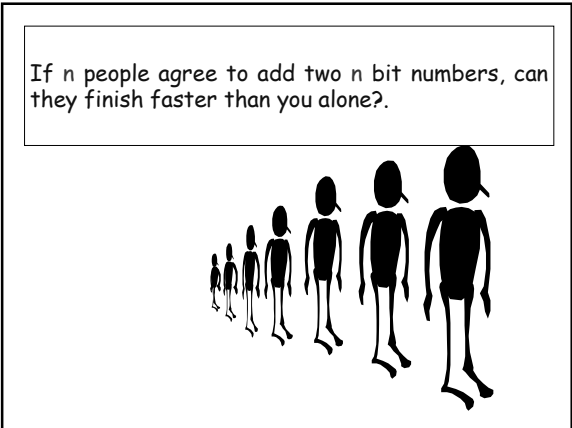
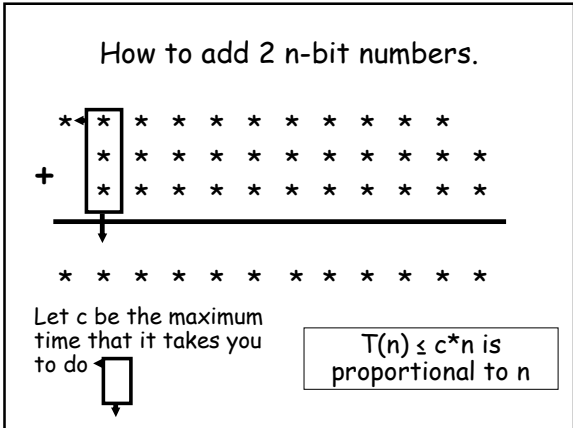



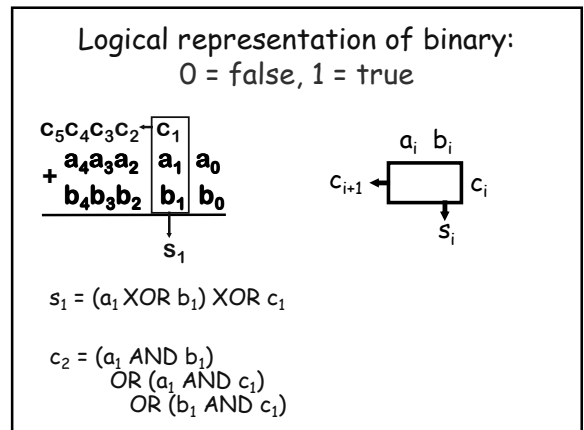

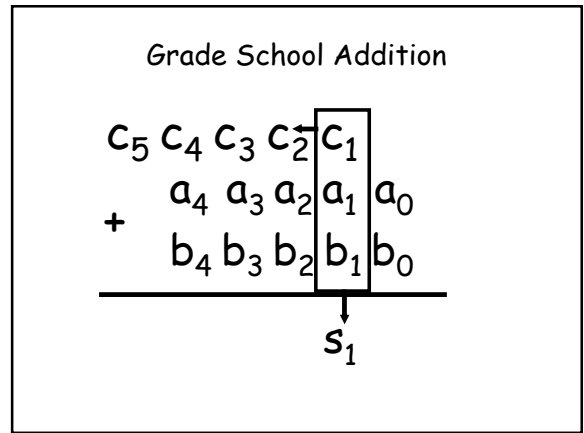
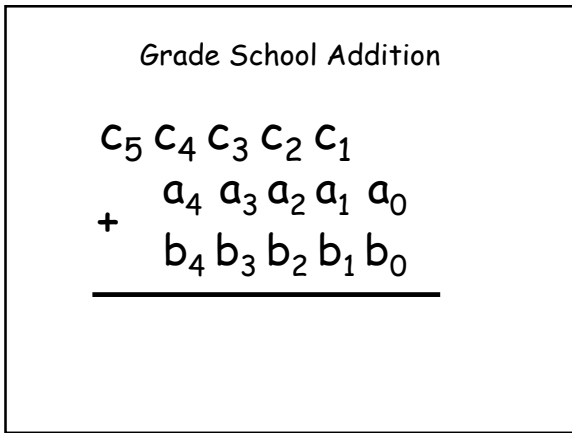
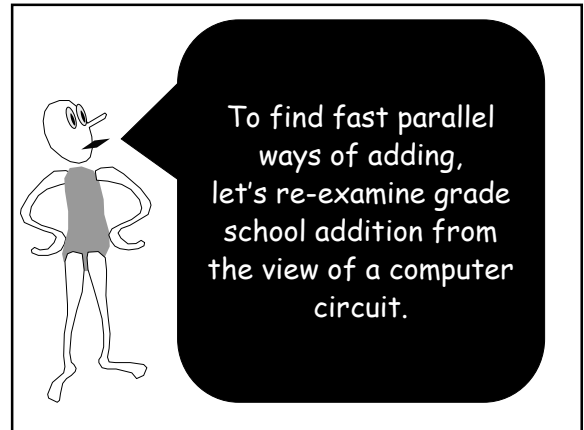

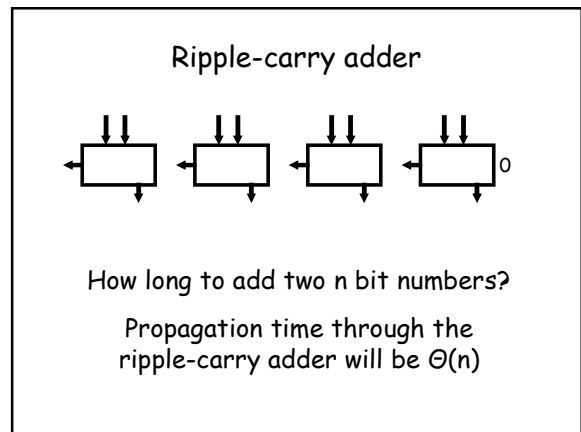
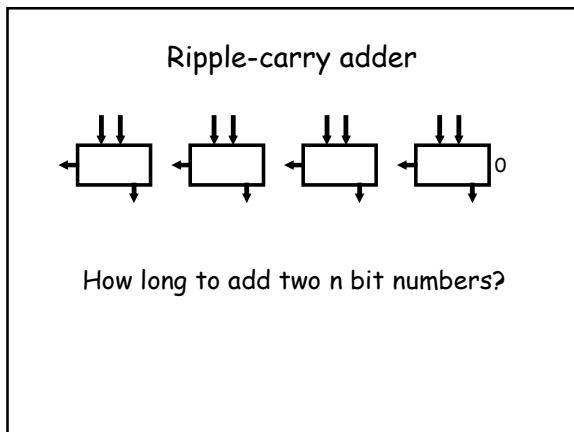
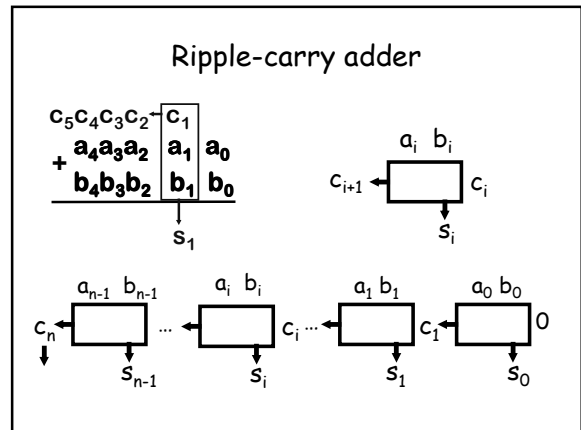



1. A Parallel Perspective
 2. Compression








Circuits compute things in parallel.

We can think of the propagation delay as **PARALLEL TIME**.

Is it possible to add two n bit numbers in less than linear parallel-time?


Darn those carries (again)!

If we knew the carries it would be very easy to do fast parallel addition




So how do we figure out the carries fast?

What do we know about the carry-out before we know the carry-in?



a	b	C_{out}
0	0	0
0	1	C_{in}
1	0	C_{in}
1	1	1

What do we know about the carry-out before we know the carry-in?




a	b	C_{out}
0	0	0
0	1	←
1	0	←
1	1	1

This is just a function of a and b. We can do this in parallel.



a	b	C_{out}
0	0	0
0	1	←
1	0	←
1	1	1


Idea: do this calculation first.



Note that this just took one step!

Now if we could only replace the ← by 0/1 values...

Idea: do this calculation first.



Idea: do this calculation first.

$$\begin{array}{r} 10\leftarrow\leftarrow\leftarrow\leftarrow\leftarrow 1000 \\ + 1011111101 \\ \hline 1000000110 \end{array}$$

Idea: do this calculation first.

$$\begin{array}{r} 10\leftarrow\leftarrow\leftarrow 111000 \\ + 1011111101 \\ \hline 1000000110 \end{array}$$

Idea: do this calculation first.

$$\begin{array}{r} 10111111000 \\ + 1011111101 \\ \hline 1000000110 \end{array}$$

Idea: do this calculation first.

$$\begin{array}{r} 10111111000 \\ + 1011111101 \\ \hline 10100000011 \end{array}$$

Once we have the carries, it takes only one more step:
 $s_i = (a_i \text{ XOR } b_i) \text{ XOR } c_i$

$$10\leftarrow\leftarrow\leftarrow\leftarrow\leftarrow 1\leftarrow\leftarrow 0$$

So, everything boils down to:
can we find a fast parallel
way to convert each position
to its final 0/1 value?

Called the "parallel prefix problem"



Prefix Sum Problem

Input: $X_{n-1}, X_{n-2}, \dots, X_1, X_0$

Output: $Y_{n-1}, Y_{n-2}, \dots, Y_1, Y_0$

where

$$Y_0 = X_0$$

$$Y_1 = X_0 + X_1$$

$$Y_2 = X_0 + X_1 + X_2$$

$$Y_3 = X_0 + X_1 + X_2 + X_3$$

$$Y_{n-1} = X_0 + X_1 + X_2 + X_3 + \dots + X_{n-1}$$

+ is Addition

Idea to get 10111111000

Can think of $10 \leftarrow \leftarrow \leftarrow \leftarrow \leftarrow 1 \leftarrow \leftarrow 0$ as

$(1 M (0 M (\leftarrow M (\leftarrow M (\leftarrow M (\leftarrow M (\leftarrow M (1 M (\leftarrow M 0))))))))$

for the operator M :

$$\leftarrow M x = x$$

$$1 M x = 1$$

$$0 M x = 0$$

M	0	1	\leftarrow
0	0	0	0
1	1	1	1
\leftarrow	0	1	\leftarrow

Associative Binary Operator


Binary Operator: an operation that takes two objects and returns a third.

$$\bullet A \spadesuit B = C$$


Associative:

$$\bullet (A \spadesuit B) \spadesuit C = A \spadesuit (B \spadesuit C)$$


Example circuitry ($n = 4$)



Divide, conquer, and glue for computing y_{n-1}



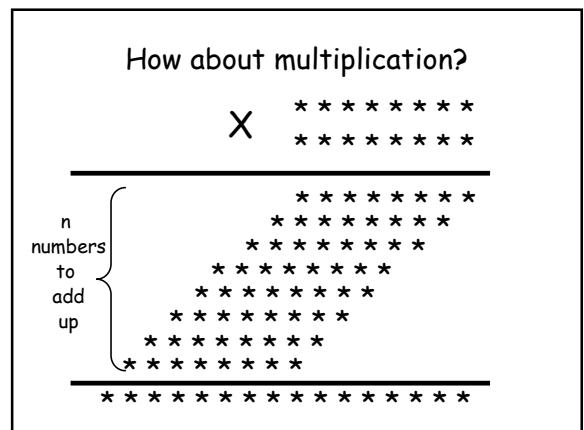
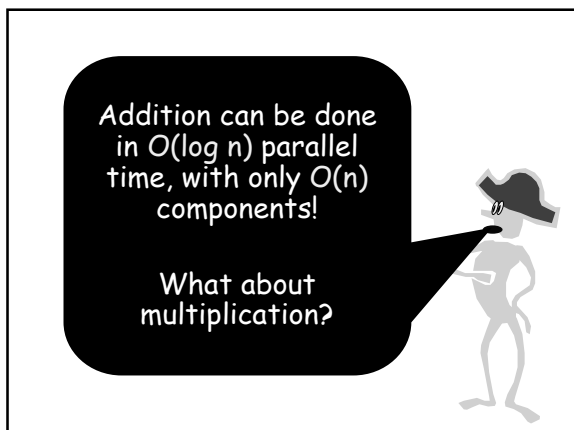
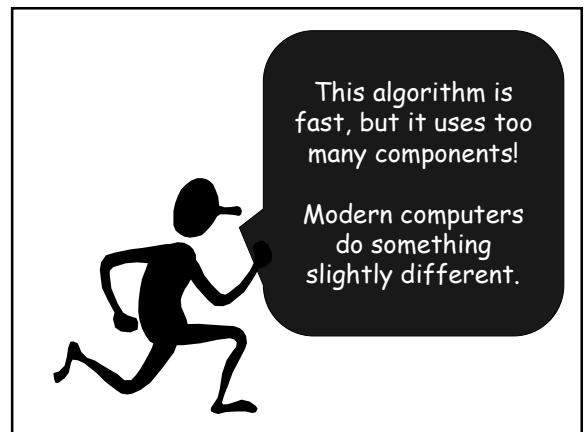
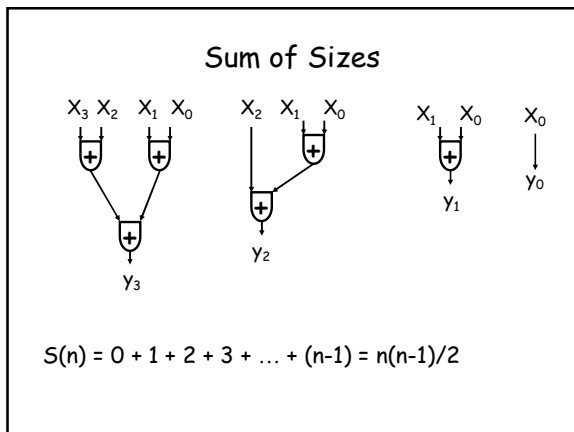
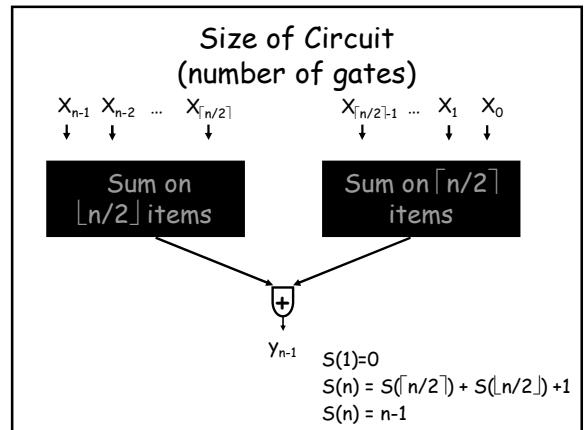
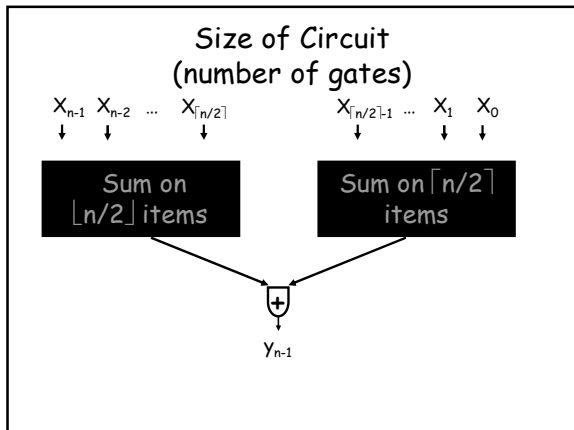
Divide, conquer, and glue for computing y_{n-1}



The parallel time taken
is $T(n) = \lceil \log_2 n \rceil$!

But how many
components does
this use? What is the
size of the circuit?



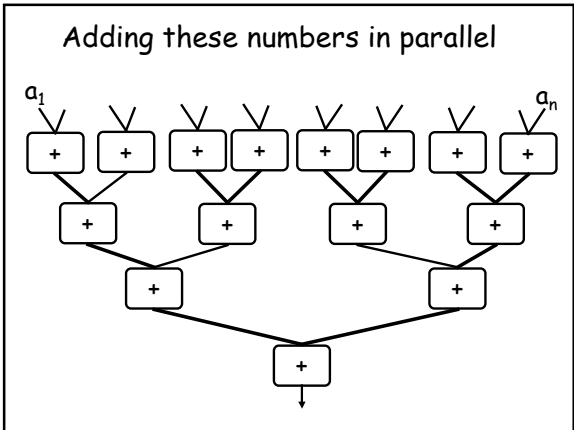


Grade School Multiplication

```

      X *****
      *****
a1  *****
a2  *****
a3  *****
.   *****
an  *****
      *****
  
```

We need to add n $2n$ -bit numbers:
 $a_1, a_2, a_3, \dots, a_n$



What is the depth of the circuit?

Each addition takes $O(\log n)$ parallel time

Depth of tree = $\log_2 n$

Total $O(\log n)^2$ parallel time

Can we do better?
 How about $O(\log n)$ parallel time?

How about multiplication?

Here's a really neat trick:

Let's think about how to add 3 numbers to make 2 numbers.

"Carry-Save Addition"

The sum of three numbers can be converted into the sum of 2 numbers in constant parallel time!

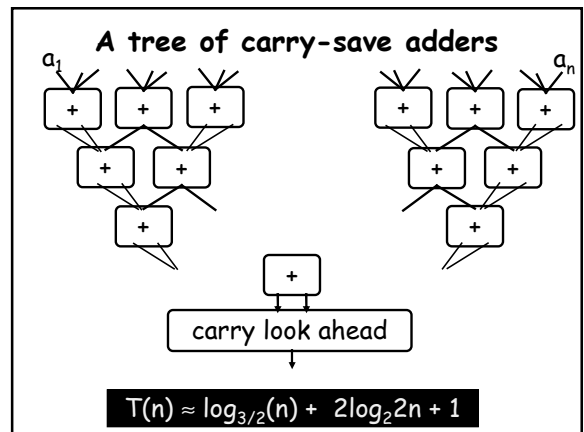
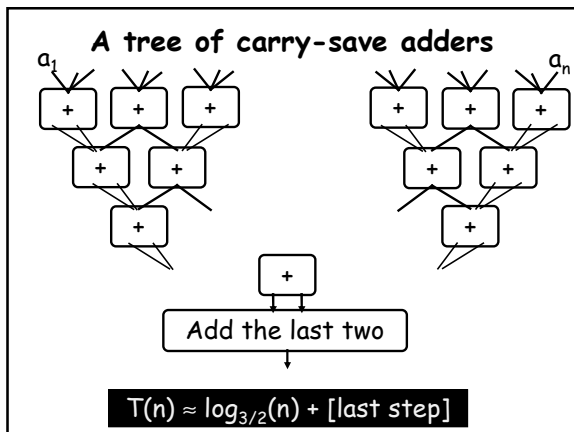
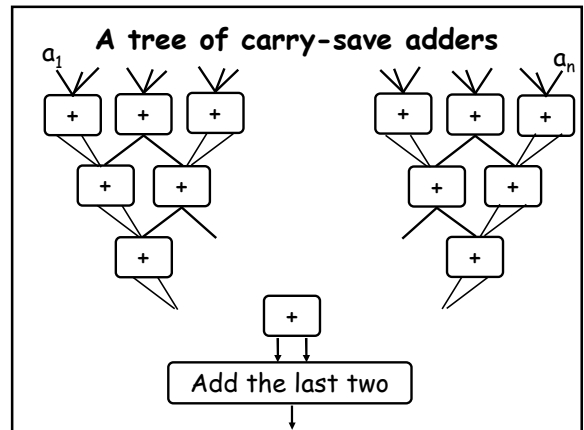
```

      1100111011
      + 1011111101
      + 1000000110
      -----
  
```


"Carry-Save Addition"

The sum of three numbers can be converted into the sum of 2 numbers in constant parallel time!

+ 1100111011	
+ 1011111101	
+ 1000000110	
+ 1111000000	XOR
1000111110	Carries



We can multiply in $O(\log n)$ parallel time too!

For a 64-bit word that works out to a parallel time of 22 for multiplication, and 13 for addition.

Addition requires linear time sequentially, but has a practical $2\log_2 n + 1$ parallel algorithm.


Multiplication (which is a lot harder sequentially) also has an $O(\log n)$ time parallel algorithm.

And this is how addition works on commercial chips.....

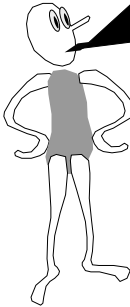
Processor	n	$2\log_2 n + 1$
80186	16	9
Pentium	32	11
Alpha	64	13

Excellent!
 Parallel time for:
 Addition = $O(\log n)$
 Multiplication = $O(\log n)$

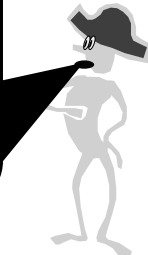
Hey, we forgot subtraction!




Most computers use two's complement representation to add and subtract integers.



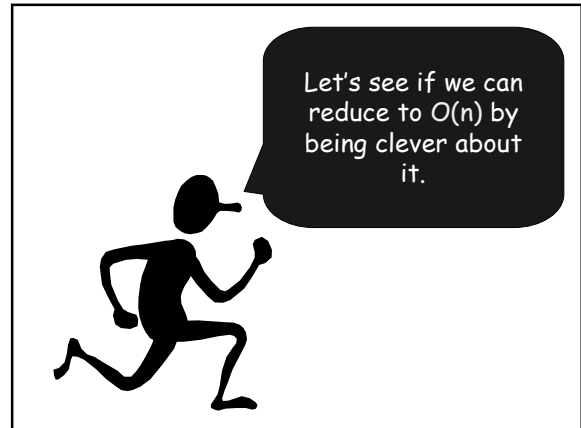
What about division?



Grade School Division



Suppose we have n bits of precision.
 Naïve method: n subtractions costing $2\log_2 n + 1$ each = $\Theta(n \log n)$ parallel time



Idea: use extended binary all through the computation!
 Then convert back at the end.



SRT division algorithm

1 0 1 1		1 1 1 0 1 1 0 1	1 1	2 ¹ r ⁶	2 ² r ⁻⁵
		-10 -1 -1			
		1 0 -1 1			
		-10 -1 -1			
		-2 0			
		-1 0 0 1			
		1 0 1 1			
		1 2			
		= 1 0 0 0			
		-1 0 -1 -1			
		0 -1 -1 1			

Rule: Each bit of quotient is determined by comparing first bit of divisor with first bit of dividend. Easy!

Time for n bits of precision in result:
 $= 3n + 2\log_2(n) + 1$

1 addition per bit Convert to standard representation by subtracting negative bits from positive.

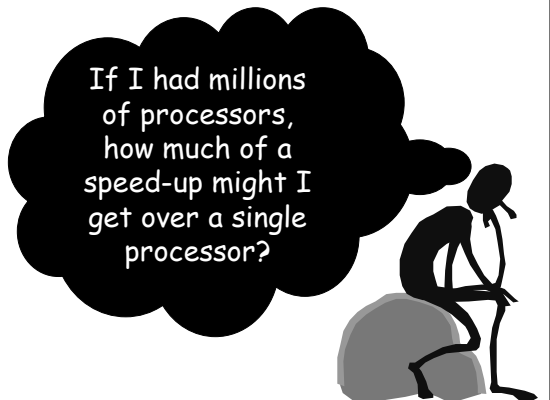
Intel Pentium division error

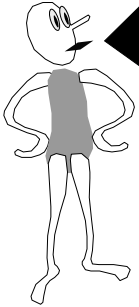
The Pentium uses essentially the same algorithm, but computes more than one bit of the result in each step.

Several leading bits of the divisor and quotient are examined at each step, and the difference is looked up in a table.

The table had several bad entries.


Ultimately Intel offered to replace any defective chip, estimating their loss at \$475 million.





Brent's Law

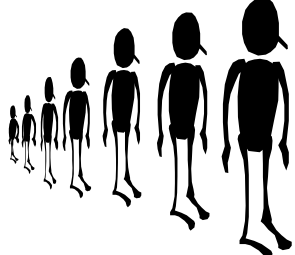
At best, p processors will give you a factor of p speedup over the time it takes on a single processor.



The traditional GCD algorithm will take linear time to operate on two n bit numbers.

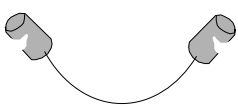
Can it be done faster in parallel?

If n^2 people agree to help you compute the GCD of two n bit numbers, it is not obvious that they can finish faster than if you had done it yourself.




Is GCD inherently sequential?

Decision Trees and Information:
A Question of Bits




20 Questions

S = set of all English nouns

Game:
I am thinking of an element of S .
You may ask up to 20 YES/NO questions.


What is a question strategy for this game?

Choice Tree



A choice tree is a rooted, directed tree with an object called a "choice" associated with each edge and a label on each leaf.


Choice Tree Representation of S



We satisfy these two conditions:

- Each leaf label is in S
- Each element from S on exactly one leaf.


Question Tree Representation of S



I am thinking of an outfit.
Ask me questions until you know which one.

What color is the beanie?
What color is the tie?

When a question tree has at most 2 choices at each node, we will call it a decision tree, or a decision strategy.



Note: Nodes with one choice represent stupid questions, but we do allow stupid questions.

20 Questions

Suppose $S = \{a_0, a_1, a_2, \dots, a_k\}$


Binary search on S.

First question will be:

"Is the word in $\{a_0, a_1, a_2, \dots, a_{(k-1)/2}\}$?"

20 Questions Decision Tree Representation

A decision tree with depth at most 20, which has the elements of S on the leaves.



Decision tree for $\{a_0, a_1, a_2, \dots, a_{(k-1)/2}\}$ Decision tree for $\{a_{(k-1)/2}, \dots, a_{k-1}, a_k\}$

Decision Tree Representation

Theorem:

The binary-search decision tree for S with $k+1$ elements $\{a_0, a_1, a_2, \dots, a_k\}$ has depth

$$\begin{aligned} & \lceil \log(k+1) \rceil \\ &= \lfloor \log k \rfloor + 1 \\ &= |k| \end{aligned}$$

"the length of k
when written
in binary"

A lower bound

Theorem: No decision tree for S (with $k+1$ elements) can have depth $d < \lfloor \log k \rfloor + 1$.

A lower bound

Theorem: No decision tree for S (with $k+1$ elements) can have depth $d < \lfloor \log k \rfloor + 1$.

Proof:

A depth d binary tree can have at most 2^d leaves.

But $d < \lfloor \log k \rfloor + 1 \Rightarrow$ number of leaves $2^d < (k+1)$

Hence some element of S is not a leaf.

Tight bounds!

The optimal-depth decision tree for any set S with $(k+1)$ elements has depth

$$\lfloor \log k \rfloor + 1 = |k|$$

Recall...

The minimum number of bits used to represent unordered 5 card poker hands

Recall...

The minimum number of bits used to represent unordered 5 card poker hands =

$$\lceil \log_2 \binom{52}{5} \rceil$$

= 22 bits

= The decision tree depth for 5 card poker hands.

Prefix-free Set

Let T be a subset of $\{0,1\}^*$.

Definition:

T is prefix-free if for any distinct $x, y \in T$,
if $|x| < |y|$, then x is not a prefix of y

Example:

$\{00, 01, 1, 11\}$ is prefix-free
 $\{0, 01, 10, 11, 101\}$ is not.

Prefix-free Code for S

Let S be any set.

Definition: A prefix-free code for S is
a prefix-free set T and
a 1-1 "encoding" function $f: S \rightarrow T$.

The inverse function f^{-1} is called the
"decoding function".

Example: $S = \{\text{buy, sell, hold}\}$.

$T = \{0, 110, 1111\}$.

$f(\text{buy}) = 0, f(\text{sell}) = 1111, f(\text{hold}) = 110$.

What is so cool
about prefix-free
codes?



Sending sequences of
elements of S over a
communications
channel

Let T be prefix-free and f be an
encoding function. Wish to send $\langle x_1,$
 $x_2, x_3, \dots \rangle$

Sender: sends $f(x_1) f(x_2) f(x_3) \dots$

Receiver: breaks bit stream into
elements
of T and decodes using f^{-1}

Sending info on a channel

Example: $S = \{\text{buy, sell, hold}\}$.

$T = \{0, 110, 1111\}$.

$f(\text{buy}) = 0, f(\text{sell}) = 1111, f(\text{hold}) = 110$.

If we see

00011011111100...

we know it must be

0 0 0 110 1111 110 0 ...

and hence

buy buy buy hold sell hold buy ...

Morse Code is not Prefix-free!

SOS encodes as ...---...

A .-.	F ..-	K -.-	P .--.	U ...-	Z ---.
B -....	G --.	L .-..	Q ---.-	V ...-	
C -.-.	H	M --	R .-.	W .--	
D -..	I ..	N -.	S ...	X -.-.	
E .	J .---	O ---	T -	Y -.-.	

Morse Code is not Prefix-free!

SOS encodes as ...---...

Could decode as: ..|.-|--|..|. = IAMIE

A.-	F...-	K-.-	P.-.-	U...-	Z---..
B-...-	G---.	L...-	Q---.-	V...-	
C-.-.	H....	M--	R.-.	W.--	
D-..	I..	N-.	S...	X-.-	
E.	J.---	O---	T-	Y-.-	

Unless you use pauses

SOS encodes as ... --- ...


A.-	F...-	K-.-	P.-.-	U...-	Z---..
B-...-	G---.	L...-	Q---.-	V...-	
C-.-.	H....	M--	R.-.	W.--	
D-..	I..	N-.	S...	X-.-	
E.	J.---	O---	T-	Y-.-	

Prefix-free codes
are also called
"self-delimiting"
codes.




Representing prefix-free codes


A = 100
B = 010
C = 101
D = 011
E = 00
F = 11




"CAFÉ" would encode as 1011001100
How do we decode 1011001100 (fast)?




If you see: 1000101000111011001100
can decode as:




If you see: 1000101000111011001100
can decode as: A




If you see: 100101000111011001100
 can decode as: AB




If you see: 10001000111011001100
 can decode as: ABA




If you see: 100010100111011001100
 can decode as: ABAD



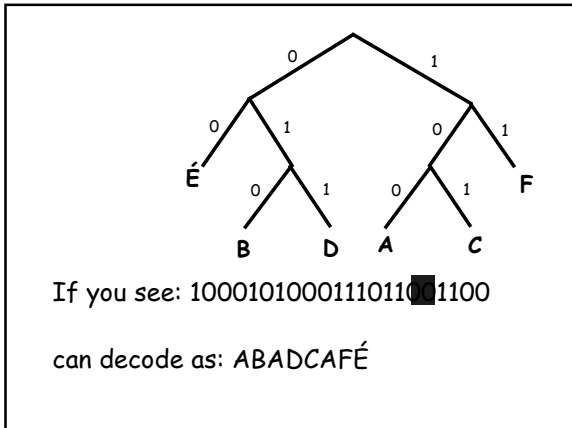
If you see: 100010100111011001100
 can decode as: ABADC



If you see: 10001010001111001100
 can decode as: ABADCA



If you see: 1000101000111011001100
 can decode as: ABADCAF



Prefix-free codes are yet another representation of a decision tree.

Theorem:

S has a decision tree of depth d

if and only if

S has a prefix-free code with all codewords bounded by length d

Extends to infinite sets

Let S is a subset of Σ^*

Theorem:

S has a decision tree where all length n elements of S have depth $\leq D(n)$

if and only if

S has a prefix-free code where all length n strings in S have encodings of length $\leq D(n)$

I am thinking of some natural number k . ask me YES/NO questions in order to determine k .

Let $d(k)$ be the number of questions that you ask when I am thinking of k .

Let $D(n) = \max \{ d(k) \text{ over } n\text{-bit numbers } k \}$.

I am thinking of some natural number k - ask me YES/NO questions in order to determine k .

Naïve strategy: Is it 0? 1? 2? 3? ...

$d(k) = k+1$

$D(n) = 2^{n+1}$ since $2^{n+1} - 1$ uses only n bits.


Effort is exponential in length of k

I am thinking of some natural number k - ask me YES/NO questions in order to determine k .

What is an efficient question strategy?

I am thinking of some natural number k ...

Does k have length 1? NO
 Does k have length 2? NO
 Does k have length 3? NO
 ...
 Does k have length n ? YES
 Do binary search on strings of length n .



$$d(k) = |k| + |k| = 2(\lfloor \log k \rfloor + 1)$$


$D(n) = 2n$ Size First/ Binary Search

Does k have length 1? NO
 Does k have length 2? NO
 Does k have length 3? NO
 ...
 Does k have length n ? YES
 Do binary search on strings of length n .

What prefix-free code corresponds to the Size First / Binary Search decision strategy?

$f(k) = (|k| - 1)$ zeros, followed by 1, and then by the binary representation of k

$|f(k)| = 2 |k|$



Another way to look at f

$k = 27 = 11011$, and hence $|k| = 5$


$f(k) = 00001 11011$

"Fat Binary" \Leftrightarrow Size First/Binary Search strategy

Is it possible to beat $2n$ questions to find a number of length n ?

Look at the prefix-free code...
 Any obvious improvement suggest itself here?

length of k in unary $\Rightarrow |k|$ bits
 k in binary $\Rightarrow |k|$ bits




In fat-binary, $D(n) \leq 2n$
 Now $D(n) \leq n + 2(\lfloor \log n \rfloor + 1)$

Can you do better?

better-than-Fat-Binary-code(k)
 concatenates

length of k in fat binary $\Rightarrow 2\lceil \log k \rceil$
 bits

k in binary $\Rightarrow \lceil \log k \rceil$
 bits




Hey, wait!

In a better prefix-free code

RecursiveCode(k) concatenates
 RecursiveCode($\lceil \log k \rceil$) & k in binary

better-than-Fat-Binary code

	better-t-FB	$\lceil \log k \rceil + 2\lceil \log k \rceil$
$\lceil \log k \rceil$ in fat binary		$\Rightarrow 2\lceil \log k \rceil$
bits		
k in binary		$\Rightarrow \lceil \log k \rceil$ bits




Oh, I need to remember how many
 levels of recursion $r(k)$

In the final code
 $F(k) = F(r(k)) \cdot \text{RecursiveCode}(k)$

$r(k) = \log^* k$

Hence, length of $F(k)$
 $= \lceil \log k \rceil + \lceil \log \lceil \log k \rceil \rceil + \lceil \log \lceil \log \lceil \log k \rceil \rceil + \dots + 1$




Good, Bonzo! I had thought you
 had fallen asleep.




Maybe I can do better...

Can I get a prefix code
 for k with length $\approx \log k$?




No!

Let me tell you why
 length $\approx \log k$
 is not possible

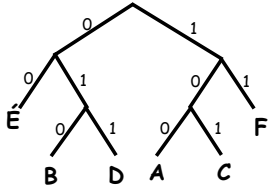


Decision trees have a natural probabilistic interpretation.

Let T be a decision tree for S . Start at the root, flip a fair coin at each decision, and stop when you get to a leaf. Each sequence w in S will be hit with probability $1/2^{|w|}$



Random walk down the tree




$\Pr(F) = \frac{1}{4}, \Pr(A) = 1/8, \Pr(C) = 1/8, \dots$

Let T be a decision tree for S (possibly countably infinite set)

The probability that some element in S is hit by a random walk down from the root is

$$\sum_{w \in S} 1/2^{|w|} \leq 1$$

Kraft Inequality




Let S be any prefix-free code.

Kraft Inequality:
 $\sum_{w \in S} 1/2^{|w|} \leq 1$

Fat Binary:
 $f(k)$ has $2|k| \approx 2 \log k$ bits

$$\sum_{k \in \mathbb{N}} \frac{1}{2^{|f(k)|}} \leq 1$$


$$\approx \sum_{k \in \mathbb{N}} 1/k^2$$


Let S be any prefix-free code.

Kraft Inequality:
 $\sum_{w \in S} 1/2^{|w|} \leq 1$

Better-than-FatB Code:
 $f(k)$ has $|k| + 2||k||$ bits

$$\sum_{k \in \mathbb{N}} \frac{1}{2^{|f(k)|}} \leq 1$$


$$\approx \sum_{k \in \mathbb{N}} 1/(k (\log k)^2)$$


Let S be any prefix-free code.

Kraft Inequality:
 $\sum_{w \in S} 1/2^{|w|} \leq 1$

Ladder Code: k is represented by $|k| + ||k|| + |||k||| + \dots$ bits


$$\sum_{k \in \mathbb{N}} \frac{1}{2^{|f(k)|}} \leq 1$$

$$\approx \sum_{k \in \mathbb{N}} 1/(k \log k \log \log k \dots)$$


Let S be any prefix-free code.

Kraft Inequality:

$$\sum_{w \in S} 1/2^{|w|} \leq 1$$



Can a code that represents k by $|k| = \log k$ bits exist?

No, since $\sum_{k \in \mathbb{N}} 1/k$ diverges !!
 So you can't get $\log n$, Bonzo...

Back to compressing words

The optimal-depth decision tree for any set S with $(k+1)$ elements has

depth
 $\lfloor \log k \rfloor + 1$

The optimal prefix-free code for $A-Z$ + "space" has length
 $\lfloor \log 26 \rfloor + 1 = 5$

English Letter Frequencies

But in English, different letters occur with different *frequencies*.

A 8.1%	F 2.3%	K 7.9%	P 1.6%	U 2.8%	Z .04%
B 1.4%	G 2.1%	L 3.7%	Q .11%	V .86%	
C 2.3%	H 6.6%	M 2.6%	R 6.2%	W 2.4%	
D 4.7%	I 6.8%	N 7.1%	S 6.3%	X .11%	
E 12%	J .11%	O 7.7%	T 9.0%	Y 2.0%	


short encodings!

Why should we try to minimize the maximum length of a codeword?

If encoding $A-Z$, we will be happy if the "average codeword" is short.

Given frequencies for $A-Z$, what is the optimal prefix-free encoding of the alphabet?

I.e., one that minimizes the average code length



Huffman Codes: Optimal Prefix-free Codes Relative to a Given Distribution

Here is a Huffman code based on the English letter frequencies given earlier:

A 1011	F 101001	K 10101000	P 111000	U 00100
B 111001	G 101000	L 11101	Q 1010100100	V 1010101
C 01010	H 1100	M 00101	R 0011	W 01011
D 0100	I 1111	N 1000	S 1101	X 1010100101
E 000	J 1010100110	O 1001	T 011	Y 101011
				Z 1010100111

References

The Mathematical Theory of Communication,
by C. Shannon and W. Weaver

Elements of Information Theory, by T. Cover
and J. Thomas