

Introduction

EE599-001 & EE699-010, Spring 2026

Hank Dietz

<http://aggregate.org/hankd/>

Course Overview

- You'll be learning about quantum computing, which is **not yet a fully-developed technology**
- This is a **computer engineering** course
 - You have some exposure to digital logic
 - You have done conventional programming
 - No quantum physics understanding required

Major Topics

- Course overview and introduction. Moore's law. The relationship between parallel and quantum computing. What is a quantum computer? Why do we care? Quantum computers as attached accelerators.
- Review of classical digital logic gates. Introduction to reversible logic. Ancilla. Logic optimization with and without fanout.
- Quantum superposition, entanglement, and measurement. Notations and models including the Bloch sphere, unitary matrix formulations, and wave models. Probability amplitude. The quantum circuit model and quantum gates.
- Quantum circuits and algorithms, the use of interference. Quantum superiority and RCS (Random Circuit Sampling). Algorithms including Deutsch-Jozsa, Grover's search, Simon's problem, QFT (Quantum Fourier Transform), and Shor's factoring algorithm.
- Quantum annealing.
- NISQ (Noisy Intermediate Scale Quantum) computer implementations. Superconducting, trapped ion, photonics, spins, etc. Qubit implementations. Qubit layout and operation scheduling issues. Error correction.
- Quantum-inspired methods. Randomized and probabilistic algorithms. Alternative representations for entangled superpositions: PBP (Parallel Bit Pattern computing, developed here at UK).
- The architecture of a computer system with an attached quantum coprocessor.
- Quantum programming above the Qubit level.
- Where is quantum going? Cryptography, machine learning, networking and teleportation. Scaling.

Schedule (very tentative!)

Lectures	Topic
3	Introduction, parallel and quantum
4	Classical & adiabatic logic (project)
3	Quantum concepts and models
3	Quantum circuits and algorithms (project)
1	Quantum annealing
1	<i>Midterm exam</i>
3	NISQ computer implementations (homework)
4	Quantum-inspired methods (project)
2	Architecture of computer with quantum coprocessor
2	Higher-level quantum programming (homework)
2	Where is quantum going?
2	<i>reserved for schedule slippage</i>
1	Review for final exam

Schedule Notes

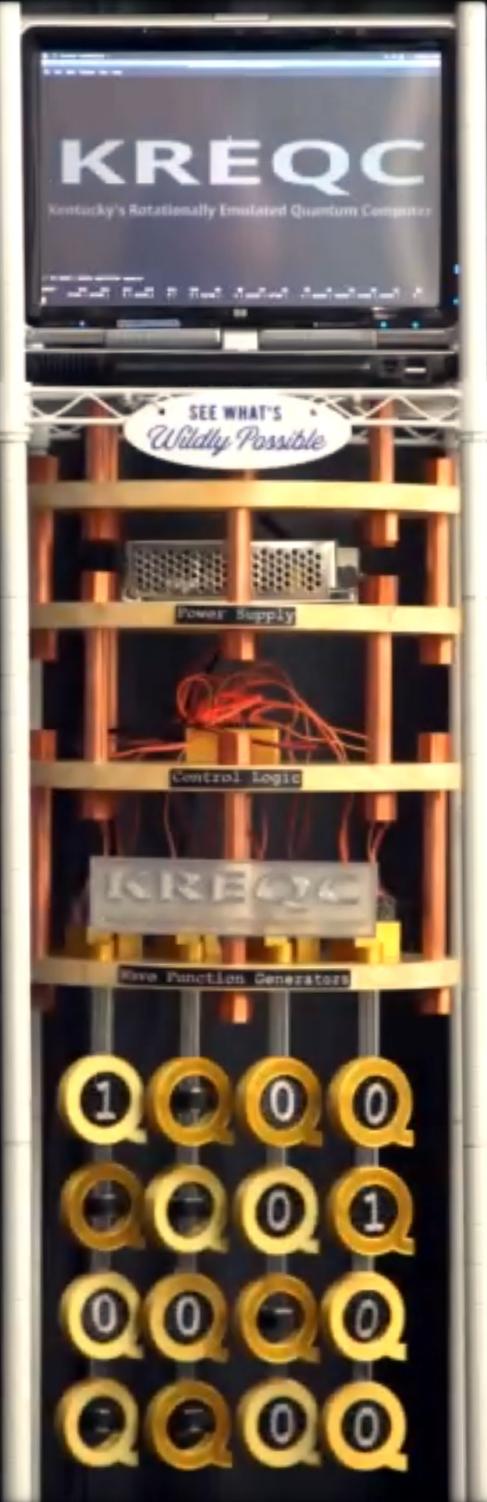
- Homeworks/projects will involve some programming using simulation environments (probably Verilog, Qiskit, and C++)
- Some topics may be given more or less time depending on how students are doing
- I will be presenting research at the IS&T Electronic Imaging conference, so we will not have regular class meetings **3/3 & 3/5**

Grading & Such

- Several **homeworks/projects**, total 50%
(a little extra on these for grad credit)
- In-person **Midterm exam**, 20%
- In-person **Final exam**, 30%
- Material from **lectures**, **canvas**, & **course URL**:
<http://aggregate.org/QC>
- You are expected to **regularly attend class**
- I try not to curve much; always in your favor

Me (and why I'm biased)

- Hank Dietz, ECE Professor and James F. Hardymon Chair in Networking
- Office: 203 Marksbury
- Research in parallel compilers & architectures:
 - Built 1st Linux PC cluster supercomputer
 - Antlr, AFNs, SWAR, FNNs, MOG, ...
 - Various awards & world records for best price/performance in supercomputing
- Lab: 108/108A Marksbury – I have **TOYS!**



Am I A Quantum Expert?

- **Nope.** This is my 1st time teaching this...
- Two main aspects to quantum computing:
 - **Quantum physics:** many A+ folks, I'm C-
 - **Computer engineering:** most D or F, I'm A+
- My background in QC:
 - Many conferences, 1st QC publication in 2017
 - Built optimizing compilers for QC
 - Designed *quantum-inspired* HW/SW
 - Committee member for several QC PhDs

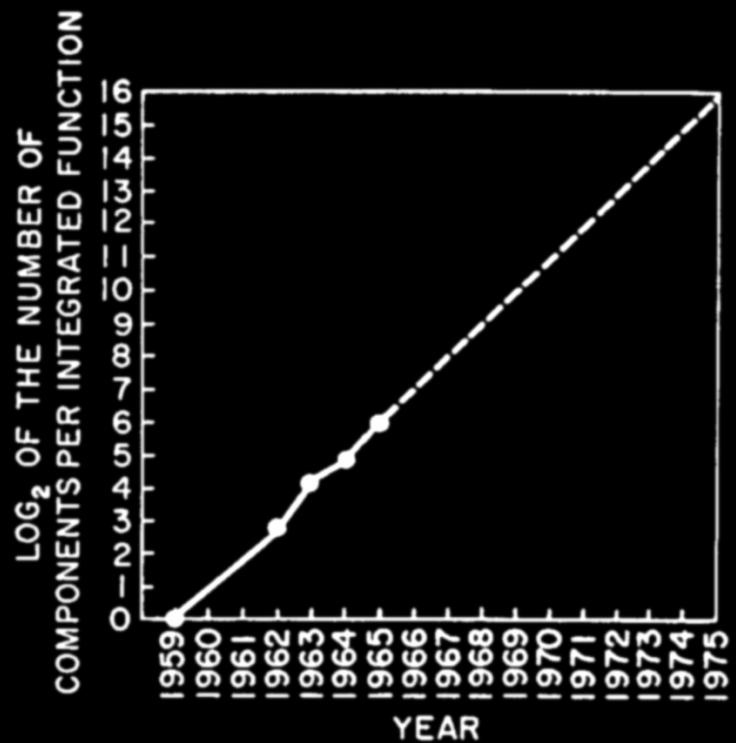
What Is This Course About?

- Using quantum phenomena in clever ways.
- Solving computational problems faster.[†]
- Using fewer resources while doing that.

[†] Especially problems that supercomputers can't feasibly solve.

How Computers Get Faster: Moore's Law

- 1965 prediction
 - Not about chip speed
 - Circuit complexity 2X every 18-24 months
- **Speedup is mostly about parallel processing**



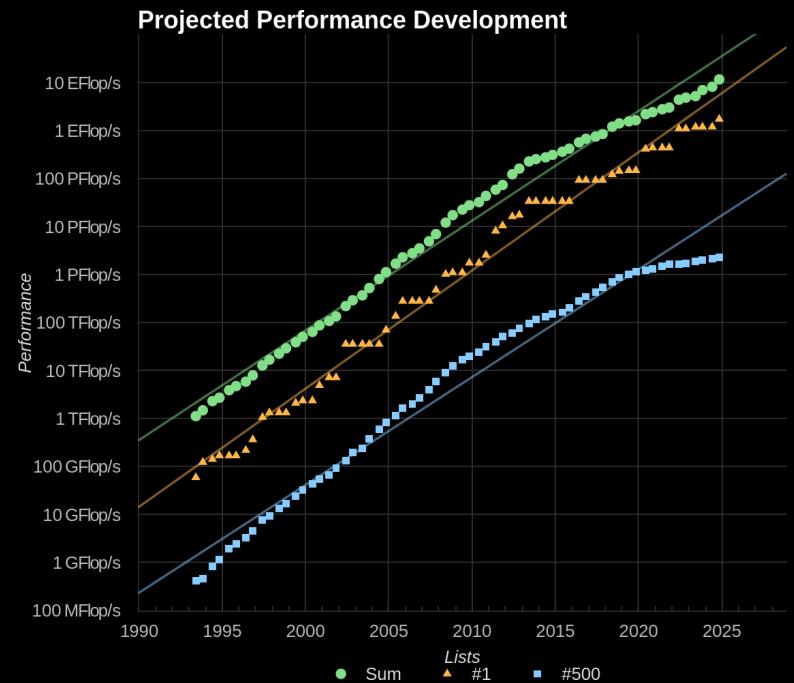
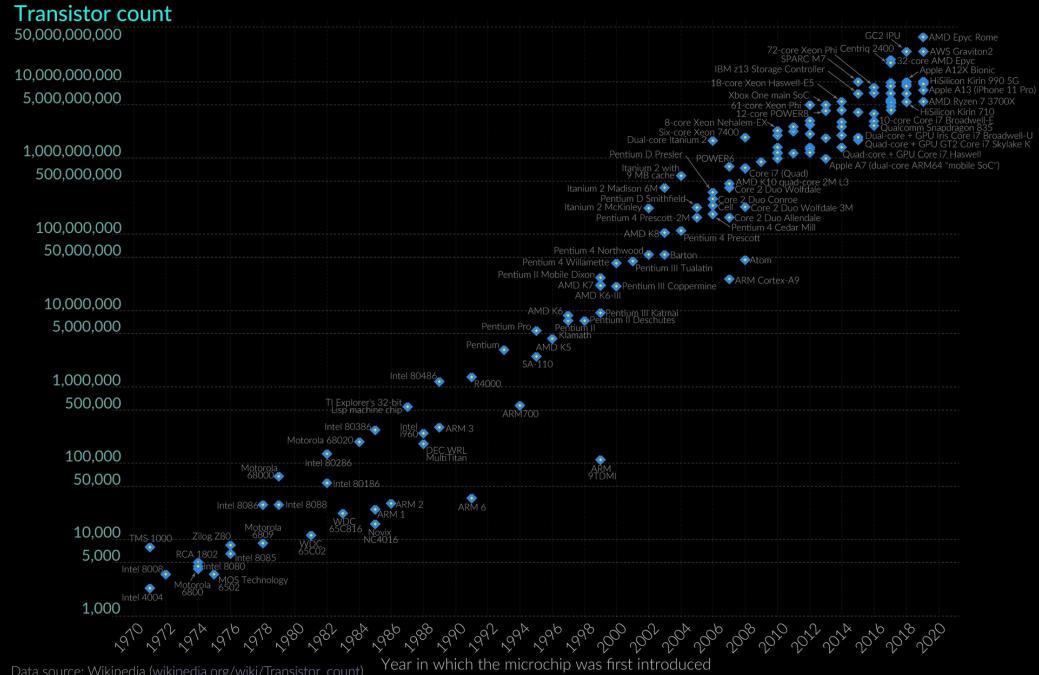
Parallel Processing

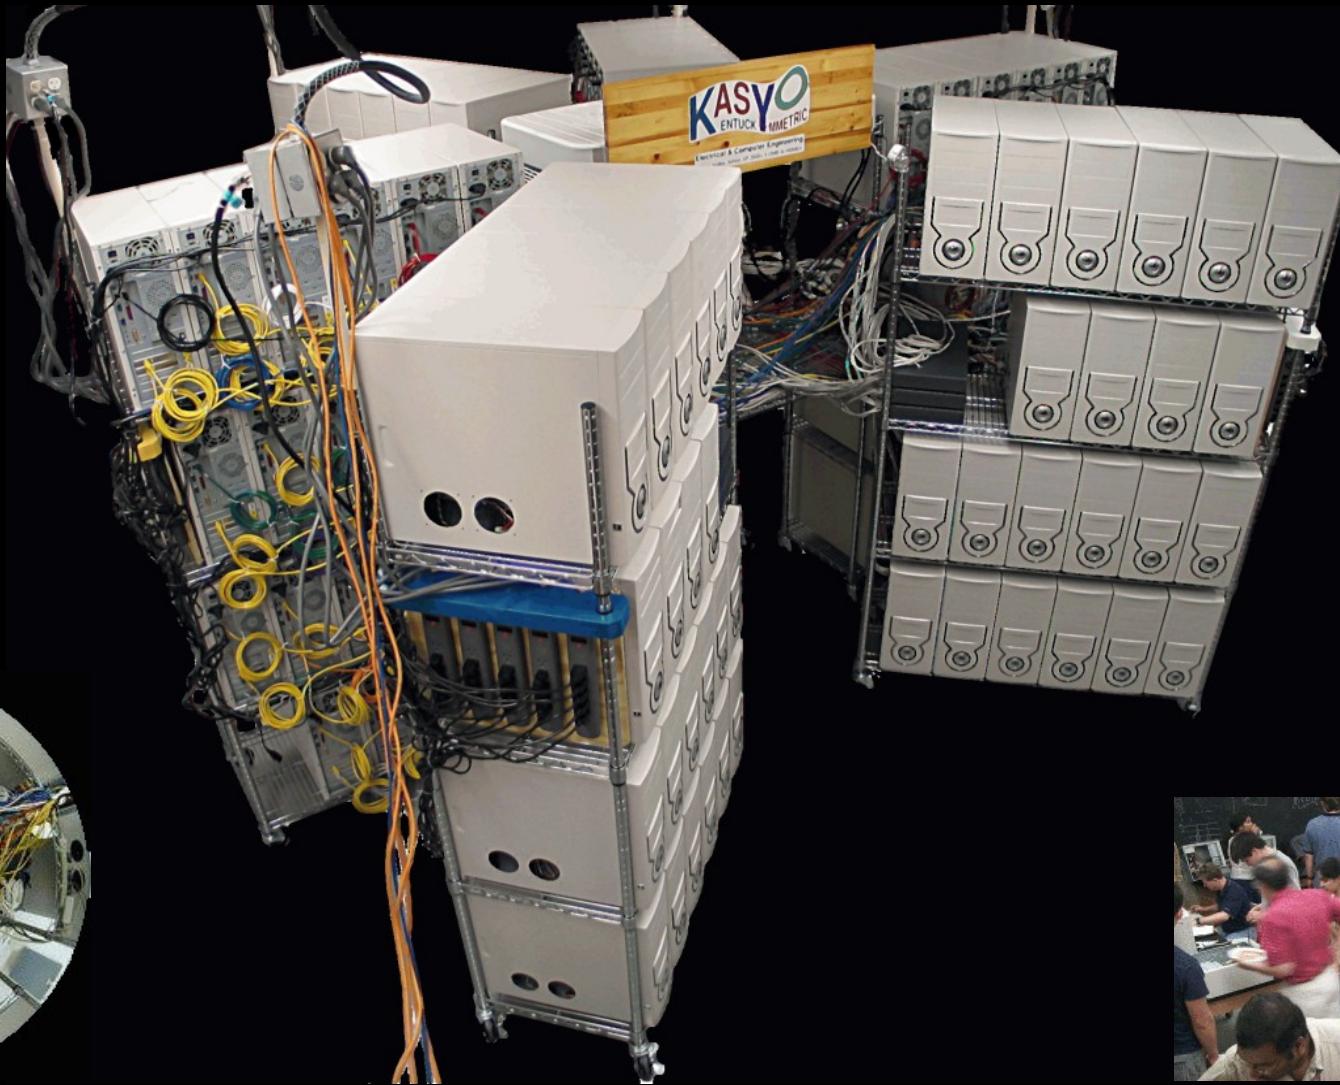
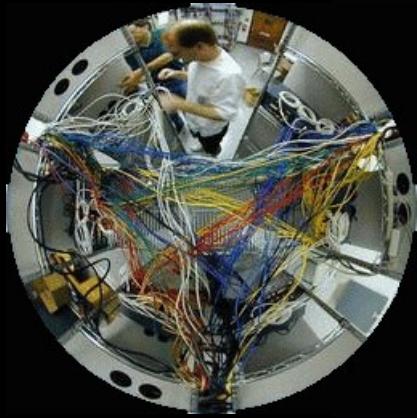
- Break program into N pieces that can execute simultaneously
 - **Scalable**: bigger N , more speedup
 - **Modular hardware**
 - Can be **fault tolerant** using redundancy
- This scales up **forever**, **right?**

Moore's Law is still sort-of OK... (using tricks like multichip modules)

Moore's Law: The number of transistors on microchips doubles every two years

Moore's law describes the empirical regularity that the number of transistors on integrated circuits doubles approximately every two years. This advancement is important for other aspects of technological progress in computing – such as processing speed or the price of computers.





Kentucky ASYmmetric Zero supercomputer
Built July 16, 2003; 1st Teraflop/s system in KY
Cost \$39K, 25kW + 5 Ton air conditioner



El Capitan supercomputer:
11,039,616 cores, 2.821 Exaflop/s
Cost approx. \$600M, 29.7 MW power



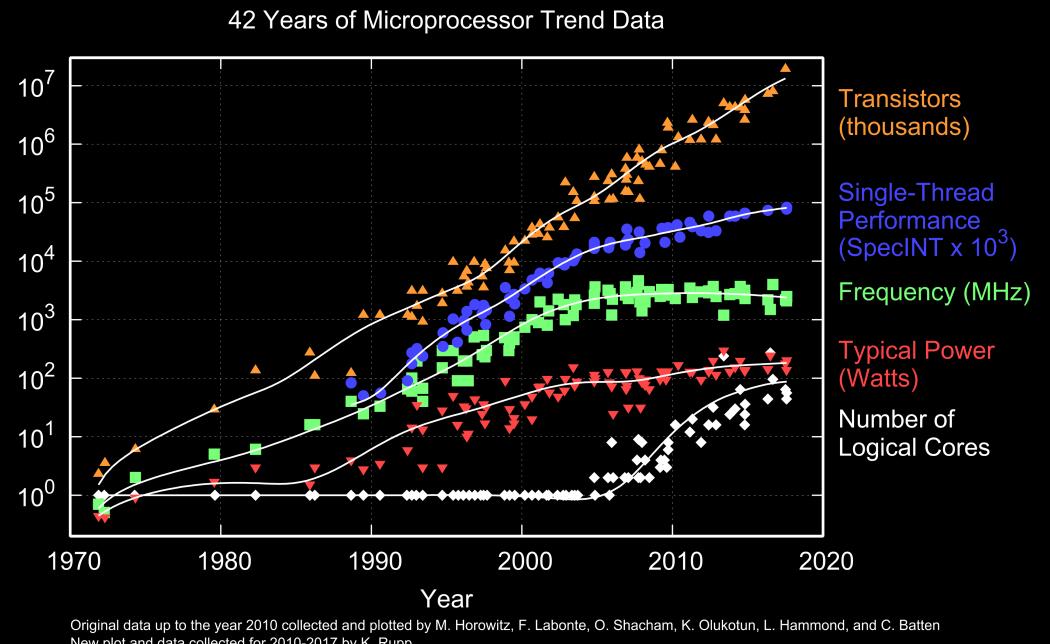
SI Terminology of Scale

1000 ¹	kilo	k	1000 ⁻¹	milli	m
1000 ²	mega	M	1000 ⁻²	micro	u
1000 ³	giga	G	1000 ⁻³	nano	n
1000 ⁴	tera	T	1000 ⁻⁴	pico	p
1000 ⁵	peta	P	1000 ⁻⁵	femto	f
1000 ⁶	exa	E			

- 1000^x vs. 1024^x ; e.g., giga vs. giba
- 1 Byte (B) is 8-10 bits (b), 4 bits in a Nybble
- Hertz (Hz) is frequency (vs. period)
- A flop is a floating-point op like add or multiply

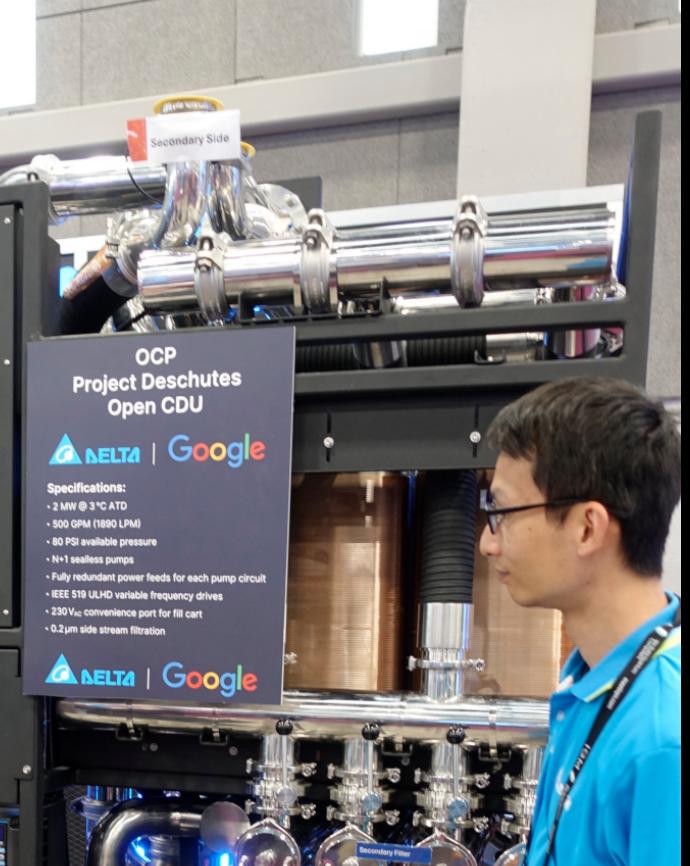
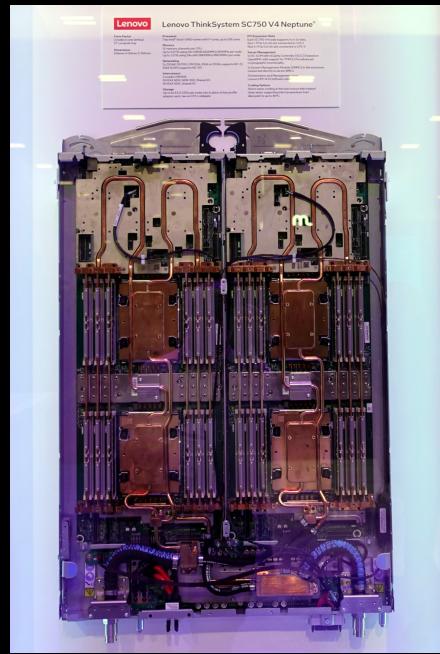
The Bad News

- Moore's Law **is slowing**
- Power/transistor  **slower than** transistors/chip 
- **Individual ops not getting much faster**



Power & Heat

- Chips can have **dark silicon**, special-purpose function units normally unpowered, but...
- Remember my machine room? It has **170kW** and **30 Tons** of air conditioning, but **there are individual racks** it couldn't power & cool!
- **Colossus**, xAI's Memphis, TN supercomputer, in April 2025 had **422MW** gas turbine power!
- Exotic water cooling systems dominated SC25!



Amdahl's Law

- In the best case, as parallelism width $N \rightarrow \infty$, runtime should go to 0, **right?**
- If $1/K$ of the program's work ***cannot*** be run in parallel, the best possible speedup is only K

If a program spends 10% of its time sequentially reading a input file, you can't get more than a 10X speedup by parallelizing just the other 90%!
You need to parallelize everything possible.

Flavors of Parallel Processing

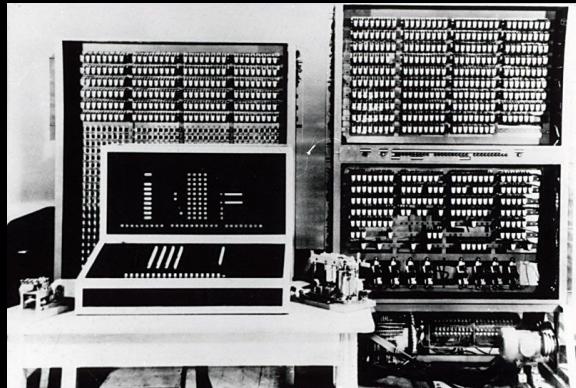
- Bit-parallel within a word
- Pipelined
- Superscalar, VLIW, EPIC
- SWAR (SIMD Within A Register), Vector
- SMP (Symmetric MultiProcessor; multi-core)
- GPU (Graphics Processing Unit)
- Clusters, Farms, Grids, and Clouds

Automatic; Semi-Auto; Explicitly Programmed

Flavors of Parallel Processing

- Bit-parallel within a word
- Pipelined
- Superscalar, VLIW, EPIC
- SWAR (SIMD Within A Register), Vector
- SMP (Symmetric MultiProcessor; multi-core)
- GPU (Graphics Processing Unit)
- ⇒ Quantum
- Clusters, Farms, Grids, and Clouds

Automatic; Semi-Auto; Explicitly Programmed



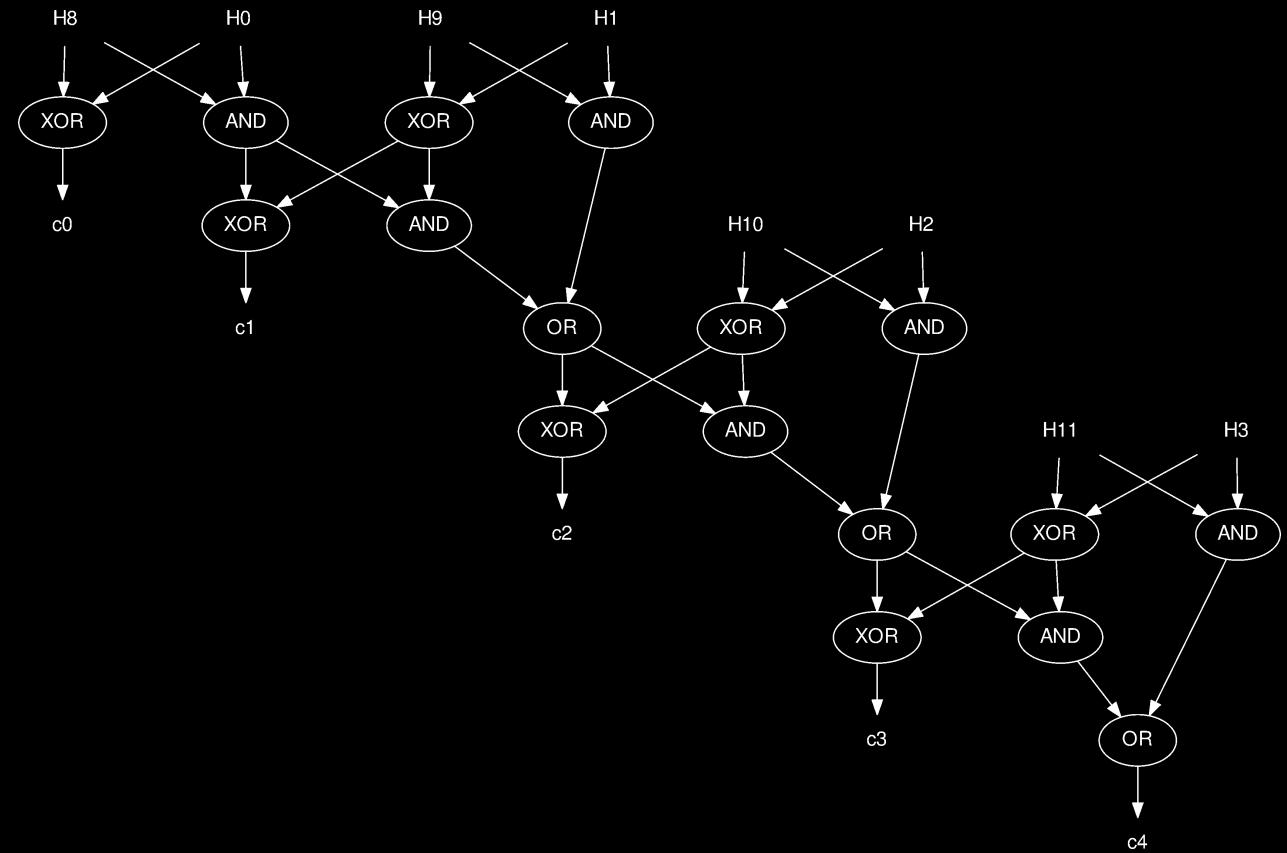
Bit-parallel within a word

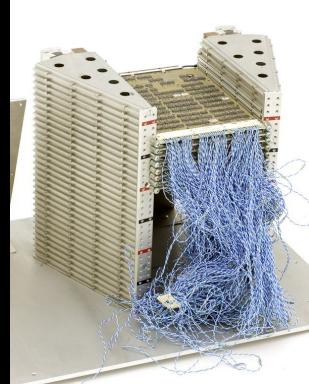
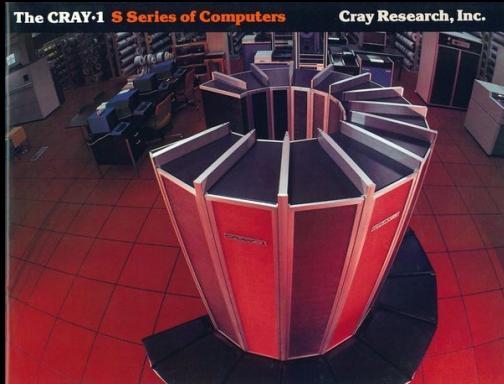
- Used by most computers (even Z3, ENIAC)
- Each machine has a **word size**
- Each operation produced a word result, with many bit-level operations happening in parallel; e.g., **A&B simultaneously $A_0 \& B_0, A_1 \& B_1, \dots$**

Word-Level Parallelism

- Consider adding 4-bit integers:

$$c = a + b$$





Vector parallelism across words

- Used by Cray 1, many others from the 1970s
- Applies **same operation** to all vector elements:
$$A[0..N-1] = B[0..N-1] + C[0..N-1]$$
- A **SIMD** (Single Instruction, Multiple Data) model, but often implemented as a **pipeline**

Vector Pipeline Processing

- Consider the previous example:

$$A[0..N-1] = B[0..N-1] + C[0..N-1]$$

- Suppose addition takes 3 clock cycles:

	Stage 0	Stage 1	Stage 2
Cycle 0:	$B[0]+C[0]$		
Cycle 1:	$B[1]+C[1]$	$B[0]+C[0]$	
Cycle 2:	$B[2]+C[2]$	$B[1]+C[1]$	$A[0]=B[0]+C[0]$
Cycle 3:	$B[3]+C[3]$	$B[2]+C[2]$	$A[1]=B[1]+C[1]$
Cycle 4:	$B[4]+C[4]$	$B[3]+C[3]$	$A[2]=B[2]+C[2]$

Modern Pipeline Processing

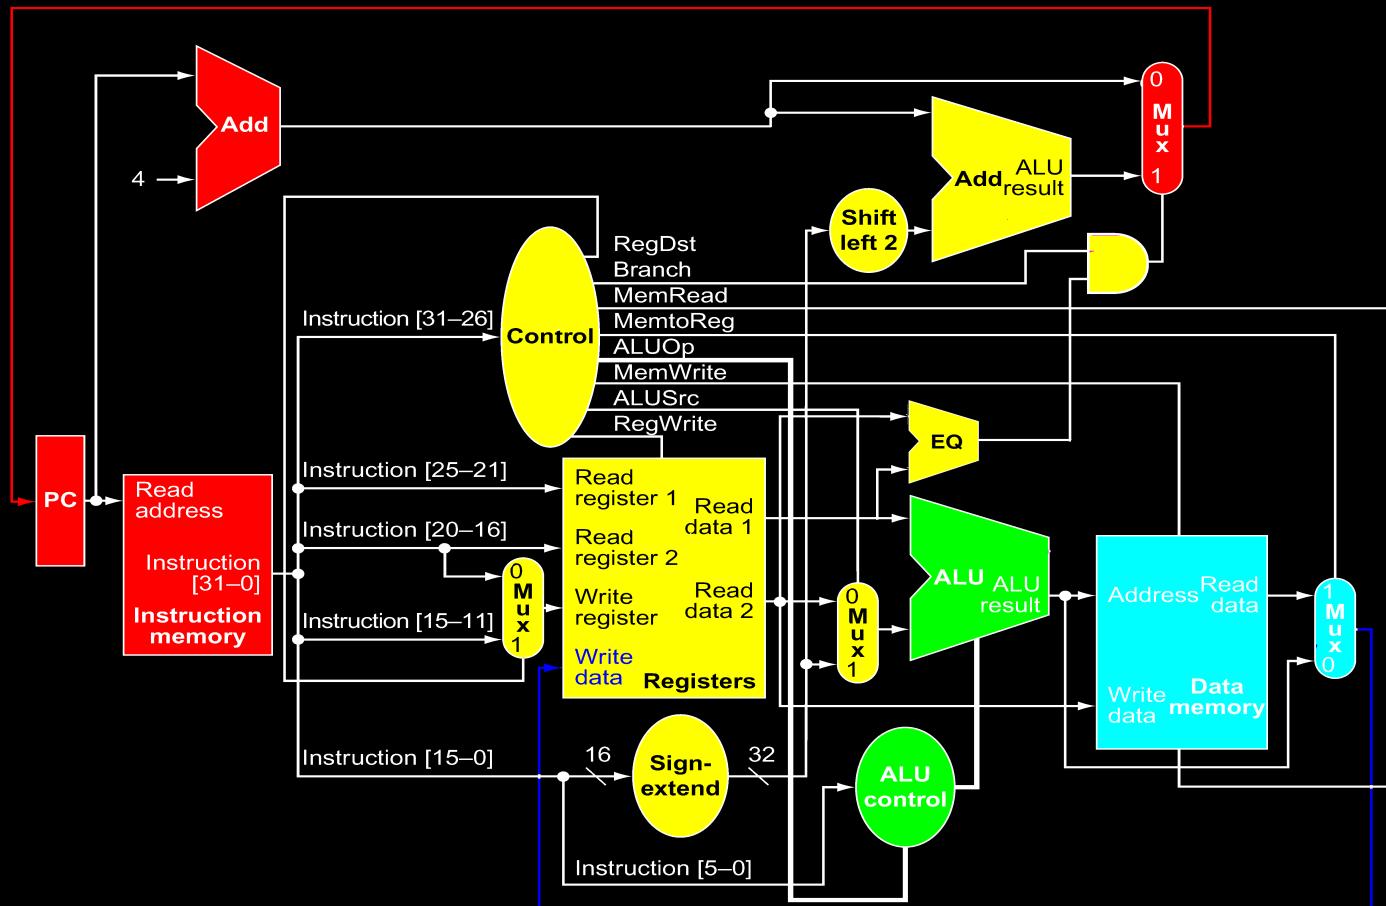
- Consider another example:

`add a,b; sub c,d; mul e,a; xor f,a`

- 3-stage pipe (with bubble):

	Stage 0	Stage 1	Stage 2
Cycle 0:	<code>add a,b</code>		
Cycle 1:	<code>sub c,d</code>	<code>add a,b</code>	
Cycle 2:	<code>mul e,a</code>	<code>sub c,d</code>	<code>a = add a,b</code>
Cycle 3:	<code>mul e,a</code>	<code>nop</code>	<code>c = sub c,d</code>
Cycle 4:	<code>xor f,a</code>	<code>mul e,a</code>	<code>nop</code>

Where Is The Circuitry?

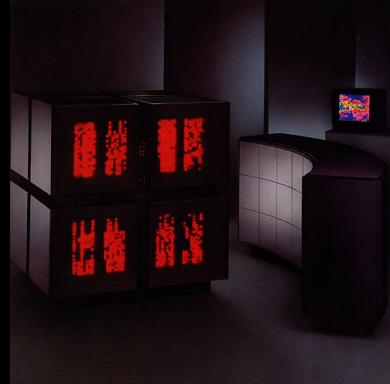


IF: Instruction Fetch
ID: Instruction Decode

EX: Execute
MEM: Memory Access

WB: Write Back

Where Is The Circuitry?

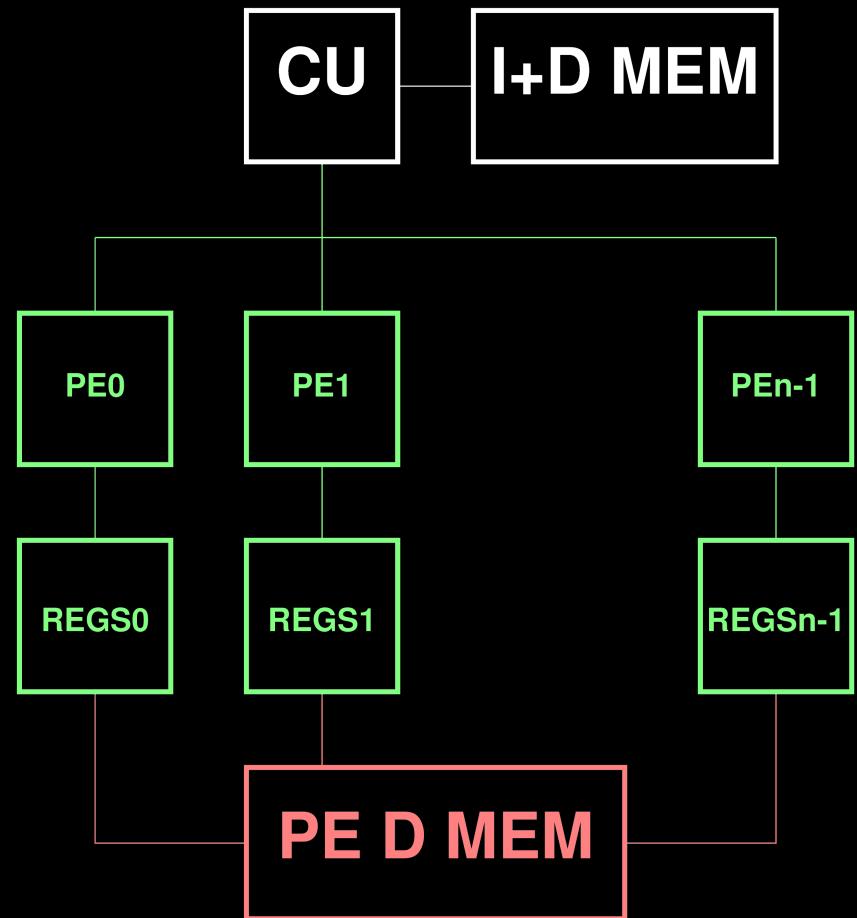


SIMD parallel across PEs (processing elements)

- E.g.: Thinking Machines, MasPar
- Each PE either applies **same operation** or is **disabled** for that operation
- 1980s examples were **bit-serial**, but parallel across massive number of PEs (e.g., 64K)

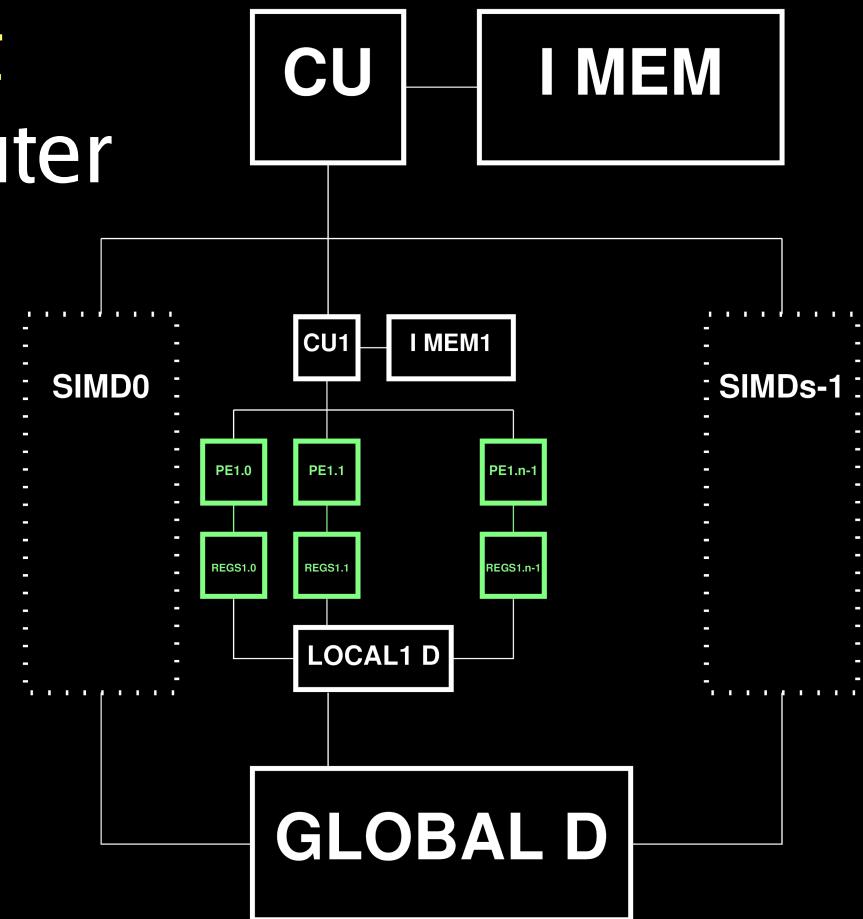
Generic SIMD Architecture

- Control Unit
conventional processor...
- Processing Elements
ALUs + local data
- PE Data Memory
Optional; coherent?
- GPUs *nest* this structure
within each PE+REGS



Generic GPU Architecture

- **Graphics Processing Unit** attached to a **Host computer**
- **Processing Elements** have 2D PE numbering
- Virtualized PEs behave like vector pipelines
- If $\text{PE}_{x,y}$ for all y are disabled, CU_x can skip



Generic SIMD Programming

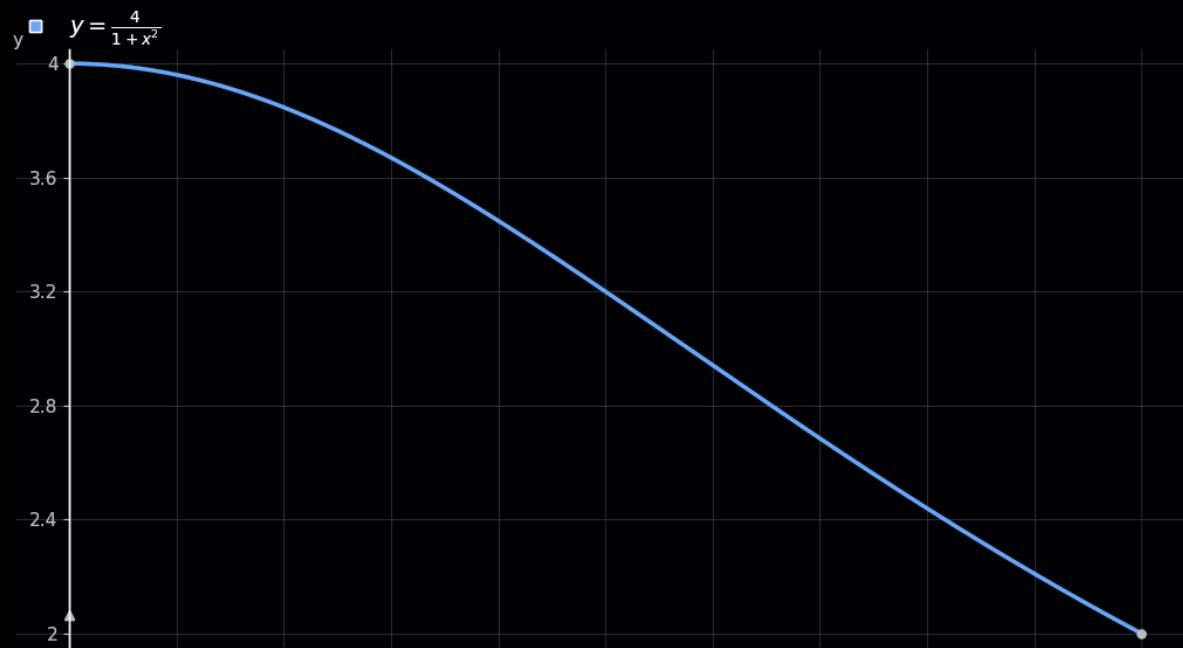
- A program has two parts:
 - Sequential (singular) part run on CU or Host
 - Parallel (plural) part run on PEs
- Parallel code says what happens to an element, hardware applies that to all elements on PEs that are not disabled
- SIMD: **inter-PE communication is cheap**
GPU: **inter-CUx communication is expensive**,
but **lower fanout allows a higher clock speed**

Sample Algorithm: Compute π

- Area of a circle is πr^2 , so $r=1$ area is π
- Area under $y=4/(1+x^2)$ over $x=0..1$ is also π

$$y = \frac{4}{1+x^2} \quad \text{for } x \in [0, 1]$$

Graph of a rational function showing a decreasing curve



π in C as an integral

```
#include <stdlib.h>
#include <stdio.h>

int main(int argc, char **argv) {
    int i, intervals = atoi(argv[1]);
    double width = 1.0 / intervals, sum = 0;
    // sum curve height at each interval
    for (i=0; i<intervals; ++i) {
        double x = (i + 0.5) * width;
        sum += 4.0 / (1.0 + x * x);
    }
    sum *= width; // multiply by width to get area
    printf("Pi is roughly %f\n", sum);
    return(0);
}
```

π in MPL (MasPar's SIMD C)

```
#include <mpl.h> // MPL support defines nproc, etc.

int main(int argc, char **argv) {
    int intervals = atoi(argv[1]);
    double width = 1.0 / intervals, sum = 0;
    plural double dsum = 0;
    plural int i;
    // sum curve height at each interval
    for (i=iproc; i<intervals; i+=nproc) {
        plural double x = (i + 0.5) * width;
        dsum += 4.0 / (1.0 + x * x);
    }
    sum = reduceAddd(dsum); // tree summation of dsum from all PEs
    sum *= width; // multiply by width to get area
    printf("Pi is roughly %f\n", sum);
    return(0);
}
```

π in MPL (MasPar's SIMD C)

```
#include <mpl.h> // MPL support defines nproc, etc.

int main(int argc, char **argv) {
    int intervals = atoi(argv[1]);
    double width = 1.0 / intervals, sum = 0;
    plural double dsum = 0;
    plural int i; int j;
    // sum curve height at each interval
    for (i=iproc; i<intervals; i+=nproc) {
        plural double x = (i + 0.5) * width;
        dsum += 4.0 / (1.0 + x * x);
    }
    // tree summation of dsum from all PEs
    for (j=nproc/2; j>0; j>>=1) dsum += router[iproc+j].dsum;
    sum = proc[0].dsum;
    sum *= width; // multiply by width to get area
    printf("Pi is roughly %f\n", sum);
    return(0);
}
```

π in CUDA (final sum in Host)

```
#include <stdio.h>
#include <cuda.h> // CUDA support
#define INTERVALS 1000000
#define BLOCKS 4 // number of thread blocks (little SIMDs)
#define THREADS 192 // number of threads per block (virtual PEs per little SIMD)

// Kernel that executes on the CUDA device (the GPU)
__global__ void summer(float *sum, int intervals, float width, int nproc) {
    float x; int i, iproc = blockIdx.x*blockDim.x+threadIdx.x; // iproc is global PE number
    for (i=iproc; i<intervals; i+=nproc) {
        x = (i+0.5)*width; sum[iproc] += 4.0/(1.0+x*x); // sum[iproc] is sum computed by PE iproc
    }
}

// Main routine that executes on the Host
int main(void) {
    dim3 dimGrid(BLOCKS,1,1); dimBlock(THREADS,1,1); // Grid and Block dimensions
    float pi = 0, width = 1.0/INTERVALS;
    size_t size = BLOCKS*THREADS*sizeof(float); // array memory size in bytes
    float *sumDev, *sumHost = (float *)malloc(size); // allocate array on Host
    cudaMalloc((void **) &sumDev, size); cudaMemset(sumDev, 0, size); // allocate & zero array on GPU
    summer <<<dimGrid, dimBlock>>> (sumDev, INTERVALS, width, THREADS*BLOCKS); // run GPU kernel
    cudaMemcpy(sumHost, sumDev, size, cudaMemcpyDeviceToHost); // copy sum array from GPU to Host
    for(int i=0; i<THREADS*BLOCKS; ++i) pi += sumHost[i]; // sum partial sums in Host (sequential!)
    pi *= width; // multiply by width to get area
    printf("Pi is roughly %f\n",pi);
    free(sumHost); cudaFree(sumDev); // free data structures (would happen at exit anyway)
    return(0);
}
```



Shared Memory MIMD (Multiple I, Multiple D)

- E.g.: SGI Origin, Multi-Core processors
- Each PE runs a process or thread
- Coherent memory doesn't easily scale...

π in C using PThread Library

```
#include <stdlib.h>
#include <stdio.h>
#include <pthread.h>
volatile double pi = 0.0; // approximation to pi (shared)
pthread_mutex_t pi_lock; // lock for unique access to pi
volatile double intervals; // how many intervals?
#define NPROC 8

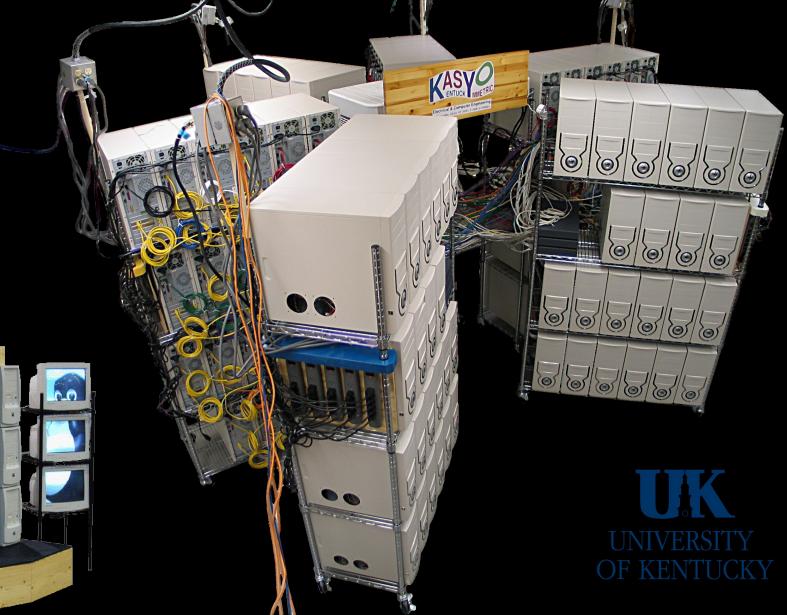
void *process(void *arg) {
    double localsum = 0, width = 1.0 / intervals;
    int i, iproc = *((int *) arg);
    for (i=iproc; i<intervals; i+=NPROC) { // partial summation on this PE
        double x = (i + 0.5) * width; localsum += 4.0 / (1.0 + x * x);
    }
    pthread_mutex_lock(&pi_lock); // wait to have unique access to pi
    pi += localsum; // I'm the only PE here now!
    pthread_mutex_unlock(&pi_lock); // we are done with pi
    return(0);
}

int main(int argc, char **argv) {
    pthread_t thread[NPROC]; void * retval; int ip[NPROC];
    intervals = atoi(argv[1]);
    pthread_mutex_init(&pi_lock, NULL); // initialize the lock on pi to unlocked
    for (int i=0; i<NPROC; ++i) ip[i] = i; // initialize iproc values for threads
    for (int i=0; i<NPROC; ++i) pthread_create(&thread[i], NULL, process, &ip[i])); // make threads
    for (int i=0; i<NPROC; ++i) pthread_join(thread[i], &retval); // join (collapse) threads
    pi *= width; // everybody joined, so multiply by width to get area
    printf("Pi is roughly %f\n", pi);
    return(0);
}
```

π in C using OMP Directives

```
#include <stdlib.h>
#include <stdio.h>
#include <omp.h> // OMP support

int main(int argc, char **argv) {
    int i, intervals = atoi(argv[1]);
    double width = 1.0 / intervals, sum = 0;
    // sum curve height at each interval
#pragma omp parallel for reduction(+: sum) schedule(static)
    for (i=0; i<intervals; ++i) {
        double x = (i + 0.5) * width;
        sum += 4.0 / (1.0 + x * x);
    }
    sum *= width; // multiply by width to get area
    printf("Pi is roughly %f\n", sum);
    return(0);
}
```



UK
UNIVERSITY
OF KENTUCKY

Message-Passing MIMD & Clusters

- E.g.: **ASCI Red, Beowulf, KLAT2, KASYO ...**
- Each PE runs a **process or thread**
- **Coherent memory doesn't easily scale...**

π in C using PVM Library

```
#include <stdlib.h>
#include <stdio.h>
#include <pvm3.h> // Parallel Virtual Machine library
#define NPROC 4 // set number of PE processes

int main(int argc, char **argv) {
    int intervals = atoi(argv[1]);
    double width = 1.0 / intervals, sum = 0;
    int tids[NPROC]; tids[0] = pvm_mytid(); // enroll in PVM
    int iproc = pvm_joingroup("pi"); // join a group
    if (iproc == 0) { // first PE process in group creates the rest
        pvm_spawn("pvm_pi", &argv[1], 0, NULL, NPROC-1, &tids[1]);
    }
    pvm_barrier("pi", NPROC); // ensure all PEs exist
    for (i=iproc; i<intervals; i+=NPROC) {
        double x = (i + 0.5) * width; sum += 4.0 / (1.0 + x * x);
    }
    // tree reduction of sum from all PEs
    pvm_reduce(PvmSum, &sum, 1, PVM_DOUBLE, 4, "pi", 0);
    sum *= width;
    if (iproc == 0) printf("Pi is roughly %f\n", sum); // only PE0 prints
    // check all PEs are here, leave group, and exit PVM
    pvm_barrier("pi", NPROC); pvm_lvgroup("pi"); pvm_exit();
    return(0);
}
```

π in C using AFAPI Library

```
#include <stdlib.h>
#include <stdio.h>
#include "afapi.h" // Aggregate Function API

int main(int argc, char **argv) {
    int intervals = atoi(argv[1]);
    double width = 1.0 / intervals, sum = 0;
    if (p_init()) exit(1); // check in with AF network hardware
    for (i=IPROC; i<intervals; i+=NPROC) {
        double x = (i + 0.5) * width; sum += 4.0 / (1.0 + x * x);
    }
    sum = p_reduceAdd64f(sum); // AF network sums from all PEs
    sum *= width;
    if (IPROC == CPROC) { // only AFAPI console PE prints
        printf("Pi is roughly %f\n", sum);
    }
    p_exit(); // check out with AF network hardware
    return(0);
}
```

π in C using MPI Messages

```
#include <stdlib.h>
#include <stdio.h>
#include <mpi.h> // MPI support

int main(int argc, char **argv) {
    int i, intervals = atoi(argv[1]);
    double width = 1.0 / intervals, sum = 0, lsum = 0;
    int nproc, iproc; MPI_Status s;
    if (MPI_Init(&argc, &argv) != MPI_SUCCESS) exit(1); // check in
    MPI_Comm_size(MPI_COMM_WORLD, &nproc); // how many PEs?
    MPI_Comm_rank(MPI_COMM_WORLD, &iproc); // who am I?
    for (i=iproc; i<intervals; i+=nproc) {
        double x = (i + 0.5) * width; lsum += 4.0 / (1.0 + x * x);
    }
    lsum *= width;
    if (iproc != 0) {
        MPI_Send(&lsum, 1, MPI_DOUBLE, 0, 0, MPI_COMM_WORLD); // send to PE0
    } else {
        sum = lsum;
        for (i=1; i<nproc; ++i) { // add lsum from each of PE1..PEnproc
            MPI_Recv(&lsum, 1, MPI_DOUBLE, MPI_ANY_SOURCE, MPI_ANY_TAG, MPI_COMM_WORLD, &s);
            sum += lsum;
        }
        printf("Pi is roughly %f\n", sum);
    }
    MPI_Finalize(); return(0); // check out and exit
}
```

π in C using MPI Collectives

```
#include <stdlib.h>
#include <stdio.h>
#include <mpi.h> // MPI support

int main(int argc, char **argv) {
    int i, intervals = atoi(argv[1]);
    double width = 1.0 / intervals, sum = 0, lsum = 0;
    int nproc, iproc; MPI_Status s;
    if (MPI_Init(&argc, &argv) != MPI_SUCCESS) exit(1); // check in
    MPI_Comm_size(MPI_COMM_WORLD, &nproc); // how many PEs?
    MPI_Comm_rank(MPI_COMM_WORLD, &iproc); // who am I?
    for (i=iproc; i<intervals; i+=nproc) {
        double x = (i + 0.5) * width; lsum += 4.0 / (1.0 + x * x);
    }
    lsum *= width;
    // collectively tree reduce lsum values into a single sum by adding
    MPI_Reduce(&lsum, &sum, 1, MPI_DOUBLE, MPI_SUM, 0, MPI_COMM_WORLD);
    if (iproc == 0) {
        printf("Pi is roughly %f\n", sum);
    }
    MPI_Finalize(); return(0); // check out and exit
}
```

What is a Quantum Computer?

Is this a Quantum Computer?



Is this a Quantum Computer?

Yup!

Google
Sycamore



Is this a Quantum Computer?



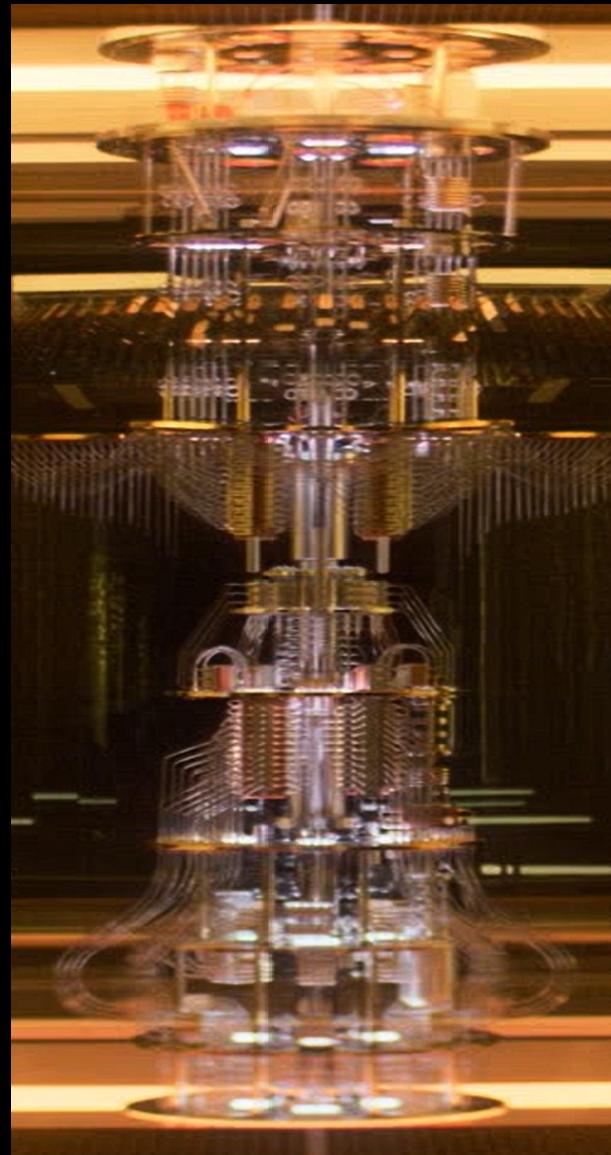
Is this a Quantum Computer?



Yup!

SpinQ
Gemini
Mini
~\$8K

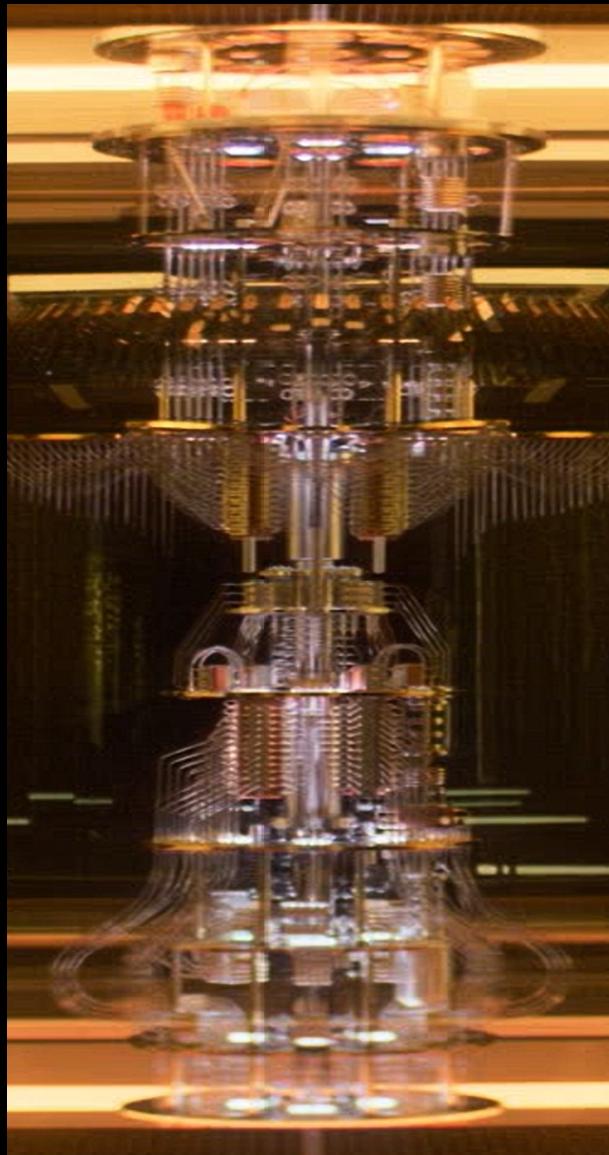
Is this a Quantum Computer?



Is this a Quantum Computer?

Nope!

It is from the
2020 TV mini
series **DEVS**



Is this a Quantum Computer?



Is this a Quantum Computer?

Yup!

It is a
D-Wave
2000Q



Is this a Quantum Computer?



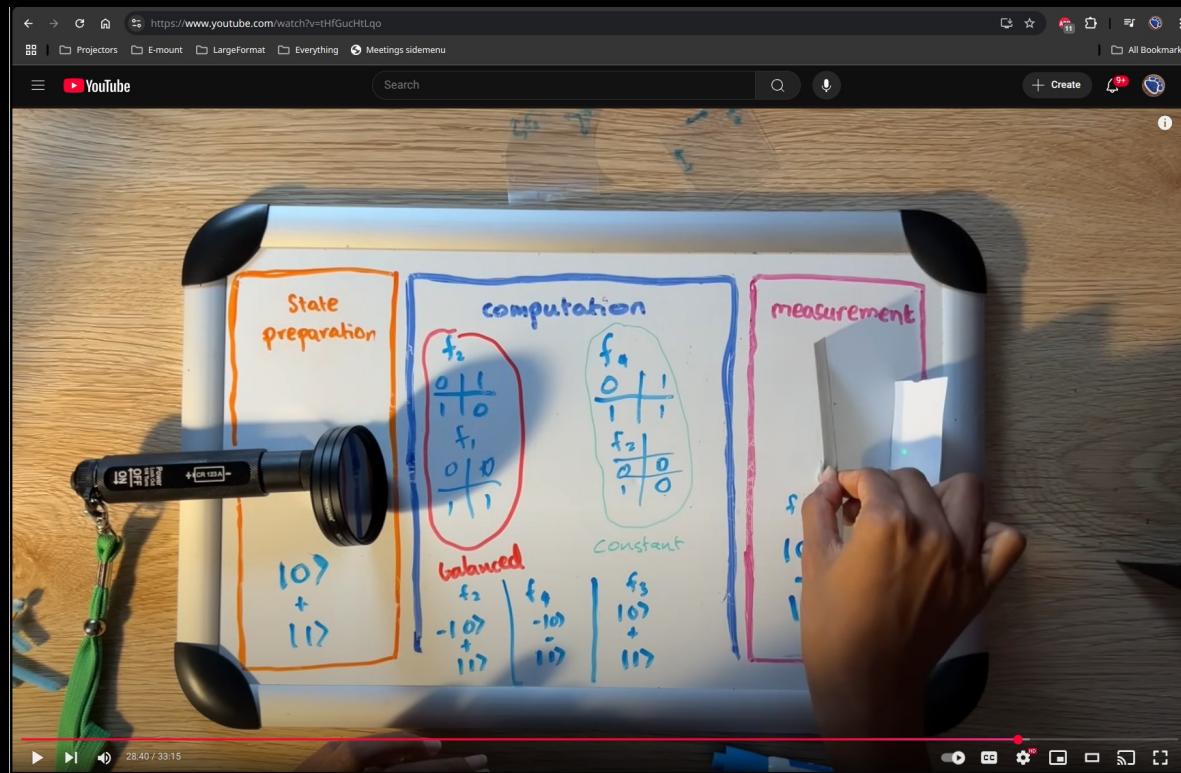
Is this a Quantum Computer?

Not quite.

1/2-scale
model of
Fujitsu



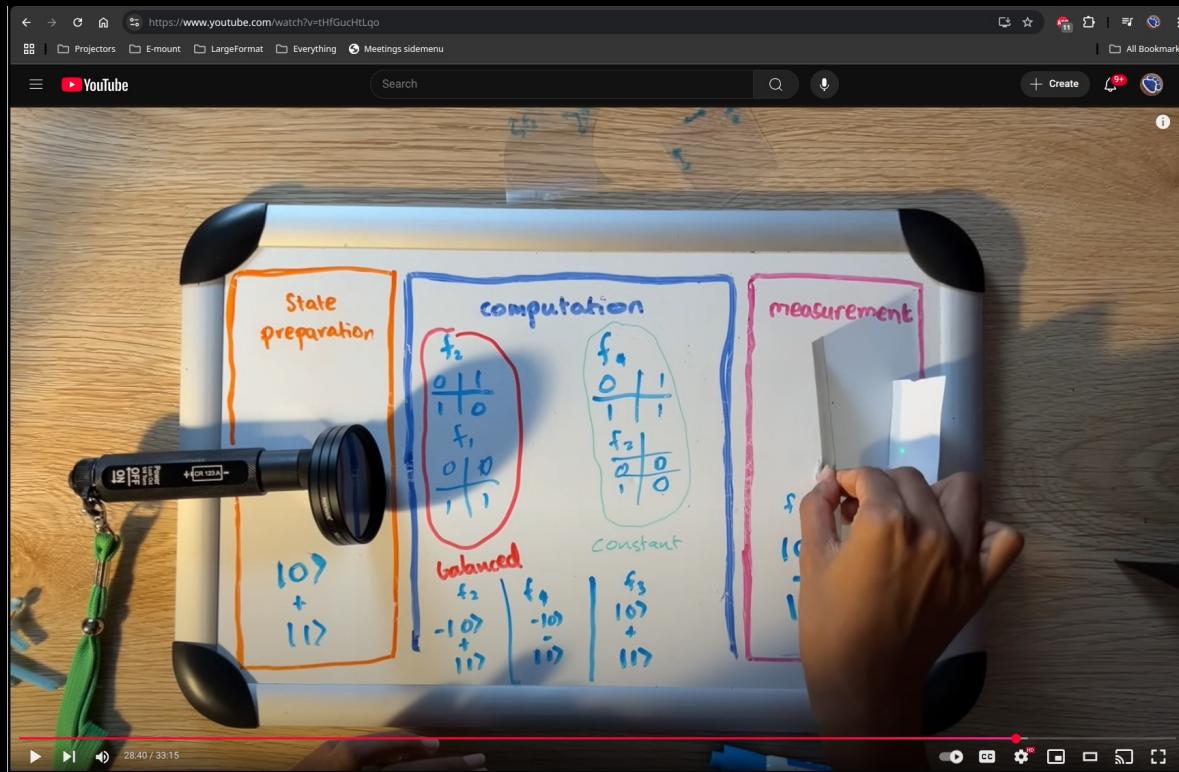
Is this a Quantum Computer?



Is this a Quantum Computer?

Yup.

A photonic qubit using polarization



Is this a Quantum Computer?



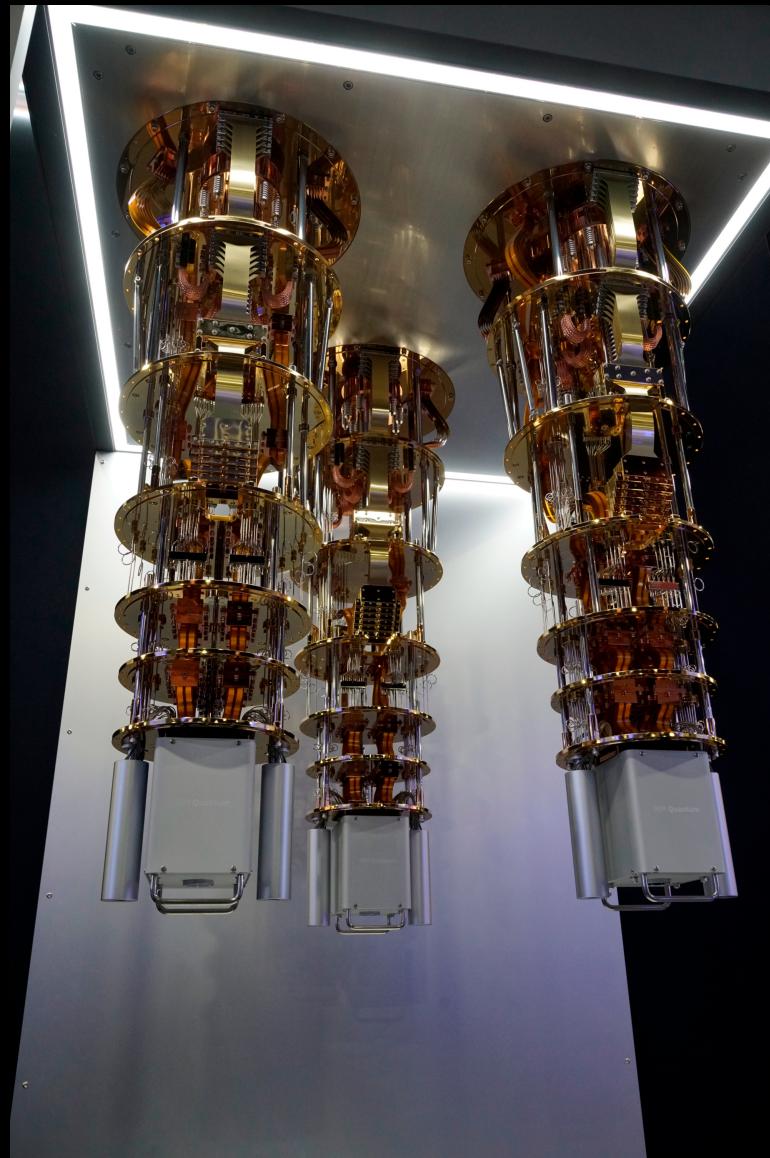
Is this a Quantum Computer?

Yup!

IBM Q



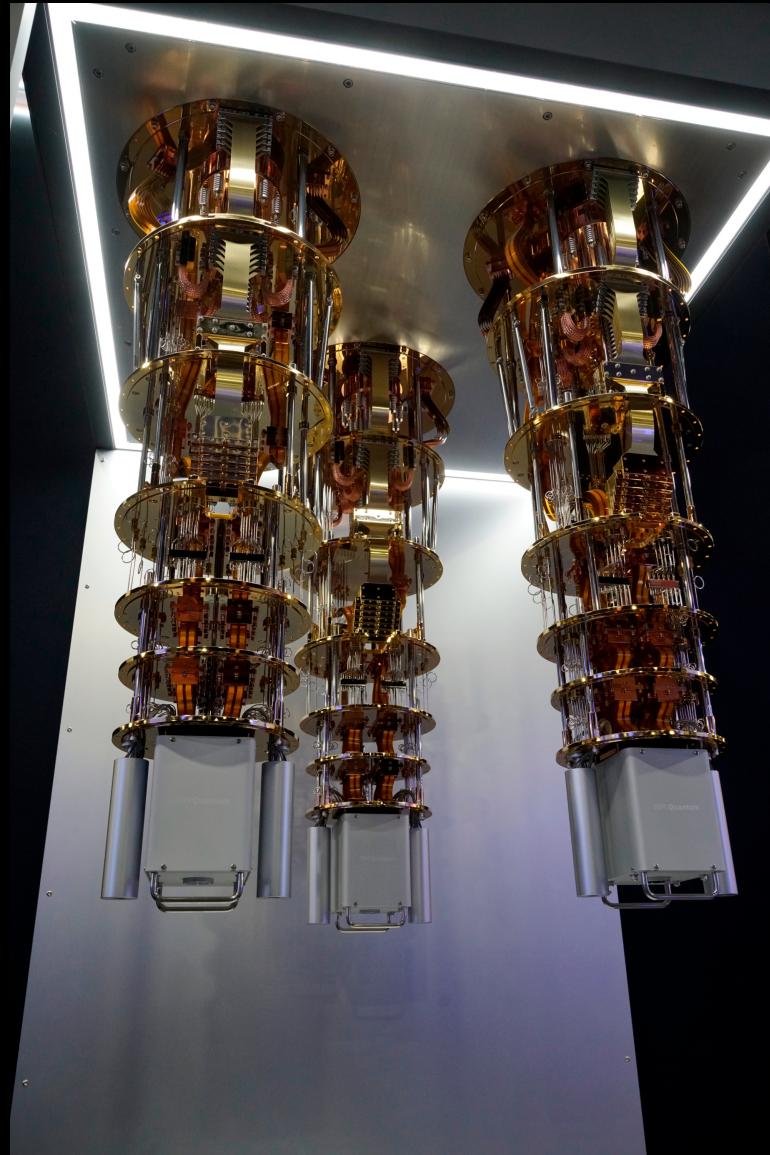
Is this 3 Quantum Computers?



Is this 3 Quantum Computers?

Not quite.

80% scale
model of
**IBM Quantum
System Two**



Is this a Quantum Computer?



Is this a Quantum Computer?

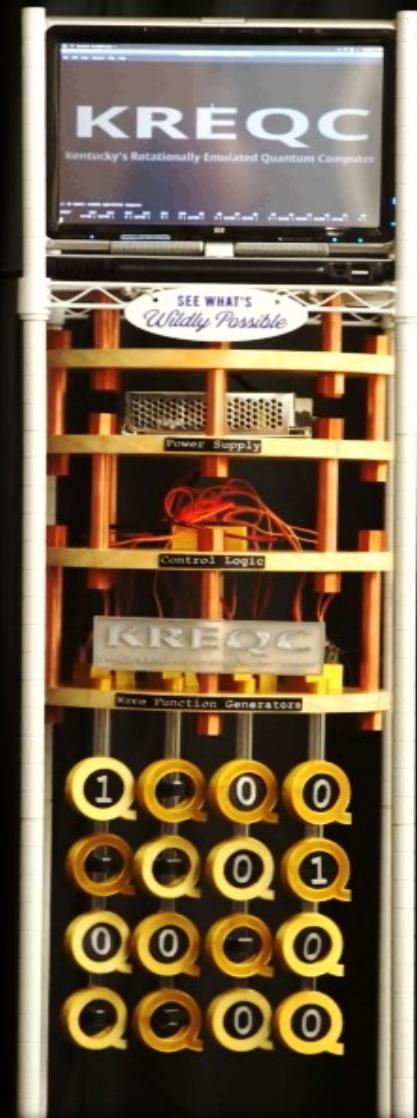
Nope!

BLUEFORS

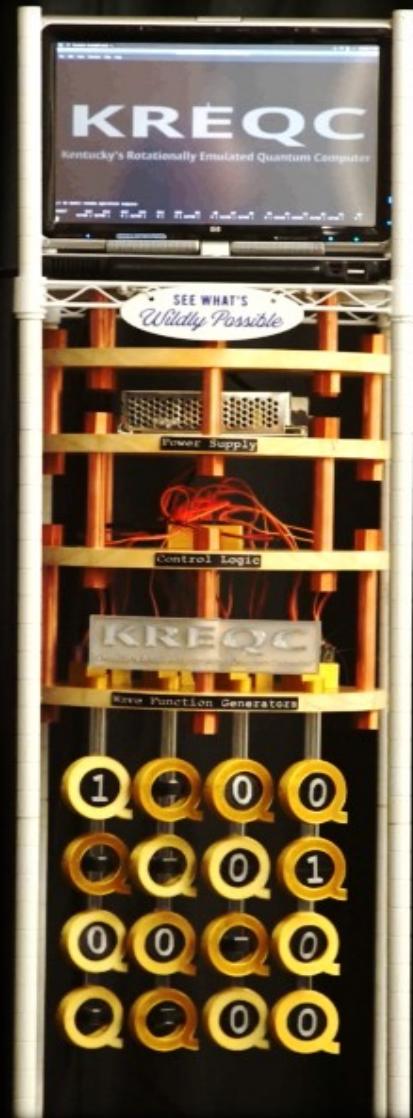
This is a
dummy; they
make cooling,
not computers



Is this a Quantum Computer?



Is this a Quantum Computer?



Not quantum;
quantum-inspired.

KREQC is
Kentucky's
Rotationally
Emulated
Quantum
Computer

What is a Quantum Computer?

- Operates on **Qubits** rather than bits
 - A bit can be either 0 or 1
 - A qubit can be 0, 1, or a **Probability Density Function** over {0,1}
 - A gate-level op on a qubit alters the PDF
- What is a **probability density function**?
 - *Encodes* probability of 0 or 1
 - Not probability because it has more than 1D
- A qubit is really a wave. Isn't everything? ;-)

What is a Quantum Computer?

- Two or more qubits can be **Entangled**
 - Entangled things have their waves aligned
 - This implies *E* entangled qubits can hold a PDF over all 2^E possible *E*-bit values
 - Entangled things are coupled, creating what Einstein called *Spooky action at a distance*
- Measuring a qubit's value always gives 0 or 1, and makes that qubit's value what you read: **measurement collapses superposition**... or does it? Maybe you just get entangled too? ;-)

What is a Quantum Computer?

- We think **everything** in the Universe operates on **quantum mechanics**, so there are **lots of ways** to implement **qubits** and **operations** on them
- The catch is that it is very difficult to keep qubits from interacting with other stuff
 - Noise corrupts/collapses superposition
 - *Performing an operation is injecting noise*
- Only **thermodynamically reversible** ops are viable, and which depends on mechanism used

What is a Quantum Computer?

- Sequential code could work if we had enough qubits, etc., but we don't
- Using superposed values to simulate vectors could compute the local sums, but then how do we sum across all of them?
- No obvious way to do different operations on individual superposed value components at the same time, so it doesn't match the MIMD model
- Need a somewhat different algorithm...

π in a Quantum Computer?

- Sequential code could work if we had enough qubits, etc., **but we don't**
- Using superposed values to simulate vectors could compute the local sums, but then **how do we sum across all of them?**
- No obvious way to do different operations on individual superposed value components at the **same time**, so it doesn't match the **MIMD** model
- Need a somewhat different algorithm...

π in C using int sampling

```
#include <stdlib.h>
#include <stdio.h>

int main(int argc, char **argv) {
    int i, j, intervals = atoi(argv[1]);
    int w = intervals * intervals, sum = 0;
    for (i=0; i<intervals; ++i) {
        int h = w / (intervals + ((i*i) / intervals));
        for (j=0; j<intervals; ++j) {
            if (h > j) ++sum;
        }
    }
    printf("Pi is roughly %f\n", (4.0 * sum) / w);
    return(0);
}
```

π in PBP, pint sampling

```
#include "pbp.h" // PBP classes and support

int main(int argc, char **argv) {
    int bits = atoi(argv[1]); // number of pbits
    pint intervals(1 << bits); // intervals in pbits
    pint w(1 << (2 * bits)); // int scaling factor
    pint x = pint(0).Had(bits); // all x values
    pint y = pint(0).Had(bits,bits); // all y values
    pint h = w / (((x * x) >> bits) + intervals);
    pint r = (h > y); // r is 1 where below curve
    // count 1s; quantum would sample probability
    double pi = (4.0 * r.Pop()) / (1 << REWAYS);
    printf("Pi is roughly %f\n", pi);
    return(0);
}
```