COMPUTER SCIENCE AND ENGINEERING

A Historical and Cultural Introduction

Abhijit Das

Department of Computer Science and Engineering
Indian Institute of Technology Kharagpur

26 July, 2017

Computing Devices

Abacus

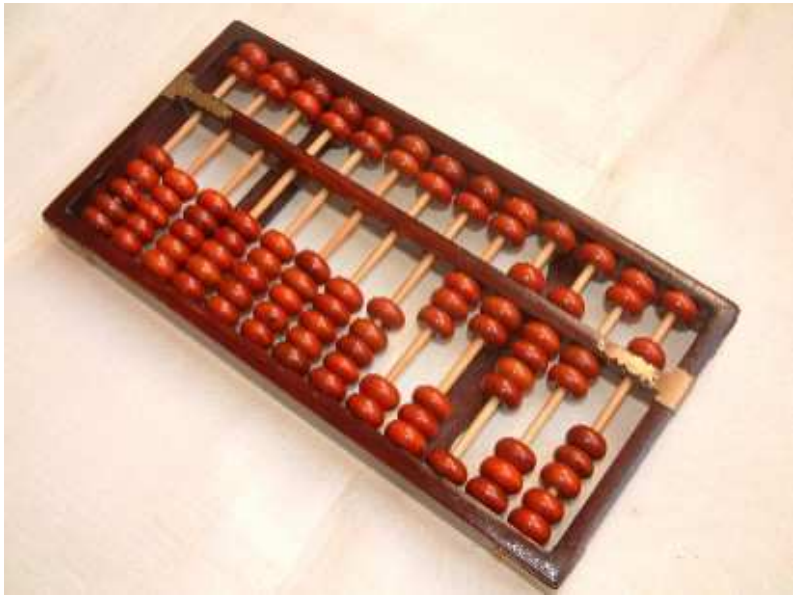

Source: Wikimedia

- Earliest known computing device
- May be invented in Babylon ca. 2700–2300 BC
- Used for counting
- Used also for simple operations like addition and subtraction

Antikythera Mechanism

Source: Wikimedia

- Oldest known mechanical analog computer
- Built ca. 100 BC
- Used to calculate the positions of celestial bodies
- Simulates the retrograde motions of planets

Astrolabe, Torquetum, . . .

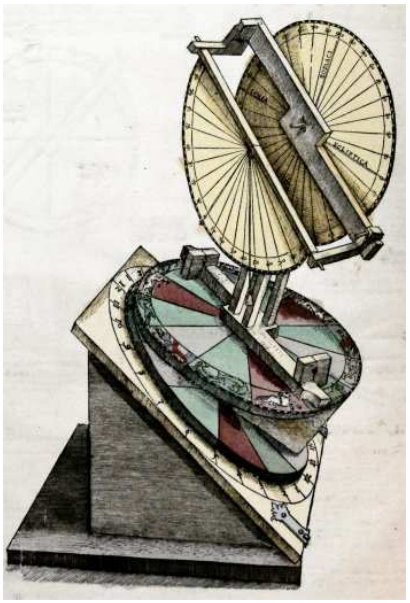

Source: Wikimedia

- Invented by the medieval Islamic world
- An astrolabe measures the inclinations of celestial bodies
- A torquetum measures and converts horizontal, equatorial and ecliptic coordinates
- Advances in cryptography and cryptanalysis

Pascaline

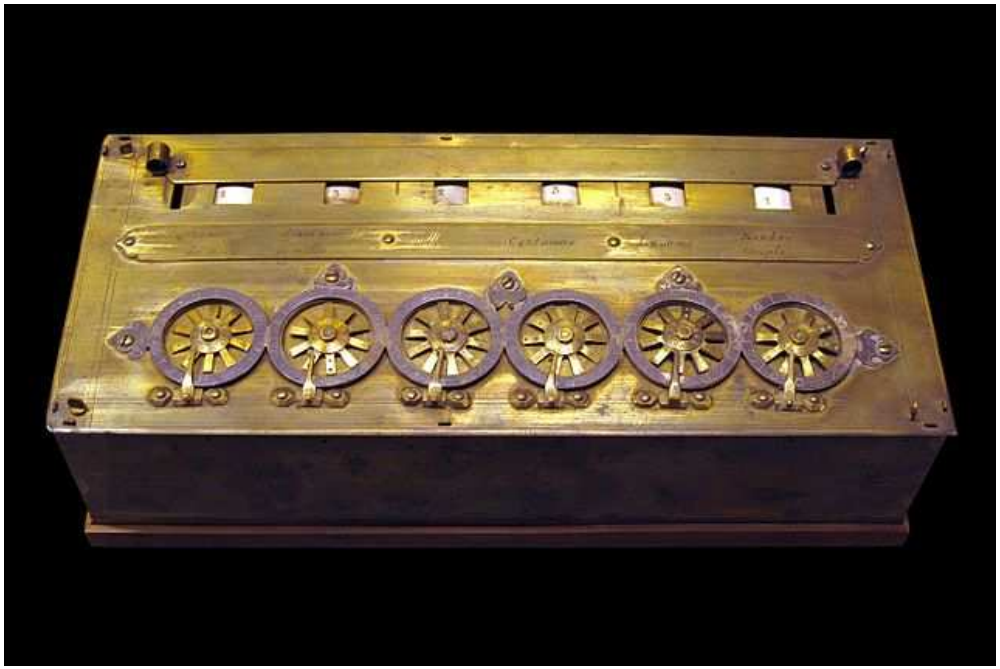

Source: Wikimedia

- Designed by Blaise Pascal in 1642
- Used to do arithmetic operations in decimal
- Handled automatic transfer of carry

Difference Engine

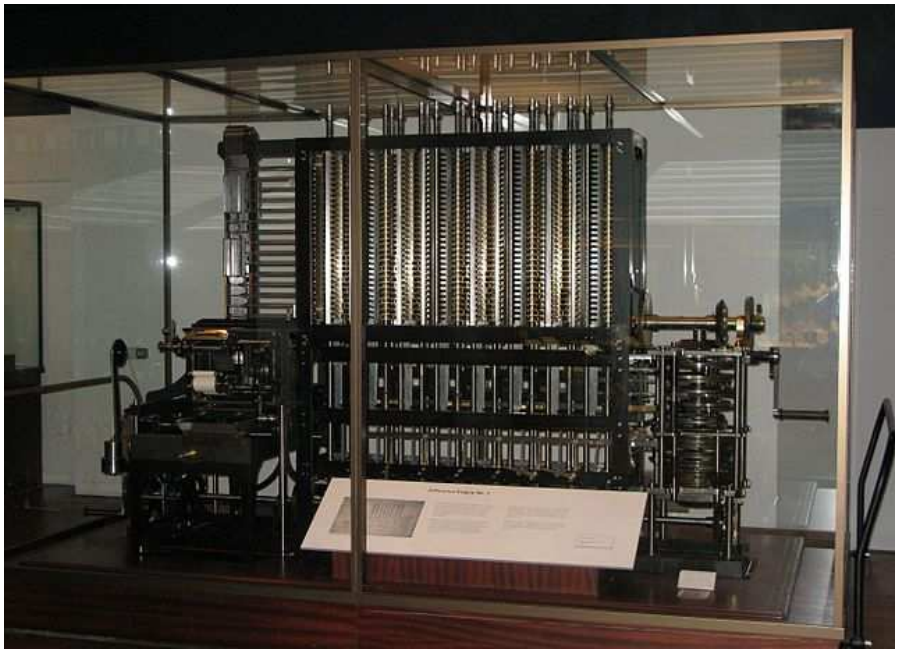

Source: Wikimedia

- Charles Babbage started designing this machine in 1823
- It was a mechanical calculator
- Used polynomial computations

Analytical Engine

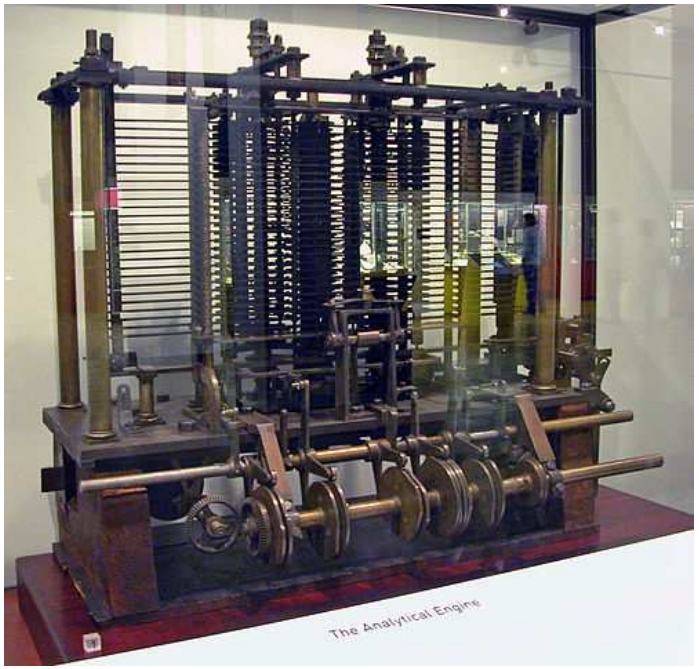

Source: Wikimedia

- Proposed by Charles Babbage in 1837
- It was a general-purpose mechanical computer
- Had arithmetic and logic unit and memory
- Could handle conditional branching and loops
- Later proved to be Turing-complete

Tide Predictor

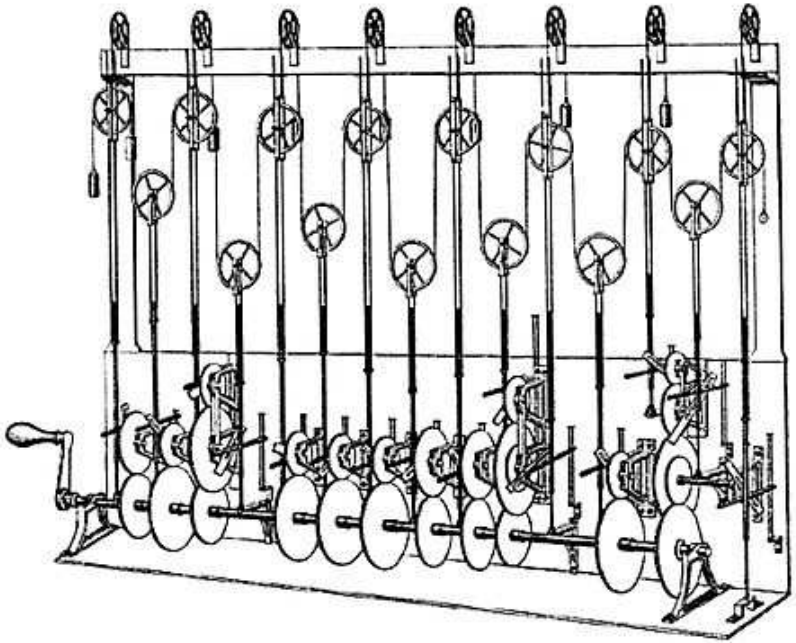

Source: Wikimedia

- First modern mechanical analog computer
- Designed by Sir William Thomson (aka Lord Kelvin)

Differential Analyzer

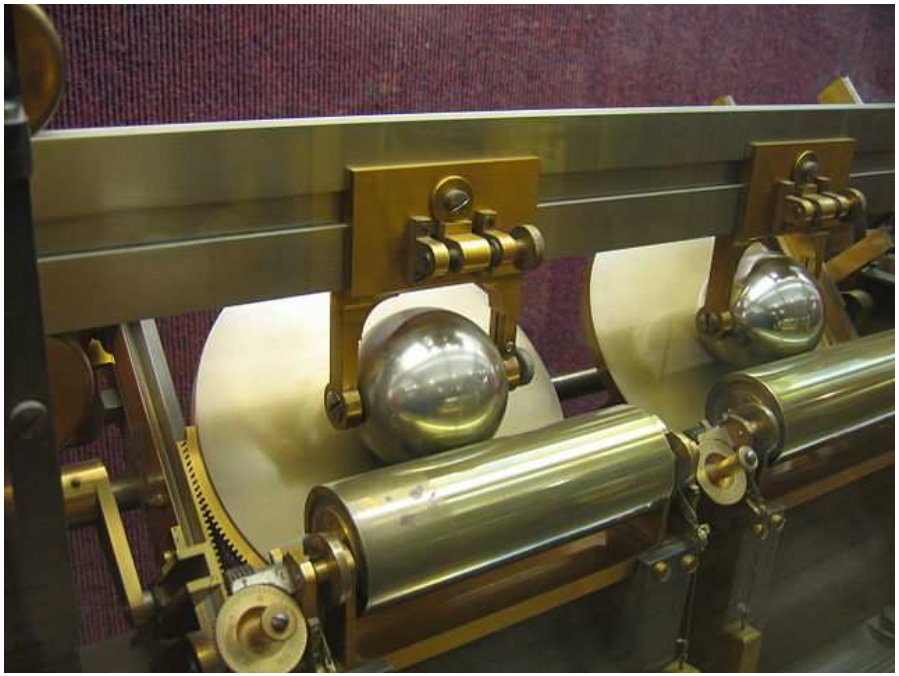

Source: Wikimedia

- Another mechanical analog computer
- Conceptualized by James Thomson (brother of Lord Kelvin)
- Built by H. L. Hazen and V. Bush at MIT starting in 1927

The Age of Digital Computers

- First conceptualized by Alan Mathison Turing in 1936
- John von Neumann proposed a popular architecture for stored-program computers in 1945
- Based on Boolean algebra (George Boole, 1854) which is connected to computation by Claude Shannon and Victor Shestakov in 1930s
- Initial designs were electromechanical
- Superseded by electrical devices (valves or vacuum tubes)
- Eventually, valves were replaced by electronic devices (transistors and diodes)
- Advent of VLSI (very large scale integration) technology helped to dramatically reduce the size of computers

Z1

Source: Wikimedia

- Mechanical computer designed (1935–36) and built (1936–38) by Konrad Zuse
- First known programmable computer
- Not very reliable in operation
- Destroyed by allied bombing in Berlin in 1943
- Superseded by Z2, Z3, Z4.

Bombe

Source: Wikimedia

- Initial design by Alan Turing in 1939, built by Harold Keen
- Electromechanical device used to decrypt German messages encrypted by Enigma during WWII

Atanasoff–Berry Computer (ABC)

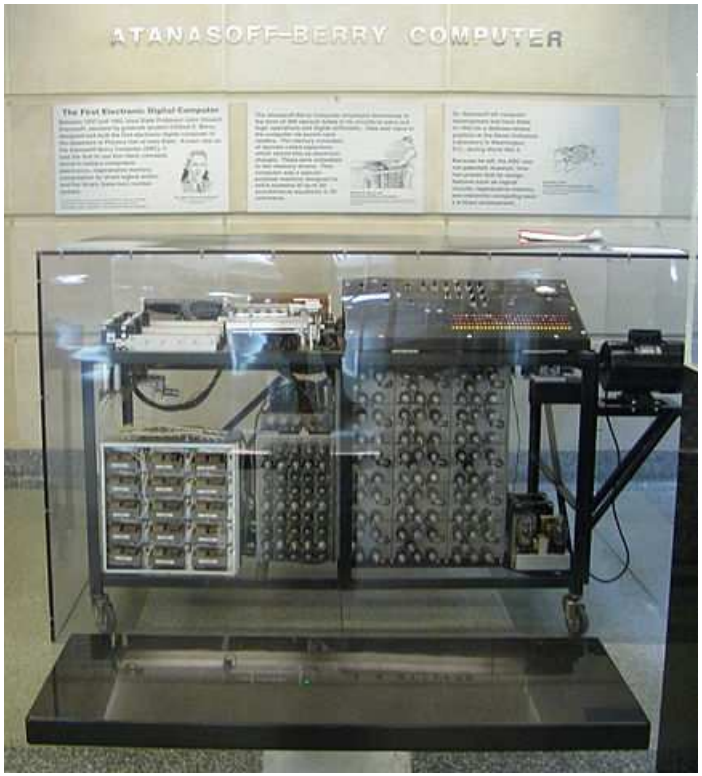

Source: Wikimedia

- First automatic electronic digital computer
- Designed by John Vincent Atanasoff and Clifford Berry in Iowa State College in 1942
- Used only to solve systems of linear equations
- Not programmable

Colossus

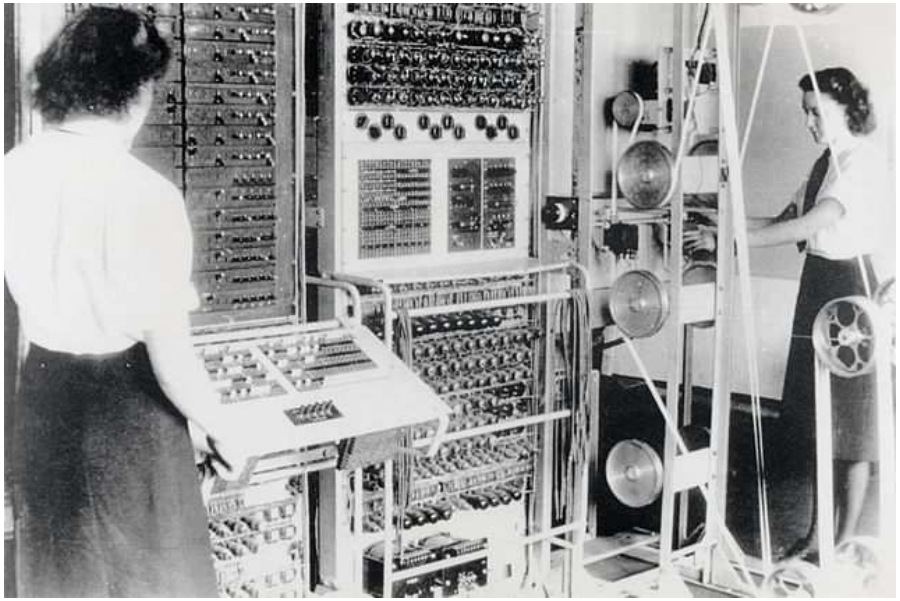

Source: Wikimedia

- World's first programmable, electronic, digital computer
- Programmed by switches and plugs
- Designed by Tommy Flowers in 1943
- Used by British code-breakers during 1943–45 for the cryptanalysis of the Lorenz cipher

ENIAC: Electronic Numerical Integrator and Computer

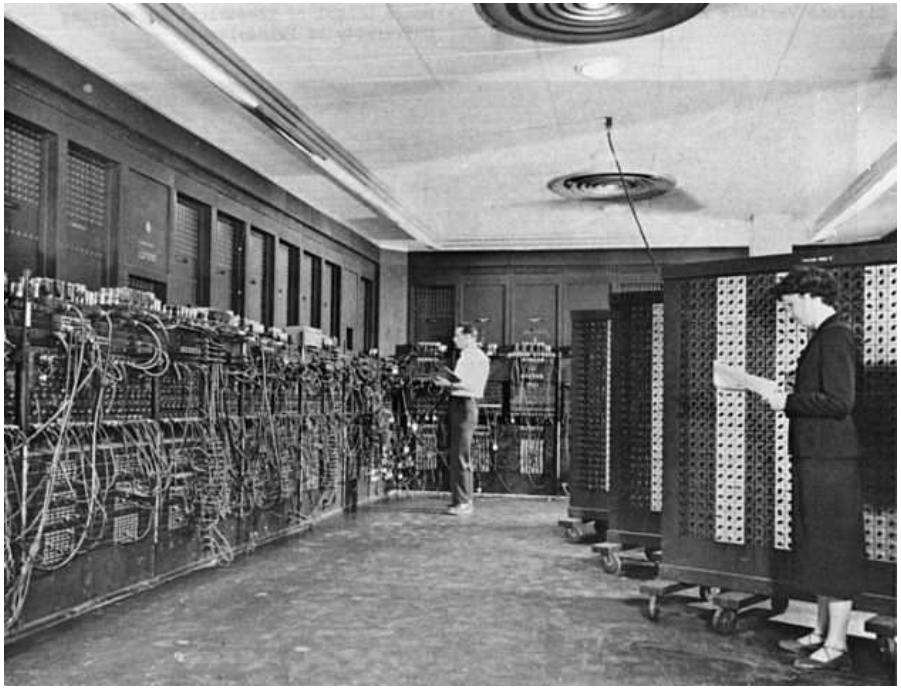

Source: Wikimedia

- First electronic programmable computer built in the US
- Built during 1943–45 by John Mauchly and J. Presper Eckert of University of Pennsylvania, funded by US Army
- Turing-complete, much faster than Colossus
- Programmed by cables and switches

EDVAC: Electronic Discrete Variable Automatic Computer

Source: Wikimedia

- Successor of ENIAC, built by Mauchly and Eckert during 1946–49 under consultation with John von Neumann
- Electronic stored-program computer
- Ran reliably until 1961 (replaced by BRLESC)

Manchester Baby

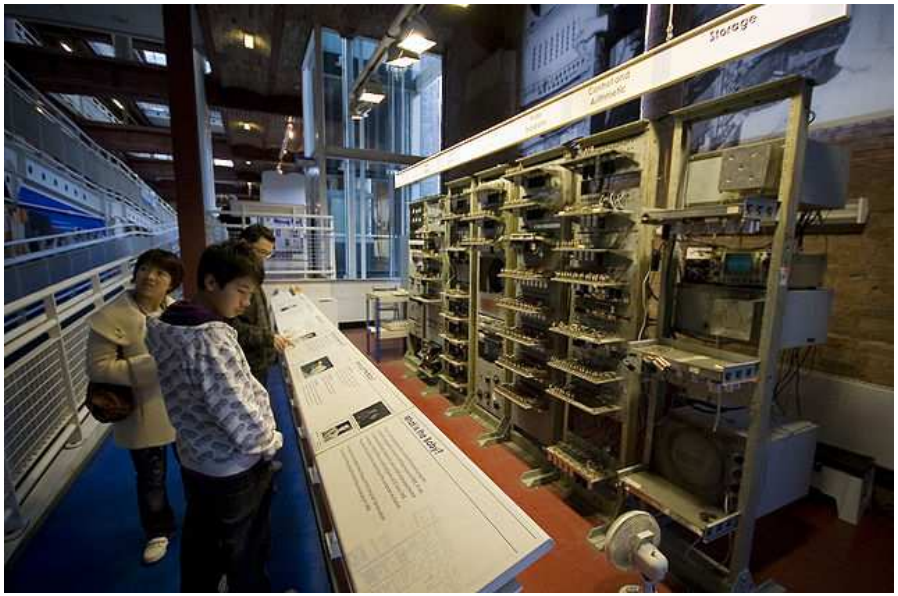

Source: Wikimedia

- Small-Scale Experimental Machine (SSEM)
- World's first stored-program computer
- Built by Frederic C. Williams, Tom Kilburn and Geoff Tootill in Victoria University of Manchester
- First program ran on 21 June 1948

Manchester Ferranti

Source: Wikimedia

- Aka Ferranti Mark 1 and Manchester Electronic Computer
- World's first commercial general-purpose electronic computer
- Designed at the University of Manchester by Freddie Williams and Tom Kilburn
- Built by UK firm Ferranti International PLC
- Sale started in 1951

UNIVAC I: UNIVersal Automatic Computer I

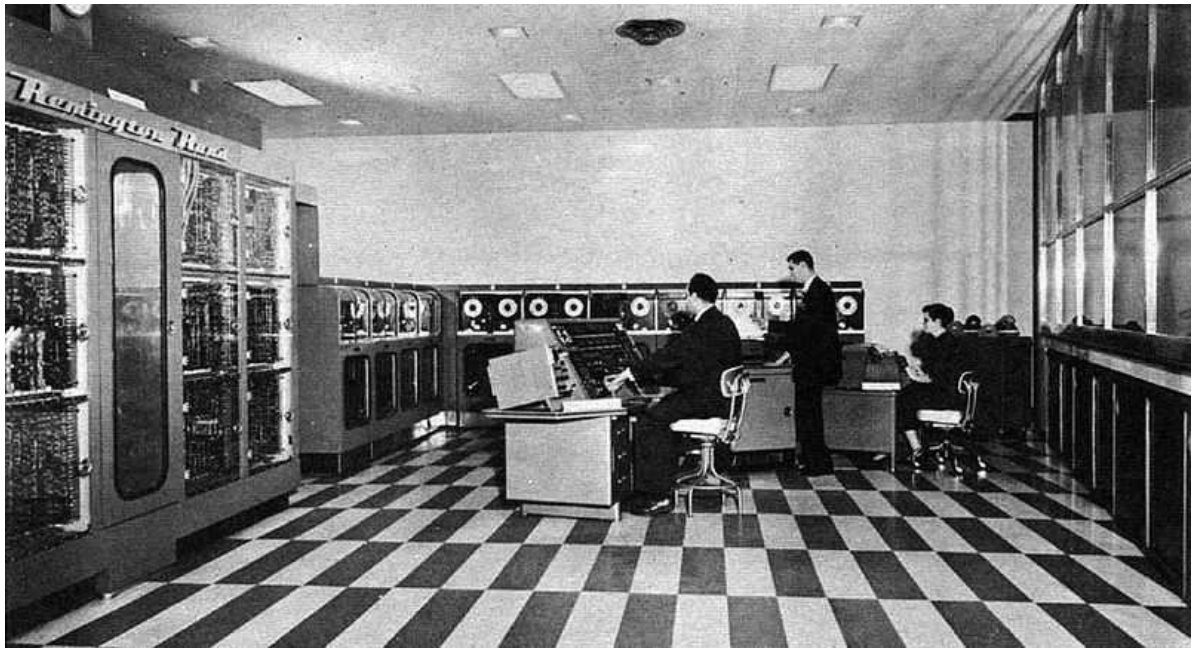

Source: Wikimedia

- First commercial computer in US
- Designed by Eckert–Mauchly Computer Corporation
- Produced from 1951

EDSAC 2: Electronic Delay Storage Automatic Calculator 2

Source: Wikimedia

- Concept of microprogramming was introduced by Sir Maurice Vincent Wilkes of University of Cambridge
- EDSAC2 was designed by Wilkes in 1958
- First computer to have a microprogrammed control unit

Harwell CADET: Transistor Electronic Digital Automatic Computer

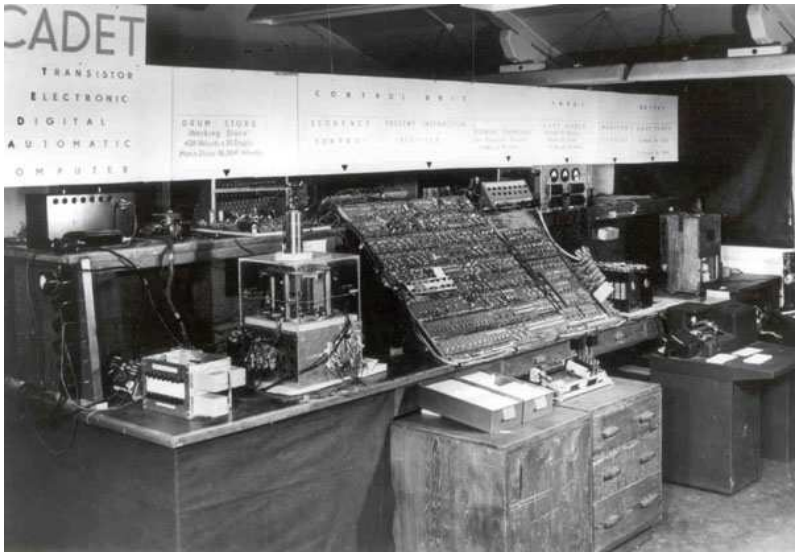

Source: Wikimedia

- Bipolar transistor was invented in 1947 by Bardeen, Brattain, and Shockley of Bell Labs
- A partially transistorized computer was built by Tim Kilburn at University of Manchester in 1953
- CADET was world's first fully transistorized computer
- Built in 1955 by Atomic Energy Research Establishment, UK

Atlas and IBM 7030 Stretch

Source: Wikimedia

- World's first supercomputers
- Atlas was developed jointly by University of Manchester, Ferranti, and Plessey
- Stretch was developed at IBM

CDC 6600

Source: Wikimedia

- World's first successful supercomputer
- Designed by Seymour Cray at Control Data Corporation
- World's fastest computer from 1964 to 1969
- Superseded by CDC 7600

Integrated Circuits

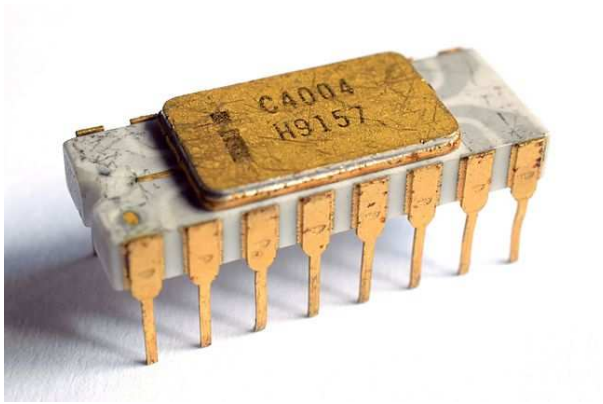

Source: Wikimedia

- Geoffrey William Arnold Dummer first conceptualized the concept of integrated circuits (microchips) and demonstrated in 1952
- In 1958, Jack Kilby at Texas Instruments and Robert Noyce at Fairchild Semiconductor developed first practical ICs
- The first single-chip microprocessor was the Intel 4004
- It was designed and realized by Ted Hoff, Federico Faggin, and Stanley Mazor at Intel in 1971
- It had a CPU clock rate of 740 kHz

Personal Computers

Source: Wikimedia

- Altair 8800: an 8-bit microcomputer, designed in 1974 by Micro Instrumentation and Telemetry Systems, the first commercially successful personal computer
- Apple II: an 8-bit home computer, designed primarily by Steve Wozniak, released in June 1977, the first highly successful mass-produced microcomputer

Laptop Computers

Source: Wikimedia

- Epson HX-20 (aka HC-20) was the first laptop computer
- Invented in July 1980 by Yukio Yokozawa at Seiko, Japan
- Introduced by Epson in North America as HX-20 in 1981
- Weight: 1.6 kg
- It was both the first notebook and the first hand-held computer

Q7: The Largest Computer Ever Built

Source: Wikimedia

- AN/FSQ-7 Combat Direction Central was built in 1958 for USAF
- Weight: 250 Tons
- 60,000 vacuum tubes
- 3 megawatts of electricity
- 75,000 instructions per second

Today: Computers are Everywhere

Source: theengineer.co.uk

Programming Languages and Compilers

How to Specify Algorithms to Computers?

- An algorithm is a recipe to solve a computational problem
- First such algorithms were written by Ada Lovelace in the 19th Century for Babbage's Analytical Engine
- Alan Turing's 1935 essay gave birth to computer science and software engineering
- Initially, programs are specified to machines by switches, punched cards, or magnetic tapes
- Then came programming in assembly languages
- Finally, high-level languages were developed

Plankalkül

- Developed for the German Z3 by Konrad Zuse between 1943 and 1945
- Not published/publicized immediately
- First compiler for it was implemented in 1975 by Joachim Hohmann

```
P1 max3 (V0[:8.0],V1[:8.0],V2[:8.0]) -> R0[:8.0]
max(V0[:8.0],V1[:8.0]) -> Z1[:8.0]
max(Z1[:8.0],V2[:8.0]) -> R0[:8.0]
END
P2 max (V0[:8.0],V1[:8.0]) -> R0[:8.0]
V0[:8.0] -> Z1[:8.0]
(Z1[:8.0] < V1[:8.0]) -> V1[:8.0] -> Z1[:8.0]
Z1[:8.0] -> R0[:8.0]
END
```

Short Code

- Proposed by John Mauchly in 1949
- William Schmitt implemented it for BINAC and UNIVAC I
- Manual conversion needed
- The converted code was interpreted
- Supported branching and some function calls

```
X3 =   (  X1 +  Y1 )  /  X1 * Y1    substitute variables
X3 03 09 X1 07 Y1 02 04 X1   Y1    substitute operators and parentheses.
                                        Note multiplication is represented
                                        by juxtaposition.
07Y10204X1Y1                        group into 12-byte words.
0000X30309X1
```

Autocode

- First compiled programming language
- The language and its compiler were developed by Alick Glennie in 1952 for the Mark 1 computer at University of Manchester

```
c@VA t@IC x@1/2C y@RC z@NC
INTEGERS +5 ->c            # Put 5 into c
      ->t                  # Load argument from lower accumulator to variable t
   +t     TESTA Z          # Put |t| into lower accumulator
   -t
          ENTRY Z
SUBROUTINE 6 ->z           # Run square root subroutine on lower accumulator
                           # value and put the result into z
  +tt ->y ->x              # Calculate t^3 and put it into x
  +tx ->y ->x
+z+cx   CLOSE WRITE 1      # Put z + (c * x) into lower accumulator and return
```

FORTRAN

- First widely used high-level general purpose programming language
- Invented at IBM by John Backus in 1954
- Still popular language for high-performance computing

```
C AREA OF A TRIANGLE WITH A STANDARD SQUARE ROOT FUNCTION
C INPUT - TAPE READER UNIT 5, INTEGER INPUT
C OUTPUT - LINE PRINTER UNIT 6, REAL OUTPUT
      READ INPUT TAPE 5, 501, IA, IB, IC
  501 FORMAT (3I5)
C CHECK FOR VALID INPUT
      IF (IA) 777, 777, 701
  701 IF (IB) 777, 777, 702
  702 IF (IC) 777, 777, 703
  703 IF (IA+IB-IC) 777, 777, 704
  704 IF (IA+IC-IB) 777, 777, 705
  705 IF (IB+IC-IA) 777, 777, 799
  777 STOP 1
C USING HERON'S FORMULA WE CALCULATE THE AREA OF THE TRIANGLE
  799 S = FLOATF (IA + IB + IC) / 2.0
      AREA = SQRTF( S * (S - FLOATF(IA)) * (S - FLOATF(IB)) *
     +      (S - FLOATF(IC)))
      WRITE OUTPUT TAPE 6, 601, IA, IB, IC, AREA
  601 FORMAT (4H A= ,I5,5H  B= ,I5,5H  C= ,I5,8H  AREA= ,F10.2,
     +        13H SQUARE UNITS)
      STOP
      END
```

LISP

- LISt Processor
- Invented by John McCarthy in MIT in 1958, Lisp is the second-oldest high-level programming language in widespread use today
- First functional programming language
- Influenced by Alonzo Church's lambda calculus
- It quickly became the favored programming language for artificial intelligence

```
(defun factorial (n)
(if (= n 0) 1
    (* n (factorial (- n 1)))))
```

FLOW-MATIC

- English-like data-processing language
- Developed for UNIVAC I under Grace Hopper during 1955–59
- Influenced the design of COBOL

```
 (0)  INPUT INVENTORY FILE-A PRICE FILE-B ; OUTPUT PRICED-INV FILE-C UNPRICED-INV
     FILE-D ; HSP D .
 (1)  COMPARE PRODUCT-NO (A) WITH PRODUCT-NO (B) ; IF GREATER GO TO OPERATION 10 ;
     IF EQUAL GO TO OPERATION 5 ; OTHERWISE GO TO OPERATION 2 .
 (2)  TRANSFER A TO D .
 (3)  WRITE-ITEM D .
 (4)  JUMP TO OPERATION 8 .
 (5)  TRANSFER A TO C .
 (6)  MOVE UNIT-PRICE (B) TO UNIT-PRICE (C) .
 (7)  WRITE-ITEM C .
 (8)  READ-ITEM A ; IF END OF DATA GO TO OPERATION 14 .
 (9)  JUMP TO OPERATION 1 .
(10)  READ-ITEM B ; IF END OF DATA GO TO OPERATION 12 .
(11)  JUMP TO OPERATION 1 .
(12)  SET OPERATION 9 TO GO TO OPERATION 2 .
(13)  JUMP TO OPERATION 2 .
(14)  TEST PRODUCT-NO (B) AGAINST ZZZZZZZZZZZZ ; IF EQUAL GO TO OPERATION 16 ;
     OTHERWISE GO TO OPERATION 15 .
(15)  REWIND B .
(16)  CLOSE-OUT FILES C ; D .
(17)  STOP . (END)
```

ALGOL

- ALGOrithmic Language, designed during 1958–68
- Strongly influences design of Simula, BCPL, B, Pascal, C
- First language to support nested functions and lexical scoping
- Sir Charles Antony Richard Hoare (inventor of Quick Sort) wrote a compiler for ALGOL 60 that supports recursion
- Presented in the Backus–Naur form

```
procedure Absmax(a) Size:(n, m) Result:(y) Subscripts:(i, k);
    value n, m; array a; integer n, m, i, k; real y;
comment The absolute greatest element of the matrix a, of size n by m
    is transferred to y, and the subscripts of this element to i and k;
begin
    integer p, q;
    y := 0; i := k := 1;
    for p := 1 step 1 until n do
        for q := 1 step 1 until m do
            if abs(a[p, q]) > y then
                begin y := abs(a[p, q]);
                    i := p; k := q
                end
end Absmax
```

B

- Developed at Bell Labs circa 1969 by Ken Thompson and Dennis Ritchie
- Introduced the currently used operator symbols (= for assignment, ==, ++, =+, and so on)

```
/* The following function will print a non-negative number, n, to
   the base b, where 2<=b<=10.  This routine uses the fact that
   in the ASCII character set, the digits 0 to 9 have sequential
   code values.  */

printn(n, b) {
        extrn putchar;
        auto a;

        if (a = n / b)          /* assignment, not test for equality */
                printn(a, b); /* recursive */
        putchar(n % b + '0');
}
```

C

Source: Wikimedia

- Developed by Dennis Ritchie during 1969–1973 at Bell Labs
- Used to re-implement the Unix operating system

Simula

- Developed in the 1960s at the Norwegian Computing Center by Ole-Johan Dahl and Kristen Nygaard.
- The first object-oriented programming language
- Introduced objects, classes, inheritance and subclasses, virtual procedures, coroutines, discrete event simulation, and garbage collection

```
Glyph Class Line (elements);
   Ref (Glyph) Array elements;
Begin
   Procedure print;
   Begin
      Integer i;
      For i:= 1 Step 1 Until UpperBound (elements, 1) Do
         elements (i).print;
      OutImage;
   End;
End;
```

C++

Source: Wikimedia

- C with Classes
- Developed by Bjarne Stroustrup during 1979–1983
- Influenced by Simula

Prolog

- General-purpose logic programming language
- Developed in 1972 by Alain Colmerauer with Philippe Roussel
- Used for theorem proving, expert systems, term rewriting, type inference, automated planning, and natural language processing

```
mother_child(trude, sally).

father_child(tom, sally).
father_child(tom, erica).
father_child(mike, tom).

sibling(X, Y)      :- parent_child(Z, X), parent_child(Z, Y).

parent_child(X, Y) :- father_child(X, Y).
parent_child(X, Y) :- mother_child(X, Y).
```

Markup Languages

- GenCode: The first well-known public presentation of markup languages in computer text processing was made by William W. Tunnicliffe in 1967
- troff and nroff: Introduced in Unix
- TEX: Created by Donald Knuth in the 1970s and '80s
- LATEX: A macro package designed by Leslie Lamport in 1985 to make TEX easy to use
- HTML: Internet-based hypertext system conceived and initially designed by Sir Tim Berners-Lee

PostScript

- Language for printers
- Created at Adobe Systems by John Warnock, Charles Geschke, Doug Brotz, Ed Taft and Bill Paxton during 1982–84
- Superseded by PDF (portable document format)

```
%!PS
/Courier              % name the desired font
20 selectfont         % choose the size in points and establish
                      % the font as the current one
72 500 moveto         % position the current point at
                      % coordinates 72, 500 (the origin is at the
                      % lower-left corner of the page)
(Hello world!) show   % stroke the text in parentheses
showpage
```

Operating Systems

Evolution

- OS = Libraries + Links to Hardware + Utilities
- First computers had no OS. Users were given exclusive access and were needed to supply full hardware specification to control the drivers
- Support code (libraries) were eventually developed
- Libraries were linked to user programs to assist I/O
- OSes started auditing access and usage information of programs
- Then came automated batch processing
- New features were introduced: GUIs, utilities, applications, configuration tools
- Eventually, kernels were developed

GM-NAA I/O: The First Operating System

Source: Wikimedia

- Created in 1956 by Robert L. Patrick of General Motors Research and Owen Mock of North American Aviation
- Built for IBM 704
- Supported batch processing and shared I/O libraries

Multiprogramming

Source: leo-computers.org.uk

- LEO III (Lyons Electronic Office III), a British computer developed in 1961, was the first computer to introduce multiprogramming
- It did batch processing
- Multiple programs are loaded to memory
- When a program reached an instruction waiting for a peripheral, the context of this program was stored away, and the next program in memory was given a chance to run

Cooperative Multitasking

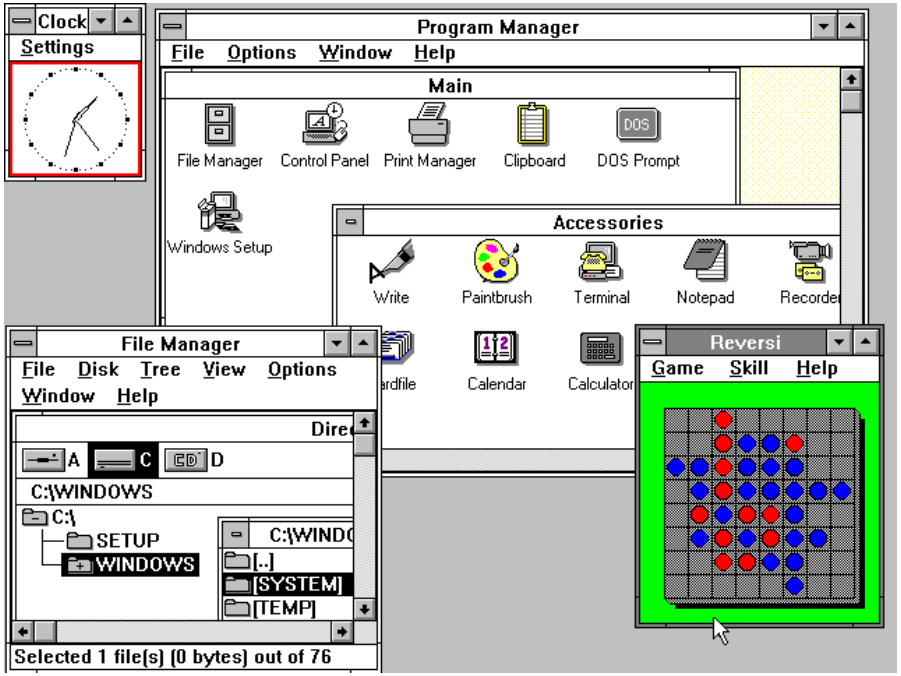

Source: Wikimedia

- Programs have embedded codes to relinquish the control of CPU to other programs
- Introduced in late 60s and early 70s in IBM TSO (Time Sharing Option) and VM/CMS (Virtual Machine/Cambridge Monitor System)
- Microsoft Windows 3.* used cooperative multitasking

Preemptive Multitasking

Source: Wikimedia

- Each process runs for a regular slice of operating time and is preempted from the CPU after the slice
- Introduced in 1964 in TOPS-10 (Total Operating System-10) from DEC and Multics (Multiplexed Information and Computing Service) jointly from MIT, GE, and Bell Labs
- Standard feature in Unix and Unix-like OSes (Linux, BSD, Solaris, MAC OS X)
- Introduced from Windows 9x and NT

Kernel

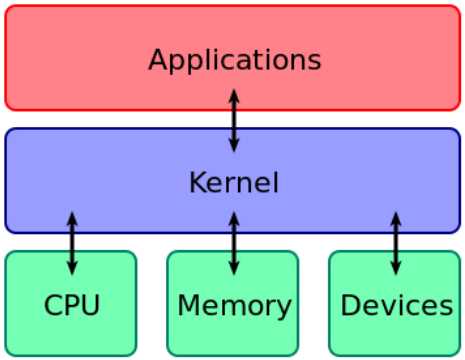

Source: Wikimedia

- Early computers used bare metal approach: no hardware abstraction or OS support
- Later, program loaders and debuggers were left in memory between runs, or loaded from ROM
- RC 4000 Multiprogramming System designed in 1969 was the first attempt to break down an operating system into a group of interacting programs communicating via a message-passing kernel (microkernel approach)
- Unix and its derivatives implemented monolithic kernels where all OS services run in the main kernel thread

Unix

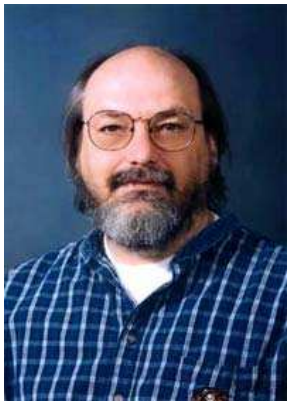

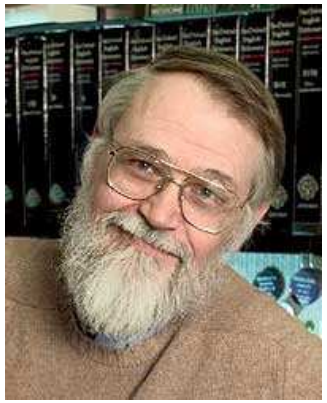

Source: computerhope.com

- Developed by Ken Thompson, Dennis Ritchie, Brian Kernighan, Douglas McIlroy, and Joe Ossanna at Bell Labs from 1969
- At present, fully compliant implementations do not exist
- Many Unix-like systems are popular including Linux and Android

Linux

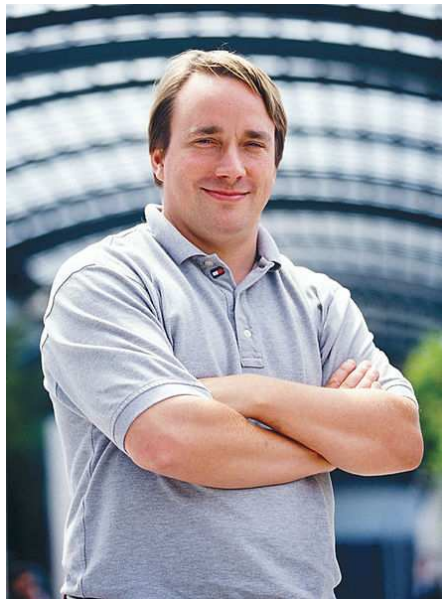

Source: Wikimedia

- In 1983, Richard Stallman founded Free Software Foundation to design a complete Unix-like operating system, called GNU, composed entirely of free software
- Independently, in 1991, Linus Torvalds released the first version of the Linux kernel
- Better name: GNU/Linux, or Lignux

Software Engineering

- Application of engineering to the development of software in a systematic method
- The discipline was created to ensure that software is built systematically, rigorously, measurably, on time, on budget, and within specification
- Research, design, develop, and test operating systems-level software, systems-level software, compilers, and network distribution software
- Edsger W. Dijkstra wrote his seminal paper *Go To Statement Considered Harmful* in 1968
- David Parnas introduced the key concept of modularity and information hiding in 1972
- The first conference on software engineering was sponsored by NATO in 1968

Computer Networks and the Internet

Evolution

- ARPAnet (Advanced Research Projects Agency network) began as a military computer network in 1969
- National Science Foundation (NSF) introduced the concept and proliferation of a global computer network
- NSFNET was born in 1986
- The Internet rapidly expanded in Europe and Australia in mid to late 1980s and to Asia in late 1980s and early 1990s
- Dedicated transatlantic communication between the NSFNET and networks in Europe was established with a low-speed satellite relay between Princeton University and Stockholm, Sweden in December 1988
- For about a decade, the Internet was primarily used by researchers and academics
- From 1995, proliferation to a large number of customers started
- Now: About 1 billion sites and about 5 billion pages

The World-Wide Web

Source: Wikimedia

By Christmas 1990, Sir Tim Berners-Lee built all the tools necessary for a working Web:

- HyperText Transfer Protocol (HTTP) 0.9
- HyperText Markup Language (HTML)
- First Web browser (with an HTML editor)
- First HTTP server software (CERN HTTPD)
- First web server (Berner-Lee's NeXT computer)
- First web pages (info.cern.ch)

The First Web Page

World Wide Web

The WorldWideWeb (W3) is a wide-area hypermedia information retrieval initiative aiming to give universal access to a large universe of documents.

Everything there is online about W3 is linked directly or indirectly to this document, including an executive summary of the project, Mailing lists , Policy , November's W3 news , Frequently Asked Questions .

What's out there?
Pointers to the world's online information, subjects , W3 servers, etc.
Help
on the browser you are using
Software Products
A list of W3 project components and their current state. (e.g. Line Mode ,X11 Viola , NeXTStep , Servers , Tools , Mail robot , Library)
Technical
Details of protocols, formats, program internals etc
Bibliography
Paper documentation on W3 and references.
People
A list of some people involved in the project.
History
A summary of the history of the project.
How can I help ?
If you would like to support the web..
Getting code
Getting the code by anonymous FTP , etc.

Source: w3.org

CERN = Conseil Européen pour la Recherche Nucléaire
(European Council for Nuclear Research)

Electronic Business

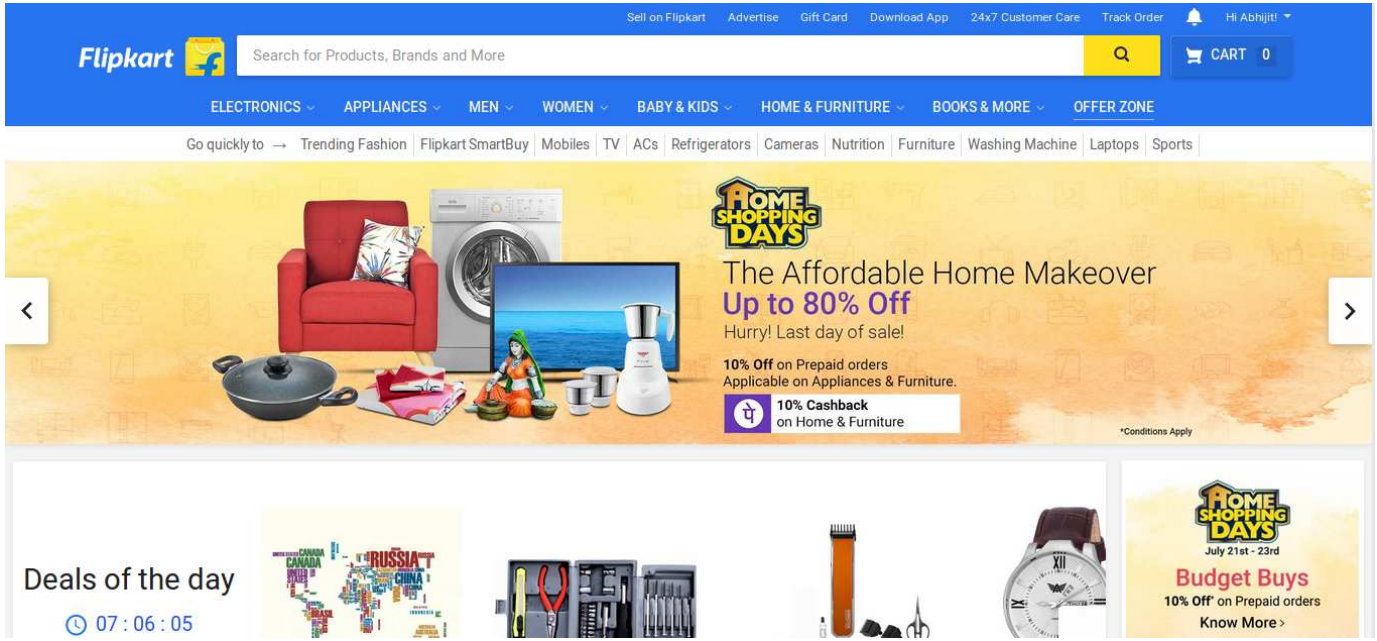

Source: flipkart.com

- Started in 1994 by IBM with its agency Ogilvy and Mather
- Applications: Mobile commerce, electronic funds transfer, supply chain management, Internet marketing, online transaction processing, electronic data interchange, inventory management systems, and automated data collection
- Concerns: Security, privacy and confidentiality, authenticity, data integrity, access control, availability, cost

Wiki

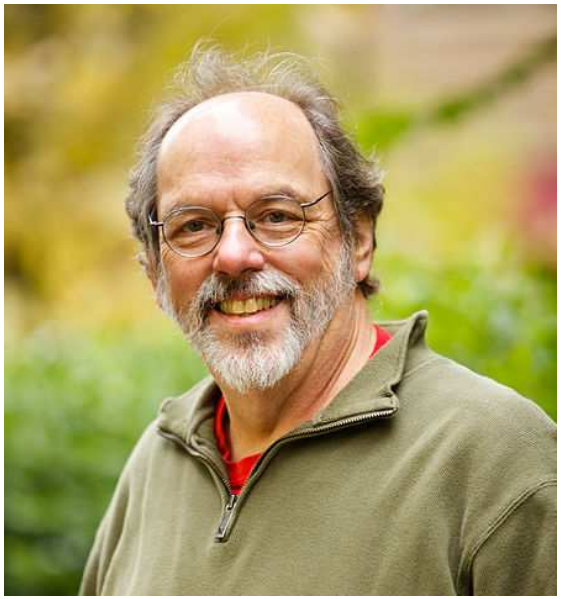

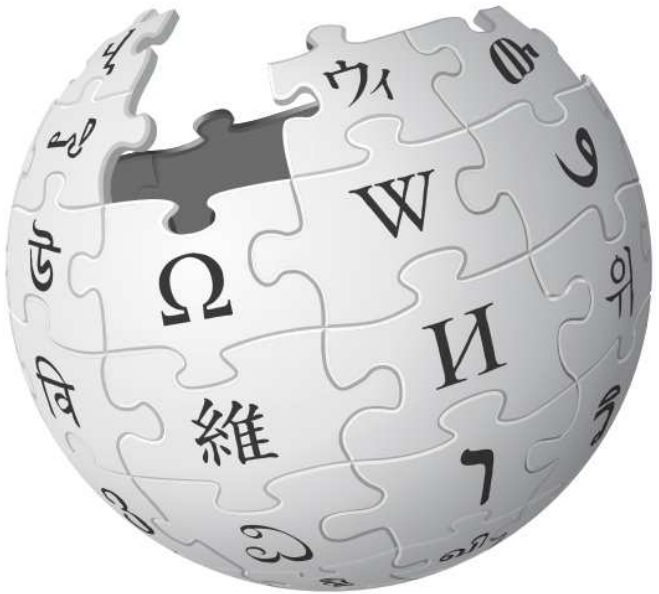

Source: Wikimedia

- A wiki is a website on which users collaboratively modify content and structure directly from the web browser
- All users—not just experts—can edit any page or create new pages within the wiki web site
- Ward Cunningham invented wikis, and launched the first wiki WikiWikiWeb in 1995

Image and Video Hosting

- Before the world wide web was developed, images were primarily downloaded manually from file servers or from binary newsgroups on Usenet
- Private hosting of images and videos was problematic even after the advent of the WWW
- Some early image hosting services still active are Webshots (launched in 1995), SmugMug (2002), and Flickr (2004)
- The first Internet video hosting site was shareyourworld.com founded in 1997 by Chase Norlin and closed in 2001 for budget and bandwidth problems
- Some early active video-hosting sites are Vimeo (2004) and YouTube (2005)

Social Networking

- Online platform which people use to build social networks or social relations with other people
- Earliest social networking sites are Geocities (1994, now defunct), TheGlobe.com (1995, now defunct), and Tripod.com (1995, active)
- New social networking methods were developed by the end of 1990s for users to find and manage friends
- Some early sites were SixDegrees (1997, now defunct), MakeOutClub (2000, defunct), Hub Culture (2002, active), Friendster (2002, paused)
- Facebook, launched in 2004, became the largest social networking site in the world in early 2009
- Facebook was first introduced as a Harvard social networking site, expanding to other universities and eventually, anyone

Teaching and Research

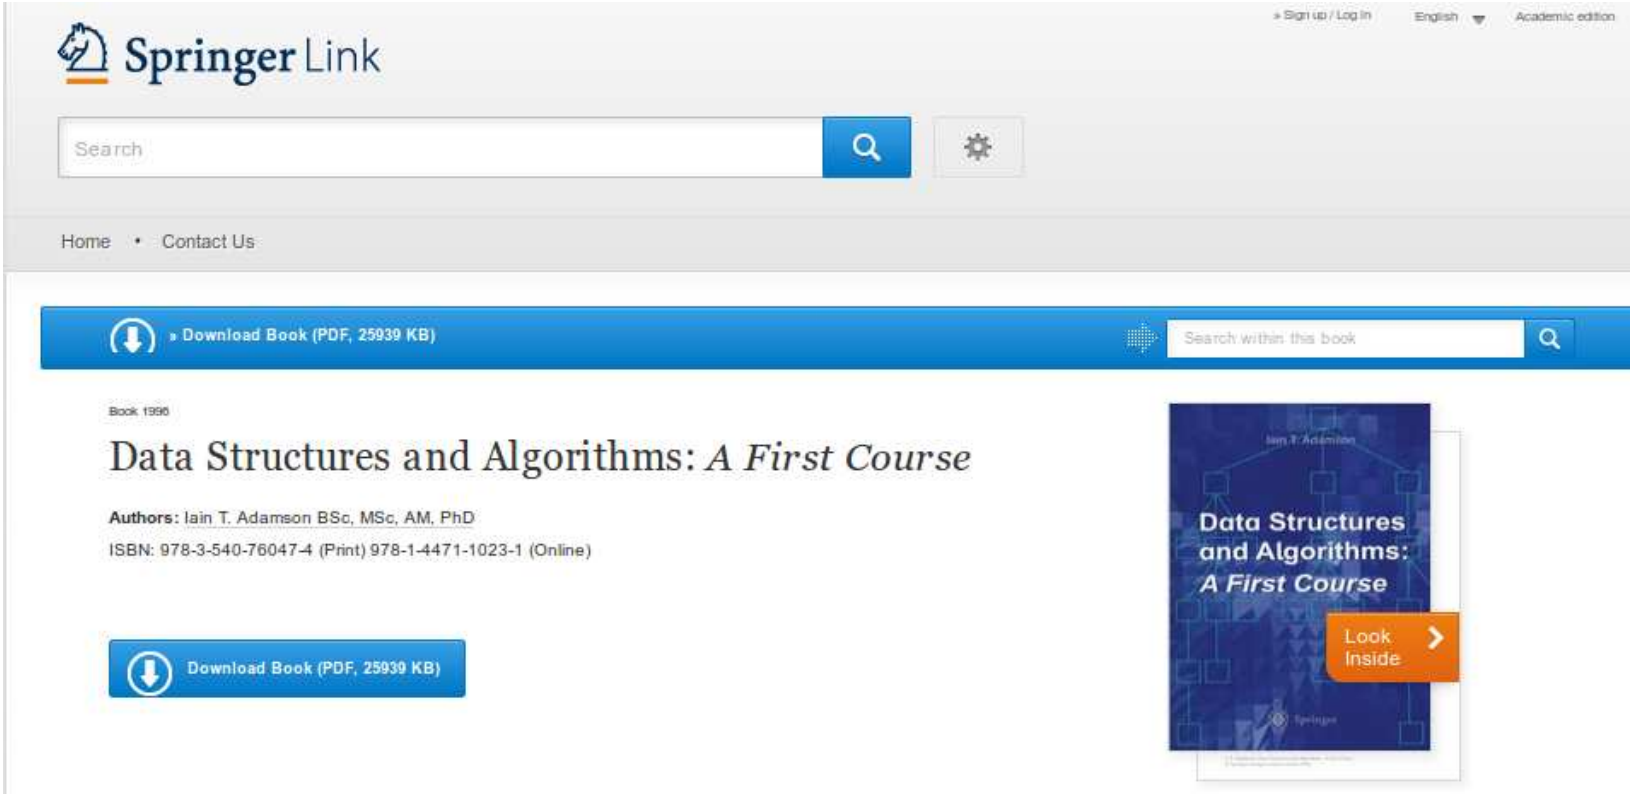

Source: link.springer.com

- Online tutoring
- MIT Open Courseware: https://ocw.mit.edu/
- Electronic books

Search Engines

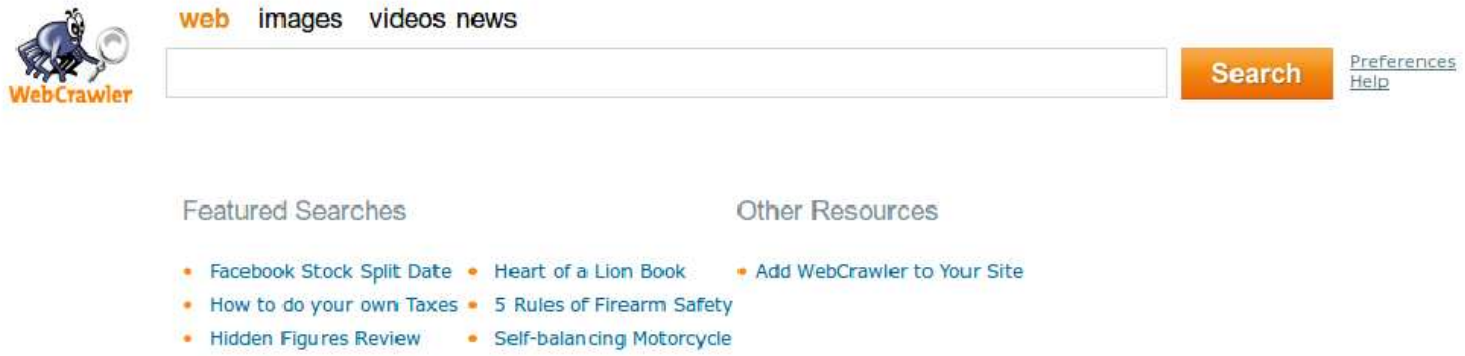

Source: webcrawler.com

- Requirements: Quick response, smart suggestions, user-friendly
- Archie (1990) was the first search engine with a database of FTP file names
- First content-based search engines: W3Catalog, Aliweb, JumpStation, WWW Worm (1993, all defunct now)
- WebCrawler and Lycos (1994) are still active
- Processes: Web crawling, indexing, and searching
- Google's search engine, launched in 1998, uses a page-ranking algorithm

Limits of Computing

What is Computation? What is an Algorithm?

- Any formula/procedure gives a recipe to solve a problem
- Examples: Sieve of Eratosthenes, Pythagoras theorem, Euclid's algorithm, Sridharacharya's formula, Panini's Sanskrit grammar
- What are the formal definitions of computation and algorithm?
- Can *everything* be computed?

Hilbert's Tenth Problem

Source: Wikimedia

- David Hilbert (1862–1943) posed 23 open problems in a Paris conference in 1900
- Tenth problem (*Entscheidungsproblem*): Find an **effective procedure** to find out whether a given multi-variate polynomial with integer coefficients has integer solutions

Lambda Calculus

Source: Wikimedia

- Proposed by Alonzo Church (1903–1995) in 1930s

Mu Recursive Functions

Source: Wikimedia

- Proposed by Kurt Friedrich Gödel (1906–1978) and Jacques Herbrand (1908–1931)

Type 0 Grammars

Source: Wikimedia

- Proposed by Avram Noam Chomsky (1928–)

Post Rewriting Systems

Source: Wikimedia

- Proposed by Emil Leon Post (1897–1954)

Combinatory Logic


Source: Wikimedia

- Proposed independently by Moses Ilyich Schönfinkel (1889–1942) in 1924 and Haskell Brooks Curry (1900–1982) in 1927

Markov String Rewriting Systems

Source: Wikimedia

- Proposed by Andrey (Andrei) Andreyevich Markov (1856–1922)

Turing Machines

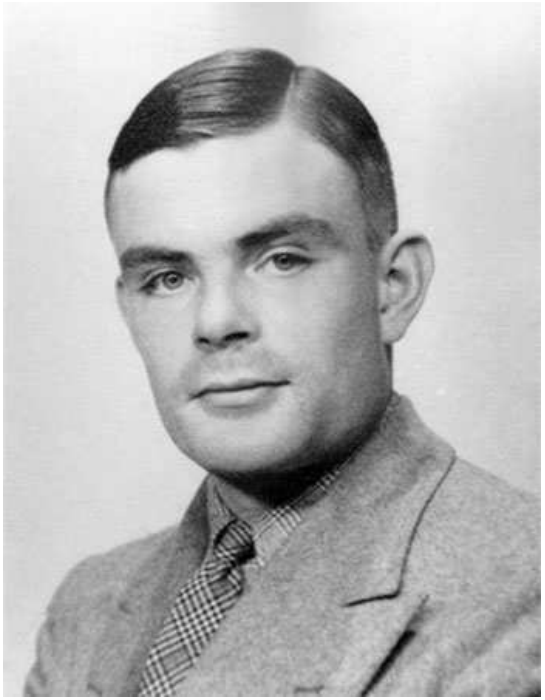

Source: Wikimedia

- Proposed by Alan Mathison Turing (1912–1954) in 1936

The Church–Turing Thesis

- All formulations of *effective computability* turned out to be equivalent
- You can add any high-level language or any Turing-complete computing device to the list
- **Conjecture:** A function on the natural numbers is computable by a human following an algorithm, ignoring resource limitations, if and only if it is computable by a Turing machine
- Algorithm = A Turing Machine that halts on every input
- Computation = The working of an algorithm

Unsolvable Problems

- **Halting Problem:** It is impossible to write a C program that, given another C program P and an input I for P, can decide whether P is going to halt on I
- Simulation does not help
- If the simulation runs for a huge time, it is impossible to say in general whether the simulation is solving a very difficult problem or is not going to halt at all

- **Hilbert's tenth problem** was proved to be unsolvable in 1970 by Martin Davis, Yuri Matiyasevich, Hilary Putnam and Julia Robinson
- The proof involved a combined effort of 21 years

- Suggested Documentary:
  Dangerous Knowledge from BBC (2007)

Complexity Theory

Source: Wikimedia

- Some problems are efficiently solvable by algorithms
- Some problems *need* hints for being efficiently solvable
- Hints are called non-determinism
- The million-dollar question of Computer Science:
  $\mathrm{P} \stackrel{?}{=} \mathrm{NP}$ (Are hints mandatory for some problems?)
- NP-Complete problems were introduced by Stephen Cook and Leonid Levin

Our Core Courses

- **Hardware:** Switching Circuits and Logic Design, Computer Organization and Architecture, High Performance Computer Architecture (for dual-degree students only)
- **Systems:** Compilers, Operating Systems, Computer Networks
- **Theory:** Discrete Structures, Algorithms–I, Algorithms–II, Formal Languages and Automata Theory, Theory of Computation
- **Others:** Software Engineering

A Look into the Future

Parallel and Distributed Computing

- Current-day semiconductor technology reached a saturation point—no significant increase in clock frequency is possible
- Hardware designers are supplying chips with more hardware
  - Multi-core processors are widely available since 2005
  - The first use of SIMD instructions dates back to the vector supercomputers of the early 1970s such as the CDC Star-100 and the Texas Instruments ASC
  - From late 1990s, SIMD registers are commonly included in machines from Intel, Motorola, and AMD
  - Graphic processing units (GPUs), initially designed for speeding up graphics applications, have been used for general-purpose computing
  - The trend of GPGPU computing may lead to wider use of SIMD in the future
- Need to learn parallel and distributed algorithms, and implementation tools targeted to specific architectures

Quantum Computing

Source: Wikimedia

- Conceived by Richard Phillips Feynman (1918–1988) in 1981
- Quantum computation does not violate the Church–Turing thesis
- Quantum computers can solve some problems faster than classical computers, like factoring integers
- Development of actual quantum computers is still in its infancy
- Hopes persist, efforts not stymied, some funding available

Embedded Systems and the Internet of Things

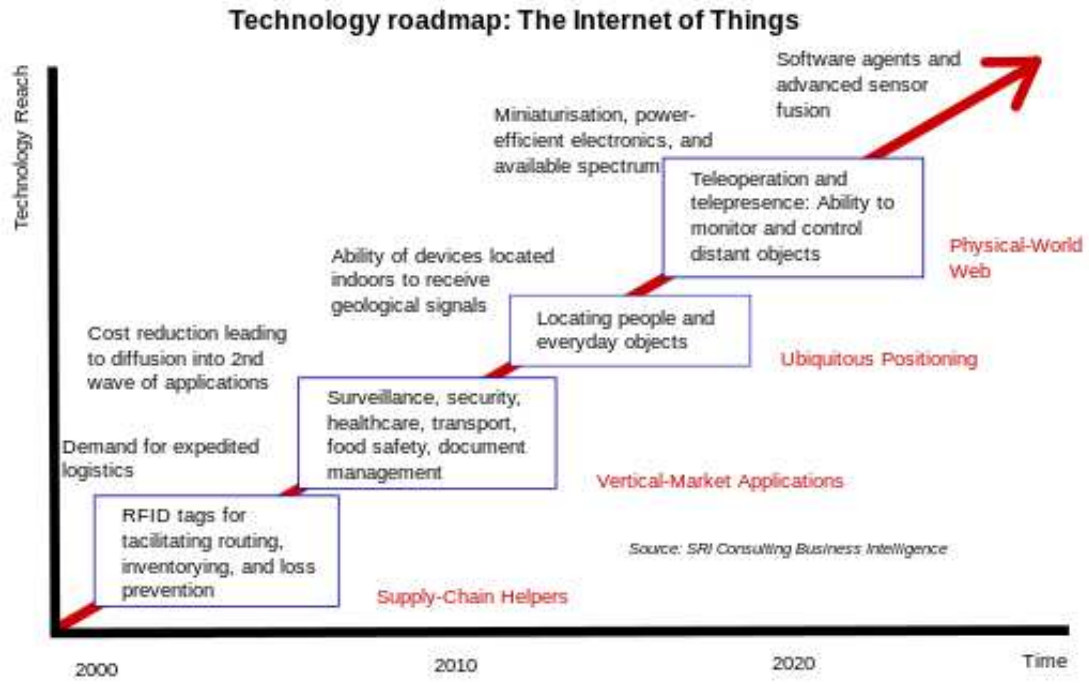

Source: Wikimedia

- Internet of things (IoT) is the connectivity among physical devices, vehicles, buildings, and other items
- Estimates suggest that 20–30 billion devices will be wirelessly connected to the IoT by 2020
- Issues: Standardization, privacy, security, data storage, impact on environment

Cryptography and Security

- Need to protect stored and communicated data from unauthorized access
- A requirement for military, electronic business, medical, and many other applications
- Future requirement: Protecting the IoT
- Cryptography deals with design of protocols to ensure confidentiality and authentication
- Cryptanalysis is the study of *directly* breaking the protocols
- Cryptography and cryptanalysis were intricately linked with the development of computers
- Data/Network/Information/Database/... security deals with other protection mechanisms, like against implementation bugs, system vulnerability, intruders, malware, misuse, denial of access

Machine Learning

- ML gives "computers the ability to learn without being explicitly programmed" (Arthur Samuel, 1959)
- Grew out of the quest for artificial intelligence (neural networks)
- Increased emphasis on logical knowledge-based approach caused a rift between AI and ML
- Applications: Computer vision, speech recognition, document classification, automated driving, computational science, and decision support
  - Filtering spam messages
  - Query completion
  - User-specific suggestions
  - Playing games against a human opponent
- Tools: Supervised, unsupervised, and reinforcement learning

Big Data Analysis

Source: Wikimedia

- Data sets too large or complex to fit in the RAM or local hard disks and to be handled by traditional software
- Examples: Multimedia files, content of the Internet, huge databases
- Challenges: Capture, storage, analysis, searching, sharing, transfer, visualization, querying, updating and privacy
- Techniques: Streaming, sketching, sampling

Clustering and Data Mining

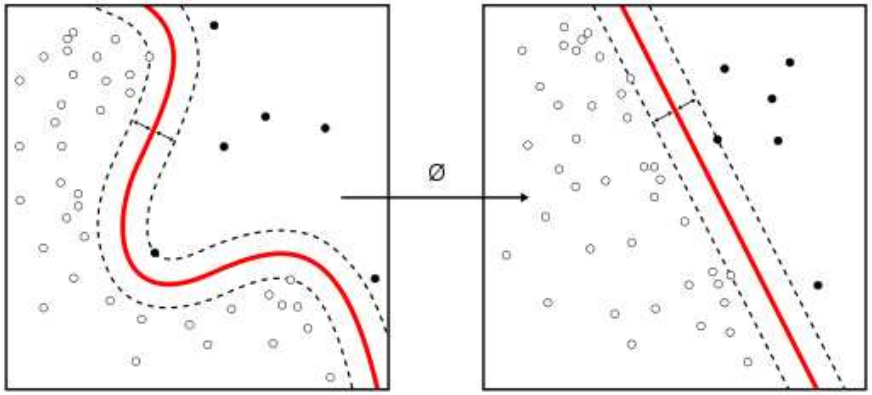

Source: Wikimedia

- Grouping a set of objects in such a way that objects in the same group are similar
- Used in machine learning, pattern recognition, image analysis, information retrieval, bioinformatics, data compression, and computer graphics
- Can be formulated as a multi-objective optimization problem
- Involves statistical data analysis

Random Graphs and Complex Networks

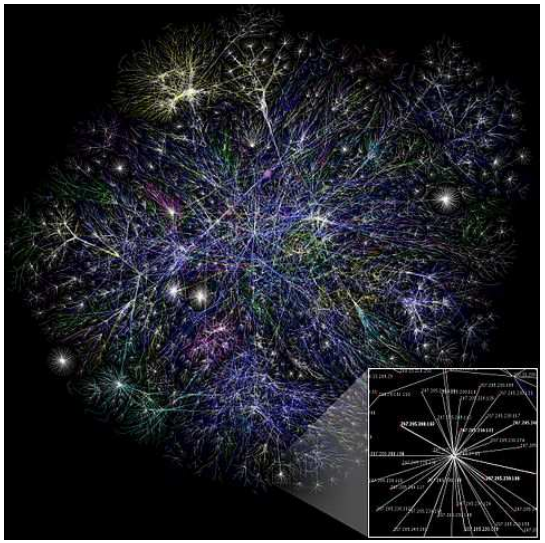

Source: Wikimedia

- Random graphs were first defined in 1959 by Paul Erdös and Alfréd Rényi and independently by Edgar Nelson Gilbert
- Examples: The Internet, IoT, WWW, social networks, paper citations, research collaborations
- These graphs are often dynamic
- Classical graph theory deals with supplying exact answers to questions on specific graphs
- Need to model random graphs and study their statistical properties

Interaction with Other Disciplines

Source: Wikimedia

- All disciplines in science and engineering use computing
  - Computer-aided design (electronic design automation, mechanical design automation)
  - Computer simulation (finite element method)
  - Biological and medical data processing and imaging
  - Genetic engineering
  - Electronic music (CSIRAC)
- Computers can be used for social and educational purposes
- Challenges: Identifying the key areas, understanding our roles

Donald Ervin Knuth (1938–)

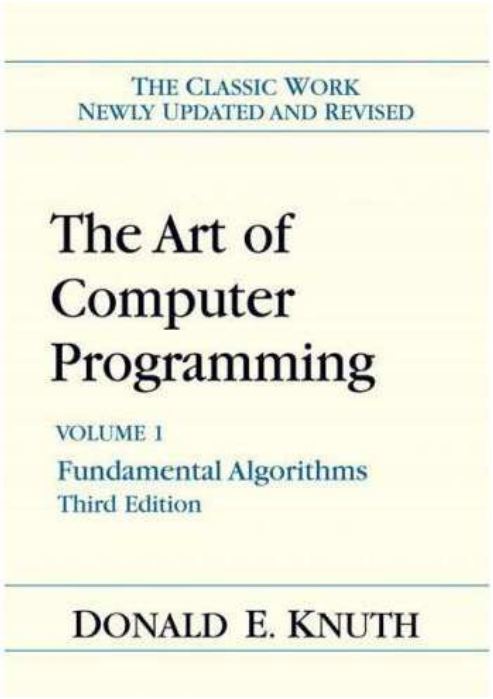

Source: Wikimedia

- CSE teaching is based on digital logic and discrete mathematics
- The phrase *Concrete Mathematics* was coined by Ronald Graham, Donald Knuth, and Oren Patashnik

Paradigm Change in CSE Teaching?



Source: Authors' official web pages

"*While traditional areas of computer science remain highly important, increasingly researchers of the future will be involved with using computers to understand and extract usable information from massive data arising in applications, not just how to make computers useful on specific well-defined problems. . . . One of the major changes is the switch from discrete mathematics to more of an emphasis on probability, statistics, and numerical methods.*"

– Avrim Blum, John Edward Hopcroft, and Ravindran Kannan
in *Foundations of Data Science*

Acknowledgment

- Most of the material has been taken (often copied verbatim) from Wikipedia
- Sources are mentioned below the pictures
- The book *Foundations of Data Science* by Blum, Hopcroft and Kannan was consulted

  https://www.cs.cornell.edu/jeh/book.pdf

- The LaTeX beamer package was used to typeset

THANK YOU