A.2 Label-Cover Hardness for Stochastic Steiner Tree with Correlated Costs

Recall that our results hold only when the edge costs on Tuesday are independent of the set of materialized demands S. We now prove that relaxing this condition makes the problem much harder to approximate.

Let us state the general problem (which we call the General Stochastic Steiner Tree problem) formally. We are given a graph G with Monday's cost function $c_M : E \to \mathcal{R}_+$, and a probability distribution π on subset of demands. Furthermore, for each subset S of demands, we are given edge costs $c_{T,S} : E \to \mathcal{R}_+$, which is the cost on Tuesday if subset S materializes. A solution now buys a subset of edges on Monday E_M, and if subset S materializes on Tuesday, it buys a set of edges $E_T(S)$ so that $E_M \cup E_T(S)$ connects S to the root r. The goal is to minimize

$$c_M(E_M) + \mathbf{E}_{S \leftarrow \pi}[\, c_{T,S}(\, E_T(S)\,)\,] \,. \tag{A.6}$$

Theorem A.2 *The two-stage General Stochastic Steiner Tree problem when Tuesday's costs are dependent on the materialized demand is at least* $\Omega(2^{\log^{1-\varepsilon} n})$ *hard for every fixed* $\varepsilon > 0$.

We will show hardness for the special case of the problem where the number of subsets S in the support of the distribution π is polynomially bounded.

A.2.1 The Reduction

We will reduce the minimization version of the Label Cover problem to the General Stochastic Steiner Tree problem: recall that the Label Cover problem has as input a bipartite graph $G = (V_1 \cup V_2, E)$, an integer K, and a partial function $f_e : [K] \to [K]$ for each edge $e \in E$. A solution to this problem assigns a subset $L(x)$ of labels to each vertex $x \in V_1 \cup V_2$, so that for each edge $e = (u, v) \in V_1 \times V_2$, there is at least one label $\ell \in L(u)$ such that $f_e(\ell) \in L(v)$. The goal is to minimize $\sum_{u \in V_1} |L(u)| + \sum_{v \in V_2} |L(v)|$ [AL96]. (Note that this version of Label Cover is special because we have a partial *function* $f_e : [K] \to [K]$ for each edge, instead of having an arbitrary relation $R_e \subseteq [K] \times [K]$; while stronger hardness results can be proved for the latter version, it does not have all the properties we need for our reduction.)

Our reduction follows the usual path when reducing from Label Cover: we build an undirected graph on the vertex set

$$A = V_1 \cup V_2 \cup (V_1 \times [K]) \cup (V_2 \times [K]) \cup \{r\}.$$

We now add several edges: the first are the "labeling edges" from each $x \in V_i$ (with $i = 1, 2$) to all $(x, \ell) \in V_i \times [K]$ having unit cost on Monday: buying an edge $(u, (u, \ell))$ can be interpreted as adding the label ℓ to $L(u)$. Moreover, there are "consistency edges" from (u, ℓ) to $(v, f_{uv}(\ell))$ for all u, v, ℓ, and "root edges" from the root r to all $u \in V_1$; both these kinds of edges have cost M on Monday, where $M > K(|V_1| + |V_2|)$, and hence no reasonable solution would ever buy these edges on Monday.

Let us now specify the scenarios: for each edge $e = (u, v) \in E$ in the Label Cover instance, we define a scenario where we want to connect the root r to the vertex v. In this scenario, the Tuesday costs of the root edge (r, u), and of the consistency edges between $(u, *)$ and $(v, *)$ are set to 0; all other edges are set to have cost M.

Note that for r to be connected to v in this scenario with cost less than M, the first edge on the path must be $r \to u$, and then from $u \to (u, \ell)$. However, since the only consistency edges with cost $< M$ until now have been between $(u, *)$ and $(v, *)$, we cannot use edges to $(v', *)$ for $v' \neq v$:

the next edge *must be* $(u, \ell) \to (v, f_{uv}(\ell))$. And then the last hop on the $r \to v$ path must be from $(v, f_{uv}(\ell)) \to v$. Moreover, taking this entire path is possible if and only if we bought the edges $(u, (u, \ell))$ and $(v, (v, f_{uv}(\ell)))$—or in other words, we picked consistent labels for u and v—on Monday.

A.2.2 Closing Note: Single-source Single-Sink

Note that we could extend the proof (by creating a new vertex r' and attaching each $v \in V_2$ to r', whose Monday costs would be M, and in scenario corresponding to (u, v), its Tuesday costs would be 0) to show the hardness holds even if the vertices to be connected on Tuesday (i.e., r and r' in this case) are specified in advance, and the randomness is only in the costs.