Reliability-Aware Power Management

Fei Li

Department of Computer Science
George Mason University
lifei@cs.gmu.edu

August 18, 2010

Lifetime Reliability

MTTP (Mean-Time-To-Failure) is used to describe the expected processor's life. It depends on the processor temperature and utilization.

1. Electromigration:

$$MTTF_{EM} \propto (J)^{-n} e^{\frac{E_{aEM}}{k\mathbf{T}}}.$$

   J is the current density in the interconnect, E_{aEM} is the activation energy for electromigration, k is Boltzmann's constant, T is absolute temperature, and n is a constant.

2. Stress migration:

$$MTTF_{SM} \propto |T_0 - \mathbf{T}|^{m} e^{\frac{E_{eSM}}{k\mathbf{T}}}.$$

   T_0 is the metal's stress-free temperature. m is a constant.

3. Thermal Cycling:

$$MTTF_{TC} \propto C_o(\Delta T - \Delta T_o)^{-q}.$$

   (Coffin-Manson equation, from JEDEC Publication JEP122E: "Failure mechanism and models for semiconductor devices"). C_o is material dependent constant, ΔT is the entire temperature cycle-range for the device, ΔT_o is the portion of the temperature range in the elastic region, and q is the coffin-Mansion exponent.

Challenges

Observation

Switching between different power consumption levels leads to thermal cycling (Mihic et al. DSD 2004).

Goal

Tradeoff between power consumption, performance, and reliability.

Settings

1. On single processor
2. On multi-processor
3. System-level reliability. Overall system reliability depends on the topology of the networks and power management strategies.

Relationship between Frequency and Temperature

For RC model (Skadron et al. TACO 2004), the relationship between frequency (that we can set for processors) and the temperature (that is related to reliability) is

$$T(t) = P \cdot R + T_A - (P \cdot R + T_A - T_i)e^{\frac{-t}{R \cdot C}},$$

where P is the power, R is the thermal resistance, C is the thermal capacitance, T_A is the ambient temperature, the air temperature of environment where computers and related equipment are kept, and T_i is the initial temperature.

Model (Single Processor)

1. A processor has multiple frequency levels $f_1, f_2, \ldots, f_n$.
2. Each job has a release time r_j, a processing time p_j, and a deadline d_j.
3. The energy cost is a cubic function of the frequency.
4. Reliability requirement
   - Option 1: In any time interval with length of T, the difference between the highest frequency and the lowest frequency is bounded by a constant F, or
   - Option 2: For each time of frequency change from f_i to f_j, a cost of C_{ij} (αC_{ij} units of energy cost) is paid.
   - ...
5. Goal: Schedule the jobs by their deadlines with the minimum energy cost.

Model (Multi-Processor)

1. There are m processors.
2. Each processor has multiple frequency levels $f_1, f_2, \ldots, f_n$.
3. Each job has a release time r_j, a processing time p_j, and a deadline d_j.
4. Jobs can (or cannot) migrate among processors.
5. The energy cost is a cubic function of the frequency.
6. Reliability requirement
   - Option 1: For a given processor P_i, in any time interval with length of T_i, the difference between the highest frequency and the lowest frequency is bounded by a constant F_i, or
   - Option 2: For a given processor P_i, each time of frequency change from f_j to f_k, a cost of C^i_{jk} (αC^i_{jk} units of energy cost) is paid.
   - ...
7. Goal: Schedule the jobs by their deadlines with the minimum energy cost.

Model (Network-on-Chip)

1. Consider a graph of m nodes, where each node represents a processor.
2. Each processor has multiple frequency levels $f_1, f_2, \ldots, f_n$.
3. Each job has a release time r_j, a processing time p_j, and a deadline d_j.
4. Jobs can (or cannot) migrate among processors.
5. For any two processors P_i and P_j sharing a link, if both processors are being active, the difference between their frequencies at the time cannot be over a constant D_{ij}.
6. The energy cost is a cubic function of the frequency.
7. Reliability requirement
    - Option 1: For a given processor P_i, in any time interval with length of T_i, the difference between the highest frequency and the lowest frequency is bounded by a constant F_i, or
    - Option 2: For a given processor P_i, each time of frequency change from f_j to f_k, a cost of C^i_{jk} (αC^i_{jk} units of energy cost) is paid.
    - ...
8. Goal: Schedule the jobs by their deadlines with the minimum energy cost.