Thermal-Aware Schedulability Analysis for Fixed-Priority Non-Preemptive Real-Time Systems

Javier Pérez Rodríguez

Patrick Meumeu Yomsi

CISTER – Research Centre in Real-Time & Embedded Computing Systems

Motivation

Safety critical systems

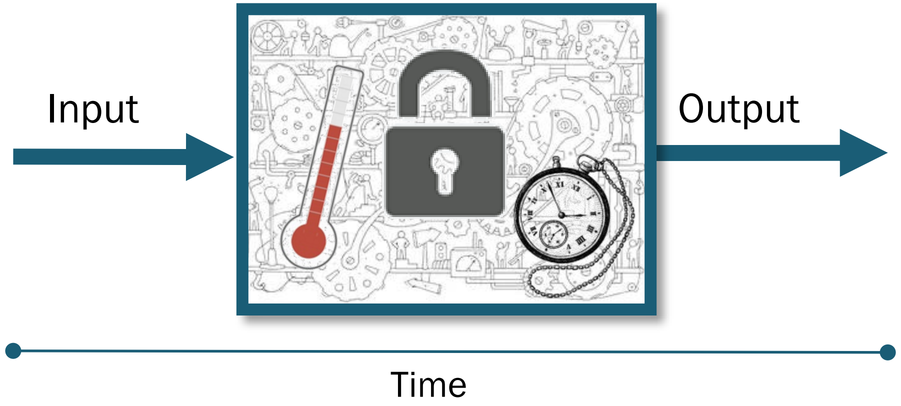

System Requirements

Functional

- Correctness
- Authentication
- Authorization levels

Non-functional

- Performance
- Security
- Temperature

Single framework

Research context

Safety critical systems 

- Lack of formal policy and/or guidance by Federal Aviation Administration and European Union Aviation Safety Agency.
- Inherent non-deterministic behavior
  - Safety
  - Performance
  - Integrity
- Detailed technical information are available only under non-disclosure agreements.

- Multi-core platform
- Single-core platform

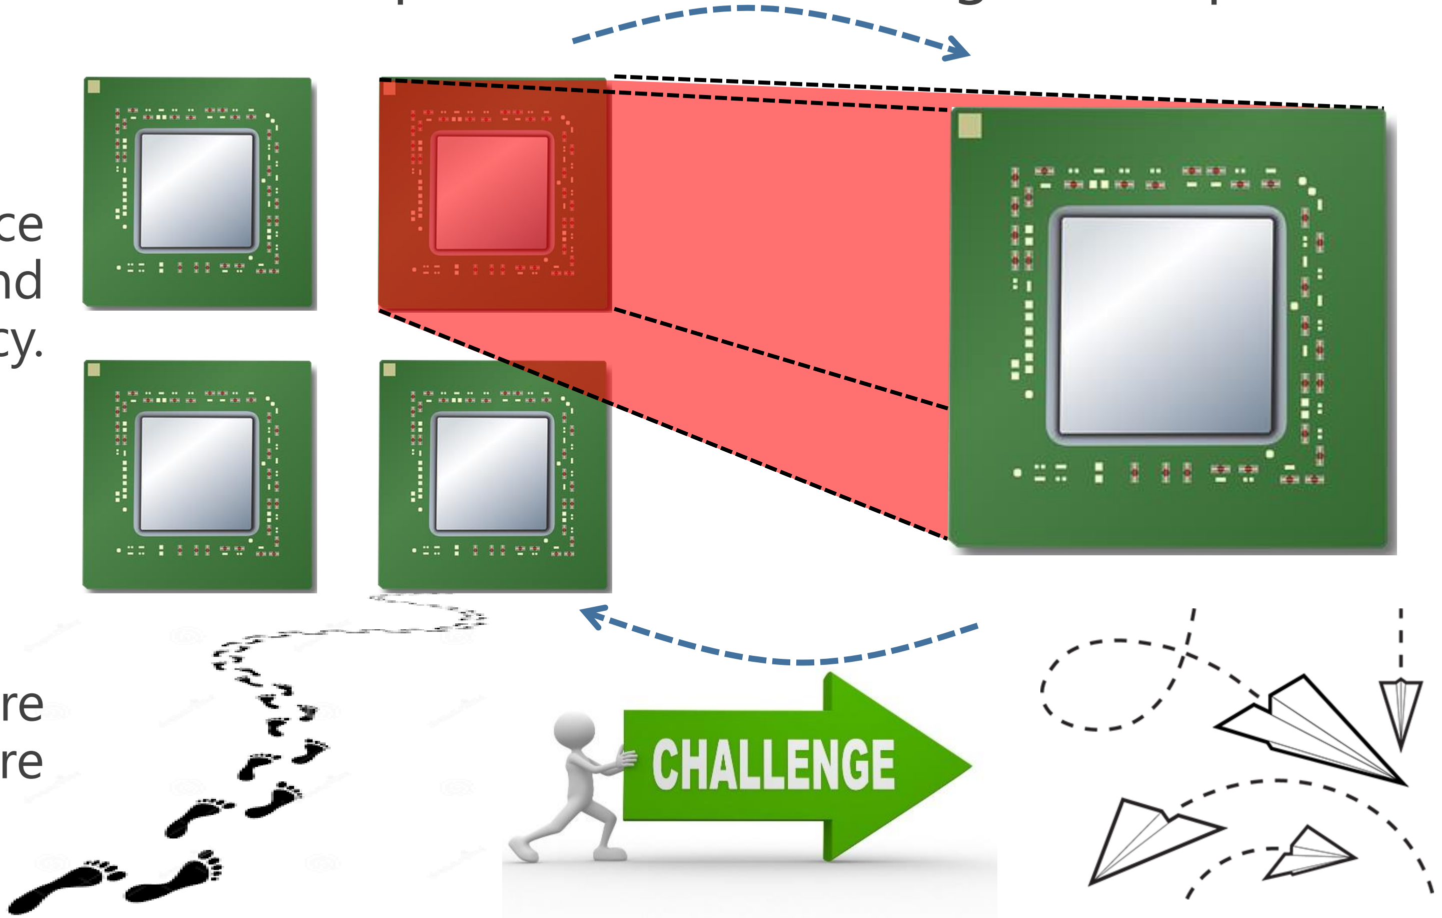

Existing strategies

Dynamic Voltage/Frequencie Scaling (DVFS)

- Discrete and limited number of frequencies
- Longer cooling periods
- Not available on all devices

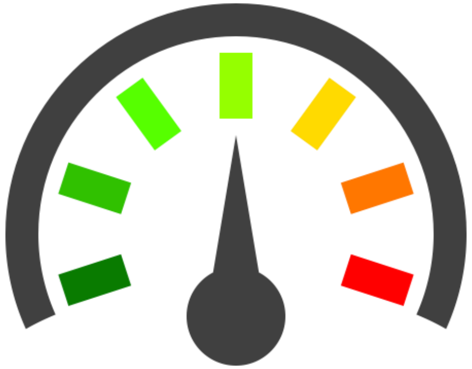

Fetch Toggling (*global clock gating*)

- Smaller cooling periods
- Available on all devices

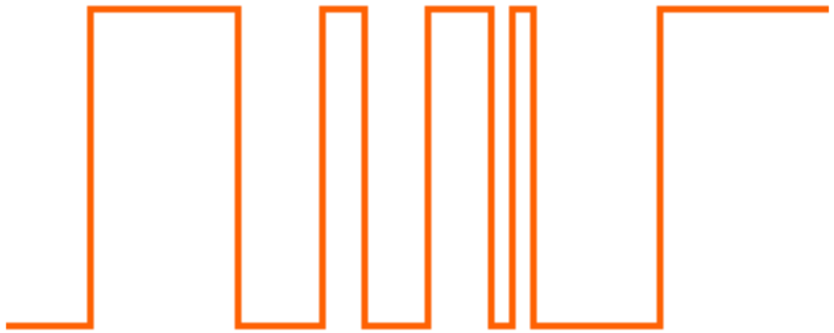

Task model

- Independent and periodic tasks: $\tau_i = (O_i, C_i, D_i, T_i)$
- Constrained deadline: $D_i \leq T_i, \forall i$
- Non-preemptive
- Same thermal profile

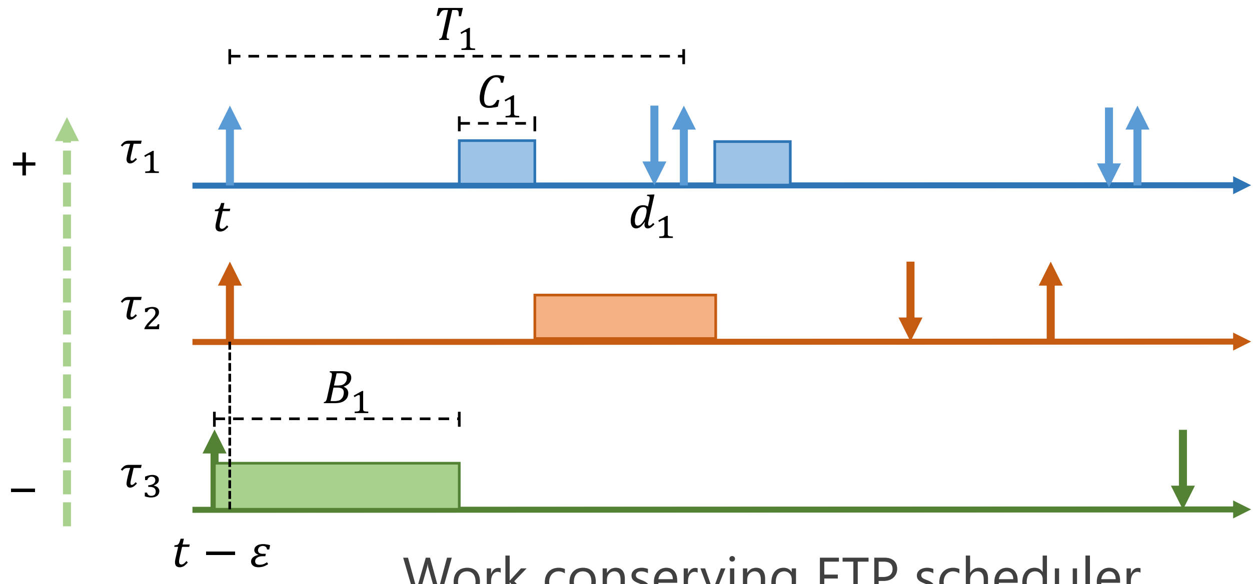

Work conserving FTP scheduler

- ⊖ Introduce additional blocking time in higher priority tasks.
- ⊕ Naturally guarantees mutual exclusion and exclusive access to shared resources.
- ⊕ Cheap to implement.
- ⊕ Limited run-time overhead.

Thermal model

- HotSpot [1]

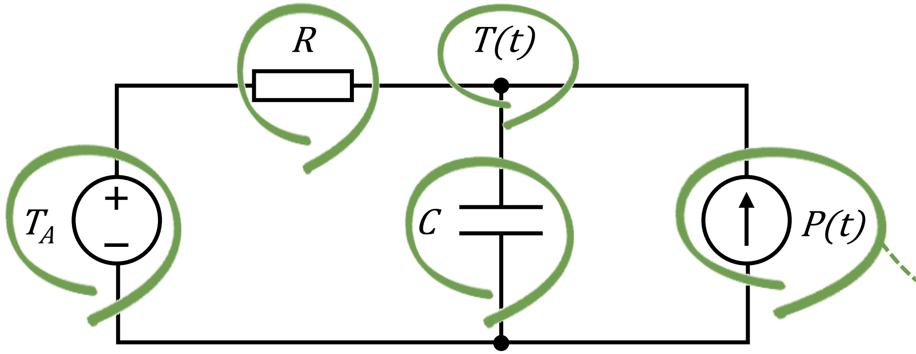

- RC serial circuit
- Fine-grained
- Block-level

Constant ambient temperature

$$\frac{dT(t)}{dt} = \frac{P(t)}{C} - \frac{T(t) - T_A}{R \cdot C}$$

$$P(t) = P_D(t) + P_L(t)$$

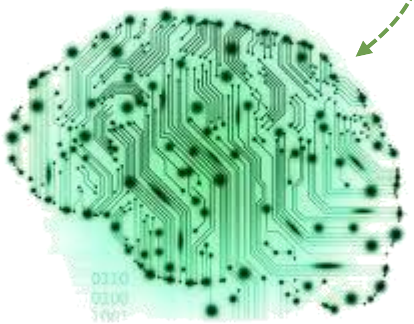

[1] K. Skadron, M. R. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan, "Temperature-aware microarchitecture: Extended discussion and results," in *In Proceedings of the 30th Annual Int. Symposium on Computer Architecture*, 2003, pp. 2–13.

Thermal behavior

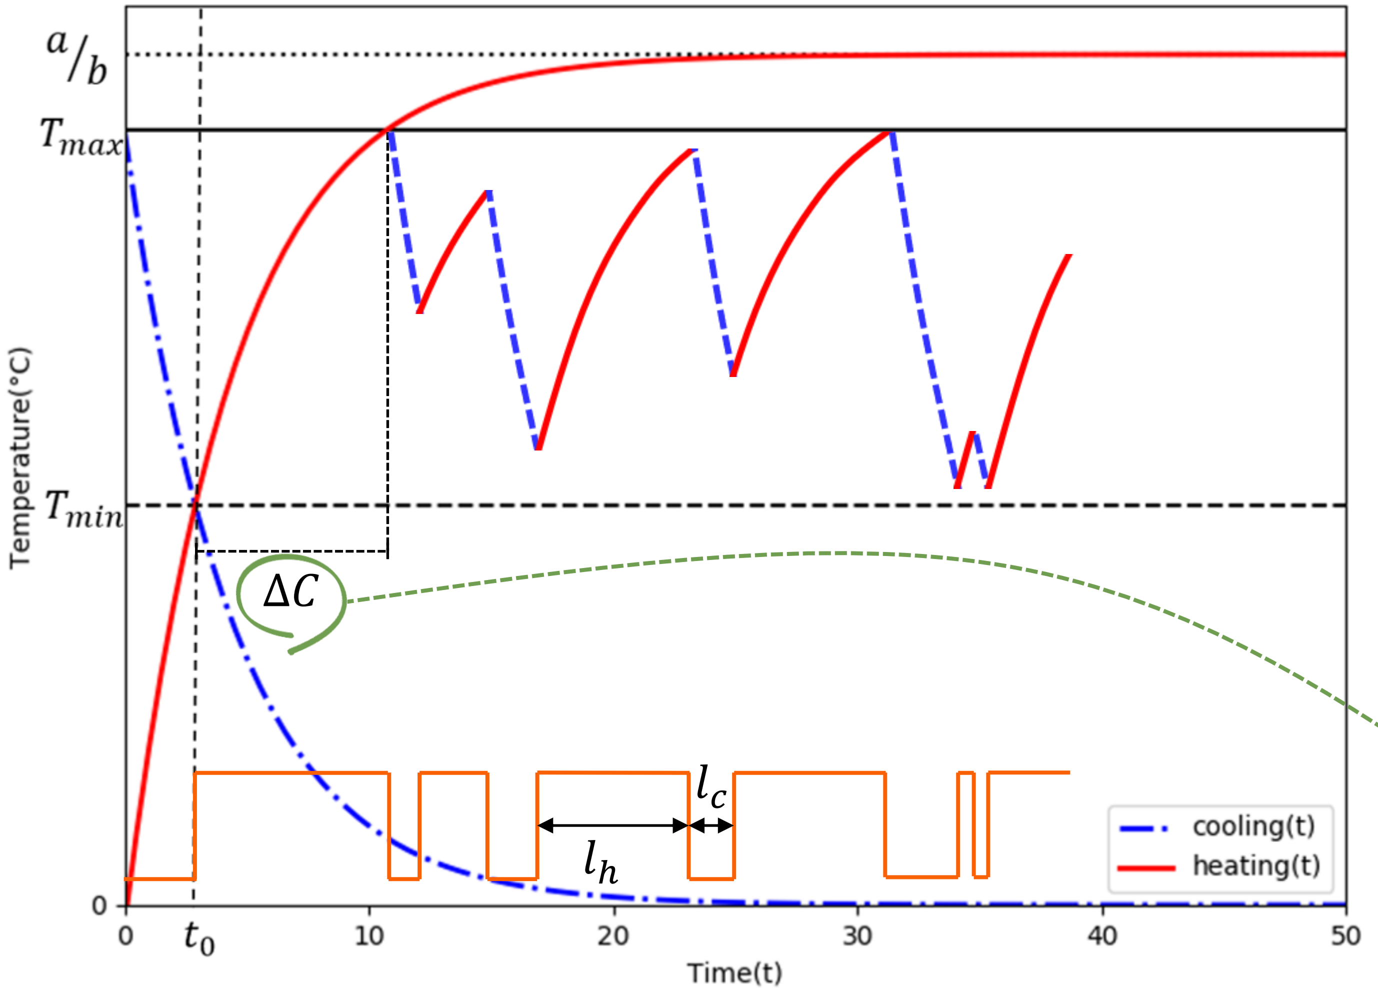

$$\frac{dT(t)}{dt} = \frac{P(t)}{C} - \frac{T(t) - T_A}{R \cdot C}$$

$$T_h(t) = \frac{a}{b} + \left(T(t_0) - \frac{a}{b}\right) \cdot e^{-b\cdot(t-t_0)}$$

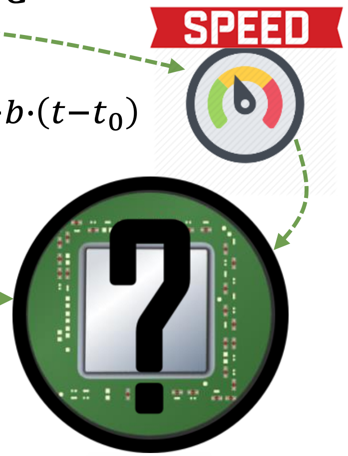

$$T_c(t) = T(t_0) \cdot e^{-b\cdot(t-t_0)}$$

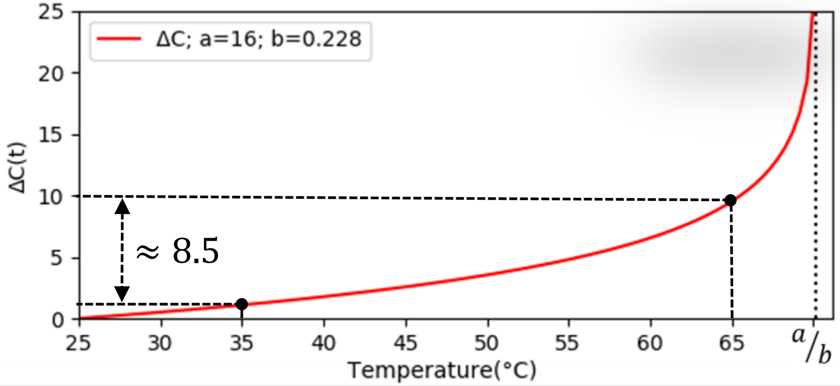

Naive solution for cooling phases

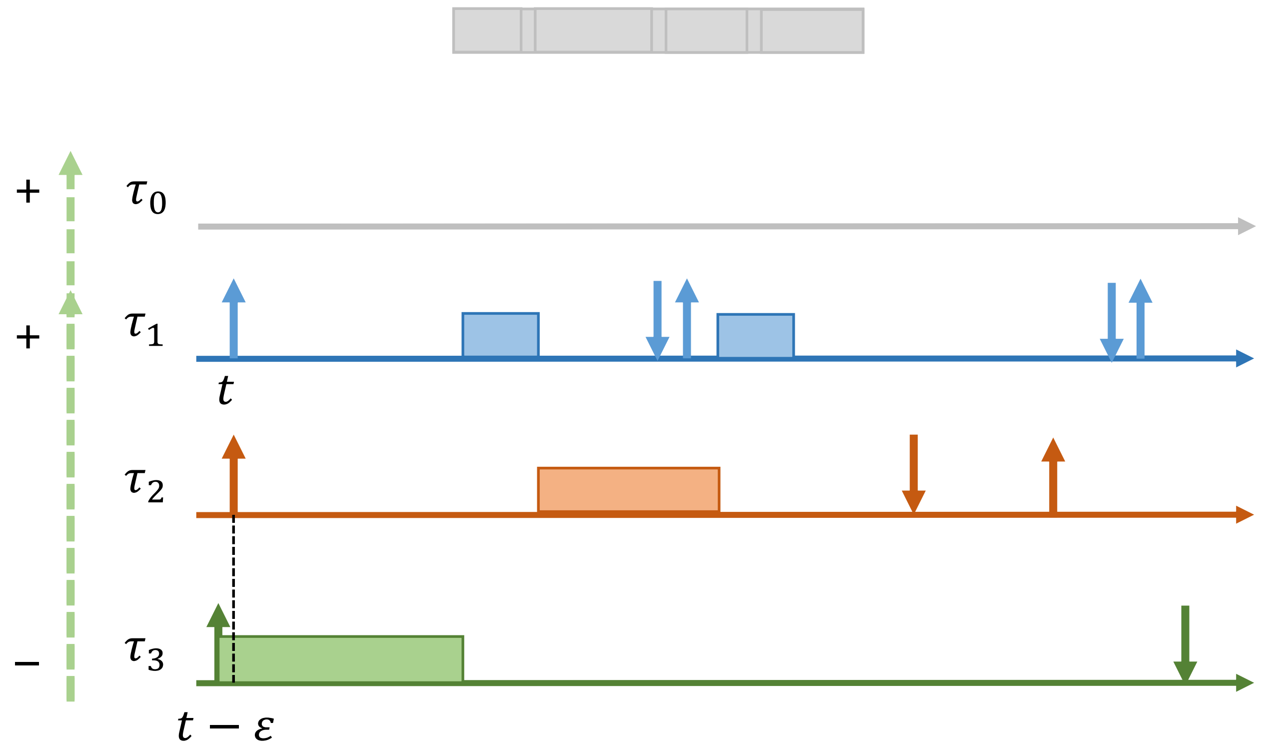

Naive solution for cooling phases

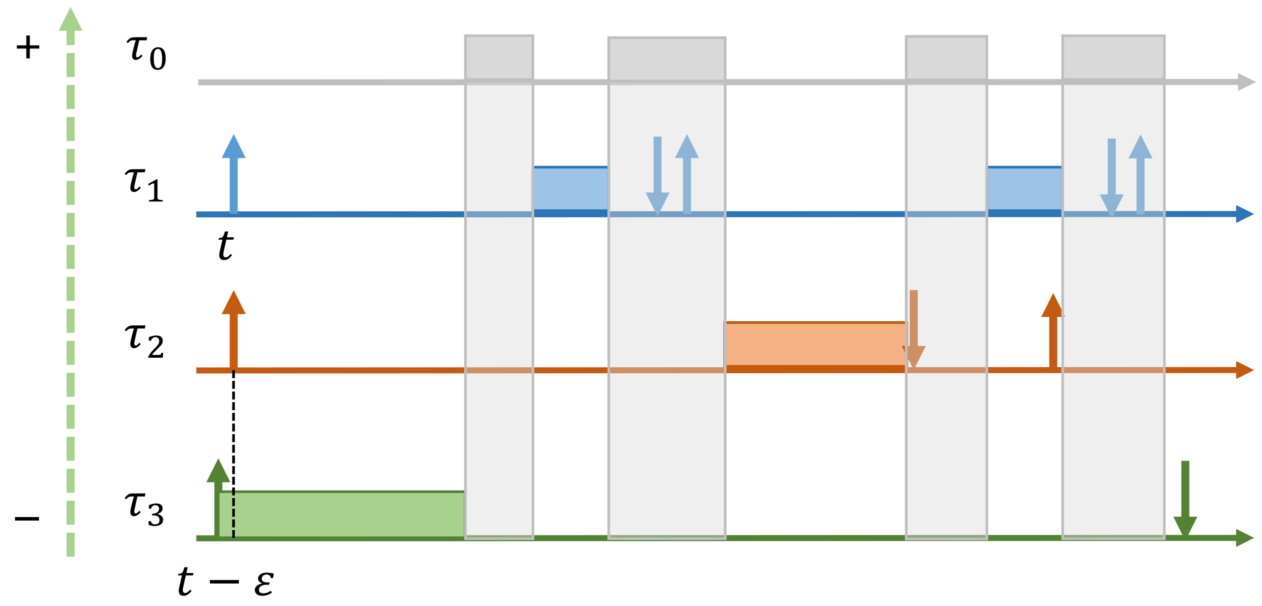

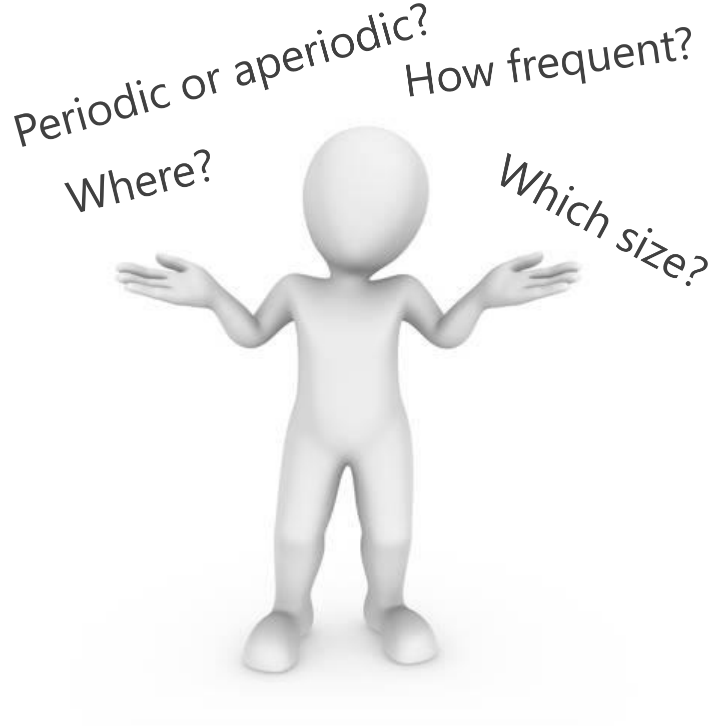

Our contributions

- Reactive scheduler: NP-HBC
  - Target low average temperatures

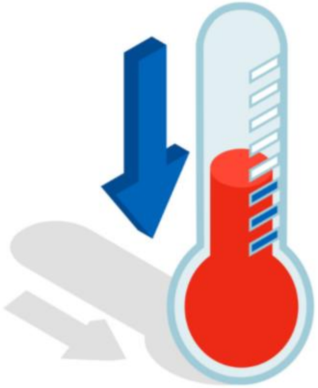

- Proactive scheduler: NP-CBH
  - Target high performances

Cooling property

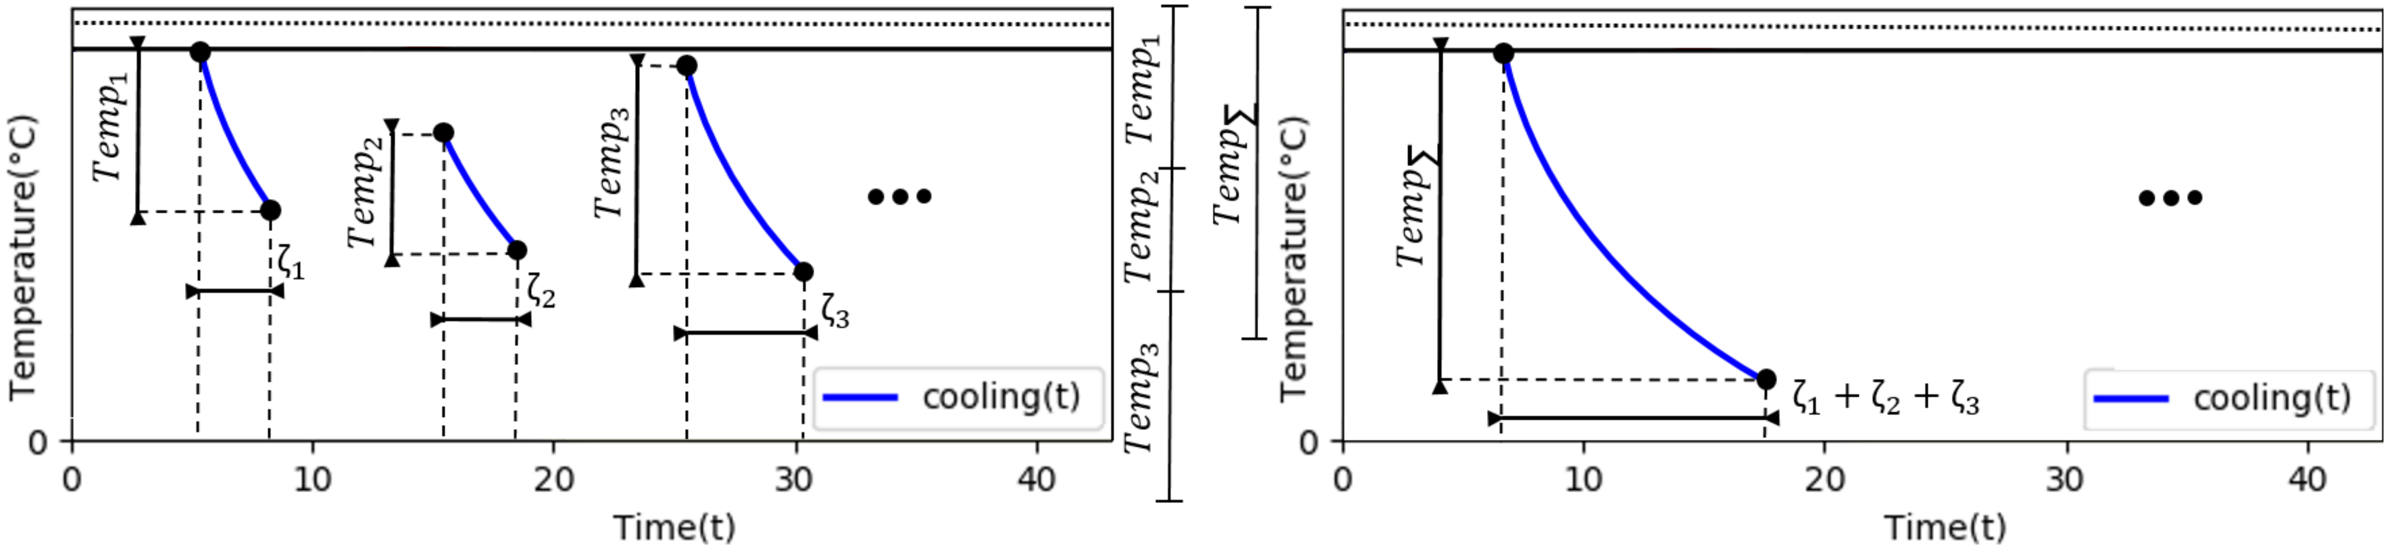

$$\sum_{j=1}^{k} T(\zeta_j) \geq T_c\left(\sum_{j=1}^{k} \zeta_j\right)$$

Average temperature reduction

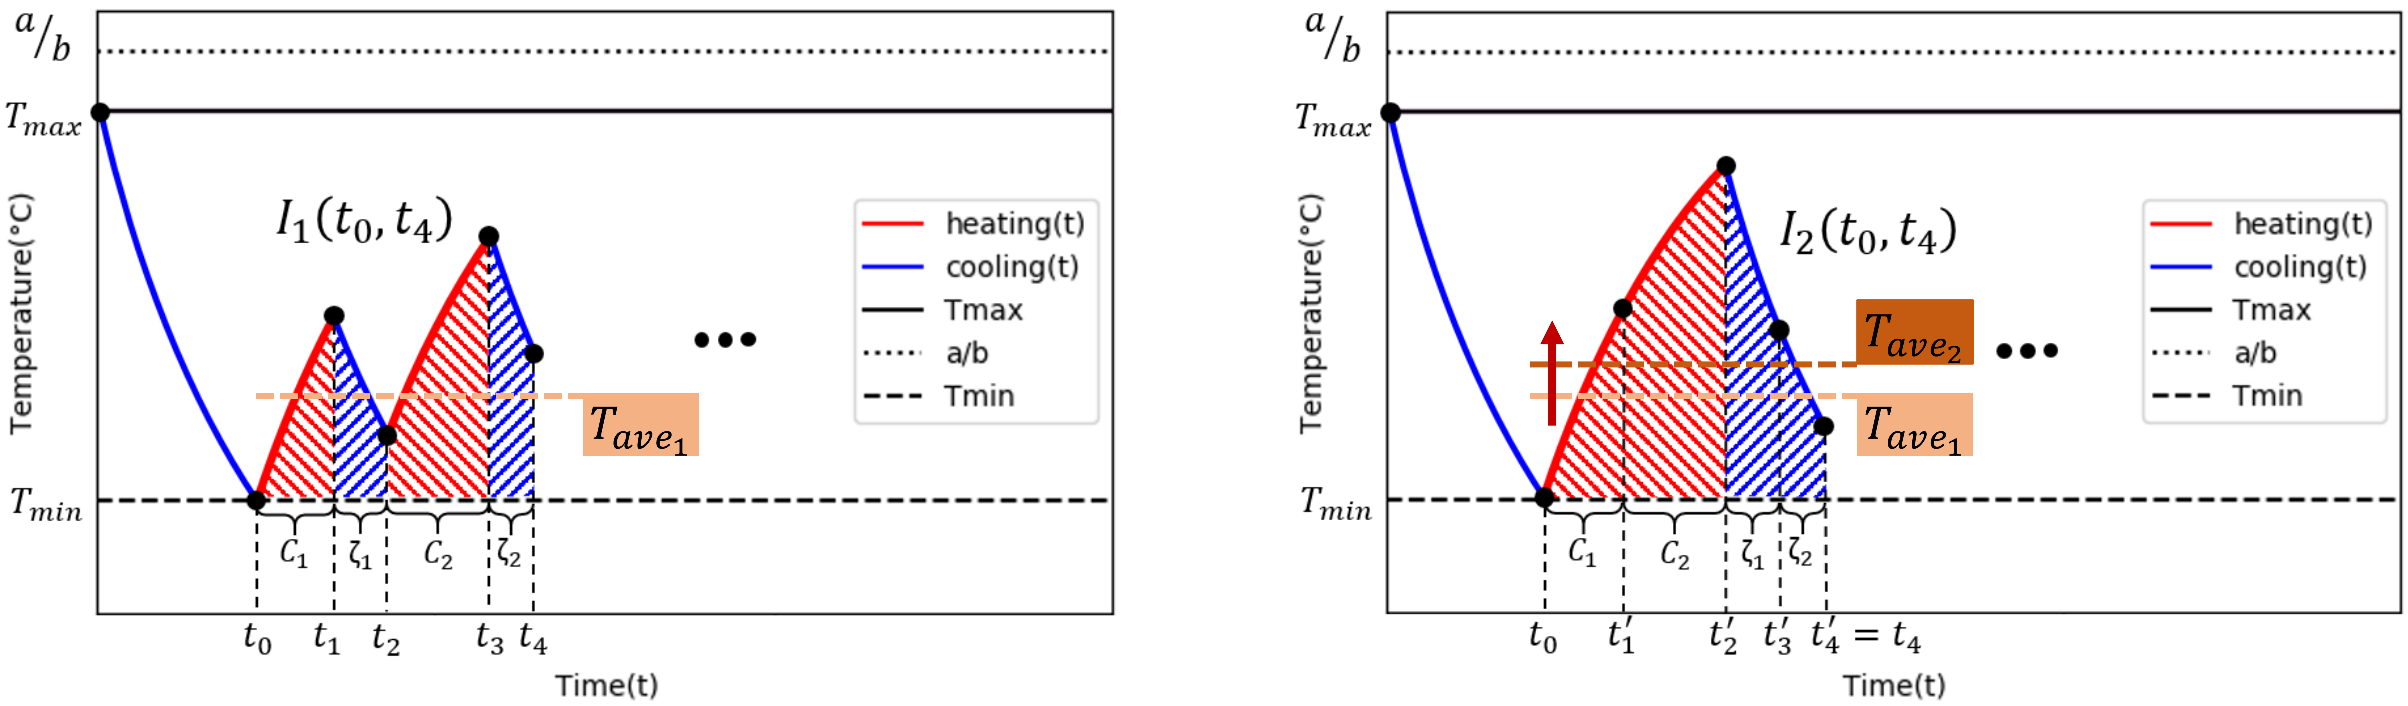

Reactive scheduler: NP-HBC (1/2)

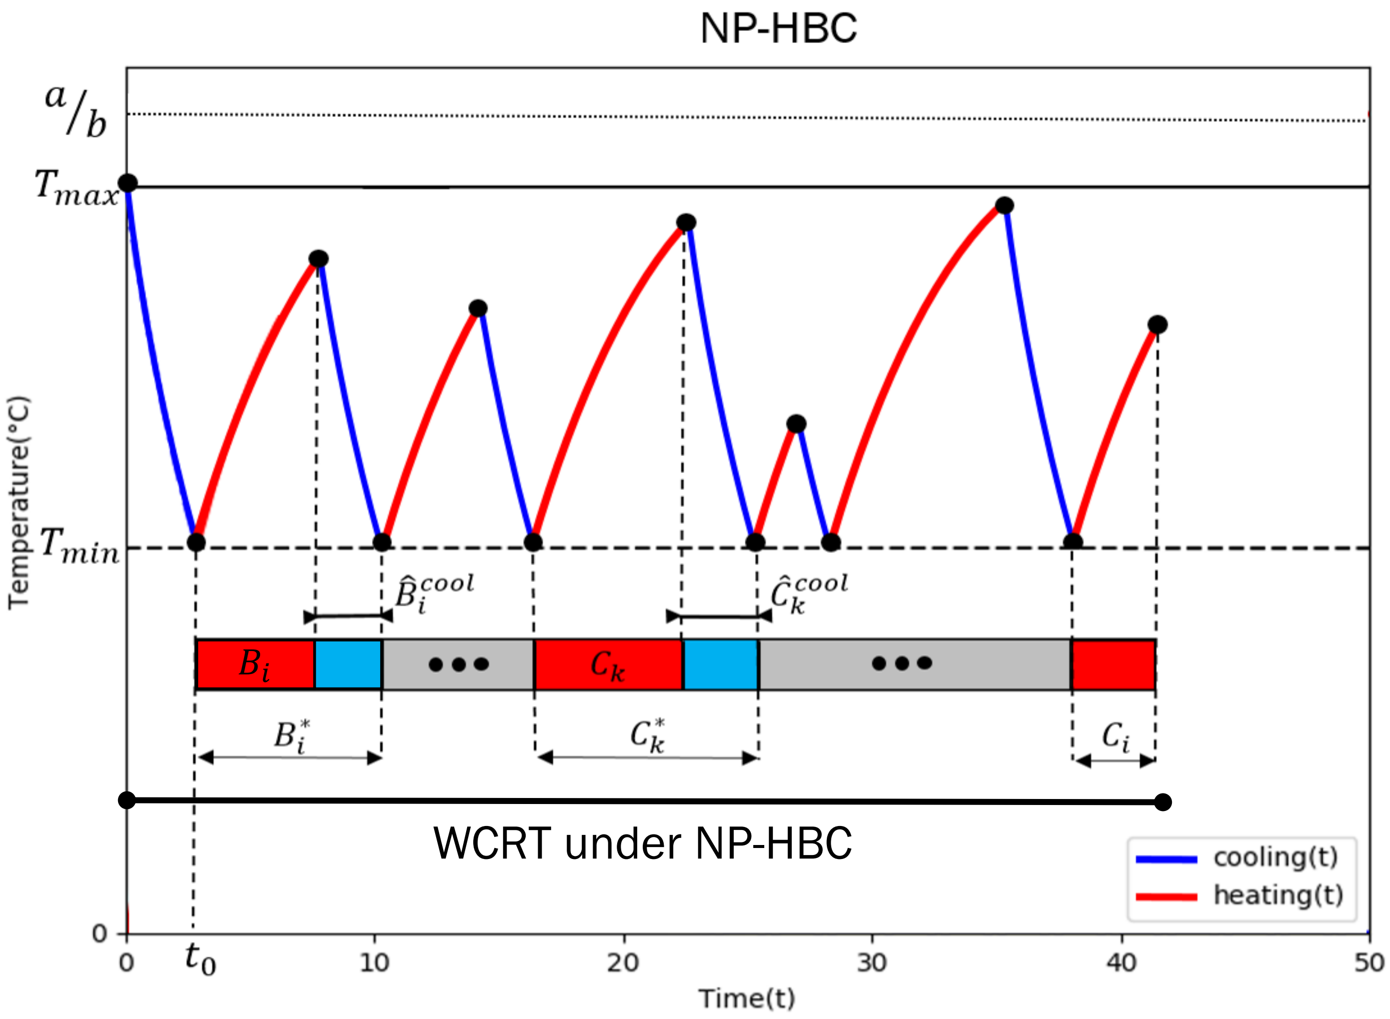

Reactive scheduler: NP-HBC (1/2)

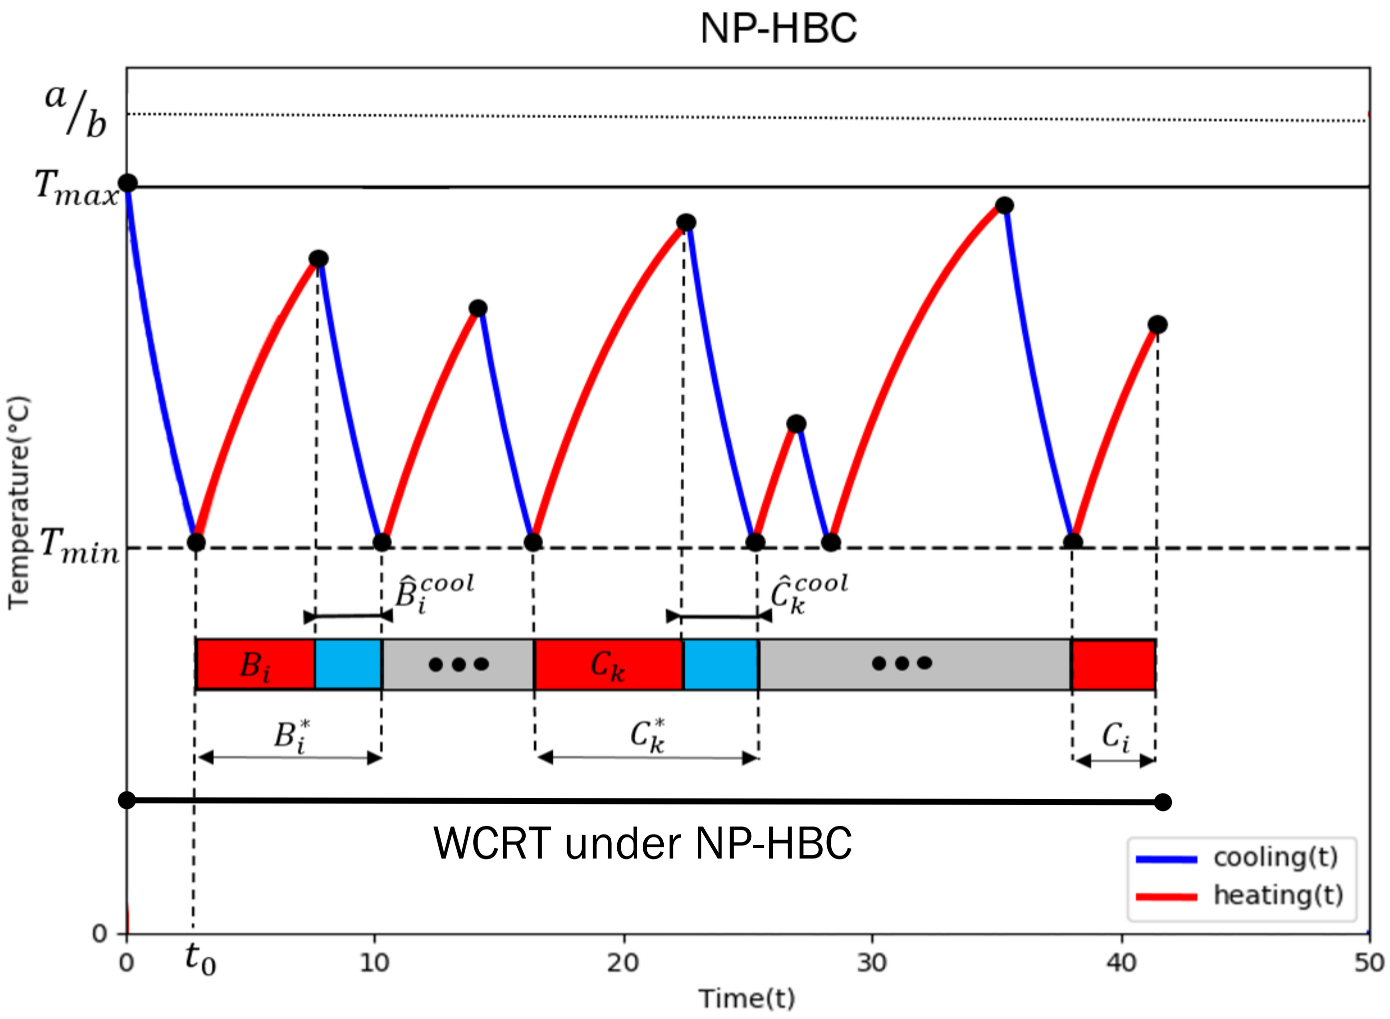

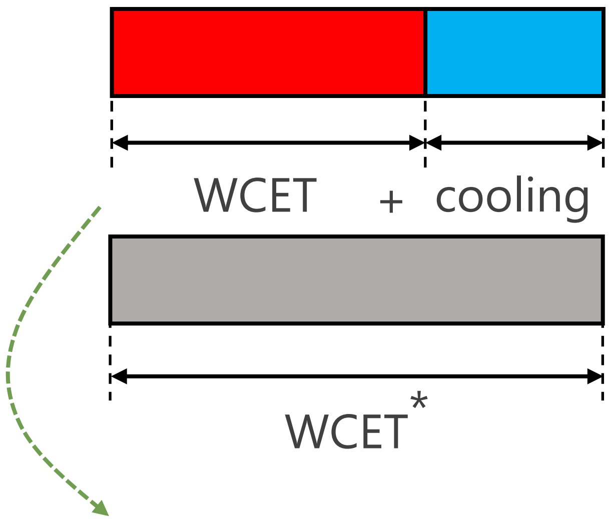

Extends the classical scheduling theory

Reactive scheduler: NP-HBC (2/2)

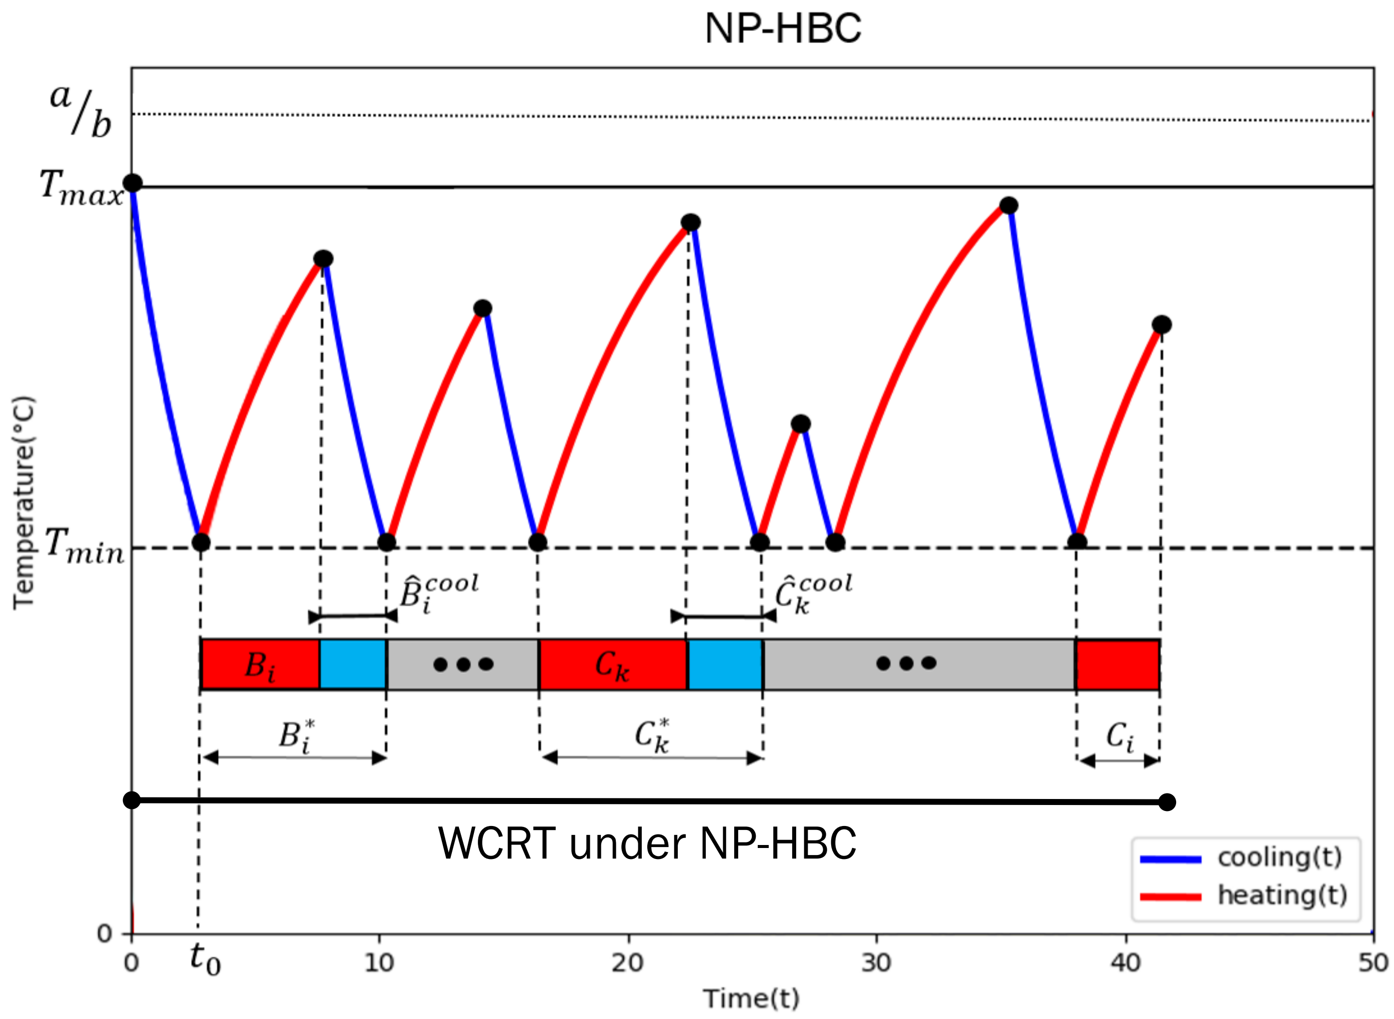

- Intuitive
- Easy to implement
- Closed-form equation
- Agnostic to the priority assigment scheme
- Pessimism in the cooling time
- Pessimism in the response time

Proactive scheduler: NP-CBH (1/3)

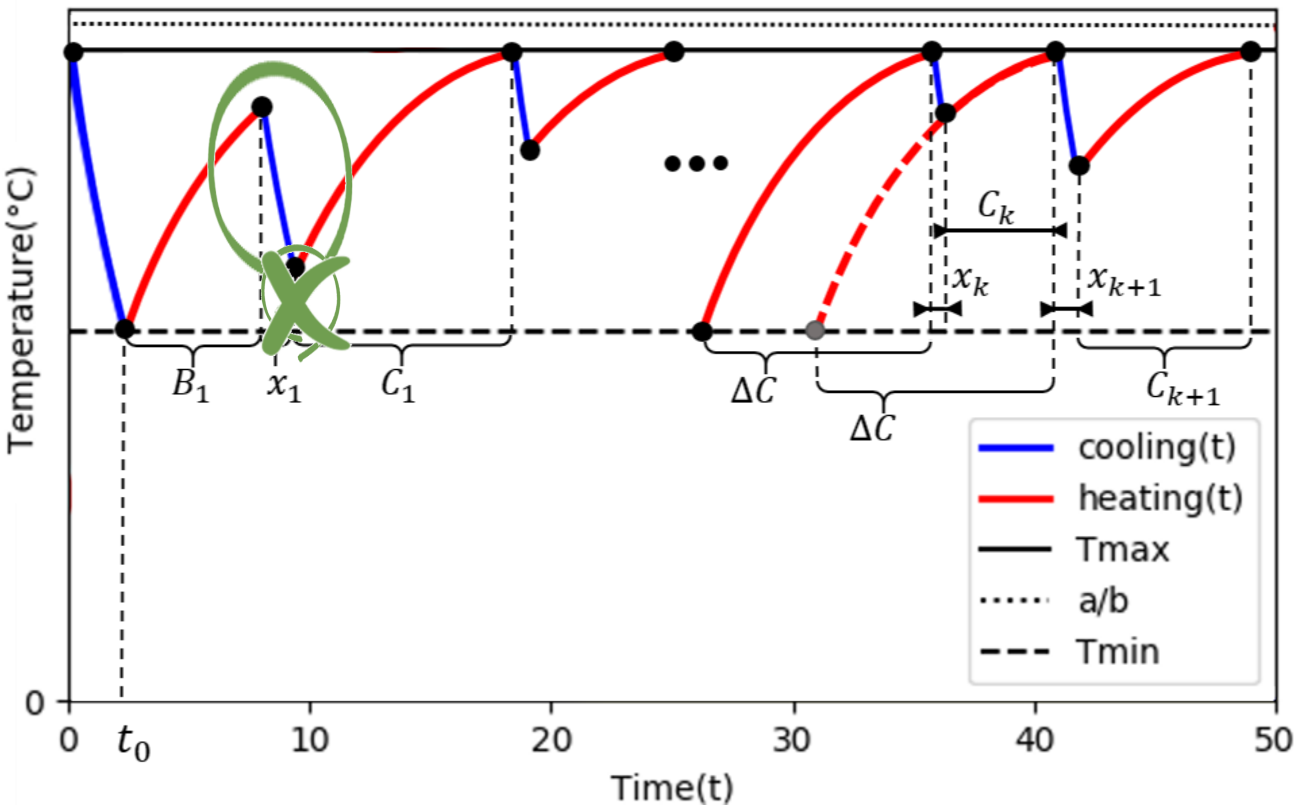

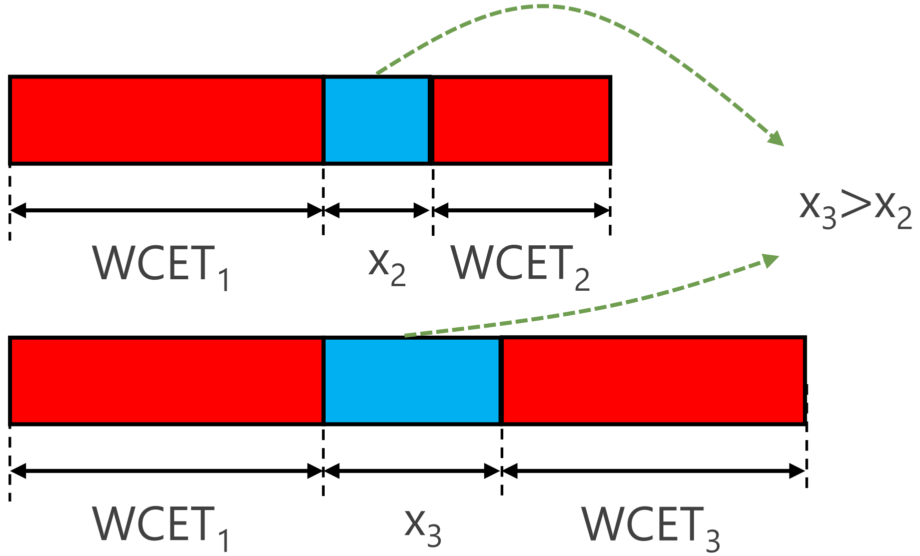

τ_1: $WCET_1$

τ_2: $WCET_2$

τ_3: $WCET_3$

$WCET_3 > WCET_2$

Proactive scheduler: NP-CBH (2/3)

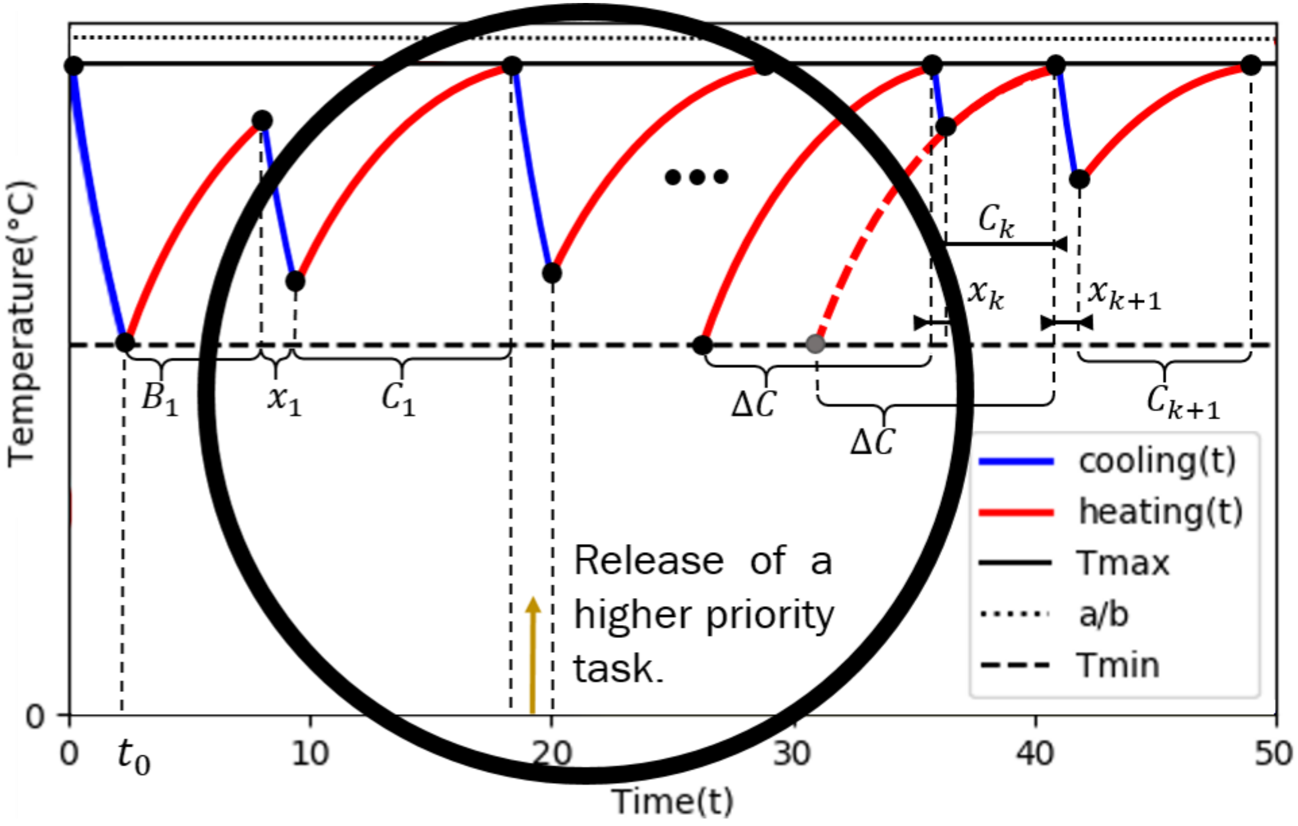

Proactive scheduler: NP-CBH (3/3)

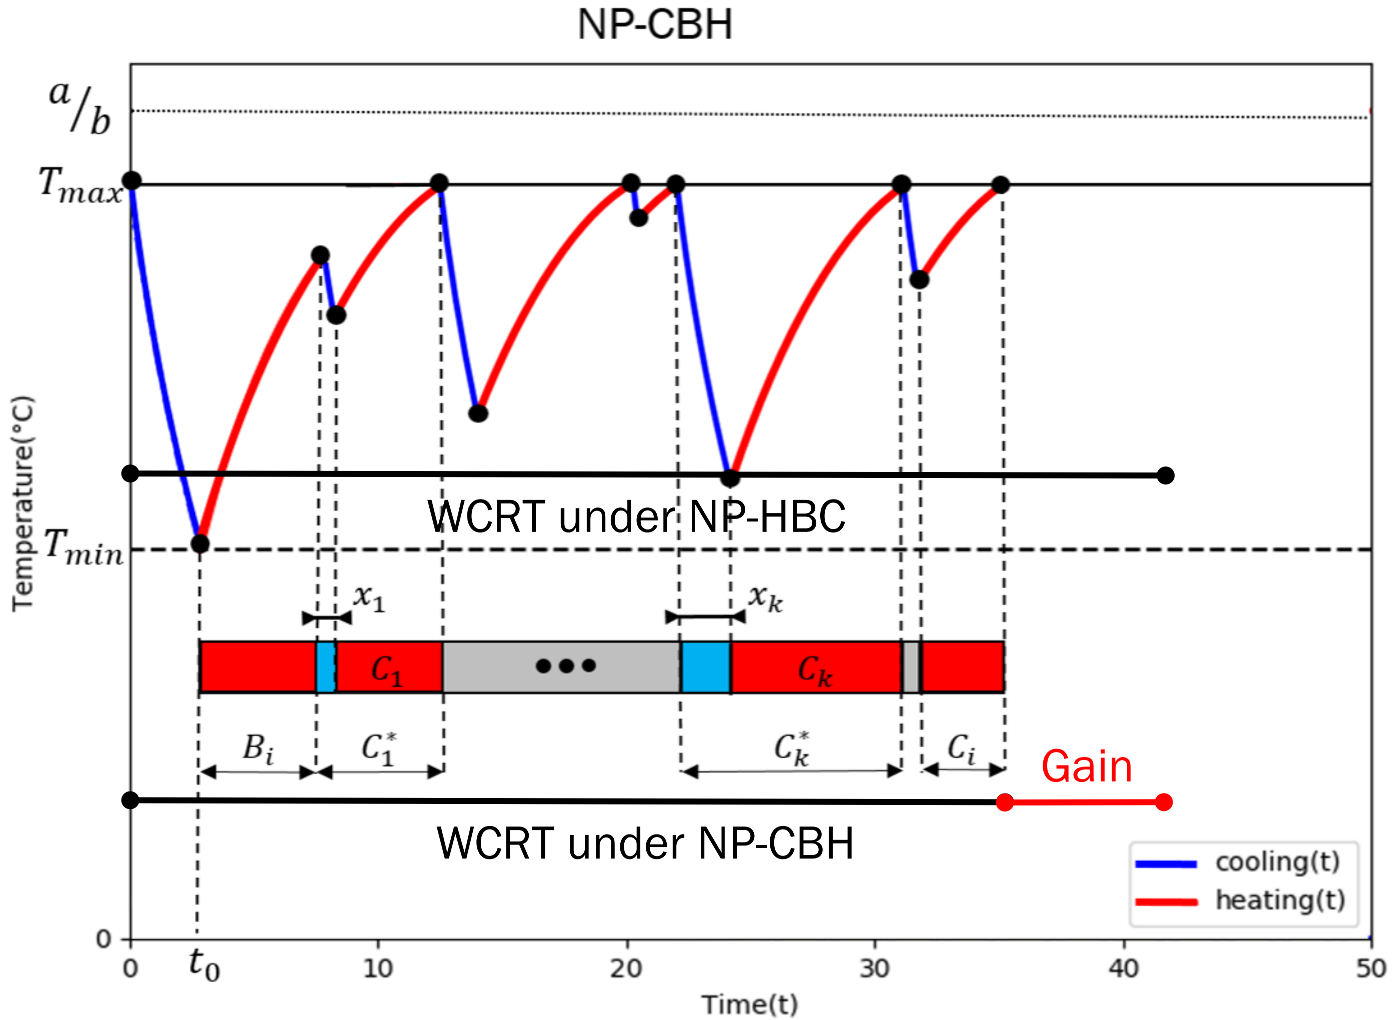

- More accurate cooling time
- More accurate response time
- Complex to implement
- Task priority matters

Results (1/2)

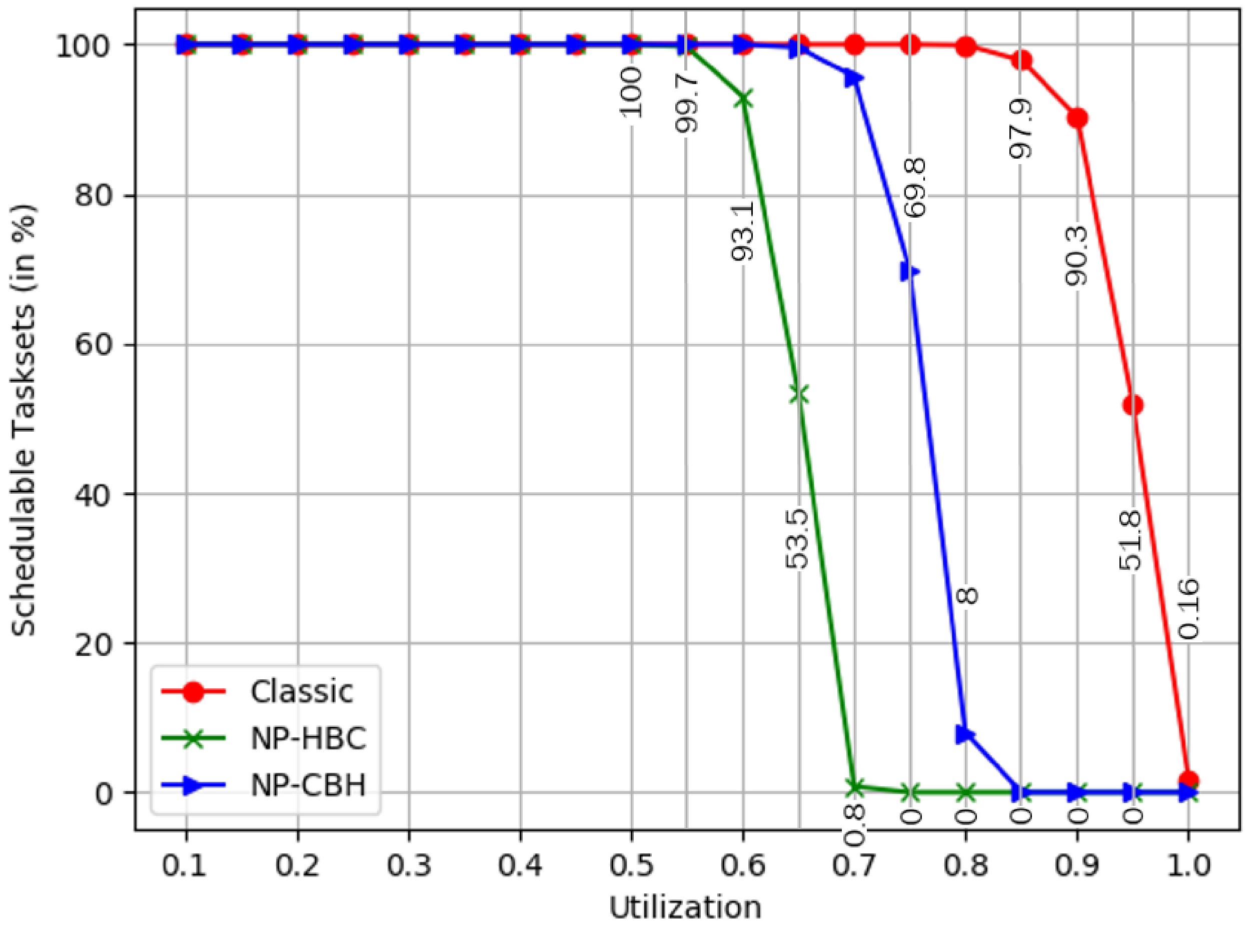

- 1000 task set per utilization:
  - Implicit deadline
  - Priority assignment: Rate Monotonic
- $a = 16$; $b = 0.228$ (silicon chip [1])
- $T_{max} = 65°C$, $T_{min} = 30°C$

Results (2/2)

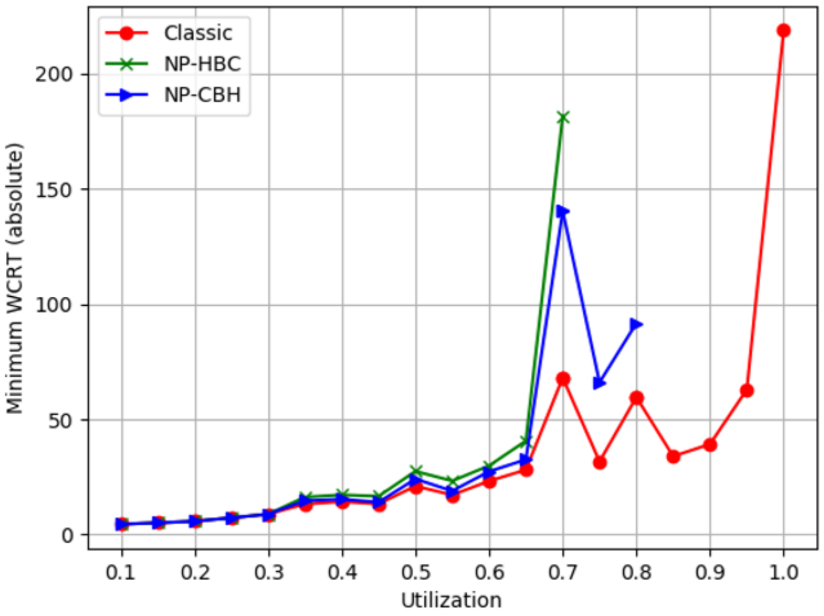

Classic
NP-HBC
NP-CBH
Minimum WCRT (absolute)
Utilization

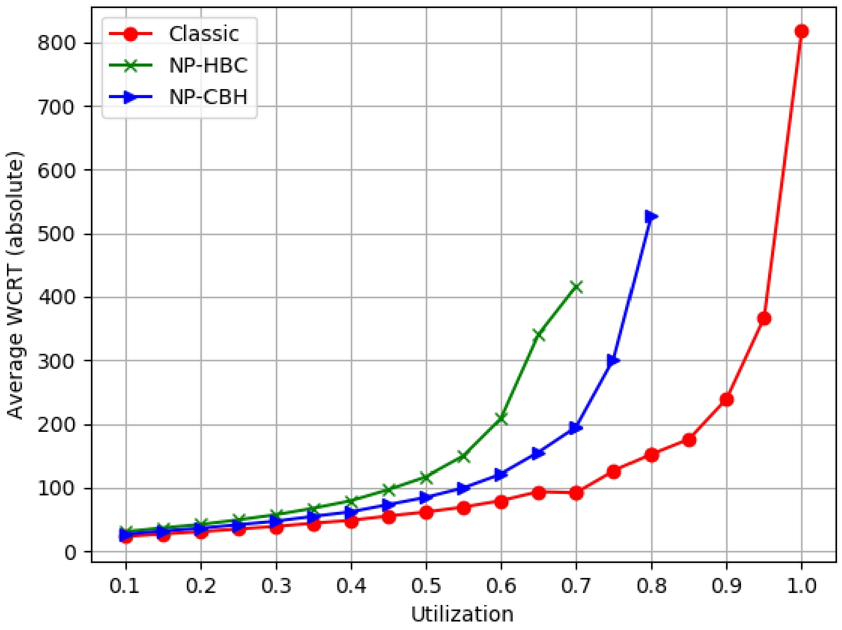

Classic
NP-HBC
NP-CBH
Average WCRT (absolute)
Utilization

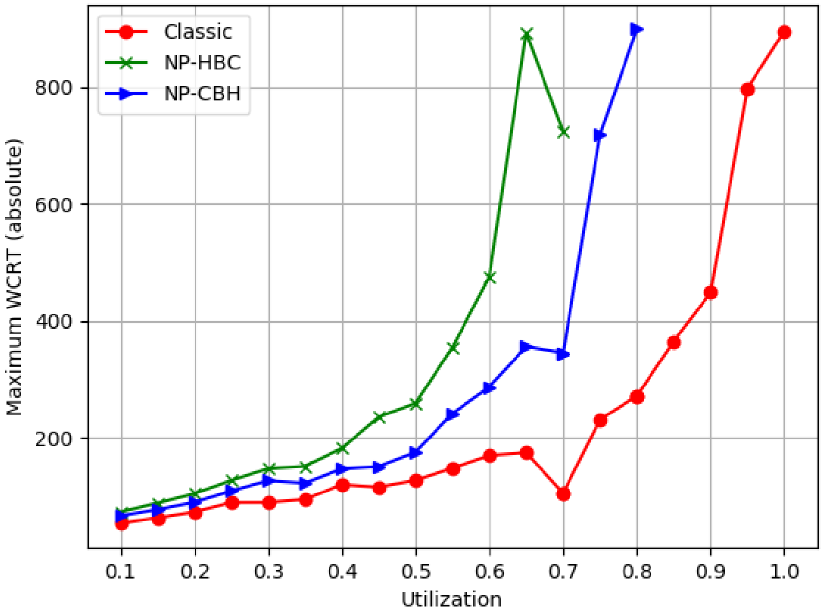

Classic
NP-HBC
NP-CBH
Maximum WCRT (absolute)
Utilization

Conclusion and future work

- Captured both the thermal and timing behaviors of the system in a single framework.
- Proposed two thermal-aware schedulers together with their schedulability analysis.
- Validated the run-time behavior of our solutions through intensive simulations.

- Future work:
  - How to extend the framework to DVFS-enabled platforms?
  - How to get around the priority assignment problem?
  - How to safely move to multi-cores?

Thermal-Aware Schedulability Analysis for Fixed-Priority Non-Preemptive Real-Time Systems

Javier Pérez Rodríguez

Patrick Meumeu Yomsi

CISTER – Research Centre in
Real-Time & Embedded Computing Systems