Past, Present, & Future Parallel Programming Paradigms and Numerical Algorithms

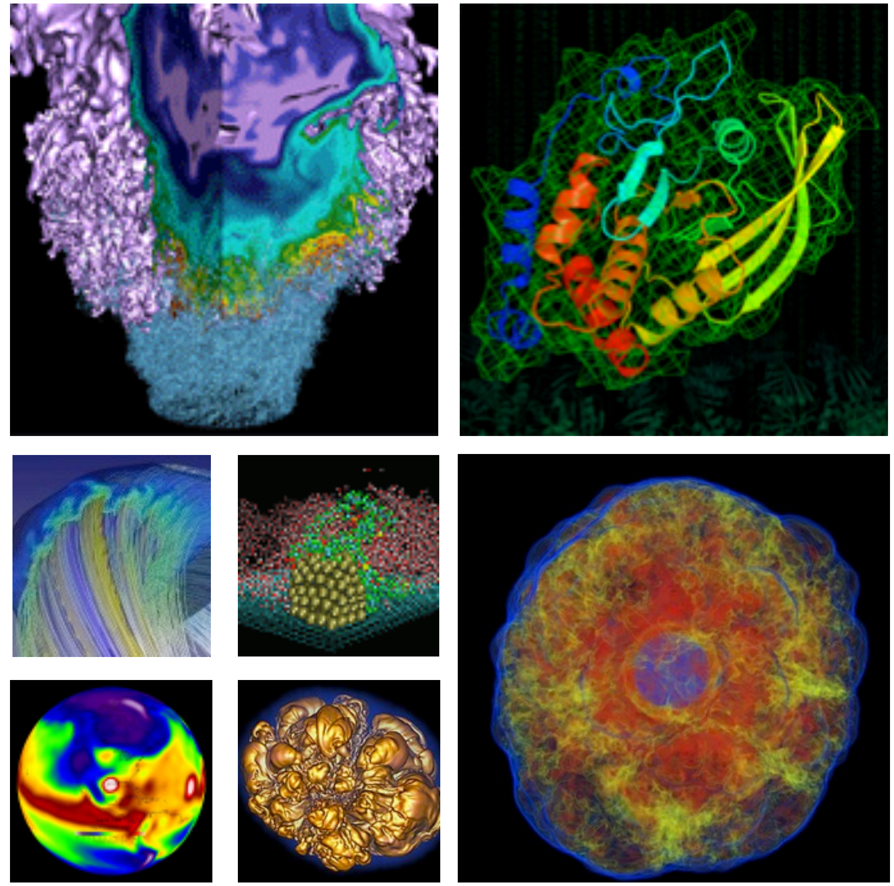

Rebecca Hartman-Baker, PhD
NERSC User Engagement Group Leader

March 1, 2017

About Me

- Group Leader, User Engagement, National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley Lab
- Worked at Pawsey Supercomputing Centre (Australia) and Oak Ridge National Laboratory
- Expertise in scalable parallel numerical algorithms
- PhD in Computer Science, University of Illinois at Urbana-Champaign
- BS in Physics, University of Kentucky
- Mom to Vincent (9) and Elena (almost 1)

Glossary

- **FLOP = Floating-Point Operation**
  – Basic arithmetic operation involving a decimal point, e.g., 1.1+1.1=2.2
- **FLOPS = Floating-Point Operations per second**

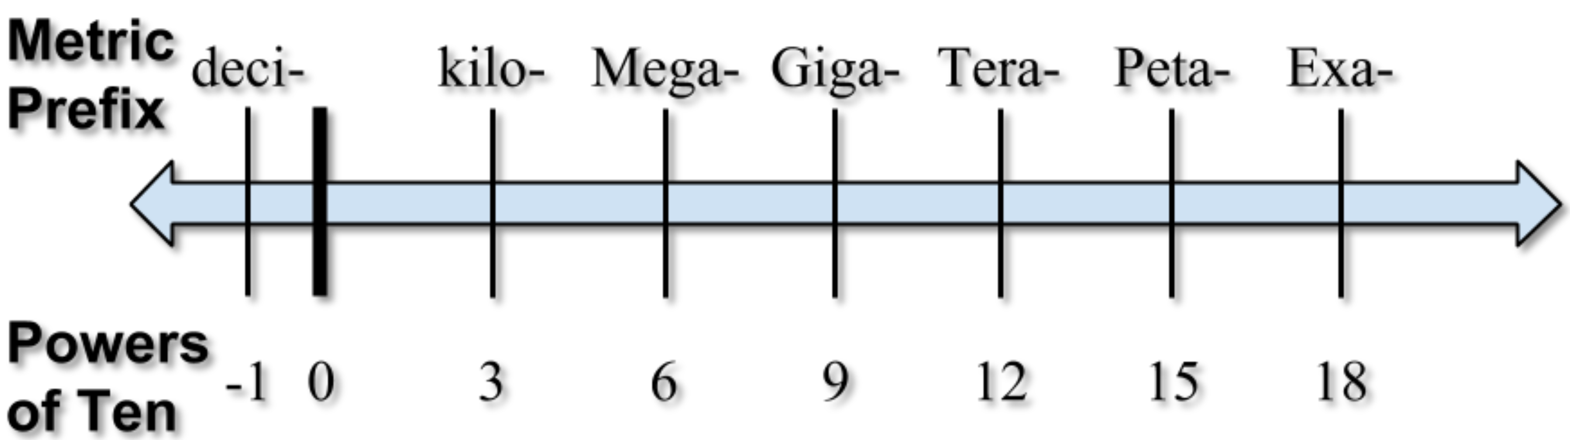

Computing History Timeline

Subflops Era

2500 BCE - 1640 AD

- Bronze Age
- Iron Age
- Post-classical

Deciflops Era 1640-1850

Flops Era 1850-1940

Digital Era 1940-1980

Parallel Era 1980-2007

Petascale Era 2008-present

The Subflops Era: 2500 BCE-1640 AD

NERSC

Subflops Era

2500 BCE - 1640 AD

- Bronze Age
- Iron Age
- Post-classical

Deciflops Era 1640-1850

Flops Era 1850-1940

Digital Era 1940-1980

Parallel Era 1980-2007

Petascale Era 2008-present

Bronze Age Computing

- **Computations preserved in tables**
  - Even basic arithmetic non-trivial in early eras
  - Cost of computation so dear, reuse it
- **Applications**
  - Commerce
  - Science

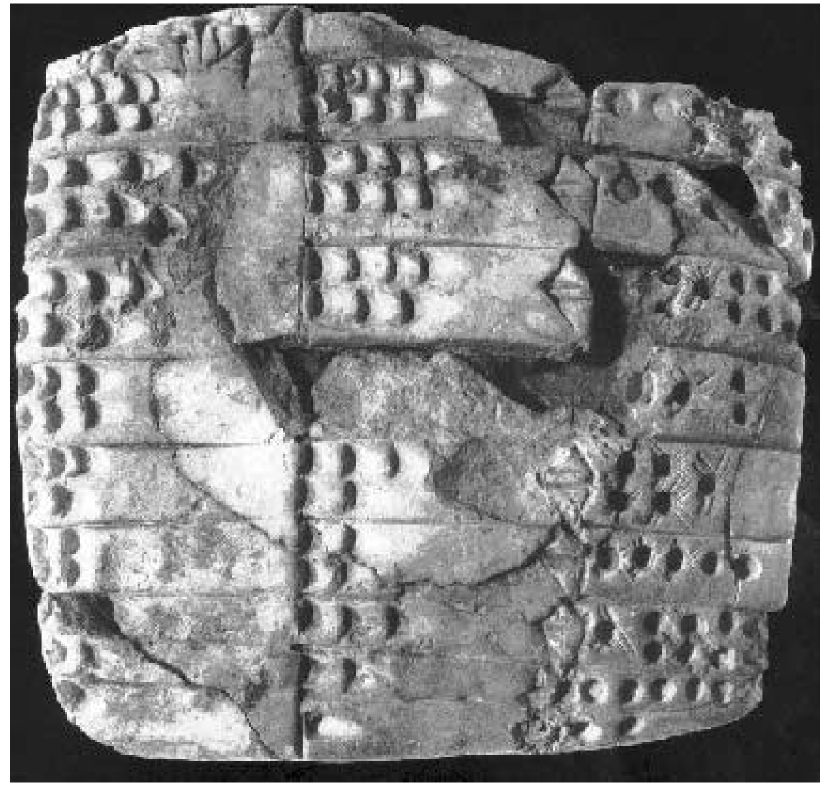

World's oldest known mathematical table, Sumerian, c. 2600 BCE, used to compute area of fields (3rd column) from length and width (1st & 2nd columns).

Tables in Ancient Sumer, Babylonia, & Assyria

- **Area of fields**
- **Accounting**
- **Metrological (unit conversions)**
- **Arithmetical (reciprocals and coefficients)**
- **Pythagorean triples**
- **Squares and square roots**

- **Astrological predictions**
  - Lunar, solar, star/planet omens
- **"Ephemerides" for predicting astronomical occurrences**
  - Also in other cultures:
    - Panchanga tables developed in India, 2nd millennium BCE
    - Chinese tables known from 1408 AD

Iron Age Computing

- **Hipparchus (c. 190-c. 120 BCE)**
  - Quantitative and accurate models for motion of sun and moon and solar eclipses
  - Trigonometric tables
- **Ptolemy (90-168 AD)**
  - *Almagest* and *Handy Tables*
- **Muḥammad ibn Mūsā al-Khwārizmī (c. 780-c. 850 AD)**
  - Astronomical tables based on Indian astronomical methods
  - Introduced decimal positional number system to Western world
  - First systematic solution of linear and quadratic equations
  - Algebra

Medieval & Renaissance Europe

- **Multiplication and Division**
  - Prosthaphaeresis – approximating multiplication and division using trigonometric formulas, c. 1500-1620s

$$\cos\alpha\cos\beta = \tfrac{1}{2}(\cos(\alpha+\beta)+\cos(\alpha-\beta))$$

  - Logarithms – developed by John Napier, 1614
    - Computed, over course of 20 years, results that allowed him to find *L*, for 5≤*N*≤ 10 million

$$N = 10^7(1-10^{-7})^L$$

    - First table of logarithms compiled by Henry Briggs in 1617 (8 digits of precision)
    - Gunter's rule (by 1620) then slide rule (1620s) developed
    - Critical to advances in science, especially astronomy; surveying, celestial navigation

Computational Methods

- **Unknown exactly how original tables were computed**
  - Probably collective activity for first tables
  - Tables often copied for memorization
- **Babylonian abacus developed c. 2400 BCE; Chinese abacus developed c. 200 BCE; ancient Mayans developed abacus called nepohualtzintzin**
- **Ephemerides produced collaboratively:**
  - Constructed from best available existing tables, by teams of computers (either astronomy researchers or hired computers)
  - Often international collaborations, particularly in later eras

Technological Innovations

- **Precursors to computers:**
  - First fully mechanical clock built by Chinese inventor Liang Lingzan, 724 AD
  - Hydropowered organ playing cylinders with raised pins on surface, & programmable automatic flute player invented by Banu Musa brothers, c. 650 AD
- **Computing devices:**
  - Astrolabe perfected by Ibn Samh, c. 1020 AD
  - Torquetum, transforming between spherical coordinate systems, developed by Jabir ibn Aflah, c. 1100 AD
  - Castle Clock, astronomical clock considered earliest programmable analog computer, by Arab engineer Al-Jazari, 1206 AD
  - Plate of Conjunctions, analog computer used to predict time of planetary conjunctions and to perform linear interpolation, Jamshid al-Kashi, c. 1400 AD

The Deciflops Era: 1640-1850

NERSC

Subflops Era

2500 BCE - 1640 AD

- Bronze Age
- Iron Age
- Post-classical

Deciflops Era 1640-1850

Flops Era 1850-1940

Digital Era 1940-1980

Parallel Era 1980-2007

Petascale Era 2008-present

Tables

- **Tables of higher math functions developed and distributed, beginning in early 1700s**
- **Nautical Almanac: from 1767**
  - Employed "freelance computers" coordinated through postal service

By Matthias Bernegger - Manuale Mathematicum, 1619.

Computing Machines (1)

- **Machine arithmétique (mechanical calculator), first machine with controlled carry mechanism, introduced by Blaise Pascal, 1642**
  - Predated by Schickard Calculator, which lacked carry mechanism
- **Non-decimal adding machine for use with English money, developed by Samuel Morland, 1668**

Computing Machines (2)

- **Portable calculator with all 4 math operations developed by Philipp Matthaus Hahn, 1774**
- **Automatic loom controlled by punch cards, developed by Joseph-Marie Jacquard, 1801**

Computing Machines (3)

- **Arithmometer, first mass-produced mechanical calculator, refined from Pascal's and Leibnitz's work, Charles Xavier Thomas de Colmar, beginning 1820, mass production starting in 1851**

Bureau du Cadastre Tables Project

- **In 1793, French Bureau du Cadastre charged with producing logarithmic & trigonometric tables to 14 and 29 decimal places, from scratch**
- **Project headed by Gaspard Riche de Prony, inspired by works of Adam Smith to produce tables by industrial process, division of labor akin to pin factory from Smith's *Wealth of Nations* (1759)**
- **Division of labor:**
  - Small advisory group of eminent analysts
  - 6-8 professional mathematicians
  - 60-90 computers, relatively unskilled workers performing basic arithmetic
- **Organized computing bureaus became new standard for large table-making projects, through 1940s**

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Charles Babbage (1791-1871)

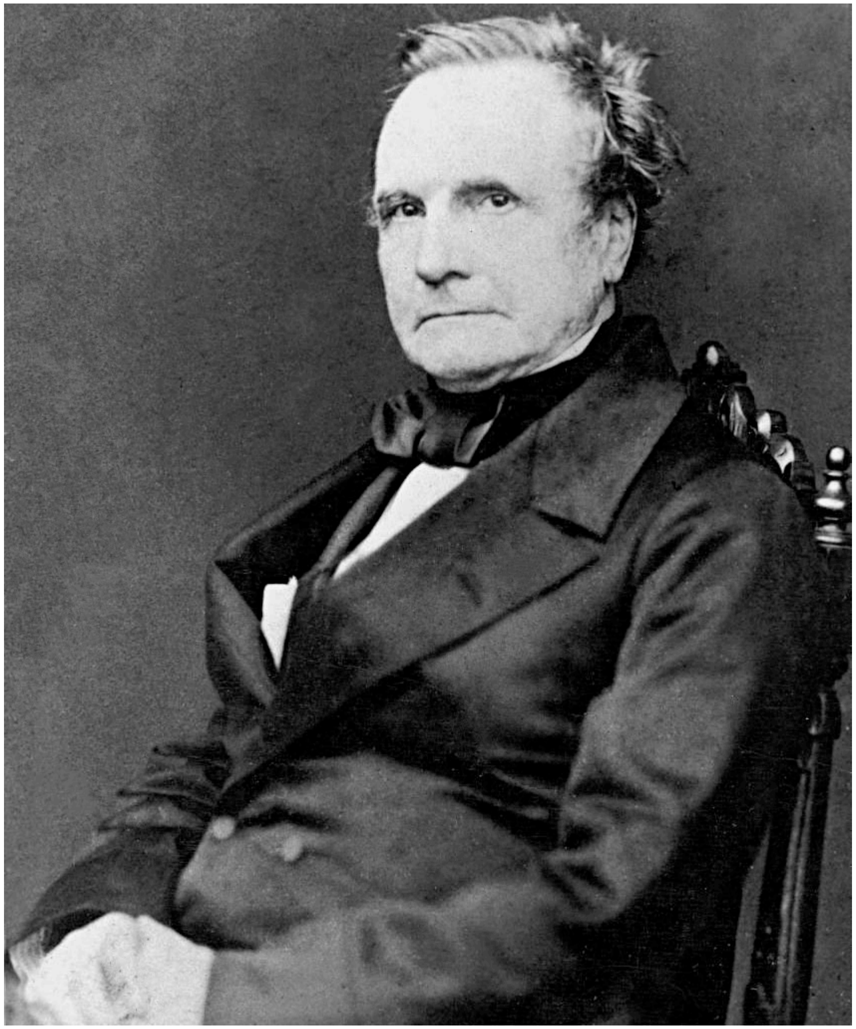

- **Originated concept of programmable computer**
- **Motivation: error-prone human table-making (& typesetting) process**
- **Difference Engine:**
  - Designed to compute values of polynomial functions via finite differences
  - Never completed (until 1991)

Ada Lovelace (1815-1852)

- **Built upon Babbage's 2nd computer design, the Analytical Engine**
  - Never completed
  - Turing-complete
  - Programmable with punch cards
  - She developed program to compute sequence of Bernoulli numbers

Numerical Methods

- **Golden age of great mathematicians: Gauss, Fourier, Jacobi, etc.**
- **Focus on science & engineering applications**
- **Did not know that their work would be instrumental 100+ years later in digital computing**
- **This presentation too brief to discuss all numerical algorithms, so focus on only a few**

Iterative Methods

- **Iterative Method: mathematical procedure that generates sequence of improving approximations to solution of a problem**
- **Example: Stationary iterative method for system of linear equations (aka relaxation method)**
  - Iterative method: solves system with an operator approximating original
  - Stationary: operator stays the same

Relaxation Method (1)

- **Solve a linear system of equations $Ax=b$ with the form**

$$x^{(k+1)} = Gx^{(k)} + c$$

Where G and c are chosen so that a fixed point of $x = Gx + c$ is a solution to $Ax=b$

Choose G by performing splitting of A, $A = M-N$, then take $G=M^{-1}N$ and $c= M^{-1}b$

Converges if

$$\rho(G) = \rho(M^{-1}N) < 1$$

(where ρ = spectral radius = max eigenvalue in absolute value)

Relaxation Method (2)

- **How to pick M,N? (For $A = L+D+U$, lower triangle, diagonal, upper triangle, respectively)**
  - Jacobi Method: $M=D, \quad N=-(L+U)$
    - Requires double storage for x because all values from step k needed at step $k+1$
    - Convergence guaranteed if matrix diagonally dominant by rows, generally slow
    - Developed by Jacobi in 1846
  - Gauss-Seidel Method: $M=D+L, \quad N=-U$
    - Duplicate storage for x not needed
    - Sequential dependence of updating unknowns
    - Convergence faster than Jacobi
    - Originated in 1823 letter from Gauss to Gerling; further refined and published in 1874 by Seidel

The Flops Era: 1850-1940

Subflops Era

2500 BCE - 1640 AD

- Bronze Age
- Iron Age
- Post-classical

Deciflops Era 1640-1850

Flops Era 1850-1940

Digital Era 1940-1980

Parallel Era 1980-2007

Petascale Era 2008-present

Mechanical Calculators

- **Tabulating machine: 1st full-scale difference engine, 15-digit numbers & 4th-order differences, completed by Swedish father & son Per Georg & Edvard Scheutz, 1853**
- **Comptometer developed by Dorr Felt of Chicago, 1884**
  - Enter operands by pressing keys (vs. dialing)
  - Manufacture begins 1887
  - Felt also invents 1st printing desk calculator, 1886

Herman Hollerith (1860-1929)

- **Herman Hollerith, US Census Department employee, develops means to process census data via punch cards, 1890**
  – Data processing finished in 4 years & far under budget
  – Goes on to found Tabulating Machine Company (later IBM)

Hollerith 1890 tabulating machine & sorting box.
By Adam Schuster - Flickr: Proto IBM, CC BY 2.0, https://commons.wikimedia.org/w/index.php?curid=13310425

IBM

- **Tabulator developed for census tabulations, 1890**
  - Specific-purpose machine, not profitable
- **Tabulator with plugboard to adapt for different applications, 1906**
- **Standardizes punch cards with 80 columns & rectangular holes, 1928 (standard lasting into 1970s)**
- **IBM 601 Multiplying Punch: electomechanical machine that read 2 numbers up to 8 digits, multiplied them, & punched product back onto card, 1931**

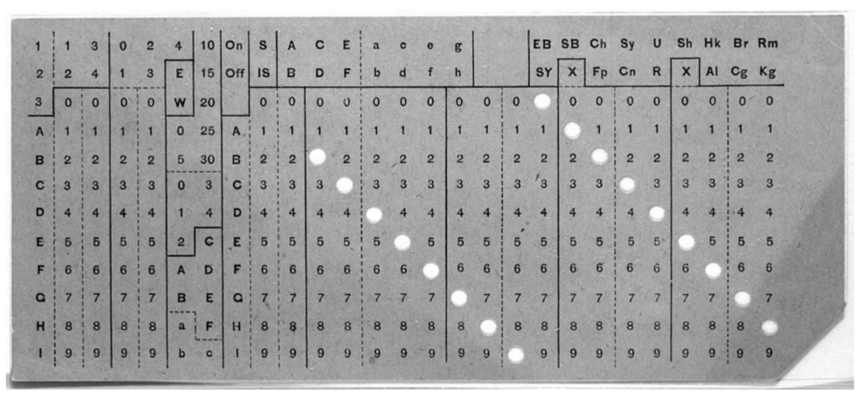

Depression-Era Computing Projects

- **WPA Mathematical Tables Project (1930s & 40s)**
  - Employed 450 previously unemployed people, most of whom had not completed high school
  - Workers assigned to either addition, subtraction, multiplication, or (if skilled enough) division
  - Produced tables of powers, trig functions, probability functions
  - Contributed to the Handbook of Mathematical Functions
- **Army ballistics tables (wartime, 1940s)**
  - Computed by 100-200 female computers at University of Pennsylvania
  - Each woman equipped with a desk calculating machine

The Digital Era: 1940-1980

Subflops Era

2500 BCE - 1640 AD

- Bronze Age
- Iron Age
- Post-classical

Deciflops Era 1640-1850

Flops Era 1850-1940

Digital Era 1940-1980

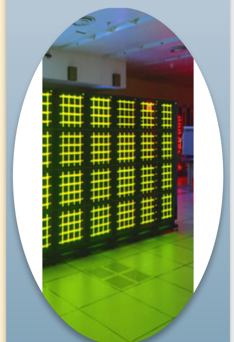

Parallel Era 1980-2007

Petascale Era 2008-present

Dawn of Modern Computers (1940-1945)

- **Computers developed for computing tables (paradoxically, obsoleting tables!)**
- **First digital computers used vacuum tubes**
  - Neon lamps as prototype memory
- **Colossus, 1943**
  - To crack German cypher
  - 2400 vacuum tubes, read 5000 chars/s from punched tape
  - 10 built, most destroyed after war

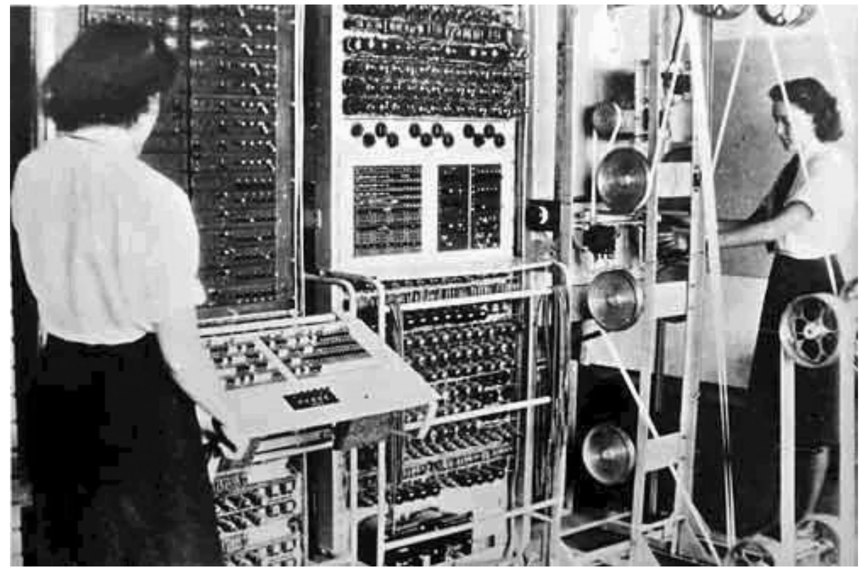

Colossus codebreaking computer in action.
By Unknown - This file is from the collections of The National Archives (United Kingdom), catalogued under document record FO850/234.

Early Computers (1945-1950)

- **ENIAC: Electrical Numerical Integrator and Computer, 1946**
  - Totally electronic, valve driven, digital, programmable
  - 5000 simple adds/subtracts per second
  - 100’ x 8’ x 3’, 60000 lbs, 150 kW
  - Turing-complete
  - Original purpose: ballistics tables, but with involvement of Los Alamos, first computations were for hydrogen bomb

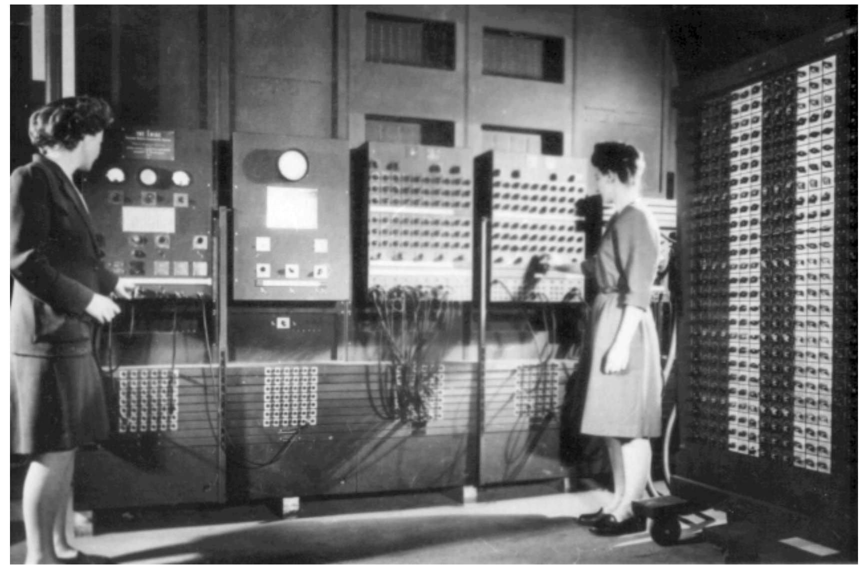

Two women operating the ENIAC's main control panel while the machine was still located at the Moore School. U.S. Army Photo from the archives of the ARL Technical Library. Left: Betty Jennings (Mrs. Bartik) Right: Frances Bilas (Mrs. Spence)

Women in Computing

- **Women played prominent role in computing**
- **Ada Lovelace developed first computer program**
- **Emma Gifford compiled 500-page book, *Natural Sines,* 1914, computing sines and cosines to unprecedented precision**
- **Women were active in early modern computing**
  - Six women did majority of ENIAC programming: Kay McNulty, Betty Jennings, Betty Snyder, Marlyn Wescoff, Fran Bilas, & Ruth Lichterman
  - Grace Hopper, remarkable, accomplished computer scientist: developed first compiler, idea of machine-independent programming languages, etc.

Numerical Renaissance

- **Rebirth of interest in numerical solution of equations**
- **(A few) Great numerical analysts of the era:**
  - Hestenes
  - Krylov
  - Lanczos
  - Stiefel
  - Young

Relaxation Methods (Review)

- **Solve a linear system of equations $Ax=b$ with the form**

$$x^{(k+1)} = Gx^{(k)} + c$$

Where G and c are chosen so that a fixed point of $x = Gx + c$ is a solution to $Ax=b$

Choose G by performing splitting of A, $A = M\text{-}N$, then take $G=M^{-1}N$ and $c= M^{-1}b$

Converges if

$$\rho(G) = \rho(M^{-1}N) < 1$$

(where ρ = spectral radius = max eigenvalue in absolute value)

Relaxation Methods

- **Successive Over-Relaxation (SOR)**
  - Young, 1950

$$M = \tfrac{1}{\omega} D + L, \quad N = (\tfrac{1}{\omega} - 1) D - U$$

  - ω: $1 < \omega < 2$ is relaxation factor (precise value depends on properties of matrix)
  - Repeated sweeps through unknowns
  - Requires only one storage vector (like Gauss-Seidel)
  - Can be order of magnitude faster than Gauss-Seidel

Birth of the Modern Computer

- **EDSAC (Electronic Delay Storage Automatic Calculator), 1949**
- **Constructed by Maurice Wilkes et al. at Cambridge**
  - Led to first PhD in Computer Science: David Wheeler, who invented concept of subroutine
- **4m x 5m x 2m, 500 kHz, 12 kW power**
- **Considered birth of modern computing: first complete, fully functional von Neumann architecture computer**

Innovations (1950-1965)

- Whirlwind computer, first to allow interactive computing, through keyboard and CRT, 1951
- A-0 high-level compiler invented by Grace Hopper, 1951
- First transistor-based computer, 1953
- FORTRAN developed, 1954-1957
- Integrated circuit invented by Jack Kilby at Texas Instruments, 1958
- ATLAS, University of Manchester, introduced many modern architectural concepts: spooling, interrupts, pipelining, interleaved memory, virtual memory, paging, 1962
- Moore's law (doubling of chip performance every 18 months), 1965

The Dawn of Supercomputers

- **Control Data CDC 6600, 1965**
- **Architecture designed by Seymour Cray**
  - CPU handled only arithmetic & logic, with additional 10 parallel functional units
  - 10 peripheral processors did the rest
- **O(1) Megaflops**
- **Plus-shaped design; size of 4 filing cabinets**
- **$8 million**

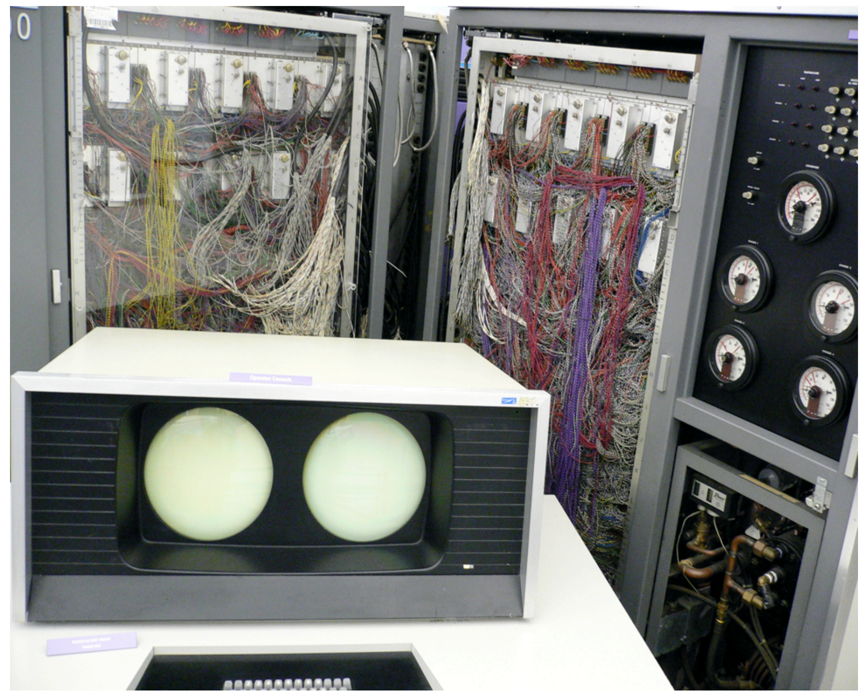

Dawn of Vector Supercomputers (1970s)

- Cray 1, 1976
- First to use integrated circuits (IC)
- Shape to assure proper timing of speed-dependent IC modules
- Freon-cooled
- 136 Megaflops
- 85 units sold 1976-1982, $5-8 million each

Vector Supercomputing

- **Vector processing: performing the same operation on an array of numbers (e.g., multiply, add)**
- **Relaxation methods ideal for vector processing**
  - Gauss-Seidel and SOR preferred due to storage constraints and superior convergence

The Parallel Era: 1980-2007

Subflops Era

2500 BCE - 1640 AD

- Bronze Age
- Iron Age
- Post-classical

Deciflops Era 1640-1850

Flops Era 1850-1940

Digital Era 1940-1980

Parallel Era 1980-2007

Petascale Era 2008-present

The Rise of the Microprocessor (1980s)

- Thinking Machines CM-2, 1987
- 1.5 m cube, divided into 8 cubes
- Each sub-cube contained 16 printed circuit boards & sequencer (main processor)
- Up to 65,536 microprocessor cores, each with 4 kbits RAM
- Up to 25 GB hard drive
- Known for visually striking design

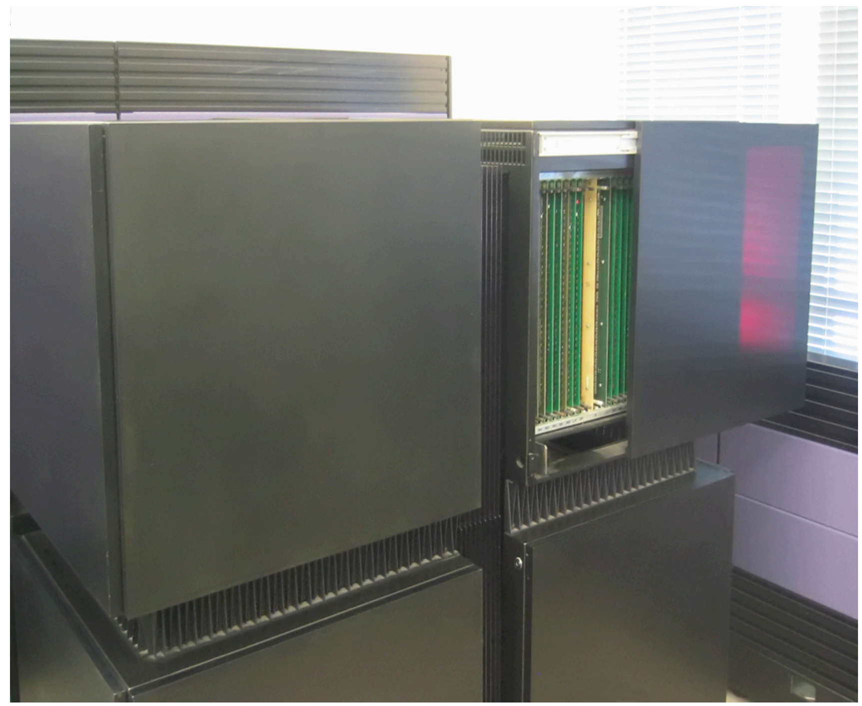

The Rise of the Clock Rates (1990s)

- **Intel Paragon series, 1992**
- **Up to 1024 nodes, 2 or 3 cores/node**
  - 75 MHz i860s
  - 1 core reserved for communication handling
- **150 Tflops (ORNL's paragon #1 on TOP500, 1995)**

NERSC

The Rise of Massive Parallelism (2000s)

- **IBM BlueGene/L, 2004**
- **Massive parallelism, low power consumption**
- **Large numbers of low-power 700 MHz PowerPC 440**
  - Node theoretical peak 5.6 GF
- **LLNL: 16 racks, 1024 compute nodes, 70.72 Tflops, #1 on TOP500, November 2004**
  - Expanded machine remained #1 until June 2008

Computing for Parallelism

- Distributed-memory architecture, in which information needing to be shared must be passed across slow interconnect
- MPI standard (communication) developed in mid-1990s
- Degree of parallelism rises as machines grow in size and processor count
- Relaxation methods: Gauss-Seidel/SOR no longer viable – communication cost outweighs inferior performance of Jacobi method

Beyond Stationary Iterative Methods

- **Stationary Iterative Methods generally superseded by Krylov subspace methods**
- **Conjugate Gradient Method (SPD matrices)**
  – Originally envisioned as direct method by Hestenes and Stiefel in 1952, popularized as iterative method by Reid, Golub, et al. in early 1970s
  – Minimizes residual over Krylov subspace
- **GMRES (non-SPD)**
  – Proposed in 1986 by Saad & Schultz
  – Iterate in Krylov subspace minimizes Euclidean norm of residual
- **Preconditioning increasingly important**

The Petascale Era: 2008-present

Subflops Era

2500 BCE - 1640 AD

- Bronze Age
- Iron Age
- Post-classical

Deciflops Era 1640-1850

Flops Era 1850-1940

Digital Era 1940-1980

Parallel Era 1980-2007

Petascale Era 2008-present

The Rise of Hybrid Supercomputers (1)

- **IBM Roadrunner, 2008**
- **First petaflop supercomputer (1.7 PF peak, 1.026 PF sustained)**
- **Hybrid dual-core Opteron + Cell accelerator**
  - 6480 Opteron processors
  - 12960 Cell processors
- **2.35 MW power**

The Rise of Hybrid Supercomputers (2)

- **Architecture gained popularity due to low power envelope**
  – Current #1 & #2 on Top500 list are hybrid supercomputers
- **Node consists of CPU (brain) + accelerator (brawn)**
- **Requires offloading programming model, in which FLOP-intensive computations fed to accelerator**
- **Exposing kernels for offloading often leads to better CPU-only performance**
- **Communication even more constrained; development of asynchronous/chaotic relaxation methods**

The Rise of Multicore Supercomputers

- **First Petaflop multicore supercomputer: Jaguar Cray XT5, 2009**
  - 200 cabinets
  - 18,688 compute nodes, 12 cores/node
  - 1.75 PF peak
  - #1 in Top 500, 11/2009
- **Most powerful supercomputer in southern hemisphere: Magnus Cray XC40, 2014**
  - 8 cabinets
  - 1536 nodes, 24 cores/node
  - 1.097 PF peak
  - #41 in Top 500, 11/2014

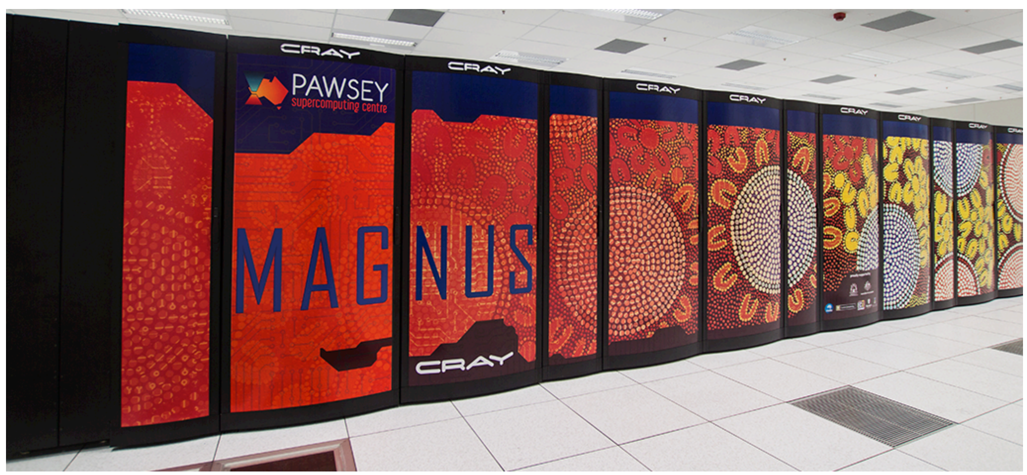

The Rise of Manycore Architectures

- **Multicore → Manycore**
  – Many (O(100)) low-power cores per node
  – Each core is low frequency, but aggregate FLOPS high
  – Requires preparation for new architecture, exposing parallelism in application and using MPI + threading
- **Arriving ~Q3 2016**

Evolution of Computer Architectures

1950s to 1960s: Architectural innovations for faster processors

1970s: Vectors to increase performance

1980s: Parallelism begins, rise of microprocessors

1990s: Era of massive parallelism and clock rate scaling

2000s: Parallelism continues, but clock rate scaling ends. Multicore processors

2010s: Multicore and accelerator based systems push parallelism from O(100K) to O(10B) way parallelism. Manycore systems appear

IBM 360/9

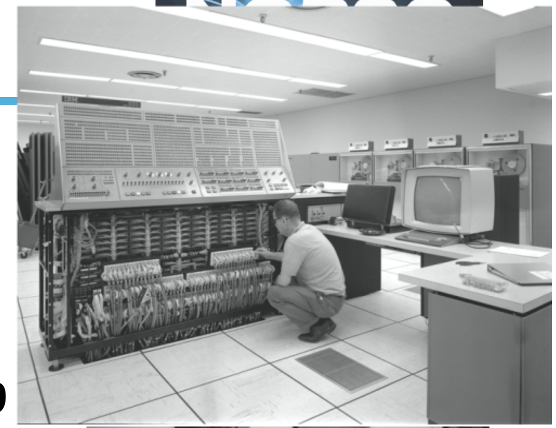

Cray X-MP

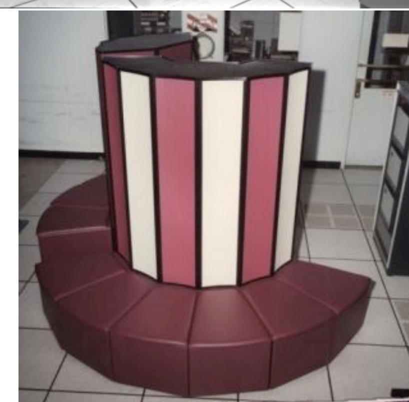

Intel Paragon

The Exascale Era: 2022-?

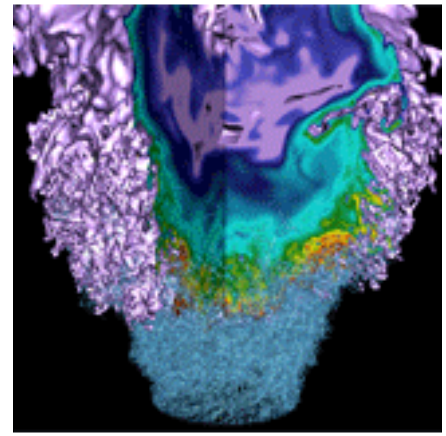

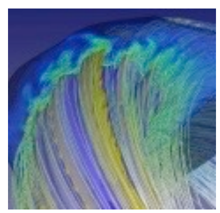

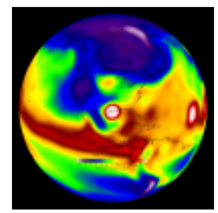

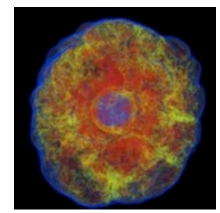

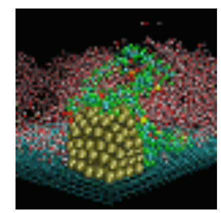

Exascale Supercomputer Architectures

- Billion-way concurrency
- 1000-10,000-way concurrency on single node
- Without new technologies, memory and interconnect bandwidth will not grow commensurately
- Dramatic overhaul of numerical methods imperative

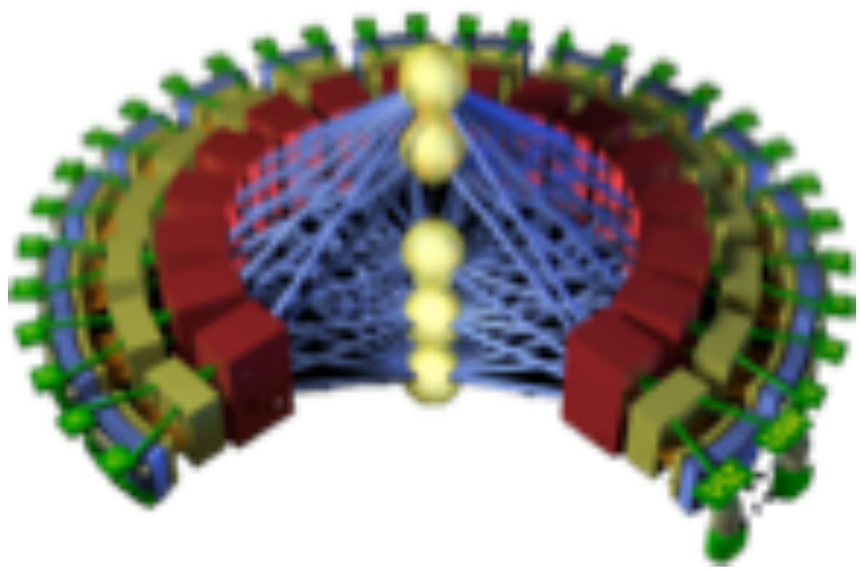

Numerical Methods in the Exascale Era

- **Flops: effectively "free"**
  - Even today's top algorithms not 100% of peak
- **Memory: precious compared to Flops**
- **Communication: avoid global syncs!**
- **What algorithms can we employ to exploit exascale architectures?**

Exascale Algorithms

- **How can we get there from here?**
  - Requires a major paradigm shift
- **Analogy: We want to get to the moon**
  - There is series of trees that we can climb to get closer to moon
  - But there is no tree tall enough to climb all the way to the moon

Another Analogy

- **Computations = music, processors = musicians**
- **Today's algorithms = compositions for marching band**
  - Band marching in lock-step, each member has predefined role
  - Band leader keeps everybody in time
  - Predefined structure and order of operations
  - Use different arrangements of score for different sizes of bands

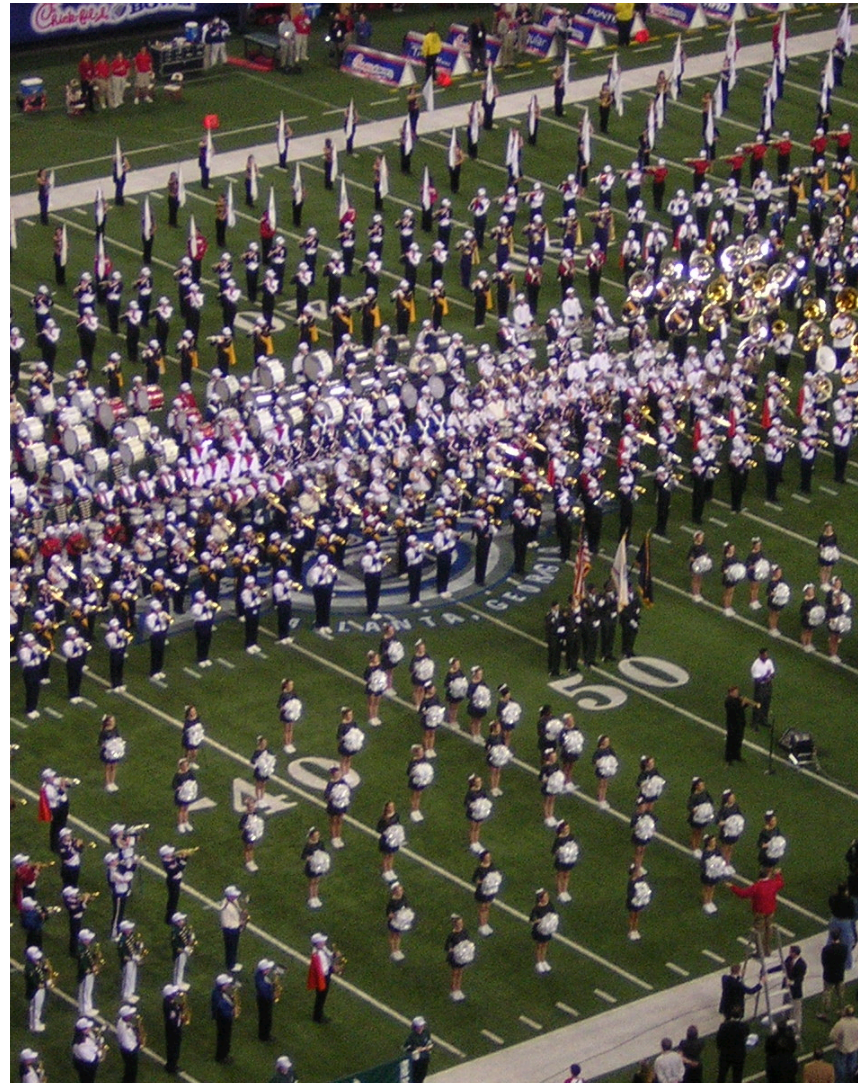

Another Analogy

- **But, marching band is not scalable**
  - Conductor can be seen by only so many musicians
  - Ultimately, can't lead a million musicians in lock-step
- **Must rethink algorithms for creating music**
- **Musical equivalent of exascale algorithms unfamiliar to Western ear**
  - Who says we need a score?
  - What if we provided certain parameters (e.g., key and starting note of scale), and let everybody improvise?
  - Perhaps some minimal synchronization between neighbors
  - A distributed, million-musician Raga?

NERSC

National Energy Research Scientific Computing Center

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

A Final Analogy

- One of my little devil's favorite books: *Hand, Hand, Fingers, Thumb* by Al Perkins
- After countless mind-numbing readings of this book, I realized it was an allegory for the future of high-performance computing!

A Final Analogy

Early Computers

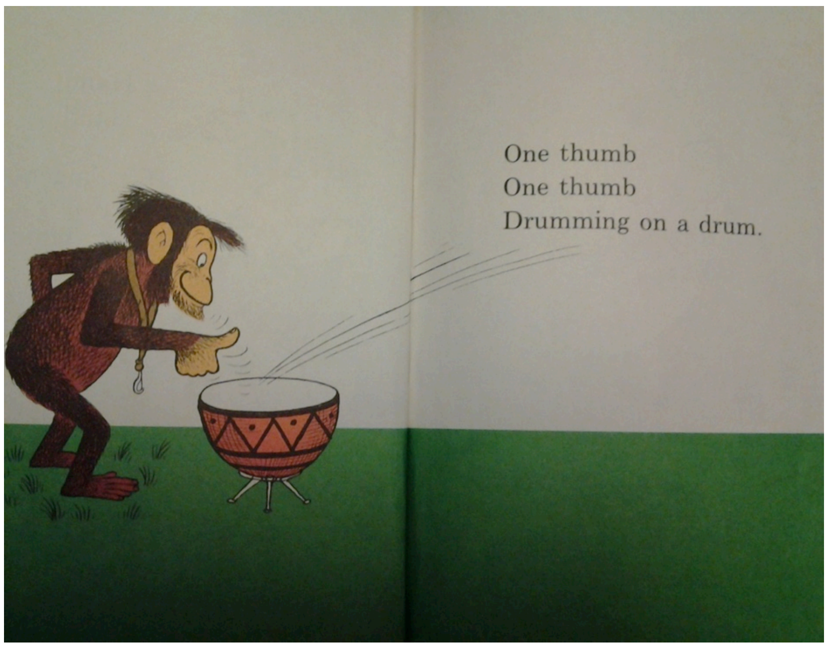

Vector Processors

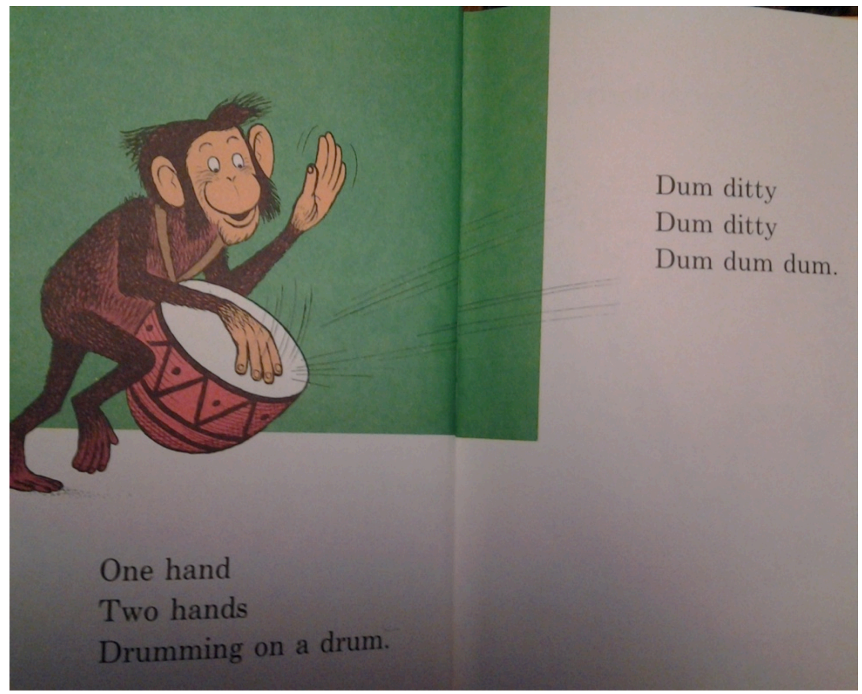

A Final Analogy

CPU + Accelerator

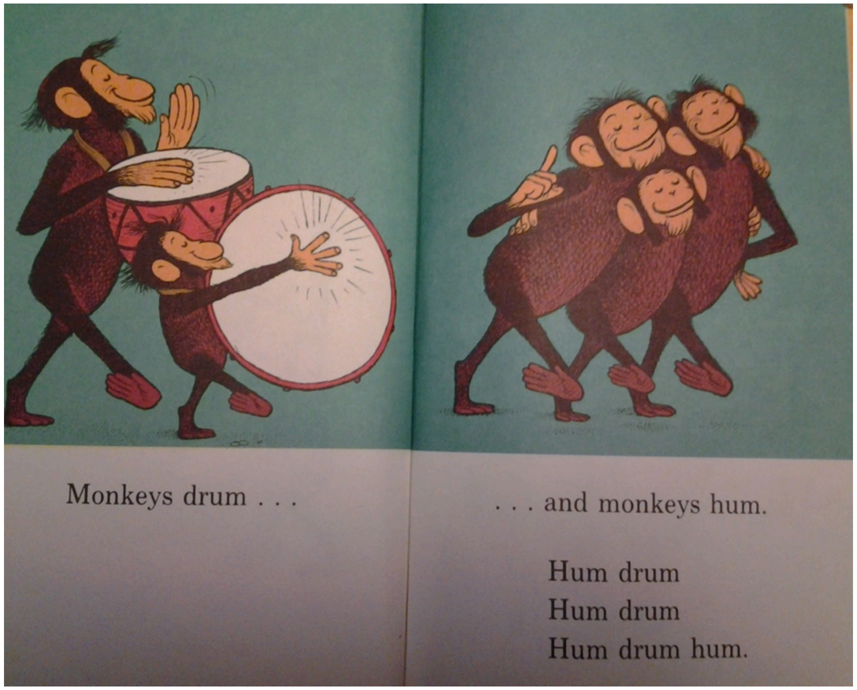

MPI Communication

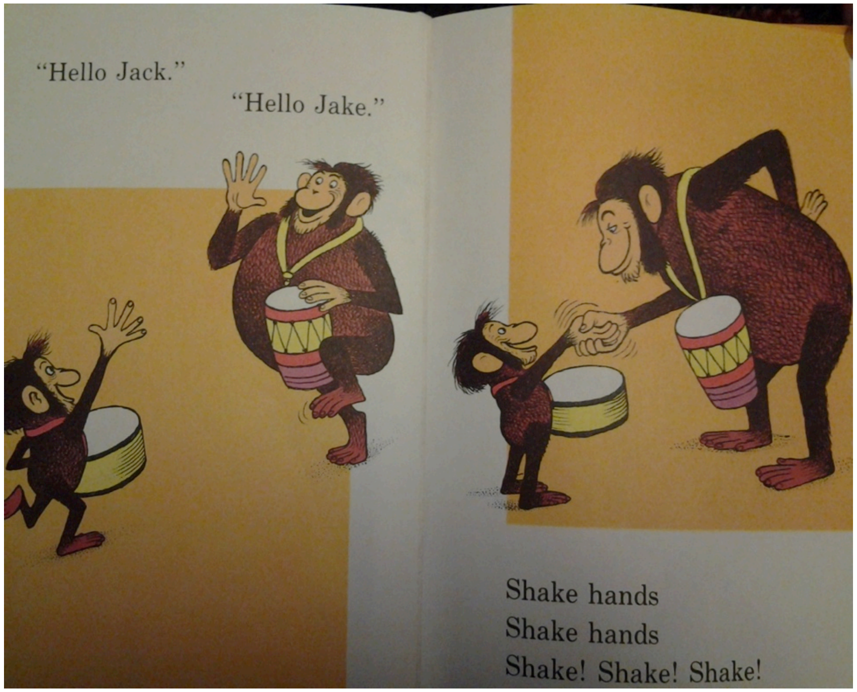

A Final Analogy

Petascale Supercomputer

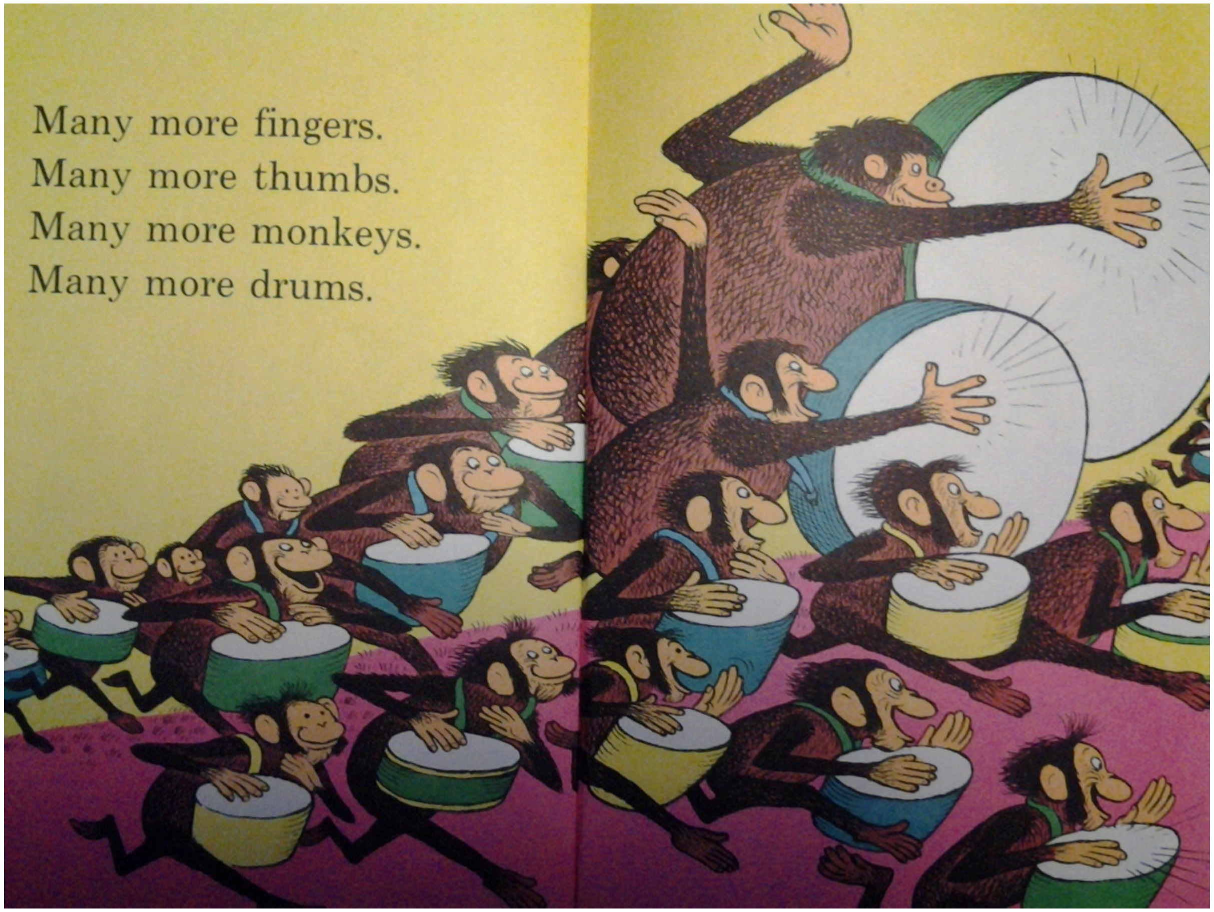

A Final Analogy

Exascale Supercomputer