Alpha CPU and Clock Design Evolution

- This lecture uses two papers that discuss the evolution of the Alpha CPU and clocking strategy over three CPU generations
  - Gronowski, Paul E., et.al., "High Performance Microprocessor Design", *IEEE Journal of Solid-State Circuits*, Vol. 33, No. 5, May 1998, pp. 676-686
  - Bailey, Daniel W. and Bradley J. Benschneider, "Clocking Design and Analysis for a 600-Mhz Alpha Microprocessor", *IEEE Journal of Solid-State Circuits*, Vol. 33, No. 11, November 1998, pp. 1627-1633
- These papers are interesting because one can see how shrinking device sizes presented new design challenges that forced design methodology changes.
- All notes in this lecture are from these two papers.

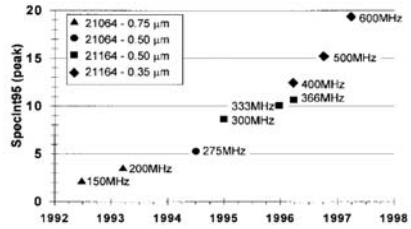

Fig. 1. Alpha performance versus time.

Integer performance graph. Clock speed increased from 150 Mhz to 600 Mhz.

Gronowski, Paul E., et.al., "High Performance Microprocessor Design", *IEEE Journal of Solid-State Circuits*, Vol. 33, No. 5, May 1998, pp. 676-686

TABLE I
MICROPROCESSOR FEATURES

	21064	21164	21264
Transistor Count (million)	1.68	9.3	15.2
Die Size (mm²)	16.8x13.9	18.1x16.5	16.7x18.8
Process Technology	0.75 µm	0.50 µm	0.35 µm
Power Supply (Volts)	3.3	3.3	2.2
Power Dissipation (Watts)	30	50	72
Target Design Frequency (MHz)	200	300	600
Typical Gate Delays/Cycle	16	14	12
On-chip Cache	8-KB I,I-I 8-KB L1-D	8-KB I,I-I 8-KB L1-D 96-KB L2	64-KB I,I-I 64-KB L1-D
Instruction Issue/Cycle	2	4	6
Execution Flow	in-order	in-order	out-of-order

Transistor count – increased by 9X (mainly due to caches)
Die size – slight increase to constant
Frequency – 3X increase
Power dissipation - 2.4X increase
Power supply – reduced by 1/3
Gate delays – indicates longest register-to-register path

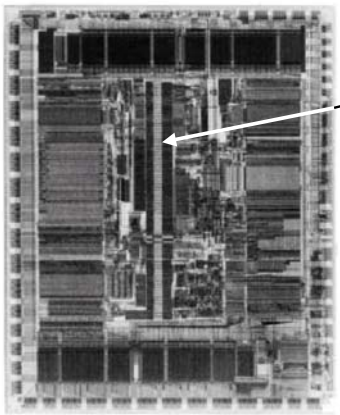

21064 Die Photo

Single Clock driver, 2 transistors for buffer visible to naked eye

Clocking scheme was 2 phase, single wire.

Clock load was 3.5 nF

Gate length of final driver was 35 cm (not a misprint, used serpentine layout to get this gate length).

21064 Clock Skew Distribution

Max clock skew approx. 180 ps (3.6% of 5 ns clock period) – 1 gate delay about 300 ps, so clock skew about 50% of a gate delay.

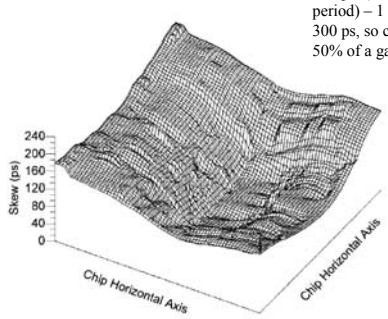

Note the skew is smallest closest to center of chip where driver is located.

Fig. 14. 21064 clock skew.

Thermal Image of 20164

76C at center of chip

46C at edges of chip

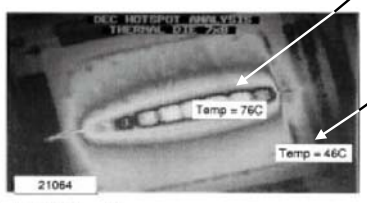

Fig. 16. 21064 thermal image.

30C thermal gradient across chip!!

On Chip Decoupling Capacitors

Large inrush current at clock edges (*di/dt*) required use of on-chip decoupling capacitors placed around clock driver.

Created a decoupling capacitor using NMOS transistor:

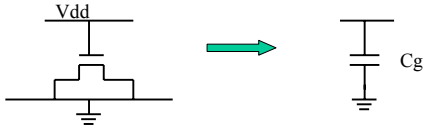

21164 Clock Distribution

Goal of 21164 Clock distribution was to reduce skew by 30% and reduce the thermal gradient.

A predriver was centered between two main clock drivers

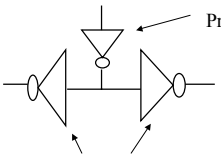

Clock skew was reduced by factor of 2, and thermal gradient was reduced.

Clock skew lowest near two main clock drivers.

Max clock skew approx. 80 ps (2.4% of 3.3 ns clock period) – 1 gate delay about 240 ps, so clock skew about 1/3 of a gate delay.

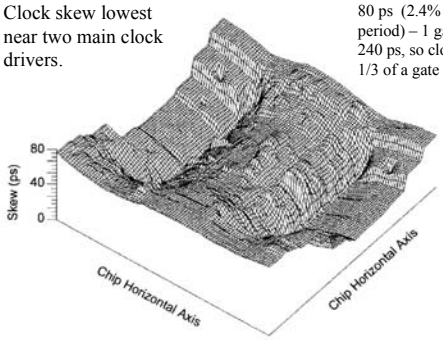

Fig. 17. 21164 clock skew.

Aside: Why a Gridded Clock?

- Both 21064, 21164 used a single global clk distributed by a metal clock grid.
- Skew is largely determined by grid interconnect density and is insensitive to gate load placement
  - Why? Because capacitance of grid wiring dominates the gate loads connected to it.
- Universal availability of clock signals
- Design teams can proceed in parallel since clock constraints well known
- Good process-variation tolerance
- The disadvantage is the extra capacitance of the grid
  - Power-performance tradeoff is determined by choice of skew target, which establishes the needed grid density, which determines the clock driver size.

21264 Clock Distribution

- 21264 clocking fundamentally different from previous Alphas because it supported a hierarchy of clocks
  - Still had a GCLK (global Clk) grid, but conditional and local clocks had several buffer stages after GCLK
- Conditional clocks used to save power
  - Clocks gated to functional units in design
  - If not executing a floating point instruction, then stop the clock to the floating point unit to save power!
- State elements and clocking points were 0 to 8 gates past Gclk
- Six major regional clocks two gain stages past GCLK with grids juxaposed with GCLK, but shielded from it.
  - Major clocks drive local clocks and conditional clocks
- Goals were to improve performance, reduce power.

Clock Hierarchy of 21264

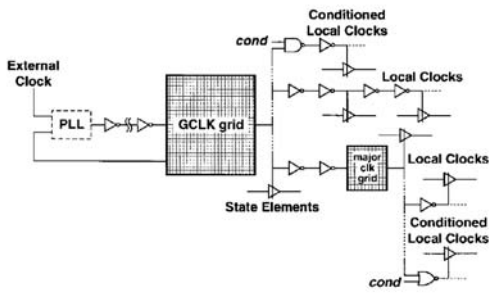

Fig. 1. Clock hierarchy.

21264 Global Clock Distribution

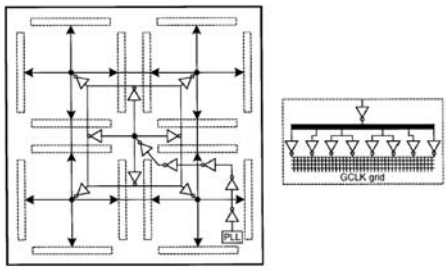

Fig. 2. Global clock distribution network.

Window pane arrangement - same skew to all 'panes' . Note redudant drive to clock nets

Phased Lock Loops (PLLs)

- PLLs and Delay Locked Loops (DLLs) are used to perform clock multiplication of an off-chip clock
  - PLLs/DLLs used to align clock edges of original clock with multiplied clocks
- PLLs are analog circuits that use a charge pump and a voltage controlled oscillator (VCO) to perform phase alignment
  - Alpha 21264 PLL used a separate, regulated 3.3 V supply and was located in the corner of the chip to minimize noise impact
  - Section 9.5.2 of Rabaey text has a block diagram of a PLL
- All high performance CPUs and most ASICs now include a PLL for internal clock generation

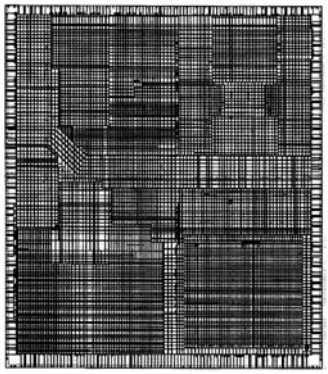

Fig. 3. Global clock grid.

Global clock grid.

Uses 3% of M3/M4 routing layers

(lines in picture are misleadingly thick).

All GCLK lines are laterally shielded by Vss/Vdd

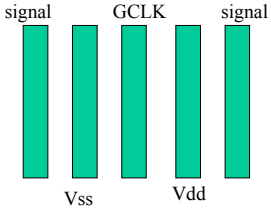

Lateral shielding via Vss/Vdd prevents clock noise from coupling into signal lines.

Clock wires and lateral shields were manually placed

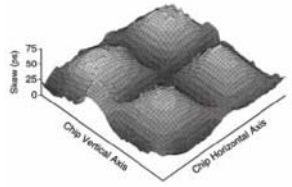

Fig. 4. Simulated global clock skew.

Simulated worst case GCLK skew was 72 ps .

Skew on M1, M2 was less than 10 ps.

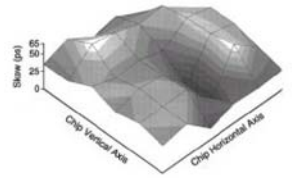

Fig. 5. Measured global clock skew.

Measured worst case GCLK skew via ebeam tester was 65 ps

TABLE I

Major Clock	Area
CCLK	bus interface unit
ECLK	integer issue and execution units
FCLK	floating point issue and execution units
JCLK	instruction fetch and branch prediction unit
MCLK	load/store unit
PCLK	pad ring

Major clocks are two inversions past GCLK

Major clocks saved power over a single global clocks because they service a lighter load and distribution area is smaller – both of these means smaller drivers are needed.

Gclk+Major clocks used 24 W @ 2.2 V, 600 Mhz. It is estimated that at least 40W would have been required if only global clocks were used. 10%-90% rise/fall times were targeted at < 320 ps.

Major Clock grids.

Densest major clock grids used up to 6% of M3/M4 routing.

White areas are serviced by local clocks, local clocks also present in major clock grids.

Major clocks also laterally shielded.

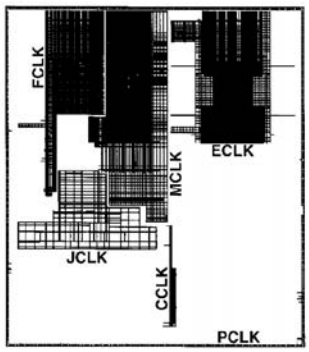

Fig. 6. Major clock grids.

Local, Conditional Clocks

- Local clocks generated from any clock – GCLK, Major clocks, other local clocks
- Local clocks were neither shielded or gridded
- Having local clocks gave freedom to move clock edges with respect to data to solve timing problems
  - This is another form of "time borrowing"
- 60,000 local clocks nodes, all were analyzed with SPICE using minimum and maximum gate capacitance estimates
  - Some local clocks had very high min/max delay variation tolerances (up to 280 ps)

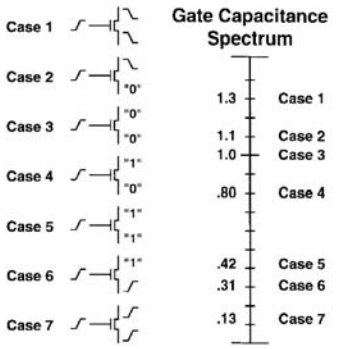

"effective" gate input capacitance is not a constant – depends on state of other inputs.

Causes timing characterization headaches!!

Fig. 8. Data-dependent gate loading.

Power Consumption 21264

- 72 W total (600 Mhz @ 2.2 V)
- Clock distribution power consumption: 46.8 W
  - Gclk 10.2 W
  - Major Clocks 24 W
  - Local unconditional clocks 7.6 W
  - Local conditional clocks 15.6
- Clocking accounted for 65% of the total power in the 21264!

Power Distribution 21064

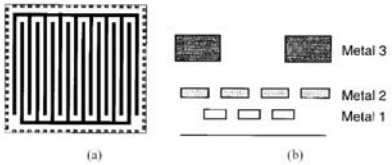

Fig. 10. (a) 21064 metal layers and (b) power distribution.

M3 vertical, M2 horizontal. M3 width was twice the process minimum. Vdd/Gnd bond pads only on top/bottom, connected to M2.

Supply current of 9A (30W @ 3.3 V).

Power Distribution 21164

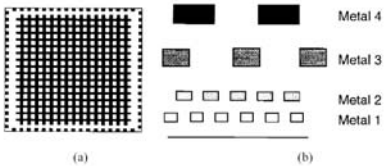

Fig. 11. (a) 21164 metal layers and (b) power distribution.

M3, M4 used for power grid.

Vdd/Gnd bond pads on all 4 sides.

Supply current was 15.2A (50W @ 3.3 V).

Power Distribution 21264

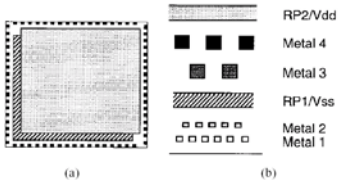

Fig. 12. (a) 21264 metal layers and (b) power distribution.

Supply current of 33A! (72 W @ 2.2 V).

Difference in current requirements between adjacent clock phases because of conditional clocking could be as high as 25A. This required two, thick, low resistance aluminum reference planes to be used for Vdd/Vss. The VSS plane reduced crosstalk between M3 and M2 signals.

15%-20% of the die area was used for on-chip decoupling capacitors.

21064, 21164 used level-based design (latch based)

- 21064 used a version of the TSPC latch (true-single phase clock)
- 21164 used a simpler latch design that had less latency
- Both embedded logic into latches to reduce gate delays

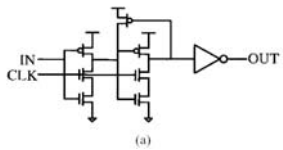

21064 latch design, dynamic

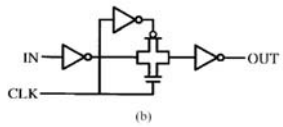

21164 latch design, dynamic.

Embedded logic in 21064, 21164 latches

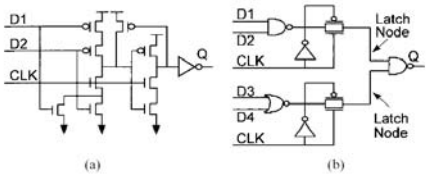

Fig. 22. Embedding logic into latches: (a) 21064 Function Latch: one level of logic; (b) 21164 Function Latch: two levels of logic.

Embedding logic in latches helps to reduce the number of gate delays between latches.

21264 used edge-triggered design

Because of gated clock usage to save power, needed static latches.

Went to edge-triggered design to simply timing verification since timing was now much more complicated due to conditional clocking

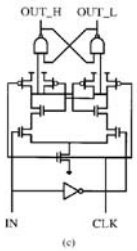

Family of edge-triggered FFs built using this static latch design.

Timing Analysis 21264

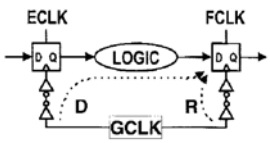

Critical Path Definition and Criteria

- Identify Common Clock, D and R
- Maximize D
- Minimize R
- $D + t_{SETUP} - R \le T_{CYCLE}$ [1 Cycle]
- $D + t_{SETUP} - R \le T_{PHASE}$ [½ Cycle]

ECLK, FCLK major clocks derived clocks from GCLK.

Need to hand values off between clock domains.

D is drive path delay time.

R is receive path delay time.

Worst case delay analysis maximizes D, minimizes R (if R arrives early, then can clock in data before it arrives).

Timing Analysis 21264 (cont.)

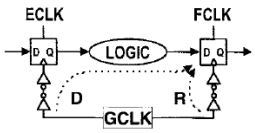

Race Definition and Criteria

- Identify Common Clock, D and R
- Minimize D
- Maximize R
- $D / (R + t_{HOLD}) \ge X$
- $X \ge 1$, X determines margin

Race analysis determines if value produced by ECLK is actually clocked in by FCLK before a new value from ECLK arrives.

Race analysis minimizes D, maximizes R.

Worst Case Analysis

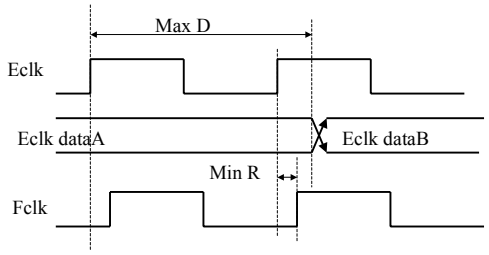

Fclk stage gets incorrect value because new data is not ready yet before Fclk edge arrives. Delay on Fclk helps in this case.

Race Analysis

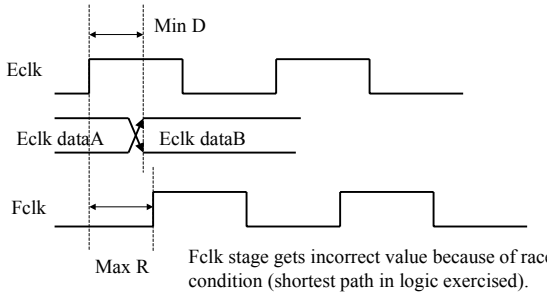

Fclk stage gets incorrect value because of race condition (shortest path in logic exercised). Delay on Fclk hurts in this case.

Conditional Clock Analysis

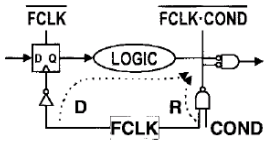

Same analysis (worst case, race) applied for Conditional Clocks.