Thermally Optimal Stop-Go Scheduling of Task Graphs with Real-Time Constraints

Pratyush Kumar Lothar Thiele Computer Engineering and Networks Laboratory ETH Zürich, Switzerland

ASP-DAC 2011, Yokohama,

26th January, 2011



Outline











The heat is on!

• Exponential rise in processor power densities



- Consequentially high on-chip temperatures
- Hot chips can cause short-term *functional errors* and long-term *reliability degradation*

What are the solutions?

- Better VLSI design
 Thermal-aware design, new materials ...
- e Hardware cooling solutions

Use better heat sinks, fans, air cooling (or even water cooling)

What are the solutions?

- Better VLSI design
 Thermal-aware design, new materials ...
- Hardware cooling solutions
 Use better heat sinks, fans, air cooling (or even water cooling)
- Voltage (and frequency) scaling (DVS)
 When temperatures are too high, reduce supply voltage V_{DD} (and slow down the processor)
- Stop-go execution

When temperatures are too high, completely turn off the processor and peripherals

Cool real-time

- Performance degrades with DVS or stop-go execution
- For real-time applications it is important to ensure that tasks *do not miss deadlines* when such techniques are used
- Clearly, this is a case for *co-design*: thermal and real-time objectives must be considered simultaneously
- This is the focus of this work

Limitations in existing work

- Existing work considers only *sporadic tasks* with individual deadlines.
 - However real-life applications are usually expressed as *task graphs*: a set of tasks constrained by precedence constraints.
 - Such applications only have a consolidated deadline, with greater scheduling freedom.

Limitations in existing work

- Existing work considers only *sporadic tasks* with individual deadlines.
 - However real-life applications are usually expressed as *task graphs*: a set of tasks constrained by precedence constraints.
 - Such applications only have a consolidated deadline, with greater scheduling freedom.
- Existing work exclusively focuses on DVS
 - DVS is becoming less important as the headroom for voltage scaling reduces
 - DVS cannot be applied to other I/O components like radio, network bus, memory controller, etc.

Limitations in existing work

- Existing work considers only *sporadic tasks* with individual deadlines.
 - However real-life applications are usually expressed as *task graphs*: a set of tasks constrained by precedence constraints.
 - Such applications only have a consolidated deadline, with greater scheduling freedom.
- Existing work exclusively focuses on DVS
 - DVS is becoming less important as the headroom for voltage scaling reduces
 - DVS cannot be applied to other I/O components like radio, network bus, memory controller, etc.
- We address these two issues in this work

Outline





System model and problem definition







System Model

- We have a system that can exist in two modes: active and idle
- In active mode the system consumes some power and processes tasks at some rate
- In idle mode the system consumes a lower amount of power and processors no tasks

System Model

- We have a system that can exist in two modes: active and idle
- In active mode the system consumes some power and processes tasks at some rate
- In idle mode the system consumes a lower amount of power and processors no tasks
- The system cannot be put in the idle mode while running a task (requires context save and restore)
- Only control: Put the system in idle mode in between task executions

Thermal Model

- Heat dissipation is modelled using the compact thermal model
- From the compact model parameters, we can derive the following closed-loop solution to the temperature of the system

$$\mathsf{T}(\mathsf{t}) = \mathsf{T}^\infty + (\mathsf{T}(\mathsf{t}_o) - \mathsf{T}^\infty) \cdot e^{-\mathfrak{a}(\mathsf{t} - \mathsf{t}_o)}$$

where t_{o} is the starting time, T^{∞} is the steady-state temperature and α is the time-constant.

Thermal Model

- Heat dissipation is modelled using the *compact thermal model*
- From the compact model parameters, we can derive the following closed-loop solution to the temperature of the system

$$T(t) = T^{\infty} + (T(t_o) - T^{\infty}) \cdot e^{-\alpha(t-t_o)}$$

where t_o is the starting time, T^∞ is the steady-state temperature and α is the time-constant.

- The power consumption is mode-dependent and thus the parameters T^{∞} and a are also mode-dependent. We use super-scripts to denote the modes
- The thermal properties of the system is given in the tuple $\mathbf{T} = \{T^{act}, T^{idl}, a^{act}, a^{idl}\}$

Application Model

- Application is modeled as a task graph G = (V, E)
- Each $\nu \in V$ is a task with an execution time τ_ν
- Each edge $e=(\nu_1,\nu_2)$ denotes that the task ν_1 must complete before the task ν_2 starts
- It is required that all tasks of the task graph complete execution with $\omega_{m\alpha x}$ units of time
- The application is characterized by $\mathbf{A} = \{G, \tau, \omega_{max}\}$

Scheduling problem

- A stop-go schedule of a task graph is characterized by the tuple ${\bf S}=\{\sigma,t^{idl}\}$
- σ is the static-ordering of tasks: σ_i is the ith task to be executed
- t^{id1} is the amount of idle times inserted: t^{id1}_j is the idle time inserted before the execution of the jth task

Scheduling problem

- A stop-go schedule of a task graph is characterized by the tuple $\mathbf{S}=\{\sigma,t^{idl}\}$
- σ is the static-ordering of tasks: σ_i is the ith task to be executed
- t^{id1} is the amount of idle times inserted: t^{id1}_j is the idle time inserted before the execution of the jth task

Problem Definition

Given is a system with thermal parameters **T** and an application **A**. To find a stop-go schedule **S** that schedules the application to complete within its makespan and optimally minimizes the peak temperature

Putting it all together

System model: $T^{act} = 395$ K, $T^{idl} = 325$ K, $a^{act} = a^{idl} = 6.667 s^{-1}$ Application model: Task graph with 5 tasks, $\tau_* = 50$ ms, $\omega = 415$ ms. Schedule: $\sigma = (CABDE)$, $t^{idl} = (0, 0, 43, 61, 61)$ ms



Outline











Step 1 : Given static-order of tasks

- Let the static-ordering σ be given
- We are to only compute optimal values of t^{idl}
- Let T_i denote the temperature of the system at the *finish* of the *j*th task

Step 1 : Given static-order of tasks

- Let the static-ordering σ be given
- We are to only compute optimal values of t^{idl}
- Let T_j denote the temperature of the system at the *finish* of the jth task
- We have the following result:

Theorem 1

Suppose an arbitrary non-preemptive stop-go schedule with given static-ordering of tasks, σ . Then, any change of the idle times such that none of T_j , $j \in \{1, ..., |V|\}$ decreases, and at least one increases, decreases the makespan.

Additional slide: Proof sketch

Interpretation of Theorem 1



Interpretation of Theorem 1





Interpretation of Theorem 1



Insight

If the maximum temperature that we reach is $T^{m\alpha x}$, we must not remain in idle longer than is required to end up with $T_j = T^{m\alpha x}$

JUst Sufficient Throttling (JUST) schedule

Informal definition

- $\mathtt{JUST}(\sigma)$ is a stop-go schedule where
- the first p^{opt} tasks are executed with no idle time before them,
- the remaining tasks are executed with an idle time before them such that $T_{j}=T^{\rm opt},$
- the makespan $\omega = \omega_{\max}$

JUst Sufficient Throttling (JUST) schedule

Informal definition

- $JUST(\sigma)$ is a stop-go schedule where
- the first p^{opt} tasks are executed with no idle time before them,
- the remaining tasks are executed with an idle time before them such that $T_{\rm j}=T^{\rm opt},$
- the makespan $\omega = \omega_{\max}$



Are p^{opt}, T^{opt} unique? How to compute them?

Yes, they are unique. Computation described in Algorithm 1 in paper.

Are p^{opt}, T^{opt} unique? How to compute them?

Yes, they are unique. Computation described in Algorithm 1 in paper.

Brief Idea

• Let T_j^{cont} be the temperature of the system when executing the first j tasks without any idle time

Are p^{opt}, T^{opt} unique? How to compute them?

Yes, they are unique. Computation described in Algorithm 1 in paper.

- Let T_j^{cont} be the temperature of the system when executing the first j tasks without any idle time
- We must have $T^{\rm cont}_{{\tt p}^{\rm opt}} \leqslant T^{\rm opt}$

Computation of p^{opt} and T^{opt}

Are p^{opt}, T^{opt} unique? How to compute them?

Yes, they are unique. Computation described in Algorithm 1 in paper.

- Let T_j^{cont} be the temperature of the system when executing the first j tasks without any idle time
- We must have $T^{\rm cont}_{p^{\rm opt}} \leqslant T^{\rm opt}$
- We must have $T_{p^{opt}+1}^{cont} > T^{opt}$

Are p^{opt}, T^{opt} unique? How to compute them?

Yes, they are unique. Computation described in Algorithm 1 in paper.

- Let T_j^{cont} be the temperature of the system when executing the first j tasks without any idle time
- We must have $T^{\rm cont}_{p^{\rm opt}} \leqslant T^{\rm opt}$
- We must have $T_{p^{opt}+1}^{cont} > T^{opt}$
- We must have have $\omega = \omega_{max}$
- Binary search over p^{opt}

Optimality of JUST schedule

Theorem 2

 $JUST(\sigma)$ optimally minimizes the peak temperature amongst all stop-go schedules following a given static-order σ and satisfying the makespan constraint.

Optimality of JUST schedule

Theorem 2

 $JUST(\sigma)$ optimally minimizes the peak temperature amongst all stop-go schedules following a given static-order σ and satisfying the makespan constraint.

Proof idea

Take any schedule S with static-order σ .

Repeatedly apply Theorem 1 to modify S to become more and more similar to the JUST schedule.

Step 2: Integrated choice of σ and t^{idl}

Skip this section

- Let σ^{opt} be the optimal static-order with a corresponding JUST(σ^{opt}) schedule with p^{opt} and T^{opt}
- Observe that
 - Changing the order of execution of the first p^{opt} tasks does not change the makespan or peak temperature
 - Changing the order of execution of the last (|V| p^{opt}) tasks does not change the makespan or peak temperature
- To characterize σ^{opt} , it is sufficient to specify tasks is one among the first p^{opt} or not
- Hence, because of the special property of a JUST schedule, we have only a *binary decision problem*!

Algorithm to compute σ^{opt}

- $\bullet~$ We start by assuming a target value of the peak temperature $T^t,$ i.e. $T^{opt}\leqslant T^t$
- For a given T^t, we can compute σ^{opt} by solving a binary integer program (BIP) with binary variable x_{ν} denoting whether task ν belongs to the first p^{opt} tasks or not
- Precedence constraints, makespan constraints and peak temperature being below the target temperature can be modelled in the BIP
- If a feasible solution exists, we reduce T^t and again solve another instance of BIP

Additional slide: BIP formulation

Periodically repeating applications

- We have considered a single run of a task graph
- But in practice applications run periodically, with period equal to maximum makespan (ω_{max})

Periodically repeating applications

- We have considered a single run of a task graph
- But in practice applications run periodically, with period equal to maximum makespan ($\omega_{max})$
- What is different? The finishing temperature of the nth iteration is the finish temperature of the (n + 1)th iteration
- We make a simple extension: In each iteration, compute the optimal JUST schedule
- Let $(T^{opt})_n$ denote the maximum temperature in the nth iteration

Periodically repeating applications

- We have considered a single run of a task graph
- But in practice applications run periodically, with period equal to maximum makespan (ω_{max})
- What is different? The finishing temperature of the nth iteration is the finish temperature of the (n + 1)th iteration
- We make a simple extension: In each iteration, compute the optimal JUST schedule
- Let $(T^{opt})_n$ denote the maximum temperature in the nth iteration
- Questions: What is the maximum temperature of the system if run iteratively? Does **JUST** schedule optimally reduce it?

Results for periodic applications

Theorem 3

The series $\{(T^{opt})_n\}$ converges for $n \to \infty$ and the limit does not depend on the static-ordering of tasks for JUST schedules.

Corollary 4

With JUST schedules the highest temperature of a system running a periodic task graph is minimal and independent of the ordering of tasks.

Inferences from results

- If we want to run a task graph periodically, using JUST schedule in each period, optimally minimizes the peak temperature
- Further, the choice of σ is irrelevant to minimize the peak temperature
- Thus, we can choose σ w.r.t. other considerations, such as buffer capacity, and use the JUST schedule for this σ to optimally minimize peak temperature

Outline











System Properties

 $\label{eq:area} \begin{array}{l} \text{ARM-like core} \\ \text{T}^{act} = 394\text{K}, \text{T}^{idl} = 325 \\ a^{act} = a^{idl} = 6.667 \text{s}^{-1} \\ \text{T}_{o} = 330\text{K} \end{array}$

Realistic example - MP3 application

H	им		Execution
		Task ID	Time
$\left(RQ0 \right)$	$\left(RQ1 \right)$		$(\times 10^{-7} s)$
Ť	Ť	HM	236070
(ROO)	RO1	RQ0	139325
\triangleleft	\sum	RQ1	139325
\succ	\prec	RO0	69385
S	TR	RO1	69385
\sim	-	STR	73618
(AR0)	(AR1)	AR0	13088
\mathbf{Y}	\bigvee	AR1	13088
		IM0	711744
	(IMI)	IM1	711744
\downarrow	\downarrow	FI0	157184
(FI0)	(FI1)	FI1	157184
\checkmark	\bigvee	SY0	1866138
(SY0)	(SY1)	SY1	1866138

MP3 application - comparison with other schedules



Figure: Temperature trace for (a) JUST schedule, (b) workload conserving execution, and (c) equally distributed idletimes

Synthetic example

(1) (2)	Task ID	Execution Time (in msec)
Å Å	1	30
$\bigcirc 4$	2	140
	3	20
Ġ Ġ	4	50
Y Y	5	50
	6	50
(7)	7	50

Synthetic example - Results



$$\begin{split} & \text{Figure: } T^t = 375 \text{K}, \sigma = (1, 3, 5, 2, 4, 6, 7) \\ & \{T^{\text{opt}}\} = (374 \text{K}, 377 \text{K}, 378 \text{K}, 378.3 \text{K}, \ldots) \end{split}$$

Synthetic example - Results



 $\{T^{opt}\} = (374K, 377K, 378K, 378.3K, ...)$

Figure: $T^{t} = 370 \text{ K}, \sigma = (2, 1, 3, 5, 4, 6, 7)$ $\{T^{opt}\} = (369K, 377K, 378K, 378.3K, ...)$

Outline











Conclusions and future work

- Stop-go execution is an effective way to manage on-chip temperature while executing real-time applications
- We proved optimality of the proposed JUST schedule for a given static-ordering
- We gave an algorithm to optimize the static-ordering as well
- For periodically repeating tasks, peak temperature is optimally minimized by JUST schedule and is independent of the static-ordering

Conclusions and future work

- Stop-go execution is an effective way to manage on-chip temperature while executing real-time applications
- We proved optimality of the proposed JUST schedule for a given static-ordering
- We gave an algorithm to optimize the static-ordering as well
- For periodically repeating tasks, peak temperature is optimally minimized by JUST schedule and is independent of the static-ordering
- Future work:
 - Extension to multi-processor distributed systems
 - Combination of DVS and stop-go scheduling

Thank you for your attention

Any questions?

Sketch of proof of theorem 1

Define F as

$$F = \prod_{j=1}^{|V|} \exp^{-\alpha^{idl} t_j^{idl}}$$

We can reduce the above to

$$F = \mathbf{c} \cdot \left(\mathsf{T}_{|\mathsf{V}|} - \mathsf{T}_{|\mathsf{V}|}' \right) \cdot \prod_{j=1}^{|\mathsf{V}|-1} \frac{\mathsf{T}_j - \mathsf{T}_j'}{\mathsf{T}_j - \mathsf{T}^{\mathsf{idl}}}$$

From which we can derive

$$\frac{\partial F}{\partial T_{|V|}} > 0 \tag{2}$$

Solution Also, we can show that $\frac{\partial F}{\partial \omega} < o$. Hence the theorem

(1)

Binary Integer Program for a given target temperature

Maximise1subject to
$$x_{\nu} \in \{0,1\}$$
 $\forall \nu \in V$ $x_{\nu 1} \leqslant x_{\nu 2}$ $\forall (\nu 1, \nu 2) \in E$ $A1^T x \geqslant b1$ $A2^T x \leqslant b2$

where
$$A_{1\nu} = \frac{T^t - T_{\nu}''}{f^{act}(t_j^{act})(T^t - T^{idl})}$$
, $b_1 = f^{idl}(\omega - \tau_{tot})$, $A_{2\nu} = \tau(\nu)$,
 $b_2 = \tau^{tot} - (f^{act})^{-1} \left(\frac{T^{act} - T^t}{T^{act} - T_o}\right)$. Back