

Topic Notes: Brute-Force Algorithms

Our first category of algorithms are called *brute-force algorithms*.

Levitin defines brute force as a straightforward approach, usually based directly on the problem statement and definitions of the concepts involved.

We have already seen a few examples:

- consecutive integer checking approach for finding a GCD
- matrix-matrix multiplication

Another is the computation of a^n by multiplying by a n times.

Brute-force algorithms are not usually clever or especially efficient, but they are worth considering for several reasons:

- The approach applies to a wide variety of problems.
- Some brute-force algorithms are quite good in practice.
- It may be more trouble than it's worth to design and implement a more clever or efficient algorithm over using a straightforward brute-force approach.

Brute-Force Sorting

One problem we will return to over and over is that of *sorting*. We will first consider some brute-force approaches.

We will usually look at sorting arrays of integer values, but the algorithms can be used for other comparable data types.

Bubble Sort

Right at the start of our class, we looked at a very intuitive sort. We just go through our array, looking at pairs of values and swapping them if they are out of order.

It takes $n - 1$ “bubble-ups”, each of which can stop sooner than the last, since we know we bubble up one more value to its correct position in each iteration. Hence the name *bubble sort*.

The version we came up with during the first lab looked like this:

```

ALGORITHM BUBBLESORT(A)
  //Input: an array A[0..n - 1]
  for i ← 0..n - 2 do
    for j ← 0..n - 2 - i do
      if A[j + 1] < A[j] then
        swap A[j + 1] and A[j]

```

The size metric is n , the size of the input array.

The basic operation is either the comparison or the swap inside the innermost loop. The comparison happens every time, while the swap only happens when necessary to reorder a pair of adjacent elements. Remember that a swap involves three assignments, which would be more expensive than the individual comparisons.

The best, average, and worst case for the number of comparisons is all the same. But for swaps, it differs. Best case, the array is already sorted and we need not make any swaps. Worst case, every comparison requires a swap. For an average case, we would need more information about the likelihood of a swap being needed to do any exact analysis.

So we proceed by counting comparisons, and the summation will also give us the worst case for the number of swaps.

$$\begin{aligned}
 C(n) &= \sum_{i=0}^{n-2} \sum_{j=0}^{n-2-i} 1 \\
 &= \sum_{i=0}^{n-2} [(n-2-i) - 0 + 1] \\
 &= \sum_{i=0}^{n-2} (n-1-i) \\
 &= \frac{(n-1)n}{2} \in \Theta(n^2).
 \end{aligned}$$

So we do $\Theta(n^2)$ comparisons. We swap, potentially, after each one of these, a worse case behavior of $\Theta(n^2)$ swaps.

Selection Sort

A simple improvement on the bubble sort is based on the observation that one pass of the bubble sort gets us closer to the answer by moving the largest unsorted element into its final position. Other elements are moved “closer” to their final position, but all we can really say for sure after a single pass is that we have positioned one more element.

So why bother with all of those intermediate swaps? We can just search through the unsorted part of the array, remembering the index of (and hence, the value of) the largest element we’ve seen so far, and when we get to the end, we swap the element in the last position with the largest element

we found. This is the *selection sort*.

ALGORITHM SELECTIONSORT(A)

//Input: an array $A[0..n - 1]$

for $i \leftarrow 0..n - 2$ **do**

$min \leftarrow i$

for $j \leftarrow i + 1..n - 1$ **do**

if $A[j] < A[min]$ **then**

$min \leftarrow j$

 swap $A[i]$ and $A[min]$

The number of comparisons is our basic operation here.

$$\begin{aligned}
 C(n) &= \sum_{i=0}^{n-2} \sum_{j=i+1}^{n-1} 1 \\
 &= \sum_{i=0}^{n-2} [(n-1) - (i+1) + 1] \\
 &= \sum_{i=0}^{n-2} (n-1-i) \\
 &= \frac{(n-1)n}{2} \in \Theta(n^2).
 \end{aligned}$$

Here, we do the same number of comparisons, but at most $n - 1 \in \Theta(n)$ swaps.

Sequential Search

We will cover this topic as an in-class exercise handout.

Brute-Force String Match

The *string matching* problem involves searching for a *pattern* (substring) in a string of *text*.

The basic procedure:

1. Align the pattern at beginning of the text
2. Moving from left to right, compare each character of the pattern to the corresponding character in the text until
 - all characters are found to match (successful search); or
 - a mismatch is detected

3. While pattern is not found and the text is not yet exhausted, realign the pattern one position to the right and repeat Step 2

The result is either the index in the text of the first occurrence of the pattern, or indices of all occurrences. We will look only for the first.

Written in pseudocode, our brute-force string match:

```

ALGORITHM BRUTEFORCESTRINGMATCH( $T, P$ )
  //Input: a text string  $T[0..n - 1]$ 
  //Input: a pattern string  $P[0..m - 1]$ 
  for  $i \leftarrow 0..n - m$  do
     $j \leftarrow 0$ 
    while  $j < m$  and  $P[j] = T[i + j]$  do
       $j \leftarrow j + 1$ 
    if  $j = m$  then
      return  $i$ 
  return  $-1$ 

```

The sizes of the input strings, m for the pattern and n for the text, are the problem size parameters and the basic operation is the element to element comparisons.

We have significant differences among the best, average, and worst cases. Best case, we find the pattern at the start of the text and we need only m comparisons: $\Theta(m)$. In the worst case, we encounter “near misses” at each starting location in the text, requiring about n searches of size m , resulting in $\Theta(nm)$. On average, however, most non-matches are likely to be detected quickly: in the first element or two, leading to an average case behavior of $\Theta(n)$.

We will consider improvements for string matching later in the semester.

In class exercise: Exercise 3.1.4, p. 102

Exhaustive Search

An *exhaustive search* is a brute force solution to a problem involving search for an element with a special property, usually among combinatorial objects such as permutations, combinations, or subsets of a set.

The typical approach involves:

1. Generation of a list of all potential solutions to the problem in a systematic manner (*e.g.*, our permutations from the first problem set – we will see others).
2. Evaluation of the potential solutions one by one, disqualifying infeasible ones and, for an optimization problem, keeping track of the best one found so far.
3. Announcing the solution when the search ends.

We will look briefly at a few examples of problems that can be solved with exhaustive search.

Traveling Salesman

The *traveling salesman problem (TSP)* asks us, given n cities with known distances between each pair, find the shortest tour that passes through all the cities exactly once before returning to the starting city.

Solutions to this kind of problem have important practical applications (delivery services, census takers, etc.)

An alternative statement of the problem is to find the shortest Hamiltonian circuit in a weighted, connected graph. A *Hamiltonian circuit* is a cycle that passes through all of the vertices of the graph exactly once.

If we think about it a bit, we can see that it does not matter where our tour begins and ends – it's a cycle – so we can choose any city/vertex as the starting point and consider all permutation of the other $n - 1$ vertices.

So an exhaustive search would allow us to consider all $(n - 1)!$ permutations, compute the cost (total distance traveled) for each, returning the lowest cost circuit found.

The text has a small example in Fig. 3.7, p. 117, and we will go over a similar one in class.

Knapsack Problem

Another well-known problem is the *knapsack problem*.

Given n items with weights w_1, w_2, \dots, w_n and values v_1, v_2, \dots, v_n , what is the most valuable subset of items that can fit into a knapsack that can hold a total weight W .

The exhaustive search here involves considering all possible subsets of items and finding the subset whose total weight is no more than W with the highest value.

For n items, this means we have to generate 2^n subsets, giving us that as a lower bound on our cost of $\Omega(2^n)$.

An example of this problem with a knapsack capacity of 16:

item	weight	value
1	2	20
2	5	30
3	10	50
4	5	10

All possible subsets (not including the empty set, which would be the answer only if all items were too heavy to be in the knapsack even alone):

subset	weight	value
{1}	2	20
{2}	5	30
{3}	10	50
{4}	5	10
{1,2}	7	50
{1,3}	12	70
{1,4}	7	30
{2,3}	15	80
{2,4}	10	40
{3,4}	15	60
{1,2,3}	17	too heavy
{1,2,4}	12	60
{1,3,4}	17	too heavy
{2,3,4}	20	too heavy
{1,2,3,4}	22	too heavy

So the winner is {2,3} with a total weight of 15 and a total value of 80.

Assignment Problem

Our final exhaustive search example is the *assignment problem*. It may be stated as the assignment of n jobs among n people, one job per person, where the cost of assigning person i to job j is given by $C[i, j]$, and we wish to minimize the total cost across all possible assignments.

For example, consider this cost matrix:

	Job 0	Job 1	Job 2	Job 3
Person 0	9	2	7	8
Person 1	6	4	3	7
Person 2	5	8	1	8
Person 3	7	6	9	4

Our exhaustive search here involves the generation of all possible assignments, computing the total cost of each, and selection of the lowest cost.

How many assignments as possible? Well, we can assign the first person any of n jobs. Then, there are $n - 1$ to choose for the second person, $n - 2$ for the third, and so on, until there is just one remaining choice for the last person: a total of $n!$.

We can list the solutions by generating all permutations of the numbers 1, 2, ..., n (sound familiar?), and treating the number in each position as the job to assign to the person at that position.

Exhaustive Search Wrapup

- Exhaustive-search algorithms run in a realistic amount of time only on very small instances.


- Sometimes, much better approaches exist (and we will see some of those soon).
- But sometimes, an exhaustive search is the only known way to be guaranteed an exact solution.

Graph Traversals

We return to graph structures for our next group of algorithms. In particular, we consider the *graph traversal* problem, that of visiting all vertices in a graph.

The two main approaches are the *depth-first search (DFS)* and the *breadth-first search (BFS)*.

We will examine the ideas using the example graph we have seen before.



In either case, we will choose a starting vertex, and visit all other vertices in the same connected component of the graph as that starting vertex. Either traversal will visit all vertices in the connected component, but each will visit the vertices in a different order.

Depth-first Traversal

A depth-first traversal (DFT) is produced by conducting a depth-first search (DFS). It proceeds by moving from the last visited vertex to an adjacent unvisited one. If no unvisited adjacent vertex is available, the traversal backtracks to a previously visited vertex which does have an unvisited adjacent vertex. When we have backtracked all the way to the starting vertex and no further adjacent vertices remain unvisited, the algorithm terminates.

For the example graph, the following is a DFT starting at A:

A, B, C, E, D, G, F

Note that we do not visit the disconnected vertex H.

If we wanted to perform a DFT that includes all vertices, we could restart the algorithm with an unvisited vertex.

Levitin has an algorithm on p. 123-124. Note the implicit use of a stack (here, the call stack to manage the recursion).

The following will perform a DFT of the vertices in the same connected component as the starting vertex with an explicit stack. The “visit” can mean to perform some computation on that vertex, which could simply be to print out or number the vertices.

```

ALGORITHM DFS( $G, s$ )
  //Input: a Graph  $G = \{V, E\}$ , a starting vertex  $s$ 
   $L \leftarrow$  a new empty stack
  for  $v : V$  do
     $v.visited \leftarrow false$ 
   $L.push(s)$ 
  while  $L$  not empty do
     $v \leftarrow L.pop()$ 
    if  $v.visited = false$  then
      // visit  $v$ 
       $v.visited \leftarrow true$ 
      for  $w : \text{adjacent to } v$  do
        if  $w.visited = false$  then
           $L.push(w)$ 

```

The efficiency class of this algorithm depends on which graph representation is used.

For an adjacency matrix, we have $\Theta(|V|^2)$, and for adjacency list, we have $\Theta(|V| + |E|)$.

The intuition behind this is that for each vertex in the graph, we have to visit each incident edge. With an adjacency matrix, this involves looking at each of $|V|$ matrix entries (including those which are null, indicating that such an edge does not exist). For the adjacency list, we have the exact list of edges readily available, so across all vertices, we visit $\Theta(|E|)$ total edges.

Breadth-first Traversal

A breadth-first traversal (BFT) is performed by breadth-first search (BFS). It proceeds by visiting all adjacent vertices of the last visited vertex before proceeding to any subsequent adjacencies.

For the example graph, the following is a BFT starting at A:

A, B, C, F, G, E, D

See the Levitin algorithm for BFS on p. 126.

Note that this proceeds by visiting the starting vertex, then the immediate neighbors of the starting vertex, then the neighbors of those neighbors, and so on.

We accomplish this with the queue structure. Note that the DFS algorithm above can be converted to a BFS simply by swapping the stack for a queue:

```

ALGORITHM BFS( $G, s$ )
  //Input: a Graph  $G = \{V, E\}$ , a starting vertex  $s$ 

```



```
L ← a new empty queue
for v : V do
    v.visited ← false
    L.enqueue(s)
while L not empty do
    v ← L.dequeue()
    if v.visited = false then
        // visit v
        v.visited ← true
        for w : adjacent to v do
            if w.visited = false then
                L.enqueue(w)
```

The efficiency classes of this algorithm are the same as those of DFS, and depend on which graph representation is used.

For an adjacency matrix, we again have $\Theta(|V|^2)$, and for adjacency list, we have $\Theta(|V| + |E|)$.