Power Mitigation Techniques in Complex MPSoCs

Olivier Franza

MPSoC'08

Agenda

- Background
- Dynamic Power
- Leakage Power
- Applications
- Conclusion

Background

- Technology
  - Voltage scaling slower than technology
  - Power density doubling every generation
  - 30% transistor capacitance reduction each generation
    - Implies reduction in power only with iso-transistor count
  - Miniaturization calling for reduction in heat dissipation
  - Traditional cooling solutions reaching physical limits
- Economy
  - Electronic device presence multiplying
  - Energy consumption increasing
  - Energy costs rising
- Sociology and environment
  - Global warming awareness

http://www.phys.ncku.edu.tw/~htsu/humor/fry_egg.html

Background

- Move to multi-core
  - Good for power and power density reduction, but not enough…

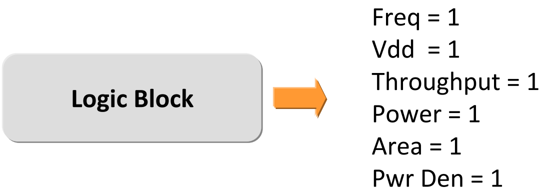

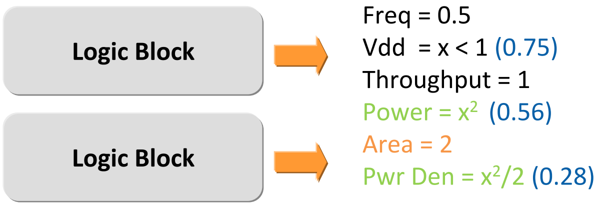

- Definition
  - Power consumed = dynamic power + leakage power

$$P = (P_{switch.} + P_{s.c.} + P_{cont.} + P_{glitch}) + P_{leak.}$$

$$P = (\alpha.C.V_{cc}^2.F + P_{s.c.} + P_{cont.} + P_{glitch}) + V_{cc}.I_{leak.}$$

α :	activity factor	$I_{leak.}$:	leakage current ~ exp(-qVt/kT)
C:	switching capacitance	Vt:	threshold voltage
V_{cc} :	power supply voltage	T:	temperature
F:	clock frequency		

J. Rattner, Keynote, DesignCon2006

Dynamic Power Reduction

$P_{dyn.} = \alpha . C . V^2 . F + P_{s.c.} + P_{cont.} + P_{glitch}$

- Power mitigation opportunities
  - Switching activity reduction
    - Conditional execution, pre-charge
    - Conditional clocking
      - Force clock-gating awareness in rtl & design
      - Improve skew management methodology
      - Evaluate clock enable logic benefits
    - Turn-off inactive blocks
    - Reduce toggling of high C nodes & busses
  - Voltage/technology scaling
    - Dynamic voltage scaling
    - Low threshold transistors
    - Multiple voltages
      - Operate as low as possible within reliability limits
  - Clock frequency reduction
    - Multiple cores, multi-threads / parallelism
    - Reduce pipeline stages
    - Use double-edged sequential elements
  - Switching capacitance reduction
    - Optimize circuit design
      - Balance power-delay trade-off: move closer to optimal power-performance design point
      - Shift to static versus dynamic logic
      - Minimize diffusion, wire and gate loading particularly in high α areas (domino, clocks)
    - Reduce bus power
      - Implement transition encoding to minimize toggles a low-voltage differential on-chip buses
      - Reduce driver capacitance with appropriate repeater insertion
    - Optimize layout design
      - Use efficient layout techniques (shielding, spacing)
    - Lower clock loading
      - Reduce local clock interconnect routing
      - Group /cluster sequential elements

Dynamic Power Reduction

$P_{dyn.} = \alpha . C . V^2 . F + P_{s.c.} + P_{cont.} + P_{glitch}$

- Power mitigation opportunities
  - Short circuit power
    - Power dissipated if both the PMOS and NMOS transistors are in an on state
      - Function of $(V_{cc} - 2V_t)^3$
      - Linearly increases with input signal slope
      - Highly sensitive to in/out slope ratio
    - Avoid large slope-in to slope-out ratios
    - Avoid power races and contention (force state, apply appropriate reset/enable/mut-ex conditions)
  - Continuous power
    - Power burnt due to tail end of signal which doesn't go to full rail for a long time
    - Design optimal transistor sizing
    - Resize overloaded paths without under-driving them
  - Glitch power
    - Caused by unequal propagation delays of input signals to gate
    - Glitches multiply as they propagate through a combinational logic
    - Size gates to avoid delays/races
    - Prohibit multiple bit transitions (00→11)
    - Increase noise robustness / decrease coupling
    - Decrease sequential "vulnerability window" (when logic propagates) using edge triggered sequential elements versus transparent latches
      - Trade-off required vs. active power

H.J.M. Veendrick, "Short Dissipation of Static CMOS Circuitry and its Impact on the Design of Buffer Circuits", IEEE JSSC1984

Leakage Power Reduction

$P_{leak.} = V_{cc} \cdot I_{leak.}$

- Sources of leakage
  - Sub-threshold leakage (SD leakage)
    - Increasing with process technology, voltage, and temperature
  - Gate-oxide leakage (direct tunneling)
    - Increasing with process technology, voltage, and temperature
    - Gate oxide thickness nearing limit
      - If T_{ox} scaling slows down, then V_{dd} scaling will have to slow down
      - Temporary relief with high k dielectric
  - Other sources of leakage include
    - Junction reverse-bias leakage
    - Gate induced drain leakage

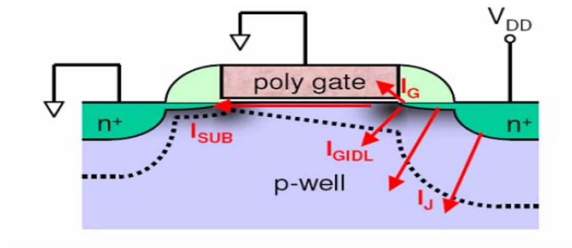

Leakage Power Reduction

$P_{leak.} = V_{cc} \cdot I_{leak.}$

- Reduction Techniques
  - Transistor level
    - Multiple transistor flavors for mutli-performance usages
      - High Vt / Long Lg (Low leakage device)
      - Nom Vt / Nom Lg (Nominal leakage device)
      - Low Vt / Short Lg (High leakage device)
    - IBM's power processors leveraging triple Vt process option
      - % high Vt device increased from 26.2% in Power4 to 33.9% in Power5
  - Block level
    - Stack forcing
      - Force one transistor into two-transistor stack with same input load
      - Trade-off between leakage and speed: can be applied to gates with timing slack
      - Leakage reduced considerably when two or more transistors are off in a stack
    - Sleep transistors for cell-based design
      - Insert sleep transistors to create virtual Vdd and Vss nodes
      - Common in cache design
      - Switching sleep transistors can cost energy

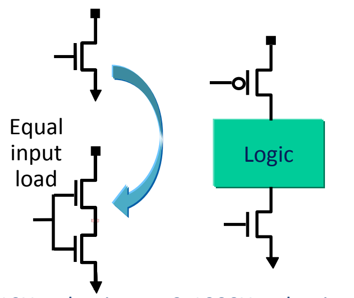

J. Clabes et al., "Design and implementation of the POWER5 microprocessor", ISSCC2004 – J. Rattner, Keynote, DesignCon2006

Leakage Power Reduction

$P_{leak.} = V_{cc} \cdot I_{leak.}$

- Reduction Techniques
  - Die level
    - Power switches
      - On die
      - On/off voltage
    - Multiple Vcc domains
      - Located off-die i.e on the board
      - Discrete voltages
    - Body bias
      - Sun's dual-core SPARC processor uses body bias to reduce leakage at burn-in
      - Effect stronger in long channel devices (Vt controlled by body bias), diminishing with short channel devices
  - Platform level
    - Lower junction temperature located off-die (on the board)
    - Drive junction temperature down for platforms to reduce leakage current

J.M. Hart et al., "Implementation of a fourth-generation 1.8-GHz dual-core SPARC V9 microprocessor", JSSC2005

Applications

- Itanium™
  - High-performance mission-critical computing Intel® Architecture processor
  - Quad-Core Itanium® processor with multi-threading (8T)
  - World's first 2 billion transistor microprocessor
  - Increased performance vs. Dual-Core Itanium® Processor 9100 series
    - 2x performance at 25% more power
  - Energy efficiency
    - High-level of system integration
      - Multi-core, QuickPath interconnects, integrated memory controllers, advanced RAS, large 30+MB cache, etc.
    - Voltage and frequency management for optimal use of power and thermal envelope
    - Multiple power supplies

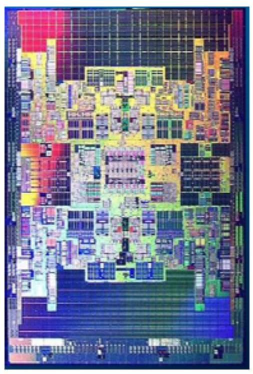

B. Stackhouse, "A 65nm 2-Billion-Transistor Quad-Core Itanium® Processor", ISSCC 2008

Applications

- Atom™
  - Low-power Intel® Architecture processor: fully Core 2 Duo ISA compatible
  - Average power consumption target in the order of a few hundred mW
  - Performance similar to mainstream Ultra-Mobile PCs
  - 47M transistors in a die size under 25mm^2 manufactured in 45nm CMOS
  - Thermal Design Power (TDP) consumption 2W @ 2GHz
  - 10x lower power than ULV Dothan
    - Low leakage transistors
    - Deep power down (C6) architecture
    - Optimized register-file and cache 6T bit cells
    - CMOS mode on quad-pumped FSB IO
    - Split IO power supply

G. Gerosa et al., "A Sub-1W to 2W Low-Power IA Processor for Mobile Internet Devices in 45nm High-K Metal-Gate CMOS",

Conclusion

- Power mitigation opportunities
  - Various power reduction techniques
  - Involvement of many different domains
    - Process, architecture, clocks, library, power management, design specification and methodology, …
  - Tailored solution required to meet actual needs and trade-offs
    - Not all solutions are good in all cases
- But that's not enough…
  - Adequate power modeling is also required!

Acknowledgments: Mondira Pant, Bill Bowhill, Jeff Pickholtz, Rajay Pai