

---

**Computing in the Next 20 Years:  
Talk Driven by a Findings Digest of  
Frontiers of Extreme Computing 2007  
Zettaflops Workshop**

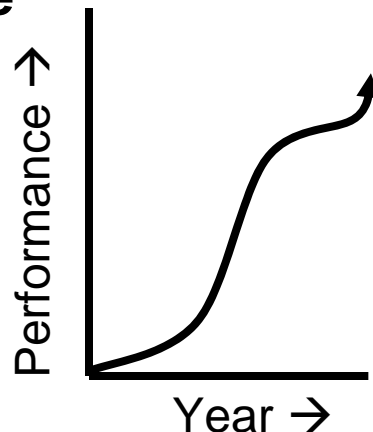
**Erik P. DeBenedictis**



# Introduction & Overview Of Ben's Questions

---

- **Phase A: Years 1 to 12±**
  - ITRS roadmap goes 14 years to 2022; Zettaflops finds it valid for only 12± years
  - Exponential growth in some parameters, but increasingly tight design space



- **Phase B: Years 12± to 20**
  - This will begin by a period optimization of a mature underlying technology
  - Maybe a new physical device will arrive towards the end
    - Quantum computing
    - Reversible logic



# Process

---

- In 2007 we held a 3 ½ day “Zettaflops” workshop  
<http://www.zettaflops.org>
- 50 people represented a cross section of Government, industry, universities, labs
- Four working groups
  - Applications
  - Software
  - Architecture
  - Technologies
- To try and separate my bias from consensus, I organized the working group notes by cross-cutting issue
- With some imagination, this created a banded matrix

# Banded Structure of Interests

	Applications	Software	Architecture	Technologies
Legacy code	Applications working group only studied existing code (legacy), which is synonymous with an MPI software base	Legacy applications must work although there is an opportunity to move beyond		[Link to informatics applications]
Parallelism	Billion-way parallelism is indicated	Billion-way parallelism is the significant challenge	Concurrency a top challenge; design points indicate half-billion-way parallelism for exaflops	Their viewgraph on clock rate "747 bubble overshoot"
Memory and Storage	At exaflops: 400 PB main memory, 4-40 EB storage	Global address space believed to be indicated and must be supported in systems software; the deeper hierarchy requires new tools.	3D packaging merges processor and memory	Memory options are available, including architecturally distinct ones such as IBM storage class memory; public sector needs seem outstrip private sector developments (i. e. terabyte nano memory on a chip per Bona's talk)
Interconnect	Bisection bandwidth .5 - 1 EB/s	(Curiously, there was no reference to interconnect in the outbrief notes.)		Optical interconnect required for exaflops and beyond, commodity market not believed to be headed to necessary interconnect solution fast enough for public sector needs
Reliability	Error checking and notification; checkpoint/restart ("other issue")	New issue coping with large systems; need fault-oblivious applications; dynamic reconfiguration	Critical to offset aggregate error rates, Error Detection and Correction (EDAC) in hw, communications, no single point of interrupt left uncovered, checkpoint/rollback, Algorithm-based Fault tolerance (ABFT) research needed (See slide 15)	
Numerical precision			Dealing with variable numerical precision needs	
Programming model	There is a 20 year lead time on new programming languages ("other issue")	New programming model required for billion-way parallelism; global address space desirable; exploit new architectural features; transactional memory	New model with better locality "ease of programming versus efficiency", SVP microthreads research, active objects, use compilation support instead of hardware support, simpler multi-core chip design	Some nano-memory options might be sufficient to drive informatics applications with a different architecture
Community process		Testbed and simulator; VMs to test OS at scale; leverage community efforts to improve parallel programming for the masses	Two point designs proposed; simulation project	Monitor industry. Also, estimate performance of a system based on fully-mature CMOS and optical interconnect
Power		Needs to be managed as a precious resource; systems software needs to expose and manage power resources; minimize data movement	Problem -- mitigate with clock rate, asynchronous logic, locality, reversible logic, variable precision, less speculation	ITRS predictions of power flatlining are valid, can be circumvented only by a new device
Accelerators		"Heterogeneous computing is quickly approaching"	Need a new microarchitecture with lower control overhead relative to computation; Vector and acceleration technologies; bio inspired	Special function units are a way to improve throughput with current logic; possibly leverage commodity processors with public-sector special function units
A new device technology (reversible, quantum, etc.)			New device needed to reach last 1 or two orders of magnitude before Zettaflops; do we know how to cope with reversible logic?	New device technology is a valid issue



# Cross-Cutting Issues

---

- In case you couldn't read the last Viewgraph, the cross cutting issues are on the right →
- Legacy code
- Parallelism
- Memory and Storage
- Interconnect
- Reliability
- Numerical precision
- Programming model
- Community process
- Power
- Accelerators
- A new device technology (reversible, quantum, etc.)



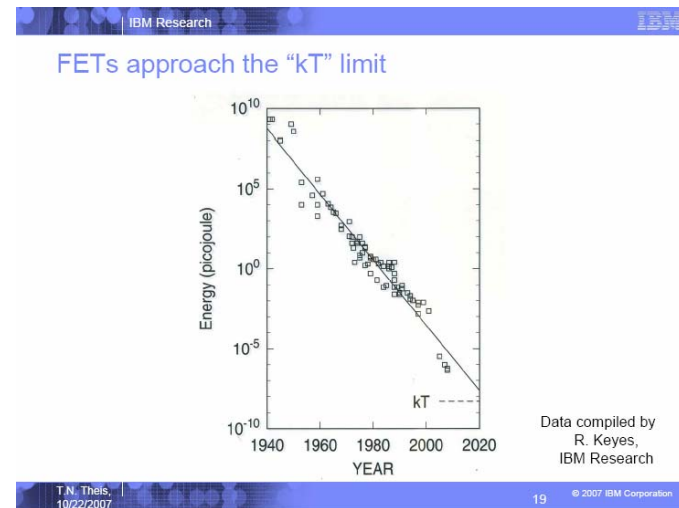
# Findings/Agenda

---

- **Transition of Eras:  $kT$  limit**
- **Flat Clock Rate**
- **Fascinating Memory Story**
- **On-Chip Optical Interconnect**
- **Architecture**
- **A New Physical Device**

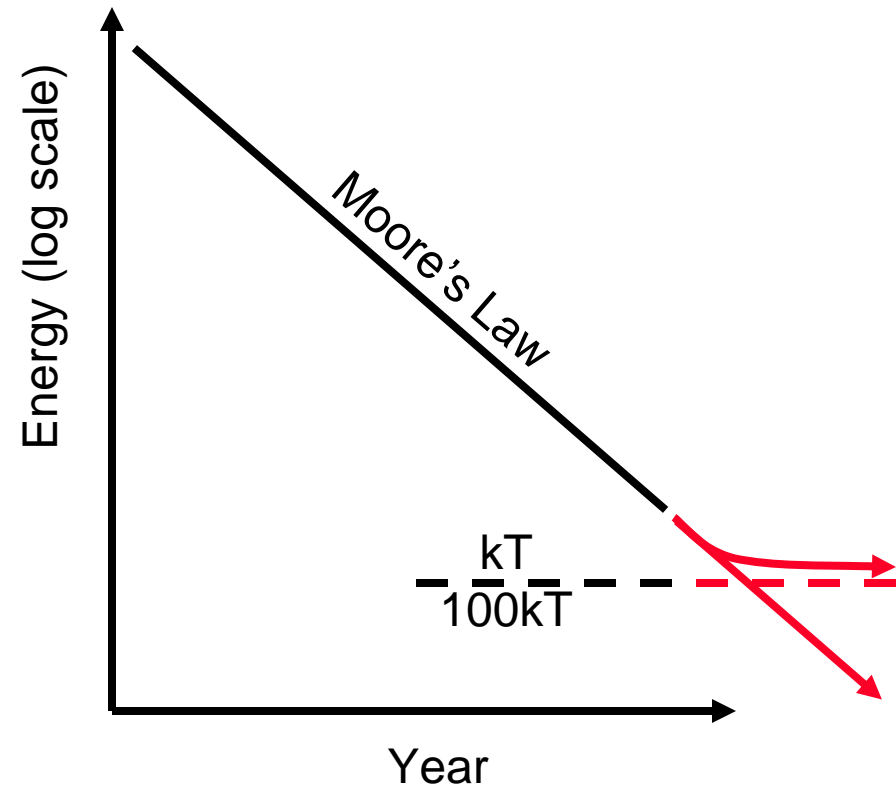
# Finding 1: kT Limit Moderates Optimism for Perpetual Exponential Growth


- In past workshops, many participants considered Moore's Law a fact of physics that can overrule finite atom size, Landauer's limit, etc.
- Zettaflops 2007 concludes that Exaflops and Zettaflops are qualitatively different



# The Timeless kT Limit Graph

- For decades, we've seen graphs like this one →
- Until recently, computer users regarded the kT barrier as an obstacle to be overcome, like soft errors (64K DRAM) or leakage current (High K dielectric)
- kT barrier now seen as invalidating most of computer engineering





# Background & Consequence

---

- **What is the physics issue?**
  - **Signals representing 1's and 0's have to have more energy than that due to random thermal motion or they spontaneously switch**
  - **AND and OR gates that are the basis of today's computers have two input signals and one output, destroying one**
- **Workshop participants divided the future into two segments**
  - **Petaflops to Exaflops**
    - **A tough road**
  - **Exaflops to Zettaflops**
    - **Reliant on a new “post transistor” device**
- **What does this mean for general computing?**
  - **Dividing line  $\frac{3}{4}$  way across ITRS (~2020)**

# ITRS Roadmap Predictions 2007-2022

Table PIDS2a High-performance Logic Technology Requirements—Near-term Years

High-performance Logic Technology Requirements—Long-term Years

Grey cells delineate one of two time periods: either before initial production ramp has started for ultra-thin body fully depleted (UTB FD) SOI or double-gate (DG) MOSFETs, or beyond when planar bulk or UTB FD MOSFETs have reached the limits of practical scaling (see the text and the table notes for further discussion).

Year of Production	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022
MPU/ASIC Metal 1 (M1) ½ Pitch (nm)(contacted)	68	59	52	45	40	36	32	28	25	22	20	18	16	14	13	11
MPU Physical Gate Length (nm)	25	22	20	18	16	14	13	11	10	9	8	7	6.3	5.6	5.0	4.5
Effective Ballistic Enhancement Factor, Kbal [12]										Kbal [12]						
Extended Planar Bulk	1	1	1	1	1	1										
UTB FD				1.05	1.1	1.16	1.2	1.24	1.28							
DG					1.17	1.25	1.31	1.37	1.53	1.67	1.87	1.99	1.97	2.11	2.11	2.11
R <sub>sd</sub> : Effective Parasitic series source/drain resistance [13]										Parasitic series source/drain resistance [13]						
Extended Planar Bulk (Ω-μm)	200	200	200	180	180	180										
UTB FD (Ω-μm)				180	180	180	170	160	160							
DG (Ω-μm)					180	180	170	160	160	155	150	145	145	145	135	135
C <sub>g,ideal</sub> : Ideal NMOS Device Gate Capacitance [14]										Ideal NMOS Device Gate Capacitance [14]						
Extended Planar Bulk (F/μm)	4.70E-16	6.30E-16	6.63E-16	6.70E-16	6.71E-16	6.33E-16										
UTB FD (F/μm)				5.65E-16	5.52E-16	5.08E-16	4.98E-16	4.22E-16	3.83E-16							
DG (F/μm)					4.60E-16	4.39E-16	4.48E-16	3.80E-16	3.45E-16	3.27E-16	2.91E-16	2.68E-16	2.30E-16	2.11E-16	1.92E-16	1.72E-16
C <sub>g,total</sub> : Total gate capacitance for calculation of CV/I [15]										Total gate capacitance for calculation of CV/I [15]						
Extended Planar Bulk (F/μm)	7.10E-16	8.40E-16	8.43E-16	8.40E-16	8.36E-16	7.93E-16										
UTB FD (F/μm)				8.08E-16	7.22E-16	6.78E-16	6.58E-16	5.82E-16	5.43E-16							
DG (F/μm)					6.50E-16	6.29E-16	6.28E-16	5.59E-16	5.25E-16	5.07E-16	4.81E-16	4.58E-16	4.10E-16	3.91E-16	3.62E-16	3.42E-16
τ = CV/I: NMOSFET intrinsic delay (ps) [16]										NMOSFET intrinsic delay (ps) [16]						
Extended Planar Bulk (ps)	0.64	0.55	0.51	0.46	0.43	0.4										
UTB FD (ps)				0.41	0.36	0.31	0.28	0.23	0.21							
DG (ps)					0.34	0.29	0.26	0.21	0.18	0.15	0.13	0.11	0.1	0.09	0.08	0.08
1/τ: NMOSFET intrinsic switching speed (GHz) [17]										NMOSFET intrinsic switching speed (GHz) [17]						
Extended Planar Bulk (GHz)	1563	1818	1961	2174	2326	2500										
UTB FD (GHz)				2439	2778	3226	3571	4348	4762							
DG (GHz)					2941	3448	3846	4762	5556	6667	7692	9091	1.00E+04	1.11E+04	1.25E+04	1.25E+04

Manufacturable solutions exist, and are being optimized  
 Manufacturable solutions are known  
 Interim solutions are known  
 Manufacturable solutions are NOT known



Manufacturable solutions exist, and are being optimized  
 Manufacturable solutions are known  
 Interim solutions are known  
 Manufacturable solutions are NOT known





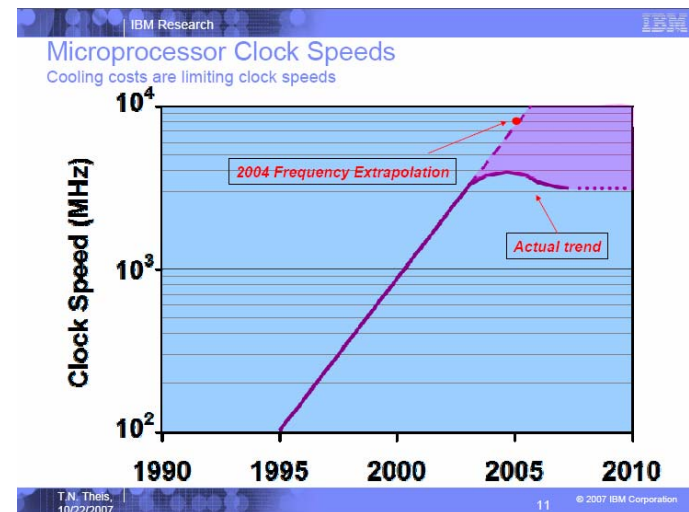
# Findings/Agenda

---

- **Transition of Eras:  $kT$  limit**
- **Flat Clock Rate**
- **Fascinating Memory Story**
- **On-Chip Optical Interconnect**
- **Architecture**
- **A New Physical Device**

## Finding 2: Clock Rate Flat Lined

- Clock rate flat lined a couple years ago, as vendors put excess resources into multiple cores
- This is a historical fact and evident to everybody, so there is little reason to comment on the cause
- However, it has profound architectural consequences (later slide)






# Findings/Agenda

---

- Transition of Eras:  $kT$  limit
- Flat Clock Rate
- Fascinating Memory Story
- On-Chip Optical Interconnect
- Architecture
- A New Physical Device

# Finding 3: The Memory and Storage Stories are Game Changing

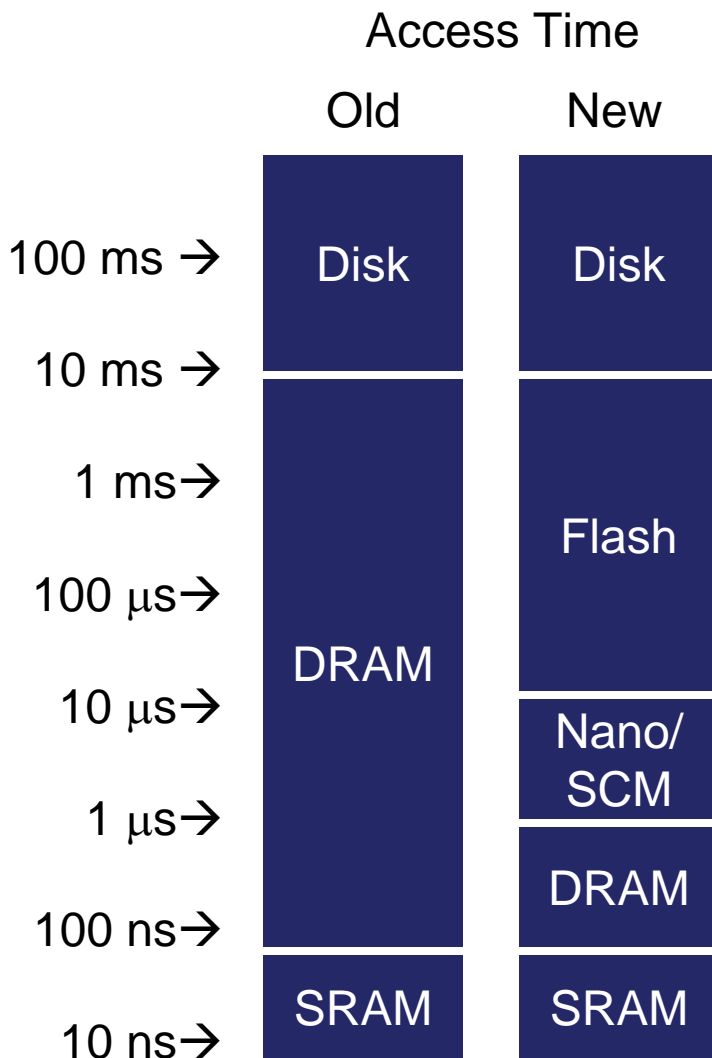
- The workshop attendees were fascinated by progress in nano memory
  - IBM presented a table of 70 memory options
- Emerging memories are seen in a different hierarchy than SRAM-DRAM-Disk
  - Removes memory wall
- 1 TByte memory with CPU seen as good thing





# Storage Hierarchy

---



- **Question and Answer**

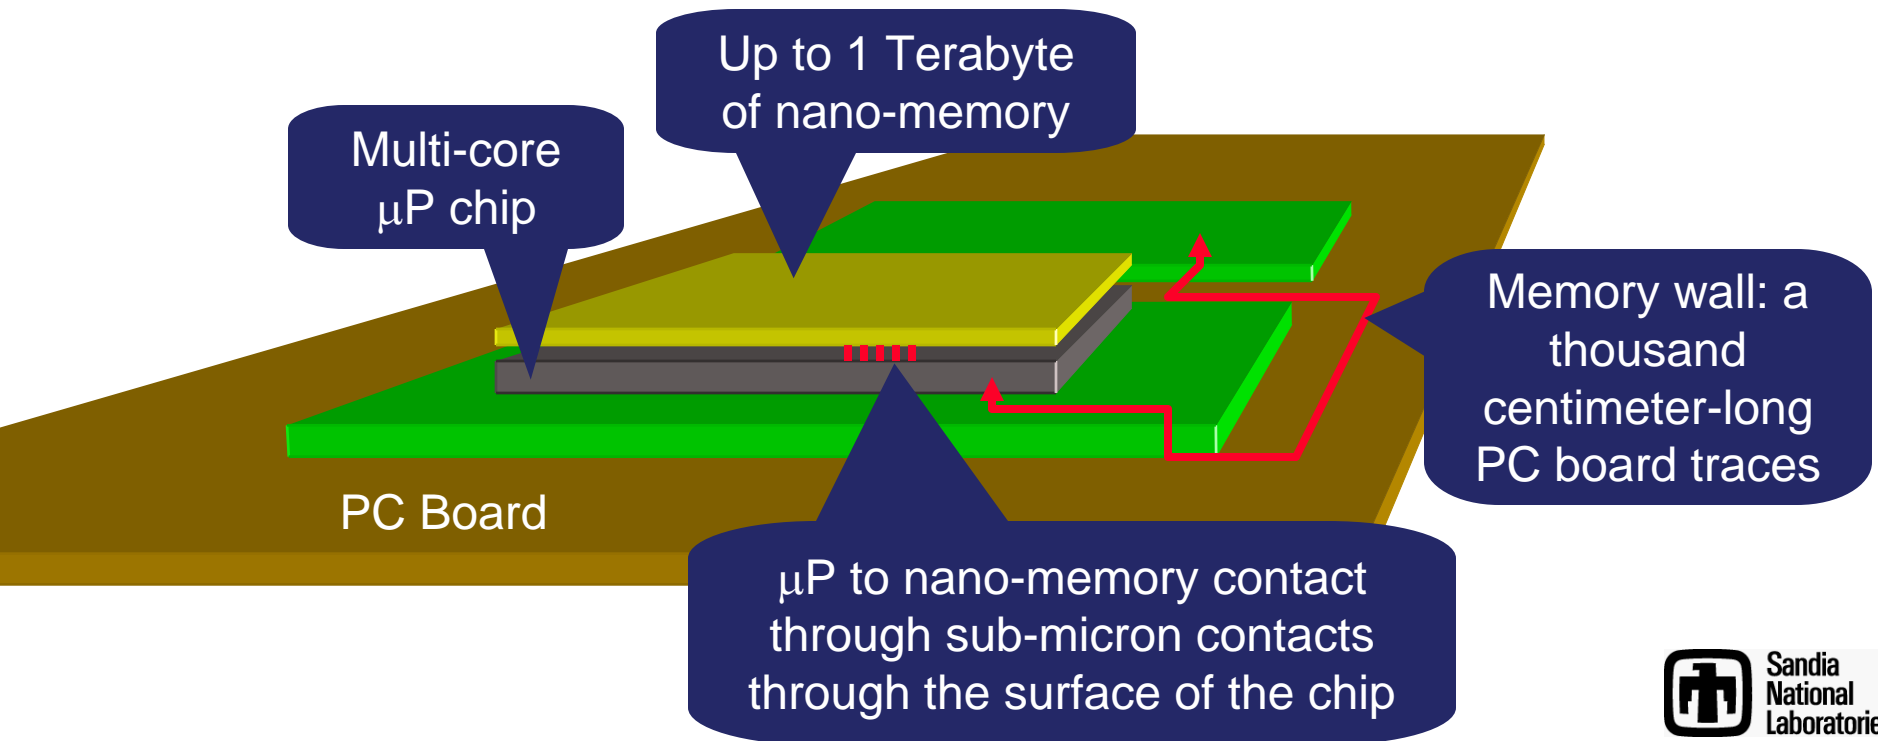
- **Q: What is the DRAM replacement?**
- **Isn't one**

- **Issue**

- **Emerging nano-memory is closer to roll-out than many users expect**
- **However, the nano memory will fill a useful by unprecedented spot in the storage hierarchy**

# Architectural Impact of Memory Wall

- Some of the nano-memory proposals offer advantages in bandwidth and latency as well as density





# Memory Observations & Conclusions

---

- **The memory and storage story is technically fascinating to High Performance Computing**
- **Industry claims they don't have market drivers to develop terabyte embedded memories for Teraflops chips**
  - **In fact, I don't see commercial applications that call for Teraflops chips irrespective of memory**
- **The net result seems to be “good news” for computing in 20 years, but there are unknown architectural consequences**



# Findings/Agenda

---

- Transition of Eras:  $kT$  limit
- Flat Clock Rate
- Fascinating Memory Story
- **On-Chip Optical Interconnect**
- Architecture
- A New Physical Device



## **Finding 4: On-Chip Optical Interconnect**

---

- **Metal (Cu) interconnect is seen as excessively high power in desirable architectures**
- **Effective architectures will need high bandwidth, low latency cross-chip communications**
- **This would require high power drive on long lines with frequent repeaters**
- **Proposed on-chip optical interconnect is seen as a promising alternative**
- **Darpa UNÍC is seen as a promising option**



# Findings/Agenda

---

- Transition of Eras:  $kT$  limit
- Flat Clock Rate
- Fascinating Memory Story
- On-Chip Optical Interconnect
- **Architecture**
- A New Physical Device

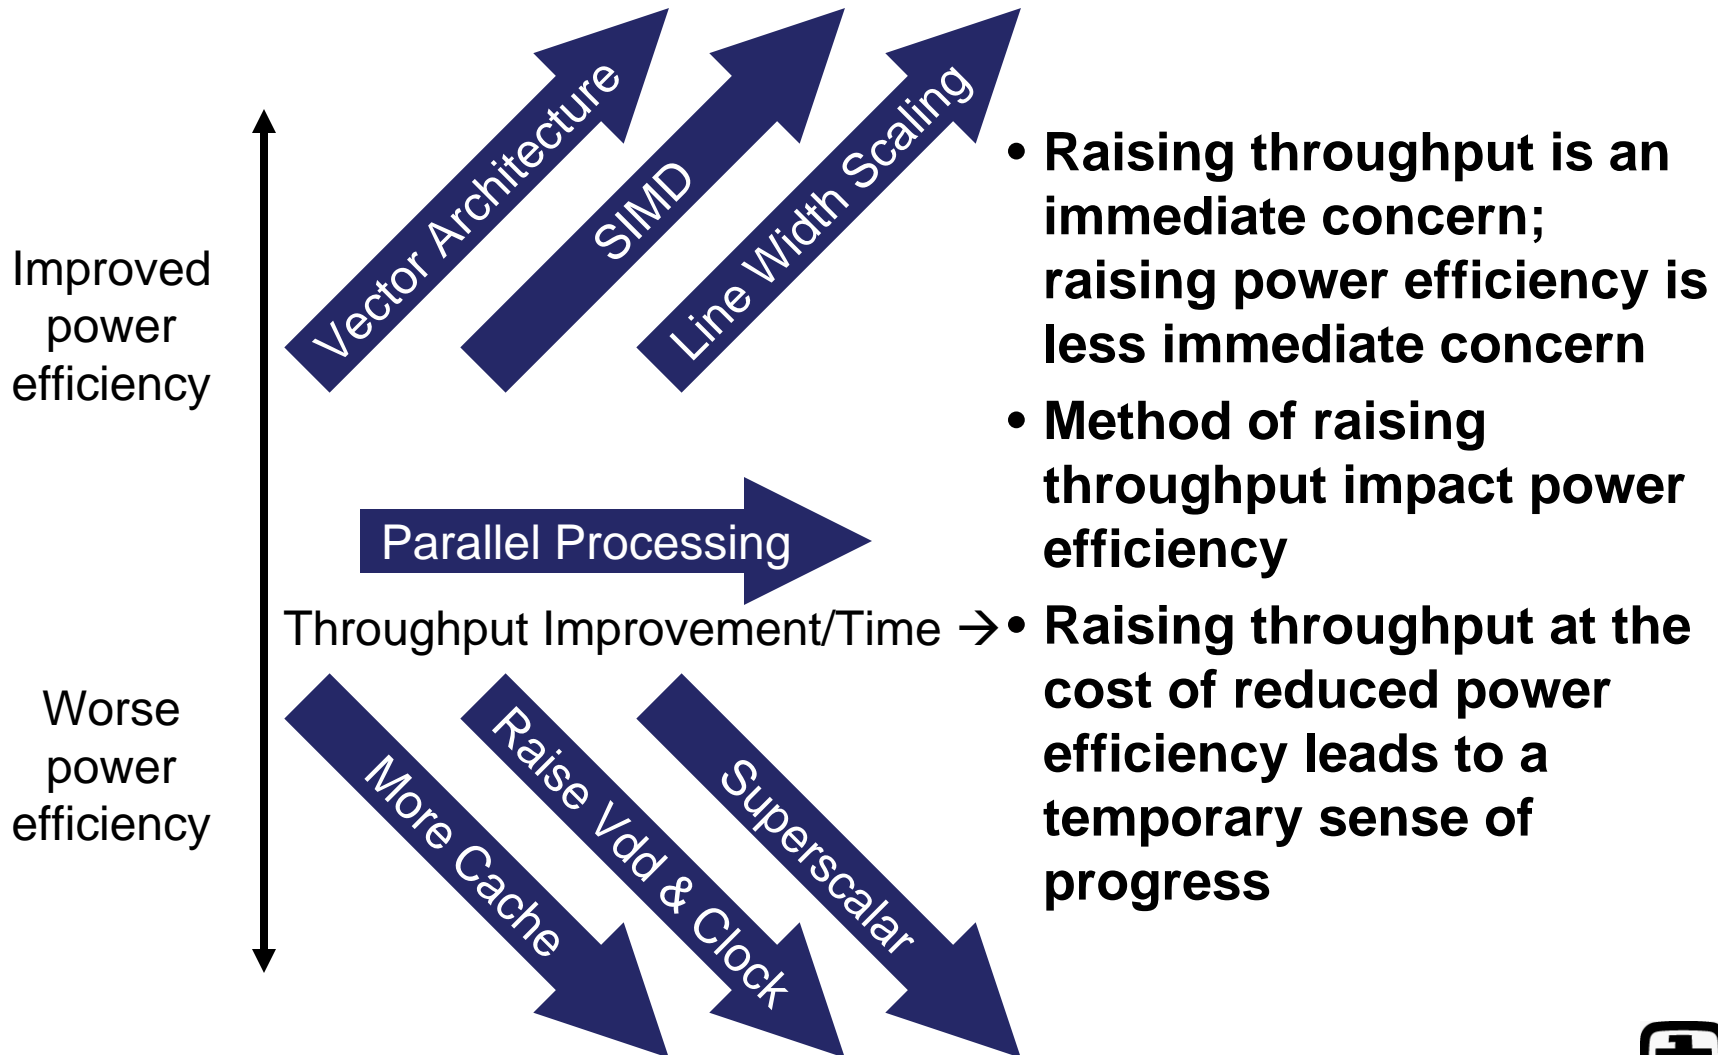


## **Finding 5: Architecture & Language Accommodations**

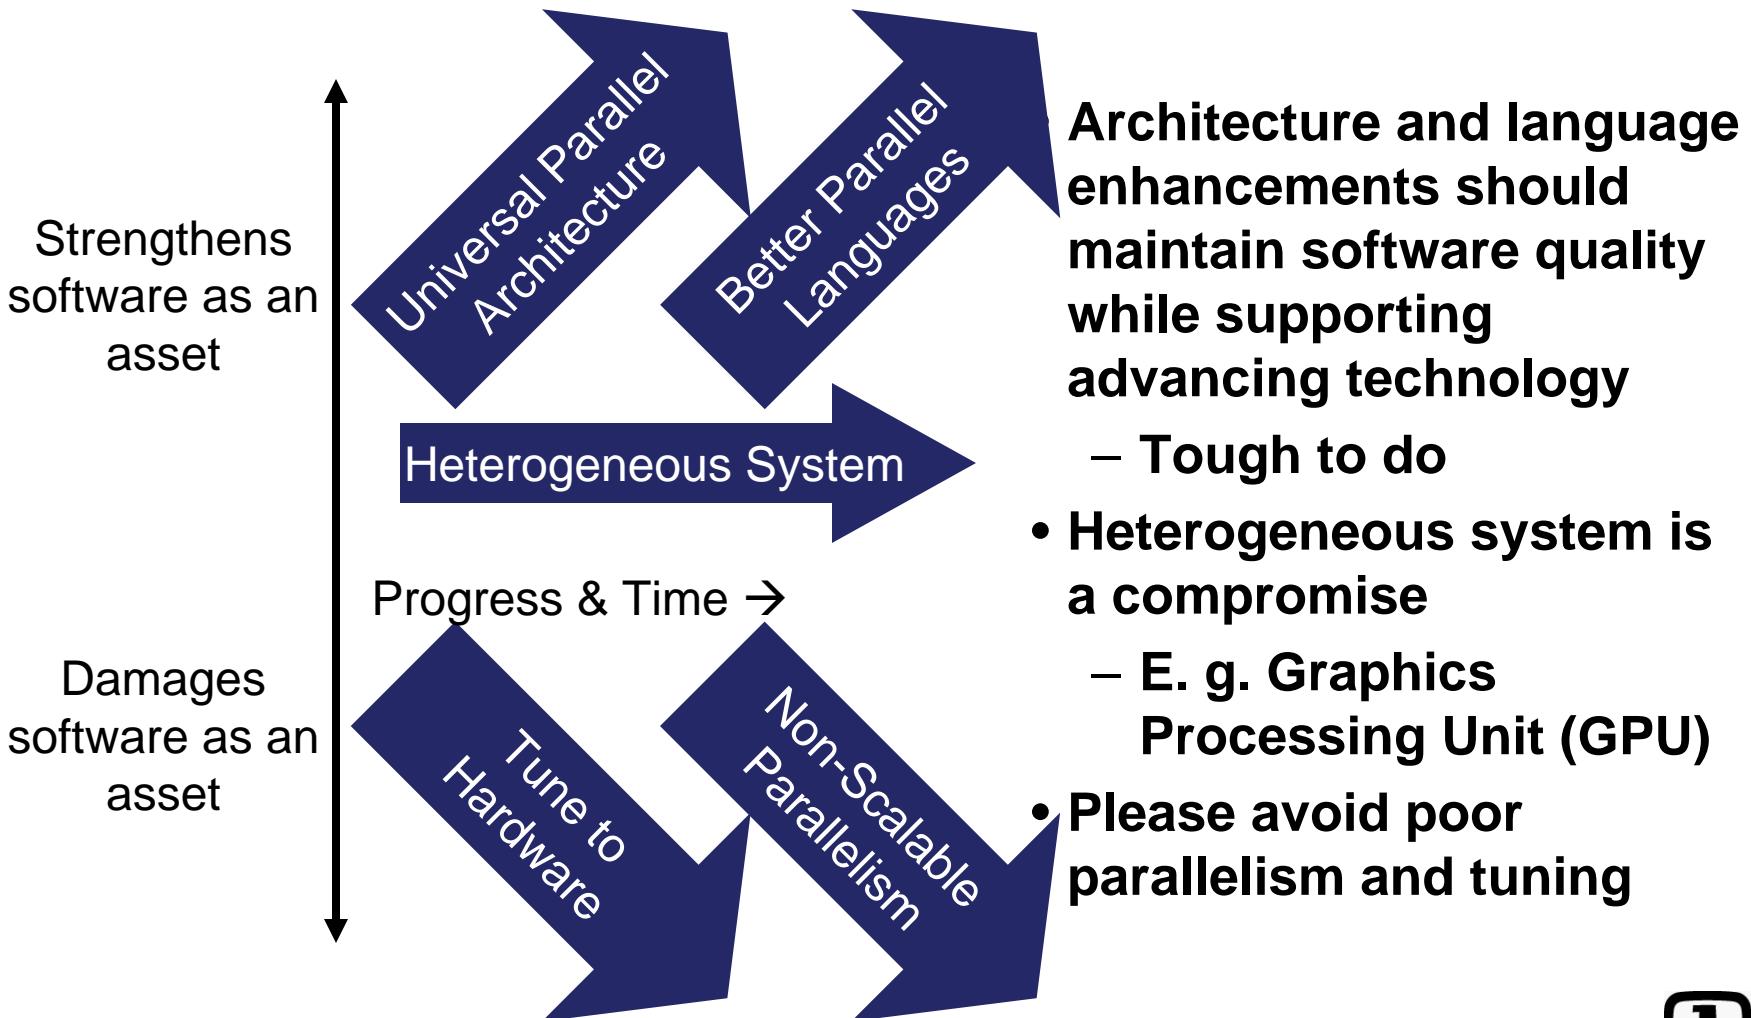
---

- **There is a non-trivial interplay between clock rate, power, and programmability**
  - Urgent need to mitigate stagnant clock rate
  - Long term need to mitigate power limit
- **Rhetorical questions**
  - What if a clock rate mitigation makes the power problem worse?
  - What if mitigating a hardware problem makes a software problem worse?

# Raising Throughput/Clock Rate Flat Lining



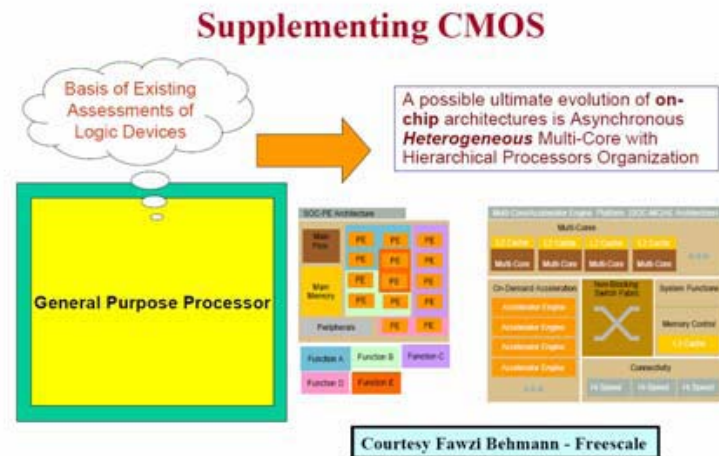
# Programming and Technology Trends



# Heterogeneous Architecture: Contested

- Heterogeneous parallel & accelerated hardware is ubiquitous (go to Frys)
- Heterogeneous programming languages are available for free on Internet (nVidia CUDA)
- However, experts are unsure whether accelerators will ever catch on
- Disconnect!

- Heterogeneous parallel accelerated architectures are in the roadmap:





# Call to Action

---

- **We don't know what the right architectures are**
- **However, a “point study” could be useful**
- **General technology trends are known (ITRS)**
  - **Use ITRS the last year predicted (currently 2022)**
  - **Do a few trial designs of a system using 2022 technology**
  - **Report the best design options and likely throughput**



# Findings/Agenda

---

- Transition of Eras:  $kT$  limit
- Flat Clock Rate
- Fascinating Memory Story
- On-Chip Optical Interconnect
- Architecture
- A New Physical Device



## Finding 6: A New Physical Device

---

- **HP, IBM, and Intel all confirm that the end of CMOS scaling is not necessarily the end of Moore's Law**
  - **IBM and HP: Reversible computing; Intel: logic “out of thermal equilibrium with the environment”**
  - **Quantum computing**
- **No workshop participants voiced disagreement**
- **Physical science research on such devices is believed to be worthy, but should not take resources away from continued CMOS scaling**

# Overall Conclusions

- The Zettaflops participants found an “era transition” at  $12\pm$  years ( $2020\pm$ )
  - No growth to slow growth in clock rate
  - Exponential power reduction to limit
  - Good news for memory
  - Architecture and programming cope with unstable hardware
- Beyond the transition
  - Physical science research into a new physical device is encouraged, but nobody has a timescale
- Moore’s Law continues!

