Constant Factor Approximation for Subset Feedback Set Problems via a new LP relaxation*

Chandra Chekuri   Vivek Madan

October 14, 2015

Abstract

We consider subset feedback edge and vertex set problems in undirected graphs. The input to these problems is an undirected graph $G = (V, E)$ and a set $S = \{s_1, s_2, \ldots, s_k\} \subset V$ of k terminals. A cycle in G is interesting if it contains a terminal. In the Subset Feedback Edge Set problem (Subset-FES) the input graph is edge-weighted and the goal is to remove a minimum weight set of edges such that no interesting cycle remains. In the Subset Feedback Vertex Set problem (Subset-FVS) the input graph is node-weighted and the goal is to remove a minimum weight set of nodes such that no interesting cycle remains.

A 2-approximation is known for Subset-FES [12] and a 8-approximation is known for Subset-FVS [13]. The algorithm and analysis for Subset-FVS is complicated. One reason for the difficulty in addressing feedback set problems in undirected graphs has been the lack of LP relaxations with constant factor integrality gaps; the natural LP has an integrality gap of $\Theta(\log n)$.

In this paper, we introduce new LP relaxations for Subset-FES and Subset-FVS and show that their integrality gap is at most 13. Our LP formulation and rounding are simple although the analysis is non-obvious.

1 Introduction

In the classical Feedback Vertex Set problem (FVS) the input is a a node-weighted graph $G = (V, E)$ and the goal is to find a minimum weight set of nodes whose removal makes the graph acyclic. FVS is interesting for its applications as well as connections to graph theory and combinatorial optimization. In this paper we restrict our attention to *undirected* graphs. FVS is easily seen to generalize the Vertex Cover problem and inherits NP-Hardness as well as the hardness of approximation bounds for Vertex Cover. We could also consider the Feedback Edge Set problem (FES) where the goal is remove a minimum weight set of edges to make it acyclic. FES is polynomial-time solvable; the complement of the edge-set of a maximum weight spanning tree in G can be easily seen to be an optimum solution. FVS and FES can also be viewed as hitting set problems where the goal is to find edges or nodes to intersect all cycles. In this paper we consider the more general *subset* feedback problems.

Subset Feedback Vertex Set (Subset-FVS): Input is an undirected graph $G = (V, E)$ along with non-negative node weights $w(v), v \in V$, and a set $S = \{s_1, \ldots, s_k\} \subset V$ of terminals. A cycle is interesting if it contains a terminal. The goal is to find a minimum weight set of nodes $V' \subset V$ that intersect all interesting cycles.

Subset Feedback Edge Set (Subset-FES): Input is an undirected graph $G = (V, E)$ along with a non-negative edge weights $w(e), e \in E$, and a set $S = \{s_1, \ldots, s_k\} \subset V$ of terminals. A cycle is interesting if it contains a terminal. The goal is to find a minimum weight set of edges $E' \subset E$ that intersect all interesting cycles.

Subset-FVS generalizes two well-known NP-complete problems. When $S = V$, we obtain FVS. When $|S| = 1$ it can be shown to be equivalent to the node-weighted Multiway Cut Problem (Node-wt-MC). In Node-wt-MC the input consists of a node-weighted graph G and a set of terminals T; the goal is to remove a minimum weight set of nodes such that there is no path left between any two terminals. Node-wt-MC can be reduced to Subset-FVS by adding a new node s of infinite weight and making it adjacent to each terminal in T. In a similar vein Subset-FES generalizes edge-weighted Multiway Cut Problem (Edge-wt-MC).

FVS, Subset-FVS and Subset-FES all admit constant factor approximation algorithms. In particular there is a 2-approximation for FVS [1, 3] and Subset-FES [12], and a 8-approximation for Subset-FVS [13]. There is a natural LP relaxation for these problems when viewed as a hitting set problem. For instance consider FVS. The relaxation has a variable $z(v) \in [0, 1]$ for each $v \in V$, and for each cycle C,

*Dept. of Computer Science, University of Illinois, Urbana, IL 61801. {chekuri,vmadan2}@illinois.edu. Supported in part by NSF grants CCF-1319376 and CCF-1526799.

a constraint $\sum_{v \in C} z(v) \geq 1$. This LP relaxation has an $\Omega(\log n)$ integrality gap [12]. Algorithms for feedback problems in undirected graphs have mainly relied on combinatorial techniques at the high-level. The non-trivial 2-approximation algorithm for FVS from [1] has been later interpreted as a primal-dual algorithm by Chudak *et al.* [8], however, the underlying LP is not known to be solvable in polynomial-time and does not generalize to Subset-FES or Subset-FVS. The 2-approximation for Subset-FES [12] is simple and combinatorial but delicate to analyze. The 8-approximation for Subset-FVS [13] is very complicated to describe and analyze; the algorithm is combinatorial at the high-level but solves a sequence of relaxed multicommodity flow LPs to optimality. Recall that Subset-FVS captures Node-wt-MC as a special case and all the known constant factor approximations for Node-wt-MC are via LP relaxations. To some extent this explains why one needs LP-type techniques for Subset-FVS. One of the open problems in Vazirani's book on approximation [23] is to find a simpler constant factor approximation algorithm for Subset-FVS with the eventual goal of finding an improved approximation ratio. Even *et al.* [12] write that "finding a linear program for Subset-FES and Subset-FVS for which the integrality gap is constant is very challenging".

In this paper we describe new LP relaxations for Subset-FES and Subset-FVS and derive constant factor approximations through them. Our results are captured by the following theorem.

Theorem 1.1. *There are polynomial-sized integer programming formulations for* Subset-FES *and* Subset-FVS *whose linear programming relaxations have an integrality gap of at most* 13.

The approximation bound of 13 that we are able to establish is weaker than the existing approximation ratios for the problems. However, we do not know of an integrality gap worse than 2 for the LP relaxations we propose. We believe that related formulations and ideas would lead to improved algorithms for Subset-FES and Subset-FVS.

Our formulation and algorithms are simple and are based on a new perspective on the problem. Although our analysis uses only elementary arguments it is not as straight forward.

1.1 The idea for the new LP formulations

We outline the key idea that allows us to develop new LP relaxations for Subset-FES and Subset-FVS. It is easier to explain it for Subset-FES. First, it is convenient to simplify the instance via well-known reductions; in the simplified instance each terminal s_i has degree 2 and is connected by infinite-weight edges to its neighbors a_i, b_i that are not adjacent to any other terminals. Thus, in any feasible solution of finite weight, the edges $s_i a_i$ and $s_i b_i$ are not cut; we think of these edges as *special* edges. Now, consider any minimal feasible solution $F \subset E$ such that the graph $H = G - F$ has no cycle containing a terminal. We can assume without loss of generality that G is connected, and hence the graph H is also connected by minimality of F. Consider the block-cut-vertex tree T of H [1]. No non-trivial block of H contains a terminal (otherwise there would be a cycle containing it). Thus, in H, each terminal is a cut-vertex and each special edge is a cut-edge. We can root T at a block r that does not contain any terminals; for simplicity assume that r is a single node. See Fig 1. Consider $k+1$ labels where terminal s_i has label i for $1 \leq i \leq k$ and the root has label $k+1$. By rooting T at r we obtain a natural label assignment for each node u of G as follows: label u by the index of the first terminal (or r) on a path in T from the block containing u to the root r. This labeling has the following property. The end points of each non-special edge in H have the same labels, and by minimality of F, the end points of edges in F receive different labels. It is important to note that the end points of some of the special edges receive different labels but they can never be cut. Thus, the problem can be viewed as finding a labeling of the nodes to minimize the cost of non-special edges whose end points receive different labels.

However, a labeling by itself does not suffice to obtain a good lower bound. One can assign the label $k+1$ to all non-terminal nodes and no non-special edges are cut. An additional property of the labeling obtained from reasoning via the block-cut-vertex tree is the following: for each terminal s_i, *exactly* one of the two neighbors a_i, b_i should be assigned the label i (the label of s_i). We can thus add this "spreading" constraint. The resulting labeling LP gets us most of the way; using just this LP, we can reduce the original instance to one in which each connected component "essentially" has only one interesting cycle which can be solved easily. To obtain a single LP we add constraints that ensure that the length of each cycle is at least one. To round the LP with assignment variables we borrow ideas from algorithms for multiway cut [5] and metric labeling [18] but we note that the rounding we have is subtly different because of the special edges and the spreading constraints; unlike those other problems we are not disconnecting the terminals.

[1] A block in a graph G is a maximal 2-node-connected component of G. The block-cut-vertex tree is a standard decomposition of a graph into its blocks and we refer the reader to books on graph theory such as [10, 24] for more information.

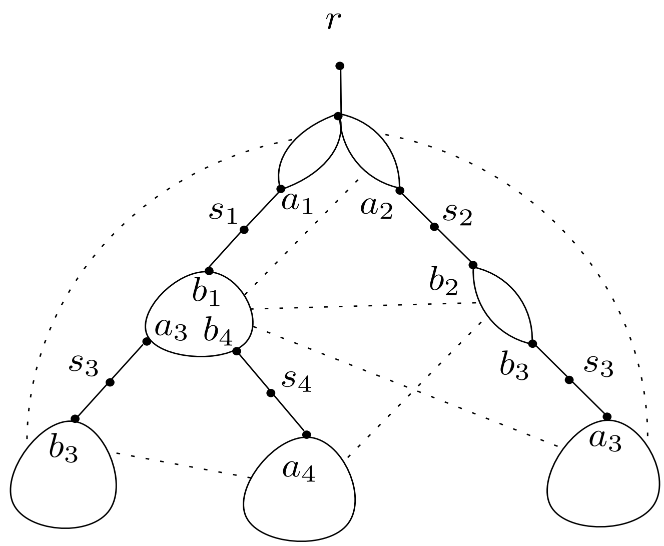

Figure 1: Block-cut-vertex tree of $H = G - F$. Dotted edges are edges in F that are cut

For Subset-FVS we use a similar labeling procedure but need additional variables to take into account node weights. It is easier to understand the formulation and analysis for Subset-FES before seeing the description and analysis for Subset-FVS. In fact some parts of the analysis for Subset-FVS rely on the analysis for Subset-FES. For this reason we first discuss Subset-FES in detail.

1.2 Other Related Work

Figure 2 taken from Vazirani's book on approximation [23] shows the relationship of Subset-FES and Subset-FVS to several well-known problems, some of which we already discussed. A natural open problem here is whether Subset-FVS has a 2-approximation.

Unweighted FVS is related to the well-known Erdos-Pósa theorem [11] which states that if k is the size of the smallest cardinality feedback vertex set in a graph G then there are $\Omega(k/\log k)$ node disjoint cycles in G. This immediately shows, via duality, that the integrality gap of the standard LP that we discussed in the introduction is $O(\log n)$ for unweighted FVS. The first constant factor approximation for the unweighted FVS problem is due to Bar-Yehuda *et al.* who obtained a 4-approximation [2]; the same paper also obtained a $2\Delta^2$ approximation for weighted FVS where Δ is the maximum degree. Subset-FVS also admits a Δ-approximation [12] which can be better than the 8-approximation [13] in some instances. Goemans and Williamson [16] considered FVS, Subset-FVS and other related problems in planar graphs and obtained a 9/4 approximation via a primal-dual algorithm with respect to the standard hitting set LP; recall that this LP has an $\Omega(\log n)$ integrality gap in general graphs.

Feedback problems are of much interest in directed graphs as well. FVS and FES are equivalent in directed graphs, and similarly Subset-FVS and Subset-FES. Leighton and Rao [19] obtained an $O(\log^2 n)$-approximation for FVS in directed graphs using their separator algorithms. Building on Seymour's work [21] that related the fractional packing of cycles to the minimum feedback vertex set[2], Even *et al.* [14] obtained a $O(\log k \log\log k)\}$ approximation for Subset-FVS. It is also known that assuming the Unique Game Conjecture FVS in directed graphs does not admit a constant factor approximation [17, 22].

Recall that Subset-FES and Subset-FVS generalize multiway cut problems Edge-wt-MC and Node-wt-MC respectively. Calinescu, Karloff and Rabani [5] developed a labeling LP for Edge-wt-MC which has since been extensively studied. The labeling LP can be motivated by viewing Edge-wt-MC as a partition problem instead of viewing it as an edge removal problem. Although it is slightly less natural, Node-wt-MC can also be viewed as a partition problem indirectly via the hypergraph-cut problem (see [7, 6] and references). Despite this indirect connection, a labeling based LP and rounding for Node-wt-MC has not explicitly been written in the literature. Our work here gives such an LP in the more general context of Subset-FVS. Label-

[2] Erdos-Pósa theorem relates integer packing of cycles to feedback vertex sets in undirected graphs. The relationship between integer packing of cycles and feedback vertex sets in directed graphs is much more difficult and is addressed by Reed *et al.* [20].

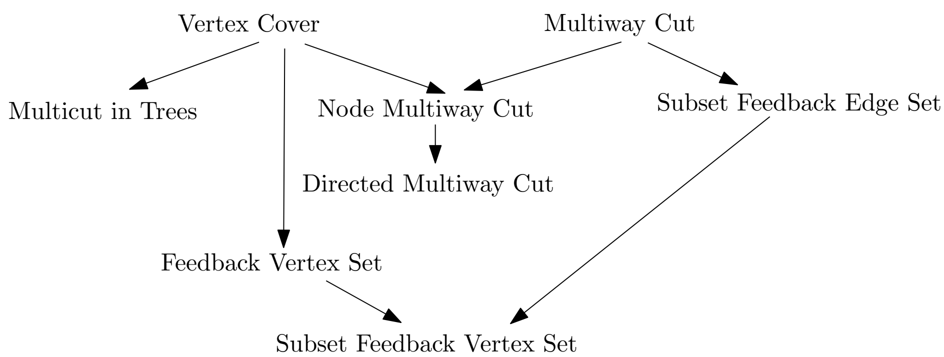

Figure 2: Approximation preserving reductions related to feedback problems. Figure is reproduced from [23]. All problems except Subset-FVS have a 2-approximation with Multiway Cut admitting an approximation better than 2.

ing problems such as metric labeling [18], zero-extension [4] and submodular-cost labeling [7] have provided powerful tools to address a variety of problems and our work here gives yet another application to an interesting class of problems.

We refer the reader to a survey on feedback set problems for additional information [15]. There has also been extensive work on fixed-parameter algorithms for feedback problems — see [9].

1.3 Organization

The rest of the paper is organized as follows. In Section 2, we start by describing the LP formulation for Subset-FES, and then discuss our two step rounding scheme. Building upon these ideas, in Section 3 we describe a similar LP formulation and rounding scheme for Subset-FVS.

2 LP-relaxation based constant factor approximation for Subset-FES

In this section we describe our LP-relaxation based algorithm and analysis for Subset-FES. First, we will assume without loss of generality that the input instance has a certain restricted structure; similar assumptions have been used previously [12] and are easy to justify. The assumptions are: (i) The input graph G is connected. (ii) Each terminal s_i has degree 2 and is connected by infinite weight edges to it's neighbors a_i, b_i. (iii) No two terminals are connected by an edge or share a neighbor. (iv) There exists a special non-terminal vertex r with a single infinite weight edge incident to it.

We briefly justify these assumptions: If G is not connected, the problem can be solved separately in each connected component. If a terminal s is not degree two, we can sub-divide each edge incident to s by adding a new node; then, remove s from the set of terminals and add the new nodes to the set of terminals. If $e = uv$ is an edge of weight $w(e)$, sub-dividing e by adding a node q and setting $w(uq) = \infty$ and $w(qv) = w(e)$ does not change the problem. This can be used to justify the other assumptions as well.

For technical reasons, we also assume that every interesting cycle contains *at two* terminals. This can also be justified by subdividing the edges incident on a terminal and adding the new nodes as terminals. We perform one more reduction step to ensure that no two terminals are connected by an edge or share a neighbor. Note that in the new instance that satisfies these assumption, the number of terminals could be much larger than in the original instance.

Remark 2.1. *We refer readers who may wonder about the need or the utility of the simplifying reductions, to Section A. There we provide some additional details and examples.*

2.1 LP formulation

Recall the structure of a minimal feasible solution discussed in Section 1.1. If $F \subset E$ is a minimal feasible solution then $H = (V, E \setminus F)$ is connected and each terminal is a cut vertex in H. Let T be the block-cut-vertex tree of H rooted at r. Each vertex u is labeled by the index of the first terminal on a path in T from the block containing u to r, or by $k+1$ if there is no such terminal. We will use E_s to denote the set of special edges $\cup_{i=1}^{k}\{(s_i, a_i), (s_i, b_i)\}$ that are incident to the terminals and have infinite weight. Any non-special edge with different labels on the end points is cut. Special edges are never cut and for each terminal, one of the incident special edge always has different labels on the end points. We formulate an integer program

$$
\begin{array}{llcr}
 & \textbf{Subset-FES-Rel} & & \\
\min & \displaystyle\sum_{e \in E \setminus E_s} w(e) z(e) & & \\
 & s.t & & \\
 & \displaystyle\sum_{i=1}^{k+1} x(u,i) & = \quad 1 & u \in V \\
 & x(s_i, i) & = \quad 1 & 1 \le i \le k \\
 & x(r, k+1) & = \quad 1 & \\
 & x(a_i, i) + x(b_i, i) & = \quad 1 & 1 \le i \le k \\
 & \displaystyle z(e) - \frac{1}{2}\sum_{i=1}^{k+1} |x(u,i) - x(v,i)| & \ge \quad 0 & e = uv \in E \setminus E_s \\
 & z(e) & = \quad 0 & e \in E_s \\
 & \displaystyle\sum_{e \in C} z(e) & \ge \quad 1 & C \in \mathbb{C} \\
 & x(u,i) & \ge \quad 0 & u \in V, 1 \le i \le k+1 \\
 & z(e) & \ge \quad 0 & e \in E
\end{array}
$$

Figure 3: LP Relaxation for SUBSET-FES

based on this structure which can then be relaxed to obtain a linear program. We have two types of binary variables, the labeling variables $x(u,i)$ for each $u \in V$ and $i \in \{1, \ldots, k+1\}$, and the edge variables $z(e)$ for each $e \in E$. $x(u,i)$ is an indicator variable for whether u is assigned label i. $z(e)$ is an indicator variable for whether e is cut or not. The following constraints explain our reasoning:

- Each node u is labeled by exactly one label: $\sum_{i=1}^{k+1} x(u,i) = 1$ for all $u \in V$.
- Terminals are labeled by their own index, $x(s_i, i) = 1$ for each i. Root r is labeled $k+1$, $x(r, k+1) = 1$.
- For each s_i, exactly one of a_i, b_i is labeled i: $x(a_i, i) + x(b_i, i) = 1$ for $1 \le i \le k$.
- Non-special edge $e = uv$ is cut (that is $z(e) = 1$) if u, v receive different labels and is not cut ($z(e) = 0$) if u, v receive same labels. Hence, $z(e) \ge \frac{1}{2}\sum_{i=1}^{k+1} |x(u,i) - x(v,i)|$. This can be written as a linear constraint with additional variables that we suppress for ease of notation.
- Special edges are not cut: $z(e) = 0$ for $e \in E_s$.

We also have an additional constraint. Let $\mathbb{C}$ be the set of interesting cycles. For any $C \in \mathbb{C}$, at least one of the edges is cut: hence $\sum_{e \in C} z(e) \ge 1$. There are an exponential number of such constraints but we can can express them compactly via triangle inequalities and it is also easy to see that we can separate over them efficiently in polynomial-time. This constraint is essential for LP to have a bounded integrality gap. See Section A for an example illustrating this fact. Objective is to minimize $\sum_{e \in E \setminus E_s} w(e) z(e)$. We can drop the constraints that upper bound the variables by 1. The full description of the LP relaxation is given in Figure 3.

2.2 Rounding scheme and analysis

THEOREM 2.1. *There is a polynomial-time algorithm that given a feasible solution* $\boldsymbol{x}, \boldsymbol{z}$ *to* SUBSET-FES-REL *outputs a feasible integral solution of weight at most* $13 \sum_{e \in E \setminus E_s} w(e) z(e)$.

Given a feasible solution $\mathbf{x}, \mathbf{z}$ to SUBSET-FES-REL, we round it in two steps. In the first step, we round the fractional solution using the labeling variables $x(u,i)$ to find a subset $E' \subset E$ of edges such that removing E' yields a graph $G' = G - E'$ that has very restricted structure. In particular each connected component of G' has essentially only one interesting cycle; more formally all interesting cycles in each component have the same *signature* which is defined formally below. Solving an instance in which all cycles have the same signature is easy; we can find an optimal solution. Letting E'' denote the edge set removed in the second step (we take

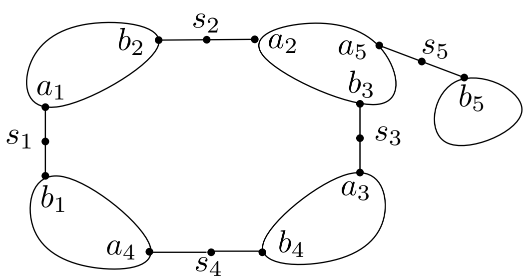

Figure 4: Structure of connected component in $G' = G - E'$

the union of the solutions from each component), the final output of the algorithm is $E' \cup E''$.

Definition 2.2. Let $C = s_{i_1}, c_{i_1}, A_1, c'_{i_2}, s_{i_2}, c_{i_2}, A_2, \ldots, c'_{i_t}, s_{i_t}, c_{i_t}, A_t, c'_{i_1}, s_{i_1}$ be an interesting cycle where $s_{i_1}, \ldots, s_{i_t}$ are terminals, and for $1 \leq j \leq t$, $c_{i_j}, c'_{i_j} \in \{a_{i_j}, b_{i_j}\}$ and $c_{i_j}, A_j, c'_{i_{j+1}}$ is a path with no terminals; here A_j can be empty. Signature of C denoted by $sig(C)$ is defined as $s_{i_1}, c_{i_1}, c'_{i_2}, s_{i_2}, c_{i_2}, \ldots, c'_{i_t}, s_{i_t}, c_{i_t}, c'_{i_1}, s_{i_1}$. Given two cycles C_1 and C_2 we say that their signatures are the *same* if the cycles $sig(C_1)$ and $sig(C_2)$ are isomorphic as labeled graphs.

The heart of the rounding and analysis is the first step which is formalized in the lemma below.

Lemma 2.3. *Given a feasible solution $\boldsymbol{x}, \boldsymbol{z}$ to* Subset-FES-Rel*, there is an efficient algorithm to find a subset of edges $E' \subset E$ with cost at most $12 \sum_{e \in E \setminus E_s} w(e)z(e)$ such that any two interesting cycles in the same connected component of $G' = G - E'$ have the same signature.*

It is useful to see Figure 4 to understand what it means for all interesting cycles to have the same signature. In particular, in each connected component H of G' that has an interesting cycle C, removing any terminal from C suffices to kill all interesting cycles in H.

Algorithm 1 Initial Cut for Subset-FES

Given: Feasible solution $\mathbf{x}, \mathbf{z}$ to Subset-FES-Rel
Pick $\theta \in (1/3, 1/2)$ uniformly at random
For $1 \leq i \leq k$, $B_i := \{u \mid x(u,i) > \theta\}$
$E' := \left(\cup_{i=1}^{k} \delta(B_i)\right) \setminus E_s$
Return E'

Algorithm 1 is a simple randomized algorithm that achieves the properties claimed by the preceding lemma. Here, $\delta(S)$ denote the edge boundary of set S, formally defined as $\{uv \in E \mid |\{u,v\} \cap S| = 1\}$. We note that although the algorithm is related to "ball-cutting" type schemes for cut problems such as multiway cut, there are subtle differences. It is important for our analysis that $\theta < 1/2$ which is counter intuitive from a cut perspective. Since $\theta \in (1/3, 1/2)$, B_i and B_j may intersect for $i \neq j$. Another subtlety is the fact that a special edge e may be in $\delta(B_i)$ for some i but is not allowed to be cut. These issues make the analysis tricky.

We define the label set for v, $L(v) = \{i \mid v \in B_i\}$; note that some nodes may not receive a label. We make some simple observations before proceeding with the proof of Lemma 2.3. In the analysis below, all statements hold for each choice of $\theta \in (1/3, 1/2)$, and randomness plays a role only in analyzing the expected cost of E'.

Lemma 2.4. *A node can have at most two labels that is, for each vertex $v \in V, |L(v)| \leq 2$.*

Proof: A label $i \in L(v)$ implies $x(v,i) > \theta > 1/3$. Since $\sum_{i=1}^{k+1} x(v,i) = 1$, $|L(v)| \leq 2$. □

The spreading constraint and the choice of θ ensures some useful properties about the labels of a_i and b_i.

Lemma 2.5. *For each $i \in \{1, \ldots, k\}$, a_i or b_i is labeled i, that is, $i \in L(a_i) \cup L(b_i)$. Moreover if $i \notin L(a_i)$, then $L(b_i) = \{i\}$. Similarly, if $i \notin L(b_i)$, then $L(a_i) = \{i\}$.*

Proof: The constraint $x(a_i,i) + x(b_i,i) = 1$ implies that $x(a_i,i) \geq 1/2 > \theta$ or $x(b_i,i) \geq 1/2 > \theta$. Hence, $a_i \in B_i$ or $b_i \in B_i$ which is equivalent to $i \in L(a_i) \cup L(b_i)$. If $i \notin L(a_i)$, $x(a_i,i) \leq \theta < 1/2$ and therefore $x(b_i,i) = 1 - x(a_i,i) \geq 1 - \theta > 1/2$. Therefore $i \in L(b_i)$. Also, since $\sum_{j=1}^{k+1} x(b_i,j) = 1$, and $x(b_i,i) \geq 1-\theta$, for all $\ell \neq i$, $x(b_i,\ell) \leq \theta$. Hence, $L(b_i) = \{i\}$. □

Lemma 2.6. *For any non-special edge $uv \in E \setminus E'$, $L(u) = L(v)$. Hence, for any two nodes u, v connected in G' by a path with only non-special edges, $L(u) = L(v)$.*

Proof: Let $uv \in E \setminus E'$ be a non-special edge with $L(u) \neq L(v)$. Then, there exists i in $L(u) \setminus L(v)$ or $L(v) \setminus L(u)$. In both cases, edge uv is in $\delta(B_i)$ and thus in E'. This leads to contradiction as $uv \in E \setminus E'$. □

Now comes an important lemma on the label set of nodes in any interesting cycle remaining in G'.

Lemma 2.7. *Let C be an interesting cycle in G' with $sig(C) = s_{i_0}, c_{i_0}, c'_{i_1}, s_{i_1}, c_{i_1}, \ldots, c'_{i_{t-1}}, s_{i_{t-1}}, c_{i_{t-1}}, c'_{i_0}, s_{i_0}$ where $s_{i_0}, \ldots, s_{i_{t-1}}$ are terminals, and $c_{i_j}, c'_{i_j} \in \{a_{i_j}, b_{i_j}\}$. If $t \geq 2$, then exactly one of the following three conditions hold. Addition and subtraction of the indices here is modulo t.*

- *For* $0 \le j \le t-1, L(c_{i_j}) = \{i_j, i_{j+1}\}, L(c'_{i_j}) = \{i_j, i_{j-1}\}$
- *For* $0 \le j \le t-1, L(c_{i_j}) = \{i_j\}, L(c'_{i_j}) = \{i_{j-1}\}$
- *For* $0 \le j \le t-1, L(c_{i_j}) = \{i_{j+1}\}, L(c'_{i_j}) = \{i_j\}$

Proof: For ease of notation, we rename the terminals and their neighbors such that $sig(C) = s_1, a_1, b_2, s_2, a_2, \ldots, b_t, s_t, a_t, b_1, s_1$ and identify $t+1$ with 1, 0 with t when adding and subtracting indices.
Case 1: $\forall j \in \{1, \ldots, t\}, j \in L(a_j) \cap L(b_j)$. Since a_j is connected by a path containing non-special edges to b_{j+1} in C (and hence G'), by Lemma 2.6, $L(a_j) = L(b_{j+1})$. From assumption $t \ge 2$, we know that $b_{j+1} \ne b_j$ and label $j+1$ is not same as label j. And, since $j \in L(a_j)$ and $j+1 \in L(b_{j+1})$, we get $j, j+1 \in L(a_j), L(b_{j+1})$. Lemma 2.4 states that $|L(v)| \le 2$ for all v which implies $L(a_j) = L(b_{j+1}) = \{j, j+1\}$. The labels in this case satisfies the first condition in the lemma.
Case 2: $\exists j \in \{1, \ldots, t\}, j \notin L(b_j)$. Since, $j \notin L(b_j)$, by Lemma 2.5, we have $L(a_j) = \{j\}$. Since a_j is connected to b_{j+1} via non-special edges, by Lemma 2.6, $L(a_j) = L(b_{j+1}) = \{j\}$. This implies that $j+1 \notin b_{j+1}$. Arguing inductively along the cycle we can see that $L(a_i) = L(b_{i+1}) = \{i\}$ for $1 \le i \le t$. The labels in this case satisfy the second condition in the lemma.
Case 3: $\exists j \in \{1, \ldots, t\}, j \notin L(a_j)$. Similar to Case 2 and labels satisfy third condition in the lemma. $\Box$

REMARK 2.8. *There may be nodes which are not labeled but are connected to terminals in G'. For example, if $v = a_i$ with $x(v, j) \le \theta$ for all j, then $L(v) = \emptyset$ but v is still connected to s_i in G' since special edges are not cut. Lemma 2.7 implies that such a vertex is not part of any interesting cycle in G'.*

Next, we prove bound on expected cost[3] of the edges cut by Algorithm 1.

LEMMA 2.9. $\Pr[e \in E'] \le 12z(e)$, *and hence the expected cost of E' is at most* $12\sum_{e \in E \setminus E_s} w(e)z(e)$.

Proof: We focus on non-special edges as no special edge is cut by the algorithm. Let $e = uv$ be any such edge. Edge e is cut if and only if $e \in \delta(B_i)$ for some i. $\Pr[e \in \delta(B_i)] = \Pr[\theta \in [\min(x(u,i), x(v,i)), \max(x(u,i), x(v,i)))]$ which is at most $6|x(u,i) - x(v,i)|$ since θ is chosen uniformly from $(1/3, 1/2)$. By union bound, $\Pr[e \in E'] \le \sum_{i=1}^{k} \Pr[e \in \delta(B_i)] \le 6\sum_{i=1}^{k} |x(u,i) - x(v,i)| \le 12z(e)$. $\Box$

[3]It is easy to derandomize the algorithm by trying "all possible" values of θ in $(1/3, 1/2)$; we only need to try all values of $x(u, i)$.

We now prove the main structural observation, namely Lemma 2.3, that in each connected component of G' all interesting cycles have the same signature.

Proof of Lemma 2.3: We prove the lemma for the edge set E' returned by Algorithm 1. The expected cost of E' is upper bounded by Lemma 2.9.

In the graph $G' = (V, E \setminus E')$, consider any two interesting cycles C_1, C_2. We will first prove that if C_1 and C_2 share a terminal, then their signature must be same. Second, we will prove that if C_1, C_2 do not share a terminal then then they must be in different connected components.

First, consider the case when C_1 and C_2 share a terminal vertex. By renaming terminal nodes and their neighbors, let the cycle C_1 be such that $sig(C_1) = s_1, a_1, b_2, s_2, a_2, \ldots, b_t, s_t, a_t, b_1, s_1$ and the shared terminal be s_1. Since each terminal has degree 2, any cycle through s_1 has to contain a_1, b_1. We can assume without loss of generality that $sig(C_2) = s_{i_1}, c'_{i_1}, c_{i_2}, s_{i_2}, c'_{i_2}, \ldots, c_{i_h}, s_{i_h}, c'_{i_h}, c_{i_1}, s_{i_1}$, where $s_{i_1} = s_1$ and $c'_{i_1} = a_1, c_{i_1} = b_1$. Since $sig(C_1) \ne sig(C_2)$ there is a smallest integer $r \ge 2$ such that $c_{i_r} \ne b_r$ and the prefix of $sig(C_2)$ till c_{i_r} agrees with $sig(C_1)$. From Lemma 2.6 we have $L(a_{r-1}) = L(b_r)$ (via C_1) and $L(a_{r-1}) = L(c_{i_r})$ (via C_2) and hence $L(b_r) = L(c_{i_r})$. Given a SUBSET-FES instance, we assumed that each interesting cycle contains at least two terminals. Hence, $t, h \ge 2$. We will consider Lemma 2.7 applied to C_1 and C_2 and the resulting consistency requirements on the labels, in particular for a_{r-1}, b_r, c_{i_r}. Note that these nodes all receive two distinct labels or all receive exactly one label. We consider three cases based on the guarantee of Lemma 2.7 applied to C_1, and in each case derive a contradiction.

- For all $j, L(a_j) = \{j, j+1\}, L(b_j) = \{j, j-1\}$. Thus $L(a_{r-1}) = L(b_r) = \{r-1, r\}$ and $L(a_r) = \{r, r+1\}$. Lemma 2.7 applied to C_2 implies that $L(a_{r-1}) = L(c_{i_r}) = \{r-1, i_r\}$. This in particular implies that $i_r = r$, and hence $s_{i_r} = s_r$. Since we have $c_{i_r} \ne b_r$ we must have that $c_{i_r} = a_r$. However, $L(a_r) = \{r, r+1\}$ via C_1 which means that $L(a_{r-1}) \ne L(c_{i_r})$, a contradiction.
- For all $j, L(a_j) = \{j+1\}, L(b_j) = \{j\}$. Thus $L(a_{r-1}) = \{r\}$ and $L(b_r) = \{r\}$ and $L(a_r) = \{r+1\}$. Lemma 2.7 applied to C_2 implies that $L(a_{r-1}) = L(c_{i_r}) = \{r-1\}$ or $L(a_{r-1}) = L(c_{i_r}) = \{i_r\}$. Since, $L(a_{r-1}) = \{r\}$, $L(a_{r-1}) = L(c_{i_r}) = \{i_r\} = \{r\}$. Thus $i_r = r$; since $c_{i_r} \ne b_r$ we have $c_{i_r} = a_r$ but then $L(c_{i_r}) = L(a_r) = \{r\}$, a contradiction.
- For all $j, L(a_j) = \{j\}, L(b_j) = \{j-1\}$. Similar to preceding case.

Next, consider the case when C_1 and C_2 do not share a terminal. By contradiction, assume that they are connected in G'. For an interesting cycle C let $S(C)$ denote the indices of the set of terminals in $V(C)$. From Lemma 2.7 and Lemma 2.6, for each $v \in V(C)$, $\emptyset \neq L(v) \subseteq S(C)$. Since, $S(C_1) \cap S(C_2) = \emptyset$, we conclude that $V(C_1) \cap V(C_2) = \emptyset$.

Let $u \in C_1, v \in C_2$ be two nodes in G' connected by a path not intersecting C_1 or C_2. Since, each terminal is a degree two node, neighbors of all terminals in $S(C_1)$ and $S(C_2)$ are part of cycle C_1 and C_2, we have that u and v are not terminals. Let the path between u and v in G' be $u, A_1, c_{i_1}, s_{i_1}, c'_{i_1}, A_2, \ldots, A_r, c_{i_r}, s_{i_r}, c'_{i_r}, A_{r+1}, v$ where $s_{i_1}, \ldots, s_{i_r}$ are terminals, $c_{i_j}, c'_{i_j} \in \{a_{i_j}, b_{i_j}\}$ and A_j is a path with no terminals. First we observe that if $r = 0$ then there is a path connecting u to v without terminals and hence $L(u) = L(v)$; this is not possible since $L(u) \subset S(C_1)$ and $L(v) \subset S(C_2)$ and $S(C_1) \cap S(C_2) = \emptyset$. Assume $r \geq 1$. We have $L(u) = L(c_{i_1})$ and since $L(u) \subset S(C_1)$ and $s_{i_1} \notin V(C_1)$, we obtain that $i_1 \notin L(c_{i_1})$. This implies, via Lemma 2.5 that $L(c'_{i_1}) = \{i_1\}$.

Claim 2.10. *For $1 \leq j \leq r$, $L(c'_{i_j}) = \{i_j\}$.*

Proof: We have already established the base case that $L(c'_{i_1}) = \{i_1\}$. Assume by induction that $L(c'_{i_j}) = \{i_j\}$ for $j = \ell - 1$. By Lemma 2.6, $L(c_{i_\ell}) = L(c'_{i_{\ell-1}}) = \{i_{\ell-1}\}$. By Lemma 2.5, $L(c_{i_\ell}) = \{i_\ell\}$. □

Using the above claim, $L(c'_{i_r}) = \{i_r\}$ which by Lemma 2.6 implies that $L(v) = \{i_r\}$. However, $L(v) \subset S(C_2)$ and $s_{i_r} \notin V(C_2)$ which is a contradiction. □

Remark 2.11. *The assumption that each interesting cycle contains at least two terminals was used in the proof of Lemma 2.7. Without this assumption, we may not get the property of Lemma 2.3 that each interesting cycle in G' has same signature. For examples and more discussion on this, see Section A.*

A part of the the analysis of Algorithm 1 will be useful for us when analyzing the algorithm for Subset-FVS. The following remark captures the necessary aspects.

Remark 2.12. *Constraints involving variables $z(e)$ are used only while bounding the expected cost of the edge set E' returned by Algorithm 1. Hence, given any vector $\boldsymbol{x}$ satisfying constraints involving labeling variables(first four and $x(u,i) \geq 0$), Algorithm 1 returns an edge set E' such that all interesting cycles in a connected components in $G - E'$ have same signature.*

2.3 Second Step of Rounding

Next, we will describe the second step of rounding for Subset-FES-Rel and finish the proof of Theorem 2.1. After the first step of the rounding we are left with a graph G'. In each connected component of G' all interesting cycles have the same signature. Consider one such component H which has an interesting cycle C and without loss of generality let $sig(C) = s_1, a_1, b_2, s_2, a_2, \ldots, s_t, a_t, b_1, s_1$. Note that t could be 1. Since the signatures of all interesting cycles are the same, any such cycle C contains the edges s_1a_1 and s_1b_1. Thus to remove all interesting cycles in H it is necessary and sufficient to disconnect a_1 from b_1 in $H' = H - \{s_1a_1, s_1b_1\}$. This can be easily solved by finding a min-cut between a_1 and b_1 in H'. The following claim charges the cost of this min-cut to the LP solution.

Claim 2.13. *The cost of the min-cut between a_1 and b_1 in H' is at most $\sum_{e \in E(H)} w(e)z(e)$.*

Proof: Consider the feasible solution $\mathbf{x}, \mathbf{z}$ to Subset-FES-Rel to the original instance and the lengths of the edges $z(e)$ induces on $E(H)$. The constraints of the relaxation imply that $\sum_{e \in C} z(e) \geq 1$ for each $C \in \mathbb{C}$. Since H is a subgraph of G these constrains hold for all interesting cycles in H as well. Since $z(s_1a_1)$ and $z(s_1b_1)$ are both 0, we have that $d_z(a_1, b_1) \geq 1$ in $H' = H - \{s_1a_1, s_1b_1\}$ where $d_z(a_1, b_1)$ is the shortest path distance from a_1 to b_1 according to edge lengths given by $\mathbf{z}$. Thus $\mathbf{z}$ restricted to H' is a feasible fractional solution to the standard distance based LP relaxation for the a_1-b_1 minimum cut problem in H'; this LP relaxation is the same as the dual of the maximum flow LP. Via the maxflow-mincut theorem, the integrality gap of the LP is one and in particular there is an a_1-b_1 mincut in H' of cost at most $\sum_{e \in E(H)} w(e)z(e)$. □

Let E'' be the union of all the minimum cuts found in different connected components in G'. It is easy to see that $E' \cup E''$ is a feasible solution. From the preceding claim and the fact that the connected components of G' are edge disjoint we have that $w(E'') \leq \sum_{e \in E \setminus E_s} w(e)z(e)$.

We now finish the proof of Theorem 2.1 that establishes a constant factor upper bound on integrality gap for Subset-FES-Rel.

Proof of Theorem 2.1: Let $\alpha = \sum_{e \in E \setminus E_s} w(e)z(e)$ be the the objective function value of a feasible solution $\mathbf{x}, \mathbf{z}$ to Subset-FES-Rel. From Lemma 2.3 we obtain a set of edges E' such that $w(E') \leq 12\alpha$ and $G' = G - E'$ satisfies the property needed for the second step of the algorithm. The set of edges E'' found in the second step satisfy the property that $w(E'') \leq \alpha$. Thus

$w(E' \cup E'') \leq 13\alpha$ and $E' \cup E''$ is a feasible integral solution to the given instance. This finishes the proof. $\square$

3 LP-relaxation based constant factor approximation for Subset-FVS

In this section we extend the ideas from Section 2 to handle Subset-FVS. Several of the ideas behind the rounding and the analysis are quite similar, however, there are some non-obvious technical differences that we point out as we go along. We will again assume that the input instance satisfies some restricted structure: (i) The input graph G is connected. (ii) Each terminal has infinite weight, is a degree two vertex with both the neighbors having infinite weight. (iii) No two terminals are connected by an edge or share a neighbor. (iv) There exists a special non-terminal degree one vertex r with infinite weight. (v) Each interesting cycle contains at least two terminals. Justification of these assumptions is very similar to the case of Subset-FES.

3.1 LP formulation

Let $A \subset V$ be a minimal feasible solution. Since, terminals and their neighbors have infinite weight, none of these nodes are in A. Each terminal is a cut vertex in $H = G - A$. But, unlike the setting of Subset-FES, H might not be connected even if A is a minimal feasible solution. For simplicity we will first assume that H is connected. As before we can now obtain a labeling by considering a block-cut-vertex tree T of H rooted at r, the block containing the special non-terminal r. This leads to a labeling of nodes in $V \setminus A$ with labels 1 to $k+1$ where a node u is labeled i if s_i is the first terminal on the path from the block containing u to r in T. Note that only nodes in $V \setminus A$ receive a label. As before the property of this labeling is that for any non-special edge uv in $E(H)$, u and v receive same labels. Now we briefly address the case when H may not be connected. For each connected component we pick an arbitrary non-terminal vertex, add a dummy edge connecting it to r and consider the labeling corresponding to block-cut vertex tree of the modified graph. Note that the addition of these dummy edges is a thought experiment to justify the validity of the existence of the labeling.

As in case of Subset-FES, we formulate an integer program based on this structure with a few changes. We have two types of binary variables, the labeling variables $x(u, i)$ for each $u \in V$ and $1 \leq i \leq k$ and node variables $z(u)$ for each node $u \in V$. $x(u, i)$ indicates whether or note u is assigned label i and $z(u)$ indicates whether or not u is cut. Here are some of the constraints, a minimal feasible solution satisfies, based on the reasoning via the structure of block-cut-vertex tree T.

- Either a node u is cut or is labeled by exactly one label: $z(u) + \sum_{i=1}^{k+1} x(u, i) = 1$ for $u \in V$.
- No terminal s_i or it's neighbors a_i, b_i are cut: $z(u) = 0$ for $u \in \cup_{i=1}^{k}\{s_i, a_i, b_i\}$.
- Terminal is labeled by it's own index, $x(s_i, i) = 1$ for $1 \leq i \leq k$. Root r is labeled $k+1$, $x(r, k+1) = 1$.
- For each terminal s_i, exactly one of its neighbors is labeled i: $x(a_i, i) + x(b_i, i) = 1$ for $1 \leq i \leq k$.
- For each non-special edge $e = uv$, either one of u, v is cut or they have the same label. This is captured by constraints: $x(u, i) + z(u) \geq x(v, i)$ and $x(v, i) + z(v) \geq x(u, i)$ for all $1 \leq i \leq k + 1$. It is important to note that this constraint holds even for the non-special edges incident on $u \in \cup_{j=1}^{k}\{a_j, b_j\}$.
- If $\mathbb{C}$ is the set of interesting cycles, then for every $C \in \mathbb{C}$, $\sum_{u \in C} z(u) \geq 1$.

The objective is to minimize $\sum_{u \in V} w(u)z(u)$. The full description of the LP relaxation is given in Fig 5.

3.2 Rounding scheme and analysis

Theorem 3.1. *There is a polynomial-time algorithm that given a feasible solution $\boldsymbol{x}, \boldsymbol{z}$ to* Subset-FVS-Rel *outputs a feasible integral solution of weight at most $13 \sum_{u \in V} w(u)z(u)$.*

The rounding scheme is similar to the case of Subset-FES but we remove nodes instead. As before, given a solution $\mathbf{x}, \mathbf{z}$, we round the solution in two steps. In first step, we find a subset $V' \subset V$ of nodes such that removing V' yields a graph $G' = G - V'$ such that interesting cycles in each connected component of G' have same signature. Each component behaves like an instance with single cycle and can be solved optimally. Final output of the algorithm is $V' \cup V''$ where V'' is union of the optimal solution over connected components of G'. Following lemma formalizes the first step of the algorithm:

Lemma 3.1. *Given a feasible solution $\boldsymbol{x}, \boldsymbol{z}$ to* Subset-FVS-Rel*, there is an efficient algorithm to find a subset of nodes $V' \subset V$ with cost at most $12 \sum_{u \in V} w(u)z(u)$ such that any two interesting cycles in the same connected component in $G' = G[V \setminus V']$ have same signature.*

Algorithm 2 shows a simple randomized procedure to find V' which achieves the properties claimed by the preceding lemma. Here, $N(B_i)$ denotes the node boundary of set B_i formally defined as $\{v \mid v \notin B_i, \exists u \in$

$$
\begin{array}{llcl r}
 & \textbf{Subset-FVS-Rel} & & & \\
\min & \displaystyle\sum_{u \in V} w(u) z(u) & & & \\
 & s.t & & & \\
 & \displaystyle z(u) + \sum_{i=1}^{k+1} x(u,i) & = & 1 & u \in V \\
 & x(s_i, i) & = & 1 & 1 \le i \le k \\
 & x(r, k+1) & = & 1 & \\
 & x(a_i, i) + x(b_i, i) & = & 1 & 1 \le i \le k \\
 & z(u) + x(u,i) - x(v,i) & \ge & 0 & 1 \le i \le k+1, uv \in E \setminus E_s \\
 & z(u) & = & 0 & u \in \cup_{i=1}^{k} \{s_i, a_i, b_i\} \\
 & \displaystyle\sum_{u \in C} z(u) & \ge & 1 & C \in \mathbb{C} \\
 & x(u,i) & \ge & 0 & u \in V, 1 \le i \le k+1 \\
 & z(u) & \ge & 0 & u \in V
\end{array}
$$

Figure 5: LP Relaxation for Subset-FVS

Algorithm 2 Initial Cut for Subset-FVS

Given a feasible solution $\mathbf{x}, \mathbf{z}$ to Subset-FVS-Rel
Pick $\theta \in (1/3, 1/2)$ uniformly at random.
For $1 \le i \le k, B_i := \{u \mid x(u,i) > \theta\}$
$V' := \cup_{i=1}^{k} (N(B_i) \setminus \{a_i, b_i\}) \setminus \cup_{i=1}^{k} \{s_i\}$
Return V'

$B_i, uv \in E\}$. And, $\delta(B_i)$ denote the edge boundary formally defined as $\{(u,v) \mid |\{u,v\} \cap B_i| = 1\}$.

Note that we cannot remove the terminals and their neighbors since they have infinite weight. It is not obvious that V', as defined by the algorithm, does not contain a neighbor of a terminal.

Lemma 3.2. *For $1 \le i, j \le k$, $N(B_i) \cap \{a_j, b_j\} = \emptyset$ if $j \ne i$. Thus V' does not contain a neighbor of any terminal.*

Proof: If $u \in \{s_j, a_j, b_j\}$ for some j, then $z(u) = 0$ and we prove that $u \notin V'$. Easy to note from Algorithm 2 that $s_j \notin V'$. And $u \in \{a_j, b_j\}$ is in V' iff $u \in N(B_i)$ for some $i \ne j$. We prove that this is not possible. Consider $uv \in E$ where $v \ne s_j$. From LP constraint $z(u) + x(u,i) \ge x(v,i)$, and since $z(u) = 0$ we have $x(u,i) \ge x(v,i)$. Also, $x(s_j, i) = 0$ for $i \ne j$. Thus, if for some $v \in N(u)$, $v \in B_i$ for $i \ne j$ then, $u \in B_i$. Equivalently, $u \notin N(B_i)$ for $i \ne j$. □

Proof of Lemma 3.1 consists of two parts. The first is to bound the expected cost of the nodes that are cut which is provided in the lemma below. The proof of this lemma is not as straight forward as the one for the case of edges in Subset-FES-Rel.

Lemma 3.3. $\Pr[u \in V'] \le 12 z(u)$*, and hence the expected cost of V' is at most $12 \sum_u w(u) z(u)$.*

Proof: From Lemma 3.2 no terminals or neighbors of terminals are in V'. Consider some other node u. Then u is in V' iff $u \in N(B_i)$ for some $1 \le i \le k$. And $u \in N(B_i)$ iff $x(u,i) \le \theta$ and there exists $v \in N(u)$ such that $x(v,i) > \theta$. Equivalently, $u \in N(B_i)$ iff $\theta \in [x(u,i), \max_{v \in N(u)} x(v,i))$. We will denote the interval $(1/3, 1/2) \cap [x(u,i), \max_{v \in N(u)} x(v,i))$ as $I_i(u)$. Thus, $\Pr[u \in V'] = \Pr[\theta \in \cup_{i=1}^{k} I_i(u)]$. If $\cup_{i=1}^{k} I_i(u)$ has length at most $2z(u)$ then, since θ is chosen uniformly at random from the range $(1/3, 1/2)$ of length $1/6$, $\Pr[u \in V'] \le 12 z(u)$. Next, we prove this fact. From LP constraint $z(u) + x(u,i) - x(v,i) \ge 0$, we conclude that $\max_{v \in N(u)} x(v,i) \le z(u) + x(u,i)$.

If there is only one index i_1 such that $x(u, i_1) > 1/3$, then for $j \ne i_1$, $I_j(u) \subset (1/3, 1/3 + z(u))$. This implies that $\cup_{i=1}^{k} I_i(u) \subset [x(u,i_1), x(u,i_1) + z(u)) \cup (1/3, 1/3 + z(u))$. This has length at most $2z(u)$. If on the other hand there are two indices i_1, i_2 such that $x(u,i_1), x(u,i_2) > 1/3$ then, by LP constraint $z(u) + \sum_{i=1}^{k+1} x(u,i) = 1$, we get $x(u,j) < 1/3 - z(u)$ for $j \ne i_1, i_2$. This implies that $I_j(u) = \emptyset$ for $j \ne i_1, i_2$ and $\cup_{i=1}^{k} I_i(u) \subset (x(u,i_1), x(u,i_1) + z(u)) \cup (x(u,i_2), x(u,i_2) + z(u))$. This range also has length at most $2z(u)$. Thus, if $u \in V \setminus \cup_{i=1}^{k} \{s_i, a_i, b_i\}$, $\Pr[u \in V'] \le 12 z(u)$. □

The second part of the proof is to show the property of the signatures in the graph $G' = G - V'$. To prove this we rely on the analysis that we did in the setting of SUBSET-FES. Let $E' = (\cup_{i=1}^{k}\delta(B_i)) \setminus E_s$. We will prove that $E' \subseteq \cup_{u \in V'}\delta(u)$. Next we will prove, via the analysis from Section 2, that in $G - E'$ each connected component has interesting cycles with the same signature; Remark 2.12 is relevant here. This will prove that $G - V'$ has the desired property.

Proof of Lemma 3.1: Given feasible solution $\mathbf{x}, \mathbf{z}$ to SUBSET-FVS-REL we define a new set of assignment values $\tilde{x}(u, i)$ as follows: $\tilde{x}(u, i) = x(u, i)$ for $u \in V$ and $1 \le i \le k$ and $\tilde{x}(u, k+1) = z(u) + x(u, k+1)$. We have for each $u \in V$, $z(u) + \sum_{i=1}^{k+1} x(u, i) = 1$ and hence $\sum_{i=1}^{k+1} \tilde{x}(u, i) = 1$. Now consider Algorithm 1 for SUBSET-FES with input $\tilde{\mathbf{x}}$. Let the edges set returned be E'. We observe that the algorithm only uses labels 1 to k and since $\mathbf{x}$ and $\tilde{\mathbf{x}}$ are identical on these labels, for each i and θ, Algorithm 1 and Algorithm 2 produce the same sets $B_1, \ldots, B_k$. For this reason we can use the analysis of Algorithm 1 about the structure of the graph $G - E'$ (Lemma 2.3). In particular we have the property that in every connected component of $G - E'$ all interesting cycles have the same signature. We also observe that this property remains true for $G - \tilde{E}$ if $\tilde{E}$ is a superset of E'. To finish the proof of Lemma 3.1 we now prove that $E' \subseteq \cup_{u \in V'}\delta(u)$, that is, removing V' removes every edge in E' and perhaps more.

CLAIM 3.4. $E' \subseteq \cup_{u \in V'}\delta(u)$.

Proof: Recall that $E' = (\cup_{i=1}^{k}\delta(B_i)) \setminus E_s$. Consider an edge $uv \in E'$. Implies that there is an i such that $v \in B_i$ and $u \notin B_i$, and also that uv is non-special edge. Note that if u is not a terminal or neighbor of a terminal then $u \in N(B_i)$ and therefore $u \in V'$ and hence $uv \in \delta(u)$. We can assume that u is not a terminal for then uv is a special edge. Thus the only case left to consider is that u is a neighbor of terminal, say $u = a_j$ for some j, and v is not a terminal, otherwise uv is again a special edge. Since $z(u) = 0$ and uv is not a special edge, from the LP constraint $z(u) + x(u, i) \ge x(v, i)$, we have that $x(u, i) \ge x(v, i)$ which implies that $u \in B_i$ if $v \in B_i$ contradicting the fact that $u \in N(B_i)$. Thus $uv \in E'$ implies that $uv \in \cup_{u \in V'}\delta(u)$. □

This finishes the proof of the lemma. □

3.3 Second step of rounding

Here, we will describe the second step of rounding for SUBSET-FVS-REL and finish the proof of Theorem 3.1. The second step of the algorithm is to process the graph $G' = G - V'$ which has very restricted structure. Consider a connected component H of G' which has an interesting cycle C; without loss of generality s_1 is a terminal on C. Since all signatures are identical in H, disconnecting a_1 from b_1 in the graph $H' = H - s_1$ is necessary and sufficient to remove all interesting cycles in H. Since we are in the node-weighted setting we need to find a minimum weight node cut between a_1 and b_1 in H'. In analogy with Claim 2.13 we have the following.

CLAIM 3.5. *The cost of the minimum node-weighted cut between a_1 and b_1 in H' is at most $\sum_{u \in V(H)} w(u)z(u)$.*

Let V'' be the union of all the minimum node cuts found in each connected component of G'. It is easy to see that $V' \cup V''$ is a feasible solution. From the preceding claim and the fact that the connected components of G' are node disjoint we have that $w(V'') \le \sum_{u \in V} w(u)z(u)$. We now finish the proof of Theorem 3.1.

Proof of Theorem 3.1: Let $\alpha = \sum_{u \in V} w(u)z(u)$ be the the objective function value of a feasible solution $\mathbf{x}, \mathbf{z}$ to SUBSET-FVS-REL. From Lemma 3.1 we obtain a set of nodes V' such that $w(V') \le 12\alpha$ and $G' = G - V'$ satisfies the property needed for the second step of the algorithm. The set of nodes V'' found in the second step satisfy the property that $w(V'') \le \alpha$. Thus $w(V' \cup V'') \le 13\alpha$ and $V' \cup V''$ is a feasible integral solution to the given instance. This finishes the proof. □

Concluding remarks: Our work opens up the possibility of obtaining improved approximations for SUBSET-FES and SUBSET-FVS. For both problems the worst-case integrality gap we know comes from corresponding gaps for EDGE-WT-MC and NODE-WT-MC respectively. However we have not tried hard to find bad examples. It will be interesting to examine the easy cases of FVS and SUBSET-FES first before considering SUBSET-FVS. Another direction that we have briefly explored is to consider SUBSET-FES and SUBSET-FVS when the number of terminals k is small. We believe that we can obtain an algorithm for SUBSET-FVS that runs in time $\exp(k, 1/\varepsilon) \cdot \mathrm{poly}(n)$ and yields a $(4+\varepsilon)$-approximation. Note that even when $k = 1$ the problem is APX-Hard.

References

[1] Vineet Bafna, Piotr Berman, and Toshihiro Fujito. A 2-approximation algorithm for the undirected feedback vertex set problem. *SIAM Journal on Discrete Mathematics*, 12(3):289–297, 1999.

[2] Reuven Bar-Yehuda, Dan Geiger, Joseph Naor, and Ron M Roth. Approximation algorithms for the feedback vertex set problem with applications to constraint satisfaction and bayesian inference. *SIAM journal on computing*, 27(4):942–959, 1998.

[3] Ann Becker and Dan Geiger. Optimization of pearl's method of conditioning and greedy-like approximation algorithms for the vertex feedback set problem. *Artificial Intelligence*, 83(1):167–188, 1996.

[4] Gruia Calinescu, Howard Karloff, and Yuval Rabani. Approximation algorithms for the 0-extension problem. *SIAM Journal on Computing*, 34(2):358–372, 2005.

[5] Gruia Calinescu, Howard J. Karloff, and Yuval Rabani. An improved approximation algorithm for multiway cut. *Journal of Computer and System Sciences*, 60(3):564–574, 2000. Preliminary version in STOC 1998.

[6] Chandra Chekuri and Alina Ene. Approximation algorithms for submodular multiway partition. In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 807–816. IEEE, 2011.

[7] Chandra Chekuri and Alina Ene. Submodular cost allocation problem and applications. In *International Colloquium on Automata, Languages and Programming (ICALP (1))*, pages 354–366, 2011. A longer version is available on the arXiv, abs/1105.2040.

[8] Fabián Chudak, Michel X Goemans, Dorit S Hochbaum, and David P Williamson. A primal–dual interpretation of two 2-approximation algorithms for the feedback vertex set problem in undirected graphs. *Operations Research Letters*, 22(4):111–118, 1998.

[9] Marek Cygan, Fedor Fomin, Łukasz Kowalik, Daniel Loksthanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2014. In print.

[10] Reinhard Diestel. *Graph Theory: Springer Graduate Text GTM 173*, volume 173. Reinhard Diestel, 2012.

[11] P. Erdos and L. Pósa. On independent circuits contained in a graph. *Canadian Journal of Mathematics. Journal Canadien de Mathématiques*, 17:347–352, 1965.

[12] Guy Even, Joseph Naor, Baruch Schieber, and Leonid Zosin. Approximating minimum subset feedback sets in undirected graphs with applications. *SIAM Journal on Discrete Mathematics*, 13(2):255–267, 2000.

[13] Guy Even, Joseph Naor, and Leonid Zosin. An 8-approximation algorithm for the subset feedback vertex set problem. *SIAM Journal on Computing*, 30(4):1231–1252, 2000. Preliminary version in *Proc. of IEEE FOCS*, 1996.

[14] Guy Even, Joseph Seffi Naor, Satish Rao, and Baruch Schieber. Divide-and-conquer approximation algorithms via spreading metrics. *Journal of ACM(JACM)*, 47:585–616, 2000.

[15] Paola Festa, Panos M Pardalos, and Mauricio GC Resende. Feedback set problems. In *Handbook of combinatorial optimization*, pages 209–258. Springer, 1999.

[16] Michel X Goemans and David P Williamson. Primal-dual approximation algorithms for feedback problems in planar graphs. *Combinatorica*, 18(1):37–59, 1998.

[17] Venkatesan Guruswami, Rajsekar Manokaran, and Prasad Raghavendra. Beating the random ordering is hard: Inapproximability of maximum acyclic subgraph. In *Foundations of Computer Science, 2008. FOCS'08. IEEE 49th Annual IEEE Symposium on*, pages 573–582. IEEE, 2008.

[18] 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.

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

[20] Bruce Reed, Neil Robertson, Paul Seymour, and Robin Thomas. Packing directed circuits. *Combinatorica*, 16(4):535–554, 1996.

[21] Paul D. Seymour. Packing directed circuits fractionally. *Combinatorica*, 15(2):281–288, 1995.

[22] Ola Svensson. Hardness of vertex deletion and project scheduling. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 301–312. Springer, 2012.

[23] Vijay V Vazirani. *Approximation algorithms*. Springer Science & Business Media, 2013.

[24] Douglas West. *Introduction to graph theory*, volume 2. Prentice hall Upper Saddle River, 2001.

A Further Remarks on the LP Relaxations

Cycle length constraint: Subset-FES-Rel with only labeling constraints does not have a bounded integrality gap. It is essential to add the constraint that each interesting cycle has length at least 1: for each interesting cycle $C \in \mathbb{C}, \sum_{e \in C} z(e) \geq 1$. Consider the third graph in Figure 6. Without this constraint, Subset-FES-Rel has cost 0 with $x(a_i, i) = x(b_{i-1}, i) = 1$ for $i \in \{1, 2, 3\}$, $x(u, 1) = 1, x(w, 3) = 1$, and $z(e) = 0$ for all $e \in E$, whereas the optimal solution for Subset-FES has cost 1.

However, we note that the cycle length constraint is only useful in the second step which can be solved via simple min-cut computations. Alternatively, we can think of the labeling approach as strengthening the naive cycle constraint based LP which has $\Omega(\log n)$ integrality gap.

Simplifying Reductions: Given a Subset-FES instance we first simplified it via reductions so that we can assume that each terminal is a degree 2 node with incident edges having infinite weight. Also, that no two terminals are connected by an edge or share a neighbor. A natural question here is whether these reductions are for convenience or whether they are essential in enabling the formulation.

Assuming that all edges incident on terminals are

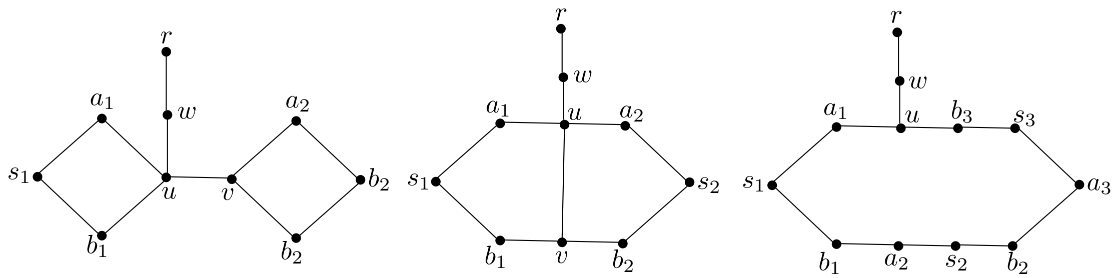

Figure 6: Edges incident on u, v have weight 1 except wu which has weight 0. Other edges have infinite weight.

of infinite weight, if a terminal has degree more than 2 then we change the spreading constraint to the following constraint: $\sum_{u \in N(s_i)} x(u, i) = |N(s_i)| - 1$. If two terminals s_i, s_j are adjacent, then we can replace the above constraint for s_i, s_j to the following constraint: $\sum_{u \in N(s_i) \setminus \{s_j\}} x(u, i) + \sum_{u \in N(s_j) \setminus \{s_i\}} x(u, j) = |N(s_i)| + |N(s_j)| - 3$. If an edge incident on a terminal does not have infinite weight, we can conceptually split the edge and then write constraints based on the virtual node inserted. Doing these reductions does not change the integrality gap of the LP but simplifies the analysis considerably. Thus, for the most part the simplifying reductions are for convenience.

However, we point out the the reduction that allows us to assume that each interesting cycle has at least two terminals, has a more direct impact in terms of our analysis. First, without this assumption, Lemma 2.3 does not hold and Lemma 2.7 does not hold true for cycle with one terminal. There may be two distinct interesting cycles in G', each containing one terminal and connected by a path. We may also have a connected component in G' containing two terminals with an interesting cycle containing both the terminals while another cycle containing just one terminal. Consider the first two graphs in Figure 6 with the following feasible solution: $x(a_1, 1) = 1/2, x(a_1, 2) = 1/2, x(w, 3) = 1$ and b_1, u, v, a_2, b_2 have the same assignment x as a_1. For any choice of $\theta \in (1/3, 1/2)$ we will get $L(a_1) = L(b_1) = L(u) = L(v) = L(a_2) = L(b_2) = \{1, 2\}, L(w) = \{3\}$ and $E' = \{wu\}$.

We can modify the algorithm to incorporate these special cases, however the analysis becomes complicated. Also, for the case when cycles in a connected component have different signatures as in the case of second graph of Figure 6, the second step of the rounding will loose a factor of 2 giving us an upper bound of 14 on integrality gap of Subset-FES-Rel. One can write additional constraints to avoid this but it is simpler for our analysis to make the assumption.