CS250
VLSI Systems Design

Lecture 4: Physical Realities: Beneath the Digital Abstraction, Part 1: Timing

Fall 2012

John Wawrzynek, Jonathan Bachach
with
John Lazzaro, Krste Asanović
and
Rimas Avizienis

What do Computer Architects need to know about physics?

- Physics effect:

Area ⇒ cost
Delay ⇒ performance
Energy ⇒ performance & cost

- Ideally, zero delay, area, and energy. However, the physical devices occupy area, take time, and consume energy.
- CMOS process lets us build transistors, wires, connections, and we get capacitors, inductors, and resistors whether or not we want them.

Physical Layout

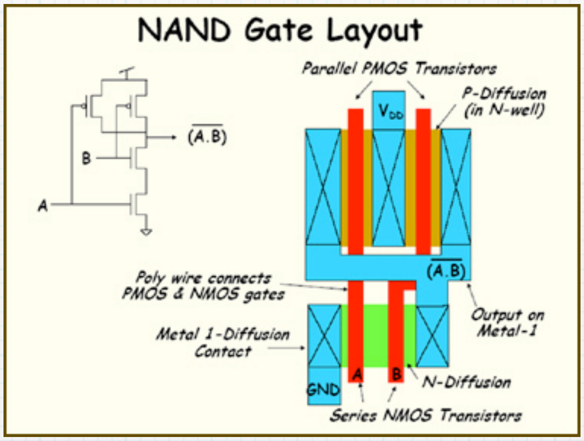

- "Switch-level" abstraction gives a good way to understand the function of a circuit.
  - nFET (g=1 ? short circuit : open)
  - pFET (g=0 ? short circuit : open)
- Understanding delay means going below the switch-level abstraction to transistor physics and layout details.

"Gate Delay"

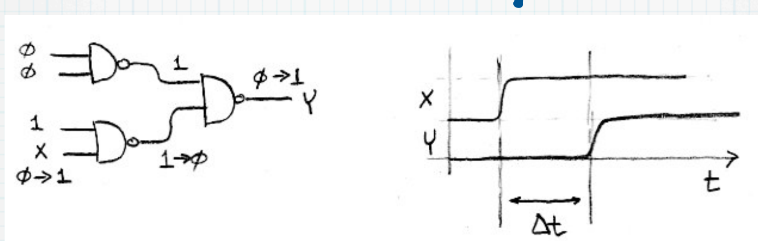

- Modern CMOS gate delays on the order of a few picoseconds. (However, highly dependent on gate context.)
- Often expressed as FO4 delays (fan-out of 4) - as a process independent delay metric:
  - the delay of an inverter, driven by an inverter 4x smaller than itself, and driving an inverter 4x larger than itself.
  - For our 90nm process FO4 is around 20ps.

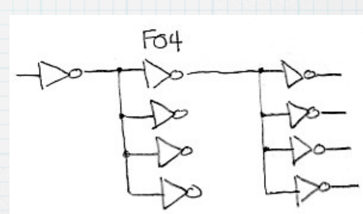

"Path Delay"

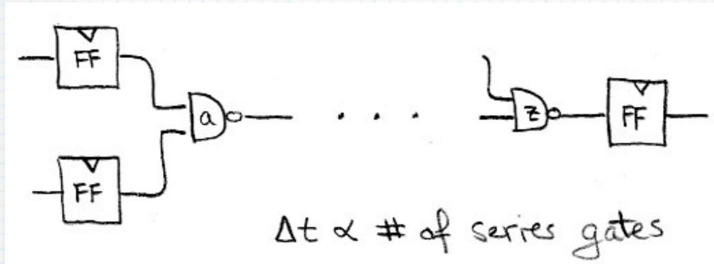

- For correct operation:
  Total Delay $\leq$ clock_period - FF_{setup_time} - $FF_{clk\_to\_q}$ - Clock_skew
  on all paths.

- High-speed processors critical paths have around 10-20 FO4 delays.

FO4 Delays per clock period

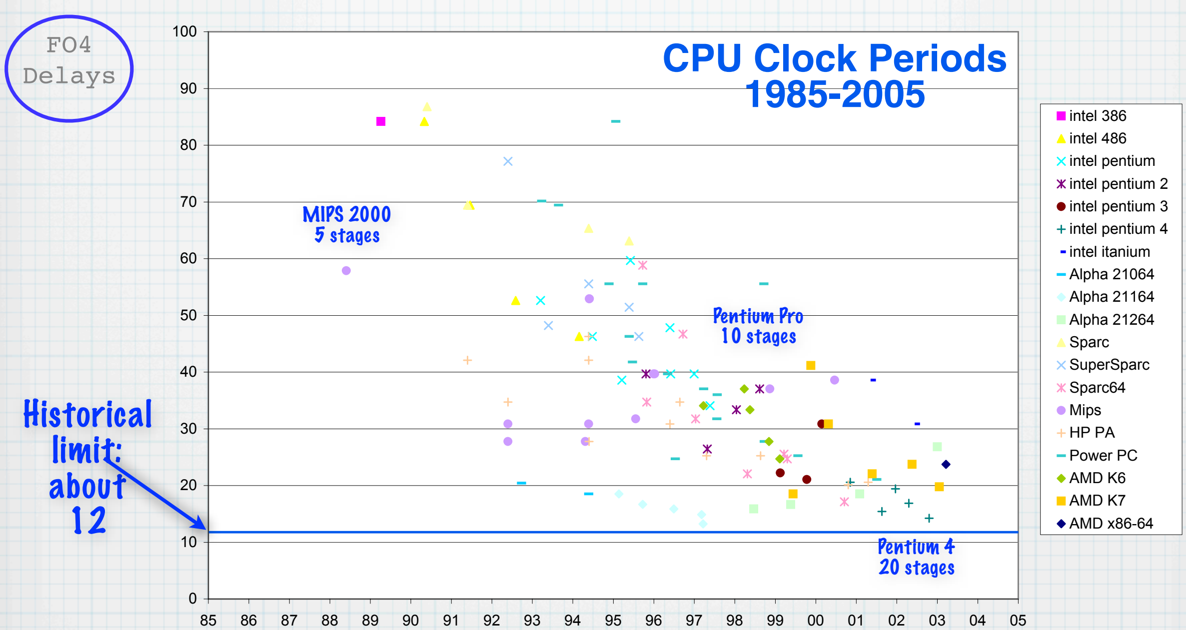

6

Thanks to Francois Labonte, Stanford

"Gate Delay"

- What determines the actual delay of a logic gate?
- Transistors are not perfect switches - cannot change terminal voltages instantaneously.
- Consider the NAND gate:

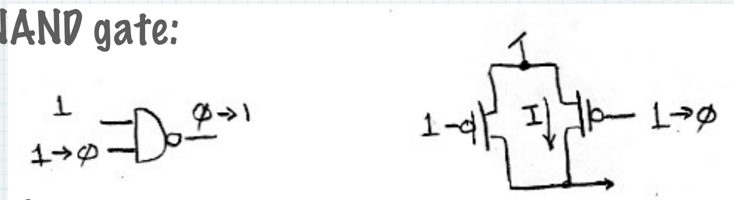

  - Current (I) value depends on: process parameters, transistor size

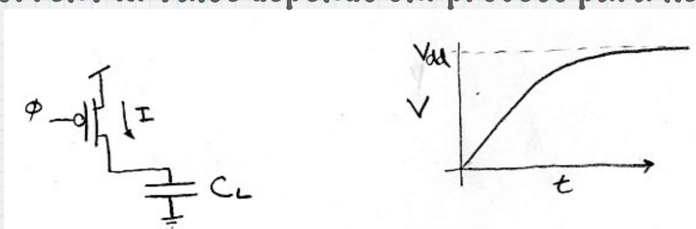

$\Delta \propto C_L / I$

- C_L models gate output, wire, inputs to next stage (Cap. of Load)
- C "integrates" I creating a voltage change at output

More on transistor Current

- Transistors act like a cross between a resistor and "current source"

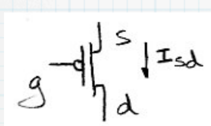

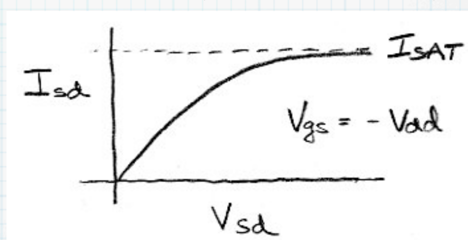

- I_{SAT} depends on process parameters (higher for nFETs than for pFETs) and transistor size (layout):

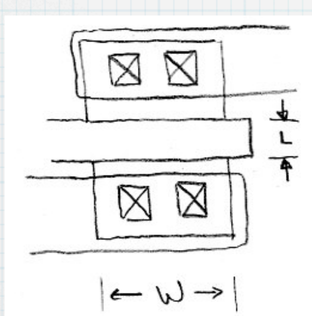

$I_{SAT} \propto W/L$

More on C_L

- Everything that connects to the output of a logic gate (or transistor) contributes capacitance:

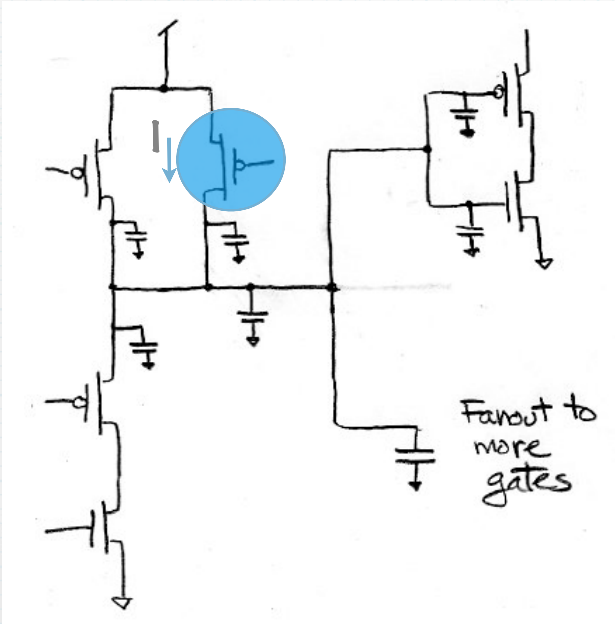

- Transistor drains
- Interconnection (wires/ contacts/vias)
- Transistor Gates

Wires

- So far, simple capacitors:

  $C \propto$ Area = width $*$ length

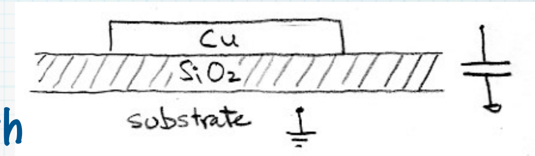

- Wires have finite resistance, so have distributed R and C:

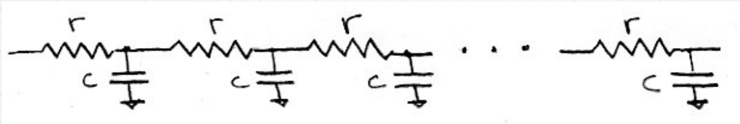

with r = res/length, c = cap/length, $\Delta \propto rcL^2 \cong rc + 2rc + 3rc + \ldots$

- For short wires (between gates) R is insignificant (total RC delay << gate delay)
- For long wires R becomes significant. Ex: busses, clocks, reset
- "rebuffering" helps

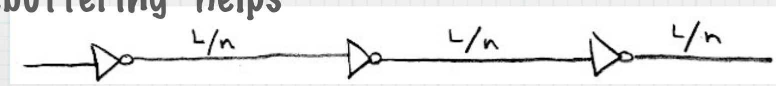

Turning Rise/Fall Delay into Gate Delay

- Cascaded gates:

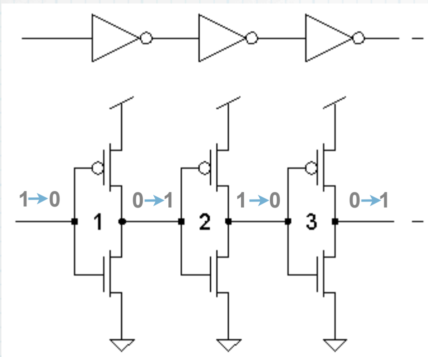

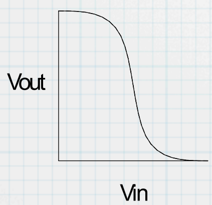

"transfer curve" for inverter.

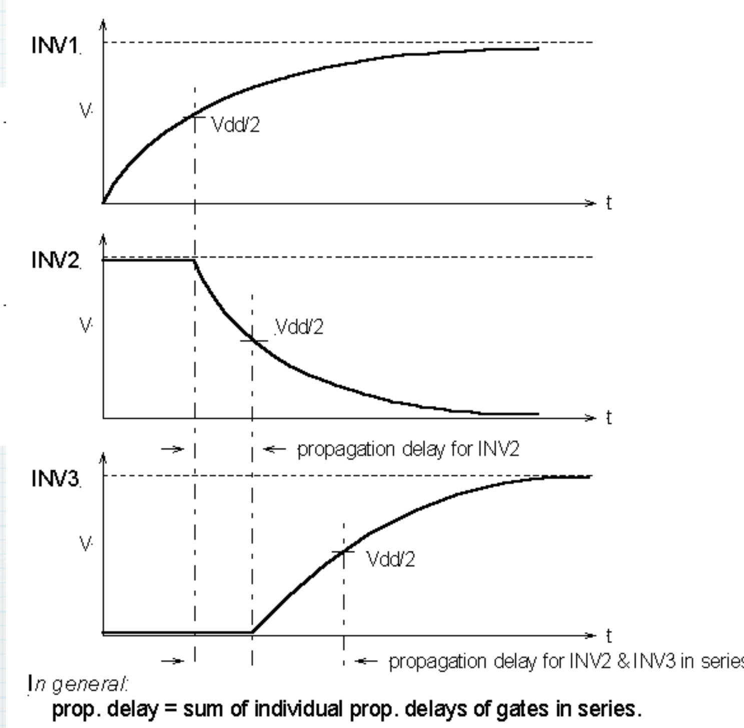

Driving Large Loads

- Large fanout nets: clocks, resets, memory bit lines, off-chip
- Relatively small driver results in long rise time (and thus large gate delay)

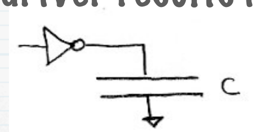

- Strategy:

Staged Buffers

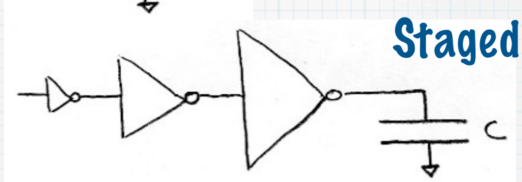

- Optimal trade-off between delay per stage and total number of stages ⇒ **fanout of ~4-6 per stage**

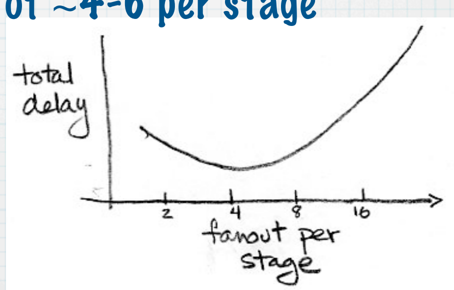

Components of Path Delay

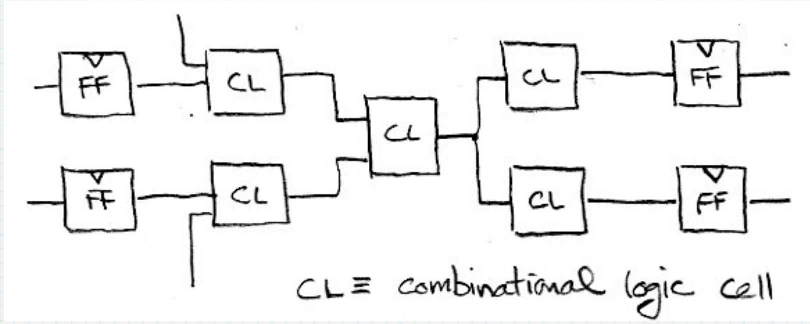

1. # of levels of logic
2. Internal cell delay
3. wire delay
4. cell input capacitance
5. cell fanout
6. cell output drive strength

Who controls the delay?

	foundary engineer (TSMC)	Library Developer (Aritsan)	CAD Tools (DC, IC Compiler)	Designer (Rimas)
1. # of levels			synthesis	RTL
2. Internal cell delay	physical parameters	cell topology, trans sizing	cell selection	
3. Wire delay	physical parameters		place & route	layout generator
4. Cell input capacitance	physical parameters	cell topology, trans sizing	cell selection	instantiation
5. Cell fanout			synthesis	RTL
6. Cell drive strength	physical parameters	transistor sizing	cell selection	instantiation

Timing Closure: Searching for and beating down the critical path

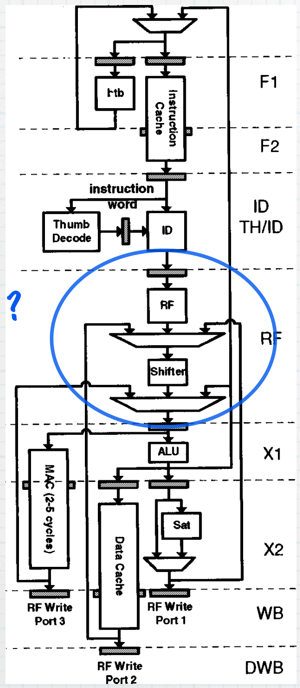

Must consider all connected register pairs, paths from input to register, register to output. Don't forget the controller.

- Design tools help in the search.
- **Synthesis tools** work to meet clock constraint, report delays on paths,
  - Special **static timing analyzers** accept a design netlist and report path delays,
  - and, of course, **simulators** can be used to determine timing performance.

Tools that are expected to **do something** about the timing behavior (such as synthesizers), also include provisions for specifying input arrival times (relative to the clock), and output requirements (set-up times of next stage).

Timing Analysis, real example

The critical path

Most paths have hundreds of picoseconds to spare.

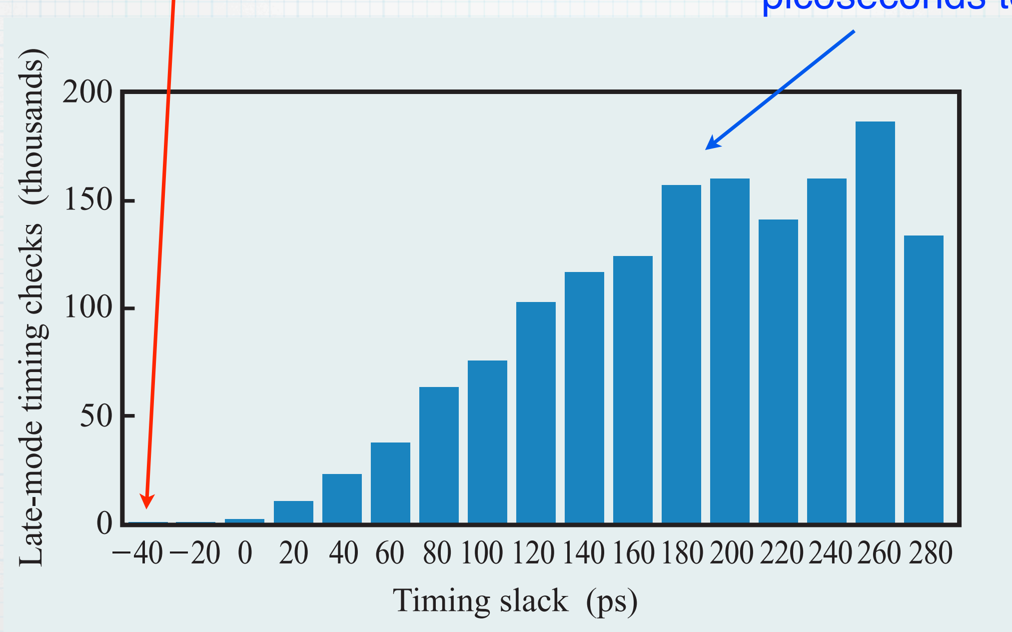

From "The circuit and physical design of the POWER4 microprocessor", IBM J Res and Dev, 46:1, Jan 2002, J.D. Warnock et al.

Timing Analysis Tools

- Static Timing Analysis: Tools use delay models for gates and interconnect. Traces through circuit paths.
  - Cell delay model capture
    - For each input/output pair, internal delay (output load independent)
    - output dependent delay
- Standalone tools (PrimeTime) and part of logic synthesis.
- Back-annotation takes information from results of place and route to improve accuracy of timing analysis.
- DC in "topographical mode" uses preliminary layout information to model interconnect parasitics.
  - Prior versions used a simple fan-out model of gate loading.

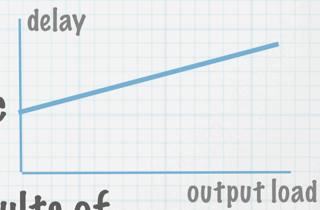

Hold-time Violations

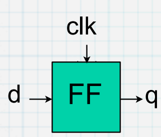

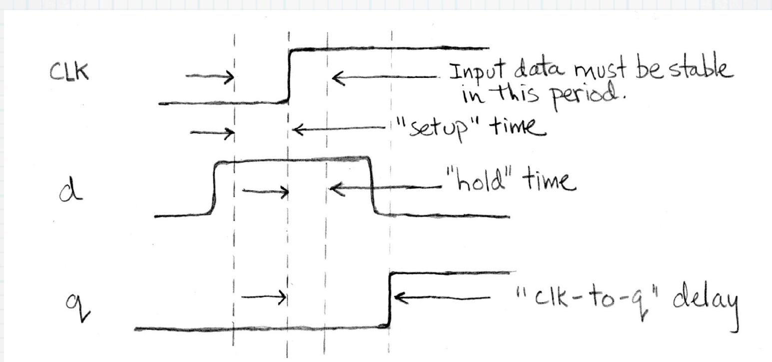

- Some state elements have positive hold time requirements.
  - How can this be?
- Fast paths from one state element to the next can create a violation. (Think about shift registers!)
- CAD tools do their best to fix violations by inserting delay (buffers).
  - Of course, if the path is delayed too much, then cycle time suffers.
  - Difficult because buffer insertion changes layout, which changes path delay.

Conclusion

- Timing Optimization: You start with a target on clock period. What control do you have?
- Biggest effect is RTL manipulation.
  - i.e., how much logic to put in each pipeline stage.
  - We will be talking later about how to manipulate RTL for better timing results.
- In most cases, the tools will do a good job at logic/circuit level:
  - Logic level manipulation
  - Transistor sizing
  - Buffer insertion
  - But some cases may be difficult and you may need to help
    - Hand instantiate cells, layout generators

End of Physical Realities part 1 Timing