

Design Challenges in Multi-GHz Microprocessors

Bill Herrick

**Director, Alpha Microprocessor
Development**

COMPAQ

www.compaq.com

Introduction

- ◆ Moore's Law (the trend that the demand for IC functions and the capability of the semiconductor industry to meet that demand, will double every 1.5 to 2 years) has worked well during the last 30 years
- ◆ Difficult challenges face the industry attempting to maintain the pace
- ◆ With collaboration, understanding, vision and innovation this trend can continue for high performance microprocessors

Topics

- ◆ Historical Trends
 - Intel
 - Alpha chips and design style
 - Observations and trends
- ◆ Technology Predictions
 - ITRS 1999
- ◆ Key Design Challenges
 - Clocking and power - how Alpha has managed
 - Clocking and power - long term solutions

Historical Trends: Then and Now

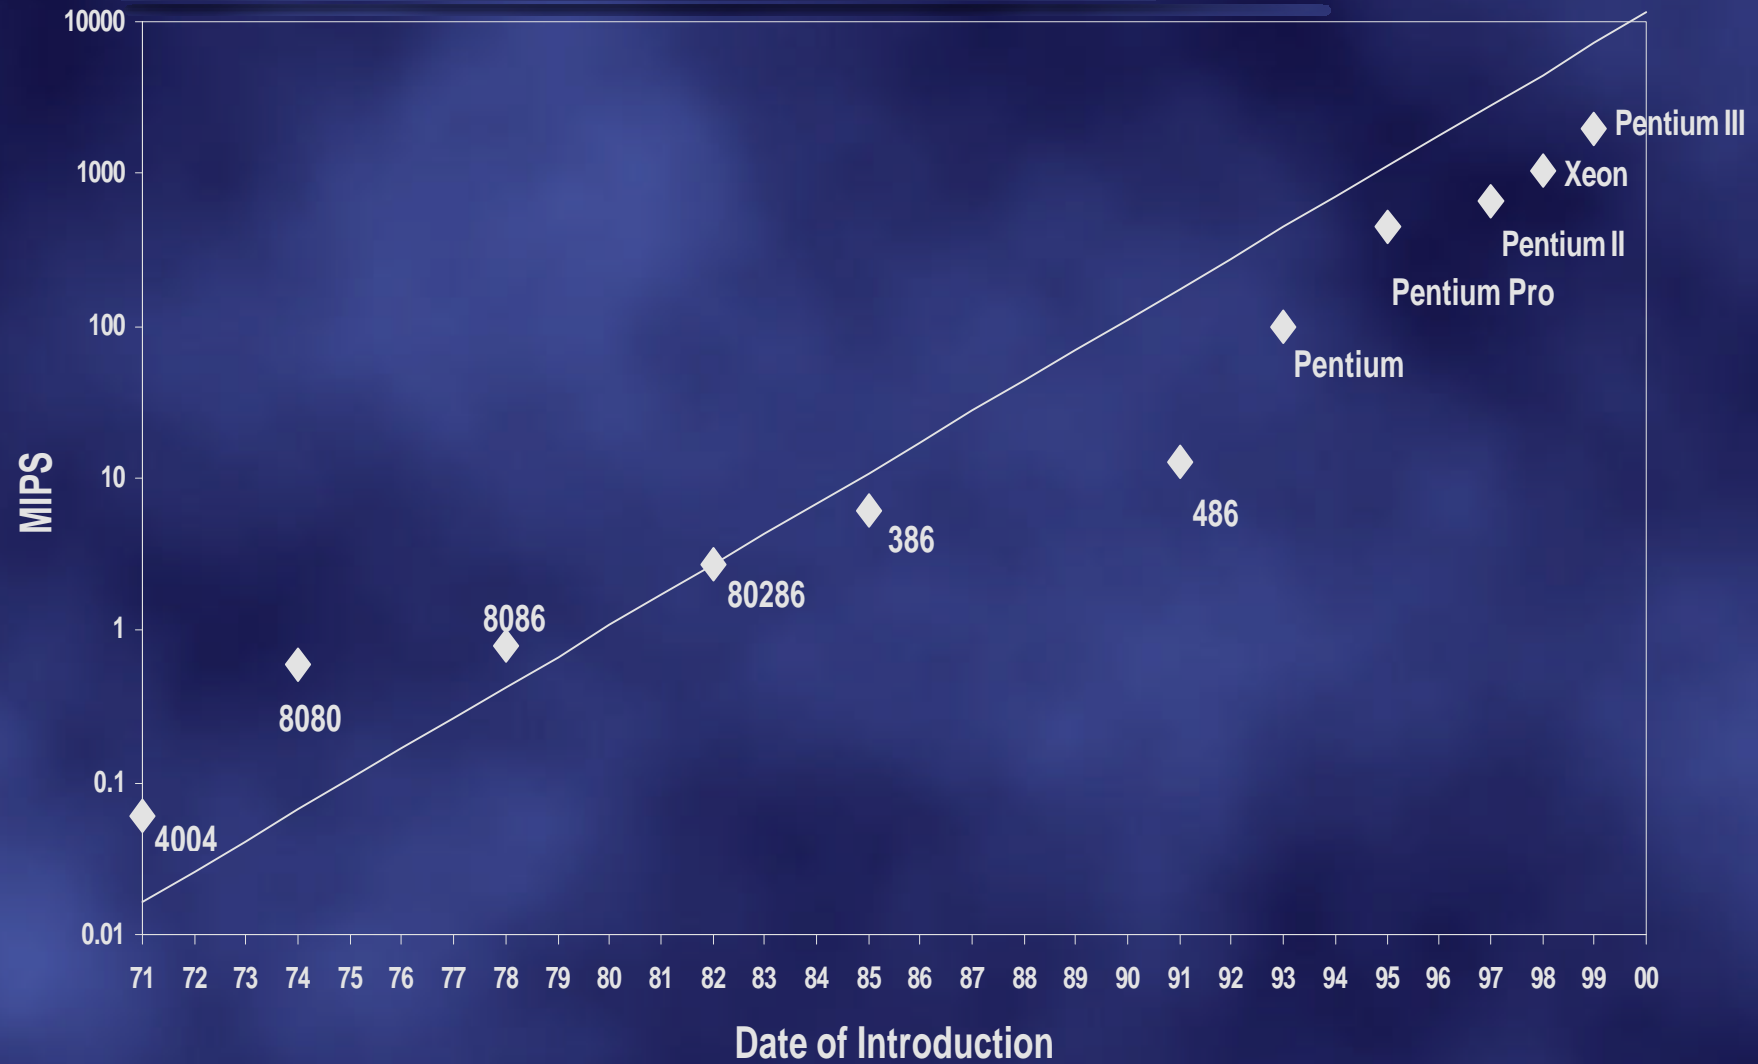
Circa 1970

- ◆ 12 μ PMOS
- ◆ 1000 transistors
- ◆ 5 - 10 mm² die size
- ◆ 10V supply
- ◆ 50 - 100 kHz frequency
- ◆ 100 - 200 mW
- ◆ 16 pin DIPs

Circa 2000

- ◆ 0.18 μ CMOS
- ◆ 10 - 100 million transistors
- ◆ 300 - 400 mm² die size
- ◆ 2.5V supply
- ◆ 500 - 1000 MHz frequency
- ◆ 50 - 100 W
- ◆ 500 - 1000 pin BGAs

Intel Performance History



COMPAQ

ASP DAC 2000

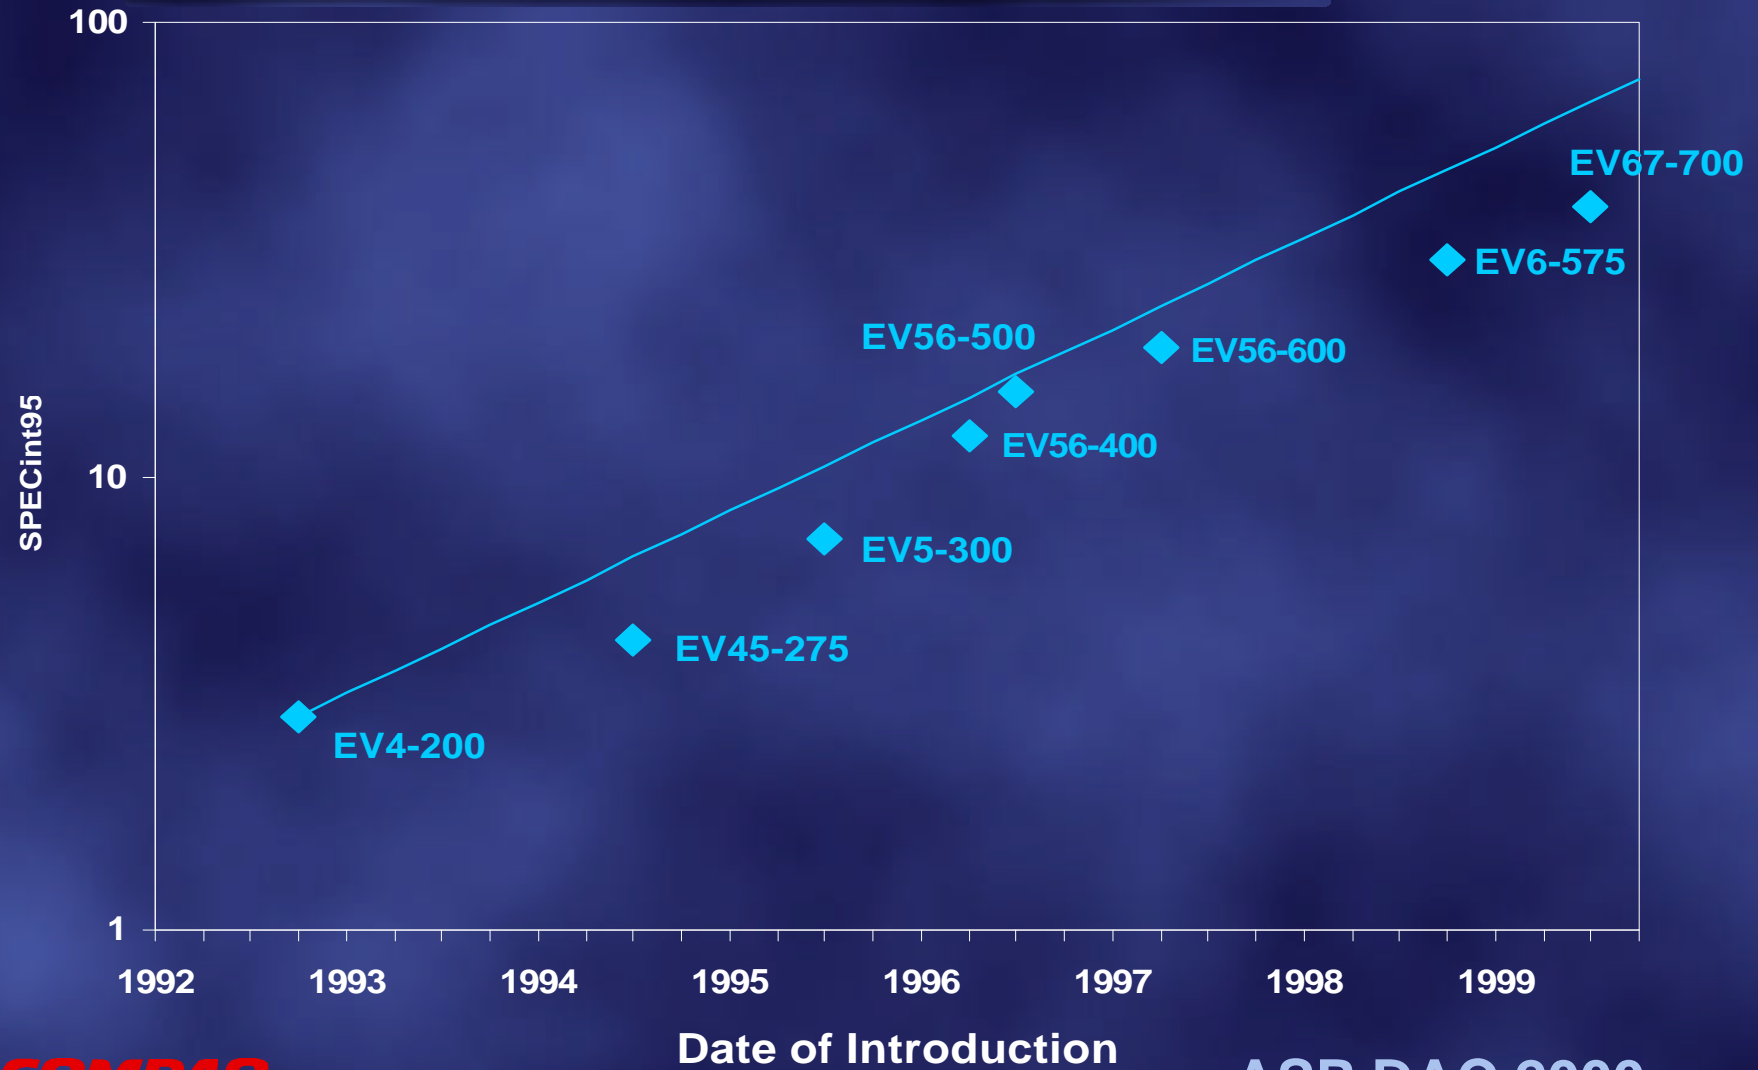
Intel Trends

- ◆ The 4004 (1971)
 - 2300 transistors in a 10u process,
 - 108kHz operation, executing 0.06 MIPs
- ◆ Pentium III (1999)
 - 28 million transistors in a 0.18u process,
 - 733MHz operation, executes 2000 MIPs
- ◆ Over nearly 30 years
 - performance has increased 30,000x,
 - transistor count has increase 10,000x
 - frequency has increased 7,000x
 - die size has increased only 25x.
 - Moore's law predicts 30,000x to 1,000,000x improvement over this period.

Alpha Architecture

- ◆ Alpha is a true 64-bit load/store RISC architecture
- ◆ Alpha is designed for high clock speed
 - Simple, fixed length (32-bit) instructions
 - Minimal instruction ordering constraints
 - No conditions codes
 - No branch delay slots
- ◆ Chip micro-architecture is carefully chosen to maximize performance without impacting cycle time


Alpha Performance History



COMPAQ

ASP DAC 2000

EV4 Chip Overview




- 0.75 μm 3LM N-well CMOS, $L_{\text{eff}}=0.5\mu\text{m}$, $T_{\text{ox}}=10.5\text{nm}$
- 3.3V Vdd
- 200MHz @100 ° C & 3.3V
- 16 gate delays per cycle
- 30W @200MHz & 3.3V
- 13.9mm x 16.8mm (233 mm²)
- 1.7 Million Transistors
~ 0.85 Million Logic Transistors
- 431 pin PGA (291 signals)

EV4 Micro-Architecture

- ◆ Dual In-Order Instruction Issue
 - single-issue Integer & single-issue FP
- ◆ Fully Pipelined (except Integer MUL and FP DIV)
 - 7-stage Integer and 10-stage FP pipelines
- ◆ 1-bit Branch Prediction: 2k-entry BHT
- ◆ 8kB direct-mapped I-Cache and 8kB direct-mapped write-through D-Cache
- ◆ 32 Integer and 32 FP Registers, 64b/entry
- ◆ Flexible external interface: shared 128b/64b data, 34b address L2 cache and system interface

EV5 Chip Overview




- 0.50 μ m 4LM N-well CMOS, $L_{\text{eff}}=0.365\mu\text{m}$, $T_{\text{ox}}=9.0\text{nm}$
- 3.3V Vdd
- 350MHz @100 ° C & 3.3V
- 14 gate delays per cycle
- 60W @350MHz & 3.3V
- 16.5mm x 18.1mm (298 mm²)
- 9.3 Million Transistors
 - ~ 2.5 Million Logic Transistors
- 499 pin PGA (294 signals)

EV5 Micro-Architecture

- ◆ Quad In-Order Instruction Issue
 - dual-issue Integer & dual-issue FP
- ◆ 7-stage Integer and 9-stage FP pipelines
 - FP latencies reduced by 2 cycles
- ◆ 2-bit Branch Prediction: 2k-entry BHT
- ◆ 8kB I-Cache and 8kB write-through D-Cache
- ◆ 96kB unified on-chip L2 Cache
- ◆ Improved external interface supports a non-blocking cache scheme

EV6 Chip Overview



- 0.35 μm 6LM N-well CMOS, $L_{\text{eff}}=0.25\mu\text{m}$, $T_{\text{ox}}=6.0\text{nm}$
- 2.2V Vdd
- 575MHz @100 ° C & 2.2V
- 12 gate delays per cycle
- 90W @575MHz & 2.2V
- 16.7mm x 18.8mm (314 mm²)
- 15.2 Million Transistors
 - ~ 6 Million Logic Transistors
- 587 pin PGA (374 signals)

EV6 Micro-Architecture

- ◆ Four-wide Instruction Fetch
- ◆ Tournament Branch Predictor
- ◆ Out-of-Order Execution Pipelines
 - Quad-speculative-issue integer pipeline
 - Dual-speculative-issue floating-point pipeline
- ◆ 80 In-flight Instructions
- ◆ Registers: 80 Integer, 72 Floating Point
- ◆ Queue Entries: 20 Integer, 15 Floating Point
- ◆ 2-Way 64KB L1 On-Chip Instruction and Data Caches
- ◆ Up to 16 outstanding off-chip memory references

EV7 Chip Overview

- ◆ 0.18 μ m CMOS technology
- ◆ 1.5V Vdd
- ◆ Clock frequency >1.0GHz
- ◆ 100W
- ◆ ~350mm²
- ◆ ~100 Million transistors
- ◆ EV6 core
- ◆ Integrated L2 Cache (1.75 MB 7-way)
- ◆ Integrated memory controller (RAMBUS)
- ◆ Integrated network interface

EV8 Chip Overview

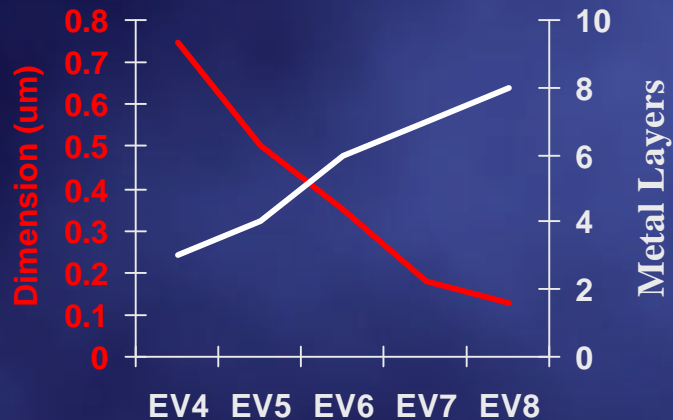
- ◆ Clock frequency range 1.0-2.0GHz
- ◆ Leading edge 0.125 μ m CMOS technology
- ◆ ~1.2V Vdd
- ◆ <150W
- ◆ ~250 Million transistors
- ◆ Enhanced out-of-order execution
- ◆ 8-wide superscalar
- ◆ 4-way simultaneous multi-threading (SMT)
- ◆ EV7 memory and system enhancements

Alpha Circuit Design Philosophy

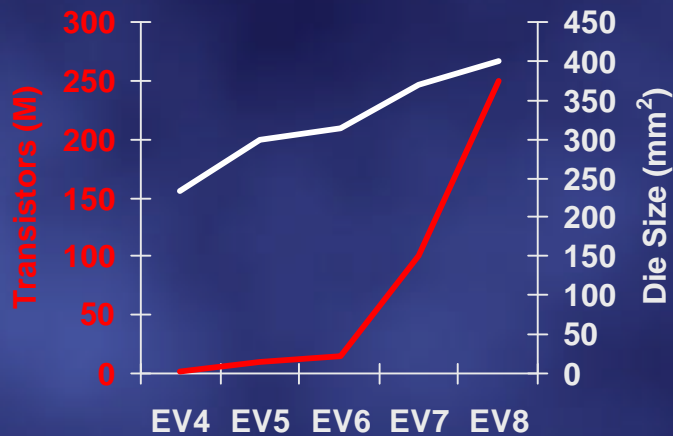
- ◆ Transistor level circuit design
- ◆ Broad range of circuit styles and logic families
 - Complementary CMOS
 - Dynamic logic
 - DCVSL (cascode)
 - Ratioed logic
- ◆ Key components to enable *high performance*
 - On-chip clock generation and distribution (low-skew, fast edge)
 - Latching (low latency)
 - Low noise on-chip power distribution
 - On-chip signal integrity management

Complexity Trends

Process Features



Chip Features




- ◆ Process scaling has continued steadily
- ◆ Planarization has enabled an increase in the number of interconnect layers
- ◆ Transistor counts have increased dramatically with the L2 cache SRAMs
- ◆ Additionally, design team size has increased ~40% per generation
- ◆ Opportunities to manage complexity and productivity
 - Fundamental understanding and modeling of process and circuit element behaviors
 - High level design methods
 - CAD
 - Design reuse
 - Micro-architecture

COMPAQ


ASP DAC 2000

Power Dissipation Trends

Power Dissipation




Supply Current



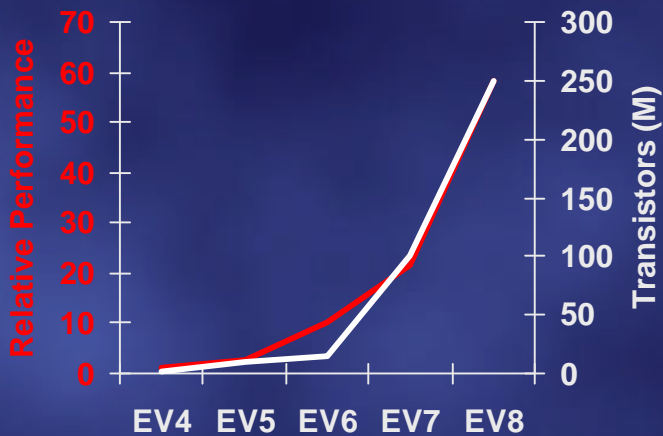
- ◆ Power consumption is increasing
 - Better cooling technology needed
- ◆ Supply current is increasing faster!
- ◆ On-chip signal integrity will be a major issue
- ◆ Power and current distribution are critical
- ◆ Opportunities to slow power growth
 - Accelerate Vdd scaling
 - Low dielectrics & thinner (Cu) interconnect
 - SOI circuit innovations
 - Clock system design
 - micro-architecture

Performance Trends

Clock Speed



Transistor Count



COMPAQ

- ◆ Performance has increased significantly (7x) faster than frequency
- ◆ Performance tracks transistor count when L2 cache ignored
 - Transistor budget has increased more than performance when L2 cache is considered but
 - benchmarks did not reflect larger applications
- ◆ Opportunities to continue performance improvements
 - Continued scaling of devices, interconnect and dielectrics
 - Clock distribution
 - Micro-architecture
 - System design


ASP DAC 2000

Micro-Architecture Trends


- ◆ Trends have included
 - Wider super-scalar machines, deep pipelines
 - Larger register, L1 caches
 - On-chip L2 caches
 - Out of order execution
 - Sophisticated branch prediction, predication, speculation
 - Integrated memory and network controllers
 - SMT
 - Less idle logic but more bookkeeping logic
- ◆ Future opportunities include
 - Floating point performance improvements
 - Vectors
 - Thread-level speculation
 - More pipelining
 - Better on-chip communications
 - Banking, replicating structures
 - Clustering functional units
 - On-chip SMP

Challenging Design Trends

Logic Levels per Cycle



Cycles Across Chip



COMPAQ

- ◆ Micro-architecture and logic design are stressed as frequency has increased faster than scaling
- ◆ Further reducing the number of gate delays per cycle will be difficult
- ◆ Cycles to communicate across chip track with frequency
- ◆ Clock edge rates are not scaling
- ◆ Opportunities to continue performance increases
 - Chip implementation design
 - Clock system design
 - Micro-architecture

ASP DAC 2000

ITRS -1999 Key Messages

- ◆ No major issues through 130 nm generation, but significant issues for 100 nm generation (2005)
- ◆ Continued technology scaling will require the introduction of new process materials and new devices
- ◆ Transistor densities will continue the historical trends $\sim 2X / 2yrs$
- ◆ Clock frequency increases will slow compared to historical trends

ITRS-1999 The Roadmap

	1999	2001	2003	2005	2008	2011	2014
Generation (nm)	180	130	130	100	70	50	35
L_{eff} (nm)	140	100	80	65	45	31	21
Devices (M)	110	220	441	882	2494	7053	19949
Chip Size (mm²)	450	450	567	622	713	817	937
Signals	768	1024	1024	1024	1280	1408	1472
Pins	1600	2007	2518	3158	4437	6234	8758

ITRS-1999 The Roadmap (continued)

	1999	2001	2003	2005	2008	2011	2014
Generation (nm)	180	130	130	100	70	50	35
Clock (MHz)	1200	1454	1724	2000	2500	3000	3600
Local Clk (MHz)	1250	1767	2490	3500	6000	10000	13500
IO (MHz)	480	722	932	1035	1285	1540	1800
Wiring Levels	7	7	8	9	9	10	10
Vdd (V)	1.8	1.5	1.5	1.2	0.9	0.6	0.6
Power (W)	90	115	140	160	170	174	183

ITRS - 1999 Highlights

◆ Transistors

- Drive currents will remain constant through 2014 at 750 $\mu\text{A}/\mu\text{m}$ for NFETs and 350 $\mu\text{A}/\mu\text{m}$ for PFETs
- Leakage currents will double every 3 years from 5 nA/ μm in 1999 to 160 nA/ μm in 2014

◆ Interconnect

- Use of Cu and low dielectrics will become standard
- Local interconnect delays will scale with gate delays
- Global interconnect delays, even with repeaters will not scale with gate delays
- Coplanar waveguides, free space RF and optical interconnect may be needed longer term

ITRS - 1999 Highlights (2)

◆ Packaging

- Maximum junction temperature must be reduced from 100°C to 85°C by 2002 for reliability concerns
- Significant θ_{ja} improvements will be required to maintain air cooling system solutions: a 50% reduction by 2002 and another 30% by 2014

◆ Modeling & Simulation

- 2D and 3D interconnect models with inductance and transmission line effects will be needed
- Transistor models of non-quasi-static effects and quantum mechanical gate effects will be needed; gate currents will become important
- OCV modeling will become necessary
- CPU efficient and accurate models will be essential

ITRS - 1999 Highlights (3)


◆ Design Productivity

- Design team sizes will not exceed 300 people
- Design cycle times will decrease from 36 months in 1999 to 30 months in 2005 to 24 months in 2014


◆ Verification and Test

- Verification has become more than 50% of the total design effort
- Use of formal verification will increase from 15% now to 30% in 2005 to 60% in 2014
- BIST coverage will increase from 20% now to 40% in 2005 to 70% in 2014

EV4 Clocking



Clock waveform




Location of clock driver on die

- ◆ 2 phase single wire clock, distributed globally
 - Low skew
 - Fast edge rate
- ◆ 1 clock driver channel
 - 3.5nF clock load
 - 35 cm final driver width

EV4 Latches


First single wire clock implementation

- Race immune latch
- Level sensitive design
- 2 latches per cycle
- Can build logic into first stage of latch
- t_{cycle} latch overhead is approximately 25%




CLK high loading latch


EV4 Thermal Gradient



EV5 Clocking



Clock waveform




Location of clock driver on die

- ◆ 2 phase single wire clock, distributed globally
- ◆ 2 distributed driver channels
 - Reduced RC delay/skew
 - Improved thermal distribution
 - 3.75nF clock load
 - 58 cm final driver width
- ◆ Local inverters for latching
- ◆ Conditional clocks in caches to reduce power
- ◆ More complex race checking
- ◆ Device variation

EV5 Latches


Reduce t_{dq} and reduce clock load

- Local clock inverter complicated race issues
- Level sensitive design
- 2 latches per cycle
- Can build logic into first and last stages of latch
- t_{cycle} latch overhead is approximately 15%
- Smaller, faster and lower power than EV4 latch

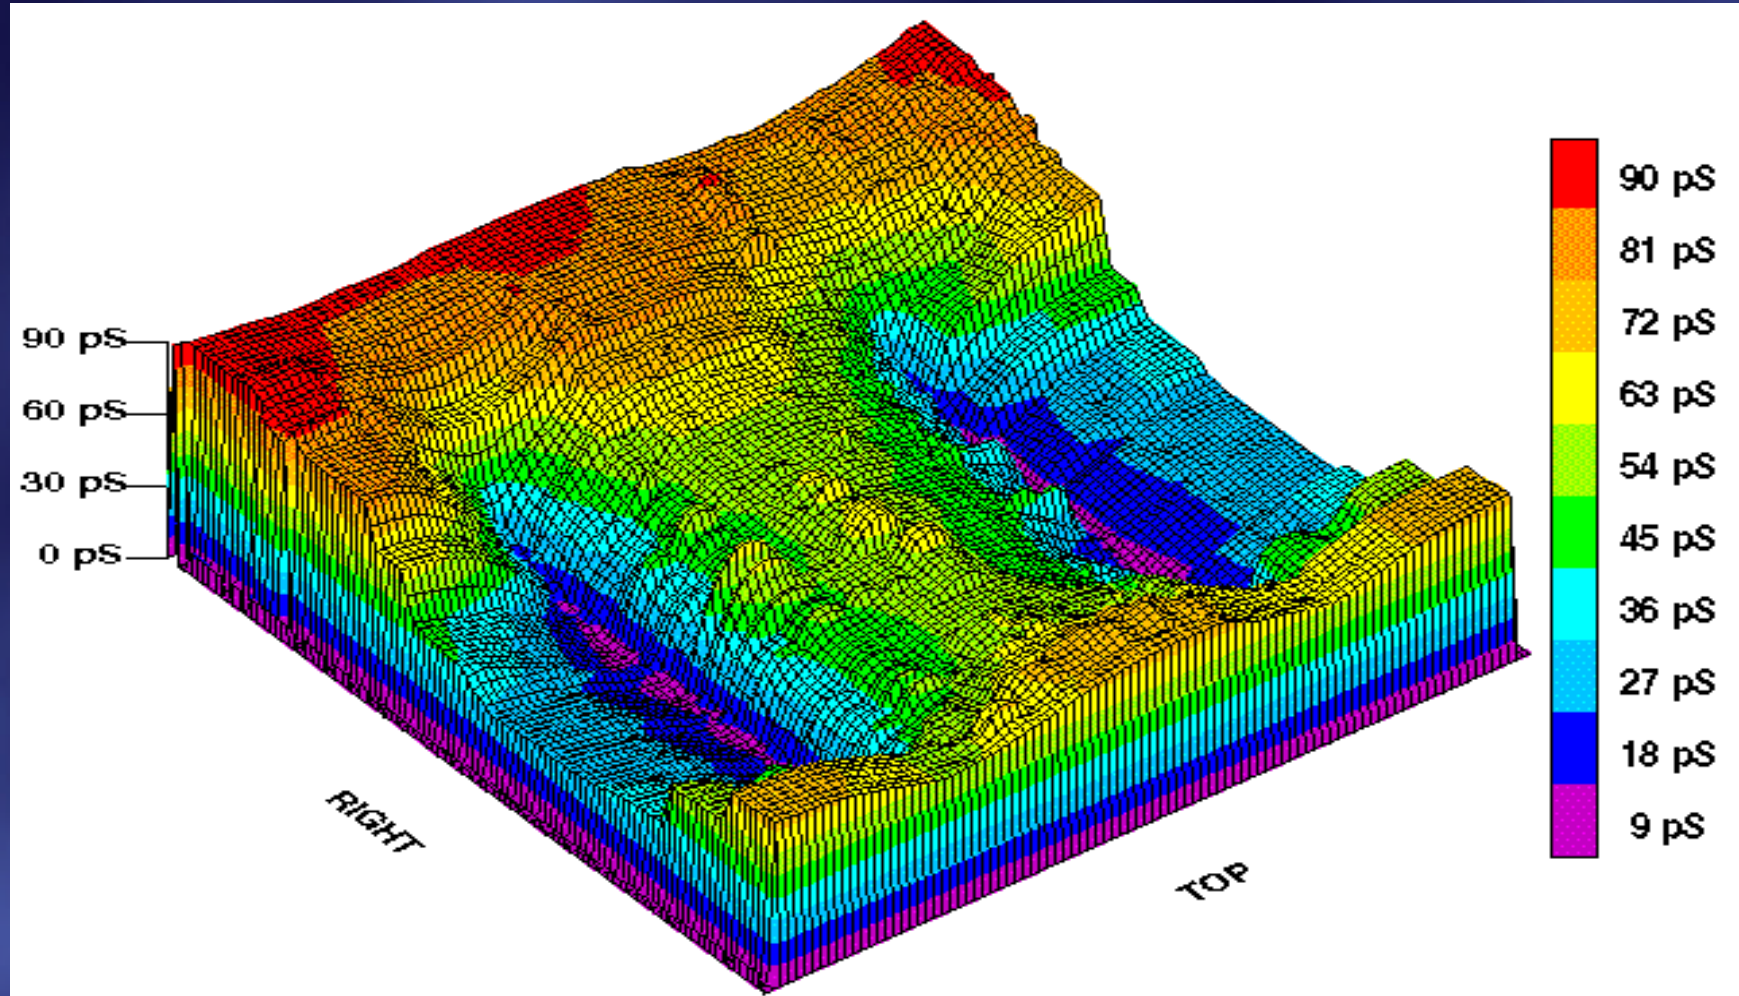


CLK high loading latch


EV5 Thermal Gradient




EV5 Global Clock Skew



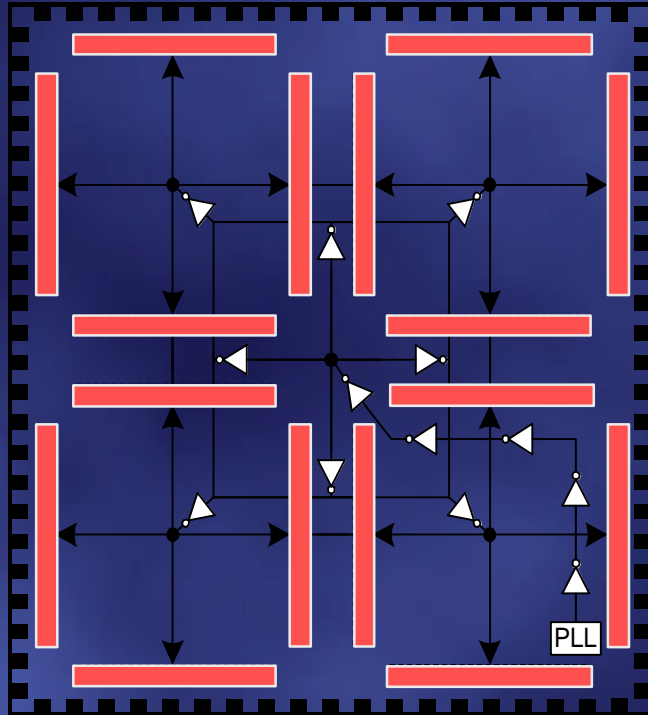
EV5 Local Clock Skew



EV6 Clcking



Global clock waveform




- ◆ 2 Phase, with multiple conditional buffered clocks
 - 2.8 nF clock load
 - 40 cm final driver width
- ◆ Local clocks can be gated “off” to save power
- ◆ Reduced load/skew
- ◆ Reduced thermal issues
- ◆ Multiple clocks complicate race checking

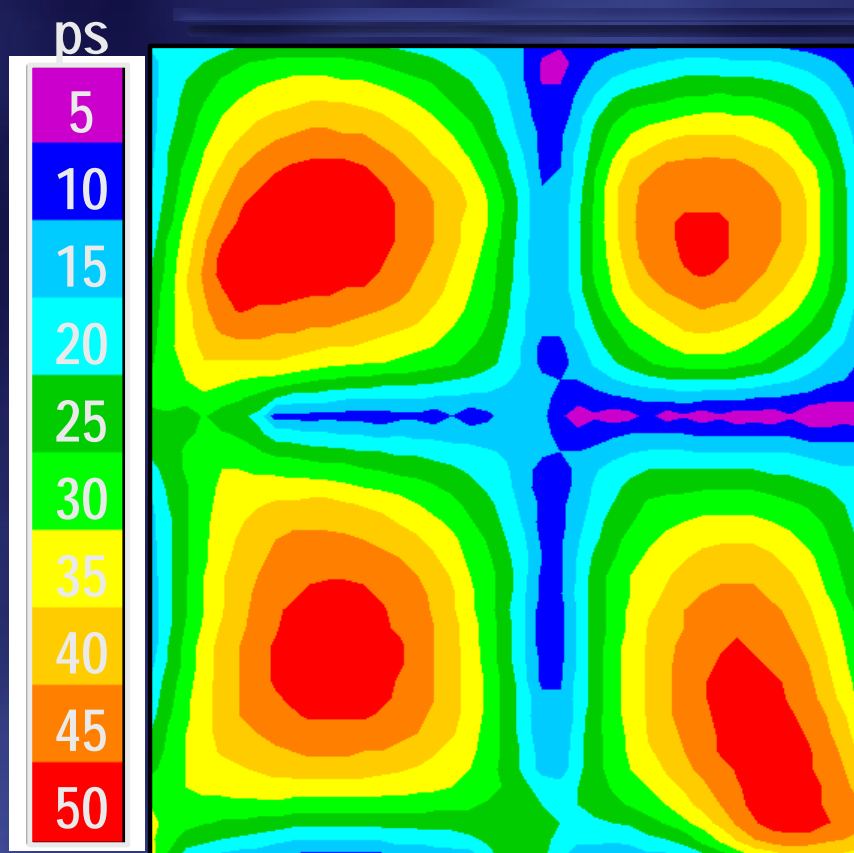
EV6 Latches

Conditional clocks to reduce power


- Static design
- 1 latch per cycle
- Edge triggered to simplify race rules
- Can build logic into latch
- t_{cycle} latch overhead is approximately 15%



EV6 Clock Results



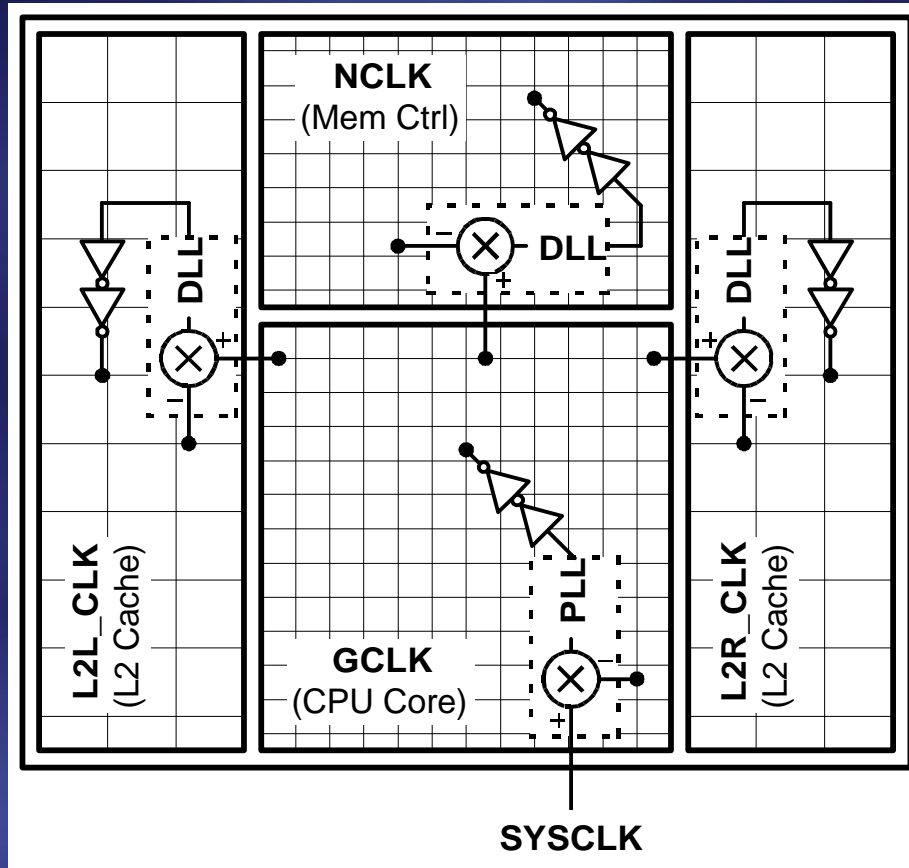
GCLK Skew
(at Vdd/2 Crossings)



GCLK Rise Times
(20% to 80% Extrapolated to 0% to 100%)

EV7 Clock Hierarchy



Active Skew Management and Multiple Clock Domains



- + widely dispersed drivers
- + DLLs compensate static and low-frequency variation
- + divides design and verification effort
- DLL design and verification is added work
- + tailored clocks

Power Consumption

- ◆ Clocks consume the largest fraction of power
- ◆ Driving inter-unit busses consumes as much power as intra-unit gates and interconnect




EV4 - 3 Metal Layers

3rd “coarse and thick” metal layer added to the technology for EV4 design

Power supplied from two sides of the die via 3rd metal layer

2nd metal layer used to form power grid

90% of 3rd metal layer used for power/clock routing



Metal 3



Metal 2



Metal 1


EV5 - 4 Metal Layers

4th “coarse and thick” metal layer added to the technology for EV5 design

Power supplied from four sides of the die

Grid strapping done all in coarse metal

90% of 3rd and 4th metals used for power/clock routing



Metal 4



Metal 3



Metal 2





Metal 1

EV6 - 6 Metal Layers

2 reference plane metal layers added to the technology for EV6 design

Solid planes dedicated to Vdd/Vss
Significantly lowers resistance of grid
Lowers on-chip inductance



Reference Plane Example

Simulation Methodology

Extract Inductance & Resistance versus Frequency

Model Skin Effect Both Vertically and Horizontally

Construct Time-Domain SPICE Model and Simulate with SPICE

Use FF Devices, High Vdd & Low Temperature to Aggravate Inductive Effects




COMPAQ

■ \dot{U} Metal 3
Victims


ASP DAC 2000

Reference Plane Example (continued)



RP2 Only

COMPAQ



RP1 & RP2

ASP DAC 2000

De-coupling Capacitor Ratios

◆ EV4

- total effective switching capacitance = 12.5nF
- 128nF of de-coupling capacitance
- de-coupling/switching capacitance ~ 10x

◆ EV5

- 13.9nF of switching capacitance
- 160nF of de-coupling capacitance

◆ EV6

- 34nF of effective switching capacitance
- 320nF of de-coupling capacitance -- not enough!


EV6 De-coupling Capacitance

Design for $\Delta I_{dd} = 25 \text{ A}$ @ $V_{dd} = 2.2 \text{ V}$, $f = 600 \text{ MHz}$

- 0.32- μF of on-chip de-coupling capacitance was added
 - Under major busses and around major gridded clock drivers
 - Occupies 15-20% of die area
- 1- μF 2- cm^2 Wirebond Attached Chip Capacitor (WACC) significantly increases “Near-Chip” de-coupling
 - 160 V_{dd}/V_{ss} bondwire pairs on the WACC minimize inductance

EV6 WACC

389 Signal - 198 VDD/VSS Pins



Clocking Futures

- ◆ Frequencies will continue to scale
- ◆ Clock edge rates are not scaling as well
- ◆ Multiple clock zones required
 - Architectures minimizing global communications
 - Adaptive and passive synchronization techniques
 - DLLs in the near term
 - Clocking schemes utilizing encoding, extraction, multi-state and local phase optimization to compensate for skew and latency
- ◆ Asynchronous or quasi-synchronous architectures

Power Futures

- ◆ Low power modes
- ◆ Power tradeoffs in the micro-architecture
- ◆ More emphasis on a low power circuits
 - Reducing clock load
 - Low swing differential clocks
 - Low swing buses
 - Adiabatic circuits, clocked drivers, retractile logic
 - Asynchronous design
- ◆ Reference plans also help to minimize inductive and wave effects

Conclusion

- ◆ Physical technology advances will enable multi-GHz chips
- ◆ Key challenges in power, clocking, complexity, and verification must be addressed
- ◆ New tools and methods will be needed
- ◆ CAD developers and chip designers must collaborate more closely than ever
- ◆ With a solid understanding of the fundamentals, a clear vision of the product and ingenuity, we will realize multi-GHz microprocessors