Power-Conscious and Embedded Computing

Krste Asanovic
Laboratory for Computer Science
M.I.T.

http://www.csg.lcs.mit.edu/6.823

Next-Generation Computing Infrastructure

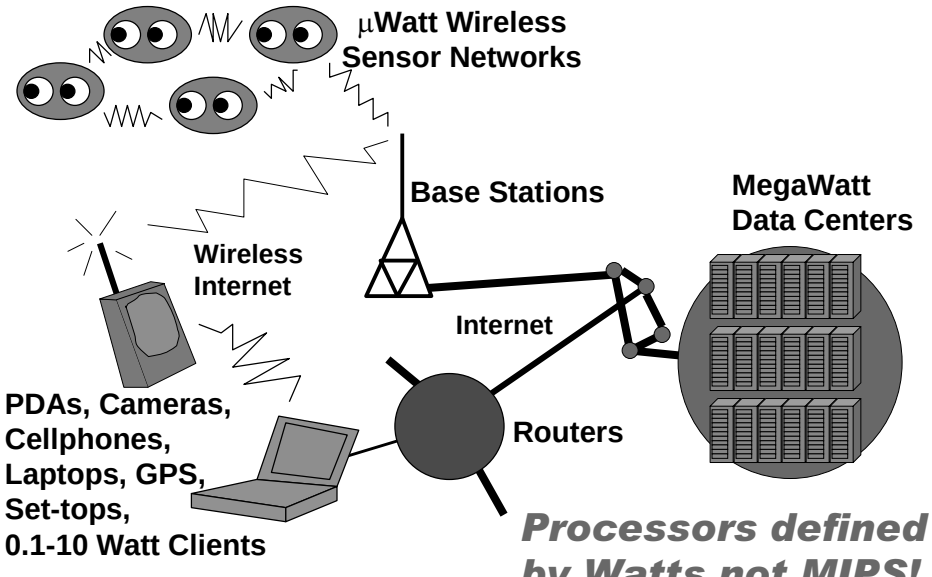

Physics Review

- **Energy measured in Joules**
- **Power is rate of energy consumption measured in Watts (Joules/second)**
- **Instantaneous power is Vdd * Idd**

Power versus Energy

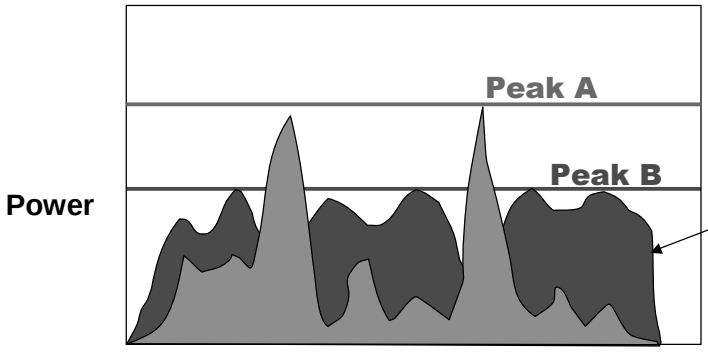

Integrate power curve to get energy

- System **A** has higher peak power, but lower total energy
- System **B** has lower peak power, but higher total energy

Impacts on Computer System

- **Energy consumed per task determines battery life**
  - **Second order effect is that higher current draws decrease effective battery energy capacity (higher power also lowers battery life)**
- **Current draw causes IR drops in power supply voltage**
  - **Requires more power/ground pins to reduce resistance R**
  - **Requires thick&wide on-chip metal wires or dedicated metal layers**
- **Switching current (dI/dt) causes inductive power supply voltage bounce $\propto$ LdI/dt**
  - **Requires more pins/shorter pins to reduce inductance L**
  - **Requires on-chip/on-package decoupling capacitance to help bypass pins during switching transients**
- **Power dissipated as heat, higher temps reduce speed and reliability**
  - **Requires more expensive packaging and cooling systems**
  - **Fan noise**
  - **Laptop temperature**

System Levels for Energy Management

Application	Export computation to server
Algorithm	Variable resolution processing
Source Code	Improved code structure
Compiler	Energy-conscious compiler
Run-Time/O.S.	Just-in-time scheduling
Instruction Set	Energy-exposed architectures
Microarchitecture	Clock gating
Circuit Design	Low voltage-swing circuits
Fabrication Technology	SOI, Low-k dielectrics

Can usually combine savings at different levels

Power Dissipation in CMOS

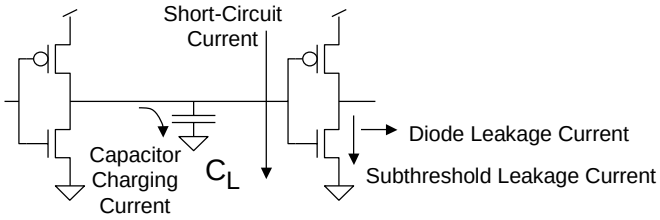

Primary Components:

- **Capacitor Charging (85-90% of active power)**
  - **Energy is ½ CV^2 per transition**
- **Short-Circuit Current (10-15% of active power)**
  - **When both p and n transistors turn on during signal transition**
- **Subthreshold Leakage (dominates when inactive)**
  - **Transistors don't turn off completely**
  - **Becoming more significant part of active power with scaling**
- **Diode Leakage (negligible)**
  - **Parasitic source and drain diodes leak to substrate**

Reducing Switching Power

Power $\propto$ activity * ½ CV^2 * frequency

- **Reduce activity**
- **Reduce switched capacitance C**
- **Reduce supply voltage V**
- **Reduce frequency**

Reducing Activity

Clock Gating

- don't clock flip-flop if not needed
- avoids transitioning downstream logic
- Pentium-4 has hundreds of gated clocks

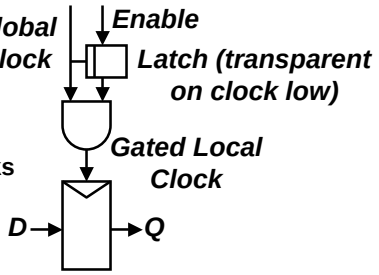

Bus Encodings

- choose encodings that minimize transitions on average (e.g., Gray code for address bus)
- compression schemes (move fewer bits)

Remove Glitches

- balance logic paths to avoid glitches during settling
- use monotonic logic (domino)

Reducing Switched Capacitance

Reduce switched capacitance C

- Different logic styles (logic, pass transistor, dynamic)
- Careful transistor sizing
- Tighter layout
- Segmented structures

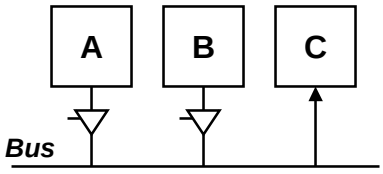

Shared bus driven by A or B when sending values to C

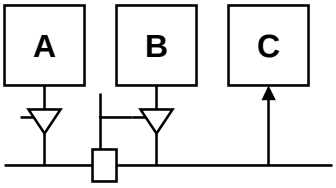

Insert switch to isolate bus segment when B sending to C

Reducing Supply Voltage

Quadratic savings in energy per transition – BIG effect

- **Circuit speed is reduced**
- **Must lower clock frequency to maintain correctness**

Reducing Frequency

- **Doesn't save energy, just reduces rate at which it is consumed**
  - **Some saving in battery life from reduction in rate of discharge**

Voltage Scaling for Reduced Energy

- **Reducing supply voltage by 0.5 improves energy per transition by 0.25**
- **Performance is reduced – need to use slower clock**
- **Can regain performance with parallel architecture**

- **Alternatively, can trade surplus performance for lower energy by reducing supply voltage until "just enough" performance**

Dynamic Voltage Scaling

"Just Enough" Performance

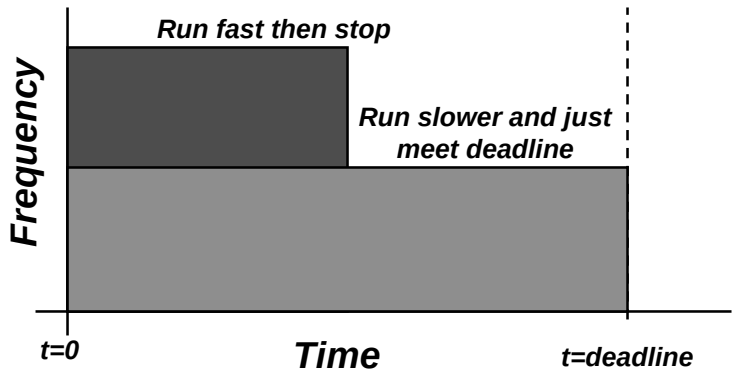

- **Save energy by reducing frequency and voltage to minimum necessary (usually done in O.S.)**

Voltage Scaling on Transmeta Crusoe TM5400

Frequency (MHz)	Relative Performance (%)	Voltage (V)	Relative Energy (%)	Relative Power (%)
700	100.0	1.65	100.0	100.0
600	85.7	1.60	94.0	80.6
500	71.4	1.50	82.6	59.0
400	57.1	1.40	72.0	41.4
300	42.9	1.25	57.4	24.6
200	28.6	1.10	44.4	12.7

Parallel Architectures Reduce Energy at Constant Throughput

- **8-bit adder/comparator**
  - **40MHz at 5V, area = 530 kμ^2**
  - **Base power Pref**
- **Two parallel interleaved adder/compare units**
  - **20MHz at 2.9V, area = 1,800 kμ^2 (3.4x)**
  - **Power = 0.36 Pref**
- **One pipelined adder/compare unit**
  - **40MHz at 2.9V, area = 690 kμ^2 (1.3x)**
  - **Power = 0.39 Pref**
- **Pipelined and parallel**
  - **20MHz at 2.0V, area = 1,961 kμ^2 (3.7x)**
  - **Power = 0.2 Pref**

Chandrakasan et. al. "Low-Power CMOS Digital Design", IEEE JSSC 27(4), April 1992

System Operating Modes

- **Fixed throughput**
  - **e.g., MP3 player**
  - **want to minimize energy at fixed throughput (equivalent to minimizing power)**
- **Maximum throughput**
  - **e.g., spreadsheet update**
  - **want to run "*as fast as possible*"??**
- **How do we trade performance and energy/operation?**

Portable and Wireless Devices

- **Lithium ion/polymer batteries 100-200 Wh/Kg**
  - **200Wh/Kg = 720 J/g**
- **Intel Xscale (StrongARM-2)**
  - **1.5W @ 1GHz, 1.75V = 1.5 nJ/cycle = 2 pg/cycle**
  - **55mW @ 200MHz, 0.7V = 0.3 nJ/cycle = 0.4 pg/cycle**
- **Wireless comms**
  - **1W for 10Mb/s (Wavelan) = 100 nJ/bit = 140 pg/bit**
  - **1W for 20Kb/s (cellphone) = 50,000 nJ/bit = 70,000 pg/bit**

Embedded Computing

What is an Embedded Computer?

- A computer not used to run general-purpose programs, but instead used as a component of a system
- Usually, user cannot change the computer program
- Example applications:
  - Car (some have >100 processors)
  - Cellphone
  - Digital camera (some have several processors)
  - Games machine
  - Router
  - Television

Early Examples

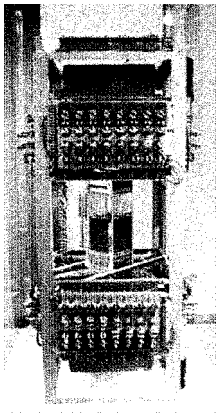

- **MIT Whirlwind, 1946-51**
  - **developed for real-time flight simulator**
- **Intel 4004, 1971**
  - **developed for Busicom 141-PF printing calculator**

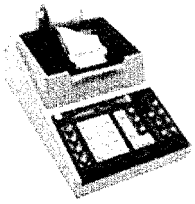

Important Metrics for Embedded Computers

- **Power**
  - **expensive package and cooling affects cost, system size, weight**
- **Manufacturing Cost**
  - **includes cost of supporting structures, particularly memory**
  - **static code size very important (cost of ROM/RAM)**
  - **often ship millions of copies (worth engineer time to optimize cost down)**
- **Real-time performance**
  - ***hard real-time*: if deadline missed system has failed (car brakes!)**
  - ***soft real-time*: missing deadline degrades performance (skipping frames on DVD playback)**
- **Real-world I/O performance**
  - **sensor and actuators require continuous I/O (can't batch process)**

What is Performance?

- **Latency (or response time or execution time)**
  - **time to complete one task**
- **Bandwidth (or throughput)**
  - **tasks completed per unit time**

Performance Guarantees

Inputs

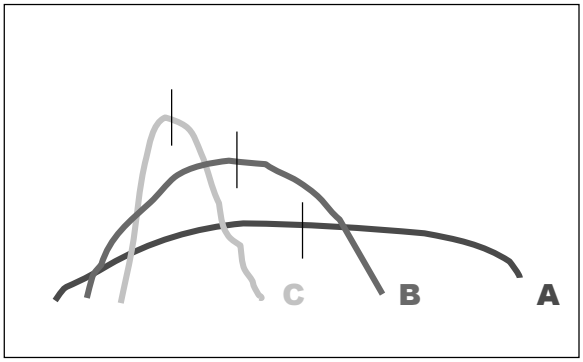

Execution Rate

Average Rate: **A** > **B** > **C**

Worst-case Rate: **A** < **B** < **C**

Types of Embedded Computer

- General Purpose Processors
  - often too expensive, too hot, too unpredictable, and require too much support logic for embedded applications
- Microcontroller
  - emphasizes bit-level operations and control-flow intensive operations (a programmable state machine)
  - usually includes on-chip memories and I/O devices
- DSP (Digital Signal Processor)
  - organized around a multiply-accumulate engine for digital signal processing applications
- FPGA (Field Programmable Gate Array)
  - reconfigurable logic can replace processors/DSPs for some applications

New Forms of Domain-Specific Processor

- Network processor
  - arrays of 8-128 processor cores on a single chip used to process Internet packets
  - used in high-end routers
- Media processor
  - conventional RISC or VLIW engine extended with media processing instructions (SIMD or Vector)
  - used in set-top boxes, DVD players, digital cameras

Programming Embedded Computers

- **Microcontrollers, DSPs, network processors, media processors usually have complex, non-orthogonal instruction sets with specialized instructions**
  - **poor compiled code quality (% peak with compiled code)**
  - **high static code efficiency**
  - **high MIPS/$ and MIPS/W**
  - **usually assembly-coded in critical loops**
- **Worth one engineer year in code development to save $1 on system that will ship 1,000,000 units**
- **Assembly coding easier than ASIC chip design**
- ***But room for improvement…***