Emerging Challenges for MP-SoC Platforms

Pierre Paulin, Director
SoC Platform Automation, AST
STMicroelectronics, Ottawa, Canada

Outline

- Market Trends
- Technology Trends
- Surviving 45 nm and beyond

Key SoC Market Trends

- Continued reduction in **time-to-market**
  - Fast changing specs, requirements

- Increasing **cost** of SoC platform design
  - 10M$ to 100 M$ for today's 90nm SoC's

- Need to increase **time-*in*-market**

- Implies higher **flexibility** is needed
  - Rising proportion of function in eS/W
  - Currently 50% to 75% of design costs

Emb. Systems Prog. Survey 2005: Number of Processors per chip

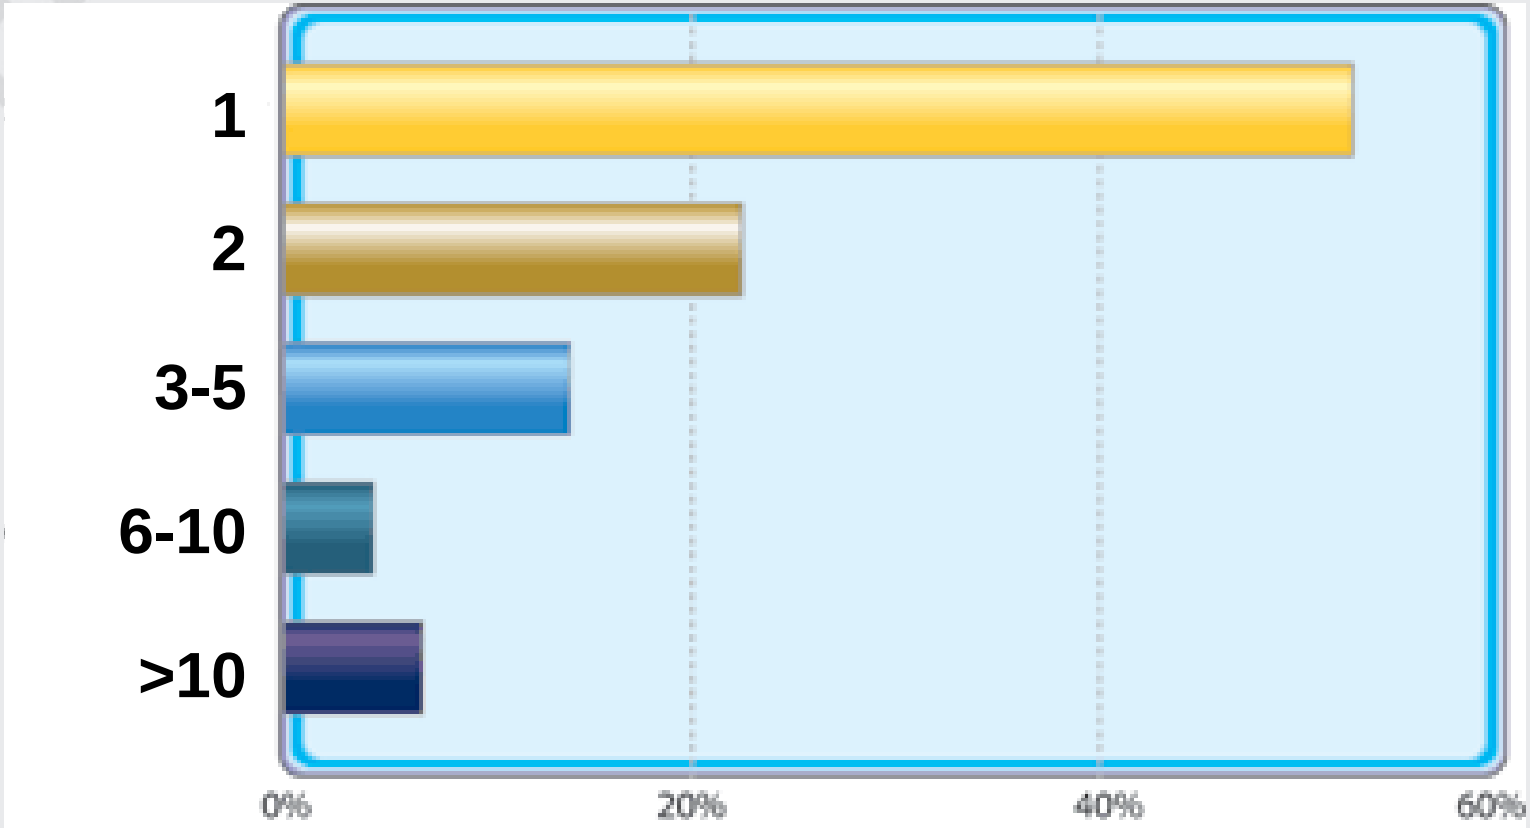

- Nearly 50% of chips use multiple processors
- Over 100 projects used >10 processors

Source: Embedded Systems Programming Magazine, 2005

Processor Heterogeneity

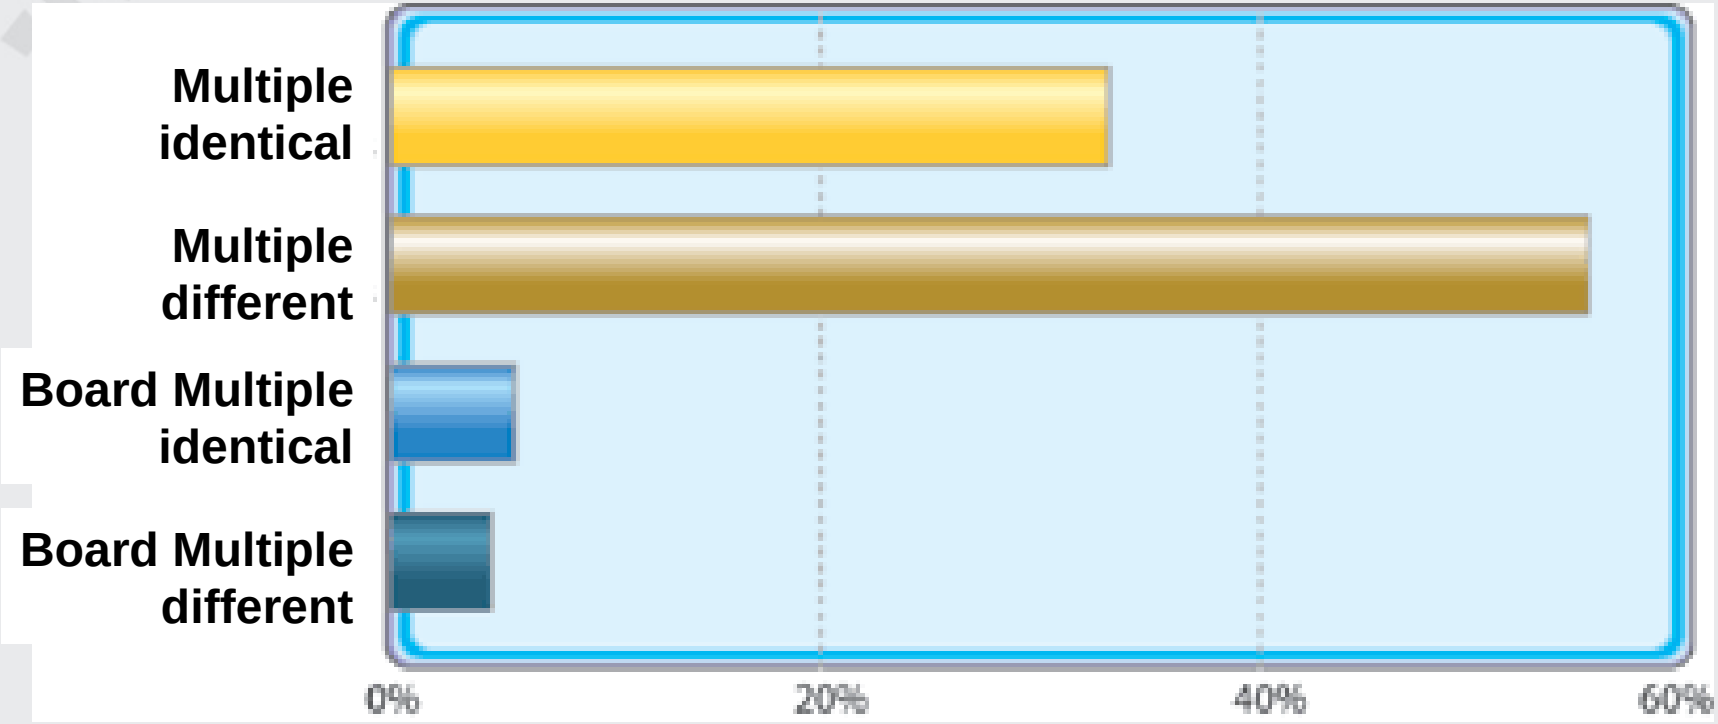

- Nearly 2/3 of SoC's are heterogeneous MP

Source: Embedded Systems Programming Magazine, 2005

Processor Selection Criteria

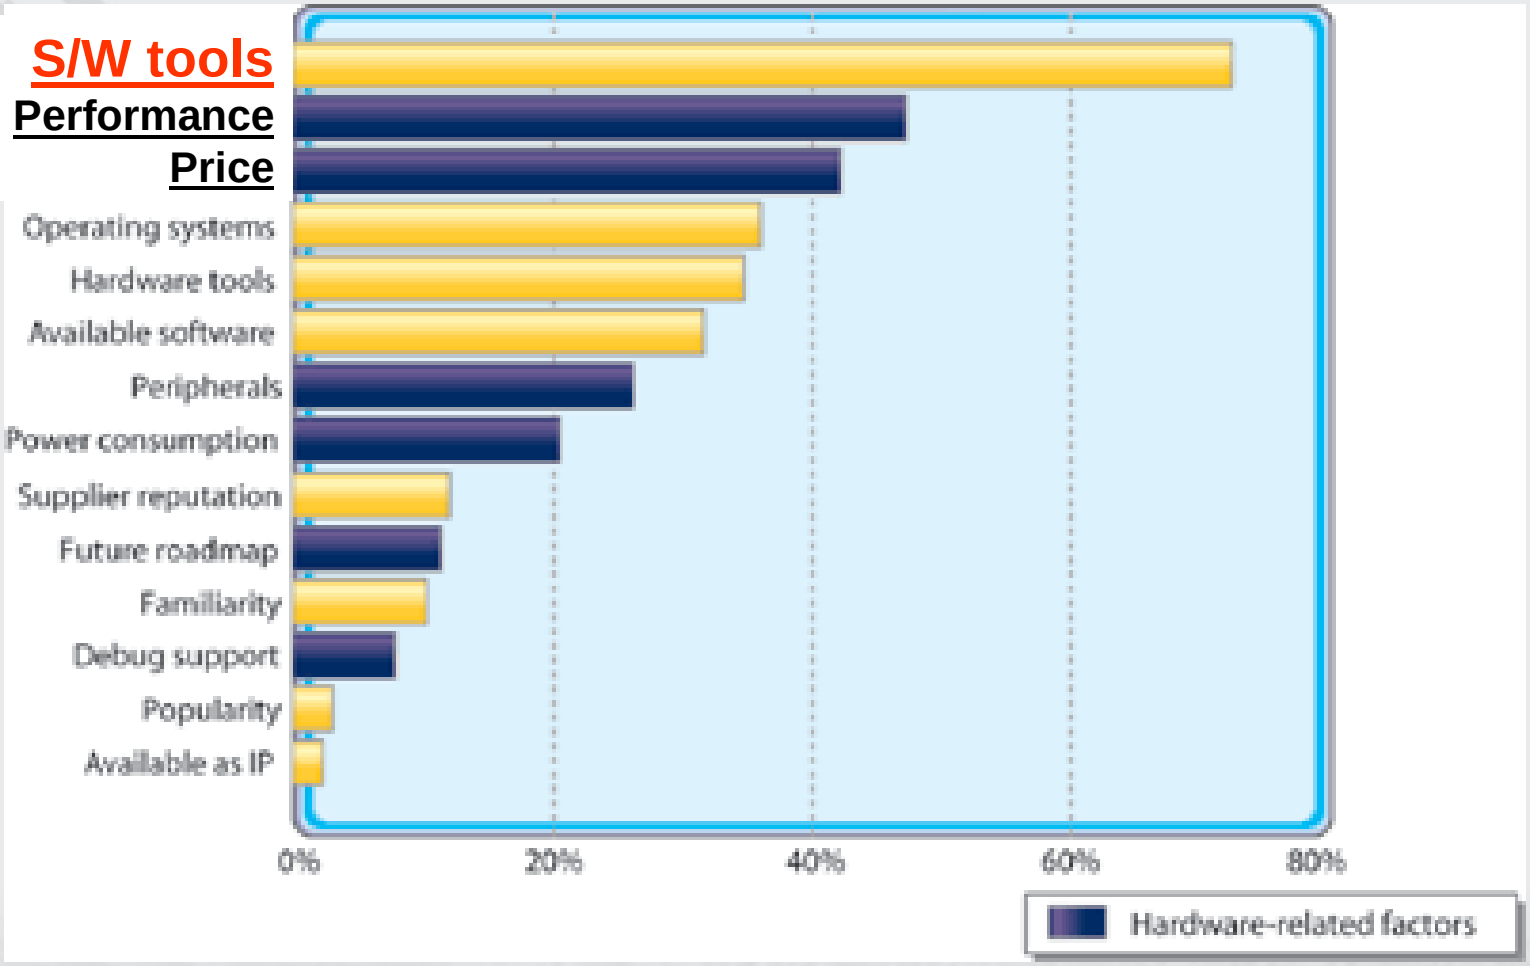

- Quality of software tools sells the processor

Source: **Embedded Systems Programming Magazine, 2005**

Outline

- Market Trends
- ***Technology Trends***
- Surviving 45 nm and beyond

Key MP-SoC Technology Trends

- Embedded processor speed wall ~1GHz
  - More function requires more parallelism
- Parallelism favors lower power solutions

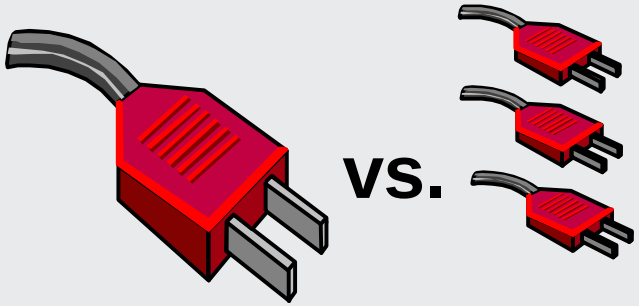

Example

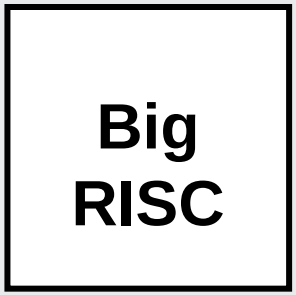

620 MHz,
900 MIPS

225 MHz,
330 MIPS x 3

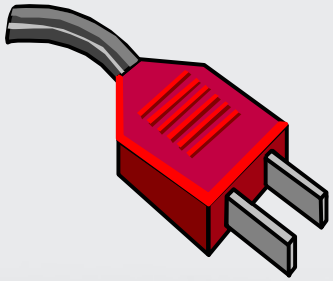

365mW

vs.

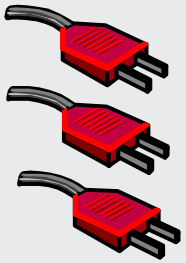

3 x 32mW
= 96 mW

4X
Power
efficiency

Power Scaling (Core Only)
Commercial Cores

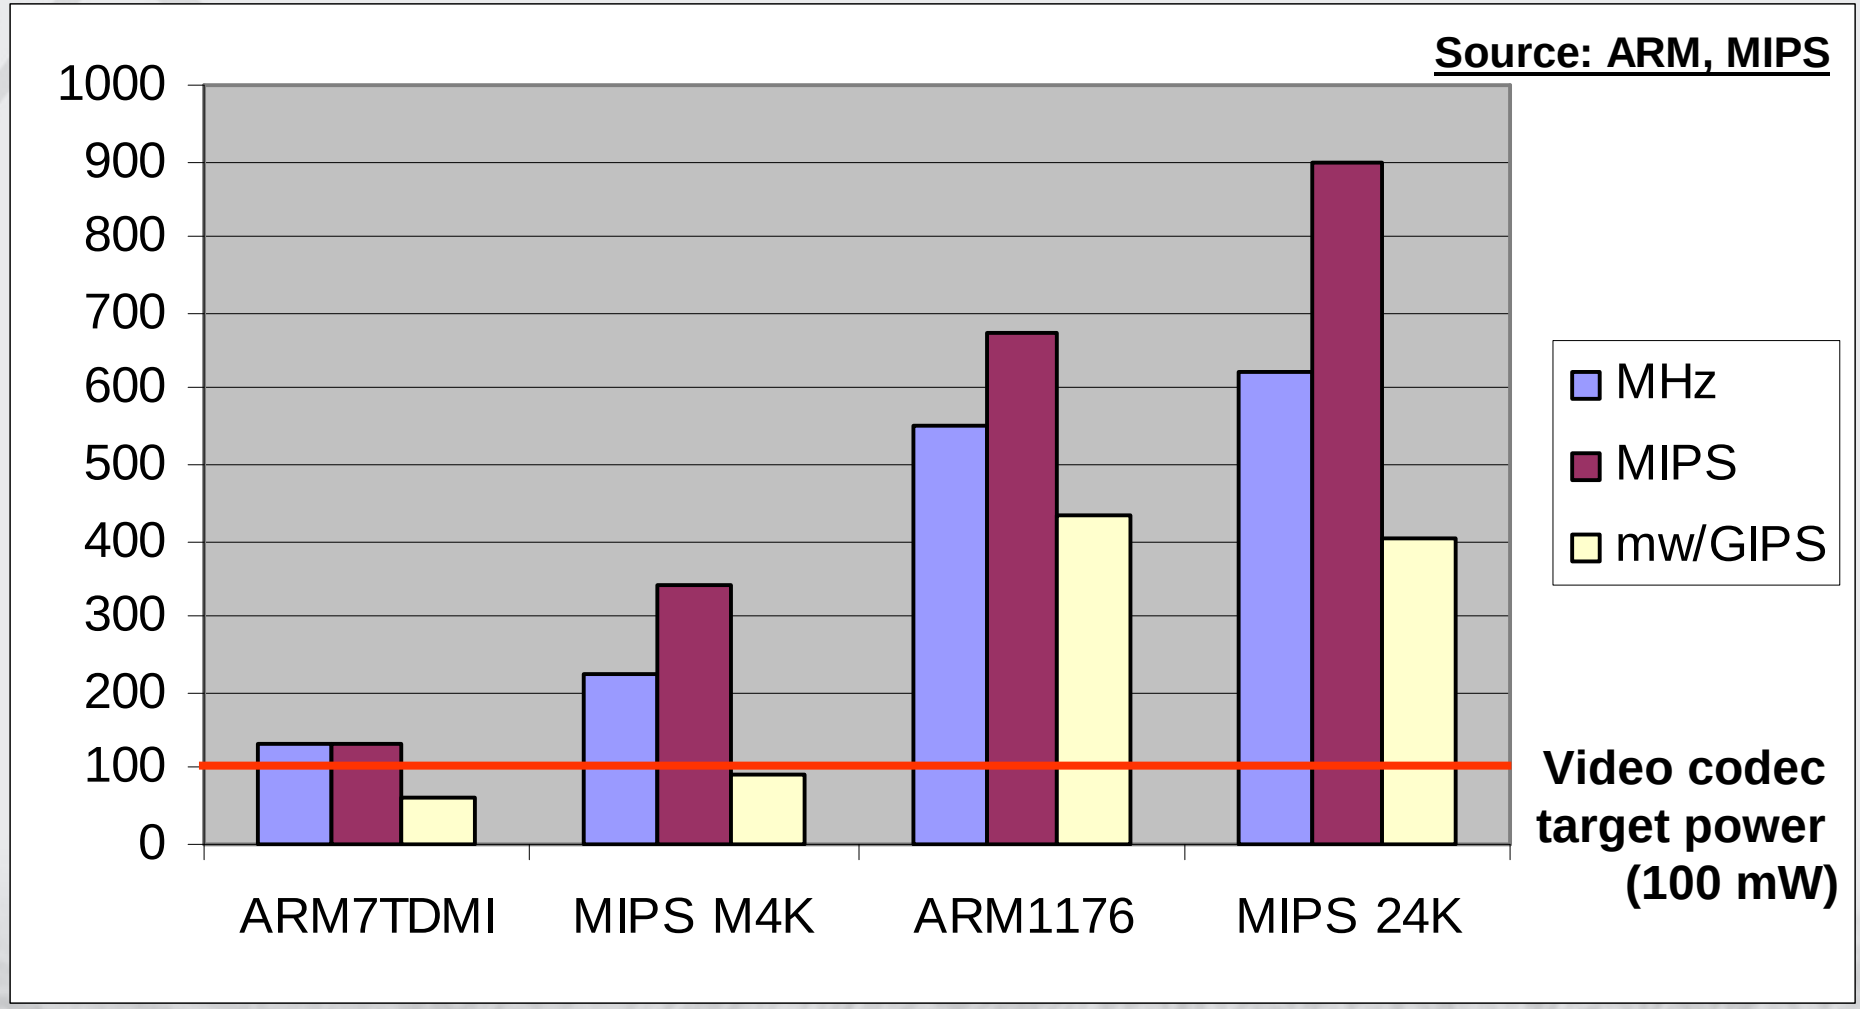

Power Scaling (Cores + Caches)

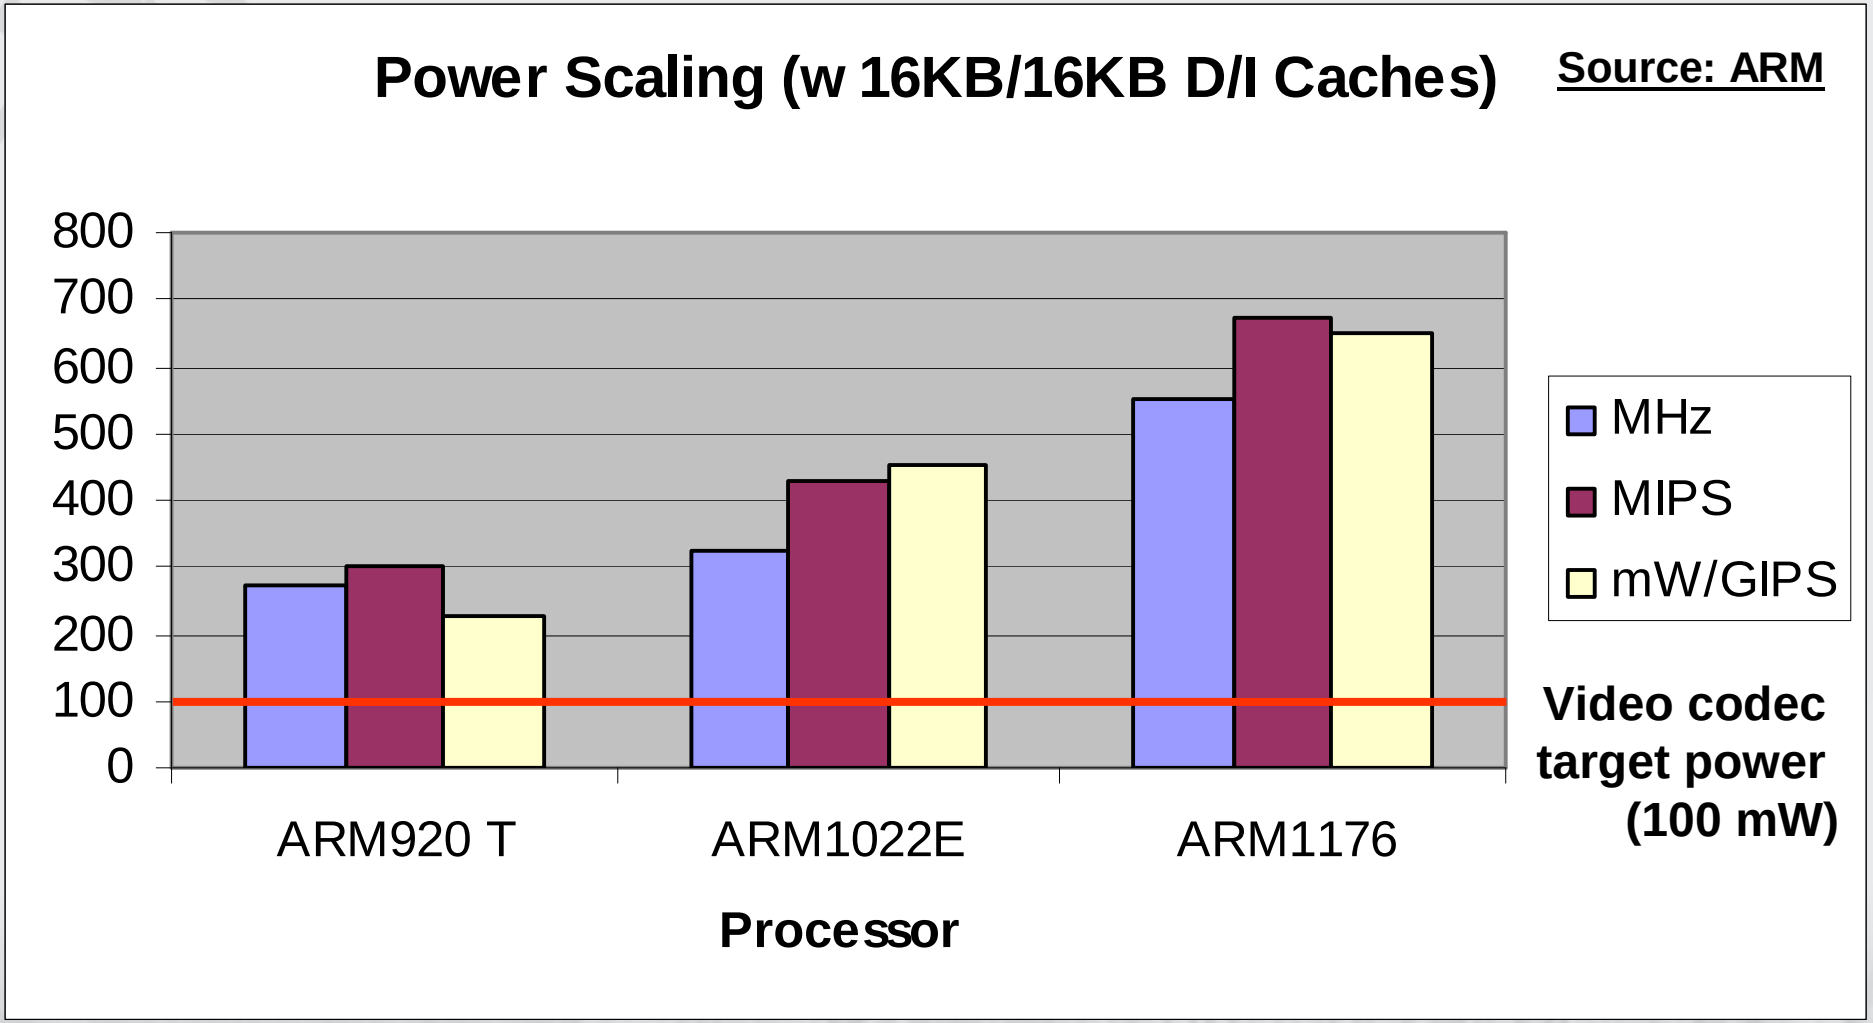

Voltage and Frequency Scaling (ADI Blackfin DSP)

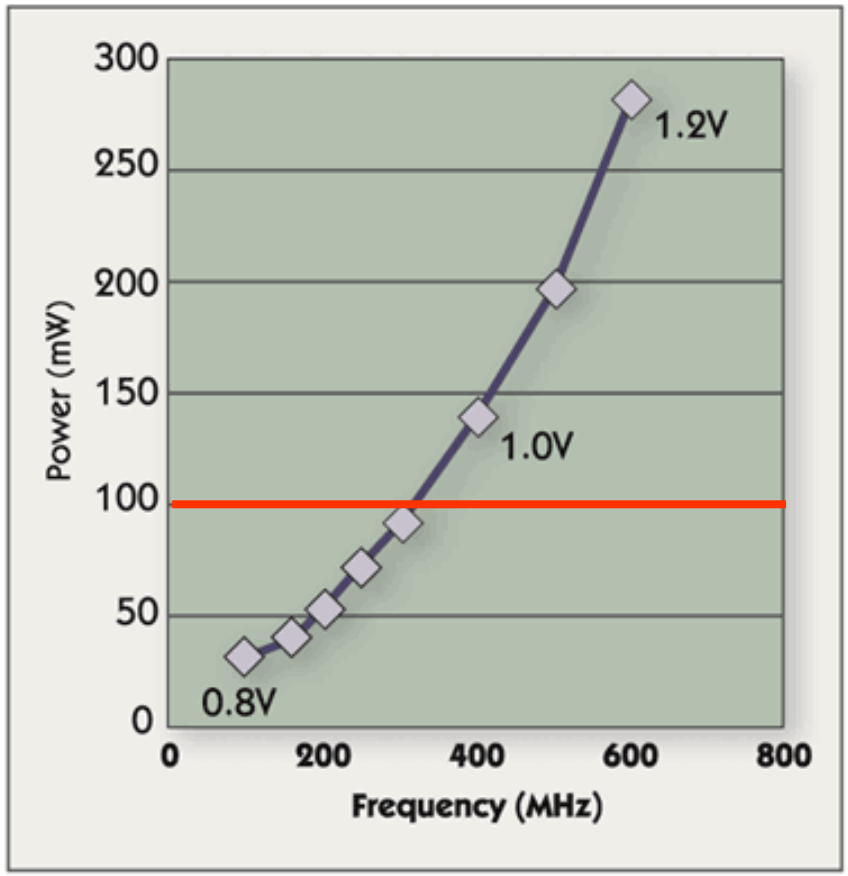

Fig. 2 Power consumption vs. operating frequency and voltage, ADI Blackfin.

Source: ADI

Voltage and Frequency Scaling (ADI Blackfin DSP)

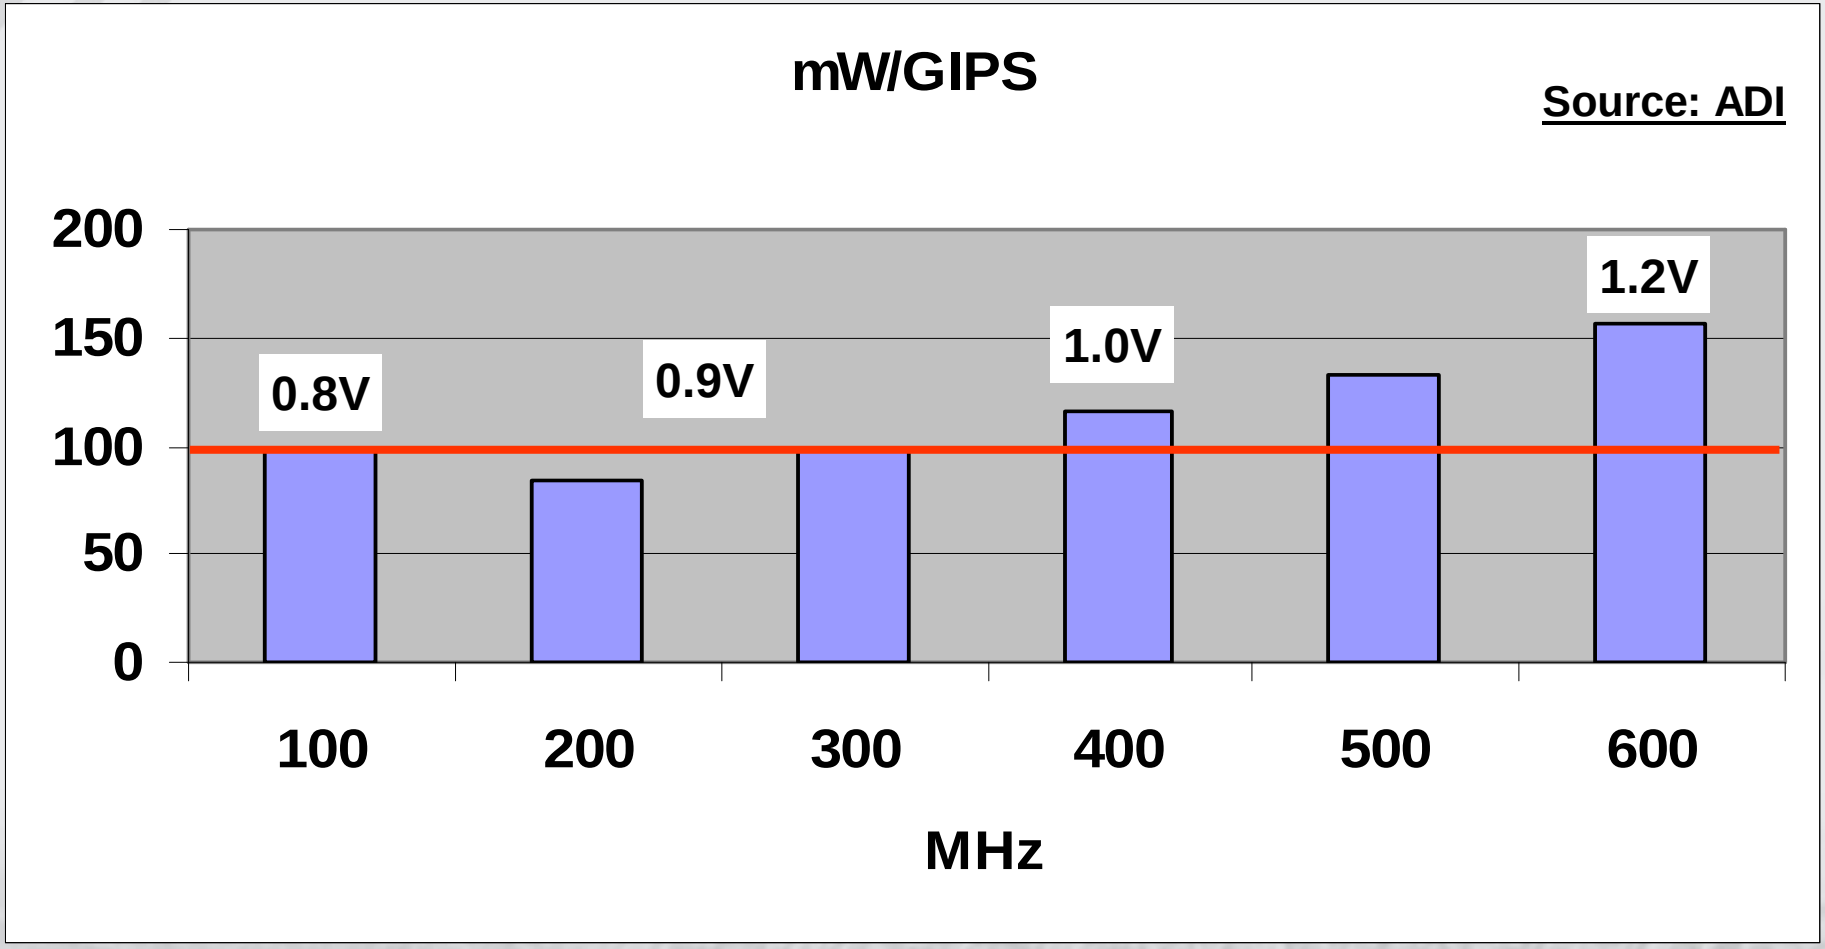

Latencies, latencies, latencies …

Increasing gap memory & processor speeds (2x / 2 years)

More parallel processing implies more IPC

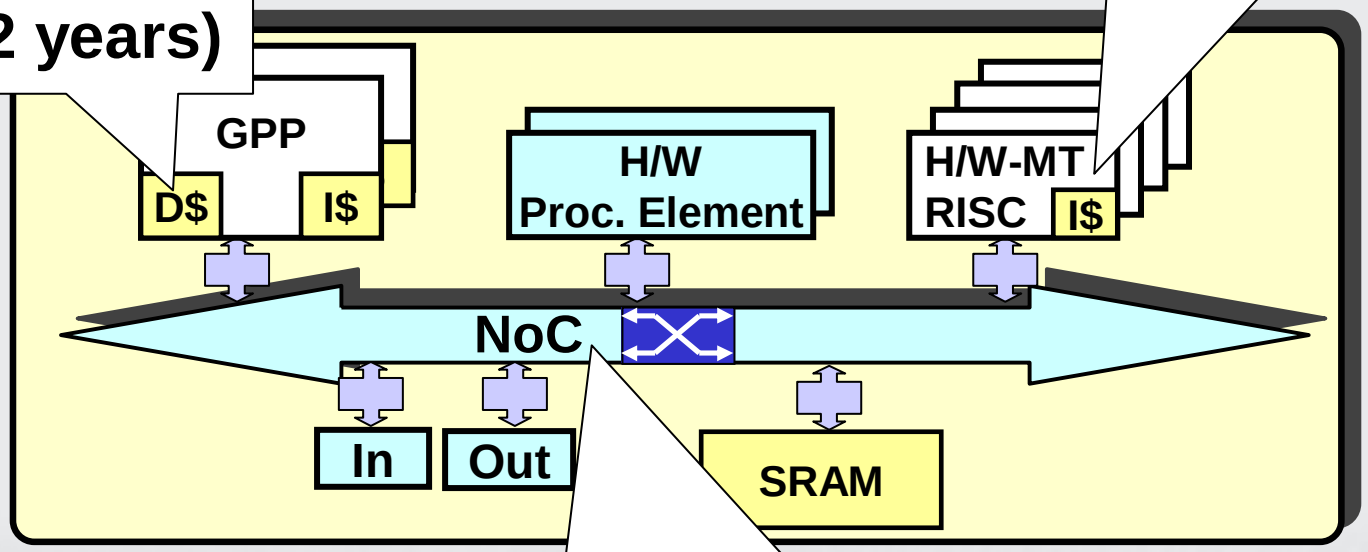

Increasing gap interconnect & gate delays (multi-clock intra-chip delay)

Outline

- Market Trends
- Technology Trends
- ***Surviving 45 nm and beyond***

Multimedia Performance Needs

- Audio:
  - ⇨High-end set top box — 800 MIPS
- Graphics (HD 720p, 30fps):
  - ⇨OpenGL 1.1 -> 240 Ops/Pixels — 7 GOPS
  - ⇨OpenGL 2.0 -> 400 Ops/Pixels — 11 GOPS
- H.264 encode (HD 720p, 30fps)
  - ⇨Video pipeline coder : — 8 GOPS
  - ⇨Bit stream processor: — 8 GOPS
  - ⇨Deblocking filter: — 8 GOPS
  - ⇨Hierarchical motion estimation: — 25~160 GOPS
- Digital TV
  - ⇨2004: 9000 Ops/Pixel — 450 GOPS
  - ⇨2008: 18000 Ops/Pixels — 900 GOPS

Subsystem Optimization

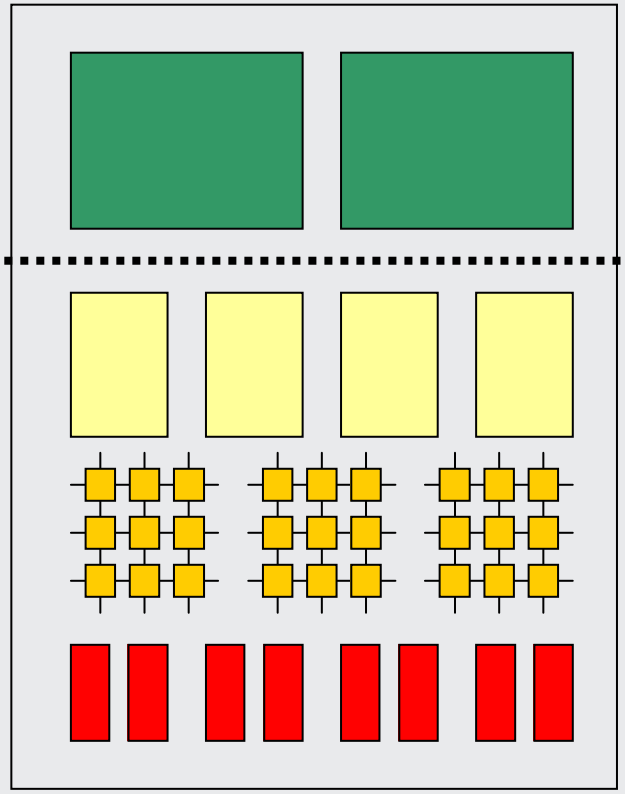

	Freq MHz	Gops/ mm^2	Efficiency %	Effective Gops
GP Core, Host	1000	1	20%	0.2
Application Oriented Cores	400	2	40%	0.8
Coarse grain Reconf.	150	5	80%	4
Specialized HW	150	15	100%	15

2 orders of magnitude

- Heterogeneity essential to obtain efficient platforms

‘Regular Heterogeneity’

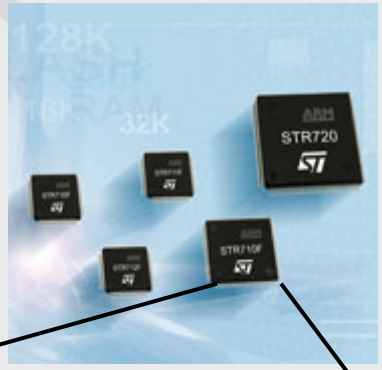

SiP

RISC
DSP
Bus
Bus
Bus
I/O
Mem
H/W

RISC
DSP
ASIP
ASIP
NoC
H/W
sASIC
I/O
FPGA
Mem

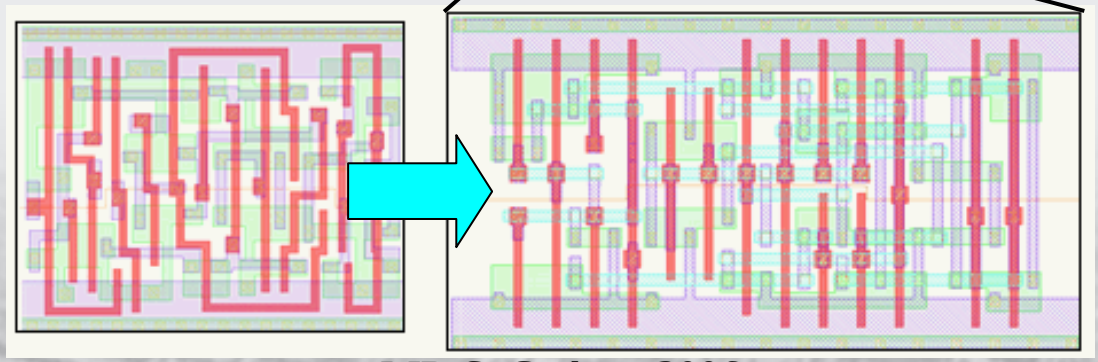

DFM, Fault-Tolerance

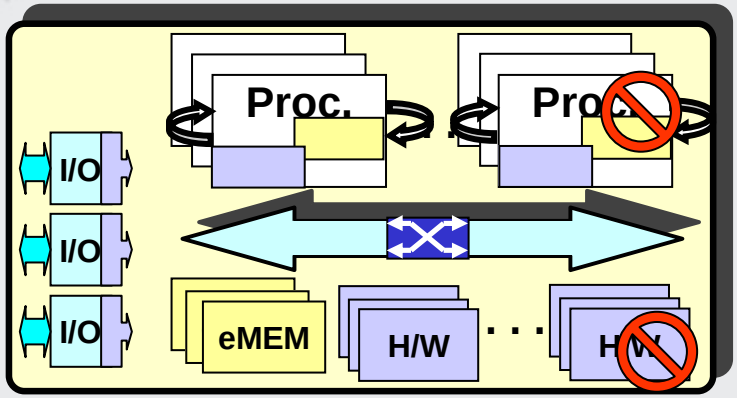

- **Full-featured, max speed**
- **Full I/O (big package)**

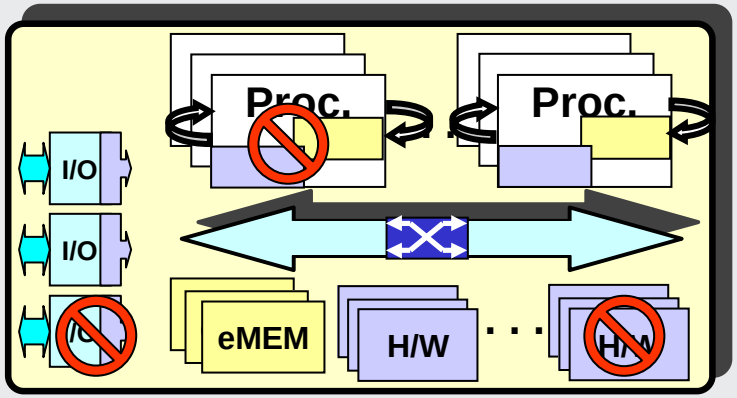

- **Subset of features, slower**
- **Low-cost package**

- Optimize yield/cost tradeoffs
  - Broadcom 14xx (2 MIPS, 4 I/F) -> 12XX (1 MIPS, 2 I/F)
  - PMC-Sierra's RM-9200 (2 MIPS) and RM-9100 (1 MIPS)
  - Motorola's PowerQuicc 8560 and 8540
- Fault-tolerance during product lifetime

Hiding latency with H/W MT

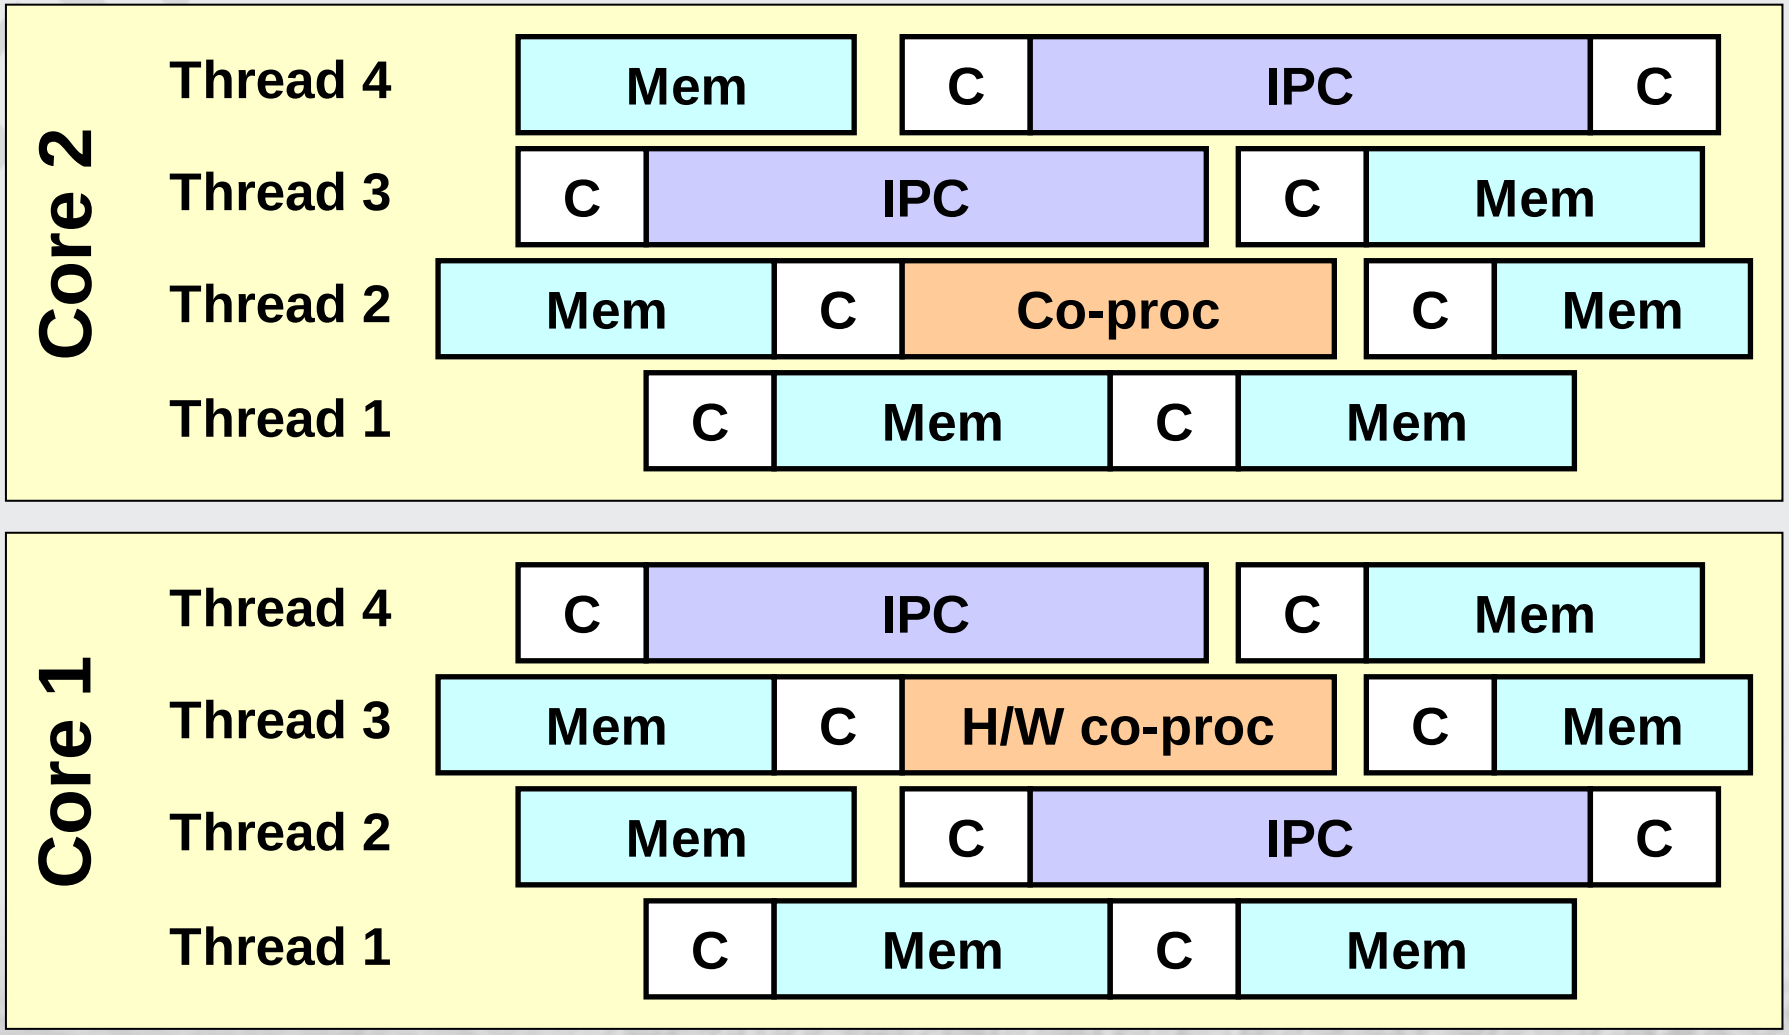

Exploitable Parallelism

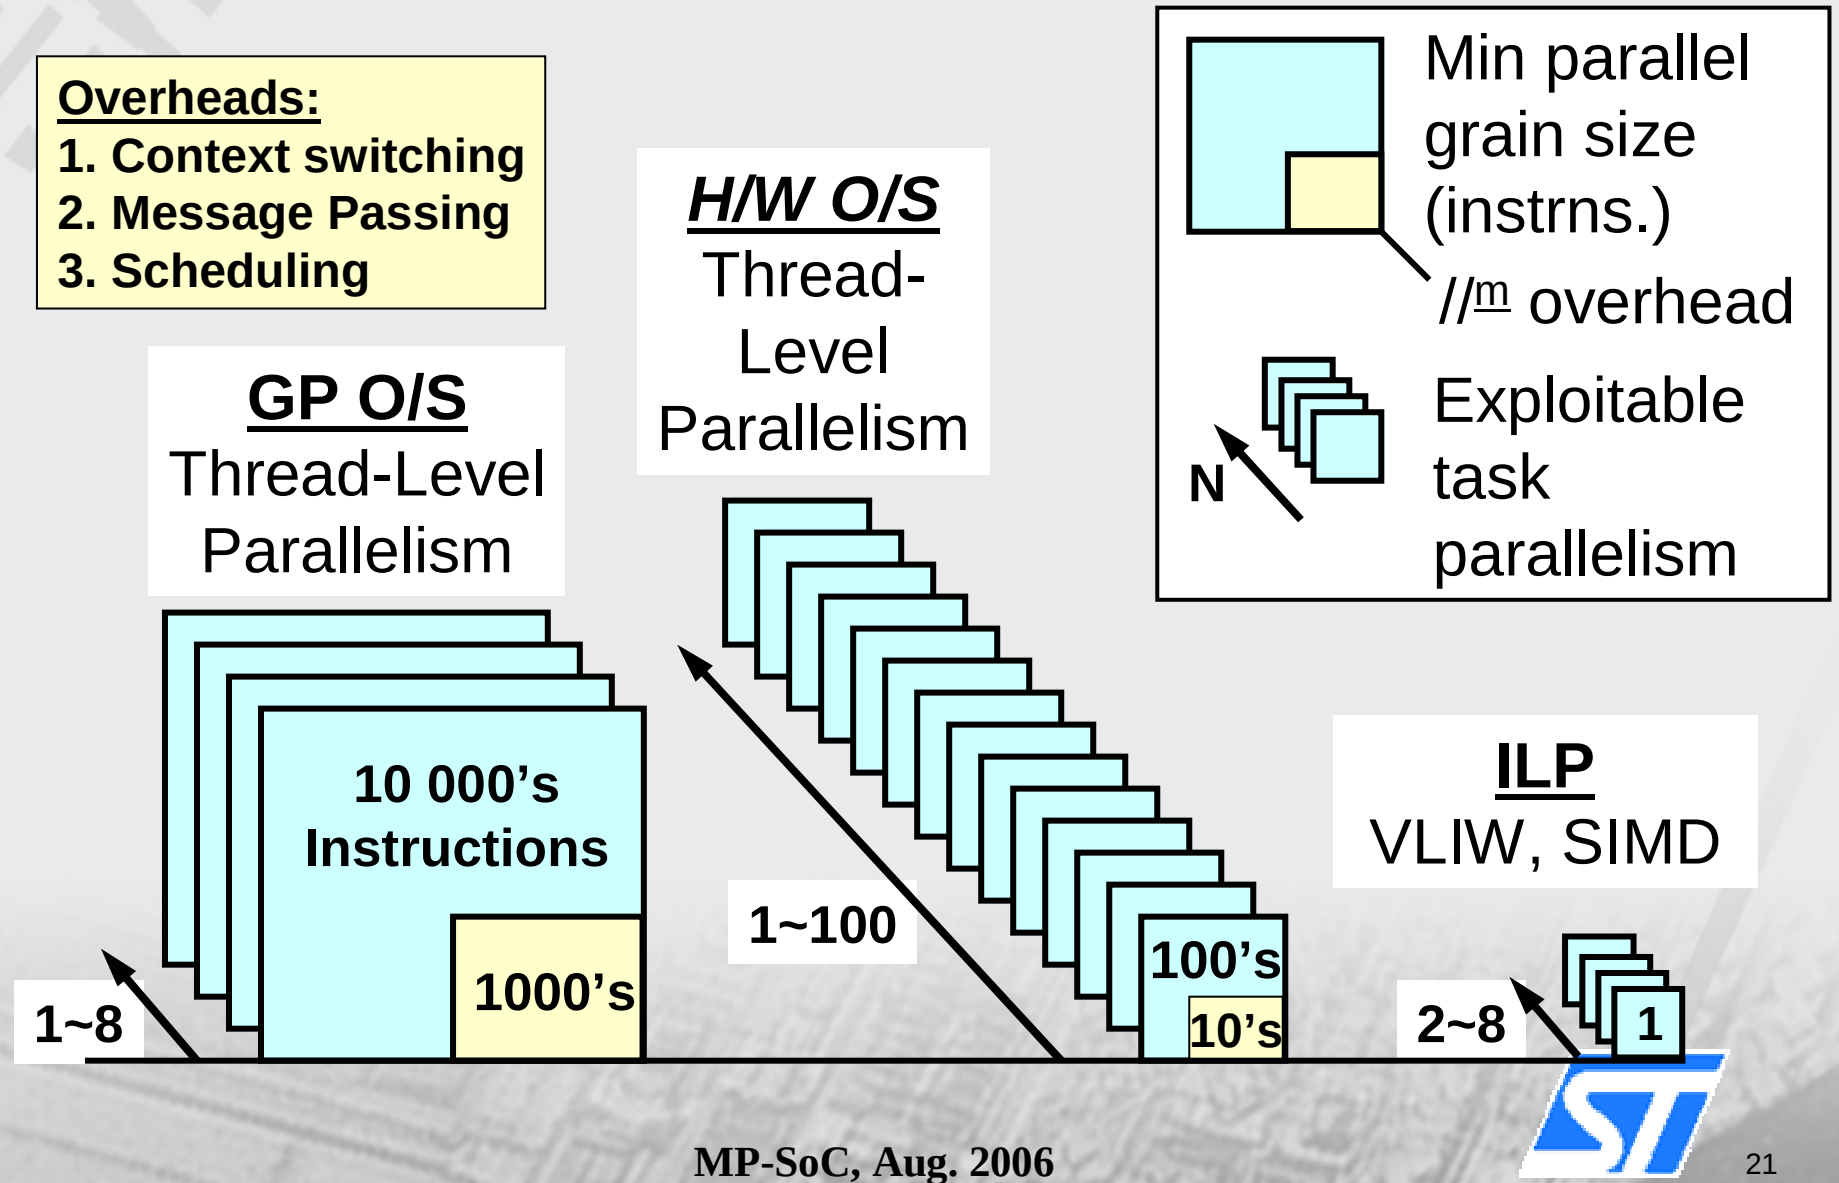

MP-SoC Programming Model

- High-level abstraction of heterogeneous SoC Platform
- Abstracts
  - Heterogeneity of components
  - Heterogeneity of tools
  - Low-level communication mechanisms
- Exposes
  - High-level parallelism supported by platform
    - ⇨OO message passing, shared-memory, streaming
  - Functions performed on *all* components
    - ⇨S/W, H/W, communication, storage, I/O

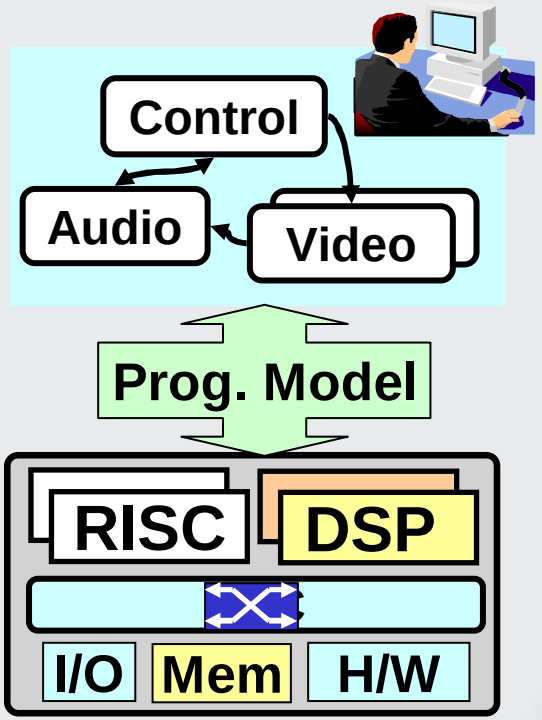

Prog. Model Transparency Objectives

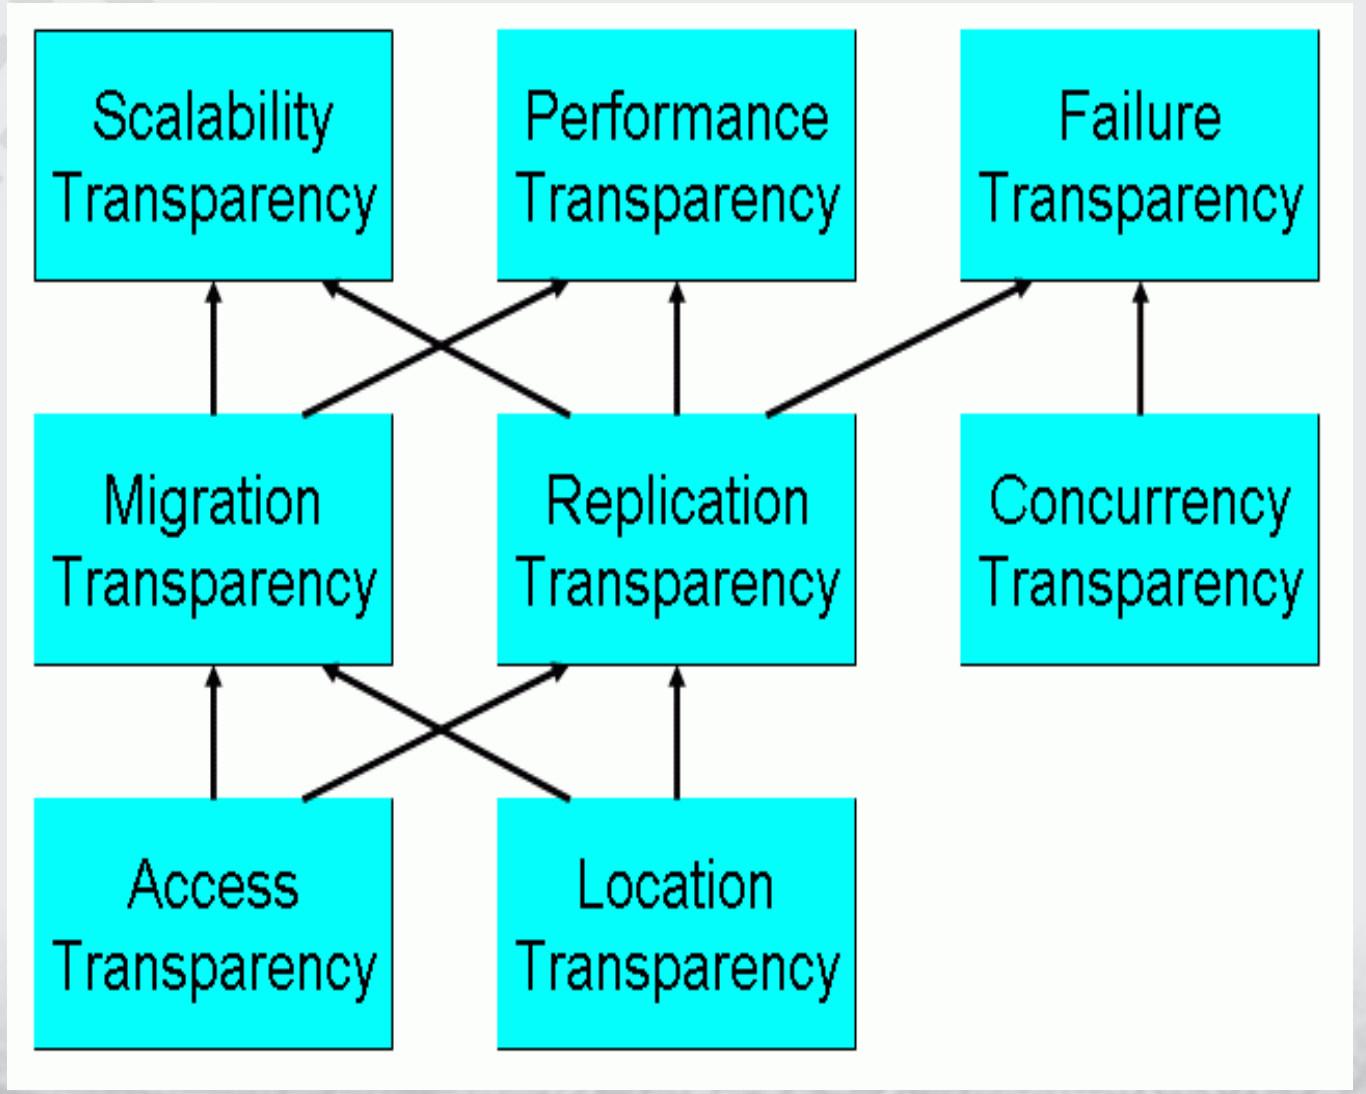

Ref:
P. Paulin et al, "Distributed Object Models for Multi-Processor SoC's, with Application to Low-power Multimedia Wireless Systems", *Proc. of DATE*, Munich, March 2006

Summary:
45 nm Survival Kit

- Exploit Parallelism for lower-power
- 'Regular heterogeneity' at all levels
  - Effective use of resources
  - DFM, fault-tolerance
- H/W assisted management of parallelism
  - H/W O/S, msg. passing engines
  - H/W multi-threading to hide latency
- Platform programming models
  - Abstract heterogeneity
  - Expose high-level platform parallelism
- Mapping tools should guide platform architecture design

What is Next SoC Paradigm Shift?

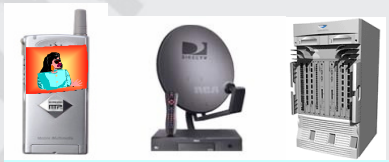

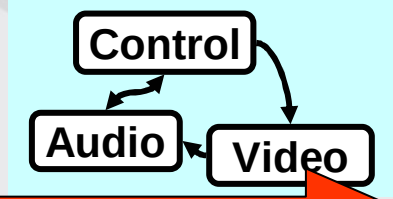

Prog. Models

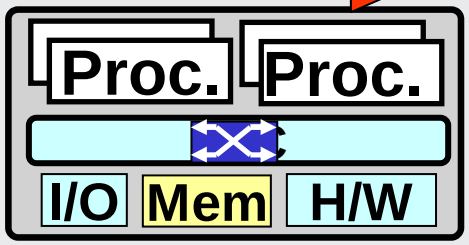

DFM

System Applications
- Video codecs
- Audio codecs
- Still image processing
- Communication stacks

months

2000's: Platform "Insulation"

Architecture Platforms
- Hard-disk drive platform
- Set-top box, DVD, HDTV
- Mobile multimedia
- Still Image processing platform

2 years

1990's: RTL, ISA "Insulation"

Component IP
- Value-add cores: ASIP, H/W IP
- Commodity cores: RISC, DSP, Bus
- Libraries: Cells, memories, I/O

2 years

1980's: Cell Library "Insulation"

References

- P. G. Paulin et al, “Parallel Programming Models for a Multiprocessor SoC Platform Applied to Networking and Multimedia”, *IEEE Transactions on VLSI Systems*, Vol. 14, No. 7, July 2006
- P. G. Paulin et al, “Distributed Object Models for Multi-Processor SoC’s, with Application to Low-power Multimedia Wireless Systems”, *Proc. of DATE*, Munich, March 2006