https://algs4.cs.princeton.edu

1.4 ANALYSIS OF ALGORITHMS

- *introduction*
- *running time (experimental analysis)*
- *running time (mathematical models)*
- *memory usage*

https://algs4.cs.princeton.edu

1.4 Analysis of Algorithms

- *introduction*
- *running time (experimental analysis)*
- *running time (mathematical models)*
- *memory usage*

Different viewpoints

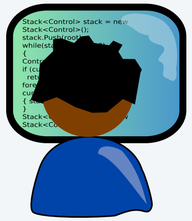

programmer needs to
build a working solution

client wants the problem
solved efficiently

theoretician seeks to understand
why things work

student (you) plays all
three roles in this course

Running time

> *" As soon as an Analytical Engine exists, it will necessarily guide the future course of the science. Whenever any result is sought by its aid, the question will then arise—By what course of calculation can these results be arrived at by the machine in the shortest time ? "* — Charles Babbage (1864)

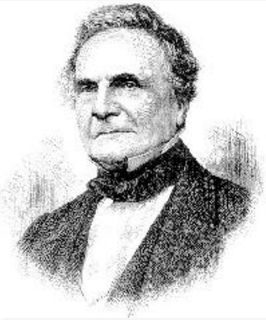

how many times do you have to turn the crank?

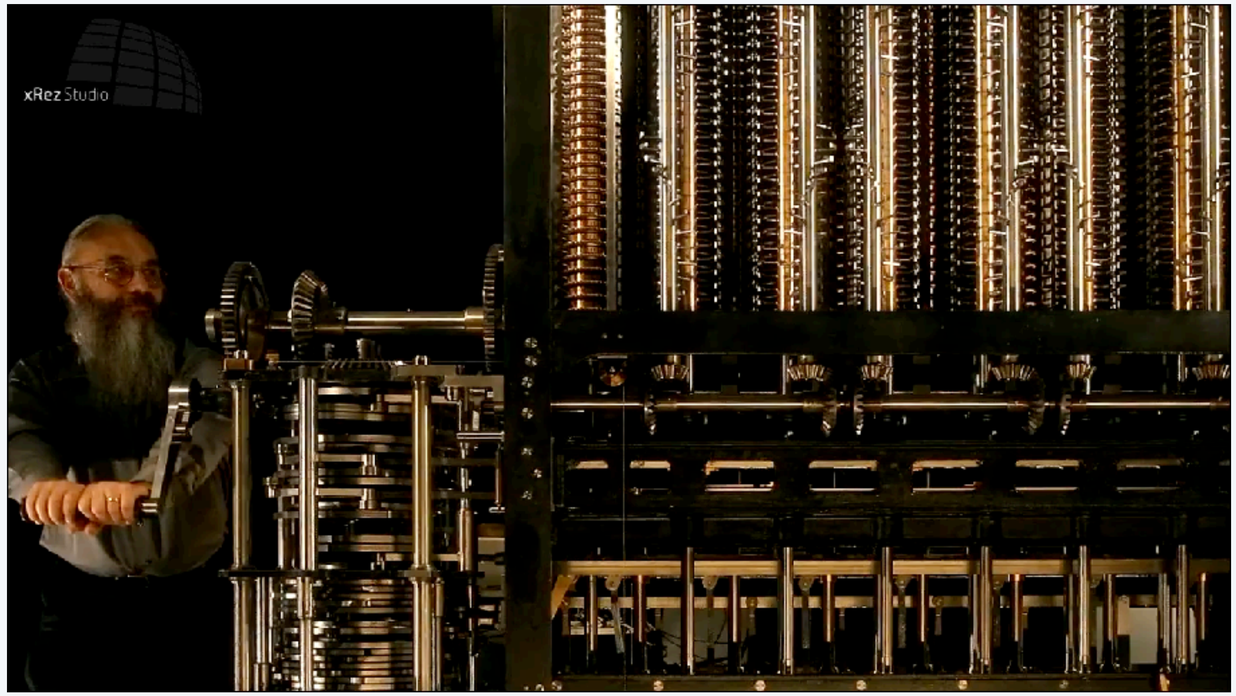

https://vimeo.com/49080293

Running time

> *“ As soon as an Analytical Engine exists, it will necessarily guide the future course of the science. Whenever any result is sought by its aid, the question will then arise—By what course of calculation can these results be arrived at by the machine in the shortest time ? ”* — Charles Babbage (1864)

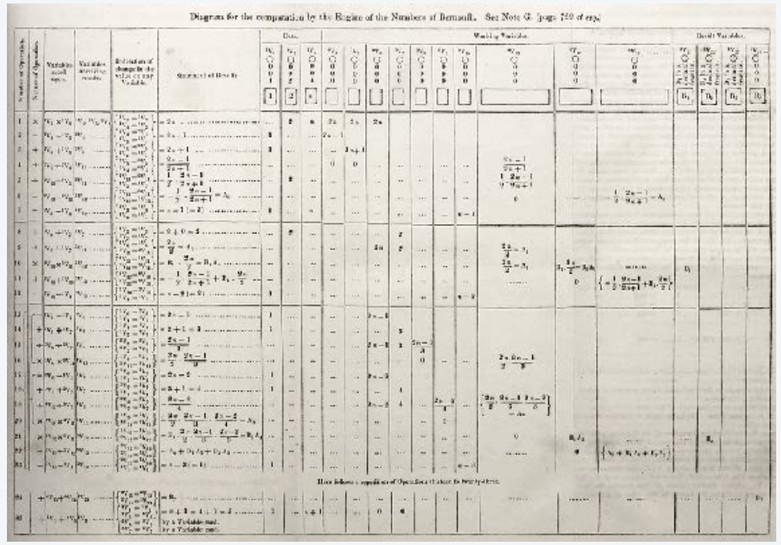

Diagram for the computation by the Engine of the Numbers of Bernoulli. See Note G. (page 722 of 1842)

Ada Lovelace’s algorithm to compute Bernoulli numbers on Analytical Engine (1843)

Rare book containing the world’s first computer algorithm earns $125,000 at auction

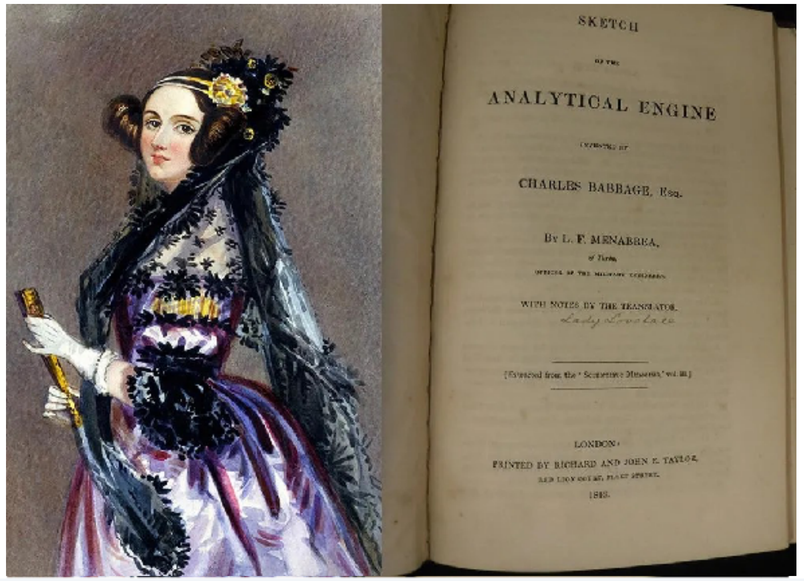

SKETCH

OF THE

ANALYTICAL ENGINE

CHARLES BABBAGE, Esq.

By L. F. MENABREA,

WITH NOTES BY THE TRANSLATOR.

Lady Lovelace

LONDON:

PRINTED BY RICHARD AND JOHN E. TAYLOR,

1843.

An algorithmic success story

Goal. Multiply two polynomials of degree n.

$$(x^3 + x^2 - 2x + 1) \cdot (3x^3 - x^2 + 2x + 1) = 3x^6 + 2x^5 - 5x^4 + 8x^3 - 4x^2 + 1$$

- Applications: JPEG compression, MRI, astrophysics, and more.
- Grade-school algorithm: $\Theta(n^2)$ operations.
- FFT algorithm: $\Theta(n \log n)$ operations, enabling modern technology.

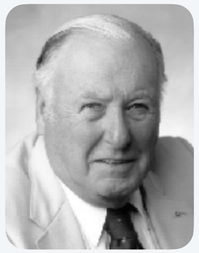

James Cooley

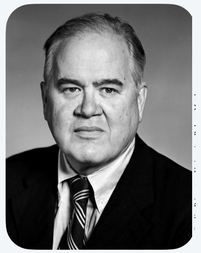

John Tukey

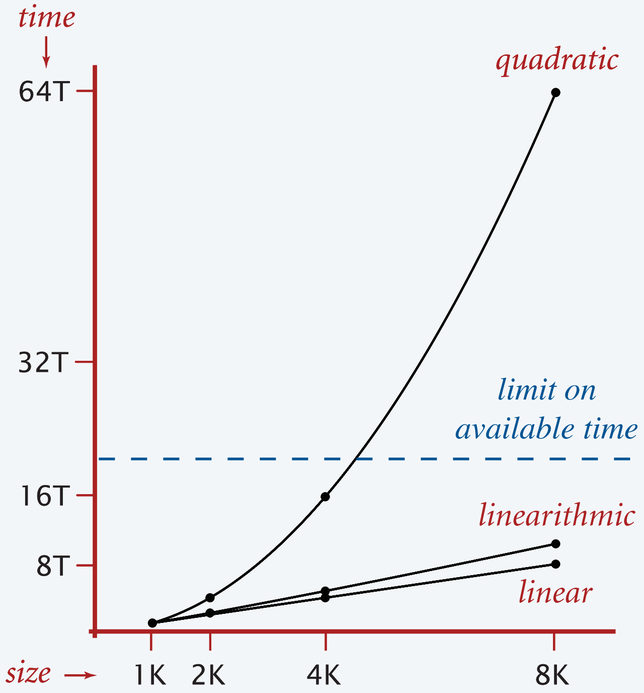

Another algorithmic success story?

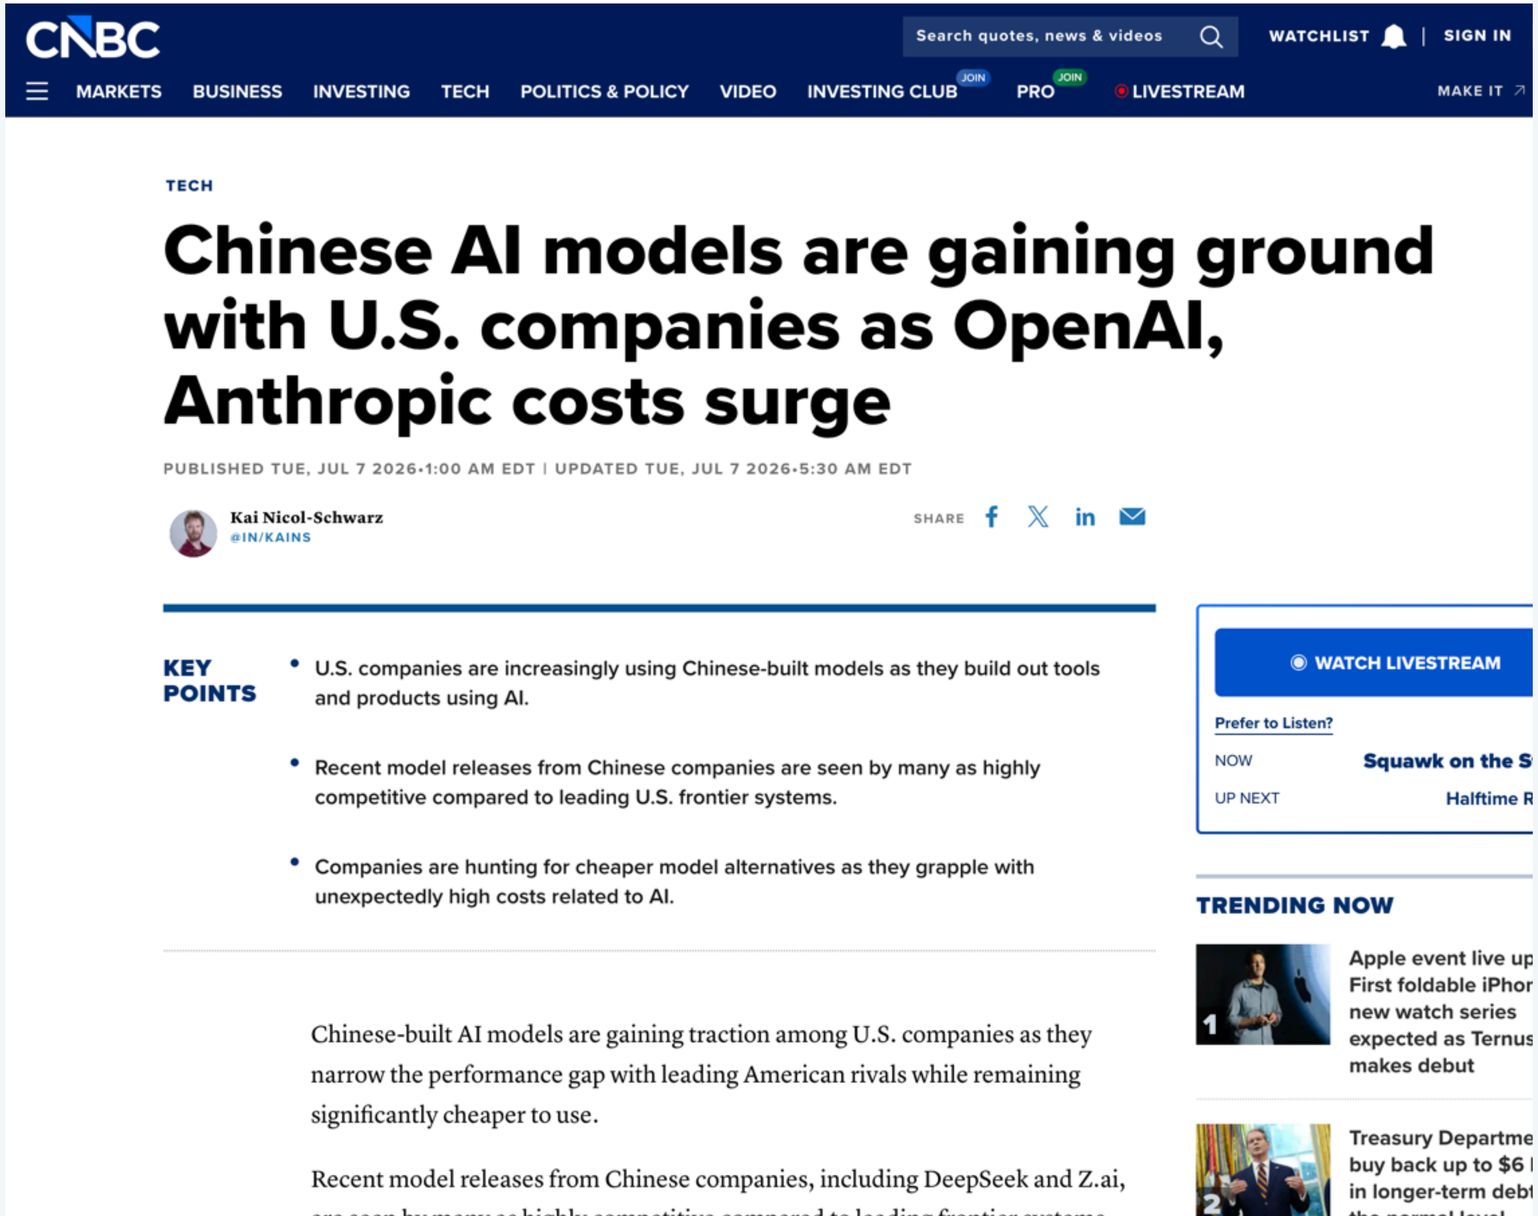

CNBC

Search quotes, news & videos | WATCHLIST | SIGN IN

MARKETS BUSINESS INVESTING TECH POLITICS & POLICY VIDEO INVESTING CLUB JOIN PRO JOIN LIVESTREAM MAKE IT

TECH

Chinese AI models are gaining ground with U.S. companies as OpenAI, Anthropic costs surge

PUBLISHED TUE, JUL 7 2026•1:00 AM EDT | UPDATED TUE, JUL 7 2026•5:30 AM EDT

Kai Nicol-Schwarz
@IN/KAINS

SHARE

KEY POINTS

- U.S. companies are increasingly using Chinese-built models as they build out tools and products using AI.
- Recent model releases from Chinese companies are seen by many as highly competitive compared to leading U.S. frontier systems.
- Companies are hunting for cheaper model alternatives as they grapple with unexpectedly high costs related to AI.

Chinese-built AI models are gaining traction among U.S. companies as they narrow the performance gap with leading American rivals while remaining significantly cheaper to use.

Recent model releases from Chinese companies, including DeepSeek and Z.ai,

WATCH LIVESTREAM

Prefer to Listen?

NOW Squawk on the S

UP NEXT Halftime R

TRENDING NOW

1 Apple event live up First foldable iPhor new watch series expected as Ternus makes debut

2 Treasury Departme buy back up to $6 in longer-term debt

The core challenge

Q1. Will my program handle on large, real-world inputs?

Q2. If not, how can I analyze and improve its performance?

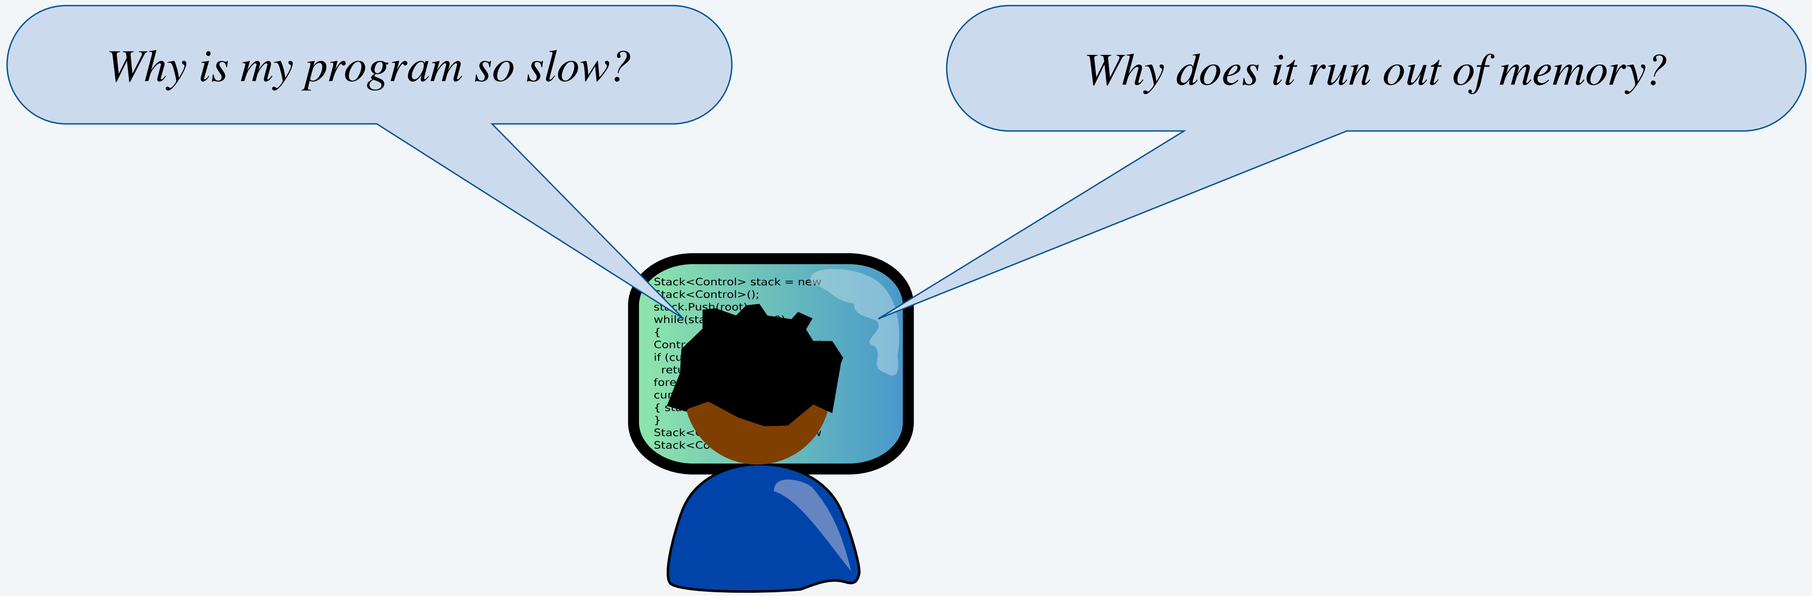

Our approach: a combination of experiments and mathematical modeling.

Example: 3-SUM problem

Goal. Given an array of n distinct integers, count triples $i < j < k$ such that $a[i] + a[j] + a[k] = 0$.

```
~/cos226/analysis> more 8ints.txt
8
30 -40 -20 -10 40 0 10 5

~/cos226/analysis> java ThreeSum 8ints.txt
4
```

	a[i]	a[j]	a[k]	sum	
1	30	–40	10	0	✔
2	30	–20	–10	0	✔
3	–40	40	0	0	✔
4	–10	0	10	0	✔

Context. Arises in in computational geometry (and even in computer games!)

Open problem. What is the optimal running time for solving 3-SUM ?

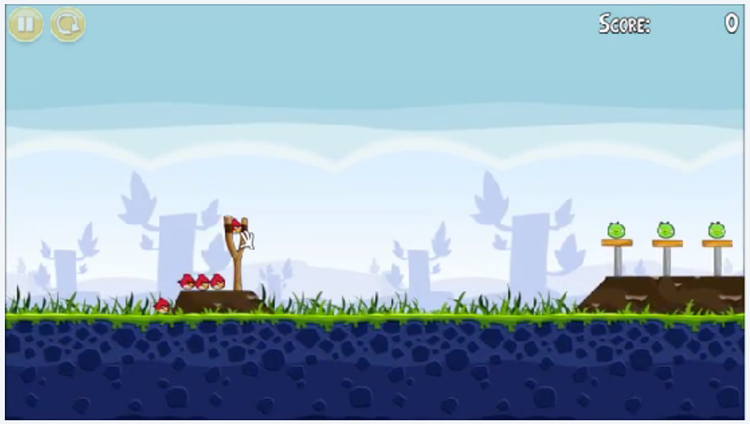

3-SUM problem: brute-force algorithm

```
public class ThreeSum {

   public static int count(int[] a) {
      int n = a.length;
      int count = 0;
      for (int i = 0; i < n; i++)
         for (int j = i+1; j < n; j++)
            for (int k = j+1; k < n; k++)
               if (a[i] + a[j] + a[k] == 0)
                  count++;
      return count;
   }

   public static void main(String[] args) {
      In in = new In(args[0]);
      int[] a = in.readAllInts();
      StdOut.println(count(a));
   }
}
```

count distinct triples that sum to 0

assume no integer overflow

https://algs4.cs.princeton.edu

1.4 Analysis of Algorithms

- *introduction*
- *running time (experimental analysis)*
- *running time (mathematical models)*
- *memory usage*

Measuring running time

Experiment. Measure the program's running time on inputs of different sizes.

Observation. The running time $T(n)$ increases with the input size n.

```
~/cos226/analysis> java ThreeSum 1Kints.txt
70
```

tick tick

```
~/cos226/analysis> java ThreeSum 2Kints.txt
528
```

tick tick tick tick tick tick tick tick
tick tick tick tick tick tick tick tick

```
~/cos226/analysis> java ThreeSum 3Kints.txt
1670
```

tick tick tick tick tick tick tick tick tick tick tick tick
tick tick tick tick tick tick tick tick tick tick tick tick
tick tick tick tick tick tick tick tick tick tick tick tick
tick tick tick tick tick tick tick tick tick tick tick tick
tick tick tick tick tick tick

Measuring running time

Experiment. Measure the program's running time on inputs of different sizes.

Observation. The running time $T(n)$ increases with the input size n.

n	time (seconds) †
1,000	0.21
1,500	0.71
2,000	1.63
2,500	3.11
3,000	5.43
4,000	12.8
5,000	25.0
7,500	84.4
10,000	199.3

† *Apple M2 Pro with 32 GB memory running OpenJDK 11 on MacOS Ventura*

Data analysis: running time vs. input size

Visualization. Plot the running time $T(n)$ versus the input size n.

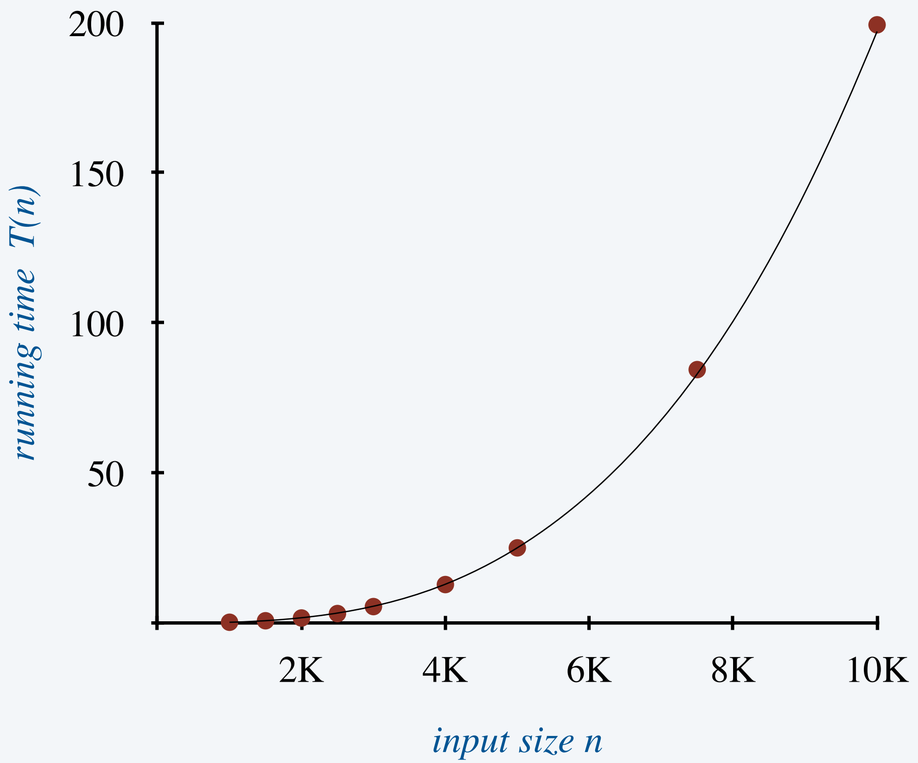

Hypothesis. The running time follows a power law: $T(n) = a \times n^b$ seconds.

Questions. How can we test this hypothesis? How can we estimate a and b ?

Doubling test: estimating the exponent b

Doubling test. Run the program, doubling the size of the input.

- Assume running time satisfies $T(n) = a \times n^b$.
- Estimate b as $\log_2(ratio)$.

n	time (seconds) †	ratio	$\log_2$(ratio)
500	0.05	–	–
1,000	0.21	4.20	2.07
2,000	1.63	7.76	2.96
4,000	12.8	7.85	2.97
8,000	103.1	8.05	3.01
16,000	819.0	7.94	2.99

$\log_2$ (103.1 / 12.8) = 3.01

seems to converge to a constant $b \approx 3.0$

$$\frac{T(n)}{T(n/2)} = \frac{an^b}{a(n/2)^b} = 2^b$$

$$\Longrightarrow \quad b = \log_2 \frac{T(n)}{T(n/2)}$$

why the $\log_2$(ratio) works

Doubling test: estimating the leading constant a

Doubling test. Run the program, doubling the size of the input.

- Assume running time satisfies $T(n) = a \times n^b$.
- Estimate b as $\log_2(ratio)$.
- Estimate a from $T(n) = a \times n^b$ using a sufficiently large value of n.

n	time (seconds) †	ratio	$\log_2$ ratio
500	0.05	–	–
1,000	0.21	4.20	2.07
2,000	1.63	7.76	2.96
4,000	12.8	7.85	2.97
8,000	103.1	8.05	3.01
16,000	819.0	7.94	2.99

$819.0 = a \times 16{,}000^3 \implies a = 2.00 \times 10^{-10}$

Hypothesis. Running time is well modeled by $2.00 \times 10^{-10} \times n^3$ seconds.

Analysis of algorithms: poll 1

Estimate the running time when $n = 96{,}000$ (assuming power law hypothesis).

A. 39 seconds

B. 52 seconds

C. 117 seconds

D. 350 seconds

n	time (seconds)
1,000	0.02
2,000	0.05
4,000	0.20
8,000	0.81
16,000	3.25
32,000	13.01

Machine invariance

Hypothesis. For a fixed algorithm, the running times on different computers are the same up to a constant factor.

Note. That constant factor can be large, sometimes several orders of magnitude.

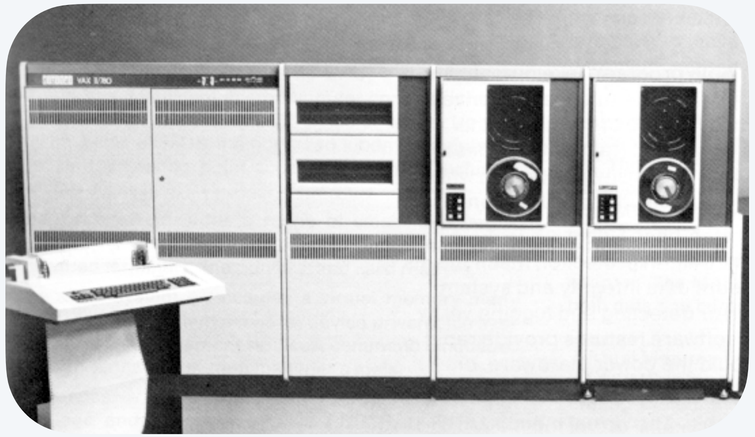

1970s
(VAX-11/780)

10,000× *faster*

2020s
(Macbook Pro M2)

What affects the running time?

System independent effects.

- Algorithm.
- Input data.

determines exponent b in power law $T(n) = a \times n^b$

System dependent effects.

- Hardware: CPU, memory, cache, …
- Software: compiler, interpreter, garbage collector, …
- System: operating system, network, other apps, …

determines leading coefficient a in power law $T(n) = a \times n^b$

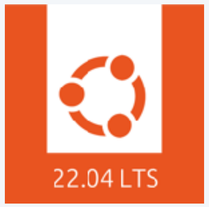

Bad news. Getting accurate timing measurements can be difficult.

system-dependent effects can introduce noise

Context: the scientific method

Experimental algorithmics follows the scientific method: observe, hypothesize, experiment, repeat.

Chemistry
(an expensive experiment)

Biology
(a slow experiment)

Physics
(a dangerous experiment)

Computer Science
(millions of cheap, fast, and safe experiments)

Good news. Computational experiments are cheap, fast, and safe.

https://algs4.cs.princeton.edu

1.4 Analysis of Algorithms

- *introduction*
- *running time (experimental analysis)*
- *running time (mathematical models)*
- *memory usage*

Mathematical model of running time

Model. The running time = $\sum$ (frequency of operation) $\times$ (cost of operation).

- Frequency of operation: depends on the algorithm and the specific input.
- Cost of operation: depends on hardware, software, system, and low-level implementation details.

The New York Times

PROFILES IN SCIENCE

The Yoda of Silicon Valley

Donald Knuth, master of algorithms, reflects on 50 years of his opus-in-progress, "The Art of Computer Programming."

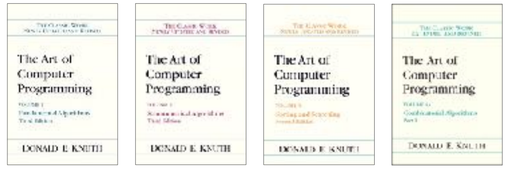

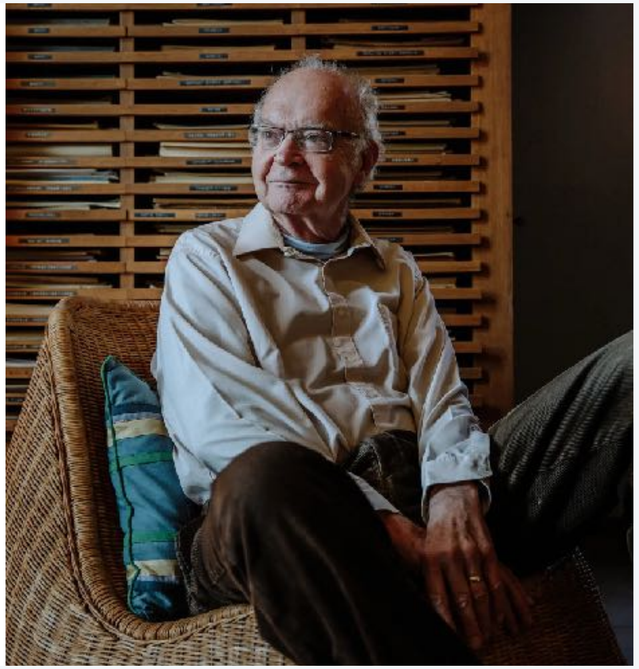

Warning. For arbitrary programs, frequencies may be impossible to determine. ⟵ *halting problem*

Example: one-sum problem

Q. How many operations does this code perform as a function of the input size n ?

```
int count = 0;
for (int i = 0; i < n; i++)
   if (a[i] == 0)
      count++;
```

operation	cost (ns) †	frequency
variable declaration	2 / 5	2
assignment statement	1 / 5	2
less than compare	1 / 5	$n + 1$
equality test	1/ 1 0	n
array read	1 / 10	n
increment	1 / 10	n to $2\,n$

in practice, depends on caching, bounds checking, … (*see* COS 217)

painful to count exactly

† *representative estimates* (*with a bit of poetic license*)

Simplification 1: cost model

Cost model. Pick one elementary operation as a proxy for running time. ← *array accesses, compares, API calls, floating-point operations, …*

```
int count = 0;
for (int i = 0; i < n; i++)
   if (a[i] == 0)
      count++;
```

← *"inner loop"*

operation	cost (ns) †	frequency
variable declaration	2 / 5	2
assignment statement	1 / 5	2
less than compare	1 / 5	$n + 1$
equality test	1 / 10	n
array read	1 / 10	n
increment	1 / 10	n to $2\,n$

← *cost model = array accesses*

Simplification 2: asymptotic notation

Tilde notation. Ignore lower-order terms.

Big Theta notation. Ignore both lower-order terms and the leading constant.

← *rigorous definitions involve limits*

function	tilde notation	big Theta ← *"order of growth"*
$4n^5 + 20n^3 + 16$	$\sim 4n^5$	$\Theta(n^5)$
$0.01n^2 + 100n^{4/3} - 56$	$\sim 0.01n^2$	$\Theta(n^2)$
$2^n + n^5$	$\sim 2^n$	$\Theta(2^n)$
$\frac{1}{6}n^3 - \frac{1}{2}n^2 + \frac{1}{3}n$	$\sim \frac{1}{6}n^3$	$\Theta(n^3)$

discard lower-order terms

(*e.g.*, n = 1,000: 166.667 *million* vs. 166.167 *million*)

Rationale.

- For large n, lower-order terms have negligible effect.
- For small n, the value is so small that we don't care.

Analysis of algorithms: poll 2

Which of the following correctly describes the function $f(n) = n \log_2 n + 3n^2 + 10n$?

A. $\sim 10\, n$

B. $\sim n \log_2 n$

C. $\sim n^2$

D. $\Theta(n \log n)$

E. $\Theta(n^2)$

Example: 2-SUM analysis

Q. Approximately how many operations as a function of input size n ?

```
int count = 0;
for (int i = 0; i < n; i++)
   for (int j = i+1; j < n; j++)
      if (a[i] + a[j] == 0)
         count++;
```

"inner loop"

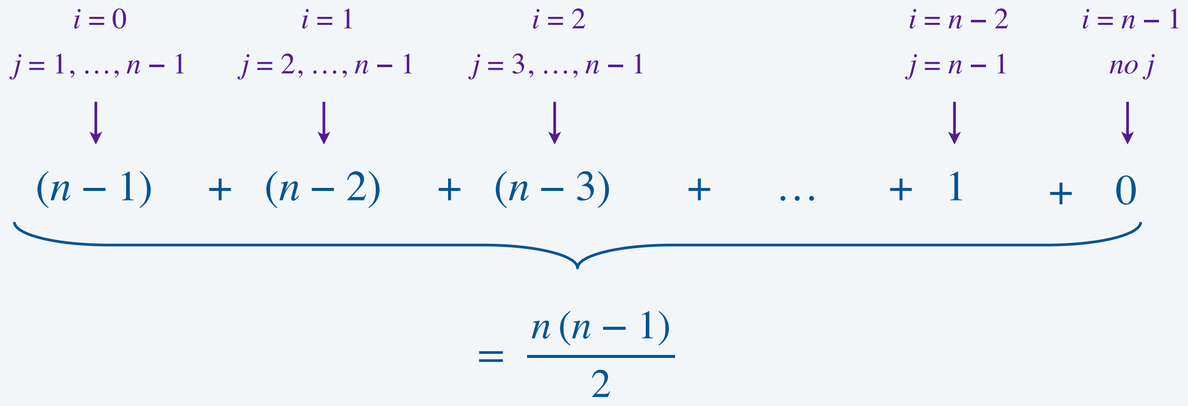

Step 1. Pick a cost model: array accesses.

Step 2. Count array accesses: $2 \times \frac{n(n-1)}{2} \sim 1 n^2$.

body inner loop makes 2 array accesses

Nested loops.

- Independent loops: analyze separately and multiply.
- Dependent loops: write a sum (and simplify).

Triangular sum

Claim. $0 + 1 + \ldots + (n-2) + (n-1) = \frac{1}{2}\, n\,(n-1)\,.$

Proof. [by picture]

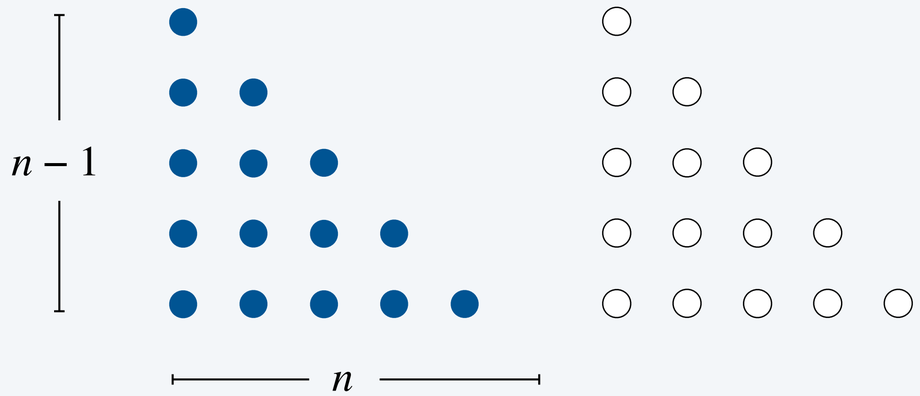

$$2 \times (\,1 + 2 + 3 + \ldots + n - 1\,) = n\,(n-1)$$

$$\Longrightarrow\ (\,1 + 2 + 3 + \ldots + n - 1\,) = \frac{1}{2}\, n\,(n-1)$$

Example: 3-SUM analysis

Q. Approximately many operations as a function of input size n ?

```
int count = 0;
for (int i = 0; i < n; i++)
   for (int j = i+1; j < n; j++)
      for (int k = j+1; k < n; k++)
         if (a[i] + a[j] + a[k] == 0)
            count++;
```

"inner loop" $\binom{n}{3} = \frac{n(n-1)(n-2)}{3!} \sim \frac{1}{6}n^3$

see COS 240

Step 1. Pick a cost model: array accesses.

Step 2. Count the number of array accesses: $\Theta(n^3)$.

Bottom line. Using a cost model and asymptotic notation makes analysis manageable.

Common orders of growth

order of growth	emoji	name	typical code pattern	description	example	$T(2n) / T(n)$
$\Theta(1)$	😍	**constant**	`a = b + c;`	statement	*add two numbers*	1
$\Theta(\log n)$	😎	**logarithmic**	`for (int i = n; i > 0; i /= 2) { ... }`	repeatedly divide in half	*binary search*	~ 1
$\Theta(n)$	😁	**linear**	`for (int i = 0; i < n; i++) { ... }`	single loop	*find the maximum*	2
$\Theta(n \log n)$	😀	**linearithmic**	*mergesort*	divide-and-conquer	*mergesort*	~ 2
$\Theta(n^2)$	😕	**quadratic**	`for (int i = 0; i < n; i++) for (int j = 0; j < n; j++) { ... }`	double loop	*check all pairs*	4
$\Theta(n^3)$	☹️	**cubic**	`for (int i = 0; i < n; i++) for (int j = 0; j < n; j++) for (int k = 0; k < n; k++) { ... }`	triple loop	*check all triples*	8
$\Theta(2^n)$	😈	**exponential**	*towers of Hanoi*	brute-force search	*check all subsets*	2^n

Useful discrete sums and approximations

Triangular sum. $1 + 2 + 3 + \ldots + n \sim \frac{1}{2} n^2$

Logarithmic sum. $\log_2 1 + \log_2 2 + \log_2 3 + \ldots + \log_2 n \sim \int_{x=1}^{n} \log_2 x \, dx = n \log_2 n$

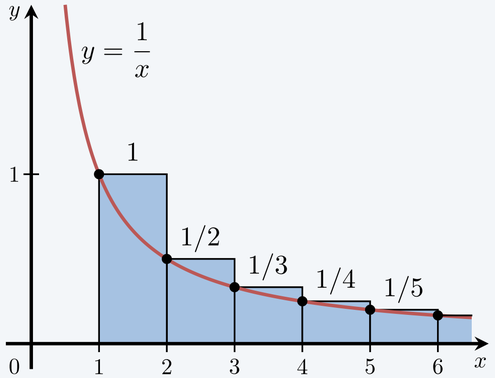

Harmonic sum. $\frac{1}{1} + \frac{1}{2} + \frac{1}{3} + \ldots + \frac{1}{n} \sim \int_{x=1}^{n} \frac{1}{x} dx = \ln n$

Geometric sum. $1 + 2 + 4 + 8 + \ldots + n = 2n - 1$

n is a power of 2

Geometric sum'. $n + \frac{n}{2} + \frac{n}{4} + \ldots + 1 = 2n - 1$

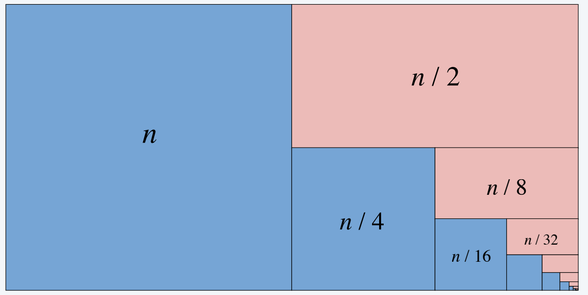

Geometric sum meme

https://marekbennett.com/2014/03/06/recursive-load

Analysis of algorithms: poll 3

Approximately how many array accesses as a function of n ?

```
int count = 0;
for (int i = 0; i < n; i++)
   for (int j = i+1; j < n; j++)
      for (int k = 1; k <= n; k = k*2)
         if (a[i] + a[j] >= a[k])
            count++;
```

A. $\sim n^2 \log_2 n$

B. $\sim \frac{3}{2} n^2 \log_2 n$

C. $\sim \frac{1}{2} n^3$

D. $\sim \frac{3}{2} n^3$

Analysis of algorithms: poll 4

What is the order of growth of the running time as a function of n ?

```
int count = 0;
for (int i = n; i >= 1; i = i/2)
   for (int j = 1; j <= i; j++)
      count++;
```

A. $\Theta(n)$

B. $\Theta(n \log n)$

C. $\Theta(n^2)$

D. $\Theta(2^n)$

https://algs4.cs.princeton.edu

1.4 Analysis of Algorithms

- *introduction*
- *running time (experimental analysis)*
- *running time (mathematical models)*
- *memory usage*

Memory basics: bits, bytes, and pointers

Bit. A single binary digits (0 or 1).

0				
1				

term	symbol	size
byte	B	8 bits
kilobyte	KB	10^3 bytes
megabyte	MB	10^6 bytes
gigabyte	GB	10^9 bytes
terabyte	TB	10^{12} bytes

some systems use powers of 2
(e.g., 1 MB = 2^{20} bytes)

Assumption. Running on a 64-bit machine with 8-byte pointers.

some JVMs "compress" pointers
to 4 bytes to avoid this cost

Typical memory usage of primitive types and arrays in Java

type	bytes
`boolean`	1
`byte`	1
`char`	2
`int`	4
`float`	4
`long`	8
`double`	8

primitive types

type	bytes
`boolean[]`	$1n + 24$
`int[]`	$4n + 24$
`double[]`	$8n + 24$

wasteful
(but ~ 36n bytes in Python 3)

array overhead = 24 bytes

one-dimensional arrays (length n)

type	bytes
`boolean[][]`	$\sim 1\ n^2$
`int[][]`	$\sim 4\ n^2$
`double[][]`	$\sim 8\ n^2$

two-dimensional arrays
(n-by-n array of arrays)

type	bytes
object reference	8

64-bit machine

Typical memory usage of objects in Java

Object header. Every Java object begins with a fixed 16-byte header. ← *class pointer, garbage collector bits, lock state, hash code, …*

Padding. Object size is rounded up to a multiple of 8 bytes.

Example. A `Date` object uses 32 bytes of memory.

```
public class Date {
   private int day;
   private int month;
   private int year;

   ...
}
```

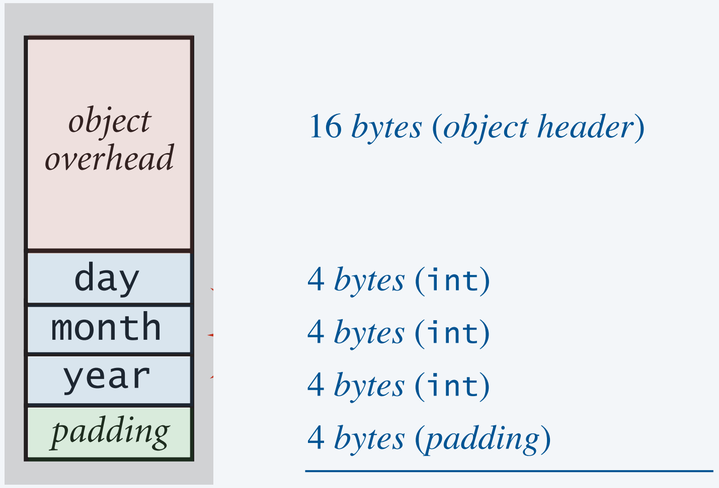

Analysis of algorithms: poll 5

How much memory does a `WeightedQuickUnionUF` object use as a function of n ?

A. $\sim 4n$ bytes

B. $\sim 8n$ bytes

C. $\sim 4n^2$ bytes

D. $\sim 8n^2$ bytes

```
public class WeightedQuickUnionUF {
   private int[] parent;
   private int[] size;
   private int count;

   public WeightedQuickUnionUF(int n) {
      parent = new int[n];
      size   = new int[n];
      count = 0;
      for (int i = 0; i < n; i++)
         parent[i] = i;
      for (int i = 0; i < n; i++)
         size[i] = 1;
   }
   ...
}
```

Analysis of algorithms: poll 5

How much memory does a `WeightedQuickUnionUF` object use as a function of n ?

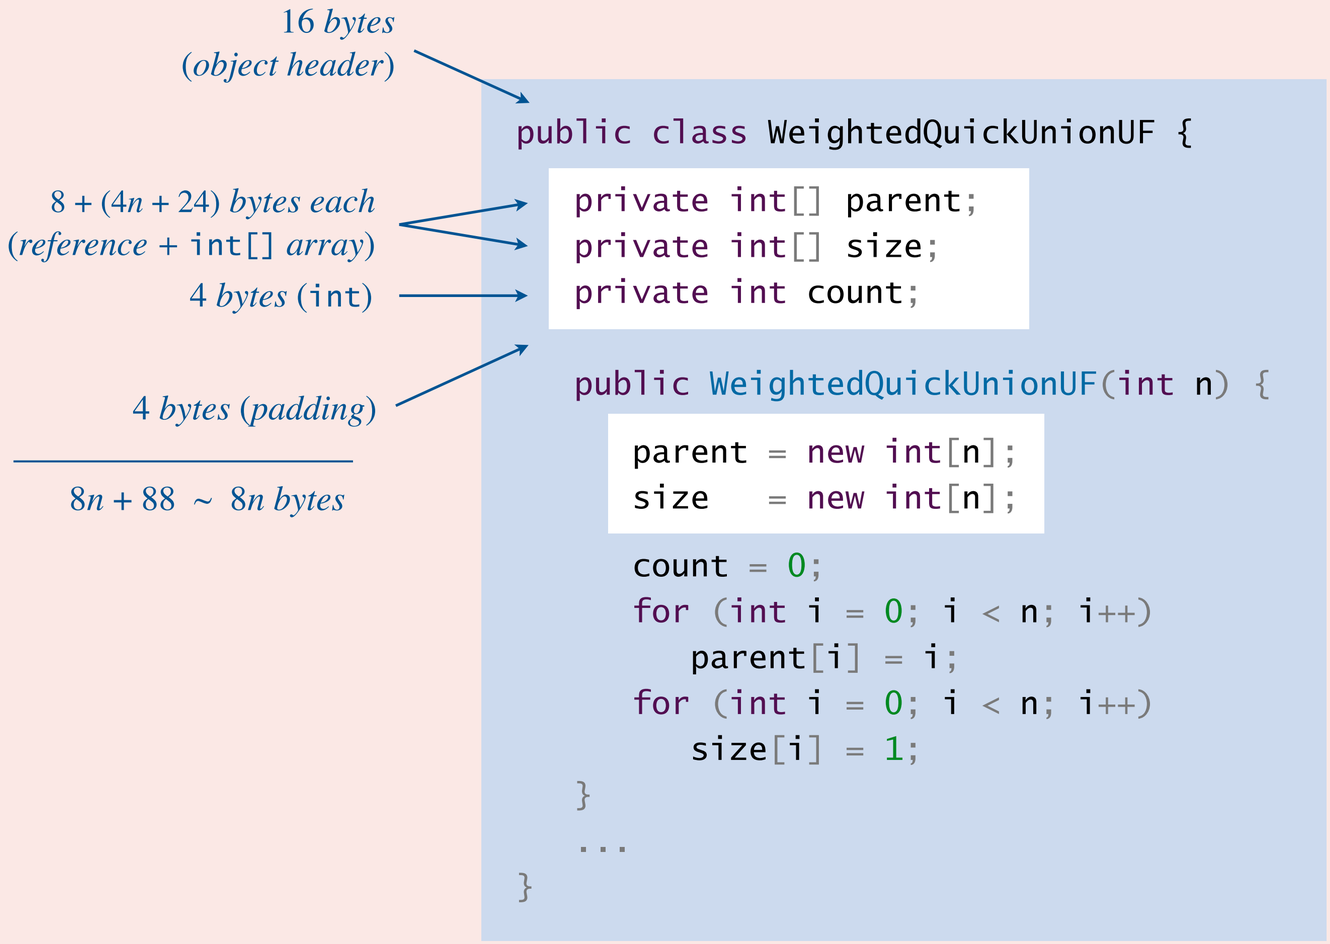

Turning the crank: summary

Goal. Analyze running time (or memory) as a function of input size.

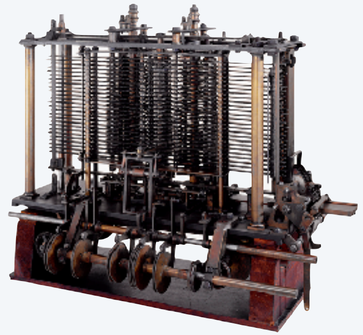

Empirical analysis.

- Run the program and measure how long it takes.
- Formulate a hypothesis for the running time.
- Use the model to predict behavior.

Mathematical analysis.

- Analyze the algorithm on an abstract machine.
- Count the frequency of its dominant operations.
- Apply a cost model and asymptotic notation to simplify analysis.
- Use the model to explain behavior.

$$\sum_{h=0}^{\lfloor \lg n \rfloor} \lceil n/2^{h+1} \rceil \; h \; \sim \; n$$

This course. You will learn to use both.

Credits

image	source	license
Charles Babbage	The Illustrated London News	public domain
Babbage Enginine in Operation	xRez Studio	
Algorithm for the Analytical Engine	Ada Lovelace	public domain
Ada Lovelace and Book	Moore Allen & Innocent	
Galaxies Colliding	SaltyMikan	
James Cooley	IEEE	
John Tukey	Princeton University	
Programmer Icon	Jaime Botero	public domain
Head in the Clouds	Ellis Nadler	education
Student Raising Hand	classroomclipart.com	educational use
Running Time	pano.si	
Analog Stopwatch	Adobe Stock	education license

Credits

image	source	license
Analog Stopwatch	Adobe Stock	education license
Apple M4 Chip	Apple	
Macbook Pro M2	Apple	
Scientific Method	Sue Cahalane	by author
Laboratory Apparatus	pixabay.com	public domain
Dissected Rat	Allen Lew	CC BY 2.0
Harmonic Integral	Wikimedia	public domain
Geometric Series	Wikimedia	CC BY-SA 3.0
Recursive Load	Marek Bennett	
The Yoda of Silicon Valley	New York Times	
Babbage's Analytical Engine	Science Museum, London	CC BY-SA 2.0
Alan Turing	Science Museum, London	

A final thought

> *“ It is convenient to have a measure of the amount of work involved in a computing process, even though it be a very crude one. We may count up the number of times that various elementary operations are applied in the whole process and then give them various weights.”* — Alan Turing (1947)