

## RESCHEDULING THE RESCHEDULED CLASS

Does May 1st work? Email if not.

May 1st is during reading period, but will let us avoid using final exam time for makeup.

# MODELS OF COMPUTATION

The Turing Machine: the first theoretical model of computation.

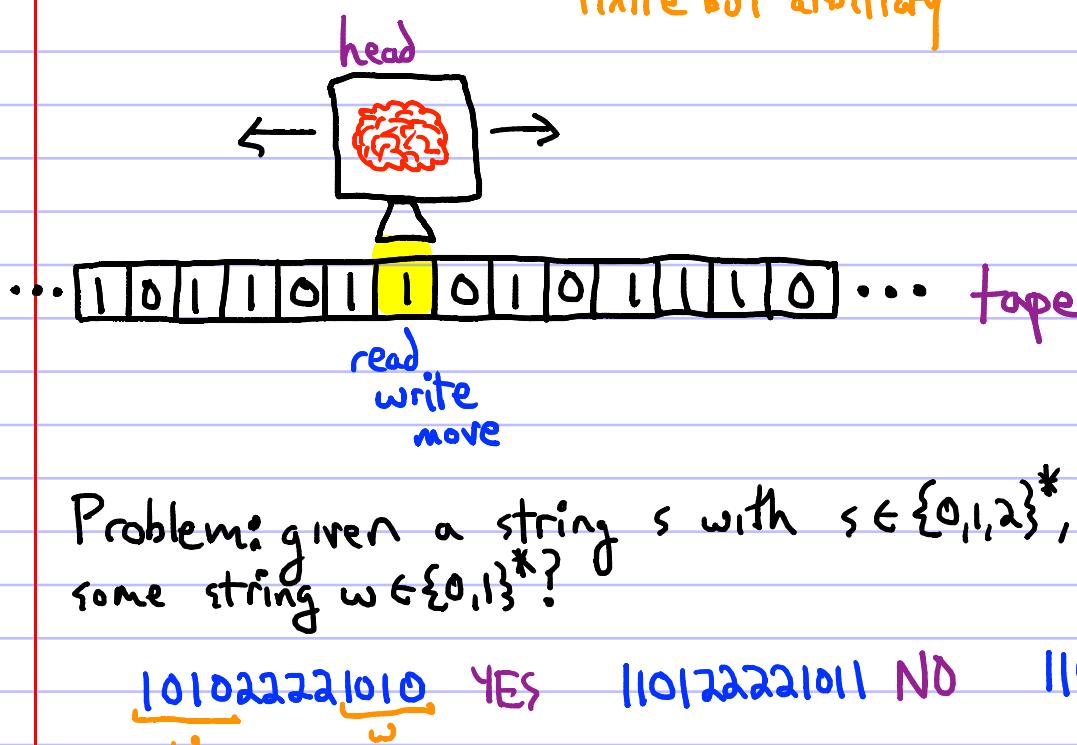
Two parts:

1. A head containing an algorithm and program counter.

2. A tape containing an input, and infinite additional space.

A fixed algorithm of  $O(1)$  size

finite but arbitrary



Problem: given a string  $s$  with  $s \in \{0,1,2\}^*$ , is  $s = w^2 w^{1\omega}$  for some string  $w \in \{0,1\}^*$ ?

101022221010 YES    110122221011 NO    1102222210 NO

(Turing machine) Algorithm: 1. while (symbols remain):

$\Theta(n^2)$  total steps

2. read first symbol  $0/1=a$ , erase it
  3. go right until first  $2$ , erase it
  4. go right until first  $0/1=b$ , erase it
  5. check  $a=b$ .
  6. if anything failed, return NO
  7. return YES
- $\left.\frac{2}{3}n\right]$

Thm [Hennie '65]: Any TM algorithm for this problem takes  $\Theta(n^2)$  steps

Proof Idea: Left & right third modifications must happen in synchrony, otherwise machine is wrong. Synchrony takes  $\Omega(n)$  time per change.

A "regular" algorithm :  $\Theta(n)$  [for  $i$  from 1 to  $\frac{n}{3}$ :

$\Theta(n)$  total time

$\Theta(1)$  [check that  $s[i] == s[\frac{2}{3}n+i]$ ]  
 $\Theta(1)$  [check  $s[i+\frac{n}{3}] == 2$ ]  
if anything failed, return NO  
return YES

A "regular" algorithm and running time analysis assumes **RAM model**:

**R**andom **A**ccess **M**achine. Look anywhere in  $O(1)$  time.

Turing machine vs. RAM: which is a better model?

DISTANCE MATTERS

DISTANCE IRRELEVANT

Goal: model that only permits reasonable operations that incur reasonable time/space costs. Reasonable = natural, realistic, simple.

RAM is a family of models: locality-independent memory access cost.

What about operation costs? How long does computing  $n!$  take?

$t=1$

for  $i \in [1 \dots n]$ :  $\Theta(n)$   
 $t = t \cdot i$  ]  $\Theta(1)$ ?

return  $t$

$\Theta(n)$  time to compute  $n!$ ?

Let  $a, b \in \mathbb{N}$ . How long does " $a * b$ " take?

It is typical to assume  $\Theta(1)$ .

$n$  can be written in  $\log(n)$  bits.

$n!$  takes  $\Theta(n \log n)$  bits.

$\Theta(n)$  time to compute an output of size  $\Theta(n \log n)$ ? Not possible.

Karatsuba algorithm:  $\tilde{\Theta}(\log n)^{1.585}$  time to multiply two integers  $i_1, i_2$  with  $n = \max(i_1, i_2)$ . Not  $\Theta(1)$

## Naive multiplication:

Given two binary numbers  $a = a_n a_{n-1} \dots a_0$  &  $b = b_n b_{n-1} \dots b_0$ :

$t = 0$

$\Theta(n)$  [ for  $i$  from 0 to  $n$ :  
 $\Theta(n)$  [ if  $a_i == 1$ :  
 $t += b \ll i$  //  $b * 2^i = b \ll i$  ]  
 return  $t$  ]  $\Theta(n^2)$  time

$$n=4, a = \underset{0}{1010} = 22, b = \underset{1}{1110} = 29$$

$$\begin{aligned} a * b &= a_0 \cdot [b \ll 0] + a_1 \cdot [b \ll 1] + a_2 \cdot [b \ll 2] + a_3 \cdot [b \ll 3] + a_4 \cdot [b \ll 4] \\ &= b \ll 1 + b \ll 2 + b \ll 4 \end{aligned}$$

$$\begin{array}{r} 111010 \\ 1110100 \\ \hline 10011110 \end{array} = 512 + 64 + 32 + 16 + 8 + 4 + 2 = 638 = 22 \cdot 29$$


---

To go faster, use following idea:

$$a = \underline{\underline{110101}} \underline{\underline{1011}} = \underline{\underline{11010}} \cdot 2^5 + \underline{\underline{1011}} = a_L \cdot 2^5 + a_R$$

$$b = \underline{\underline{111000101}} = \underline{\underline{1110}} \cdot 2^5 + \underline{\underline{00101}} = b_L \cdot 2^5 + b_R$$

$$a * b = (a_L \cdot 2^5 + a_R)(b_L \cdot 2^5 + b_R) = a_L b_L \cdot 2^{10} + (a_L b_R + a_R b_R) \cdot 2^5 + a_R b_R$$

$$a_L b_R + a_R b_R = (a_L + a_R)(b_L + b_R) - a_L b_L - a_R b_R$$

$$x = a_L b_L \quad y = (a_L + a_R)(b_L + b_R) - z - x \quad z = a_R b_R$$

$$a * b = x \cdot 2^{10} + y \cdot 2^5 + z \quad ] 3 \text{ mults.}$$

## Karatsuba Multiplication Algorithm: [Karatsuba 1960]

Given two binary numbers  $a = a_n \dots a_1 a_0$ ,  $b = b_n \dots b_1 b_0$ :

- ① Let  $a_L = a_n \dots a_{n/2}$ ,  $a_R = a_{n/2-1} \dots a_0$ .  
 ② Let  $b_L = b_n \dots b_{n/2}$ ,  $b_R = b_{n/2-1} \dots b_0$ .

③ Compute  $x = a_L b_L$ ,  $y = a_R b_R$ ,  $z = (a_L + a_R)(b_L + b_R)$  using Karatsuba's.

④ Return  $x \ll n + y \ll \frac{n}{2} + z$ .

Running time:  $T(n) = 3T(\frac{n}{2}) + O(n) = O(n^{\log 3})$ .

Compare to:  $T(n) = 4T(\frac{n}{2}) + O(n) = O(n^2)$ .

Generalized Karatsuba (Toom-Cook): [Toom 1963], [Cook 1966]

1. Break  $a, b$  into  $k$  pieces each.
2. Compute pairwise products of  $k$  pieces.
3. Do  $2k-1$  multiplications and some additions to compute  $a \times b$ .

Running time:  $O(f(k) \cdot n^{\log_k(2k-1)})$ . ( $f(k)$  is addition time)

Karatsuba:  $k=2$ . Why not let  $k \rightarrow n$ ?  $f(k)$  grows fast.

(Can do even better:  $O(n \log n \log \log n)$  [Schönhage, Strassen 1971] hard

$O(n \log n \cdot 2^{O(\log^* n)})$  [Fürer 2007] not implemented

Conjecture: any algorithm takes  $\Omega(n \log n)$  time. [SS 1971]

Given two  $n$ -bit numbers:

- Multiplication takes  $\tilde{\Theta}(n \log n)$
- Addition takes  $\Theta(n)$
- Equality,  $>$ ,  $<$  take  $\Theta(n)$

All basic operations take about linear time in  $n$  bits.

How many bits are allowed in an operand? Many options.

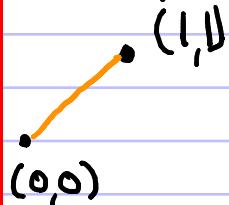
- Arbitrary (pay  $\Theta(1)$ ) real RAM [Blum, Shub, Smale 1989], [Shamos 1978]
- Arbitrary (pay  $\Theta(\# \text{bits})$ ) bit-level RAM von Emde Boas, Willard
- $\log(\# \text{input bits})$  (pay  $\Theta(1)$ ) word RAM

As expected, real RAM is simple but not realistic:

If  $+, -, *, /, \lfloor \cdot \rfloor$  are allowed operations, can solve QBF [As hard as any problem in polynomial time. [Schönhage 1979]] using polynomial space. Up to  $2^{(n^c)}$  time.

Also true if only  $+, -, *$  and  $\wedge, \vee, \neg$  are allowed. [Hartmanis, Simon 1974]

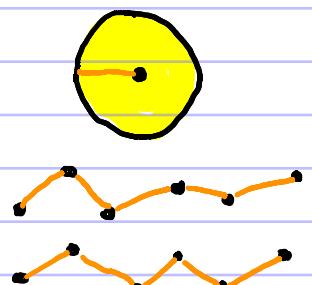
So why use real RAM model ever? Geometric problems.



Compute distance between two points?

Compute area of a circle?

Decide which of two paths is shorter?



Sum-of-squares problem: given two sequences of integers  $a_1, a_2, \dots, a_n$  &  $b_1, b_2, \dots, b_n$ , is  $\sum_{i=1}^n a_i^2 < \sum_{i=1}^n b_i^2$ ?

Solvable in  $O(n)$  time in real RAM (used in shortest path w/ obstacles)

Not known to be in NP in word RAM or bit-level RAM.

Common in computational geometry to use real RAM [Shamos 1978] but with the caveat that there is some imprecision, and with unspoken agreement not to cheat & do PSPACE-hard stuff. realistic by self-regulation

**Bit-level RAM**: reasonable but annoying. As input gets bigger, indexing into arrays changes, etc.:

```
for (i=0; i<n; ++i)
{
    A[i] = A[i] + 1;
}
```

i grows in size. total work of all "++i":

$$\sum_{i=1}^n \log(i) = \Theta(n \log n)$$

for-loops no longer run in linear time. bleh

**Word RAM**: Constant-time to do operations on a "word".

A word has size  $\omega \geq \log n$  for input of size  $n$ .

Also assume input numbers each fit into a word.  $(0, 1, \dots, 2^\omega - 1)$

Word RAM tends to allow faster algorithms and data structures than real RAM because of bounded universe & word-level parallelism.

Sorting:  $O(n \cdot 2^\omega)$  using counting sort

$O(n \cdot \frac{\omega}{\log n})$  using radix sort ( $\frac{\omega}{\log n}$  iterations of counting sort)

$O(n \cdot \log \omega)$  using dynamic tree with  $O(\log \omega)$  insert/delete  
Von Emde Boas

Counting Sort:  $O(n)$  for  $\omega = \log n$

Also  $O(n)$  for all  $\omega = \lceil 2(\log^{2+\epsilon} n) \rceil$  possible

Radix Sort:  $O(n)$  for  $\omega = O(\log n)$

Von Emde Boas Tree Sort:  $O(n \lg \lg n)$  for  $\omega = \log^c n$ .

## Word RAM and other models with words

Cell probe

More powerful

Transdichotomous RAM

Word RAM

Less powerful

All three models can do the same computations (all) but have different costs associated/different sets of  $O(1)$ -bit operations.

Cell probe: Pay  $O(1)$  for read/write of a word.  
All other computation is free.

Transdichotomous RAM: Pay  $O(1)$  for read/write/<sup>finite set</sup>any operation on a word.

Word RAM: Pay  $O(1)$  for read/write/simple operation on a word.  
<sup>probably not \*</sup>

Lower bound in cell probe  $\Rightarrow$  same lower bound in 2 other models.

E.g. finding max element in array takes  $\Omega(n)$ .

## External Memory Model [Aggarwal, Vitter 1988]

Two levels of memory: fast and bounded cache  
slow and unbounded disk

Cache consists of blocks of  $B$  words  
with total space  $M \gg B$ .

Can only compute on cache memory, running time is  
# blocks written & read. Computation is free!

Like cell probe, but with a "pile of words".

Can find max/min in  $O(\lceil \frac{N}{B} \rceil)$  time.

Can build dynamic BST with insert()/delete()/find() in  $O(\log_{B+1} n)$ .

Lower bound of  $\Omega(\log_{B+1} n)$  for search():

1. Insertion location requires  $\log(n+1)$  bits to specify, info content

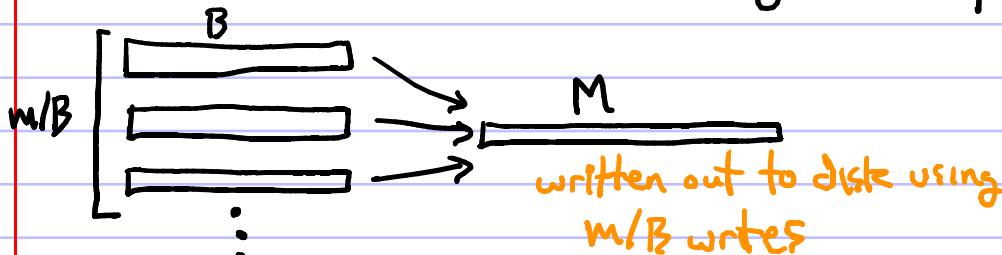
2. Each block read gives at most  $\log(B+1)$  bits. Where query fits

3. So at least  $\frac{\log(n+1)}{\log(B+1)} = \Omega(\log_{B+1} n)$  reads needed.

Can sort in  $O(\frac{N}{B} \log(\frac{N}{B}))$  using MergeSort()

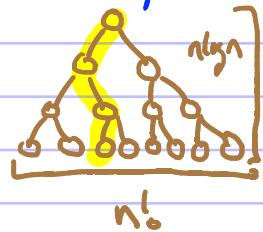
with a base case of size  $B$ .

Can sort in  $O(\frac{N}{B} \log_{M/B} \frac{N}{B})$  using  $\frac{M}{B}$ -way MergeSort()



A lower bound of  $\Omega\left(\frac{n}{B} \log_{M/B}\left(\frac{M}{B}\right)\right)$  for sorting:

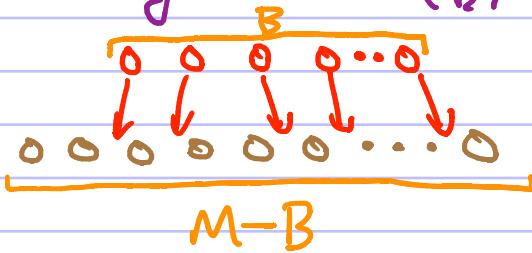
Recall that lower bound of  $\Omega(n \log n)$  for regular sorting says:  
 $n!$  permutations, each operation reduces these by a factor of  $\Sigma$ ,  
 $\Sigma \geq \log_2(n!) = \Omega(n \log n)$  operations for some permutation



Use the same idea w/special considerations.

1. Assume each block is sorted,  $O(n)$  time & reduces permutations by a factor of  $(B!)^{n/B}$

2. Each read of a block lets us reduce the # permutations by a factor of  $\binom{M}{B}$ , since new block can only be permuted with existing blocks in  $\binom{M}{B}$  ways:



$\binom{M}{B}$  ways to insert  $B$  things between  $M-B$  things

$$\text{So time is } \geq \log_{\binom{M}{B}}\left(\frac{n!}{(B!)^{n/B}}\right) \geq \log_{(M^B/B!)}\left(\frac{n!}{(B!)^{n/B}}\right)$$

$$\geq \frac{\log(n!) - \log((B!)^{n/B})}{\log(M^B) - \log(B!)} = \Omega\left(\frac{n \log n - \cancel{n} \log B}{B \log(M) - B \log B}\right)$$

$$= \Omega\left(\frac{n \log\left(\frac{M}{B}\right)}{B \log\left(\frac{M}{B}\right)}\right) = \Omega\left(\frac{n}{B} \log\left(\frac{M}{B}\right)\right) \text{ tight}$$

## Cache-oblivious Model [Prokop 1999]

External memory model **except**  $M, B$  are not known by algorithm.

Realistic:  $M, B$  are low-level & separation of concerns may hide them.

$M, B$  may be dynamic, code may run in VM moving around metal.  
Many  $M_s, B_s$ .

Analysis goal: optimality (e.g.  $O(\frac{A}{B} \log_{M/B} (\frac{A}{B}))$  sorting) **for all  $M, B$** .

How does block read/write happen? Automatically using an optimal replacement strategy.

Increasing  $M$  by  $O(1)$  factor makes strategy optimal.

Sorting still possible in  $O(\frac{A}{B} \log_{M/B} (\frac{A}{B}))$  [Frigo et al. 1999]  
assuming  $M = \Sigma(B^{1+\epsilon})$  for some  $\epsilon$ .