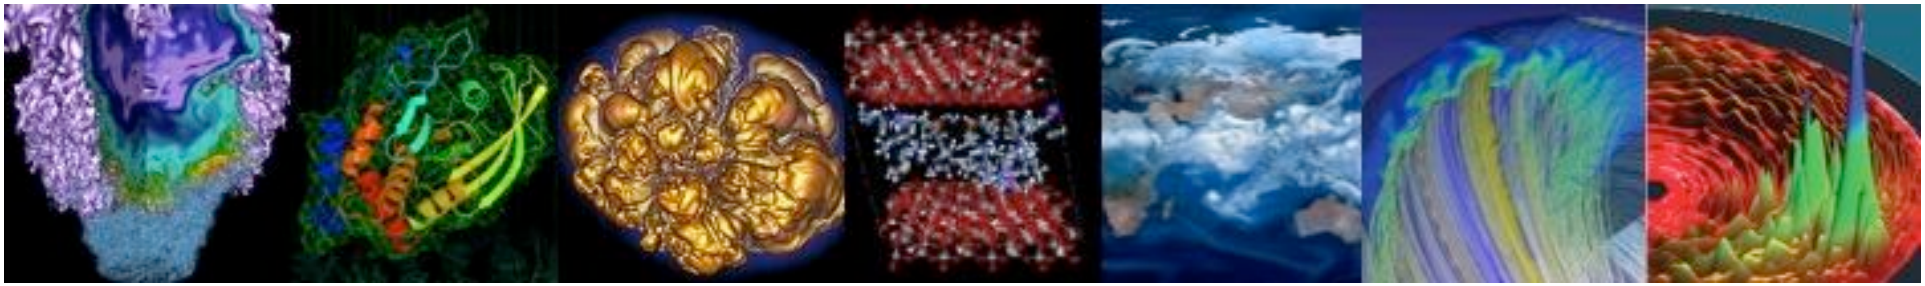

More and Moore: Growing Computing Performance for Scientific Discovery

Kathy Yelick
Associate Laboratory Director
Acting NERSC Director
Lawrence Berkeley National Laboratory

EECS Professor, UC Berkeley

National Energy Research
Scientific Computing Center

High Performance Computing in Science

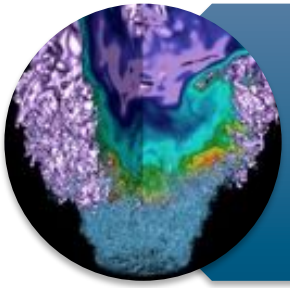

Science at Scale

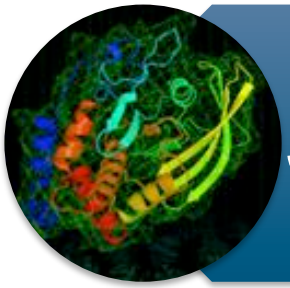

Science at Volume

Science in Data

Science at Scale: Petascale Simulations Aid in the Energy Efficient Devices

- **Combustion simulations improve future designs**
  - **Model fluid flow, burning and chemistry**
  - **Uses advanced math algorithms**
  - **Requires petascale systems today**

Simulations reveal features not visible in lab experiments

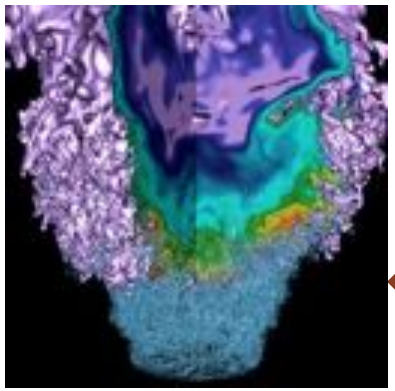

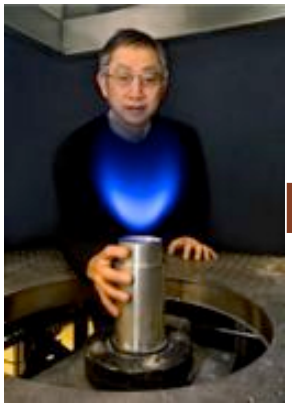

Energy efficient, low emissions technology licensed by industry

- **Need exascale computing to design for alternative fuels, new devices**

Science in Computing Volume: Screening from Diseases to Batteries

- **Large number of simulations covering a variety of related materials, chemicals, proteins,…**

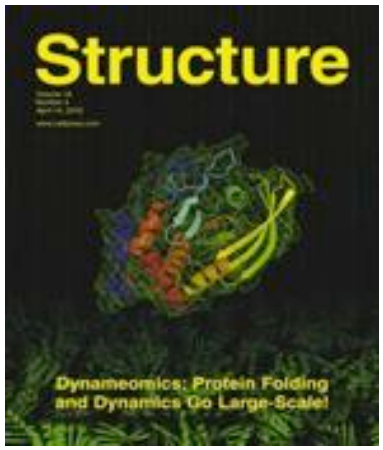

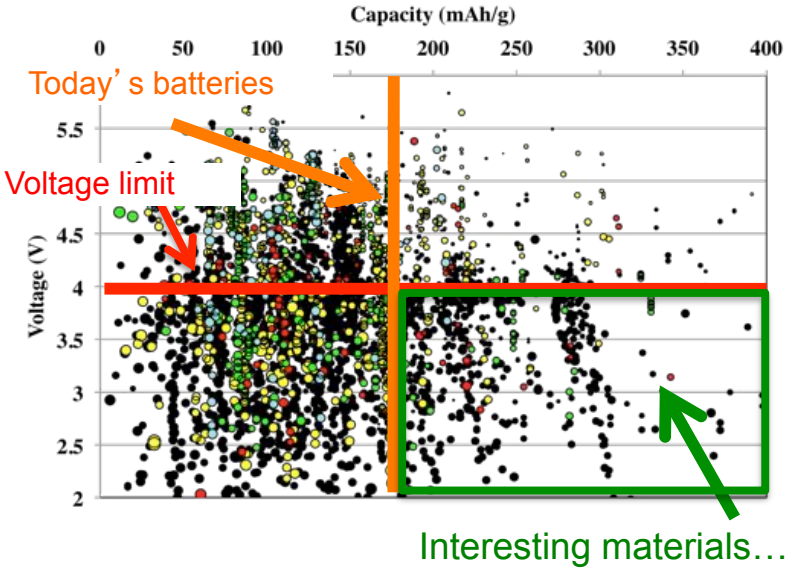

Dynameomics Database
Improve understanding of disease and drug design, e.g., 11,000 protein unfolding simulations stored in a public database.

Materials Genome
Cut in half the 18 years from design to manufacturing, e.g., 20,000 potential battery materials stored in a database

Science in the Data: From Simulation to Image Analysis in Astronomy

HPC used in 2011 Nobel Prize

- Type Ia supernovae used as “standard candles” to measure distance.
- Simulations at NERSC in late 90s modeled the appearance from Earth.

The research shows that the universe is expanding at an accelerating rate. The nature of the dark energy force behind this may be the most important problem in 21st century physics.

More recently: astrophysics discover early nearby supernova.

- Palamor Transient Factory runs machine learning algorithms on ~300GB/night delivered by ESnet “science network”
- Rare glimpse of a supernova within hours of explosion, 20M light years away
- Telescopes world-wide redirected to catch images

NITRD Has Moved Scientists through Difficult Technology Transitions

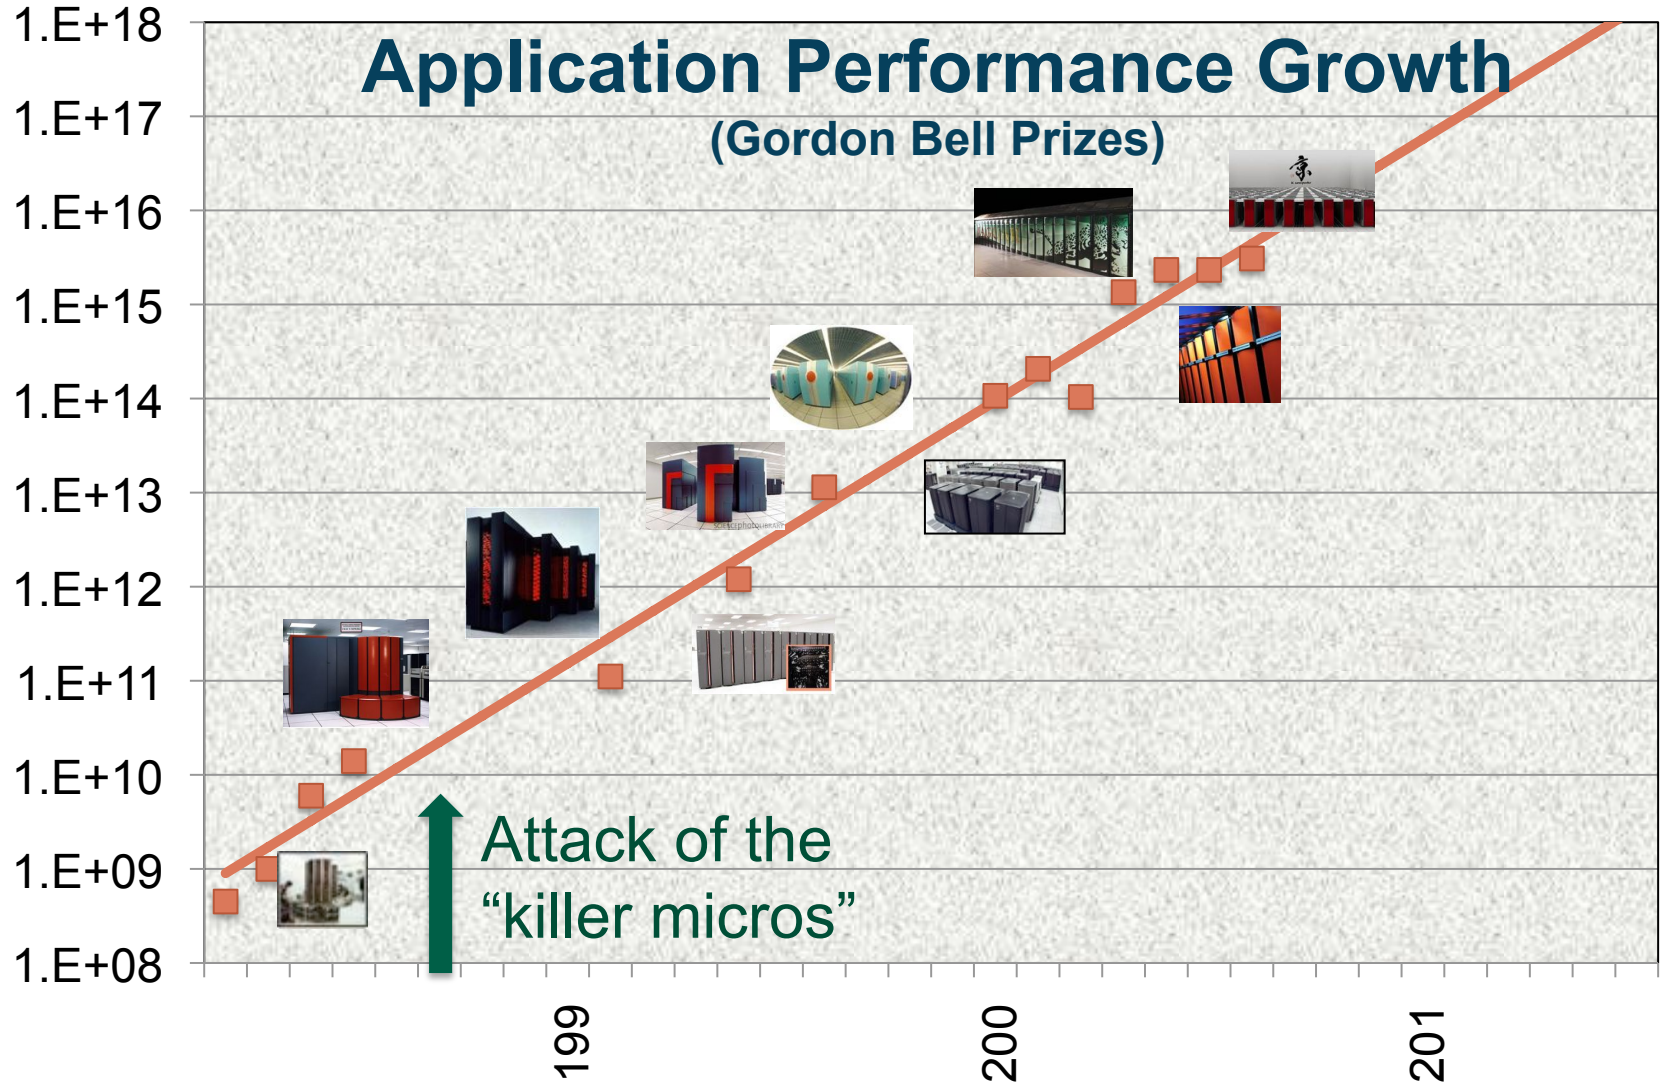

HPC: From Vector Supercomputers to Massively Parallel Systems

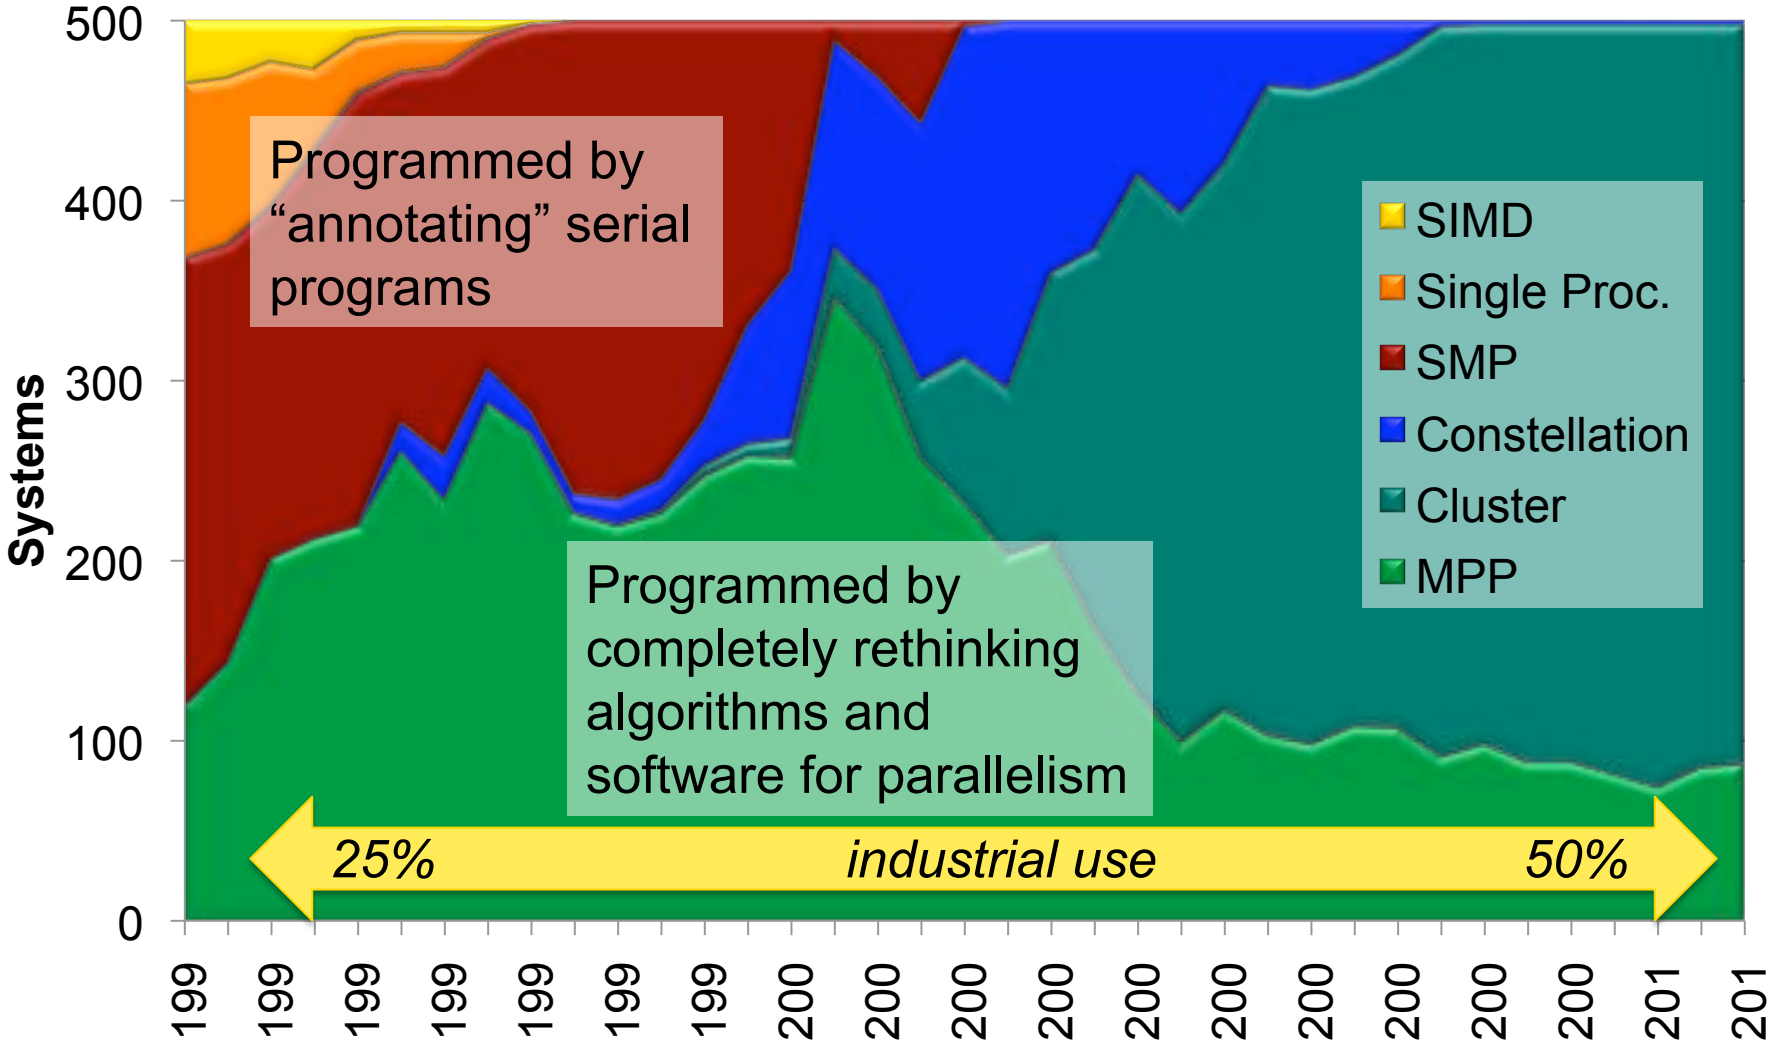

NITRD agency success story: MPI (1992-)

- The message passing interface (MPI) is a standard library
- MPI Forum first met April 1992, released MPI in June 1994
- Involved 80 people from 40 organizations (industry, academia, government labs) supported by NITRD projects and funded centrally by ARPA and NSF
- Scales to millions of processors with separate memory spaces.
- Hardware-portable, multi-language communication library
- Enabled billions of dollars of applications
- MPI still under development as hardware and applications evolve


MPI Forum, March 2008, Chicago

Programming Challenges and NITRD Solutions

- *Application complexity grew due to parallelism and more ambitious science problems (e.g., multiphysics, multiscale)*
- *Scientific libraries enable these applications*

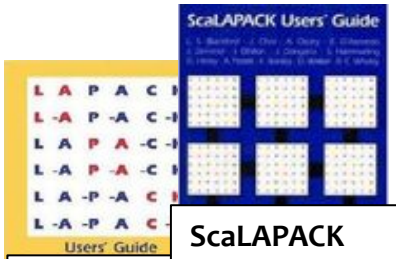

LAPACK
35% of apps

ScaLAPACK
20% of apps

netCDF:
~12% of apps

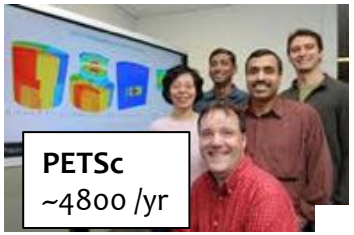

PETSc
~4800 /yr

FFTW: ~25% of apps

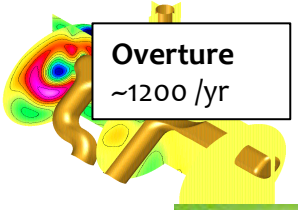

Overture
~1200 /yr

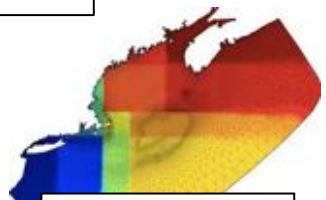

METIS
4% of apps

Trilinos
21,000 total

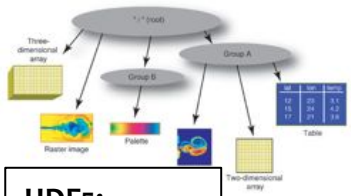

HDF5:
~11% of apps

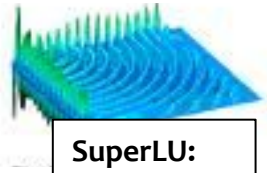

SuperLU:
~4% of apps

FastBit
6,300 total

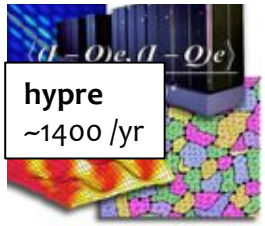

hypre
~1400 /yr

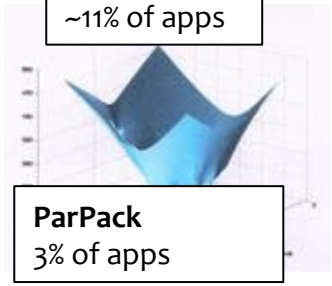

ParPack
3% of apps

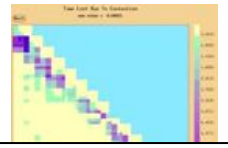

GlobalArrays
28% of apps

Numbers show downloads per year or total; percentages are based on the percentage of NERSC projects that use this library

NITRD Projects Addressed Programmer Productivity of Irregular Problems

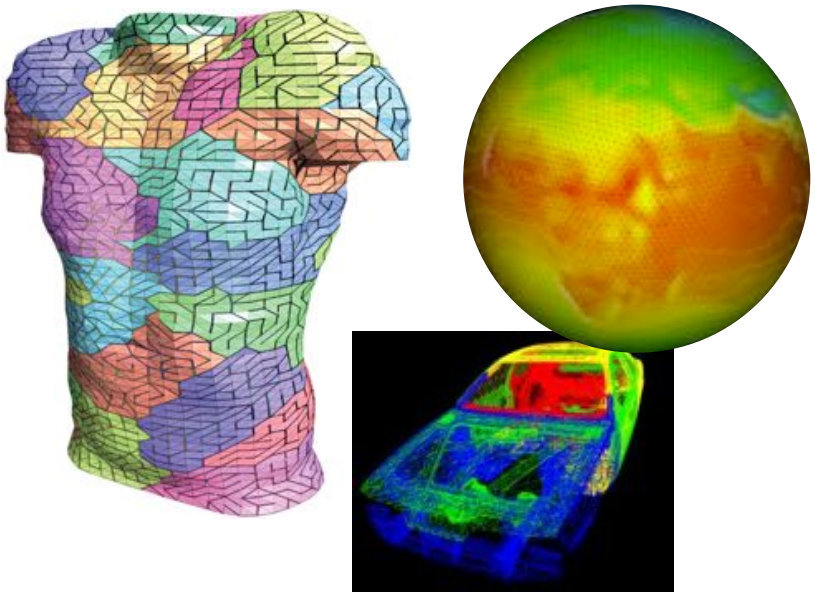

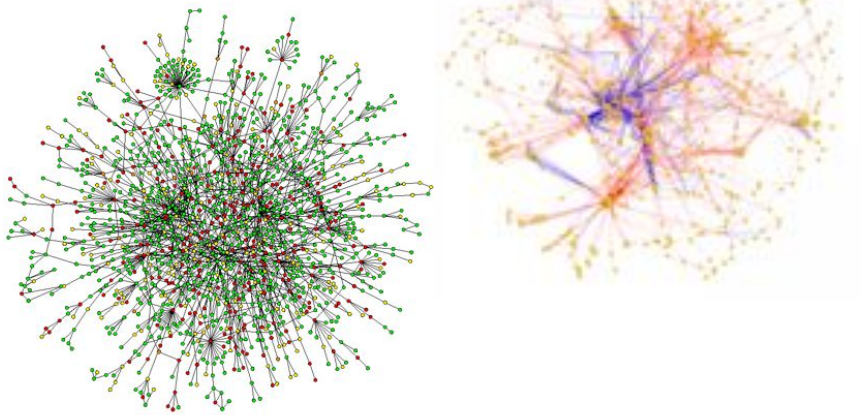

Message Passing Programming

Divide up domain in pieces

Compute one piece and exchange

PVM, MPI, and many libraries

Global Address Space Programming

Each start computing

Grab whatever / whenever

UPC, CAF, X10, Chapel, Fortress, Titanium, GA,

Scientists Need to Undertake another Difficult Technology Transitions

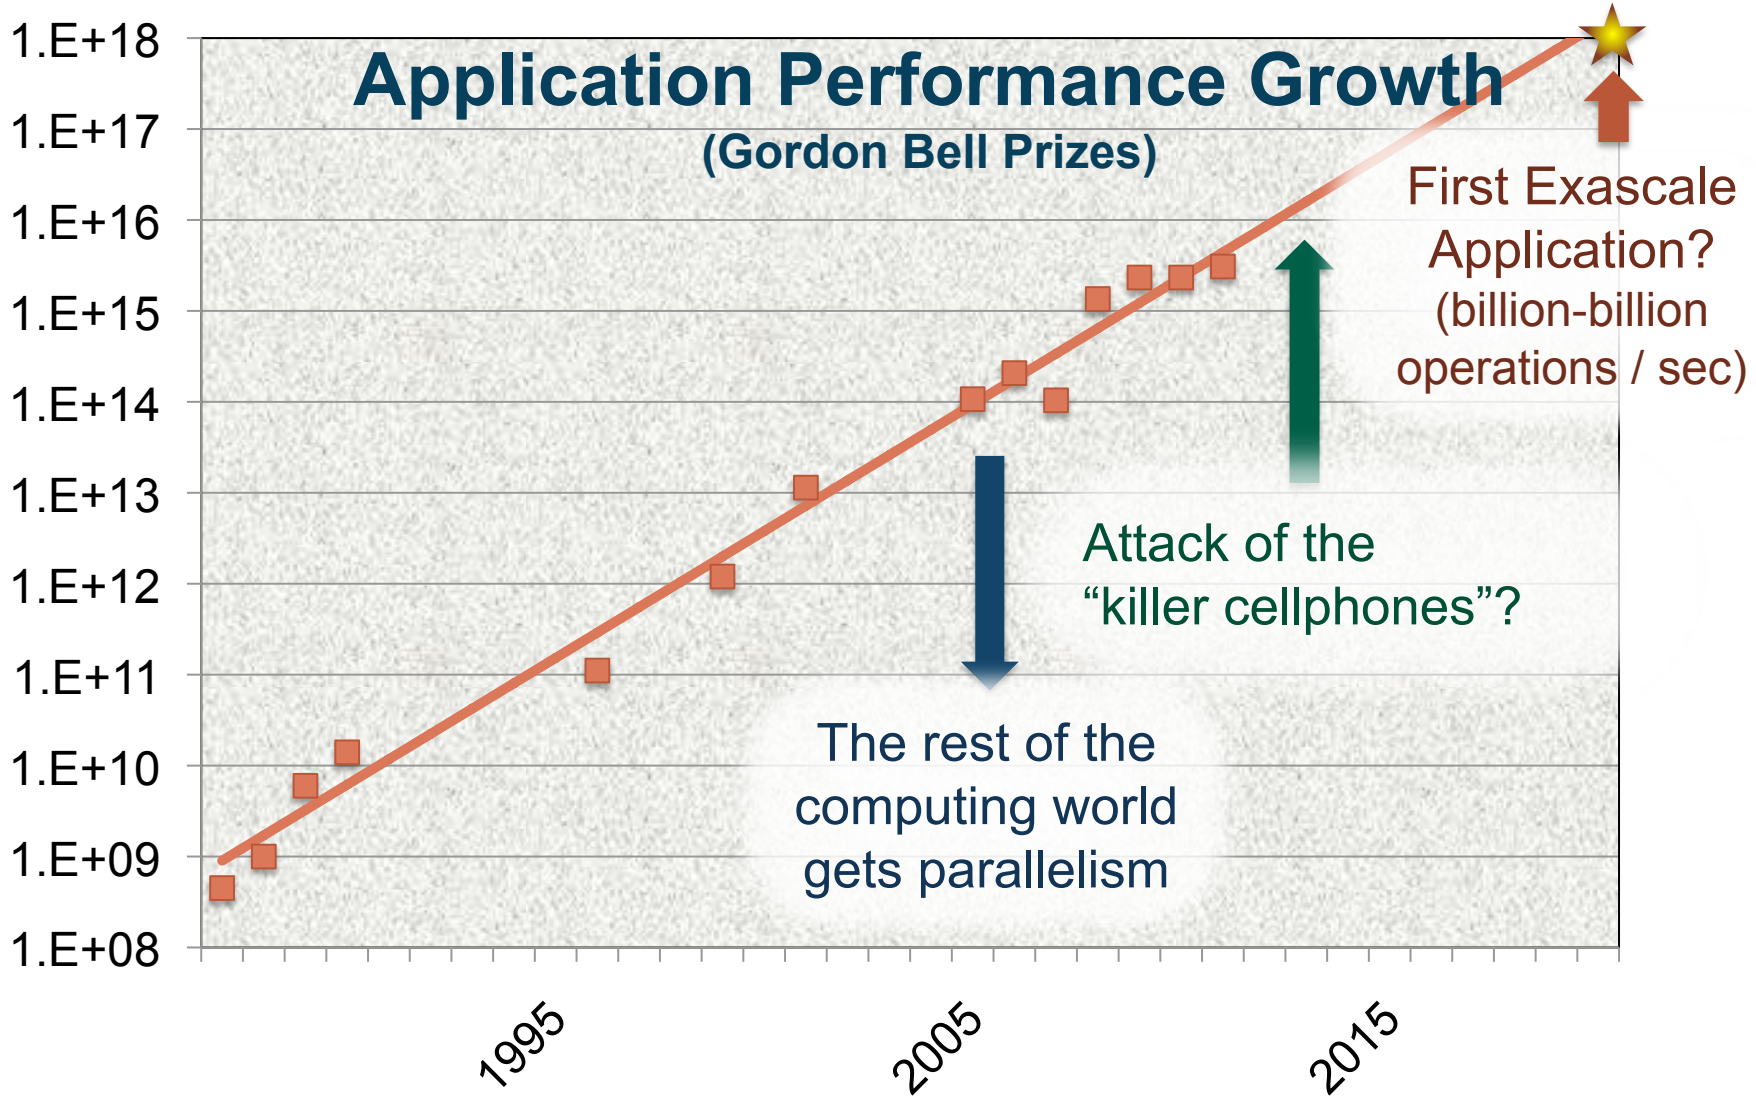

Computing Performance Improvements will be Harder than Ever

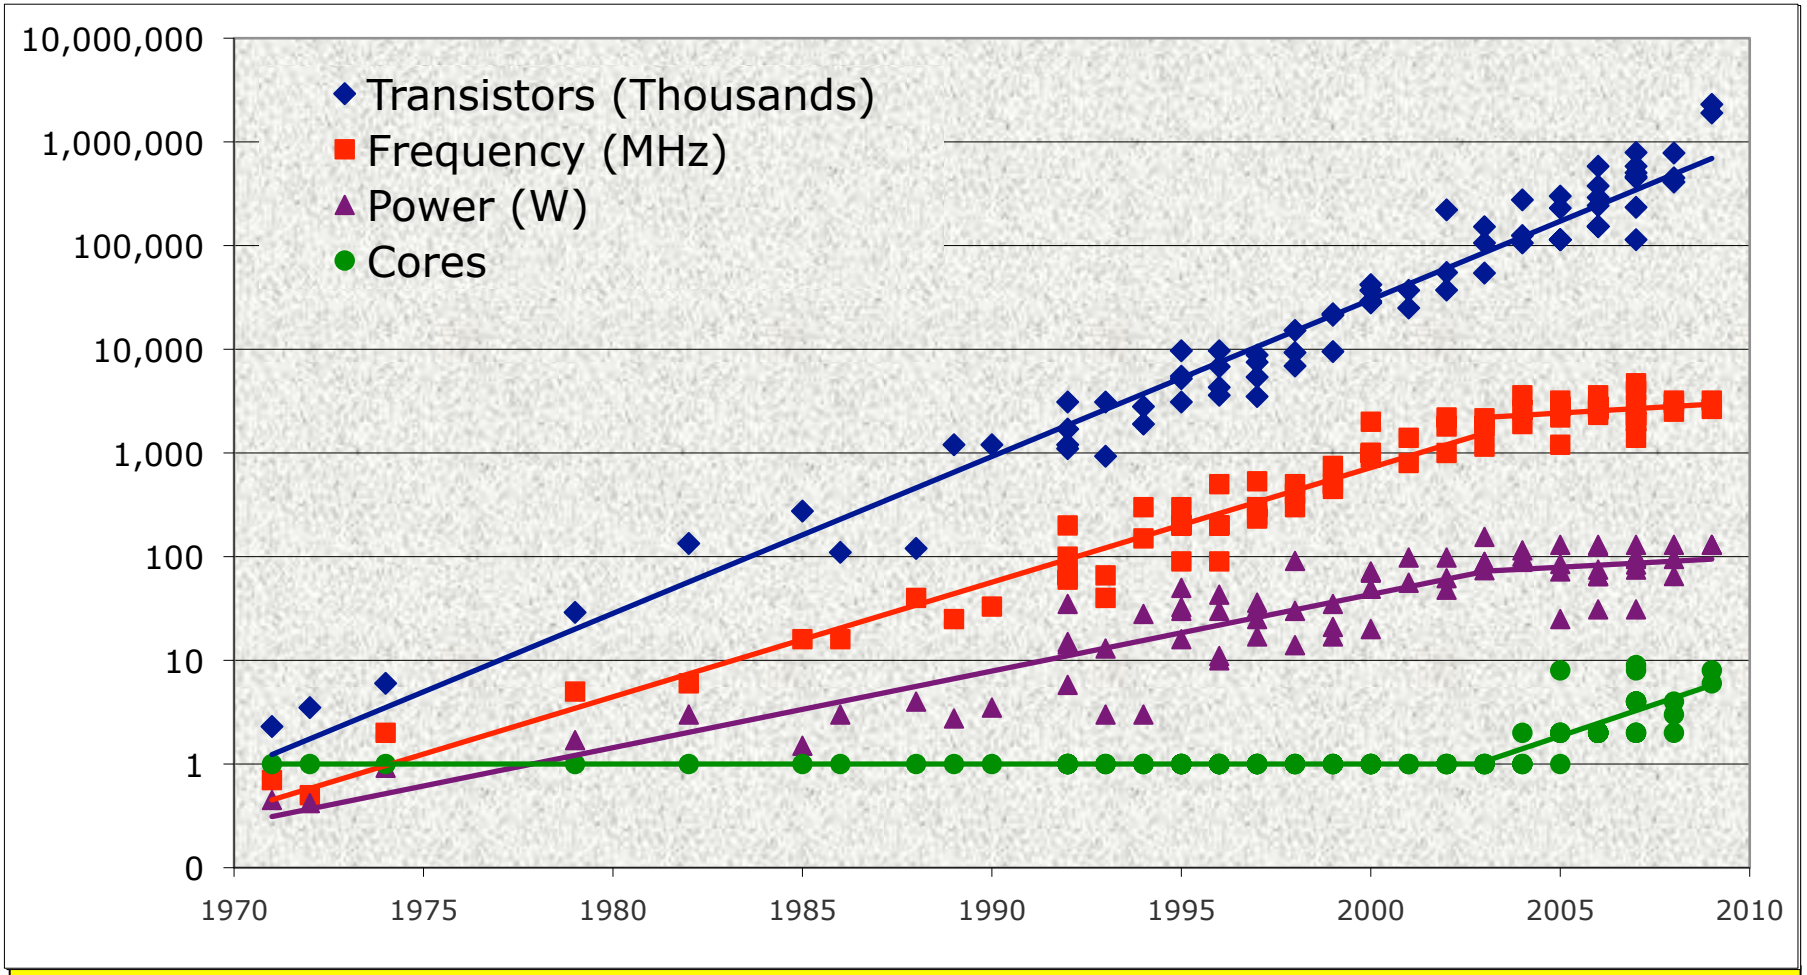

Moore's Law continues, but power limits performance growth. Parallelism is used instead.

Energy Efficient Computing is Key to Performance Growth

At $1M per MW, energy costs are substantial

- **1 petaflop in 2010 used 3 MW**
- **1 exaflop in 2018 would use 130 MW with "Moore's Law" scaling**

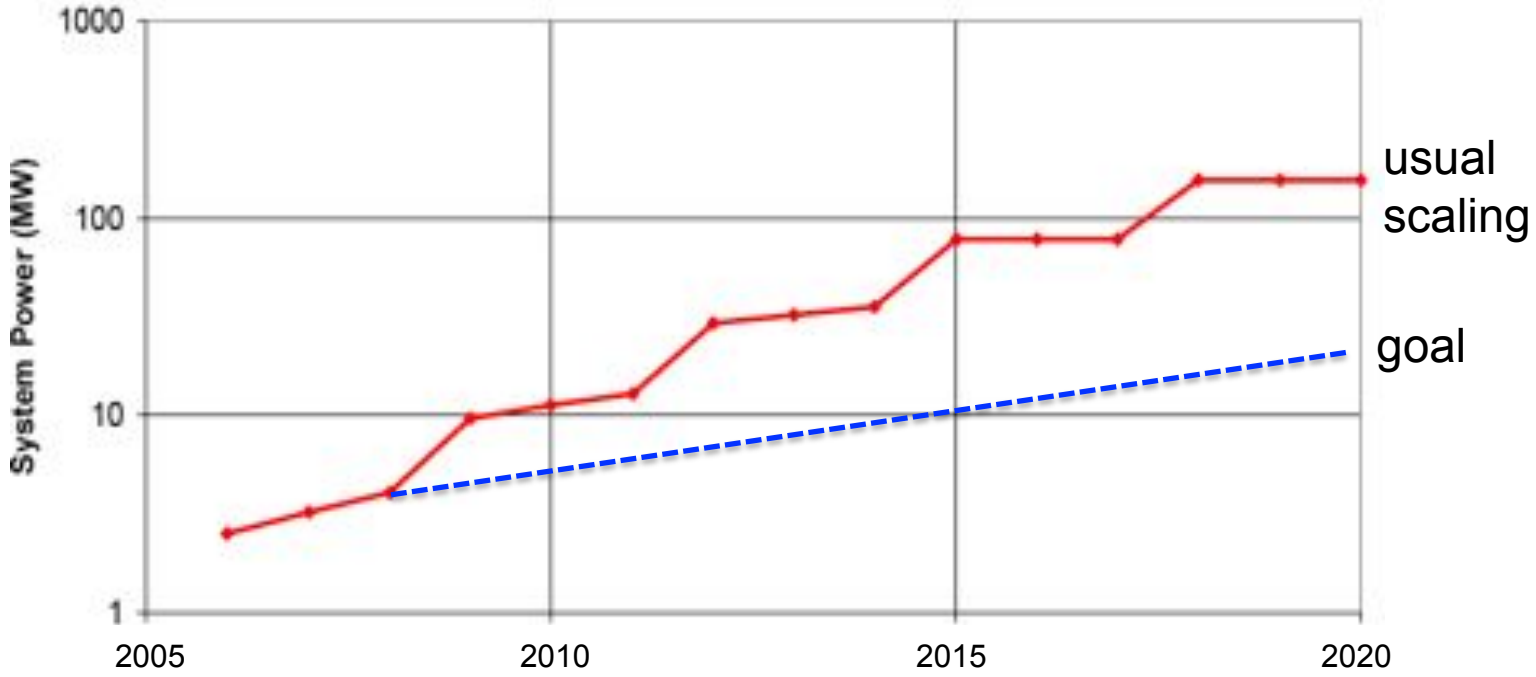

This problem doesn't change if we were to build 1000 1-Petaflop machines instead of 1 Exasflop machine. It affects every university department cluster and cloud data center.

New Processor Designs are Needed to Save Energy

Cell phone processor
(0.1 Watt, 4 Gflop/s)

Server processor
(100 Watts, 50 Gflop/s)

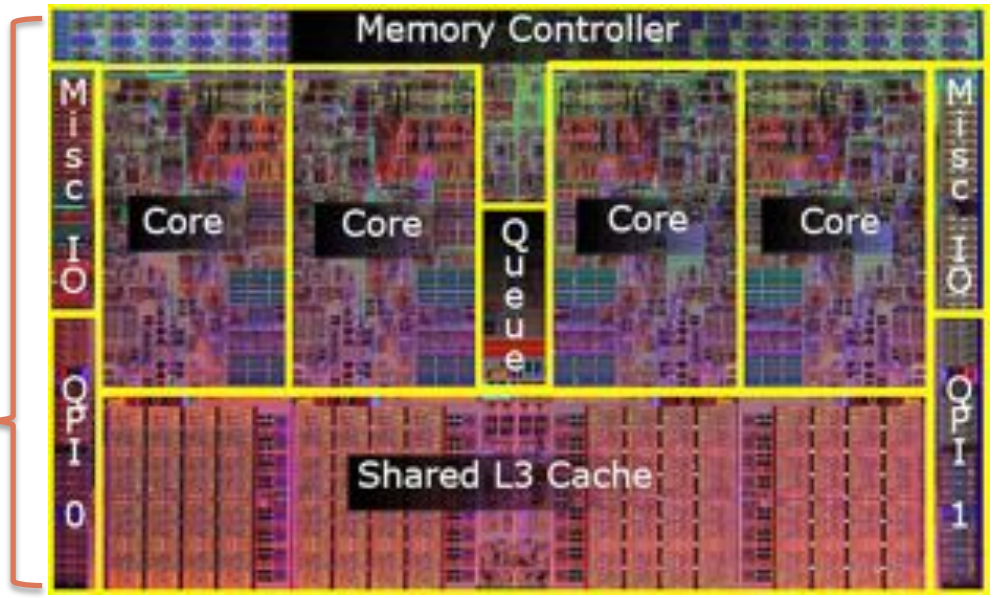

- **Server processors** designed for performance
- **Embedded and graphics processors** use simple low-power processors → good performance/Watt
- **New processor architecture and software** for future HPC systems

Need to Redo Software and Algorithms

- **Communication is expensive:**
  - 10-100x in time/energy
  - Even to memory
- **Computation is almost free!**

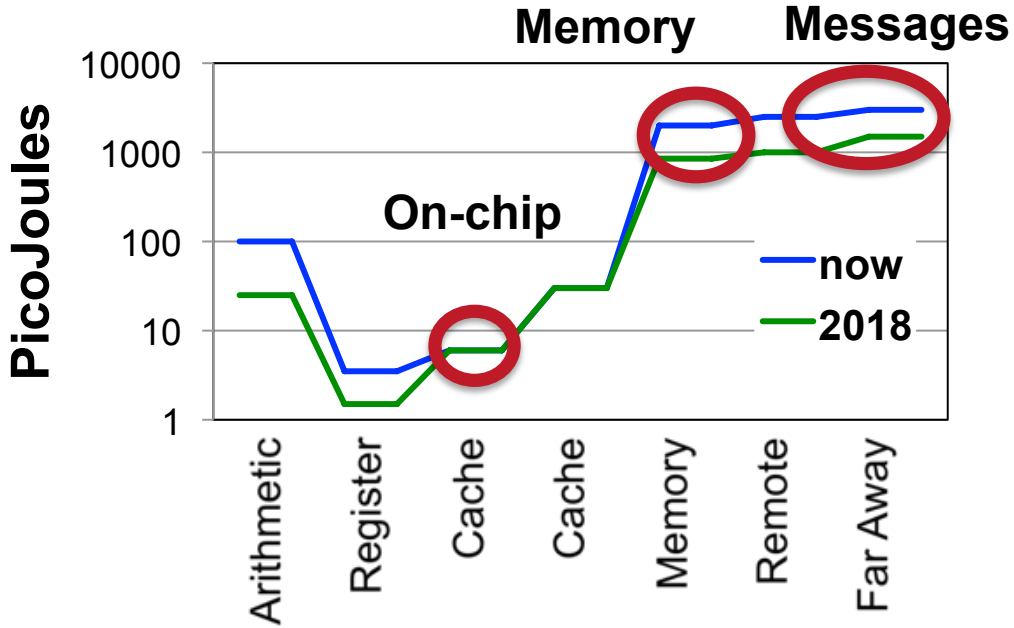

- **Software and algorithms**
  - **Software will control data movement to avoid waste**
  - **Algorithms should minimize data movement**
  - **Exascale science involves more complex interactions**

Stop Communicating!

Computing Crisis is Not Just about High Performance Computing

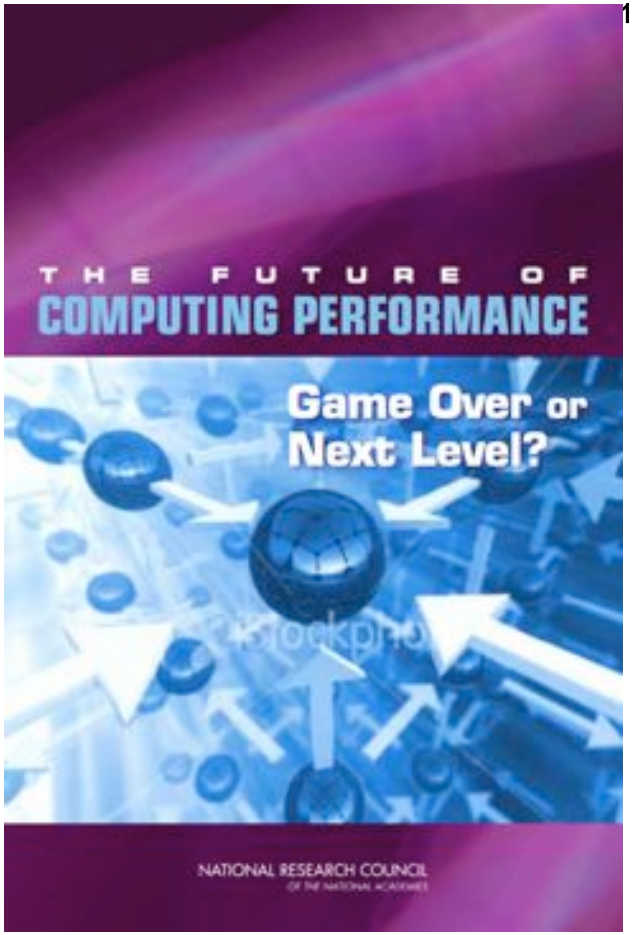

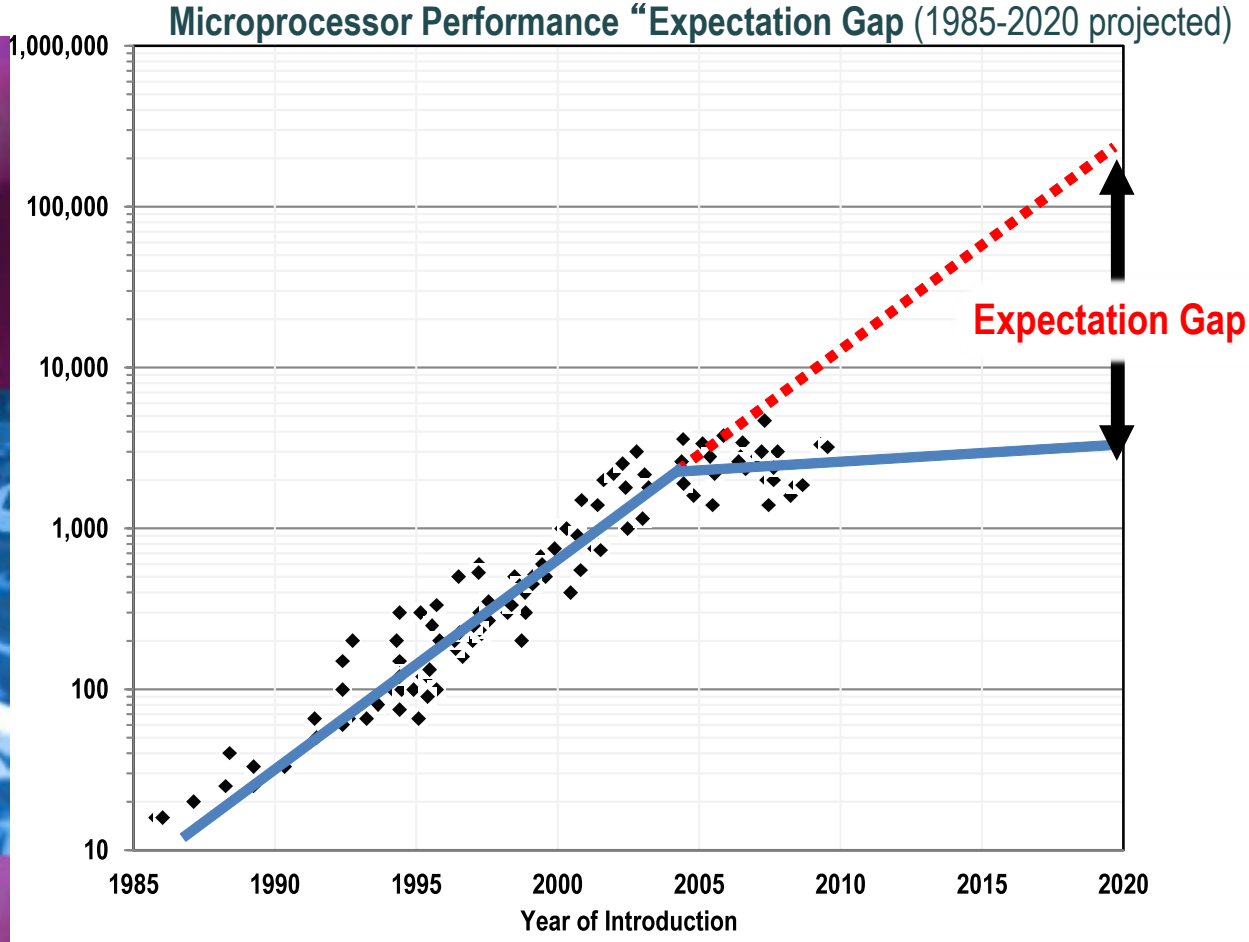

Industry is motivated; HPC can lead the field.