IBM Research

Where is the EDA Research Going?

- How to Run an EDA Contest?

Cliff Sze
Research Staff Member
IBM Austin Research Laboratory

Disclaimer

- **These slides are my own personal opinions only and do not necessarily reflect the positions or opinions of my employer (IBM) or their affiliates. All comments are based upon my current knowledge and my own personal experiences.**

Where is the EDA Research Going?

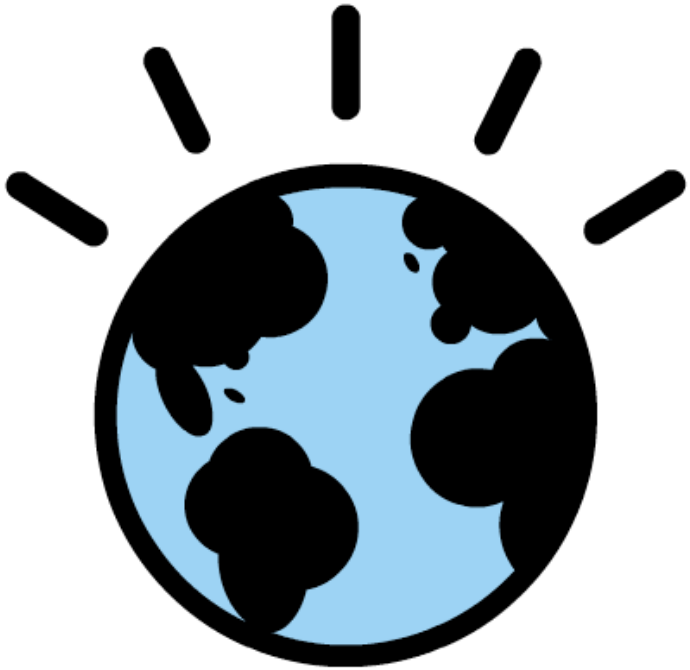

EDA veterans

Professors

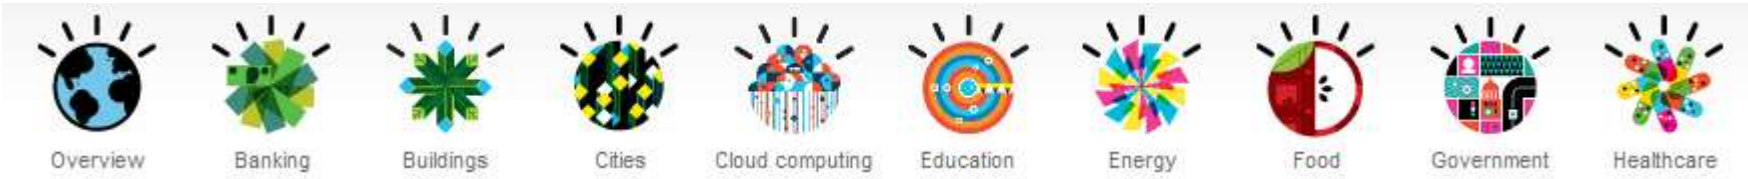

Cause?

- **No More Moore's!**
- **"Exponential" hits brick wall**
  - Power
  - Yield
  - Variability
- **End of Scaling**
- **ASIC cost is too high**

TECHWORLD

IT Jobs Newsletters Videos

News Reviews Features Blogs

Security Virtualisation Operating Systems Networking

Home > News

Operating Systems

In this channel:
News

Moore's Law is dead, says Gordon Moore

Key predictor of IT will end sometime, r
its progenitor.

CMP EE Times

EE Times:

Gartner: ASIC design starts to fall by 22% in '09

Dylan McGrath
(03/30/2009 3:01 AM EDT)
URL: http://www.eetimes.com/showArticle.jhtml?articleID=216401584

SAN FRANCISCO—FPGAs are displacing ASICs—a trend

ASIC design starts are expected to drop by 22 percent in 2

ASIC design starts—completed designs that have taped ou
push-outs of design projects into 2009.

ore's Law is dead, according
ore, its inventor.

e extrapolation of a trend th
ng ag
the
end,

SUMMARY FIGURE
PROJECTED GLOBAL MARKET FOR ASICS, 2007-2014
($ MILLIONS)

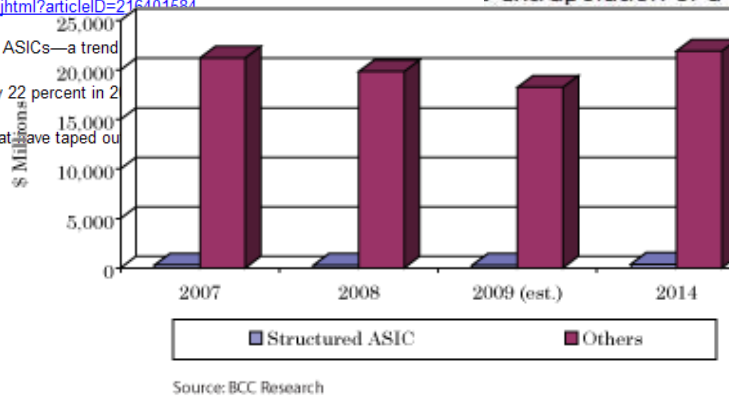

Source: BCC Research

Who's in Charge of EDA Research?

- **Industry**
  - Restricted by time and resources
  - Productivity
  - Quick and dirty
- **Academia**
  - Professors and students go after funding and publication
  - Focuses on "new" problems
    - (1) "discover" a new problem & provide simple solutions
      - Visionary, guide the direction
    - (2) explain an existing problem & provide robust solutions
      - Have to be really practical
  - Most papers are elegant, but….

EDA Research Problem Space

- **Why core problems are left untouched?**
  - "Marketing"
  - Infrastructures and realistic benchmarks

(1) How EDA contests can help the community focus on the core problems?

(2) Example of core EDA problems in IBM

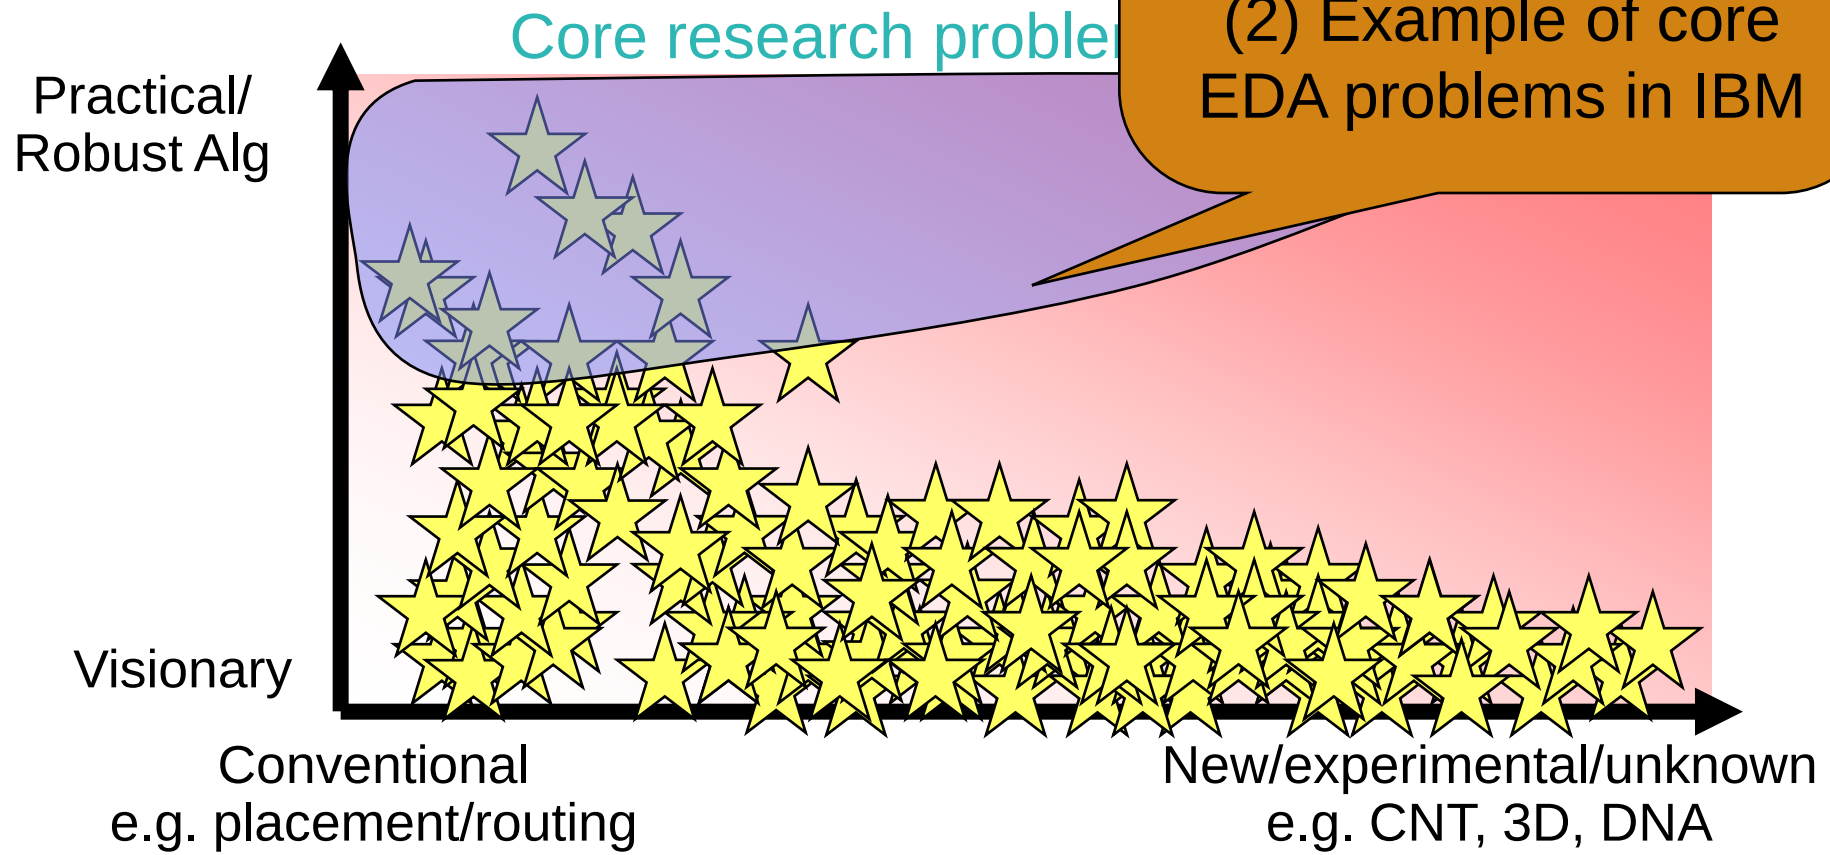

International Symposium on Physical Design Contests

2005	**Placement**	**- HPWL was the sole quality metric** **- 8 new industrial benchmarks released**
2006	**Placement**	**- scale HPWL with runtime and density** **- 8 new industrial benchmarks released**
2007	**Routing**	**- 2D/3D versions, Overflow as sole quality metric** **- 8 benchmarks derived from 05/06 placements**
2008	**Routing**	**- Realistic via formulation, scale WL with CPU** **- 8 new benchmarks released**
2009	**Clock Synthesis**	**- Clock Skew Range over 2 SPICE with diff Vdd** **- 7 new benchmarks are used**
2010	**Clock Synthesis**	**- sub-8ps clock skew target w/ Vdd/wire variations** **- Real benchmarks released from IBM/Intel**

ISPD 2005 Placement Contest

- **First ISPD contest**
  - Placement algorithms all far from optimal
- **9 academic placement tools participated**
  - Good coverage of placement tools
- **8 new placement benchmarks were released.**
  - All were derived from real industrial ASIC designs
  - Extensively being used in placement research
- **HPWL was used as sole quality metric**
  - No routability estimation
  - No timing analysis
  - No runtime measurement
- **Analytic placement tools dominated**

ISPD 2006 Placement contest

- **Total 16 new placement benchmarks**
  - All derived from real ASIC designs
  - Variety of floorplans
  - Big Block Placement
  - 5 benchmarks with more than million objec
- **ISPD 2006 Contest**
  - Indirectly address routability issue
  - Turn-around time (Productivity)
  - Huge Improvements from ISPD 2005 results
- **Still not timing-driven placement**
- **Impact**
  - Significantly more publication on placement
  - Greatly improve our placer

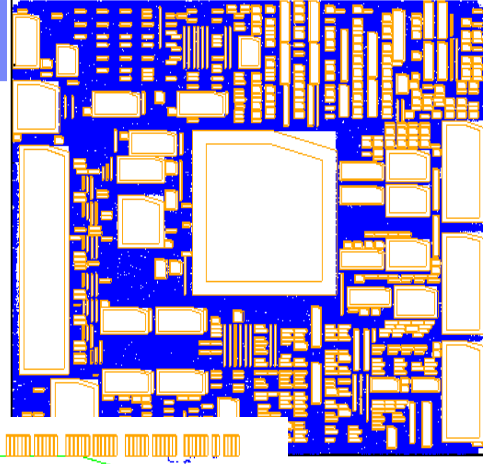

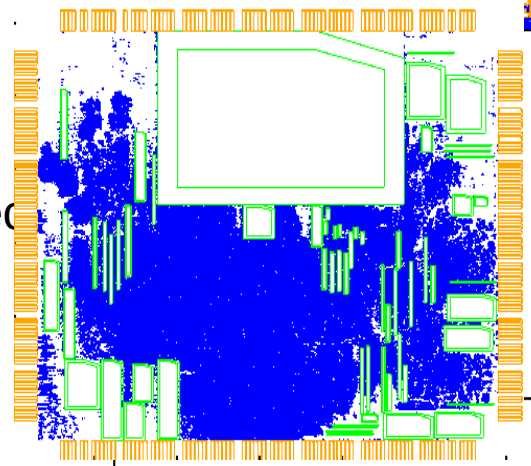

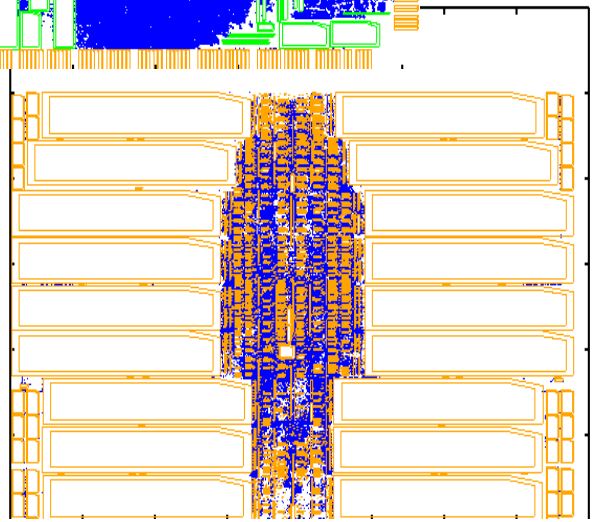

ISPD 2007 Global Routing Contest

- **3 initial teams from industry**
- **11 final entries**
- **8 new global routing benchmarks are released**
  - All derived from ISPD 2005/2006 placement benchmark solutions
- **Contestants had about 2 weeks to run their global router on benchmarks**
  - Organizer verified all global routing solutions with an official script
- **Quality metrics**
  - Minimizing overflows
  - Routed wire length - second objective
  - No CPU time limits
- **Winner: Mike Moffitt (an AI guy)**
  - Doesn't know much about VLSI

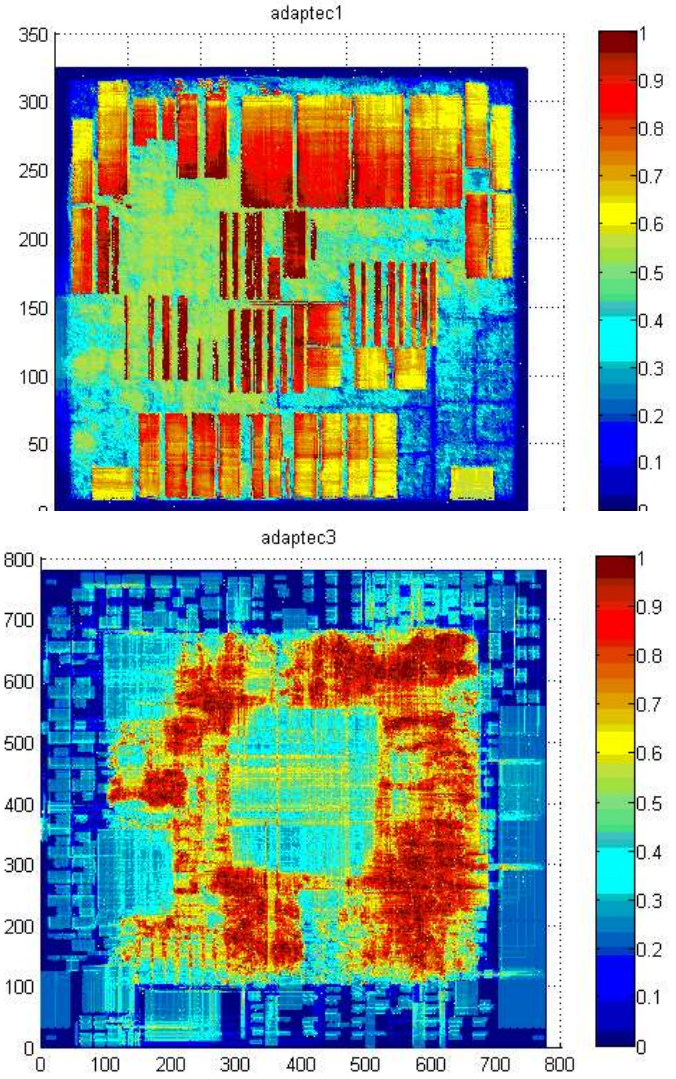

ISPD 2008 Global Routing Contest

- **Huge improvements from 2008 Results**
- **Discussion with Cadence, Synopsys and Magma routing experts**
- **A good set of global routing benchmarks**
  - Overflow minimization
  - Hard CPU limit of 24 hours
  - CPU-weighted wire length minimization
  - One via connecting two consecutive metal layer = WL of one g-cell
    - Based on resistance matching
- **Limitation**
  - Ignore other metrics
    - 90% congestion
    - Average 20% worst nets
  - Information inside a G-cell is ignored

ISPD 2009 Clock Network Synthesis Contest

- **Clock network Synthesis**
  - "Easy" problem with complicated rules…
  - Very different from Placement and global routing
    - hard problems with "simple" rules
- **Typical high performance clock network synthesis**
  - More latency near the source
    - Hard to control skew
    - Trade power for robustness
  - Minimize power near clock sinks (e.g. latches, registers, FF)
- **Contest formulations**
  - Realistic data for 45nm technology node
  - Accurate delay calculation by SPICE
  - Power limit
  - Real clock skew considering Vdd variations
  - 7 benchmarks are created roughly based on real clocking problems

2009 CNS Contest Details

- **Ngspice release 18**
- **Predictive Technology Model (PTM 45nm HP)**
  - Matches IBM model for HP uP
- **Two inverters**
  - Mid-sized inverter
    - 10um nmos, 14.6um pmos (for similar R/F delay)
    - input cap = 35fF resistance = 61.2Ohm output parasitic cap = 80fF
  - Small inverter
    - 1.37um nmos, 2um pmos
    - input cap = 4.2fF, resistance = 440Ohm, output parasitic cap = 6.1fF

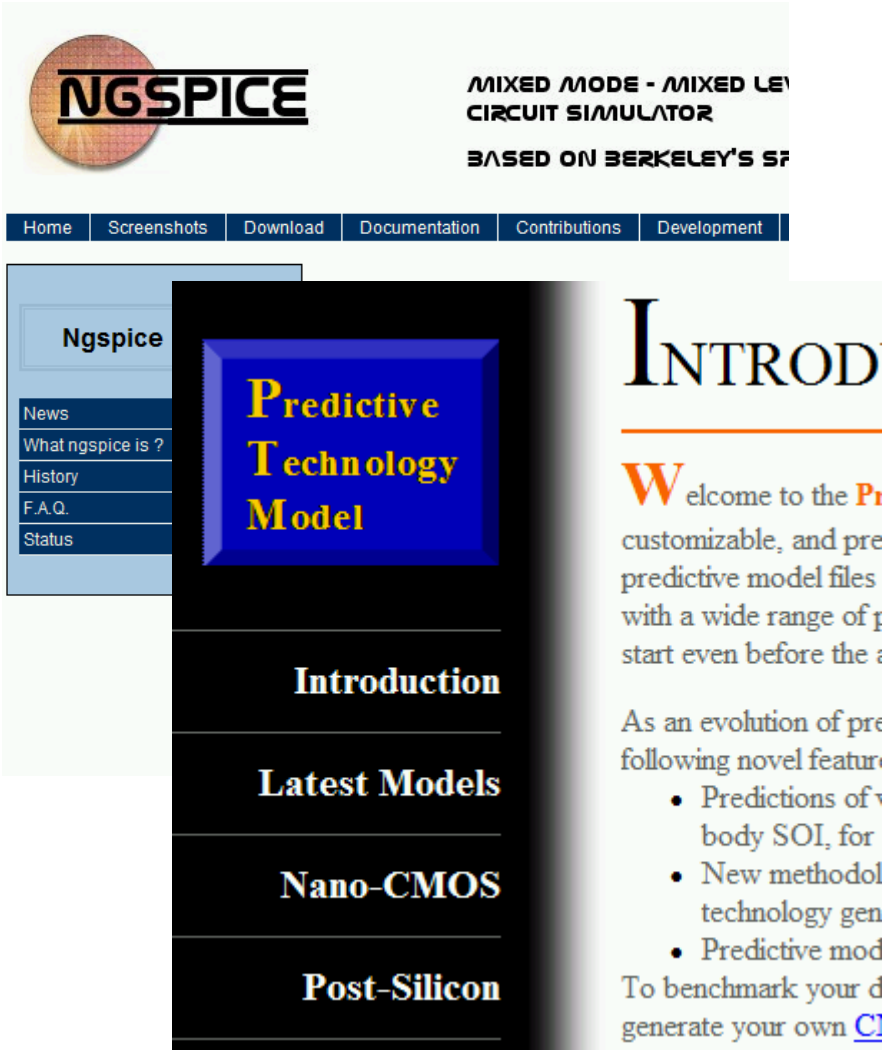

2009 CNS Contest Details (2)


Inverters connected in parallel

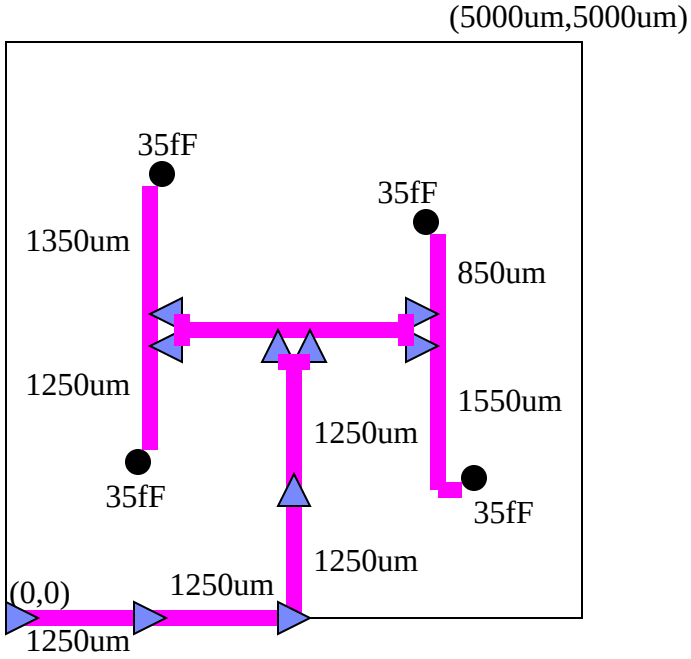

- **Two wire types**
  - Loosely based on IBM 45nm technology data
  - Wide 0.1 Ohm/um 0.2 fF/um
  - Narrow 0.3 Ohm/um 0.16 fF/um
- **Slew (10%-90%) limit = 100ps**
- **The source is directly driving the mid-sized inverter.**
- **The input slew to this inverter is 100ps.**
- **Clock source is at (0,0).**
- **Vdd = 1V**
- **Clock frequency = 2GHz**
- **Clock period = 500ps**

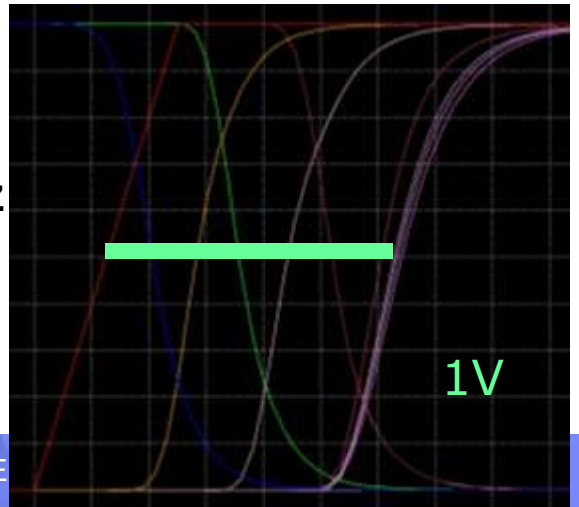

Lessons Learned from 2009

- **Clock latency range (CLR) is not practical**
  - This upper bound is too loose
  - What if we use CLR for MCMM?
- **No wire variation is considered**
  - Encourage more wire delay
- **Not challenging enough**
  - Minimize CLR with power limit
    - All teams use clock trees
      - Best nominal skew: ~5ps
      - Best CLR: ~30ps
  - Skew requirement should be tighter
- **The contest was a mixture of ASIC methodology and server methodology**
  - ASIC needs very fast algorithm (clock tree)
    - 20k sinks for 5 mins (parallel programming?)
    - No SPICE, Elmore or moment matching
  - Microprocessor demands high robustness (grid)
    - Skew with OCV < 10ps
    - SPICE simulation with greatest accuracy

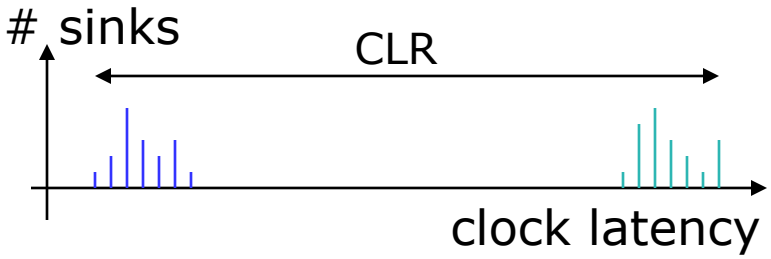

What's New for 2010?

- **A local clock skew limit**
- **To minimize total clock capacitance**
- **Much more clock sinks than 2009**
  - From 981 to 2249 (compare to 91 - 330 in 2009)
  - New ngspice (version 20) is much faster
- **Variations on inverter supply voltage and wire width**
- **Benchmarks from real IBM and Intel microprocessor designs**

Some results from 2010 CNS contest

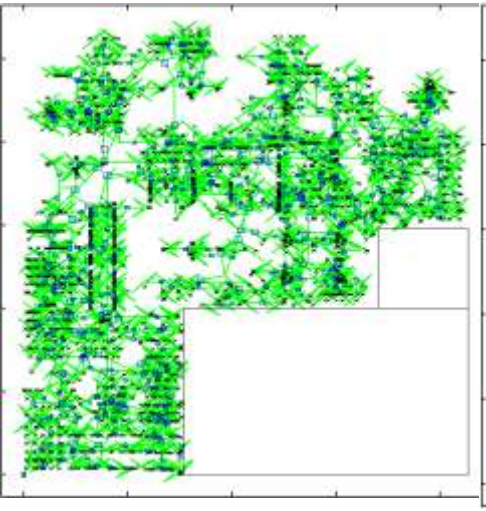

08

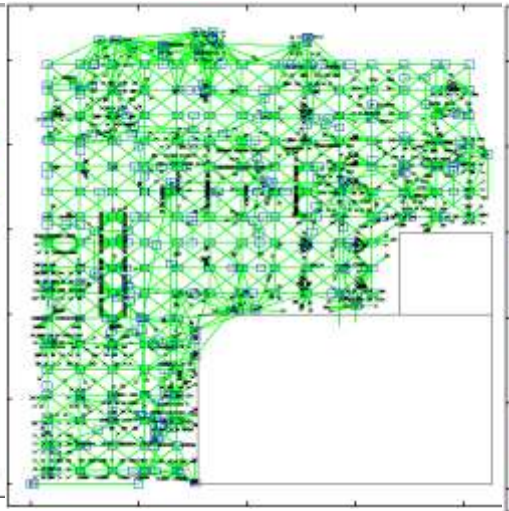

09

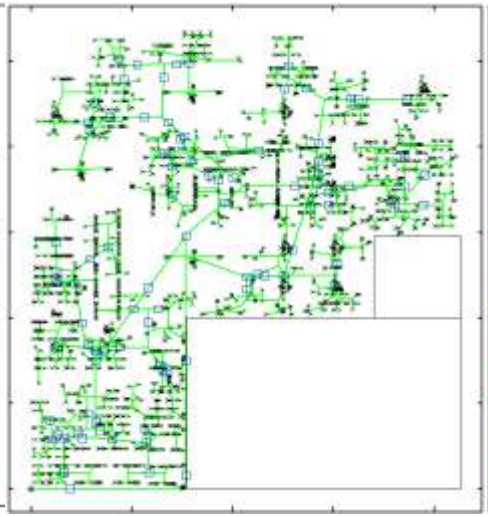

15

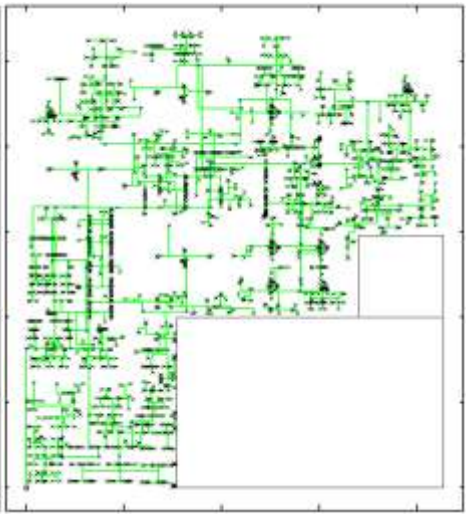

20

Local Clock Skew

	cpu/s	mean	min	med	95th	max	cap	nom	sink-c	inv-c	wire-c	rank
08	58	7.45	4.85	7.36	9.55	11.36	325206	4.58	12346	181649	131211	3
09	32	2.83	1.65	2.71	3.98	5.74	277151	3.90	12346	177215	87590	2
15	6075	3.26	2.04	3.16	4.46	6.06	71843	1.01	12346	40326	19170	1
20	3051	7.75	4.60	7.51	10.53	13.04	71035	1.00	12346	39790	18899	3

After All, How to Run an EDA Contest?

- **Identify the core problem in your own area**
- **Talk to other companies and professors**
- **Extract the key problem formulation and try to simplify other second-order factors**
- **Create an infrastructure**
  - Collaboration between Industry and Academia

EDA Research Problem Space

- **Why core problems are left untouched?**
  - Marketing
  - Infrastructures and realistic benchmarks

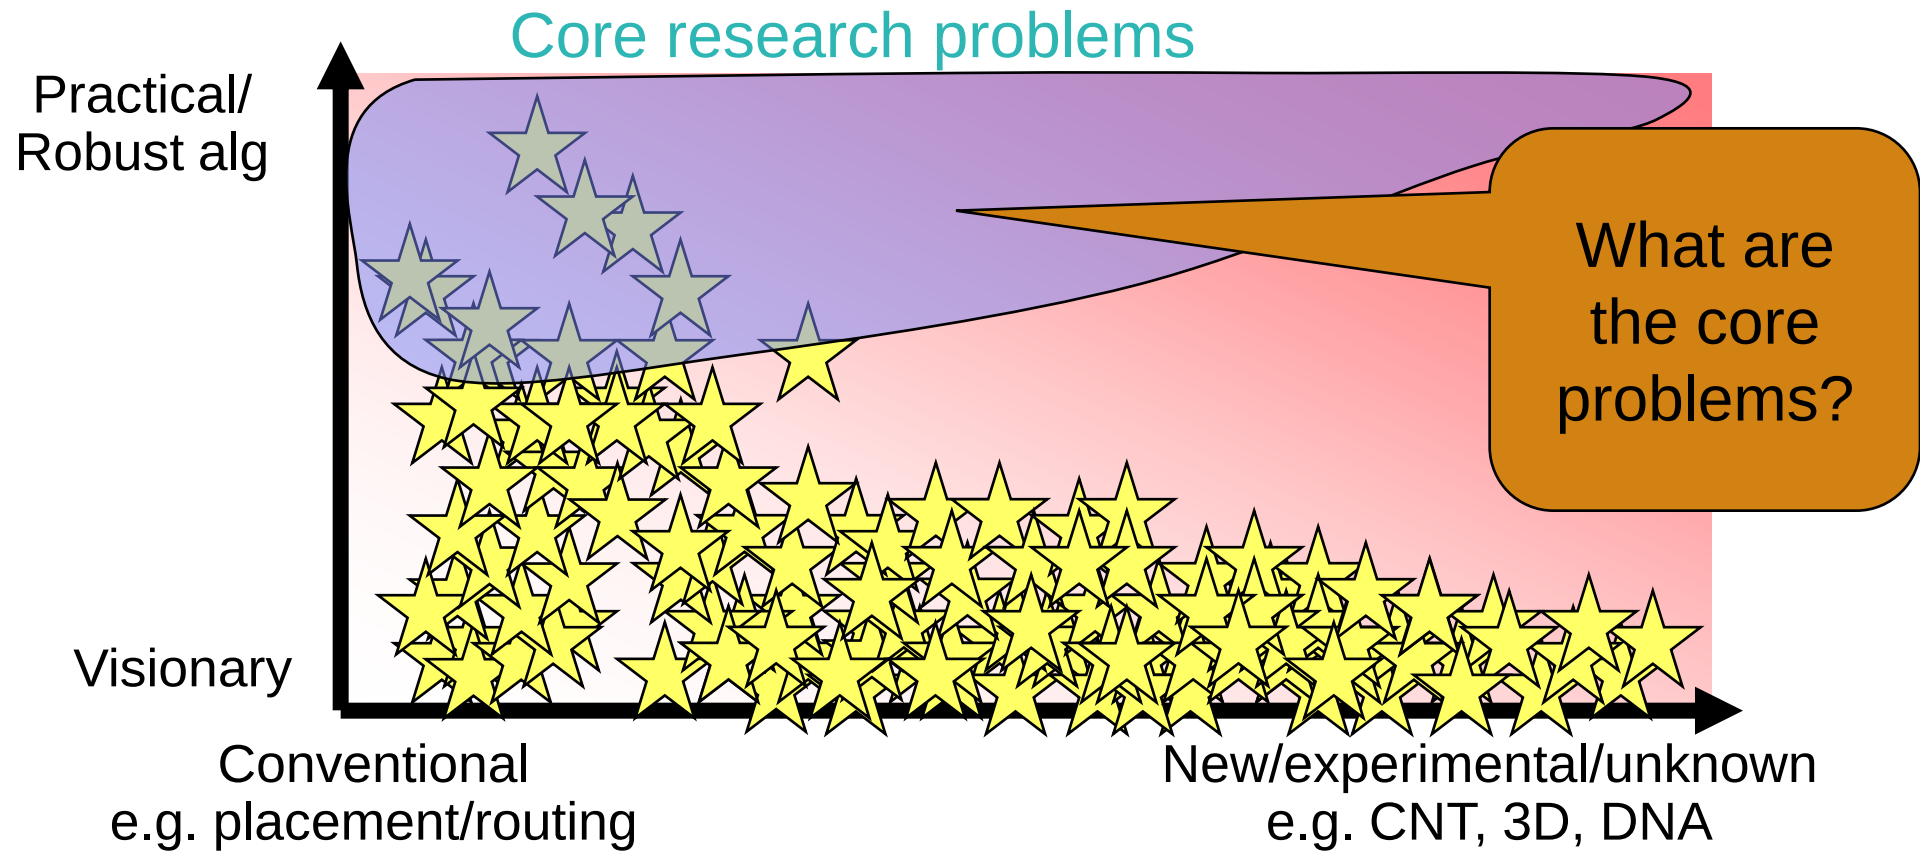

Trends and Opportunities

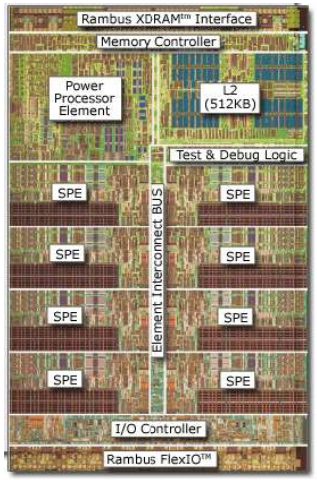

Heterogeneous cores

High performance ASIC

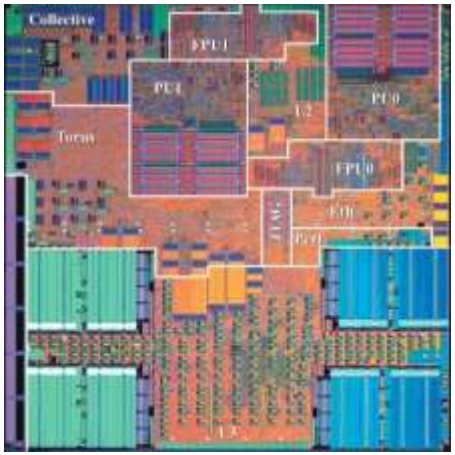

Microprocessor

ASIC

Full-custom

Full-automated

Bigger
Higher
Lower
Longer

Design Team
Performance
Power
Time-to-Market

Smaller
Lower
Higher
Shorter

Very expensive!

COST

A lot cheaper!

Physical Synthesis for Timing Closure

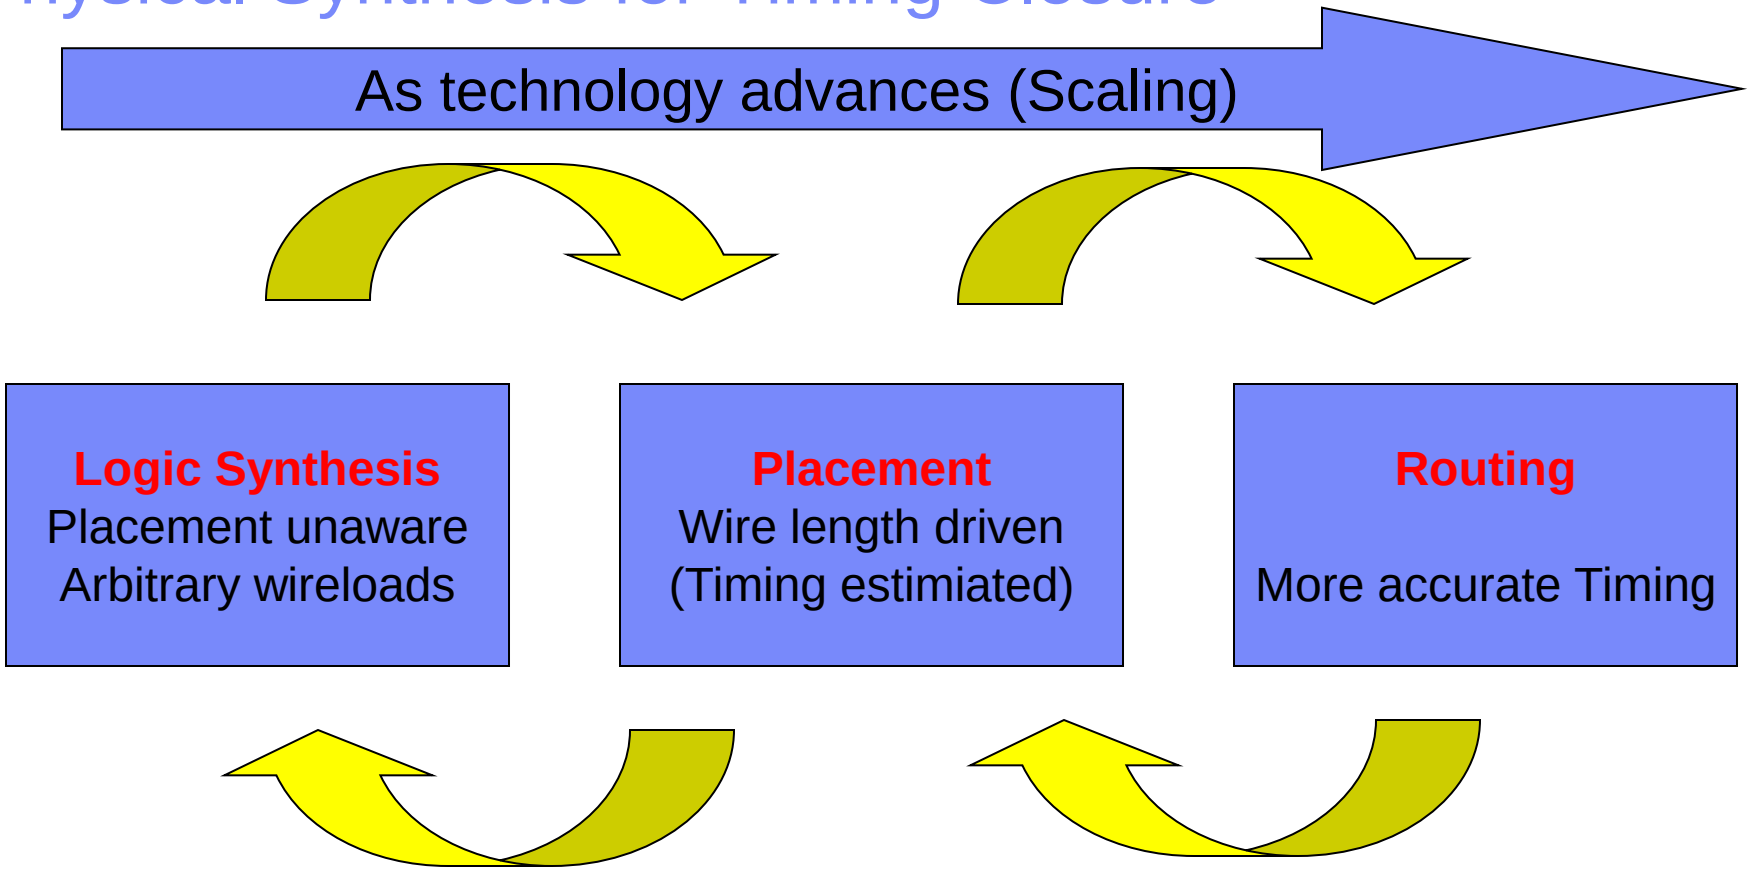

Convergence???

It is a bigger problem than you see.

Are you still talking about Scaling?

- **“Scaling” is so 90’s**
  - It is coming to an end and now we have multi-cores
- **Core problems have not been solved!!!**
- **Scaling kills global interconnects**
  - Hold on, it is also killing local interconnects

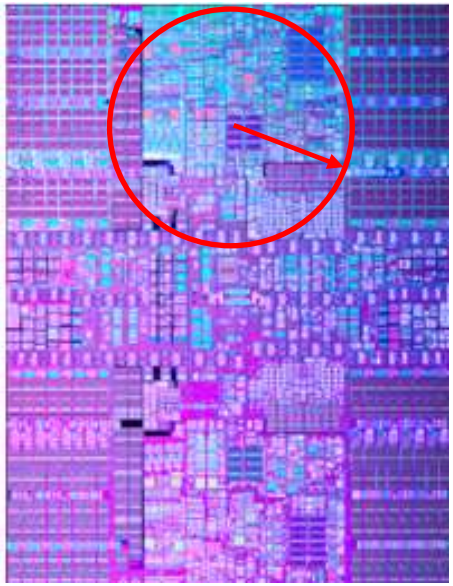

P6, ~core cycle reach

65nm, ~5.2 GHz

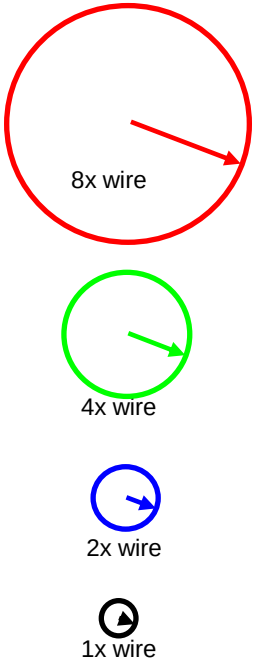

Puri, ISPD 09

The Wire Stack

Physical Synthesis for ~~Timing~~ Congestion Closure

Placement & Optimization
Wire length driven
(Timing estimiated)
(Know nothing about routes)

Congestion Evaluator

Routing
More accurate Timing
However,
detour mess up prediction

Congestion Evaluator

- Guide the placer and only spread out congested area
- Guide optimization (e.g. buffering) for timing closure
- Guide Timing Analysis for more consistent timing

Probabilistic

Route Fast

No Detour

Additional Detour

Completely Detour

Over Blockages

Around Blockages

Obey layer constraints

May violate layer constraints

The key infrastructure

...if we want to have a

congestion driven

place and route

contest

... is not a Router

- **Complete routing with as few overflow as possible**

Route Carefully

Completely Detour

Around Blockages

May violate layer constraints

What's for ISPD Contests 2011 and Beyond?

- **We haven't decided yet**
- **Core problems in the Physical Design Area**
  - Congestion driven / timing driven placement
    - Congestion estimator
    - Timing analysis tools
    - Gate timing model
  - Gate sizing/buffering/physical synthesis
    - Netlist information
    - Gate timing model
  - Detailed Routing or Detailed Placement
    - Rules
  - DFM, inverse lithography
    - Lithographic simulation infrastructure
    - Source-mask model
  - Etc…

EDA Research is Going to

- **Focus on core problems**
- **Be close collaboration between industry and academia**
  - EDA contests is just one way to help