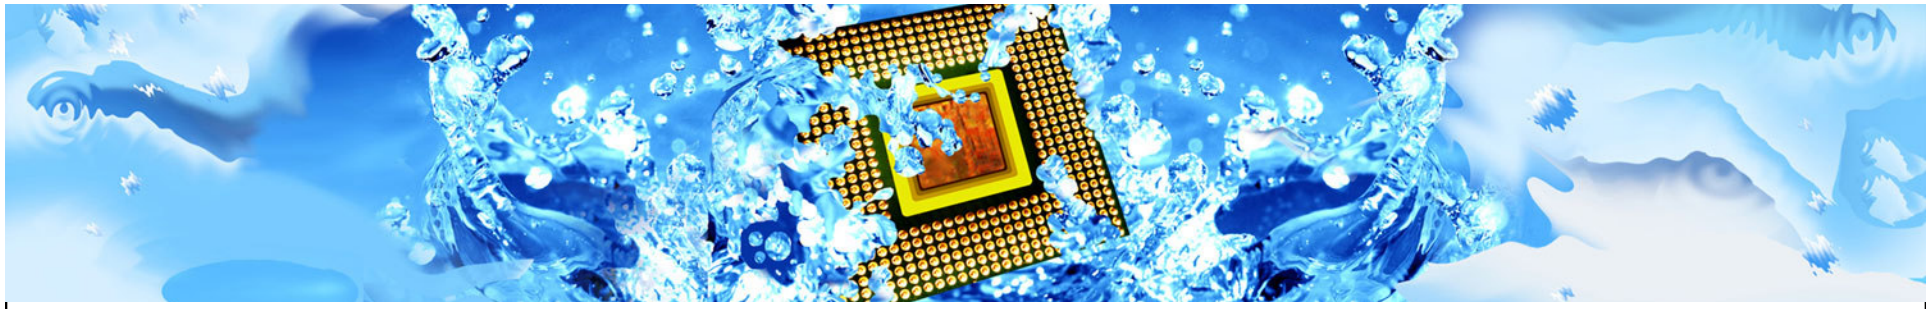

Power Management Early in the Design Flow: Exploration to Implementation

April, 2007

Holly Stump, VP Marketing

SEQUENCE

Enabling Power-Aware SoC Design℠

Power Management Early in the Design Flow: Exploration to Implementation

- Managing Power at the Architectural Level
- RTL Power Management
- Power Debug Environment
- Silicon-Aware Power Management
- Software and Mode-Dependent Stimulus
- Identifying and Eliminating Wasted Power at RTL
- Popular Power Reduction Techniques
- Clock Power and Clock Gating
- Multi-Vt
- Voltage Islands
- Power Gating
- Power Regression Testing
- Power Metrics
- Summary

SEQUENCE
Enabling Power-Aware SoC Design℠

Design Investigations and Power Requirements

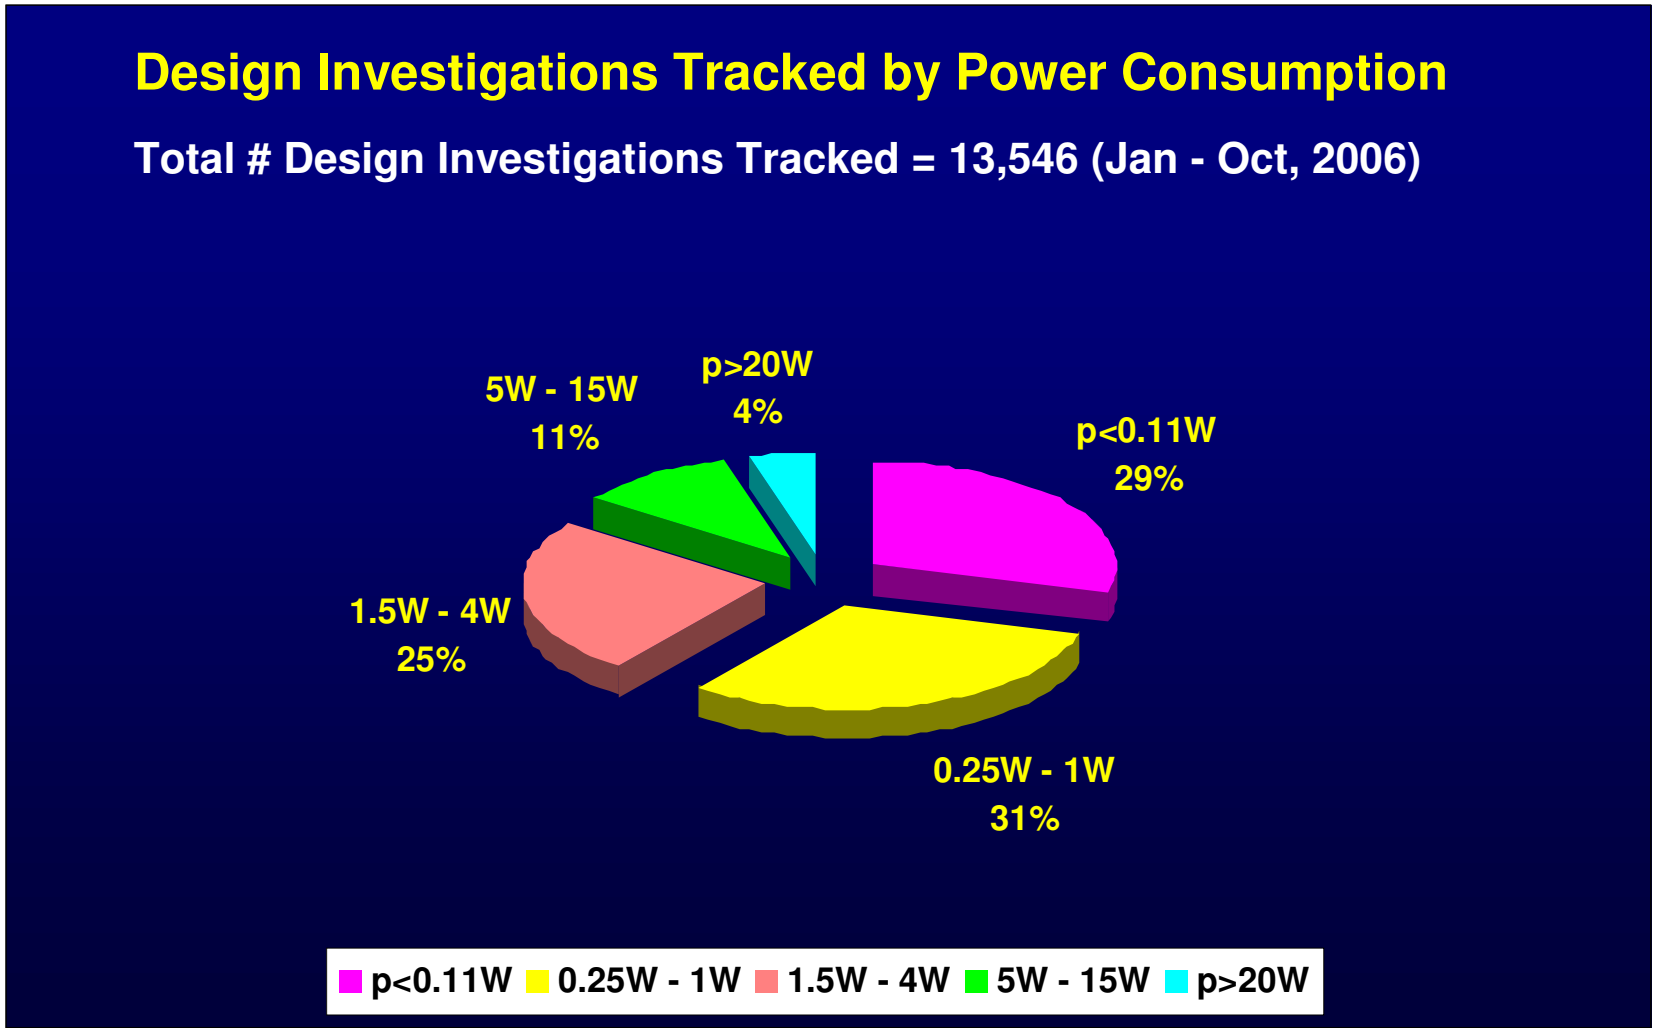

Ref: Chip Design Trends Newsletter, John Blyler, Dec 2006

SEQUENCE
Enabling Power-Aware SoC Design℠

Power Is The New Performance!

Low Power Is Critical Due To:

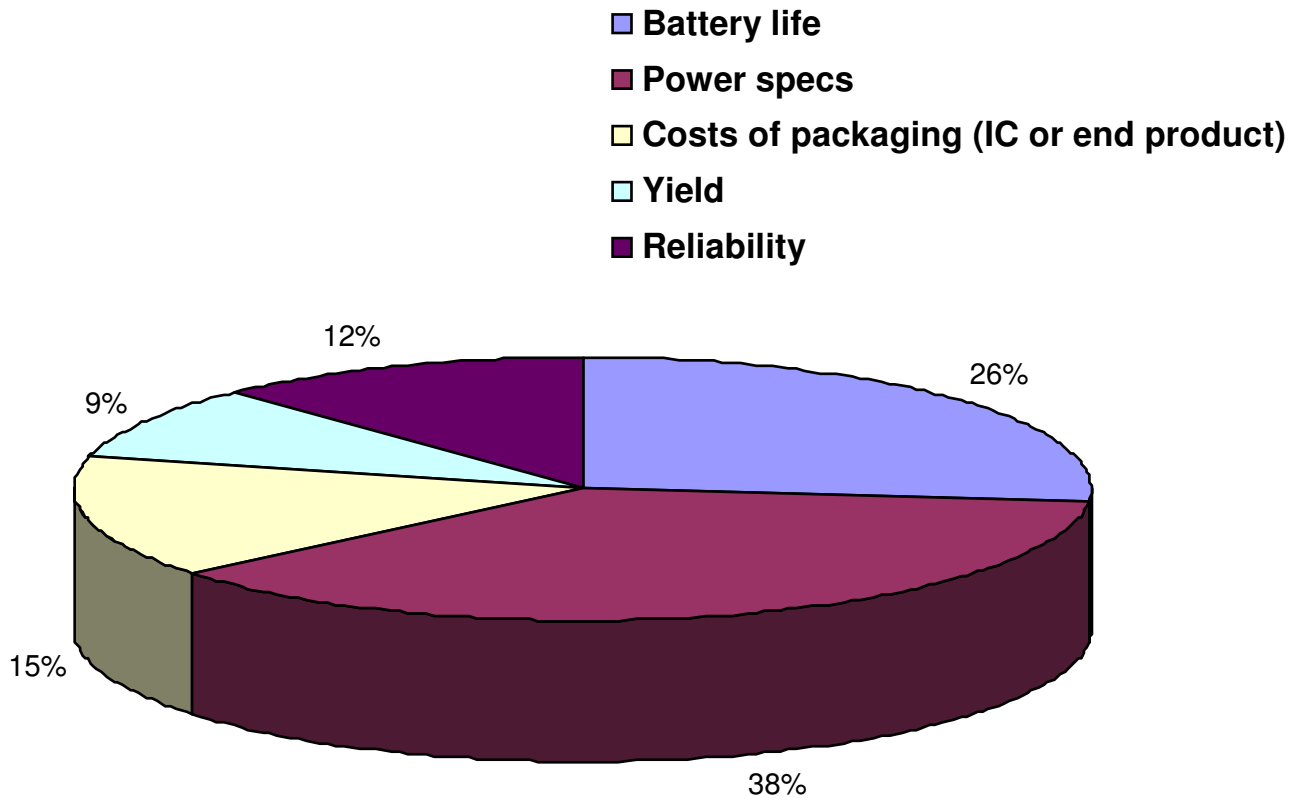

Survey Summary of SoC Designers DAC 2006
Sample size = 115

SEQUENCE
Enabling Power-Aware SoC Design℠

Managing Power at the Architectural Level

- **Are you an architect?**
  - What-if analysis for micro-architectures
  - Optimization for power, performance, area
- **ESL and other explorations**
- **Intelligent debug environment**
  - RTL power estimation
  - RTL power management

SEQUENCE
Enabling Power-Aware SoC Design℠

Power-Aware ESL Synthesis Flow

Without ESL Power Analysis

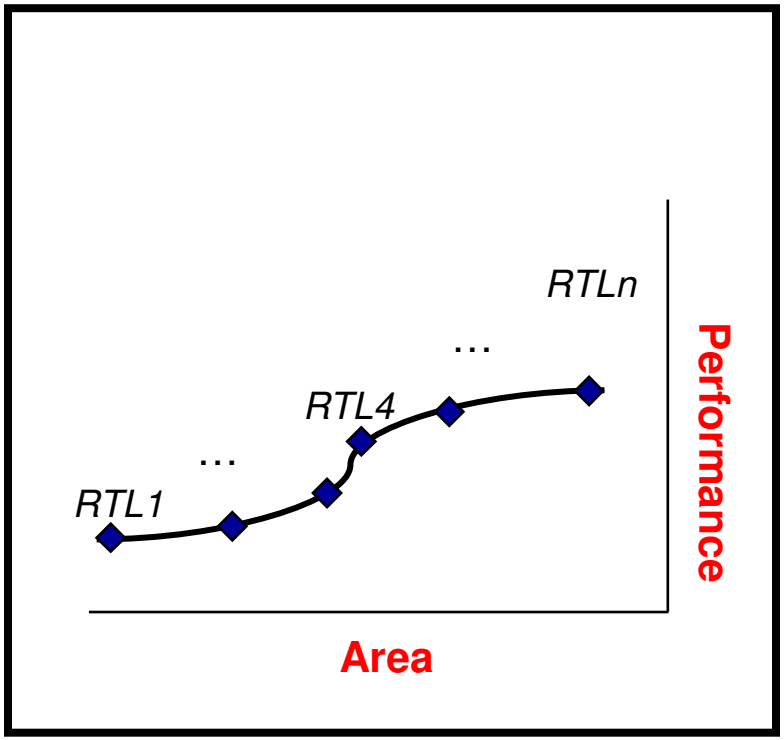

Area vs. performance tradeoff

With ESL Power Analysis

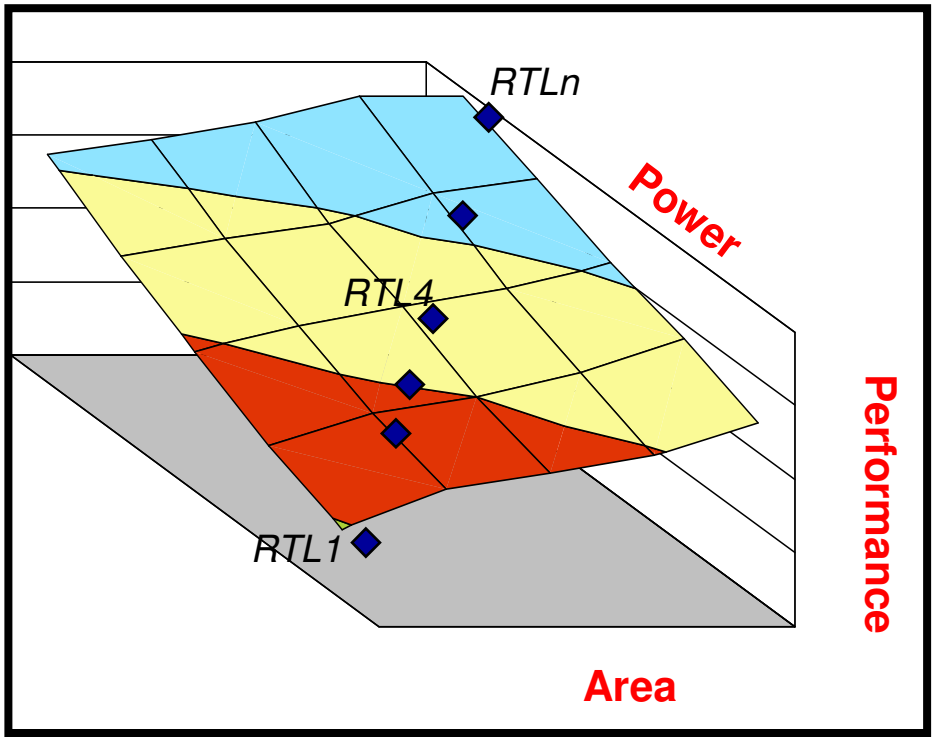

*Area vs. performance vs. **power** tradeoff*

SEQUENCE
Enabling Power-Aware SoC Design℠

802.11a: Optimized for power, area, performance

- **802.11a Wi-Fi transmitter**
- **7 candidate micro-architectures**
  - Push-button tool flow:
    - Bluespec for design
    - Sequence for power
    - Synopsys for RTL synthesis
- **Final design: 4 milliwatts**

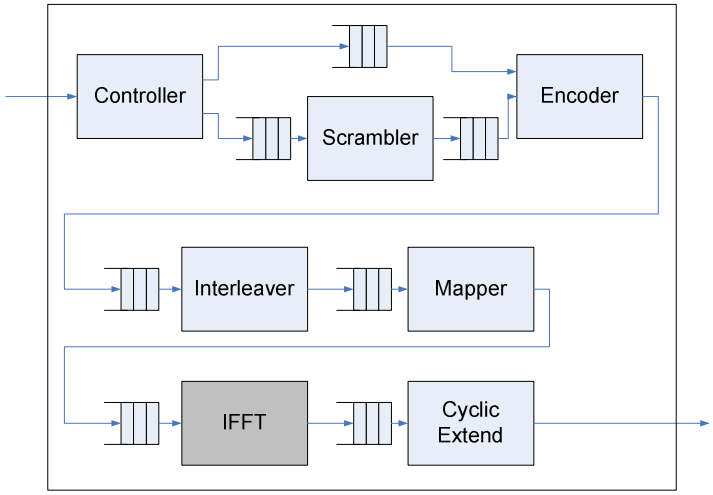

7 micro-architectures implemented and explored in only 5 engineer-days

Source: Dave, Pellauer, Gerding & Arvind
Courtesy: MIT

bluespec

SEQUENCE
Enabling Power-Aware SoC Design℠

802.11a Transmitter Overview

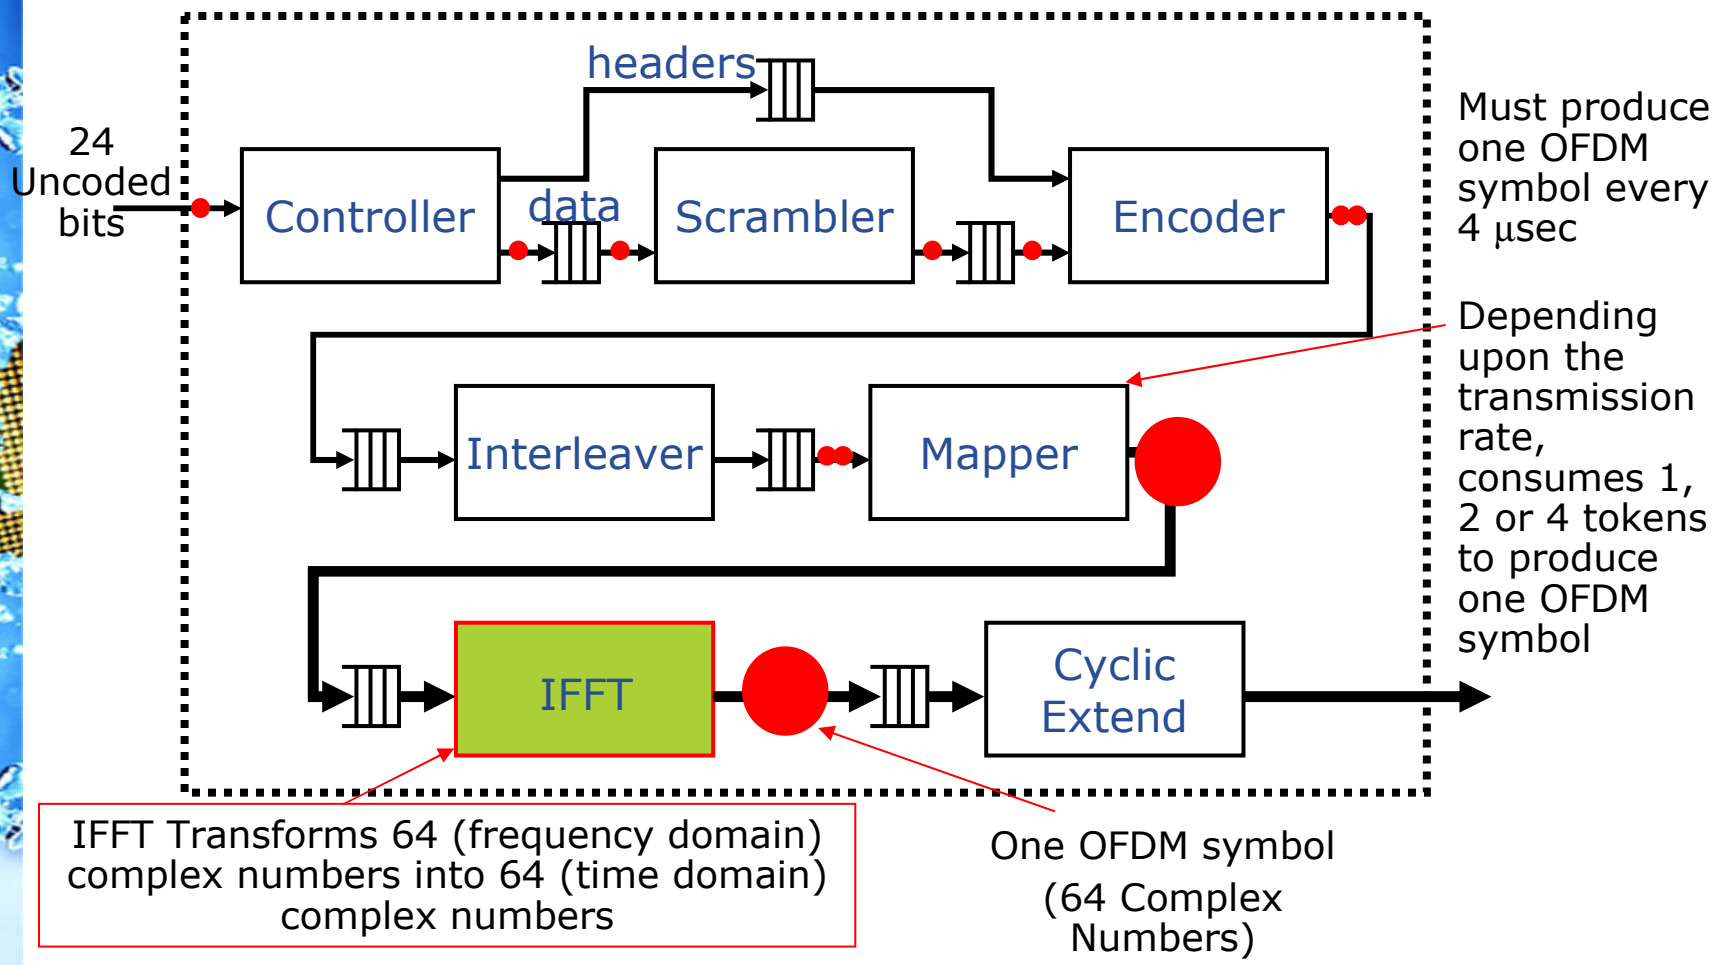

bluespec

SEQUENCE
Enabling Power-Aware SoC Design℠

IFFT module (combinational)

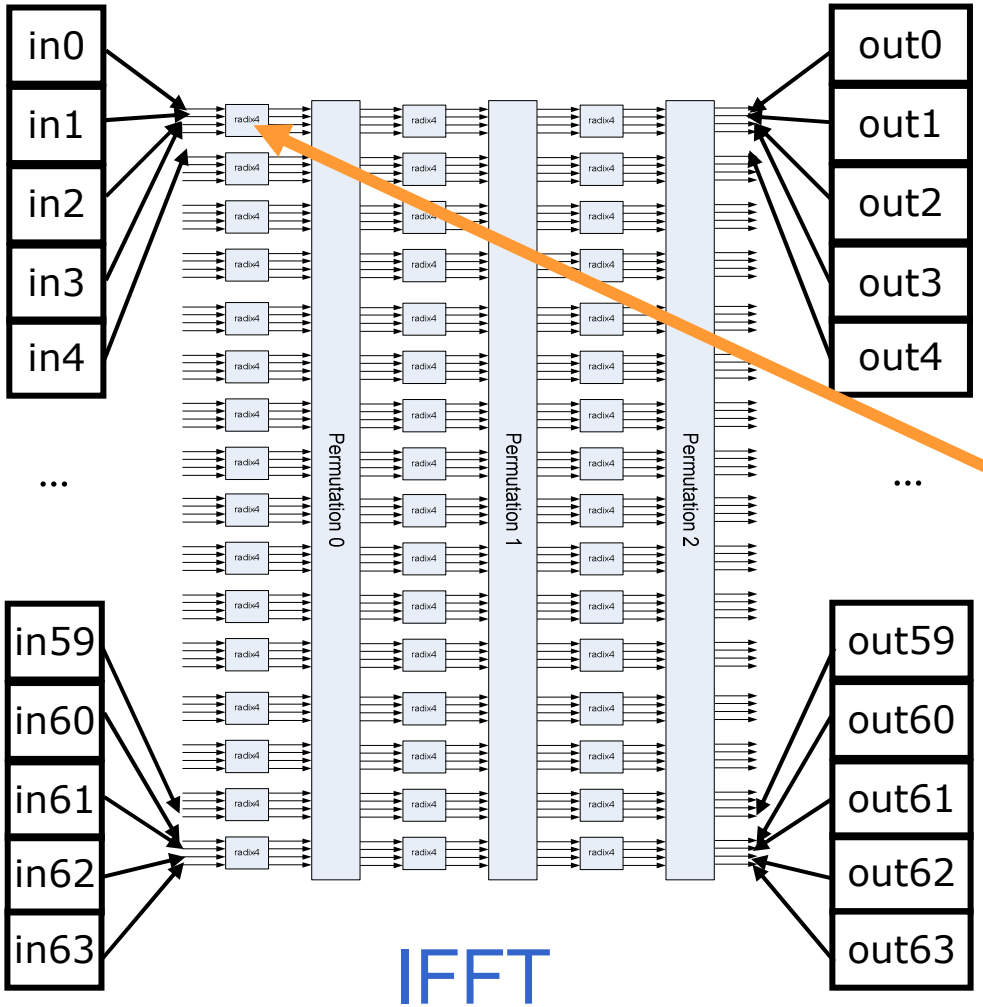

IFFT

Each of the 48 radix4 blocks looks like this:

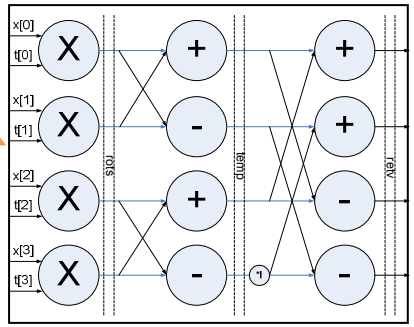

All numbers are complex and represented as two sixteen bit quantities. Fixed-point arithmetic is used to reduce area, power, ...

bluespec

SEQUENCE
Enabling Power-Aware SoC Design℠

IFFT: Micro-architectural exploration

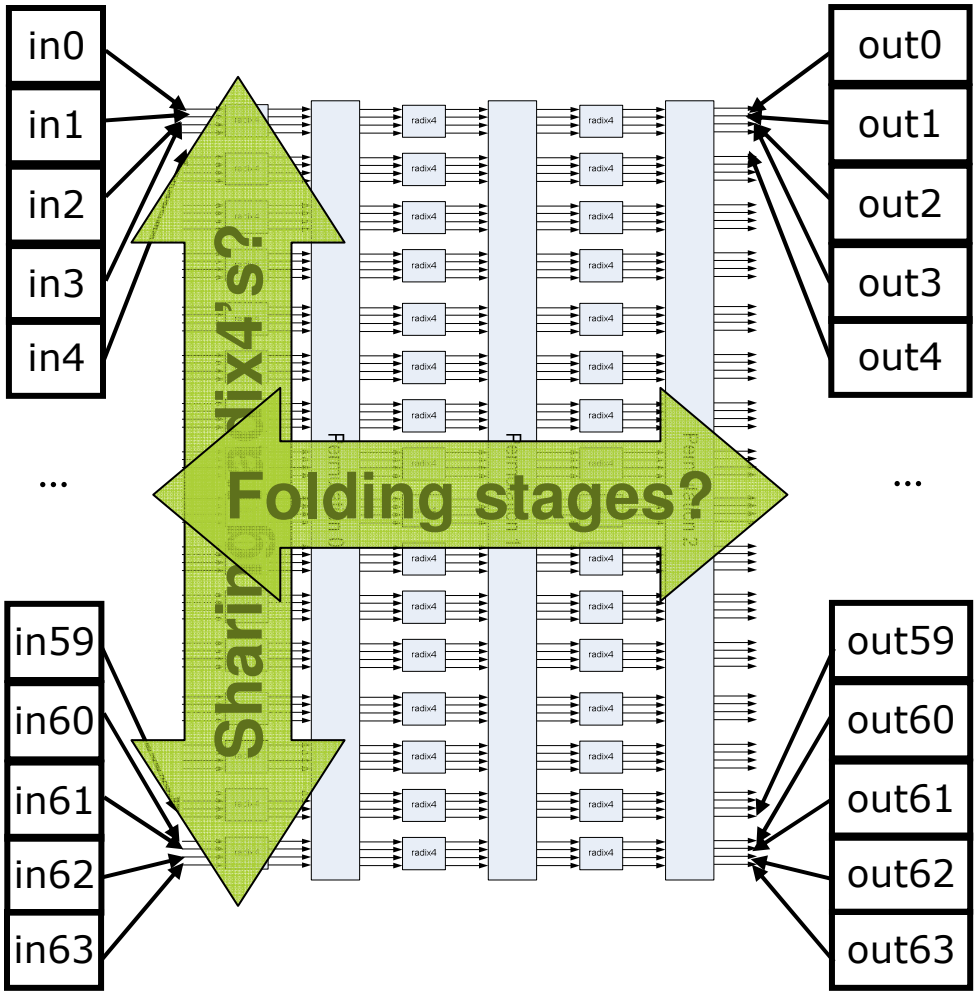

Each stage's 16 radix4 blocks could be also implemented with 8, 4, 2 or 1 radix4 block(s) used over multiple cycles

Each stage is almost identical, why not fold and re-use what you can?

bluespec

SEQUENCE
Enabling Power-Aware SoC Design℠

Superfolded circular pipeline: Just one Radix-4 node!

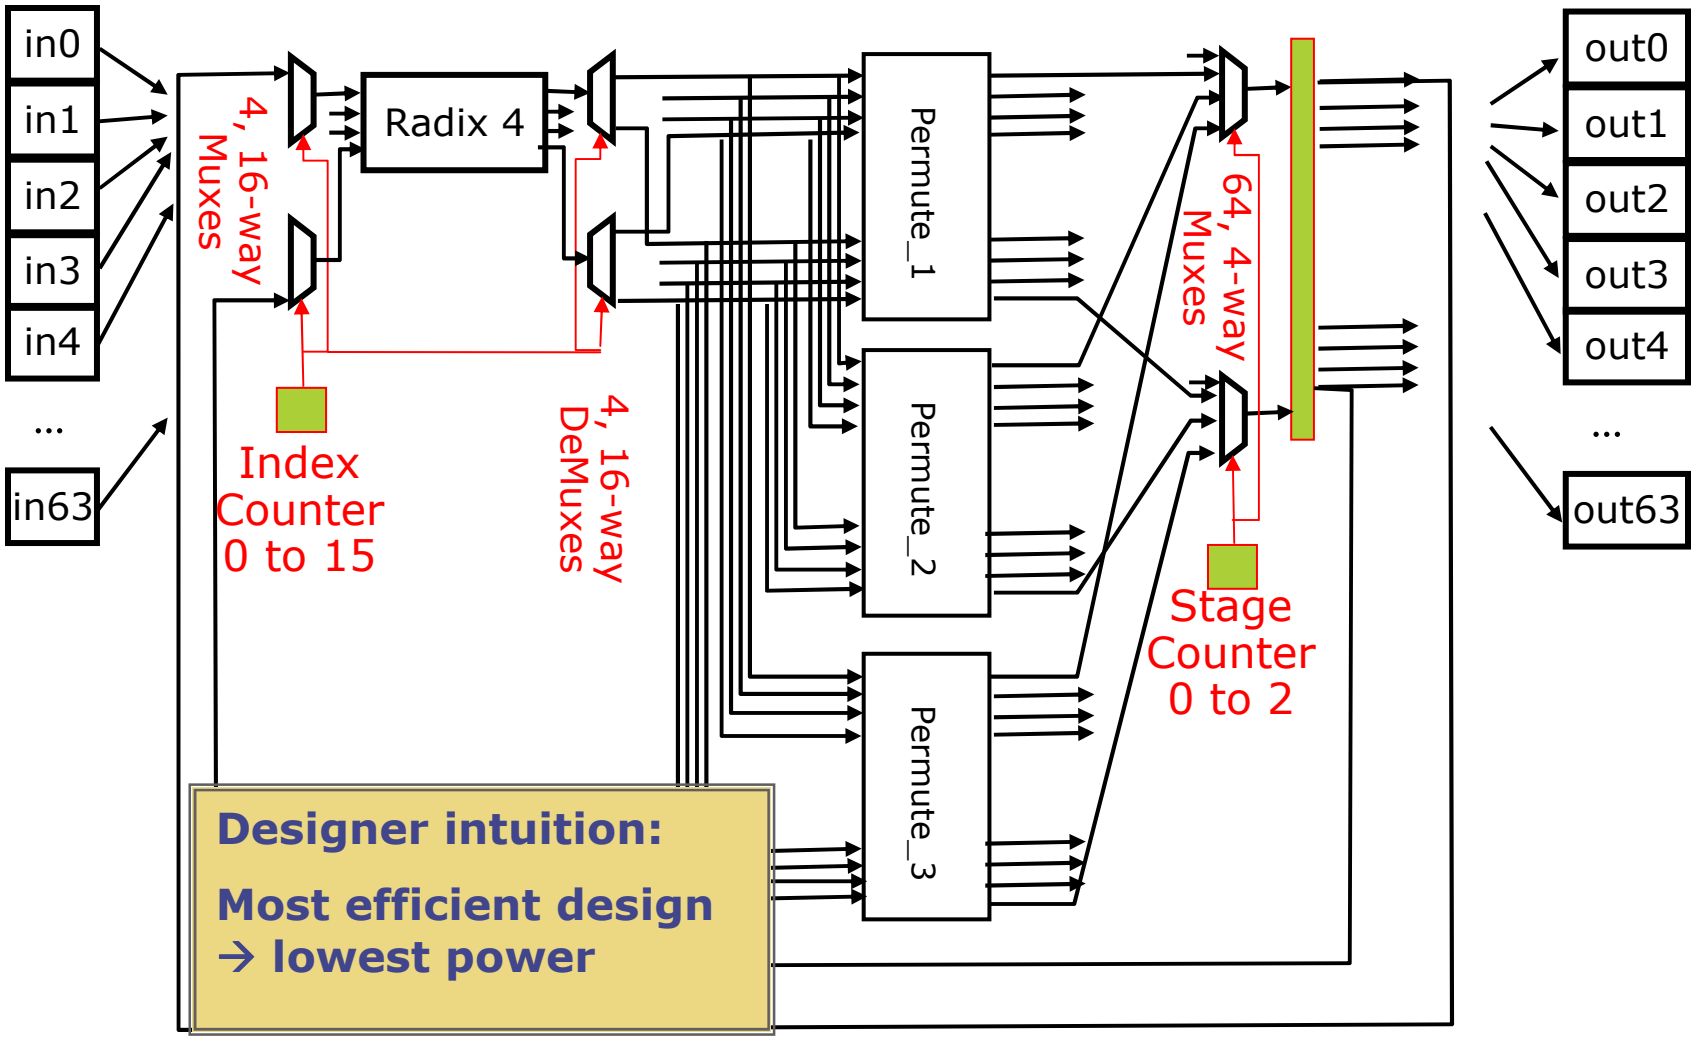

bluespec

SEQUENCE
Enabling Power-Aware SoC Design℠

Performance Results

All the combinations created and explored within five days

Designers were astounded to find that their intuitions were wrong and that the critical areas for reducing power were not where they suspected

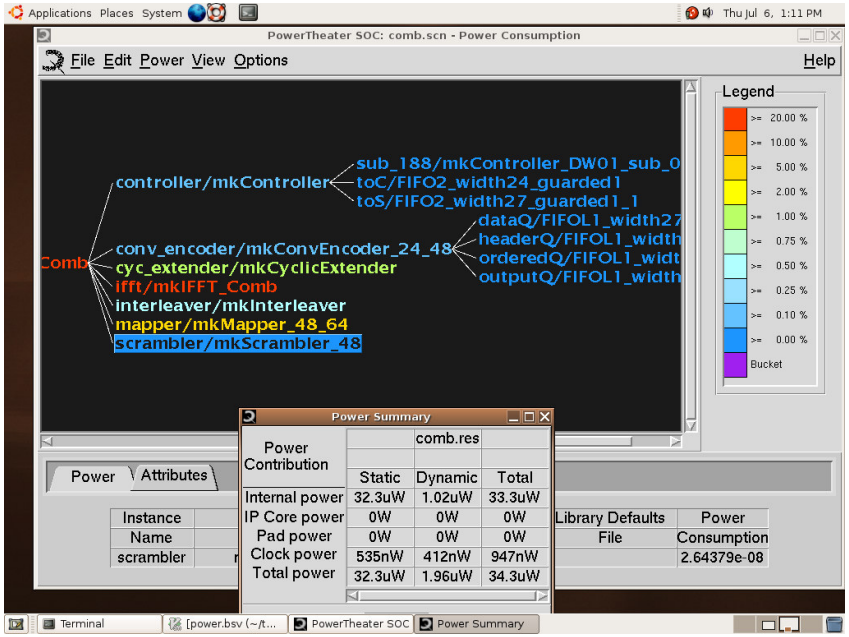

802.11a Design (by IFFT block type)	Area (um^2)	Symbol Latency (cycles)	Throughput (clks/ symbol)	Min frequency required (MHz)	Average Power (mW)
Combinational	4.91	10	4	1.0	3.99
Pipelined	5.25	12	4	1.0	4.92
Folded - 16 radix4	3.97	12	4	1.0	7.27
Folded - 8 radix4	3.69	15	6	1.5	10.9
Folded - 4 radix4	2.45	21	12	3.0	14.4
Folded - 2 radix4	1.84	33	24	6.0	21.1
Folded - 1 radix4	1.52	57	48	12.0	34.6

Optimal power

Original designer intuition

bluespec

SEQUENCE
Enabling Power-Aware SoC Design℠

ESL Study in Video Decoding

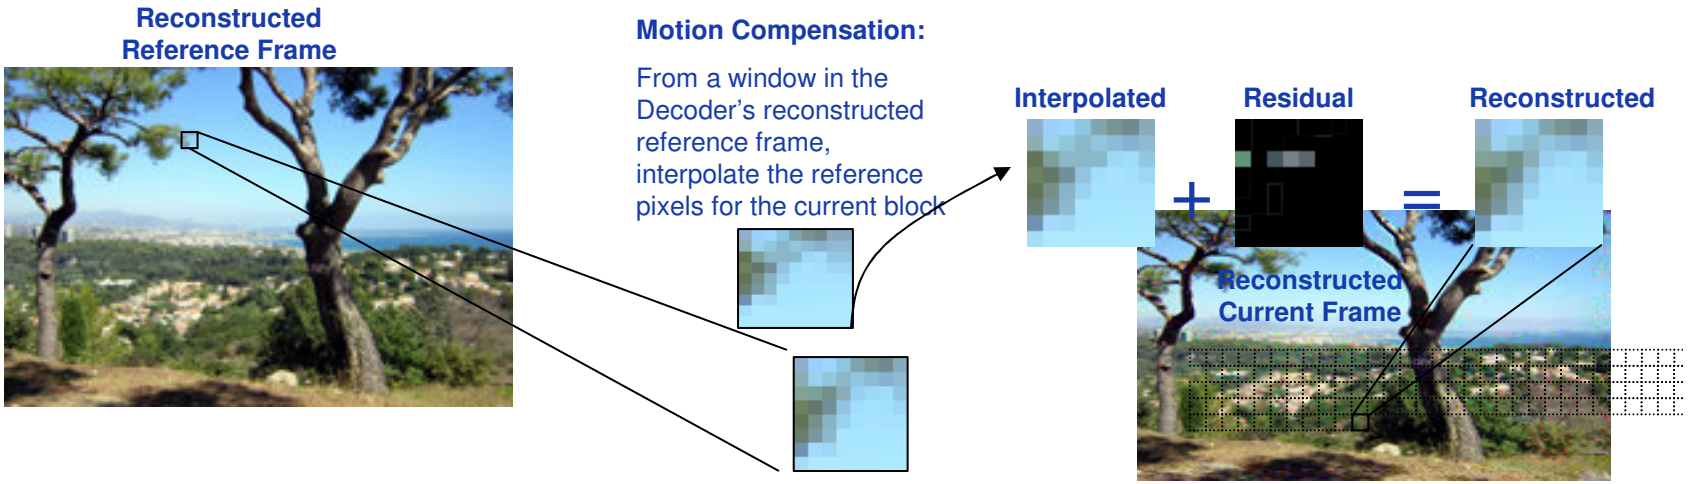

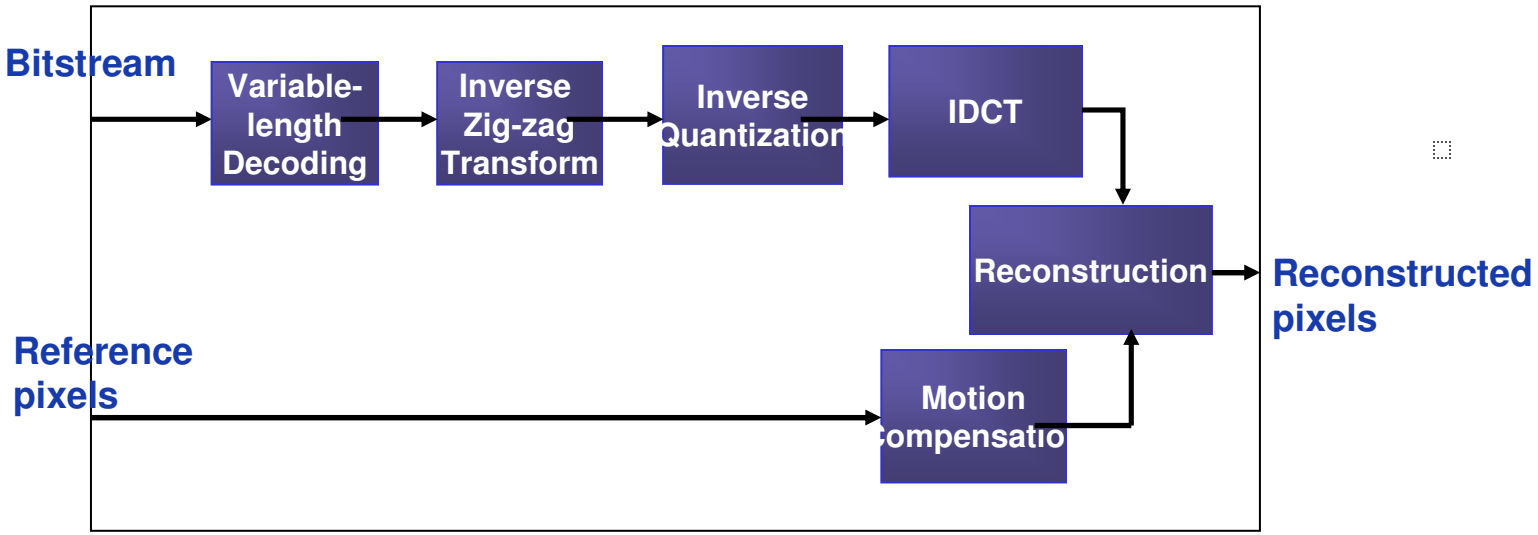

Synfora

SEQUENCE
Enabling Power-Aware SoC Design℠

ESL Synthesis, Estimation: Decoder Hardware

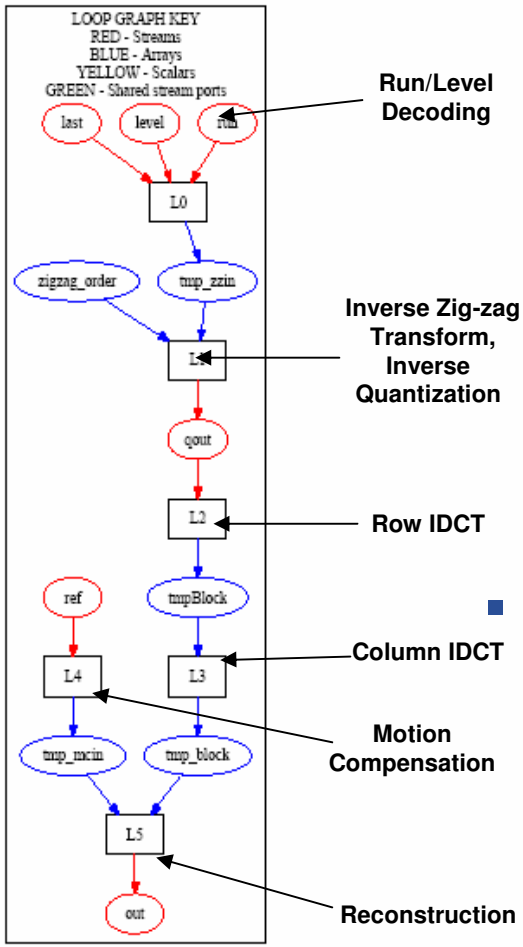

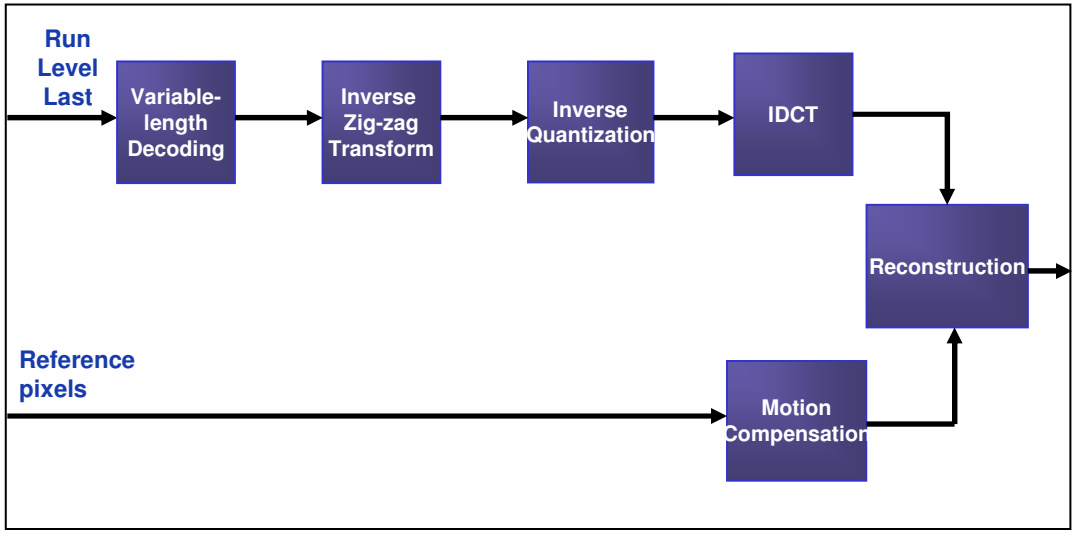

- RTL Estimate Results for “video_demo” application:
  - 40.4k gates
  - 100MHz
  - Library: TSMC 90G – Worst Case
  - 0.9V power supply
- Power analysis
  - RTL estimate = 4.79mW
  - Gate-level results = 4.40mW

Synfora

SEQUENCE
Enabling Power-Aware SoC Design℠

RTL Power Management

- **80% of chip power is determined at RTL (or earlier)**
  - **SoC power must be dealt with at RTL**
  - **Gate level appropriate for high-accuracy *verification***

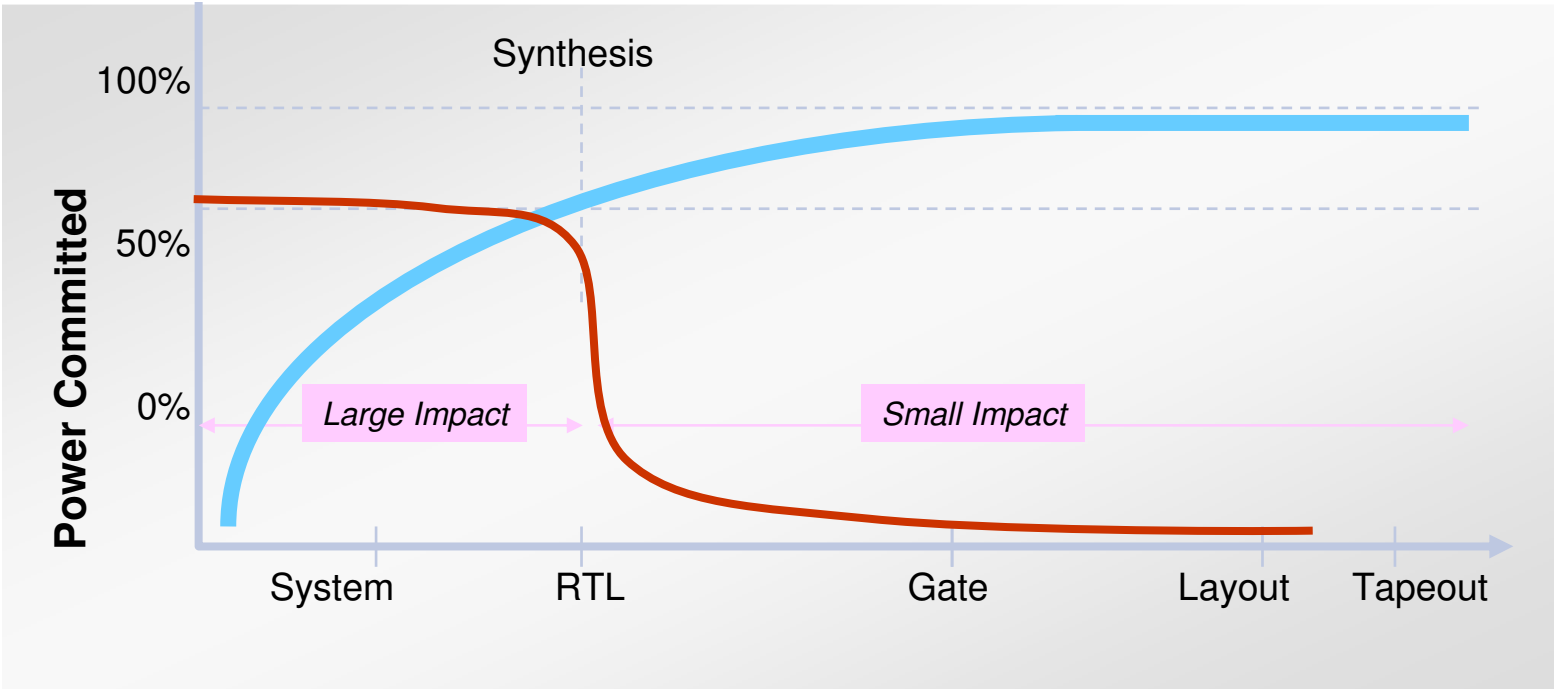

SEQUENCE
Enabling Power-Aware SoC Design℠

Power Debug Environment

- **What is critical?**
  - Architectural trade-offs not available at gate
  - Estimate block, IP and full chip power
  - Vector and vectorless modes
- **Performance and Capacity**
  - RTL: 10X gate level throughput
  - RTL abstraction / capacity
- **Accuracy**
  - Within 20% of gate
    - Clock power algorithms
    - Macro-level power modeling
    - Library, memory, IP power attributes
    - Robust power arc matching
- **Fast Power Debug**
  - Visibility and prioritization
    - Thermal map on design hierarchy tree highlights areas to investigate
    - Cross probing to source: isolate power problems
    - Detailed visual and textual reports

SEQUENCE
Enabling Power-Aware SoC Design℠

Silicon-Aware Power Management

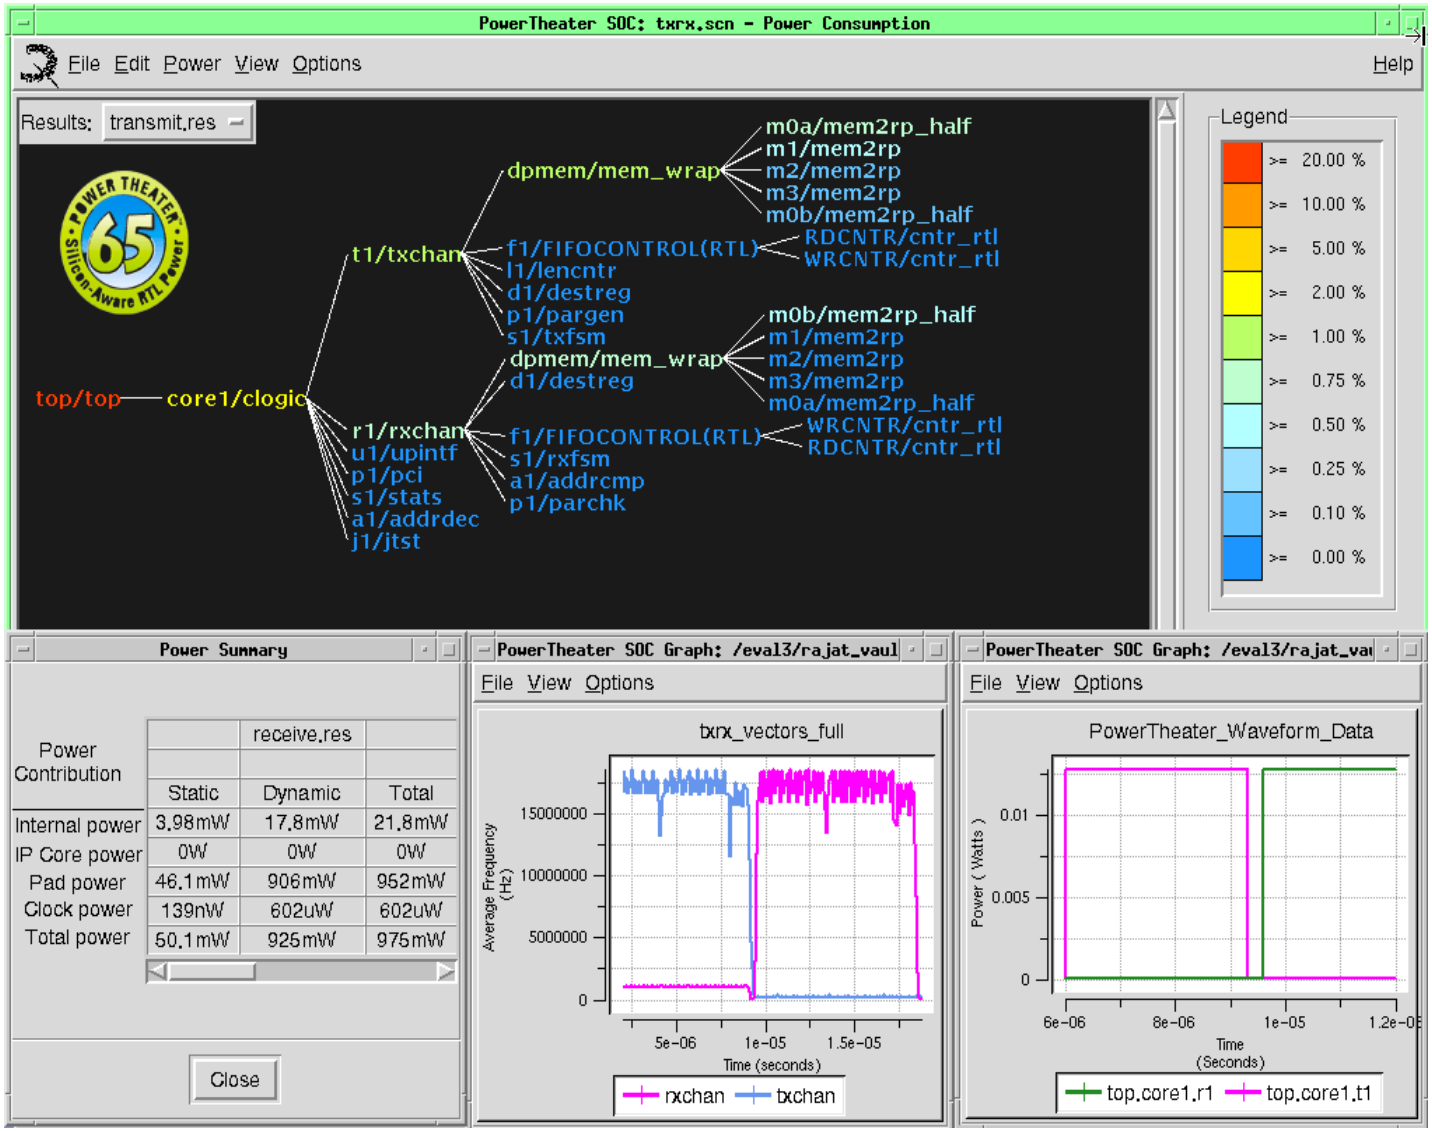

SEQUENCE
Enabling Power-Aware SoC Design℠

Software / Vectors / Modal Power Analysis

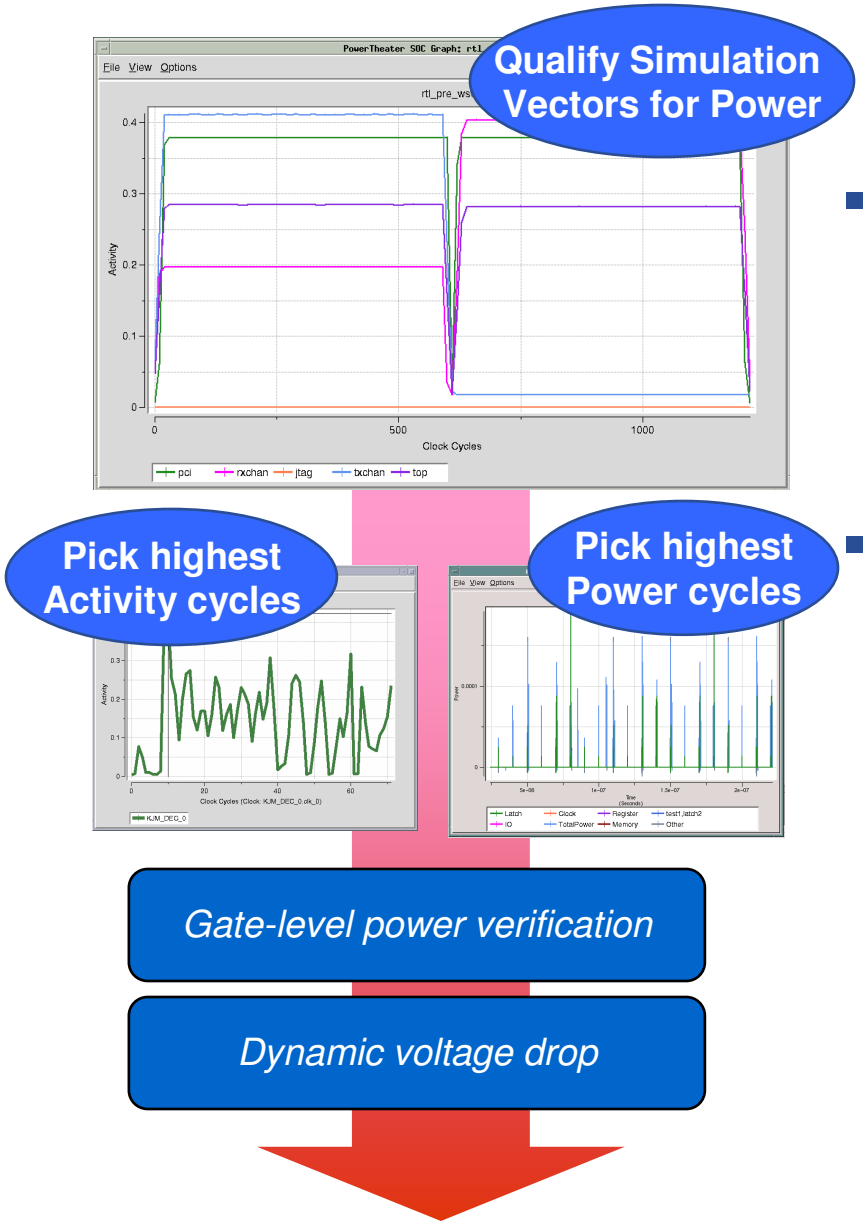

- **Modal Power Analysis at RTL**
  - Run simulation vectors for all critical modes of operation
  - Analyze power and activity per mode
  - Find modal power bugs
- **Power Vector Forward**
  - Identify and qualify worst case power cycles
  - Feed forward to gate-level power analysis
  - Feed forward to implementation DVD analysis

SEQUENCE
Enabling Power-Aware SoC Design℠

Identifying, Eliminating Wasted Power

- **Common errors, mishaps and wasted power**
  - Enabled clock toggles while data is inactive (shown)
  - Data toggles on register input while clock is inactive (register power)
  - Wasted (un-gated) clock toggles while data is inactive
  - Use clock gating cell with local explicit clock enable instead of feedback mux
  - Enable active, data / clock not
  - Mux select active, data inactive….
  - Memory splitting advisory
- **A variety of errors pertaining to**
  - Clock, datapath, control, memory, I/O
  - Muxes, clock-gated registers, memories

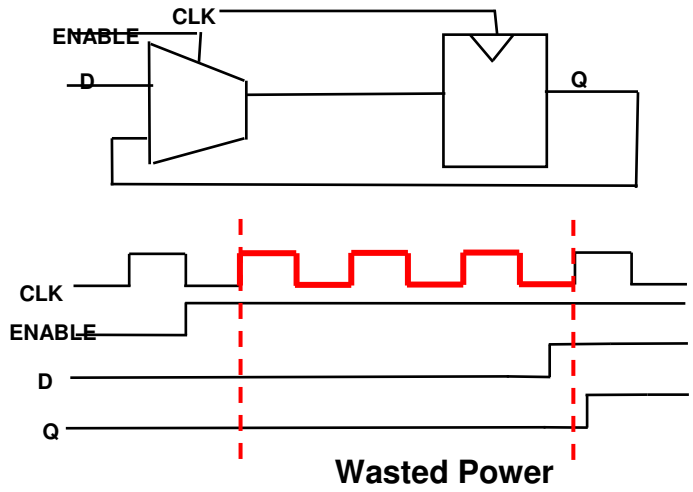

SEQUENCE
Enabling Power-Aware SoC Design℠

Eliminate Wasted Power

- **Pre-defined topology rules based on activity**
  - Reporting where, how and how much power can be reduced
    - Clock, datapath, control, memory, I/O
    - Muxes, clock-gated registers, memories

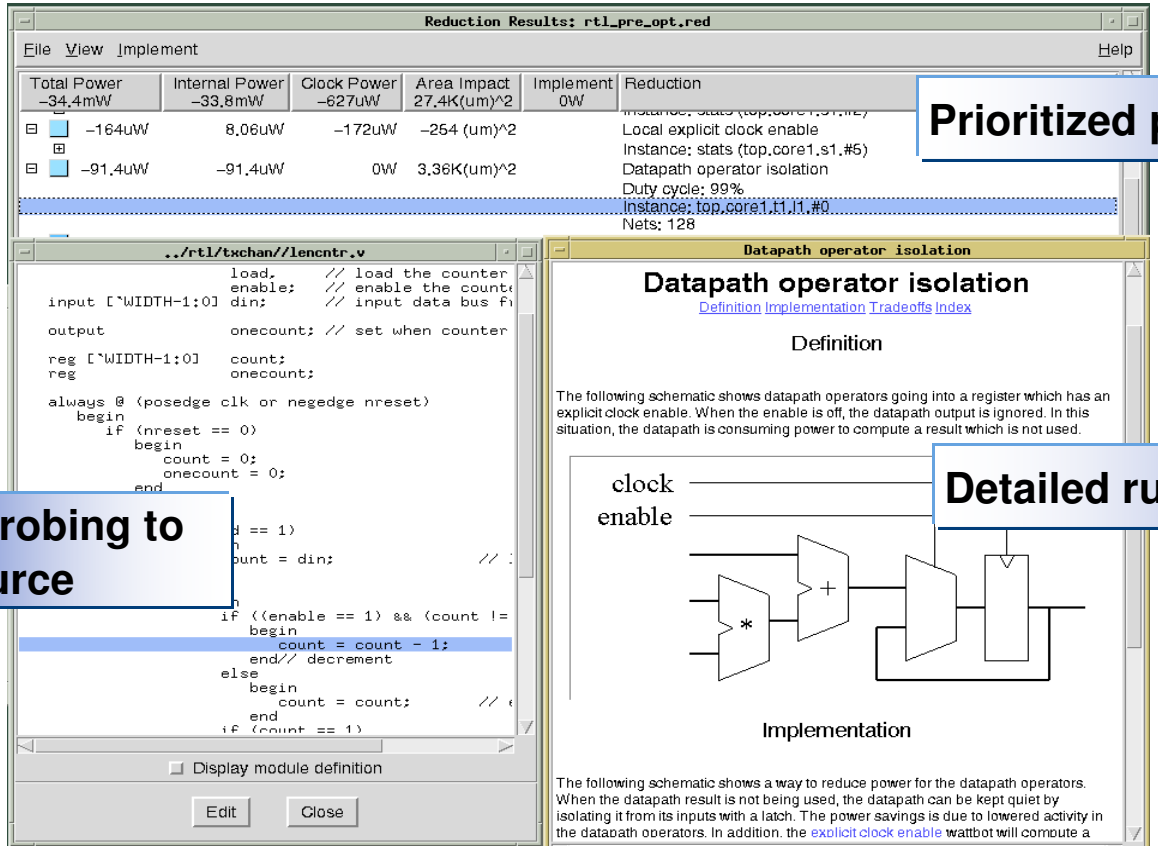

Prioritized power savings

Detailed rule description

Cross probing to RTL source

SEQUENCE
Enabling Power-Aware SoC Design℠

Ubicom Slashes Power 25%

- **Communications and Media Processor standout Ubicom**
  - StreamEngine 5000™ family
  - Chip includes
    - 10 MPUs
    - Commercial IP
    - Memory
    - 350K gates of standard cell logic
- **Reduced power in multi-core IC logic arrays by 25%**
  - RTL power analysis and optimization
  - RTL clock power analysis
  - Automated power reduction (wasted power)
- **RTL power analysis correlated to within 5% of gates**

SEQUENCE
Enabling Power-Aware SoC Design℠

Airgo Networks (now Qualcomm)

- **Single chipset that supports 802.11 a/b/g**
- **802.11 MIMO chip sets (baseband and RF)**
  - First commercial/consumer MIMO systems
- **First cost-effective True MIMO products with 2x max data rate**
  - 108Mbps in one RF channel
  - 6x to 8x rate/range performance
  - 2x3 MIMO System Architecture
- **Target: Acceptable power consumption levels**

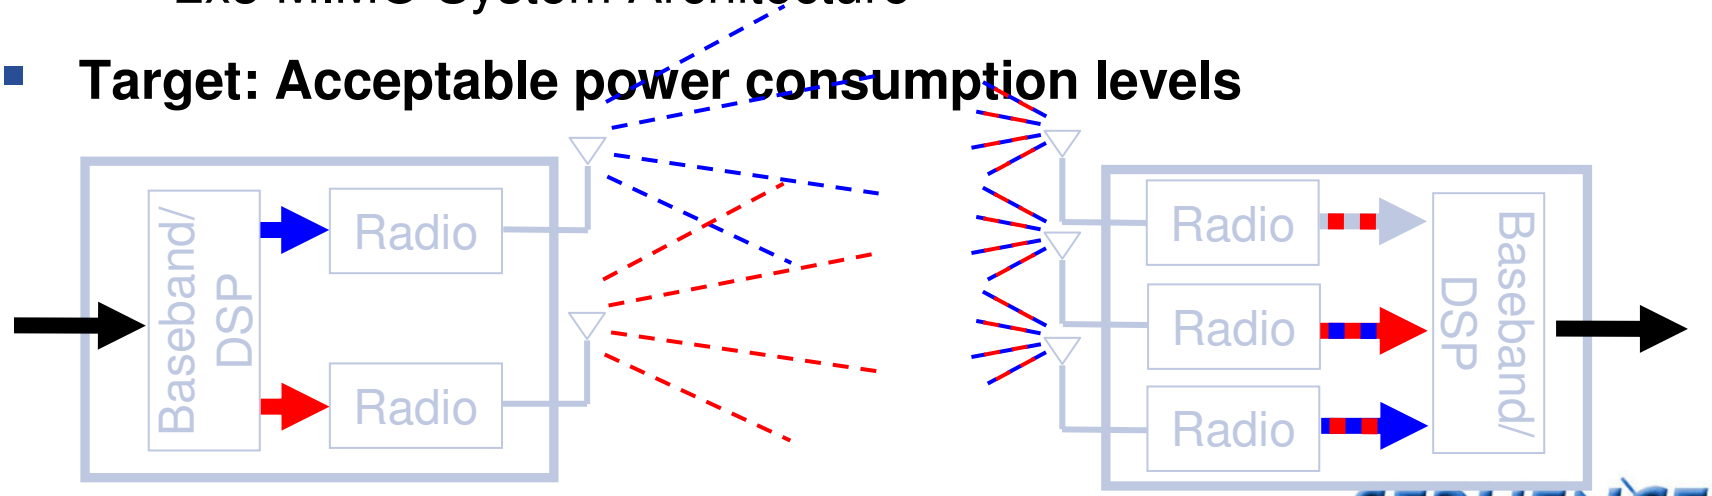

SEQUENCE
Enabling Power-Aware SoC Design℠

Wireless

"PowerTheater predicted power usage within 10% of silicon for our AGN100BB chip. We achieved this level of accuracy by leveraging PowerTheater's innovative RTL and gate-level analysis capabilities."

- Derrick Lin
Senior Director
ASIC engineering
Airgo Networks

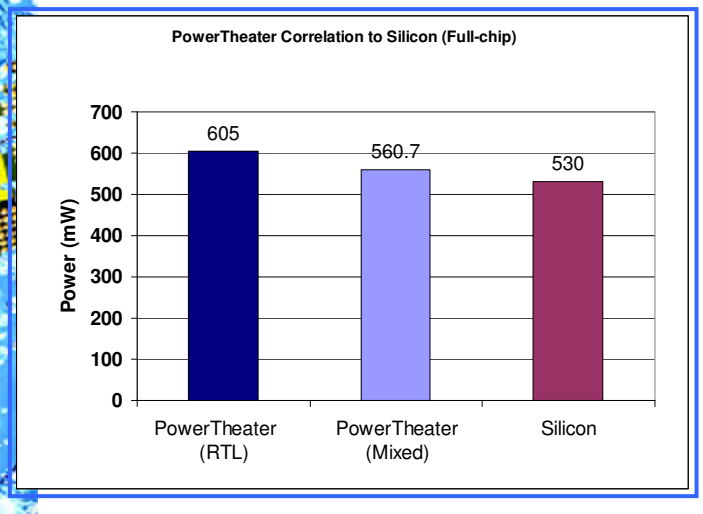

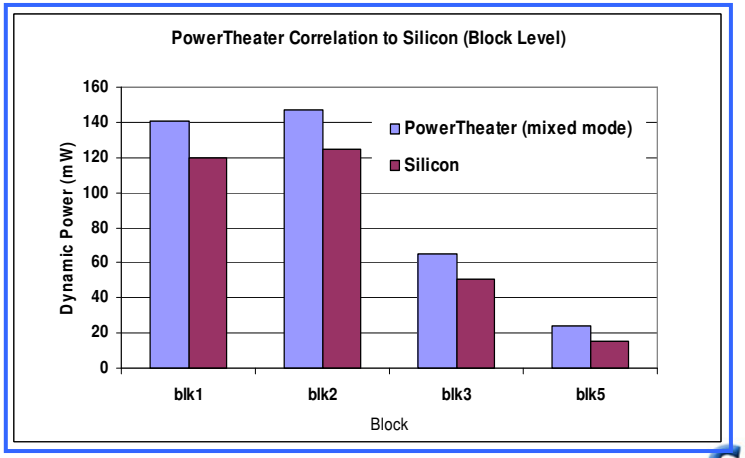

RTL Data, RTL Clock 12% of silicon
RTL Data, Gate clock ~5% of silicon

SEQUENCE
Enabling Power-Aware SoC Design℠

Airgo Methodology

- **RTL: fastest, but not the most accurate**
- **Gate: good accuracy, but too slow and a new flow for projects that don't do gate level sims**
- **40% of dynamic power consumed in the clock tree**
  - Estimate clock tree accurately (SDF back annotated simulations along with SPEF data)
  - Estimate "data" power using RTL simulations
  - Add the two numbers up (using clock activity factors from RTL) to get the "average" dynamic power
- **Determine a good "power vector" based on design knowledge**
- **Modeling requirements**
  - ASIC library has to be characterized for power
  - RAMs and IOs also have to be characterized for power
  - Wireload models increase data power accuracy
- **Taped out several chips with comfortable margins**

SEQUENCE
Enabling Power-Aware SoC Design℠

Popular Power Management Techniques

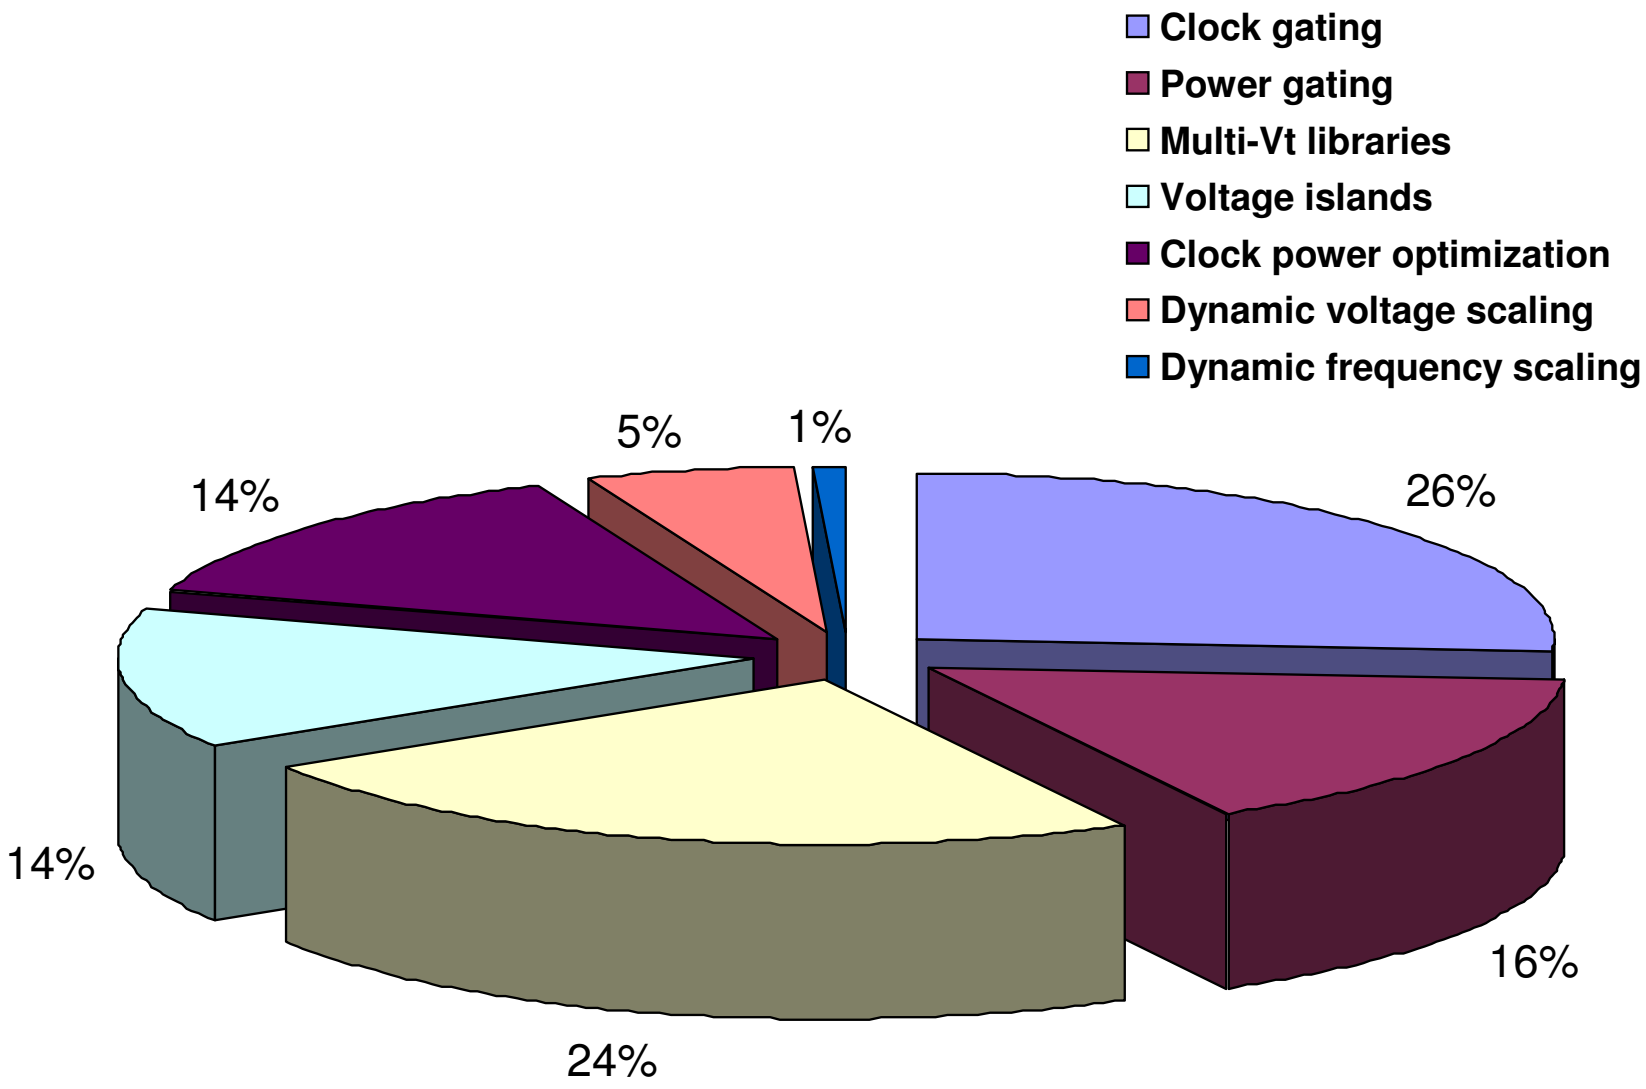

Survey Summary of End Users to DAC Suite 2006
Sample size = 115

SEQUENCE
Enabling Power-Aware SoC Design℠

Clock Power and Clock Gating

- **Clock power is significant**
  - Frequently 40-50% of total active power
  - Clock and clock tree
- **Clock gating**
  - Explore Power Savings: >25% of clock power

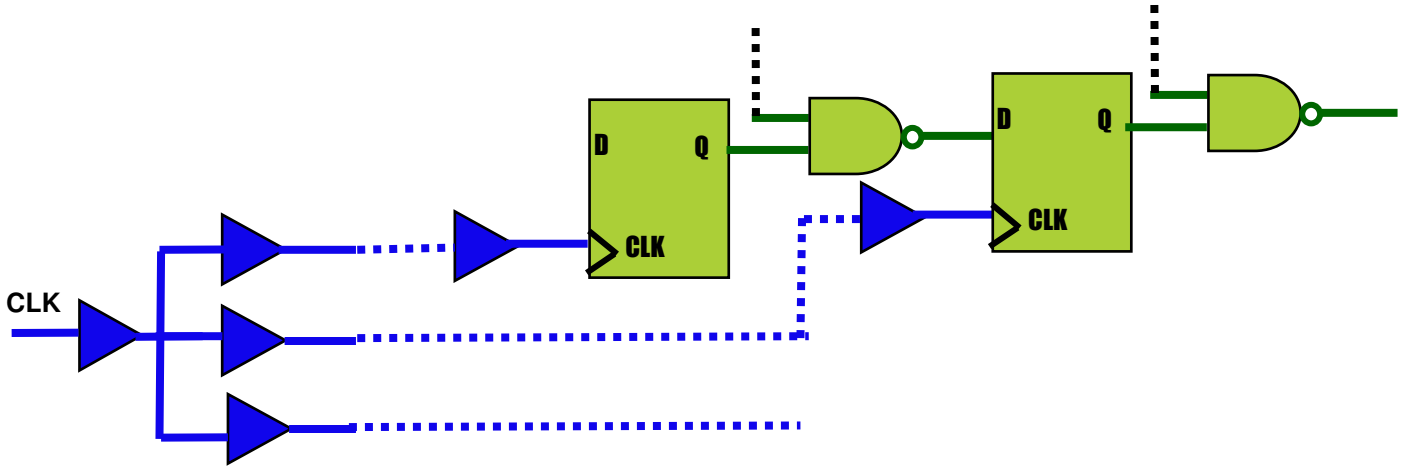

Clock tree

Clock domain without clock tree

Clock domain with clock tree

SEQUENCE
Enabling Power-Aware SoC Design℠

Reporting Clock Domain Power

VIM - ~/processor.rpt

```
7. Clock domain power consumption
=================================

1.  Clock Net: myprc_platform_256k.processor_clk1

    Instance Summary:
    ================

                                    Power(Watts)
                                Static  Dynamic    Total
                                ------  -------    -----
    Register power              34.7uW   1.72mW   1.75mW
    Latch power                 2.44uW     13uW   15.4uW
    Memory power                82.3uW    441uW    524uW
    IP core power                   0W       0W       0W
    Pad power                       0W       0W       0W
    Clock power                 916nW   1.28mW   1.29mW
    Other power                 161uW   4.91mW   5.07mW
    Internal load power             0W   33.7mW   33.7mW

    Total power                 281uW   42.1mW   42.4mW


    Traced instances:
    Library                 Clock  Driven By  Driven Net Driven        Power (Watts)          Traced
    Model                   Level  Net Number  Numbers   Loads   Static  Dynamic    Total    Instance
    -----                   -----  ----------  -------   -----   ------  -------    -----    --------
    sdffpr_4                1        3          16319     1      14.1nW    4.28uW     4.3uW  myprc_platform_256k.myprc_pl
atform.myprc_core.Proc1136JFS.uClkGates_uWFIPendSync_Sync1_reg
    sdffpr_a11p_1           1        3          16318     1      9.35nW    4.35uW    4.36uW  myprc_platform_256k.myprc_pl
atform.myprc_core.Proc1136JFS.uClkGates_uIROWFISync_Sync1_reg
```

- Power-per-category of the clock domain
- Power-per-gate-instance of the traced clock domain

SEQUENCE
Enabling Power-Aware SoC Design℠

Hierarchical Clock Gating

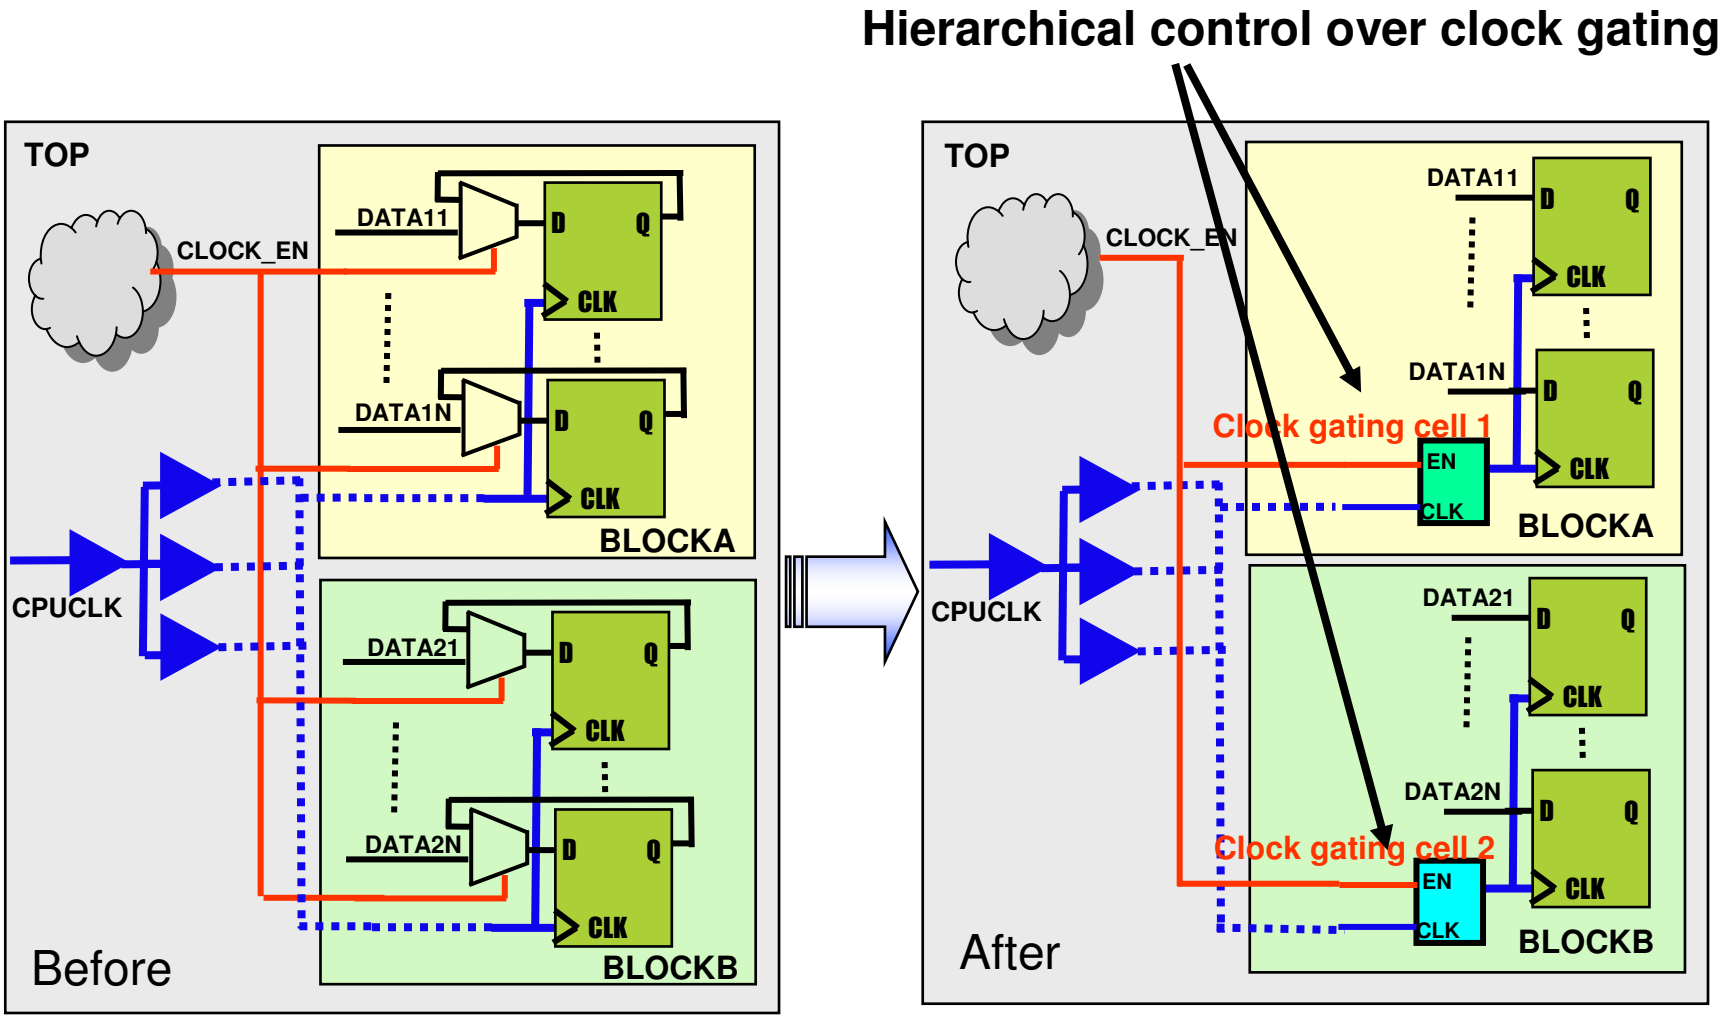

SEQUENCE
Enabling Power-Aware SoC Design℠

Multi-Vt

- **The good:**
  - Multi-Vt libraries can save approximately 5-15%
  - Multi-Vt is not a challenge; easy to do
- **The bad:**
  - Cost of multi-Vt libraries
  - Timing issues
  - Signal integrity
- **The tradeoffs are generally obvious…and tools exist…**

Voltage Islands

- **The good**
  - Voltage islands can save approximately 10-50%
- **The bad**
  - Cost of characterization of libraries
  - Area penalties
    - Power grid
    - Level shifters
  - Performance degradation
    - Timing
  - Complexity
    - Timing, SI, voltage drop
- **The ugly**
  - Design, verification, implementation tools are not integrated or robust

Power Gating...Power vs Penalty

- **The good**
  - Power gating can save approximately 10-1000X leakage
- **The bad**
  - Increases complexity
    - Rush current problems; wake-up time, switch sizing and sequencing
    - Timing, SI, voltage drop
  - Area penalties
    - Power grid
    - Level shifters, isolation cells
  - Impacts performance and creates timing issues
- **The ugly**
  - Again, design, verification, implementation tools are not integrated or robust

"Leakage will become a major industry crisis, threatening the survival of CMOS itself"
ITRS2005 Executive Summary

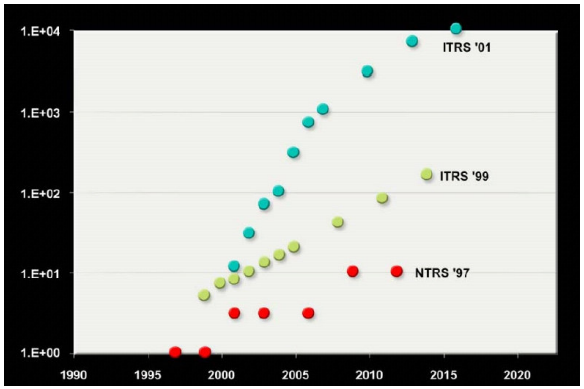

SEQUENCE
Enabling Power-Aware SoC Design℠

Cradle Architecture

- **Capable of real-time encoding**
  - 16 channels of MPEG4 SP@L3
  - 4 channels of MPEG4 ASP@L5
  - or 1 channel of H.264 Main Profile D1
- **55 Million Transistors (In-House RISC and DSP processor design), 180K Flops, 4 clock domains**
- **24 processing cores**
  - 16 DSPs and 8 General-Purpose Processors (GPPs)
- **A smart I/O subsystem providing up to 144 fully programmable I/O pins**
- **A DDR SDRAM interface to support high data throughputs for high-definition video processing**

Cradle Architecture

- **Low power consumption across the family**
  - As low as 1.5 W
- **Loosely coupled multiprocessor architecture enables more efficient system level performance:**
  - Megapixel sensor interfaces, image enhancement, data encryption, video/audio Codecs, complex network stacks and system

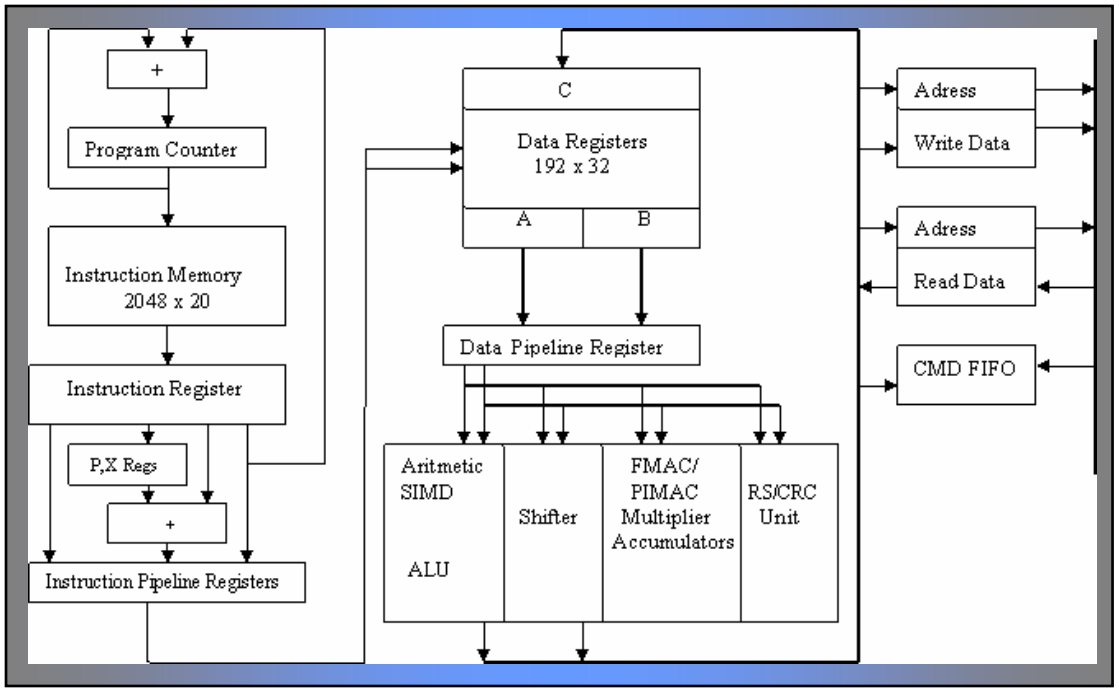

Cradle DSP (DSE) Processor

SEQUENCE
Enabling Power-Aware SoC Design℠

Power DSE Results

Test	Total/Clock (mW)	sub-module power consumption (mW)			
		IF	MAC	RF	DF
1 Directed Test	174 / 61	35.09 (20.13%)	27.5 (15.8%)	20.92 (12.45%)	8.26 (4.91%)
2 Directed Test	173 / 61	34.42	27.3	20.68	8.61
3 Directed Test	169 / 61	34.65	23.1	20.88	8.38
4 MAC usage	194 / 61	34.7	41.5	21.29	11.05
5 MAC power down mode	138 / 40	25.01	4.84	21.99	7.15
6 Application	159 / 61	33.76	19.9	20.09	6.64

Power Estimation at RTL: DSP at Cradle

PT Accuracy

CRADLE TECHNOLOGIES

Design	PT Power	Actual Silicon Power	
DSE Block	160 mW*	150 mW*	(8% less)
Full Chip	5.3 W**	4.9 W**	(8% Less)
Clock Tree	2.17 W	1.95 W	(10% less)

* Block power depends on test and activity. It ranges from 138mW to 194mW
**Full Chip power is based on specific vectors and does NOT represent overall power in applications

55 Million Transistors (In-House RISC and DSP processor design) , 180K Flops, 4 clock domains
24 processing cores, including 16 DSPs and 8 General-Purpose Processors (GPPs)

SEQUENCE
Enabling Power-Aware SoC Design℠

Design for Power!
Simple Power Saving Techniques

- **Power debug**
  - Apply power reduction schemes first to the sub-modules that consume power mostly
  - Determine and eliminate “hot spots”
- **Clocks**
  - Clock tree consumes 40 – 50% of total power; reduction scheme is very important
- **Vectors**
  - Develop accurate power vectors that exercise all possible nodes
- **Write RTL and optimize for power**
  - Shut off data switching when in idle state
  - Provide chip enables and output enables effectively
  - Use gated clocks to all data flops
  - Use multiplexed flops to change data when enables are set
  - Use “Grey Code” scheme for FIFO pointers and memory addresses
- **Memories**
  - Power-efficient memory selection is key
  - Replace flop cluster with custom or standard compiled memories
  - Try different memory configurations (1 bank, 2 bank, different aspect ratio etc).
  - Shut off clocks to all memories when not in use: Memory Chip Enable
  - Splitting memory can reduce power as much as 30%
- **Multi-VT libraries reduce power and leakage**

Power Regression Testing Methodology

Tensilica Configurable and extensible microprocessor cores for embedded SOC designs

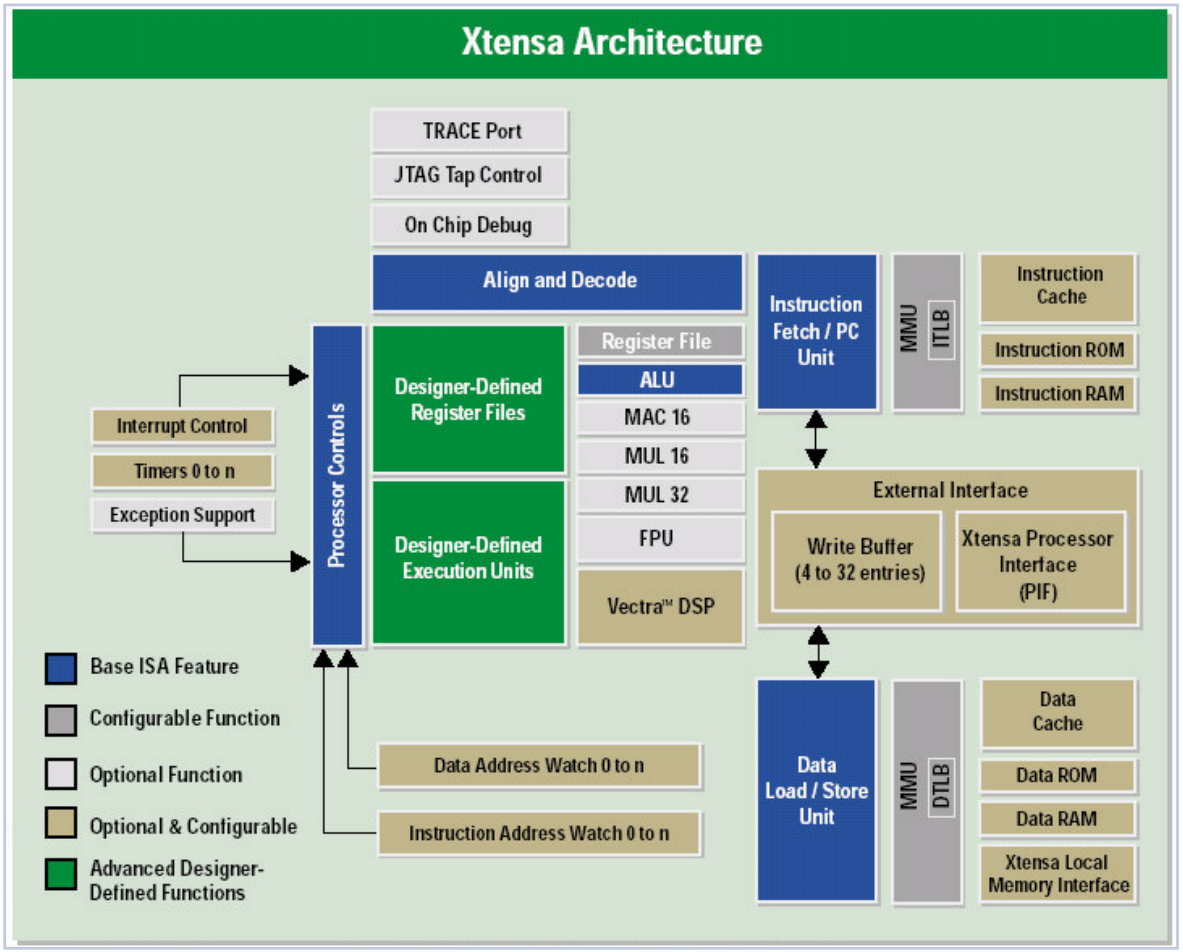

You Can't Fix What You Can't Measure!

SEQUENCE
Enabling Power-Aware SoC Design℠

Xtensa-LX: Low Power a Primary Goal

- **Design Goal**
  - Reduce power dissipation by 25% compared to the previous generation design
- **Previous generation already optimized for low power operation**
  - Must work on lowering power during early design phase
- **Special challenges for configurable IP cores**
  - Configurable cores have numerous combinations to test
  - Soft IP characterized for various fabrication processes
  - Requires database of area, timing, and power numbers
- **Need a methodology for monitoring power dissipation on a regular basis, with meaningful feedback to designers**
- **RTL advantages over gate**
  - Debug visibility
  - Performance and capacity for long simulations

SEQUENCE
Enabling Power-Aware SoC Design℠

Power Regressions

- **Measure effectiveness of clock gating**
  - Xtensa processor employs global & functional clock gating
- **Tune DSP extensions for low power**
  - Profile power dissipated executing common DSP kernels
  - Tune assembly code and hardware implementation to meet aggressive power goals
- **Guard against any undue increase in power**
- **Generate characterization data for the Xtensa "Processor Generator"**

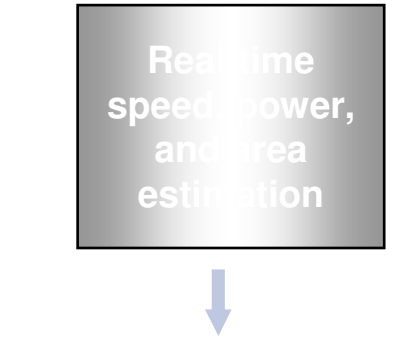

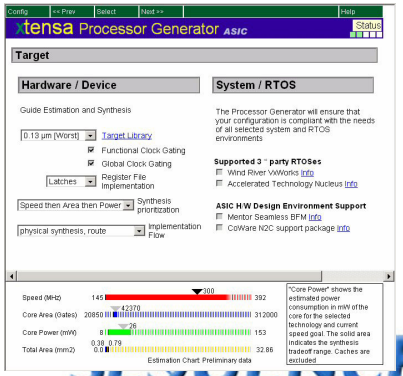

Enabling Power-Aware SoC Design℠

Weekly Power Regression Tracking

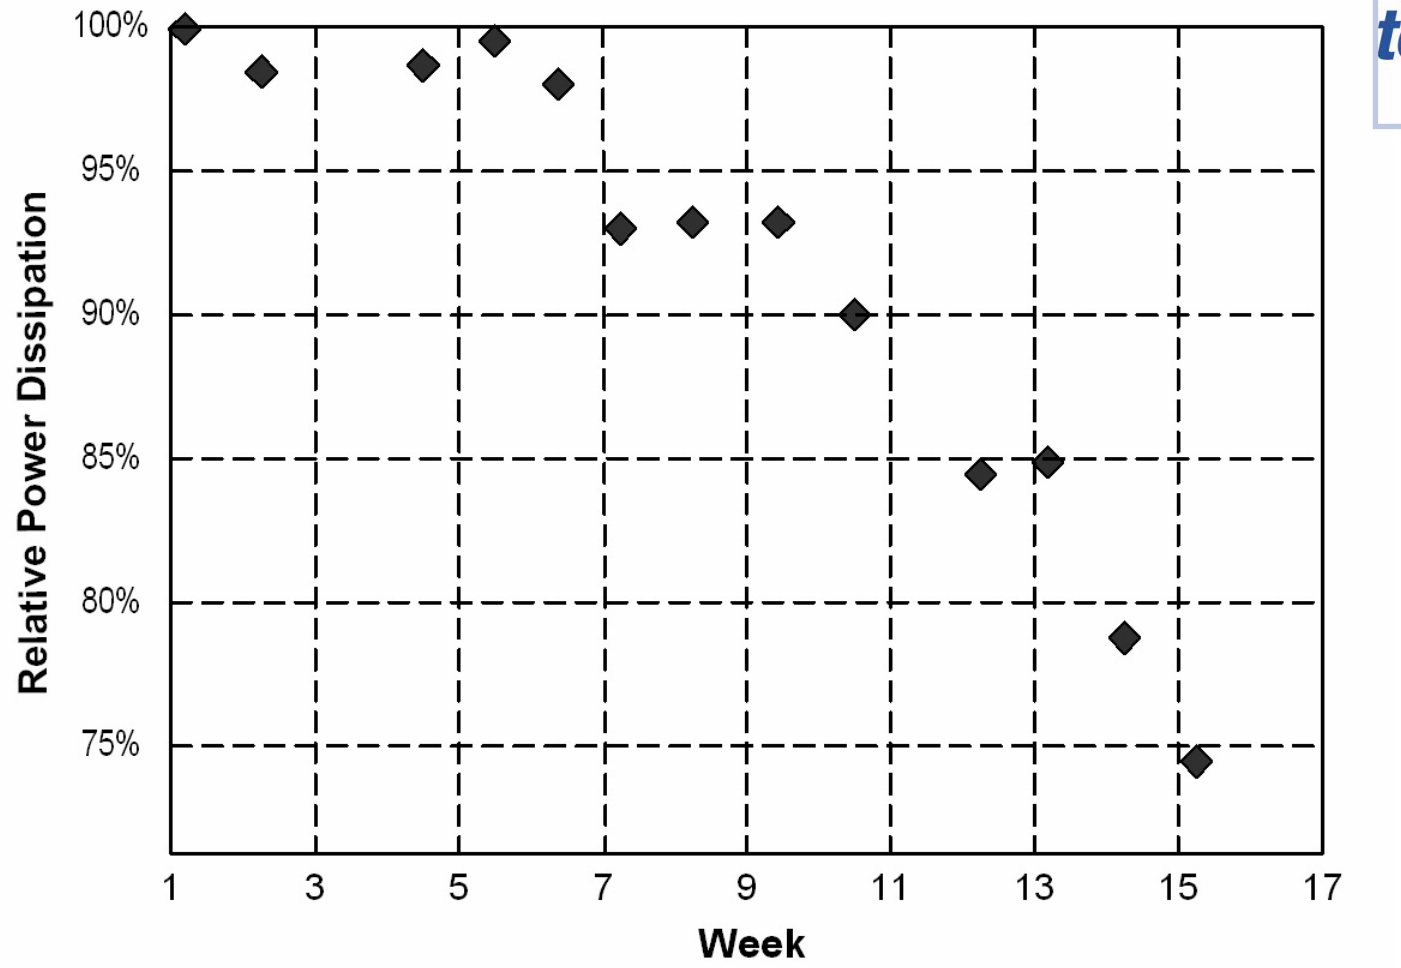

Goal met: 25% power reduction in 15 weeks

Configurable, extensible processor cores for embedded SOC designs

Verify and Refine Power at Gate Level

- **Eliminate power creep**
  - Enhance productivity and ease of debug
  - Gate-level power analysis using RTL simulations
  - Power vs time waveforms correlated to event waveforms

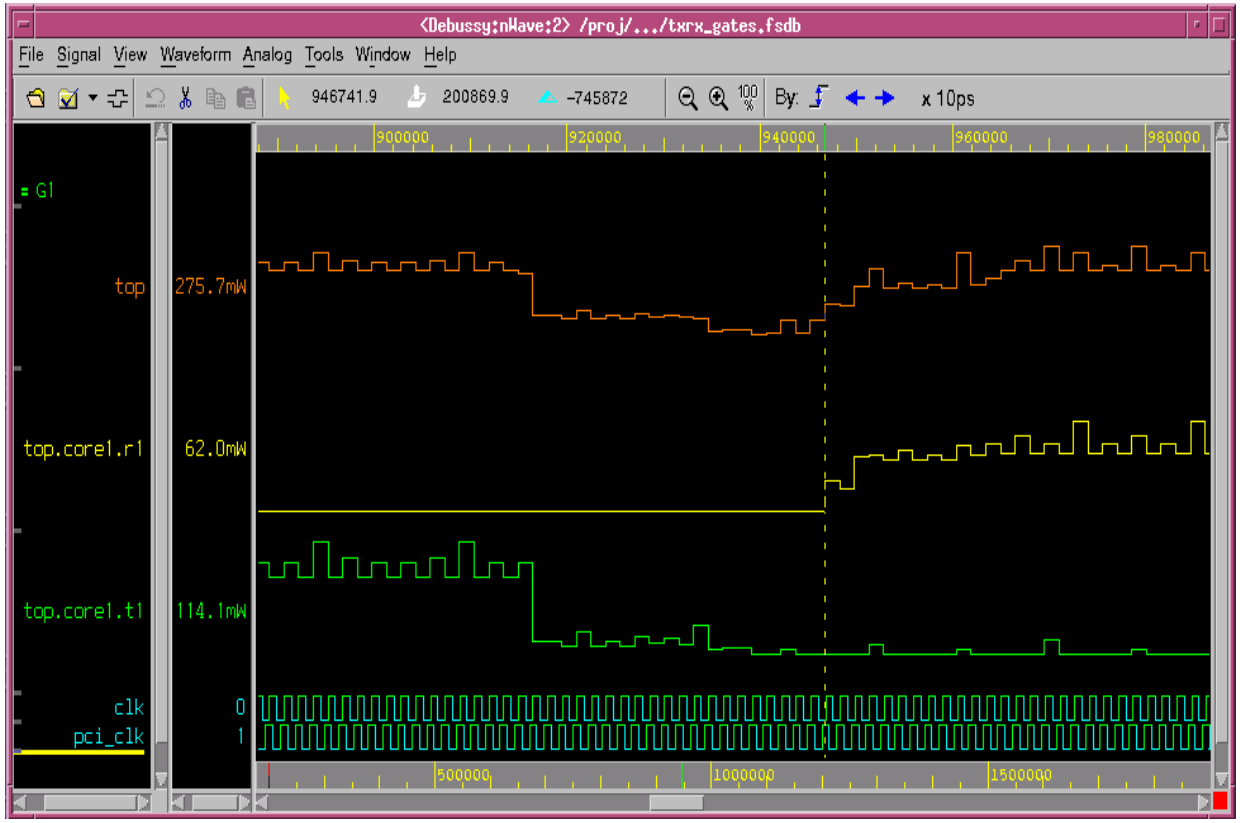

SEQUENCE
Enabling Power-Aware SoC Design℠

Power Reduction Through Silicon-Aware Techniques

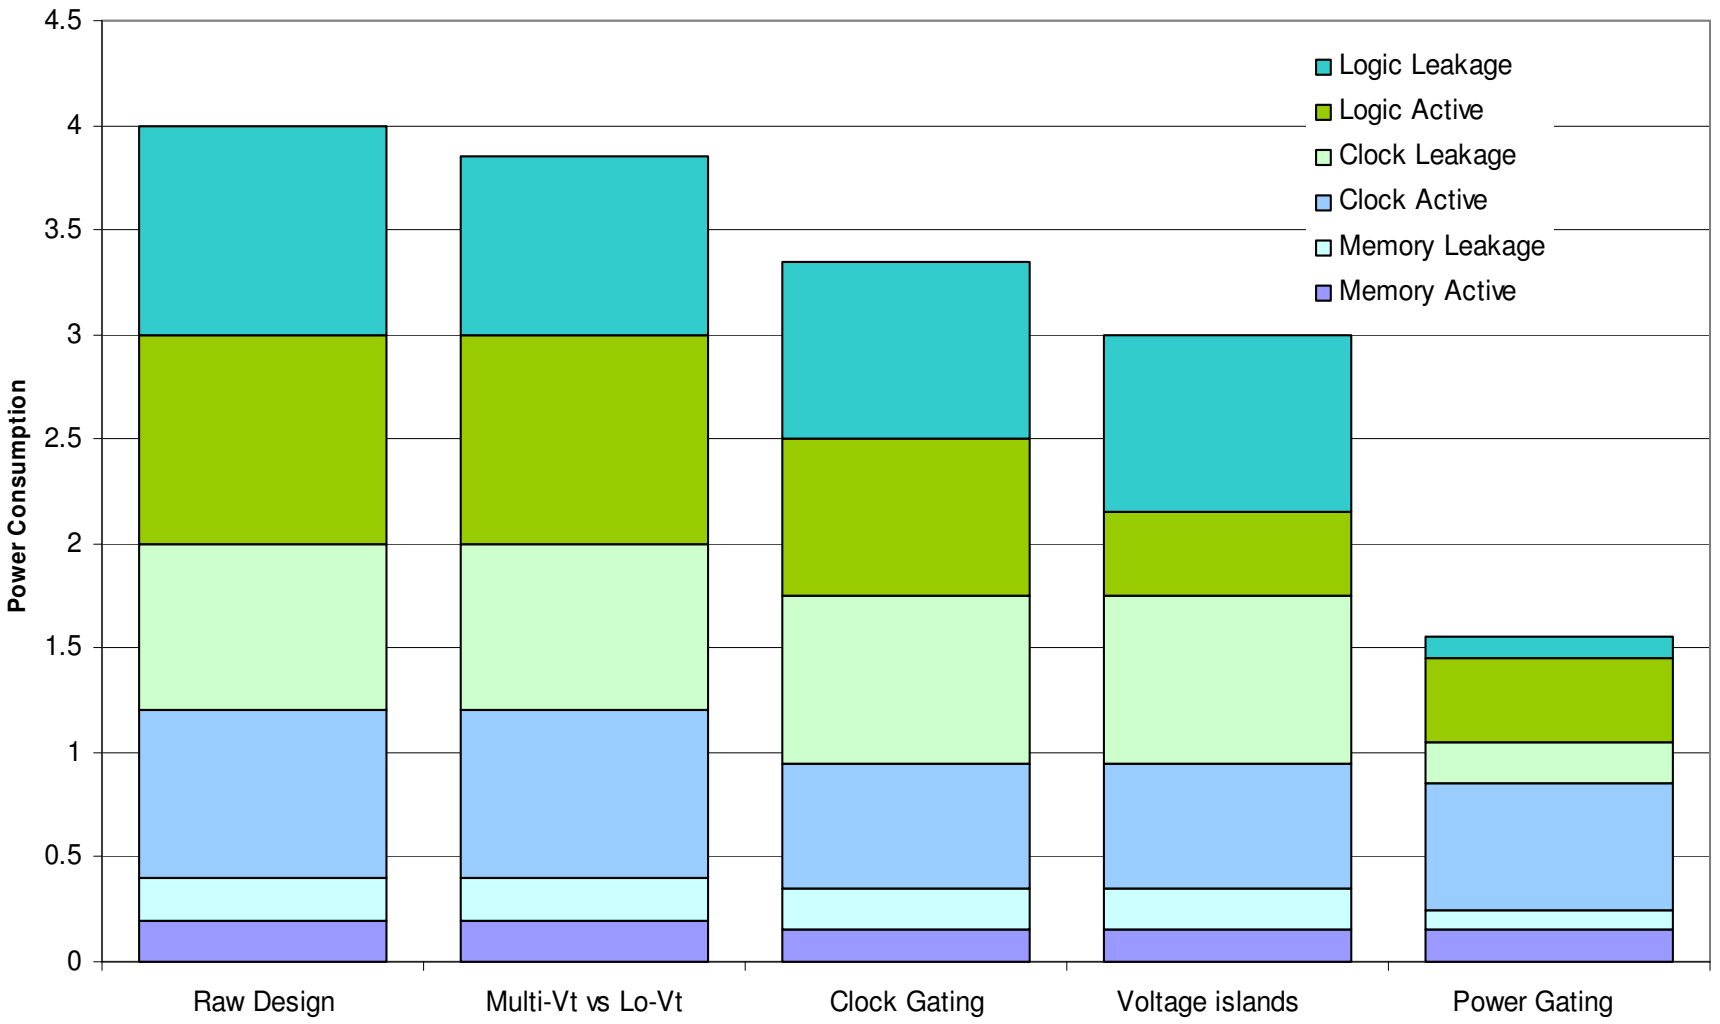

SEQUENCE
Enabling Power-Aware SoC Design℠

Emerging Low Power Standards

- Design for power intent
- Enabling innovation

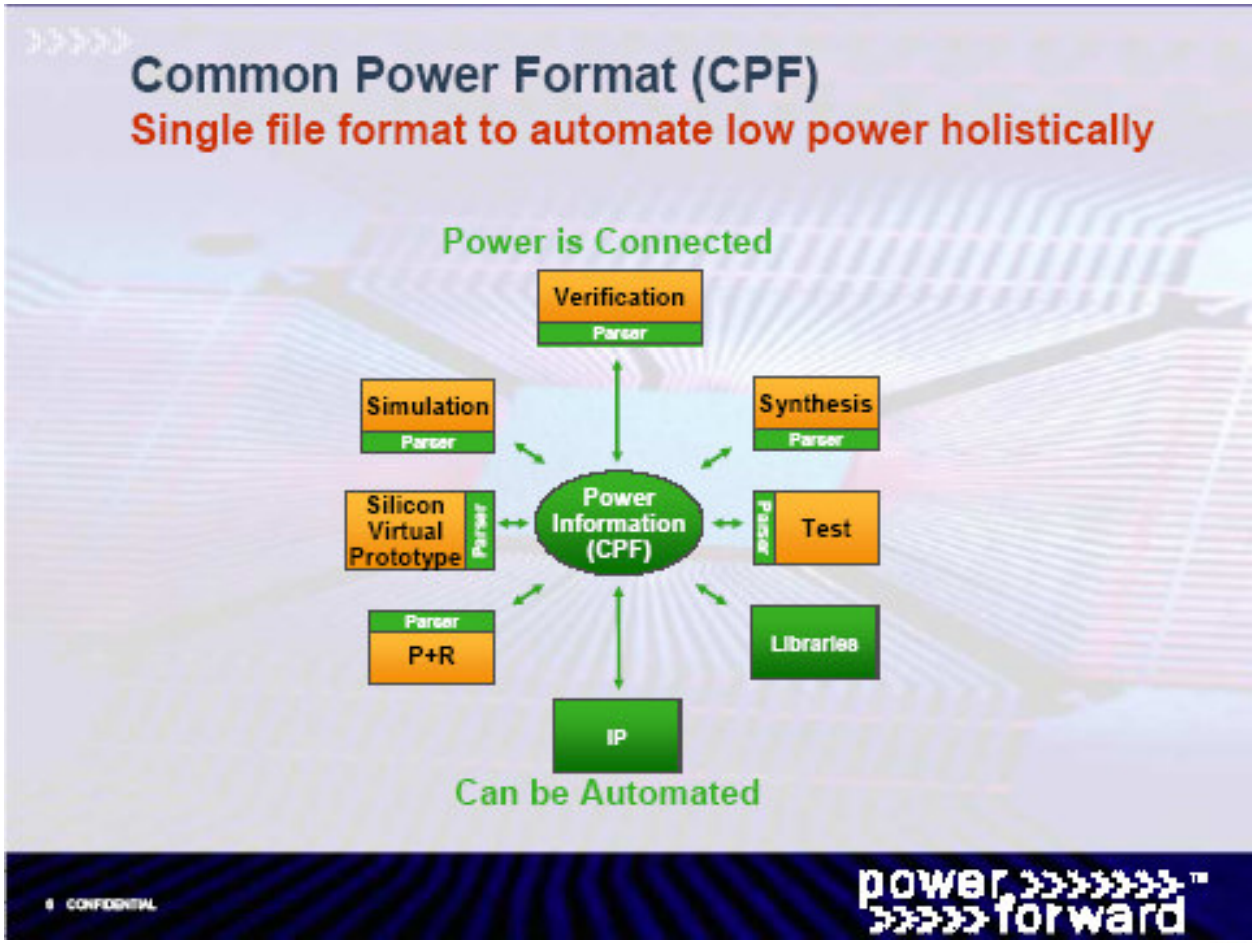

SEQUENCE
Enabling Power-Aware SoC Design℠

Best Practices Summary

- **Design for Low Power Intent!**
- **Architectural exploration has great impact**
- **Verify power early and often**
- **Prioritize "power offenders" and take corrective action during the RTL design phase**
- **Choose worst case power vectors from each mode of operation**
- **Eliminate "wasted" power**
- **Run "power regressions" throughout – RTL to tapeout**
- **A single low power standard: to unify the methodology flow**

- **Control the true costs of power…..**

SEQUENCE
Enabling Power-Aware SoC Design℠