Given that the expected facility cost is at most $\sum_{i \in F} f_i y_i^*$, we get that the overall expected cost is at most

$$\sum_{i \in F} f_i y_i^* + \sum_{j \in D} C_j^* + \frac{2}{e} \sum_{j \in D} v_j^* \leq \text{OPT} + \frac{2}{e} \text{OPT}.$$

This yields the following theorem.

Theorem 12.5. *Algorithm 12.1, modified to choose* j_k *to minimize* $v_j^* + C_j^*$, *is a* $(1 + \frac{2}{e})$*-approximation algorithm for the uncapacitated facility location problem, where* $1 + \frac{2}{e} \approx 1.736$.

12.2 The Single-Source Rent-or-Buy Problem

In this section, we consider the *single-source rent-or-buy problem*. The input for the problem is an undirected graph $G = (V, E)$ with edge costs $c_e \geq 0$ for all $e \in E$, a root vertex $r \in V$, a set of terminals $X \subseteq V$, and a parameter $M > 1$. We need to design a network connecting all terminals to the root; for each terminal we specify a path of edges from the terminal to the root. We say that a terminal *uses* an edge if the edge is on the terminal's path to the root. To build the paths, we can both buy and rent edges. We can buy edges at cost Mc_e, and once bought, any terminal can use the edge. We can also rent edges at cost c_e, but then we need to pay the rental cost for each terminal using the edge. The goal is to find a feasible network that minimizes the total cost. We can formalize this by letting $B \subseteq E$ be the set of edges that are bought, and letting R_t be the set of edges that are rented by terminal $t \in X$. Then for each $t \in X$, the set of edges $B \cup R_t$ must contain a path from t to the root r. Let $c(F) = \sum_{e \in F} c_e$ for any $F \subseteq E$. Then the total cost of the solution is $Mc(B) + \sum_{t \in X} c(R_t)$. We must find edges B to buy and R_t to rent that minimize this overall cost.

We will give a randomized approximation algorithm for the problem that cleverly trades off the cost of buying versus renting. The *sample-and-augment algorithm* draws a sample of terminals by marking each terminal t with probability $1/M$ independently. Let D be the random set of marked terminals. We then find a Steiner tree T on the set of terminals D plus the root, and buy the edges of T. To find a Steiner tree, we use the 2-approximation algorithm of Exercise 2.5 that computes a minimum spanning tree on the metric completion of the graph. We then augment the solution by renting paths from the unmarked terminals to the tree T. To do this, we find the shortest path from each unmarked t to the closest vertex in T, and rent these edges.

The analysis of the sample-and-augment algorithm begins by observing that the expected cost of buying the edges in the tree T is at most twice the cost of an optimal solution to the rent-or-buy problem.

Lemma 12.6.

$$E[Mc(T)] \leq 2\,\text{OPT}.$$

Proof. To prove the lemma, we demonstrate a Steiner tree T^* on the set of marked terminals such that the expected cost of buying the edges of T^* is at most OPT. Since we are using a 2-approximation algorithm to find T, the lemma statement then follows

We consider an optimal solution to the problem: let B^* be the set of bought edges, and let R_t^* be the edges rented by terminal t. Consider the edges from B^* together with the union of edges of R_t^* over the marked terminals t. Note that this set of edges certainly contains some Steiner tree T^* on the set of marked terminals plus the root. We now want to analyze the cost of buying this set of edges. The essential idea of the analysis is that although we now have to pay M times the cost of the rented edges in each R_t^* for marked t, since we marked t with probability $1/M$, the expected cost of these edges will be the same as the renting cost of the optimal solution. To see this formally, if D is the random set of marked terminals, then

$$\begin{aligned} E[Mc(T^*)] &\leq Mc(B^*) + E[M \sum_{t \in D} c(R_t^*)] \\ &= Mc(B^*) + M \sum_{t \in X} c(R_t^*) \Pr[t \in D] \\ &= Mc(B^*) + \sum_{t \in X} c(R_t^*) \\ &= \mathrm{OPT}. \end{aligned}$$

□

To complete the analysis, we show that the expected renting cost is no more than the expected buying cost.

Lemma 12.7.

$$E\left[\sum_{t \in X} c(R_t)\right] \leq E[Mc(T)].$$

Proof. To prove this, let us be a bit more precise about the algorithm; then we will alter the algorithm to give an equivalent algorithm, and prove the statement for the equivalent algorithm.

Let D be the (random) set of marked terminals. We run Prim's minimum spanning tree algorithm on the metric completion of the graph, starting with the root r (see Section 2.4 for a discussion of Prim's algorithm). Prim's algorithm maintains a set $S \subseteq D \cup \{r\}$ of vertices in the spanning tree, and chooses the cheapest edge e that has one endpoint in S and the other in $D - S$ to add next to the spanning tree; the endpoint of e in $D - S$ is then added to S.

We now alter the algorithm by not choosing D in advance. Rather, we choose the cheapest edge e with one endpoint in S and the other from the set of all terminals whose marking status has not been determined. Let t be the endpoint of e whose marking status is not determined. At this point, we decide, with probability $1/M$, whether to mark t. If t is marked, then we add t to D and to S, and add the edge e to the tree. If t is not marked, then we do not add t to D or S, and edge e is not added to the tree. Note that we get the same tree T on D via this process as in the case that the random set D is drawn before running the algorithm.

We let β_t for a terminal t be a random variable associated with the cost of connecting t to the tree via bought edges; we call it the buying cost. We let β_t of a marked terminal t be M times the cost of the edge that first connects it to the tree, and we let β_t be 0 if t is not marked. In our modified algorithm above, β_t is the cost of the edge that connects t to S when t is marked. The total cost of the tree is then the sum of the buying costs of all the marked terminals in the tree, so that $\sum_{t \in D} \beta_t = Mc(T)$. In a similar way, we let ρ_t for a terminal t be a random variable giving the cost of renting edges to connect t to the tree.

Now consider a given terminal t at the time we decide whether to mark t or not. Let S be the set of vertices already selected by Prim's algorithm at this point in time, and let e be the edge chosen by Prim's algorithm with t as one endpoint and with the other endpoint in S. If we mark t, we buy edge e at cost Mc_e. If we do not mark t, then we could rent edge e at cost c_e, and since all the vertices in S are marked, this will connect t to the root; thus, $\rho_t \le c_e$. Hence, the expected buying cost of t is $E[\beta_t] = \frac{1}{M} \cdot Mc_e = c_e$, whereas its expected cost of renting edges to connect t to the root is $E[\rho_t] \le (1 - \frac{1}{M}) \cdot c_e \le c_e$. Observe that this is true no matter what point in the algorithm t is considered. Thus,

$$E\left[\sum_{t \in X} c(R_t)\right] = E\left[\sum_{t \in X} \rho_t\right] \le E\left[\sum_{t \in X} \beta_t\right] = E\left[\sum_{t \in D} \beta_t\right] = E[Mc(T)]. \quad \square$$

The following theorem is then immediate.

Theorem 12.8. *The sample-and-augment algorithm described above is a randomized 4-approximation algorithm for the single-source rent-or-buy problem.*

Proof. For the solution $B = T$ and R_t computed by the randomized algorithm, we have that

$$E\left[Mc(T) + \sum_{t \in X} c(R_t)\right] \le 2 \cdot E[Mc(T)] \le 4\,\mathrm{OPT}. \quad \square$$

In Exercise 12.2, we consider the multicommodity rent-or-buy problem, which is an extension of the single-source rent-or-buy problem to multiple source-sink pairs. We see that the sample-and-augment algorithm also leads to a good approximation algorithm for this problem.

12.3 The Steiner Tree Problem

The Steiner tree problem provides an excellent example of a problem for which our understanding of its combinatorial structure has worked hand in hand with the design and analysis of a linear programming-based approach to approximation algorithm design. Furthermore, we will combine two techniques already discussed for rounding LP solutions, by relying on an iterative use of randomized rounding. The *Steiner tree problem*, as discussed in Exercise 2.5, is as follows: Given an undirected graph $G = (V, E)$, and a subset of nodes $R \subseteq V$, along with a nonnegative edge cost $c_e \ge 0$ for each edge $e \in E$, find a minimum-cost subset of edges $F \subseteq E$, such that $G = (V, F)$ contains a path between each pair of nodes in R. As discussed in that exercise, by considering the metric completion of the graph, we may assume without loss of generality that the input graph G is complete, and that the costs satisfy the triangle inequality. In Section 7.4, we showed that the primal-dual method, based on a relati