Low Power AMD Athlon™ 64 and AMD Opteron™ Processors

Hot Chips 2004
Presenter: Marius Evers

Block Diagram of AMD Athlon™ 64 and AMD Opteron™

- Based on AMD's 8th generation architecture
  - AMD Athlon™ 64 and AMD Opteron™ differ in the availability of some features
- Based on the x86-64 instruction set
- Out-of-order, 9-issue superscalar processor
- Integrated Northbridge, DDR memory controller, and HyperTransport™ links
- Supports glue-less multiprocessing

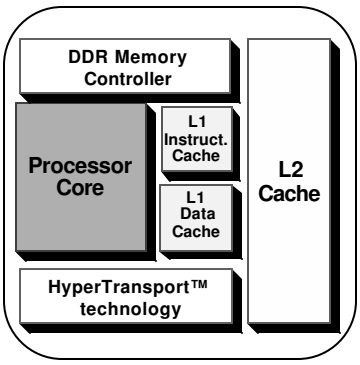

Objectives for Low Power AMD Athlon™ 64 and AMD Opteron™

- High performance in power constrained environment
  - 55W and 30W for AMD Opteron™ server processor
  - 62W, 35W and below for Mobile AMD Athlon™ 64 processor
- Improved performance at low voltage
- Reduced power consumption in low power modes
- Low power goals align well between server and mobile markets

High Performance, High Efficiency!

Performance in a Power Constrained Environment

- A processor can perform at several power / performance points
- Different points are found by changing
  - Frequency
  - Voltage
  - Process targeting (if available)
- In a power constrained environment: power efficiency leads to performance
- Many market segments are becoming power constrained because processor power keeps increasing

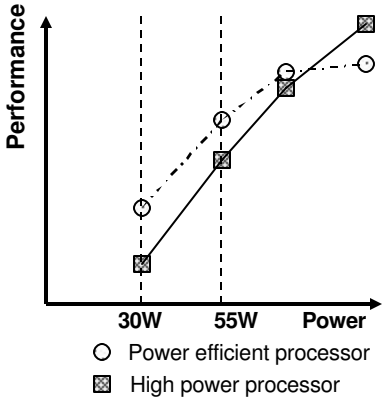

Performance in a Power Constrained Environment

- We believe performance within a given power envelope is more important than just raw performance
  - Cost-effective cooling solutions across all segments
  - Power delivery to and cooling of server racks and data centers
  - Battery life for mobile applications
- Blindly pushing MHz as a metric for performance is bad for power
  - Power efficient processors have high performance at a lower frequency
  - AMD has spearheaded efforts to use application performance instead of MHz for product naming
- Power efficiency must be addressed at all levels
  - Process technology
  - Circuits / implementation
  - Architecture
  - OS (support for low power states)
  - Systems

Outline

- Low Power Process Technology
- Improved Low Power Modes
- Reducing Static Power
- Reducing Dynamic Power
- Summary: What Does This Enable?

Low Power Process Technology

Low Power Process Technology: Silicon on Insulator Process

- No bottom capacitor in source/drain
- Overlap cap still dominates
- Tighter layout rules reduce wire parasitics
  - No N-wells/body ties
  - No antenna diodes required
  - N-active and P-active can abut

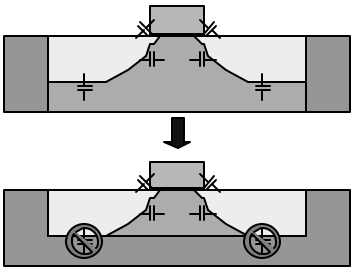

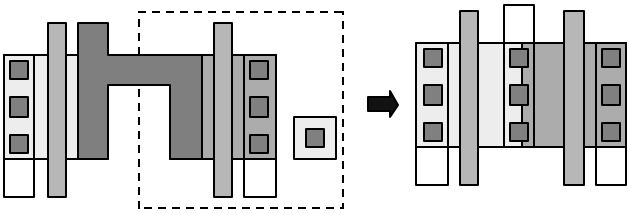

Low Power Process Technology: Process Targeting for Low Power Processors

AMD

- Thicker gate oxides (GOX)
  - Reduced static gate leakage caused by tunneling
  - Reduced gate capacitance
  - More reliable (Great for servers)
- Longer nominal channel length reduces static leakage
- Dual Gate Oxide : two gate oxide thicknesses on die
  - Thinner GOX for core transistors – not as thin as desktop process target
  - Thicker GOX for on-chip decoupling capacitors – reduces gate leakage
  - Thicker GOX transistors in I/O power domain
- Three threshold voltages
  - High Vt transistor has low leakage, low performance
  - Medium Vt has higher leakage and 10% higher drive current
  - Low Vt has very high leakage and another 10% higher drive current
  - Vt levels (and therefore leakage) are targeted for low power

AMD

Improved Low Power Modes

Improved Low Power Modes: Overview of Improvements

- Low power modes were already optimized on previous revisions
  - Clocks and voltage were ramped down
  - Other optimizations outside the scope of this talk
- Every mW counts in these modes, so several new improvements were done:
- Disable I/O circuitry in low power states
  - Tri-state DDR interface
  - Disable I/O compensation circuit
    - Pulls static current through a resistor on the board
    - Sets I/O drive strength as temperature varies over time
  - Disable termination structures in "bypass clock" input
    - Only used in test modes, so disable all other times
- Disable I/O level shifters in low power states
  - Level shifters from core to I/O power domain drive static current
  - Disable during low power states – the pins aren't wiggling
- Remove static pre-chargers
- Architectural and micro-architectural improvements

Improved Low Power Modes: Remove Static Pre-chargers

- Remove static pre-chargers for arrays that operate in low-power modes
- In high frequency operation, identical power consumption
- In low frequency operation, full-swing bit lines fully evaluate and stop drawing current

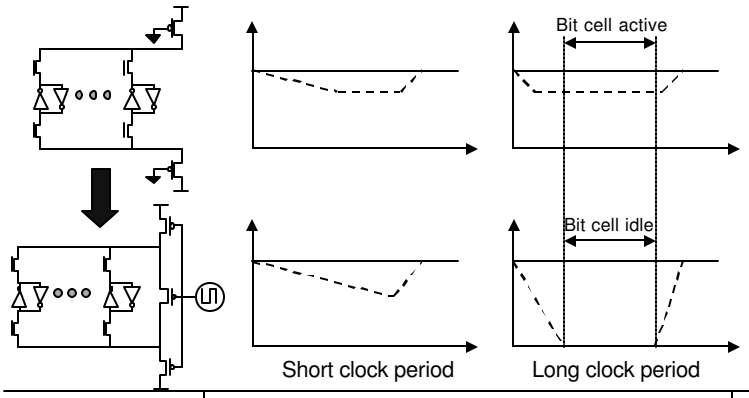

Improved Low Power Modes: Architectural Power Savings Feature

- AltVID (Alternate Voltage ID)
  - Low power modes already lower VDD in low power modes by driving a 5-bit VID code to the voltage regulator (VR) on the motherboard
  - AltVID is a method for ramping the voltage down further after the clocks have ramped down.
  - This programmable code is sent to the VR to reduce the processor voltage to the minimum operational level
  - Reduces static power by ~ 50-100 mW

Improved Low Power Modes: Micro-architectural Power Savings Feature

- Processor HALT state is optimized to save power in an innovative way
- HALT is a microcode loop, waiting for interrupts
- Clocks are automatically ramped down
- But because microcode is running, the Register File, reorder buffer, reservation stations and microcode ROM macros are still being used.
- Microcode sets a control bit that disables instruction retirement
  - Causing the instruction reorder buffer to fill up
  - Which in turn stalls instruction dispatch
  - Which in turns stalls the microcode engine
  - Thus disabling macro accesses without adding logic to timing critical paths
- Saves about 300 mW

AMD

Reducing Static Power

Reducing Static Power: Voltage Threshold (Vt) Selection

AMD

- Vt selection is a large lever for both performance and static leakage
  - High Vt cells are slow and have very low leakage
  - Low Vt cells are fast and have high leakage
  - We select a small set of cells to make low Vt to improve performance
  - Most cells must remain medium or high Vt to maintain low leakage
- Vt selection for low power microprocessors has unique issues
  - Speed improvement from lowering Vt is larger at low voltages
    - Critical paths should be all low Vt to improve frequency
    - Therefore more low Vt cells are needed to optimize low voltage operation compared to high voltage operation
  - Vt choices must be appropriate for all expected operating conditions and process variations
    - Mobile processors need to be optimized for a larger spectrum of voltages than desktop processors

Reducing Static Power: Vt Selection Algorithm

- We developed a new iterative Vt selection algorithm
- Multiple voltage/process corners were considered
- Vt selection was based on bottleneck detection
  - The timing and leakage impact of lowering the Vt of a cell was also considered
- Low Vt usage was reduced by 25-35% compared to our previous path based algorithm

Reducing Static Power: Timing Optimizations !

- If a path meets timing with medium Vt cells, it does not require low Vt
- We focused a large design effort on improving sub-critical timing paths to decrease low Vt usage
  - Special timing reports highlighted paths that used the most low Vt cells so these could be fixed first
  - Bottleneck information was also used to target fixes to the most important areas
- Low Vt usage was reduced by 20-30% due to these improvements

Reducing Static Power: Cache Optimizations

- Reduce leakage and increase cell stability in the L1 and L2 caches by using “wimpy-length” transistors
  - These remain long-channel as the process CD is pushed
- (Almost) all transistors in L2 cache are HVT
  - Considerable reduction in static leakage
- The L2 cache accounts for half the die, so these advantages are very significant

Reducing Dynamic Power

Reducing Dynamic Power: Attention to Detail

- We identified high power areas using pattern-based dynamic power simulations
  - Target these areas first using general and custom designed solutions
  - Several patterns were used to identify power hot-spots for different workloads
- In other areas, we reduced power a little at a time
  - Clock gating for flip-flops
  - Improvement of edge rates
  - Capacitance reduction
  - Logic minimization
  - … and several localized approaches
- Reducing dynamic power requires paying attention to detail

Reducing Dynamic Power: Improved Low Voltage Performance

- Dynamic and static power dissipation are strongly related to voltage
  - Lowering the voltage reduces power
- We focused timing work on low voltage operation
  - Improved timing by 5-10% relative to high voltage operation
  - Improved low voltage roll-off
  - Get same performance at lower voltage
- NB and I/O functionality was optimized to run at the lowest supported voltage
  - Avoid raising voltage for minimal workloads
- Support lower minimum voltages
  - Improves sleep state power

Reducing Dynamic Power: Lower Minimum VDD

- Minimum VDD is a functionality constraint driven primarily by writeability
- Can only lower chip voltage until memory nodes can't be written
  - Flops and latches
  - Bit cells
  - Dynamic nodes with keepers
  - Level shifters
- "One-sided" writes won't work

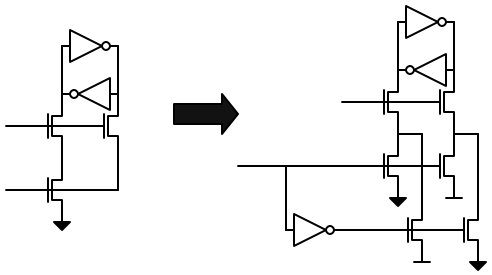

Reducing Dynamic Power: Cache Optimizations

- Don't distribute global clock over the L2 cache
  - Send low-skew clock to the edge of the L2 cache
  - Different cache banks receive timing signals with different delays
  - Pay in latency to synchronize skewed cache results to low-skew global clock
- Compare cache read data to write data and ignore the write if they are equal

AMD

Summary
What Does This Enable?

Summary

AMD

- Improved on chip power
  - Process technology
  - Circuits / implementation
  - Architecture
- Work with industry partners to deliver low power solutions
  - OS
  - Systems
- Low power design enables leadership products
  - Great mobile part
  - Power efficient server parts
  - Dual core solution within normal power envelopes !

Dual Core AMD Opteron™ Processor

- 2 CPU cores on 1 die
- Shared on-chip Northbridge
  - Coherent HT links allow glue-less connection of multiple dual core chips
  - 128 bit DDR interface
- AMD Opteron™ was designed as CMP from the beginning
  - Incremental changes for dual core design
- Programming model is SMP rather than SMT
  - 2 complete CPUs rather than extra 'virtual CPU'
  - OS manages chip as 2 processor SMP system

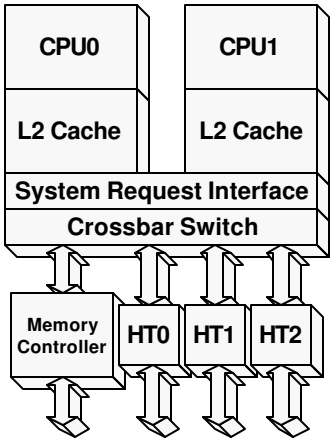

Authors

Marius Evers

Michael Golden

Mike Leary

Kevin McGrath

Chuck Moore

Trademark Attribution

AMD, the AMD Arrow Logo, AMD Athlon, AMD Opteron, and combinations thereof are trademarks of Advanced Micro Devices, Inc. HyperTransport is a licensed trademark of the HyperTransport Technology Consortium. Other product names used in this presentation are for identification purposes only and may be trademarks of their respective companies.