Greedy Algorithms

Suppose we are planning a trip from Kingston to Montreal (and suppose there is only one road so we don't have a choice of which route to take). Our car will need several stops to fill up the gas tank. We know the location of all the service stations on the highway between Kingston and Montreal, and we know exactly how far we can travel on a tank of gas. We want to minimize the number of stops we make along the way.

The intuitive solution, of course, is to go as far as possible on each tank - that is, stop at a service station iff we can't make it to the next one. It turns out that in this case intuition can be trusted - this is exactly the right solution.

But how can we ...

- prove that this algorithm always finds the correct solution?
- determine the computational complexity of this algorithm?

We can formalize this idea into an algorithm as follows. I'll show a recursive formulation first – here we will define a function RT(i) which is supposed to find an optimal solution when we are leaving station s_i with a full tank of gas ... RT(0) will give us the complete solution.

We will assume that the problem has a solution (ie that there is no gap between consecutive stations that exceeds the distance we can travel on one tank of gas ... so when we fill up, we know we can reach the next station). It is easy to test to make sure this condition is satisfied. If it is not, we can immediately conclude there is no solution – we are done.

We find the solution by calling RT(0)

To complete our consideration of this problem, we need to do several things:

1. Show that the problem does have an optimal solution (some problems don't, either because there is no solution or there is no bound on the value of the potential solutions).

- 2. Show that the greedy algorithm finds an optimal solution.
- 3. Determine the complexity of the algorithm.

First, once the stations are sorted we can easily determine if there are any feasible solutions: the trip is possible iff there is no gap between stations too long for us to cover on one tank of gas. Assuming then that there are feasible solutions, it is clear that there is at least one optimal solution since we can rank all the feasible solutions according to their cardinality and choose one with minimum cardinality.

It is also clear that the algorithm will find some solution to the problem: it never lets us run out of gas, and it gets us to the destination.

Let's also dispose of the complexity question: sorting the n stations takes $O(n^*\log n)$ time. Choosing the stations where we stop takes O(n) time since we look at each station exactly once, and decide in constant time whether or not to stop there. Thus the entire algorithm takes $O(n^*\log n)$ time ... or just O(n) if the stations are given to us already sorted.

Now the proof of correctness. This is the interesting part. We use Proof by Induction. This proof is slightly more formal than the presentation in class, but the logic is identical.

Let n be the number of stations between Kingston and Montreal.

Base Case: If n = 0, there are no possible stops between Kingston and Montreal. If there is any feasible solution (ie if we can reach Montreal on a single tank), then the algorithm's action is correct (it makes no stops before Montreal). Thus for n = 1 the algorithm finds an optimal solution.

Inductive Hypothesis

We assume that the algorithm finds an optimal solution when we have a full tank and the number of service stations still ahead of us is $\leq k$, for some $k \geq 0$

Inductive Step

Suppose there are k+1 service stations still ahead of us and we have a full tank when we leave Kingston.

Let the algorithm's solution be $S_A = \{s_i, s_j, \dots, s_{n+1}\}$... the solution lists the stations chosen for stops. So s_i is the first stop and s_{n+1} (Montreal) is the last.

First we show that there exists at least one optimal solution that starts with the algorithm's first choice. To do this, let O be some optimal solution to the problem, i.e. $O = \{o_1, o_2, \dots o_t\}$ where each element is a station where we should fill up with gas, in order. Note that $o_t = s_{n+1}$ since the trip must end in Montreal.

Consider s_i : the algorithm's first stop. From the definition of the algorithm we know it is not possible to reach s_{i+1} on the first tank of gas. Therefore $o_1 \leq s_i$

But now consider o_2 , the second stop in O. It must be reachable from o_1 ... which means it is also reachable from s_i . Thus we can create a new feasible solution:

Let $O^* = \{s_i, o_2, ..., o_t\}$

Note that $|O^*| = |O|$, so O^{*} is also optimal. Thus there is an optimal solution that contains s_i .

Now we know that the algorithm's first decision is 'safe' - ie there is at least one optimal solution that matches this decision.

When we make a decision about the first stop, the problem reduces in size to $\leq k$ (i.e. the number of remaining stations between us and Montreal is reduced, and our tank is full again). By the IH, the algorithm will find an optimal solution to this reduced problem.

The last step of our proof is to show that when we combine the algorithm's first choice with its solution to the reduced problem, we get an optimal solution to the original problem.

We know that there is at least one optimal solution that matches the algorithm's first choice s_i – we identified O* above as one such, so let's use that one again: O* = { s_i , o_2 , o_3 , ... o_t }.

Let the algorithm's solution be $S_A = \{s_i, s_j, \dots, s_{n+1}\}$, as before.

Due to the structure of the algorithm, we know S_A is a feasible solution.

Observe that $\{s_j, \ldots, s_{n+1}\}$ and $\{o_2, \ldots, o_t\}$ are both solutions to the problem of getting from station s_i to Montreal, based on the decision made to fill up at station s_i . In fact, we know $\{s_j, \ldots, s_{n+1}\}$ is an **optimal** solution to this reduced problem of getting from station s_i to Montreal.

Thus $|\{s_j, \ldots s_{n+1}\}| \leq |\{o_2, \ldots o_t\}|$. Since adding s_i to both these sets gives S_A and O^* , it is clear that $|S_A| \leq |O^*|$. Thus S_A is also an optimal solution.

This completes the proof that RT(0) finds an optimal solution to the Road Trip problem.

Based on this example we can state the general Greedy Algorithm paradigm:

The fundamental principle of greedy algorithms is "take the best thing first".

More formally, suppose we are trying to solve an **optimization problem** that involves choosing objects from a set, subject to some feasibility constraint. We want to choose the objects that will maximize (or minimize) the total value of the solution, while satisfying the constraint. The **Greedy Paradigm** looks like this:

- 1. Sort the objects according to some criterion
- 2. repeat

select the $next_{s_j,...,s_{n+1}}$ in the list if it does not violate the feasibility constraint until no more selections can be made

For the Road Trip problem, the algorithm can be phrased as

- 1. Sort the objects according to the distance from the start
- 2. repeat

skip the next station if that doesn't cause us to run out of gas until we reach the destination

Our original goal was to minimize the number of stations we stop at ... but this is exactly equivalent to maximizing the number stations we skip.

The Greedy Algorithm is such a simple idea that we cannot expect it to find optimal solutions for all problems (for example, it is very easy to define greedy algorithms that find solutions to some instances of NP-Complete problems ... but there will be instances of the problem for which the algorithm's solutions won't be optimal).

However, there is a huge class of problems for which simple greedy algorithms **will** always find the optimal solution. This means that our big job with respect to Greedy Algorithms will not be coming up with the algorithm, but finding a proof that the algorithm gives the optimal solution for all instances of the problem.