Reality-Driven Physical Synthesis

Patrick Groeneveld

Chief Technologist, Magma Design Automation, San Jose
(soon: Synopsys inc., Mountain View)

Chair, 49th Design Automation Conference, San Francisco

ISPD 2012, Napa

Kevin Trudeau, the king of Quacks

MAGMA

Physical Design of Apple processors

- Common technology:
  - 45nm Samsung
- **A4**: 2010
  - iPhone 4 & iPad 1
  - 7.3mm x 7.3mm
- **A5**: 2011
  - iPhone 4s & iPad 2
  - 10.0mm x 12.5mm
- **A5x**: 2012
  - iPad 3
  - 12.9mm x 12.7mm
    = 3x as big as the A4

"The perceived usefulness of any new product is an underlying technology's logarithmic function"
(Theo Claasen, 2003)

A closer look at the Apple's physical design style

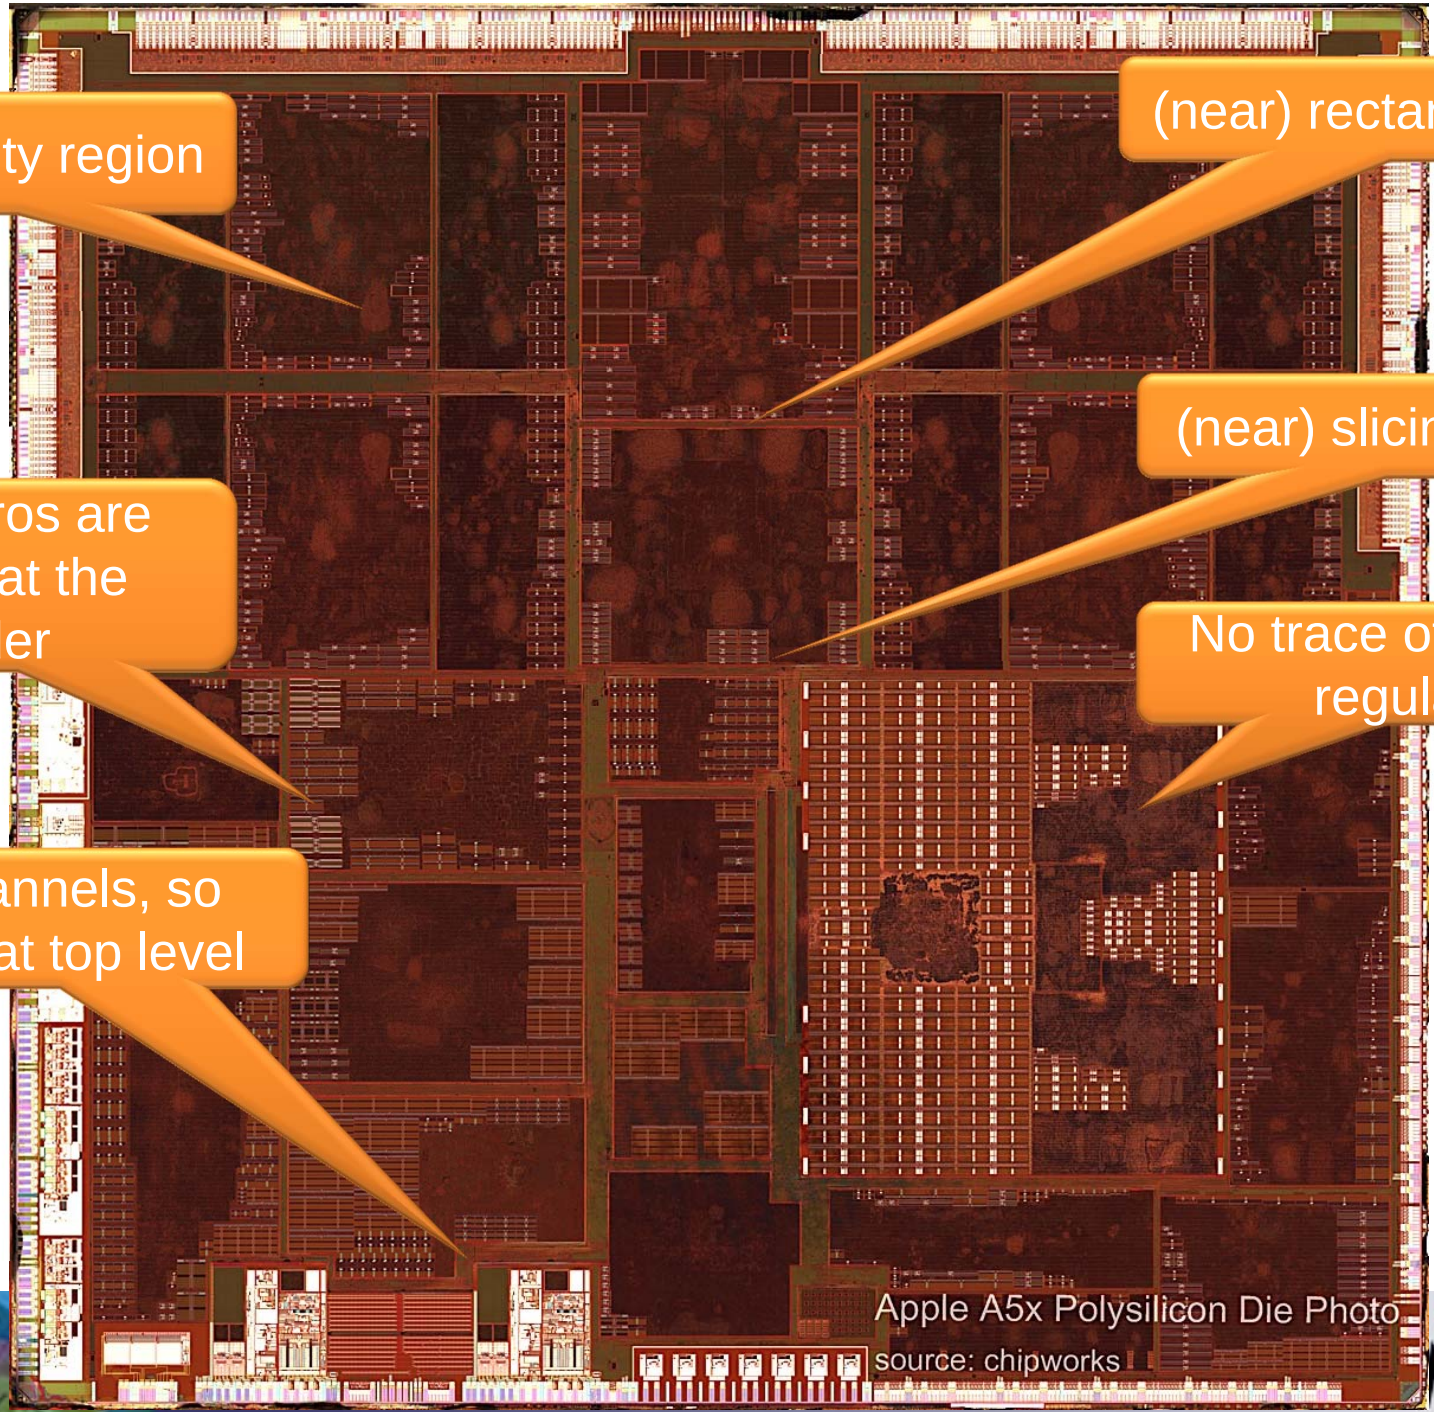

MAGMA

PD: *Many* Objectives Simultaneously

- **Correct** & manufacturable mask pattern
  - Congestion control
  - Big chip = good
- Meets **timing** & electrical requirements
  - Battle parasitics: timing, voltage drop
  - Big gates = good, compact chip = good & a little bad
- Low **power**
  - Leakage control, multi-voltage, sleep, etc
  - Small gates = good, complex floorplan = necessary evil
- Low part **cost**
  - Compact chip, dense wires = good
- Low design **effort**
  - Robust design, short tool run times, re-use
  - Simple = good, pushbutton = good

Intricate trade-off

Magma Flow: guided by ‘best available’ data

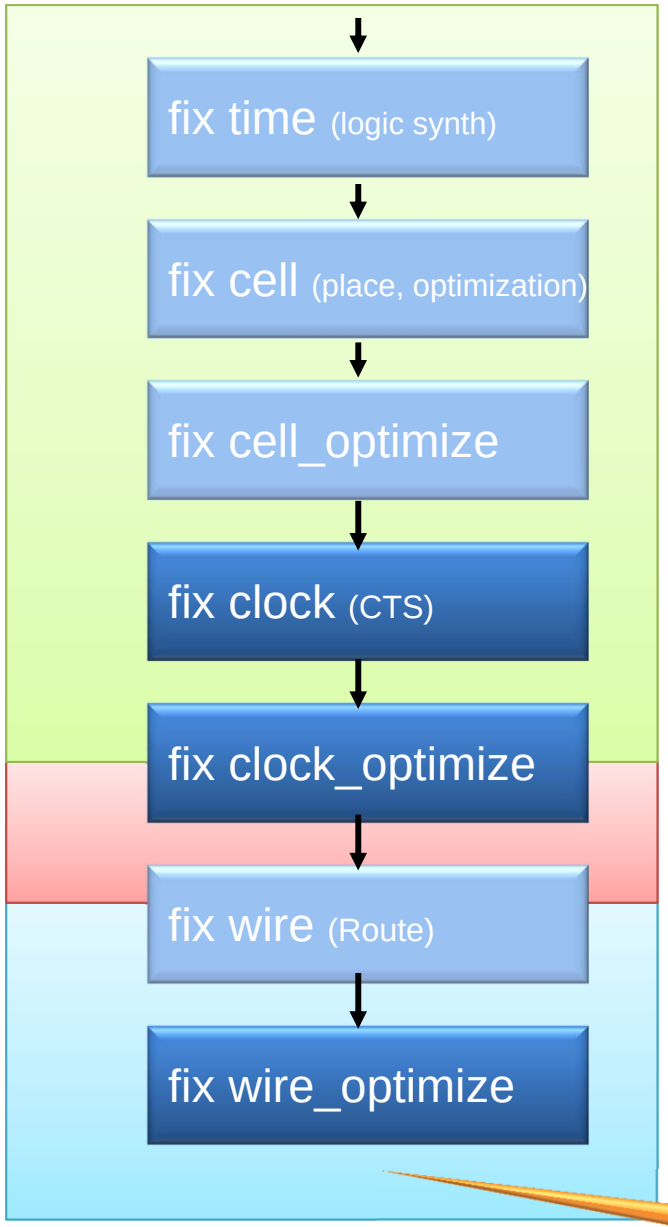

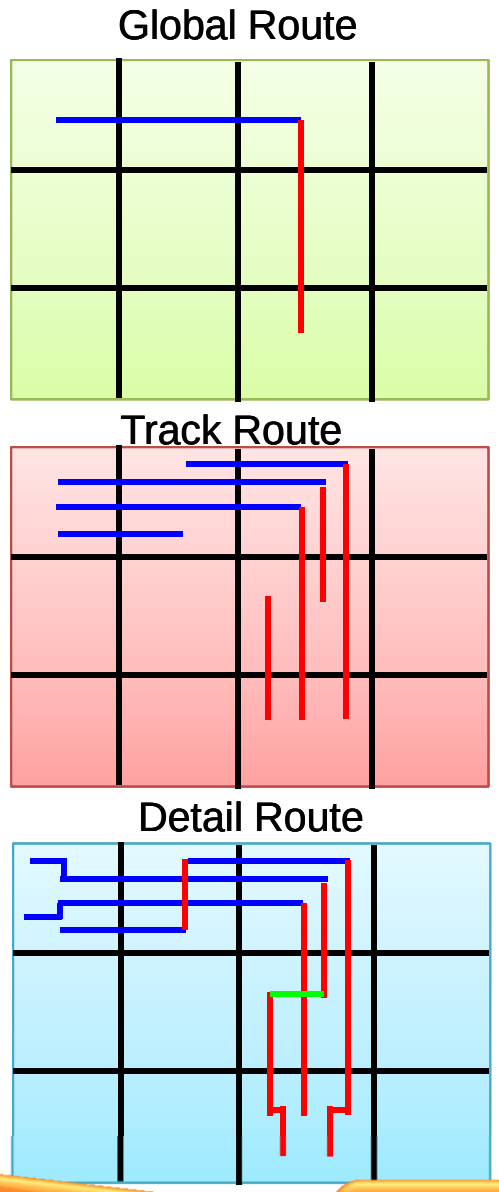

- Global route:
  - Layer assignment
  - Congestion
  - Resource contention
  - Detours
- Track route:
  - Refines global route
- Detail route
  - Copies track route
  - Fixes opens
  - Ripup & Reroute

The only thing that matters is the quality at the end!

MAGMA

Layout Design at different levels of abstraction

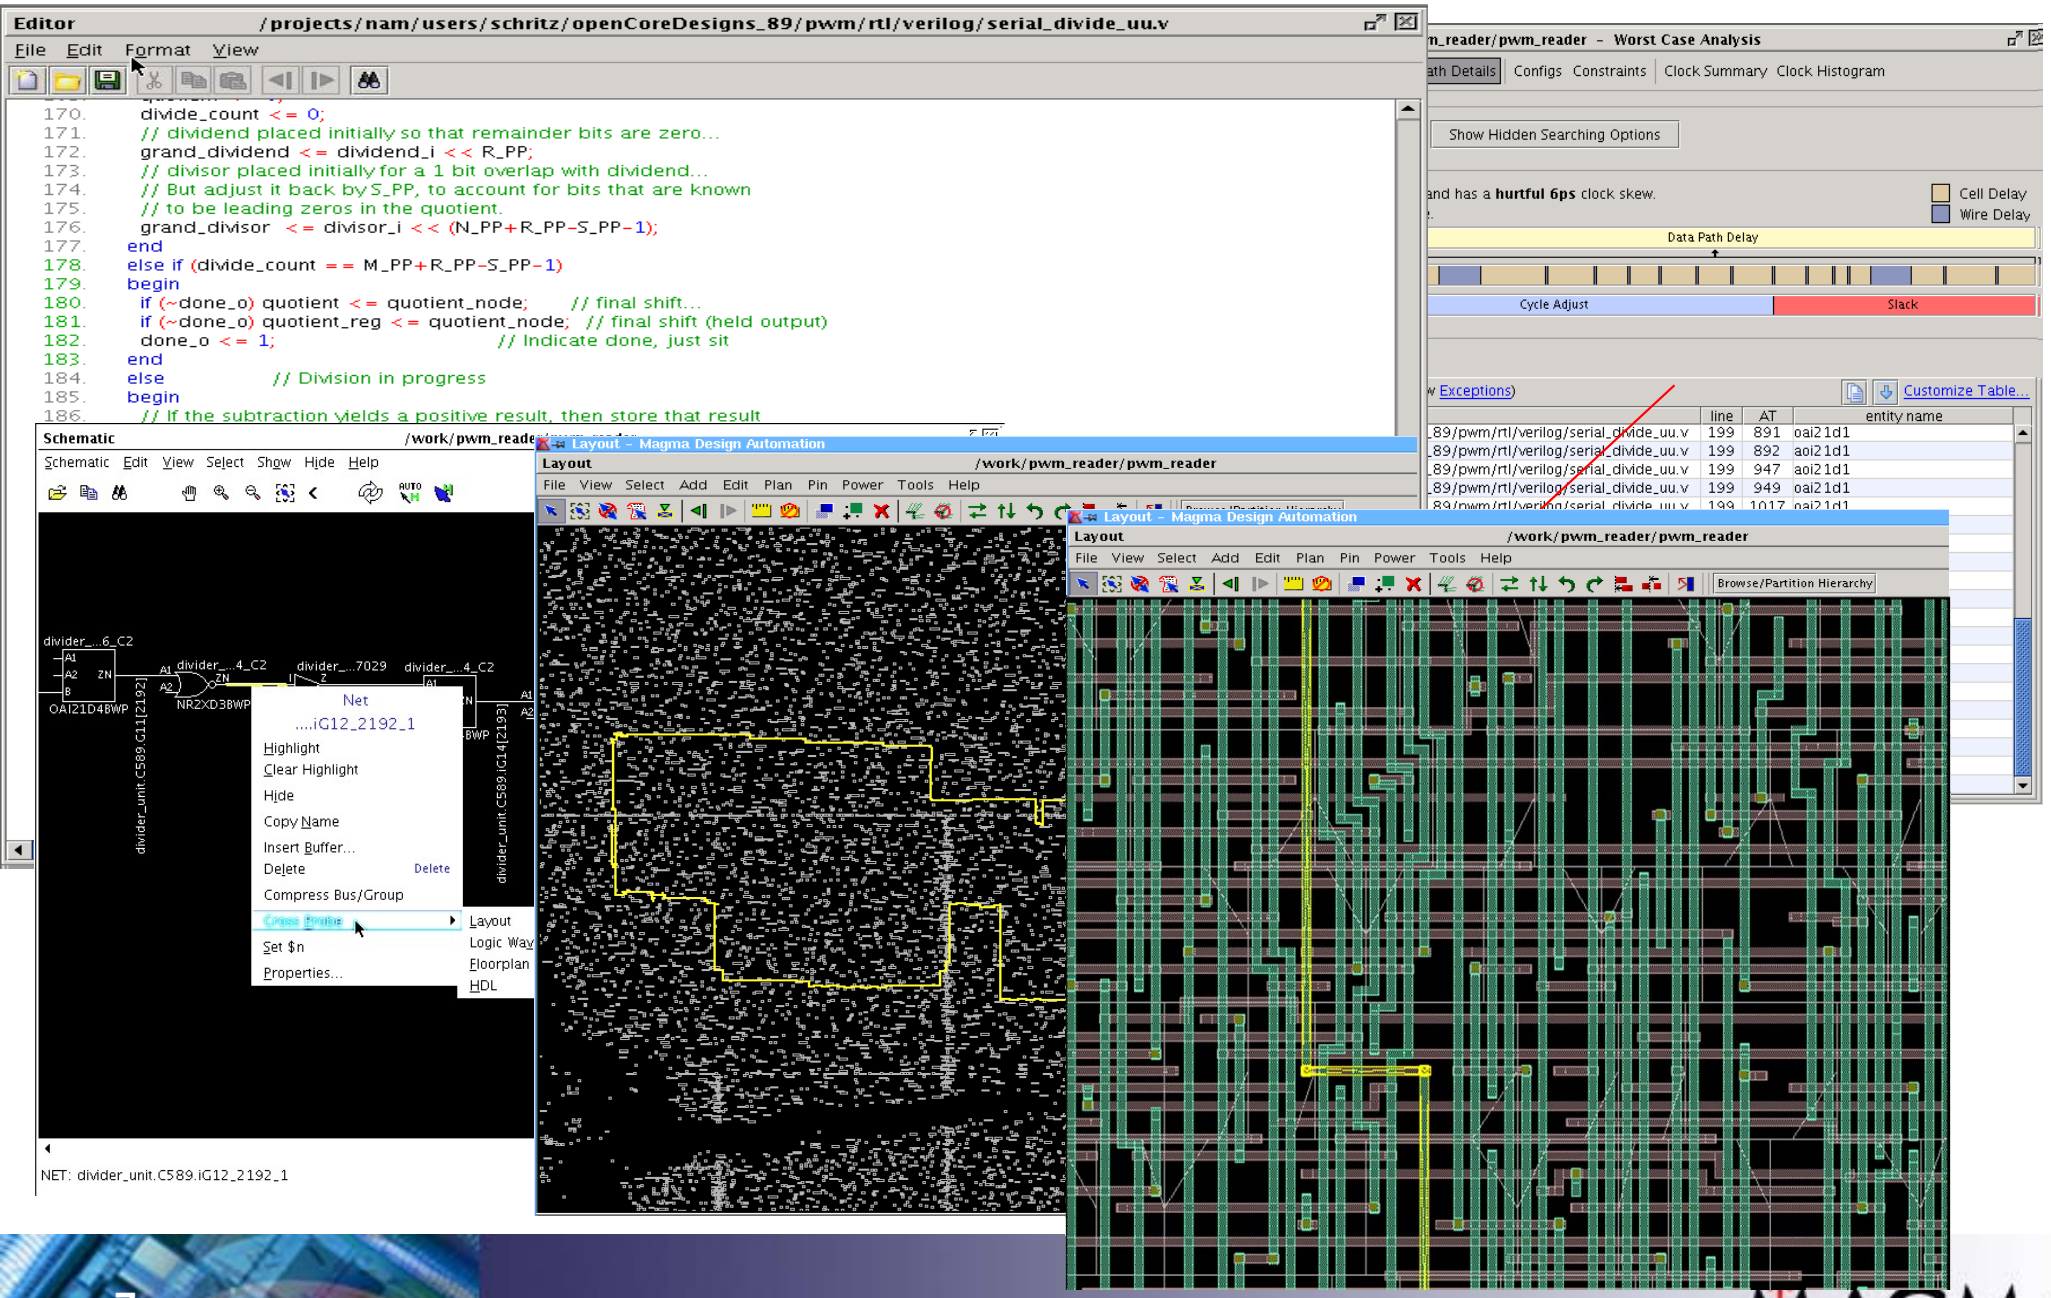

MAGMA

What is the timing accuracy?

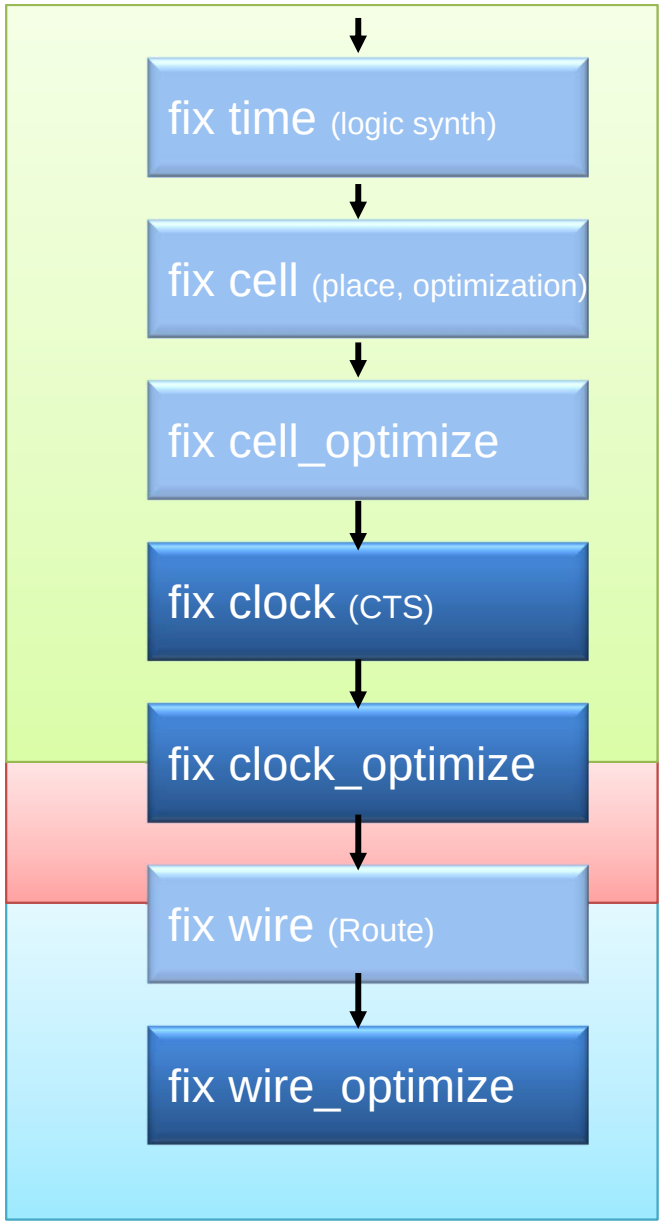

Global Route

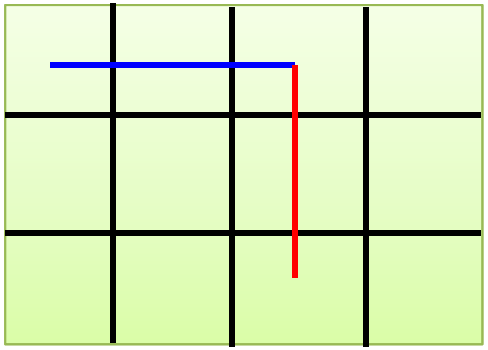

Extract glr segments
Delay calculator
Timer

GR-DR
Timing correlation?

Detail Route

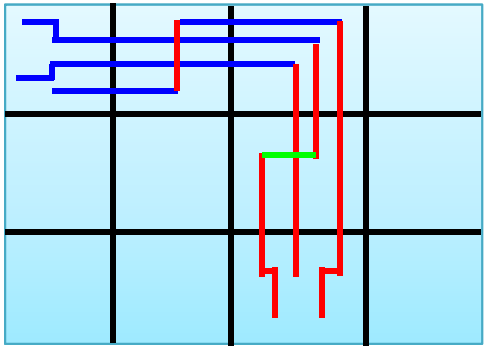

Extract detailed wires
Delay calculator
Timer

MAGMA

Measuring correlation error: Experimental set-up

- Take routed design:
  - Segments – time in global mode, CCT
  - Wires – time in final mode,. Xtalk on = golden
- Only compare 2-pin nets, > 40um length

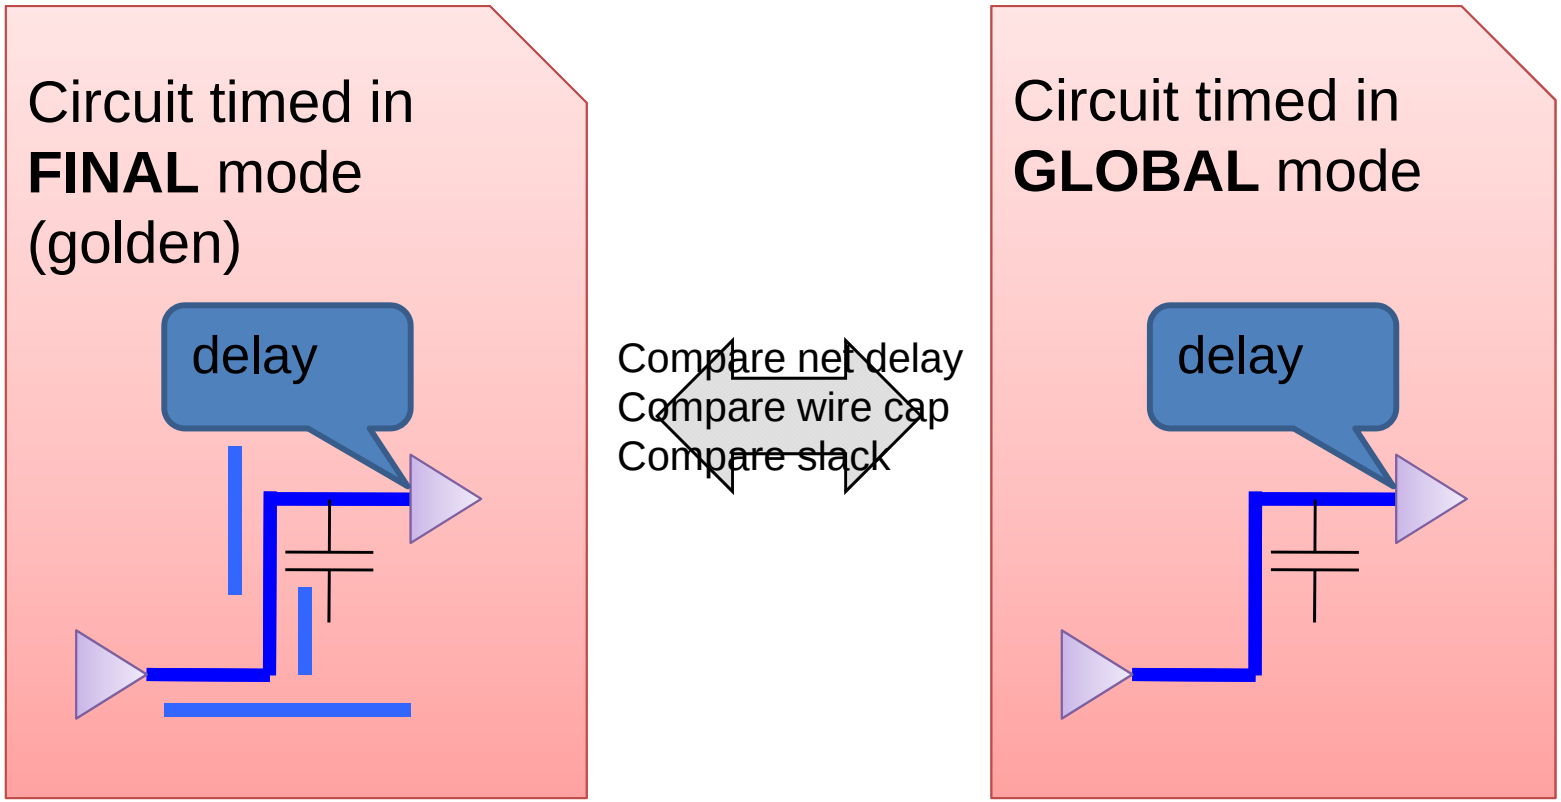

Observations on Global vs Final delay correlation

- Over 7 real designs, net delay miscorrelates badly between global and final:
  - Average = roughly OK
  - 88% standard deviation
  - So 33% of the net delays are off by more than 88%
  - 97% of nets are worse than +-5% accurate

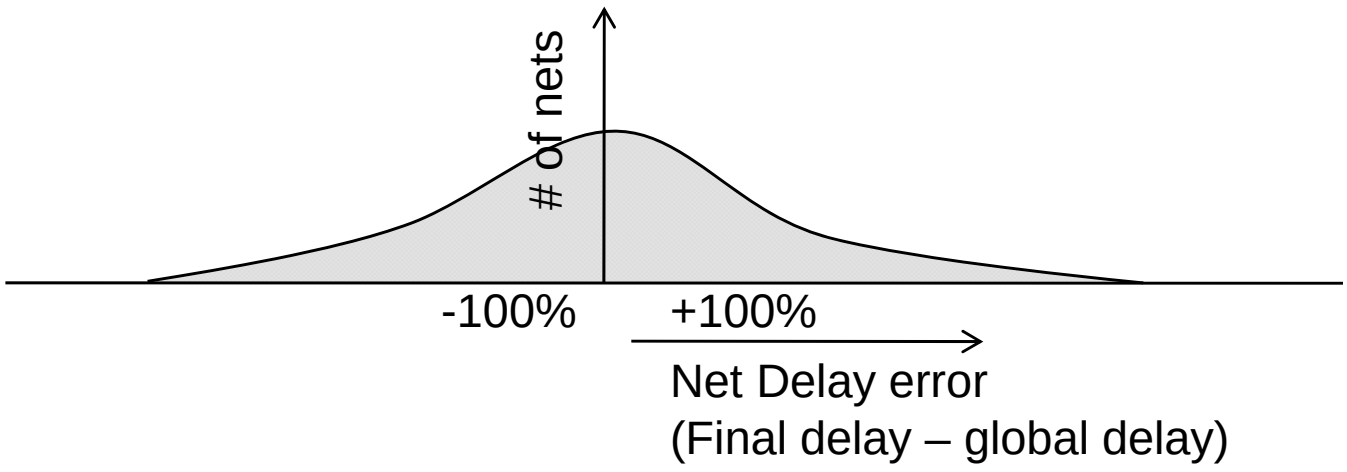

Garbage in – Garbage out ?

- Modeling inaccuracies, causes earlier opto to work on the wrong parts
- Crosstalk noise could seriously randomize results.

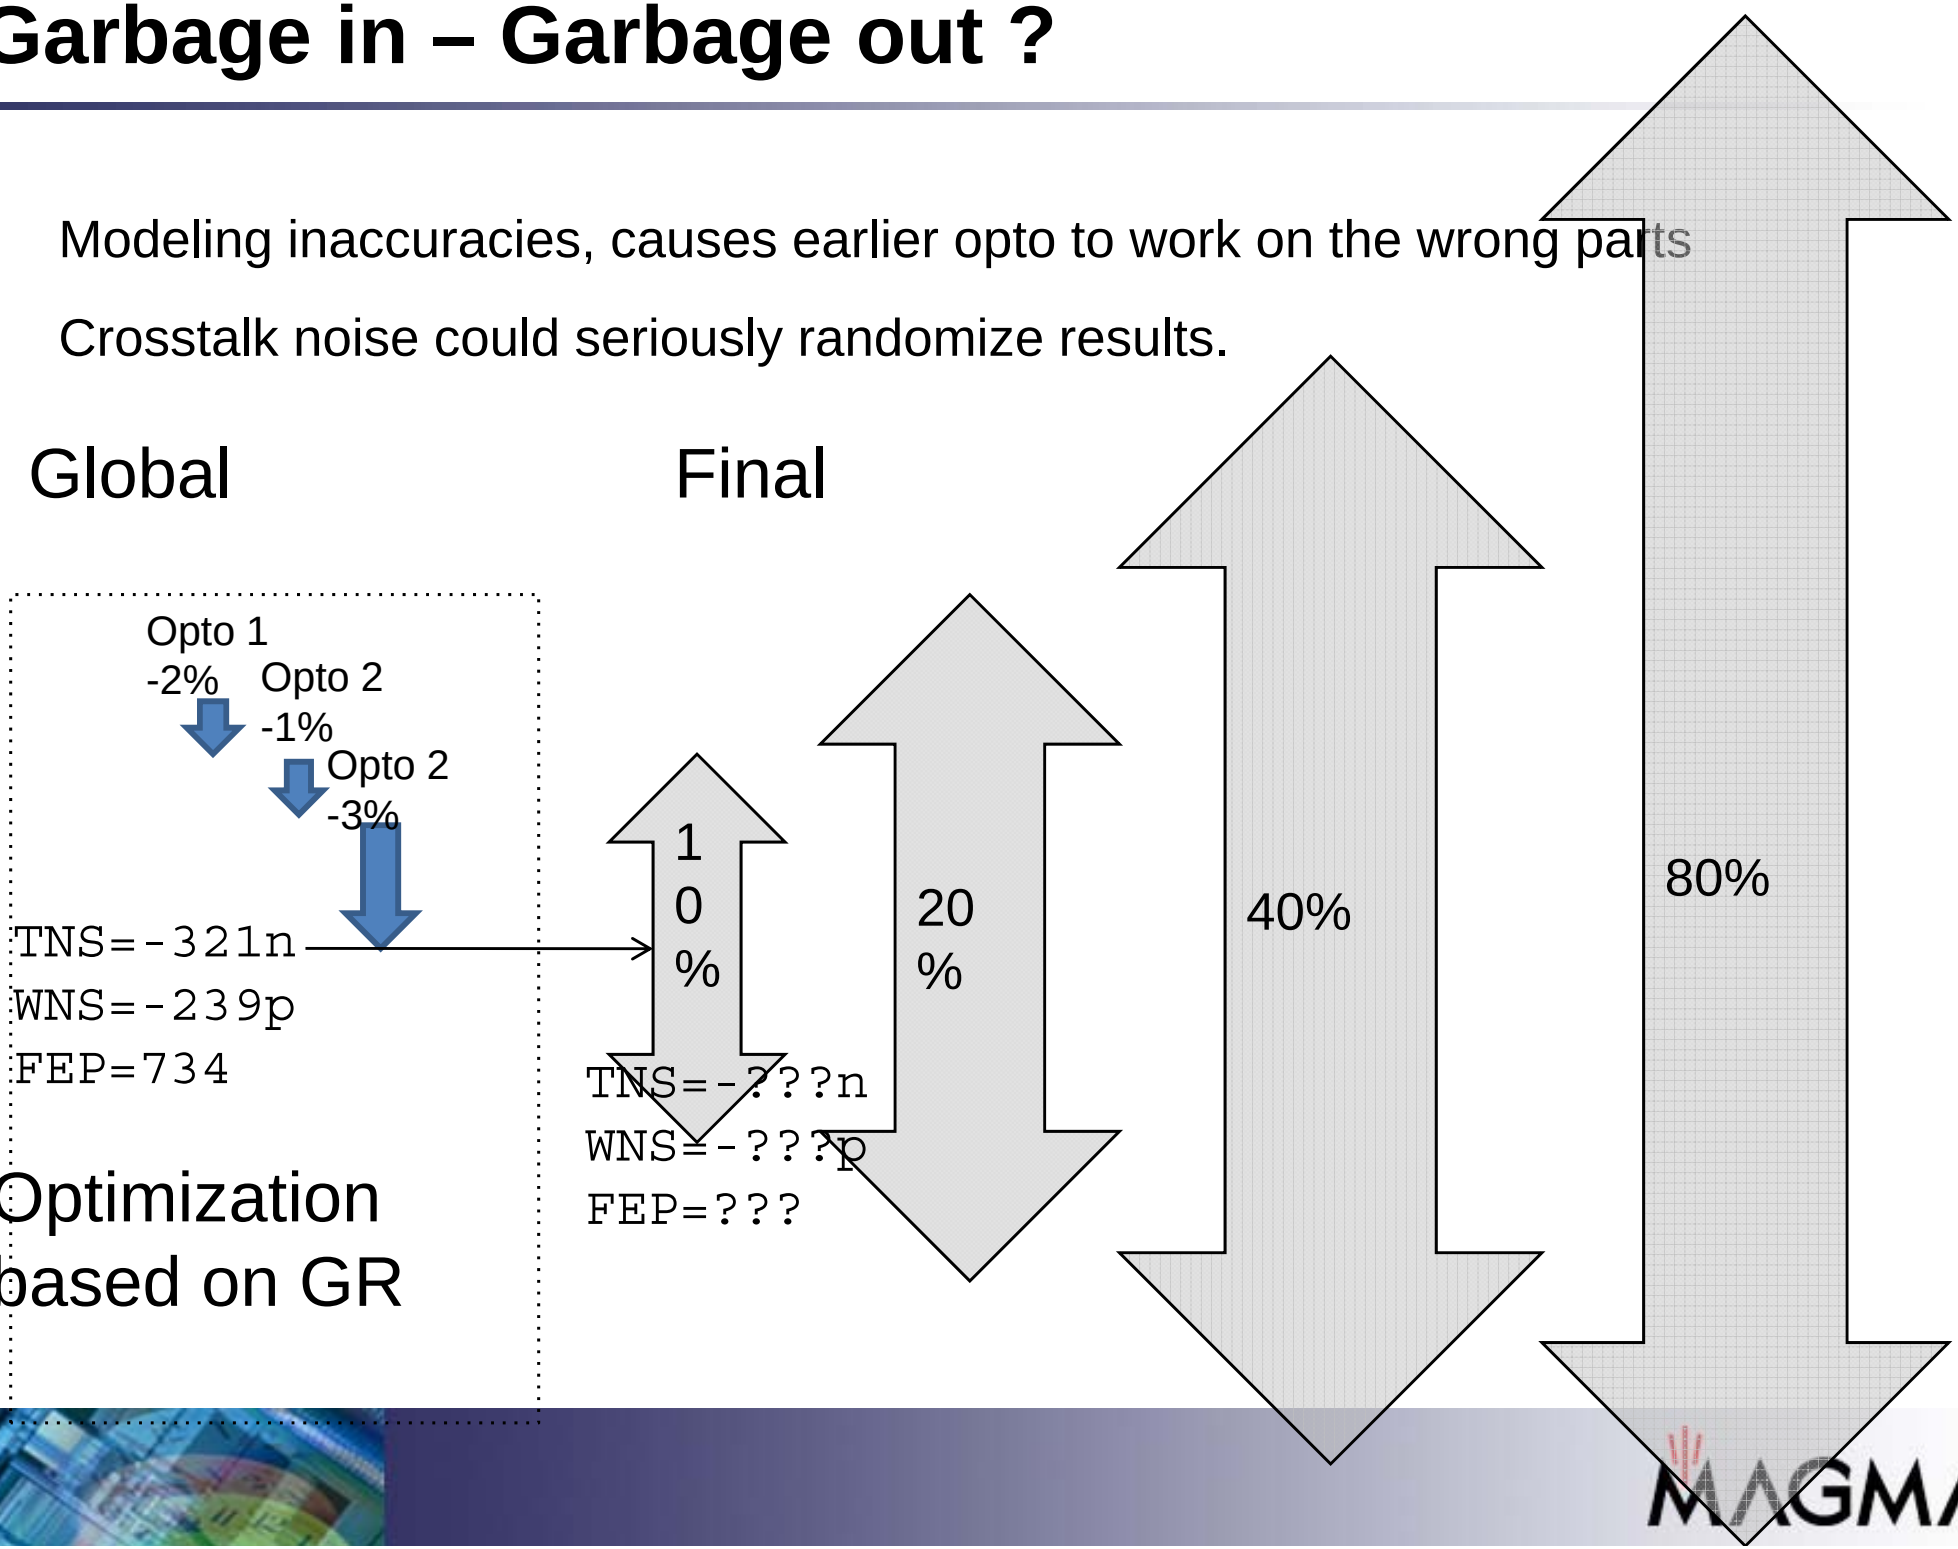

What can we do?

- Attempt to increase accuracy of early timing:
  - Add xtalk estimate during Global Route Extraction
  - Perform track routing as well
- And/or:
- Live with the problem:
  - Spend less effort on early optimization…
  - Carefully examine statistics of optimization effectiveness
  - Have a good way to patch up xtalk at the end

Building a Layout Design Flow

Observation 1:
Need gradual refinement flow using many algorithms

Observation 2:
Synthesis algorithms need highly simplified models of reality

Observation 3:
Synthesis algorithms cannot deliver good multi-objective trade-offs

Observation 4:
Optimizing a single objective often makes other objectives worse.

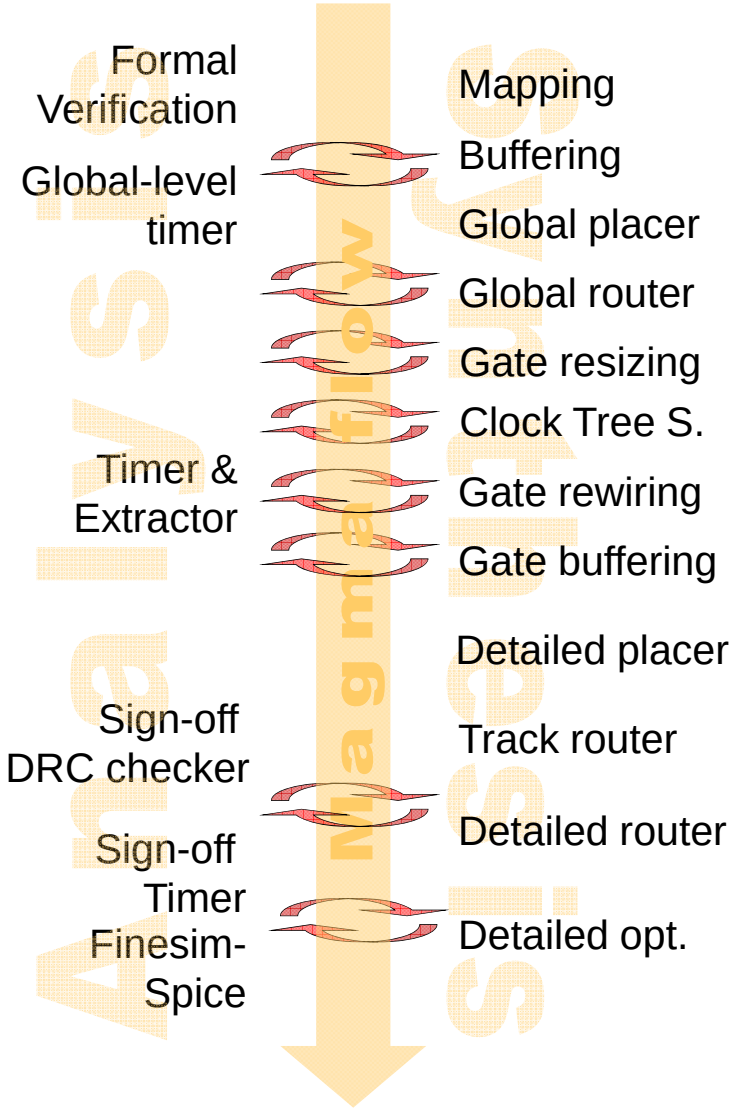

Optimal is not Optimal!

MAGMA

The ABC of a solid EDA Design Flow

A: Avoid

Use pessimism to make problem unlikely, 'Correct by Construction'

More avoidance = worse results...

B: Build

Synthesize using an algorithm

Synthesis is from Mars...

C: Correct

Fix each failure by incremental modifications (ECOs).

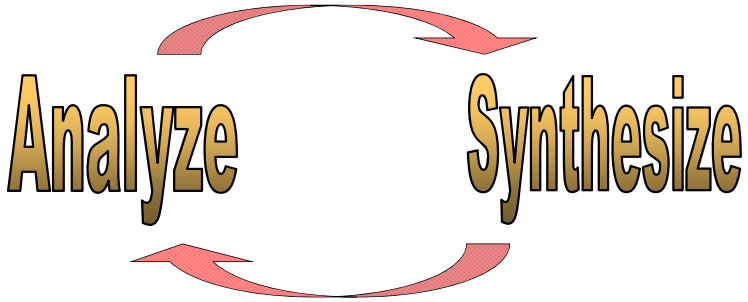

Goal: Living on the edge

- **A**void as little as possible
- … Such that the remaining failures can be **C**orrected incrementally

- And accept the reality that **B**uild algorithms offer little control

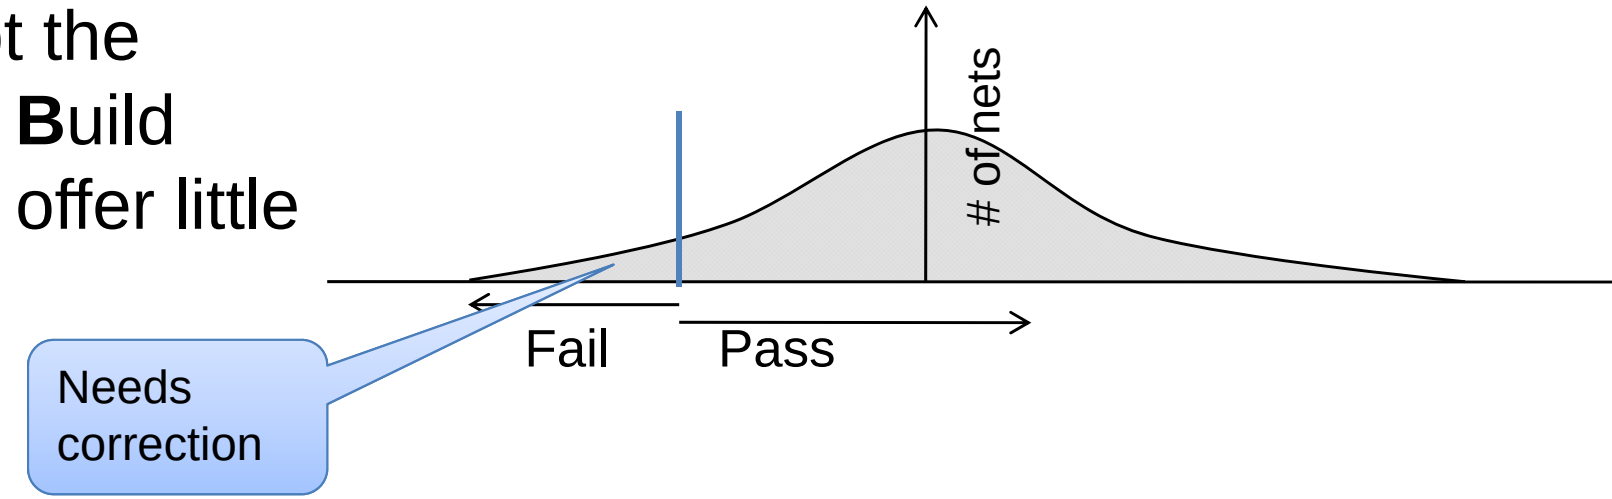

ABC in action: Combating crosstalk delay

- Avoid: using pessimism:
  - Size up all drivers: Costs cell area and power
  - Force double spacing NDR on many nets: Costs congestion = area
- Build:
  - Some routing tricks to spread & jog wires
- Correct using ECO:
  - gate re-sizing, buffering
  - Re-routing

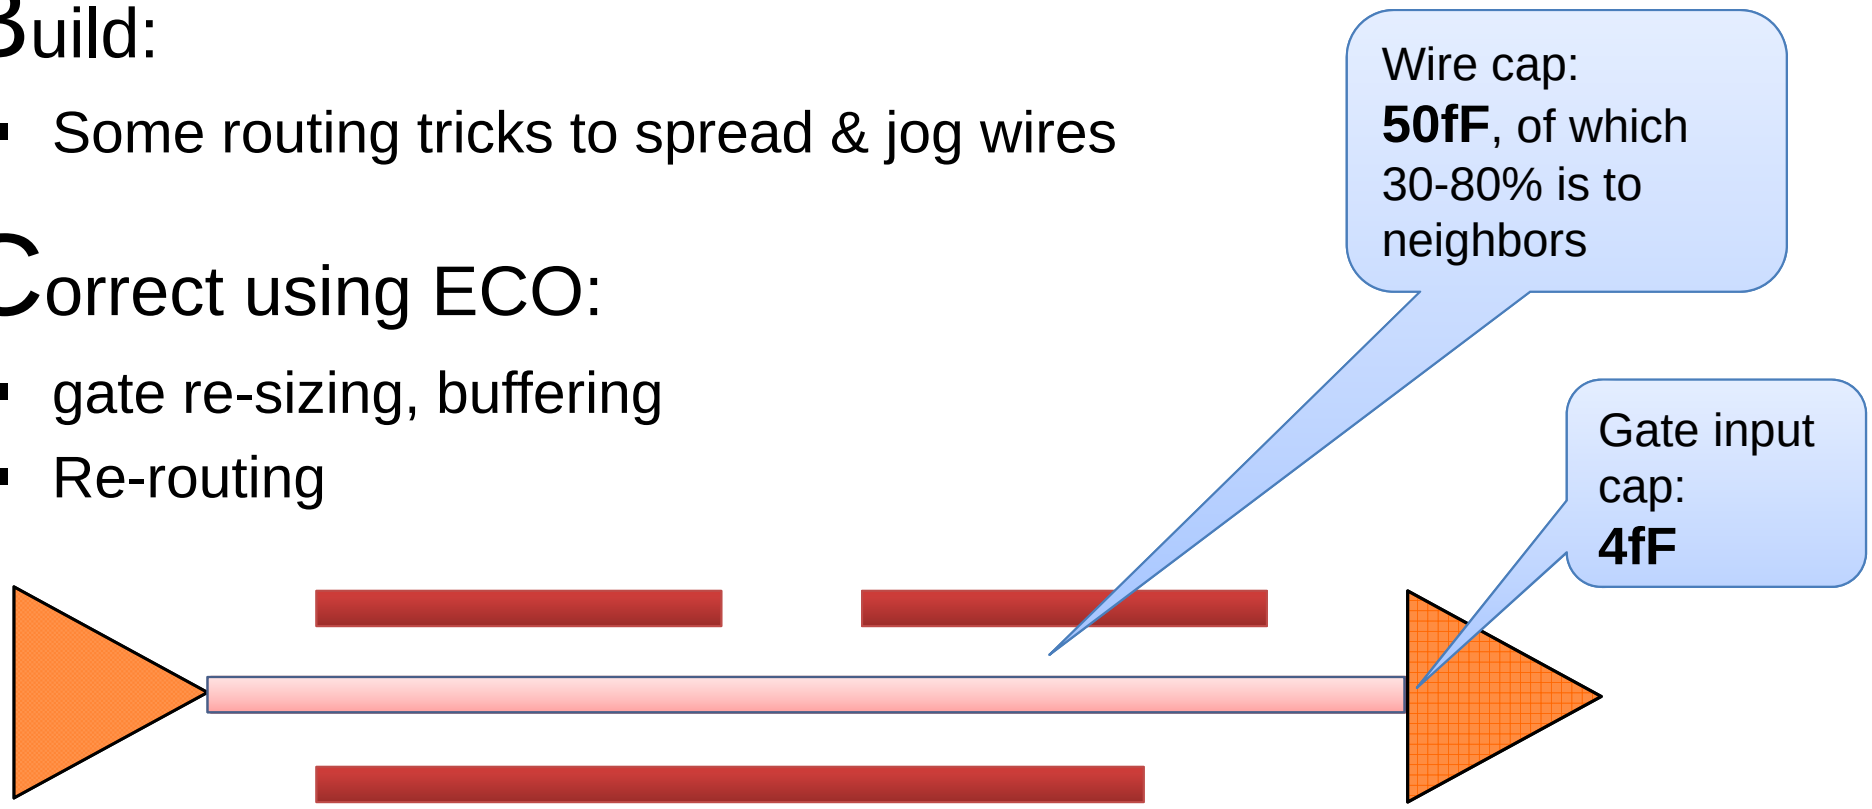

‘C’ routing improvement: pushing neighbors away

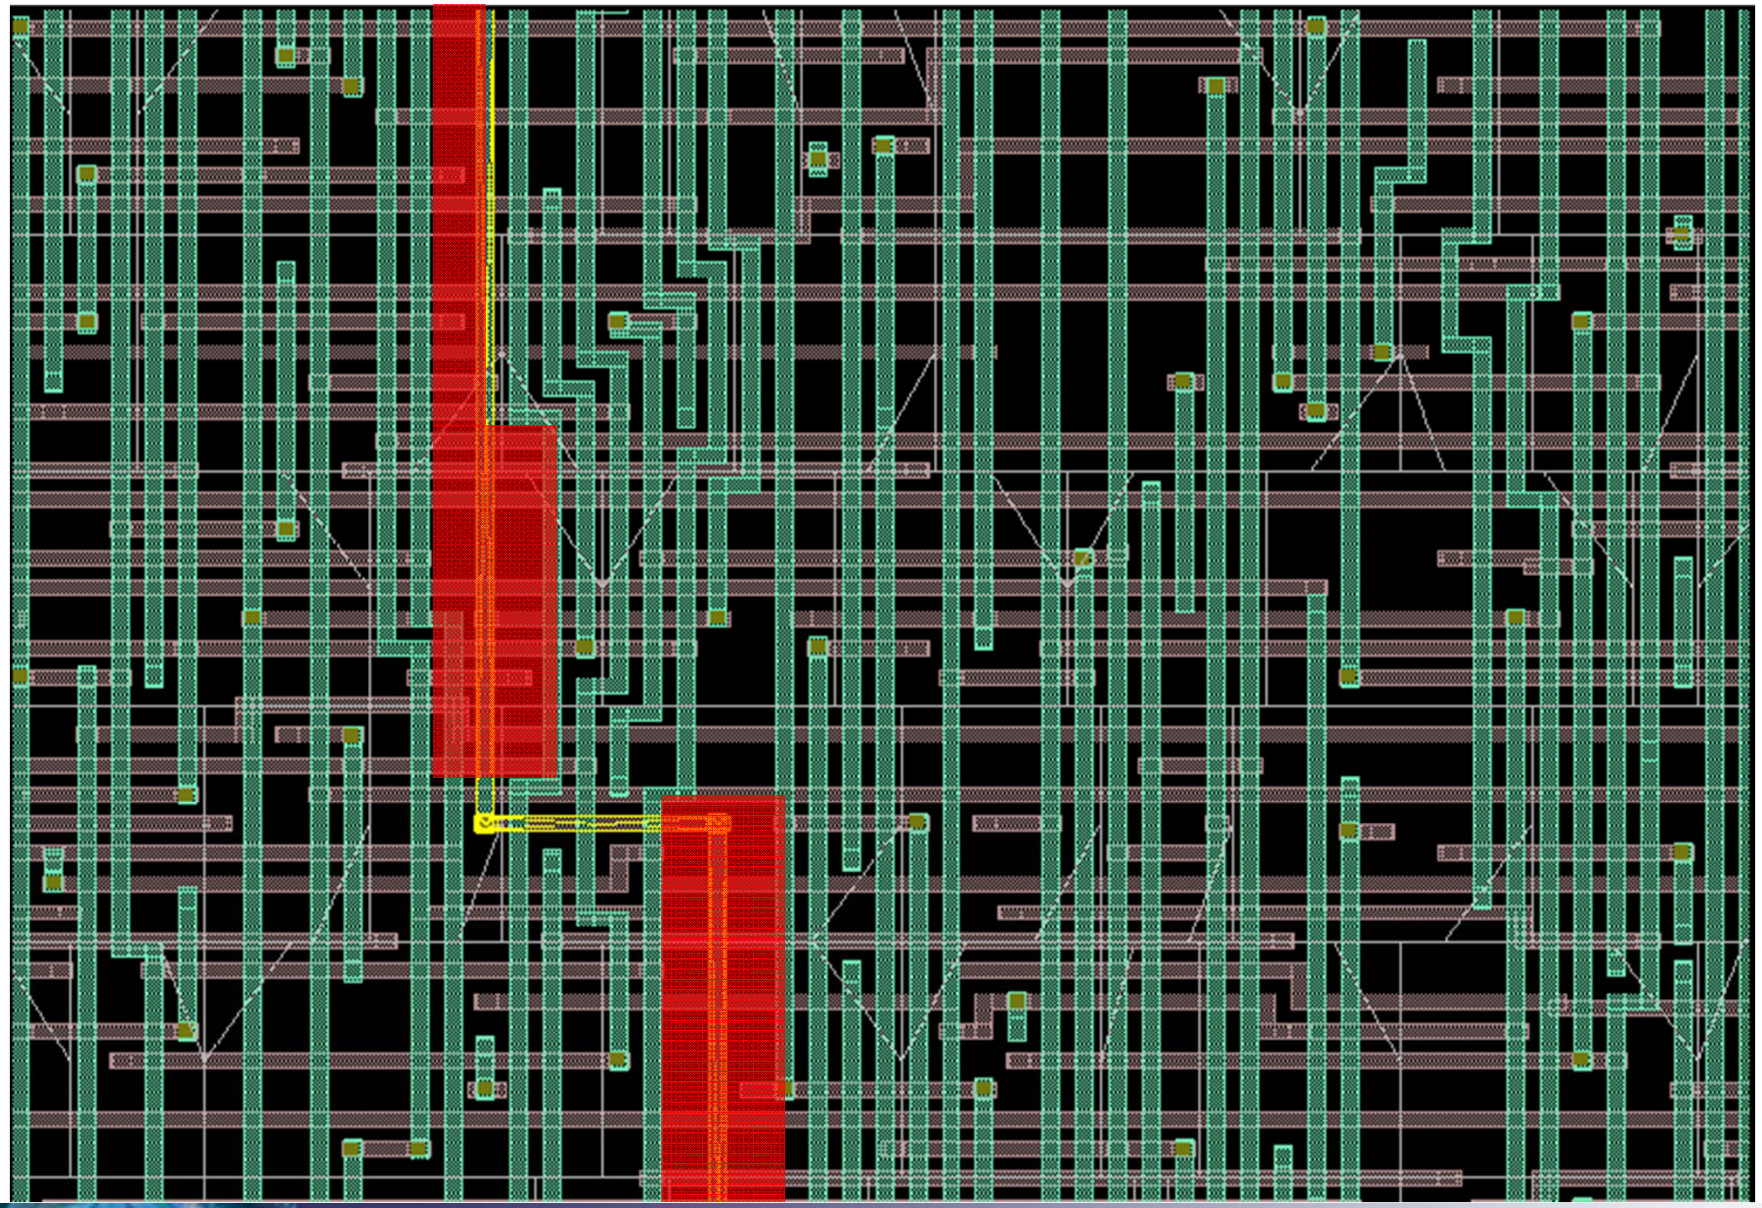

MAGMA

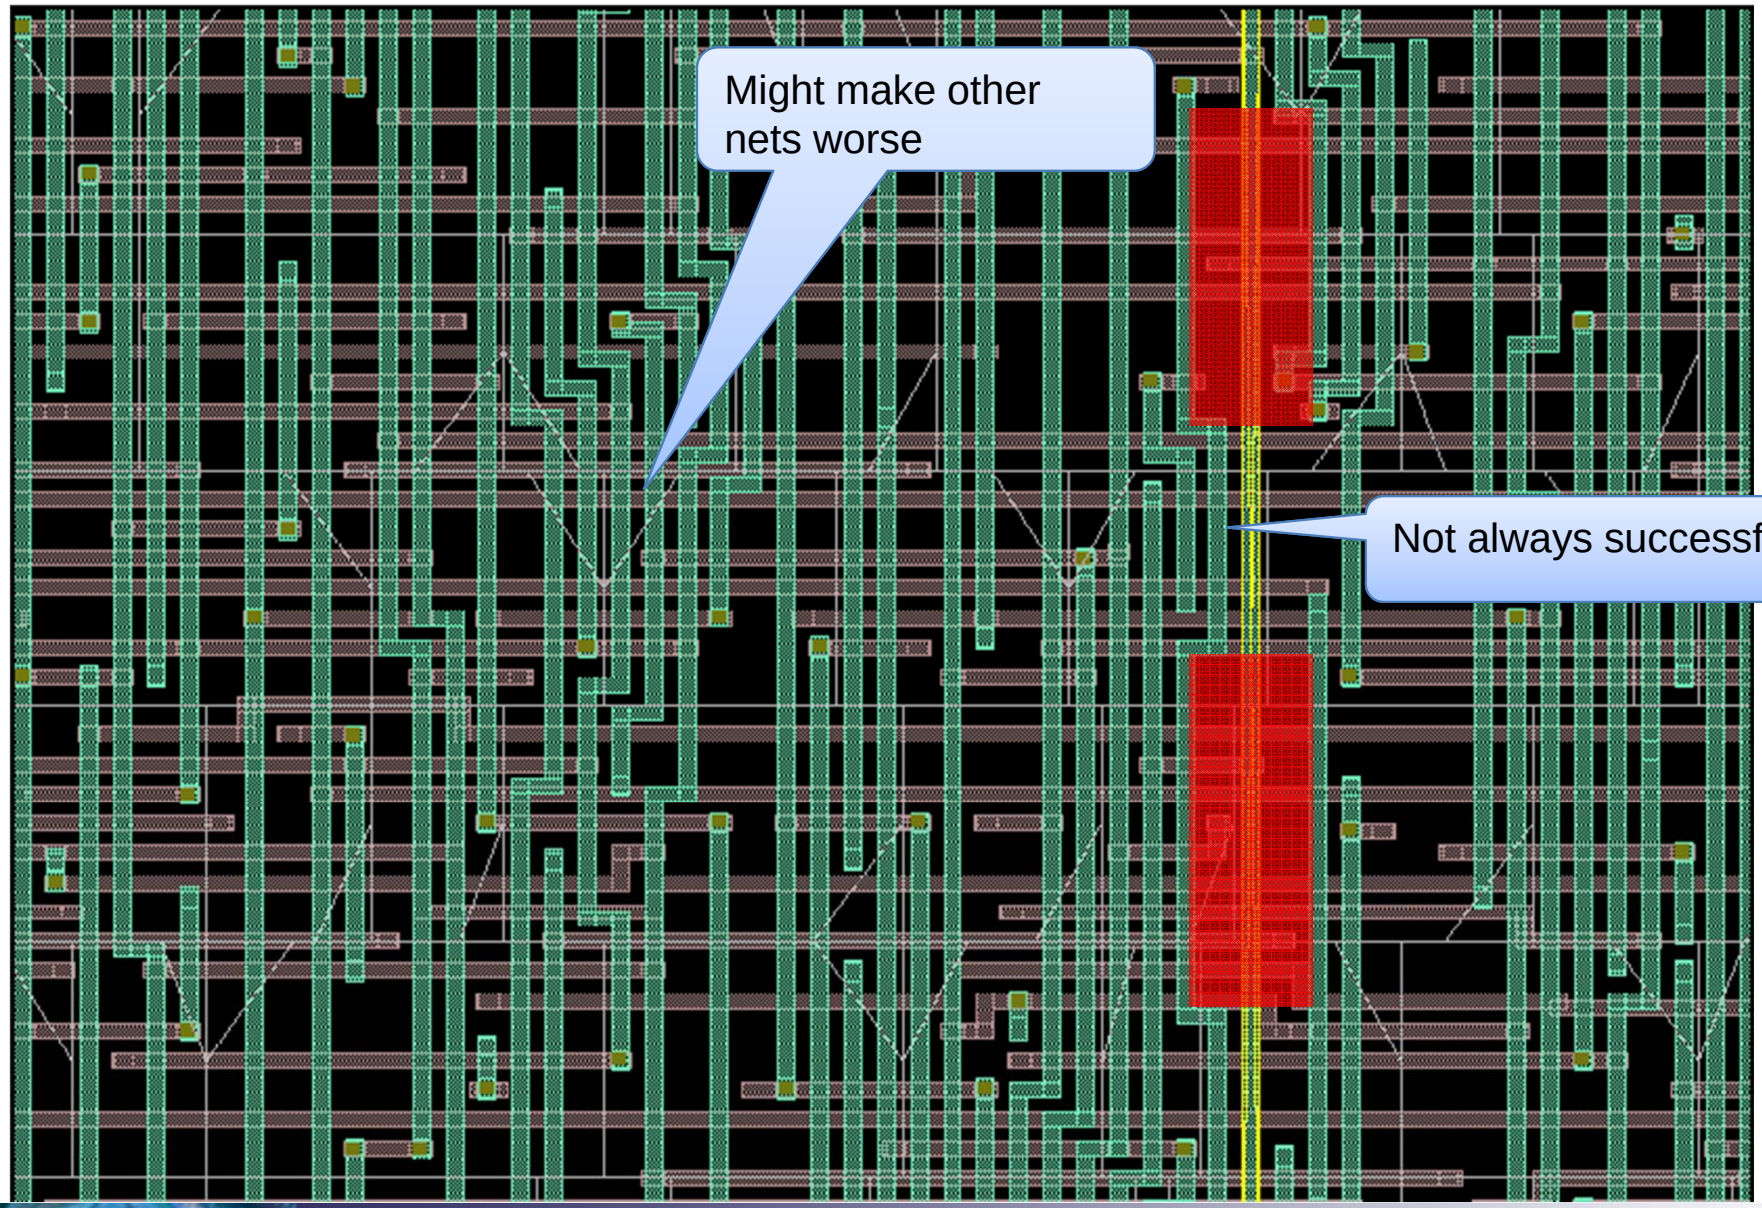
Might make other nets worse
Not always successful

Effect of this layout push on timing

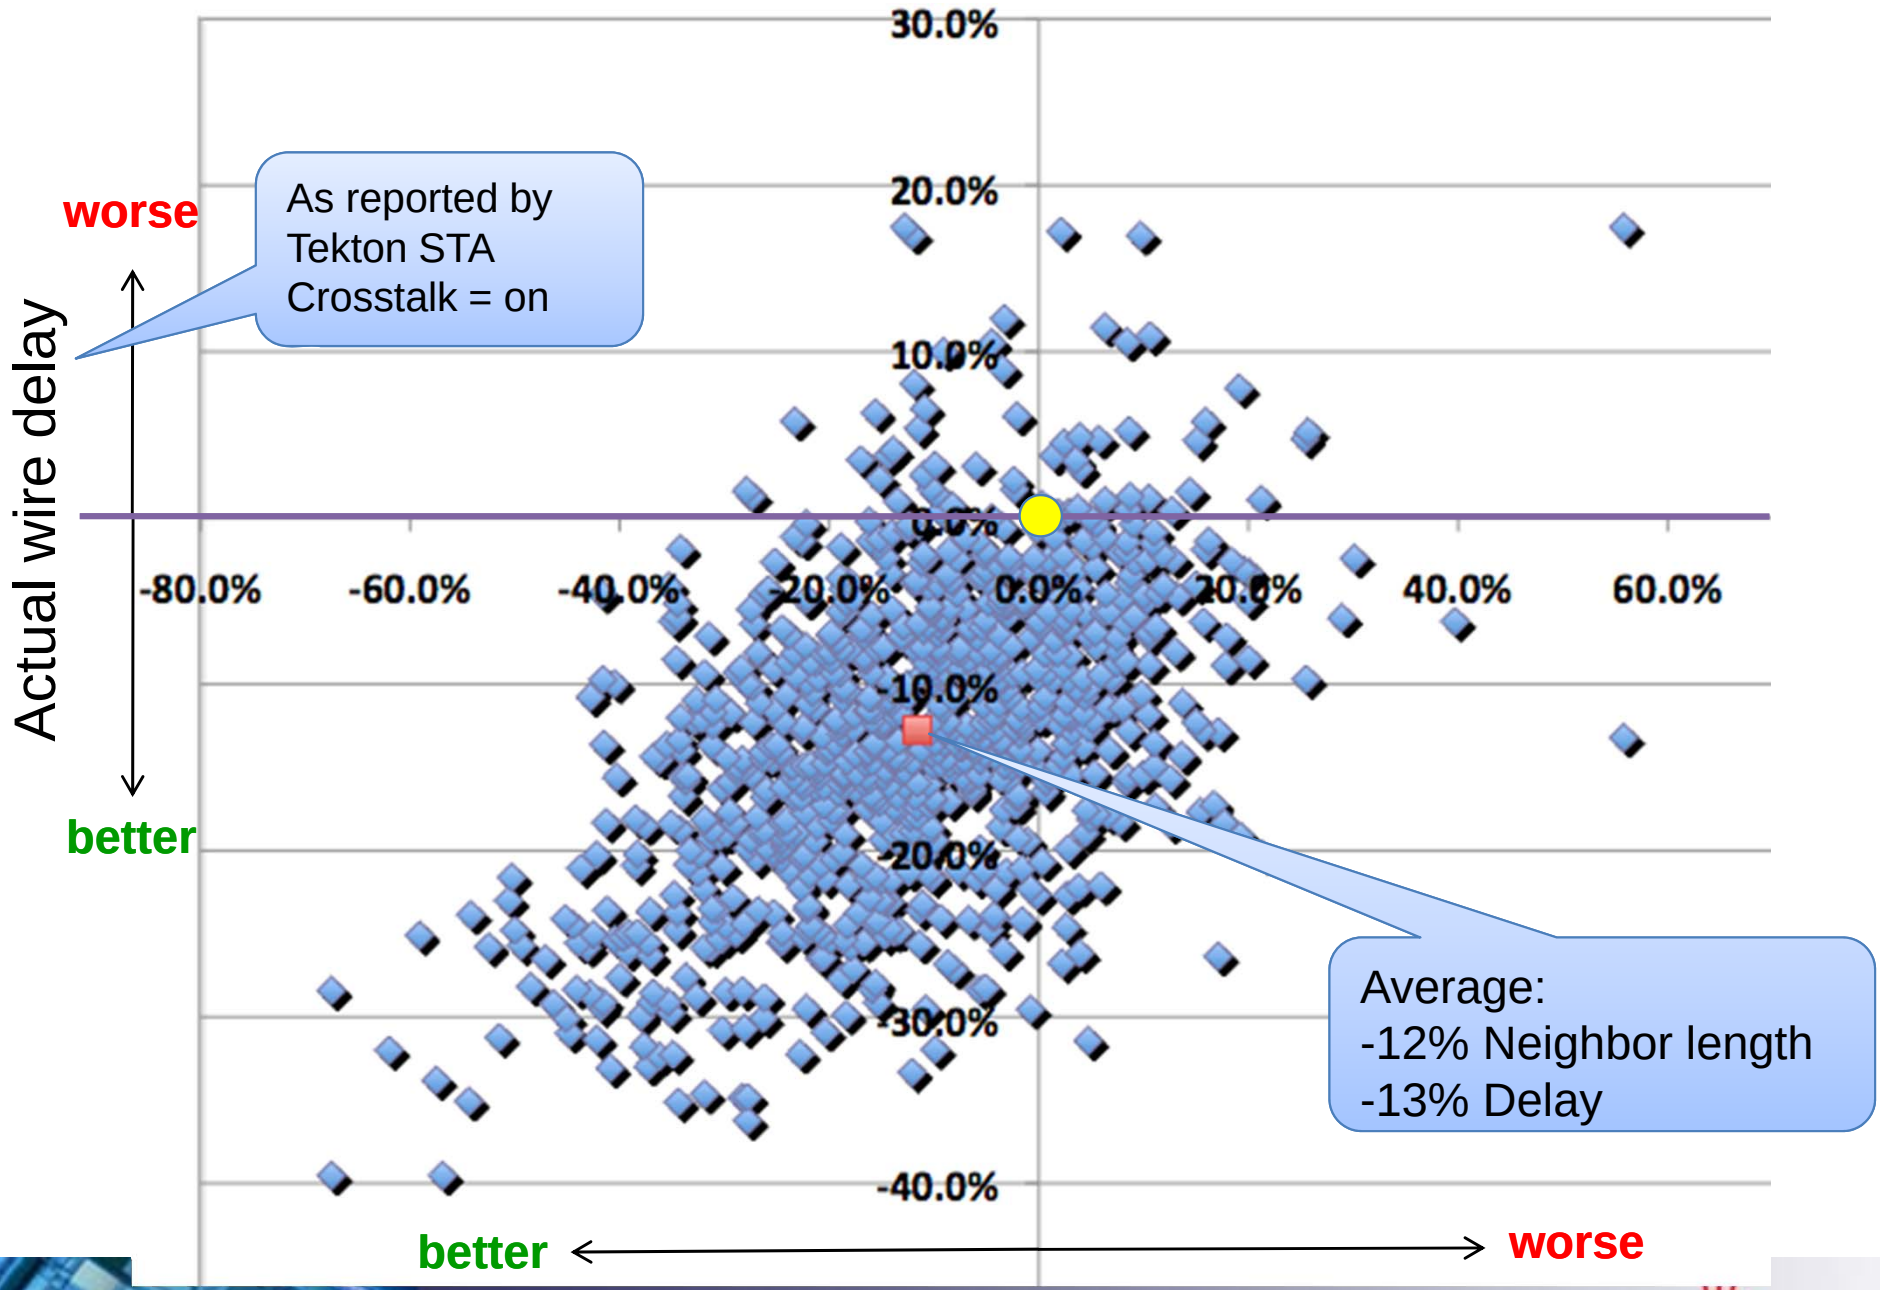

Medical tools vs. EDA tools

- New drug
  - Biological model of cause, actions and side-effects
- Develop it
- Test tube test
- Test on animals
  - Efficacy,
  - side effects
- Clinical trials
  - Large double-blind placebo controlled tests
- FDA-approval
- Deployment

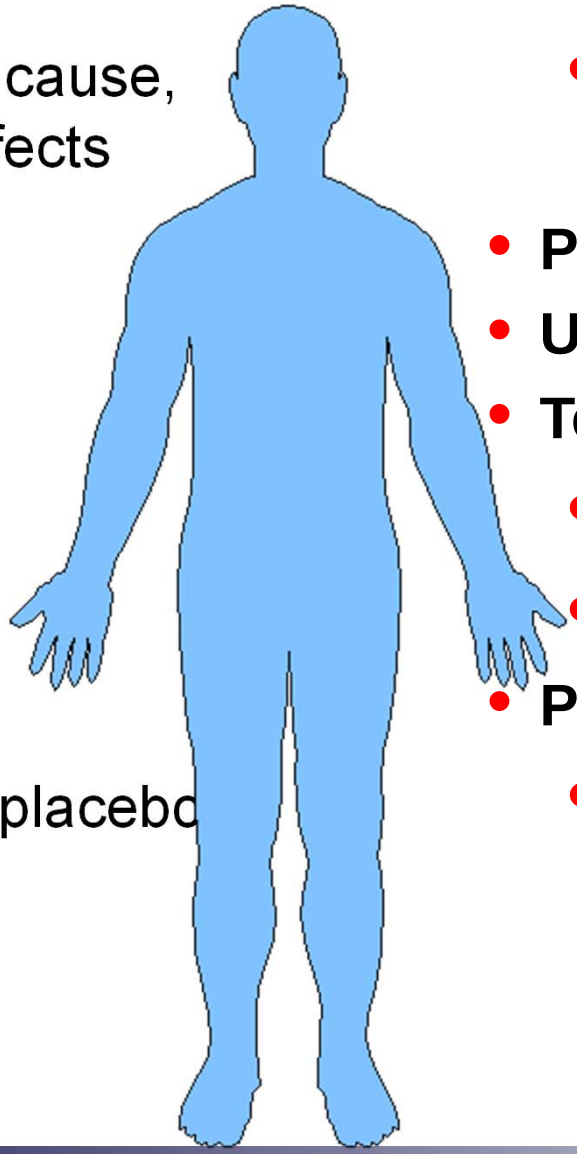

- **New Method/Algorithm**
  - Based on electrical/ physical plausibility
- **Program it (C++/TCL)**
- **Unit test**
- **Test on small testcases**
  - Debug program
  - Get a results table
- **Publish at ISPD**
  - Go for it!

Lack of Evidence = Quackery

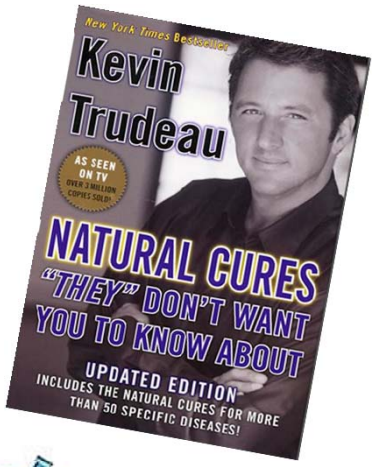

EDA
is not exempt:

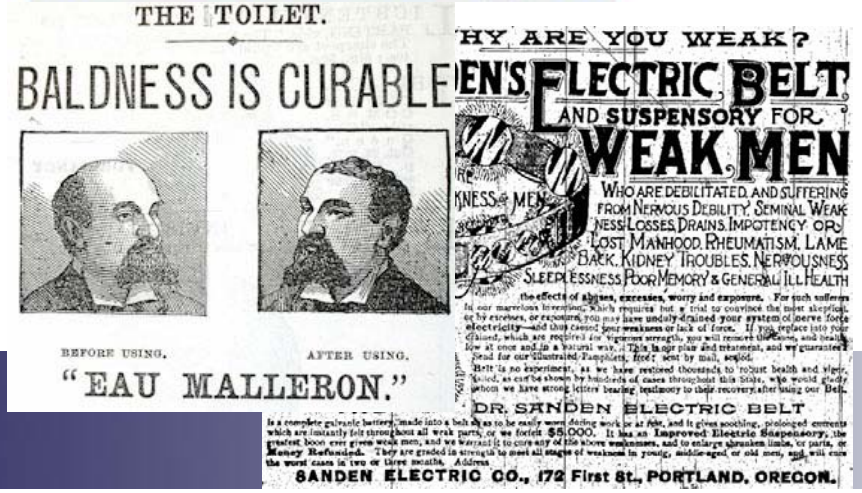

- Datapath placement
- Thermal-driven placement
- DFM-driven design
- Plug ‘n play tool interoperability
- Hybrid GPU/CPU EDA tools.
- Gridless routing
- X-Architecture

MAGMA

Skeptical wisdom for Electronic Design

- "Humans are amazingly good at self-deception"
  - This looks soooo good, therefore this *must* work
- "If it has no side effects, it probably has no effects either"
  - Example: improving temperature gradients will cost timing you! Are you really willing to pay based on the evidence?
- "Do not confuse association with causation"
  - "I took this airborne pill, and I did not get sick"
  - "I used this DFM optimizer, and the chip yields!
- "The plural of 'anecdote' is 'anecdotes', *not* 'data'"
  - Result could be a random effect, or another side effect
  - No substitute for unbiased placebo-controlled tests
  - Only large data sets are statistically relevant

Summary: observations from practice

- Layout is a multi-objective optimization problem
  - DRC, Manufacturability, timing, power, cost, design effort
- Timing is poorly predictable early in the flow
- The only thing that counts is the result at the end
  - Intermediate data is a poor indicator
  - Need hard evidence that trade off is worthwhile
- Beware of XX-driven synthesis/place/route
  - Is the gain worth the side effects?
- Optimal is irrelevant, while greedy is pretty good
- Simple A-B-C flows are proven in practice