

Analysis of Preemptively Scheduled Hard Real-Time Systems

Sebastian Altmeyer

October 25th, 2012





Hard Real-Time Systems

- embedded systems with stringent timing constraints (deadlines)
- missing a deadline is a complete system failure
- examples are airbags, anti-lock breaking system (ABS), flight control systems etc.



Correct timing behavior must be guaranteed

Still valid in Presence of Caches and Preemption?



- Cache-related preemption delay (CRPD)
 - denotes overall cost of additional reloads due to preemption
 - may increase execution time significantly.

⇒

traditional interface can be considered outdated
(execution times are bounded for non-preempted execution)

¹memory blocks a,b,c,d, * indicates cache miss, direct-mapped cache

Aim of this Thesis

A precise and sound

- a) computation of the preemption cost, and
- b) integration with schedulability analysis.

① Prior Work on Bounding Cache-Related Preemption Delay

② Cache-Related Preemption Delay

Definition & Limitations

Definitely-Cached Useful Cache Blocks

CRPD for Set-Associative Caches (Resilience Analysis)

③ CRPD-Aware Schedulability Analysis

ECB-Union


Multiset Approaches

④ Conclusions

Evicting Cache Blocks

A memory block of the preempting task is called an evicting cache block, if it may be accessed during the execution of the preempting task.


- $|ECB| \rightarrow$ global CRPD bound (based on preempting task)



$$ECB = \{x, y, z\}$$

Combining UCBs/ECBs

Only memory blocks that are both **useful** and will be **evicted** due to preemption lead to additional cache misses due to preemption.



$$UCB \cap ECB = \{b, c\}$$

① Prior Work on Bounding Cache-Related Preemption Delay

② Cache-Related Preemption Delay

Definition & Limitations

Definitely-Cached Useful Cache Blocks

CRPD for Set-Associative Caches (Resilience Analysis)

③ CRPD-Aware Schedulability Analysis

ECB-Union

Multiset Approaches

④ Conclusions

Cache-Related Preemption Delay - Definition

Given an execution trace, a set of preemption points with preempting access sequence, initial cache state, and block reload time (BRT):

$$CRPD = (\# \text{misses on preempted execution} - \# \text{misses on non-preempted execution}) \cdot \text{BRT}$$

Difference to former definitions:

- based on concrete execution trace
- $CRPD \hat{=}$ total cost of all preemptions, not just for one preemption
- preemption cost defined with respect to the execution time

Cache-Related Preemption Delay - Limitations (1)

$$CRPD = (\# \text{misses on preempted execution} \\ - \# \text{misses on non-preempted execution}) \cdot BRT$$

CRPD depends on block reload time BRT:

⇒

BRT must be bounded

⇒

Separate CRPD-computation not possible for
non-compositional architectures
e.g. architectures with timing anomalies.

Redefining Useful Cache Blocks

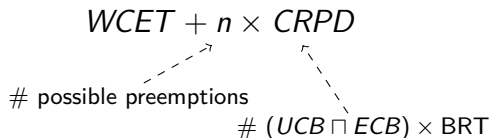
CRPD defined with respect to non-preempted execution time.
Why?

Schedulability analysis always computes:

$$WCET + n \times CRPD$$

possible preemptions

$(UCB \cap ECB) \times BRT$



However,
Timing analysis and UCB analysis are treated separately

Timing Analysis vs. UCB Analysis

Timing analysis:

- uses **underapproximation** of cache-content (**must-cache**)
- only predicts cache-hit, if accessed block \in **must-cache**

UCB analysis:


- uses **overapproximation** of cache-content (**may-cache**)
- predicts additional cache-miss, if accessed block \in **may-cache**

Hence, some cache-misses counted twice
(each access m with $m \in (\text{may-cache} \setminus \text{must-cache})$)

Useful Cache Blocks

A memory block m at program point p is called a useful cache block, if

- a) m may be **cached at p** , and
- b) m may be **reused at program point p'** that may be reached from p with no eviction of m on this path




$$UCB_p = \{a, b, c, d\}$$

Definitely-Cached Useful Cache Blocks

A memory block m at program point p is called a definitely-cached useful cache block, if


- m must be cached at p ,
- m may be reused at program point p' that may be reached from p with no eviction of m on this path, and
- m is considered a cache hit at p' by the timing analysis.




$$\text{DC-UCB}_p = \emptyset$$

Evaluation - DCUCBs

Mälardalen Benchmark Suite:



Papabench:



(ARM 7 with direct-mapped cache, 256 cache sets)

DC-UCBs and Timing Analysis

- DC-UCBs \subseteq UCBs
- DC-UCBs do not bound CRPD in isolation
- but: DC-UCBs bound CRPD with respect to timing analysis

\Rightarrow

Concept of DC-UCBs strongly improve prior CRPD bounds

CRPD for Set-Associative Caches

What was already published?

CRPD bounds for

- LRU caches solely based on ECBs
- LRU/FIFO/PLRU caches solely based on UCBs
- LRU/FIFO/PLRU caches based on UCBs and ECBs

So everything solved?

Unfortunately, only LRU solely based on UCBs correct.

Cache-Related Preemption Delay - Limitations (2)

In case of FIFO and PLRU, CRPD (for a single preemption) is not bounded by the associativity of the cache but may be proportional to the total execution time.



Concept of UCBs/ECBs is not sufficient for FIFO/PLRU


LRU Caches

In LRU caches, 2 different states converge after at most associativity-many distinct accesses

→

associativity is always a valid upper bound.

Problem: one ECB suffices to remove all UCBs




Only one ECB, but four additional misses
(former approaches fail here)

Simple CRPD Analysis for LRU

$$\text{CRPD} = \begin{cases} \text{BRT} \cdot |\text{UCB}(p)| & \text{if } \text{ECB} \neq \emptyset \\ 0 & \text{if } \text{ECB} = \emptyset \end{cases}$$

Sound but imprecise:



One ECB, three UCBs, but no additional misses.

In-depth Combination: Resilience Analysis

Resilience of UCB $m \hat{=}$

amount of disturbance, s.t. m remains useful


Resilience of m at p is given by associativity k - age of m at next reuse p' of m .

$$\text{res}_p(m) = \min(k - \text{age}(m)_{p'})$$

$$\text{CRPD} = \text{BRT} \times \left| \underbrace{\text{UCB}}_{\text{useful}} \setminus \underbrace{\{m \mid \text{res}(m) \geq |\text{ECB}|\}}_{\text{remain useful}} \right|$$

blocks contributing to CRPD

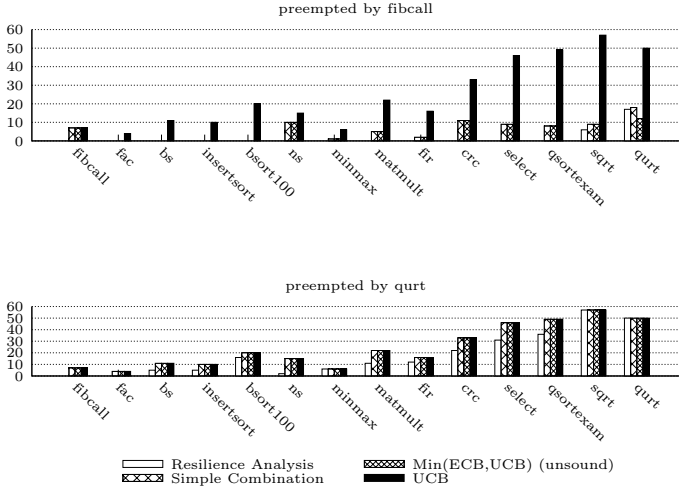
Resilience Analysis - Example



Blocks b, c, d are all 1-resilient,
i.e., survive preemption with up to one ECB

$$\begin{aligned} \text{CRPD} &= \text{BRT} \times |\text{UCB} \setminus \{m \mid \text{res}(m) \geq |\text{ECB}|\}| \\ &= \text{BRT} \times |\{b, c, d\} \setminus \{m \mid \text{res}(m) \geq 1\}| \\ &= \text{BRT} \times |\{b, c, d\} \setminus \{b, c, d\}| = 0 \end{aligned}$$

Evaluation - Mälardalen Benchmark Suite



(ARM 7, 4-way LRU, 64 cache sets)

CRPD for Set-Associative Caches

All former approaches either imprecise or incorrect or both.

UCB/ECB approach not applicable to PLRU/FIFO Caches

Resilience Analysis provides

- provably correct CRPD bounds for LRU Caches
- precise combination of UCBs/ECBs

- ① Prior Work on Bounding Cache-Related Preemption Delay
- ② Cache-Related Preemption Delay
 - Definition & Limitations
 - Definitely-Cached Useful Cache Blocks
 - CRPD for Set-Associative Caches (Resilience Analysis)
- ③ CRPD-Aware Schedulability Analysis
 - ECB-Union
 - Multiset Approaches
- ④ Conclusions

We now have precise CRPD bounds,
but how to use them?

In the following:

- Fixed priority scheduling
- Schedulability check via response time analysis (RTA)

Response Time Analysis

$$R_i = C_i + \sum_{\forall j \in \text{hp}(i)} \left\lceil \frac{R_i}{T_j} \right\rceil (C_j)$$

Response Time $R_i =$ finishing time – activation time
no deadline miss $\Leftrightarrow \forall i : R_i \leq D_i$

(exec. time C_i , period T_i , deadline D_i , tasks with higher priority $\text{hp}(i)$)

Response Time Analysis with CRPD


$$R_i = C_i + \sum_{\forall j \in \text{hp}(i)} \left\lceil \frac{R_i}{T_j} \right\rceil (C_j + \gamma_{i,j})$$

$\gamma_{i,j}$ denotes the preemption cost

How to compute $\gamma_{i,j}$?

- Only counting UCBs of preempted task
- Only counting ECBs of preempting task
- Simple combination $|\text{UCB}_i \cap \text{ECB}_j|$ possible?


Pitfall 1: Preemption of Intermediate Tasks



When computing R_3 :

τ_1 preempting τ_2 causes higher costs (1) than τ_1 preempting τ_3 (0)

Pitfall 2: Nested Preemption



When computing R_3 :

Nested preemption causes higher costs (2) than any non-nested (1)

New Approach: ECB-Union


$$\gamma_{i,j}^{\text{new}} = \text{BRT} \cdot \underbrace{\max_{k \in \text{aff}(i,j)} \left\{ \text{UCB}_k \cap \left(\bigcup_{h \in \text{hp}(j) \cup \{j\}} \text{ECB}_h \right) \right\}}_{\text{affected task with highest CRPD}} \quad \text{impact of all higher priority tasks}$$

- safe combination of ECBs and UCBs
- dominates UCB-Only
- resilience analysis fits well into ECB-Union approach (resilience needs ECB-Union of all prior preemptions)

Pessimism in basic RTA Formula

$$R_i = C_i + \sum_{\forall j \in \text{hp}(i)} \left\lceil \frac{R_i}{T_j} \right\rceil (C_j + \gamma_{i,j})$$

$\gamma_{i,j}$ = preemption cost of **one job** of τ_j during R_i




τ_1 preempting τ_2 considered 3 times within R_3

Pessimism in basic RTA Formula

$$R_i = C_i + \sum_{\forall j \in \text{hp}(i)} \left(\left\lceil \frac{R_i}{T_j} \right\rceil C_j + \gamma_{i,j} \right)$$

$\gamma_{i,j}$ = preemption cost of **all jobs** of τ_j during R_i



τ_1 preempting τ_2 considered only once within R_3

Multiset Approaches


$$R_i = C_i + \sum_{\forall j \in \text{hp}(i)} \left(\left\lceil \frac{R_i}{T_j} \right\rceil C_j + \gamma_{i,j} \right)$$

$\gamma_{i,j}$ = preemption cost of **all jobs** of τ_j during R_i

$$M = \bigcup_{k \in \text{aff}(i,j)} \left(\underbrace{\bigcup_{\lceil R_k/T_j \rceil \lceil R_i/T_k \rceil} \text{Cost}_{k,j}}_{\# \tau_j \text{ preempts } \tau_k \text{ during } R_i} \right) \quad \gamma_{i,j} = \text{BRT} \cdot \underbrace{\sum_{l=1}^{\lceil R_i/T_j \rceil} M^l}_{\text{highest values in } M}$$

M is a multiset, $\text{Cost}_{k,j}$ can be computed by any correct former approach

Evaluation



Randomly generated task sets with 10 tasks, implicit deadline, 256 cache sets, $CRT = 8\mu s$, cache utilization of 10

- ① Prior Work on Bounding Cache-Related Preemption Delay
- ② Cache-Related Preemption Delay
 - Definition & Limitations
 - Definitely-Cached Useful Cache Blocks
 - CRPD for Set-Associative Caches (Resilience Analysis)
- ③ CRPD-Aware Schedulability Analysis
 - ECB-Union
 - Multiset Approaches
- ④ Conclusions

Conclusions

Aim: A precise and sound

- a) computation of the preemption cost, and
- b) integration with schedulability analysis.

Conclusions

Aim: A precise and sound

- a) computation of the preemption cost, and
- b) integration with schedulability analysis.

Results:

- a) Preemption cost computation:
 - formal definition of CRPD
 - description of limitations of the CRPD approach
 - concept of Definitely-Cache Useful Cache Blocks
 - CRPD for LRU Caches (Resilience Analysis)

Conclusions

Aim: A precise and sound

- a) computation of the preemption cost, and
- b) integration with schedulability analysis.


Results:

- a) Preemption cost computation:
 - formal definition of CRPD
 - description of limitations of the CRPD approach
 - concept of Definitely-Cache Useful Cache Blocks
 - CRPD for LRU Caches (Resilience Analysis)
- b) CRPD-aware schedulability analysis:
 - ECB-Union approach
 - Multiset extensions
 - correct consideration of the blocking time


Questions?

CRPD - Limitations (2) - PLRU

Preempted execution ($2 * n$ misses):




Non-preempted execution (0 misses):




CRPD - Limitations (2) - FIFO

Preempted execution ($2 + 5 * n$ misses):



Non-preempted execution ($3 * n$ misses):



DC-UCB Analysis - Property on Concrete Paths

For Each memory block $m \in \text{DC-UCB}(p)$, exists a path s.t. m is reused at its next access p_n and m is classified as always hit at p_n :

$$m \in \text{DC-UCB}(p_i) \Leftrightarrow \exists [p_i, \dots, p_n] \in \Pi : \\ \#(p_n) = m \wedge \text{Classify}(m, p_n) = ah \wedge \forall_{j=i}^{n-1} \#(p_j) \neq m \quad (1)$$

$$\text{Els} : V \times M \rightarrow \{\text{true}, \text{false}\}$$

$$\text{Els}(\pi, m) = \begin{cases} \text{true} & \text{if } \pi = [p_1, \dots, p_n] \cdot \pi' \wedge \#(p_n) = m \\ & \wedge \text{Classify}(m, p_n) = ah \wedge \forall_{j=i}^{n-1} \#(p_j) \neq m \\ \text{false} & \text{otherwise} \end{cases} \quad (2)$$

DC-UCB - Abstract Transformer

$$tf : V \rightarrow (2^M \rightarrow 2^M)$$
$$tf(p)(S) = \begin{cases} S & \#(p) = \perp \\ S \cup \{\#(p)\} & \text{Classify}(\#(p), p) = ah \\ S \setminus \{\#(p)\} & \text{Classify}(\#(p), p) \neq ah \end{cases} \quad (3)$$

DC-UCB Analysis - Concretization/Abstraction

$$\gamma : 2^M \rightarrow 2^\Pi$$


$$\gamma(S) := \{\pi \mid \pi \in \Pi : \forall m \in (M \setminus S) : \neg \text{Els}(\pi, m)\} \quad (4)$$

$$\alpha : 2^\Pi \rightarrow 2^M$$

$$\alpha(S) := \{m \mid \exists \pi \in S : \text{Els}(\pi, m)\} \quad (5)$$

Resilience Analysis - Interacting Preemptions

Problem: preemptions may interact



$$\text{CRPD}_{p_1} = \text{CRPD}_{p_2} = 0, \text{ but total CRPD} = 2$$

Solution: Compute ECB-Union of all prior preemptions PR


$$\text{CRPD} = \text{BRT} \times |\text{UCB} \setminus \{m \mid \text{res}(m) \geq |\bigcup_{PR} \text{ECB}|\}|$$

Example:

$$\text{CRPD}_{p_1} = \text{BRT} \times |\text{UCB} \setminus \{m \mid \text{res}(m) \geq |\{x\}|\}| = 0$$

$$\text{CRPD}_{p_2} = \text{BRT} \times |\text{UCB} \setminus \{m \mid \text{res}(m) \geq |\{x, y\}|\}| = 2$$

Resilience under Different Paths



Resilience under different paths, assuming a 4 way LRU cache.

Resilience Analysis - Constrained/Unconstrained Age

$$tf_{ua} : V \rightarrow \text{Age} \rightarrow \text{Age}$$

$$tf_{ua}(p)(ua) :=$$

$$\lambda m. \begin{cases} 0 & m = \#(p) \\ ua(m) & ua(m) \geq ua(\#(p)) \\ ua(m) + 1 & ua(m) < ua(\#(p)) \wedge ua(m) < ua - 1 \\ \infty & \text{otherwise} \end{cases} \quad (6)$$

The accessed element $\#(p)$ is assigned age zero, all younger elements (as the accessed element) age by one.

$$tf_{ca} : V \rightarrow (\text{Age} \times \text{Age} \rightarrow \text{Age})$$

$$tf_{ca}(p)(ca, ua) :=$$

$$\lambda m. \begin{cases} 0 & m = \#(p) \wedge m \notin \text{UCB}(p) \\ ca(m) & ca(m) \geq ua(\#(p)) \vee ca(m') = k - 1 \\ ca(m) + 1 & ca(m) < ua(\#(p)) \end{cases} \quad (7)$$

Resilience - Concretization/Abstraction

$$\gamma^{\rightarrow}((ua, ca)) = \{(p \cdot \pi) \in \Pi \mid \forall m \in M : (\widehat{age}^{\rightarrow}(p \cdot \pi)(m) \leq ca^{\rightarrow}(m)) \\ \vee (\widehat{age}^{\rightarrow}(p \cdot \pi)(m) \leq ua^{\rightarrow}(m) \wedge m \notin \text{UCB}_p(m))\} \quad (8)$$


$$\alpha^{\rightarrow}(S) = (\lambda m. \max\{\widehat{age}^{\rightarrow}(p \cdot \pi)(m) \mid (p \cdot \pi) \in S\}, \\ \lambda m. \max\{\widehat{age}^{\rightarrow}(p \cdot \pi)(m) \mid (p \cdot \pi) \in S \wedge m \in \text{UCB}(p)\}) \cup \{0\} \quad (9)$$

Prior Work: UCB only or ECB only

$$R_i = C_i + \sum_{\forall j \in \text{hp}(i)} \left\lceil \frac{R_j}{T_j} \right\rceil (C_j + \gamma_{i,j})$$

$$\gamma_{i,j}^{\text{ecb}} = \text{BRT} \cdot |\text{ECB}_j| \quad \text{or} \quad \gamma_{i,j}^{\text{ucb}} = \text{BRT} \cdot \max_{\forall k \in \text{aff}(i,j)} \{|\text{UCB}_k|\}$$

$$(\text{aff}(i,j) = \text{hep}(i) \cap \text{lp}(j))$$



$$\gamma_{2,1}^{\text{ucb}} = \gamma_{2,1}^{\text{ecb}} = 2$$

actual cost = 0

Prior Work: UCB-Union


$$\gamma_{i,j}^{\text{tan}} = \text{BRT} \cdot \underbrace{\left| \left(\bigcup_{\forall k \in \text{aff}(i,j)} \text{UCB}_k \right) \cap \text{ECB}_j \right|}_{\text{that are evicted by task } \tau_j}$$

all possibly affected UCBs

-
- safe combination of ECBs and UCBs
 - dominates ECB-Only ($\gamma_{i,j}^{\text{ecb}} = \text{BRT} \cdot |\text{ECB}_j|$)
 - poor precision: set UCBs per task, not per program point

Blocking Time (Stack Resource Protocol)

$$R_i = C_i + B_i + \sum_{\forall j \in \text{hp}(i)} \left\lceil \frac{R_j + J_j}{T_j} \right\rceil (C_j + \gamma_{i,j})$$



Task τ_2 can be blocked by execution of τ_3 and pre-emption delay (τ_1 pre-empting τ_3)

ECB-Only accounts for this implicitly
all others must be extended

Evaluation - Cache Configurations

	Cache Type	Sets	Line Size	Total Size
Config 1:	direct-mapped	256	16 Byte	4kB
Config 2:	4-way LRU	64	16 Byte	4kB
Config 3:	8-way LRU	64	16 Byte	8kB



Mälardalen Benchmark Suite

Task	Description	Code Size	WCET
minmax	Derives min/max of set of integers	608B	298 μ s
insertsort	Insertion sort on array of size 10.	384B	223 μ s
fibcall	Iterative Fibonacci calculation).	256B	117 μ s
fac	Non-recursive Faculty Computation	256B	67 μ s
bs	Binary search (array of size 15).	320B	1.46ms
bsort100	Bubblesort program.	544B	11.3ms
ns	Test for deeply nested loops	576B	637 μ s
matmult	Matrix multiplication (20x20).	864B	8.5ms
fir	Finite impulse response filter.	928B	465 μ s
crc	Cyclic redundancy check .	1216B	2.7ms
select	Selects Nth largest number.	1280B	757 μ s
qsort-exam	Non-recursive quick sort algorithm.	1440B	800 μ s
sqrt	Square root by Taylor series.	3680B	2.1ms
qurt	Computation of quadr. equations.	4160B	2.9ms

Papabench



Task	Description	Period	Priority	Code Size	WCET
I5	interrupt_spi_1	50ms	1	304B	129 μ s
I6	interrupt_spi_2	50ms	2	144B	68 μ s
T12	stabilization	50ms	3	2976B	3.2ms
I4	interrupt_modem	100ms	4	288B	148 μ s
T11	reporting_task	100ms	5	5360B	5.9ms
T10	receive_gps_data	250ms	6	3008B	3ms
T7	link_fbw_send	250ms	7	192B	105 μ s
T6	climb_control	250ms	8	3296B	3.4ms
T5	altitude_control	250ms	9	1232B	826 μ s

Evaluation - Papabench



(ARM 7, 4-way LRU, 64 cache sets)


Evaluation - Mälardalen Benchmark Suite




(ARM 7, 4-way LRU, 64 cache sets)

Evaluation - DCUCBs

Mälardalen Benchmark Suite:



Papabench:



(ARM 7 with direct-mapped cache, 256 cache sets)

Evaluation - Setting


Task set:

- 10 tasks
- periods T_i range from 5ms to 500ms (log-uniform distribution)
- task utilization U_i generated using UUnifast
- execution times $C_i = U_i \cdot T_i$
- implicit deadlines, priorities in deadline monotonic order

Preemption costs:

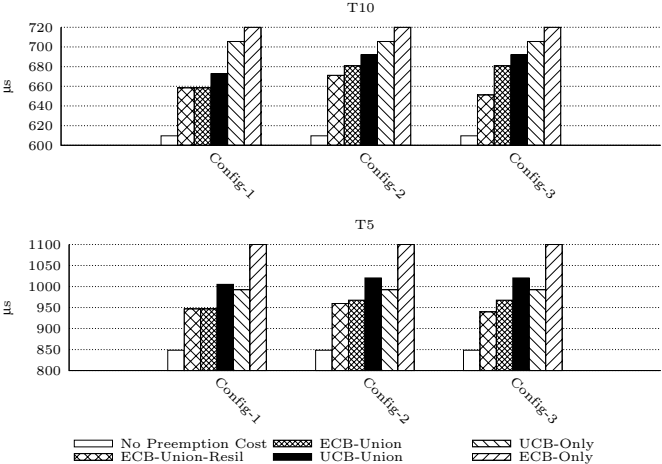
- number of cache sets ($CS = 256$)
- block-reload time ($CRT = 8\mu s$)
- cache usage using UUnifast ($CU = 10$)
- reuse factor (UCBs), uniform distribution $[0; |ECB|]$

Evaluation



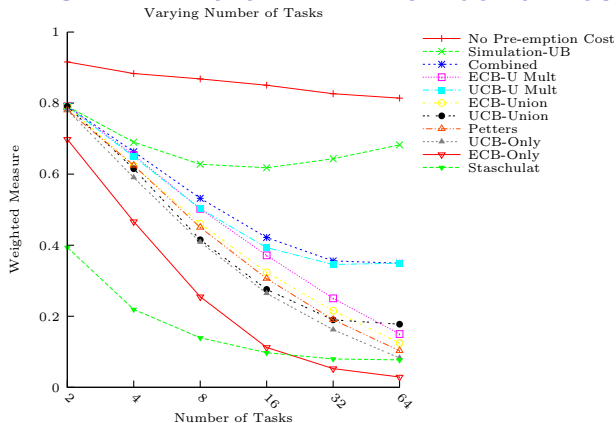
Randomly generated task sets with 10 tasks, implicit deadline, 256 cache sets, $CRT = 8\mu s$, cache utilization of 10

Evaluation - CRPD-Aware RTA



Bounds on the response time of T5 and T10 (in μs)

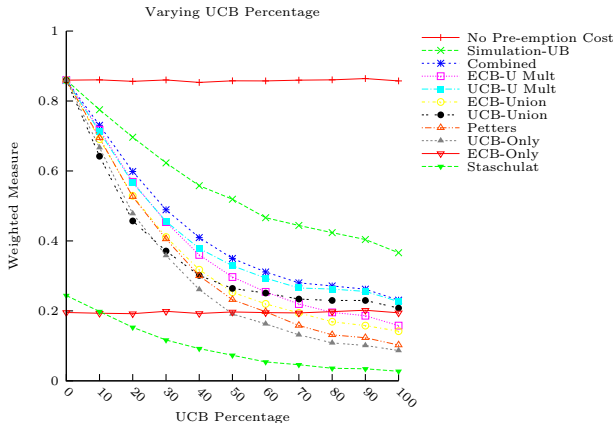
Evaluation - CRPD-Aware RTA - Number of Tasks



Weighted schedulability measure; varying number of tasks from $2 = 2^1$ to $2^6 = 64$.

$$W_y(p) = \left(\sum_{\forall \tau} u(\tau) \cdot S_y(\tau, p) \right) / \sum_{\forall \tau} u(\tau) \quad (10)$$

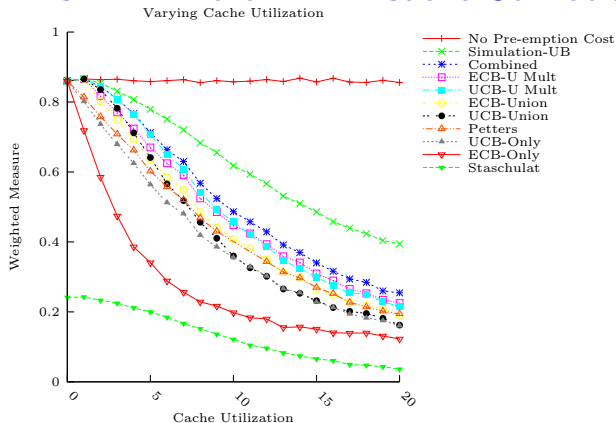
Evaluation - CRPD-Aware RTA - Reuse Factor



Weighted schedulability measure; varying reuse factor from 0% to 100%, in steps of 10%.

$$W_y(p) = \left(\sum_{\forall \tau} u(\tau) \cdot S_y(\tau, p) \right) / \sum_{\forall \tau} u(\tau) \quad (11)$$

Evaluation - CRPD-Aware RTA - Cache Utilization



Weighted schedulability measure; varying cache utilization from 0 to 20, in steps of 2.

$$W_y(p) = \left(\sum_{\forall \tau} u(\tau) \cdot S_y(\tau, p) \right) / \sum_{\forall \tau} u(\tau) \quad (12)$$