LEAKAGE CONSCIOUS DVS SCHEDULING FOR PEAK TEMPERATURE MINIMIZATION

Vivek Chaturvedi   Gang Quan
Florida International University
Miami, FL, USA

OUTLINE

- Introduction
- Related work
- System models
- Fundamental principles on peak temperature minimization
- Summary

THE EXPONENTIALLY INCREASED POWER DENSITY

Sources: C. Isci, ISAC, 2007

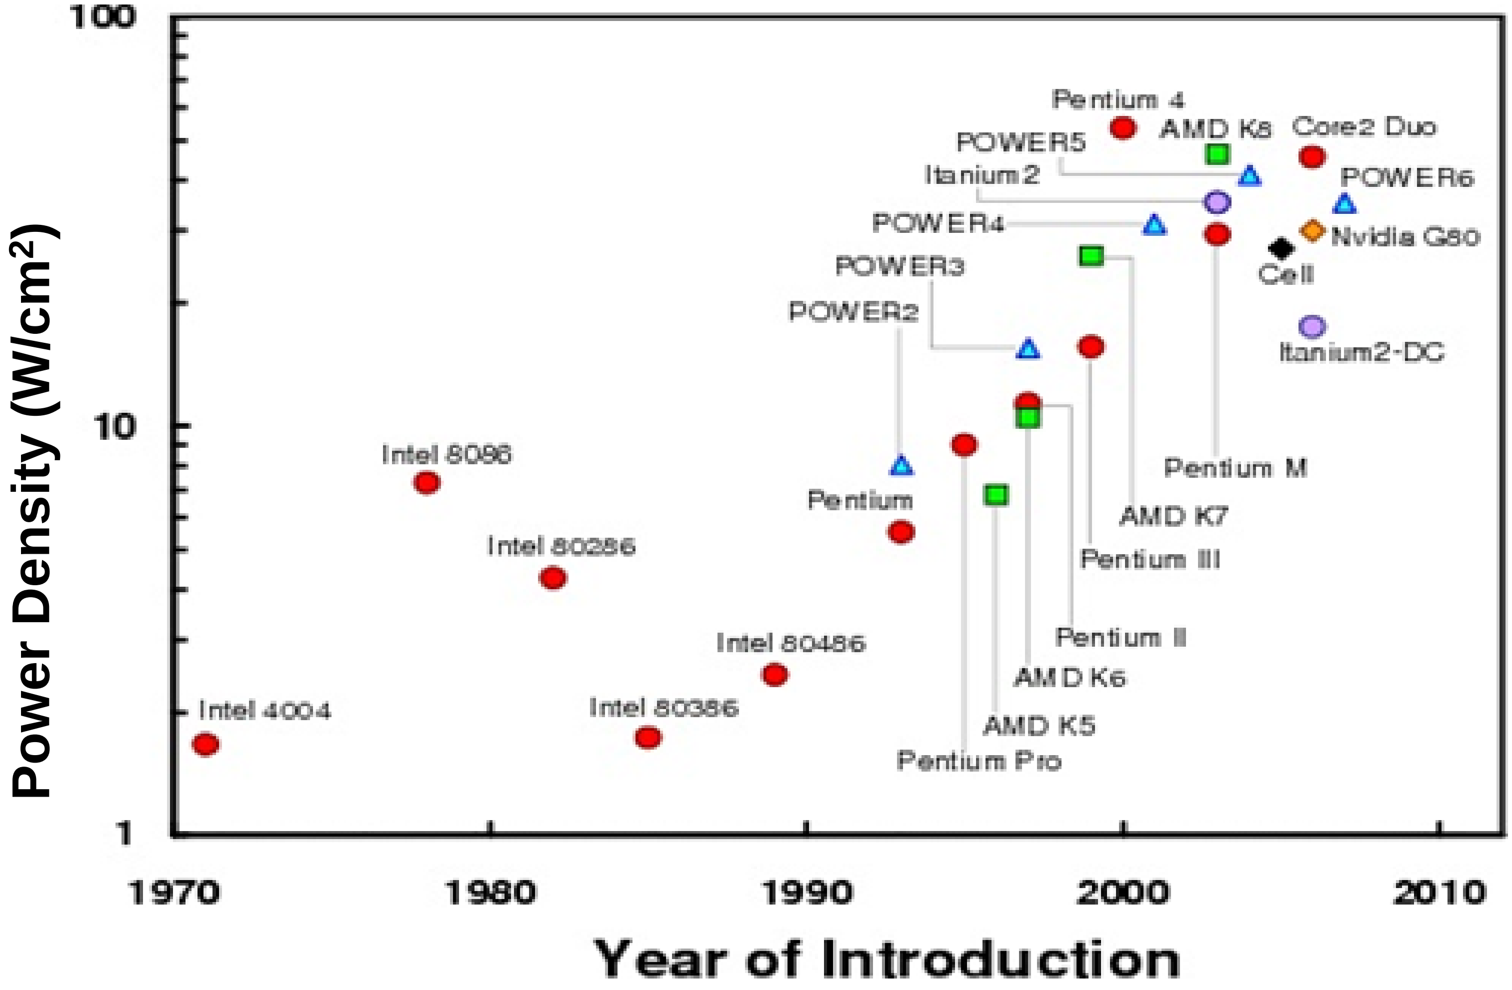

WHY TEMPERATURE MATTERS

- High cooling/packaging cost
  - 1-3 dollar/watt more packaging/cooling
  - In data center, for every one watt computing, 1 to 1½ watt for cooling
- Reliability:
  - 10^{o}C rise in temperature can result in 50% reduction in system life span
- Performance:
  - 15^{o}C rise in temperature can add approximately 10 - 15% circuit delay
- Other issues
  - Increase leakage power consumption

THE LEAKAGE POWER CONSUMPTION

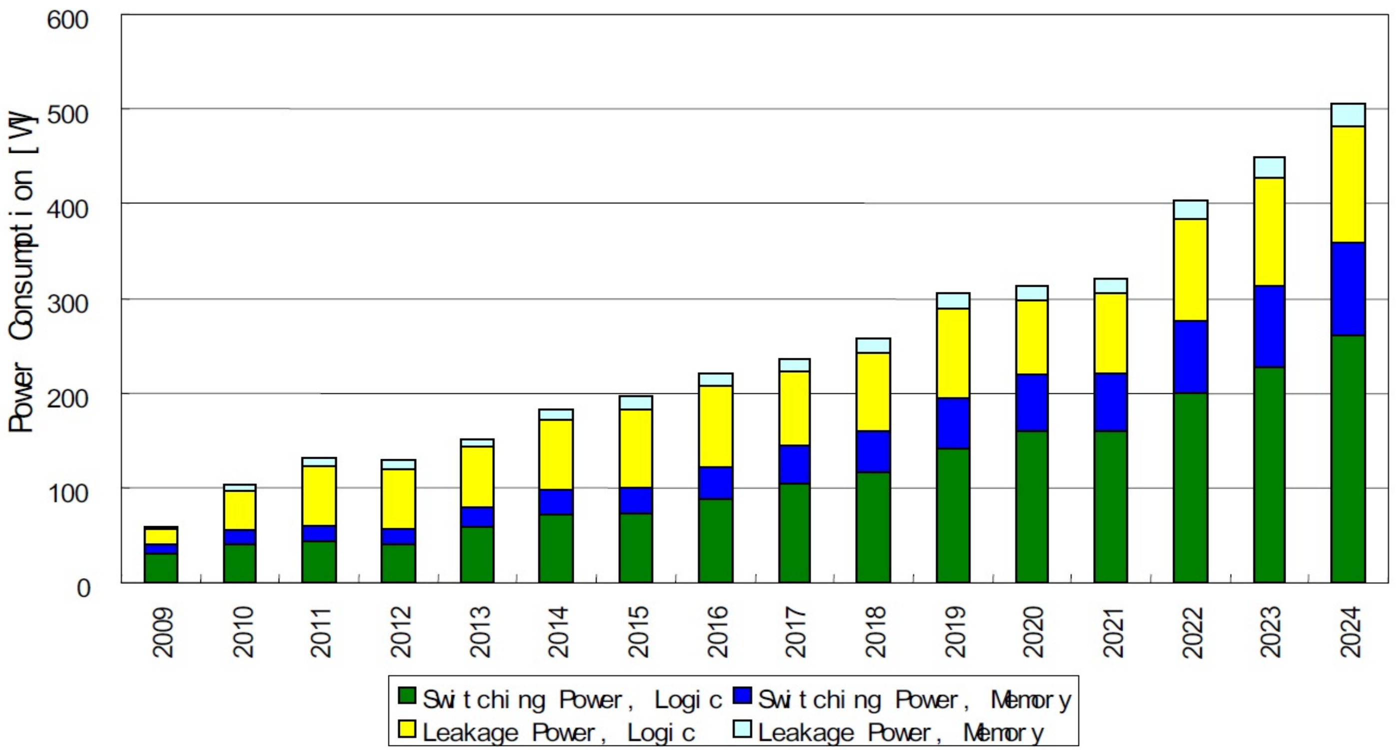

Figure SYSD11 SOC Consumer Stationary Power Consumption Trends

Leakage power consumption becomes a significant component in the overall power consumption at the deep submicron domain

THE LEAKAGE TEMPERATURE INTERPLAY

$$I_{leak} = I_s(\mathcal{A} \cdot T^2 \cdot e^{((\alpha V + \beta)/T)} + \mathcal{B} \cdot e^{(\gamma V + \delta)})$$

T : Temperature,
V : Supply voltage,
I_s, $\mathcal{A}$, $\mathcal{B}$, α, β, γ, δ: Technology dependent constants

- The positive feedback loop
  - high power consumption → high temperature → high leakage power → high power consumption

The leakage/temperature dependency becomes critical in thermal aware computing !

THERMAL AWARE VS. POWER AWARE COMPUTING

- Closely related
  - Low power consumption → low heat generation → low temperature

- Distinct different
  - Optimal power aware solutions are not necessarily optimal for temperature reduction
    - ref. Skadron et al. 2003, Bansal et al. 2007

The extensive power aware techniques cannot be readily used for thermal aware computing. New techniques need to be developed.

THE PROBLEM

- How do we adjust the system's performance so that the peak temperature can be minimized?
  - A real-time job within a given interval
  - A periodic real-time task

RELATED WORK

- Dynamic power consumption reduction (e.g. Yao et al. 1995, Ishihara et al. 1998, Pillai et al., 2001)
- Two fundamental principles
    - Principle 1: Using the lowest constant speed is the schedule that consumes the minimum dynamic energy
    - Principle 2: If a single lowest constant speed is not available, then using the two closest neighboring speeds is the optimal solution in dynamic energy reduction

RELATED WORK

- Overall power reduction assuming constant leakage (e.g. Jejurikar et al. 2004, Quan et al. 2005)
  - Reducing both dynamic and leakage power consumption
  - Constant leakage
  - No temperature/leakage dependency

RELATED WORK

- Thermal aware scheduling with no temperature /leakage dependency (e.g. Wang et al, 2006, Zhang et al. 2007)
  - No leakage power or constant leakage power consumption

- Power/thermal aware scheduling with leakage/temperature dependency
  - Leakage power changes with only temperature (e.g. Chen et al. 2009, Chantem et al. 2009)
  - Leakage power changes with both temperature and supply voltage (e.g. Quan et al 2009)

SYSTEM MODELS

- Power model
  - $P_{Total} = P_{dyn} + P_{leak}$
  - $P_{dyn} = C_2 V^3_{dd}(k)$
  - P_{leak}
    - $P_{leak} = C_0(k) . V_{dd}(k) + C_1 T$
    - Varies with both temperatures and supply voltages

SYSTEM MODELS

- Thermal model

$$\frac{dT(t)}{dt} = aP(t) - bT(t)$$

  - T(t): temperature
  - *P(t)*: the power consumption
  - a,b: cooling constants

- Commonly used chip level thermal model (e.g. Chantem et al. 2009, Chen et al. 2009, Quan et al. 2009)

MOTIVATIONS

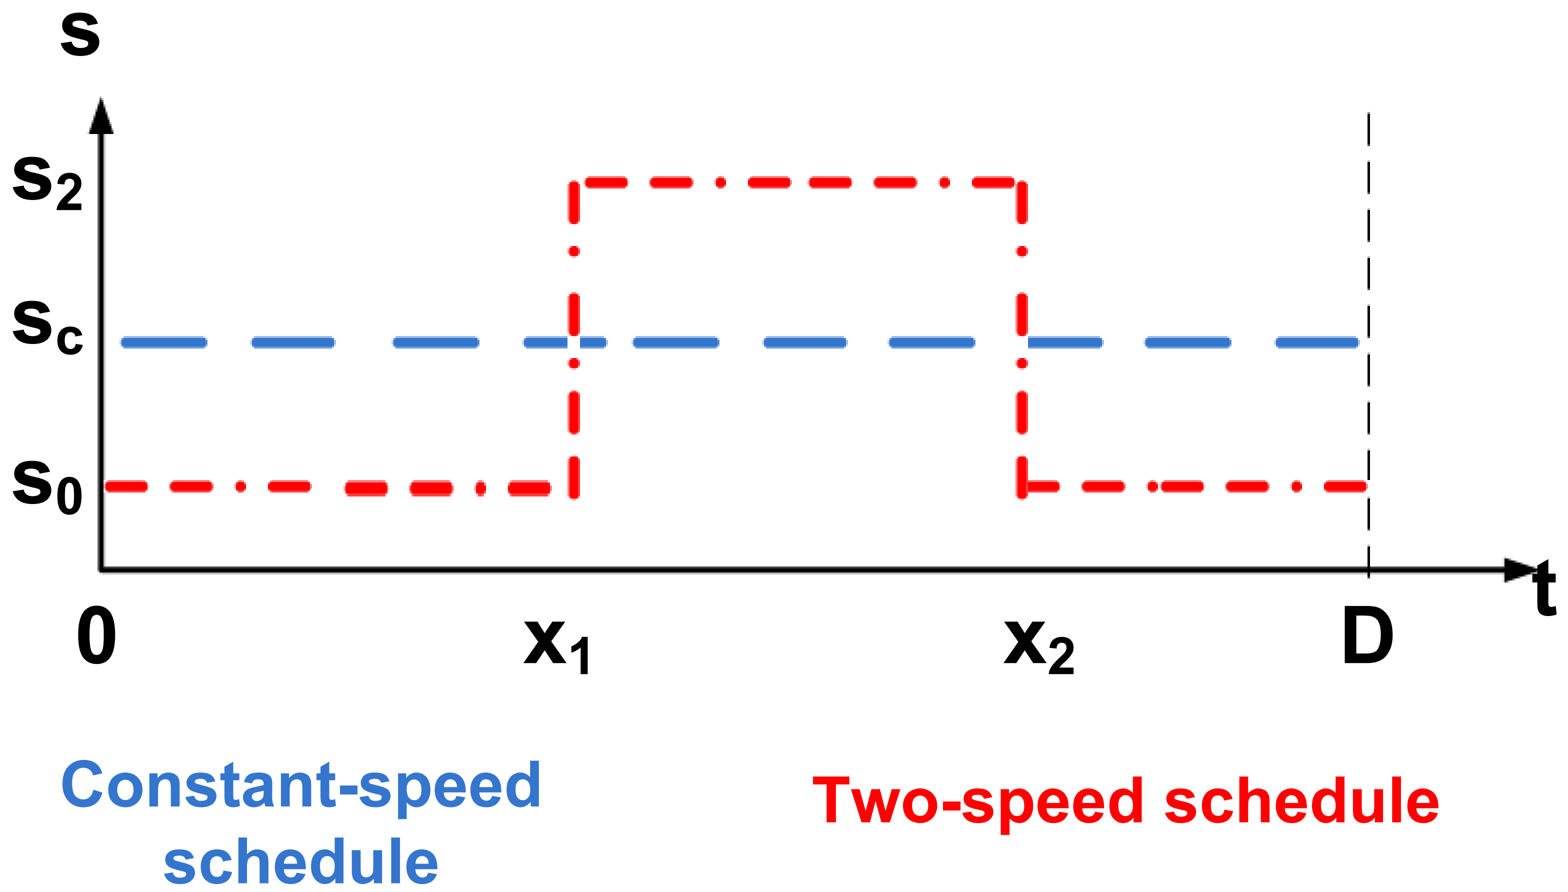

Is the constant speed schedule or the neighboring two-speed schedule still the optimal choice in peak temperature reduction?

DIFFERENT TWO-SPEED SCHEDULES

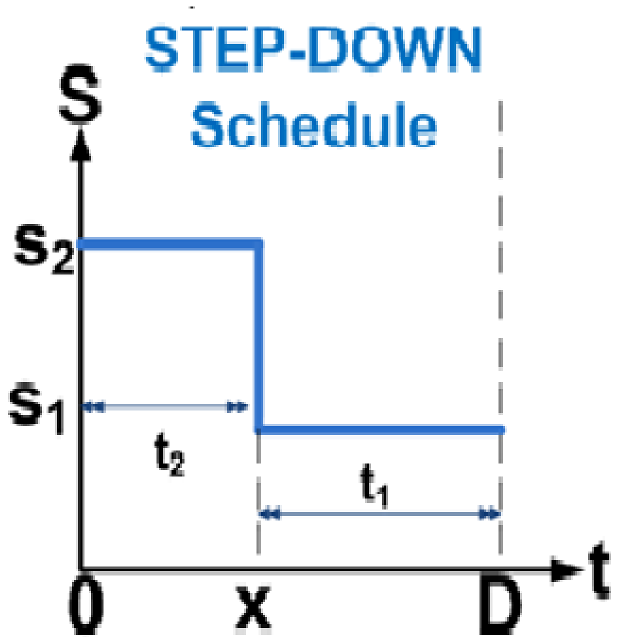

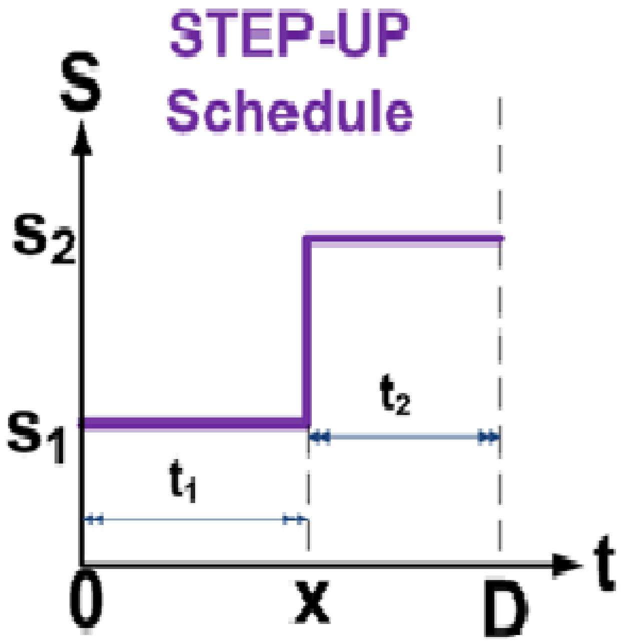

EMPIRICAL STUDIES

- Setup
  - Based on UC Berkley's BSIM device model
  - 65nm technology
  - Conventional air cooling
  - Ambient temperature 25^{o}C
  - Available supply level: 0.6v : 0.05v : 1.3v

PEAK TEMPERATURES FOR INTERVALS WITH DIFFERENT LENGTHS

- Constant speed schedule
  - **0.80 V**
- Two two-speed schedules
  - **(0.75V, 0.85 V)**
  - Step-down schedule
  - Step-up schedule

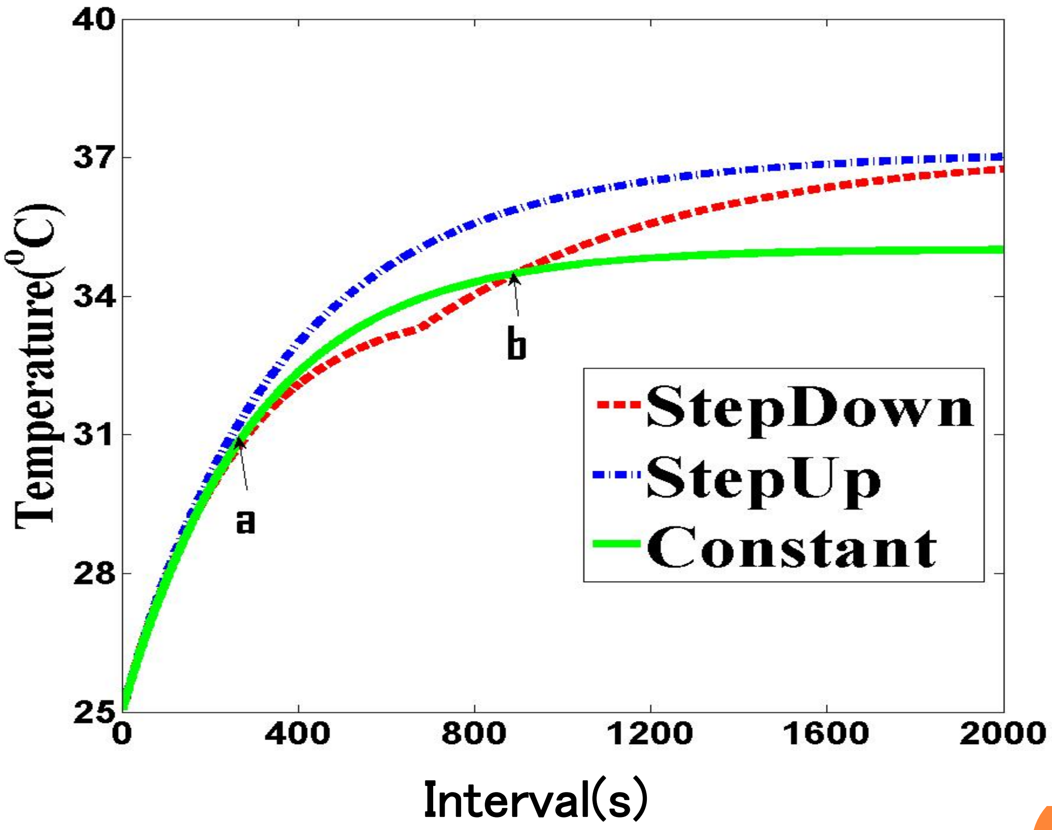

Constant schedule is not always the best choice anymore!

PEAK TEMPERATURES FOR INTERVALS WITH DIFFERENT LENGTHS

- Neighboring two-speed schedules Sa
  - **(0.75V, 0.85 V)**
  - Step-down schedule
- Non-neighboring two-speed schedules Sb
  - **(0.75V, 0.9 V)**
  - Step-down schedule

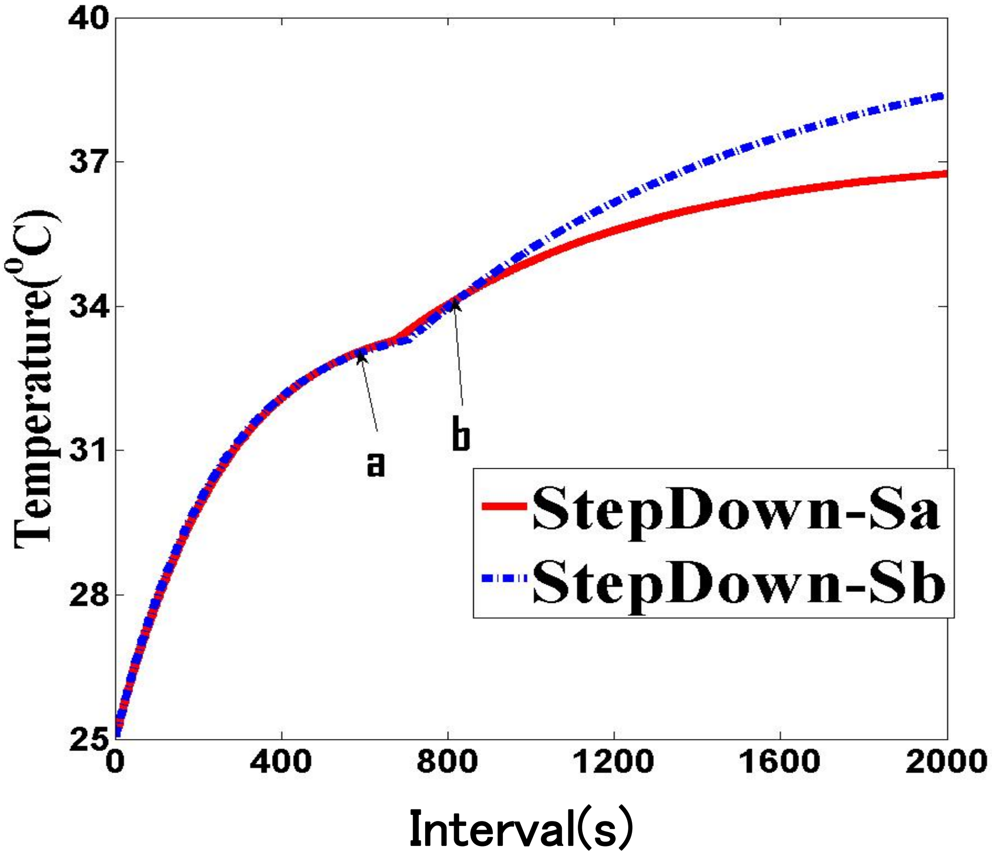

The neighboring two-speed schedule does not always outperform the non-neighboring two-speed schedule!

FUNDAMENTALS ON PEAK TEMPERATURE REDUCTION WITHIN AN INTERVAL

- **Theorems formally formulated and proved**
  - **When the lowest constant speed can still outperform other two-speed schedules in reducing the peak temperature**
  - **Step-up schedule always results in the highest peak temperature among all two-speed schedules**
  - **A two-speed schedule with two neighboring speeds is not always better than two non-neighboring speeds schedule**

PEAK STABLE TEMPERATURE FOR TASKS WITH DIFFERENT PERIODS

- Constant speed schedule
  - **0.80 V**
- Two two-speed schedules
  - **(0.75V, 0.85 V)**
  - Step-down schedule
  - Step-up schedule

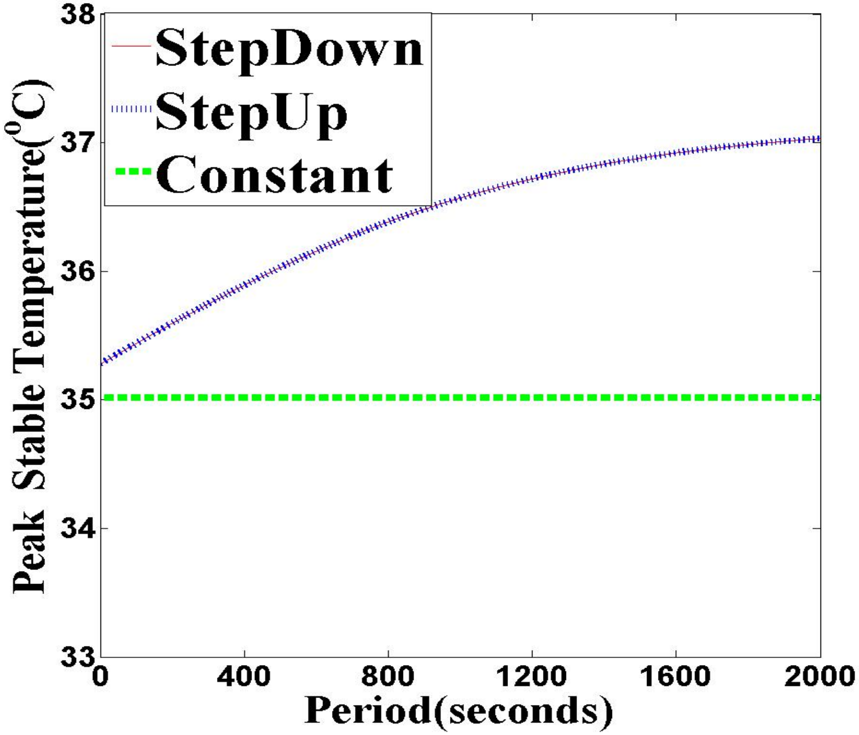

Constant schedule seems to still be the best choice !

PEAK STABLE TEMPERATURES FOR TASKS WITH DIFFERENT PERIODS

- Neighboring two-speed schedules Sa
  - **(0.75V, 0.85 V)**
  - Step-down schedule
- Non-neighboring two-speed schedules Sb
  - **(0.75V, 0.9 V)**
  - Step-down schedule

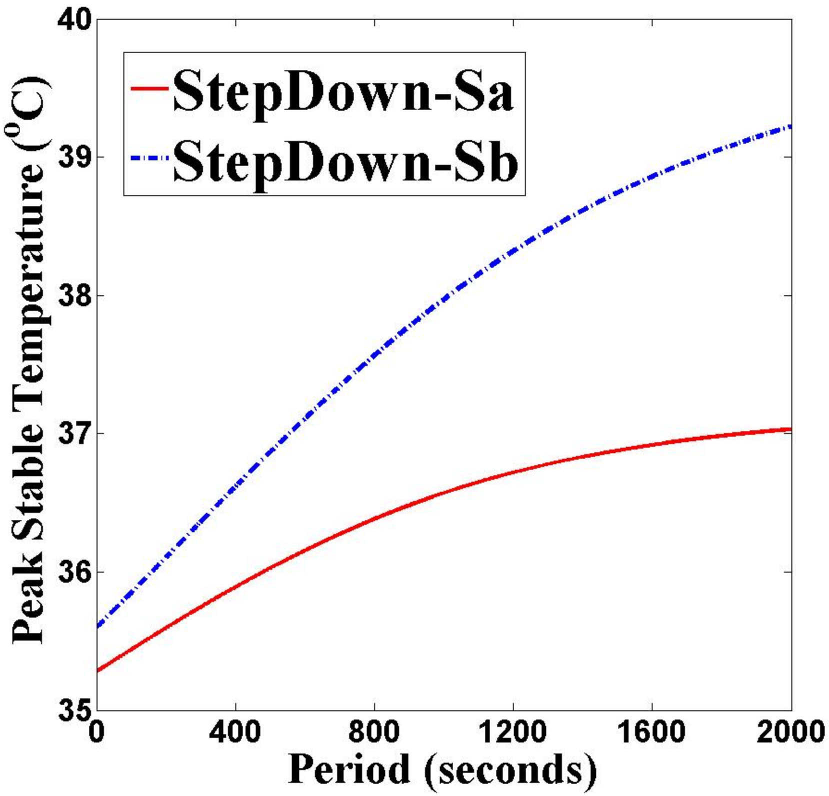

The neighboring two-speed schedule seems to consistently outperform the non-neighboring two-speed schedule

FUNDAMENTALS ON PEAK STABLE TEMPERATURE MINIMIZATION

- Theorems formally formulated and proved
  - The lowest constant speed schedule outperforms any two-speed schedule in minimizing the peak temperature at the stable status.
  - The peak temperature at the stable status by the neighboring two-speed schedule is no more than that by a non-neighboring two-speed schedule

SUMMARY

- Temperature does matter !
- Power aware and thermal aware computing related but distinctly different
- Leakage/temperature dependency is critical in thermal aware design at the deep submicron domain
- Establish several fundamental principles and guidelines to minimize the peak temperature
  - Within a given interval
  - Stable status

Thank You!