

# Bridging the Gap: Digital Architecture at the End of Scaling

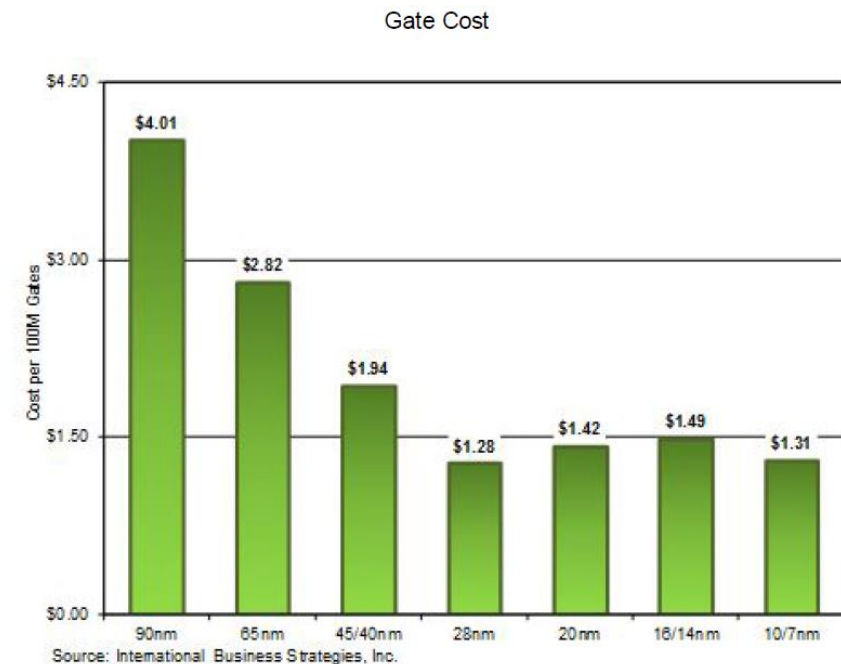
Ravi Nair  
IBM Thomas J. Watson Research Center  
Yorktown Heights, NY 10598

March 24, 2016

Special thanks for discussions with Wilfried Haensch,  
Changhoan Kim and Paul Coteus

# Scaling will continue in the short term but with reduced economic viability

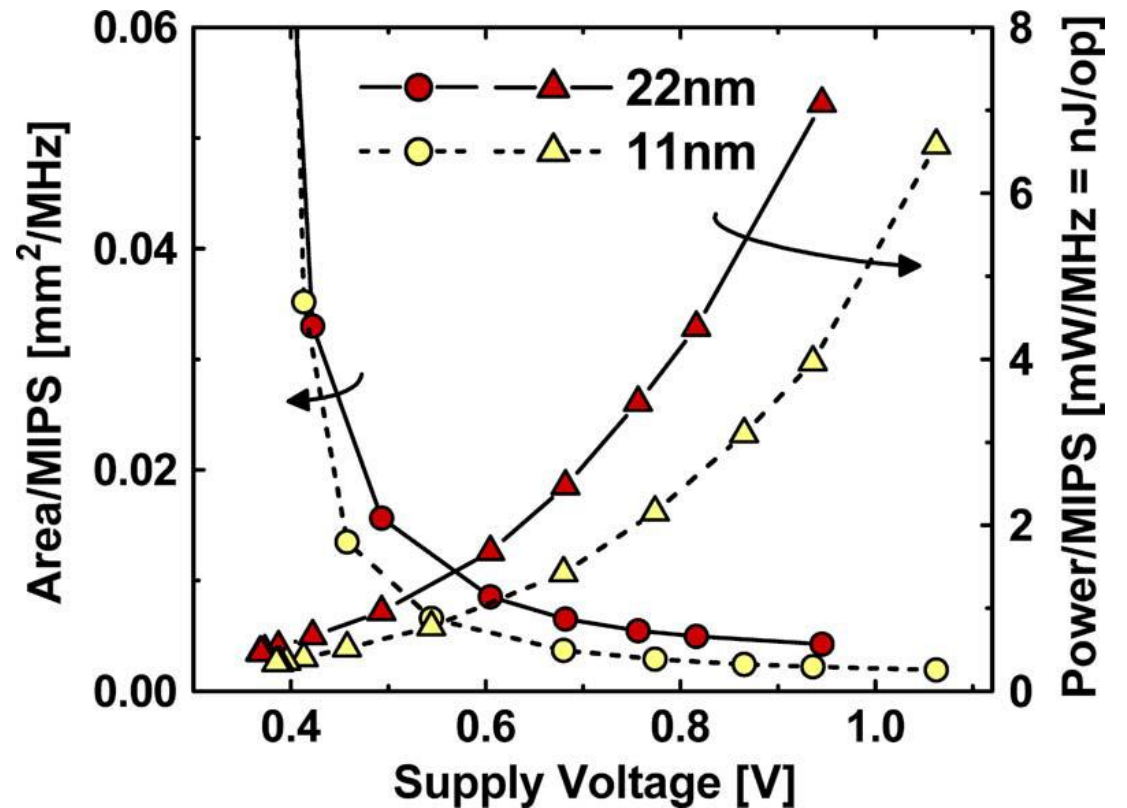
- Innovative silicon processing techniques will push down geometry of features
- Manufacturing cost per gate may not be dropping as we are used to
- Variability, power and yield limitations will affect cost of designs



The gate cost projection for 10/7nm is based on the ability to obtain high parametric and systemic yields.

# Energy consumption will force architectures to compromise in area and/or performance

- Fast and small: good if energy is free
- Slow and large: good if area is free
- Optimum solution determined by cost of energy and cost of area



Chang et al, Proc. IEEE, Feb. 2010

# Innovation in hardware must be directed towards helping software combat data movement problem

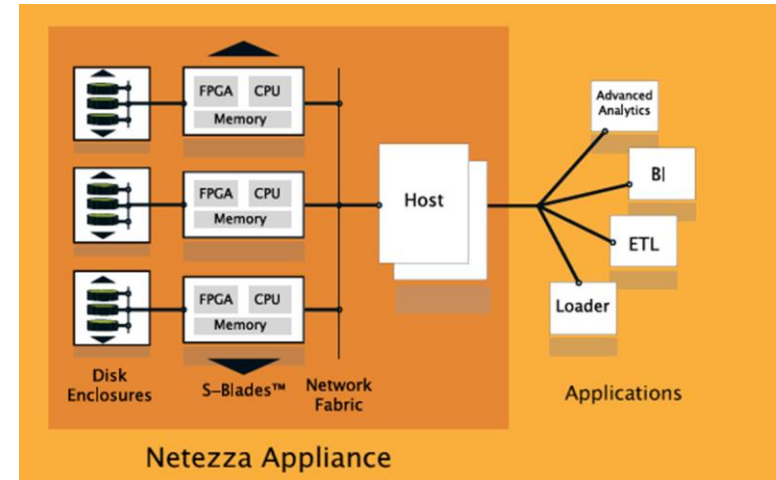
- 3D packaging
- Photonics
- Function specialization
- Use of non-volatile memory
- Accurate and timely prefetching
- Processing near memory/storage
- Remote memory transactions

Operation	Operation Energy Cost (nJ)	Equivalent ADD
ADD	0.64	-
L1->REG	1.11	1.8x
L2->REG	2.21	3.5x
L3->REG	9.80	15.4x
MEM->REG	63.64	99.7x
Stall	1.43	-
Prefetching	65.08	-

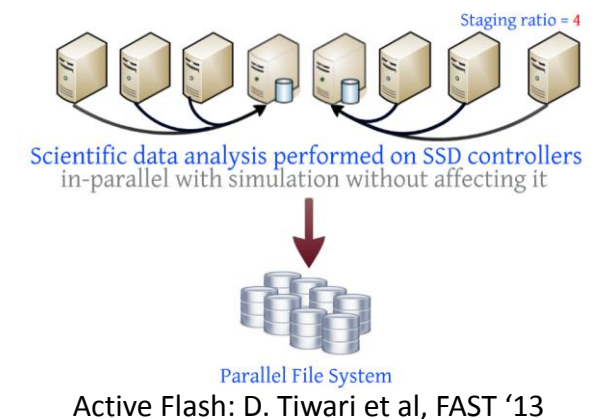
Kestor et al, Workshop on Modeling & Simulation of Systems and Applications, 2014

# Storage architectures are evolving to reduce data movement

- Offload activity to processor near storage
- Low complexity of compute element near storage




## Active Computation on SSDs



# Processing near-memory and in-memory offer approaches to reducing energy

- Large bandwidth/compute ratio
- Energy saved in moving data
  - Proximity
  - No chip crossing
  - Smaller data granularity
- Energy saved in computation
  - Domain-specific implementation
  - Atomics



**Active Memory Cube**

Nair et al, IBM Journal R&D, March/May 2015

# New technologies will not be ready for high-end applications in time to meet the end of scaling

- Tunnel FET closest (10 years?)
  - $I_{on}/I_{off}$  ratio at required  $I_{on}$  not high enough for high-performance applications
- Carbon nanotube next closest
  - Volume production problems
- Other technologies further away
  - Different problems, e.g.
    - Variability in dipole devices
    - Bit error rate not low enough in Spin Logic
    - Insufficient resistance change in voltage-controlled magnetic devices
    - Insufficient integration density in spin wave

Device name	acronym	class	subclass
Si MOSFET high perf.	CMOS HP	electronic	barrier
Si MOSFET low power	CMOS LP	electronic	barrier
Homojunction TFET	HomJTFET	electronic	tunneling
Heterojunction TFET	HetJTFET	electronic	tunneling
Interlayer tunneling FET	ITFET	electronic	tunneling
Graphene nanoribbon TFET	gnrFET	electronic	tunneling
Graphene pn-junction	GpnJ	electronic	refraction
Ferroelectric FET	FEFET	ferroelectric	hysteresis
Negative capacitance FET	NCFET	ferroelectric	non-hysteresis
Piezoelectric FET	PiezoFET	straintronic	polarization
Bilayer pseudospin FET	BisFET	orbitronic	exciton
Metal-insulator transistor	MITFET	orbitronic	bandstructure
SpinFET (Sugbara-Tanaka)	SpinFET	spintronic	spin drift
Spin torque domain wall	STT/DW	spintronic	domain wall
Spin majority gate	SMG	spintronic	domain wall
Spin torque triad	STTtriad	spintronic	nanomagnet
Spin torque oscillator	STO	spintronic	nanomagnet
All-spin logic	ASL	spintronic	spin diffusion
Charge-spin logic	CSL	spintronic	spin Hall
Spin wave device	SWD	spintronic	spin wave
Nanomagnetic logic	NML	spintronic	nanomagnet

Nikonov, Beyond CMOS benchmarking, APS 2014



# Low-end applications are best positioned to exploit technology at the end of scaling

- Low power requirement
- Tolerance for imperfection
  - Sensors
  - Analog circuits
  - MEMS
- Modest constraints
  - Performance
  - Area
- Volume provides better ROI

## Mixed Tunnel-FET/MOSFET Level Shifters: A New Proposal to Extend the Tunnel-FET Application Domain

Marco Lanuzza, Member, IEEE, Sebastiano Strangio, Felice Crupi, Senior Member, IEEE, Pierpaolo Palestri, Senior Member, IEEE, and David Esseni, Fellow, IEEE

**Abstract**—In this paper, we identify the level shifter (LS) for voltage up-conversion from the ultralow-voltage regime as a key application domain of tunnel FETs (TFETs). We propose a mixed TFET-MOSFET complete Simulink dynamic model compared with the conventional MOSFET and pure TFET solutions. The conventional voltage signal of the device with the solution variations of higher four different TFET-MOSFET design of regime compared with the conventional MOSFET and pure TFET solutions.

**Index Terms**—Level shifter (LS), technology computer-aided design (TCAD), tunnel FET (TFET).

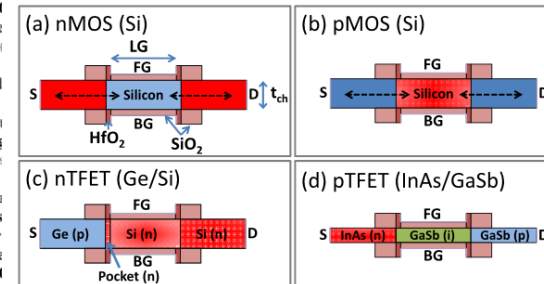
### I. INTRODUCTION

WITH THE growing interest in low-energy-budget electronic applications, the tunnel FET (TFET) is playing a major role as a new device concept, featuring a better performance/leakage tradeoff than the conventional MOSFET at scaled power supply voltage ( $V_{DD}$ ) levels.

characteristics due to band-to-band tunneling (BIBT) current [1]. Several mixed device-circuit studies as well as redesign strategies have investigated the device concept at both device and circuit

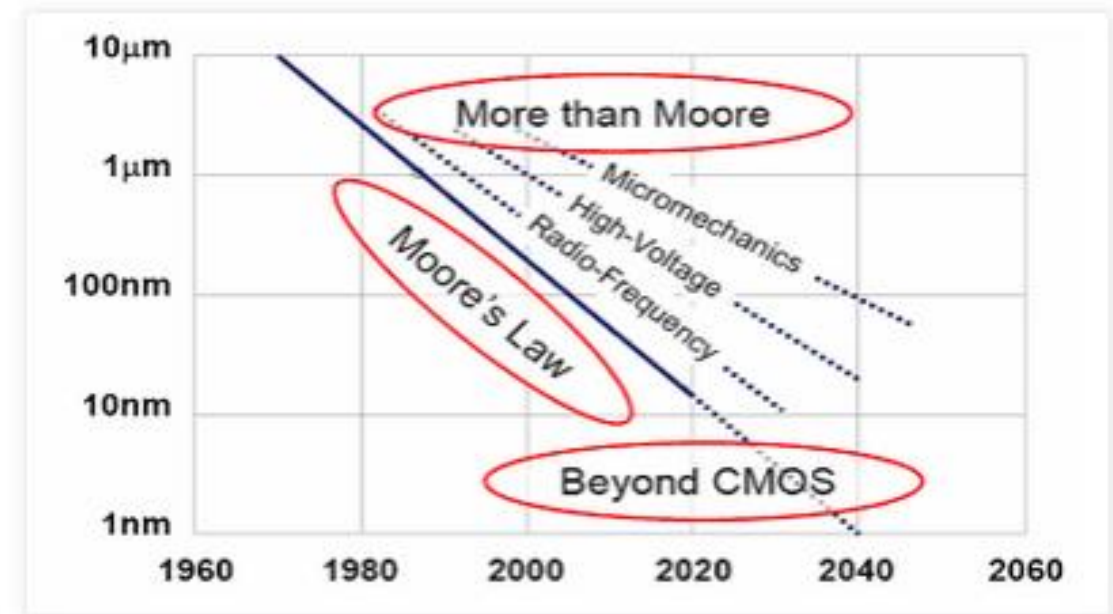
level system-on-chips (SoCs), the multisupply (MSVD) technique [14] is emerging as an effective design strategy to improve energy efficiency. The MSVD of partitioning the design into separate voltage islands, each operating at a different voltage level depending on its timing requirements, allows different domains run at higher power supply voltage to maximize the speed, whereas noncritical domains run at lower supply voltage ( $V_{DDL}$ ) to optimize energy consumption, thus effectively managing tasks that require different performances. Minimizing the delay and energy overhead of level conversion between different voltage domains is a key challenge in the design of effective multisupply SoCs, becoming particularly critical when the number of power domains and/or the data width in the SoC increases [15]. Within this context, several level shifter (LS) circuit topologies were recently proposed for speed- and energy-efficient wide-range conversion from the deep subthreshold regime up to the nominal supply voltage level [16]–[21].

In this paper, we propose a mixed TFET-MOSFET LS design methodology, which exploits the complementary



# A gap is looming in large systems between the end of scaling and the adoption of new technologies

- No significantly new technology available for large digital computers during this gap
- Traditional CMOS technology will get commoditized
  - Difference in capabilities of foundries will reduce at the end of scaling



IEEE Panel, 2014

# Traditional computing systems moved data, computed on data; they seldom made decisions

## Computers

Transaction Processing

Physical Simulation

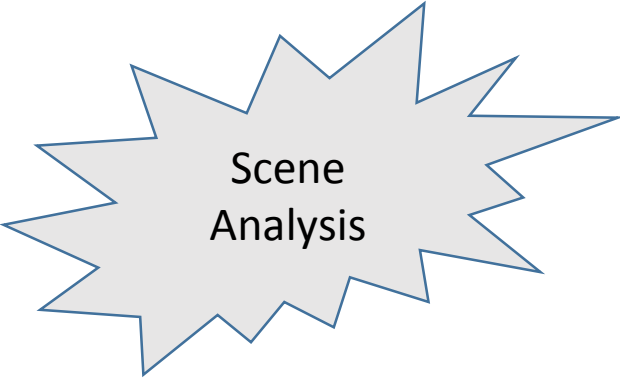
Media Rendering

## Decision Makers

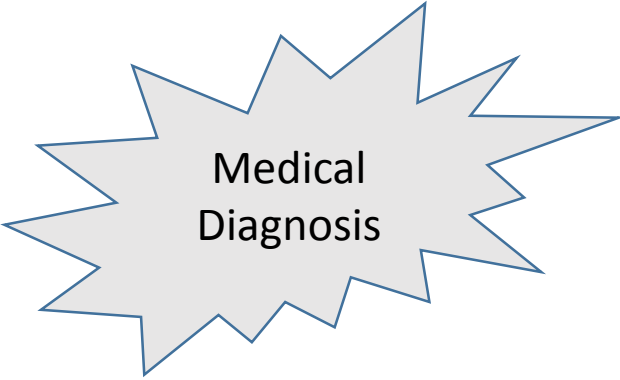


Illustration by Leo Cullum, The New Yorker

# Computers increasingly will be called upon to make rapid decisions with limited resources



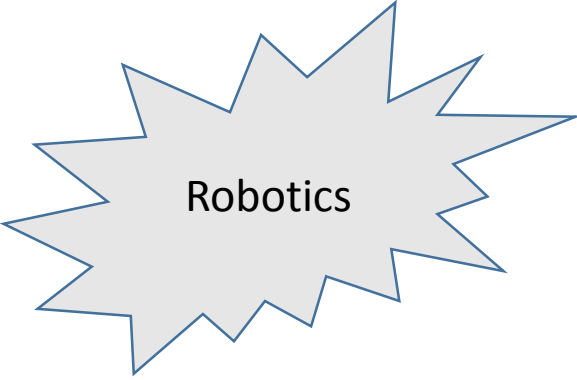
Scene  
Analysis



Medical  
Diagnosis



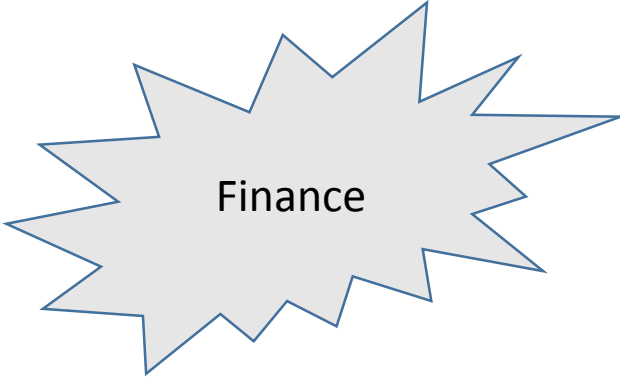
National  
Security



Robotics



Advertising



Finance

# Long-accepted practices in digital computer design will need to be revisited

## • Revisit

- Focus on performance
- Fixed precision
- Determinism
- Perfect hardware
- Perfect input data
- Even digital computing

### Approximating the input

- Sampling systematically
  - Sketching, compressed sensing
- Dropping inputs due to latency pressure
  - Media, e.g. video streams
- Noise in input stream
  - Stochastic representation of numbers
- Inputs not arriving
  - Timeout and recovery
- Precision of input

Compressed sensing is used in mobile phone camera sensors to reduce image acquisition energy by a factor of up to 15.

### Approximating the system

- Drop packets
  - That are late
  - That are corrupted
- Race-and-Repair
  - Allow stale values to be used
  - Employ fresheners at regular intervals

### Approximating the problem

- Heuristics
  - Reducing the complexity of the problem
- Empirical formulation
  - Equation may be approximate
  - Parameters may be approximate
- Precision of calculation
  - Approximate data-types, EnerJ, Uncertain-C<sup>T</sup>
- Probabilistic programming

Annotating approximate data-types in EnerJ saves 10-50% energy with little loss of accuracy in result.

*Approximate  
Computing is helping  
mitigate resource  
limitations*

### Approximating the algorithm

- Multi-precision solutions
  - Dongarra's group
  - Rinard's group
- Loop perforation
  - Rinard's group
- Control approximation
  - Yetim, Malik, Martonosi

Mixed precision iterative refinement methods have been shown to be about 2x better in energy efficiency compared to double-precision solutions.

### Approximating the program

- Changing convergence criterion
- Relaxed synchronization
- Eventual consistency
- Load value approximation

Balls et al show 5x improvement in latency of read accesses from data stores using eventual consistency techniques when 99.9% consistency was acceptable instead of 100%.

Renganarayanan et al show up to 15x improvement in performance of iterative convergence solutions using relaxed synchronization.

### Approximating the hardware

- Use analog circuits
  - Analog neural nets
- Unplanned approximation
  - Hardware failure, e.g. soft error, multi-level memory inaccuracy, refresh (Spark, Flickr)
- Algorithmic noise tolerance
- Scalable stochastic processors

St. Amant et al show up to 3.7x speedup and 6.3x energy savings with quality loss less than 10% using analog neural acceleration.

### Approximating the output

- Visualization
- Decision making
  - Jeopardy
- Public service
  - Weather, health

Specifying the result in definite and quantitative terms, rather than as a set of probabilistic outcomes.

# Custom designs form the next big opportunity

- General purpose has a cost
- Power efficiency could force narrowing of desired scope for platforms
- FPGAs already more power-efficient (but slow)
  - Substrate is still general-purpose
- Customization of substrate will yield further power-efficiency and performance

GRAPE-8 — An Accelerator for Gravitational  $N$ -body Simulation with 20.5Gflops/W Performance

Jun Makino, TiTech,  
Hiroshi Daisaka, Hitotsubashi Univ.

SC12 Nov 15, 2012

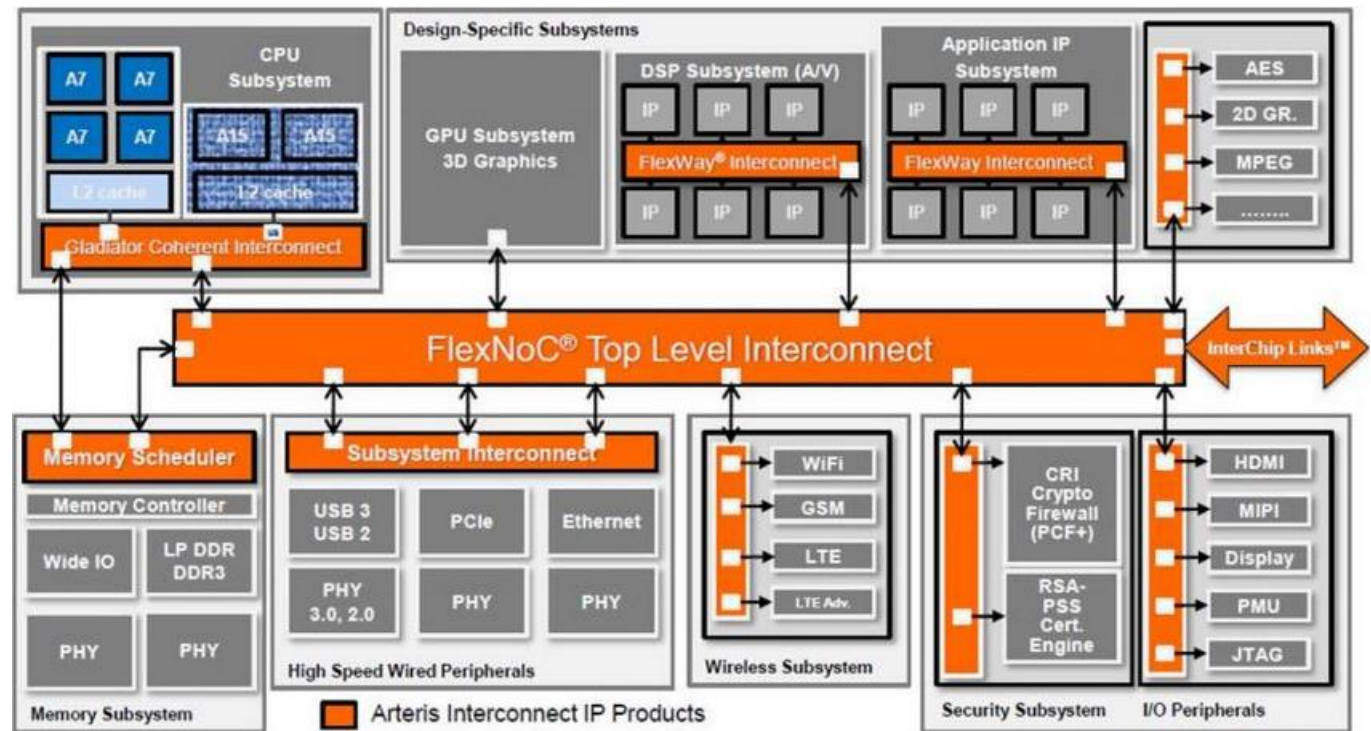
GRAPE-8 board



Two GRAPE-8 chips and one FPGA for PCIe interface  
960Gflops peak, 46W (G8 chips:26W, FPGA: 20W)

# Large systems will take a leaf out of the embedded systems book

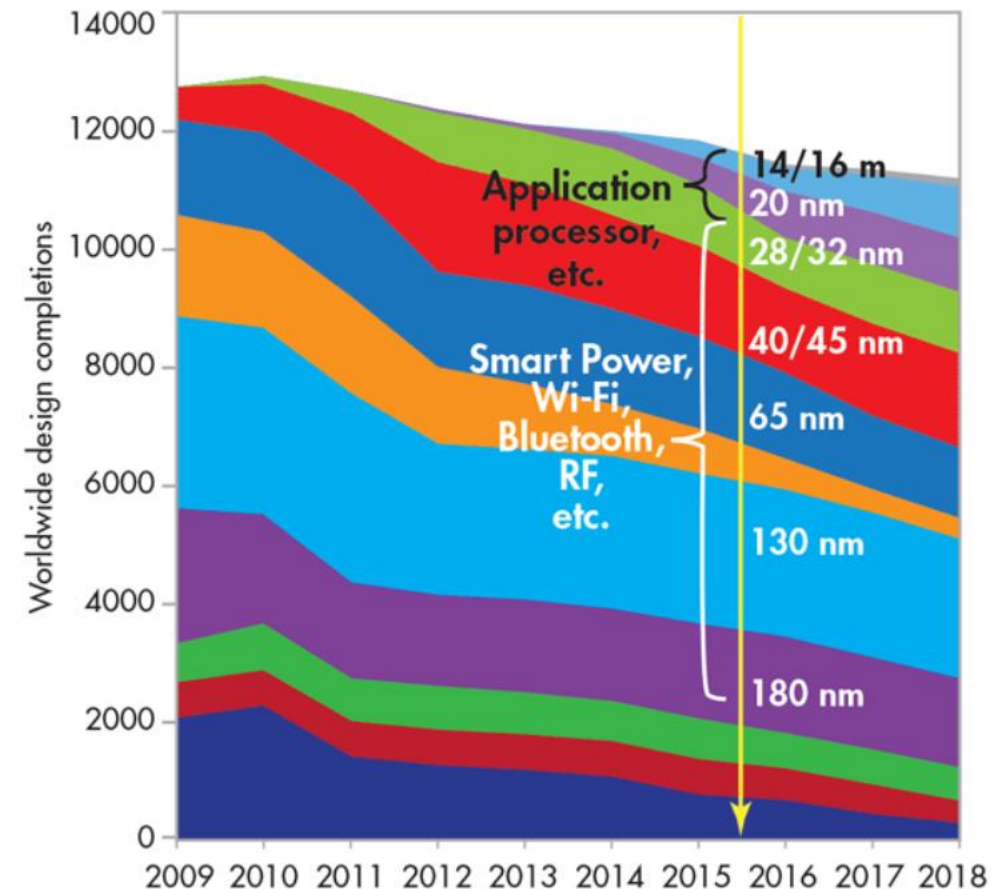
- Energy constraints will force even large systems to be built from interconnectable IP bricks



<https://lazure2.wordpress.com/soc/>

# Low-end design completions skew heavily towards established process nodes

- High-end designs tend to be early adopters of technology
  - Highly skilled designers
  - Long development cycles
  - Push on all fronts
  - Typically tend to need advances in design tools as well as programming tools




<http://electronicdesign.com/eda/iot-cost-transistor-extend-lifetimes-established-technology-nodes>  
and VLSI Research



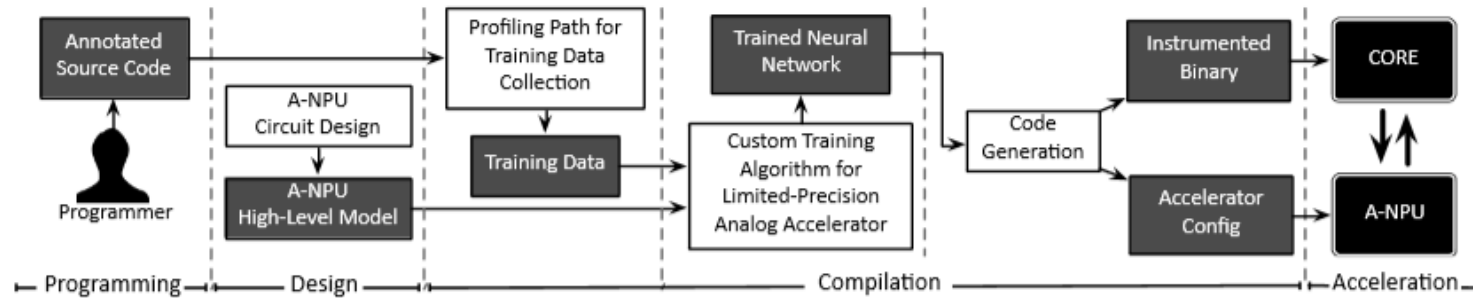
# Mature technologies offer economically viable substrates for re-thinking large systems

- Cost
- Yield
- Reliability
- Tools
  
- Retrofit advanced microarchitecture and compilation techniques into mature technologies
- Possible approach for negotiating the gap



<http://anysilicon.com/semiconductor-technology-nodes/>

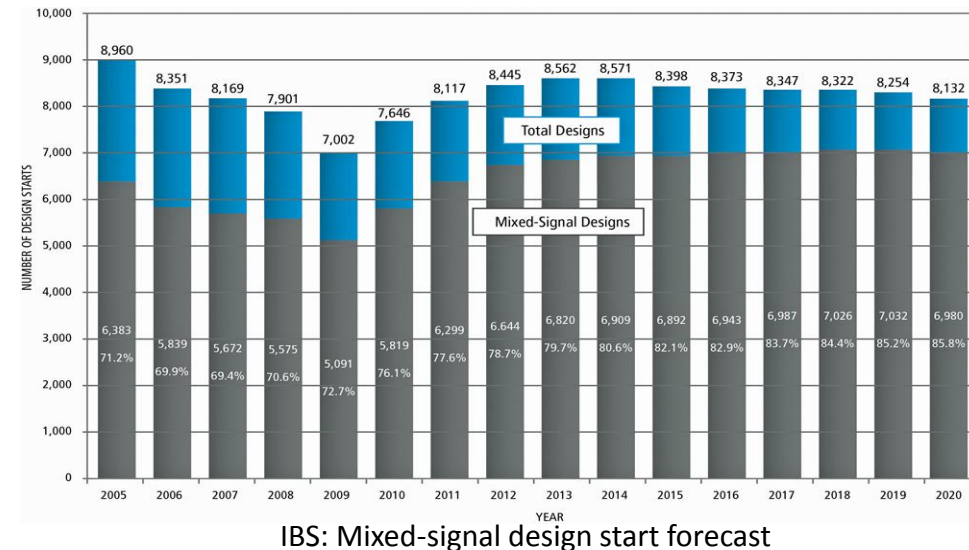
# Application specific designs can exploit variety in technology



Framework for using limited-precision analog computation to accelerate code written in conventional languages.

Amir Yazdanbakhsh et al, ASPLOS 2014

- Scaling era: Technology adapted to needs of high-end systems
  - Drove innovation in technology
- Post-scaling era: Systems must adapt to menu of technologies
  - Will drive innovation in systems



# Summary

- End of scaling offers an opportunity to re-think established practices
- The density available, even in current chips, offers a large canvas to create energy-efficient, custom compositions
- Work must be done to develop tools for making such compositions economically viable
- Effective development platforms will help ease in new technologies as they are developed