Bridging the gap between early physical and electrical wiring projections

Eli Chiprout
Strategic CAD, Intel Labs
Chandler AZ/Hillsboro OR
eli.chiprout@intel.com

April 22, 2003

Implications of design decisions on wires

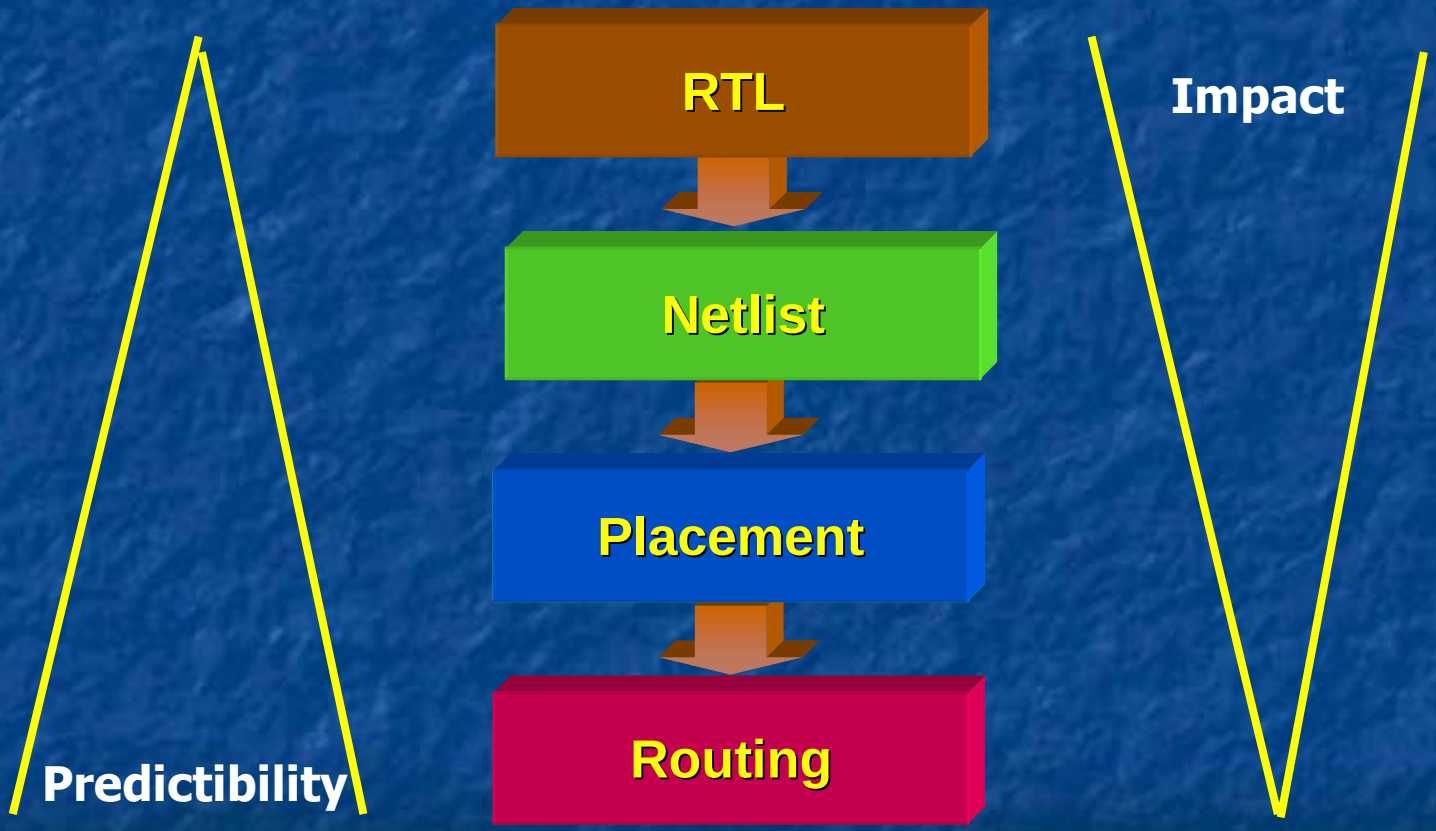

Implications of design decisions on wires

Implications of design decisions on wires

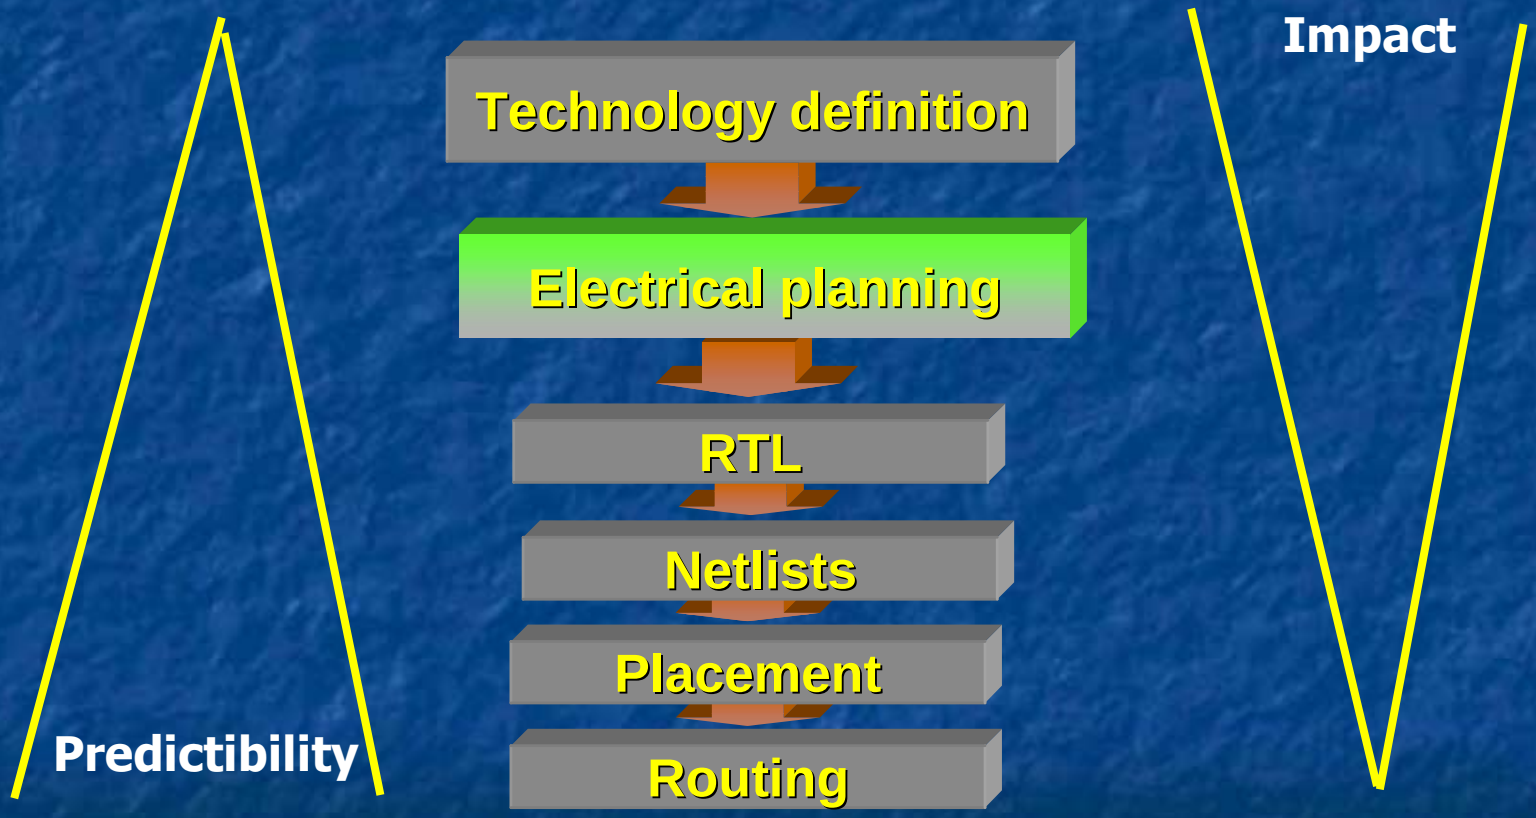

Pre-defined electrical constraints

- We generally think in terms of placement of our basic blocks and then a projection of the wiring via early wiring estimates (e.g. Rent's rule)
- Increasingly massive early (pre-placement) electrical global constraints are consuming wiring resources and are defining what we can do with the wires
- We have no clear methods to project the impact of early electrical planning on routability, congestion and tradeoff weights between wiring and electrical constraints
- Up front constraints are in power grid (current carrying capacity, droop, noise immunity), clock (noise immunity, signal loss, variational immunity), signals (signal loss, noise immunity, traversal length), etc.
- Design variables: width/pitch of signals, width/pitch of power, signal to rail ratio, repeater bays & via stacks, area assigned to clock, wrong way wires…
- **Up-front electrical wire planning will be the major wiring constraint for the next generations of microprocessors.**

Top level constraints

- Mantra: power, power, power
- Resistance increasing each process generation
- Frequency increasing
- Cell-based design CBD forces early fixing of regular structures
  - **power delivery**
  - **clock delivery**
  - **repeater bays**
  - **etc.**
- CBD used to reduce manpower and time to market

Moore's Law Continues

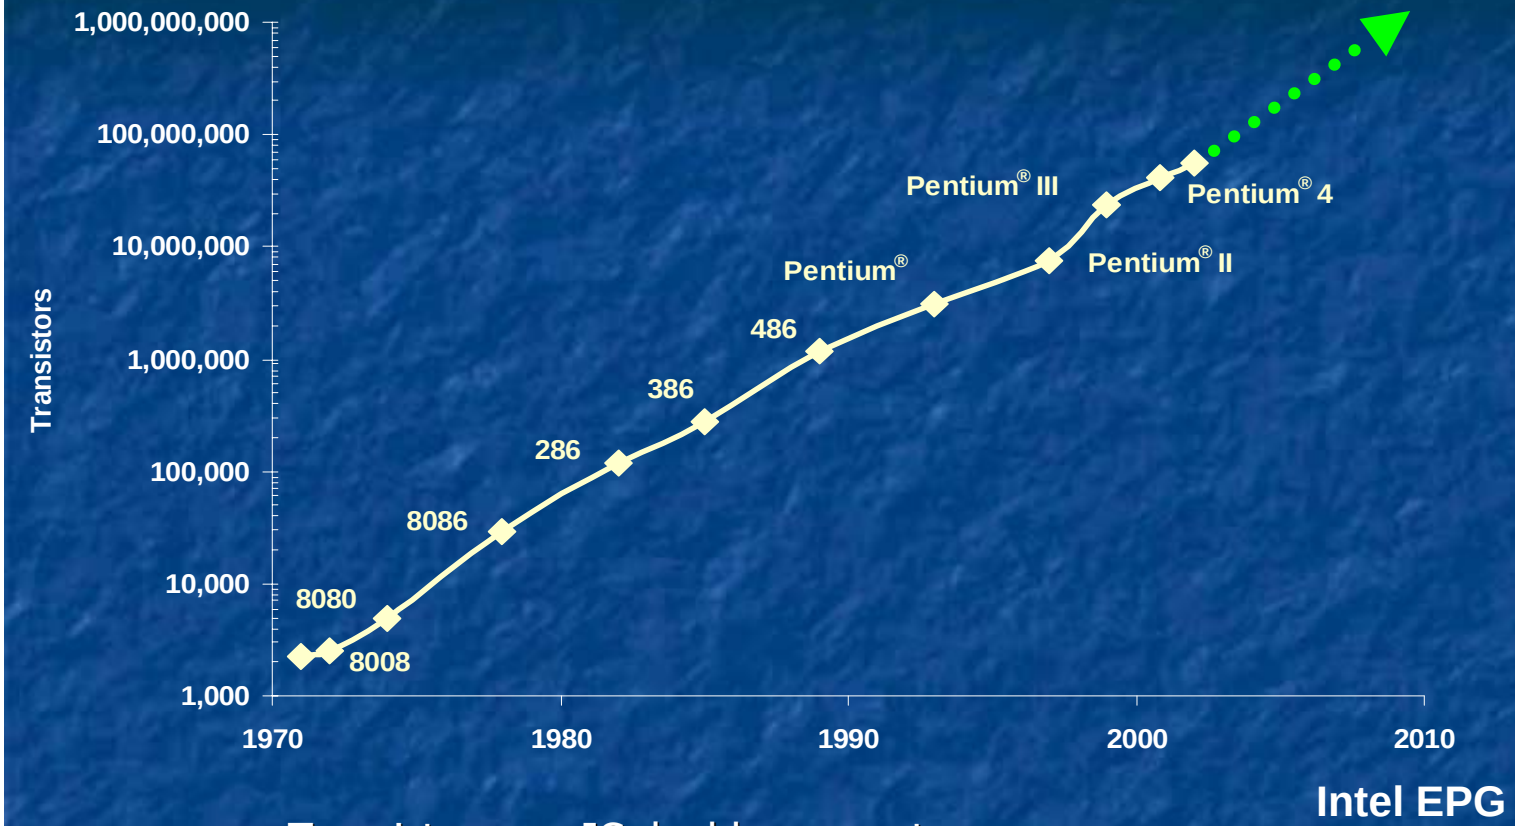

- Transistors per IC doubles every two years
- In less than 30 years
  - 1,000X decrease in size
  - 10,000X increase in performance
  - 10,000,000X reduction in cost
- Heading toward 1 billion transistors before end of this decade

Processor frequency trend

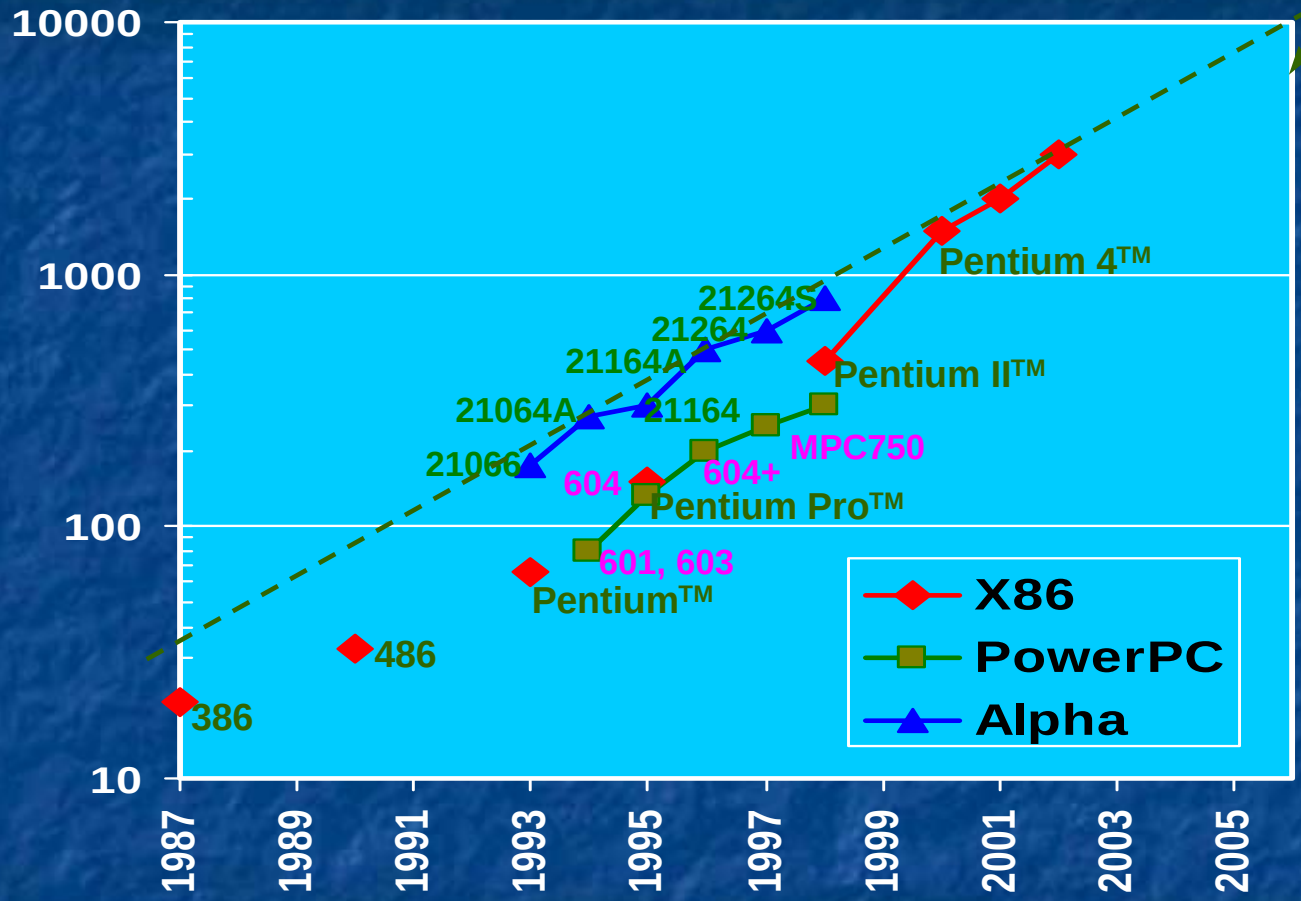

$P \propto CV^2f$

Frequency doubles each generation

Processor die size trend

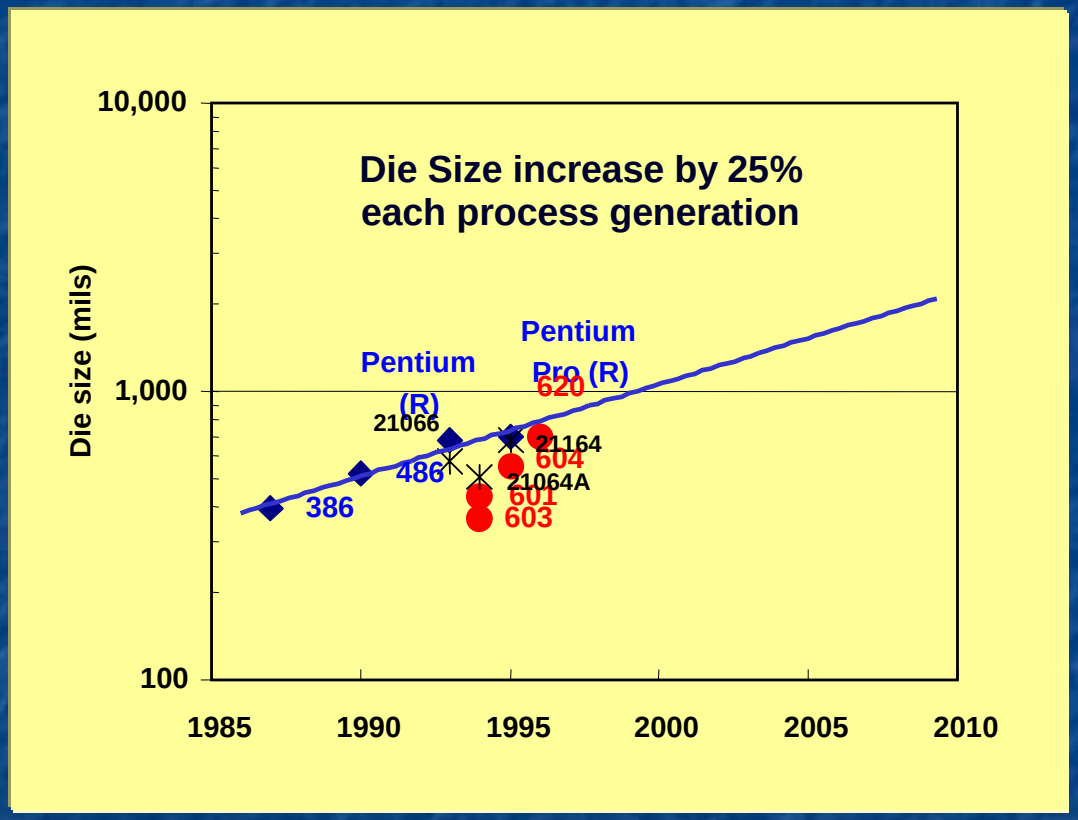

Intel labs

Processor power trend

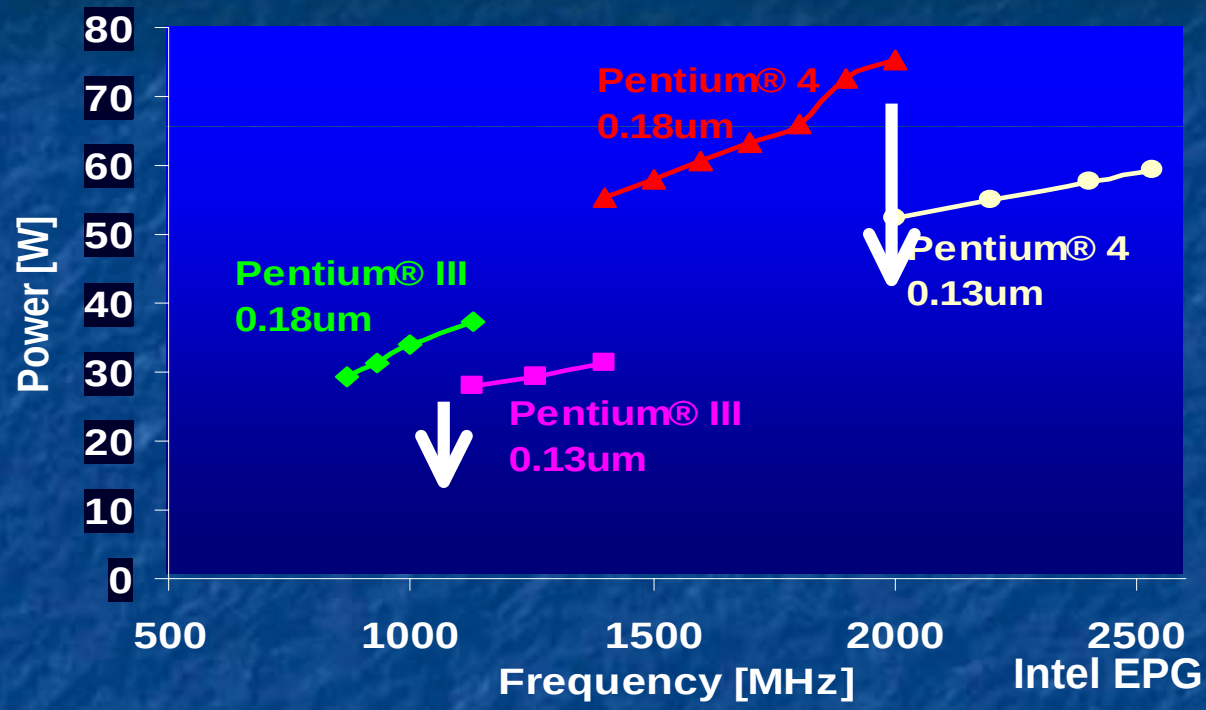

- Process scaling provides higher performance at lower power but number of devices swamps out effect for next generation. Also increasing:
  - Leakage
  - Thermal envelopes
- **Lead** processor power increases every generation
- Vcc will scale by only 0.85 (not 0.7)
- Active power density will increase by ~30-80% (not constant)
- Leakage power will make it worse as process shrinks

Processor power density trend

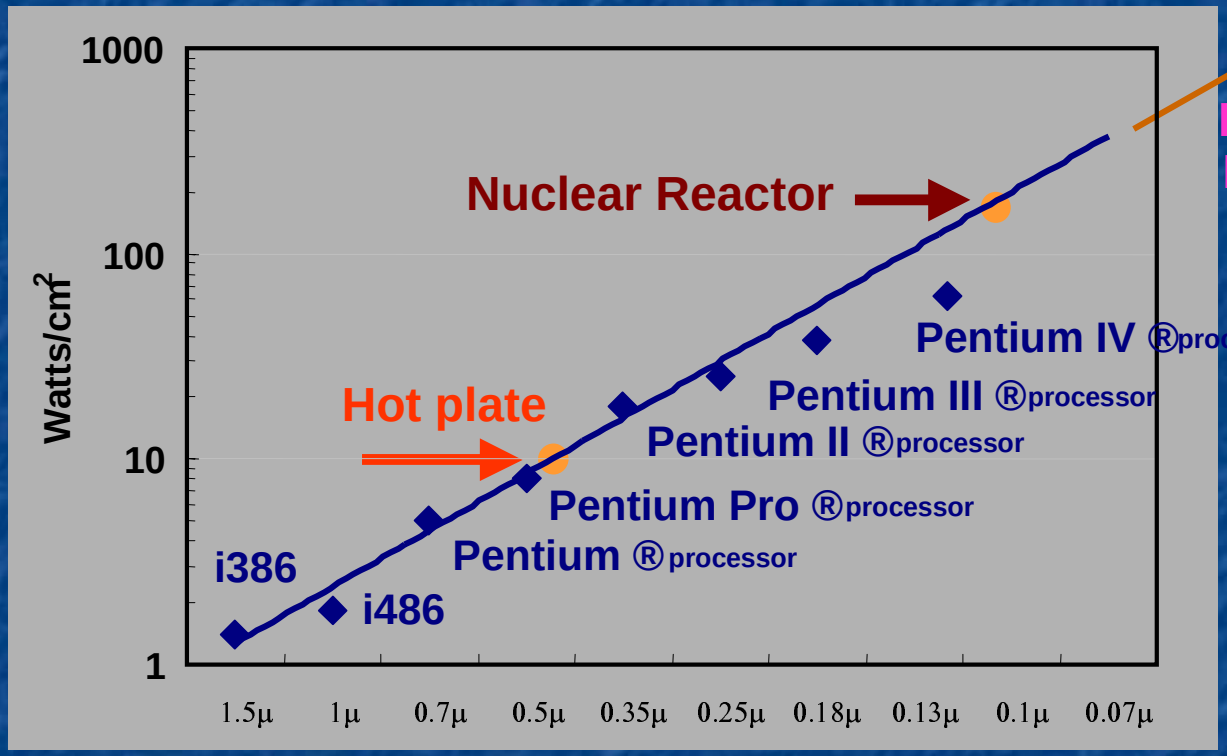

- 60Watt/cm^2 today may rise to 200Watts/cm^2 for 45nm by 2010.

Current scaling

- Power density increases 80% per generation
- Vdd scales by 0.85X per generation
- Power translates into current: doubles every generation
- Add to this effects of Leakage
- The cost to maintain the same relative IR drop is high

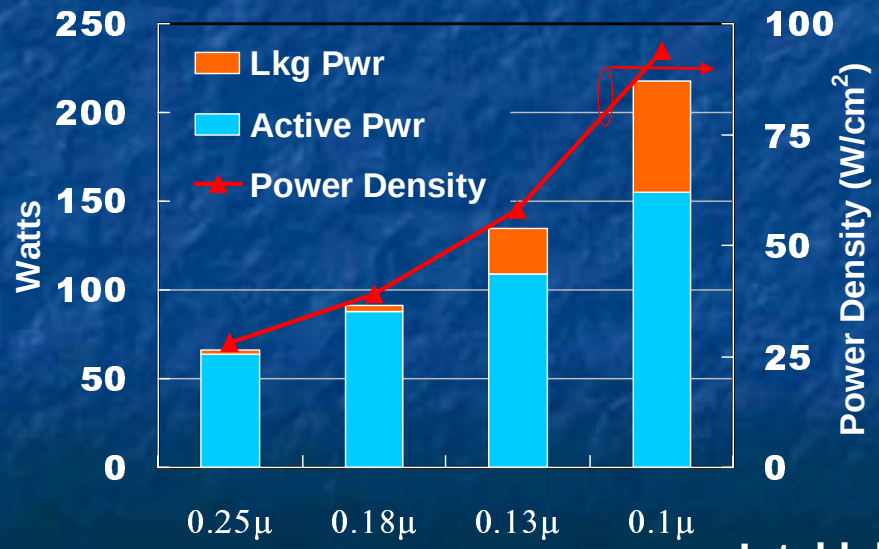

Current (I) Constraints on Wiring

- More power hungry devices mean more current being drawn through the power grid
  - Increase in IR drop
  - Electromigration and self-heating
- Power grid must be made wider or take up more area as the current gets larger
  - especially as resistance goes up
- C4 density and capacity goes up and power grid must also match frequency and current limits
- A difficult Intel design constraint happened when early wide power rails caused a routability problem later on
- Extremely difficult to change power grid later on when congestion constraints become available (CBD)

Power density and temperature

Power Map

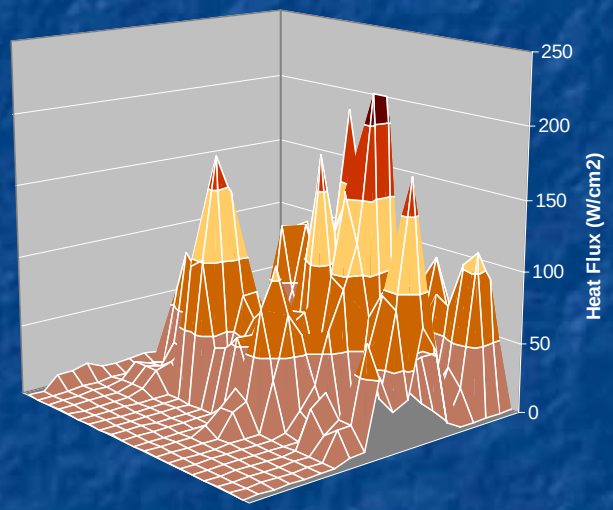

On-Die Temperature

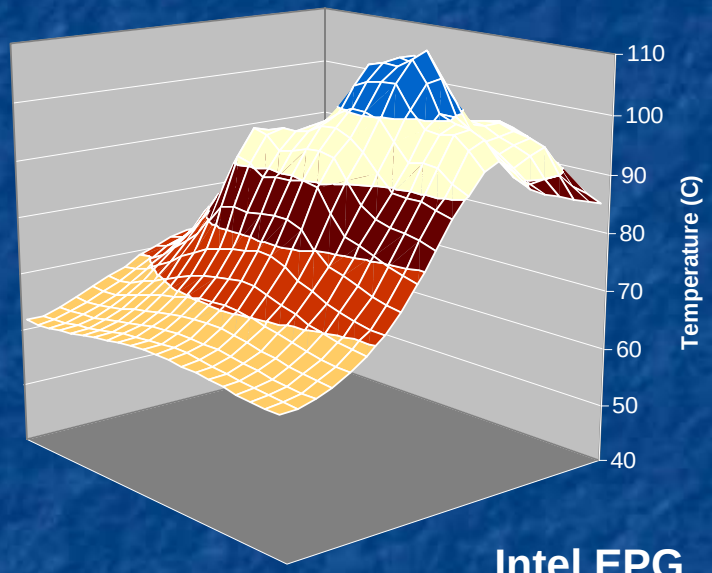

Intel EPG

- With high power density, cannot assume uniformity
  - As die temperature increases stress on wires goes up
  - At high die temp., long-term reliability can be compromised

From early power projections to estimated current/thermal impact

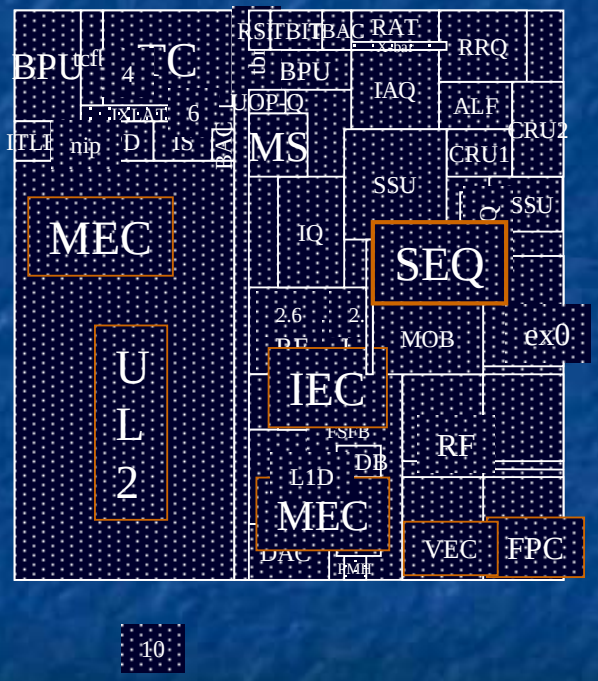

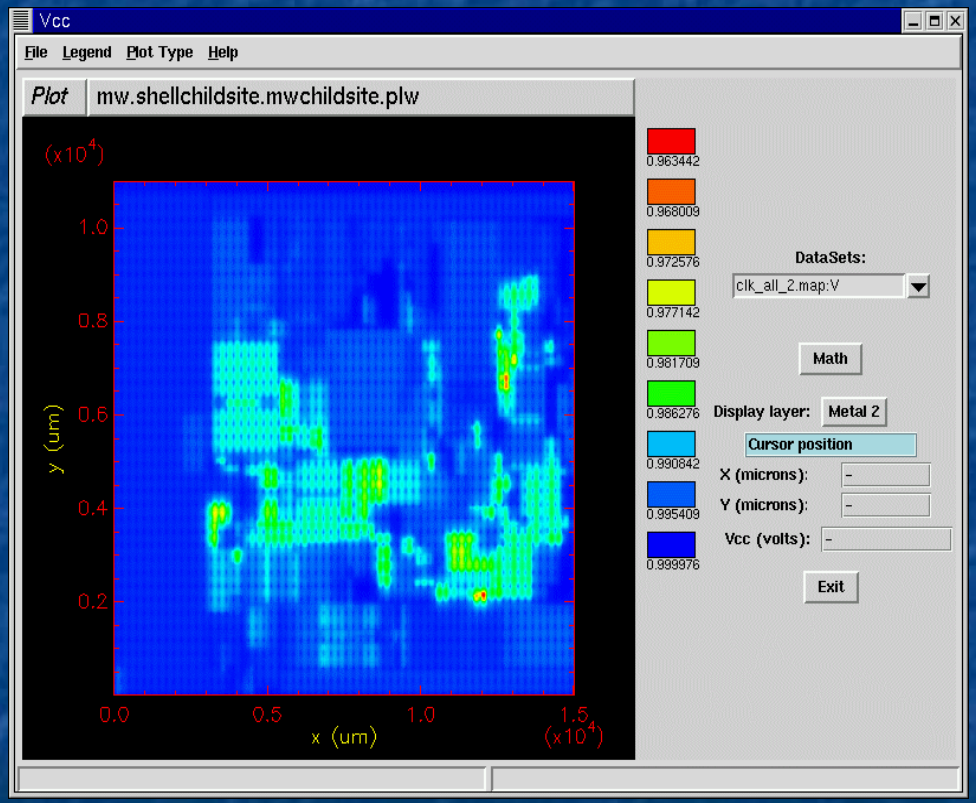

Power grid and temperature constraint on wiring

- More power hungry devices mean hot spot areas on die
- Power grid must also satisfy electromigration constraints, especially in hot spot areas
- Grid made wider to address EM (area increase won't address problem) taking up more routing resources
- Electromigration constraints near the C4 bumps can mean wrong way wiring!
- Need to understand impact on routability and congestion as these **early** decisions are being made
- Routing limitations later on have forced us to change the early electrical specifications

Frequency advances faster than gate delay reduces

- Extra frequency scaling
  - Reduce the levels of logic between flops

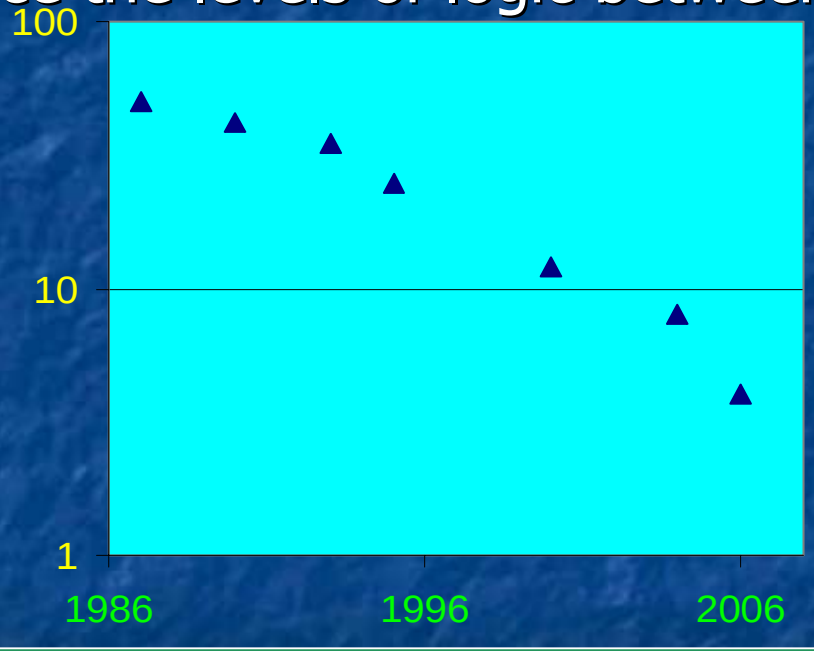

Number of gates/clock reduces by 25% per generation

High-frequency implication

- How many stages of logic will fit in a cycle for 65nm technology?
  - very few!

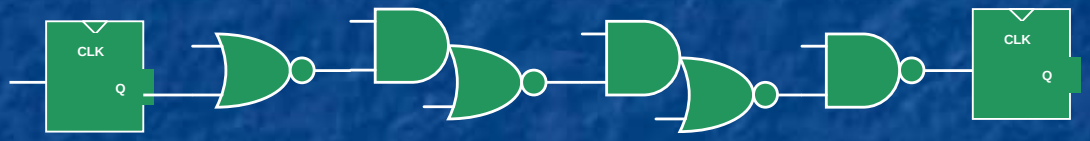

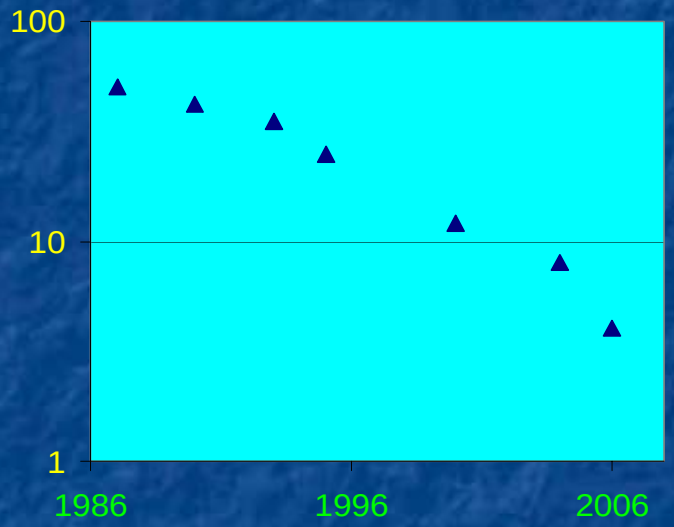

- Paths with several stages will temporarily spread out activity compared to paths with fewer stages

Frequency and first droop

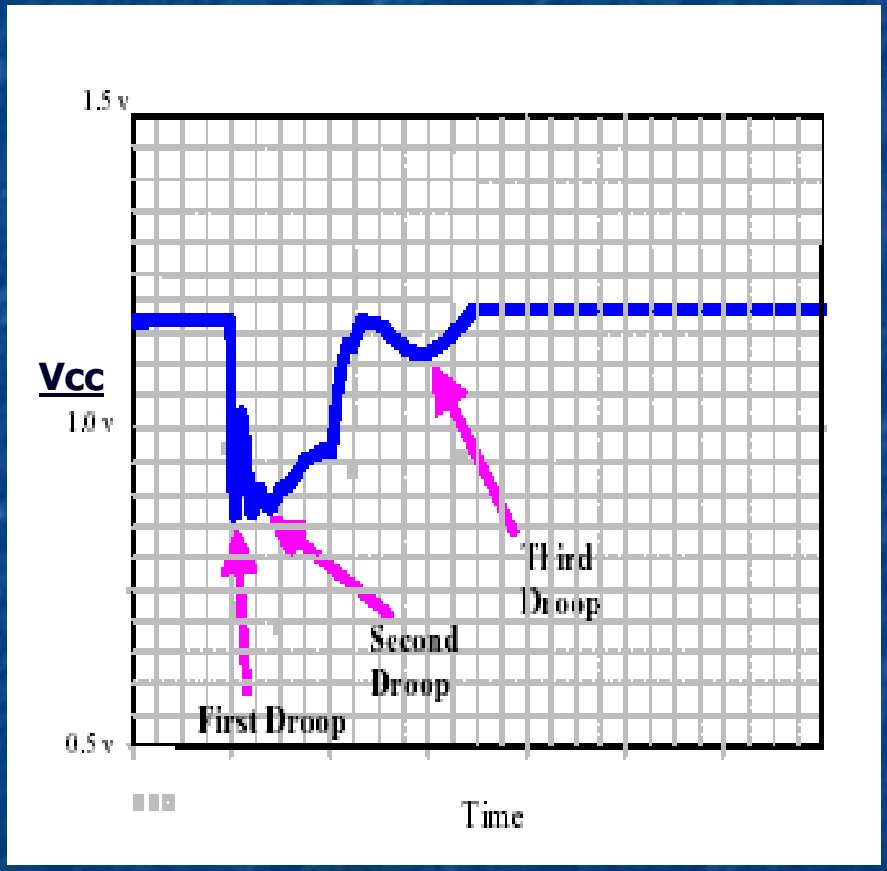

- More **simultaneous** events cause more droop effects (clock + logic)
- First droop and general droop is a serious design constraint
- Power grid must be able to tackle first droop
- Possibly enhanced power grid around first droop areas of die at the cost of non-uniform wiring resources

Ideal scaling

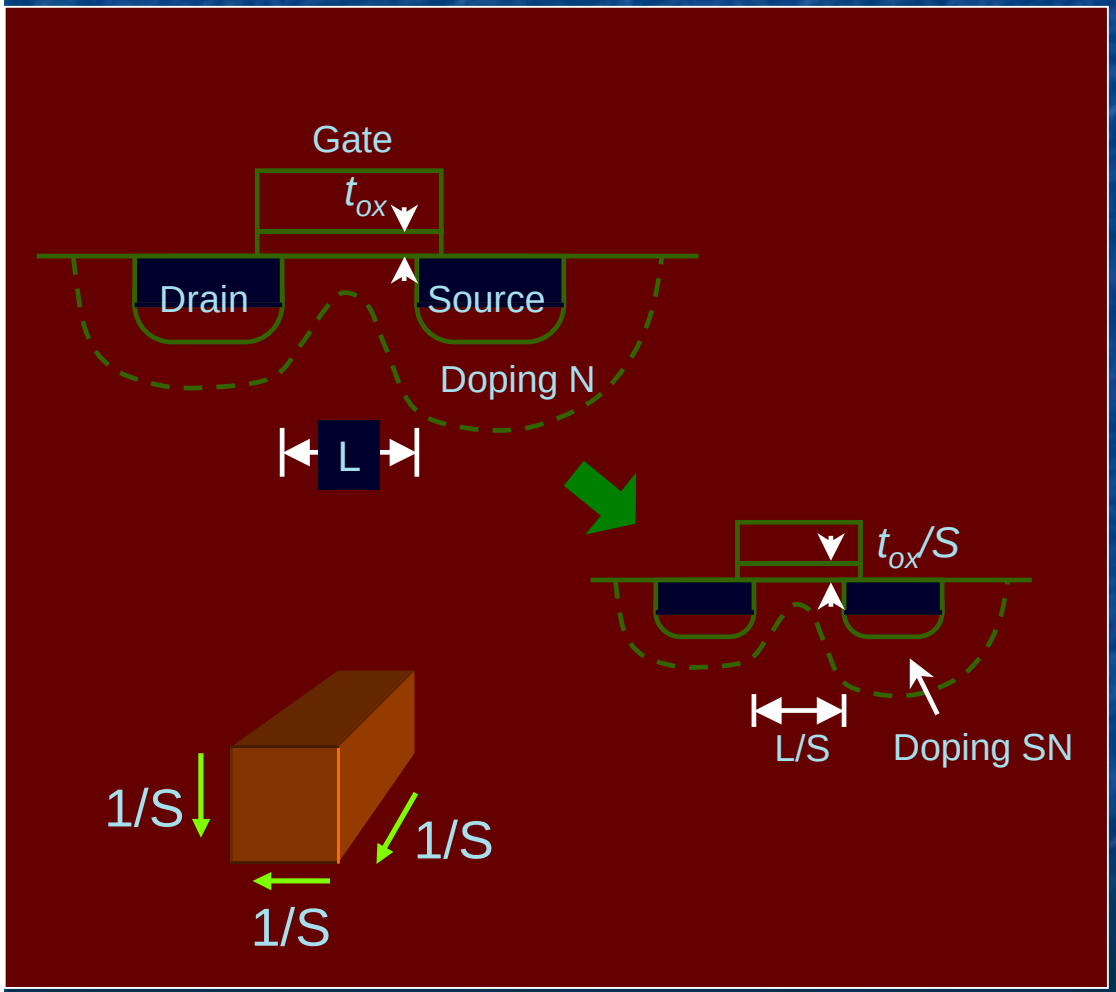

Dimension	1/S	
Die size	S_c	
V_{cc}	1/S	Devices
V_T	1/S	
C_{gate}	1/S	
R_{eff}	1	
τ_{gate}	1/S	
C_{int}	1	
R_{int}	S^2	Interconnect
τ_{int}	S^2	
L_{loc}	1/S	
τ_{loc}	S	
L_{global}	S_c	
τ_{global}	$S^2S_c^2$	
P_{gate}	$1/S^2$	
P_{total}	S_c	

Global wire length scaling

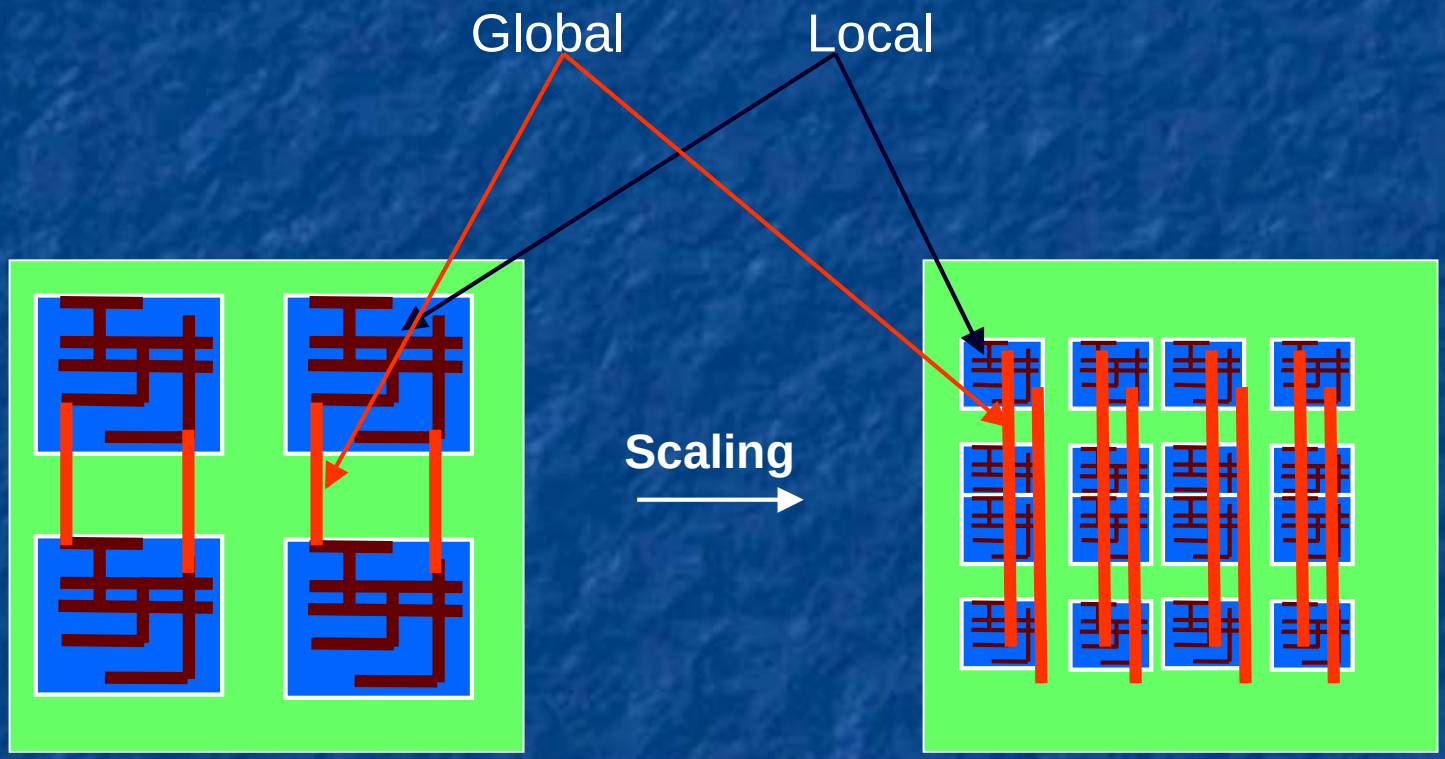

Interconnect scaling

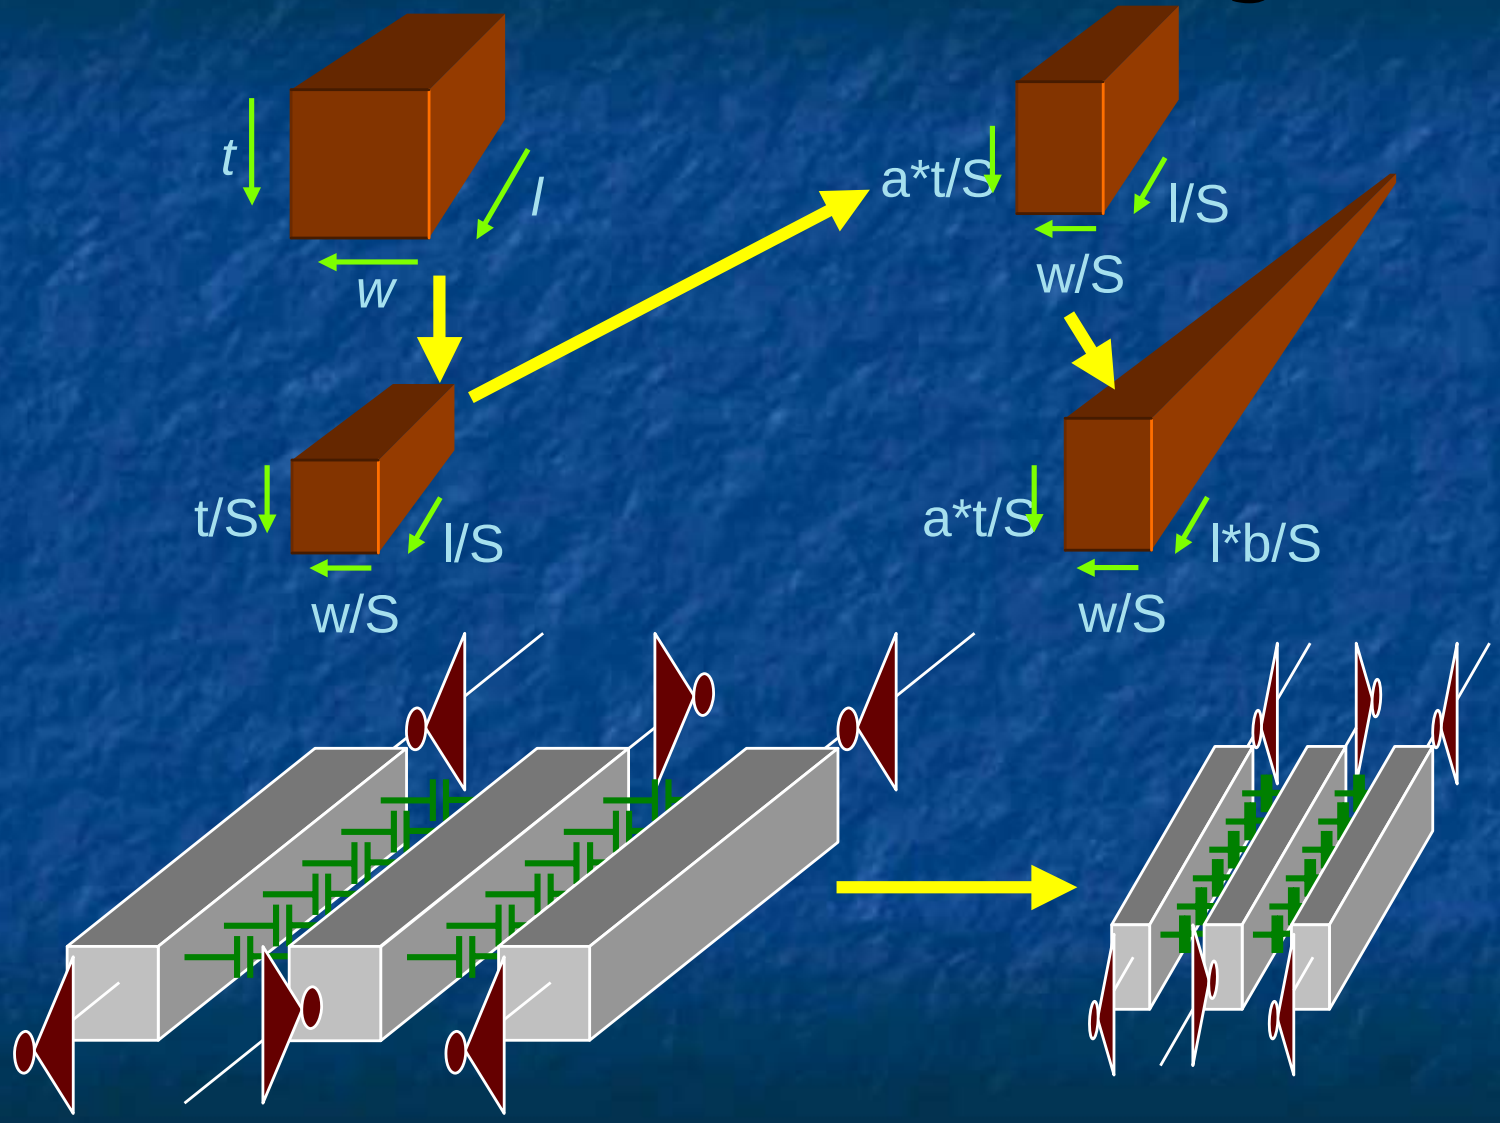

Resistivity scaling

- Ideally goes up with the square of the scaling (actually worse)
- Wires are harder to drive: need more repeaters
  - More repeaters per generation
    - inter-repeater length scales faster than gates ($S\sqrt{S}$) due to RC
    - number of repeaters increases greater than quadratically
    - increasing number of clocked repeaters due to loss of cycle time
    - More planing up front to repeater farm mesh structures
  - **shorter wires**
  - **more via stacks**
- Only exacerbates Power delivery problem already described
  - Wider rails
  - Higher density of C4 bumps and therefore wide upper metals
  - Rails more frequent generally

Inter-repeater scaling

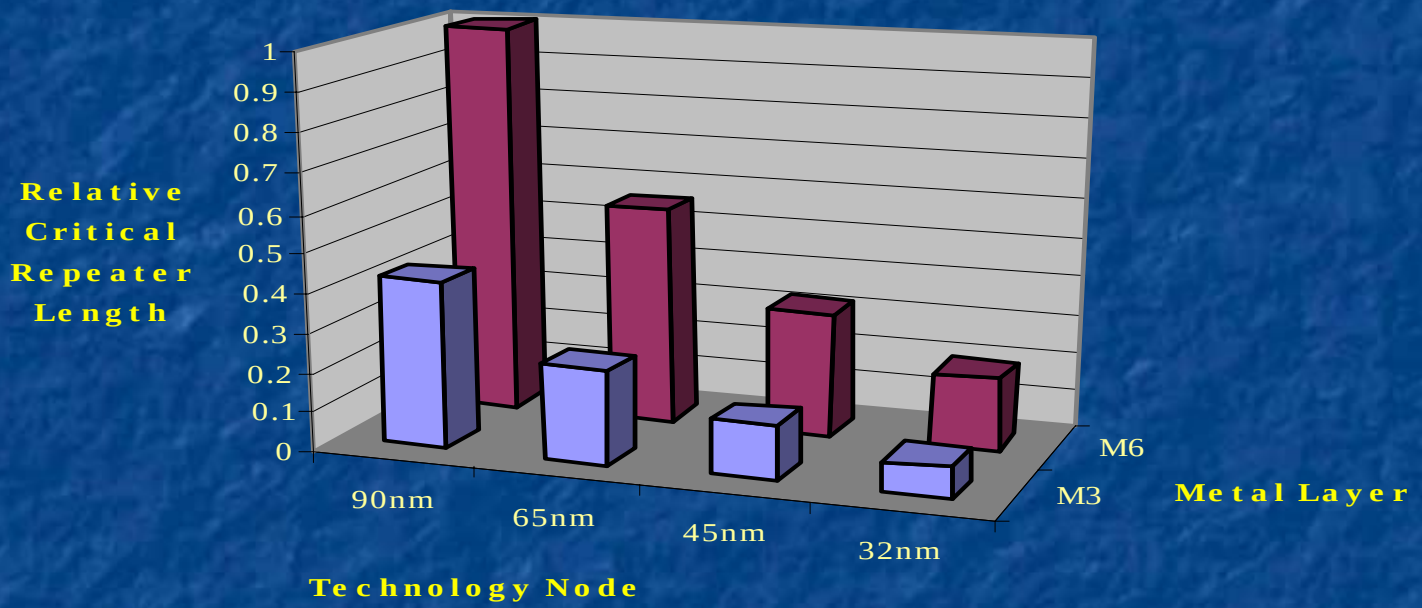

	90nm	65nm	45nm	32nm
M3	0.43	0.24	0.14	0.08
M6	1	0.56	0.32	0.19

P. Saxena/Intel labs

Repeater bays pre-planned

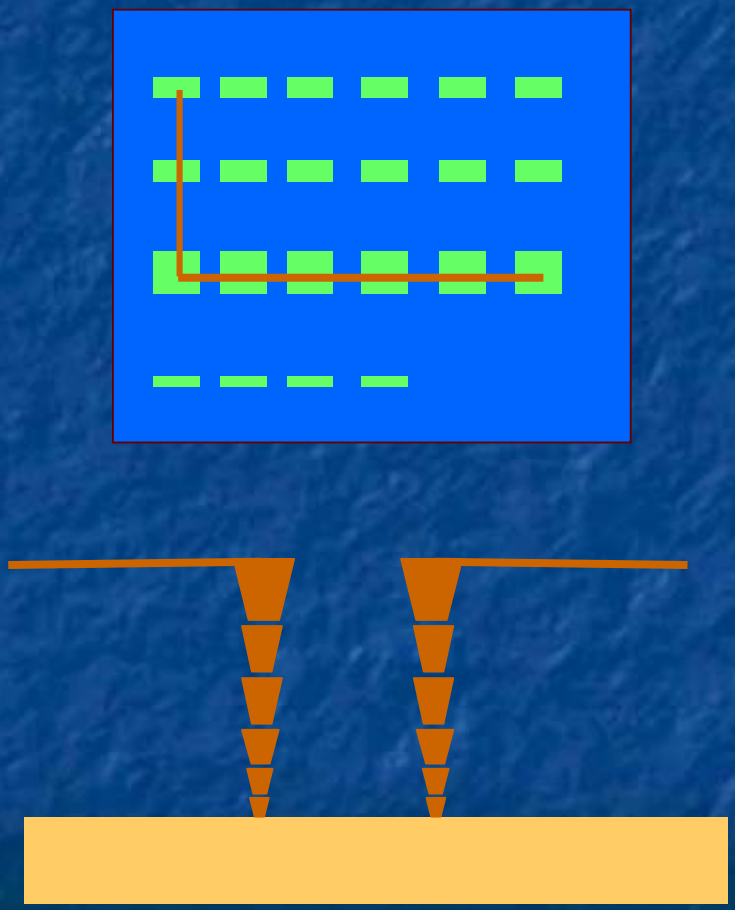

- Increasingly global pre-placement will be necessary
- This will be based on knowledge of the technology layers and what the realistic expected inter-repeater distances should be
- Via stacks will be uniformly known in advance and need to be accounted for in pre-routing estimates

Noise trends

- Switching speeds going up
- Cap coupling going up
- Inductive coupling because of synchronicity
- As process scales it becomes more susceptible to power droop impact on delay: 1%/1% -> 1%/3%

Trend of interconnect noise

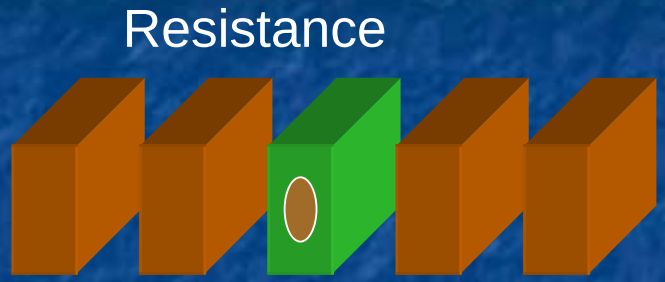

- Resistance
  - Very local
  - Only varies at high F

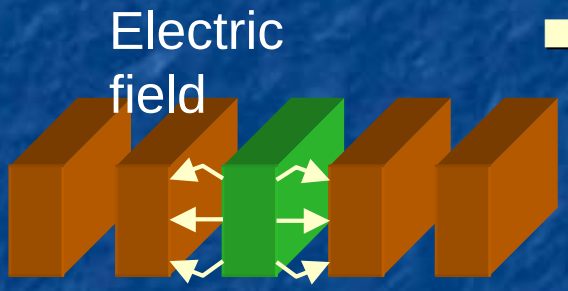

- Capacitance:
  - Electric field coupling
  - Very small and well defined interaction zone

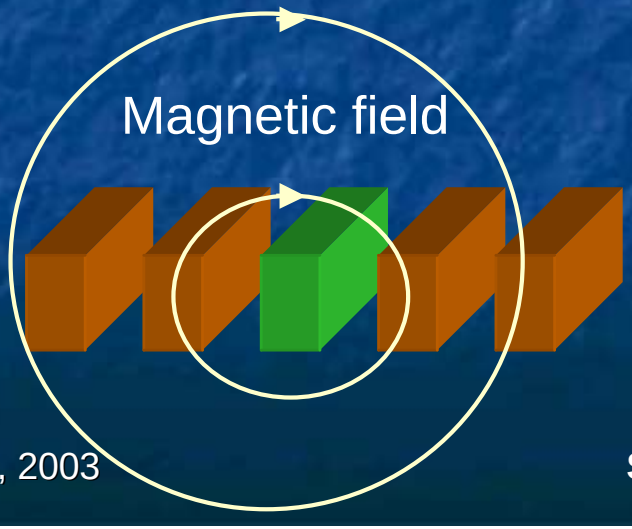

- Inductance:
  - Magnetic field coupling
  - global interaction zone

Opposed to digital design

CL Noise Due to Coupling

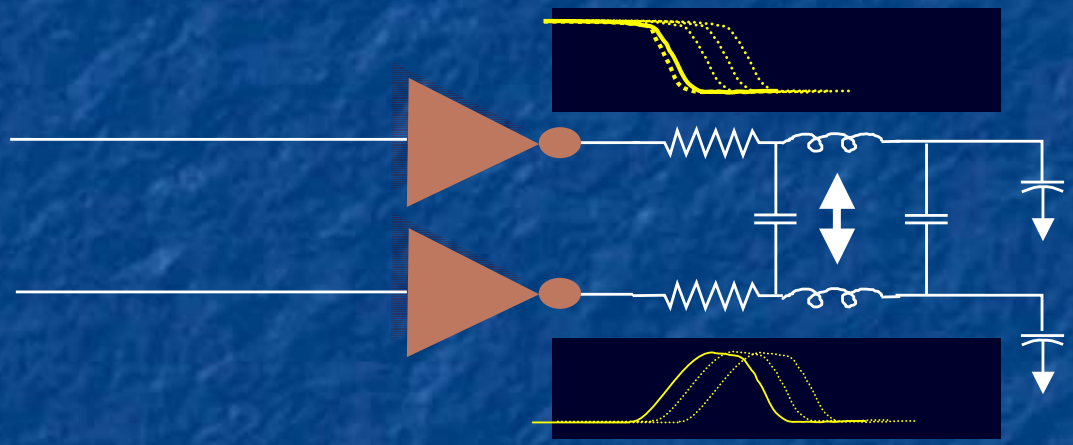

- More coupling means more noise
  - Moves design away from being digital
  - Coupling can be capacitive (easy to rectify during design and more probable) or inductive (more serious, less probable, difficult during design so plan for up front)

Timing Variation Due to CL Coupling

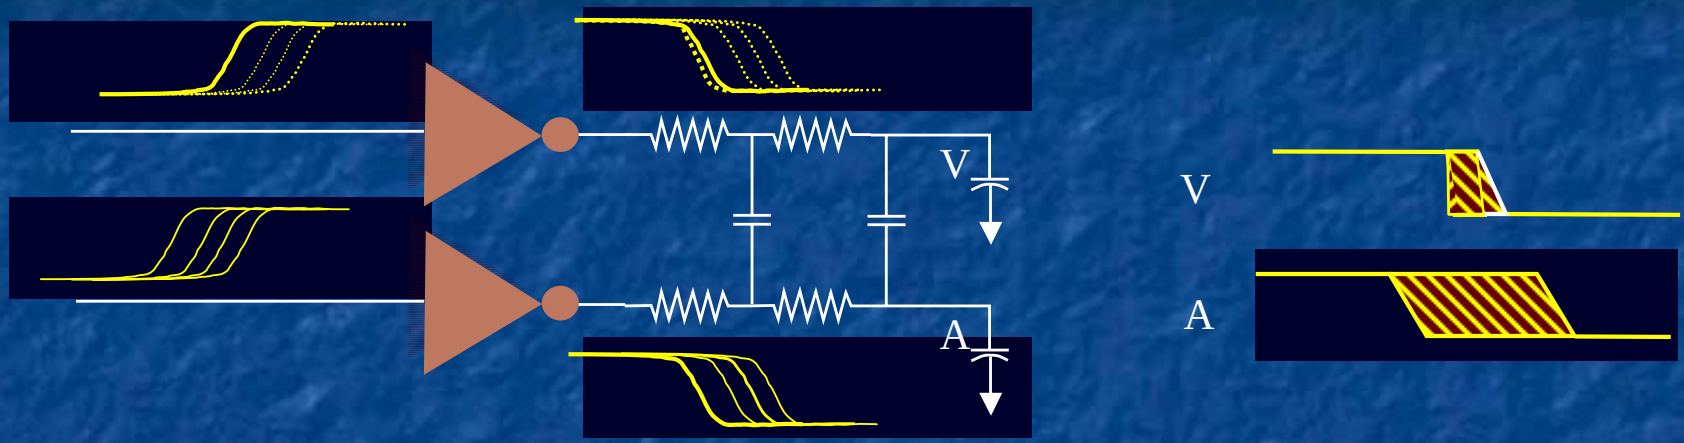

- Coupling increases delay variation (noise-on-timing)
  - Environmental variation
    - Interconnect coupling is deterministic
    - Patterns are non-deterministic generally
    - Difficult to introduce correctly into static timing flows
    - Inductive noise on timing may be non-negligible

How On-chip Inductive Noise Occurs

Induced current can itself induce current that reverses effects

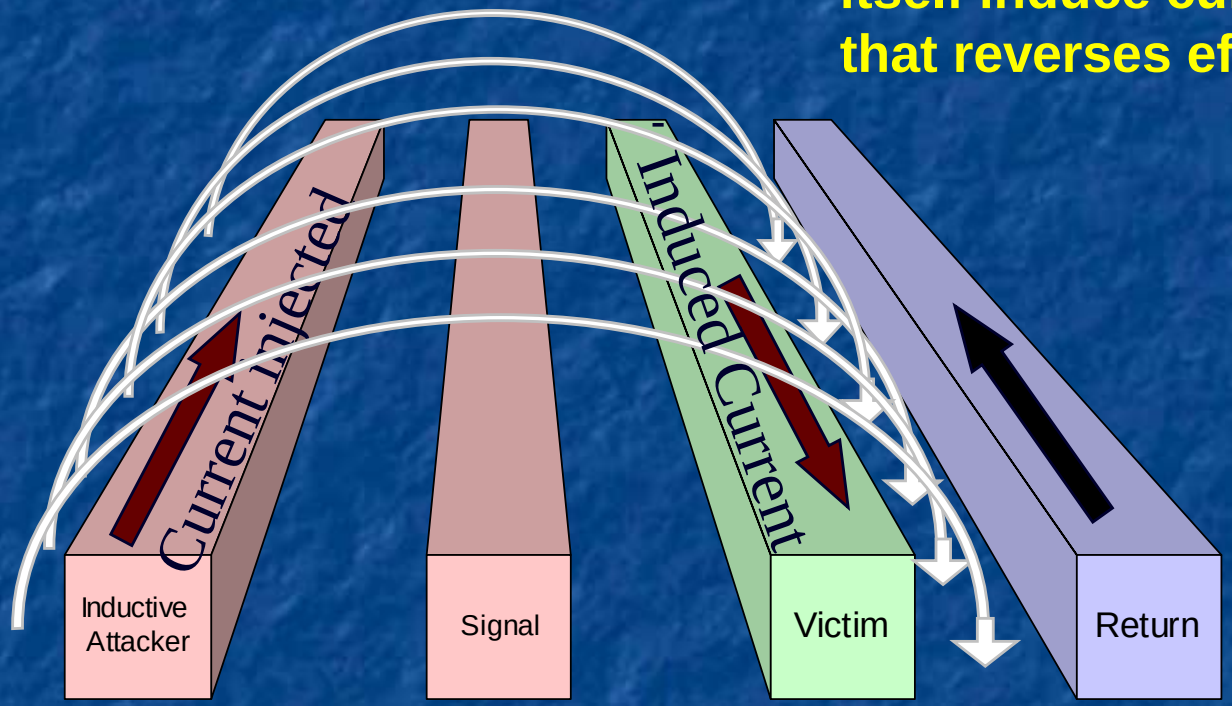

Magnetic fields can act for longer distances and die off logarithmically compared to electrical fields

RC/RLC noise difference

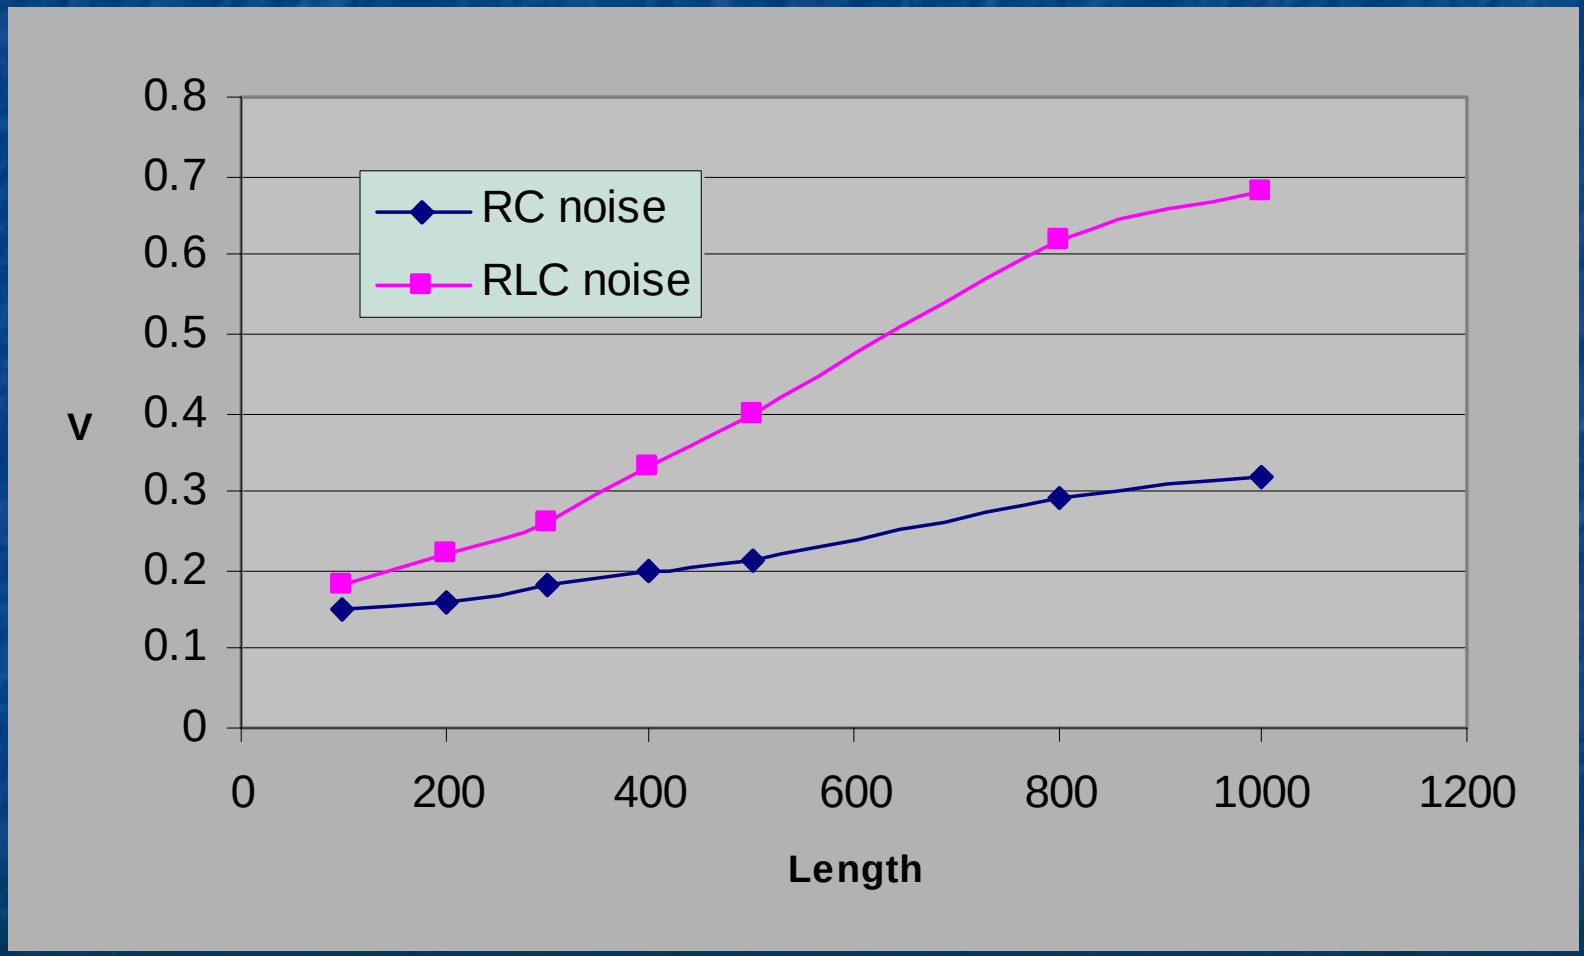

Intel labs

On-chip interaction: complex attack

Individual attacking noise effects

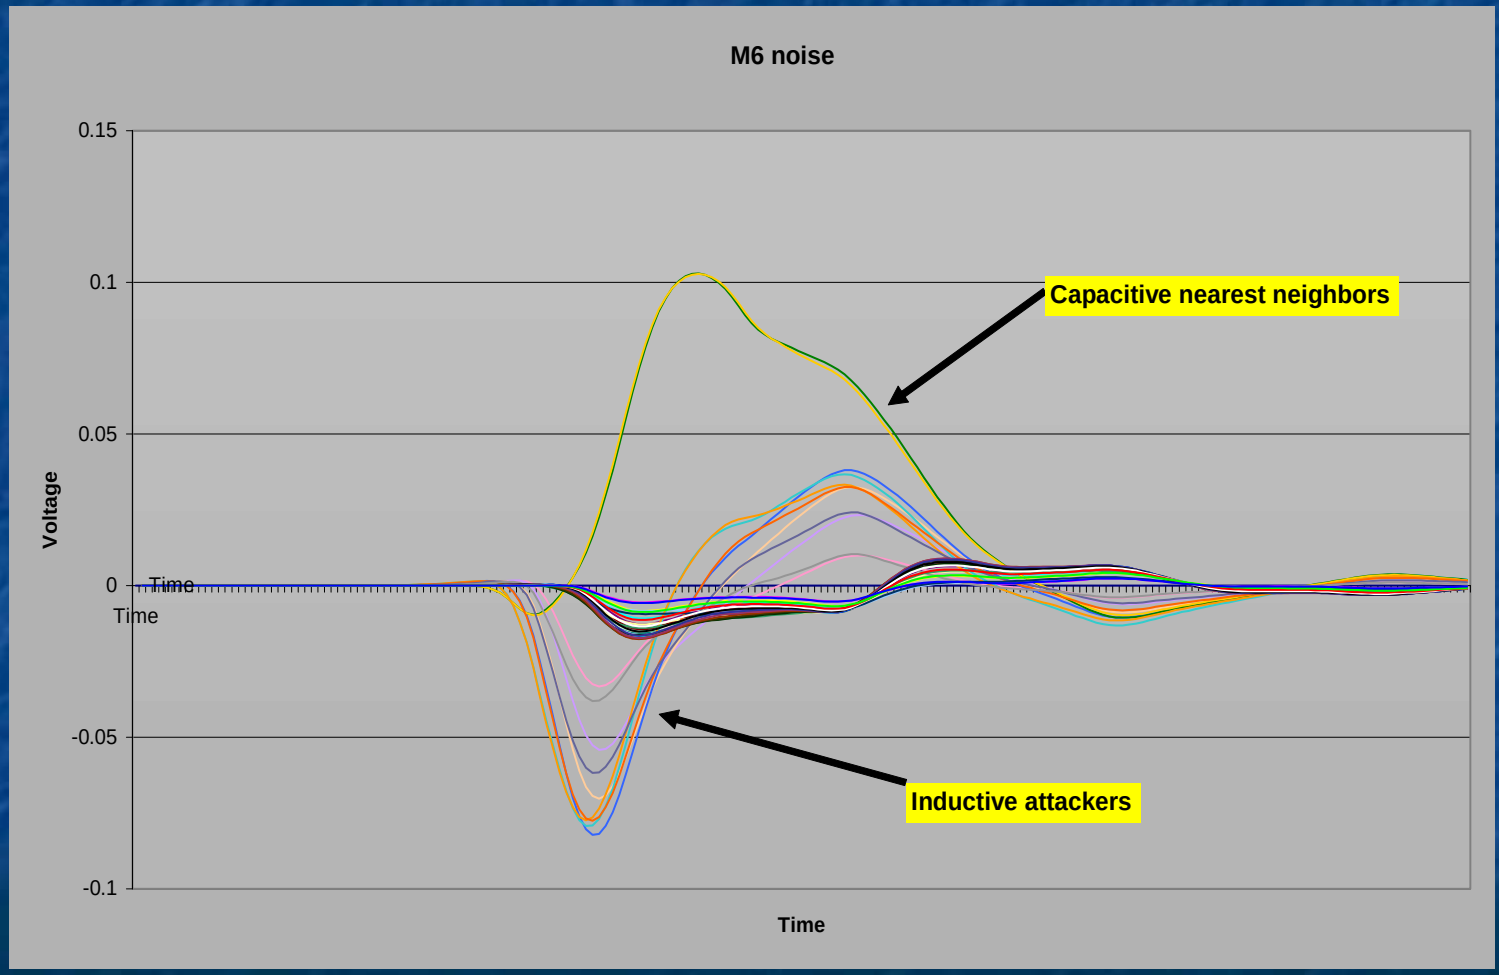

Timing

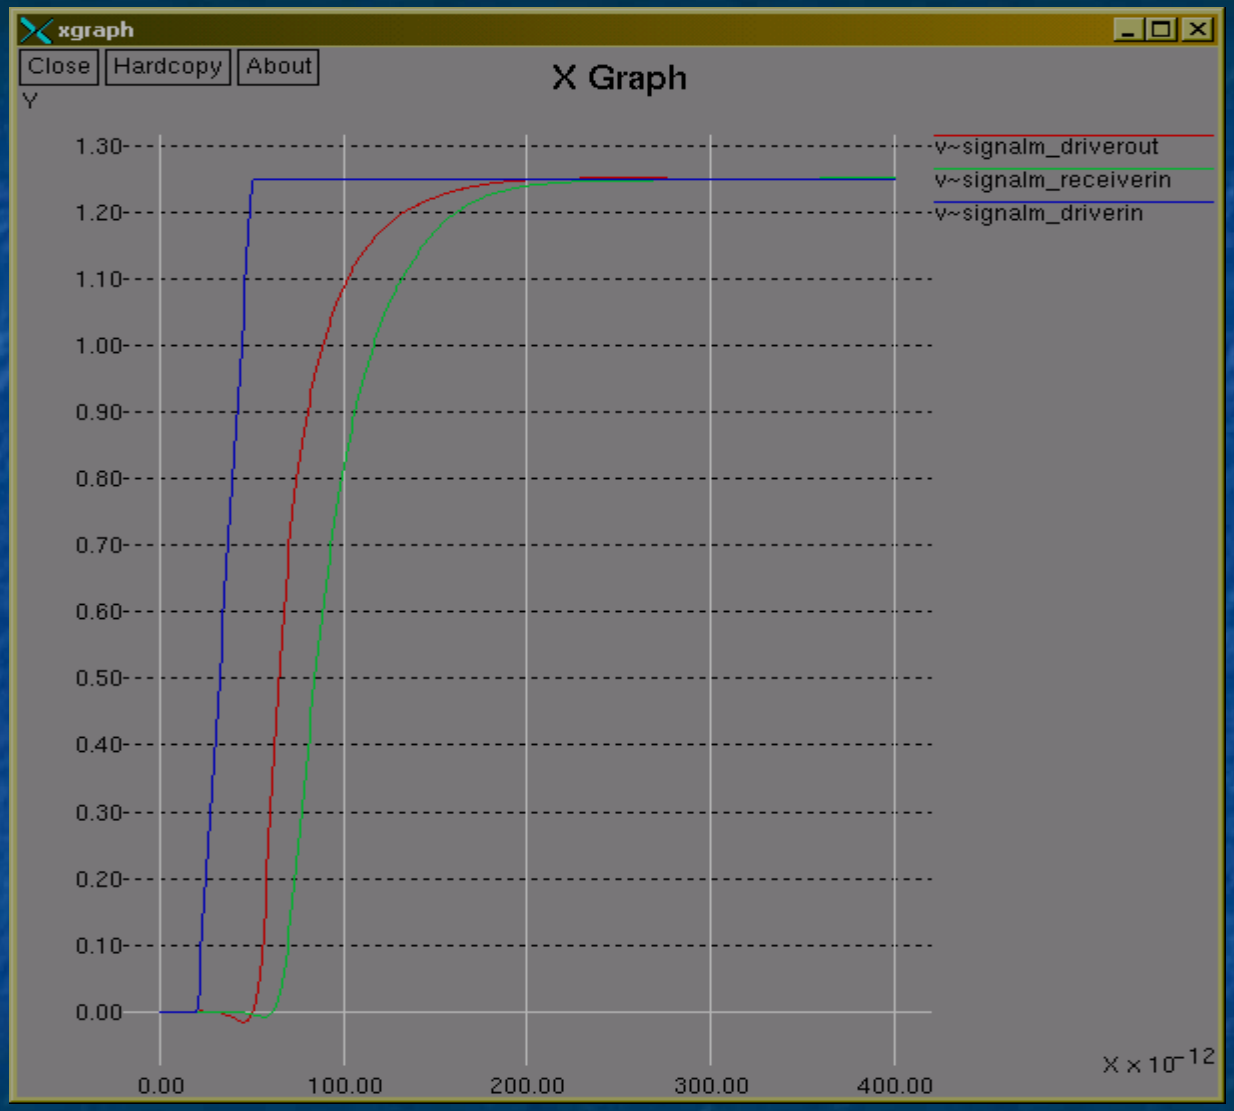

1000u

25 signals

Victim up

NO RLC noise

Intel labs

Timing

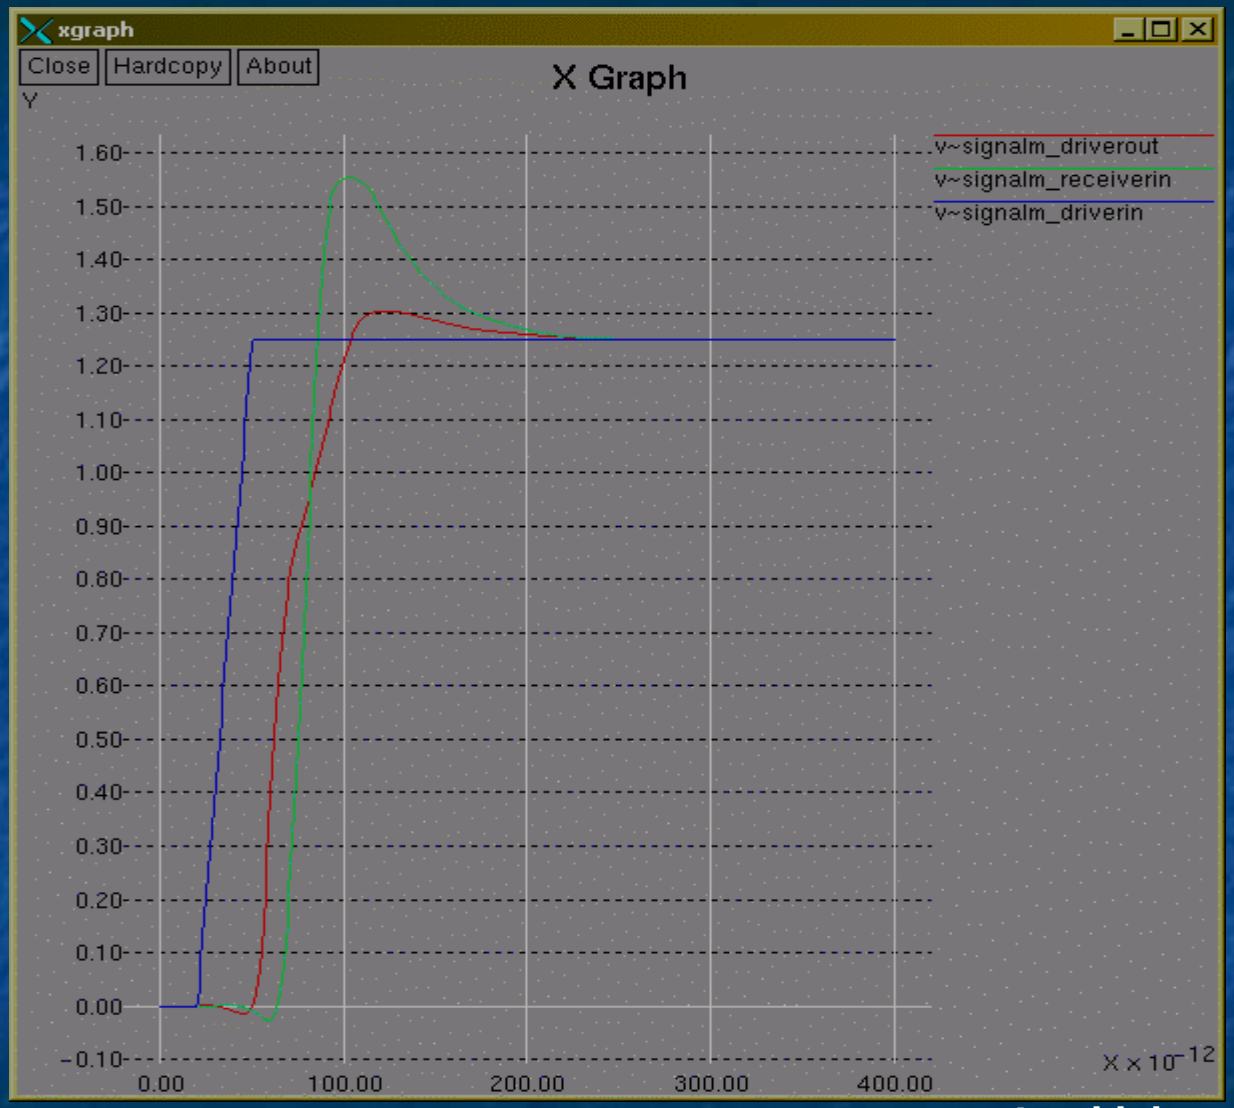

1000u

25 signals

Victim up

WC RLC noise (+)

Intel labs

Timing

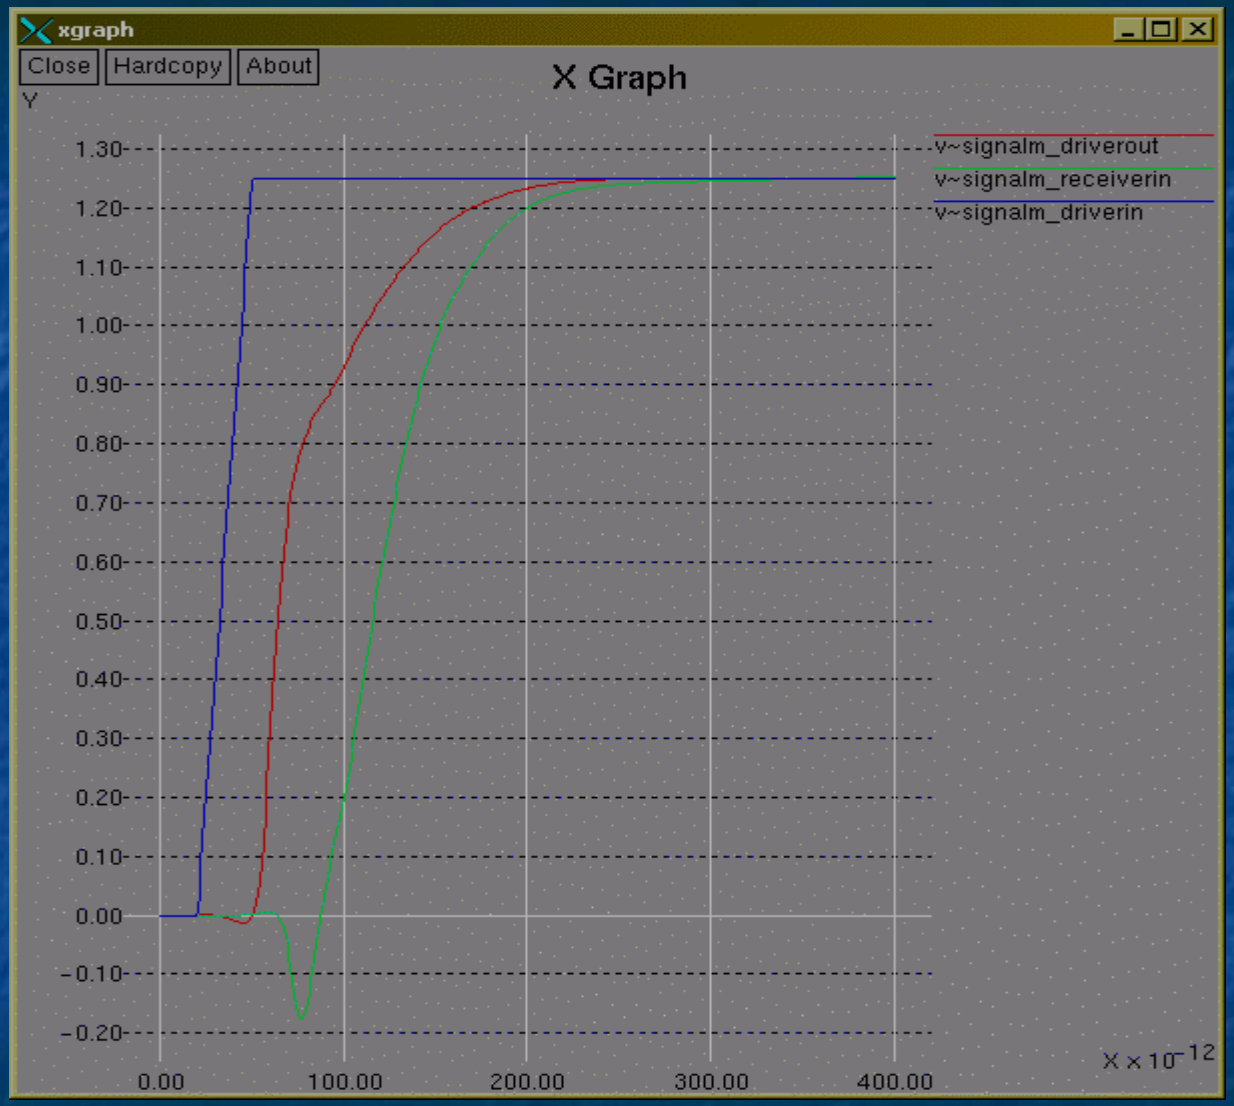

1000u

25 signals

Victim up

WC RLC noise (-)

Window of influence

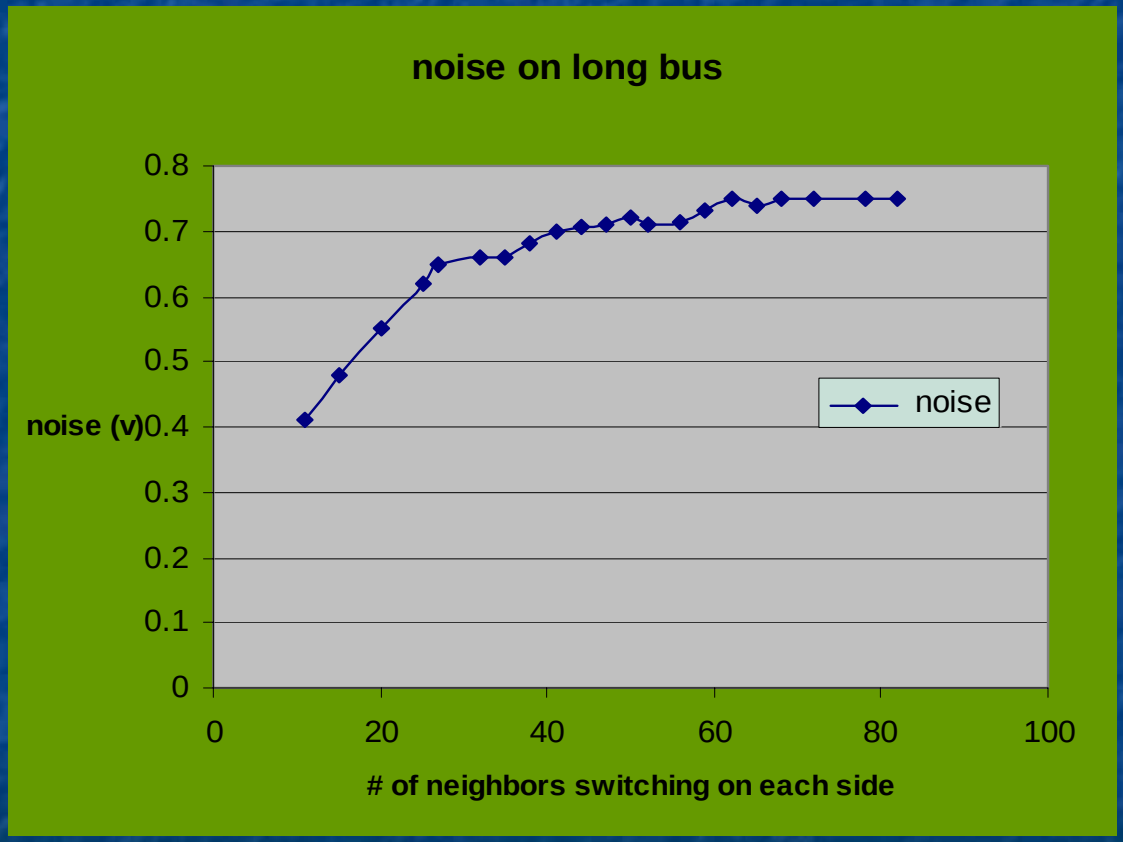

Intel labs

But nets look like this…

…not just this…

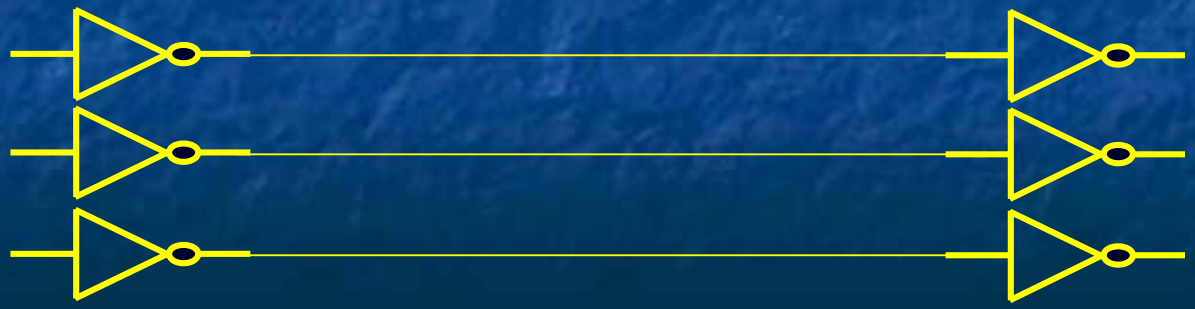

Multiple attackers

- Worst case scenario is pretty bad!
- If we used this, design could not be done
- Probability of worst case is almost zero
- **This choice of probability window can be the source of inaccuracy greater than inductive modeling!**
- Reasonable noise window and shielding returns must be chosen up front
  - the more returns the better!

Impact of noise on wiring

- Need a robust (low) signal to rail ratio
- Routability constraints nonlinear function of this ratio
- Need more and staggered repeaters to disrupt the noise

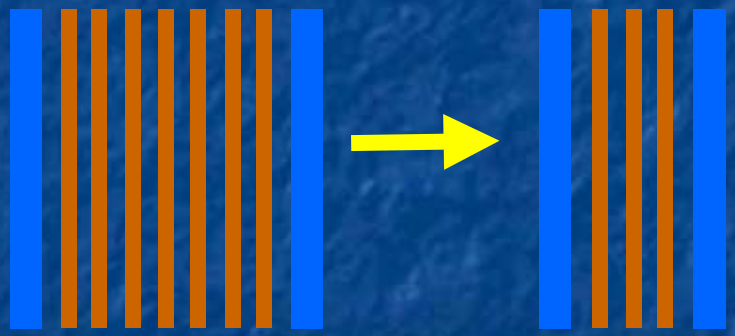

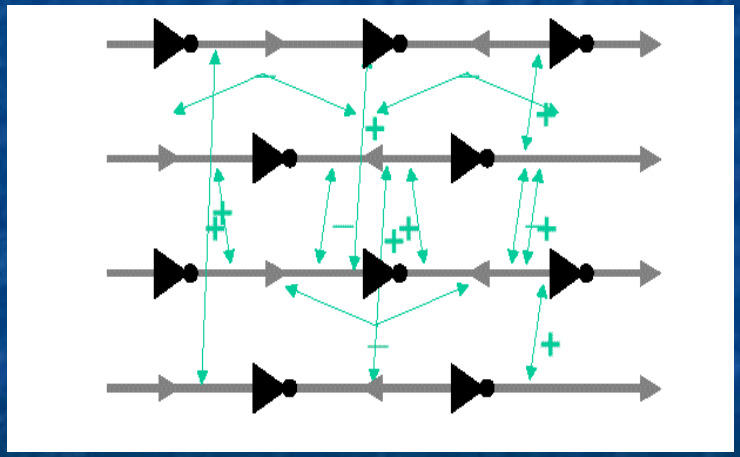

Typical early wiring studies

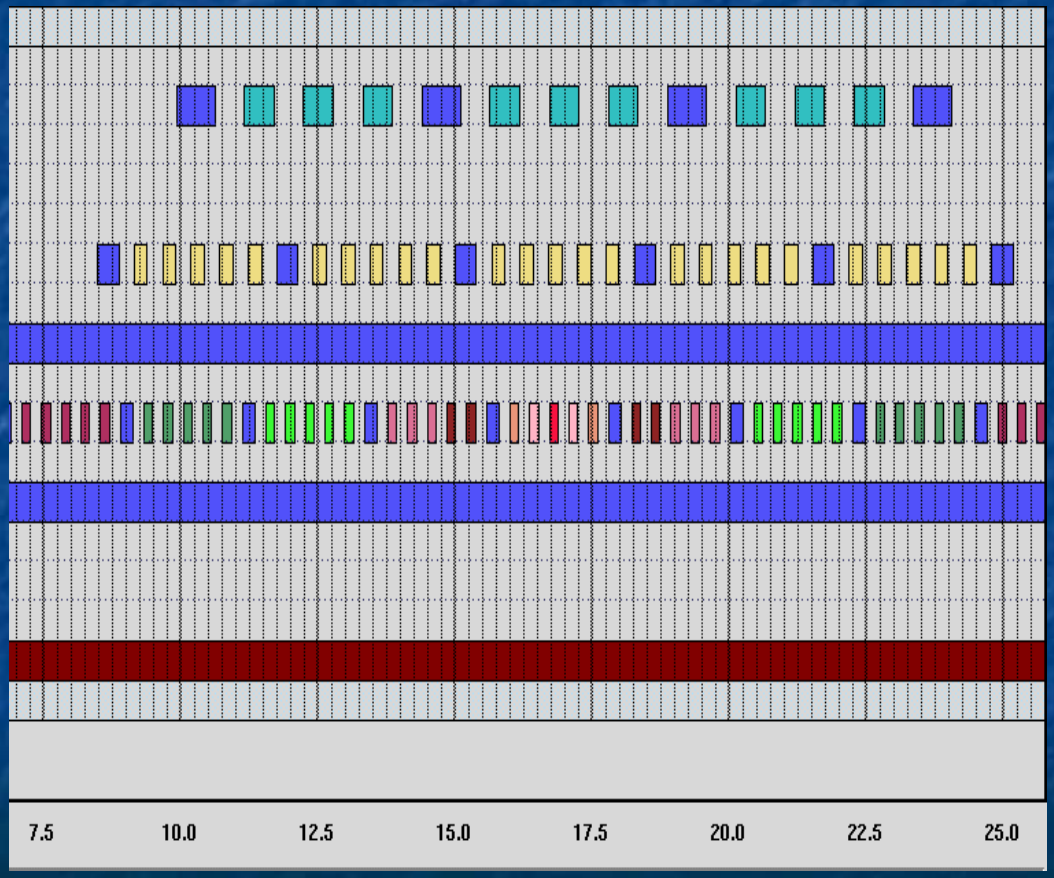

- 2D only
- freq. range
- inter-repeater length
- drivers/receivers on each layer
- simultaneous f-dept. R/L/C models
- Optimization in this domain only considers IR drop, C noise, L noise, repeater distance
- Decisions fix design

Global noise studies with variable power grid signal/rail ratio

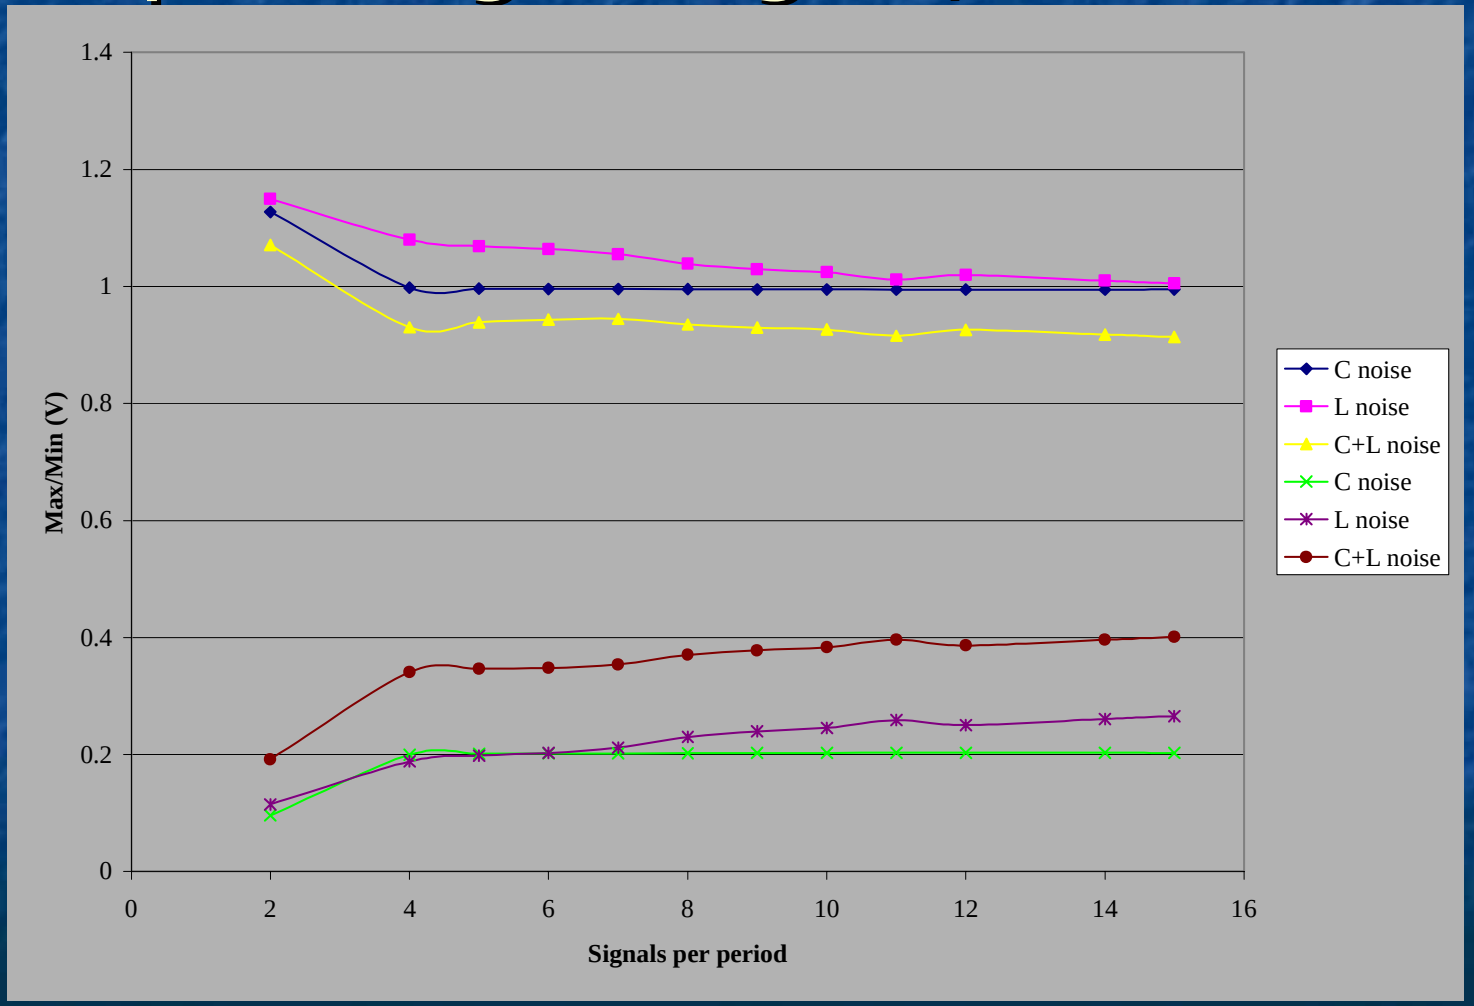

Global delay studies with variable power grid signal/rail ratio

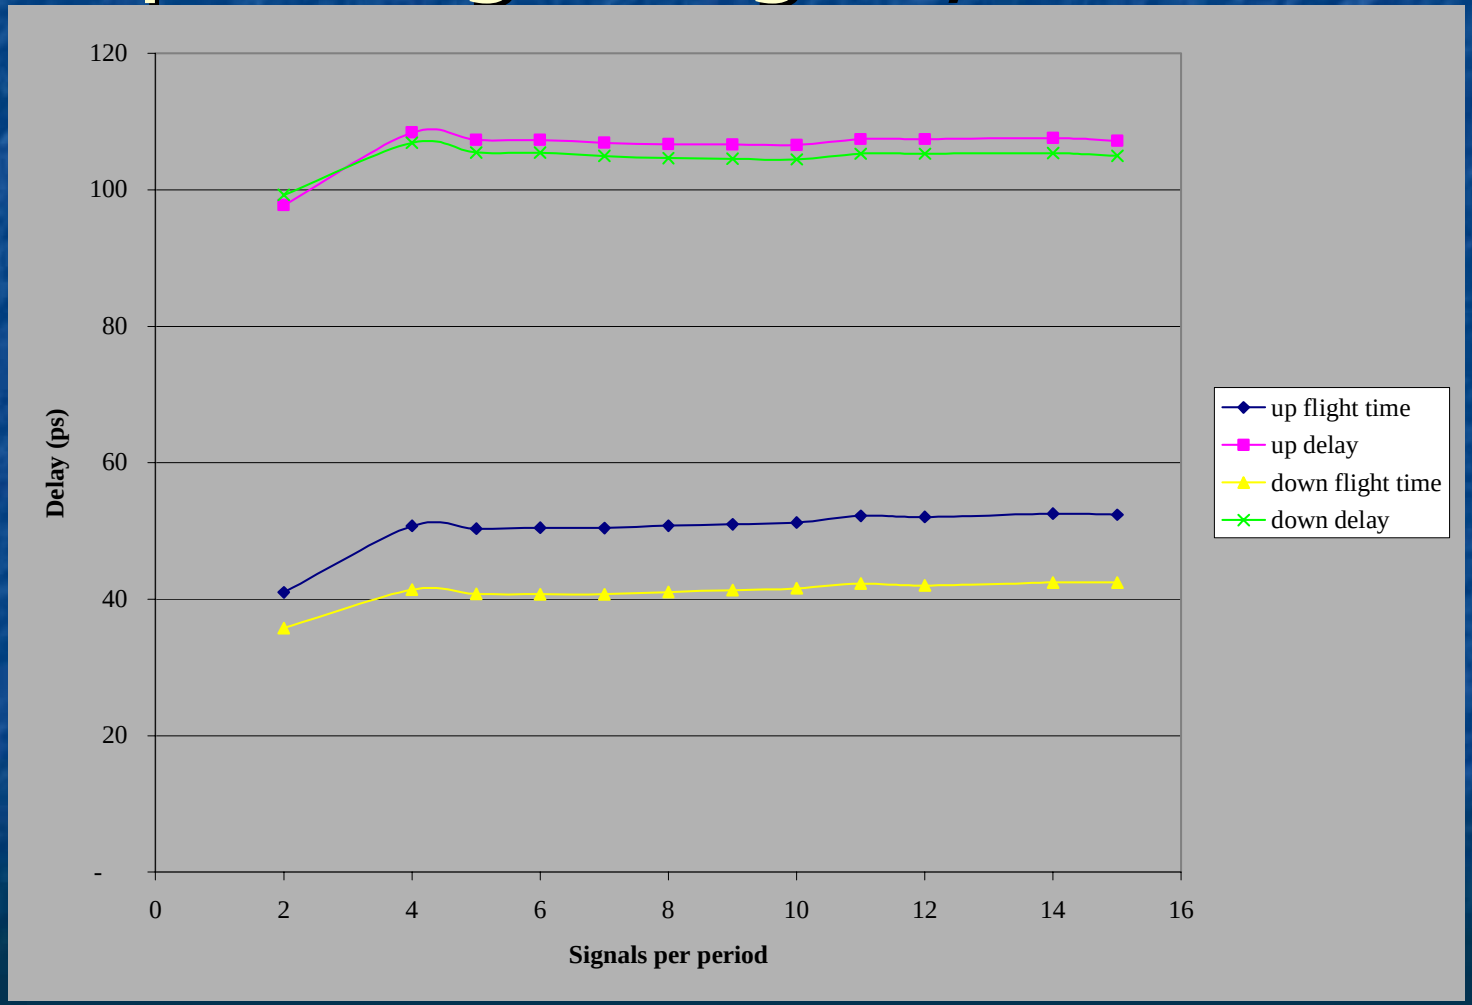

Global noise studies with fixed power grid and varying driver size

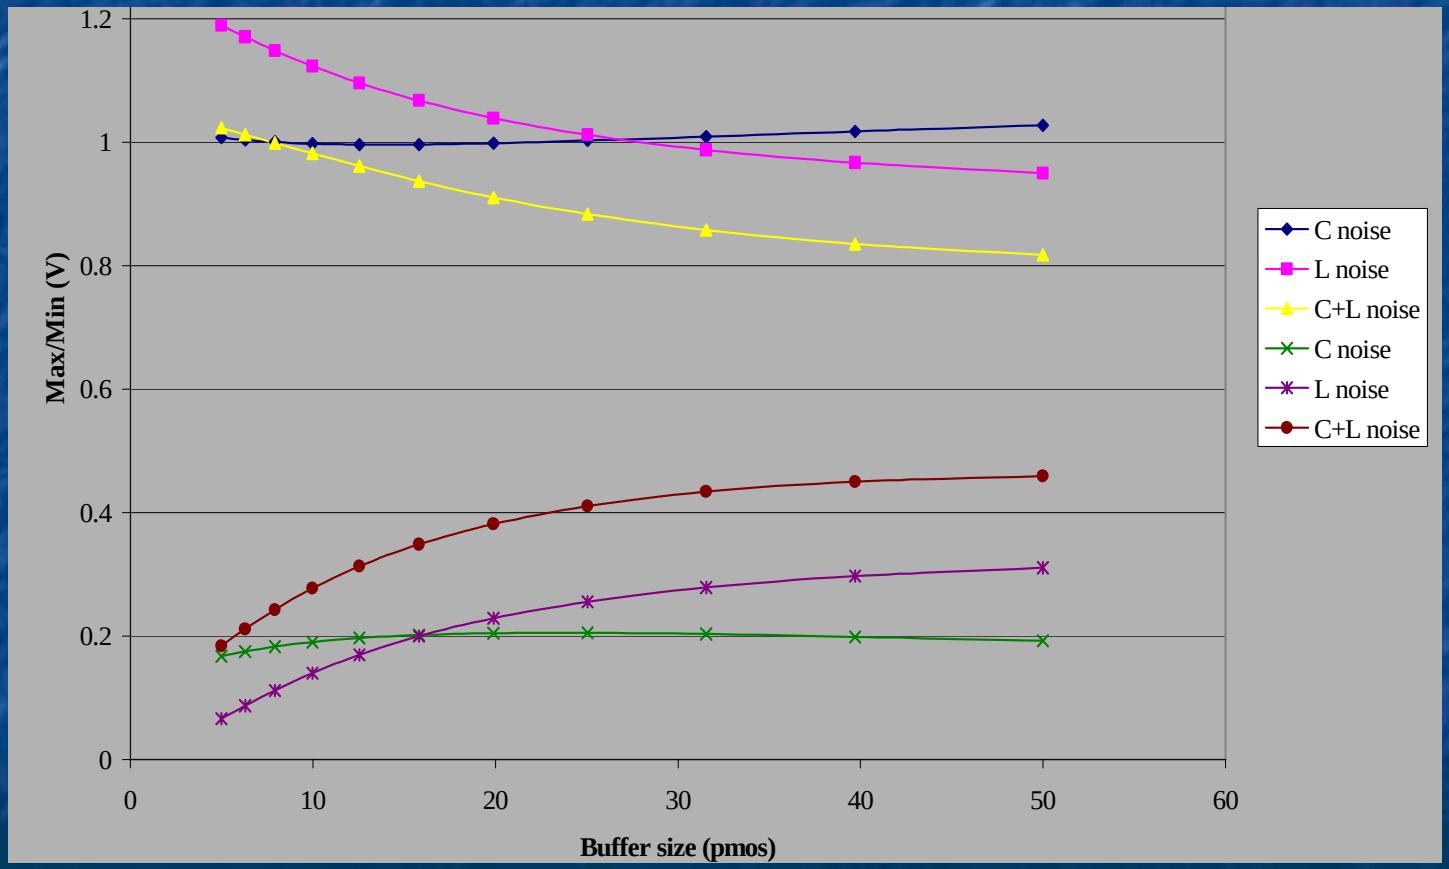

Global delay studies with fixed power grid and varying driver size

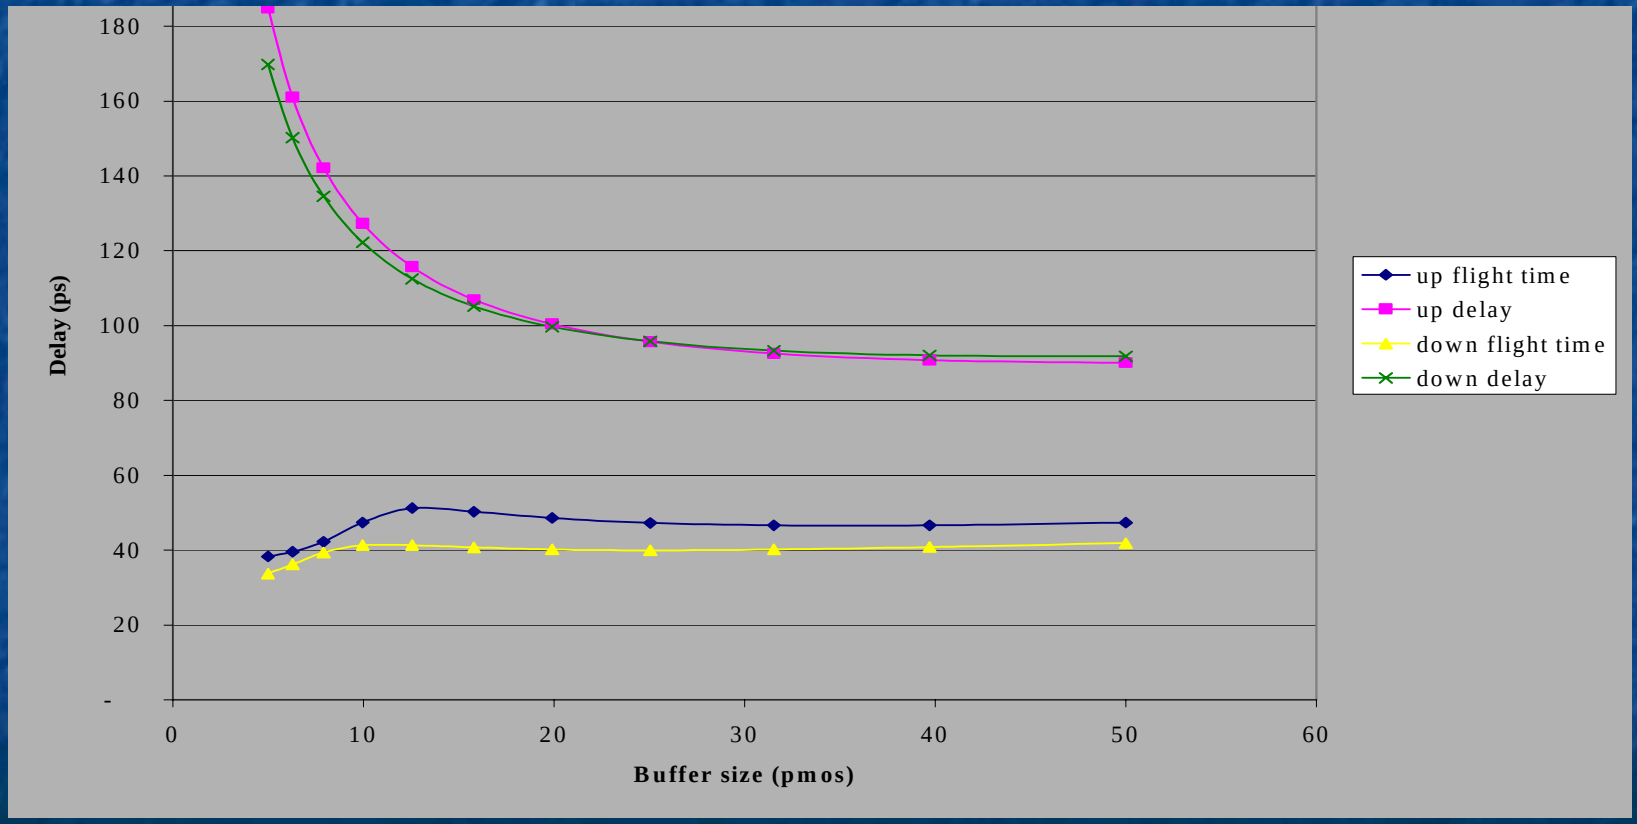

Global noise studies with fixed power grid and varying line length

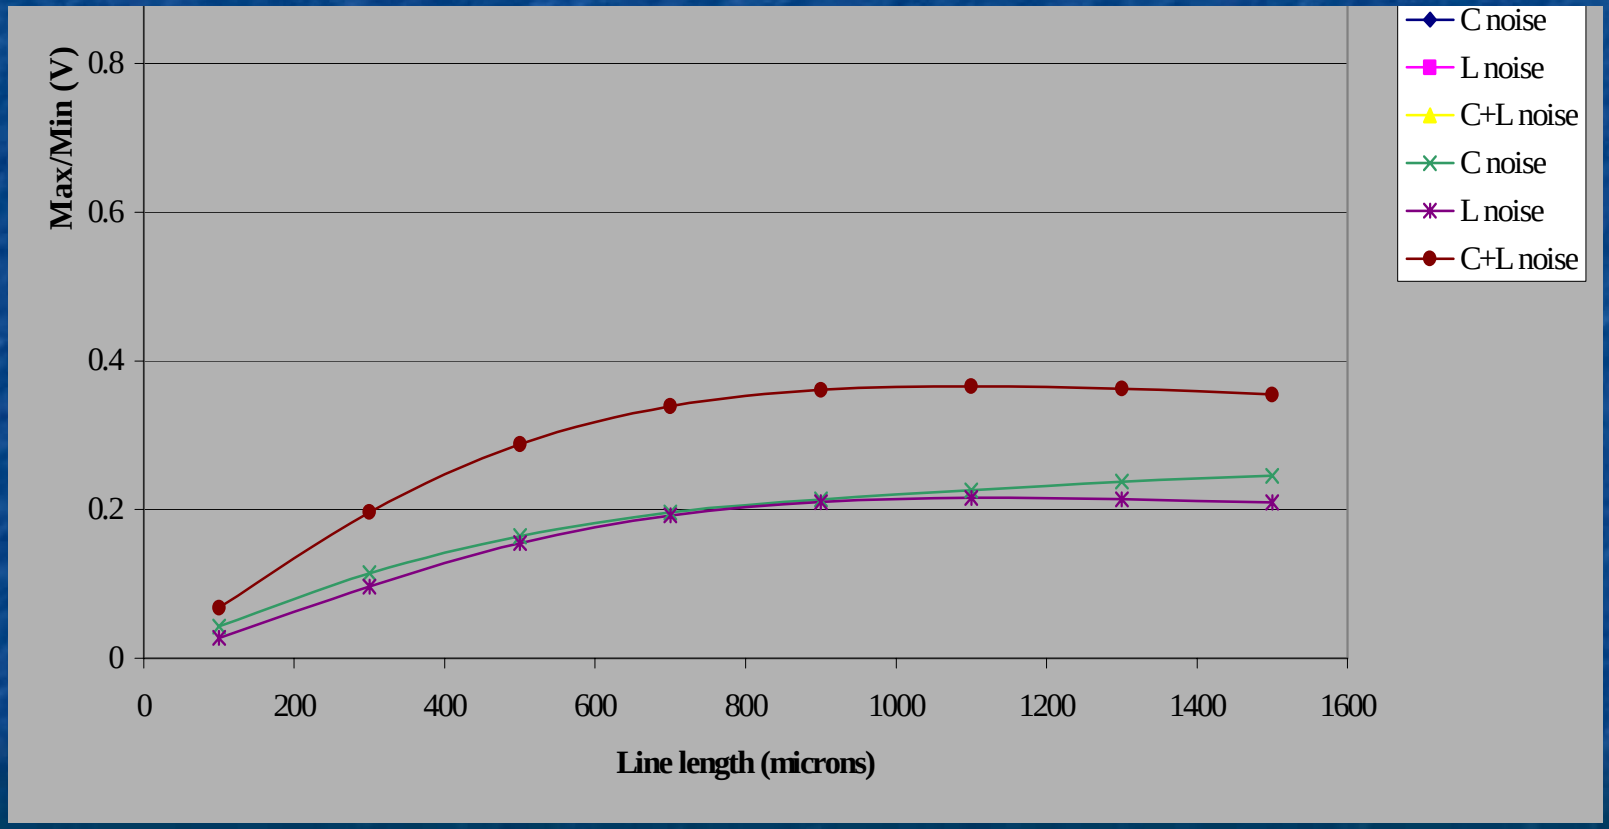

Global delay studies with fixed power grid, varying length, wc noise

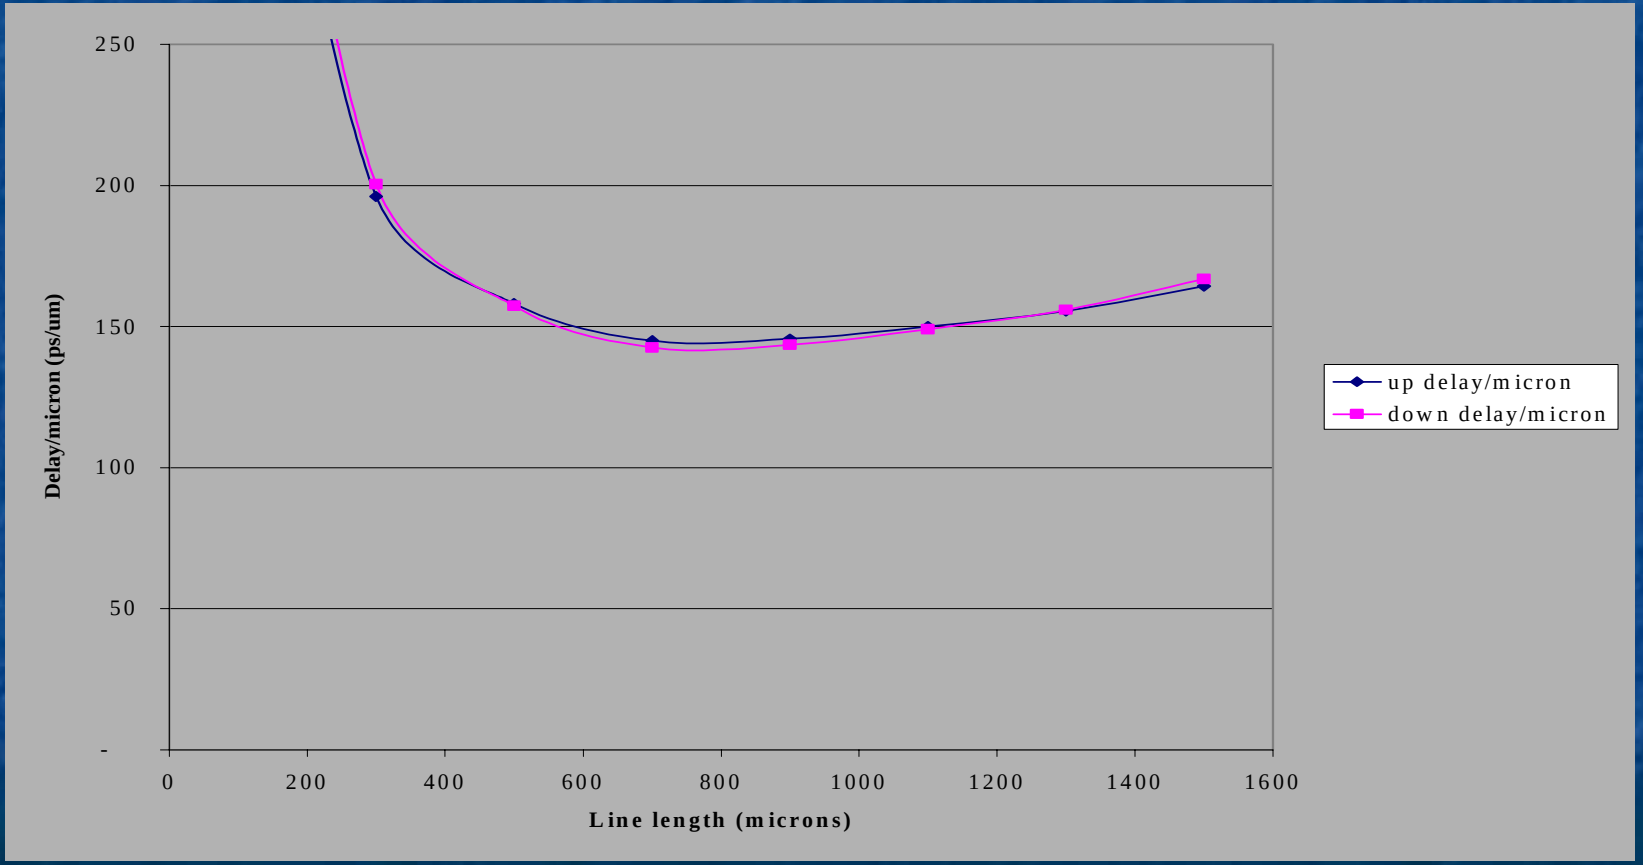

Decoupling capacitance

- Increasing dI/dt causing more need for de-cap pre-placement
- Decap planning needs to take place early and generically even before there are placeable blocks
- what is the impact of these pre-placed decap farms?
- Probably not much since contact from Vcc to Vss can take place at any layer without via stacks

Clock issues

- Clock distribution also fixed early
  - Process variation and environmental variation fix clock tracks and shields early
  - Grid increasingly used due to variability
  - Not covered in this talk
- Between clock and power most upper tracks are heavily defined

Most electrical wiring relationships fixed by the time we get to placement

- Not much consideration of physical wiring constraints
- In the past this was not a problem
  - Now, wiring resources not optimized will be consumed
- Later feedback "solutions" from routing are strange and difficult to accomplish

What does this mean for wires?

- There used to be maneuvering room for not thinking about wires until placement/routing
- Increasingly, power grid must satisfy electrical constraints early on: current, EM and noise
- Clock grid also defined early (variability, skew, inductance)
- Repeater insertion gets pre-defined (fixed bays, via stacks)
- What is the impact on the wires? No necessarily rent's rule first!
- We need a bigger picture…

Metrics needed

- Maximum power distribution wiring area allowed
- Maximum clock distribution wiring area
  - Per layer analysis and tradeoff for each
- Global noise verses routability tradeoff
- Repeater farms per area vs. via stack blockage
- Generic repeater distances verses congestion
- Number of repeater farms necessary with power grid definition and wiring needs
- Etc.

Example for one layer

- $r_{s/p}$ signal to power ratio
- w_p power rail width
- d_r inter-repeater distance

} calculated for noise/delay constraints

- w_s minimum signal width/space
- A_d Area of die
- N_p number of power rails

  $= A_D / [(2r_{s/p} + 1)w_s + w_p]$
- A_p area of power rails = $\sqrt{A_d w_p N_p}$

Example for one layer (con'd)

N_{vs}

- number of via stacks = $\sqrt{A_D}(r_{s/p}N_p)/d_r$
- assumes full wires and no L distribution
- gives number of via obstructions on layer below

n^k_{vs}

- number of via obstructions per wire length on layer k below

$$[N_{vs}]/[(2r^{(k)}_{s/p}+1)N^{(k)}_p\sqrt{A_d}]$$

  - where $N^{(k)}_p$ is calculated based on layer k power grid noise/delay constraints

Conclusion

- Electrical constraints for high-performance designs are increasingly taking over the nature of early wiring constraints
- We need new measures to tie in early electrical planning for noise/delay/power/em/de-cap/repeater numbers and wiring implications for the subsequent design flows
- A simple example given
- Would like to see academic thrust in this area