Chapter 9 Graph Algorithms

ntroductio

-graph theory

- -useful in practice
- -represent many real-life problems
- -can be slow if not careful with data structures



1

Definitions

- -an undirected graph G = (V, E) is a finite set V of <u>vertices</u> together with a set E of <u>edges</u>
- an edge is a pair (v, w), where v and w are vertices
- -this definition allows
 - -<u>self-loops</u>, or edges that connect vertices to themselves -<u>parallel edges</u>, or multiple edges that connect the same pair of vertices
- -a graph without self-loops is a simple graph
- -a graph with parallel edges is sometimes called a multigraph

3

Definitions

- -two vertices are adjacent if there is an edge between them
 -the edge is said to be incident to the two vertices
- if there are no parallel edges, the degree of a vertex is the number of edges incident to it
- -self-loops add only 1 to the degree
- -a subgraph of a graph *G* is a subset of *G*'s edges together with the incident vertices

4

Definitions

- -a path in a graph is a sequence of vertices connected by edges
- a simple path is a path with no repeated vertices, except possibly the first and last
- -a cycle is a path of at least one edge whose first and last vertices are the same
 - a simple cycle is a cycle with no repeated edges of vertices other than the first and last

-the length of a path is the number of edges in the path

Definitions

- -a graph is connected if every vertex is connected to every other vertex by a path through the graph
- a connected component G' of a graph G is a maximal connected subgraph of G: if G' is a subset of F and F is a connected subgraph of G, then F = G'
- a graph that is not connected consists of a set of connected components
- -a graph without cycles is called acyclic

Definitions

- -a tree is a connected, acyclic undirected graph
- -a forest is a disjoint set of trees
- -a spanning tree of a connected graph is a subgraph that is a tree and also contains all of the graph's <u>vertices</u>
- -a spanning forest of a graph is the union of spanning trees of its connected components

Definitions

- -if |V| is the number of vertices and |E| is the number of edges, then, in a graph without self-loops and parallel edges, there are |V| (|V| 1)/2 possible <u>edges</u>
- -a graph is complete if there is an edge between every pair of vertices
- the density of a graph refers to the proportion of possible pairs of vertices that are connected
- -a sparse graph is one for which $|E| \ll |V|(|V| 1)/2$
- -a dense graph is a graph that is not sparse
- -a bipartite graph is one whose vertices can be divided into two sets so that every vertex in one set is connected to at least one vertex in the other set

8

Definitions

7

- -in a directed graph or digraph the pairs (v, w) indicating edges are <u>ordered</u>: the edge (v, w) goes from v (the tail) to w (the head)
- -since edges have a direction, we use the notation $v \to w$ to denote an edge from v to w
- -edges in digraphs are frequently called arcs
- -the <u>indegree</u> of a vertex w is the number of arcs $v \rightarrow w$ (i.e., the number of arcs coming into w), while the outdegree of w is the number of arcs $w \rightarrow v$ (i.e., the number of arcs exiting w)
- -we will call w a source if its indegree is 0
- -an aborescence is a directed graph with a distinguished vertex u (the root) such that for every other vertex v there is a unique directed path from u to v

9



Definitions

- in a directed graph, two vertices *v* and *w* are <u>strongly</u> connected if there is a directed path from *v* to *w* and a directed path from *w* to *v*
- a digraph is strongly connected if all its vertices are strongly connected
- if a digraph is not strongly connected but the underlying undirected graph is connected, then the digraph is called weakly connected
- -a weighted graph has weights or costs associated with each edge
- -weighted graphs can be directed or undirected
- -a road map with mileage is the prototypical example

10

Graph Representation

- -two concerns: memory and speed
- -we'll consider directed graphs, though undirected graphs are similar
- -the following graph has 7 vertices and 12 edges





4

5

Graph Representation

-adjacency matrix

-alternatively, could use costs ∞ or $-\infty$ for <u>non-edges</u>

- -not efficient if the graph is <u>sparse</u> (number of edges small) -matrix $O(|V|^2)$
 - -e.g., street map with 3,000 streets results in intersection matrix with 9,000,000 elements
- -adjacency list
 - -standard way to represent graphs
- -undirected graph edges appear twice in list
- -more efficient if the graph is sparse (number of edges small)
 - -matrix O(|E| + |V|)

14



Topological Sort

15

Graph Representation

-adjacency list

- -a directed acyclic graph (DAG) is a digraph with no directed cycles
 - -a DAG always has at least one vertex
- -topological sort
 - an ordering of the vertices in a directed graph such that if there is a path from v to w, then v appears before w in the ordering
 - -not possible if graph has a cycle



Topological Sort

-topological sort

- determine the indegree for every $v \in V$ -place all source vertices in a queue
- -while there remains a $v \in V$
 - -find a source vertex
 - -append the source vertex to the topological sort
 - -delete the source and its adjacent arcs from G
 - -update the indegrees of the remaining vertices in G
 - -place any new source vertices in the queue
- -when no vertices remain, we have our ordering, or, if we are missing vertices from the output list, the graph has no topological sort

19



21

Shortest-Path Algorithms

-shortest-path problems

- input is a weighted graph with a <u>cost</u> on each edge weighted path length: $\sum_{i=1}^{N-1} c_{i,i+1}$
- .

-single-source shortest-path problem

-given as input a weighted graph, G = (V, E) and a <u>distinguished</u> vertex s, find the shortest weighted path from s to every other vertex in G

Topological Sort

```
void Graph::topsort()
{
    for{ int counter = 0; counter < NUM_VERTICES; counter++ )
    {
        for{ int counter = 0; counter < NUM_VERTICES; counter++ )
        {
            Vertex v = findlew/vertex0fIndegreeZero();
            tf(v == NUL_A_VERTEX)
            tf(v = NUL_A_VERTEX)
           tf(v = NUL_A_VERTEX)
           tf(v = NUL_A_VERTEX)
```

20



22

Shortest-Path Algorithms

-example

-shortest weighted path from v_1 to v_6 has cost of $\underline{6}$ -no path from v_6 to v_1



Shortest-Path Algorithms

- -negative edges can cause problems
- path from υ_5 to υ_4 has cost of 1, but a shorter path exists by following the negative loop, which has cost -5
- -shortest paths thus undefined



25

Shortest-Path Algorithms

-four problems

- -unweighted shortest-path
- -weighted shortest-path with no negative edges
- -weighted shortest-path with negative edges
- -weighted shortest-path in acyclic graphs

Shortest-Path Algorithms

- many examples where we want to find shortest paths
 if vertices represent computers and edges connections, the cost represents <u>communication</u> costs, delay costs, or combination of costs
- -if vertices represent airports and edges costs to travel between them, shortest path is <u>cheapest</u> route
- we find paths from one vertex to <u>all</u> others since no algorithm exists that finds shortest path from one vertex to one other faster

26

28

Unweighted Shortest Paths

- -unweighted shortest-path
 - -find shortest paths from *s* to all other vertices -only concerned with number of <u>edges</u> in path
 - -we will not actually record the path elements









<section-header><section-header><section-header><section-header><section-header>







Unweighted Shortest Paths	
void Graph::unweighted(Vertex s) { Queue=Vertex> q;	
for each Vertex v v.dist = INFINITY;	
<pre>s.dist = 0; q.enqueue(s);</pre>	
<pre>while(lq.isEmpty()) { Vertex v = q.dequeue(); </pre>	
for each Vertex w adjacent to v if(w.dist == INFINITY)	
v.dist = v.dist + 1; w.path = v;	
q.enqueue(w); }	
1	07
	- 37

Dijkstra's Algorithm

- -weighted shortest-path Dijkstra's algorithm
- -more difficult, but ideas from unweighted algorithm can be used
- -keep information as before for each vertex -known
- -set distance $d_w = d_v + c_{v,w}$ if $d_w = \infty$ using only known vertices
- $-p_v$ the last vertex to cause a change to d_v -greedy algorithm
 - -does what appears to be best thing at each stage
 - -e.g., counting money: count quarters first, then dimes,
 - nickels, pennies
 - -gives change with least number of coins

39

Dijkstra's Algorithm -pseudocode (cont.) $\begin{array}{l} \mbox{foreach vertex } v \ \{ \\ \mbox{dist} \left[v \right] = +\infty \\ \mbox{known} \left[v \right] = \mbox{false} \end{array}$ dist[s] = 0 } } }

0 1 v₃

Unweighted Shortest Paths Initial State

v ₂	F	∞	0	F	∞	0	F	2	v_1	F	2	v_1	
V3	F	0	0	т	0	0	Т	0	0	т	0	0	꼬꼬
$V_{-\frac{1}{2}}$	F	∞	0	F	∞	0	F	2	v1	F	2	v_1	$\langle \cdot \rangle \langle \cdot \rangle$
¥5	F	∞	0	F	∞	0	F	∞	0	F	∞	0	X X
v ₆	F	∞	0	F	1	v_3	F	1	v3	т	1	v3	- <u>K</u> _K
v_7	F	∞	0	F	∞	0	F	∞	0	F	∞	0	
Q:		v_3		1	v_1, v_6		v ₆ ,	v_2, v_4		1	2, V4		
	ν_2 Dequeued			v+ Dequeued			ν_5 Dequeued			v7 Dequeued			
v	known	$d_{\rm v}$	pv	known	d_v	p_v	known	d_{τ}	p_{v}	known	$d_{\rm v}$	p _v	
v1	Т	1	V3	т	1	v_3	т	1	¥3	т	1	¥3	
v2	Т	2	v ₁	т	2	v_1	Т	2	v1	Т	2	v_1	
v3	Т	0	0	т	0	0	т	0	0	Т	0	0	
V.4	F	2	v ₁	т	2	v_1	Т	2	v1	Т	2	v_1	
V5	F	3	V2	F	3	v2	т	3	V2	Т	3	v2	
¥6	Т	1	V3	т	1	¥3	Т	1	V3	Т	1	¥3	
v7	F	∞	0	F	3	v_{\pm}	F	3	ν_{+}	Т	3	v_{\pm}	
Q:	,	4.V5		,	5. 17			v7		e	mpty		

1 V3

v3 Dequeued v1 Dequeued

known $d_{\rm v}$ $p_{\rm v}$ known $d_{\rm v}$ $p_{\rm v}$ known $d_{\rm v}$ $p_{\rm v}$ known $d_{\rm v}$ $p_{\rm v}$ т

v₆ Dequeued

38

Dijkstra's Algorithm

- -pseudocode -assumption: no negative weights
 - -origin s is given

Initialization: $S \leftarrow \{s\}$ and $D \leftarrow V - \{s\}$. Set dist[s] $\leftarrow 0$ and dist[v] $\leftarrow \infty$ for all other v.

While there remains a $v \in D$:

- $\textcircled{\ }$ Select a vertex $v \in D$ which has the shortest path length from s to v using only vertices in S (e.g., known vertices).
- $S \leftarrow S \cup \{v\} \text{ and } D \leftarrow D \{v\}.$









45

Dijkstra's Algorithm: Correctness

Proposition. Dijkstra's algorithm solves the single-origin shortest-paths problem in a weight digraph with nonnegative weights.

Proof. If v is reachable from s, then every edge v $\rightarrow \! w$ is relaxed exactly once, when v is relaxed, resulting in

 $dist[w] \le dist[v] + weight[v \rightarrow w].$

This inequality holds until the algorithm terminates, since

- O dist[w] can only decrease, because relaxations can only decrease a dist[] value, and
- @ dist[v] never changes, because edge weights are nonnegative and we choose the lowest dist[] value at each step, so no later relaxation can reduce dist[v].

Thus, after all vertices reachable from s have been added to the shortest paths tree, the shortest paths optimality conditions hold.

Dijkstra's Algorithm

-complexity

46

- -sequentially scanning vertices to find minimum d_v takes O(|V|), which results in $O(|V|^2)$ overall
- -at most one update per edge, for a total of $O(|E| + |V|^2) =$ $O(|V|^2)$
 - -if graph is dense, with $|E| = \Theta(|V|^2)$, algorithm is close to optimal

-if graph is sparse, with $|E| = \Theta(|V|)$, algorithm is too slow -distances could be kept in a priority queue that reduces running time to $O(|E| + |V| \lg |V|)$





Dijkstra's Algorithm -implementation (cont.) void Graph::dijkstra(Vertex s) while(there is an unknown distance vertex) for each Vertex v Vertex v = smallest unknown distance vertex; v.dist = INFINITY: v.known = true; v.known = false; for each Vertex w adjacent to v if(!w.known) s.dist = 0; DistType cvw = cost of edge from v to w; if(v.dist + cvw < w.dist) // Update w
decrease(w.dist to v.dist + cvw); w.path = v; } }

51

Graphs with Negative Edges

- -possible solution: <u>add</u> a delta value to all weights such that none are negative
- -calculate shortest path on new graph
- -apply path to original graph
- -does not work: longer paths become weightier
- combination of algorithms for weighted graphs and unweighted graphs can work
 - -drastic increase in running time: $O(|E| \cdot |V|)$

Graphs with Negative Edges

-try to apply Dijkstra's algorithm to graph with <u>negative</u> edges Label each node with best known distance from origin *a*. Relax the edges adjacent to *a*. Select *b*, the closest node to *S*, and add it to *S*. Relax the outgoing edges adjacent to *b*. Select *c*, the closest node to *S*, and add it to *S*. We've now assimilated all nodes into *S*, so we're done.

52

All-Pairs Shortest Paths

- given a weighted digraph, find the shortest paths between $\underline{\text{all}}$ vertices in the graph
- one approach: apply Dijkstra's algorithm repeatedly
 results in O(|V|³)
- -another approach: apply Floyd-Warshall algorithm
 - -uses dynamic programming
 - -also results in $O(|V|^3)$

Minimum Spanning Tree

-assumptions

- -graph is connected
- -edge weights are not necessarily Euclidean distances
- -edge weights need not be all the same
- -edge weights may be zero or negative
- -minimum spanning tree (MST)
 - -also called minimum-weight spanning tree of a weighted graph
 - spanning tree whose weight (the sum of the weights of the edges in the tree) is the <u>smallest</u> among all spanning trees

55

57

Minimum Spanning Tree

- two algorithms to find the minimum spanning tree - Prim's Algorithm

- -R. C. Prim, Shortest Connection Networks and Some Generalizations, Bell System Technical Journal (1957)
- -Kruskal's Algorithm
 - -J. B. Kruskal, Jr., On the Shortest Spanning Subtree of a Graph and the Traveling Salesman Problem, Proc. of the American Mathematical Society (1956)

Minimum Spanning Tree

Minimum Spanning Tree

example

-Prim's algorithm
-grows tree in successive stages
Initialization: S ← {s} and D ← V - {s}.
While there remains a v ∈ D:
Find an edge with minimum weight (u, v) such that u ∈ S and v ∈ D.
S ← S ∪ {v} and D ← D - {v}.

58











Minimum Spanning Tree

-Prim's algorithm

- -runs on undirected graphs
- running time: $\mathcal{O}(|V|^2)$ without heaps, which is optimal for dense graphs
- –running time: $O(|E| \lg |V|)$ using binary heaps, which is good for sparse graphs



62

64

Minimum Spanning Tree

- -Kruskal's algorithm
 - -continually select the edges in order of <u>smallest</u> weight
 -accept the edge if it does not cause a <u>cycle</u> with already accepted edges





Minimum Spanning Tree

-Kruskal's algorithm: example 2 (cont.)



67

NP-Complete Problems

- -What problems can we solve algorithmically? which problems are easy? which problems are hard?
- Eulerian circuit: Given a vertex *s*, start at *s* and find a cycle that visits every edge exactly once -easy: solvable in O(|E| + |V|) using depth-first search
- -Hamiltonian circuit: Given a vertex *s*, start at *s* and find a cycle that visits each remaining <u>vertex</u> exactly once
- -really, really hard!

69

NP-Complete Problems

-halting problem

- suppose we had such an algorithm *terminates(p,x)*- create a new program:

```
program evil (z) {
   1: if terminates(z,z) goto 1
}
```

 -program evil() terminates if and only if the program z does not terminate when given its own code as input
 -no such algorithm can exist



68

NP-Complete Problems

-halting problem

- -in 1936, A. Church and A. Turing independently proved the non-solvability of the halting problem:
 - -is there an algorithm *terminates*(*p*,*x*) that takes an arbitrary program *p* and input *x* and returns True if *p* terminates when given input *x* and False otherwise?

-difficult: try to run it on itself

70

NP-Complete Problems

-decision problem

- -has a yes or no answer
- -undecidable if it is impossible to construct a single
- algorithm that will solve all instances of the problem
- -the halting problem is undecidable

NP-Complete Problems

-the class P

- set of problems for which there exists a <u>polynomial</u> time algorithm for their solution
- the runtime is bounded by a polynomial function of the \underline{size} of the problem
- -the class NP
 - set of decision problems for which the certification of a candidate <u>solution</u> as being correct can be performed in polynomial time
 - -non-deterministic polynomial time

73

NP-Complete Problems

-reductions

- -problem A reduces to problem B if the solvability of B implies the solvability of A
- -if A is reducible to B, then B is at least as hard to solve as A
- -in the context of algorithms, reducibility means an algorithm that solves *B* can be <u>converted</u> into an algorithm to solve *A*
 - -example: if we can sort a set of numbers, we can find the median, so finding the median reduces to sorting

75

NP-Complete Problems

-reductions

- decision version of traveling salesperson problem (TSP): -given a complete weighted graph and an integer *K*, does there exist a simple cycle that <u>visits</u> all vertices (tour) with total weight $\leq K$? -clearly, this is in *NP* - Hamiltonian circuit: given a graph G = (V, E), find a simple cycle that visits all the vertices
- -construct a new graph *G'* with the same vertices as *G* but which is <u>complete</u>; if an edge in *G'* is in *G*, give it weight 1; otherwise, give it weight 2
 - -construction requires O(|E| + |V|) work
 - –apply $\underline{\mathsf{TSP}}$ to see if there exists a tour with total weight |V|

NP-Complete Problems

-the class NP

- -for problems in *NP*, certifying a solution may not be difficult, but <u>finding</u> a solution may be very difficult
- -example: Hamiltonian circuit
 - -given a graph G, is there a simple cycle in G that includes every <u>vertex</u>?
 - -given a candidate solution, we can check whether it is a simple cycle in time $\propto |V|$, simply by <u>traversing</u> the path
- -however, finding a Hamiltonian circuit is hard!

74

NP-Complete Problems

-reductions

-problem *A* can be <u>polynomially</u> reduced to *B* if we can solve problem *A* using an algorithm for problem *B* such that the cost of solving *A* is

cost of solving B + a polynomial function of the problem size

-example: once we have sorted an array a[] of N numbers, we can find the median in <u>constant</u> time by computing N/2 and accessing a[N/2]



NP-Complete Problems

-NP-complete

- a problem *A* is *NP*-complete if it is in *NP* and all other problems in *NP* can be reduced to *A* in polynomial time
- Boolean satisfiablity (SAT): given a set of *N* boolean variables and *M* logical statements built from the variables using *and* and *not*, can you choose values for the variables so that all the statements are <u>true</u>? ($x_1 AND ! x_2 AND x_3$), (! $x_1 AND x_7$), ($x_1 AND x_{42}$), ...
- -SAT is NP-complete

79

NP-Complete Problems

-NP-complete problems

- -traveling salesperson
- -bin packing
- -knapsack
- -graph coloring
- -longest-path

81

Bin Packing

We are given n items of lengths ℓ_1, \ldots, ℓ_n , where $0 < \ell_i \leq 1$ for all i.

The items must be packed in bins of length 1, and they must be placed end-to-end. Once placed in a bin, items cannot be moved.

How do we pack them in a way that uses the fewest number of bins? This is an NP-hard problem!

Input: 0.2, 0.5, 0.4, 0.7, 0.1, 0.3, 0.8.

Optimal packing: B₁: 0.2, 0.8 B₂: 0.7, 0.3 B₃: 0.4, 0.1, 0.5

Online: item must be put in a bin before the next item is considered.

Offline: all the items are available for consideration at one time.

NP-Complete Problems

-NP-complete

- if we restrict attention to sets of boolean statements involving <u>3</u> variables, the problem is known as 3-SAT
 -3-SAT is *NP*-complete
 - -so, if you can solve 3-SAT in polynomial time, you can solve <u>all</u> problems in *NP* in polynomial time
- -meanwhile, 2-SAT is solvable in linear time!

80

NP-Complete Problems

-NP-hard problems

- -a problem *A* is *NP*-hard if there exists a polynomial-time reduction from an *NP*-complete problem to *A*
- an *NP*-hard problem is at <u>least</u> as hard as an *NP*-complete problem
- -optimization versions of *NP*-complete problems are typically *NP*-hard
 - -optimization version of TSP: given a weighted graph, find a minimum cost Hamiltonian circuit
 - -if we can solve TSP, we can solve Hamiltonian circuit

82

Bin Packing

 $\ensuremath{\textbf{Theorem}}$. No algorithm for online bin packing can always give an optimal solution.

Proof. Let $\varepsilon>0$ be small (say, $\varepsilon=1/64$), and consider an input sequence of m items of length $\frac{1}{2}-\varepsilon$ followed by m items of length $\frac{1}{2}+\varepsilon$.

Clearly, the optimal packing requires m bins.

Suppose online algorithm A yields this optimal solution. Since it is online, A must place each of the first m items in a separate bin.

Now give A an input sequence of just m items of length $\frac{1}{2}-\varepsilon.$ A will behave as before, placing each in a separate bin, thus using m bins.

However, the optimal packing in this case requires only $\lceil \, m/2\,\rceil$ bins.

Bin Packing

Since an online algorithm never knows when the input might end, any performance guarantee for the algorithm must hold whenever an item is binned.

Theorem. There are inputs that cause any online bin packing algorithm to use at least $\frac{4}{3}$ the optimal number of bins.

 ${\rm Proof.}$ Suppose not. Then there is an online algorithm A that always uses less than $\frac{4}{3}$ the optimal number of bins.

Let m be even. Apply A to an input sequence of m items of length $\frac{1}{2}-\varepsilon$ followed by m items of length $\frac{1}{2}+\varepsilon.$

Consider the situation after A has processed item m (the last of the smaller items). Suppose A has used b bins at this point.

85

Bin Packing

Thus we have two inequalities that hold:

$$\frac{2b}{m} < \frac{4}{3} \quad \text{after the first } m \text{ items,} \\ \frac{2m-b}{m} < \frac{4}{3} \quad \text{after the last } m \text{ items.}$$

From the first inequality we obtain $b/m < \frac{2}{3}$, while from the second we obtain $b/m > \frac{2}{3}$, which is a contradiction.

87

Bin Packing

Theorem

Let m be the optimal number of bins required to pack a list of items. Then next fit never uses more than 2m bins. There exist sequences for which next fit uses 2m - 2 bins.

This is an example of an approximation result.

Next fit is an approximation algorithm; it is not guaranteed to produce optimal solutions, but there exists a bound on how poorly it can do, even though we do not know the optimal solution.

Bin Packing

We know that the optimal number of bins for the first m items is m/2, so our performance assumption means $b<\frac{4}{3}\frac{m}{2},$ or $2b/m<\frac{4}{3}.$

Now consider the situation when A is finished packing all 2m items. All the bins used after bin b can only contain one item since the inputs are the longer items.

The first b bins can have at most 2 items each, and the remaining bins have one item each, so packing all 2m items requires at least 2m-b bins.

Since the optimal packing requires m bins, the performance assumption means $2m-b<\frac{4}{3}m,$ or $(2m-b)/m<\frac{4}{3}.$

86

Bin Packing

The next fit heuristic: when binning an item, check to see if it fits in the bin with the last item binned. If it does, place the new item there; otherwise, start a new bin.

Input: 0.2, 0.5, 0.4, 0.7, 0.1, 0.3, 0.8.

 $B_1: 0.2, 0.5 \\ B_2: 0.4 \\ B_3: 0.7, 0.1 \\ B_4: 0.3 \\ B_5: 0.8$

88

Bin Packing

Proof: Suppose next fit uses b bins. Let L_{j} be the total length of items in bin j; then

 $L_1 + L_2 + \cdots + L_b = \text{total length of items} \leq m.$

Consider any adjacent bins B_j and $B_{j+1}.$ The combined lengths of the items in these two bins must be larger than 1; otherwise all of these items would have been placed in B_j :

$$L_j + L_{j+1} > 1.$$

For simplicity, assume b is even. Then

$$\begin{array}{rcl}m&\geq&\underbrace{(L_1+L_2)}_1+\underbrace{(L_3+L_4)}_2+\cdots+\underbrace{(L_{b-1}+L_{b-2}}_{b/2}\\\\&>&\underbrace{1+1+\cdots+1}_2=\frac{b}{2},\end{array}$$

2m.

Bin Packing

For the second part of the theorem, suppose $n,\,{\rm the}$ number of items, is divisible by 4, and choose

$$\ell_i = \left\{ \begin{array}{ll} 1/2 & \text{if } i \text{ is odd,} \\ 2/n & \text{if } i \text{ is even.} \end{array} \right.$$

The optimal packing requires n/4 bins containing 2 items of length 1/2 and one bin containing the n/2 items of length 2/n.

Next fit, on the other hand, uses n/2 bins, with each bin containing one long and one short item.

91

Bin Packing

Theorem

Let m be the optimal number of bins required to pack a list of items. Then first fit never uses more than $\frac{17}{10}m+\frac{7}{10}$ bins. There exist sequences for which next fit uses $\frac{17}{10}(m-1)$ bins.

Consider a sequence of 6N items of size $\frac{1}{7} + \varepsilon$, followed by 6N items of size $\frac{1}{3} + \varepsilon$, followed by 6N items of size $\frac{1}{2} + \varepsilon$.

First fit will require 10N bins. However, if we apply next fit, we obtain a packing that requires only 6N bins.

93

Bin Packing

In offline bin packing, we can first sort the items, longest to shortest, and then pack them according to first fit or best fit.

The corresponding heuristics are called first fit decreasing and best fit decreasing (more properly they should be first fit nonincreasing and best fit nonincreasing).

First fit decreasing for: 0.8, 0.7, 0.5, 0.4, 0.3, 0.2, 0.1.

 B_1 : 0.8, 0.2 B_2 : 0.7, 0.3 B_3 : 0.5, 0.4, 0.1

In this case we get the optimal solution.



92

Bin Packing

The best fit heuristic: when binning an item, scan over all the bins in order and place the new item in the tightest spot.

Input: 0.2, 0.5, 0.4, 0.7, 0.1, 0.3, 0.8.

 $B_1: 0.2, 0.5, 0.1$ $B_2: 0.4$ $B_3: 0.7, 0.3$ $B_4: 0.8$

In general, the worst-case behavior is the same as first fit.

94

Bin Packing

Theorem

Let m be the optimal number of bins required to pack a list of items. Then first fit decreasing never uses more than $\frac{11}{9}m + \frac{6}{9}$ bins. There exist sequences for which first fit decreasing uses $\frac{11}{9}m + \frac{6}{9}$ bins.