



Charles Thangaraj




Early design phase design trade-offs using in-situ
macro models




Ph.D. preliminary examination




July 7th 2008

Committee Members

-  Dr. Tom Chen
 -  Electrical and Computer Engineering
 -  Advisor

-  Dr. Anthony Maciejewski
 -  Electrical and Computer Engineering
 -  Member

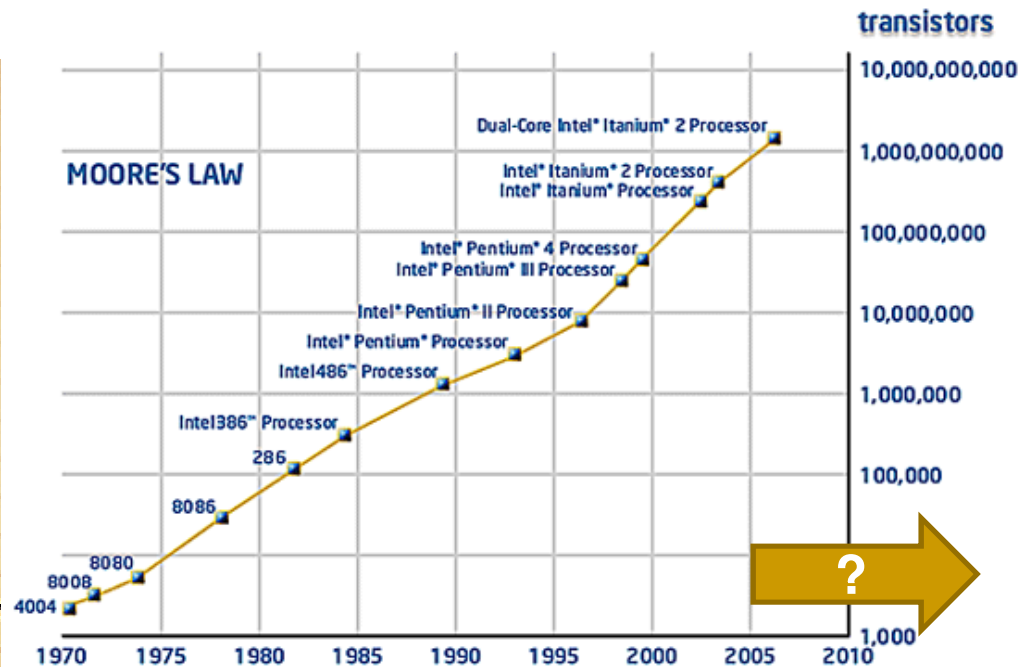

-  Dr. George Collins
 -  Electrical and Computer Engineering
 -  Member

-  Dr. Phillip Chapman
 -  Statistics
 -  Outside Member

Presentation Outline

- Background information
 - CMOS technology scaling and high-performance VLSI design
- Industry challenges
 - Business
 - Technology
 - Design automation
- Problem Statement
 - Motivating factors and objective of this work
- Existing research
 - An overview of what is out there
- Proposed Approach
 - Modeling methodology
 - In-situ macro model generation
- Experimental Validation
 - Experimental setup, assumptions and analysis
 - Software architecture
 - Results discussions
- Future work


Moore's Law - 1965




- No. of devices and operating frequency double every 18 months
 - faster FETs, higher device density, higher performance, lower cost
 - Will this trend continue?
- Moore's law is expected to continue but with **challenges** in
 - Process technology
 - EDA tools
 - Business related

Examples...

Microprocessor trend



Intel® Itanium® 2 processor
Introduced 2002
Initial clock speed
1 GHz
Number of transistors
220,000,000
Manufacturing technology
0.13μ



Intel® Pentium® D processor
Introduced 2005
Initial clock speed
3.2 GHz
Number of transistors
291,000,000
Manufacturing technology
65nm



Intel® Core™2 Duo processor
Intel® Core™2 Extreme processor
Dual-Core Intel® Xeon® processor
Introduced 2006
Initial clock speed
2.93 GHz
Number of transistors
291,000,000
Manufacturing technology
65nm



Dual-Core Intel® Itanium® 2 processor 9000 series
Introduced 2006
Initial clock speed
1.66 GHz
Number of transistors
1,720,000,000
Manufacturing technology
90nm



Quad-Core Intel® Xeon® processor
Quad-Core Intel® Core™2 Extreme processor
Introduced 2006
Intel® Core™2 Quad processors
Introduced 2007
Initial clock speed
2.66 GHz
Number of transistors
582,000,000
Manufacturing technology
65nm



Quad-Core Intel® Xeon® processor (Pentium)
Dual-Core Intel® Xeon® processor (Pentium)
Quad-Core Intel® Core™2 Extreme processor (Pentium)
Introduced 2007
Initial clock speed
> 3 GHz
Number of transistors
820,000,000
Manufacturing technology
45nm

Consumer electronics trend



Moore's law is expected to continue but with **challenges** in

- Process technology
- EDA tools
- Business related

Images from Intel and Motorola

Presentation Outline

- Background information
 - CMOS technology scaling and high-performance VLSI design
- Industry challenges
 - Business
 - Technology
 - Design automation
- Problem Statement
 - Motivating factors and objective of this work
- Existing research
 - An overview of what is out there
- Proposed Approach
 - Modeling methodology
 - In-situ macro model generation
- Experimental Validation
 - Experimental setup, assumptions and analysis
 - Software architecture
 - Results discussions
- Future work

Semiconductor Industry Trends & Challenges: Business

Business trends

- 10X revenue growth in the past decade
- Growth to US \$700 B from US \$200 B in the next decade
- Newer & emerging markets are opening up
 - ◆ Auto electronics, medical, GPS, intelligent appliances etc.

Business challenges






- Capital overhead - exponentially increasing fab costs
- R&D overhead - process development expenditure increasing
- Design cost reduction to max profit - min NRE (non recurring engineering expenditure)
- Competition – price point, product quality and time-to-market schedule pressures
- Opportunity cost – need to maintain market segment share in high volume low cost segment
- Need to innovate to stay in business

Engineering decisions are being increasingly influenced by business needs !!



- Optimize designs for specific market segments i.e. design target trade-offs
 - ◆ Cost/area–performance–power–reliability–yield

Semiconductor Industry Trends & Challenges: Tech & EDA

Technology challenges influencing design convergence

-  Vdd – Vt gap closes -- leakage power, SRAM stability, computation reliability, noise etc
-  Increase in transistor density -- thermal instabilities, yield issues
-  Interconnect bottleneck
-  Manufacturing process variation (Vt, L, tox) – design guard banding
-  Sub-wavelength lithography
 - ◆ DFM rules explosion, OPC, PSM, OAI, I-litho, EUV-litho, SRAF

EDA tool challenges influencing design optimization

-  Existing EDA tools
 - ◆ Aimed at speeding conventional design process - obsolete methodologies
 - ◆ accurate, very specific and detailed -- less flexibility
-  Limited use of system level design modeling & optimization
 - ◆ Concoct such a tool with existing tools
 - ◆ computationally expensive
 - ◆ very limited opportunity for design space exploration
 - ◆ Interoperability overhead, time consuming approach


A design team should

-  Meet the business need, overcome challenges and achieve design convergence on time !!

Presentation Outline

- Background information
 - CMOS technology scaling and high-performance VLSI design
- Industry challenges
 - Business
 - Technology
 - Design automation
- **Problem Statement**
 - **Motivating factors and objective of this work**
- Existing research
 - An overview of what is out there
- Proposed Approach
 - Modeling methodology
 - In-situ macro model generation
- Experimental Validation
 - Experimental setup, assumptions and analysis
 - Software architecture
 - Results discussions
- Future work

Competing Design Goals



Power $\downarrow \rightarrow$ supply $\downarrow \rightarrow V_t \uparrow \rightarrow$ FET size \downarrow

Deploy low power design techniques

Performance $\uparrow \rightarrow$ supply $\uparrow \rightarrow V_t \downarrow \rightarrow$ FET size \uparrow

Deploy high performance design techniques

- Nano-meter CMOS

- Leaky FETs - power consumption
- Poor critical dimension control – L variation
- Poor BEOL dimension control – wire R & C variation
- Exponential mask costs – fewer re-spins

- **Get it right the first time** – key to success of modern high performance designs

- Correct early stage design plan


- Early design phase, design target tradeoff analysis considering low level implementation details

- Improves design convergence.
- Guarantees time-to-market
- Helps in avoiding costly redesigns in later design stages

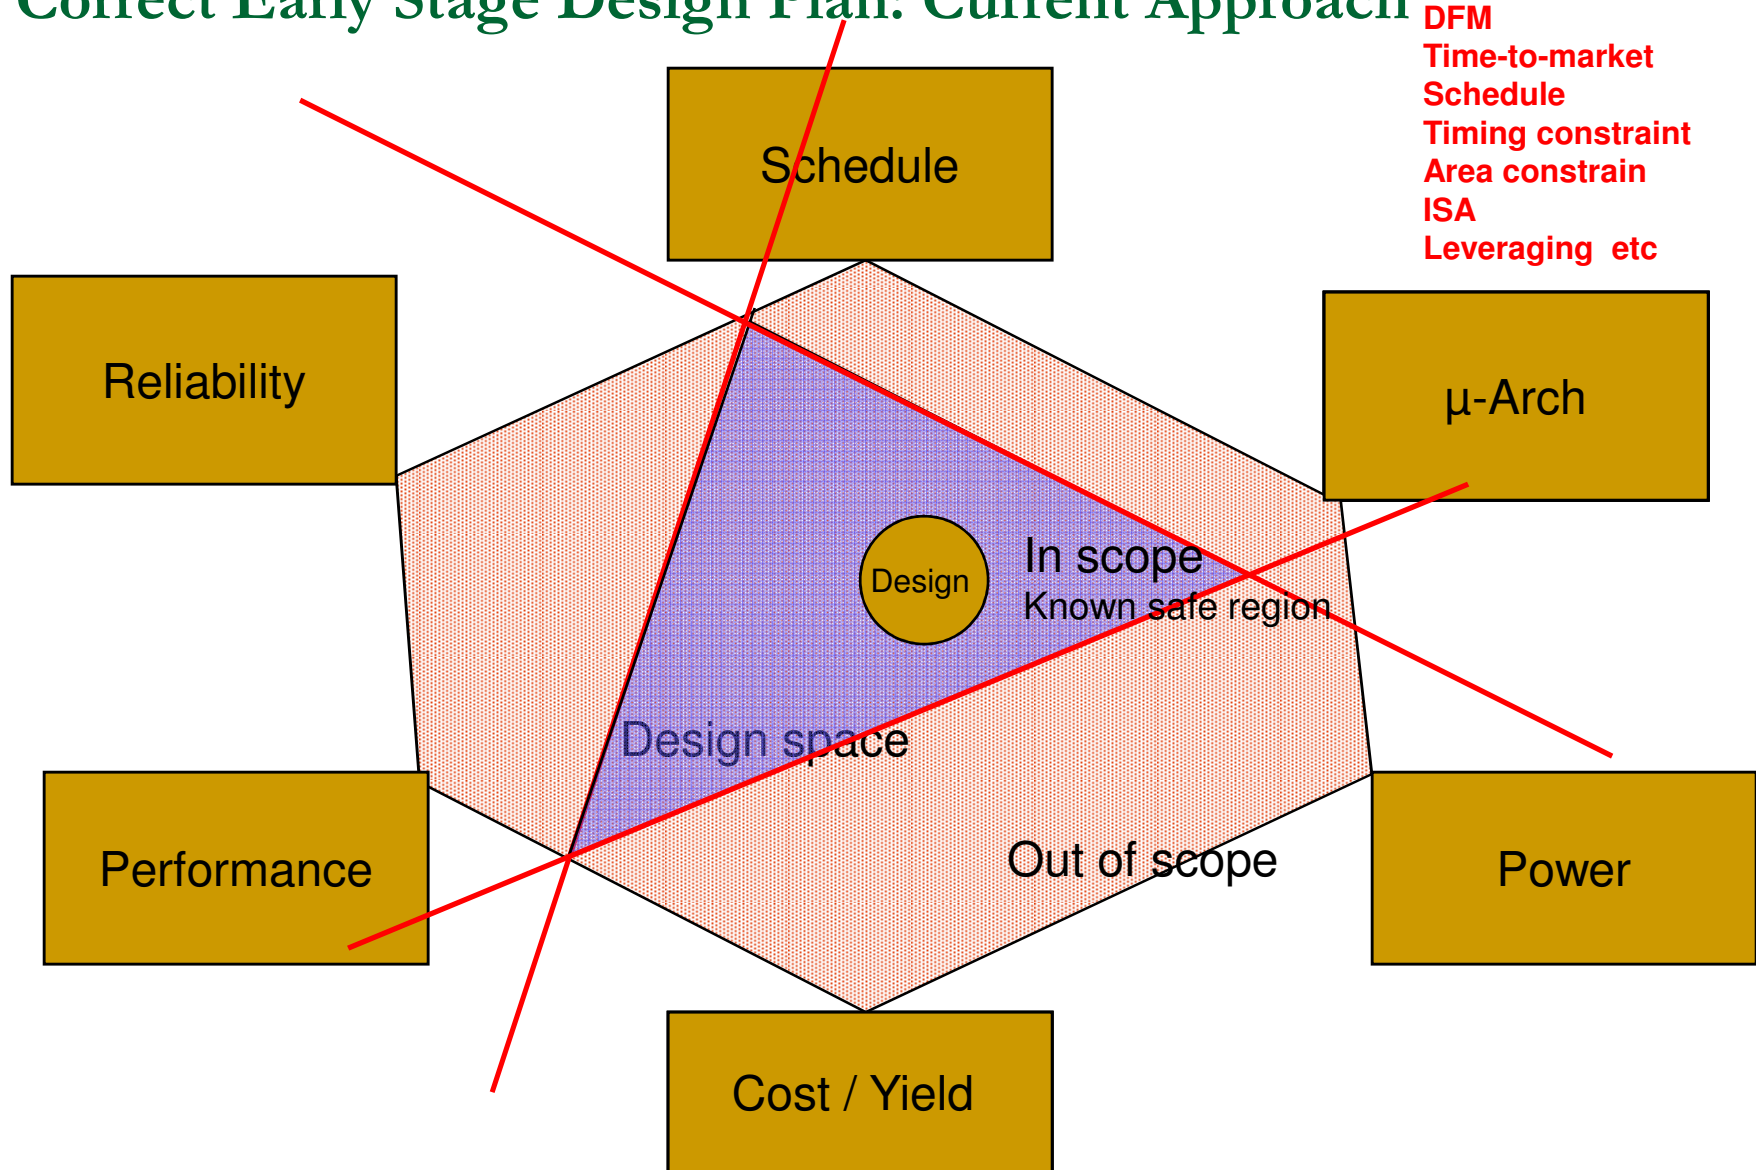
Correct Early Stage Design Plan

System level optimization considering
low level implementation details

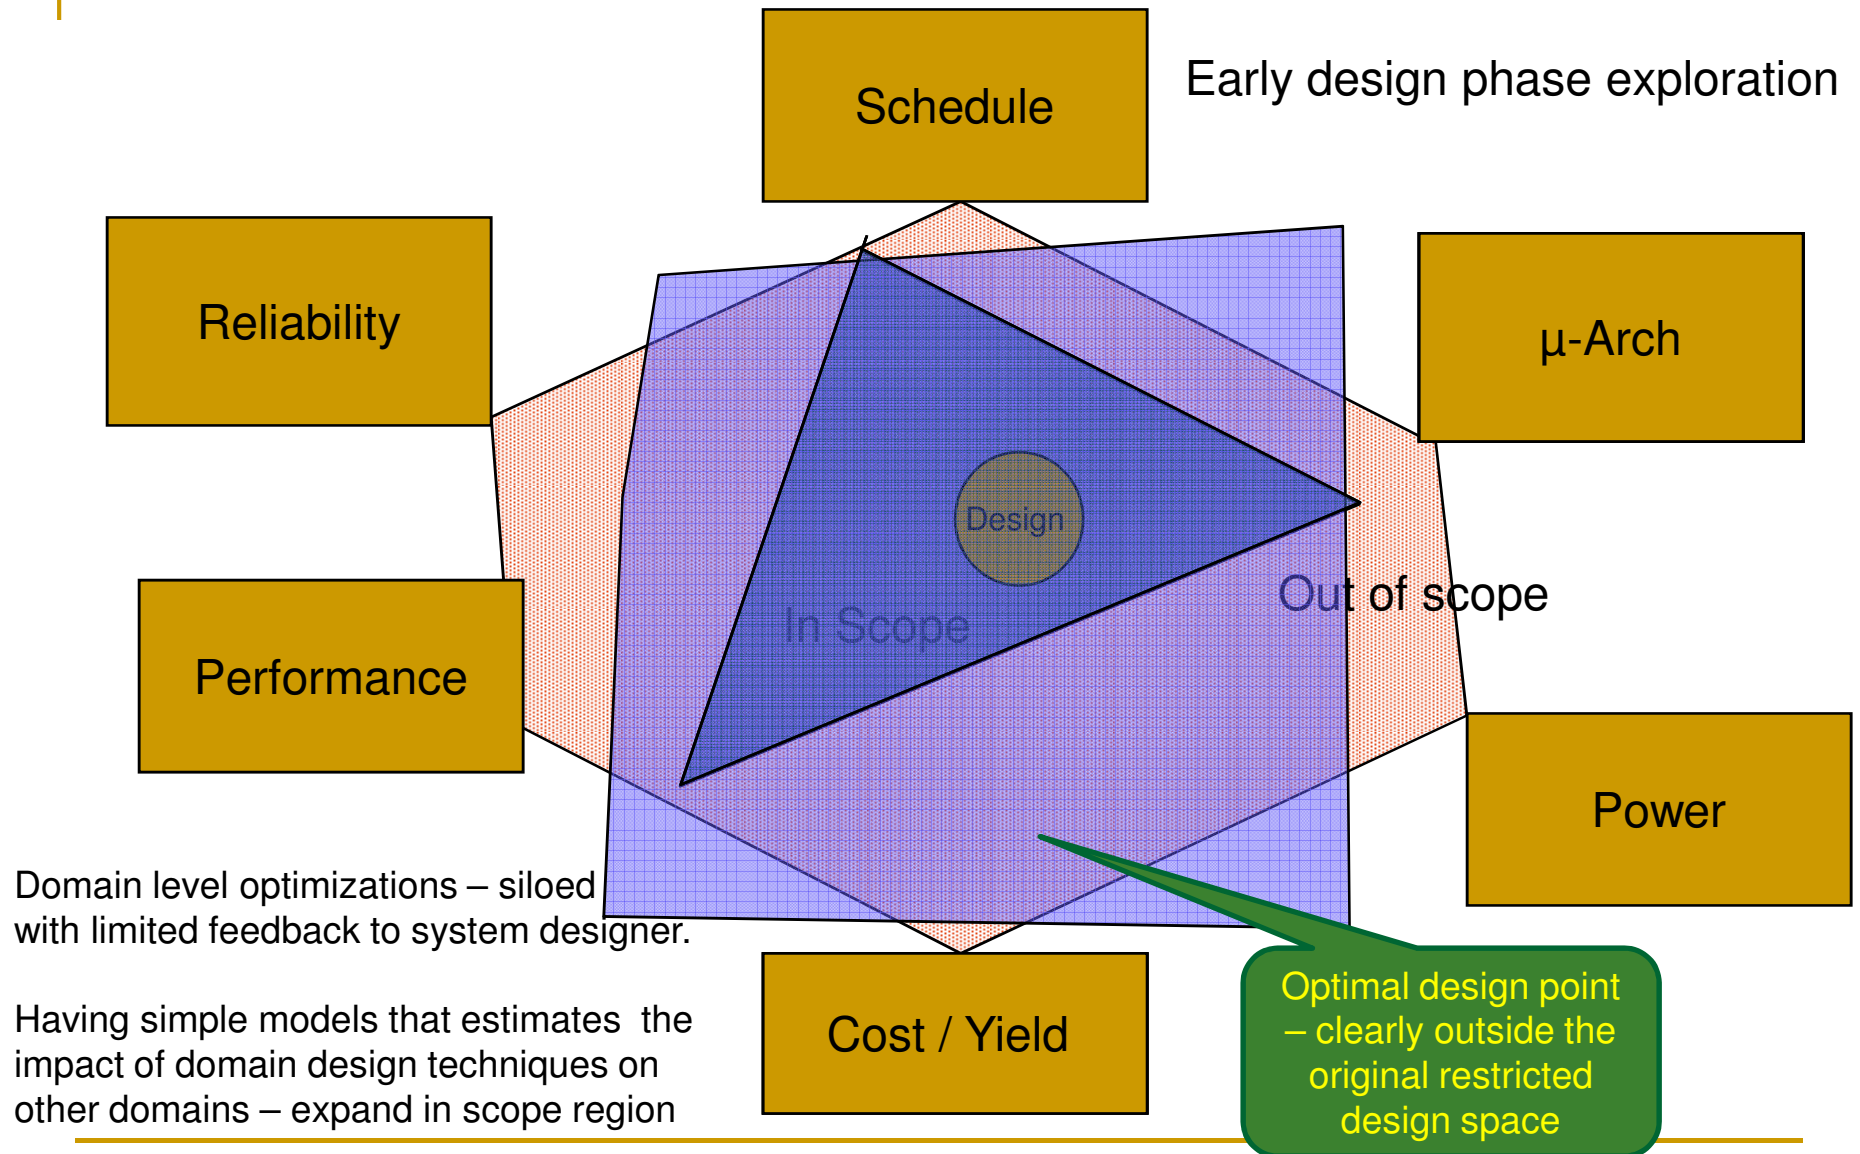
System level optimization NOT considering
low level implementation details



Correct Early Stage Design Plan: Current Approach



Correct Early Stage Design Plan: A Better Approach



Problem Statement And Research Objective

- To develop a modular system level modeling methodology
 - ▣ Simple yet effective for trade-off analysis and design optimization
 - ▣ Considers low level implementation details
 - ▣ Able to handle large designs i.e. scalable

- To develop analytical design target prediction models and module descriptors with inter-domain (design target) impact estimation capability
 - ▣ Estimate system dynamic and leakage power
 - ▣ Estimate system performance
 - ▣ Estimate system reliability
 - ▣ Estimate die-size and yield

- Software development to implement the methodology and tool flow
 - ▣ Enable **quick** design evaluation and design space exploration
 - ▣ Utilizing macro-model generation and in-situ SPICE simulations

- Experimentally demonstrate & validate the proposed methodology

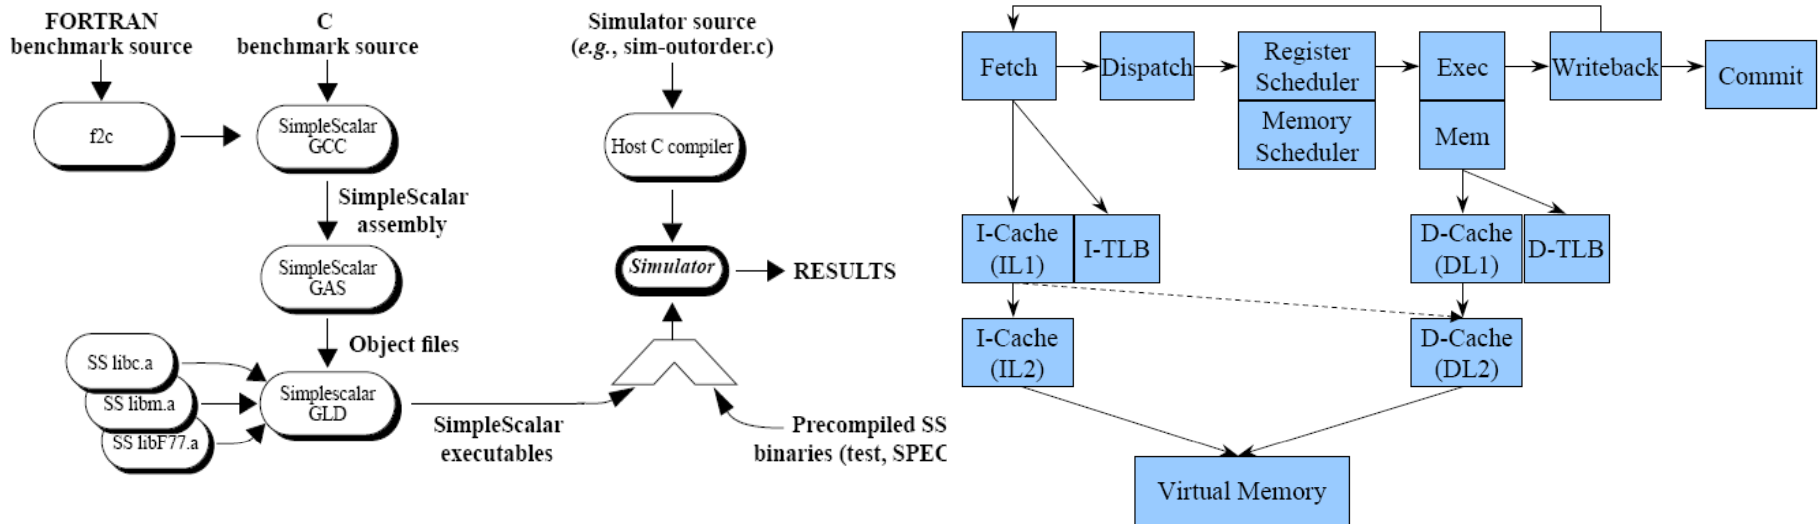
Presentation Outline

- Background information
 - CMOS technology scaling and high-performance VLSI design
- Industry challenges
 - Business
 - Technology
 - Design automation
- Problem Statement
 - Motivating factors and objective of this work
- Existing research
 - An overview of what is out there
- Proposed Approach
 - Modeling methodology
 - In-situ macro model generation
- Experimental Validation
 - Experimental setup, assumptions and analysis
 - Software architecture
 - Results discussions
- Future work

Existing Tools For System Level Design Optimization

- Two broad categories
 - Computation engine focused
 - * SimplePower, SimpleScalar, Wattch, AccuPower and PowerTimer
 - ASIC focused
 - * BACPAC – most widely know and other developmental in-house tools
- Computation engine focused tools, model an underlying computation architecture and emulate code execution on the modeled architecture.
- They collect activity rates, instruction execution rates, miss rates, and prediction efficiency to estimate cycles per instruction (CPI) and other performance metrics
- ASIC focused tools, estimate power and performance based on low level physical design parameters and key process parameters.
- BACPAC which is based on low-level physical design parameters attempts to “re-create” the design in a bottom-up manner for analysis.

Computation Engine Focused Tools - SimpleScalar



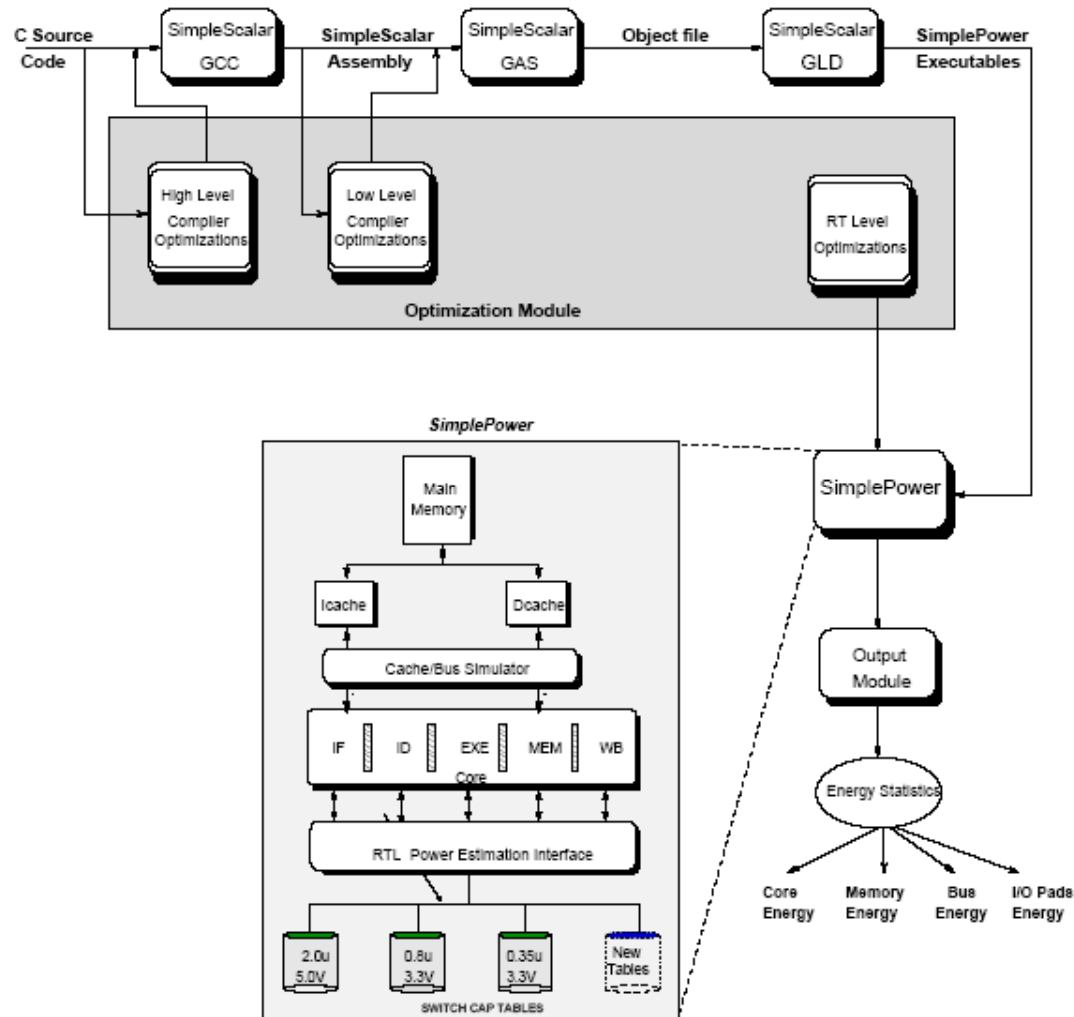
Tool block diagram.

Pipeline model

5 stage OOO pipeline with microarchitectural blocks modeled

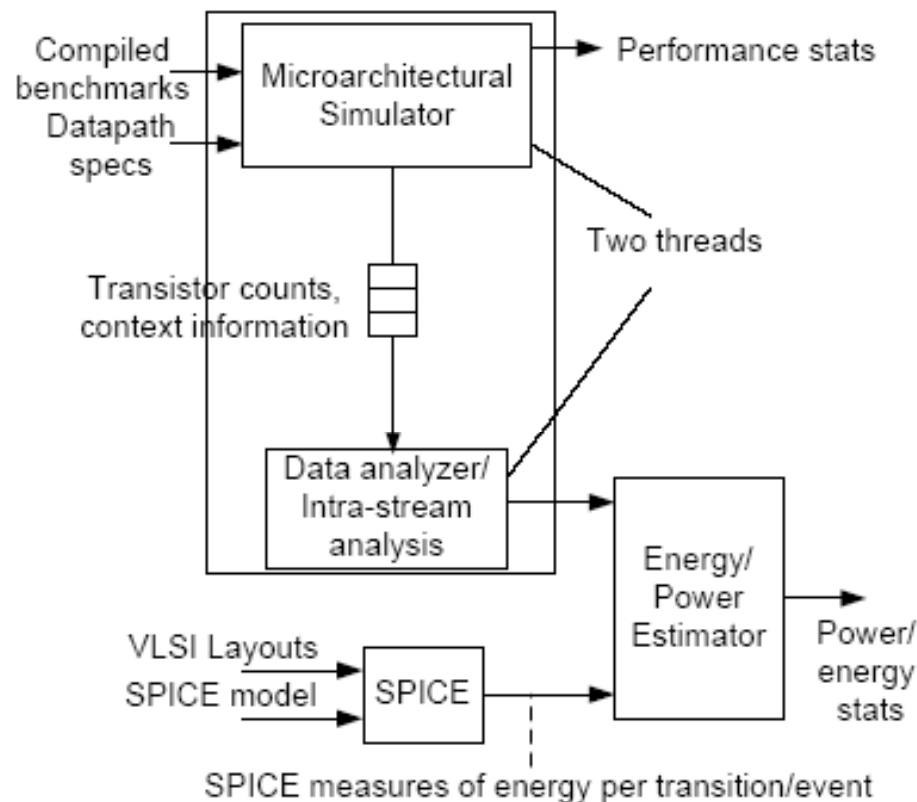
1. SimpleScalar is a microarchitectural simulator (microprocessor emulator)
2. Pure computation architecture performance and optimization tool
3. No power estimation capability

Computation Engine Focused Tools - SimplePower



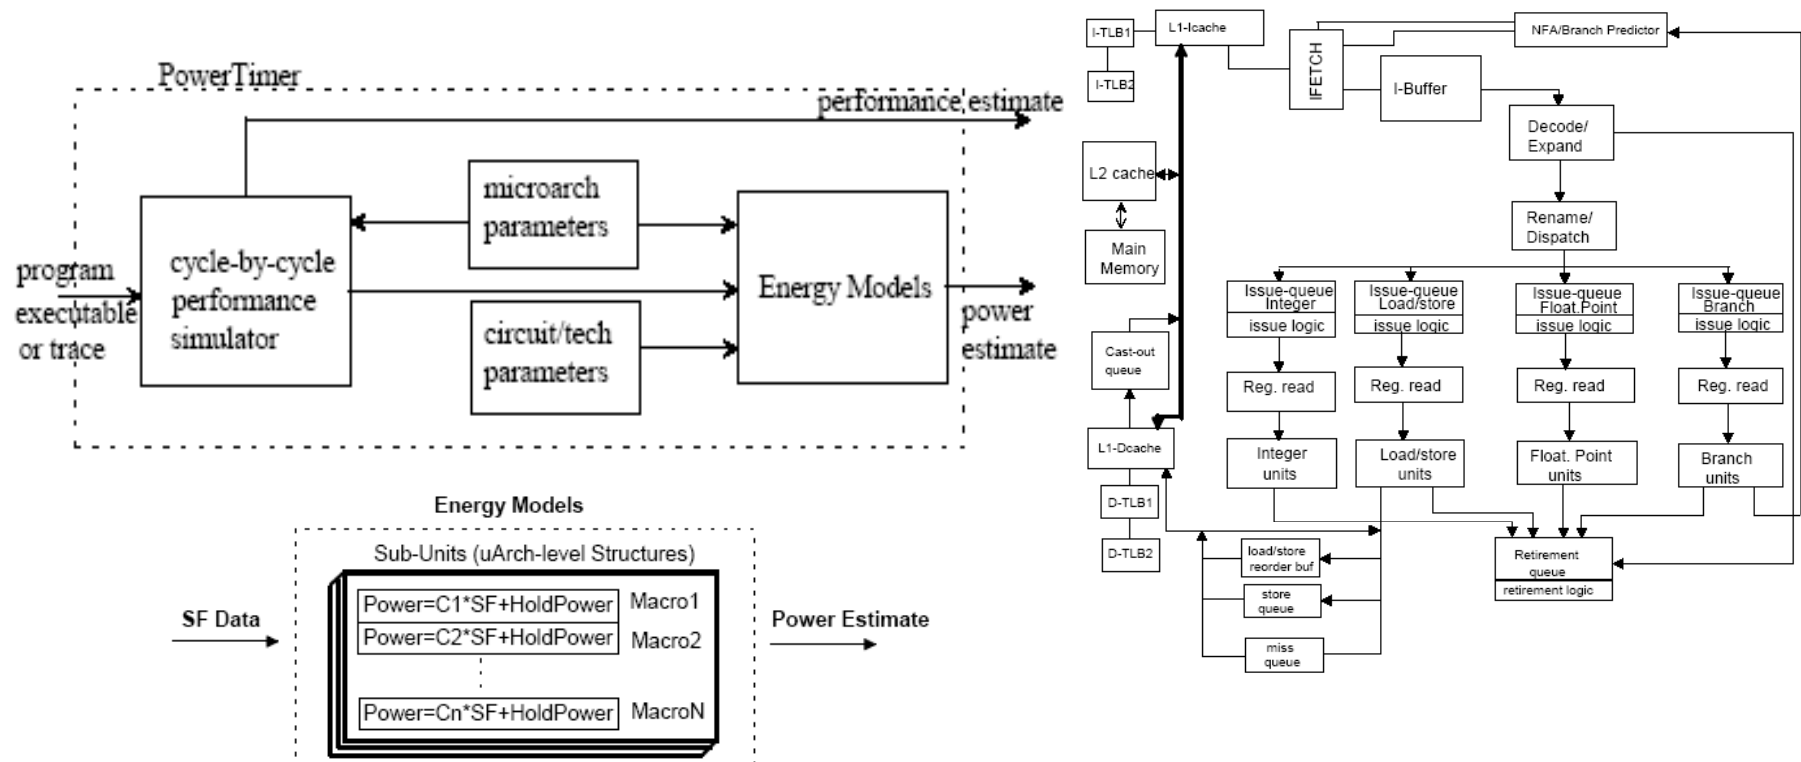
1. SimplePower is a SimpleScalar addon
2. Snoops microarchitectural block activation while executing code.
3. Power is calculated using per-transition energy tables for each microarchitectural block and the activation information.
4. Similarly bus activation is snooped to calculate power.

Computation Engine Focused Tools - AccuPower



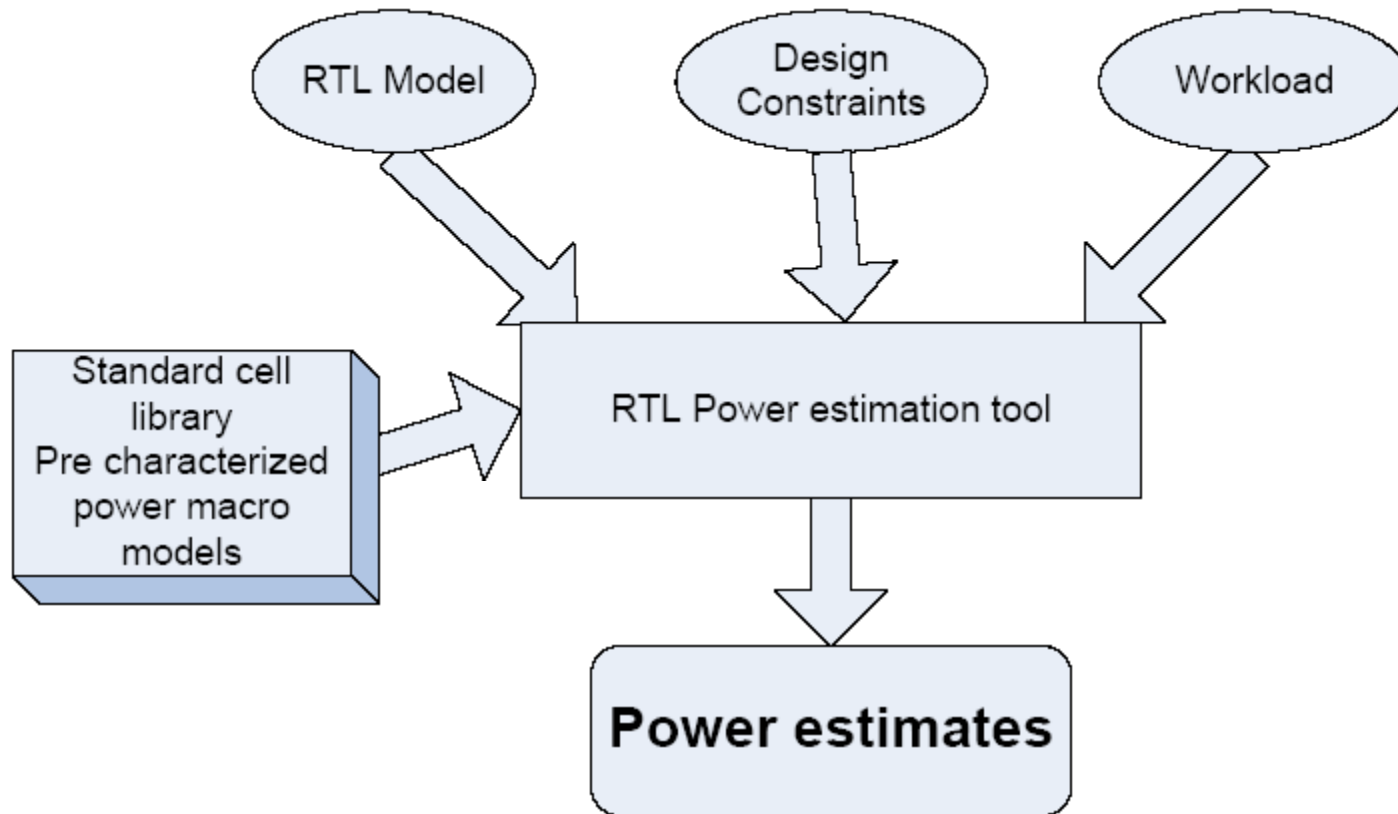
1. Based on SimpleScalar toolset.
2. Similar to SimplePower, but per transition/event power is obtained from detailed SPICE simulation.

Computation Engine Focused Tools - Wattch & Powertimer



1. More detailed pipeline model is used. – more microarchitectural blocks modeled.
2. Power models are built bottom-up using detailed simulation results in a hierarchical manner.
3. Leakage power is modeled in this tool.

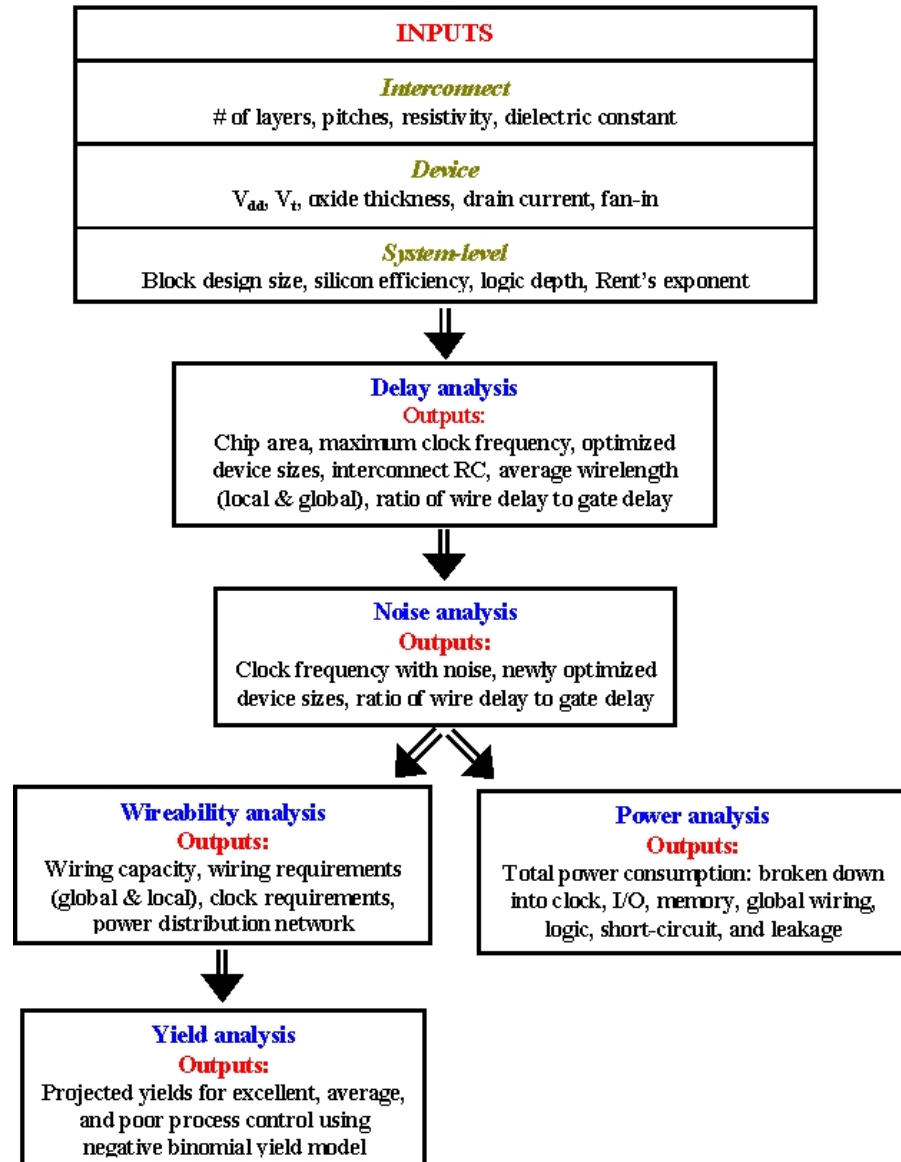
Computation Engine Focused Tools – Industrial Tools



1. Computation emulation is done at the RTL level – complete RTL model available.
2. Very accurate power estimates using detailed simulations

ASIC focused BACPAC

1. Uses a set of empirical analytical models.
2. Model parameters are physical design parameters.
3. Uses a customizable critical path model
4. Can estimate circuit performance based on the critical path delay.
5. Leakage power is estimated from total device width estimate
6. Switching cap estimates are used to calculate dynamic power for wires, devices, clock, pads and memory.



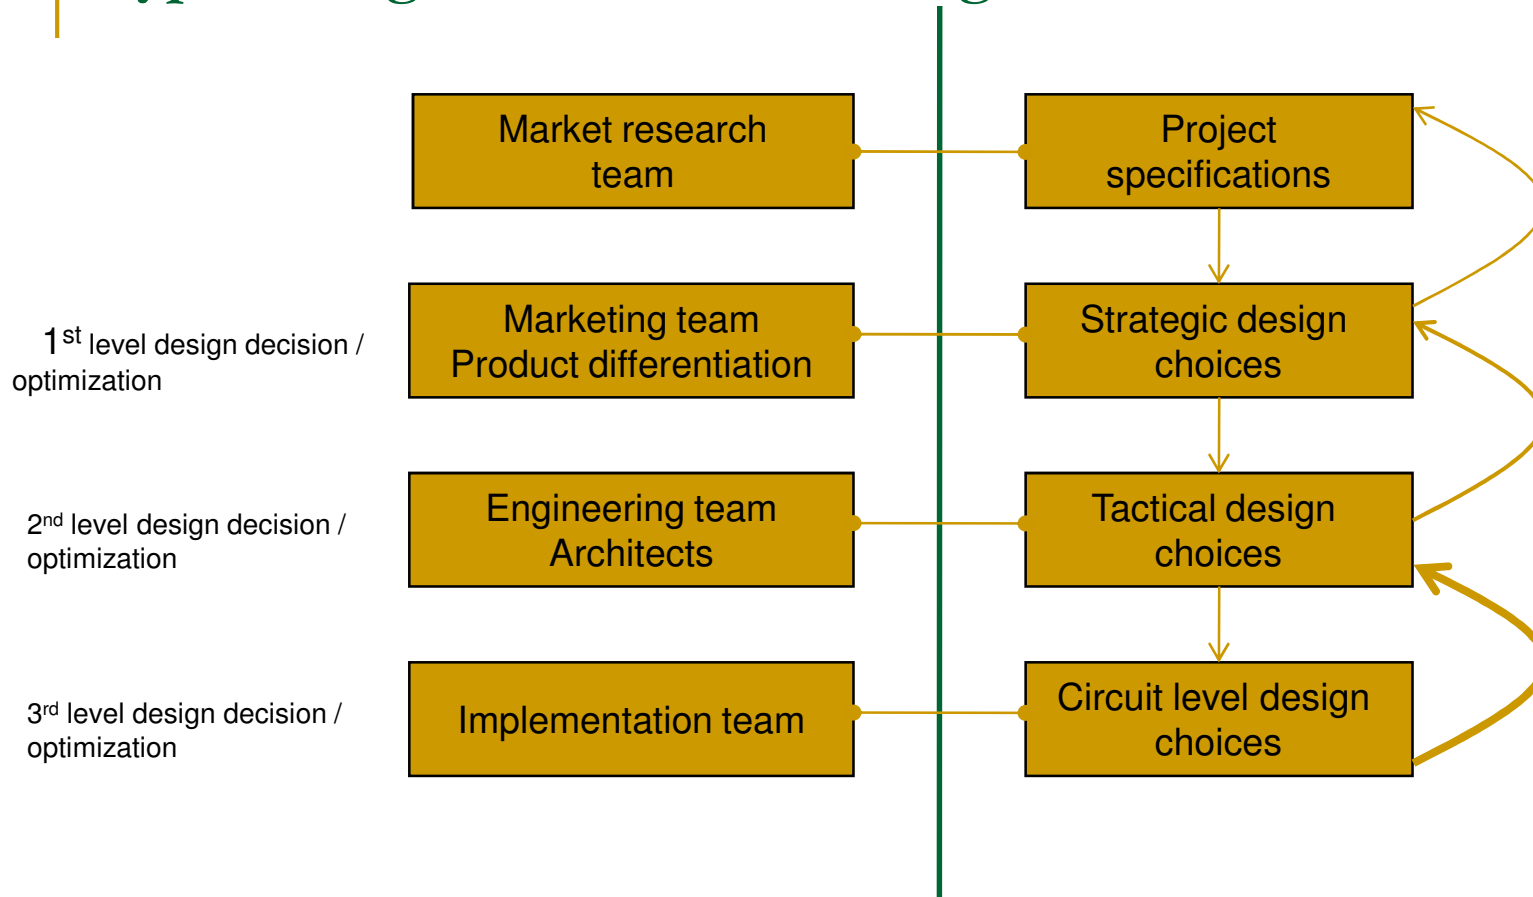
Existing Tools – Shortcomings

- Tools that perform computational architecture optimization without physical implementation constraints do not offer any benefits in achieving design convergence in power, performance, area etc.
 - SimplePower uses pre-characterized power tables to estimate power.
 - ◆ unsuitable for design space exploration with changing low-level implementation details
 - Clock network power, which can account for up to 30% of the total dynamic power is not included in SimplePower.
 - AccuPower requires complete layout information – not available during early design phase
 - Wattch & PowerTimer have a constant “hold” power equation to estimate leakage power.
 - ◆ leakage power is increasing exponentially in nano-meter CMOS
 - PowerTimer is very slow
 - ◆ Power macros need to be re-characterized with every change in low-level implementation
- Optimization tools for physical implementation based on low-level physical design parameters are limited in their ability to explore the system design space.
 - The ASIC focused BACPAC’s equi-partition approach is not suitable for modeling highly modular designs such as microprocessors and other high performance designs.

Presentation Outline

- Background information
 - CMOS technology scaling and high-performance VLSI design
- Industry challenges
 - Business
 - Technology
 - Design automation
- Problem Statement
 - Motivating factors and objective of this work
- Existing research
 - An overview of what is out there
- Proposed Approach
 - Modeling methodology
 - In-situ macro model generation
- Experimental Validation
 - Experimental setup, assumptions and analysis
 - Software architecture
 - Results discussions
- Future work

Typical High Performance Design Flow



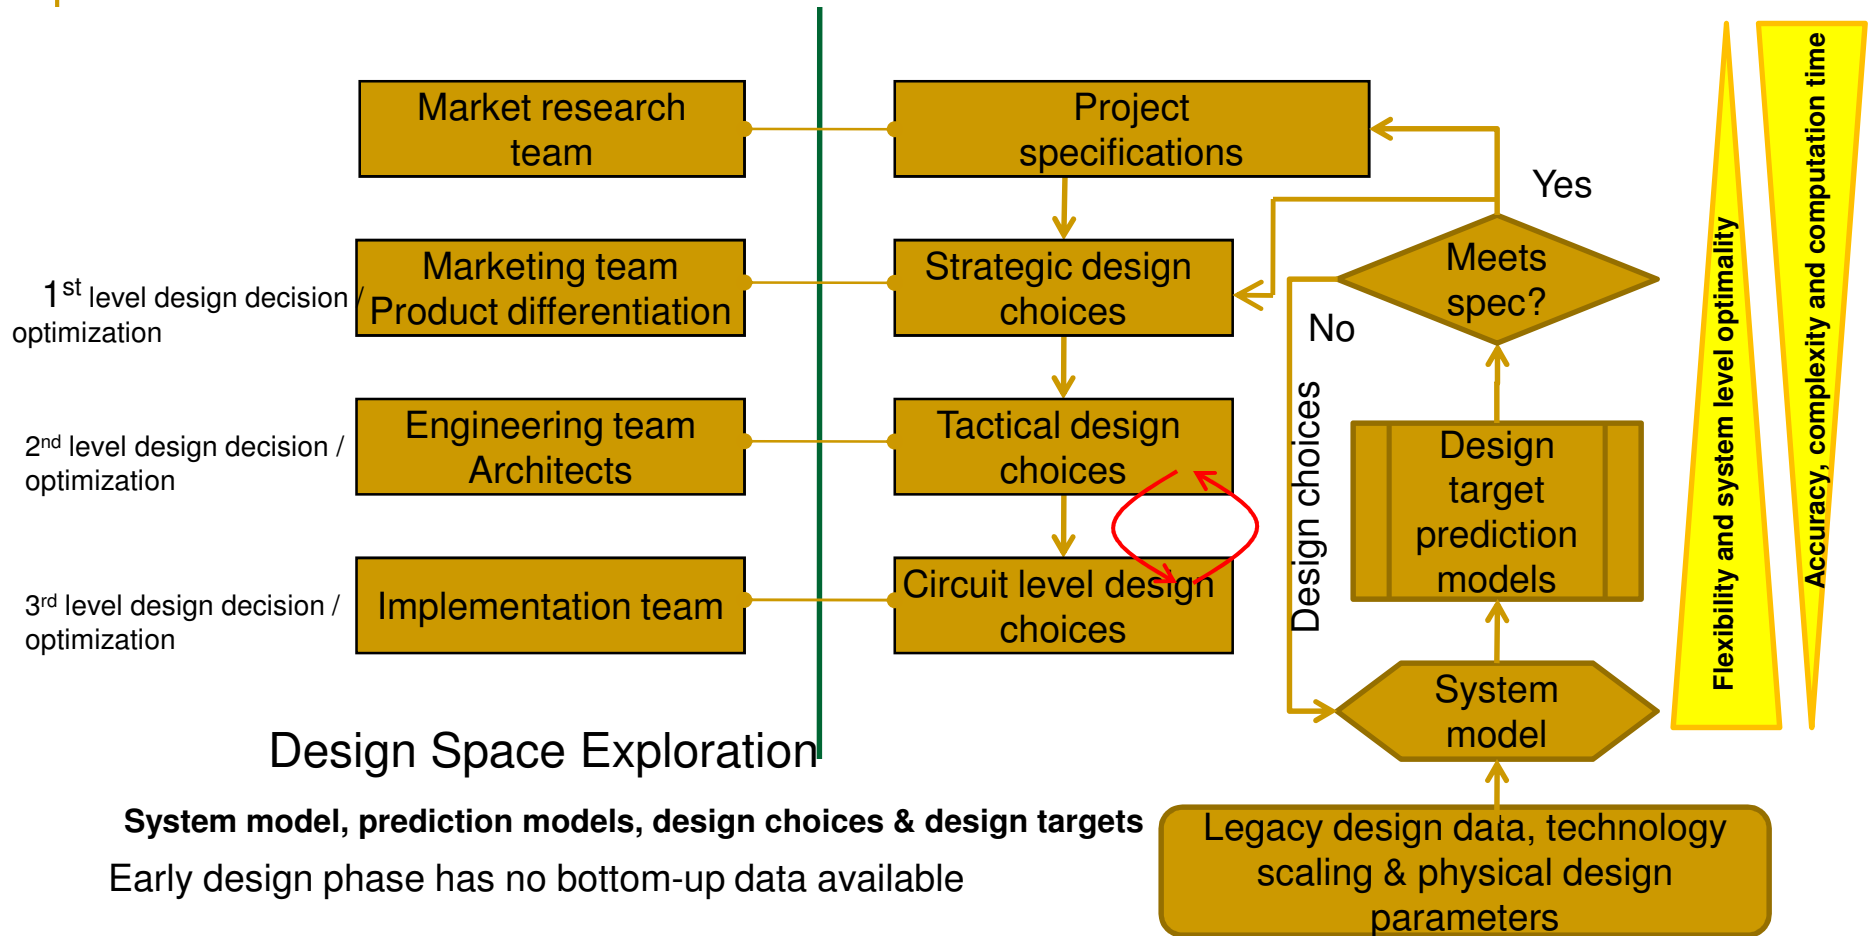
Early design phase has no low level implementation (bottom-up) data available

Early design decision are made without bottom-up data for validation

Increasing probability of a) circuit level design changes in the later design stages

b) product delivery delay or compromising product competitiveness

Proposed Design Flow



Design Space Exploration

System model, prediction models, design choices & design targets


Early design phase has no bottom-up data available

Perform design target exploration to set project specification


based on legacy design data, scaling and physical design data (abstracted low level implementation details)

Leads to **Correct Early Stage Design Plan**

System Modeling Methodology



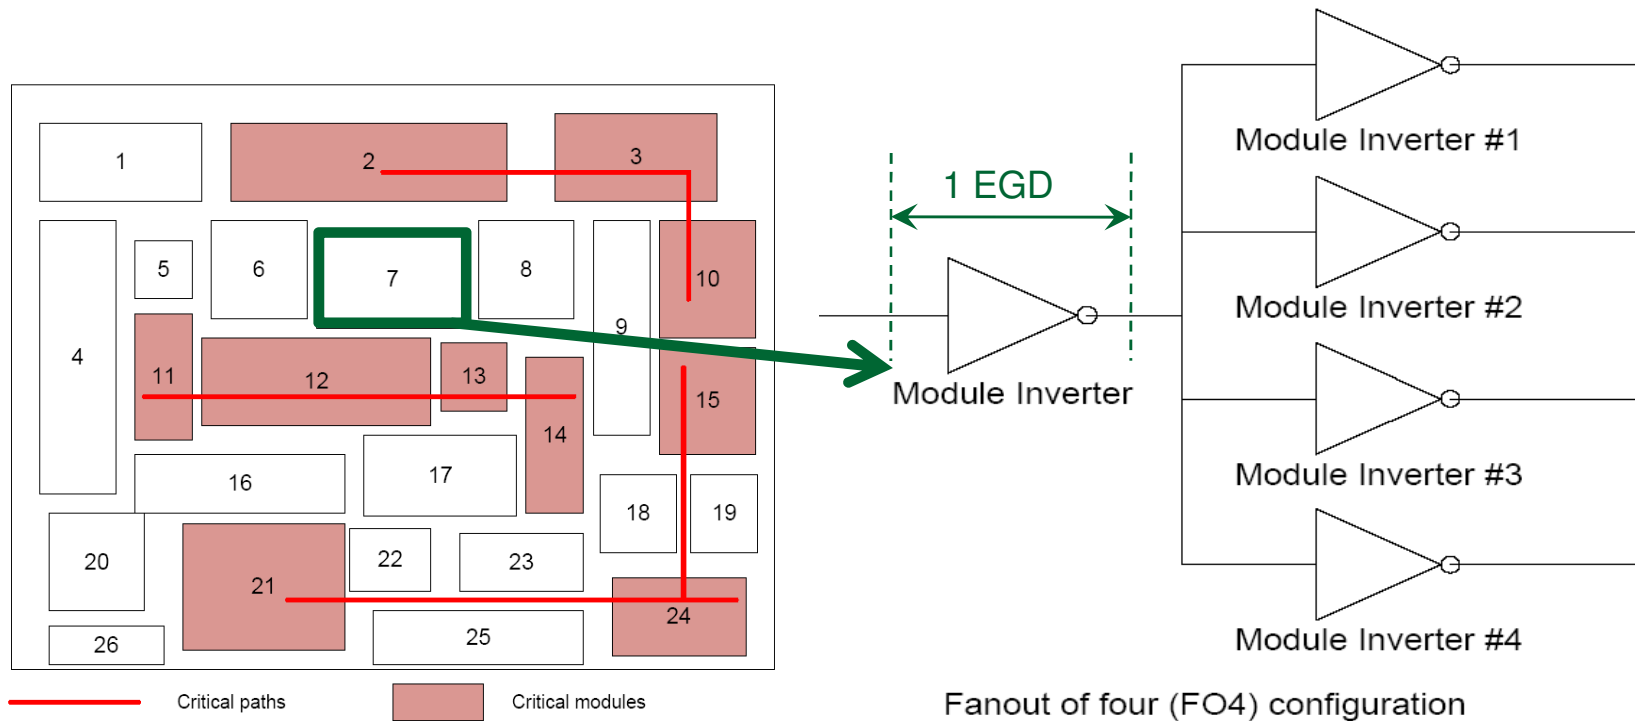
A generic system with 26 functional units or modules.



Die picture of an actual microprocessor implemented in a 100 nm process. Shown at 15th Int. Conf. on VLSI Design, 2002

Therefore a system is a collection on modules.

System Modeling Methodology

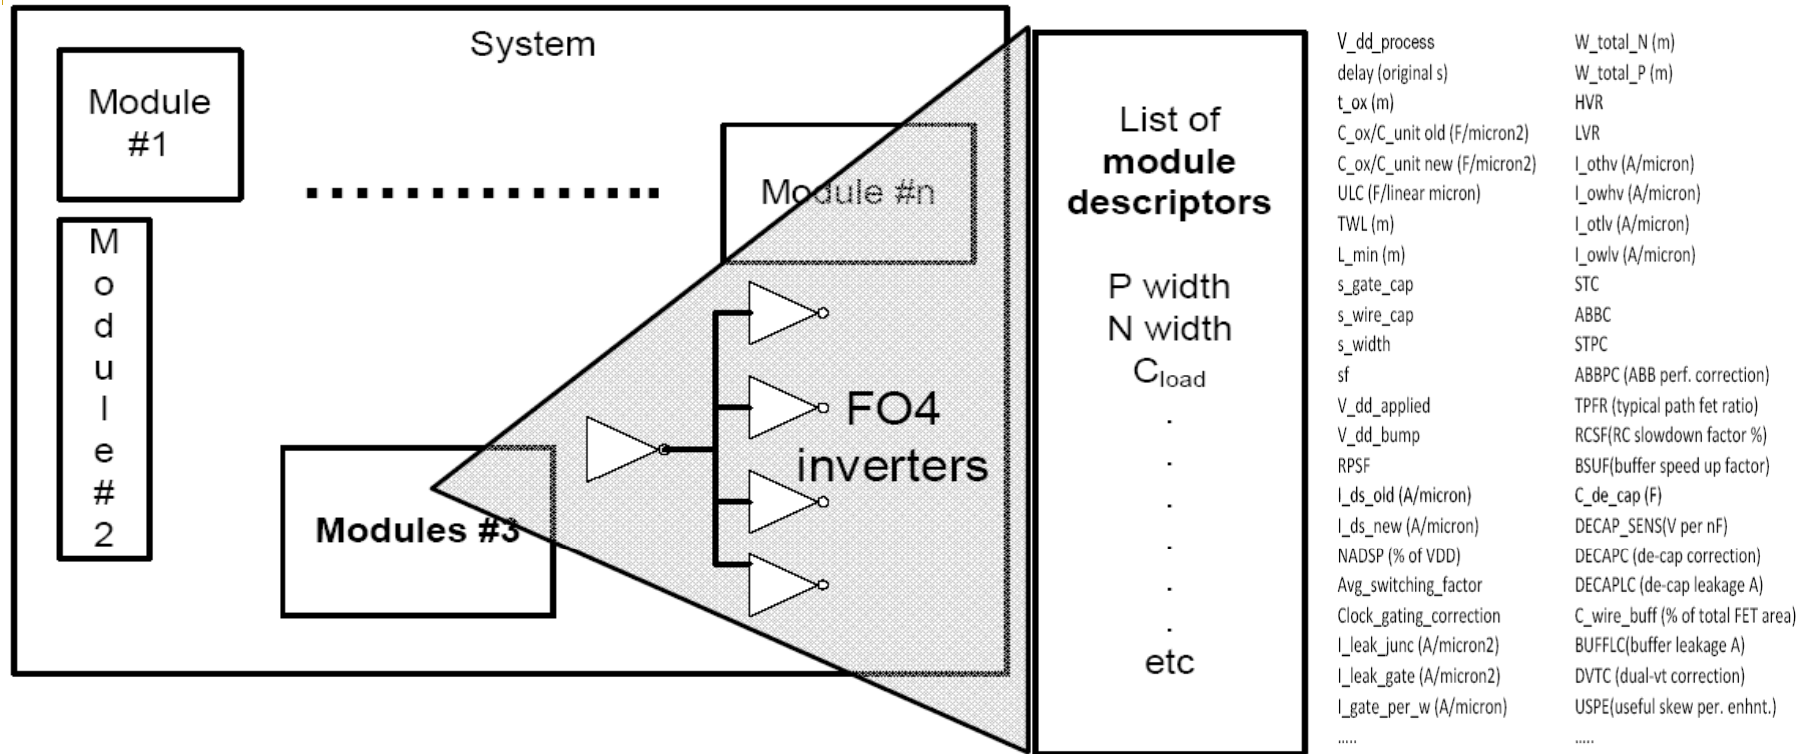


Module critical path delay = X EGDs

Each module is modeled as a fanout-of-4 equivalent logic gate configuration

- Simple yet effective abstraction, here standard inverters are used
- Module inverter size is proportional to the total N & PFET sizes in a module

System Modeling Methodology

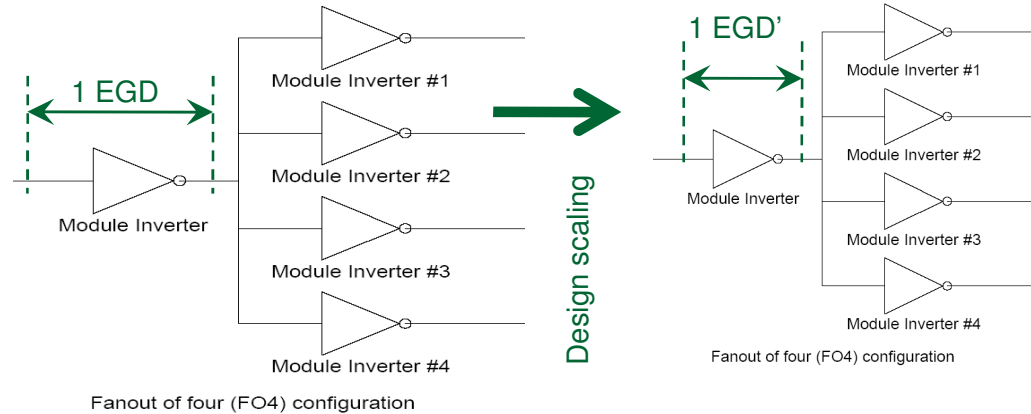



Each module is described by a set of descriptors (descriptor vector) ~ 60

- Process constants
- Physical sizes and other legacy data
- Simulations based estimates

A collection of descriptor vectors, forms the system model for early design space exploration.

Module Performance & System Performance



Module performance is given in EGDs

Normalized critical path delay_{legacy} = $EGD_{11} + EGD_{12} + EGD_{13} + EGD_{14}$

- Example $EGD_{11} = x \text{ NAND} + y \text{ NOR} + z \text{ INV} + p \text{ BUFF}$ delays expressed in EGD, say 100 EGD

When scaled, normalized critical path delay_{new} is still $EGD_{11} + EGD_{12} + EGD_{13} + EGD_{14}$

- But absolute delay is obtained from EGD' which is calculated from in-situ SPICE simulation on the FO4 model

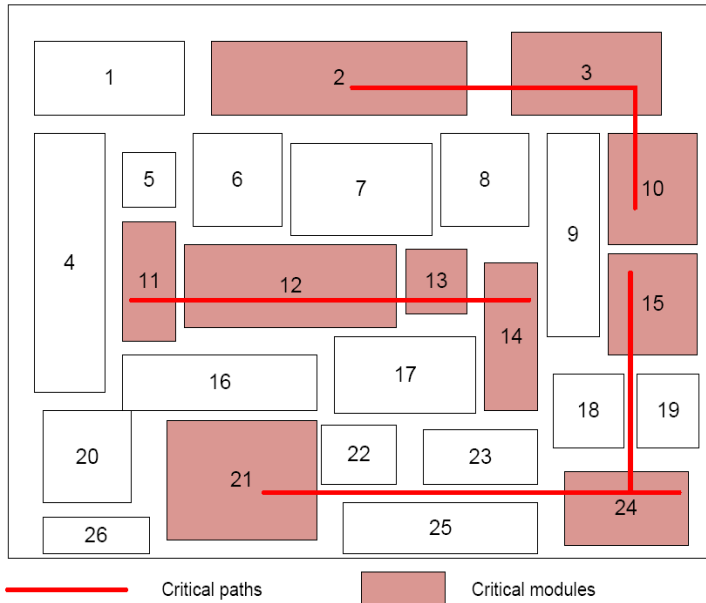
- absolute critical path delay of the scaled design = $(EGD_{11} + EGD_{12} + EGD_{13} + EGD_{14}) * EGD' = 100 * EGD'$

Modules with no legacy design uses estimated EGD

System critical path delay is the max path delay among all identified critical paths

$$f_{\text{new}} = 1 / \text{system critical path delay}$$

Module Power & System Power



For each module (1..26) :

W_N, W_P -- estimated from legacy data

$$C_{\text{dyn}} = C_{\text{gate}} + C_{\text{interconnect}}$$

$$P_{\text{dynamic}} = V^2 f_{\text{new}} C_{\text{dyn}}$$

$$P_{\text{leakage}} = P_{\text{gate}} + P_{\text{junc}} + P_{\text{sub-threshold}}$$

$$P_{\text{module}} = P_{\text{leakage}} + P_{\text{dynamic}}$$

End Loop

Total power = estimated dynamic + estimated leakage power.

Legacy design data is used to estimate C_{gate} & C_{dyn} . Interconnect switching cap is added to C_{dyn}

Total power estimated is the sum of the individual power estimates


$$P_{\text{system_dynamic}} = \text{Sum of all } P_{\text{dynamic}}$$

$$P_{\text{system_leakage}} = \text{Sum of all } P_{\text{leakage}}$$

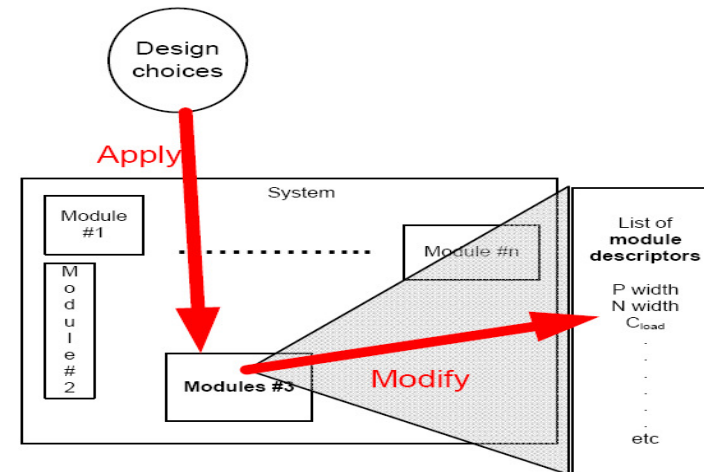
$$P_{\text{system_total}} = \text{Sum of all } P_{\text{module}}$$

Iterative Design Space Exploration


List of N circuit level design choice applicable to modules
choice₁ – choice_N



Each module has been assigned one (of N) design choice

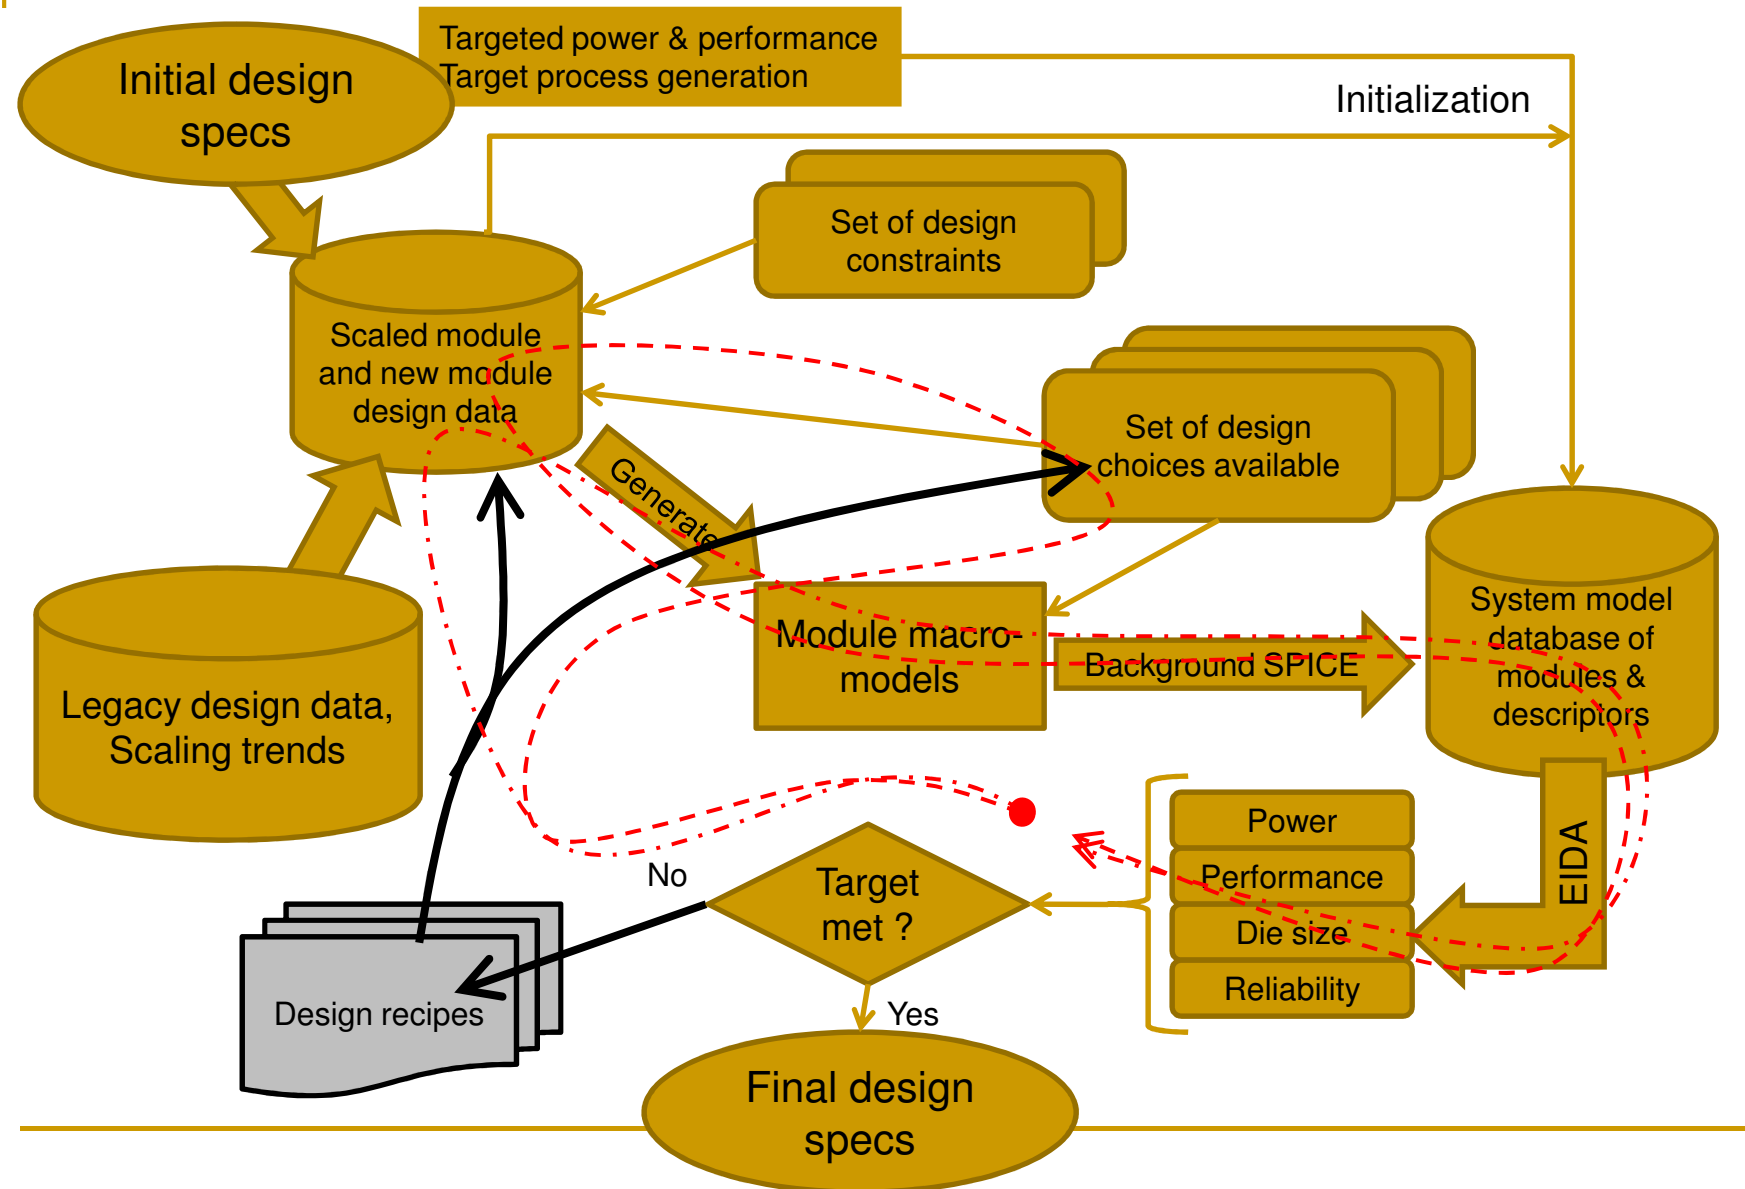


Change module design choices until necessary

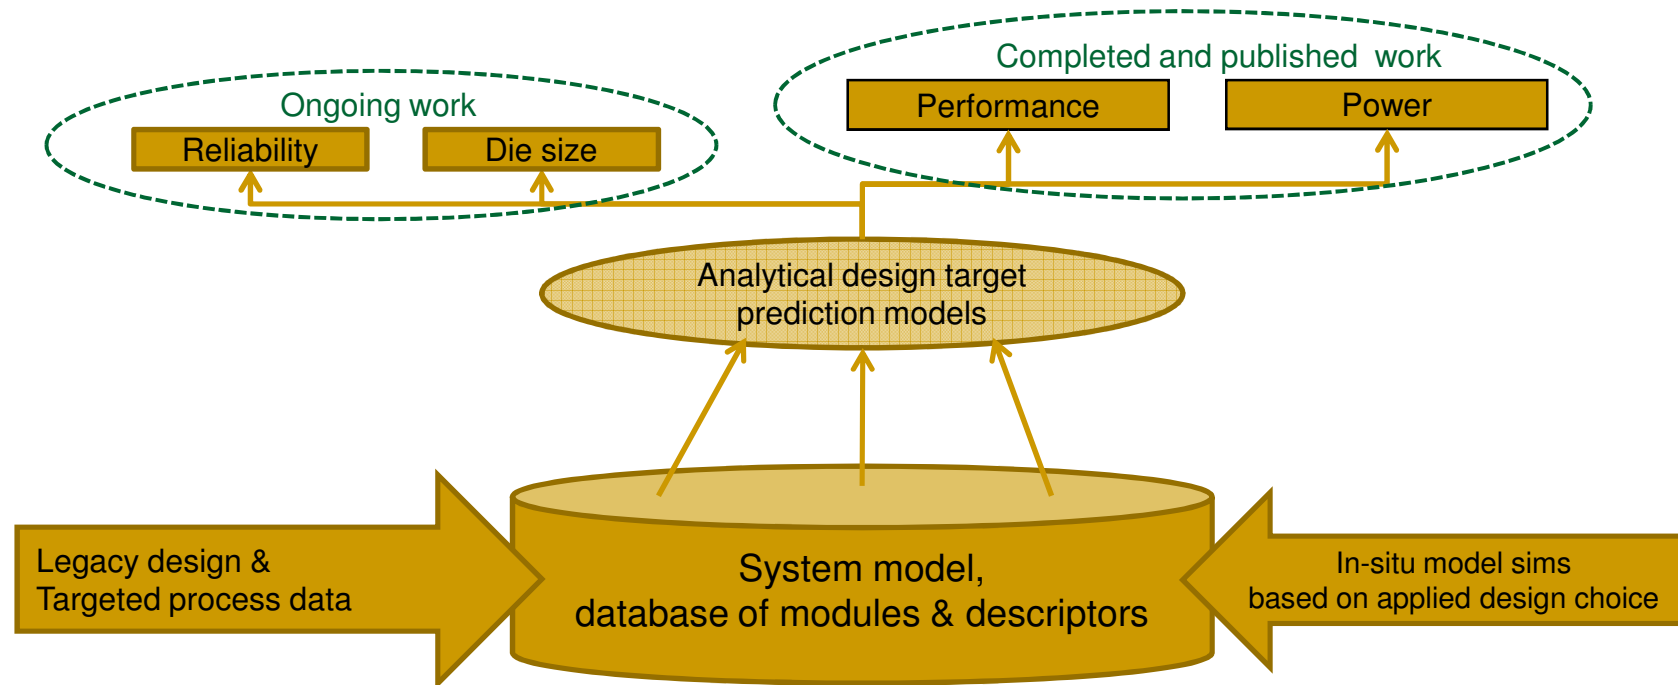


Module descriptor(s) is(are) updated, system model is updated as a result
Background SPICE simulations may be needed to generate the new set of descriptors.
When multiple modules are modified, corresponding descriptors are updated

Iterative Design Space Exploration - cont



Connecting System Design Targets And Physical Parameters



Design choices (13):

- 1) Increasing V_{dd} to improve perf.
- 2) Decreasing V_{dd} to improve pwr.
- 3) Adding sleep FETs to improve pwr.
- 4) Using low- V_t FETs to improve perf.
- 5) Applying back bias to improve pwr.
- 6) Applying back bias to improve perf.
- 7) Valid combinations of the above ~ 7

Analytical Design Target Prediction Models (more later)

Module power equations

$$P_{dyn} = C_{new} \times V_{dd_new}^2 \times f_{new} \times RPSF$$

$$C_{new} = \{C'_{orig} \times ((s_{gate_cap} \times frac_{fet}) + ((1 - frac_{fet}) \times s_{wire_cap}))\} + C_{wire_buff}$$

$$f_{new} = (f_{predict}) \times ASF \times CGF$$

$$P_{leak} = \{[(V_{dd_bump} \times I_{off}) + (V_{dd_bump} \times (I_{gate} + BUFLC)) + (V_{dd_bump} \times I_{junc})] \times ((1 - ASF) \times \varphi)\} + [V_{dd_bump} \times DECAPLC]$$

Where,

- RPSF – library redesign power saving factor
- ASF – average switching factor
- CGF – clock gating factor
- BUFLC – buffer leakage correction factor
- DECAPLC – decap leakage correction factor
- φ – sleep transistor and back bias correction

Module performance equations

$$f_{predict} = \frac{((f_{old} \times HVR) + (f_{old} \times LVR \times DVTC)) \times DIF}{(TPFR \times FSF) + ((1 - TPFR) \times RCSF_{\#})} \times X$$

$$RCSF_{\#} = \frac{RCSF}{BSUF}$$

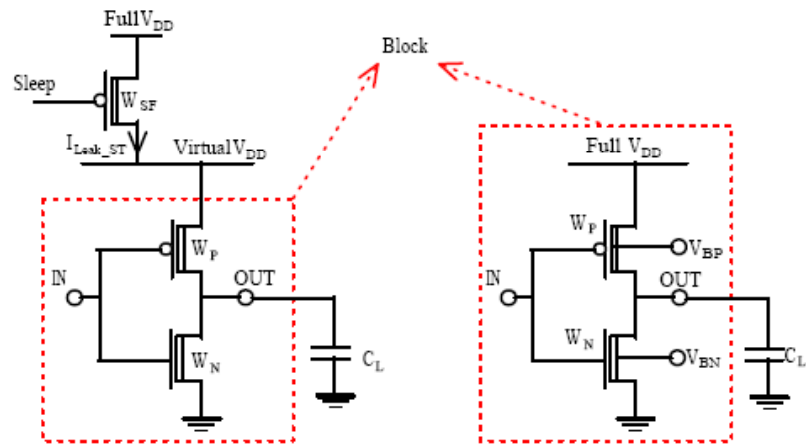
$$X = STPC \times ABBPC \times USPE$$

Where,


- HVR – high- V_t FET ratio
- LVR – low- V_t FET ratio
- f_{old} – 1/ original module critical path delay
- DVTC – dual V_t performance correction
- DIF – device (I_{ds}) improvement factor
- TPFR – typical path FET ratio (delay)
- FSF – FET slowdown factor (due to environment)
- RCSF – RC (interconnect) slowdown factor
- BSUF – buffer speed up factor (wire delay)
- STPC – sleep transistor perf. correction
- ABBPC – adaptive body bias perf. correction
- USPE – useful skew performance correction

In-situ Macromodel Generation

- The macromodels are generated when a design choice is applied to a module.
 - Sleep transistors insertion
 - Back body biasing
 - Supply voltage- V_{dd} changes
 - Low- V_t transistors for critical path speedup
 - Process technology changes through SPICE model



Experiment to obtain ABBC, ABBPC, STC and STPC factors



Experiment to obtain DVTC factor

- The effect of V_{dd} scaling on leakage currents and FET performance is estimated using a-priori analytical approximations on the target process technology.

Presentation Outline

- Background information
 - CMOS technology scaling and high-performance VLSI design
- Industry challenges
 - Business
 - Technology
 - Design automation
- Problem Statement
 - Motivating factors and objective of this work
- Existing research
 - An overview of what is out there
- Proposed Approach
 - Modeling methodology
 - In-situ macro model generation
- **Experimental Validation**
 - **Experimental setup, assumptions and analysis**
 - **Software architecture**
 - **Results discussions**
- Future work

Experimental Validation

Methodology verification


- Technology migration from 180 nm TSMC process to 130 nm PTM process
 - ◆ ISCAS85 C5315 ~ 178 inputs, 123 outputs, 2406 logic gates
 - ◆ ISCAS85 C6288 ~ 32 inputs, 32 outputs, 2406 logic gates
 - ◆ ISCAS85 C7552 ~ 207 inputs, 108 outputs, 3512 logic gates
 - ◆ ISCAS89 S9234 ~ 36 inputs, 39 outputs, 211 DFF, 5597 logic gates
 - ◆ ISCAS89 S13207 ~ 62 inputs, 152 outputs, 638 DFF, 7951 logic gates
 - ◆ ISCAS89 S15850 ~ 77 inputs, 150 outputs, 534 DFF, 9772 logic gates
 - ◆ ISCAS89 S38584 ~ 38 inputs, 304 outputs, 1426 DFF, 19253 logic gates
 - ◆ ISCAS89 S38417 ~ 28 inputs, 106 outputs, 1636 DFF, 22179 logic gates
- Accuracy validation with SPICE
 - ◆ Test circuit of manageable size for SPICE simulation setup

Design Space Exploration Using Pareto Analysis

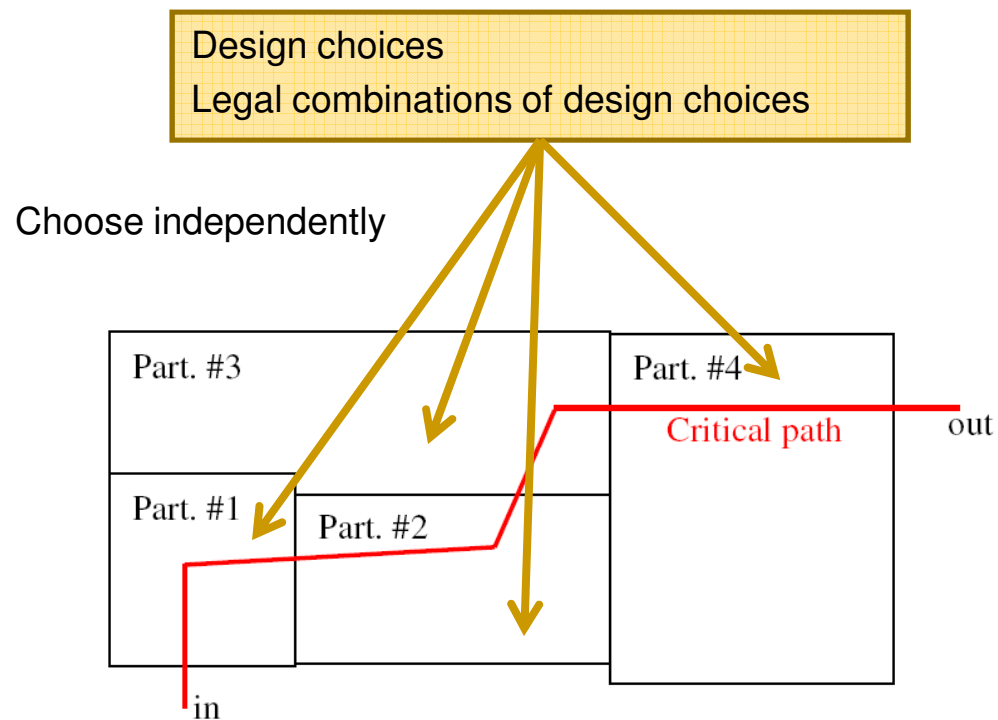
- On a larger 32 nm microprocessor based design

Methodology Verification – Tech Migration of ISCAS Circuits

- Circuits partitioned into four modules as shown (generic partition diagram)
- A portion of critical path falls in each of the modules (critical path determined using pathmill tool)
- Design choices
 - Lowering V_{dd} by 200 mV to reduce power consumption
 - Elevating V_{dd} by 200 mV to improve performance
 - Using Low V_t FETs in critical path to improve performance
 - Sleep transistor insertion to reduce idle or leakage power
 - Forward body biasing of critical path transistors to improve performance
 - Reverse body biasing of transistors to reduce leakage power
- Apply the above design choices to the partitions and evaluate the relative merit of the various design choice assignments or recipes on the design.






Methodology Verification – Tech Migration of ISCAS Circuits



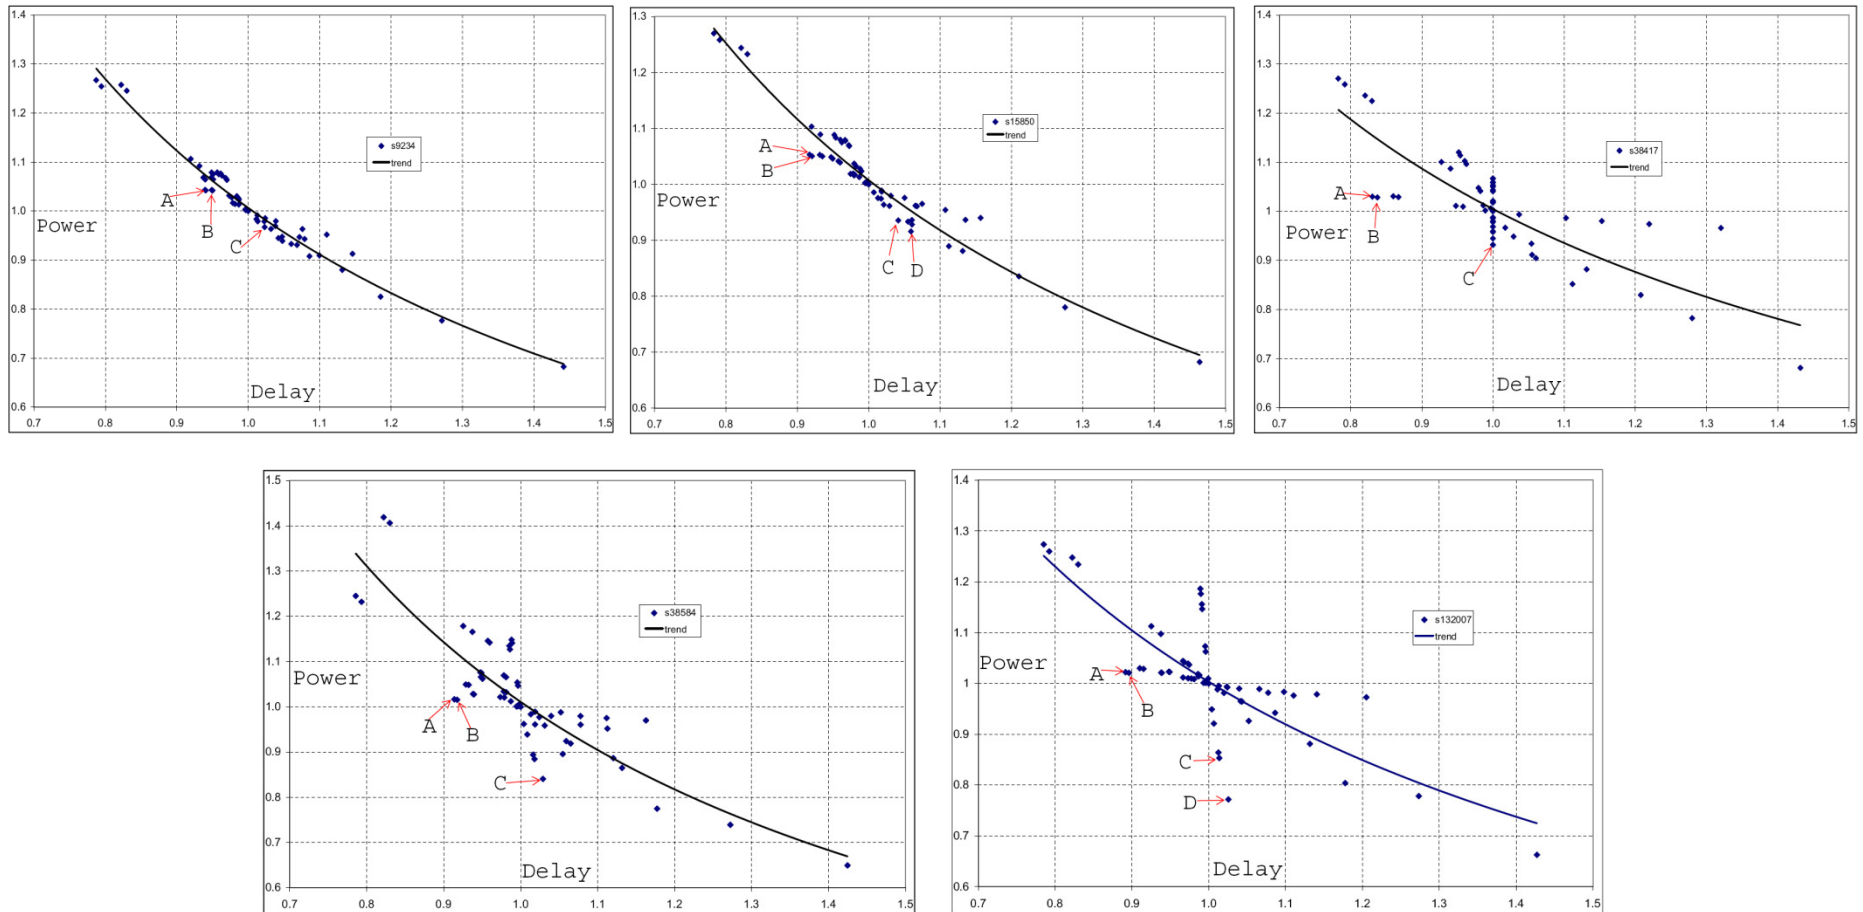
Following design choice assignment, EIDA tool predicts system power and performance.

Methodology Verification – Tech Migration of ISCAS Circuits




- ISCAS 85 circuit (Note : the change in axis)
- A,B performance centric & C, D (E) power centric

Methodology Verification – Tech Migration of ISCAS Circuits


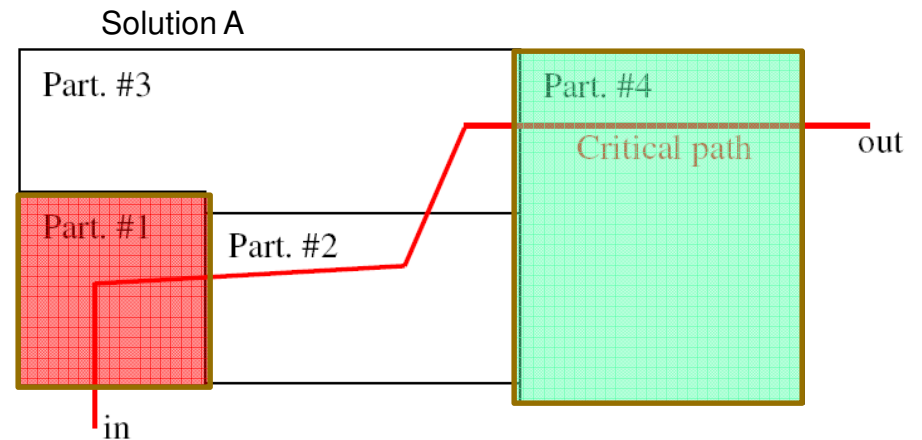


- ISCAS 89 circuit (Note : the change in axis)
- A,B performance centric & C (D) power centric

Methodology Verification – Tech Migration of ISCAS Circuits



- A - performance-centric solution
- 2.43% \uparrow power for 10.31% \uparrow performance
- E - power-centric solution
- 4.3% \downarrow performance for 17% \downarrow power




Low V_t

ABB-FB

$\downarrow V_{dd}$

ST

Methodology Verification – Tech Migration of ISCAS Circuits



- B - performance-centric solution
- 2% \uparrow power for 9% \uparrow performance
- C - power-centric solution
- 3% \downarrow performance for 16% \downarrow power
- Note – different circuit illustrated using pervious partition diagram





ABB-FB & $\uparrow V_{dd}$

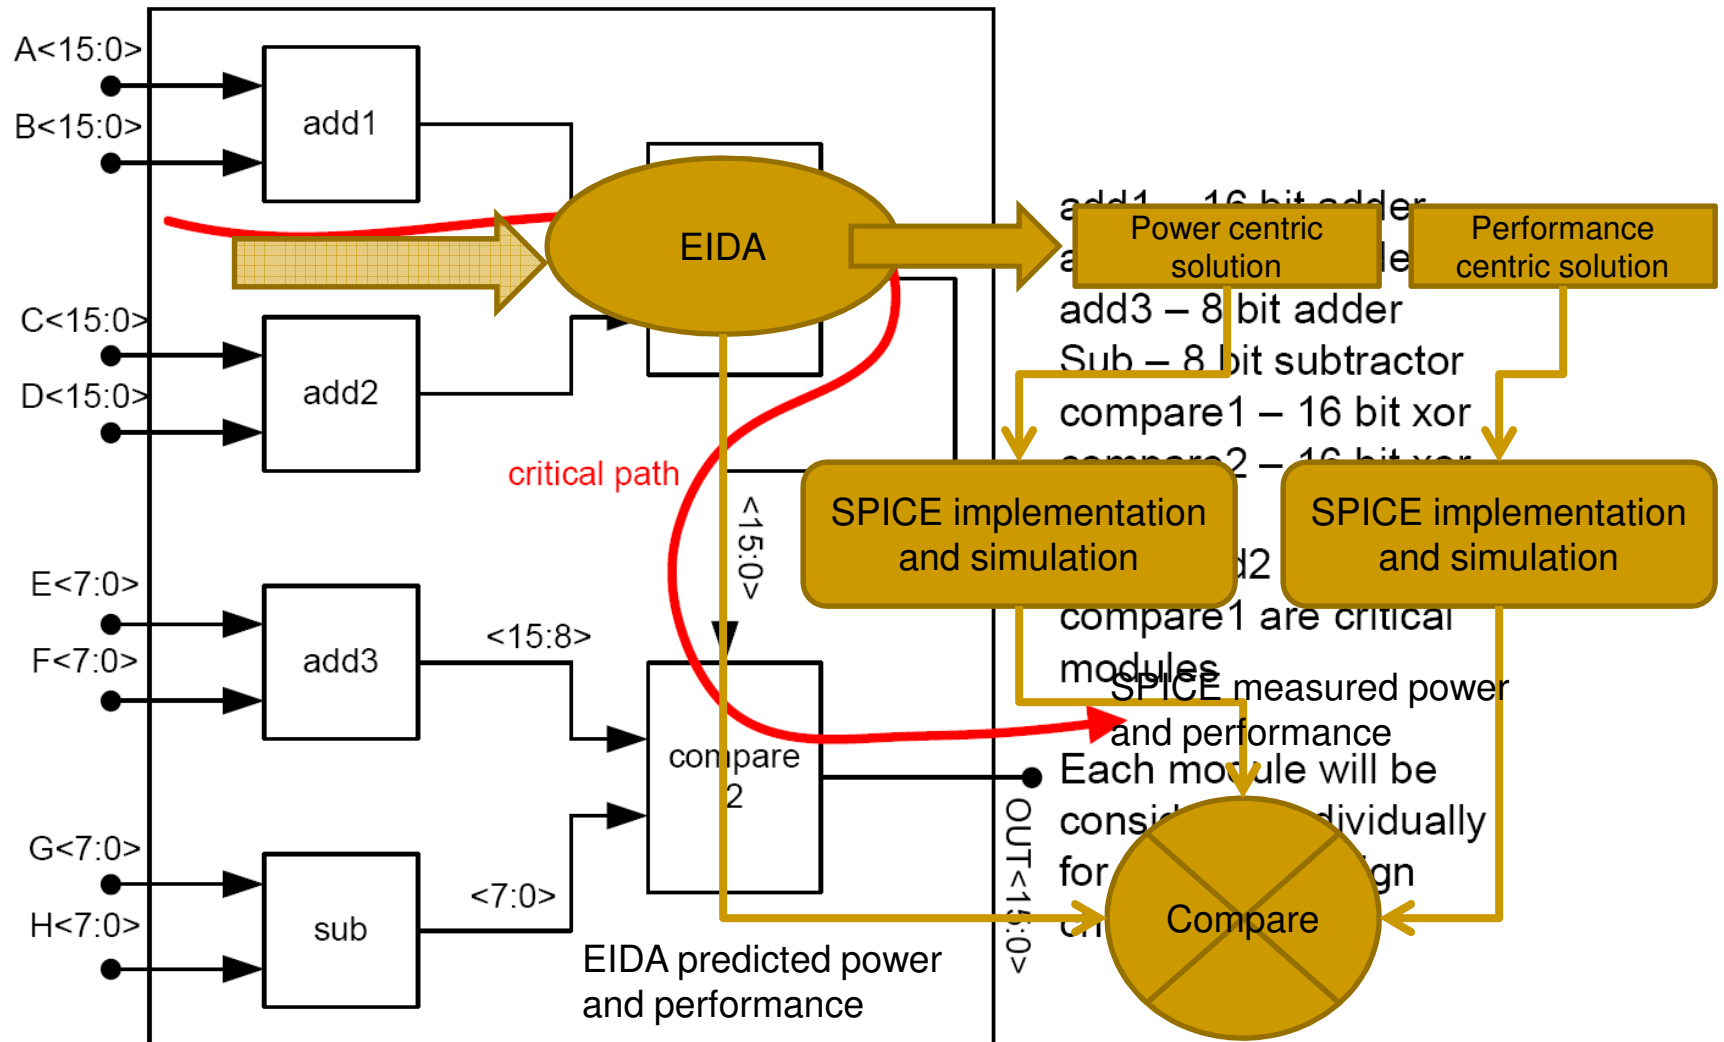
$\downarrow V_{dd}$

ST

Methodology Verification – Tech Migration of ISCAS Circuits, summary.

- Power – performance trends were as expected.
 - There is evidence for the existence of a kneeling point solution that optimizes both power and performance.
- Power – performance spread
 - exposes underlying circuit conditions that makes optimizing power and performance difficult
 - Allows us to choose power or performance centric designs
- The proposed methodology's feasibility for performing early design space exploration is established.
- How accurate are the estimates?

Methodology Verification – SPICE validation.



Methodology Verification – SPICE validation, design choices.

- Applied design choices
 - No changes to the original design
 - Reduce supply voltage by 100 mV
 - Reduce supply voltage by 100 mV & Sleep Transistors for power gating
 - Reduce supply voltage by 100 mV & ABB RB
 - Increase supply voltage by 100 mV
 - Increase supply voltage by 100 mV & dual Vt FETs in critical path
 - Increase supply voltage by 100 mV & ABB FB
 - Increase supply voltage by 100 mV, Dual Vt FETs in critical path & ABB FB
 - Dual Vt FETs in critical path
 - Dual Vt FETs in critical path & ABB FB
 - ABB FB
 - ABB RB
 - Sleep Transistors for power gating

- Use EIDA to find power & performance centric solution
 - Implemented in 32 nm process
 - Simulate the SPICE netlist, measure power and performance.
 - Compare SPICE measurements with EIDA predicted power and performance.

Methodology Verification – SPICE validation, results.

EIDA AND SPICE COMPARISON

Measurement	EIDA	SPICE	% error
just-ported pwr	1.7924 mW	1.6084 mW	-11.4%
just-ported perf	392 MHz	451 MHz	13.1%
lo-pwr assgn. pwr	1.0456 mW	0.937 mW	-11.5%
lo-pwr assgn. perf	406 MHz	368 MHz	-10.3%
hi-perf assgn. pwr	1.5523 mW	1.375 mW	-12.9%
hi-perf assgn. perf	442 MHz	513 MHz	13.8%

No bottoms up data;
Sufficient for early
design exploration

Perf. centric solution(modules 1 – 6)

Low-Vt FETs & ABB-FB

↓ Vdd & ST

↓ Vdd & ST

↓ Vdd & ABB-RB

None

None

Power centric solution(modules 1 – 6)

↑ Vdd & Low-Vt FETs & ABB-FB

↓ Vdd & ABB-RB

↓ Vdd & ABB-RB

↓ Vdd & ST

↓ Vdd & ST

↓ Vdd & ST

Design Space Exploration Using Pareto Analysis

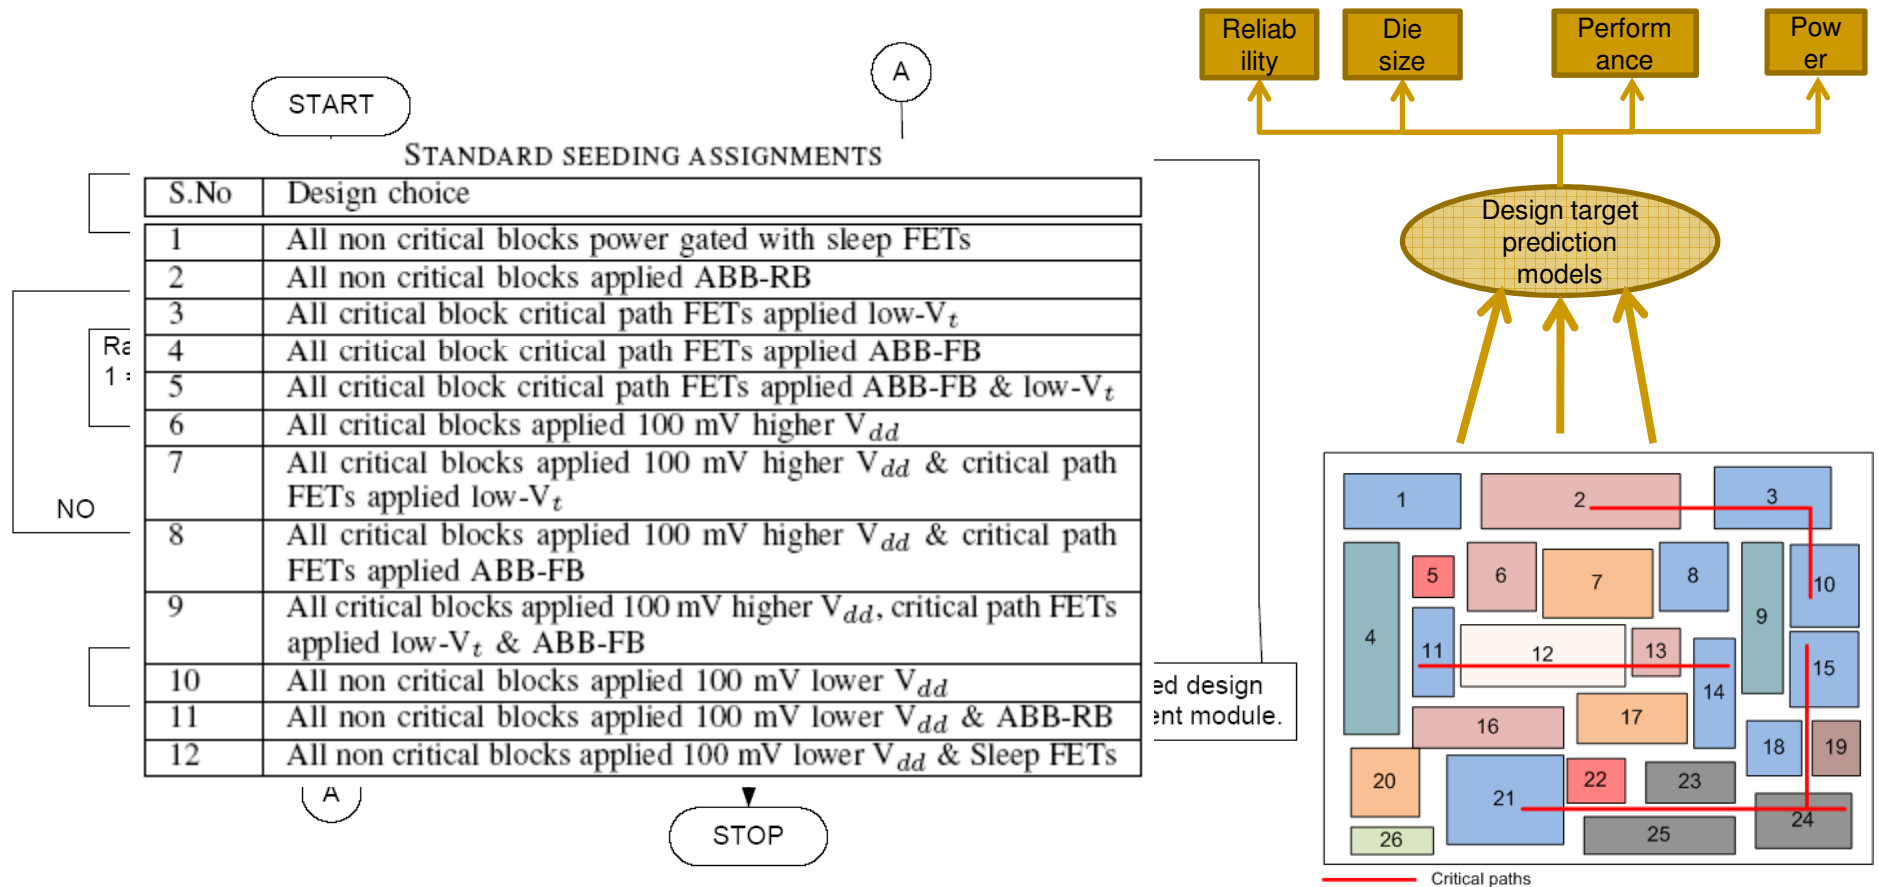
- A larger, microprocessor based design
 - 65 nm design
 - Design partitioned into 26 modules
- Critical modules and path
 - 10 critical modules out of 26 modules
- Experiment 1 -- Leveraged design scaling
 - Scaled and ported the 65 nm design to a 32 nm process
- Experiment 2 -- Design space exploration by applying various (13 possible) design choices
 - Goal is to find the optimal design choice assignment
 - Optimize power and performance
 - Pareto-front analysis

Design Space Exploration Using Pareto Analysis cont...

26 modules, 13 design choices => 13^{26} different assignments possible !!

Large solution space – random generation, seed the randomizer with standard assignments

Two stage randomizer – generate (by mutation) additional design assignments



Design Space Exploration Using Pareto Analysis – Results


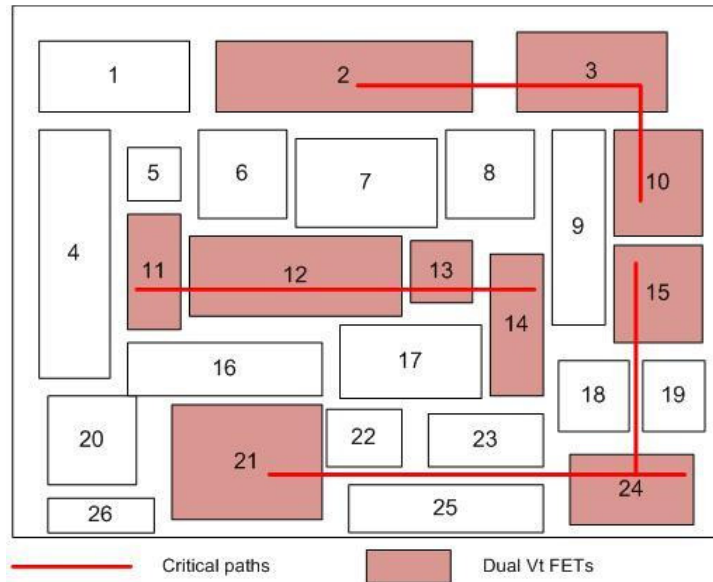



Illustration Examples



#14 – standard design practice to improve performance

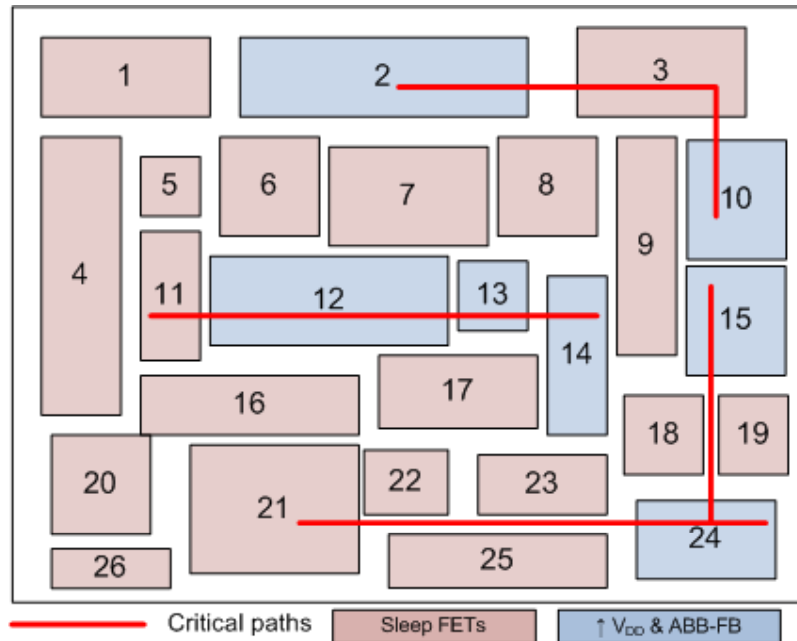
This solution improves performance by 12% while incurring a 3% power penalty



#13 – a generated design solution that optimizes power and performance.


This solution achieves 11.7% increase in performance **while saving 2% power**

Illustration Examples cont...



#4 – a generated/slightly altered standard design solution that optimizes power and performance


This solutions reduces power by 22.2% while there was 0.3 % improvement in performance



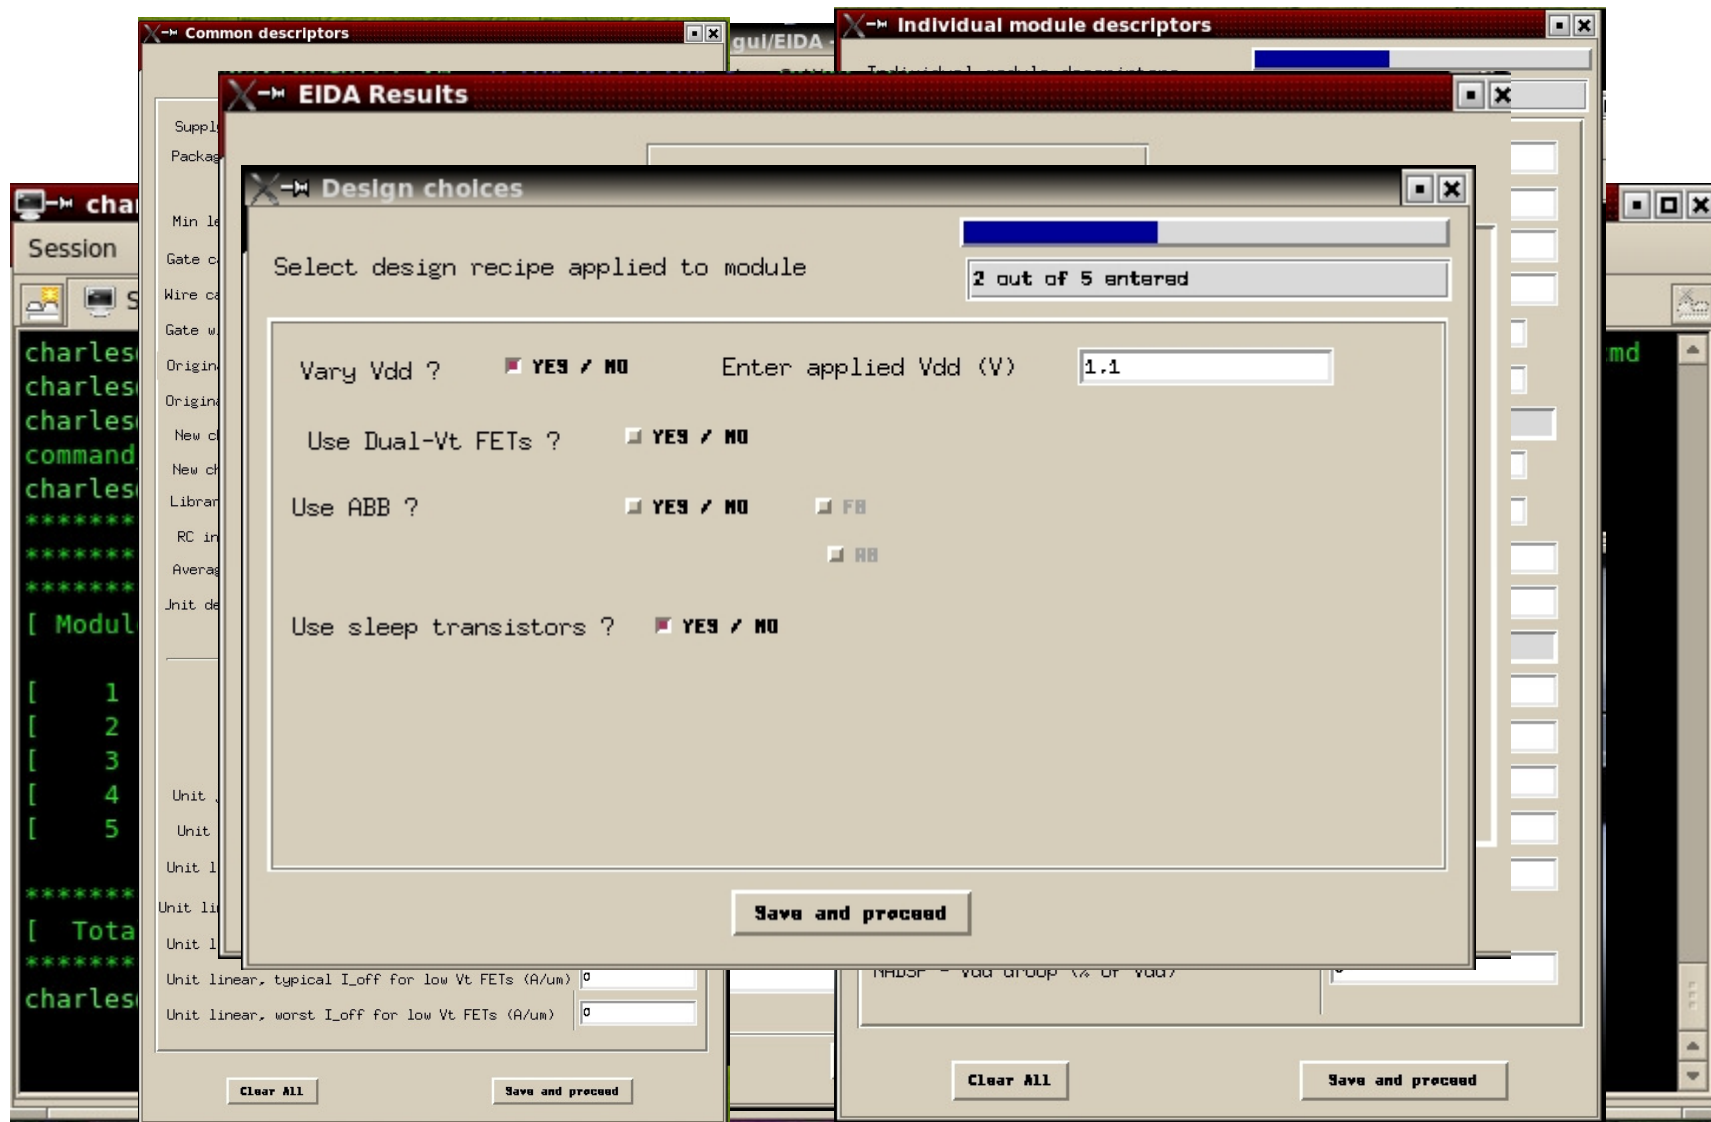
#11 – a generated design solution that optimizes power and performance.

This solution achieves 19.6% power savings **while increasing performance by 6.2%**

Impact of ABB: Complexity reduction



Software Implementation



Research Summary

- The proposed methodology's feasibility for performing early design space exploration was verified by performing technology migration experiment on ISCAS circuits
- A leveraged design space exploration from 65 to 32 nm PTM CMOS was performed on a test circuit to determine the accuracy of the proposed methodology.
- A Pareto-front analysis was applied to a larger design in 65 nm PTM CMOS
- Pareto-front analysis yielded solutions that
 - Achieves 11.7% increase in performance while saving 2% power
 - Achieves 19.6% power savings while increasing performance by 6.2%
- Solution complexity and power-performance optimality can be traded-off by considering solutions in the pareto solution region.
- Developed backend software and GUI
- The main contribution of the work are
 - Knowing what design choices to apply to meet a specific goal very early in the design phase
 - Help uncover non-intuitive design choices that meet or improve design goals
 - Help in power-performance convergence in later design stages
 - Help generate implementation design specs with higher probability of globally design convergence.

Presentation Outline

- Background information
 - CMOS technology scaling and high-performance VLSI design
- Industry challenges
 - Business
 - Technology
 - Design automation
- Problem Statement
 - Motivating factors and objective of this work
- Existing research
 - An overview of what is out there
- Proposed Approach
 - Modeling methodology
 - In-situ macro model generation
- Experimental Validation
 - Experimental setup, assumptions and analysis
 - Software architecture
 - Results discussions
- Future work

Future Research Plan

	Time
<ul style="list-style-type: none"> Design target prediction models <ul style="list-style-type: none"> Area, reliability and yield Improving interconnect delay calculation method <ul style="list-style-type: none"> To include V_{dd} scaling effects on interconnect delay Buffer insertion descriptors (BSUF & BUFFLC) Including de-cap leakage dependence on V_{dd} EGD (equivalent gate delay) formulation to define critical path delay. <ul style="list-style-type: none"> Replace current external calculation method with In-situ calculation of EGD 	1.5 month
<ul style="list-style-type: none"> Explore options for modeling analog modules, special modules that do not follow regular trends 	2.5 months
<ul style="list-style-type: none"> Exploration of better ways to perform design space exploration and pareto analysis. 	2 months
<ul style="list-style-type: none"> Software upgrades, GUI enhancements and documentation. 	1 month
<ul style="list-style-type: none"> Experiments to demonstrate multi-objective optimization <ul style="list-style-type: none"> Performed on full chip data in 32 and/or 22 nm CMOS technologies Publish results in journal and conferences 	1.5 months
<ul style="list-style-type: none"> Dissertation writing 	1.5 months

Past Publications

Conference papers

- A Fully CMOS-Compatible Optical H-Tree & Clock Recovery System, VLSI-SoC 2008.
- Using Early Design Phase In-situ Macro Models for Design Convergence, ISCAS 2008.
- Early Design Phase Power Performance Trade-offs Using In-situ Macro Models, DELTA 2008.
- Power & performance analysis for early design space exploration, ISVLSI 2007.
- Design of clock recovery circuits for optical clocking in DSM CMOS, IEEE SPIE Conf. May 2007.
- Optical Characterization of a Leaky-Mode Polysilicon Photodetector Using Near-Field Scanning Optical Microscopy, CLEO/QELS May 2006.
- Characterization of CMOS compatible, waveguide coupled leaky-mode photodetectors, Photonic Technology Letters, Aug. 2006.
- “Waveguide Coupled CMOS photodetector for on-chip optical interconnects”, Proc. of SPIE -- Volume 5556 Photonic Devices and Algorithms for Computing VI, November 2004, pp. 27-33
- “Truly CMOS compatible waveguide coupled photodetector for on-chip optical interconnects” IEEE LEOS, Puerto Rico, November 2004.

Journal papers

- Rapid Design Space Exploration Using Legacy Design Data And Technology Scaling Trend, IEEE CAD. – **Submitted**
- Fully CMOS-Compatible On-Chip Optical Clock Distribution & Recovery, IEEE Transactions on VLSI Systems – **Submitted**

THANK YOU

QUESTIONS