Panel Session:
Paving the Way for Multicore Open Systems Architectures

James C. Anderson
MIT Lincoln Laboratory

HPEC08
Wednesday, 24 September 2008

This work was sponsored by the Department of the Air Force under Air Force Contract #FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the author, and are not necessarily endorsed by the United States Government.

Objective & Schedule

- **Objective: Assess the infrastructure (hardware, software & support) that enables use of multicore open systems architectures**
  - **Where are we now?**
  - **What needs to be done?**
- **Schedule**
  - **1525: Overview**
  - **1540: Guest speaker: Mr. Markus Levy**
  - **1600: Introduction of the panelists**
  - **1605: Previously submitted questions for the panel**
  - **1635: Open forum**
  - **1655: Conclusions & the way ahead**
  - **1700: Closing remarks & adjourn**

Paving the Way for Multicore Open Systems Architectures

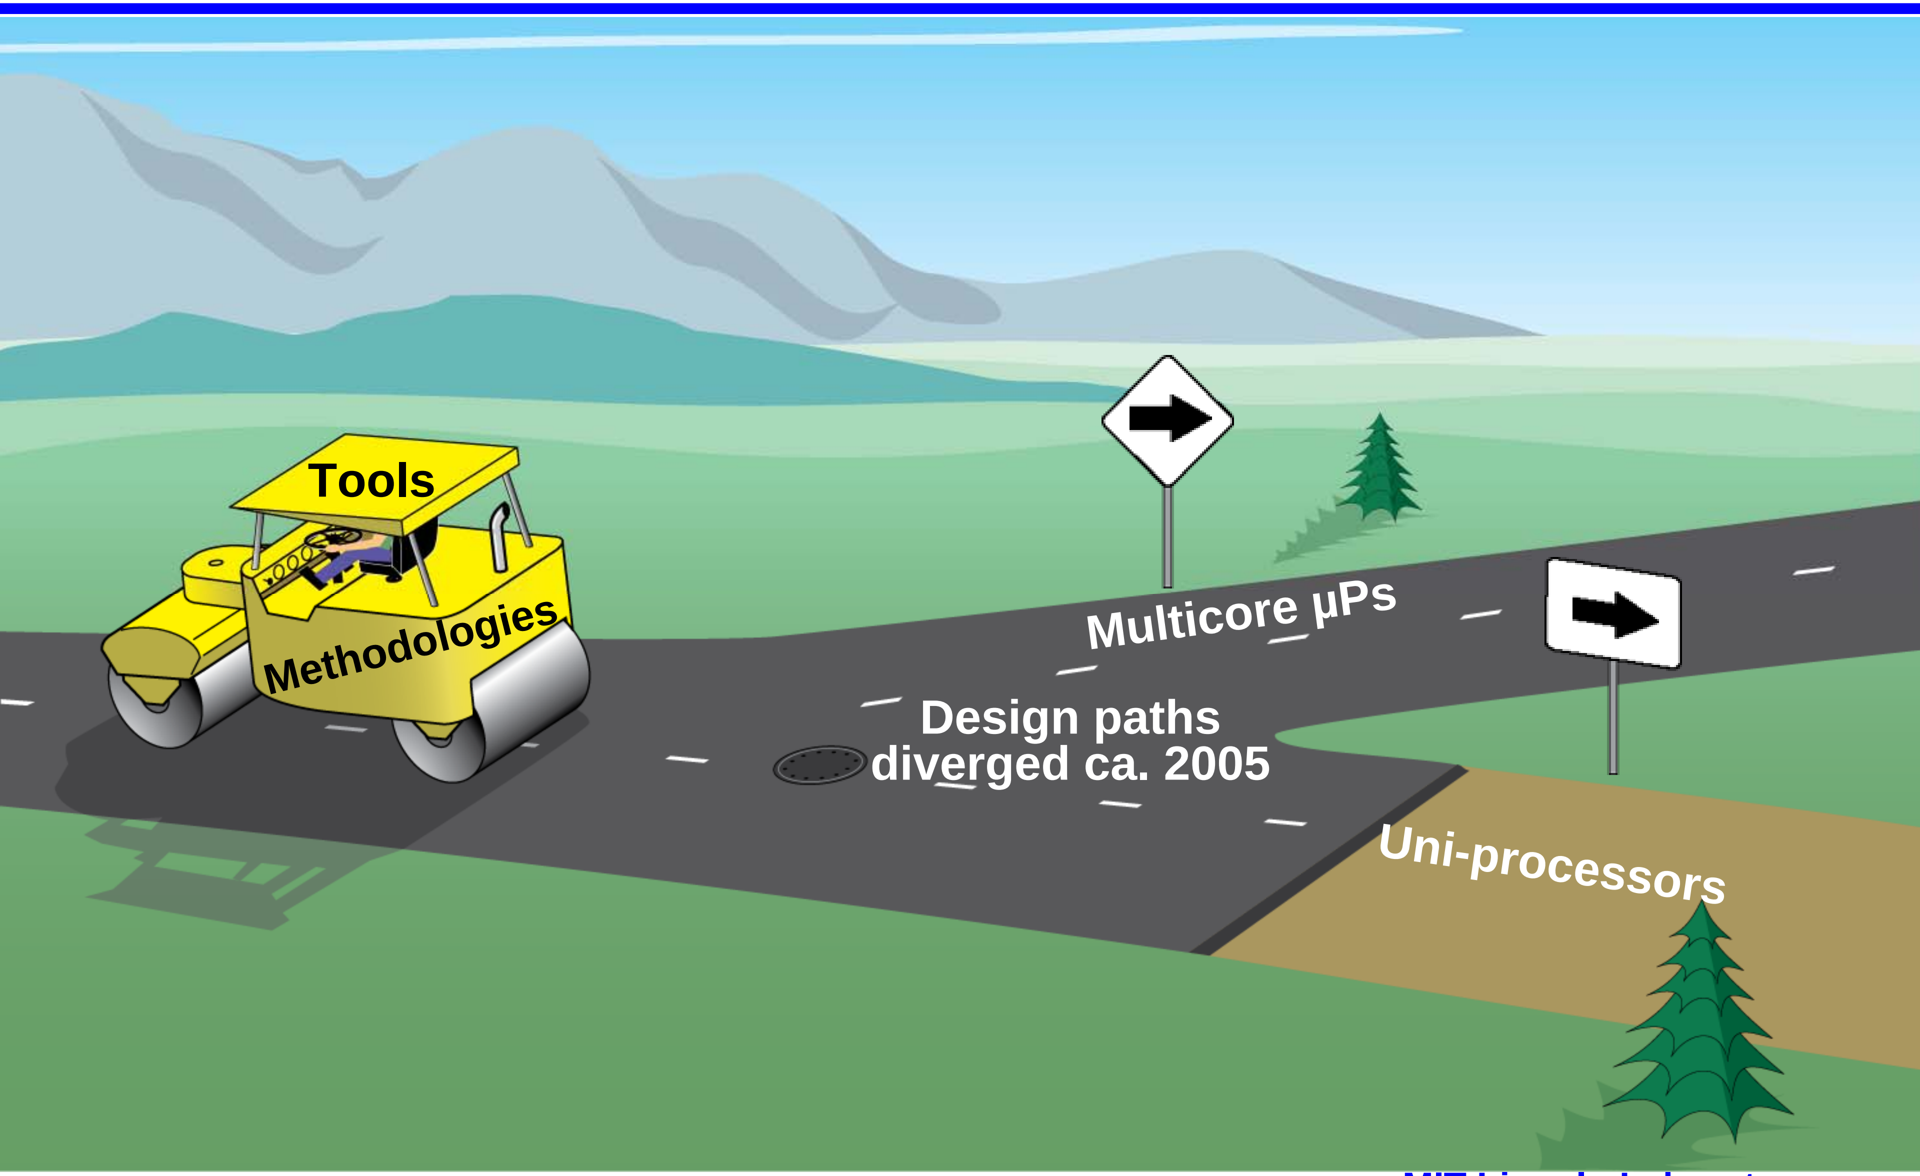

But First, A Few Infrastructure Issues

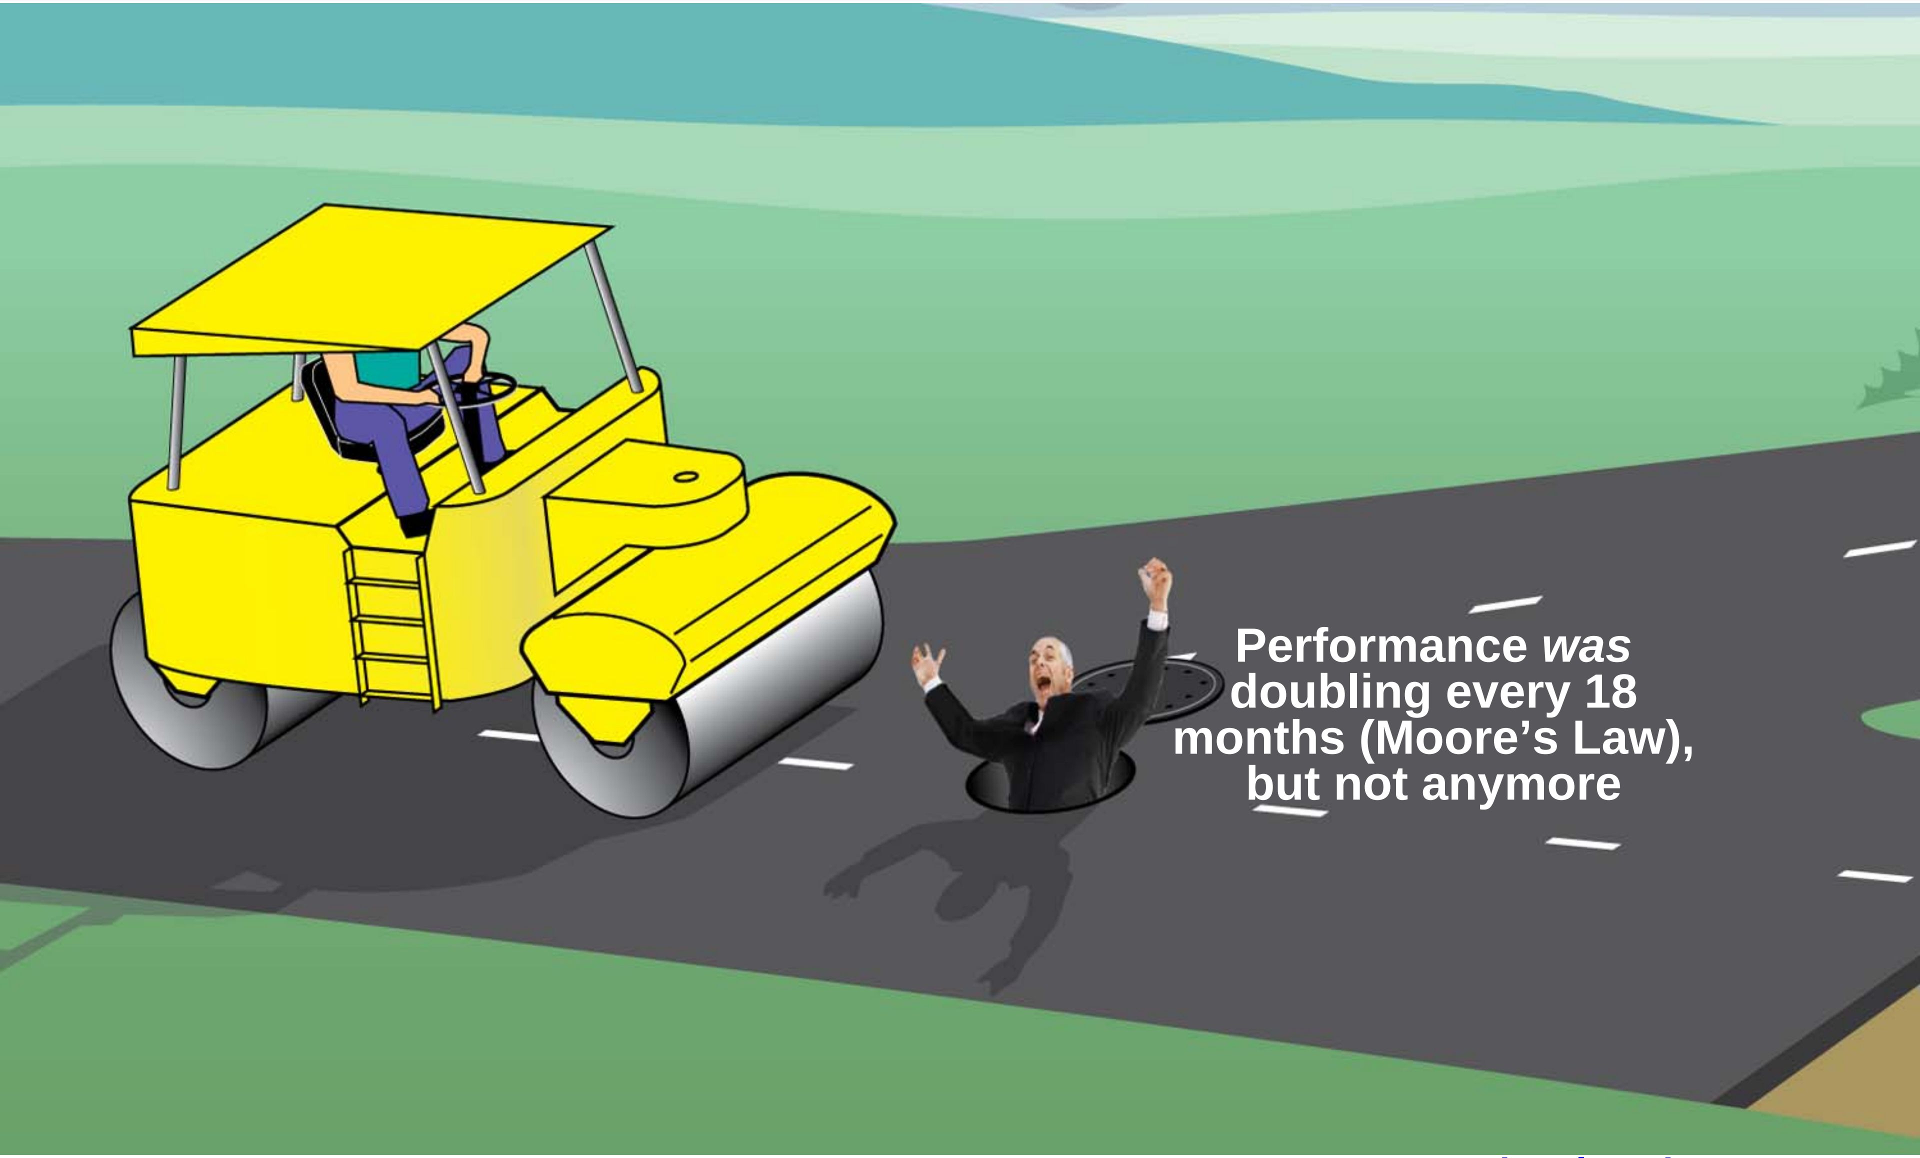

2000 International Technology Roadmap for Semiconductors (ITRS00)

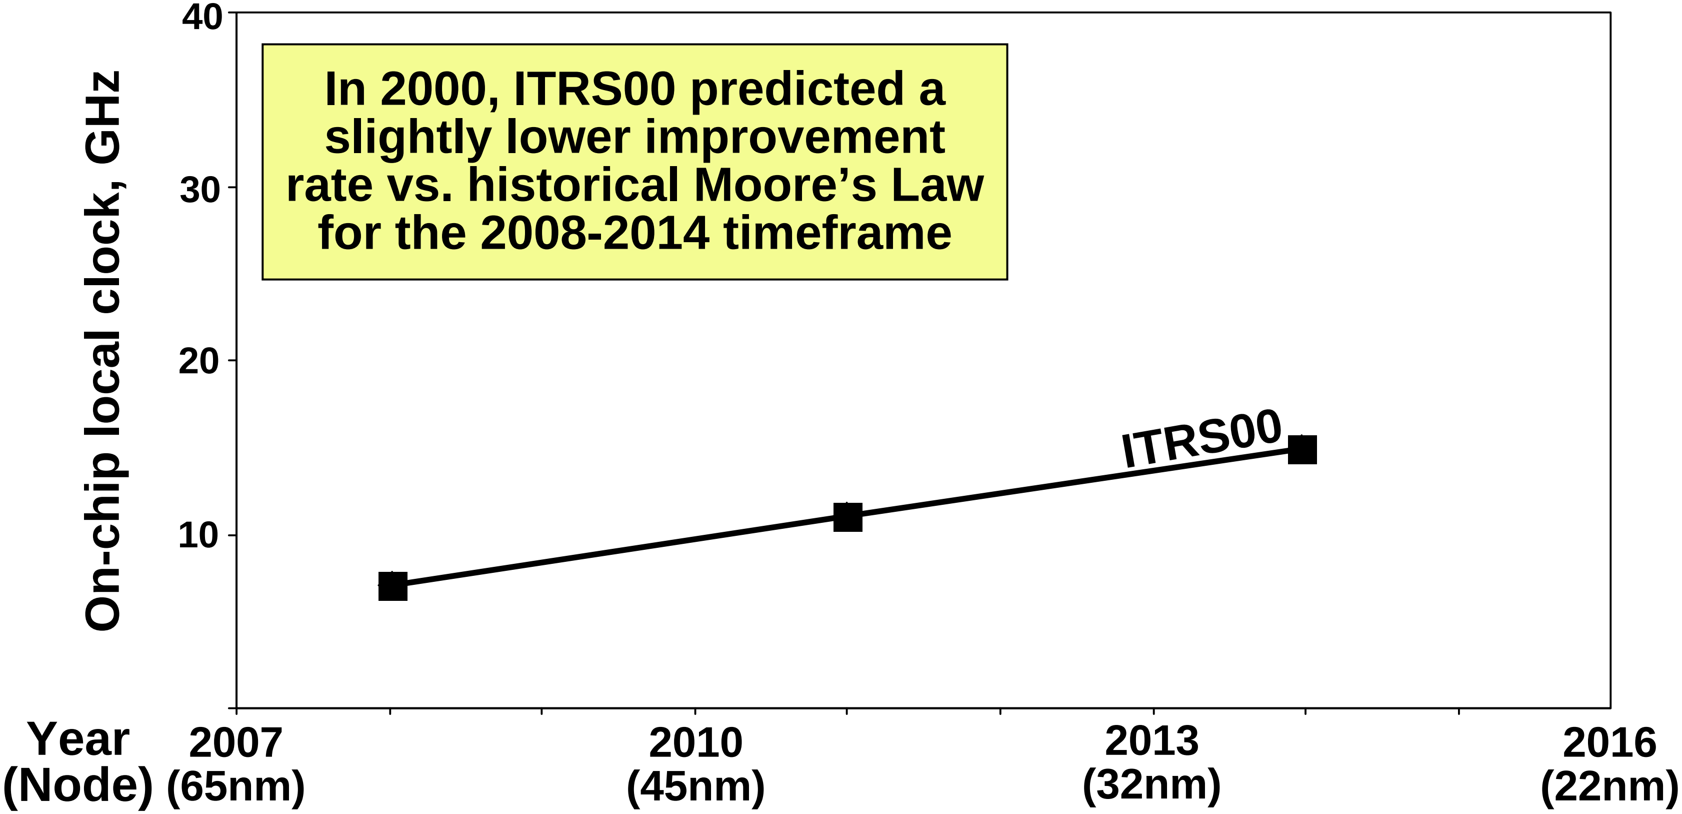

- **~3.5X throughput every 3 yrs predicted for multiple independent cores (~same as 4X every 3 yrs for historical Moore's Law)**
  - **1.4X clock speed every 3 yrs for constant power**
  - **2.5X transistors/chip every 3 yrs (partially driven by economics) for constant chip size (chip size growth ended ~1998)**

2001-2002 International Technology Roadmap for Semiconductors (ITRS01-02)

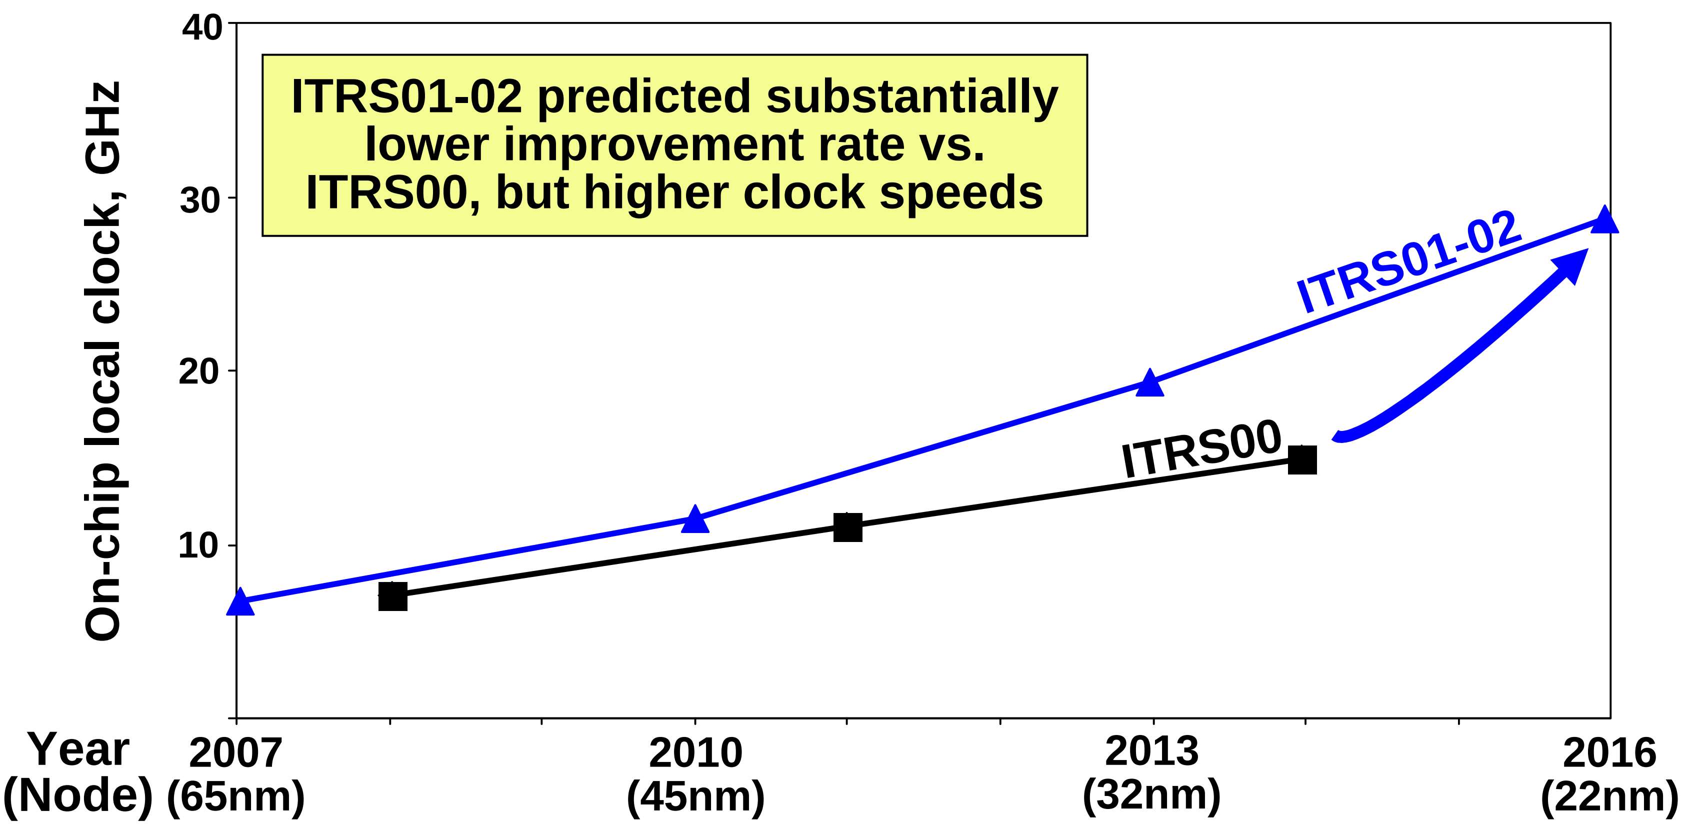

- **2.8X throughput every 3 yrs predicted for multiple independent cores**
  - **1.4X clock speed every 3 yrs for constant power (same as ITRS00)**
  - **2X transistors/chip every 3 yrs for constant chip size (less than ITRS00)**

2003-2006 International Technology Roadmap for Semiconductors (ITRS03-06)

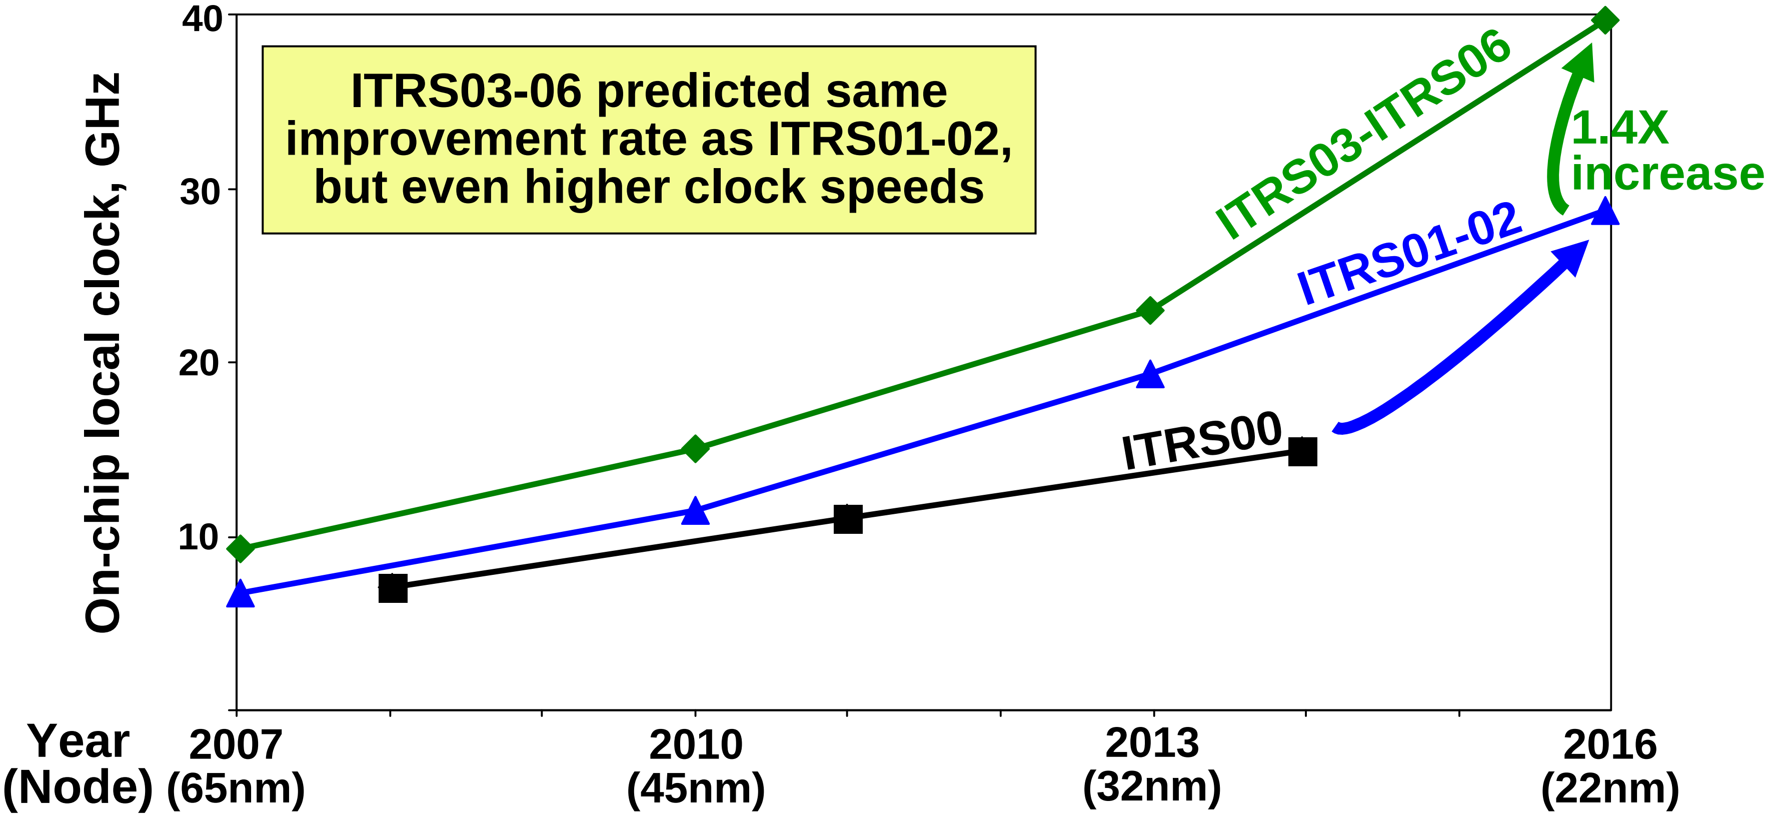

- 2.8X throughput every 3 yrs predicted for multiple independent cores
  - 1.4X clock speed every 3 yrs for constant power (same as ITRS00-02)
  - 2X transistors/chip every 3 yrs for constant chip size (same as ITRS01-02)

2007 International Technology Roadmap for Semiconductors (ITRS07)

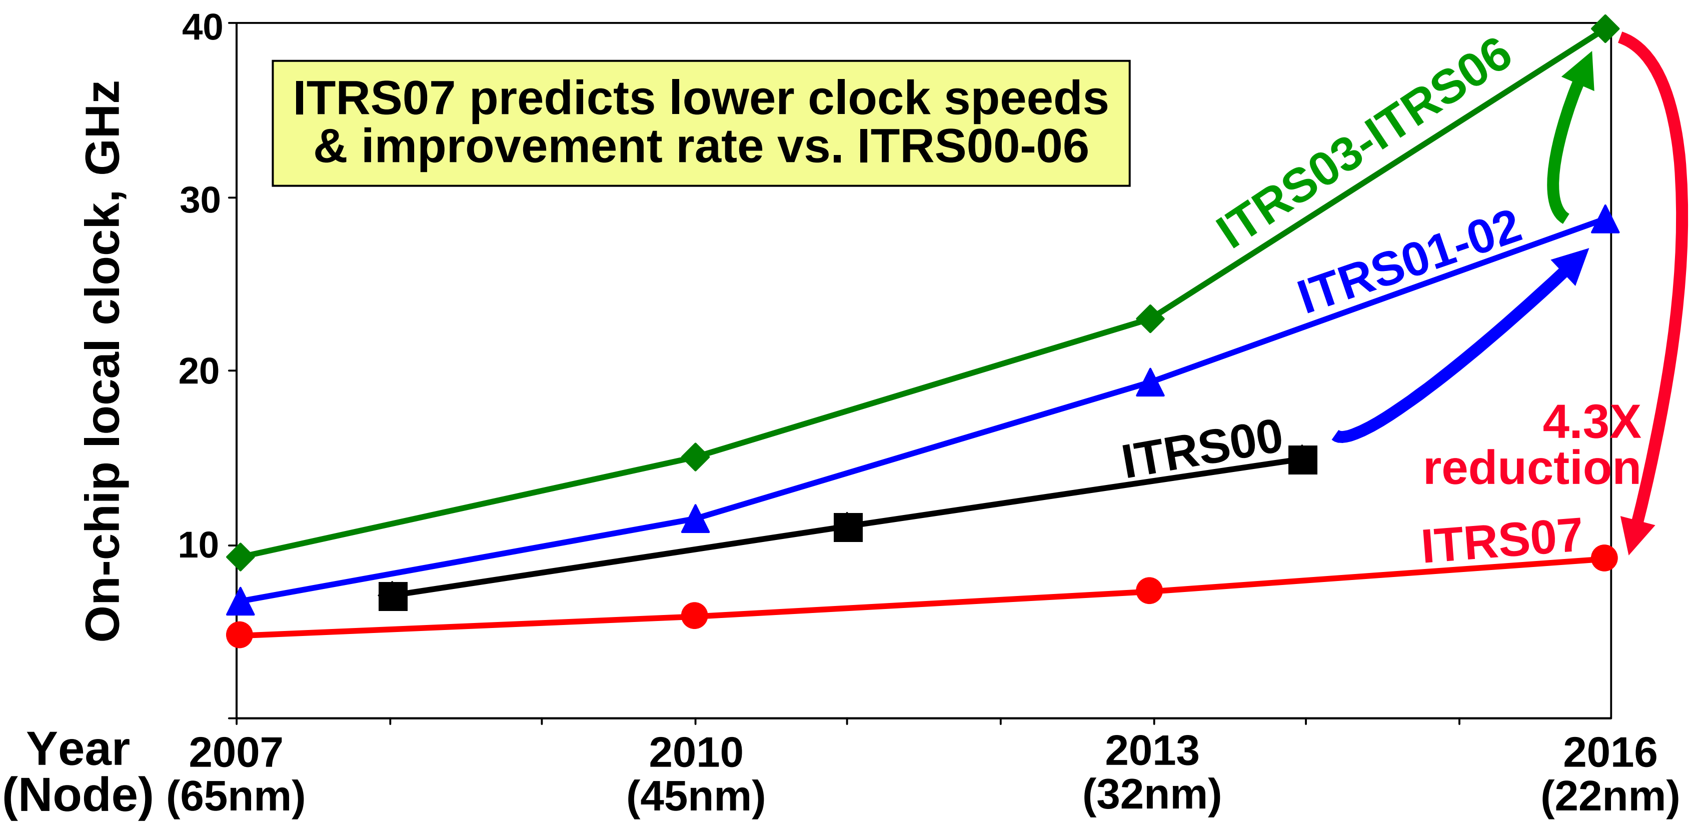

- ~2.5X throughput every 3 yrs predicted for multiple independent cores
  - 1.23X clock speed every 3 yrs for constant power (less than ITRS00-06)
  - 2X transistors/chip every 3 yrs for constant chip size (same as ITRS01-06)

COTS Compute Node (processor, memory & I/O) Performance History & Projections (2Q08)

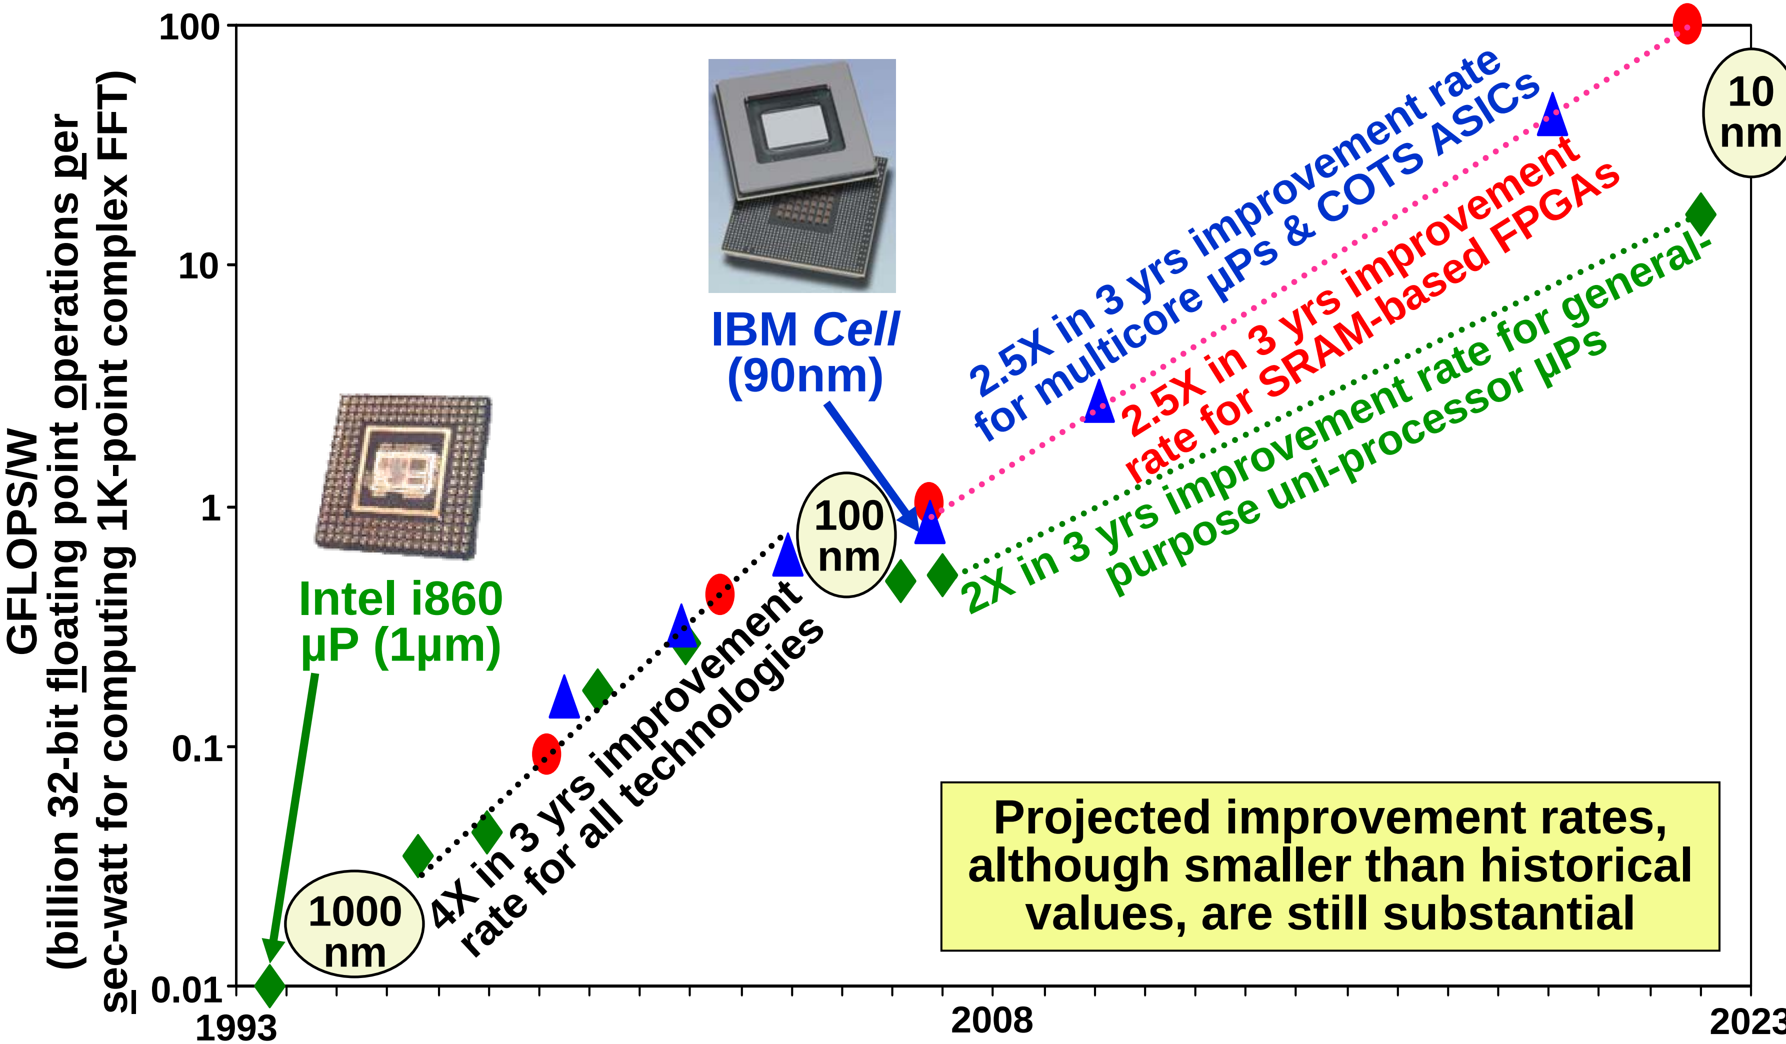

Notional Cost (cumulative) & Schedule for COTS 90nm Cell Broadband Engine

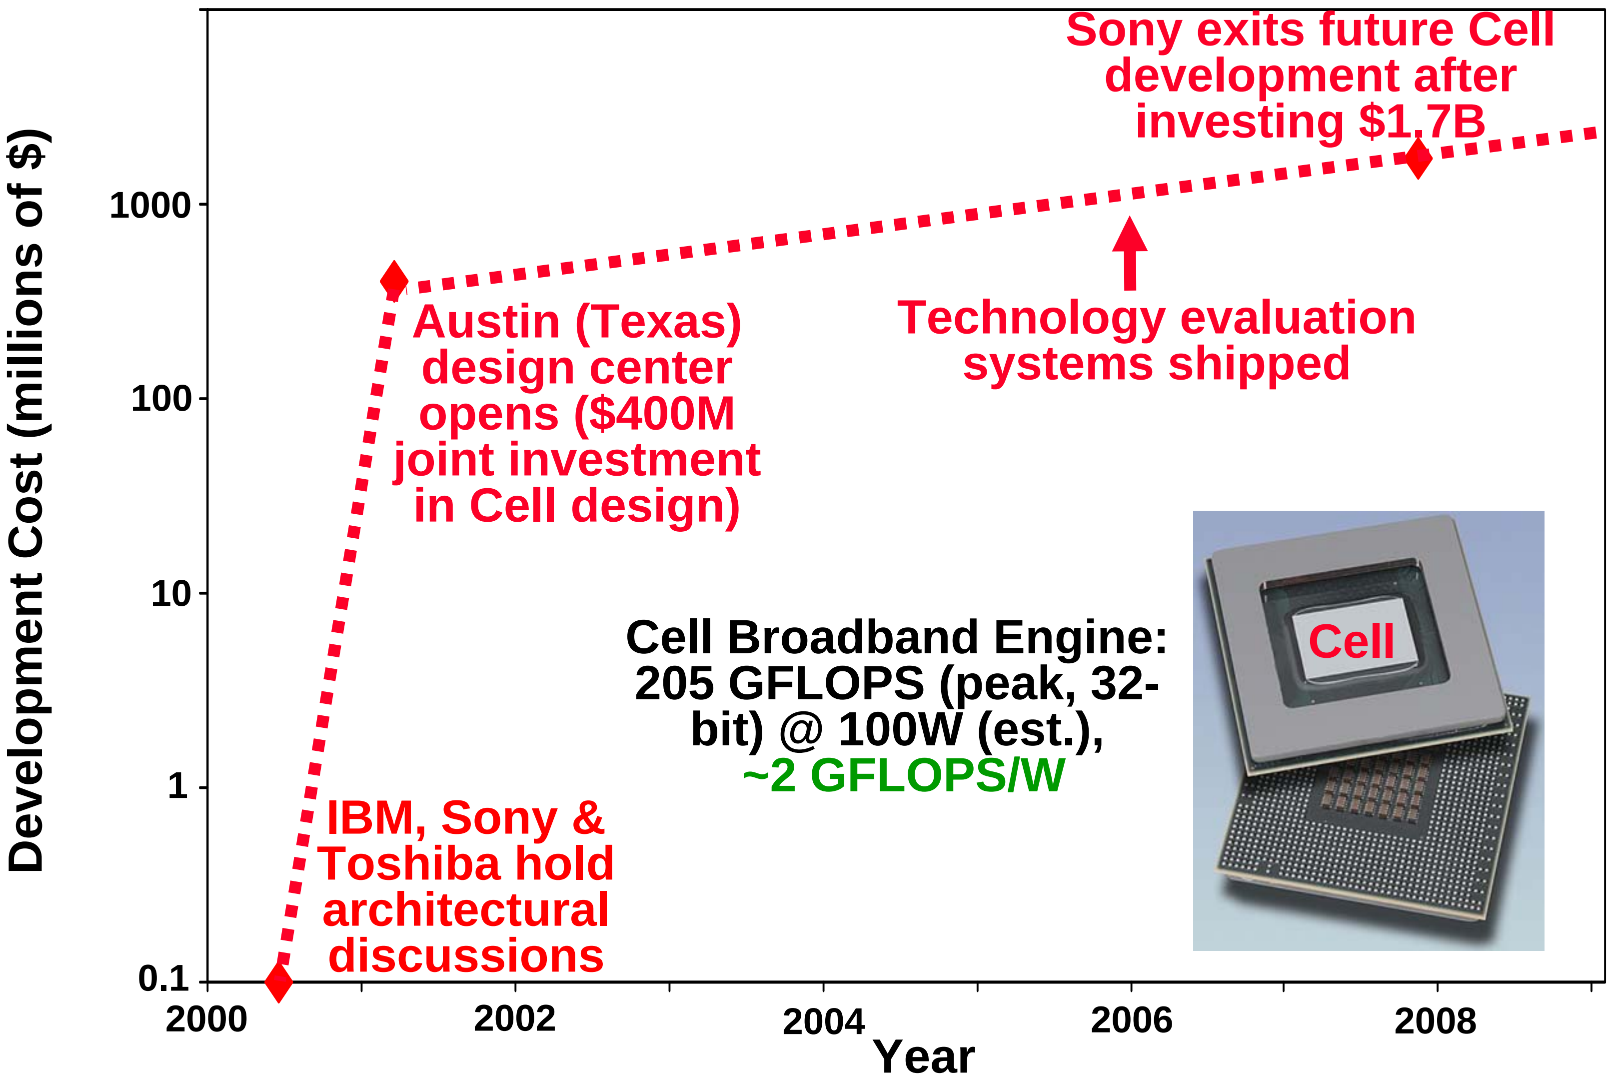

Multicore Open Systems Architecture Example

- **LEON3**
  - 32-bit SPARC V8 processor developed by Gaisler Research (Aeroflex as of 7/14/08) for the European Space Agency
  - Synthesizable VHDL (GNU general public license) & documentation downloadable from www.gaisler.com
  - Open source software support (embedded Linux, C/C++ cross-compiler, simulator & symbolic debugger)

- **0.25µm LEON3FT**
  - *Commercial* fault-tolerant implementation of LEON3
  - 75 MFLOPS/W (150 MIPS & 30 MFLOPS @ 150 MHz for 0.4W)

- **90nm quad-core LEON3FT**
  - System emulated with a single SRAM-based FPGA
  - 133 MFLOPS/W (4x500 MIPS & 4x100 MFLOPS for 3W)
  - Each core occupies <1mm^2 including caches
  - MOSIS fabricates 65nm & 90nm die up to 360mm^2 (IBM process)

How can we improve performance (FLOPS/W), which lags COTS by up to 9 yrs (15X) in this example?

Notional Cost (cumulative) & Schedule for 90nm LEON3FT Multicore Processor

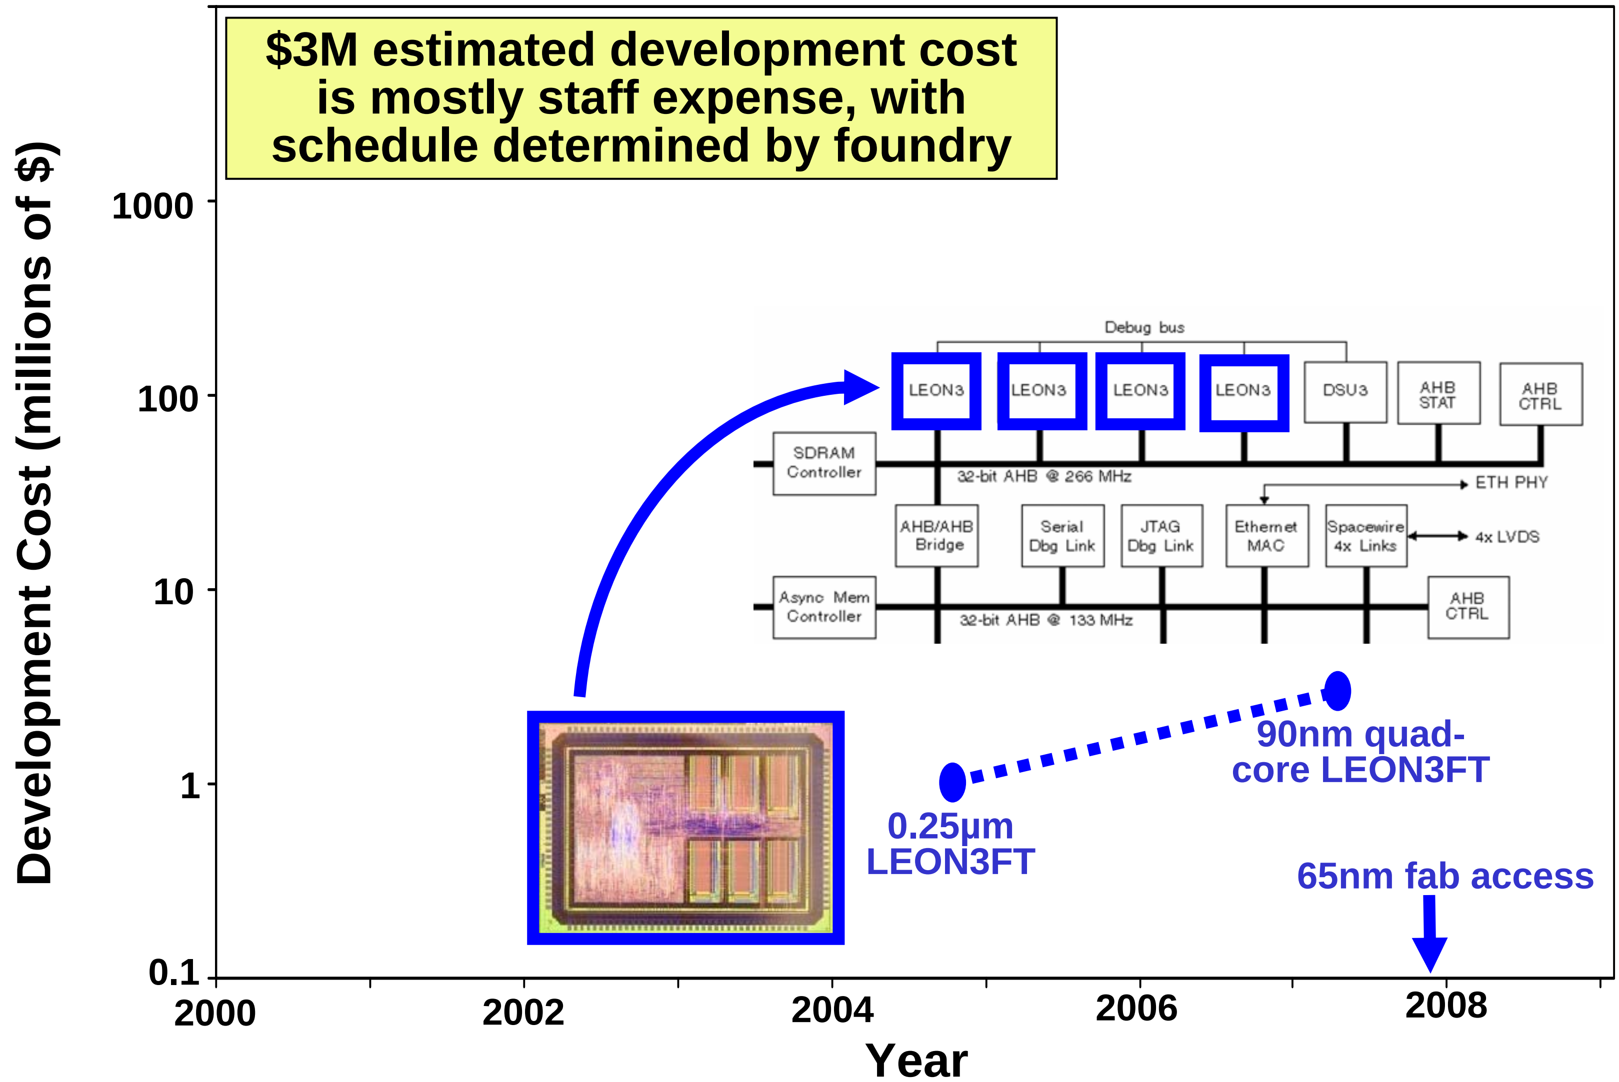

Objective & Schedule

- **Objective: Assess the infrastructure (hardware, software & support) that enables use of multicore open systems architectures**
  - **Where are we now?**
  - **What needs to be done?**

- **Schedule**
  - **1525: Overview**
  - **1540: Guest speaker: Mr. Markus Levy**
  - **1600: Introduction of the panelists**
  - **1605: Previously submitted questions for the panel**
  - **1635: Open forum**
  - **1655: Conclusions & the way ahead**
  - **1700: Closing remarks & adjourn**

Panel Session: *Paving the Way for Multicore Open Systems Architectures*

Moderator: Dr. James C. Anderson
MIT Lincoln Laboratory

Prof. Saman Amarasinghe
MIT Computer Science & Artificial Intelligence Laboratory (CSAIL)

Mr. Markus Levy
The Multicore Association &
The Embedded Microprocessor Benchmark Consortium (EEMBC)

Dr. Steve Muir
Chief Technology Officer
Vanu, Inc.

Dr. Matthew Reilly
Chief Engineer
SiCortex, Inc.

Mr. John Rooks
Air Force Research Laboratory (AFRL/RITC)
Emerging Computing Technology

Panel members & audience may hold diverse, evolving opinions

Objective & Schedule

- **Objective: Assess the infrastructure (hardware, software & support) that enables use of multicore open systems architectures**
  - **Where are we now?**
  - **What needs to be done?**

- **Schedule**
  - **1525: Overview**
  - **1540: Guest speaker: Mr. Markus Levy**
  - **1600: Introduction of the panelists**
  - **1605: Previously submitted questions for the panel**
  - **1635: Open forum**
  - **1655: Conclusions & the way ahead**
  - **1700: Closing remarks & adjourn**

Conclusions & The Way Ahead

- **Despite industry slowdown, embedded processors are still improving exponentially (2/3 of historical Moore's Law rate)**
- **Although performance improvements in multicore designs (2.5X every 3 yrs) continue to outpace those of uni-processors (2X every 3 yrs), the "performance gap" is less than previously projected**
- **New tools and methodologies will be needed to maximize the benefits of using multicore open systems architectures**
  - **Power & packaging issues**
  - **Cost & availability issues**
  - **Training & ease-of-use issues**
  - **Platform independence issues**
- **Although many challenges remain in reducing the performance gap between highly specialized systems vs. multicore open systems architectures, the latter will help insulate users from manufacturer-specific issues**

Success still depends on ability of foundries to provide smaller geometries & increasing speed for constant power (driven by large-scale COTS product economics)

Backup Slides

COTS ASIC: 90nm IBM *Cell Broadband Engine* (4Q06)

- 100W (est.) @ 3.2 GHz
- 170 GFLOPS sustained for 32-bit flt pt 1K cmplx FFT (83% of peak)
- 16 Gbyte memory options (~10 FLOPS/byte)
  - COTS Rambus XDR DRAM (Cell is designed to use only this memory)
    - 256 chips
    - 690W (note: Rambus devices may not be 3D stackable due to 2.7W/chip power consumption)
  - Non-COTS solution: Design a bridge chip ASIC (10W est.) to allow use of 128 DDR2 SDRAM devices (32W)
    - 128 chips in 3D stacks to save space (0.25W/chip)
    - Operate many memory chips in parallel
    - Buffer to support Rambus speeds
    - Increased latency vs. Rambus
- 40W budget for external 27 Gbytes/sec simultaneous I&O (using same non-COTS bridge chip to handle I/O with Cell)
- Single non-COTS CN (compute node) using DDR2 SDRAM
  - 170 GFLOPS sustained for 200W (182W est. for CN plus 18W for 91% efficient DC-to-DC converter)
  - 0.85 GFLOPS/W & 56 GFLOPS/L

COTS Compute Node Performance History & Projections (2Q08)

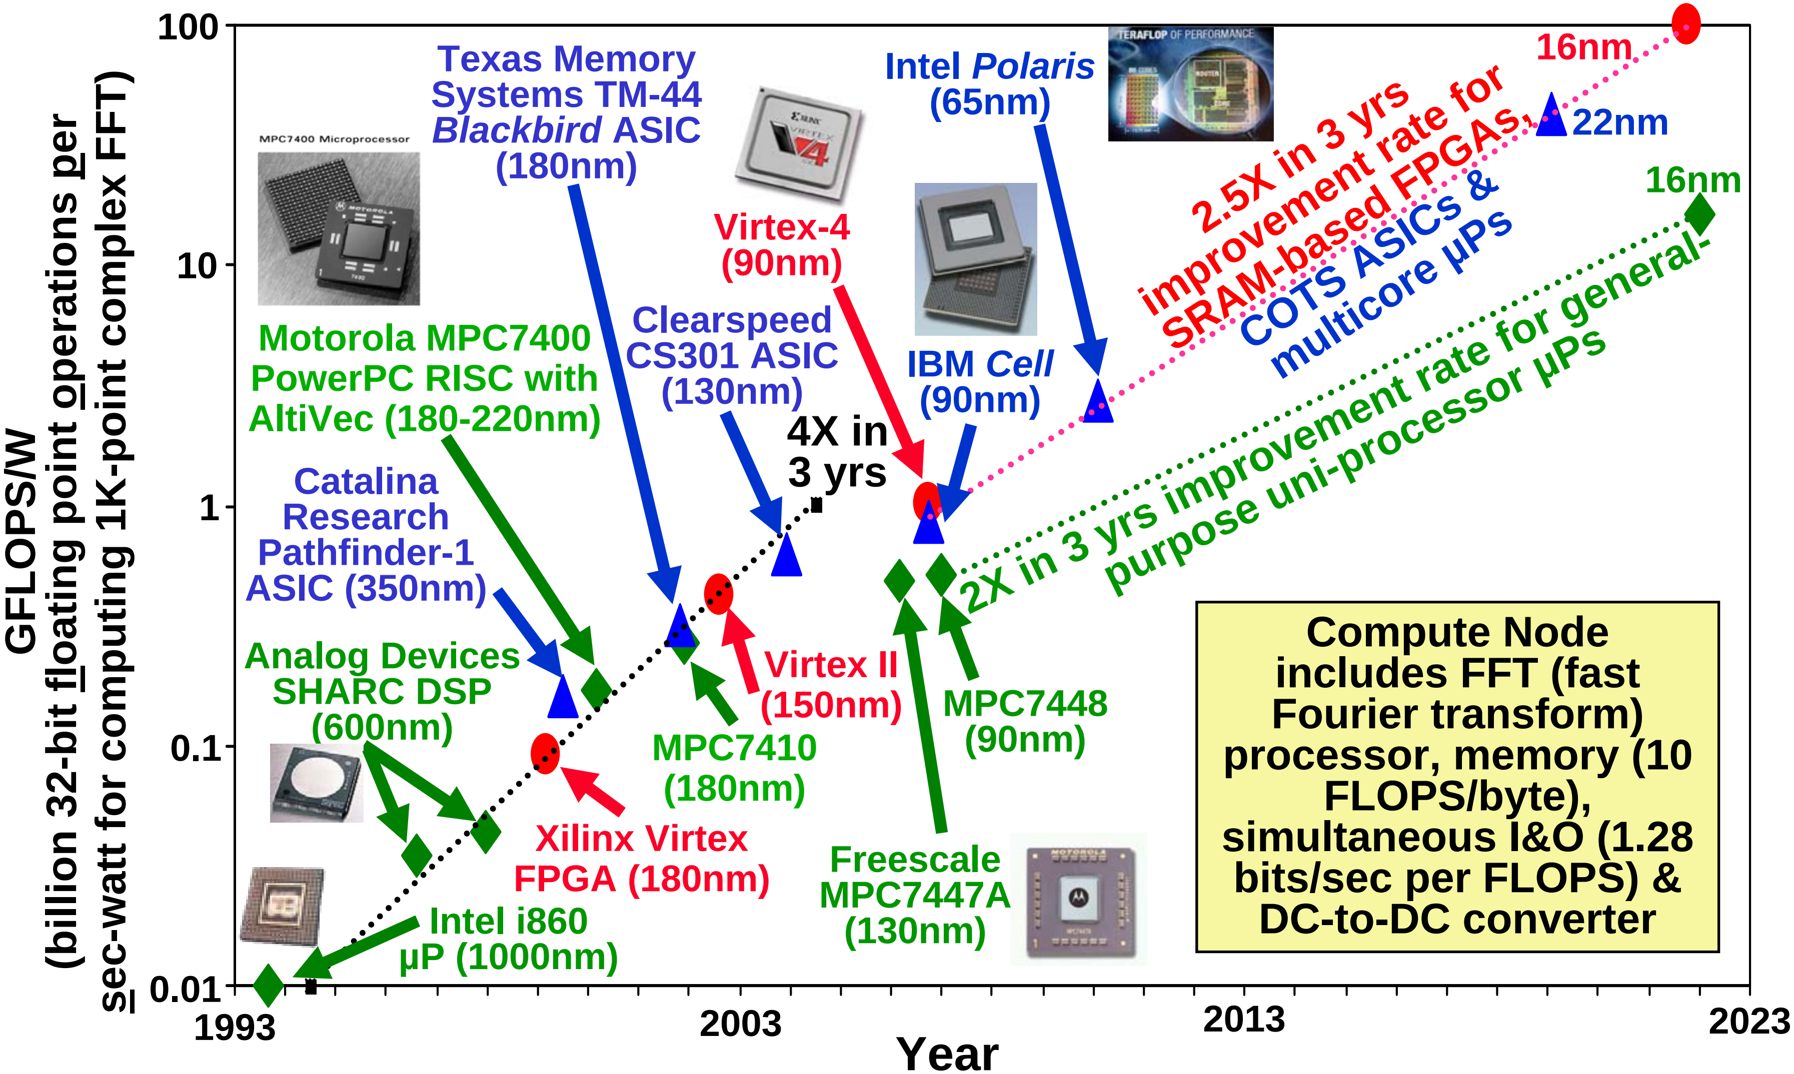

World’s Largest Economies: 2000 vs. 2024

2000 GDP*

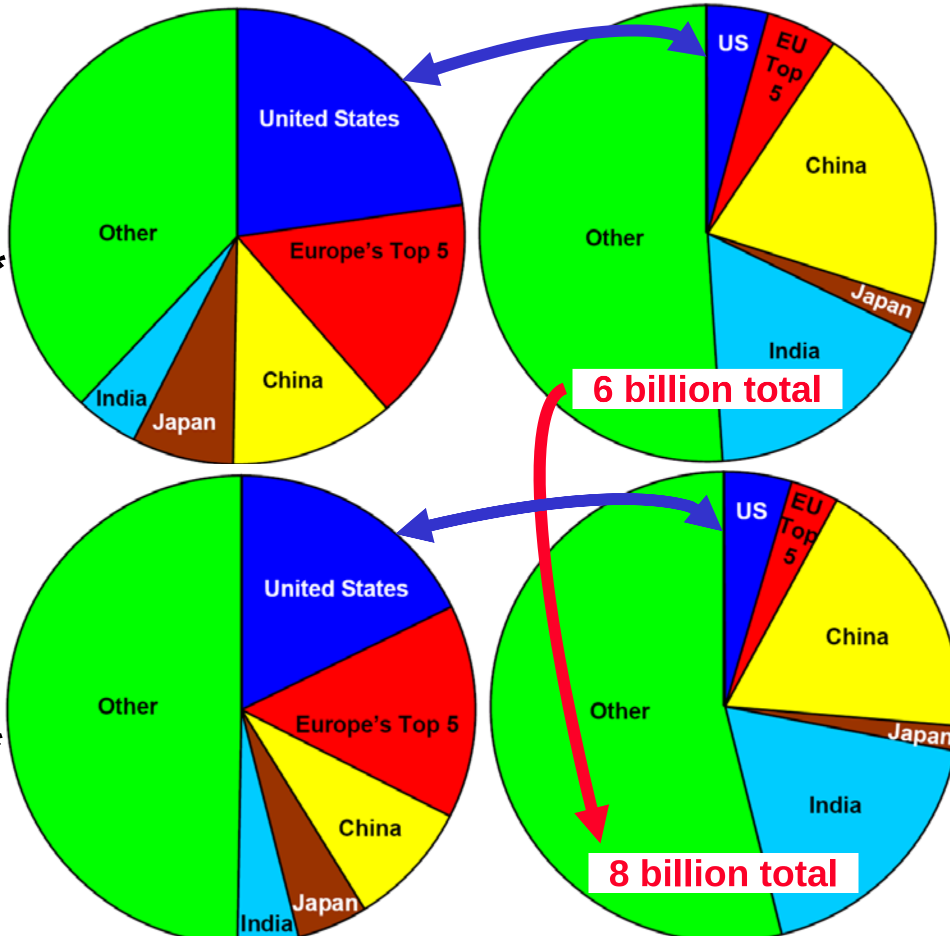

2000 population

2024 GDP*

2024 population

* Gross domestic product (purchasing power parity)

“Europe’s Top 5” are Germany, Great Britain, France, Italy & Spain

U.S. population grows by 1/3 & income shrinks from 5X to <4X world average

Highest-performance COTS (commercial off-the-shelf) ADCs (analog-to-digital converters), 3Q08

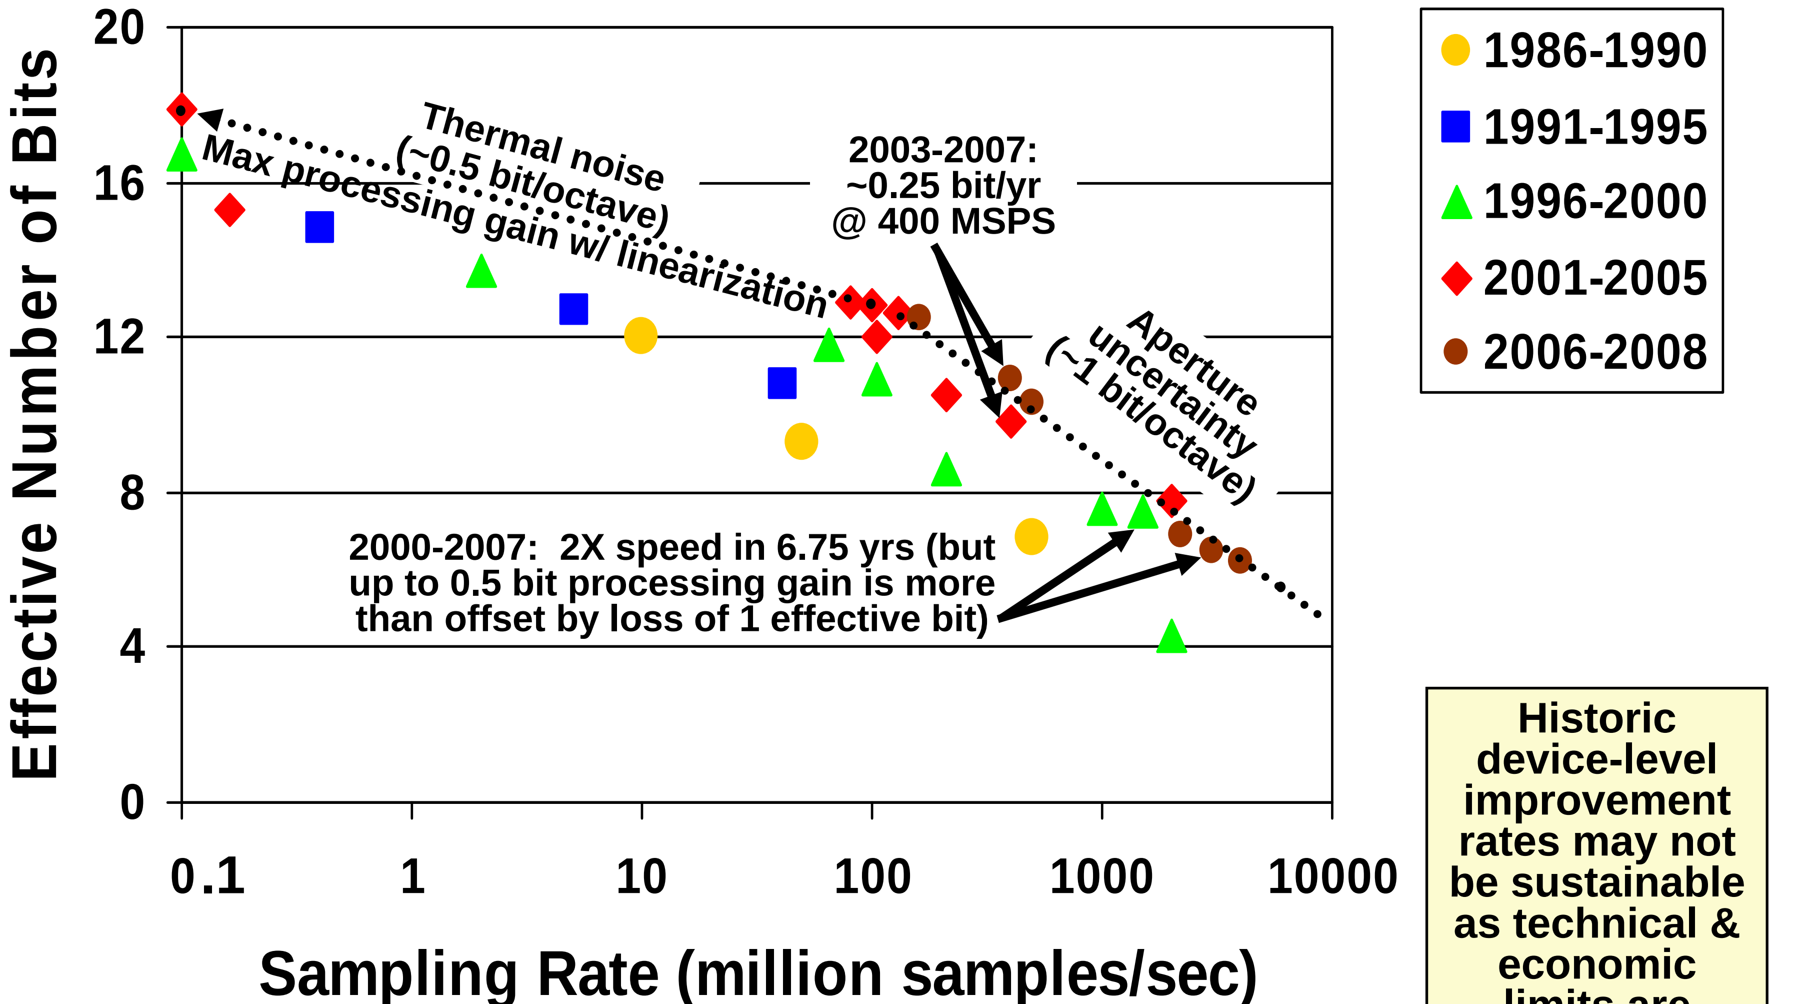

Historic device-level improvement rates may not be sustainable as technical & economic limits are approached

SFDR (spur-free dynamic range) for Highest-performance COTS ADCs, 3Q08

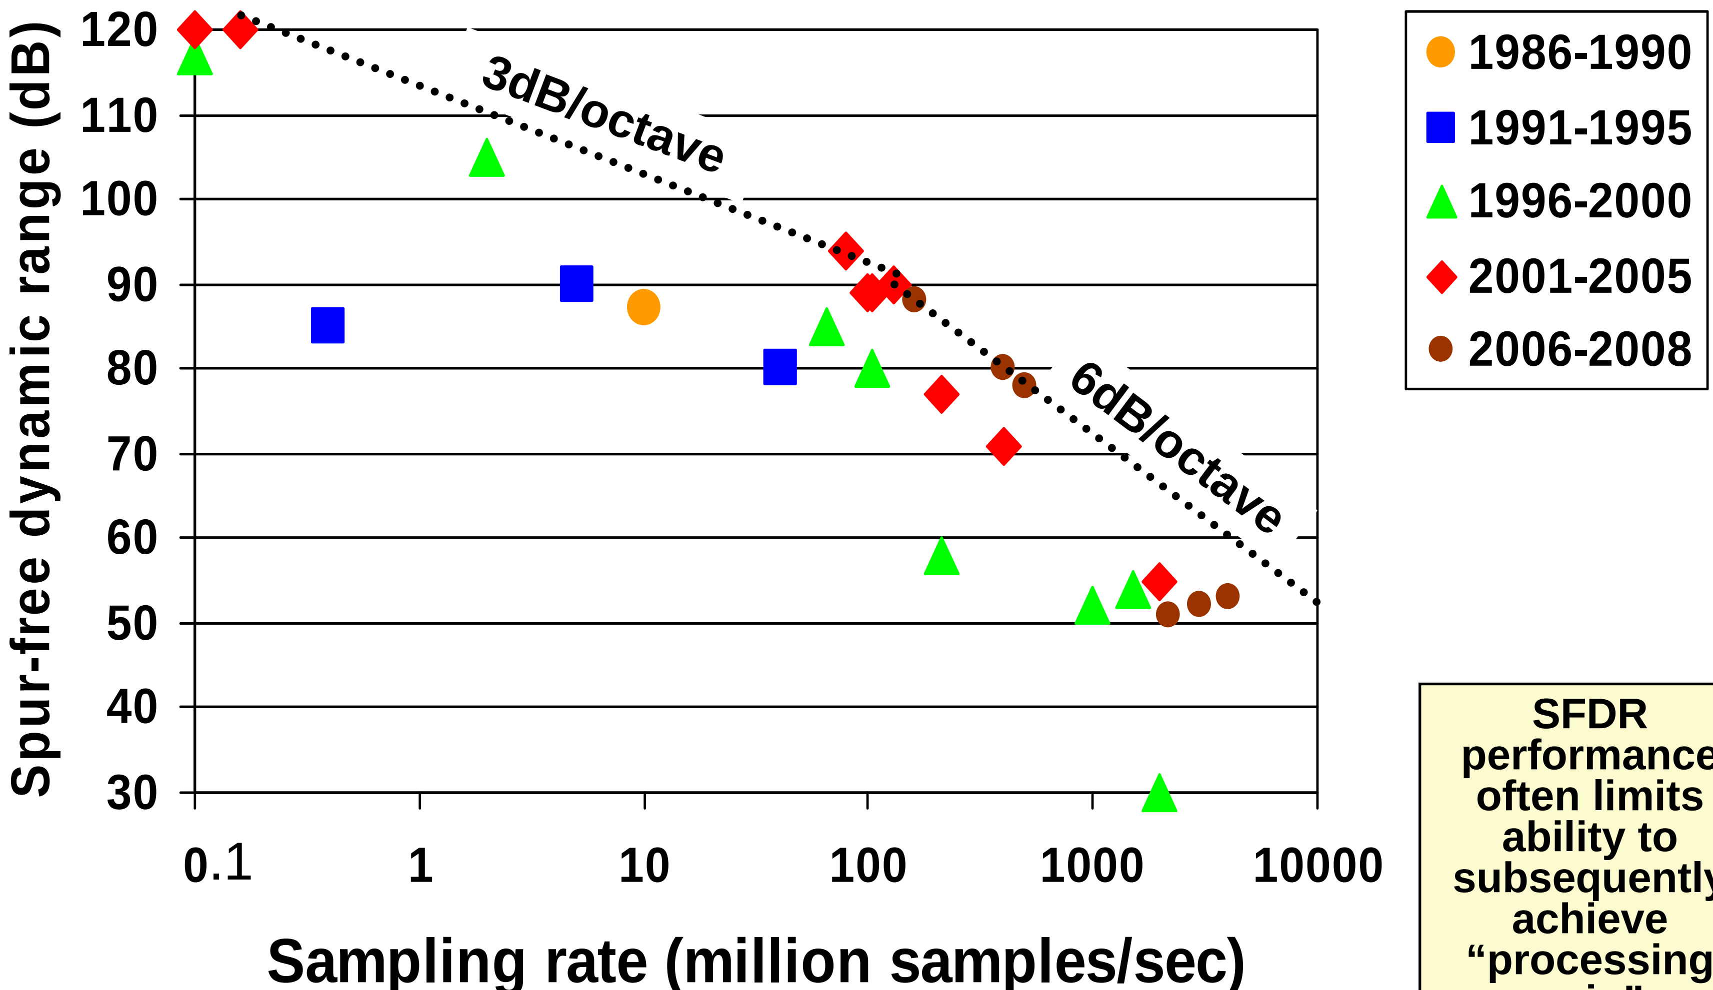

SFDR performance often limits ability to subsequently achieve "processing gain"

Energy per Effective Quantization Level for Highest-performance COTS ADCs, 3Q08

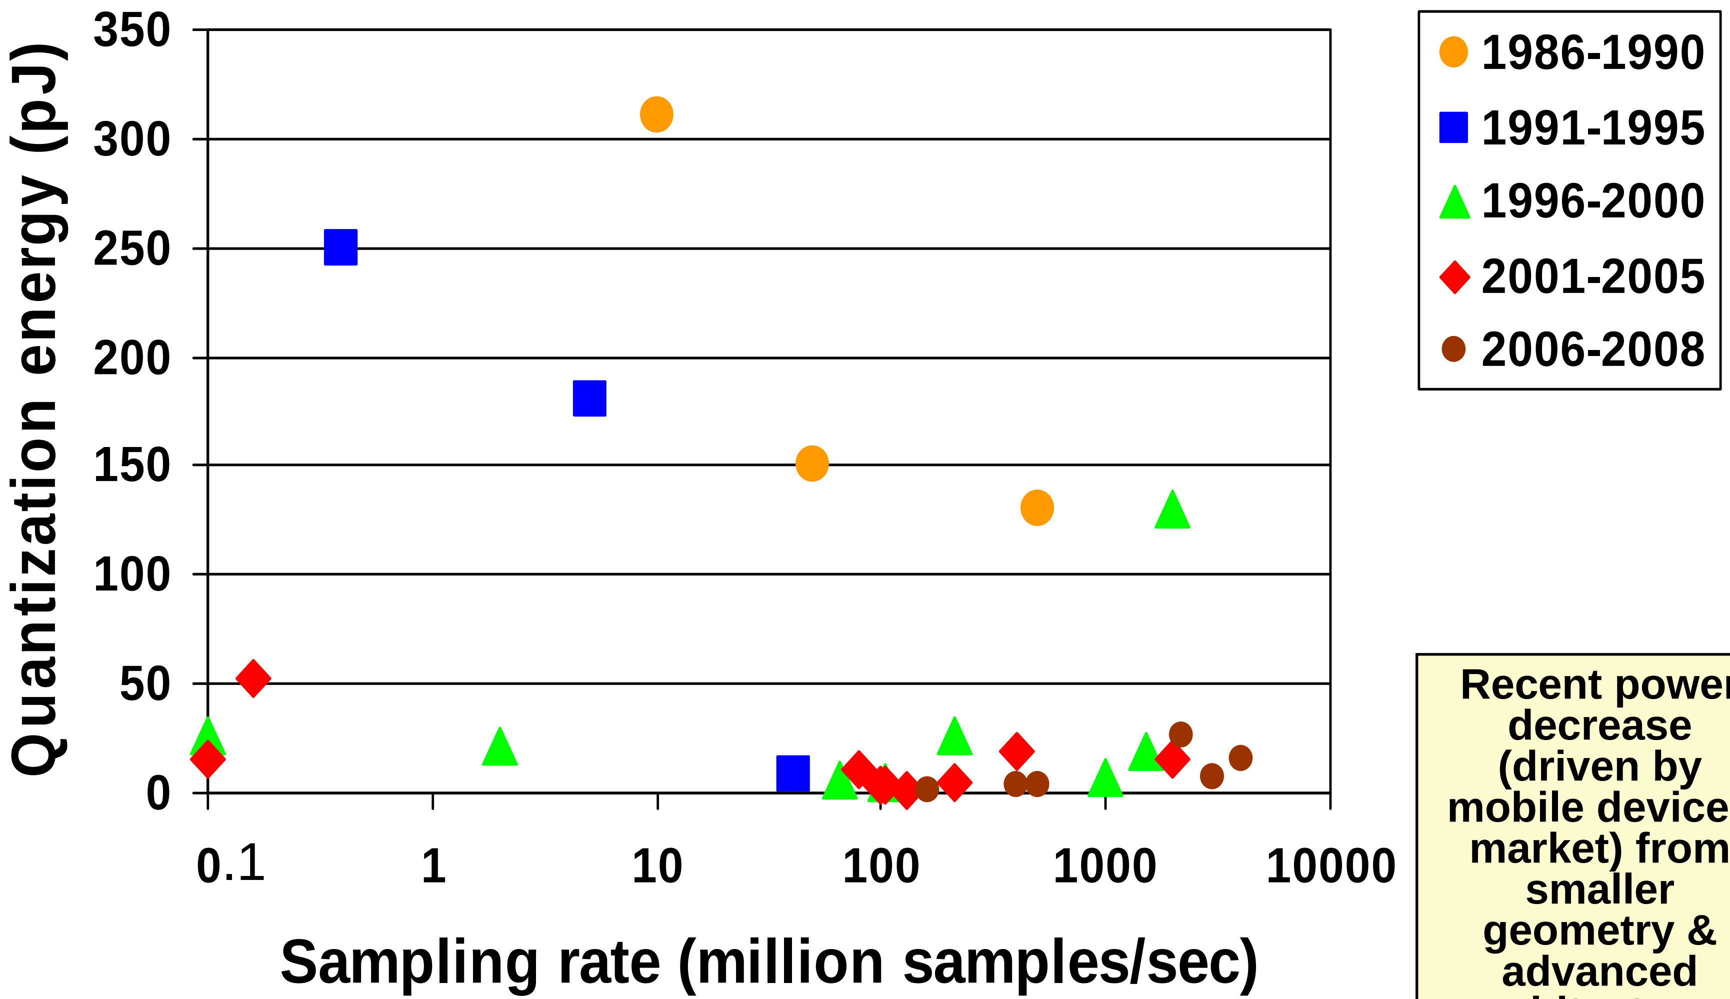

Recent power decrease (driven by mobile devices market) from smaller geometry & advanced architectures

Resolution Improvement Timeline for Highest-performance COTS ADCs, 1986-2008