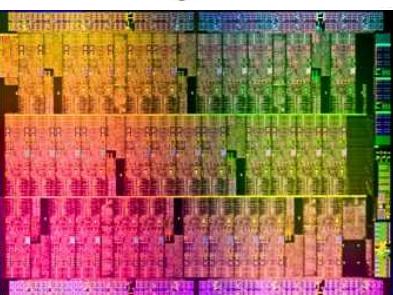


# Emerging HPC Systems and Next Generation FE Engineering Analysis Applications



Jim Ang, Ph.D.  
Acting Senior Manager,  
Extreme-scale Computing Group  
Sandia National Laboratories  
Albuquerque, NM

**Pacific Rim Workshop on  
Innovations in Civil Infrastructure Engineering  
National Taiwan University of Science and Technology  
Taipei, Taiwan  
January 9-11, 2013**



**Sandia  
National  
Laboratories**

*Exceptional  
service  
in the  
national  
interest*



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

# Sandia National Laboratories is the U.S. Department of Energy's Engineering Lab

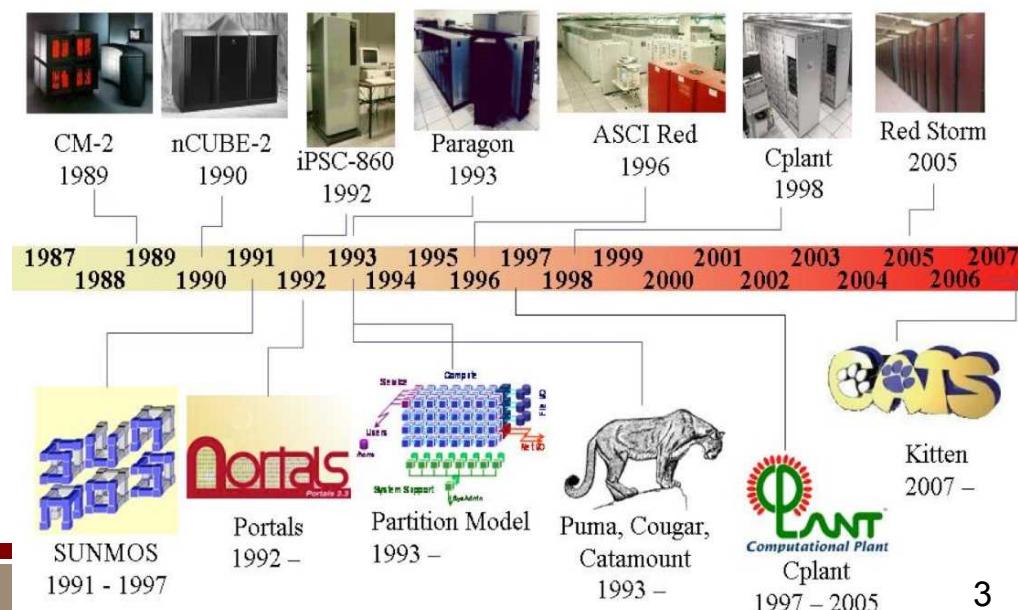
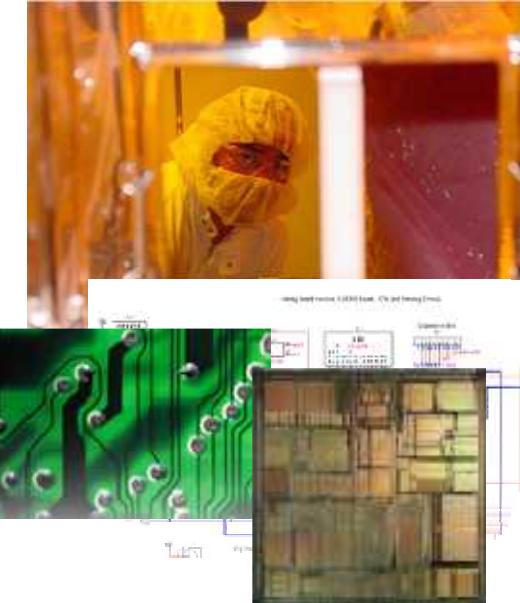
- **Sandia Uses HPC for Traditional Engineering Analysis Applications**
  - Structural mechanics
  - Structural dynamics
  - Shockwave physics
  - Computational fluid dynamics, combustion, turbulence, heat transfer
  - Circuit simulations, device physics, materials science
  - Design Optimization, uncertainty quantification

# Sandia has Unique Capabilities for HPC Development

- ⑩ Microsystems and Engineering Sciences Applications (MESA) Complex – sends more unique designs to IBM Trusted Foundry than any other institution
- ⑩ Largest concentration of HPC computer architects outside of industry
- ⑩ History of LWK operating system/runtime system software spans the MPP era
- ⑩ Helped establish the MPPs with seminal work on bypassing the limits of Amdahl's Law with Weak Scaling

The MPP Era is likely on its last legs

Our goal is to help define the next paradigm for HPC

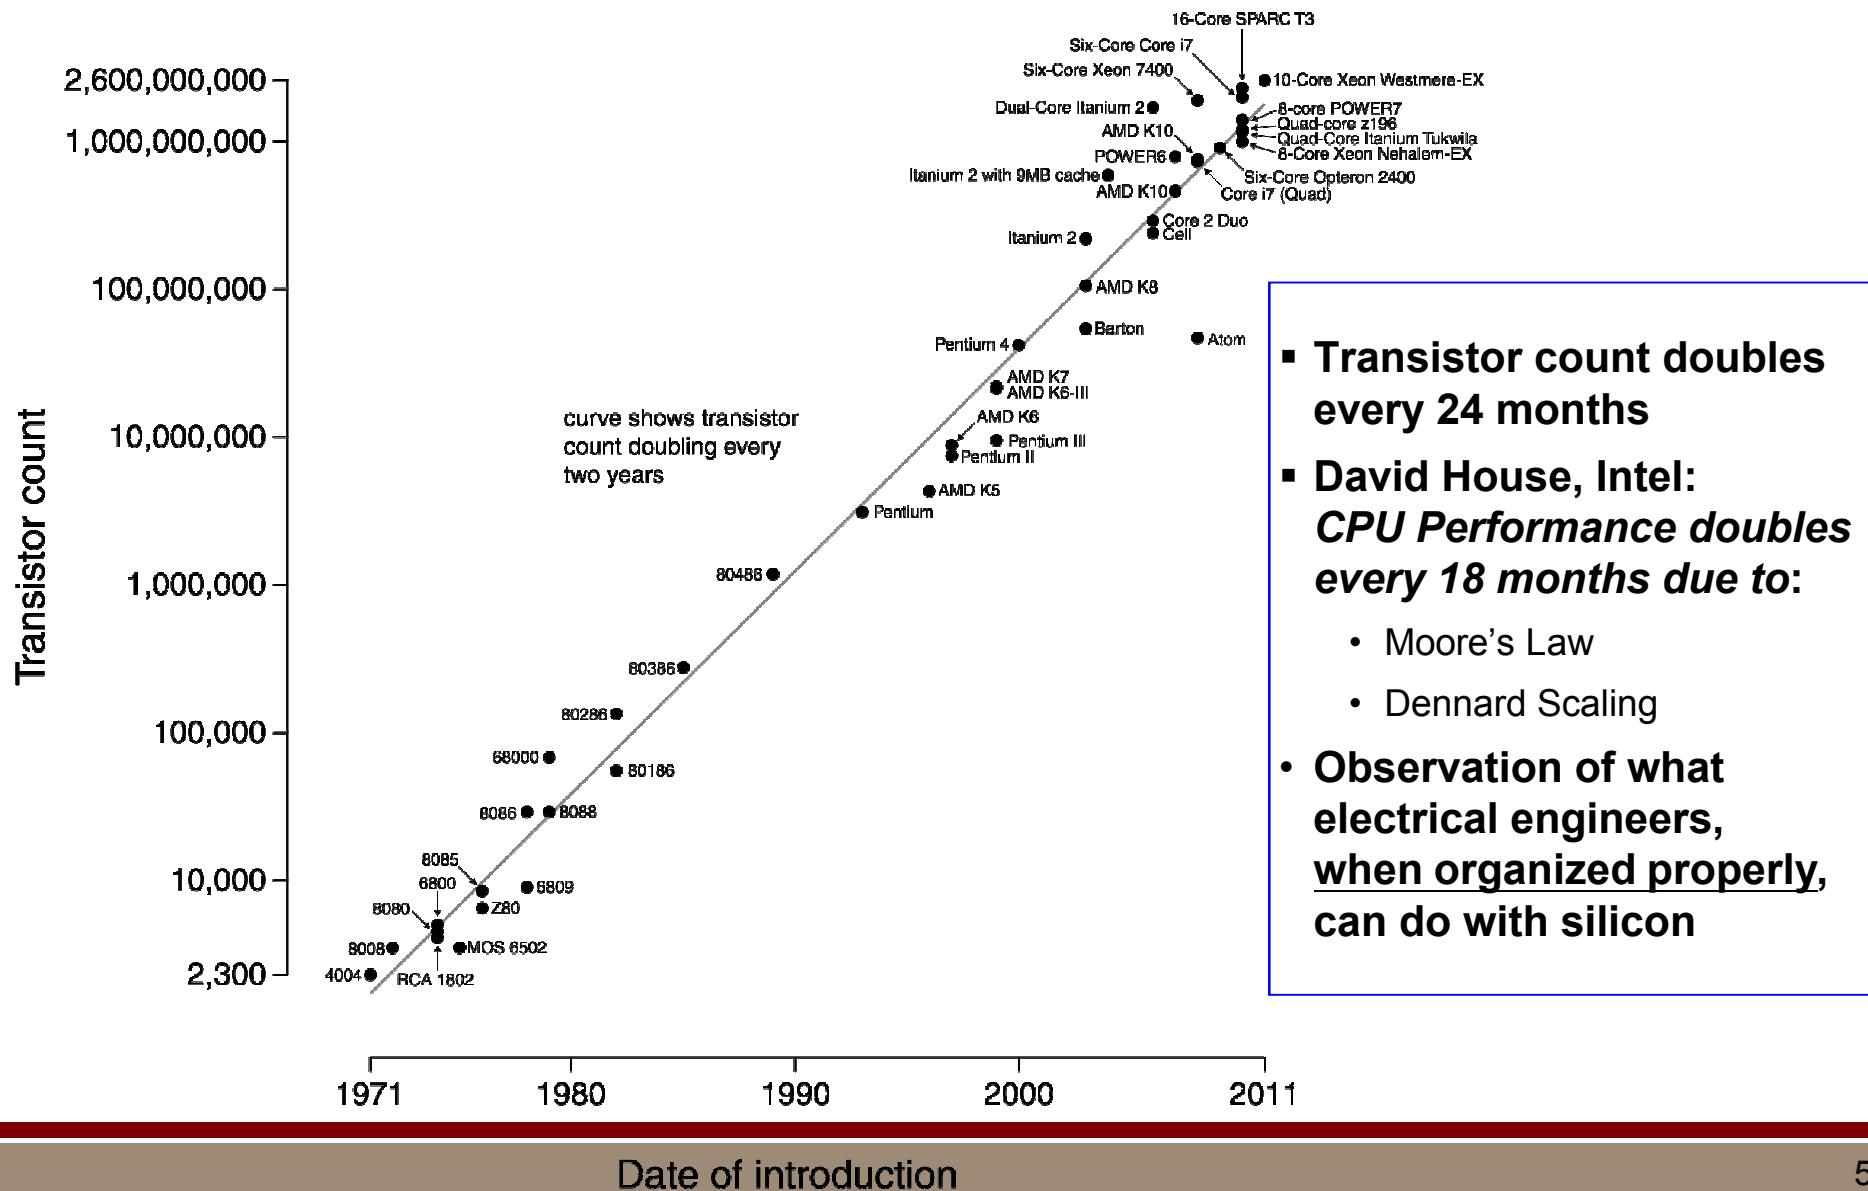


# Infrastructure versus Computers

- **Different Time and Length Scales:**
- **Infrastructure “Products”**
  - **Typical Design Lifetime is on the order of 40-50 years**
  - **Typical Sizes on the Order of  $10^1$  to  $10^4$  meters, but can be much larger, e.g. a Nation’s Power Grid**
- **Computer “Products”**
  - **Typical Design Lifetime is ~4 years for a CPU, ~2 years for a GP-GPU or Cell Phone**
  - **Typical Sizes on the Order of  $10^{-8}$  meters, as of 2012, feature sizes are 22nm =  $2.2 \times 10^{-8}$**

# Moore's Law 1971-2011:

## Growth in Transistor Count



# Moore's Law and the Loss of Dennard Scaling

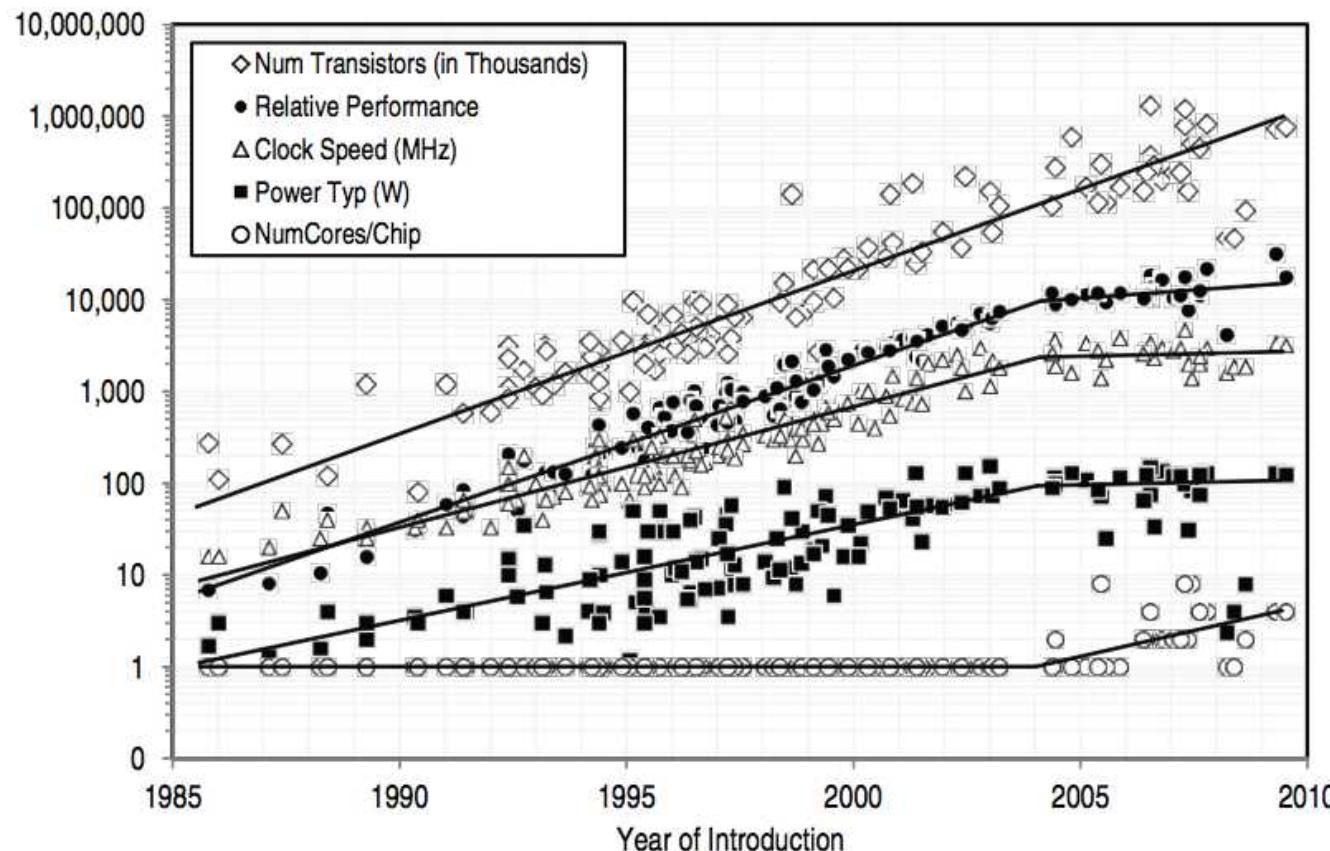


FIGURE 2.1 Transistors, frequency, power, performance, and cores over time (1985-2010). The vertical scale is logarithmic. Data curated by Mark Horowitz with input from Kunle Olukotun, Lance Hammond, Herb Sutter, Burton Smith, Chris Batten, and Krste Asanović.

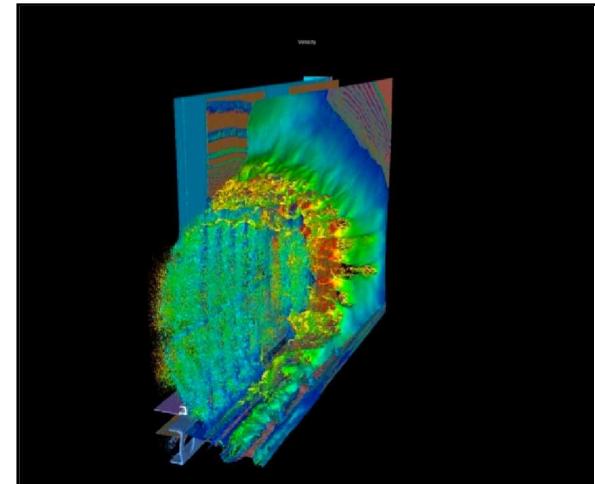
*The Future of Computing Performance: Game Over or Next Level,*  
Samuel Fuller and Lynette Millet, Eds., National Academy Press, 2011

# Classes of Computing Platforms

- **Workstations: SMP systems with ~4CPUs**
  - Typical
  - Cloud Computing
- **Clusters and MPP Integrated Systems**
  - Typical
  - Advanced Architectures

# Cielo, the ASC Program's Capability Computing Platform

Operational Time Frame	2011
Theoretical Peak Performance	1,374 TF
HPL (Linpack) Performance	1,110 TF using 142,272 cores
Cabinets	96
# Compute Nodes	8,944
# Compute Cores	143,104
Compute Processor	Dual AMD Opteron™ 6136 eight-core "Magny-Cours" Socket G34 @ 2.4 GHz
Compute Memory	286 TB DDR3 @ 1333 MHz
Compute Memory BW	763 TB/s
Service Nodes	272 AMD Opteron™ 2427 six-core "Istanbul" Socket F @ 2.2 GHz
External Login Nodes	Qty 4 Dell PowerEdge R815 Servers
User Disk Storage	7.6 PB User Available Capacity
Parallel File System	Cray DVS and Lustre
Parallel File System BW	~160 GB/s
High Speed Interconnect	Cray Gemini 3D Torus in a 16 x 12 x 24 (XYZ) Topology
Bi-section BW	6.57 x 4.38 x 4.38 (XYZ) TB/s
System Foot Print	~3,000 sq ft including Storage
Power Requirement	3,980 KW running HPL
Operating System	Cray Linux Environment



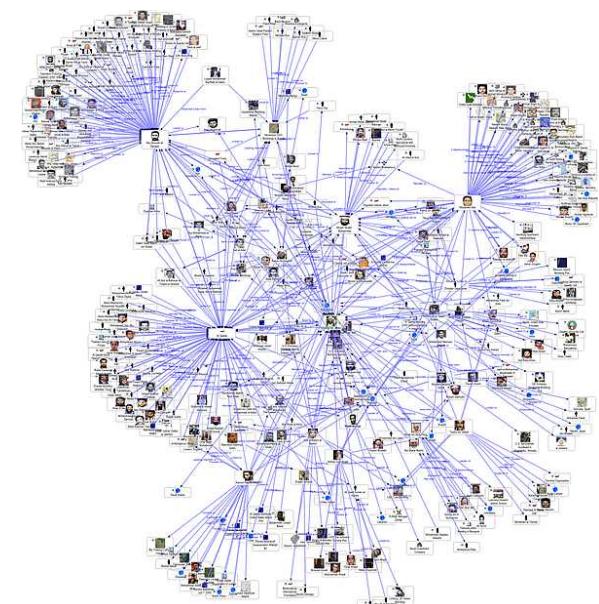
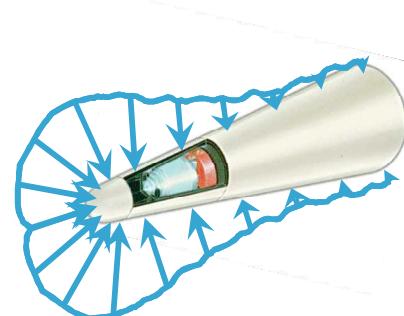
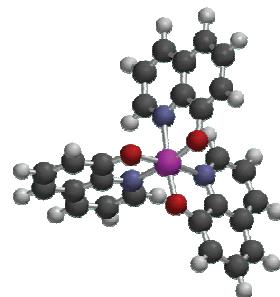
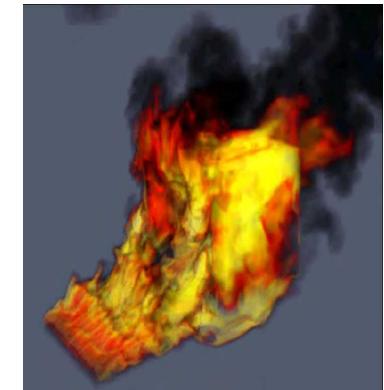
A 32,768 CTH simulation run on Cielo helps designers understand the response of structures under severe blast loading conditions

# Engineering Analysis Applications vs Materials Science Applications

- Differences in demands on and requirements from computer architectures
- Materials Science applications can benefit from special purpose computer architectures, e.g. MD-Grape, MyAnton
- Engineering Analysis applications, e.g. Finite Element Mechanics, CFD, Combustion, Nuclear Reactor design Applications, etc. stress data movement local - memory, and global - interconnect

# What is Different about Exascale?

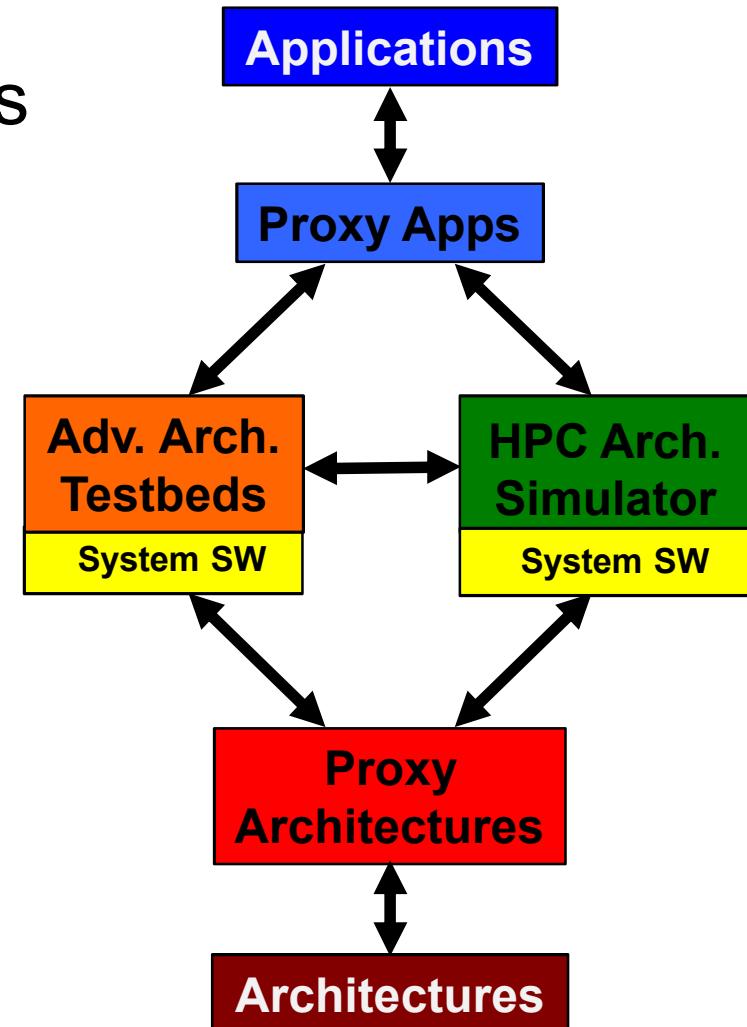
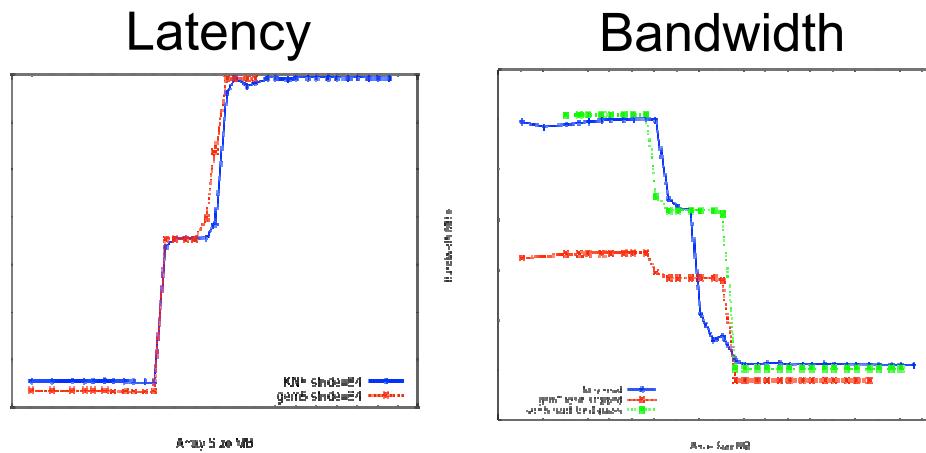
- Exascale is ~ a decade out in time
- Exascale is ~ 5 Moore's Law generations out
- With this many generations of evolution, Exascale hardware can be radically different from Petascale hardware
- Given this longer time-frame we have an opportunity for true Co-design
- In contrast, if our focus was our 2015 system, hardware is largely locked in



# Define and Develop the Co-design Methodology for HPC

## ■ Key Co-design Capabilities

- Simulators
- Proxy Applications
- Agile system software
- Testbeds
- Proxy Architectures



# Backup Slides

# Applicability of Heterogeneous Architectures to our Application Portfolio

- ⑩ HPC community focused on heterogeneous architectures with COTS processors and accelerators to solve the energy/performance challenges; e.g. China's *Tianhe-1a*, ORNL's *Titan*
- ⑩ Future heterogeneous architectures will be more tightly integrated and have unified memory systems, but limited memory capacity and bandwidth & latency performance
- ⑩ We need to understand what fraction of our NW workload can use these heterogeneous architectures
  - ⑩ Likely good match for PEM and single physics applications; e.g. LAMMPS
  - ⑩ Likely poor match to EPIC applications with multiphysics; e.g. Reentry, Combustion
  - ⑩ Likely poor match to our Cybersecurity, Graph applications

