Greedy Algorithms II

Correctness of Kruskal's Algorithm cont'd

To show that the spanning tree resulting from the algorithm is in fact a minimum cost spanning tree, we reason that after each execution of the loop, the set T of edges can be expanded to a minimum cost spanning tree using edges that have not yet been considered. Hence after termination, all edges have been considered, so T must itself be a minimum cost spanning tree. We can formalize this reasoning as follows:

Definition 1 A set T of edges of G is *promising* if T can be expanded to a minimum cost spanning tree for G.

Lemma 1 The following is a loop invariant for Kruskal's algorithm: *T is promising.*

PROOF: The proof is by induction on the number of iterations of the main loop of Kruskal's Algorithm. **Basis Case:** (at this stage the algorithm has gone through the loop zero (0) times): Initially T is the empty set, which is obviously promising.

Induction Step: We assume that T is promising, and show that T continues to be promising after one more execution of the loop.

Notice that the edges used to expand T to a spanning tree must come from those not yet considered, because the edges that have been considered are either in T already, or have been rejected because they form a cycle. We examine by cases what happens after edge e_i has been considered:

case(I): e_i is rejected. T remains unchanged, and so it is still promising.

case(II): e_i is accepted. We must show that $T \cup \{e_i\}$ is still promising. Since T is promising, there is a minimum cost spanning tree T_1 such that $T \subseteq T_1$.

subcase(a): $e_i \in T_1$. Then obviously $T \cup \{e_i\}$ is promising.

subcase(b): $e_i \notin T_1$. Then, according to the Exchange Lemma below, there is an edge e_j in $T_1 - T_2$ (where T_2 is the spanning tree resulting from the algorithm) such that $T_3 = (T_1 \cup \{e_i\}) - \{e_j\}$ is a spanning tree. Notice that $i < j$, since otherwise e_j would have been rejected from T and thus would form a cycle in T and so also in T_1. Therefore $c(e_i) \leq c(e_j)$, so $c(T_3) \leq c(T_1)$, so T_3 must also be a minimum cost spanning tree. Since $T \cup \{e_i\} \subseteq T_3$, it follows that $T \cup \{e_i\}$ is promising.

This finishes the proof of the Induction Step. □

Lemma 2 (Exchange Lemma) Let G be a connected graph, and let T_1 and T_2 be any two spanning trees for G. For every edge e in $T_2 - T_1$ there is an edge e' in $T_1 - T_2$ such that $T_1 \cup \{e\} - \{e'\}$ is a spanning tree for G.

Proof: Exercise (Hint: let e be an edge in $T_1 - T_2$. Then $T_2 \cup \{e\}$ contains a cycle – can all the edges in this cycle belong to T_1?).

A greedy algorithm for job scheduling with deadlines profits

We have n jobs, each of which takes unit time, and a processor on which we would like to schedule them in as profitable a manner as possible. Each job has a profit associated with it, as well as a deadline; if the job is not scheduled by its deadline, then we don't get its profit. Because each job takes the same amount of time, we will think of a schedule S as consisting of a sequence of job "slots" $1, 2, 3, \ldots$ where $S(t)$ is the job scheduled in slot t.

Formally, the input to the problem is a sequence of pairs $(d_1, g_1), (d_2, g_2), \ldots, (d_n, g_n)$ where g_i is a non-negative real number representing the profit obtainable from job i, and $d_i \in \mathbb{N}$ is the deadline for job i.

Definition 2 A *schedule* is an array $S(1), S(2), \ldots, S(d)$ where $d = \max d_i$ (i.e. the latest deadline, beyond which no jobs can be scheduled), such that if $S(t) = i$, then job i is scheduled at time t, $1 \leq t \leq d$. If $S(t) = 0$, then no job is scheduled at time t.

Definition 3 A schedule S is *feasible* if (a) If $S(t) = i > 0$, then $t \leq d_i$ (i.e. every scheduled job meets its deadline), and (b) If $t_1 \neq t_2$ and $S(t_1) \neq 0$, then $S(t_1) \neq S(t_2)$ (i.e. each job is scheduled at most once).

Definition 4 Let the total profit of schedule S be $P(S) = g_{S(1)} + g_{S(2)} + \ldots + g_{S(d)}$, where $g_0 = 0$.

We want to find a feasible schedule S whose profit $P(S)$ is as large as possible. This can be accomplished with the following greedy algorithm:

Sort the jobs in non-increasing order of profits: $g_1 \geq g_2 \geq \ldots \geq g_n$
$d \longleftarrow \max_i d_i$
for $t : 1..d$
    $S(t) \longleftarrow 0$
end for
for $i : 1..n$
    Find the largest t such that ($S(t) = 0$ and $t \leq d_i$), and let $S(t) \longleftarrow i$
    [find the latest possible free slot meeting the deadline]
end for

Theorem 1 The greedy solution above is optimal (i.e. the profit $P(S)$ of the schedule S computed by this greedy algorithm is as large as possible).

Definition 5 A schedule is *promising* if it can be extended to an optimal schedule. Schedule S' extends schedule S iff for all $1 \leq t \leq d$, if $S(t) \neq 0$, then $S(t) = S'(t)$.

Lemma 3 The following is a loop invariant of the above greedy algorithm: *S is promising.*

Exercise: Why does Lemma 3 imply Theorem 1? (Hint: simple observation).

PROOF: (of Lemma 3) The proof is by induction. **Basis Case:** Initially $S = (0, 0, \ldots, 0)$ is promising (take *any* optimal schedule; it is an extension).

Induction Step: Suppose that S is promising, and let S_{opt} be *some* optimal schedule that extends S. Let S' be the result of one more iteration through the loop where job i is considered. We must prove that S' continues to be promising, so the goal is to show that there is an optimal schedule S'_{opt} that extends S'. We consider two cases:

case (I) job i cannot be scheduled. Then $S' = S$, so we let $S'_{\text{opt}} = S_{\text{opt}}$, and we are done.

case (II) job i is scheduled at time t_0, so $S'(t_0) = i$ (whereas $S(t_0) = 0$) and t_0 is the latest possible time for job i in the schedule S.

subcase (a) job i is scheduled in S_{opt} at time t_1:

If $t_1 = t_0$ then, as in case (I), just let $S'_{\text{opt}} = S_{\text{opt}}$.

If $t_1 < t_0$ then let S'_{opt} be S_{opt} except that we interchange t_0 and t_1, that is: $S'_{\text{opt}}(t_0) = S_{\text{opt}}(t_1) = i$ and $S'_{\text{opt}}(t_1) = S_{\text{opt}}(t_0)$. Then S'_{opt} is feasible (why?), it extends S' (why?), and $P(S'_{\text{opt}}) = P(S_{\text{opt}})$ (why?).

The case $t_1 > t_0$ is not possible (why?).

subcase (b) job i is not scheduled in S_{opt}. Then we simply define S'_{opt} to be the same as S_{opt}, except $S'_{\text{opt}}(t_0) = i$. Since S_{opt} is feasible, so is S'_{opt}, and since S'_{opt} extends S', we only have to show that $P(S'_{\text{opt}}) = P(S_{\text{opt}})$. This follows from the following claim:

Claim: Let $S_{\text{opt}}(t_0) = j$. Then $g_j \leq g_i$.

We prove the claim by contradiction: assume that $g_j > g_i$ (note that in this case $j \neq 0$). Then job j was considered before job i. Since job i was scheduled at time t_0, job j must have been scheduled at time $t_2 \neq t_0$ (we know that job j was scheduled in S since $S(t_0) = 0$, and $t_0 \leq d_j$, so there was a slot for job j, and therefore it was scheduled). But S_{opt} extends S, and $S(t_2) = j \neq S_{\text{opt}}(t_2)$ — contradiction.

This finishes the proof of the Induction Step. □

Exercise: Make sure you can answer all the "why's" in the above proof. Also, where in the proof of the claim you use the fact that $j \neq 0$?