
Lab Worksheet 5: Greedy Algorithms

Interval Scheduling Problem

Imagine you are in charge of a single shared resource, like a lecture hall. You have received n requests to book this hall for various events. Each request i is for a specific time interval, starting at s_i and finishing at f_i . Because there is only one hall, you can only approve requests for intervals that do not overlap. We say two intervals, i and j , are **compatible** if the time period for i is completely over before the period for j begins, or vice-versa. Formally, they are compatible if $f_i \leq s_j$ or $f_j \leq s_i$.

Our Goal: To maximize the use of the hall, you want to select the largest possible subset of these requests that are all mutually compatible. Here is an example:




Figure 1: A set of 4 interval requests. The optimal solution is the set {A, D}, with size 2.

Designing a Greedy Algorithm

We will solve this problem using a **greedy algorithm**. The idea is to build the solution piece by piece. At each step, we will make a choice that seems best at that moment, without worrying about future consequences.

Before we explore specific strategies, let's think about the structure of our greedy algorithm more formally. Let's identify the following pieces:

- **State:** What defines the current state of the problem we are trying to solve?
- **Action:** What is the single action our algorithm takes at each step?
- **New State:** After we take an action, how does the state of the problem change? In other words, how do we define the remaining subproblem?

Well... here are the answers: in this context, the state of the problem is the set of available intervals, and our action is to select one. This leads to the following general process:

1. Start with the set of all n requested intervals.
2. Select one interval based on **some “greedy” strategy**. Add it to our solution set.
3. Remove the selected interval and all intervals that conflict with it from the set of requests.
4. Repeat until no more intervals are left to consider.

The crucial part is **the “greedy” strategy** for selecting an interval. A good strategy leads to an optimal solution, while a bad one may not. The state of our problem at any point is simply the set of intervals that are still available to be chosen.

Exploring some (incorrect) greedy strategies

Let's explore a few greedy strategies that seem reasonable at first glance. For each one, your task is to find a **counterexample**. That is provide a small set of intervals where the strategy fails to produce the largest possible set of compatible intervals. Show which intervals the greedy algorithm selects and which intervals form the true optimal solution.

Earliest Start Time: The greedy strategy is that at each step, pick the available interval with the earliest start time.

Your counter example:

Shortest Duration: The greedy strategy is that at each step, pick the available interval with the smallest duration ($f_i - s_i$).

Your counter example:

A Correct Greedy Strategy: Earliest Finish Time

Here is a strategy that does work. At each step, pick the available interval with the earliest finish time (smallest f_i).

Let's trace the algorithm. Consider the set of intervals shown below. First, list the intervals (by their letter) in sorted order of their finish times. Then, trace the *Earliest Finish Time* algorithm: state which interval is chosen at each step and which intervals are removed. What is the final set of intervals in your solution?





Figure 2: An example to trace the *Earliest Finish Time* algorithm.



Proof of Optimality

Now, let's focus on proving that this strategy is always optimal. Here is the statement we want to prove:

Theorem 1. *The greedy algorithm that, at each step, selects an available interval with the earliest finish time produces an optimal solution for the Interval Scheduling problem.*

To prove this, we will use an **exchange argument**. The idea is to show that the greedy algorithm's choices are always *safe* or *correct*. We will demonstrate that for any optimal

solution, we can transform it step-by-step into the greedy solution without decreasing its size. Formally, we will build an **inductive argument** on the size of the problem. It is natural to assume that the number of available intervals is a good measure of size.

Write the predicate (or statement) you would like to prove via induction:

$P(n)$:

Base case: Show the statement for $n = 1$.

Induction hypothesis: For the rest of this proof, we assume that $P(n')$ holds for any $n' < n$.¹

Induction step: We will show that if $P(n')$ holds for all $n' < n$, then $P(n)$ also holds. In the following, we have an sketch of the argument. Complete the missing parts of the proof.

We start off by showing that we can bring the greedy solution closer to the optimal one by proving that at least one of the intervals selected by the greedy algorithm must appear in some optimal solution. This is the *exchange* part of the argument.

Lemma 1. Let g_1 denote the interval with the earliest finish time among all available intervals. Let $O = \{o_1, o_2, \dots, o_m\}$ be an optimal solution, sorted by finish times. There exists an optimal solution O' that includes g_1 .

Hint: Consider a good candidate to remove from O and replace with g_1 .

¹This is called **strong induction**. Instead of assuming only that $P(n - 1)$ holds, we assume that $P(1), P(2), \dots, P(n - 1)$ all hold. This approach sometimes makes the proof easier.

Lemma 2. Suppose we are given n intervals. Assume that the greedy algorithm produces an optimal solution for any set of intervals of size smaller than n . Let $G = \{g_1, \dots, g_k\}$ be the solution produced by the greedy algorithm, and let $O' = \{g_1, o_2, \dots, o_m\}$ be an optimal solution that contains the first greedy choice g_1 . Then, G is an optimal solution.

Greedy Algorithms: The Interval Stabbing Problem

Imagine you are a biologist studying migratory birds. You know the time intervals during which different species of birds visit a particular watering hole. You want to place automated cameras to record these visits. Each camera can only be placed at a single point in time, but it will record any bird species whose visiting interval includes that point in time. Your cameras are expensive, so you want to use as few as possible to ensure that every species is observed.

This is an instance of the **Interval Stabbing Problem**. You have a set of n intervals on a line, and you want to find the smallest set of points such that every interval is “stabbed” – that is, every interval contains at least one of the points.

Our Goal: Given n intervals, find a minimum-sized set of points that stabs all of them. Here is an example:




Figure 3: A set of 4 intervals. The optimal solution is to place two points, for example at times 2.5 and 7.5. Point p_1 stabs intervals A and B. Point p_2 stabs intervals C and D.

Designing a Greedy Algorithm

We will solve this with a greedy algorithm. The idea is to repeatedly place points until all intervals are covered. The “greedy” part is deciding *where* to place the next point.

Here is the general process:

1. Start with the set of all n intervals, none of which are stabbed yet.
2. While there are still unstabbed intervals:
 - 2.1 Choose a location to place a point based on *some “greedy” strategy*. Add this point to our solution set.
 - 2.2 Remove all intervals stabbed by this new point from the set of unstabbed intervals.
3. Return the set of points.

The crucial part is the strategy for placing a point. A good strategy will lead to an optimal (minimum size) set of points.

Exploring an incorrect greedy strategies

Let's explore a greedy strategy that seems plausible. Your task is to find a **counterexample**: a small set of intervals where the strategy fails to produce the smallest possible set of stabbing points.

Most Stabbing Point: The greedy strategy is to find a point on the line that stabs the maximum number of currently unstabbed intervals. Place a point there.

Your counterexample:



A Correct Greedy Strategy: Earliest Finish Time

Here is a strategy that does work. Find the unstabbed interval with the **earliest finish time**. Place a point at its finish time.

Let's trace this algorithm. First, sort the intervals below by their finish times. Then, trace the algorithm: state which point is chosen at each step and which intervals are stabbed. What is the final set of points?





Figure 4: An example to trace the *Earliest Finish Time* algorithm.



Proof of Optimality

Let's prove this strategy is always optimal using an **exchange argument**. The core idea is to show that the greedy choice is always “safe.” We will show that we can transform any optimal solution into the greedy solution without increasing its size.

Theorem 2. *The greedy algorithm that repeatedly finds the unstabbed interval with the earliest finish time and places a point at that finish time produces an optimal solution for the Interval Stabbing problem.*

Let $G = \{p_1, p_2, \dots, p_k\}$ be the solution produced by our greedy algorithm, sorted by position. Let $O = \{q_1, q_2, \dots, q_m\}$ be any optimal solution, also sorted by position. We want to prove that $k = m$.

Our proof will show that the greedy algorithm “stays ahead.” Specifically, we will prove by induction that for all $i \geq 1$, the point p_i is at least as far to the right as q_i .

Claim: For all $i = 1, \dots, m$, we have $p_i \geq q_i$.

Base case ($i = 1$): Show that $p_1 \geq q_1$.

Hint: Let i_1 be the first interval chosen by the greedy algorithm (the one with the earliest finish time). Where does the greedy algorithm place p_1 ? Where must an optimal solution place a point q_j to stab i_1 ? How do these positions relate?

Inductive Step: Assume that $p_{i-1} \geq q_{i-1}$ for some $i > 1$. We want to show that $p_i \geq q_i$. Before we prove the inductive step, let's first show the following lemma:

Lemma 3. Suppose the greedy algorithm, at step j , selects i^* , the unstabbed interval with the earliest finish time, and places a point $p_j = f_{i^*}$ at that finish time of i^* . If an interval i' remains unstabbed after placing p_j , then i' must start strictly after p_j ; that is, $p_j < s_{i'}$.

Using this lemma, complete the argument for the inductive step.

Conclusion of the Proof: We have shown that $p_i \geq q_i$ for all $i = 1, \dots, m$. How does this prove that the greedy solution is optimal (i.e., that $k = m$)?

Hint: Assume for contradiction that the greedy algorithm is not optimal. What would that imply about the relationship between k and m ? Can you use the claim to find a contradiction?