in the other direction and hence we omit it.

This finishes the proof of Theorem 3.3.