## Lecture 6

## All-Pairs Shortest Paths I

Supplemental reading in CLRS: Chapter 25 intro; Section 25.2

### 6.1 Dynamic Programming

Like the greedy and divide-and-conquer paradigms, dynamic programming is an algorithmic paradigm in which one solves a problem by combining the solutions to smaller subproblems. In dynamic programming, the subproblems overlap, and the solutions to "inner" problems are stored in memory (see Figure 6.1). This avoids the work of repeatedly solving the innermost problem. Dynamic programming is often used in optimization problems (e.g., finding the maximum or minimum solution to a problem).

- The key feature that a problem must have in order to be amenable to dynamic programming is that of optimal substructure: the optimal solution to the problem must contain optimal solutions to subproblems.
- Greedy algorithms are similar to dynamic programming, except that in greedy algorithms, the solution to an inner subproblem does not affect the way in which that solution is augmented to the solution of the full problem. In dynamic programming the combining process is more sophisticated, and depends on the solution to the inner problem.
- In divide-and-conquer, the subproblems are disjoint.


Figure 6.1. Schematic diagrams for dynammic programming and divide-and-conquer.

### 6.2 Shortest-Path Problems

Given a directed graph $G=(V, E, w)$ with real-valued edge weights, it is very natural to ask what the shortest (i.e., lightest) path between two vertices is.

Input: A weighted, directed graph $G=(V, E, w)$ with real-valued edge weights, a start vertex $u$, and an end vertex $v$

Output: The minimal weight of a path from $u$ to $v$. That is, the value of

$$
\delta(u, v)= \begin{cases}\min \{w(p): \text { paths } u \stackrel{p}{\rightsquigarrow} v\} & \text { if such a path exists } \\ \infty & \text { otherwise },\end{cases}
$$

where the weight of a path $p=\left\langle v_{0}, v_{1}, \ldots, v_{k}\right\rangle$ is

$$
w(p)=\sum_{i=1}^{k} w\left(v_{i-1}, v_{i}\right) .
$$

(Often we will also want an example of a path which achieves this minimal weight.)
Shortest paths exhibit an optimal-substructure property:
Proposition 6.1 (Theorem 24.1 of CLRS). Let $G=(V, E, w)$ be a weighted, directed graph with realvalued edge weights. Let $p=\left\langle v_{0}, v_{1}, \ldots, v_{k}\right\rangle$ be a shortest path from $v_{0}$ to $v_{k}$. Let $i, j$ be indices with $0 \leq i \leq j \leq k$, and let $p_{i j}$ be the subpath $\left\langle v_{i}, v_{i+1}, \ldots, v_{j}\right\rangle$. Then $p_{i j}$ is a shortest path from $v_{i}$ to $v_{j}$.

Proof. We are given that

$$
p: v_{0} \xrightarrow{p_{0 i}} v_{i} \xrightarrow{p_{i j}} v_{j} \stackrel{p_{j k}}{\sim} v_{k}
$$

is a shortest path. If there were a shorter path from $v_{i}$ to $v_{j}$, say $v_{i} \stackrel{p_{i j}^{\prime}}{\rightsquigarrow} v_{j}$, then we could patch it in to obtain a shorter path from $v_{0}$ to $v_{k}$ :

$$
p^{\prime}: v_{0} \xrightarrow{p_{0 i}} v_{i} \stackrel{p_{i j}^{\prime}}{\sim} v_{j} \stackrel{p_{j k}}{\sim} v_{k},
$$

which is impossible.

### 6.2.1 All Pairs

In recitation, we saw Dijkstra's algorithm (a greedy algorithm) for finding all shortest paths from a single source, for graphs with nonnegative edge weights. In this lecture, we will solve the problem of finding the shortest paths between all pairs of vertices. This information is useful in many contexts, such as routing tables for courier services, airlines, navigation software, Internet traffic, etc.

The simplest way to solve the all-pairs shortest path problem is to run Dijkstra's algorithm $|V|$ times, once with each vertex as the source. This would take time $|V| \cdot T_{\text {DiJKSTRA }}$, which, depending on the implementation of the min-priority queue data structure, would be

$$
\begin{array}{ll}
\text { Linear array: } & O\left(V^{3}+V E\right)=O\left(V^{3}\right) \\
\text { Binary min-heap: } & O(V E \lg V) \\
\text { Fibonacci heap: } & O\left(V^{2} \lg V+V E\right) .
\end{array}
$$

However, Dijkstra's algorithm only works if the edge weights are nonnegative. If we wanted to allow negative edge weights, we could instead use the slower Bellman-Ford algorithm once per vertex:

$$
O\left(V^{2} E\right) \xrightarrow{\text { dense graph }} O\left(V^{4}\right)
$$

We had better be able to beat this!

### 6.2.2 Formulating the All-Pairs Problem

When we solved the single-source shortest paths problem, the shortest paths were represented on the actual graph, which was possible because subpaths of shortest paths are shortest paths and because there was only one source. This time, since we are considering all possible source vertices at once, it is difficult to conceive of a graphical representation. Instead, we will use matrices.

Let $n=|V|$, and arbitrarily label the vertices $1,2, \ldots, n$. We will format the output of the all-pairs algorithm as an $n \times n$ distance matrix $D=\left(d_{i j}\right)$, where

$$
d_{i j}=\delta(i, j)= \begin{cases}\text { weight of a shortest path from } i \text { to } j & \text { if a path exists } \\ \infty & \text { if no path exists }\end{cases}
$$

together with an $n \times n$ predecessor matrix $\Pi=\left(\pi_{i j}\right)$, where

$$
\pi_{i j}= \begin{cases}\text { NIL, } & \text { if } i=j \text { or there is no path from } i \text { to } j \\ \text { predecessor of } j \text { in our shortest path } i \rightsquigarrow j & \text { otherwise. }\end{cases}
$$

Thus, the $i$ th row of $\Pi$ represents the sinlge-source shortest paths starting from vertex $i$, and the $i$ th row of $D$ gives the weights of these paths. We can also define the weight matrix $W=\left(w_{i j}\right)$, where

$$
w_{i j}= \begin{cases}\text { weight of edge }(i, j) & \text { if edge exists } \\ \infty & \text { otherwise }\end{cases}
$$

Because $G$ is a directed graph, the matrices $D, \Pi$ and $W$ are not necessarily symmetric. The existence of a path from $i$ to $j$ does not tell us anything about the existence of a path from $j$ to $i$.

### 6.3 The Floyd-Warshall Algorithm

The Floyd-Warshall algorithm solves the all-pairs shortest path problem in $\Theta\left(V^{3}\right)$ time. It allows negative edge weights, but assumes that there are no cycles with negative total weight. ${ }_{-}^{1}$ The FloydWarshall algorithm uses dynamic programming based on the following subproblem:

What are $D$ and $\Pi$ if we require all paths to have intermediate vertices ${ }^{2}$ taken from the set $\{1, \ldots, k\}$, rather than the full $\{1, \ldots, n\}$ ?

[^0]

Figure 6.2. Either $p$ passes through vertex $k$, or $p$ doesn't pass through vertex $k$.

As $k$ increases from 0 to $n$, we build up the solution to the all-pairs shortest path problem. The base case $k=0$ (where no intermediate vertices are allowed) is easy:

$$
D^{(0)}=\left(d_{i j}^{(0)}\right) \quad \text { and } \quad \Pi^{(0)}=\left(\pi_{i j}^{(0)}\right),
$$

where

$$
d_{i j}^{(0)}=\left\{\begin{array}{ll}
0 & \text { if } i=j \\
w_{i j} & \text { otherwise }
\end{array}\right\} \quad \text { and } \quad \pi_{i j}^{(0)}=\left\{\begin{array}{ll}
\mathrm{NIL} & \text { if } i=j \text { or } w_{i j}=\infty \\
i & \text { if } i \neq j \text { and there is an edge from } i \text { to } j
\end{array}\right\} .
$$

In general, let $D^{(k)}$ and $\Pi^{(k)}$ be the corresponding matrices when paths are required to have all their intermediate vertices taken from the set $\{1, \ldots, k\}$. (Thus $D=D^{(n)}$ and $\Pi=\Pi^{(n)}$.) What we need is a way of computing $D^{(k)}$ and $\Pi^{(k)}$ from $D^{(k-1)}$ and $\Pi^{(k-1)}$. The following observation provides just that.

Observation. Let $p$ be a shortest path from $i$ to $j$ when we require all intermediate vertices to come from $\{1, \ldots, k\}$. Then either $k$ is an intermediate vertex or $k$ is not an intermediate vertex. Thus, one of the following:

1. $k$ is not an intermediate vertex of $p$. Thus, $p$ is also a shortest path from $i$ to $j$ when we require all intermediate vertices to come from $\{1, \ldots, k-1\}$.
2. $k$ is an intermediate vertex for $p$. Thus, we can decompose $p$ as

$$
p: i \stackrel{p_{i k}}{\sim} k \stackrel{p_{k j}}{\sim} j,
$$

where $p_{i k}$ and $p_{k j}$ are shortest paths when we require all intermediate vertices to come from $\{1, \ldots, k-1\}$, by (a slightly modified version of Proposition 6.1.

This tells us precisely that

$$
d_{i j}^{(k)}= \begin{cases}w_{i j} & \text { if } k=0 \\ \min \left(d_{i j}^{(k-1)}, d_{i k}^{(k-1)}+d_{k j}^{(k-1)}\right) & \text { if } k \geq 1 .\end{cases}
$$

Thus we may compute $D$ by the following procedure:

```
Algorithm: FLOYD-WARSHALL( \(W\) )
    \(n \leftarrow|W . r o w s|\)
    \(D^{(0)} \leftarrow W\)
    \(\Pi^{(0)} \leftarrow\) an \(n \times n\) matrix with all entries initially NIL
    for \(k \leftarrow 1\) to \(n\) do
        Let \(D^{(k)}=\left(d_{i j}^{(k)}\right)\) and \(\Pi^{(k)}=\left(\pi_{i j}^{(k)}\right)\) be new \(n \times n\) matrices
        for \(i \leftarrow 1\) to \(n\) do
            for \(j \leftarrow 1\) to \(n\) do
            \(d_{i j}^{(k)} \leftarrow \min \left(d_{i j}^{(k-1)}, d_{i k}^{(k-1)}+d_{k j}^{(k-1)}\right)\)
            if \(d_{i j}^{(k)} \leq d_{i k}^{(k-1)}+d_{k j}^{(k-1)}\) then
                \(\pi_{i j}^{(k)} \leftarrow \pi_{i j}^{(k-1)}\)
                    else
                    \(\pi_{i j}^{(k)} \leftarrow \pi_{k j}^{(k-1)}\)
```

Being in a triply nested loop, lines $8-12$ are executed $n^{3}$ times. Thus, the algorithm runs in $\Theta\left(n^{3}\right)=$ $\Theta\left(V^{3}\right)$ time. Note that while the algorithm as written requires $\Theta\left(V^{3}\right)$ space, reasonable implementations will only require $\Theta\left(V^{2}\right)$ space: we only ever need to memorize four matrices at a time, namely $D^{(k)}, D^{(k-1)}, \Pi^{(k)}, \Pi^{(k-1)}$. In fact, we can get away with memorizing even less-see Exercise 25.2-4 of CLRS.

MIT OpenCourseWare
http://ocw.mit.edu

### 6.046J / 18.410J Design and Analysis of Algorithms

Spring 2012

For information about citing these materials or our Terms of Use, visit: http://ocw.mit.edu/terms.


[^0]:    ${ }^{1}$ If there were a negative-weight cycle, then there would exist paths of arbitrarily low weight, so some vertex pairs would have a distance of $-\infty$. In particular, the distance from a vertex to itself would not necessarily be zero, so our algorithms would need some emendations.
    ${ }^{2}$ An intermediate vertex in a path $p=\left\langle v_{0}, \ldots, v_{n}\right\rangle$ is one of the vertices $v_{1}, \ldots, v_{n-1}$, i.e., not an endpoint.

