

ECE 374 B Complete: Cheatsheet

1 Languages and strings

Languages

- An *alphabet* Σ is a **finite** set of symbols.

Definitions A *string* in Σ^* is a **finite** sequence of symbols in Σ .

- A *language* is L is a set of strings over some alphabet.

All languages represent mathematical problems.

Example: multiplication of two integers:

$$L_{MULT_2} = \left\{ \begin{array}{ccc} 1 \times 1|1, & 1 \times 2|2, & 1 \times 3|3, \dots \\ 2 \times 1|2, & 2 \times 2|4, & 2 \times 3|6, \dots \\ \vdots & \vdots & \vdots \\ n \times 1|n, & n \times 2|2n, & n \times 3|3n, \dots \end{array} \right\} \quad (1)$$

- For languages A, B the *concatenation* of A, B is $AB = \{xy \mid x \in A, y \in B\}$.
- For languages A, B , their *union* is $A \cup B$, *intersection* is $A \cap B$, and *difference* is $A \setminus B$ (also written as $A - B$).
- For language $A \subseteq \Sigma^*$ the *complement* of A is $\bar{A} = \Sigma^* \setminus A$.
- Σ^n is the set of all strings of length n .
- $\Sigma^* = \cup_{n \geq 0} \Sigma^n$ is the set of all strings over Σ .
- $\Sigma^+ = \cup_{n \geq 1} \Sigma^n$ is the set of non-empty strings over Σ .

Language operations

Strings

- The *length* of a string w (denoted by $|w|$) is the number of symbols in w .

- For integer $n \geq 0$, Σ^n is set of all strings over Σ of length n . Σ^* is the set of all strings over Σ .

Definitions

- Σ^* is the set of all strings of all lengths including empty string.
- ϵ is a *string* containing no symbols.
- \emptyset is the *empty set*. It contains no strings.

- If x and y are strings then xy denotes their concatenation. Recursively:

- $xy = y$ if $x = \epsilon$
- $xy = a(wy)$ if $x = aw$

- v is *substring* of $w \iff$ there exist strings x, y such that $w = xvy$.

- If $x = \epsilon$ then v is a *prefix* of w
- If $y = \epsilon$ then v is a *suffix* of w


- A *subsequence* of a string $w = w_1w_2 \dots w_n$ is either a subsequence of $w_2 \dots w_n$ or w_1 followed by a subsequence of $w_2 \dots w_n$.

- If w is a string then w^n is defined inductively as follows: $w^n = \epsilon$ if $n = 0$ or $w^n = ww^{n-1}$ if $n > 0$

String operations

2 Overview of language complexity

Overview



Grammar	Languages	Production Rules	Automaton	Examples
Type-0	recursively enumerable	$\gamma \rightarrow \alpha$ (no constraints)	Turing machine	$L = \{w \mid w \text{ is a TM which halts}\}$
Type-1	context-sensitive	$\alpha A \beta \rightarrow \alpha \gamma \beta$	linear bounded nondeterministic Turing machine	$L = \{a^n b^n c^n \mid n > 0\}$
Type-2	context-free	$A \rightarrow \alpha$	nondeterministic pushdown automata	$L = \{a^n b^n \mid n > 0\}$
Type-3	regular	$A \rightarrow aB$	finite state machine	$L = \{a^n \mid n > 0\}$

Meaning of symbols:

- a - terminal
- A, B - variables
- α, β, γ - strings in $\{a \cup A\}^*$ where α, β are maybe empty, γ is never empty

^aTable borrowed from Wikipedia: https://en.wikipedia.org/wiki/Chomsky_hierarchy

3 Regular languages

Regular language - overview

A language is regular if and only if it can be obtained from finite languages by applying

- union,
- concatenation or
- Kleene star

finitely many times. All regular languages are representable by regular grammars, DFAs, NFAs and regular expressions.

Regular expressions

Useful shorthand to denotes a language.

A regular expression r over an alphabet Σ is one of the following:

Base cases:

- \emptyset the language \emptyset
- ϵ denotes the language $\{\epsilon\}$
- a denote the language $\{a\}$

Inductive cases: If r_1 and r_2 are regular expressions denoting languages L_1 and L_2 respectively (i.e., $L(r_1) = L_1$ and $L(r_2) = L_2$) then,

- $r_1 + r_2$ denotes the language $L_1 \cup L_2$
- $r_1 \cdot r_2$ denotes the language $L_1 L_2$
- r_1^* denotes the language L_1^*

Examples:

- 0^* - the set of all strings of 0s, including the empty string
- $(00000)^*$ - set of all strings of 0s with length a multiple of 5
- $(0 + 1)^*$ - set of all binary strings

Nondeterministic finite automata

NFAs are similar to DFAs, but may have more than one transition destination for a given state/character pair.

An NFA N accepts a string w iff some accepting state is reached by N from the start state on input w .


The language accepted (or recognized) by an NFA N is denoted $L(N)$ and defined as $L(N) = \{w \mid N \text{ accepts } w\}$.

A nondeterministic finite automaton (NFA) $N = (Q, \Sigma, s, A, \delta)$ is a five tuple where

- Q is a finite set whose elements are called states
- Σ is a finite set called the input alphabet
- $\delta : Q \times \Sigma \cup \{\epsilon\} \rightarrow \mathcal{P}(Q)$ is the transition function (here $\mathcal{P}(Q)$ is the power set of Q)
- s and Σ are the same as in DFAs

Example:

- $Q = \{q_0, q_1, q_2, q_3\}$
- $\Sigma = \{0, 1\}$



For NFA $N = (Q, \Sigma, \delta, s, A)$ and $q \in Q$, the ϵ -reach(q) is the set of all states that q can reach using only ϵ -transitions.

Inductive definition of $\delta^* : Q \times \Sigma^* \rightarrow \mathcal{P}(Q)$:

- if $w = \epsilon$, $\delta^*(q, w) = \epsilon\text{-reach}(q)$
- if $w = a$ for $a \in \Sigma$, $\delta^*(q, a) = \epsilon\text{reach}\left(\bigcup_{p \in \epsilon\text{-reach}(q)} \delta(p, a)\right)$
- if $w = ax$ for $a \in \Sigma, x \in \Sigma^*$: $\delta^*(q, w) = \epsilon\text{reach}\left(\bigcup_{p \in \epsilon\text{-reach}(q)} \left(\bigcup_{r \in \delta^*(p, a)} \delta^*(r, x)\right)\right)$

Regular closure

Regular languages are closed under union, intersection, complement, difference, reversal, Kleene star, concatenation, etc.

Deterministic finite automata


DFAs are finite state machines that can be represented as a directed graph or in terms of a tuple.

The language accepted (or recognized) by a DFA M is denoted by $L(M)$ and defined as $L(M) = \{w \mid M \text{ accepts } w\}$.

A deterministic finite automaton (DFA) $M = (Q, \Sigma, s, A, \delta)$ is a five tuple where

- Q is a finite set whose elements are called states
- Σ is a finite set called the input alphabet
- $\delta : Q \times \Sigma \rightarrow Q$ is the transition function
- $s \in Q$ is the start state
- $A \subseteq Q$ is the set of accepting/final states

Example:



Every string has a unique walk along a DFA. We define the extended transition function as $\delta^* : Q \times \Sigma^* \rightarrow Q$ defined inductively as follows:

- $\delta^*(q, w) = q$ if $w = \epsilon$
- $\delta^*(q, w) = \delta^*(\delta(q, a), x)$ if $w = ax$.

Can create a larger DFA from multiple smaller DFAs. Suppose

- $L(M_0) = \{w \text{ has an even number of 0s}\}$ (pictured above) and
- $L(M_1) = \{w \text{ has an even number of 1s}\}$.

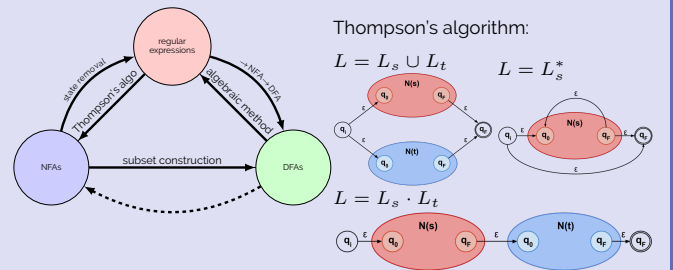
$L(M_C) = \{w \text{ has even number of 0s and 1s}\}$

Suppose $M_0 = (Q_0, \Sigma, s_0, A_0, \delta_0)$ and $M_1 = (Q_1, \Sigma, s_1, A_1, \delta_1)$. Then

- $Q = Q_0 \times Q_1 = \{(q_0, q_1) \mid q_0 \in Q_0, q_1 \in Q_1\}$
- $s = (s_0, s_1)$
- $\delta : Q \times \Sigma \rightarrow Q$, where $\delta((q_0, q_1), a) = (\delta_0(q_0, a), \delta_1(q_1, a))$
- $A = \{(q_0, q_1) \mid q_0 \in A_0 \text{ and } q_1 \in A_1\}$

Regular language equivalences

A regular language can be represented by a regular expression, regular grammar, DFA and NFA.



Arden's rule: If $R = Q + RP$ then $R = QP^*$.

Fooling sets

Some languages are not regular (Ex. $L = \{0^n 1^n \mid n \geq 0\}$).

Two states $p, q \in Q$ are distinguishable if there exists a string $w \in \Sigma^*$, such that

$$\delta^*(p, w) \in A \text{ and } \delta^*(q, w) \notin A.$$

or

Two states $p, q \in Q$ are equivalent if for all strings $w \in \Sigma^*$, we have that

$$\delta^*(p, w) \in A \iff \delta^*(q, w) \in A.$$

$$\delta^*(p, w) \notin A \text{ and } \delta^*(q, w) \in A.$$

For a language L over Σ a set of strings F (could be infinite) is a fooling set or distinguishing set for L if every two distinct strings $x, y \in F$ are distinguishable.

4 Context-free languages

Context-free languages

A language is context-free if it can be generated by a context-free grammar. A context-free grammar is a quadruple $G = (V, T, P, S)$

- V is a finite set of *nonterminal (variable) symbols*
- T is a finite set of *terminal symbols* (alphabet)
- P is a finite set of *productions*, each of the form $A \rightarrow \alpha$ where $A \in V$ and α is a string in $(V \cup T)^*$. Formally, $P \subseteq V \times (V \cup T)^*$.
- $S \in V$ is the *start symbol*

Example: $L = \{ww^R \mid w \in \{0, 1\}^*\}$ is described by $G = (V, T, P, S)$ where V, T, P and S are defined as follows:

- $V = \{S\}$
- $T = \{0, 1\}$
- $P = \{S \rightarrow \epsilon \mid 0S0 \mid 1S1\}$
(abbreviation for $S \rightarrow \epsilon, S \rightarrow 0S0, S \rightarrow 1S1$)
- $S = S$

Pushdown automata

A pushdown automaton is an NFA with a stack.


The language $L = \{0^n 1^n \mid n \geq 0\}$ is recognized by the pushdown automaton:

A *nondeterministic pushdown automaton (PDA)* $P = (Q, \Sigma, \Gamma, \delta, s, A)$ is a **six** tuple where

- Q is a finite set whose elements are called *states*
- Σ is a finite set called the *input alphabet*
- Γ is a finite set called the *stack alphabet*
- $\delta : Q \times (\Sigma \cup \{\epsilon\}) \times (\Gamma \cup \{\epsilon\}) \rightarrow \mathcal{P}(Q \times (\Gamma \cup \{\epsilon\}))$ is the *transition function*
- s is the start state
- A is the set of accepting states

In the graphical representation of a PDA, transitions are typically written as $(\text{input read}), (\text{stack pop}) \rightarrow (\text{stack push})$.

A CFG can be converted to a pushdown automaton.



The PDA to the right recognizes the language described by the following grammar:

$$S \rightarrow 0S1 \mid \epsilon$$

Context-free closure

Context-free languages are closed under union, concatenation, and Kleene star.


They are **not** closed under intersection or complement.

5 Recursively enumerable languages

Turing Machines

Turing machine is the simplest model of computation.


- Input written on (infinite) one sided tape.
- Special blank characters.
- Finite state control (similar to DFA).
- Every step: Read character under head, write character out, move the head right or left (or stay).
- Every TM M can be encoded as a string $\langle M \rangle$



Transition Function: $\delta : Q \times \Gamma \rightarrow Q \times \Gamma \times \{\leftarrow, \rightarrow, \square\}$

$$\delta(q, c) = (p, d, \leftarrow)$$

- q : current state.
- c : character under tape head.
- p : new state.
- d : character to write under tape head
- \leftarrow : Move tape head left.



6 Recursion

Simple recursion

- **Reduction:** solve one problem using the solution to another.
- **Recursion:** a special case of reduction - reduce problem to a *smaller* instance of *itself* (self-reduction).

Definitions

- Problem instance of size n is reduced to *one or more* instances of size $n - 1$ or less.
- For termination, problem instances of small size are solved by some other method as *base cases*

Arguably the most famous example of recursion. The goal is to move n disks one at a time from the first peg to the last peg.

Pseudocode: Tower of Hanoi

```
Hanoi(n, src, dest, tmp):
  if (n > 0) then
    Hanoi(n - 1, src, tmp, dest)
    Move disk n from src to dest
    Hanoi(n - 1, tmp, dest, src)
```

Tower of Hanoi

Recurrences

Suppose you have a recurrence of the form $T(n) = rT(n/c) + f(n)$.

The *master theorem* gives a good asymptotic estimate of the recurrence. If the work at each level is:

Decreasing: $rf(n/c) = \kappa f(n)$ where $\kappa < 1$ $T(n) = O(f(n))$
 Equal: $rf(n/c) = f(n)$ $T(n) = O(f(n) \cdot \log_c n)$
 Increasing: $rf(n/c) = Kf(n)$ where $K > 1$ $T(n) = O(n^{\log_c r})$

Some useful identities:

- Sum of integers: $\sum_{k=1}^n k = \frac{n(n+1)}{2}$
- Geometric series closed-form formula: $\sum_{k=0}^n ar^k = \frac{1-r^{n+1}}{1-r}$
- Logarithmic identities: $\log(ab) = \log a + \log b$, $\log(a/b) = \log a - \log b$, $a^{\log_c b} = b^{\log_c a}$ ($a, b, c > 1$).

Backtracking

Backtracking is the algorithm paradigm involving guessing the solution to a single step in some multi-step process and recursing backwards if it doesn't lead to a solution. For instance, consider the longest increasing subsequence (LIS) problem. You can either check all possible subsequences:

Pseudocode: LIS - Naive enumeration

```
algLISNaive(A[1..n]):
  maxmax = 0
  for each subsequence B of A do
    if B is increasing and |B| > max then
      max = |B|
  return max
```

On the other hand, we don't need to generate every subsequence; we only need to generate the subsequences that are increasing:

Pseudocode: LIS - Backtracking

```
LIS_smaller(A[1..n], x):
  if n = 0 then return 0
  max = LIS_smaller(A[1..n-1], x)
  if A[n] < x then
    max = max {max, 1 + LIS_smaller(A[1..(n-1)], A[n])}
  return max
```

Divide and conquer

Divide and conquer is an algorithm paradigm involving the decomposition of a problem into the same subproblem, solving them separately and combining their results to get a solution for the original problem.

Algorithm	Runtime	Space
Sorting algorithms	Mergesort	$O(n \log n)$ $O(n)$ (if optimized)
	Quicksort	$O(n^2)$ $O(n \log n)$ if using MoM

We can divide and conquer multiplication like so:

$$bc = 10^n b_L c_L + 10^{n/2} (b_L c_R + b_R c_L) + b_R c_R.$$

We can rewrite the equation as:

$$bc = b(x)c(x) = (b_L x + b_R)(c_L x + c_R) = (b_L c_L)x^2 + ((b_L + b_R)(c_L + c_R) - b_L c_L - b_R c_R)x + b_R c_R,$$

Karatsuba's algorithm

Its running time is $O(n^{\log_2 3}) = O(n^{1.585})$.

Linear time selection

The *median of medians* (MoM) algorithms give a element that is larger than $\frac{3}{10}$'s and smaller than $\frac{7}{10}$'s of the array elements. This is used in the linear time selection algorithm to find element of rank k .

Pseudocode: Quickselect with median of medians

```
Median-of-medians(A, i):
  sublists = [A[jj+5] for j ← 0, 5, ..., len(A)]
  medians = sorted(sublist)[len(sublist)/2]
  for sublist ∈ sublists

  // Base case
  if len(A) ≤ 5 return sorted(a)[i]

  // Find median of medians
  if len(medians) ≤ 5
    pivot = sorted(medians)[len(medians)/2]
  else
    pivot = Median-of-medians(medians, len/2)

  // Partitioning step
  low = l; for j ∈ A if j < pivot
  high = l; for j ∈ A if j > pivot

  k = len(low)
  if i < k
    return Median-of-medians(low, i)
  else if i > k
    return Median-of-medians(low, i-k-1)
  else
    return pivot
```

Dynamic programming

Dynamic programming (DP) is the algorithm paradigm involving the computation of a recursive backtracking algorithm iteratively to avoid the recomputation of any particular subproblem.

Longest increasing subsequence

The longest increasing subsequence problem asks for the length of a longest increasing subsequence in a unordered sequence, where the sequence is assumed to be given as an array. The recurrence can be written as:

$$LIS(i, j) = \begin{cases} 0 & \text{if } i = 0 \\ LIS(i-1, j) & \text{if } A[i] \geq A[j] \\ \max \begin{cases} LIS(i-1, j) \\ 1 + LIS(i-1, i) \end{cases} & \text{else} \end{cases}$$

Pseudocode: LIS - DP

LIS-Iterative($A[1..n]$):

```

A[n + 1] = ∞
for j ← 0 to n
    if A[i] ≤ A[j] then LIS[0][j] = 1

for i ← 1 to n - 1 do
    for j ← i to n - 1 do
        if A[i] ≥ A[j]
            LIS[i, j] = LIS[i - 1, j]
        else
            LIS[i, j] = max { LIS[i - 1, j],
                             1 + LIS[i - 1, i] }
return LIS[n, n + 1]
```

Edit distance

The edit distance problem asks how many edits we need to make to a sequence for it to become another one. The recurrence is given as:

$$Opt(i, j) = \min \begin{cases} \alpha_{x_i y_j} + Opt(i-1, j-1), \\ \delta + Opt(i-1, j), \\ \delta + Opt(i, j-1) \end{cases}$$

Base cases: $Opt(i, 0) = \delta \cdot i$ and $Opt(0, j) = \delta \cdot j$

Pseudocode: Edit distance - DP

EDIST($A[1..m], B[1..n]$)

```

for i ← 1 to m do M[i, 0] = iδ
for j ← 1 to n do M[0, j] = jδ
```

```

for i = 1 to m do
    for j = 1 to n do
```

$$M[i][j] = \min \begin{cases} COST[A[i]][B[j]] \\ \quad + M[i-1][j-1], \\ \delta + M[i-1][j], \\ \delta + M[i][j-1] \end{cases}$$

7 Graph algorithms

Graph basics

A graph is defined by a tuple $G = (V, E)$ and we typically define $n = |V|$ and $m = |E|$. We define (u, v) as the edge from u to v . Graphs can be represented as **adjacency lists**, or **adjacency matrices** though the former is more commonly used.

- path:** sequence of *distinct* vertices v_1, v_2, \dots, v_k such that $v_i v_{i+1} \in E$ for $1 \leq i \leq k-1$. The length of the path is $k-1$ (the number of edges in the path).
Note: a single vertex u is a path of length 0.
- cycle:** sequence of *distinct* vertices v_1, v_2, \dots, v_k such that $(v_i, v_{i+1}) \in E$ for $1 \leq i \leq k-1$ and $(v_k, v_1) \in E$. A single vertex is not a cycle according to this definition.
Caveat: Sometimes people use the term cycle to also allow vertices to be repeated; we will use the term *tour*.
- A vertex u is *connected* to v if there is a path from u to v .
- The *connected component* of u , $con(u)$, is the set of all vertices connected to u .
- A vertex u can *reach* v if there is a path from u to v . Alternatively v can be reached from u . Let $rch(u)$ be the set of all vertices reachable from u .

Directed acyclic graphs

Directed acyclic graphs (dags) have an intrinsic ordering of the vertices that enables dynamic programming algorithms to be used on them.

A *topological ordering* of a dag $G = (V, E)$ is an ordering \prec on V such that if $(u, v) \in E$ then $u \prec v$.

Pseudocode: Kahn's algorithm

```


Kahn( $G(V, E), u$ ):
toposort ← empty list
for  $v \in V$ :
    in(v) ← |{u | u → v ∈ E}|
while  $v \in V$  that has in(v) = 0:
    Add v to end of toposort
    Remove v from V
    for  $v$  in  $u \rightarrow v \in E$ :
        in(v) ← in(v) - 1
return toposort
```

Running time: $O(n + m)$

- A dag may have multiple topological sorts.
- A topological sort can be computed by DFS, in particular by listing the vertices in decreasing post-visit order.

Strongly connected components

- Given G , u is *strongly connected* to v if $v \in rch(u)$ and $u \in rch(v)$.
- A *maximal* group of G : vertices that are all strongly connected to one another is called a strong component.



Pseudocode: Metagraph - linear time

Metagraph($G(V, E)$):

```

Compute rev(G) by brute force
ordering ← reverse postordering of V in rev(G)
by DFS(rev(G), s) for any vertex s
Mark all nodes as unvisited
for each u in ordering do
    if u is not visited and u ∈ V then
        S_u ← nodes reachable by u by DFS(G, u)
        Output S_u as a strong connected component
G(V, E) ← G - S_u
```

DFS and BFS

Pseudocode: Explore (DFS/BFS)

```

Explore( $G, u$ ):
  for  $i \leftarrow 1$  to  $n$ :
    Visited[ $i$ ]  $\leftarrow$  False
  Add  $u$  to ToExplore and to  $S$ 
  Visited[ $u$ ]  $\leftarrow$  True
  Make tree  $T$  with root as  $u$ 
  while  $B$  is non-empty do
    Remove node  $x$  from  $B$ 
    for each edge  $(x, y)$  in  $Adj(x)$  do
      if Visited[ $y$ ] = False
        Visited[ $y$ ]  $\leftarrow$  True
        Add  $y$  to  $B, S, T$  (with  $x$  as parent)
  
```

Note:

- If B is a queue, *Explore* becomes BFS.
- If B is a stack, *Explore* becomes DFS.

Pre/post numbering

Pre and post numbering aids in analyzing the graph structure. By looking at the numbering we can tell if a edge (u, v) is a:

- *Forward edge*: $pre(u) < pre(v) < post(v) < post(u)$
- *Backward edge*: $pre(v) < pre(u) < post(u) < post(v)$
- *Cross edge*: $pre(u) < post(u) < pre(v) < post(v)$

Minimum Spanning Tress

Some notes on minimum spanning trees:

- Tree = undirected graph in which any two vertices are connected by exactly one path.
- Tree = a connected graph with no cycles.
- Sub-graph H of G is *spanning* for G , if G and H have same connected components.
- A minimum spanning tree is composed of all the safe edges in the graph
- An edge $e = (u, v)$ is a *safe* edge if there is some partition of V into S and $V \setminus S$ and e is the unique minimum cost edge crossing S (one end in S and the other in $V \setminus S$).
- An edge $e = (u, v)$ is an *unsafe* edge if there is some cycle C such that e is the unique maximum cost edge in C .
- All edges are safe or unsafe.

Pseudocode: Boruvka's algorithm: $O(m \log(n))$

```

 $T$  is  $\emptyset$  (*  $T$  will store edges of a MST *)
while  $T$  is not spanning do
   $X \leftarrow \emptyset$ 
  for each connected component  $S$  of  $T$  do
    add to  $X$  the cheapest edge between  $S$  and  $V \setminus S$ 
  Add edges in  $X$  to  $T$ 
return the set  $T$ 
  
```

Pseudocode: Kruskal's algorithm: $(m + n) \log(m)$ (using Union-Find structure)

```

Sort edges in  $E$  based on cost
 $T$  is empty (*  $T$  will store edges of a MST *)
each vertex  $u$  is placed in a set by itself
while  $E$  is not empty do
  pick  $e = (u, v) \in E$  of minimum cost
  if  $u$  and  $v$  belong to different sets
    add  $e$  to  $T$ 
    merge the sets containing  $u$  and  $v$ 
return the set  $T$ 
  
```

Pseudocode: Prim's algorithm: $(n) \log(n) + m$ (using Priority Queue)

```

 $T \leftarrow \emptyset, S \leftarrow \emptyset, s \leftarrow 1$ 
 $\forall v \in V(G) : d(v) \leftarrow \infty, p(v) \leftarrow \emptyset$ 
 $d(s) \leftarrow 0$ 
while  $S \neq V$  do
   $v = \arg \min_{u \in V \setminus S} d(u)$ 
   $T = T \cup \{vp(v)\}$ 
   $S = S \cup \{v\}$ 
  for each  $u$  in  $Adj(v)$  do
     $d(u) \leftarrow \min \begin{cases} d(u) \\ c(vu) \end{cases}$ 
    if  $d(u) = c(vu)$  then
       $p(u) \leftarrow v$ 
return  $T$ 
  
```

Shortest paths

Dijkstra's algorithm:

Find minimum distance from vertex s to **all** other vertices in graphs *without* negative weight edges.

Pseudocode: Dijkstra

```

for  $v \in V$  do
   $d(v) \leftarrow \infty$ 
 $X \leftarrow \emptyset$ 
 $d(s, s) \leftarrow 0$ 
for  $i \leftarrow 1$  to  $n$  do
   $v \leftarrow \arg \min_{u \in V - X} d(u)$ 
   $X = X \cup \{v\}$ 
  for  $u$  in  $Adj(v)$  do
     $d(u) \leftarrow \min \{d(u), d(v) + \ell(v, u)\}$ 
return  $d$ 
  
```

Running time: $O(m + n \log n)$ (if using a Fibonacci heap as the priority queue)

Bellman-Ford algorithm:

Find minimum distance from vertex s to **all** other vertices in graphs *without* negative cycles. It is a DP algorithm with the following recurrence:

$$d(v, k) = \begin{cases} 0 & \text{if } v = s \text{ and } k = 0 \\ \infty & \text{if } v \neq s \text{ and } k = 0 \\ \min \begin{cases} \min_{u \in E} \{d(u, k-1) + \ell(u, v)\} \\ d(v, k-1) \end{cases} & \text{else} \end{cases}$$

Base cases: $d(s, 0) = 0$ and $d(v, 0) = \infty$ for all $v \neq s$.

Pseudocode: Bellman-Ford

```

for each  $v \in V$  do
   $d(v) \leftarrow \infty$ 
 $d(s) \leftarrow 0$ 

for  $k \leftarrow 1$  to  $n - 1$  do
  for each  $v \in V$  do
    for each edge  $(u, v) \in E$  do
       $d(v) \leftarrow \min \{d(v), d(u) + \ell(u, v)\}$ 

return  $d$ 
  
```

Running time: $O(nm)$

Floyd-Warshall algorithm:

Find minimum distance from *every* vertex to *every* vertex in a graph *without* negative cycles. It is a DP algorithm with the following recurrence:

$$d(i, j, k) = \begin{cases} 0 & \text{if } i = j \\ \infty & \text{if } (i, j) \notin E \text{ and } k = 0 \\ \min \begin{cases} d(i, j, k-1) \\ d(i, k, k-1) + d(k, j, k-1) \end{cases} & \text{else} \end{cases}$$

Then $d(i, j, n - 1)$ will give the shortest-path distance from i to j .

Pseudocode: Floyd-Warshall

```

Metagraph( $G(V, E)$ ):
  for  $i \in V$  do
    for  $j \in V$  do
       $d(i, j, 0) \leftarrow \ell(i, j)$ 
      (*  $\ell(i, j) \leftarrow \infty$  if  $(i, j) \notin E, 0$  if  $i = j$  *)

  for  $k \leftarrow 0$  to  $n - 1$  do
    for  $i \in V$  do
      for  $j \in V$  do
         $d(i, j, k) \leftarrow \min \begin{cases} d(i, j, k-1), \\ d(i, k, k-1) + d(k, j, k-1) \end{cases}$ 

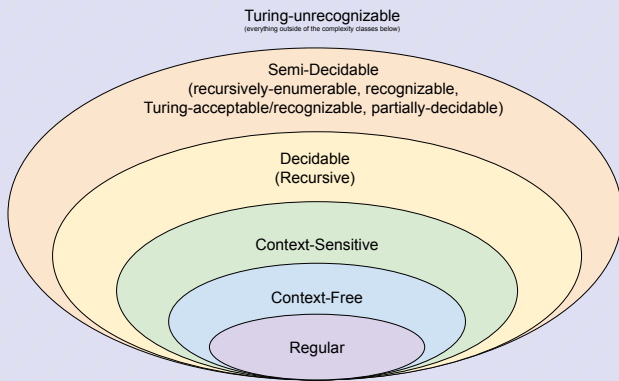
  for  $v \in V$  do
    if  $d(i, i, n - 1) < 0$  then
      return  $\exists$  negative cycle in  $G$ 

  return  $d(\cdot, \cdot, n - 1)$ 
  
```


Running time: $\Theta(n^3)$

Complexity Classes

Computational Complexity Classes



Algorithmic Complexity Classes (assuming $P \neq NP$)




Reductions

A general methodology to prove impossibility results.

- Start with some *known* hard problem X
- Reduce X to your favorite problem Y

If Y can be solved then so can $X \implies Y$. But we know X is hard so Y has to be hard too. On the other hand if we know Y is easy, then X has to be easy too.

The Karp reduction, $X \leq_P Y$ suggests that there is a polynomial time reduction from X to Y .



Assuming

- $R(n)$: running time of R
- $Q(n)$: running time of A_Y

Running time of A_X is $O(Q(R(n)))$

Sample NP-complete problems

CIRCUITSAT: Given a boolean circuit, are there any input values that make the circuit output TRUE?

3SAT: Given a boolean formula in conjunctive normal form, with exactly three distinct literals per clause, does the formula have a satisfying assignment?

INDEPENDENTSET: Given an undirected graph G and integer k , what is there a subset of vertices $\geq k$ in G that have no edges among them?

CLIQUE: Given an undirected graph G and integer k , is there a complete complete subgraph of G with more than k vertices?

kPARTITION: Given a set X of kn positive integers and an integer k , can X be partitioned into n , k -element subsets, all with the same sum?

3COLOR: Given an undirected graph G , can its vertices be colored with three colors, so that every edge touches vertices with two different colors?

HAMILTONIANPATH: Given graph G (either directed or undirected), is there a path in G that visits every vertex exactly once?

HAMILTONIANCYCLE: Given a graph G (either directed or undirected), is there a cycle in G that visits every vertex exactly once?

LONGESTPATH: Given a graph G (either directed or undirected, possibly with weighted edges) and an integer k , does G have a path $\geq k$ length?

• Remember a **path** is a sequence of distinct vertices $[v_1, v_2, \dots, v_k]$ such that an edge exists between any two vertices in the sequence. A **cycle** is the same with the addition of an edge $(v_k, v_1) \in E$. A **walk** is a path except the vertices can be repeated.

• A formula is in conjunction normal form if variables are or'ed together inside a clause and then clauses are and'ed together: $((x_1 \vee x_2 \vee x_3) \wedge (\overline{x_2} \vee x_4 \vee x_5))$. Disjunctive normal form is the opposite $((x_1 \wedge x_2 \wedge x_3) \vee (\overline{x_2} \wedge x_4 \wedge x_5))$.

Sample undecidable problems

ACCEPTONINPUT: $A_{TM} = \{ \langle M, w \rangle \mid M \text{ is a TM and } M \text{ accepts on } w \}$

HALTSONINPUT: $Halt_{TM} = \{ \langle M, w \rangle \mid M \text{ is a TM and halts on input } w \}$

HALTONBLANK: $Halt_{B_{TM}} = \{ \langle M \rangle \mid M \text{ is a TM \& } M \text{ halts on blank input} \}$

EMPTINESS: $E_{TM} = \{ \langle M \rangle \mid M \text{ is a TM and } L(M) = \emptyset \}$

EQUALITY: $EQ_{TM} = \left\{ \langle M_A, M_B \rangle \mid \begin{array}{l} M_A \text{ and } M_B \text{ are TM's} \\ \text{and } L(M_A) = L(M_B) \end{array} \right\}$