



A 32 nm, 3.1 billion transistor, 12 wide issue Itanium® Processor for Mission-Critical Servers


Reid Riedlinger¹, Rohit Bhatia¹, Larry Biro²,
Bill Bowhill², Eric Fetzer¹,
Paul Gronowski², Tom Grutkowski¹

¹Intel® Corporation, Fort Collins, CO

²Intel® Corporation, Hudson, MA

Agenda

- Poulson Overview
- Core enhancements
- System interface overview
- Server challenges
 - Minimum voltage operation
 - Frequency improvements
 - Power reductions
 - Core asymmetries
- RAS Improvements



Poulson in a system




Poulson Package


Processor Highlights

- New chip micro-architecture
 - Enhanced Power measurement system
 - Socket compatible with Tukwila
- 8 Hyper-Threaded 64 bit cores
 - Significant architectural enhancements
- 32 MB Last Level Cache
 - Intel® Cache Safe Technology
 - 54MB on-board SRAM and Register File Storage
- Improved Memory and System I/O
 - 33% bandwidth improvement
- On die Ring interconnect
- Improved RAS with twice the cores

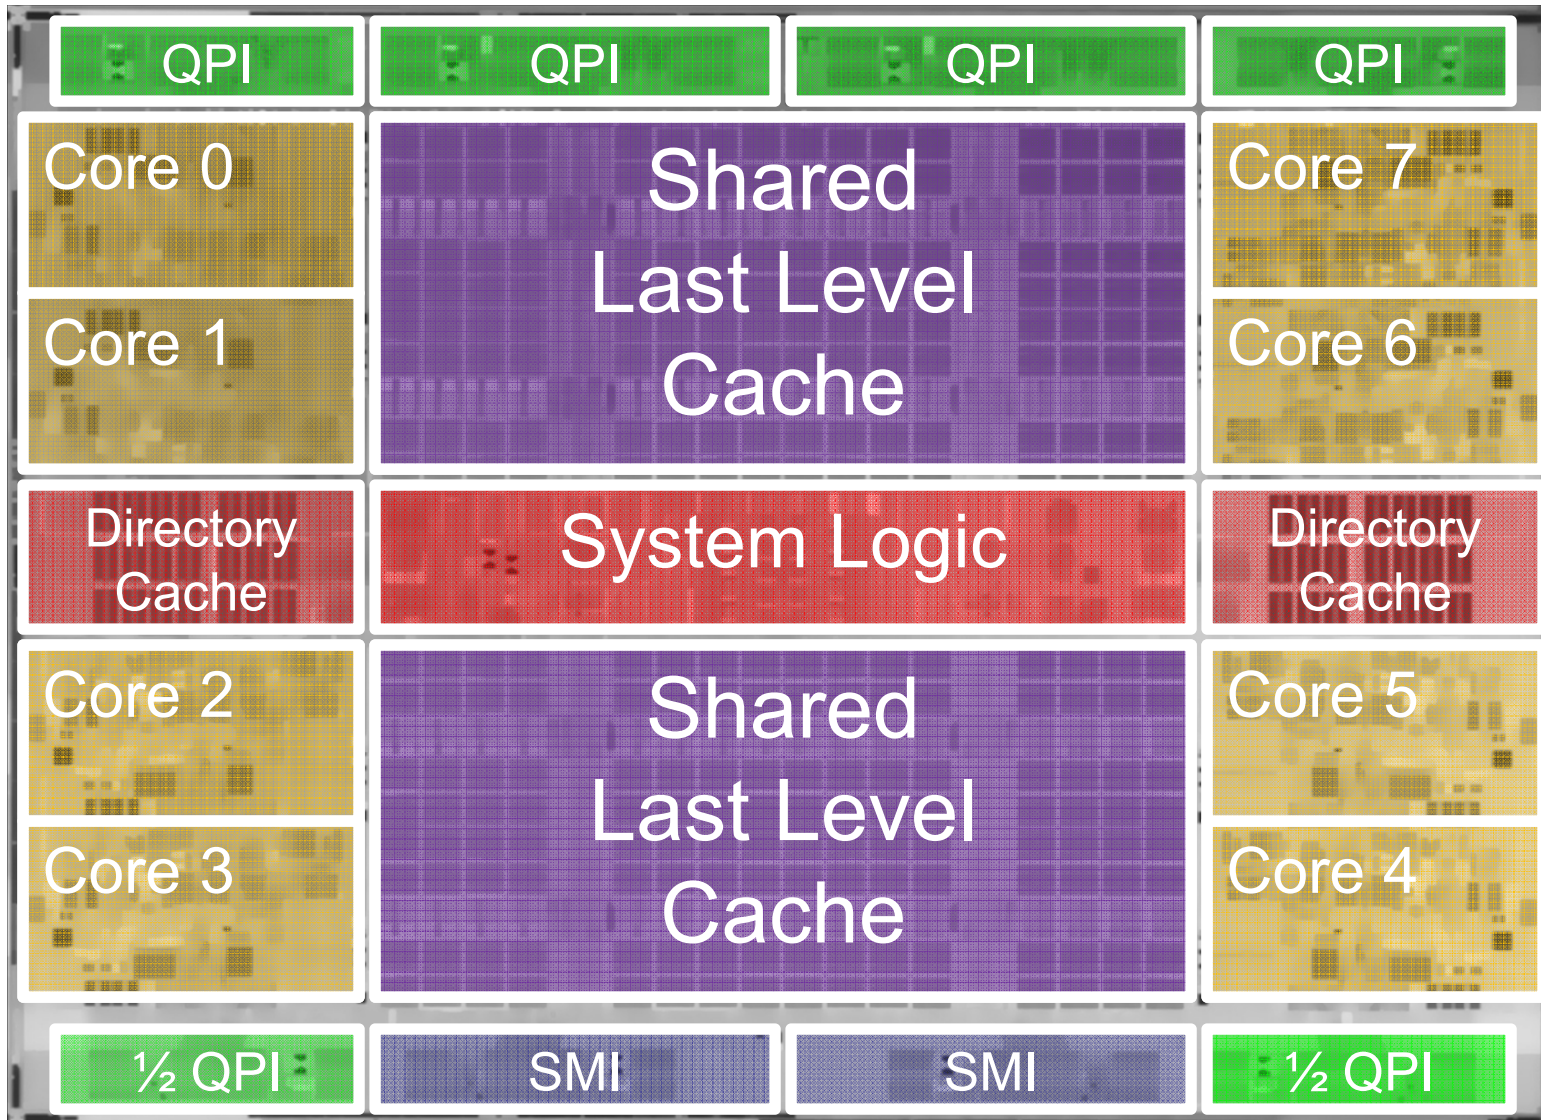
Poulson



Tukwila



Processor Overview



New Core Architecture with Ring based system interface

Chip Statistics

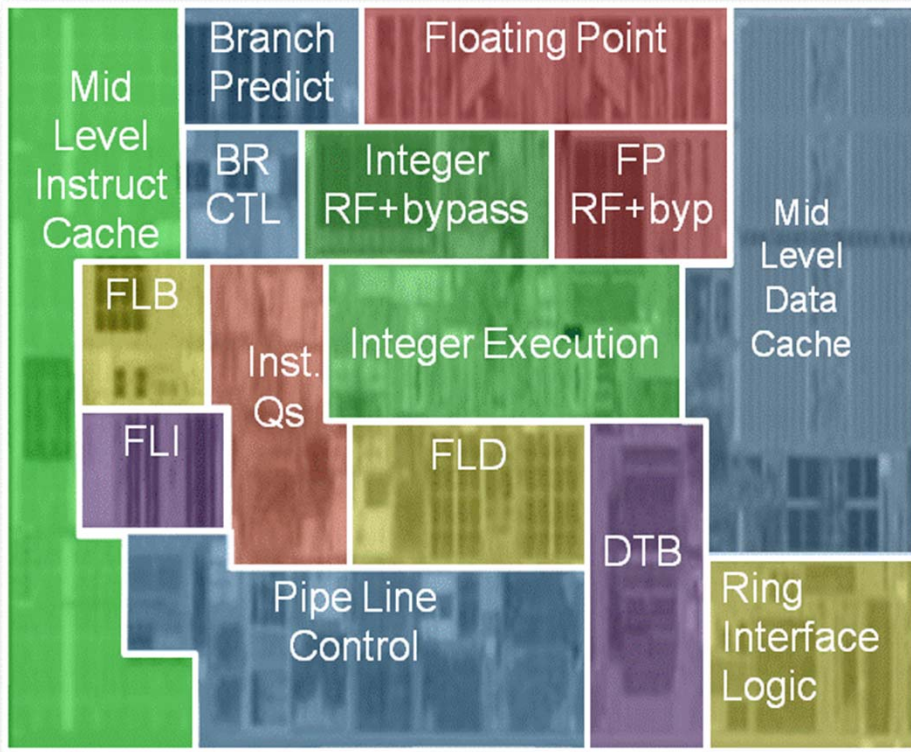
- 32nm bulk CMOS, 9 layer Cu interconnect
- 8 cores with 3.1 billion transistors
 - 29.9 mm X 18.2 mm = 544 mm² die size
 - 170 Watt max TDP
- 6 voltage and 4 frequency domains

Circuit	Devices (million)	Area (mm ²)	Voltage (Volts)	Power (Watts)
Core logic	712	158	0.85 -- 1.2	95
LLC cache	2,173	163	0.90 -- 1.1	5
SysInt logic	224	137	0.90 -- 1.1	50
IO logic	44	68	1.05 -- 1.1	20

Core Design

- Micro-architecture and floor plan optimized for future process generations
 - First comprehensive redesign of IPF core since Itanium 2 design (McKinley)
 - RC minimization of critical core signals
- Design methodology that enables process scaling
 - Emphasis on power reduction, higher frequency, low voltage operation, and high yield
 - Elimination of dynamic logic outside of RF topologies
- Decoupling buffer between Instruction fetch and execution
 - Holds 96 instructions – replicated per thread
- Replay versus Stall design
 - Significant power reduction across all work loads
 - Instruction buffer acts as replay point for backend execution
 - Commit, exception and stall timing made easier

Poulson Core Architecture



Individual Poulson Core

Key Architectural Advances

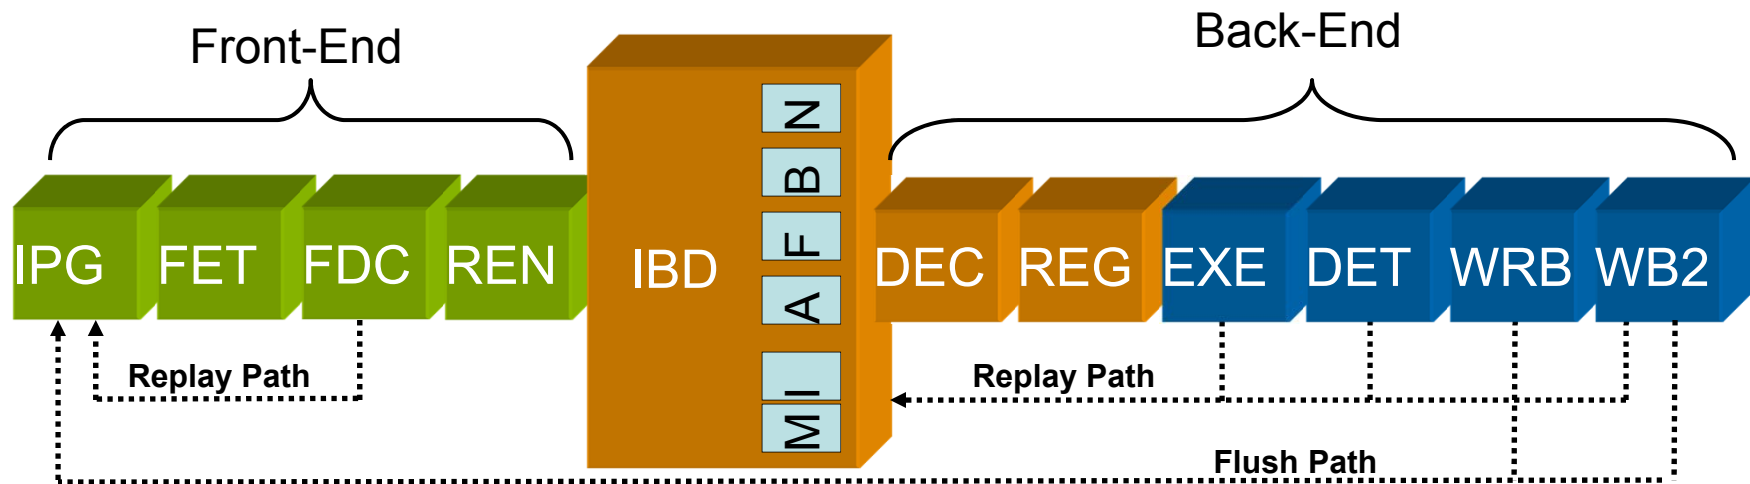
- *New Data and Instruction Pipelines*
- *New Floating Point Pipeline*
- *New Instruction Buffer*
- *Double max execution width*
 - *From 6 to 12 wide*

Derived Benefits

- *Increased instruction throughput*
- *Improved performance/watt*
- *Improved RAS coverage*
- *Core optimized for future technologies*


Increased performance, power reduction and reliability
Core architecture enables futures IPF processor designs

Main Core Pipeline



Front End: IBD Queue per thread

- instruction fetch
- branch prediction
- register renaming
- 2 bundles per cycle
- sort into queues

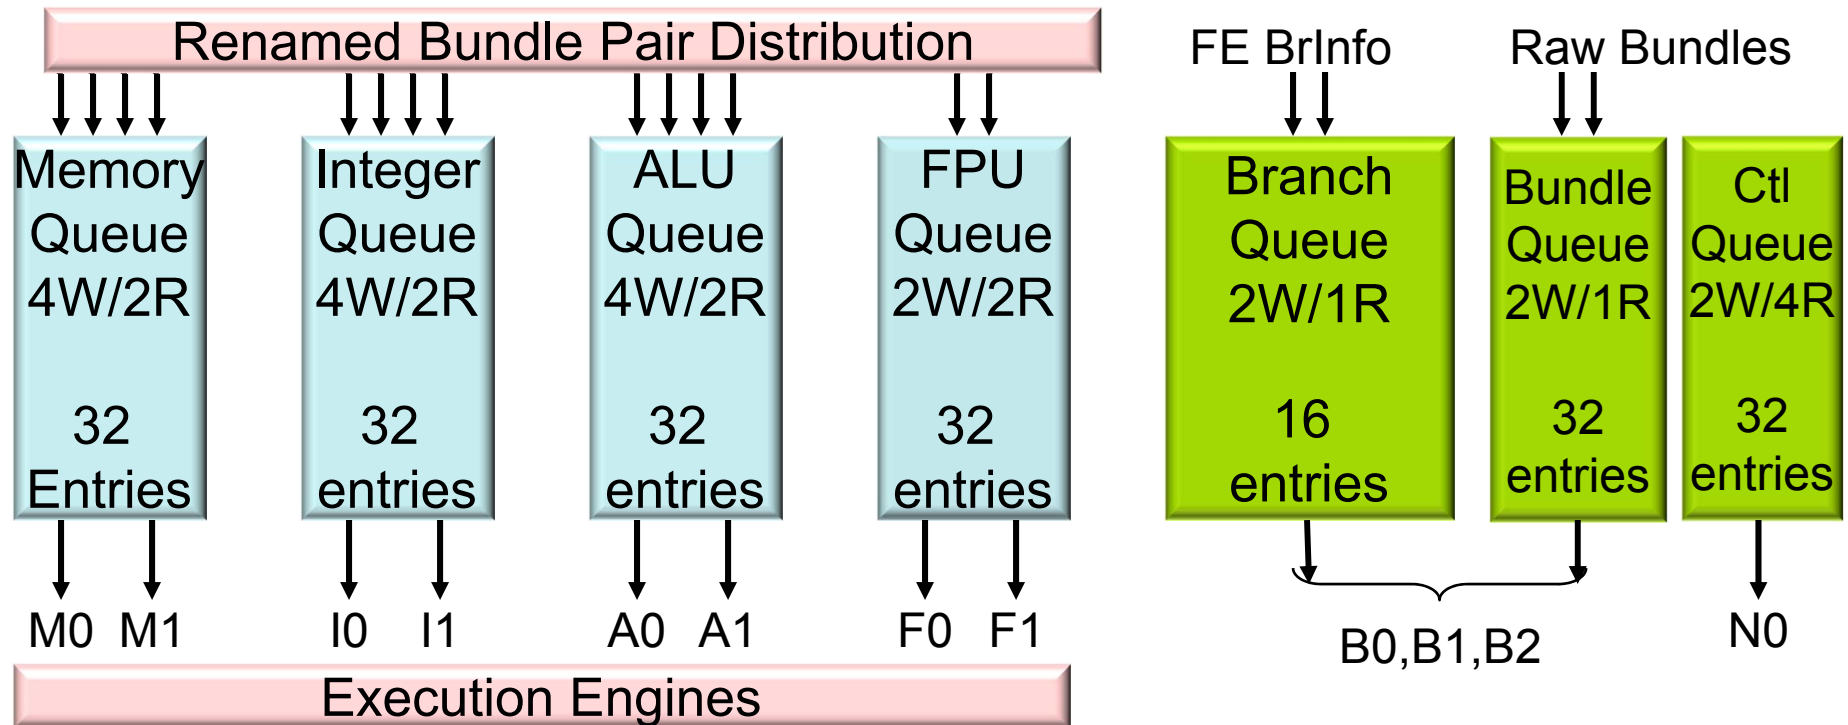


Back End:

- read 4 bundles per cycle
- execute instructions
- access FLD, or send to MLD
- replay back to queues
- instruction in queue until retirement



Instruction Buffer Logic - 12 Wide issue



- Instruction based, not bundle based
- Program order within each queue
- No order between queues
- Control queue keeps order
- NOPs only in Control Queue
 - Squashing contributed to power reduction
- Separate read pointer per queue

Max 12 wide issue including NOPs

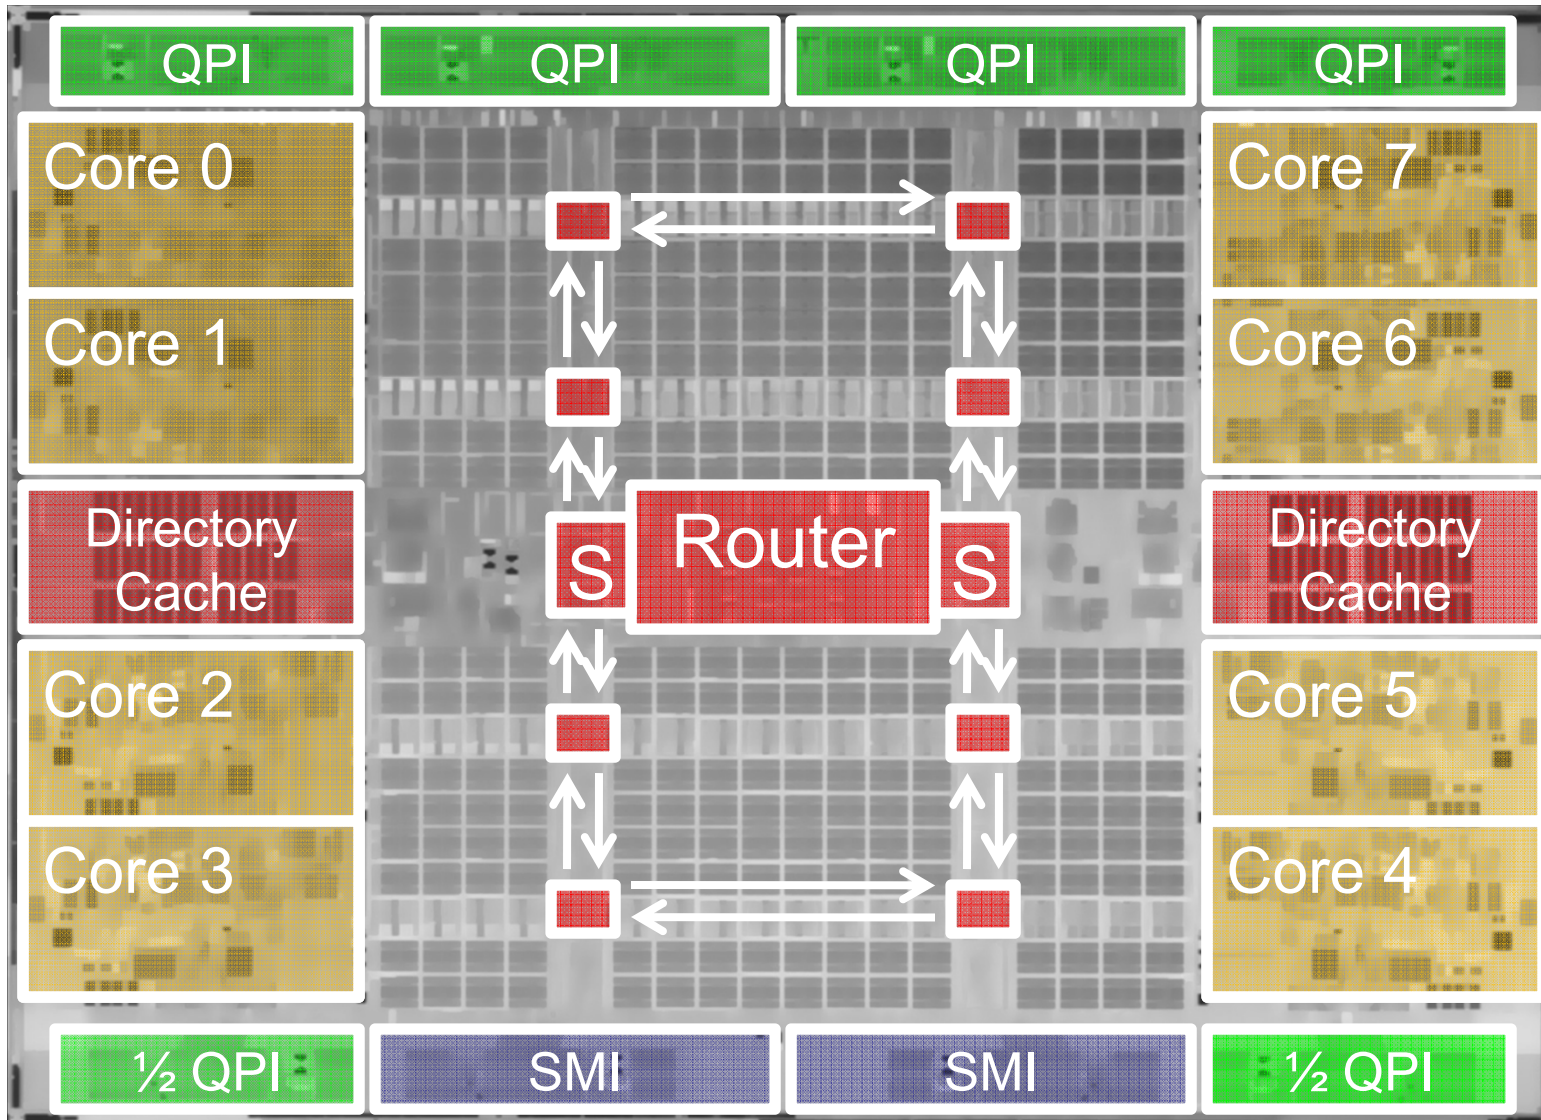
System Interface Design

- Enables socket compatibility with Tukwila design
- Ring Based system interface
 - Provides high bandwidth low latency access to cache
- Two home agents
 - Directory based cache coherence protocol
- 10 port crossbar router for IO and memory traffic
- Improved RAS capabilities
- System management bus interfaces
- Power Control Unit
- Clock delivery and configuration unit

System Interconnect

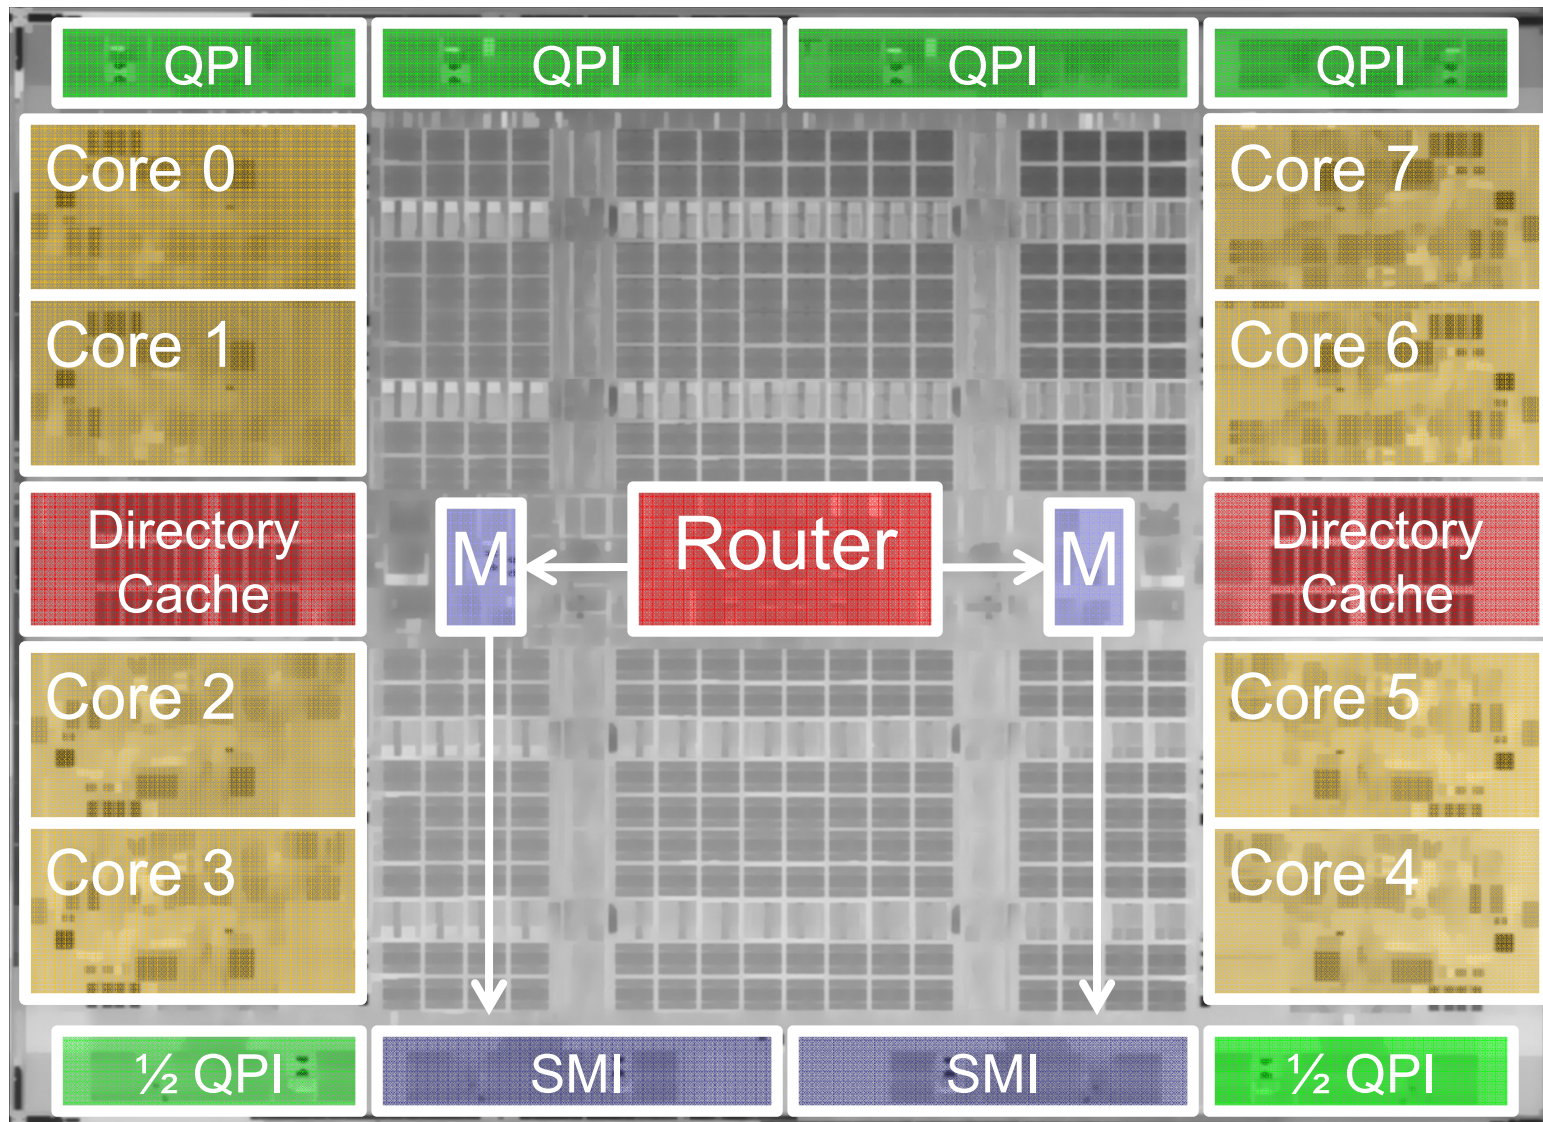
- All interconnects are double pumped at a maximum transfer rate of 6.4 GT/s. (4.8 GT/s on TKW)
- Poulson implements four full-width and two half-width QuickPath™ Interconnects (QPI)
- Four full-duplex Scalable Memory Interconnects (SMI) for processor-to-memory traffic
- Dual integrated memory controllers with Double Device Data Correction
- Enhanced DIMM clock gating to reduce system power consumption
- The IO circuit area consumes 66 mm² and contains 44M transistors

System Interface Overview




700 GB/s bandwidth provided by Ring based interconnect

System Interface Overview



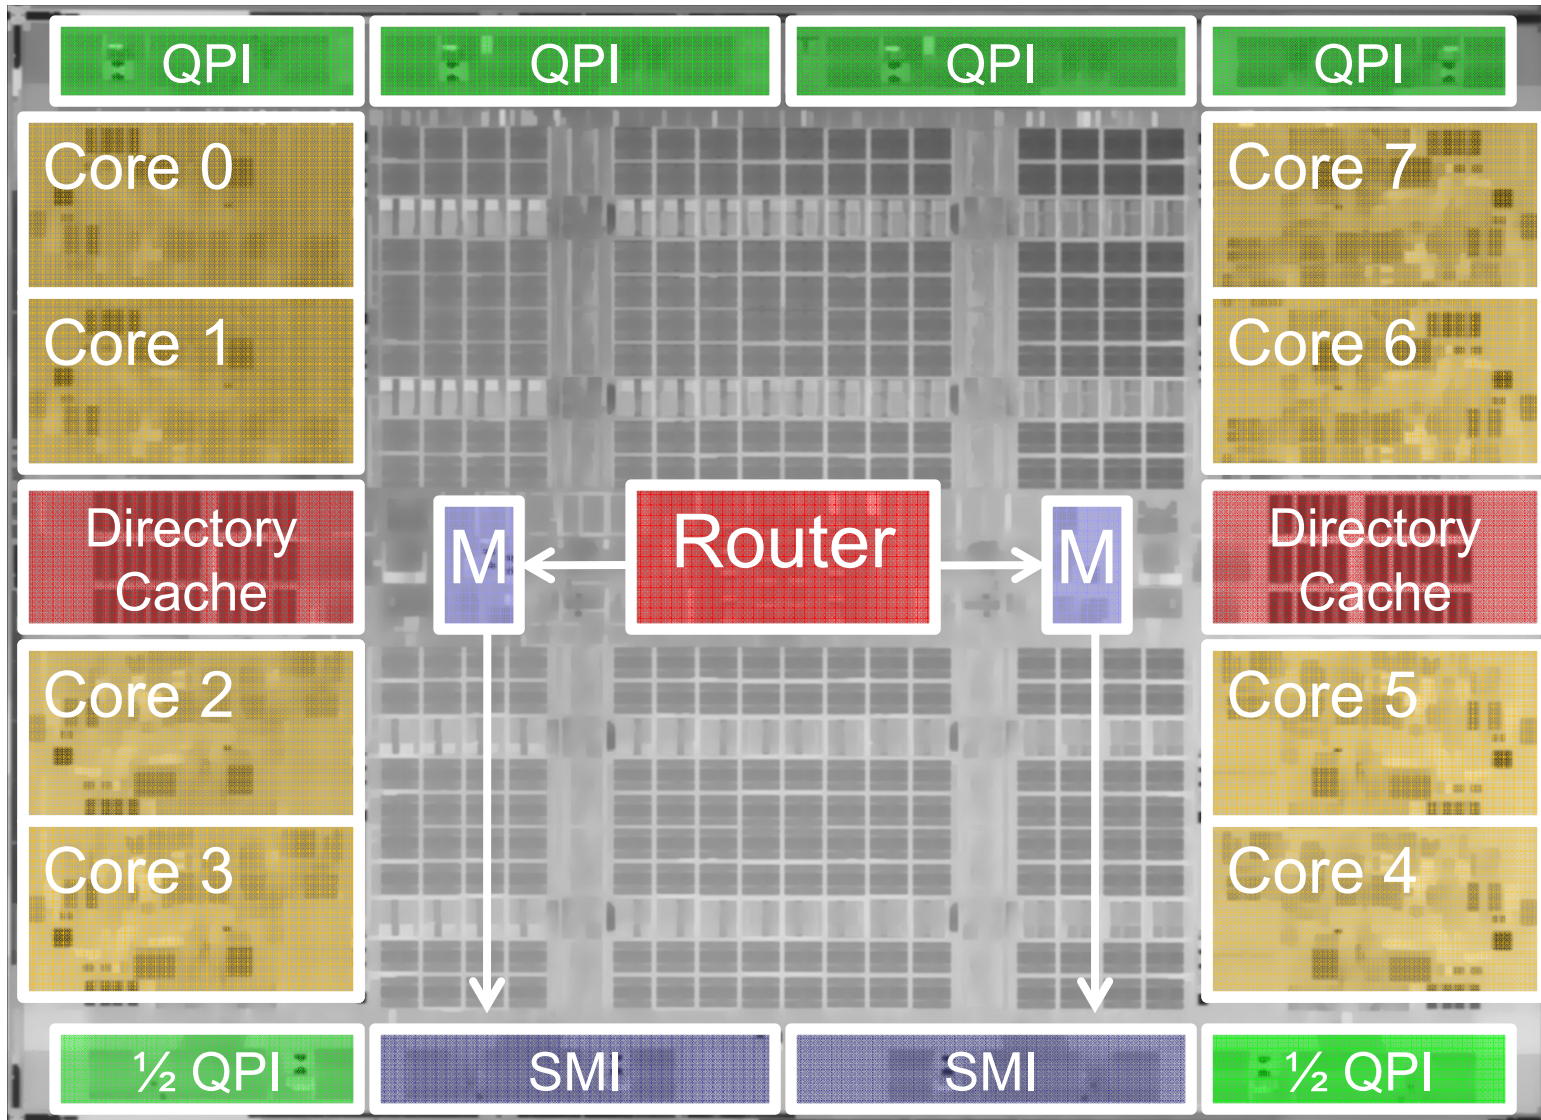
45 GB/s bandwidth provided by Scalable Memory interconnect

System Interface Overview



128 GB/s bandwidth provided by QuickPath™ interconnect

System Interface Overview




700 GB/s bandwidth provided by Ring based interconnect

SRAM Cache Summary


Structure	Logical Size (MB)	Local Bit (row)	Access	Redundancy	ECC Protection
Last Level Cache	32	256	Cycle	Column/Row/Way	DECTED
Last Level Tag/LRU	3.6	64	Phase	Column	SECEDED
Directory	2.2	128	Cycle	Column/Row/Way	SECEDED
Mid Level Data	2.0	64	Phase	Column/Row	SECEDED
Mid Level Instruction	4.0	256	Cycle	Column/Row	SECEDED
Mid Level Inst Tags	0.165	64	Phase	Row	SECEDED

All SRAM arrays protected by Intel's Cache Safe Technology


Cache Overview



SRAM bit cell




Tukwila SRAM bit cell



65 nm, 0.570 μm^2

Poulson SRAM bit cell

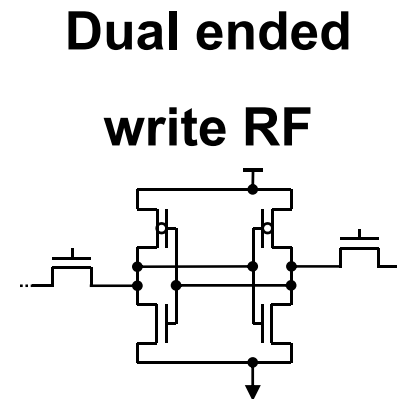
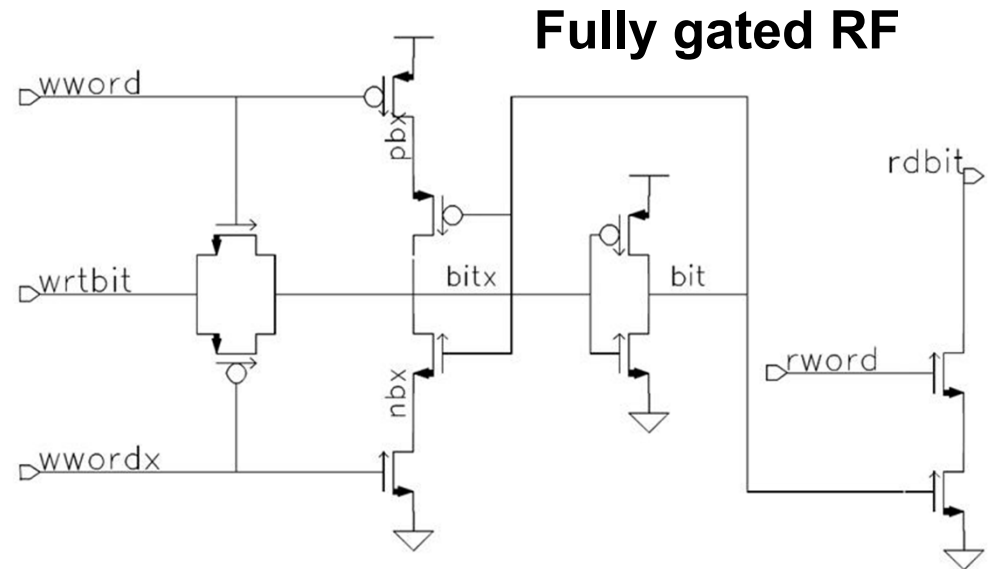


32 nm, 0.171 μm^2

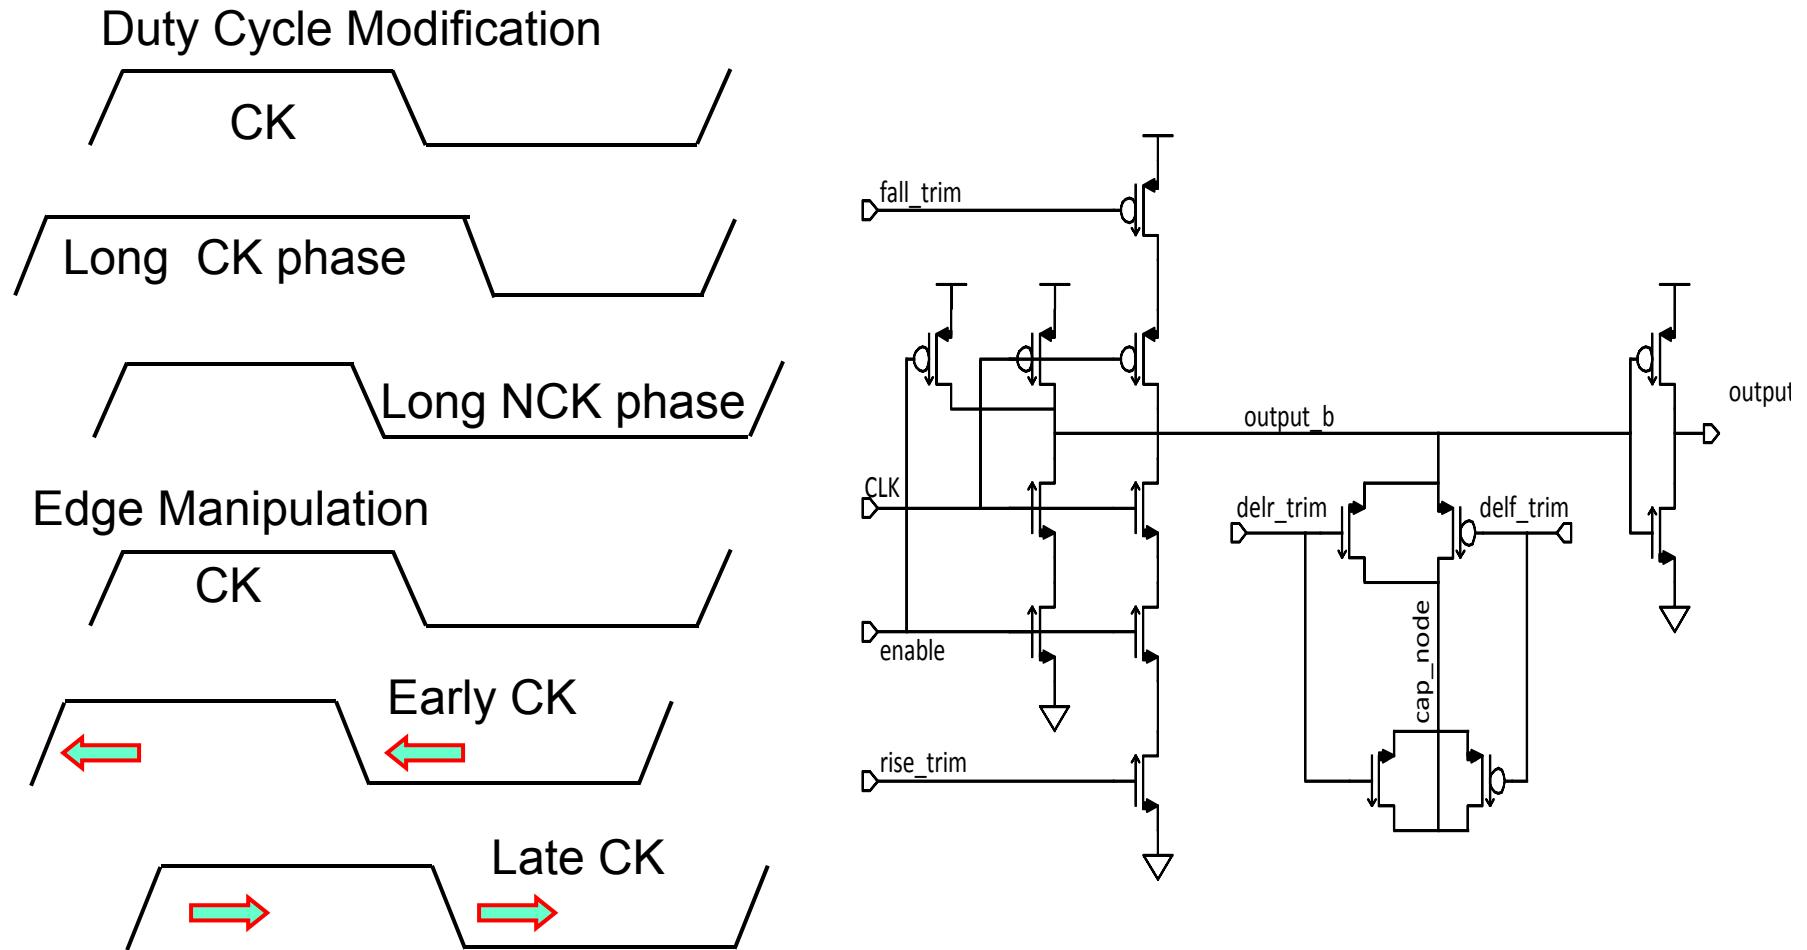
Intel's 32 nm process technology enables the integration of 50 MB of on die SRAM

Operation at minimum Voltage

- 32nm process presented challenges
 - increased fet variation
 - increased number of fets
 - Higher core count
 - More cache
- Core implemented a fully gated Register File (RF) bit cell
 - Improved write performance at V_{ccmin} (contention free)
 - Lower power design
 - Word line power less than bit lines
- Same size as other RF topologies



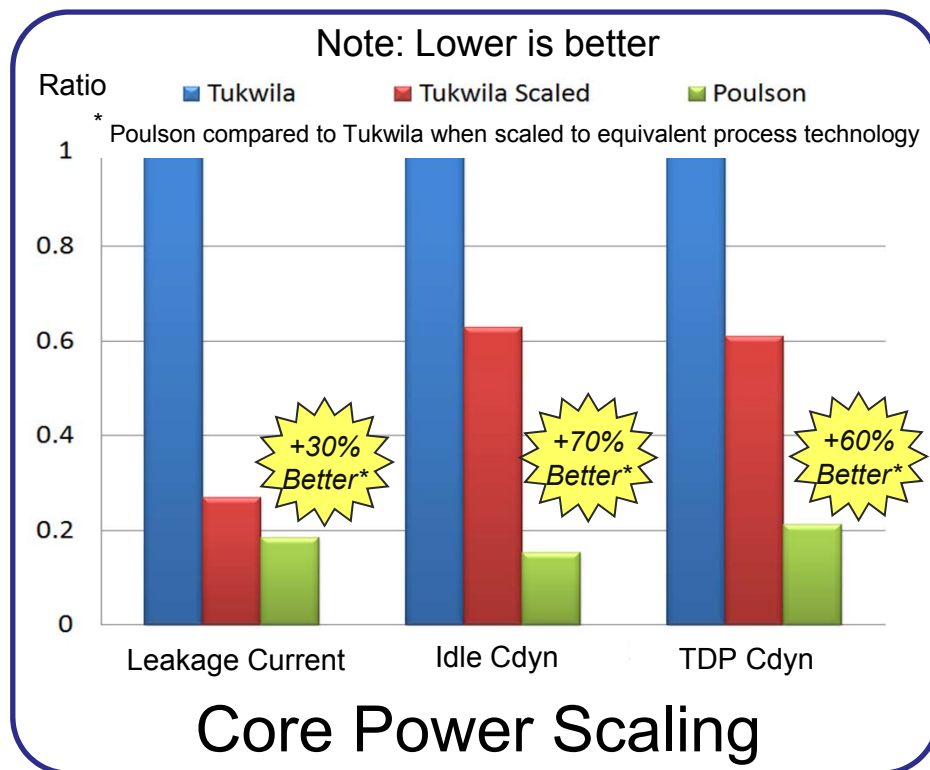
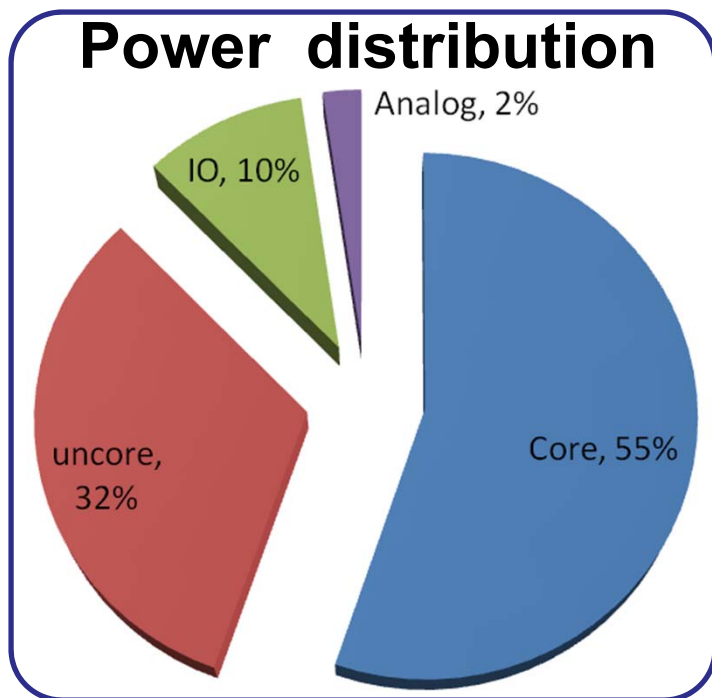
Clock Manipulation Circuitry



- Clock tuning enables 400 MHz Improvement
- Duty cycle correction / Independent edge control

Power reductions

- Methods
 - Removal of dynamic logic
 - Stall -> Replay architecture
 - Aggressive clock gating
 - FET width reduction via algorithmic tools



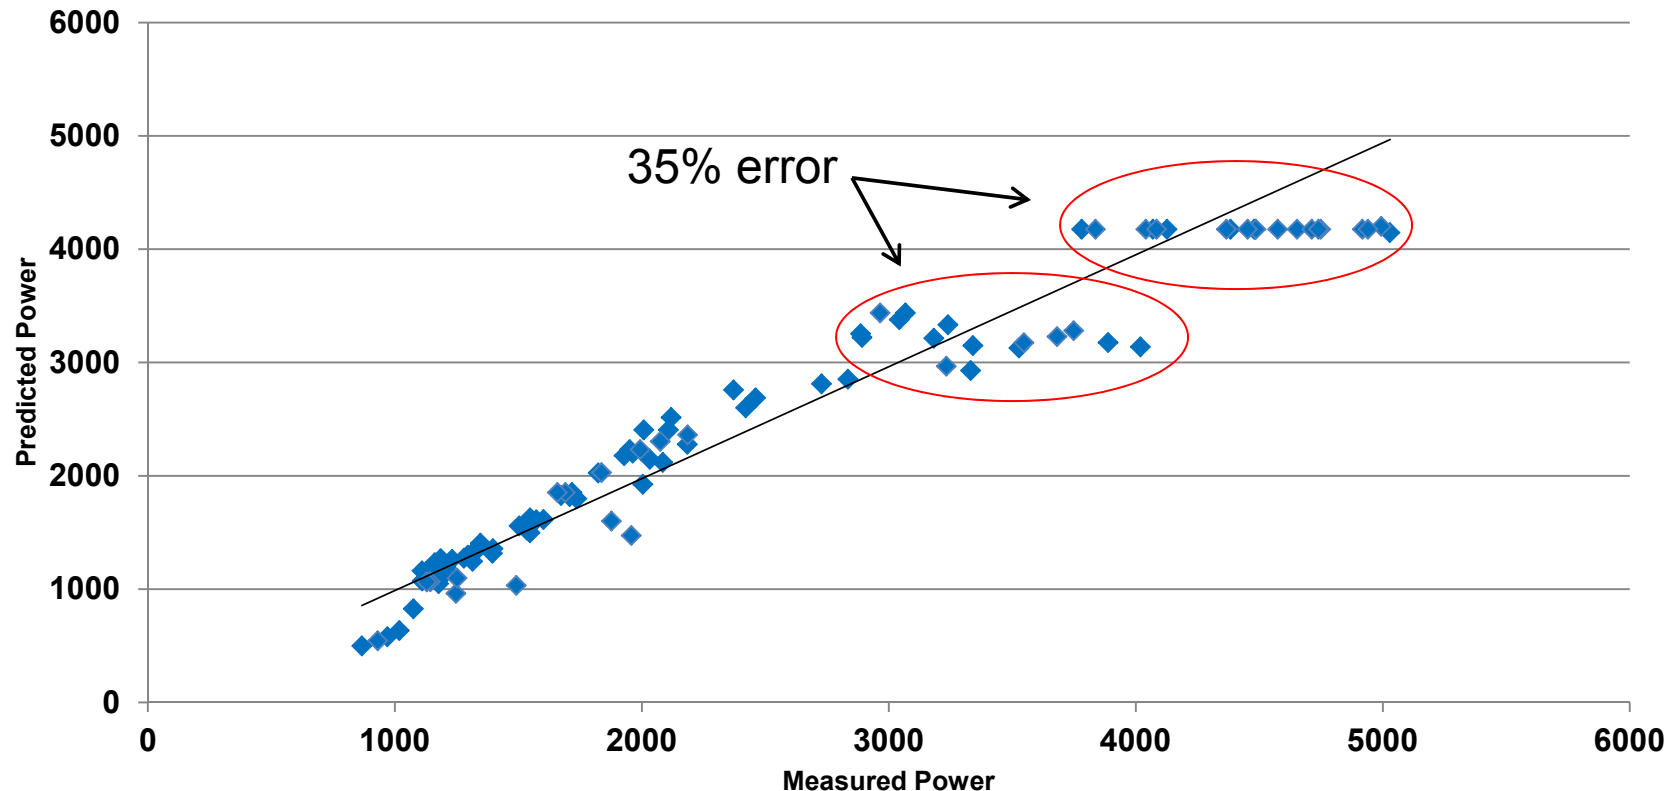
- Focus on power reduction to meet socket compatibility demands
 - (15 Watt socket power reduction)
- Core leakage and Cdyn reduction
 - Low Leakage insertion 82% vs. 70% in Tukwila
 - Idle / TDP power reduction vs. Tukwila

Significant reductions in Leakage, Idle, and TDP power translate to improved Perf

Poulson Dynamic power

Power Prediction Without Data Support

$R^2 = 0.9349$

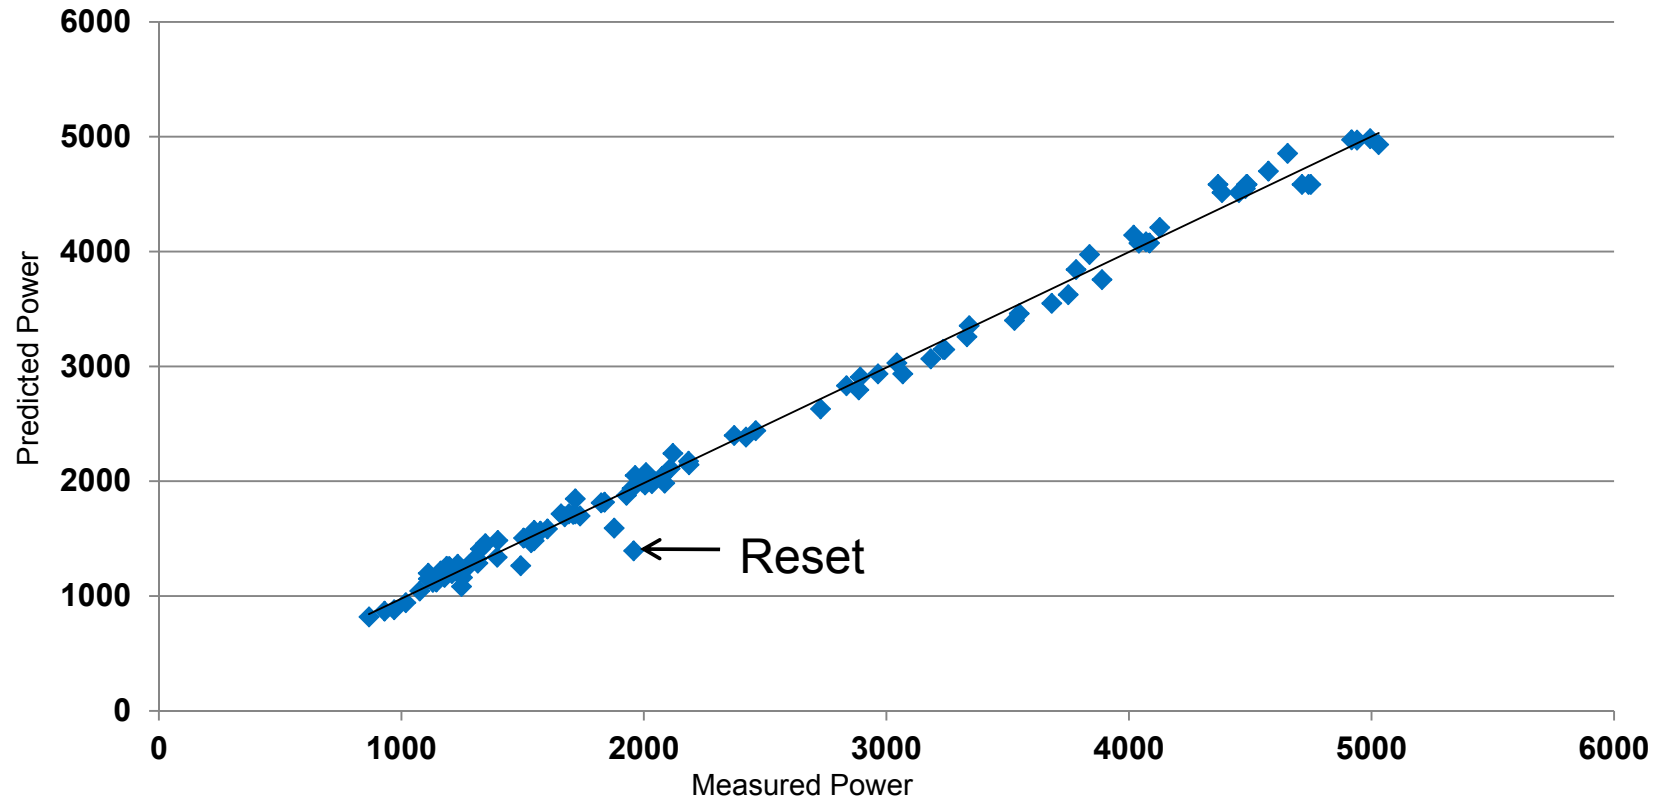


- Tukwila introduced instruction level power prediction
- Data activity represents up to 35% of dynamic power

Poulson Dynamic power


Power Prediction With data support

$R^2 = 0.9938$




- Poulson introduces data level power prediction
- Digital Power measurement now within 2% of measured values in 50 nanoseconds

Core Asymmetries



Thermal hot spots

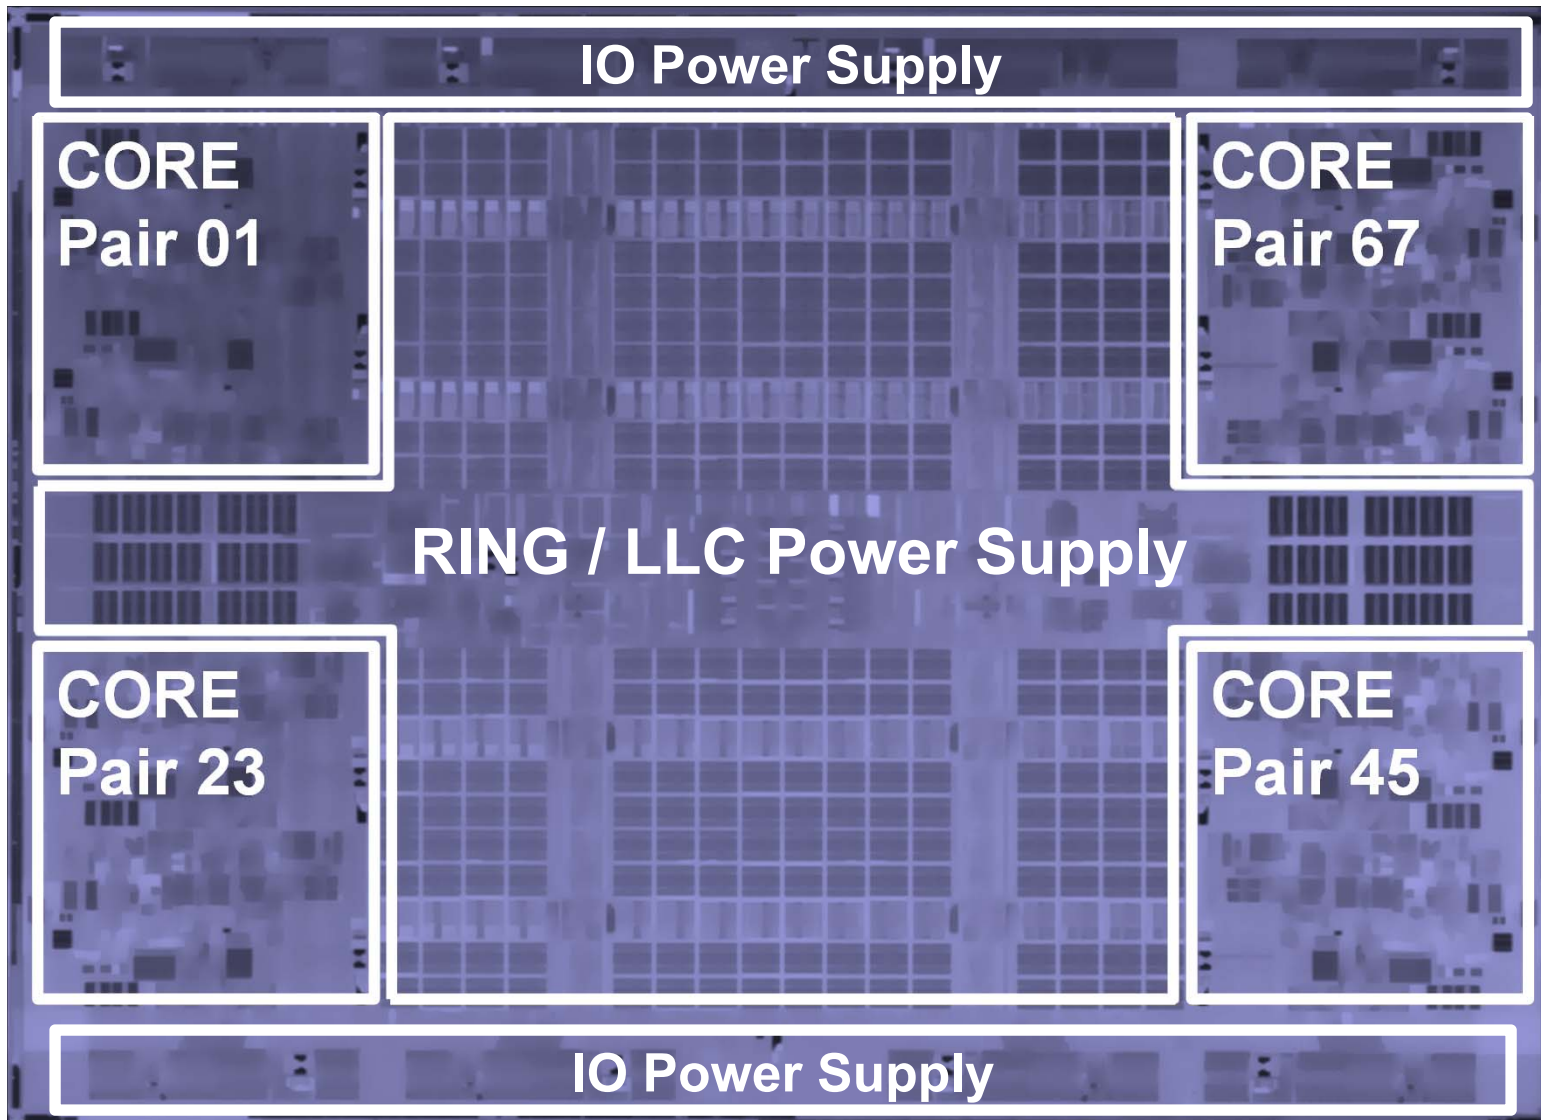
- Poulson has 10 thermal diodes
- Located in hot and cold spots on the design
- Active system to respond to thermal changes



Process Variation

- L_{eff} varies across the die
- Slow core limits operating frequency
 - Fast cores are higher power
 - Effects are stepper and mask dependent

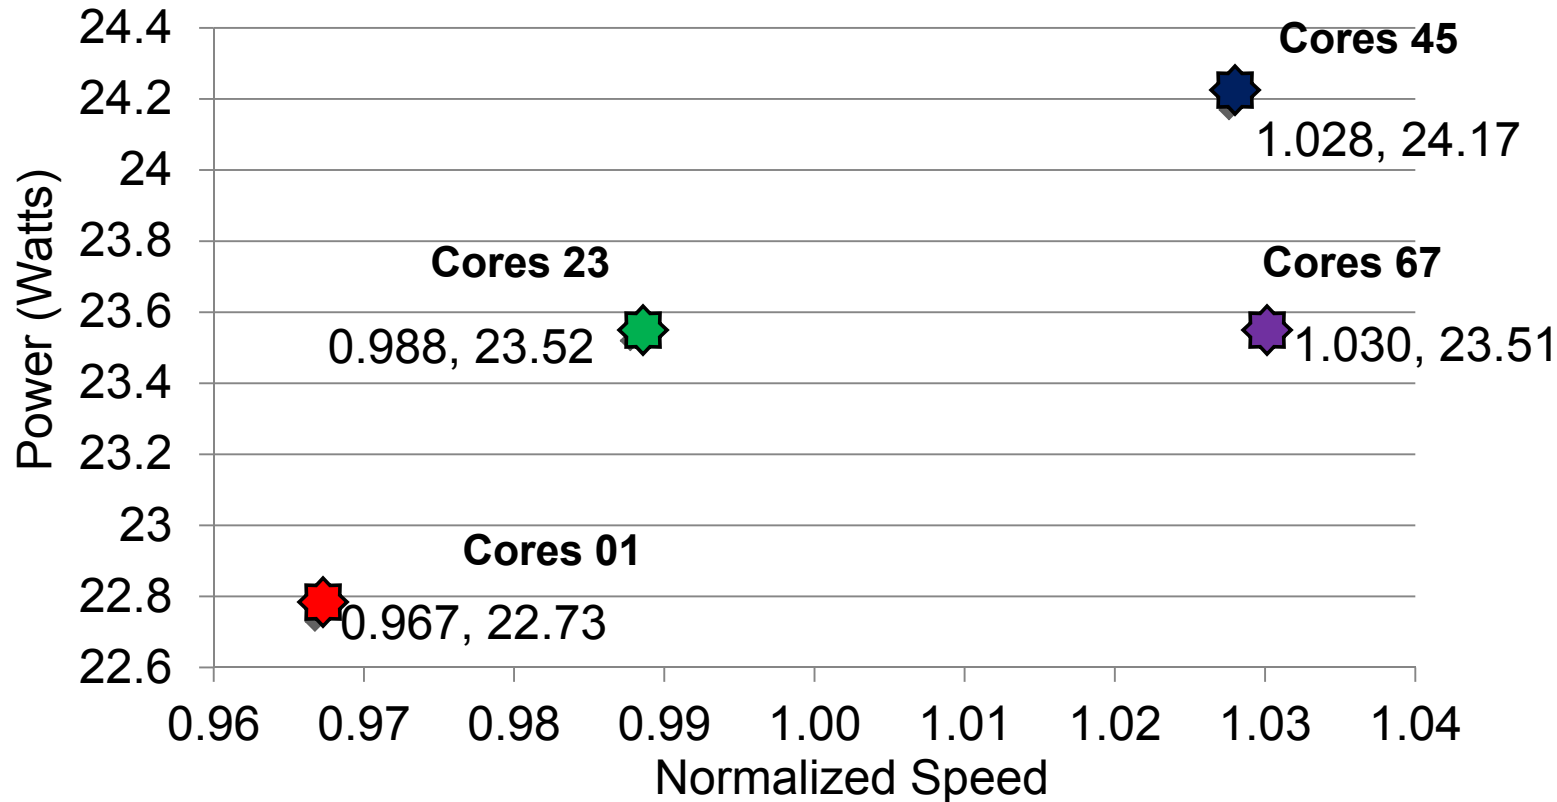
Processor Power Planes



Power is optimized across the 6 voltage domains

Core Pair Optimization

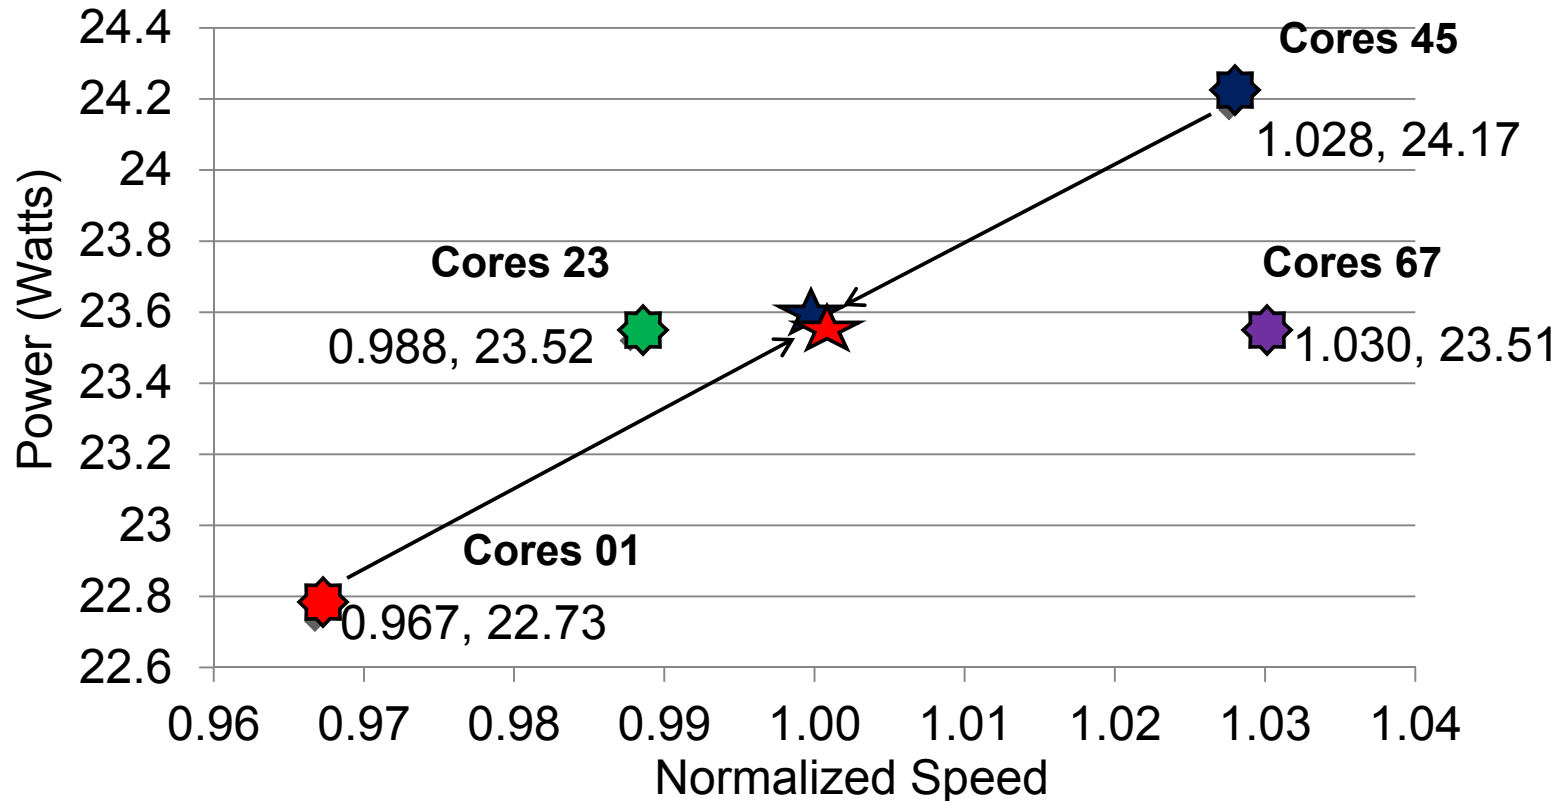
Speed and Power @ Uniform Voltage



Speed and Power can be optimized for each Voltage domain for improved performance

Core Pair Optimization

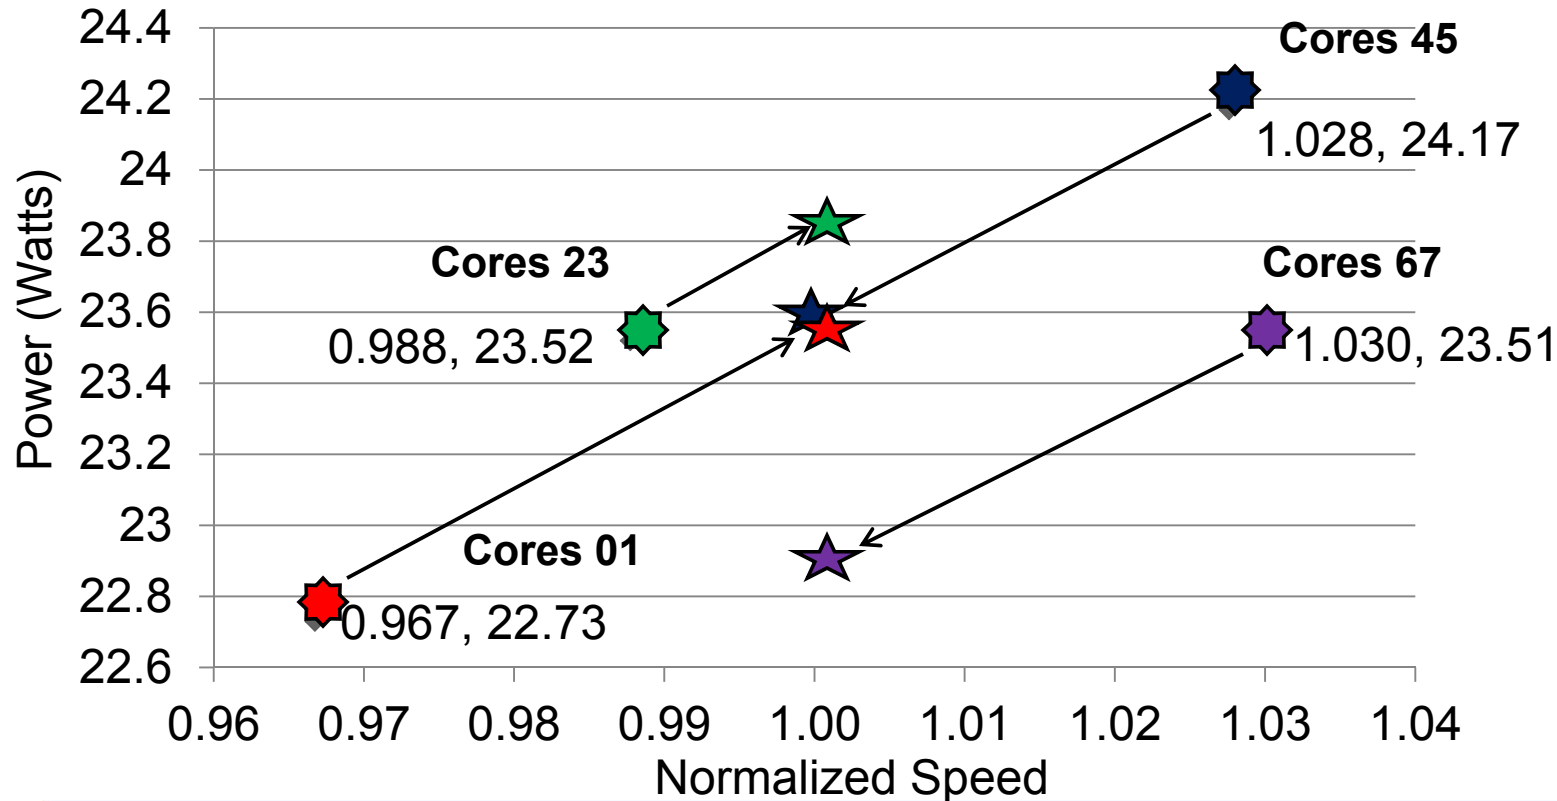
Speed and Power @ Optimized Voltage



- Cores 01 are slower and lower power
 - Increase voltage and power to improve frequency
- Cores 45 are faster and higher power
 - Decrease voltage and frequency to recover power

Core Pair Optimization


Speed and Power @ Optimized Voltage



- Cores 23 are slower and lower power
 - Increase voltage and power to improve frequency
- Cores 67 are fast and low power
 - Decrease voltage and frequency to recover power

Core Pair Optimization

Speed and Power @ Optimized Voltage



Independent supply optimization improves frequency up to 5% with no impact to power!

Poulson RAS enhancements

- Last Level Cache now utilizes inline DECTED and Intel cache safe technology
- Core Cache designs now have inline SECDED protection
- Integer and Floating Point Register Files have SECDED
- All other Register File arrays have error protection
 - Hardware/Software mechanisms that enable parity errors to be corrected
 - On Tukwila this would have resulted in Design Uncorrectable Errors
- End to End protection on many internal buses
- Residual error protection on FPU Adders and Multiplier
- PSN with 2X the cores improves RAS capabilities of IPF

Poulson enables even higher levels of Reliability

Summary

- Poulson design builds on Itanium's strengths:
 - High performance cores
 - Industry leading cache design and density
 - High levels of integration enabling mission critical RAS capability.
- Adds new features to deliver increased performance
 - 2X the number of cores
 - Power reductions translate to performance
 - High bandwidth low latency system interface
- Demonstrates innovative engineering work
 - First reported 3.1 billion transistor microprocessor
 - Deterministic adaptive power-frequency management
 - Core micro architecture optimized for 32nm process and beyond

Poulson builds a foundation for future Itanium designs