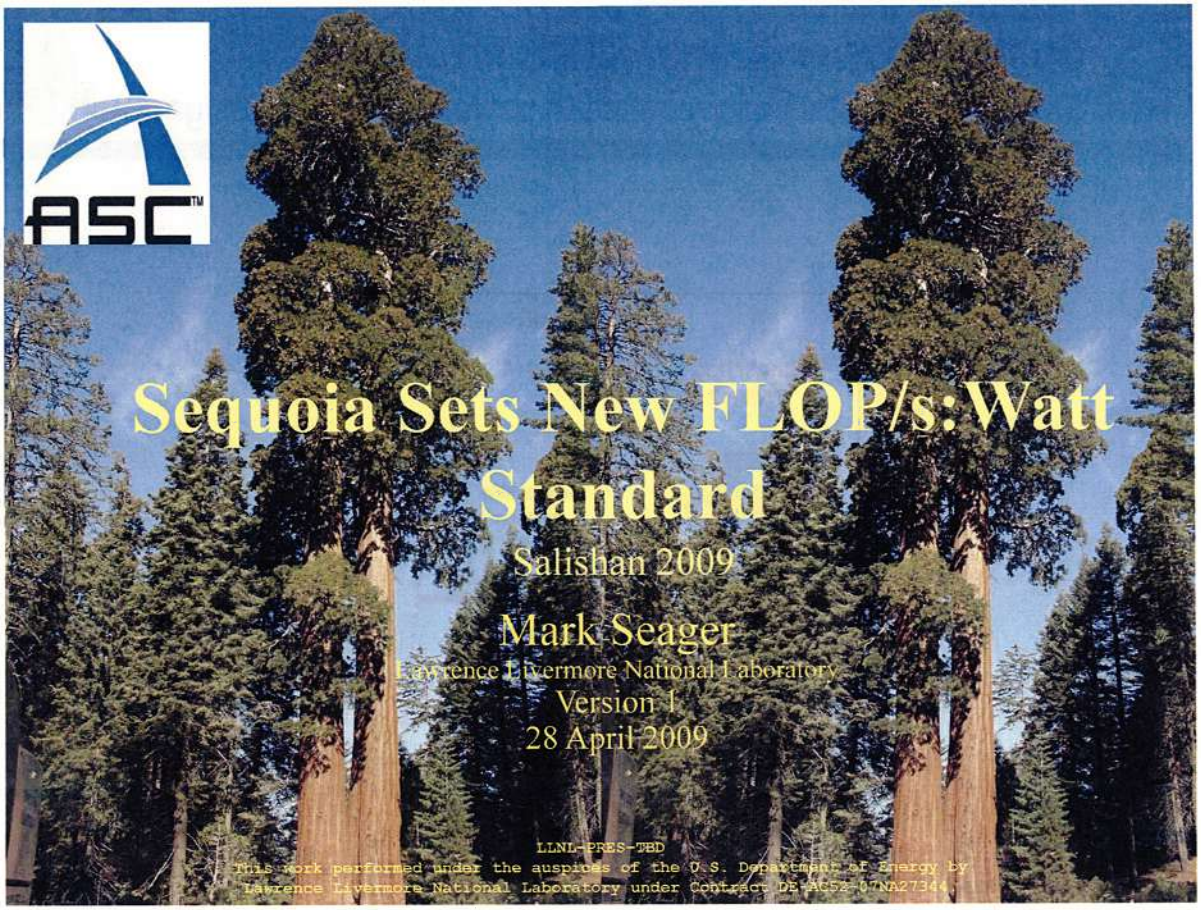

ASC Overview

- The case for a low power approach
- High level Sequoia programmatic drivers
- A high level overview of the Sequoia target architecture and multi-petascale applications strategy
- Applications work has already catalyzed new ways of thinking about parallelism and applications development model changes

System power is again THE problem

Single thread focus has resulted in power inefficient design

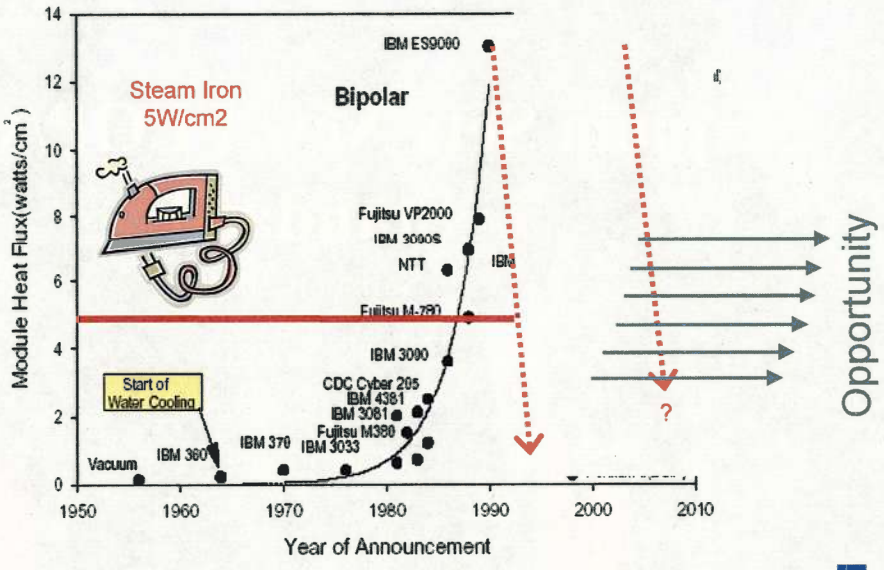

How does one know that classical CMOS scaling is really dead?

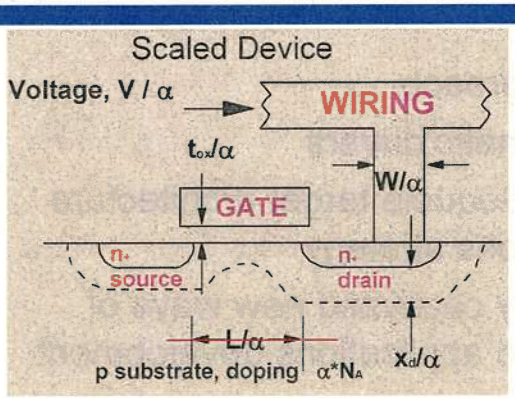

SCALING:		*RESULTS:*	
Voltage:	**V/α**	Higher Density:	~α^2
Oxide:	t_{ox}/α	Higher Speed:	~α
Wire width:	W/α	Power/ckt:	~$1/\alpha^2$
Gate width:	L/α	**Power Density:**	**~Constant**
Diffusion:	x_d/α		
Substrate:	$\alpha * N_A$		

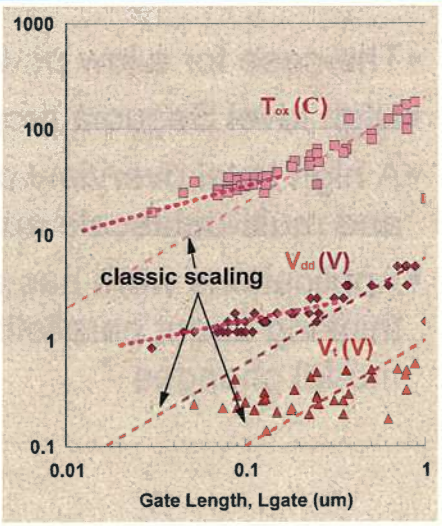

- Why deviate from "ideal" scaling
  - unacceptable gate leakage/reliability
  - additional performance at higher voltages
- What is the consequence of this deviation?
  - **a dramatic rise in power density**

In the face of adverse changes in the base silicon technology, BlueGene dramatically improves power efficiency

	Rack TFLOP /s	Width (ft)	Depth (ft)	Height (ft)	Density GFLOP/s ft^3	Linpack Perf / Watt (MFlops/Watt)
Earth Simulator	0.12	3.2	4.5	6.5	1.33	2.7
Purple	0.73	4	6	7	4.34	15.7
Dell Xanadu Nehalem	5.5	2	4	6	115	181.
Blue Gene/L	5.7	3	3	7	91	208.
Blue Gene/P	13.9	4	3	7	166	371.
Sequoia	209.7	4	3	7	2,497	2000.§

General Approach to Low Power

- High level of integration
  - fewer parts each with significant power overhead, also impacts RAS
- Focus on simplicity and address power-performance trade-offs at all levels
  - Simple power efficient processing core.
  - SIMD floating point to exploit data reuse optimally from a power perspective
  - Utilize concurrency to avoid climbing non-linear power-performance curve.
  - Large low power on chip cache based on eDRAM
  - Glueless network design to reduce chip count and increase scalability (no layer-like limits , predictable, repeatable performance
  - On-board memory controller, direct attached commodity lowest power memory
  - Power distribution focus on reducing system droop allowing for lower supply voltage.
  - Water cooled design (Sequoia) holds temperatures lower (reduces CMOS leakage currents), increases compute density (allows for shorter signaling distances requiring lower power) and improves supply efficiency.

IBM Research PDSOI optimization results indicate lower power approaches provide better power efficiency

Optimizing process technology knobs for maximum performance for each core

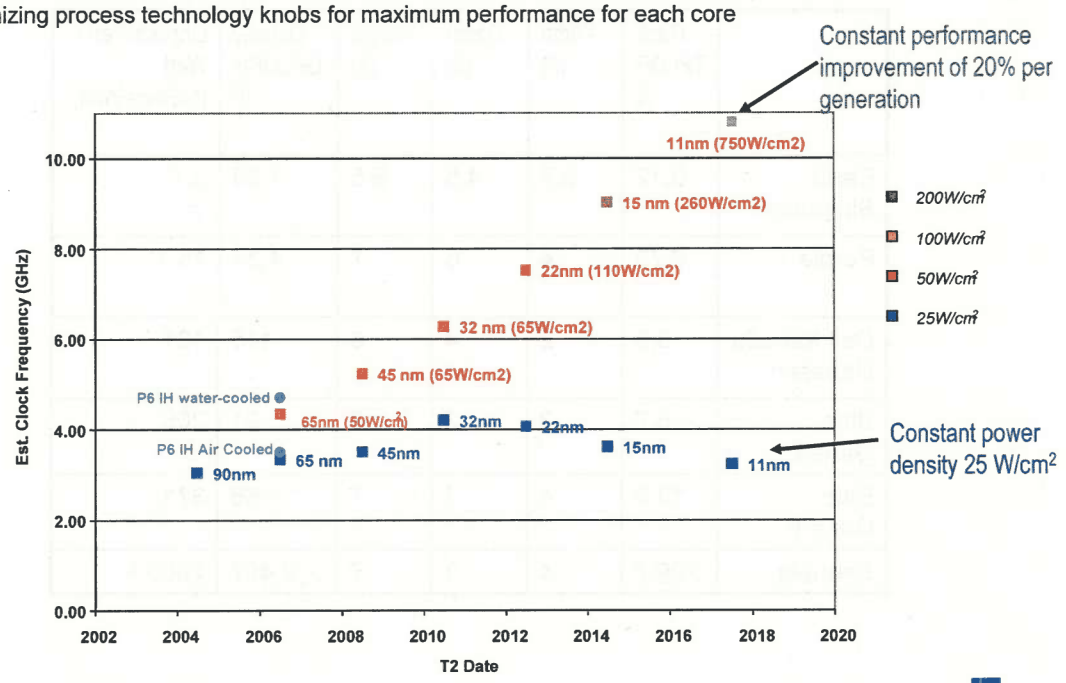

D. Frank, C. Tyberg
Source: IBM Research

High power processors are even less attractive for HPC systems in the future...

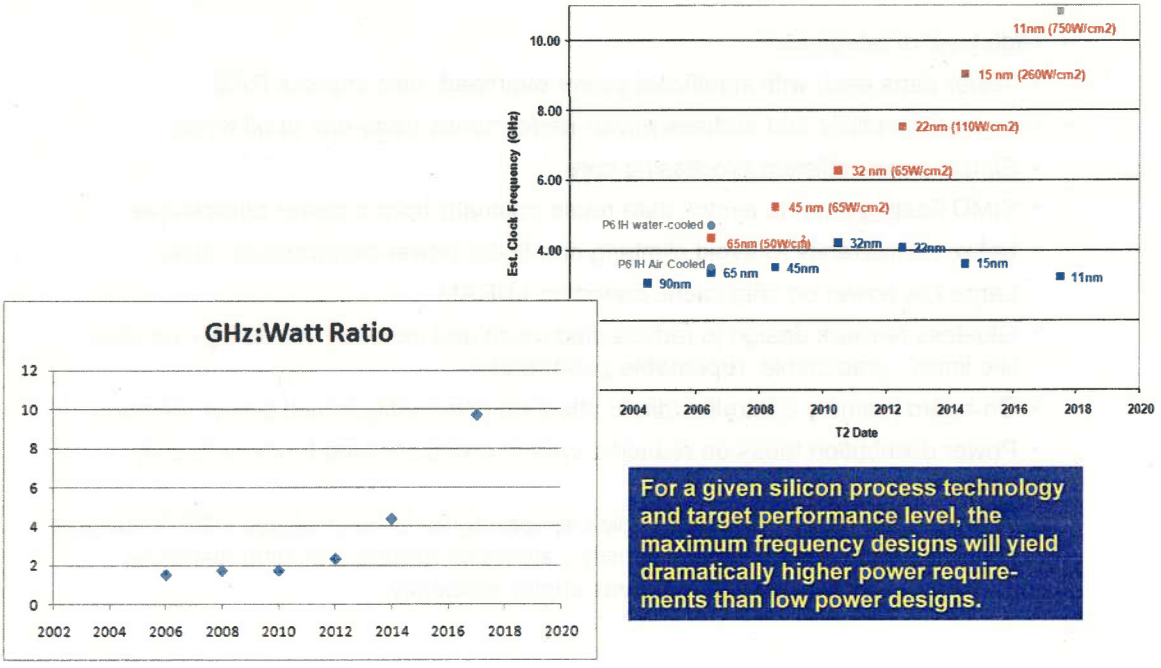

For a given silicon process technology and target performance level, the maximum frequency designs will yield dramatically higher power requirements than low power designs.

Source: IBM Research

Sequoia will be a key simulation tool for keystones and uncertainty quantification for stockpile stewardship

- ***ASC Strategy*** **and** ***ASC Roadmap*** **provide a vision for and keystones leading to "predictive simulation" or prediction with quantified uncertainties**
- ***Thermonuclear Burn Initiative, National Boost Initiative*** **and** ***Predictive Capability Framework*** **represent Stockpile Stewardship Program (SSP) planning to coordinate on the key issues impeding predictive simulation**
- **Sequoia is intended to address requirements coming from this planning in the period between 2012 - 2017, focusing on UQ and materials science, related to boost and certification**
- **To demonstrate it can meet these objectives, Sequoia will:**
  1. **Achieve 12X-24X Purple throughput for integrated weapons calculations related to Uncertainty Quantification (stretch goal >> 24X)**
  2. **Achieve 20X BG/L (stretch goal 50X) on a science materials effort**
  3. **Single RFP mandatory was Peak + Sustained ≥ 40**

Predictive simulation requires hero runs, but also large ensemble of calculations

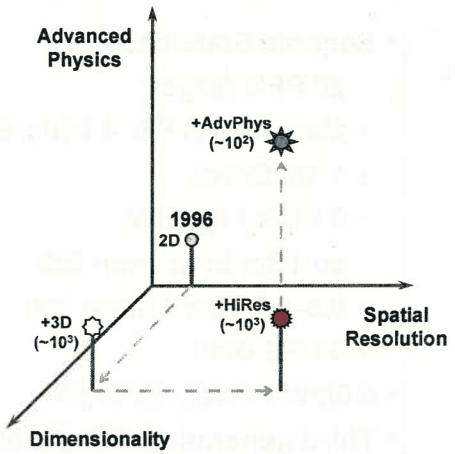

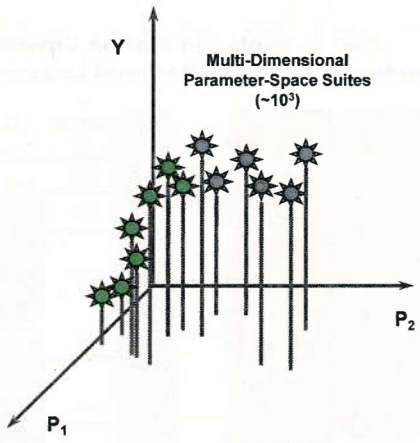

A path to a single 3D high-resolution high-fidelity calculation

Assessing the uncertainty in our predictions requires a suite of 3D high-resolution high-fidelity calculations

Sequoia will initiate the transition to 3D UQ suites with advanced models for energy balance and boost

Sequoia Timeline Delivers Petascale Resources to the Program

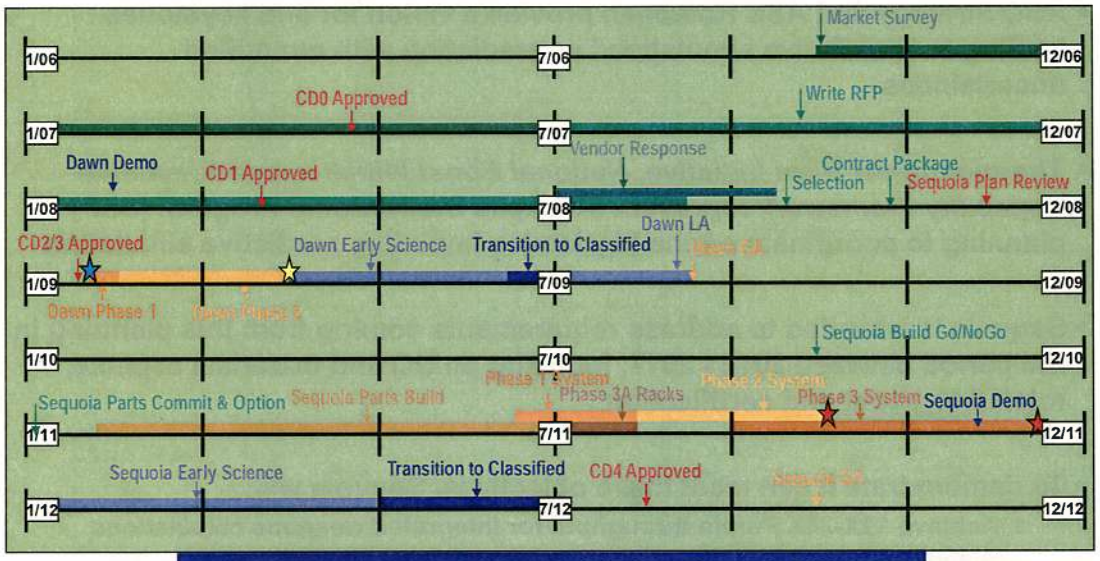

Sequoia Five Years Planned Lifetime Through CY17

★ Sequoia contract award ☆ Dawn system acceptance ★ Sequoia phase 2 & final system acceptance

Sequoia Hierarchal Hardware Architecture in Integrated Simulation Environment

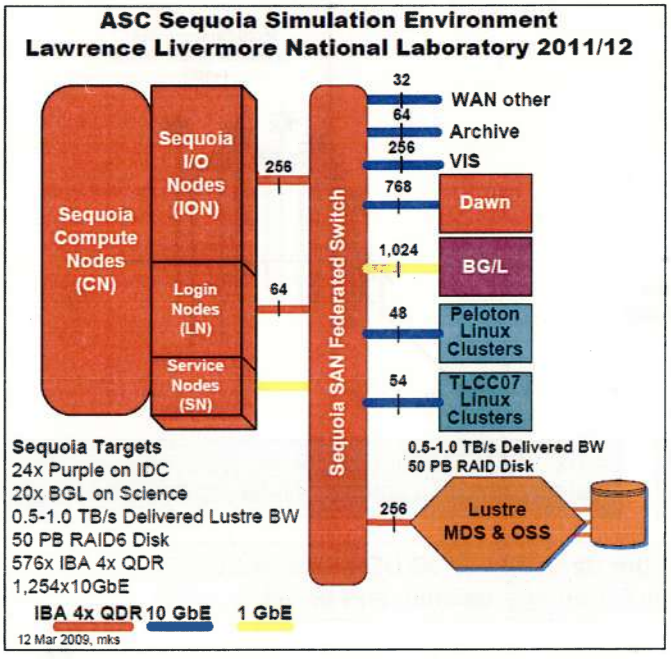

- **Sequoia Statistics**
  - **20 PF/s target**
  - **Memory 1.6 PB, 4 PB/s BW**
  - **1.5M Cores**
  - **3 PB/s Link BW**
  - **60 TB/s bi-section BW**
  - **0.5-1.0 TB/s Lustre BW**
  - **50 PB Disk**
- **6.0MW Power, 3,500 ft²**
- **Third generation IBM BlueGene**
- **Challenges**
  - **Hardware Scalability**
  - **Software Scalability**
  - **Applications Scalability**