Necessary and Sufficient Conditions for Thermal Feasibility of Periodic Real-Time Tasks

Rehan Ahmed, Parameswaran Ramanathan and Kewal K.Saluja
Department of Electrical and Computer Engineering
University of Wisconsin Madison
Madison WI 53706

Outline

- Introduction
  - What is a real-time system?
  - Why are thermal constraints important?
- Proposed solution
  - Thermally optimal scheduling algorithm for unicore
  - Extensions to multi-core
  - Thermally optimal partitioned scheduling solution for multicore

Introduction: Why Thermal Constraints are Important

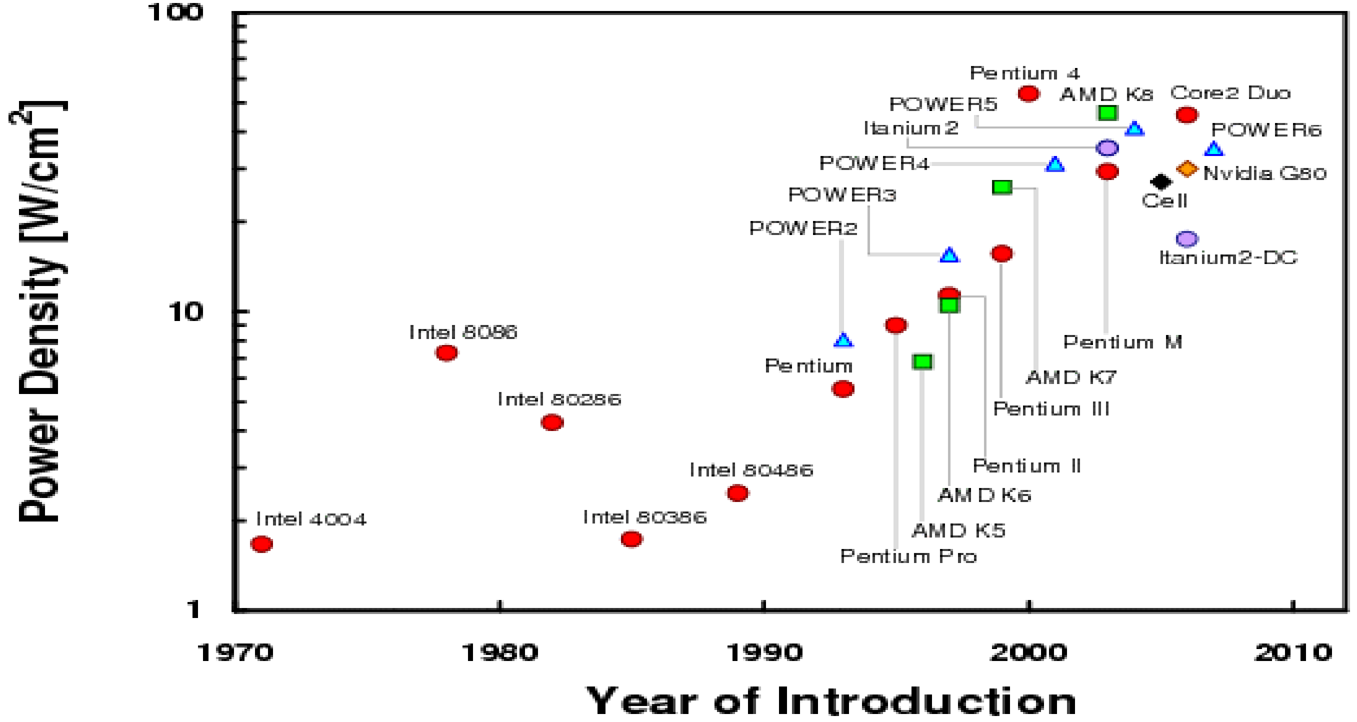

- Rapid rise in Power densities of Integrated Circuits
- Localized power densities 2 orders of magnitude higher than average power densities
- High power densities cause thermal hotspots decreasing reliability/performance

Our Work

- This work considers the scheduling of periodic tasks with thermal and timing constraints
  - All periodic task deadlines have to be met
  - System temperature has to be less than Δ
- Propose a thermally optimal scheduling strategy for uni-core
- Propose a thermally optimal partitioned scheduling algorithm for execution on multi-core

Challenges

- Processor temperature typically reduced by employing DVFS
- DVFS causes reduction in processor performance
  - Lead to deadline miss in real-time systems
- Additional challenges exist for multi-core systems

Thermal Models

Duality between heat transfer and electrical phenomenon:

- Heat transfer modeled as current passing through resistance
- Delay in heat increase modeled as a thermal capacitance
- RC pair for each architectural Unit

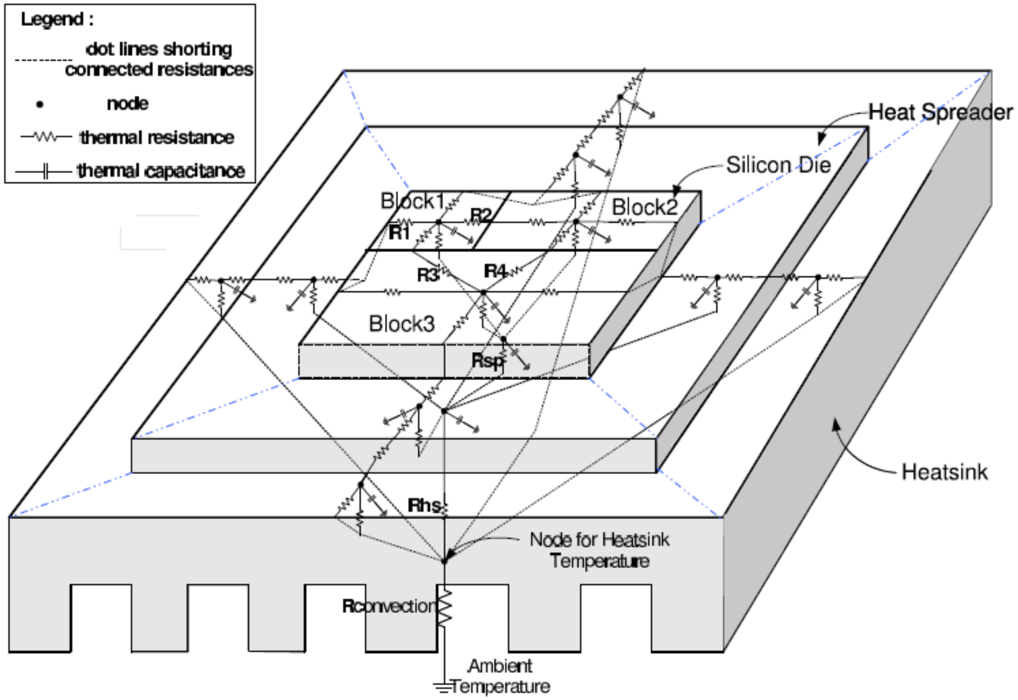

Total Thermal Impact

- Example
  - Task τ with computation time of $C_\tau = 0.5\sec$
  - Power consumption of τ is constant = P_τ

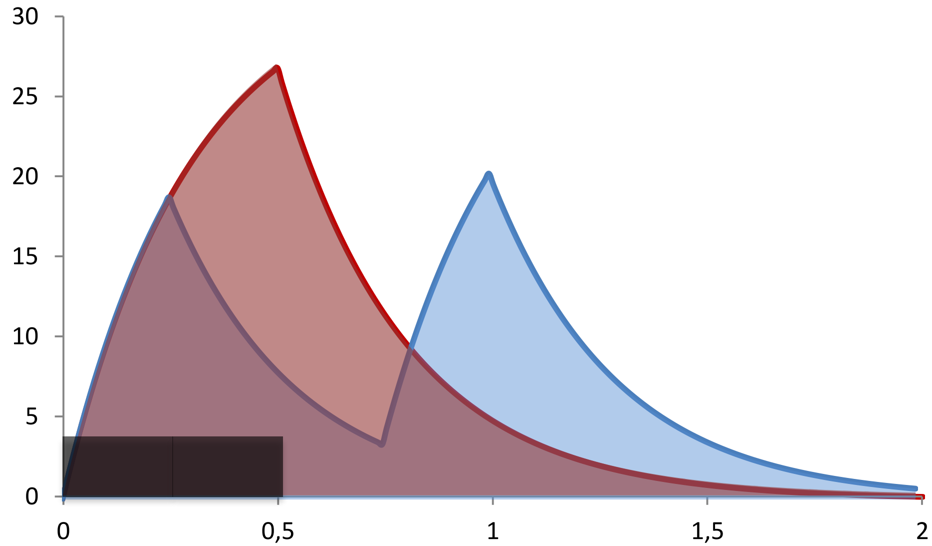

  - TTI of τ is the area under its thermal profile (Red region)
  - TTI is only dependant on the energy consumption of τ

Thermal Steady State

- Temperature at the start of successive hyperperiods increases monotonically

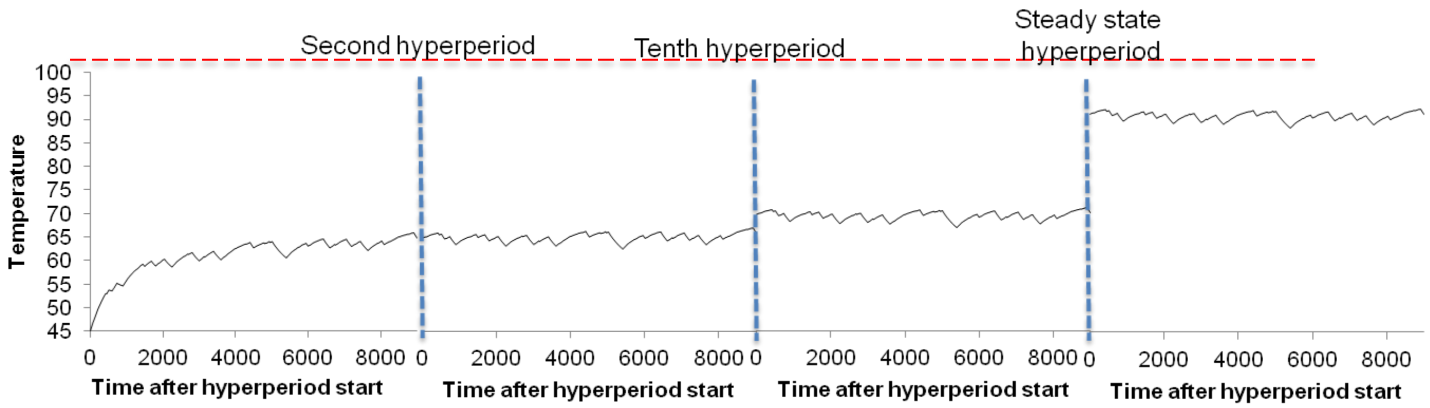

- Thermal Steady State
  - Temperature at the start and end of hyperperiod are equal

Theoretical Results

- Integral of processor temperature at thermal steady state is independent of the schedule. (RTCSA 2013)
  - Integral is equal to the TTI of all periodic task instances within the hyperperiod
  - Integral is only a function of the periodic taskset.

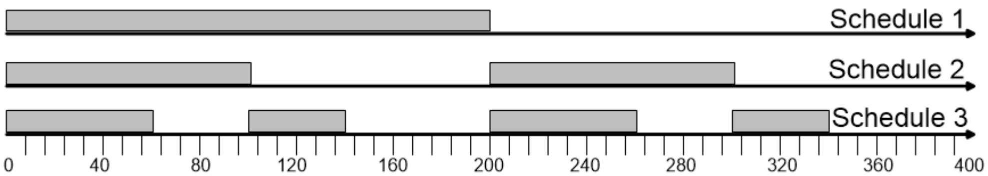

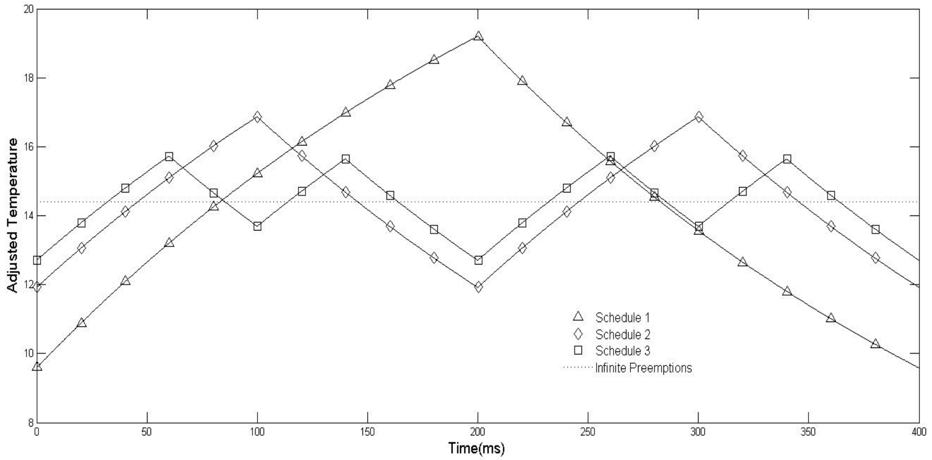

Thermal Utilization: Single Core

- Thermal Utilization of a taskset defined as

$$\frac{LB(\theta_{\max})}{\Delta}$$

- where $LB(\theta_{\max})$ is a lower bound on maximum temperature over all possible scheduling policies
- Hence, tasksets with thermal utilization >1 are not thermally feasible

Scheduling Strategy (GPS)

- Each periodic task is executed at a constant rate equal to

$$\frac{C_\tau}{T_\tau}$$

  - Distributes power uniformly across the hyperperiod
  - Optimal in the sense that if GPS cannot meet thermal and deadline constraints of a taskset, then no other scheduling policy can meet the same constraints
  - On unicore processors, GPS guarantees thermal feasibility if thermal utilization ≤1
    - Ideal policy which cannot be implemented in practice
    - Can be approximated well by Worst-case Fair Weighted Fair Queueing (WF2Q+)

GPS: Unicore Results

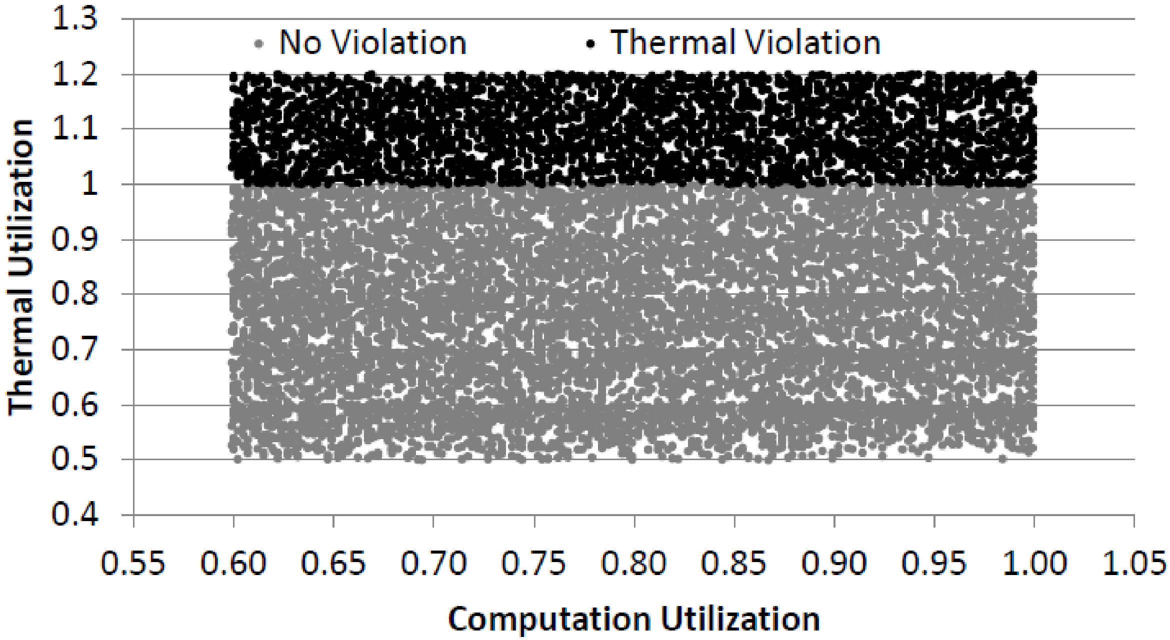

- Empirically validates that thermal utilization less than or equal to 1 is necessary and sufficient for thermal feasibility

Additional Unicore Results

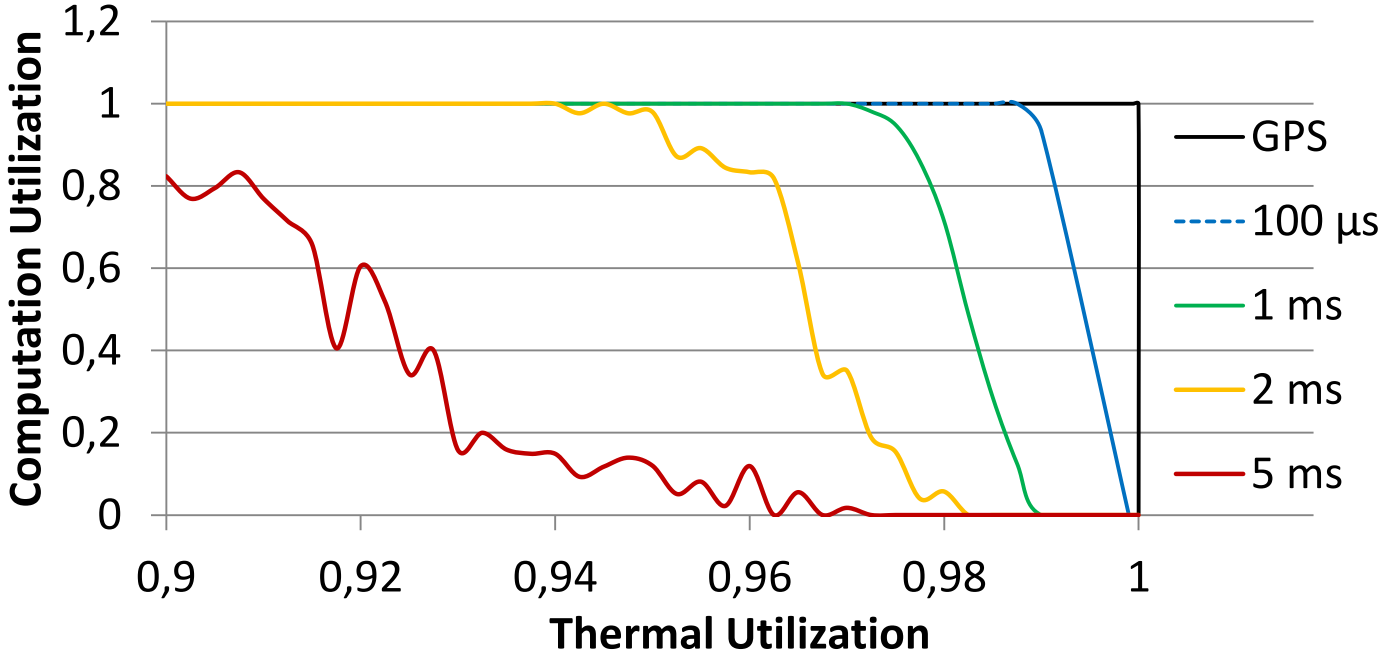

WF2Q+ successfully schedules more tasksets when minimum preemption duration decreases

WF2Q+ performs very well for reasonable minimum preemption duration

Multicore Extension

- Unit Thermal Impact ζ is a square matrix with dimensions equal to number of core
  - $\zeta_{i,j}$ is the TTI of Core *i* when core *j* consumes unit power.
- Consider partitioned scheduling of periodic tasks on multicore
  - If each task τ is assigned to core j:

$$\int_0^L \Theta(t)dt = L \cdot \sum_{\tau \in \Gamma} \frac{P_\tau C_\tau}{T_\tau} \zeta e_j$$

  - Thermal utilization of core *i* is defined as:

$$\frac{1}{L\Delta} \cdot \int_0^L \Theta(t)_i dt$$

Thermal Utilization Minimization using Partitioned Scheduling (TRuMPS)

- Partitioned scheduling
  - All instances of a periodic task execute on the same core
- TRuMPS
  - Formulate task assignment to core as a Mixed Integer Program (MIP) with an objective of minimizing the maximum thermal utilization across all cores
  - Use GPS/WF2Q+ to execute task instances on each core

TRuMPS: Multicore Results

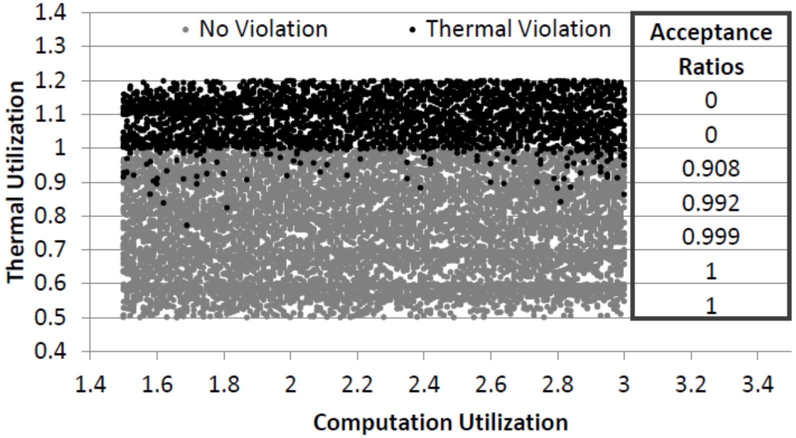

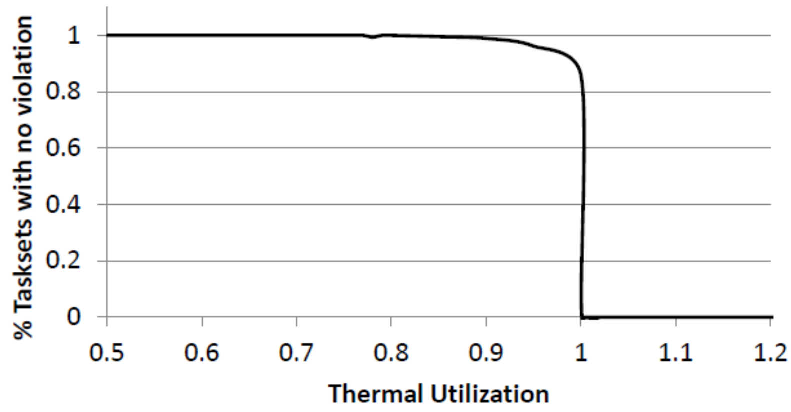

- No taskset with thermal utilization >1 is schedulable
- Some tasksets with thermal utilization less than or equal to 1 are not schedulable, possibly due to partitioned scheduling

Multicore Results Cont'd

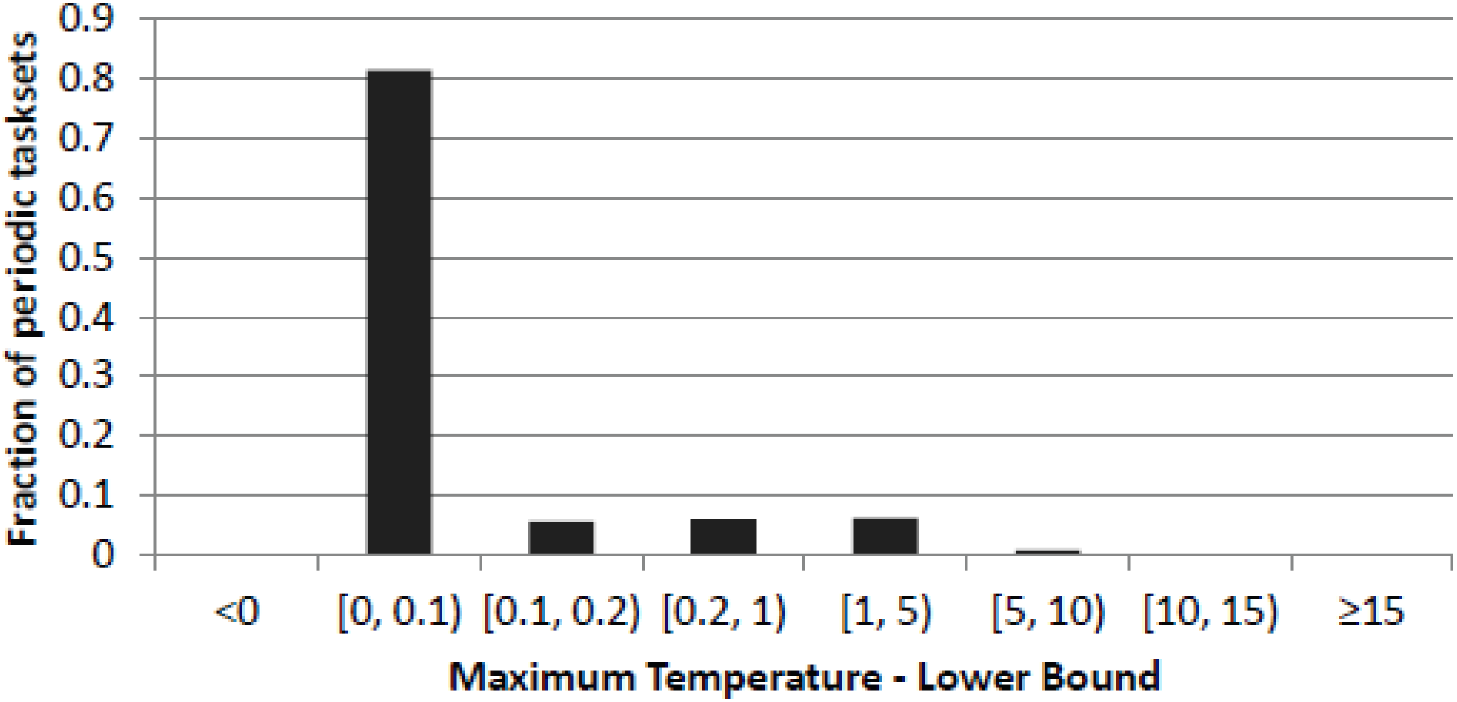

- Approach close to the lower bound on temperature in most cases
- Some tasksets have large temperature difference due to partitioned scheduling approach

Multicore Results Cont'd

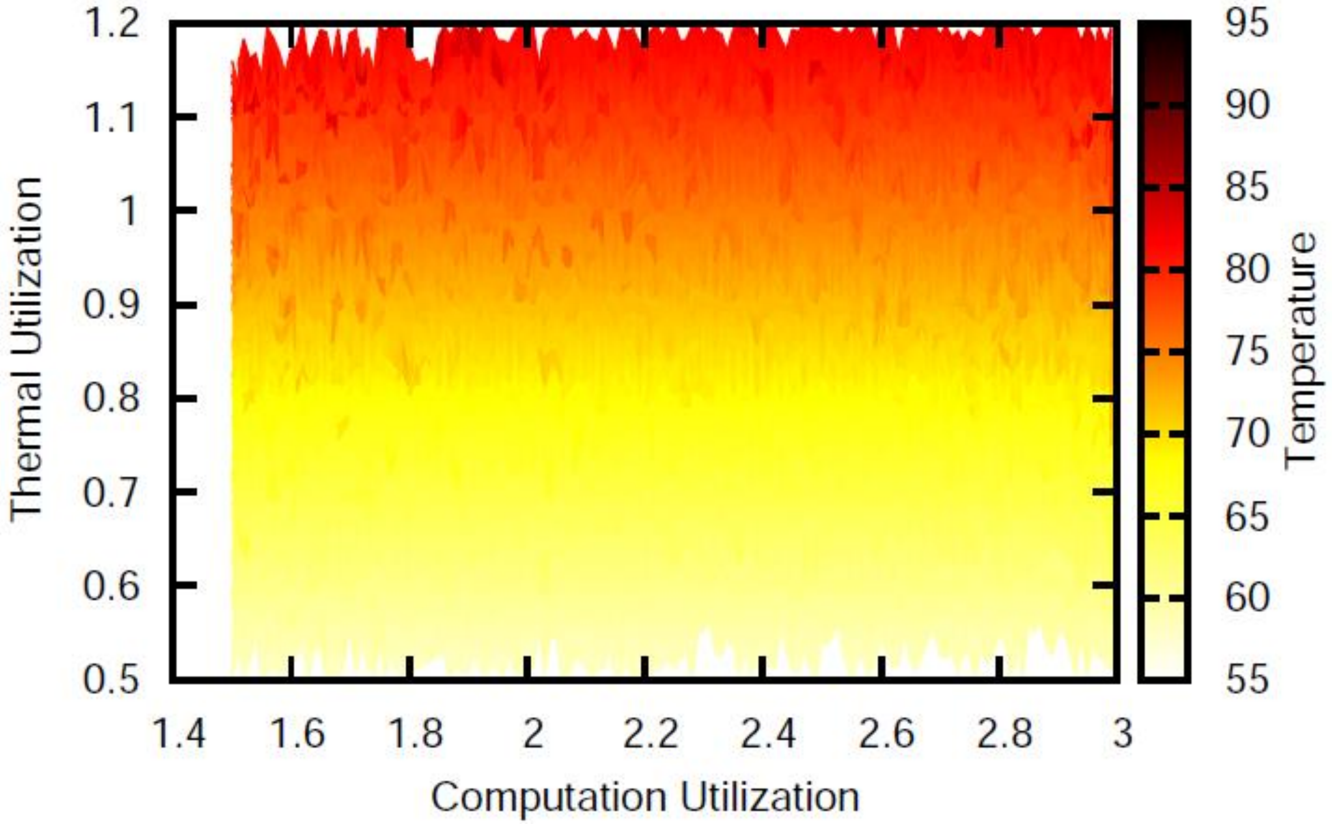

- Maximum temperature and thermal utilization have strong correlation

Presentation Takeaways

- Necessary and Sufficient Conditions for Thermal Schedulability
- GPS is a thermally optimal scheduling scheme for unicore
  - No scheduling algorithm can have lower maximum temperature
- TrUMPS performs very well for multi-core
  - Optimal partitioned scheduling scheme
  - Performance loss due to No-Migration task execution model

Total Thermal Impact contd..

- Unit Thermal Impact ζ is defined as the TTI when 1W power is consumed for 1 second
- Due to linearity of RC network TTI of τ is the $\zeta \cdot P_\tau \cdot C_\tau$
- TTI is only dependant on energy consumption NOT how that energy is consumed.

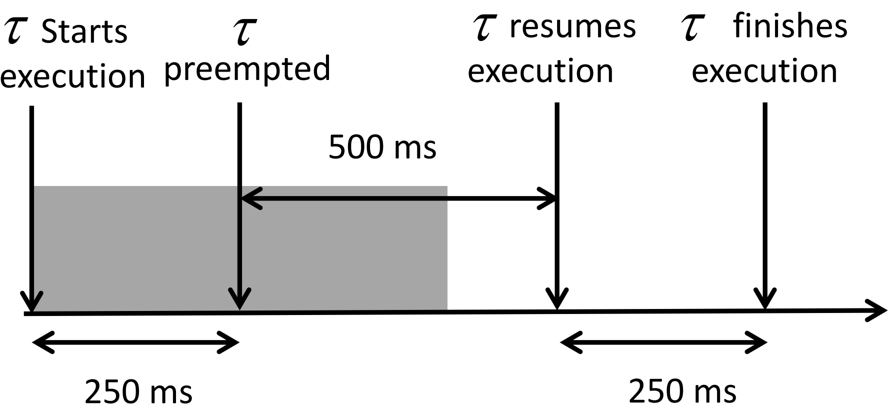

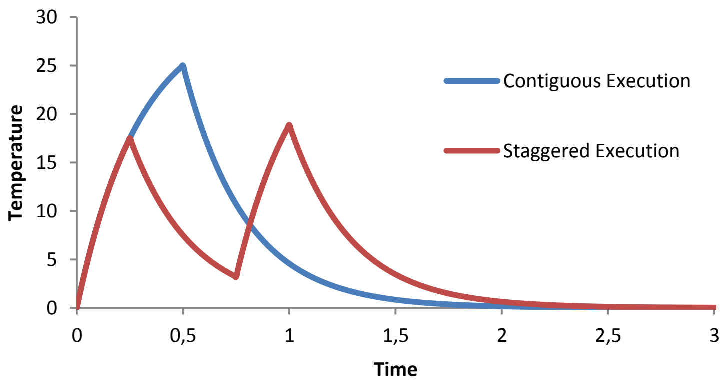

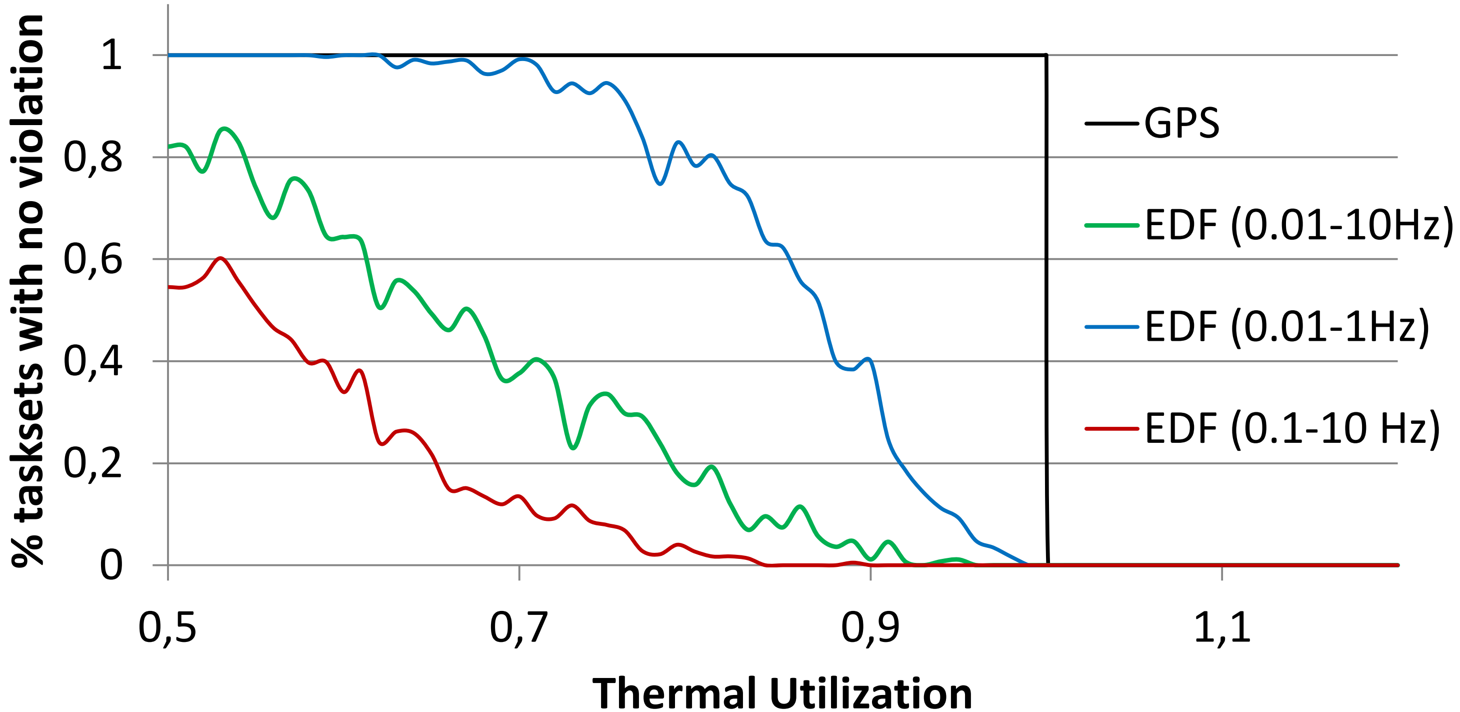

- High Frequency periodic tasks less likely to cause thermal violations
- CRS performance does not depend on frequency of periodic tasks

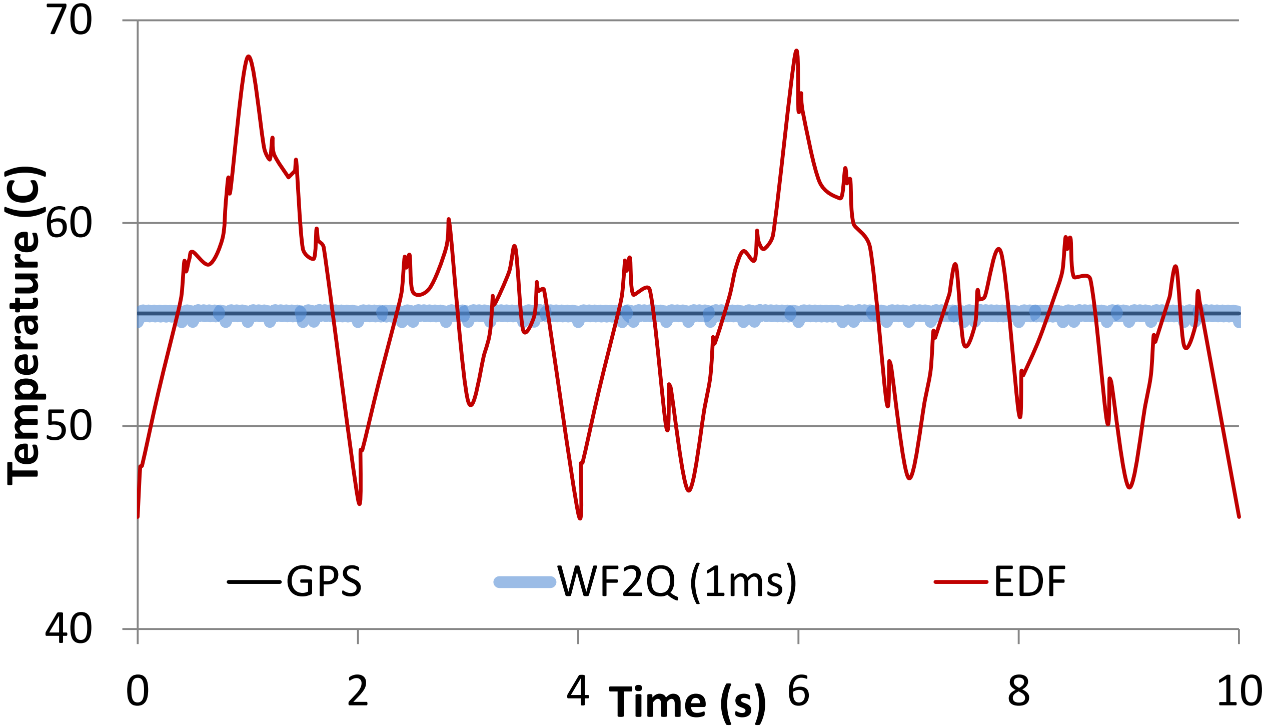

- EDF has large variations in temperature
- DCRS and CRS have similar performance

Multicore Extension

- Challenges
  - Thermal Interaction between cores adds additionally complexity to the thermal scheduling problem.

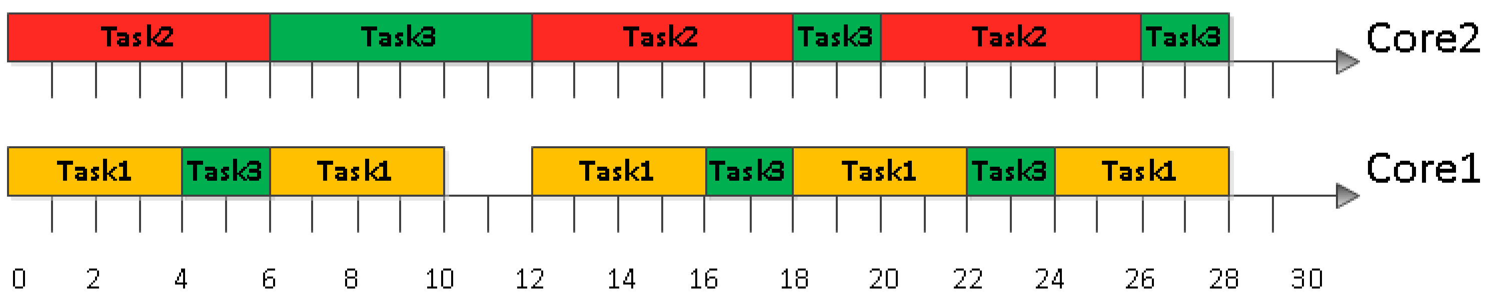

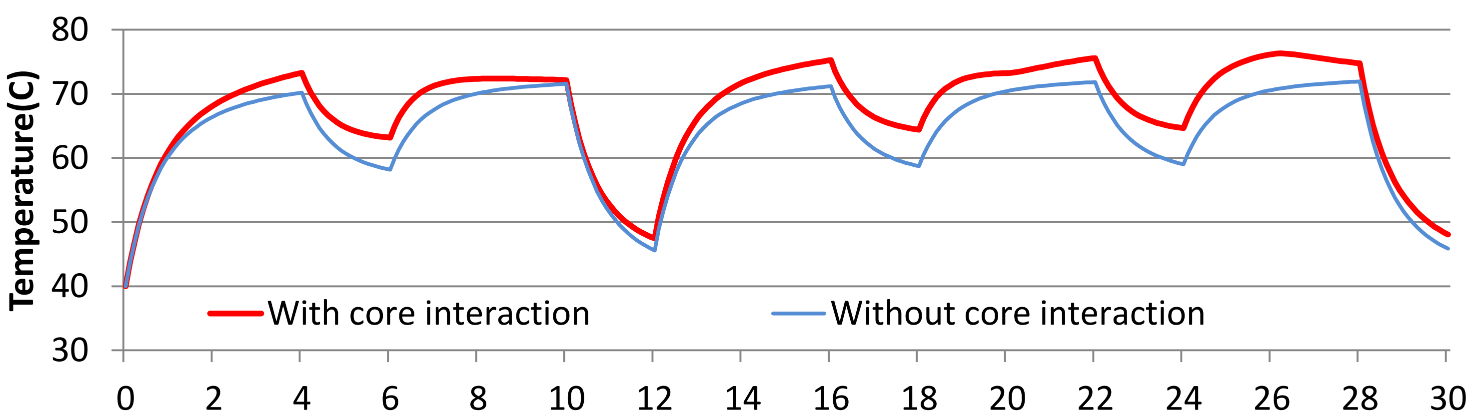

On going research

- Work on Global scheduling approaches on multi-core platform
  - Formulation of better scheduling schemes for multi-core
  - Schemes consider task-migration/instance migration model