Architectures for the Future

Doug Burger	University of Texas
Roy Ju	Intel
Ravi Nair	IBM
Kunle Olukotun	Stanford
Eric Rotenberg	North Carolina State
Erik Altman	IBM

ISPASS Panel *March 21, 2005*

Outline

- Trends
- Questions
- Provocations

Growth of Die Size

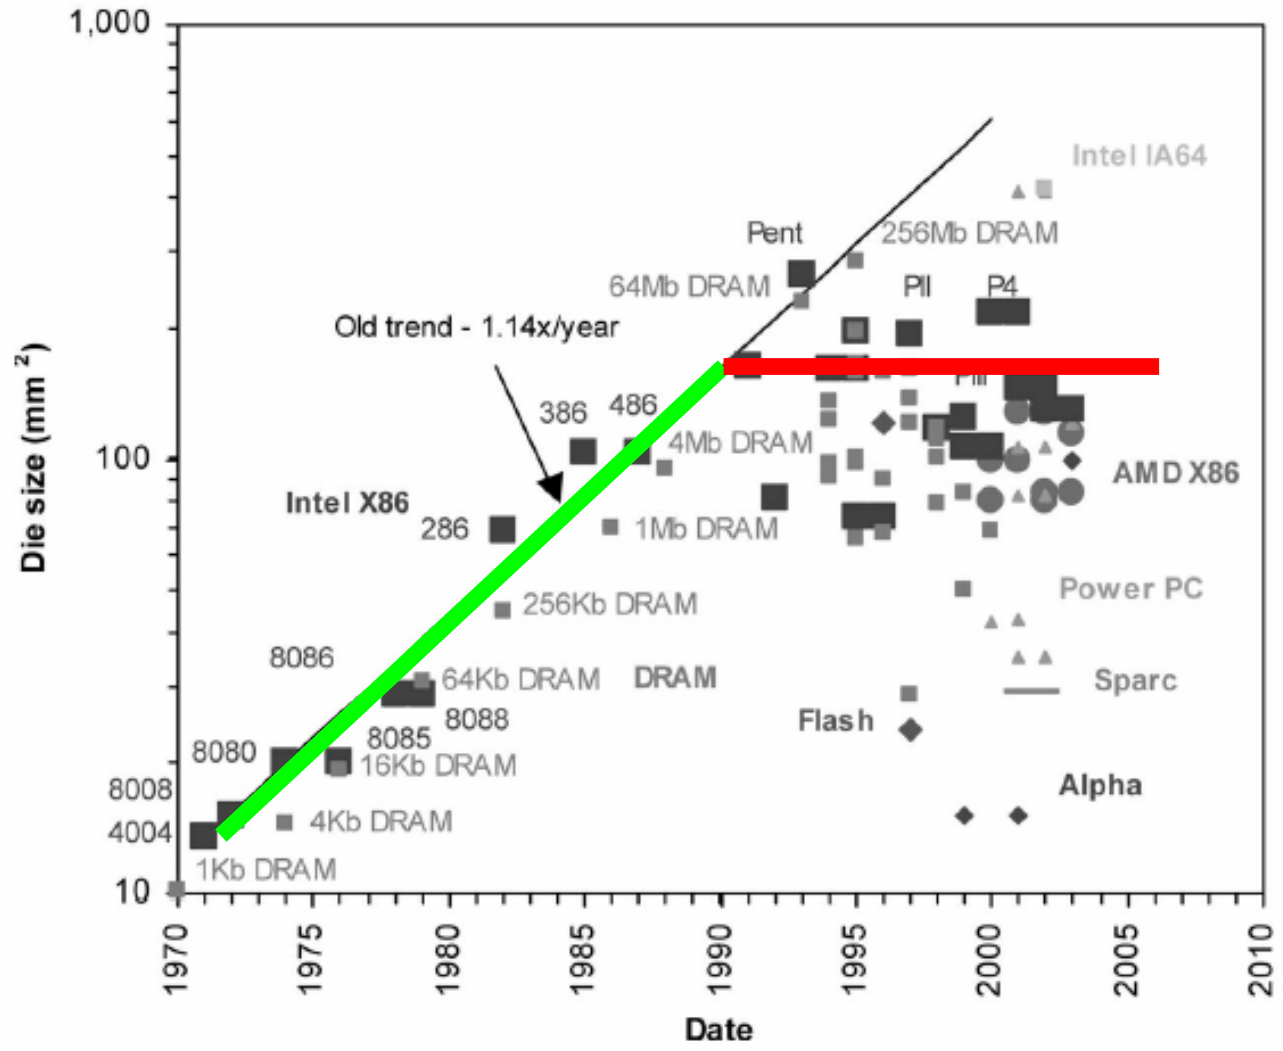

Fig. 4. Growth in die sizes over time

http://www.taoli.ece.ufl.edu/teaching/spring05/eel6935_schedule_pdf/Stojcev.pdf

Growth in Frequency

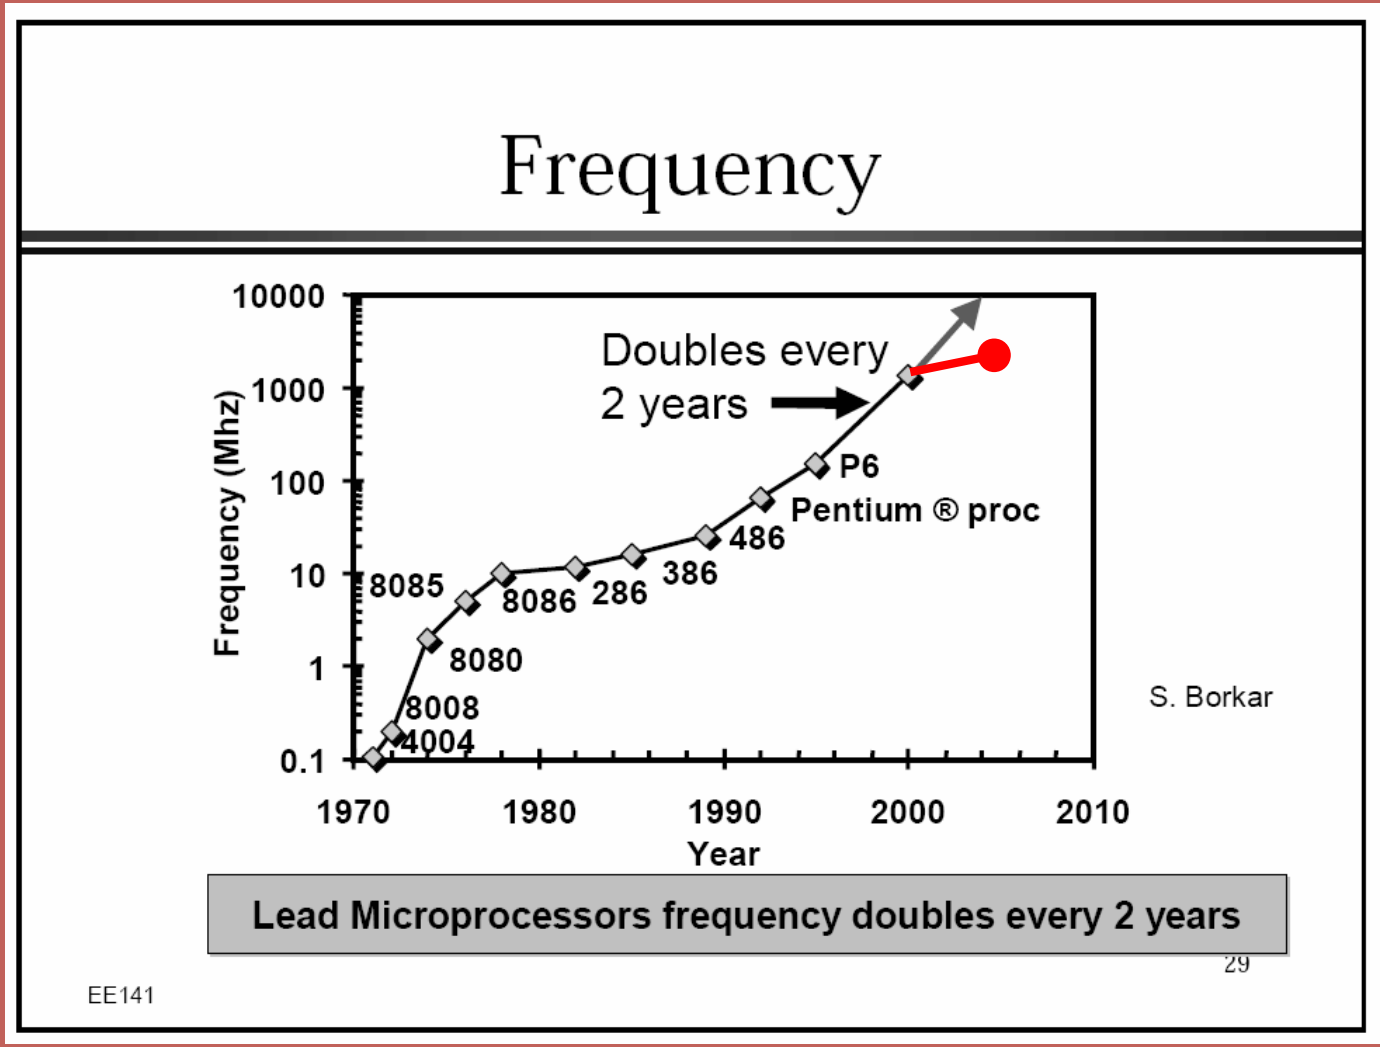

http://bwrc.eecs.berkeley.edu/Classes/icdesign/ee141_s03/Lectures/Lecture1_Introduction.pdf

Power and Performance vs F04

Srinivasan et al. MICRO-35, 11/2002

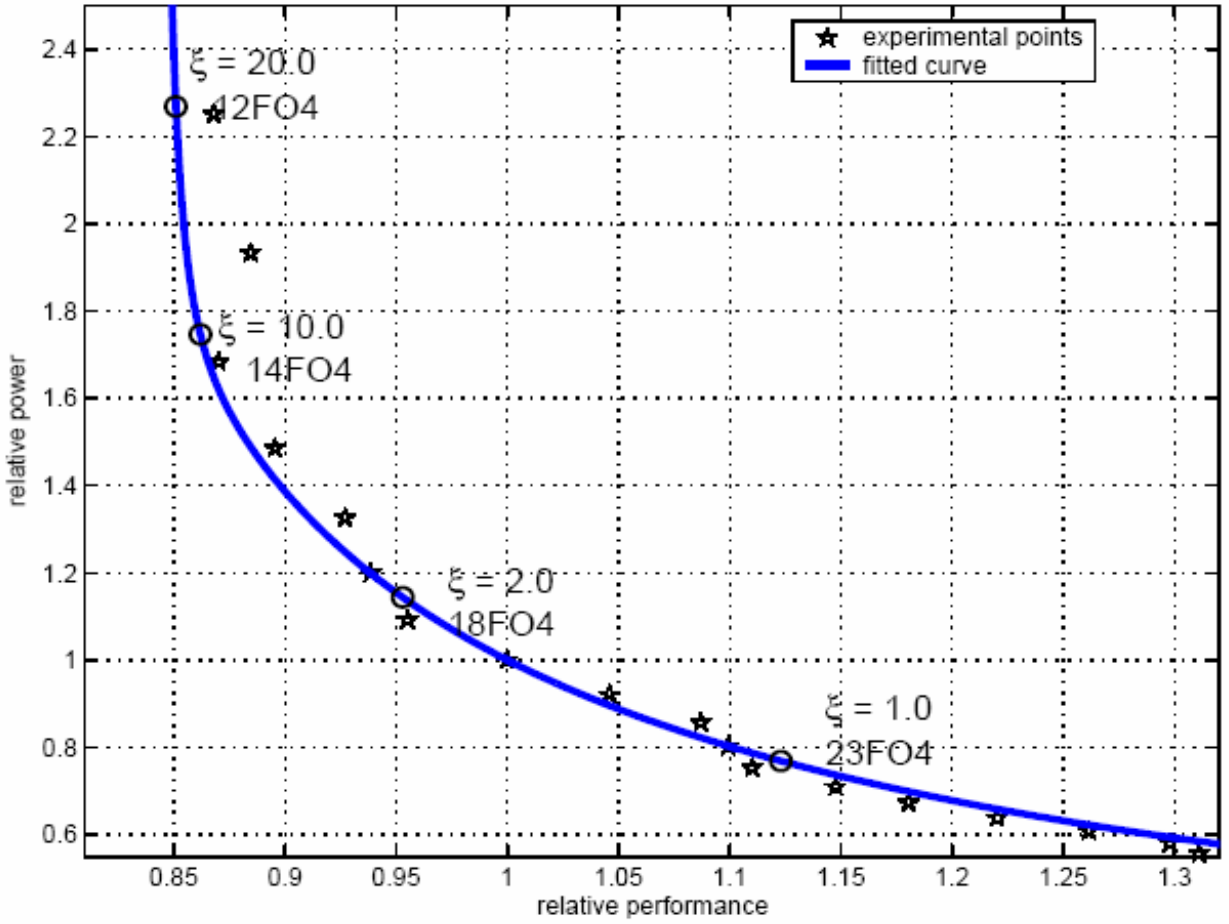

Performance

http://www.eecs.harvard.edu/~dbrooks/micro2003-tutorial-final.pdf

Effect of Critical Paths

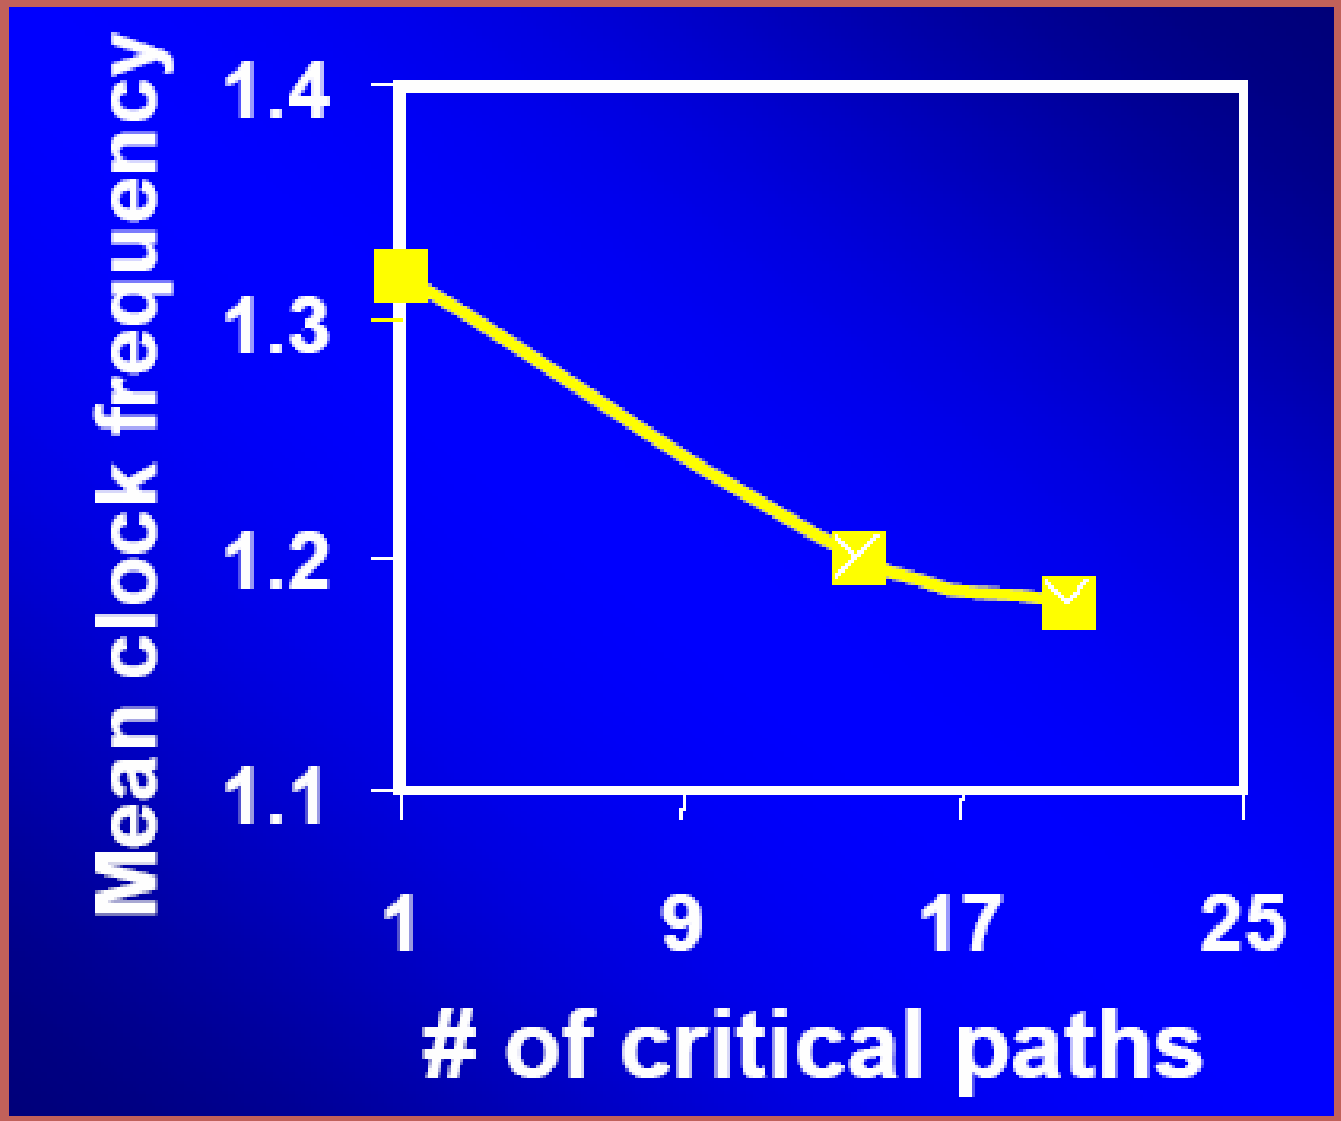

http://www.microarch.org/micro37/presentations/MICRO37%20Sborkar.pdf

Growth in Complexity & Productivity

Productivity Trends

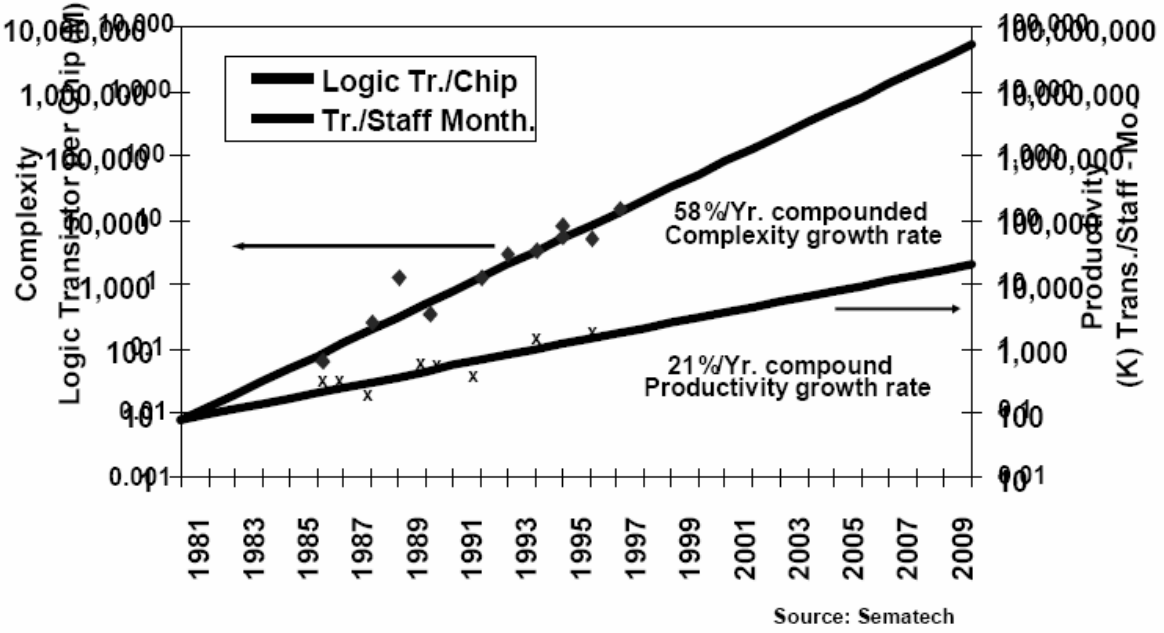

Source: Sematech

Complexity outpaces design productivity

http://bwrc.eecs.berkeley.edu/Classes/icdesign/ee141_s03/Lectures/Lecture1_Introduction.pdf

Growth of Stepper Cost

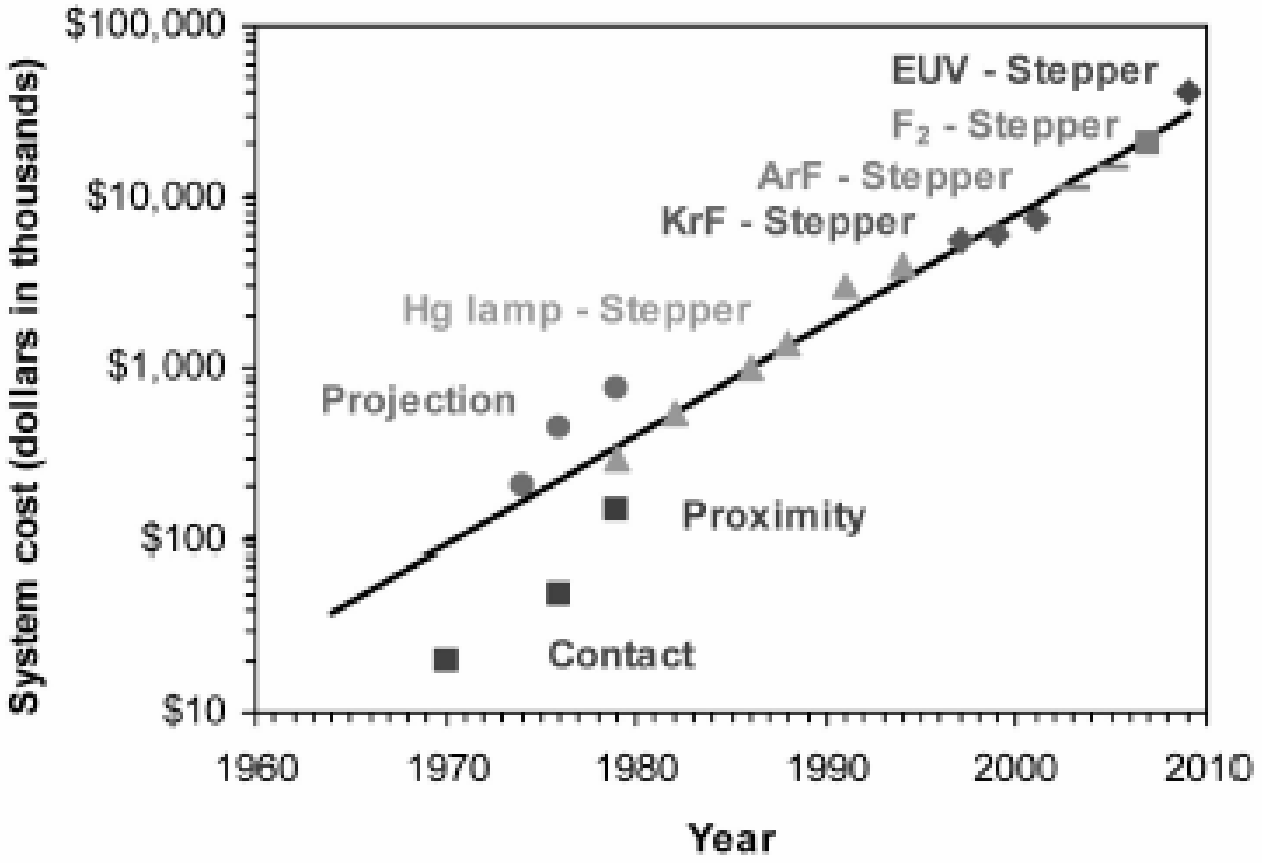

Fig. 6. System cost over time

http://www.taoli.ece.ufl.edu/teaching/spring05/eel6935_schedule_pdf/Stojcev.pdf

Growth of Onchip Cache

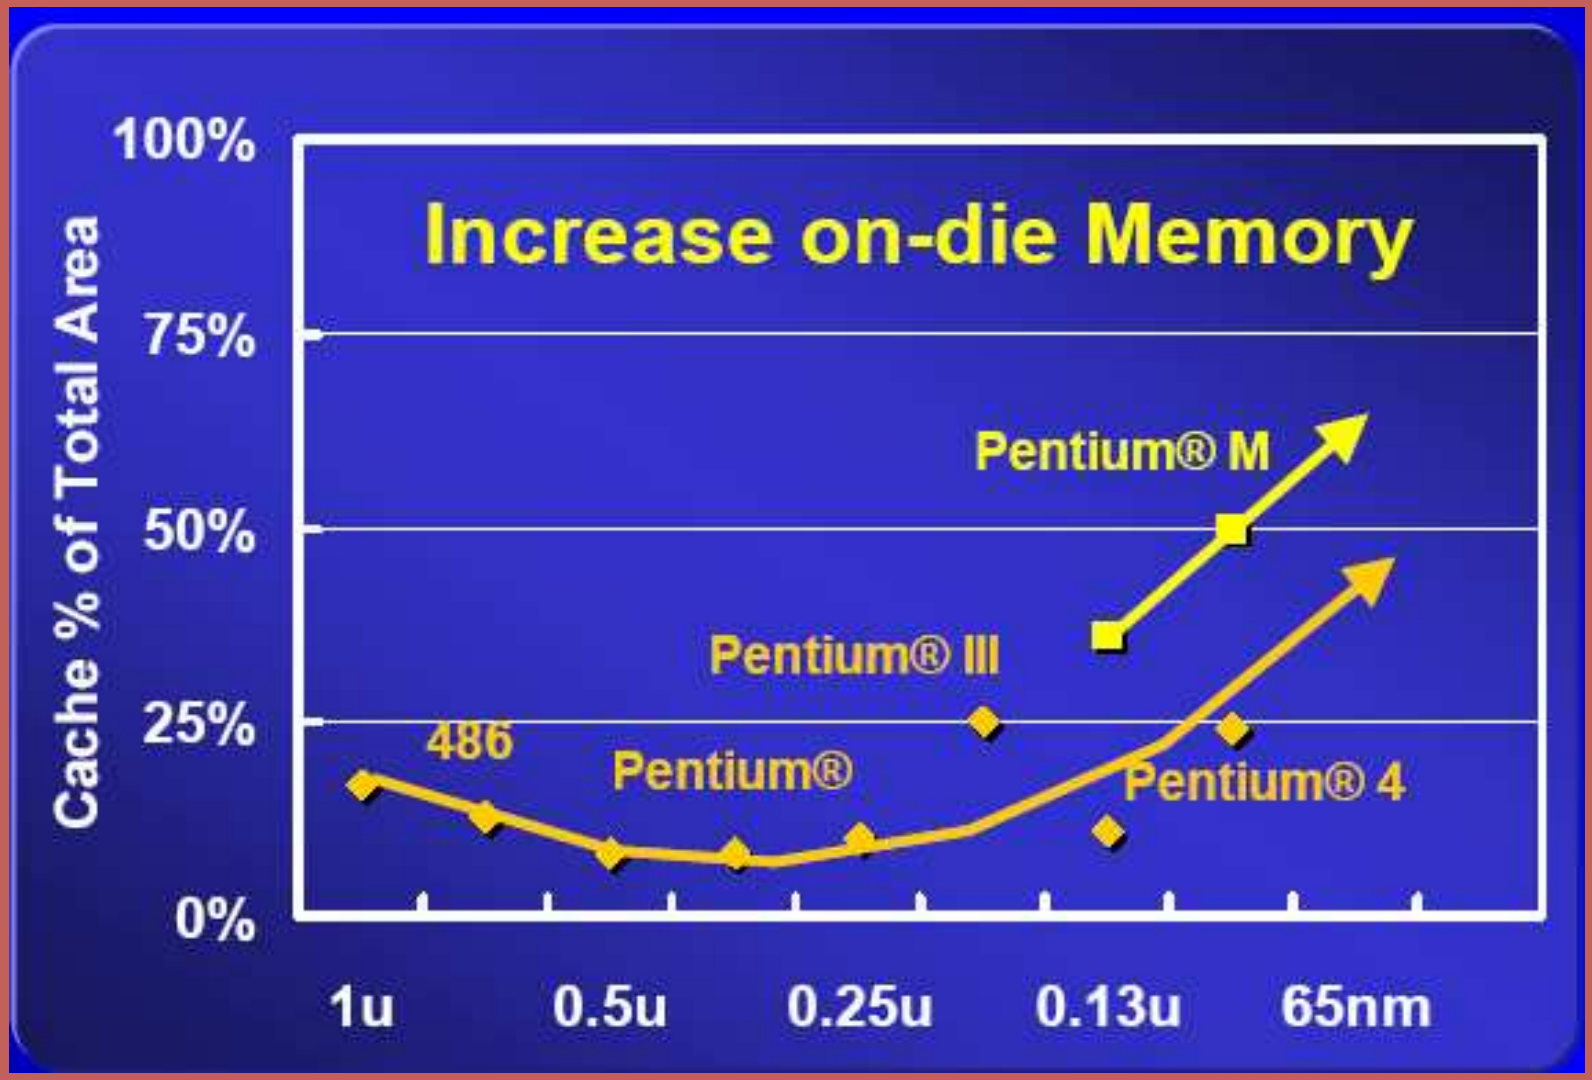

http://www.microarch.org/micro37/presentations/MICRO37%20Sborkar.pdf

Impact of Security?

- Can architecture help improve security and reliability?

Impact of Apps on Architecture?

- What is general purpose computing?
  - Should ISA / microarchitecture be different for DSP, embedded, desktop, games, scientific, commercial transaction processing?
- Are there any new killer apps emerging that will drive new architecture?
  - Speech recognition
  - Games
  - Pervasive / Sensor based: Everything connected

Role for Specialization?

- Will mask costs *decrease* the number of ASICs, or will energy efficiency *increase* the number of ASICs?
- Will processors employ specialized accelerators to help with increasing power/performance problems?
  - TCP/IP
  - Encryption
  - XML parsing
  - Speech Recognition
  - SPAM Filters
- Will reconfigurable hardware be used in mainstream processors?

Impact of Compiler and OS?

- Will compilers and operating systems be able to support more irregular architectural features:
  - Accelerators
  - SIMD
  - XY memories
  - Circular buffers
- Will compilers substantially improve performance over current optimization levels?

Impact of Multiple Threads?

- Can SMT/CMP improve single thread performance?
- What types of computing naturally have multiple threads?

Provocations

- Uniprocessor performance does not need to get any better.
- Uniprocessor performance cannot get much better.
- Current programming practices make it too hard to generate parallel code from a single thread of execution.
- Parallelism will always be the next big thing.
- Given power limitations and the memory wall, we would get better system performance if we went back to 500 MHz Chips.
- Architecture lives: Accelerators for new applications, control independence, polymorphism, virtual machines, thread level speculation, etc. point to a resurgence of architecture.