

Greedy Algorithms

TOPICS

- Greedy Strategy
- Activity Selection
- Minimum Spanning Tree
- Shortest Paths
- Huffman Codes
- Fractional Knapsack

Chapter 5

Algorithm Design *Kleinberg and Tardos*

The Greedy Principle

- **The problem:** We are required to find a feasible solution that either maximizes or minimizes a given objective solution.
- It is easy to determine a feasible solution but not necessarily an optimal solution.
- The greedy method solves this problem in stages, at each stage, a decision is made considering inputs in an order determined by the selection procedure which may be based on an optimization measure.
- The greedy algorithm always makes the choice that looks best at the moment.
 - For each decision point in the greedy algorithm, the choice that seems best at the moment is chosen
- It makes a local optimal choice that may lead to a global optimal choice.

Activity Selection Problem

- Scheduling a resource among several competing activities.
- $S = \{1, 2, 3, \dots, n\}$ is the set of n proposed activities
- The activities share a resource, which can be used by only one activity at a time -a Tennis Court, a Lecture Hall etc.,
- Each activity i has a start time, s_i and a finish time f_i , where $s_i \leq f_i$.
- When selected, the activity takes place during time (s_i, f_i)
- Activities i and j are compatible if $s_i \geq f_j$ or $s_j \geq f_i$
- The activity-selection problem selects the maximum-size set of mutually compatible activities
- The input activities are in order by increasing finishing times.
- $f_1 \leq f_2 \leq f_3 \dots \leq f_n$; Can be sorted in $O(n \log n)$ time

Procedure for activity selection (from CLRS)

```
Procedure GREEDY_ACTIVITY_SELECTOR( $s, f$ )  
 $n \leftarrow$  length  $[S]$ ; in order of increasing finishing times;  
 $A \leftarrow \{1\}$ ; first job to finish  
 $j \leftarrow 1$ ;  
for  $i \leftarrow 2$  to  $n$   
    do if  $s_i \geq f_j$   
        then  $A \leftarrow A \cup \{i\}$ ;  
             $j \leftarrow i$ ;
```

i	s_i	f_i
1	1	4
2	3	5
3	0	6
4	5	7
5	3	8
6	5	9
7	6	10
8	8	11
9	8	12
10	2	13
11	12	14

- Initially we choose activity 1 as it has the least finish time.
- Then, activities 2 and 3 are not compatible as $s_2 < f_1$ and $s_3 < f_1$.
- We choose activity 4, $s_4 > f_1$, and add activity 4 to the set A.
- $A = \{1, 4\}$
- Activities 5, 6, and 7 are incompatible and activity 8 is chosen
- $A = \{1, 4, 8\}$
- Finally activity 10 is incompatible and activity 11 is chosen
- $A = \{1, 4, 8, 11\}$
- The algorithm can schedule a set of n activities in $\Theta(n)$ time.

Greedy Algorithms

- Minimum Cost Spanning Tree
 - Kruskal's algorithm
 - Prim's Algorithm
- Single Source Shortest Path
- Huffman Codes

Minimum-Cost Spanning Trees

Consider a network of computers connected through bidirectional links. Each link is associated with a positive cost: the cost of sending a message on each link.

This network can be represented by an undirected graph with positive costs on each edge.

In bidirectional networks we can assume that the cost of sending a message on a link does not depend on the **direction**.


Suppose we want to broadcast a message to all the computers from an arbitrary computer.

The cost of the broadcast is the sum of the costs of links used to forward the message.


Minimum-Cost Spanning Trees

- Find a fixed connected subgraph, containing all the vertices such that the sum of the costs of the edges in the subgraph is minimum. This subgraph is a tree as it does not contain any cycles.
- Such a tree is called the spanning tree since it spans the entire graph G .
- A given graph may have more than one spanning tree
- The minimum-cost spanning tree (MCST) is one whose edge weights add up to the least among all the spanning trees

MCST



A Local Area Network



The equivalent Graph and the MCST

MCST

- **The Problem:** Given an undirected connected weighted graph $G = (V, E)$, find a spanning tree T of G of minimum cost.
- **Greedy Algorithm for finding the Minimum Spanning Tree of a Graph $G = (V, E)$**

The algorithm is also called **Kruskal's** algorithm.

- At each step of the algorithm, one of several possible choices must be made,
- The greedy strategy: make the choice that is the best at the moment

Kruskal's Algorithm

- Procedure **MCST_G(V,E)**
- (Kruskal's Algorithm)
- **Input:** An undirected graph $G(V,E)$ with a cost function c on the edges
- **Output:** T the minimum cost spanning tree for G
- $T \leftarrow 0;$
- $VS \leftarrow 0;$
- **for** each vertex $v \in V$ **do**
- $VS = VS \cup \{v\};$
- sort the edges of E in nondecreasing order of weight
- **while** $|VS| > 1$ **do**
- choose (v,w) an edge E of lowest cost;
- delete (v,w) from $E;$
- **if** v and w are in different sets $W1$ and $W2$ in VS **do**
- $W1 = W1 \cup W2;$
- $VS = VS - W2;$
- $T \leftarrow T \cup (v,w);$
- return T

Kruskals_MCST

- The algorithm maintains a collection VS of disjoint sets of vertices
- Each set W in VS represents a connected set of vertices forming a spanning tree
- Initially, each vertex is in a set by itself in VS
- Edges are chosen from E in order of increasing cost, we consider each edge (v, w) in turn; $v, w \in V$.
- If v and w are already in the same set (say W) of VS , we discard the edge
- If v and w are in distinct sets $W1$ and $W2$ (meaning v and/or w not in T) we merge $W1$ with $W2$ and add (v, w) to T .

Kruskals_ MCST

Consider the example graph shown earlier,

The edges in nondecreasing order

[(A,D),1],[(C,D),1],[(C,F),2],[(E,F),2],[(A,F),3],[(A,B),3],

[(B,E),4],[(D,E),5],[(B,C),6]

EdgeActionSets in VSSpanning Tree, T
=[{A},{B},{C},{D},{E},{F}]{0}{(A,D)merge

[{A,D},{B},{C},{E},{F}]{(A,D)}{(C,D)merge

[{A,C,D},{B},{E},{F}]{(A,D),(C,D)}{(C,F)merge

[{A,C,D,F},{B},{E}]{(A,D),(C,D),(C,F)}{(E,F)merge


[{A,C,D,E,F},{B}]{(A,D),(C,D),(C,F),(E,F)}{(A,F)reject

[{A,C,D,E,F},{B}]{(A,D),(C,D),(C,F),(E,F)}{(A,B)merge

[{A,B,C,D,E,F}]{(A,D),(A,B),(C,D),(C,F),(E,F)}{(B,E)reject

(D,E) reject

(B,C) reject






The equivalent Graph and the MCST




Kruskals_ MCST Complexity

- Steps 1 thru 4 take time $O(V)$
- Step 5 sorts the edges in nondecreasing order in $O(E \log E)$ time
- Steps 6 through 13 take $O(E)$ time
- The total time for the algorithm is therefore given by $O(E \log E)$
- The edges can be maintained in a heap data structure with the property,
- $E[\text{PARENT}(i)] \leq E[i]$
- remember, this property is the opposite of the one used in the heapsort algorithm earlier. This property can be used to sort data elements in nonincreasing order.
- Construct a heap of the edge weights, the edge with lowest cost is at the root
- During each step of edge removal, delete the root (minimum element) from the heap and rearrange the heap.
- The use of heap data structure reduces the time taken because at every step we are only picking up the minimum or root element rather than sorting the edge weights.

Prim's Algorithm




The equivalent Graph and the MCST



Single-Source Shortest Paths

A motorist wishes to find the shortest possible route from Perth to Brisbane. Given the map of Australia on which the distance between each pair of cities is marked, how can we determine the shortest route?



Single Source Shortest Path

- In a shortest-paths problem, we are given a weighted, directed graph $G = (V, E)$, with weights assigned to each edge in the graph. The weight of the path $p = (v_0, v_1, v_2, \dots, v_k)$ is the sum of the weights of its constituent edges:
 - $v_0 \rightarrow v_1 \rightarrow v_2 \dots \rightarrow v_{k-1} \rightarrow v_k$
 -
- The shortest-path from u to v is given by
- $d(u, v) = \min \{\text{weight}(p) : \text{if there are one or more paths from } u \text{ to } v$
- $= \infty \text{ otherwise}$

The single-source shortest paths problem

Given $G (V,E)$, find the shortest path from a given vertex $u \in V$ to every vertex $v \in V (u \neq v)$.

For each vertex $v \in V$ in the weighted directed graph, $d[v]$ represents the distance from u to v .

Initially, $d[v] = 0$ when $u = v$.


$d[v] = \infty$ if (u,v) is not an edge

$d[v] = \text{weight of edge } (u,v)$ if (u,v) exists.


Dijkstra's Algorithm : At every step of the algorithm, we compute,
 $d[y] = \min \{d[y], d[x] + w(x,y)\}$, where $x,y \in V$.

Dijkstra's algorithm is based on the greedy principle because at every step we pick the path of least weight.

- Dijkstra's Algorithm : At every step of the algorithm, we compute,
 $d[y] = \min \{d[y], d[x] + w(x,y)\}$, where $x,y \in V$.
- Dijkstra's algorithm is based on the greedy principle because at every step we pick the path of least path.



Example:



Step #	Vertex to be marked	Distance to vertex									Unmarked vertices
		u	a	b	c	d	e	f	g	h	
0	u	0	1	5	∞	9	∞	∞	∞	∞	a,b,c,d,e,f,g,h
1	a	0	1	5	3	9	∞	∞	∞	∞	b,c,d,e,f,g,h
2	c	0	1	5	3	7	∞	12	∞	∞	b,d,e,f,g,h
3	b	0	1	5	3	7	8	12	∞	∞	d,e,f,g,h
4	d	0	1	5	3	7	8	12	11	∞	e,f,g,h
5	e	0	1	5	3	7	8	12	11	9	f,g,h
6	h	0	1	5	3	7	8	12	11	9	g,h
7	g	0	1	5	3	7	8	12	11	9	h
8	f	0	1	5	3	7	8	12	11	9	--

Dijkstra's Single-source shortest path

- Procedure **Dijkstra's Single-source shortest path_G(V,E,u)**
- Input: $G=(V,E)$, the weighted directed graph and v the source vertex
- Output: for each vertex, v , $d[v]$ is the length of the shortest path from u to v .
- mark vertex u ;
- $d[u] \leftarrow 0$;
- **for** each unmarked vertex $v \in V$ **do**
- **if** edge (u,v) exists $d[v] \leftarrow \text{weight}(u,v)$;
- **else** $d[v] \leftarrow \infty$;
- **while** there exists an unmarked vertex **do**
- let v be an unmarked vertex such that $d[v]$ is minimal;
- mark vertex v ;
- **for** all edges (v,x) such that x is unmarked **do**
- **if** $d[x] > d[v] + \text{weight}[v,x]$ **then**
- $d[x] \leftarrow d[v] + \text{weight}[v,x]$

- Complexity of Dijkstra's algorithm:
- Steps 1 and 2 take $\Theta(1)$ time
- Steps 3 to 5 take $O(|V|)$ time
- The vertices are arranged in a heap in order of their paths from u
- Updating the length of a path takes $O(\log V)$ time.
- There are $|V|$ iterations, and at most $|E|$ updates
- Therefore the algorithm takes $O((|E| + |V|) \log |V|)$ time.

Huffman codes

Huffman codes are used to compress data. We will study Huffman's greedy algorithm for encoding compressed data.

Data Compression

- **A given file can be considered as a string of characters.**
- **The work involved in compressing and uncompressing should justify the savings in terms of storage area and/or communication costs.**
- **In ASCII all characters are represented by bit strings of size 7.**
- **For example if we had 100000 characters in a file then we need 700000 bits to store the file using ASCII.**

Example

The file consists of only 6 characters as shown in the table below.

Using the fixed-length binary code, the whole file can be encoded in 300,000 bits.

However using the variable-length code , the file can be encoded in 224,000 bits.

	a	b	c	d	e	f
Frequency (in thousands)	45	13	12	16	9	5
Fixed-length codeword	000	001	010	011	100	101
Variable-length codeword	0	101	100	111	1101	1100

A variable length coding scheme assigns frequent characters, short code words and infrequent characters, long code words.

In the above variable-length code, 1-bit string represents the most frequent character *a*, and a 4-bit string represents the most infrequent character *f*.

Let us denote the characters by C_1, C_2, \dots, C_n and denote their frequencies by f_1, f_2, \dots, f_n . Suppose there is an encoding E in which a bit string S_i of length s_i represents C_i , the length of the file compressed by using encoding E is

$$L(E, F) = \sum_{i=1}^n s_i \cdot f_i$$


Prefix Codes

- The prefixes of an encoding of **one character must not be equal to a complete encoding of another character.**
 - 1100 and 11001 are not valid codes
 - because 1100 is a prefix of 11001
- This constraint is called the prefix constraint.
- Codes in which no codeword is also a prefix of some other code word are called prefix codes.
- Shortening the encoding of one character may lengthen the encodings of others.
- **To find an encoding E that satisfies the prefix constraint and minimizes $L(E,F)$.**

0 101 100 111 1101 1100

The prefix code for file can be represented by a binary tree in which every non leaf node has two children. Consider the variable-length code of the table above, a tree corresponding to the variable-length code of the table is shown below.

Note that the length of the code for a character is equal to the depth of the character in the tree shown.



Greedy Algorithm for Constructing a Huffman Code

The algorithm builds the tree corresponding to the optimal code in a bottom-up manner.

The algorithm begins with a set of $|C|$ leaves and performs a sequence of 'merging' operations to create the tree.

C is the set of characters in the alphabet.

Procedure **Huffman_Encoding(S,f);**

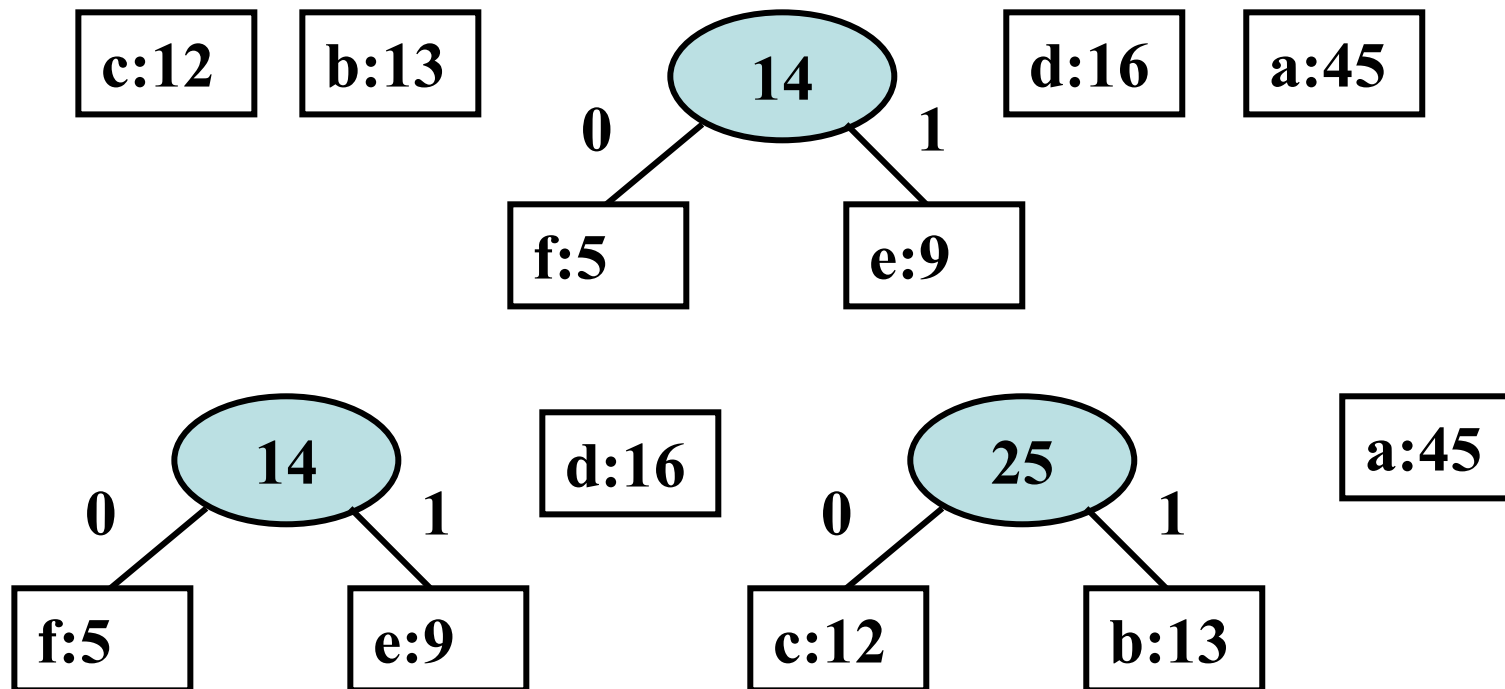
Input : **S** (a string of characters) and **f** (an array of frequencies).


Output : **T** (the Huffman tree for **S**)

1. insert all characters into a heap **H** according to their frequencies;
2. **while** **H** is not empty **do**
3. **if** **H** contains only one character **x** **then**
4. $x \leftarrow \text{root}(T)$;
5. **else**
6. $z \leftarrow \text{ALLOCATE_NODE}()$;
7. $x \leftarrow \text{left}[T,z] \leftarrow \text{EXTRACT_MIN}(H)$;
8. $y \leftarrow \text{right}[T,z] \leftarrow \text{EXTRACT_MIN}(H)$;
9. $f_z \leftarrow f_x + f_y$;
10. **INSERT**(**H**,**z**);

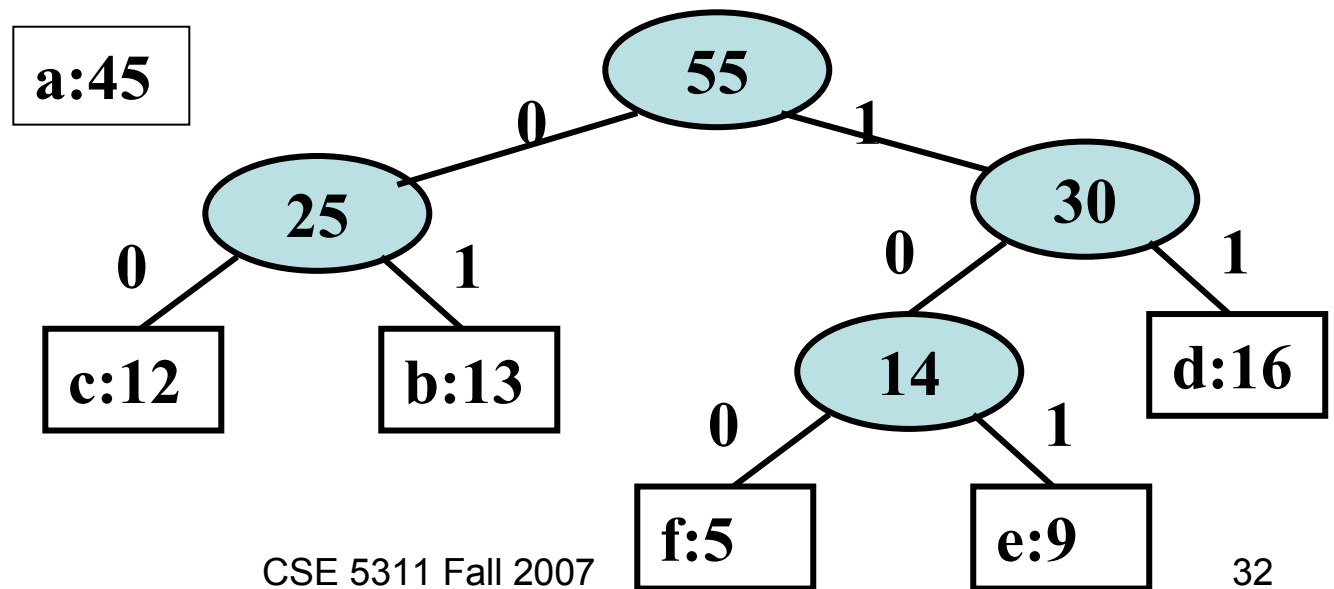



- The algorithm is based on a reduction of a problem with n characters to a problem with $n-1$ characters.
- A new character replaces two existing ones.





Suppose C_i and C_j are two characters with minimal frequency, there exists a tree that minimizes $L(E,F)$ in which these characters correspond to leaves with the maximal distance from the root.





0 101 100 111 1101 1100

Complexity of the algorithm

Building a heap in step 1 takes $O(n)$ time

Insertions (steps 7 and 8) and
deletions (step 10) on H

take $O(\log n)$ time each

Therefore Steps 2 through 10 take $O(n \log n)$ time

Thus the overall complexity of the algorithm is
 $O(n \log n)$.

- The fractional knapsack problem
 - Limited supply of each item
 - Each item has a size and a value per unit (e.g., Pound)
- greedy strategy
 - Compute value per Pound for each item
 - Arrange these in non-increasing order
 - Fill sack with the item of greatest value per pound until either the item is exhausted or the sack is full
 - If sack is not full, fill the remainder with the next item in the list
 - Repeat until sack is full

How about a 0-1 Knapsack?? Can we use Greedy strategy?