Many-Core Timing Analysis of Real-Time Systems

and its application to an industrial processor

Hamza Rihani

Université Grenoble Alpes / Verimag

December 1, 2017

Jury:

Pr. Jan Reineke	Saarland University	*Reviewer*
Pr. Christine Rochange	Université de Toulouse	*Reviewer*
Dr. Robert I. Davis	University of York	*Examiner*
Dr. Benoît de Dinechin	Kalray SA	*Examiner*
Dr. Claire Maïza	Université Grenoble Alpes	*Supervisor*
Dr. Matthieu Moy	Université Claude Bernard - Lyon 1	*Advisor*

Introduction: Real-Time Systems

Many-Core Timing Analysis of **Real-Time Systems**

Definition (Real-Time Systems)

A system that must produce **valid** outputs before a **deadline**.

Many-Core Timing Analysis of **Real-Time Systems**

Definition (Real-Time Systems)

A system that must produce **valid** outputs before a **deadline**.

- **Soft Real-Time**
  - Global Positioning System device
  - Smartphones
- **Hard Real-Time**
  - Automatic Braking System
  - Flight Management System

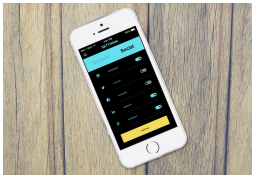

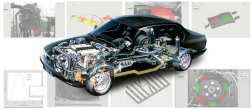

Many-Core Timing Analysis of **Real-Time Systems**

Definition (Real-Time Systems)

A system that must produce **valid** outputs before a **deadline**.

- **Soft Real-Time**
  - Global Positioning System device
  - Smartphones
- **Hard Real-Time** ★
  - Automatic Braking System
  - Flight Management System

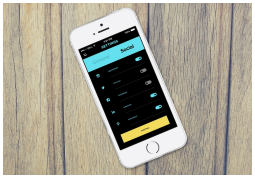

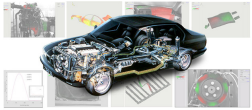

Introduction: Timing Analysis of Arbitration Policies

Many-Core **Timing Analysis** of Real-Time Systems

How long will the truck wait to cross the road?

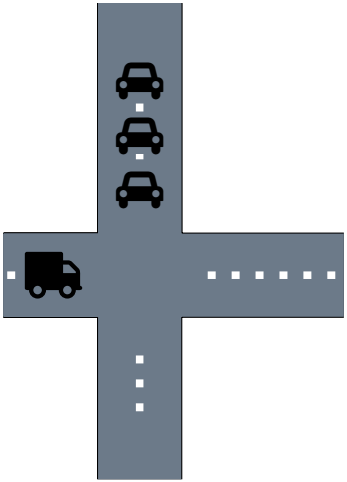

Introduction: Timing Analysis of Arbitration Policies

Many-Core **Timing Analysis** of Real-Time Systems

How long will the truck wait to cross the road?

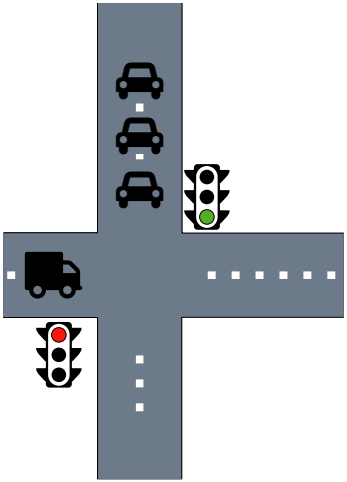

waits for the green light

Introduction: Timing Analysis of Arbitration Policies

Many-Core **Timing Analysis** of Real-Time Systems

How long will the truck wait to cross the road?

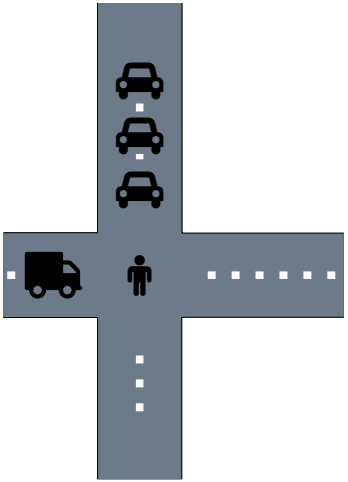

waits for the green light

grants each direction at a time

Introduction: Timing Analysis of Arbitration Policies

Many-Core **Timing Analysis** of Real-Time Systems

How long will the truck wait to cross the road?

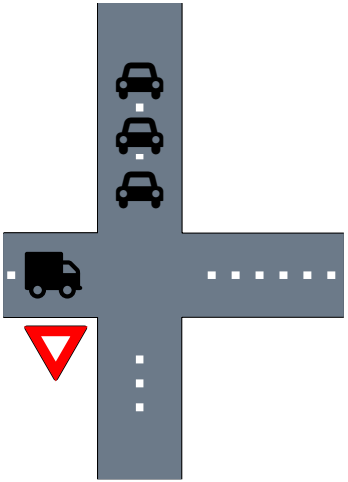

waits for the green light

grants each direction at a time

gives a priority to the cars

Introduction: Timing Analysis of Arbitration Policies

Many-Core **Timing Analysis** of Real-Time Systems

How long will the truck wait to cross the road?

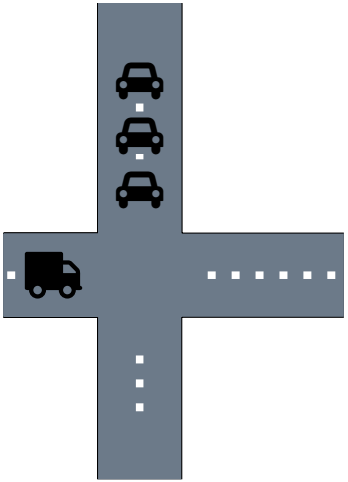

- Crossroad is a shared resource
- Vehicles request accesses to pass

- **Arbitration Policies:**
  - *Time Division Multiple Access*
  - *Round Robin*
  - *Fixed Priority*

Introduction: Many-Cores in Real Time Systems

Many-Core Timing Analysis of Real-Time Systems

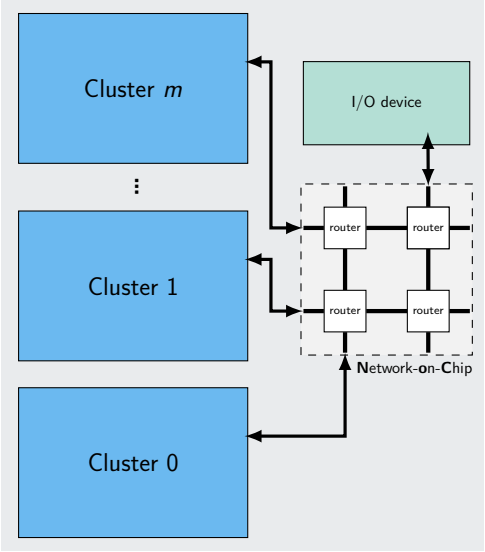

Many-Core Timing Analysis of Real-Time Systems

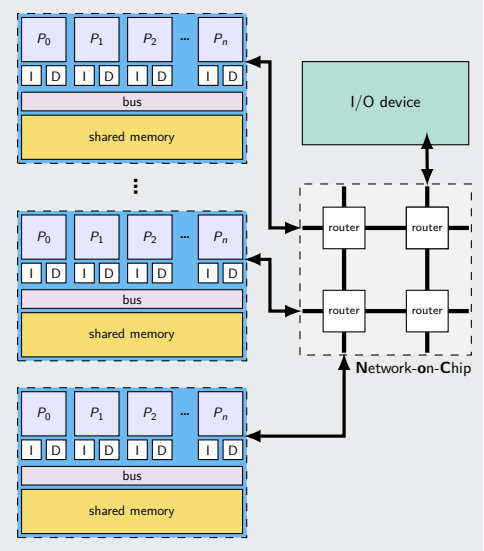

Timing analysis of cores

- Existing tools for pipeline and cache analyses

Introduction: Many-Cores in Real Time Systems

Many-Core Timing Analysis of Real-Time Systems

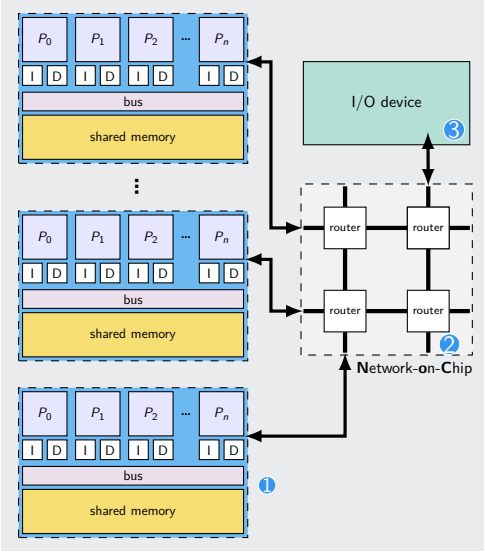

Timing analysis of cores

- Existing tools for pipeline and cache analyses

Where is the potential interference?

1. Shared buses and memory ★
2. NoC routing
3. Shared I/O controllers

Contributions

Contribution 1

Analysis of Time Division Multiple Access policy

- Approach based on Satisfiability Modulo Theory

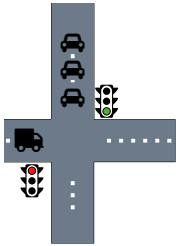

Contributions

Contribution 1

Analysis of Time Division Multiple Access policy

- Approach based on Satisfiability Modulo Theory

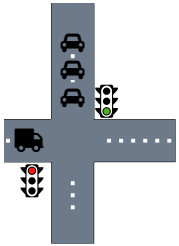

Contribution 2

Response time analysis of a many-core processor

- Synchronous Data Flow programs
- Model of the shared bus arbiter

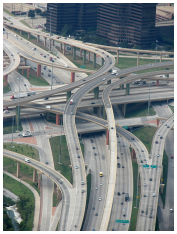

The High Five, Dallas, Texas, USA

Outline

TDMA Bus Timing Analysis

Definition: Time Division Multiple Access

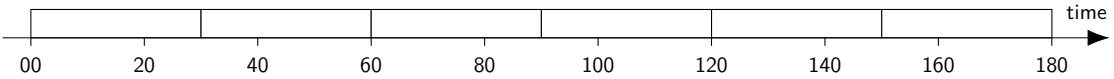

Definition: Time Division Multiple Access

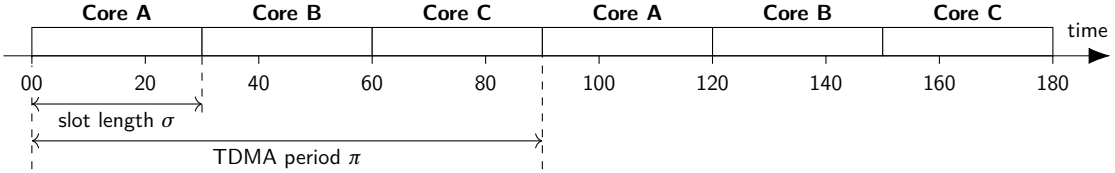

Definition: Time Division Multiple Access


Core A viewpoint:

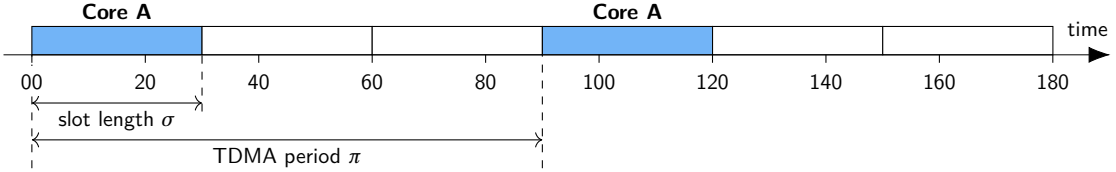

Definition: Time Division Multiple Access

Core A viewpoint:

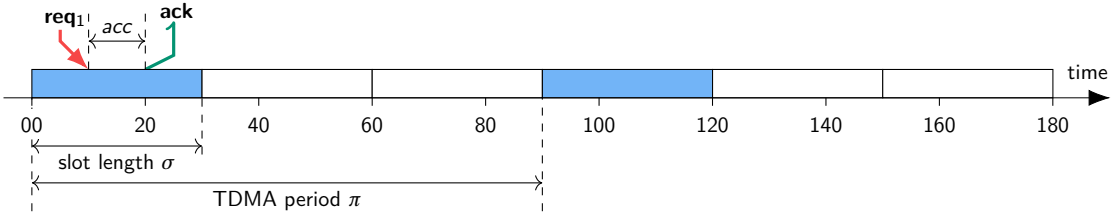

Definition: Time Division Multiple Access

Core A viewpoint:

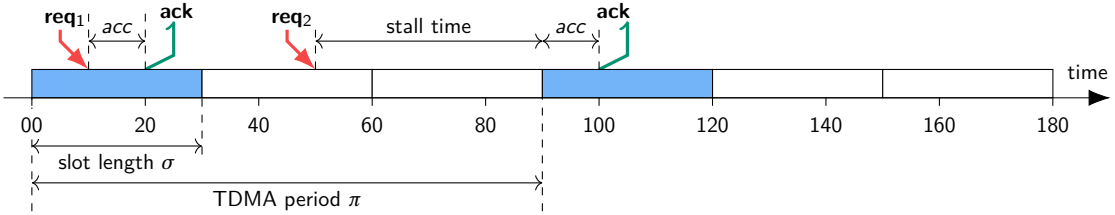

Definition: Time Division Multiple Access

Core A viewpoint:

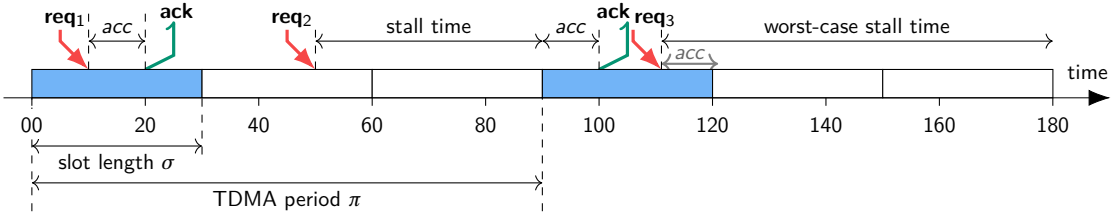

Worst-Case Stall Time = π - (σ - acc)

Definition: Time Division Multiple Access

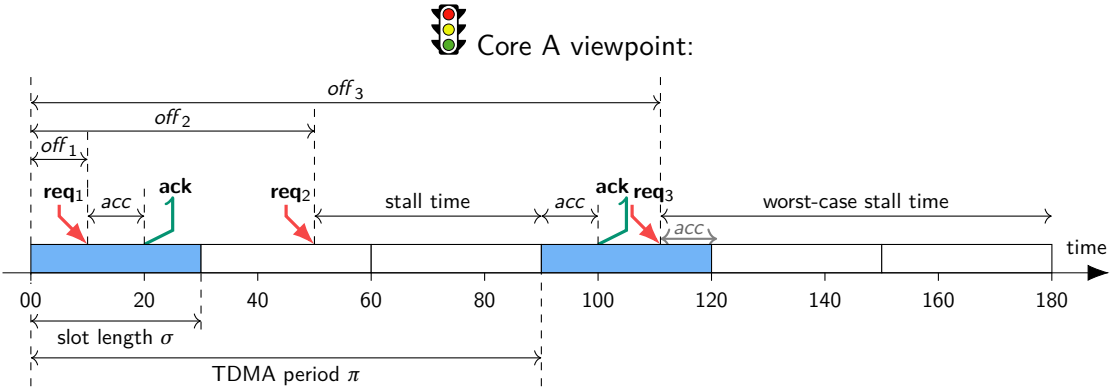

Worst-Case Stall Time $= \pi$ - (σ - *acc*)

- **Offsets** off_1, off_2, off_3 relative to the TDMA period:

$$off_{\{1,2,3\}} = time\_instant(req_{\{1,2,3\}}) \mod \pi$$

Outline: TDMA Bus Timing Analysis

Outline: TDMA Bus Timing Analysis

Worst-Case Execution Time (WCET) Analysis of TDMA

```
int f(int x){
  /* 3 cycles */
  if(cond)
  {
    /* 10 cycles */
  }
  if(!cond)
  {
    /*bus access */
  }
  return;
}
```

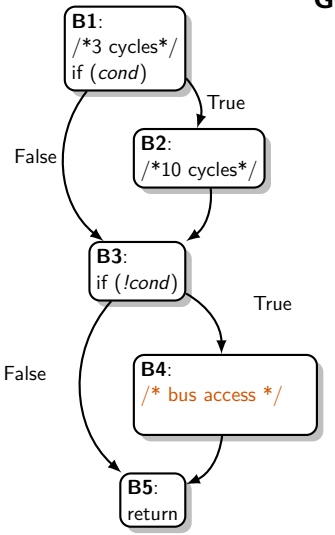

Goal: Estimate the WCET

Worst-Case Execution Time (WCET) Analysis of TDMA

```
int f(int x){
  /* 3 cycles */
  if(cond)
  {
    /* 10 cycles */
  }
  if(!cond)
  {
    /*bus access */
  }
  return;
}
```

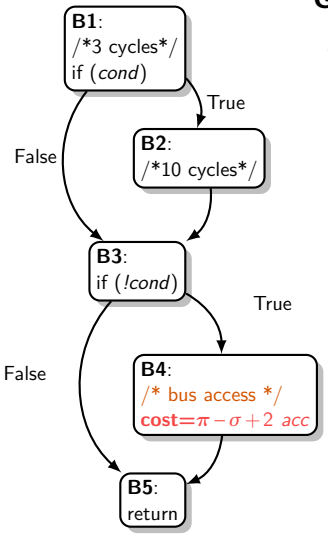

Goal: Estimate the WCET

↪ **Existing approaches:**

1. Worst-case everywhere

[Altmeyer et al., 2015; Rosèn et al., 2007...]

Worst-Case Execution Time (WCET) Analysis of TDMA

```
int f(int x){
  /* 3 cycles */
  if(cond)
  {
    /* 10 cycles */
  }
  if(!cond)
  {
    /*bus access */
  }
  return;
}
```

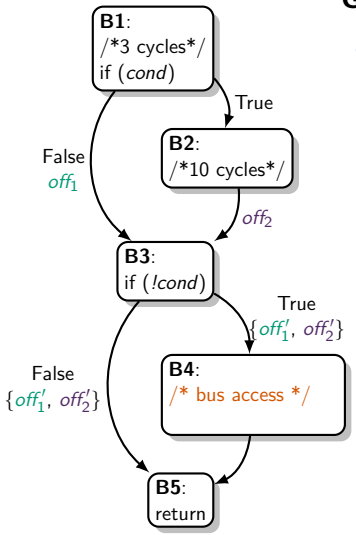

Goal: Estimate the WCET

↪ **Existing approaches:**

1. Worst-case everywhere

   [Altmeyer et al., 2015; Rosèn et al., 2007...]

2. Capture all possible offsets

   [Chattopadhyay et al., 2010; Kelter et al., 2014...]

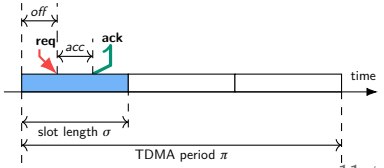

Worst-Case Execution Time (WCET) Analysis of TDMA

```
int f(int x){
  /* 3 cycles */
  if(cond)
  {
    /* 10 cycles */
  }
  if(!cond)
  {
    /*bus access */
  }
  return;
}
```

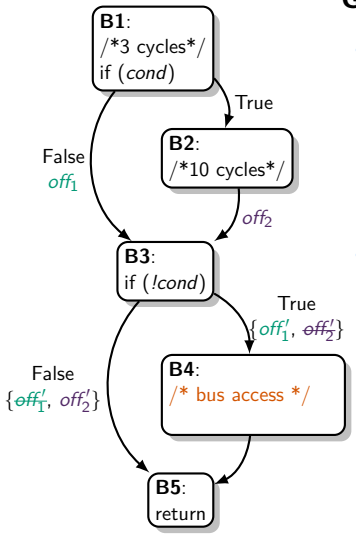

Goal: Estimate the WCET

↪ **Existing approaches:**

1. Worst-case everywhere

   [Altmeyer et al., 2015; Rosèn et al., 2007...]

2. Capture all possible offsets

   [Chattopadhyay et al., 2010; Kelter et al., 2014...]

↪ **Combined with:**

3. Feasible Path Analysis

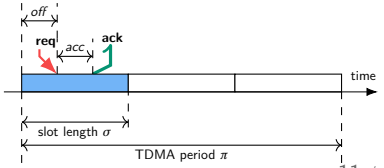

Approaches in WCET Analysis of TDMA

❷

Capture all possible offsets

[Kelter et al., 2014]

[Chattopadhyay et al., 2010]

❸

Feasible Path Analysis with SMT

[Henry et al., 2014]

Approaches in WCET Analysis of TDMA

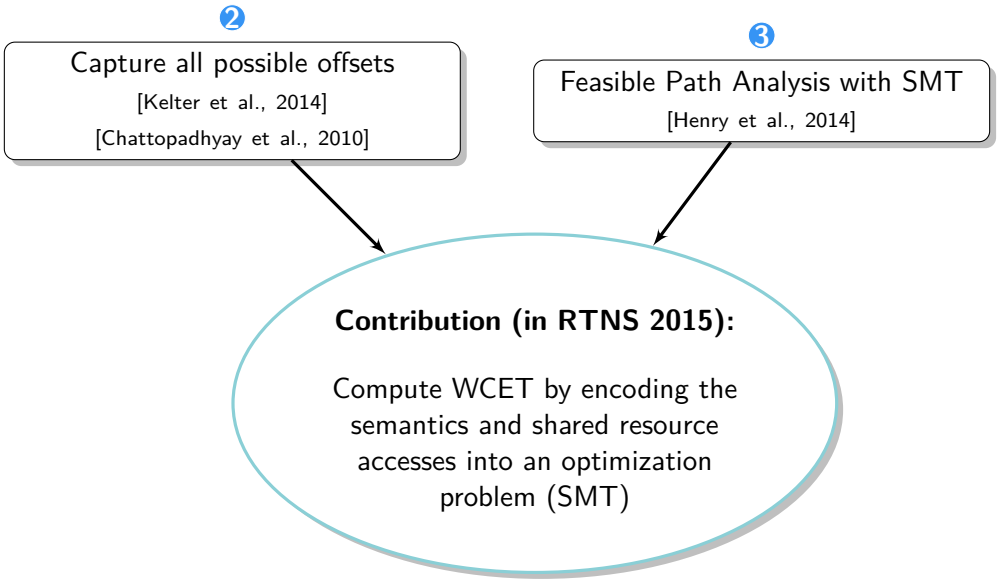

Outline: TDMA Bus Timing Analysis

WCET Analysis by SMT Encoding

- Bounded Model Checking
  - Encode the semantics into a Satisfiability Modulo Theory problem

WCET Analysis by SMT Encoding

- Bounded Model Checking
  - Encode the semantics into a Satisfiability Modulo Theory problem

$$\underbrace{\text{SMT query}}_{\text{assert}(\wedge \text{expr})} = \text{“Is there a trace with a feasible path?”}$$

- SMT-solver response:
  - SAT: There is a feasible execution path
  - UNSAT: There is no feasible execution path

WCET Analysis by SMT Encoding

- Bounded Model Checking
  - Encode the semantics into a Satisfiability Modulo Theory problem
- **Add execution times on the paths**

$$\underbrace{\text{SMT query}}_{\text{assert}(\wedge \text{expr})} = \text{“Is there a trace with a feasible path}$$

...such that the execution time is greater than X?”

- SMT-solver response:
  - SAT: **There is a feasible path with an execution time $> X$**
  - UNSAT: **X is an upper-bound on WCET**

WCET Analysis by SMT Encoding

- Bounded Model Checking
  - Encode the semantics into a Satisfiability Modulo Theory problem
- **Add execution times on the paths**

$$\underbrace{\text{SMT query}}_{\text{assert}(\wedge \text{expr})} = \text{“Is there a trace with a feasible path}$$

...such that the execution time is greater than X?”

- SMT-solver response:
  - SAT: **There is a feasible path with an execution time $> X$**
  - UNSAT: **X is an upper-bound on WCET**

Goal

Find the smallest X, such that Execution Time $> X$ is UNSAT

Example: Semantics and Timing Encoding

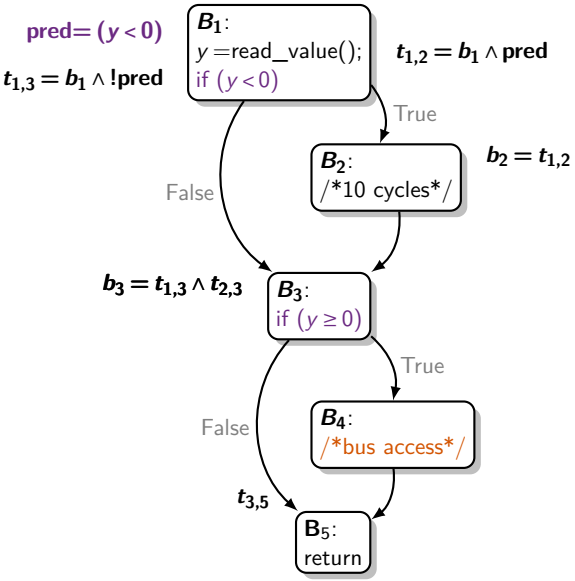

Previous work in [Henry et al., 2014]

- b_i *"true"* $\overset{\text{def}}{\iff}$ B_i executed
- $t_{i,j}$ *"true"* $\overset{\text{def}}{\iff}$ $B_i \rightarrow B_j$ taken

Example: Semantics and Timing Encoding

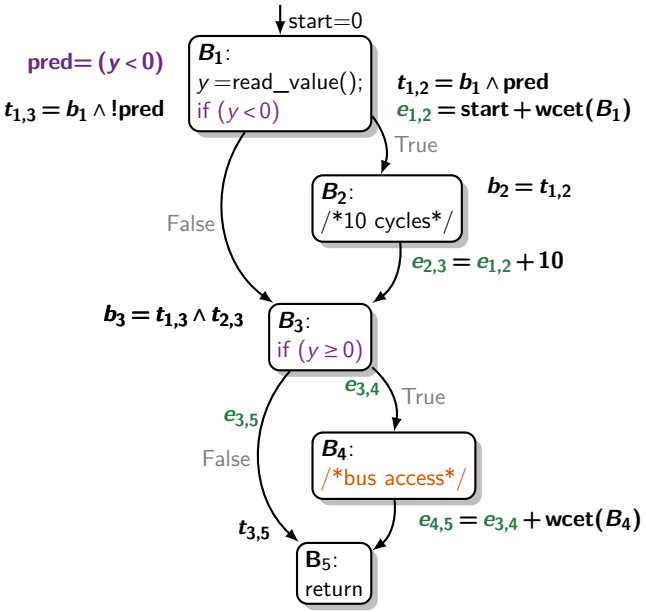

Previous work in [Henry et al., 2014]

- b_i *"true"* $\overset{\text{def}}{\Longleftrightarrow}$ B_i executed
- $t_{i,j}$ *"true"* $\overset{\text{def}}{\Longleftrightarrow}$ $B_i \rightarrow B_j$ taken
- $e_{i,j}$ *execution time* at transition $B_i \rightarrow B_j$

Example: Semantics and Timing Encoding

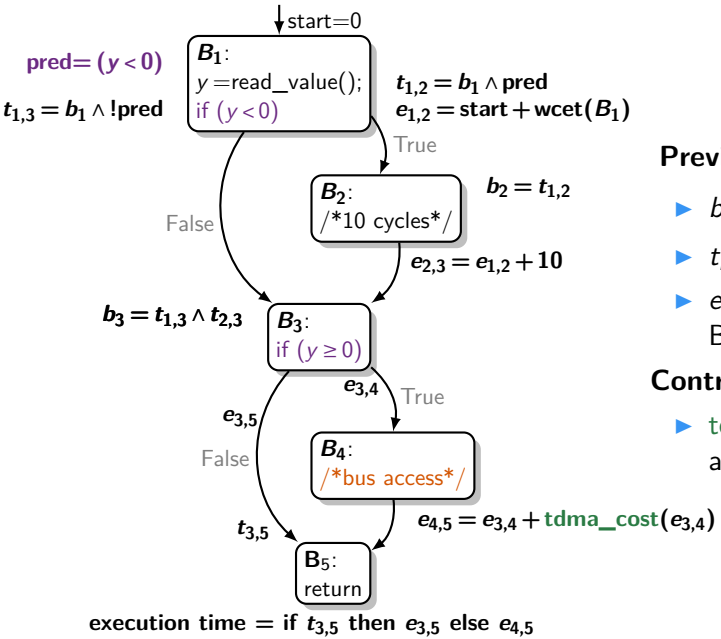

Previous work in [Henry et al., 2014]

- b_i "true" $\overset{\text{def}}{\iff}$ B_i executed
- $t_{i,j}$ "true" $\overset{\text{def}}{\iff}$ $B_i \rightarrow B_j$ taken
- $e_{i,j}$ *execution time* at transition $B_i \rightarrow B_j$

Contribution

- tdma_cost() execution time of a bus access

Naive SMT Encoding

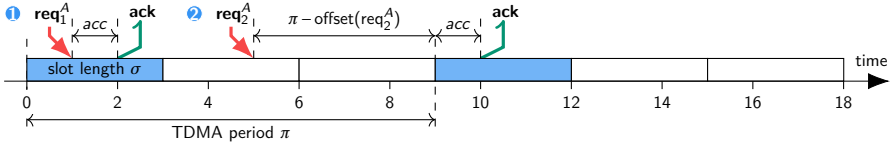

tdma_cost(e_{entry} ***)***: returns the execution time of a bus access

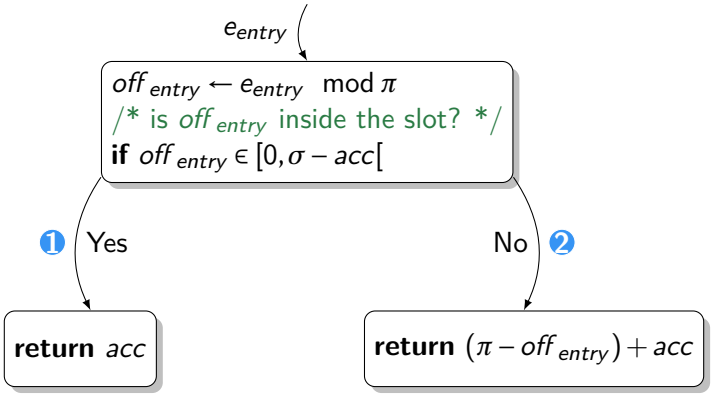

Performance of the Naive Encoding

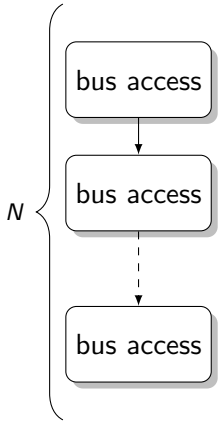

Performance of the Naive Encoding

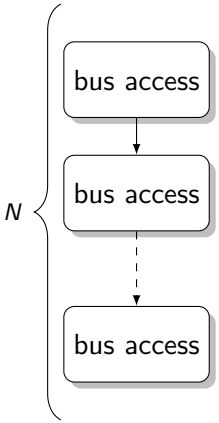

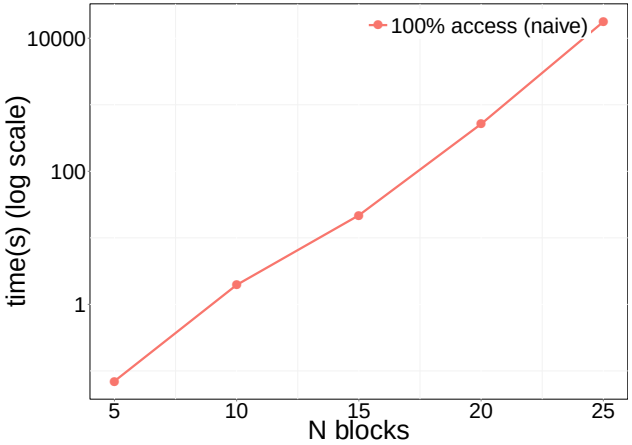

Naive SMT Encoding

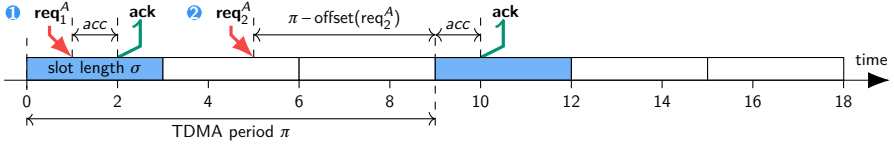

tdma_cost: returns the execution time of a bus access

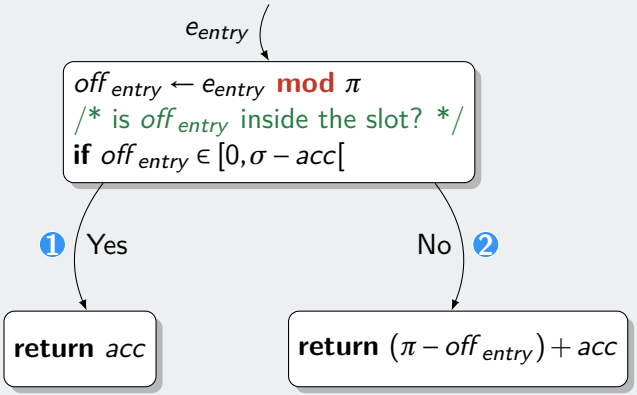

Offset-based SMT Encoding

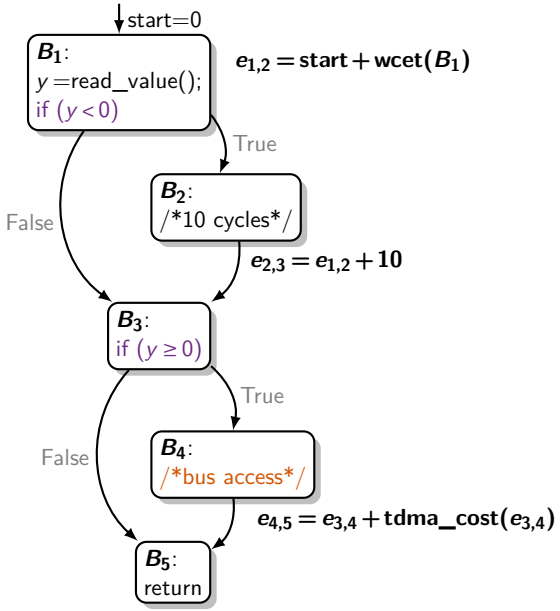

- $\mathrm{off}_{i,j} = e_{i,j} \mod \pi$

Offset-based SMT Encoding

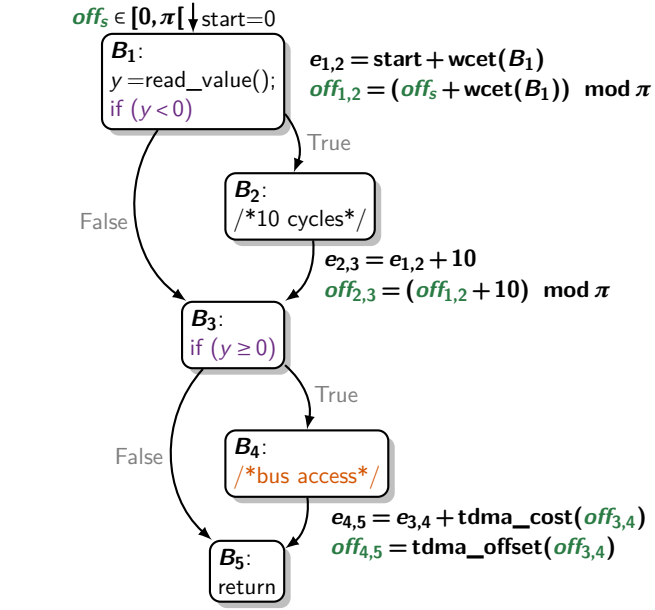

execution time $=$ **if** $t_{3,5}$ **then** $e_{3,5}$ **else** $e_{4,5}$

- $\text{off}_{i,j} = e_{i,j} \mod \pi$
- $\text{off}_{i,j}$ offset at transition $B_i \to B_j$

Offset-based SMT Encoding

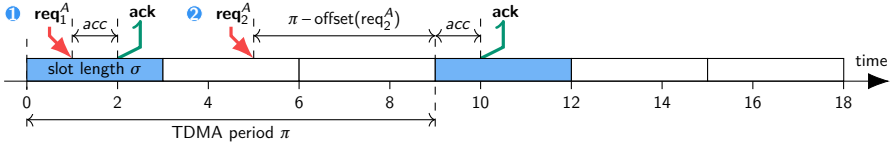

tdma_cost: returns the time after a bus access

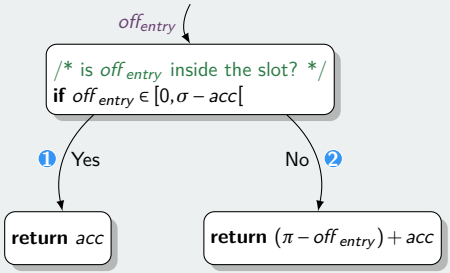

tdma_offset: returns the offset after a bus access

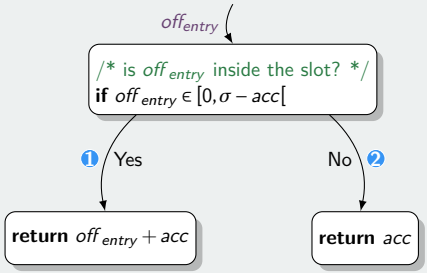

Performance of the Offset-based Encoding

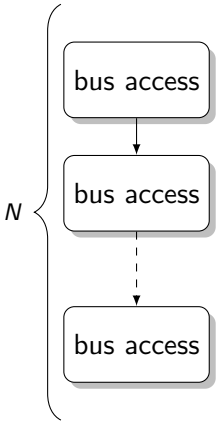

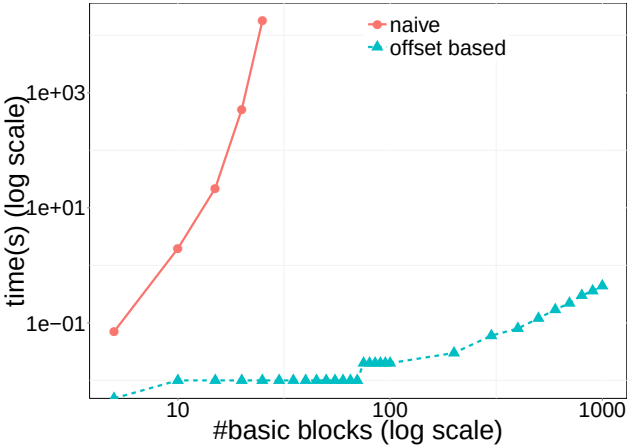

Outline: TDMA Bus Timing Analysis

Proof-of-Concept Implementation

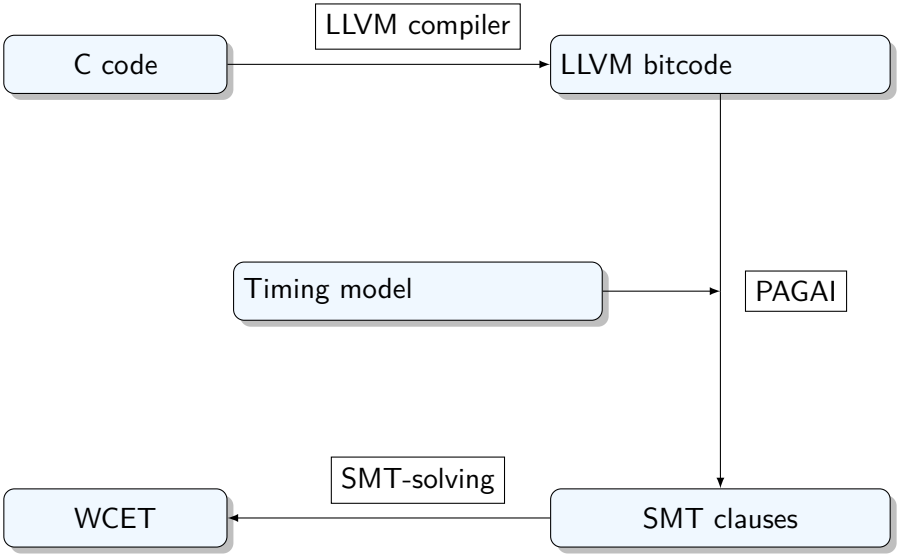

Proof-of-Concept Implementation

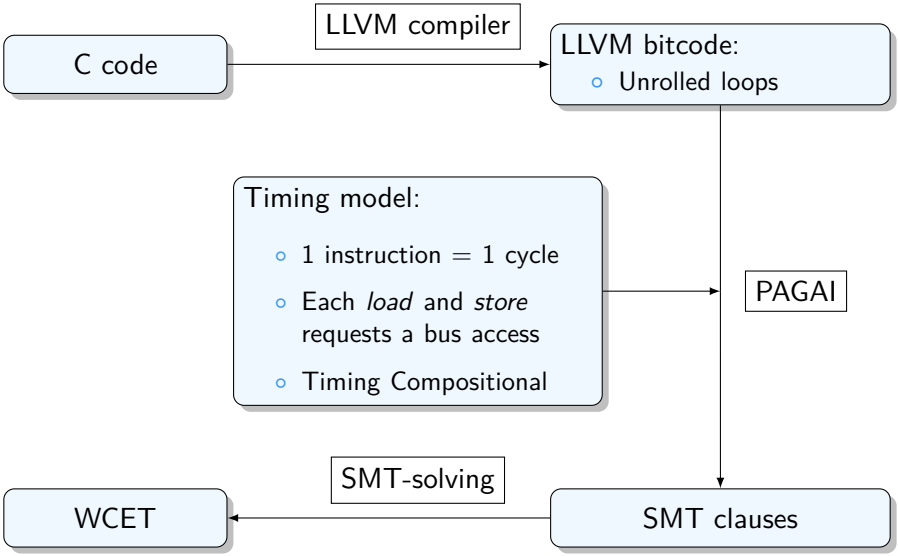

Evaluation: Benchmark Descriptions

Benchmark from TACLEBench suite [1]

Name	Description	#LLVM instr.	#bus access
bs	Binary search	231	12
insertsort	Insertion sort on a reversed array	493	65
jfdctint	Discrete Cosine Transformation	2334	448
fdct	Fast Discrete Cosine Transform	2502	385
compressdata	Data compression program adopted from SPEC95	674	131
fly-by-wire	UAV fly-by-wire software	2815	515

[1] `https://github.com/tacle/tacle-bench`

Evaluation: Experiments

Comparison between estimated WCET and pessimistic WCET

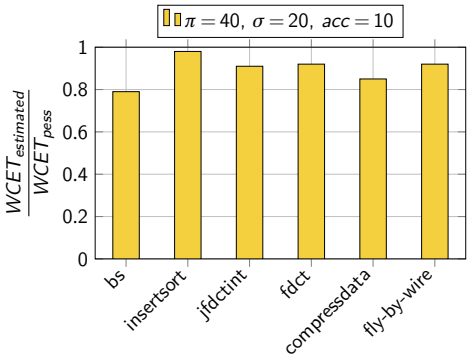

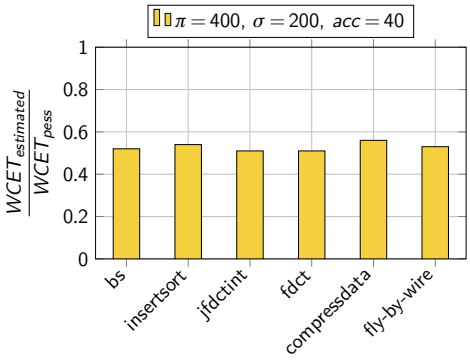

Best-case of gain: All requests are within the TDMA slots
Worst-case of gain: All requests have worst-case delay

Outline: TDMA Bus Timing Analysis

Summary of Part I

- SMT encodings for TDMA access
- Feasible path analysis combined with the WCET computation
- Comparison between different encodings
- Validation with small but relevant benchmarks

Find in the manuscript:

- Linearization of SMT encoding (modulo operators)
- Other possible SMT encodings

Future Work of Part I

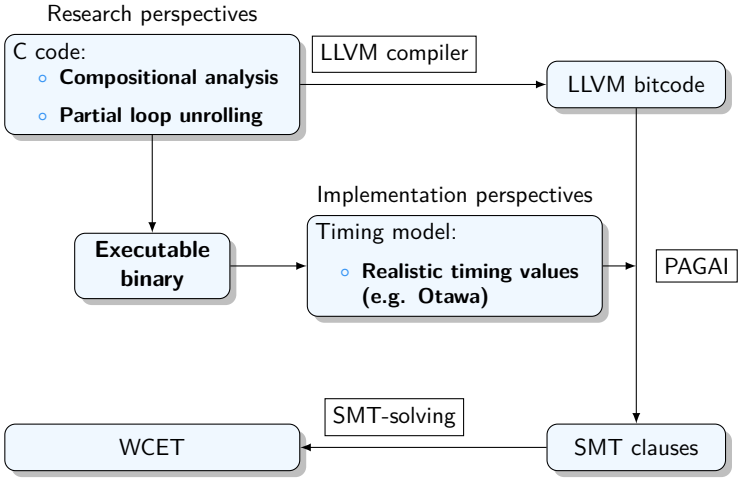

Future Work of Part I

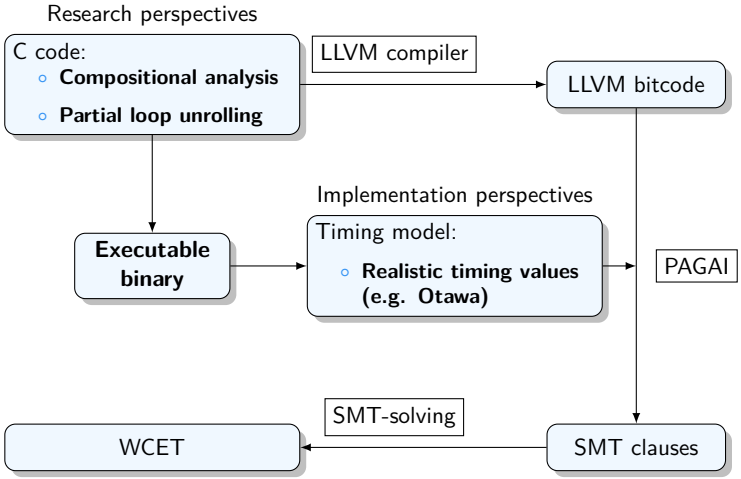

$\rightsquigarrow$ **SMT is an interesting research direction for WCET Analysis**

From TDMA to Other Arbitration Policy

Full isolation with TDMA

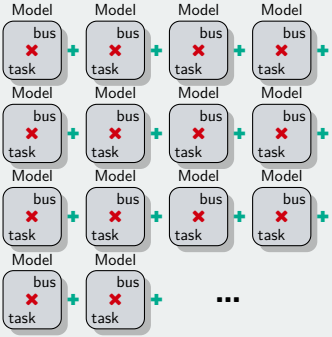

Linear complexity
in the number of tasks

From TDMA to Other Arbitration Policy

Full isolation with TDMA

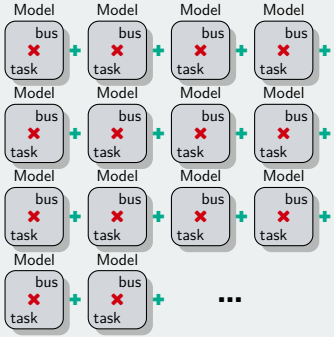

Linear complexity
in the number of tasks

No isolation

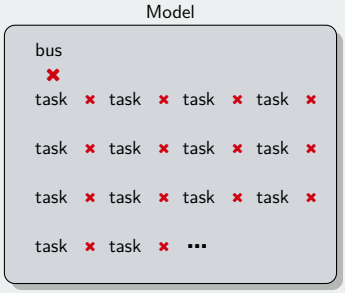

States explosion!

Analysis of Large Multi and Many-Cores

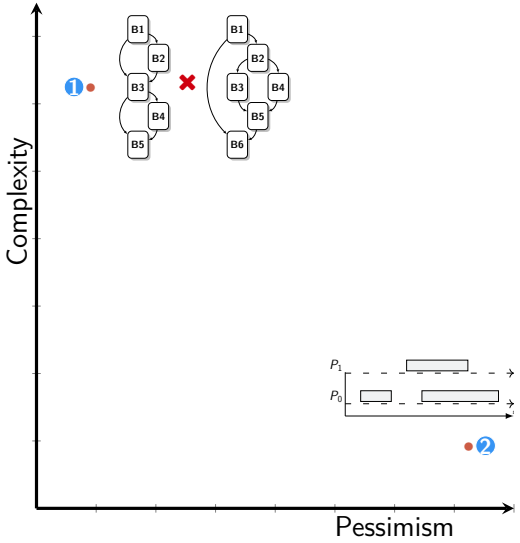

1. Exact analysis
2. Account for any interference **globally** during the task's execution

Analysis of Large Multi and Many-Cores

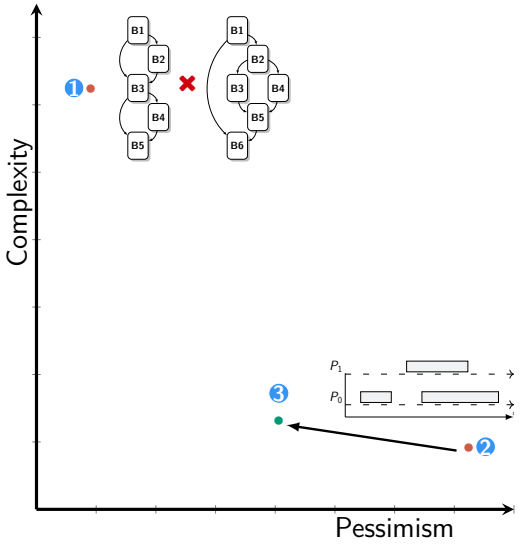

1. Exact analysis
2. Account for any interference **globally** during the task's execution

3. Exploit any information about:
   - The target architecture Kalray MPPA2
     - Reduce the interference
     - Model precisely the shared resources
   - The target application model Synchronous Data Flow

Many-Core Response Time Analysis

Outline: Many-Core Response Time Analysis

Outline: Many-Core Response Time Analysis

Implementation Choices: SDF on Multi/Many-Cores

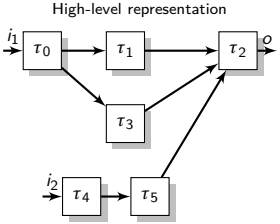

Multi/Many-Core

code generation

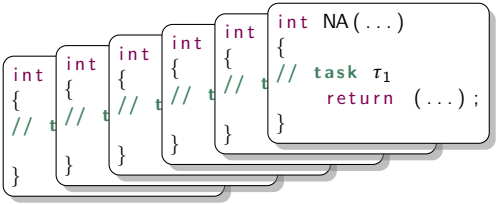

Static, time triggered, non-preemptive scheduling

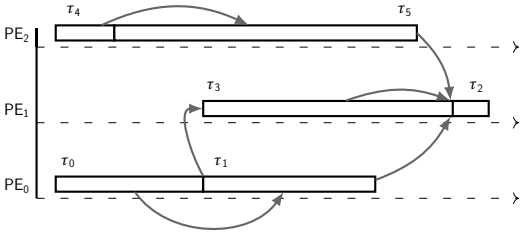

Implementation Choices: SDF on Multi/Many-Cores

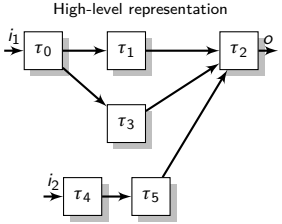

Multi/Many-Core
code generation

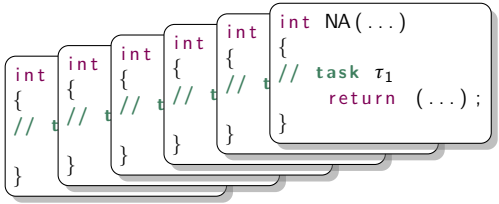

Static, time triggered, non-preemptive scheduling

✓ Respect the dependency constraints

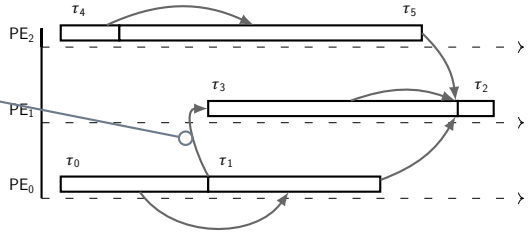

Implementation Choices: SDF on Multi/Many-Cores

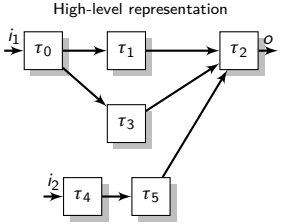

Multi/Many-Core
code generation

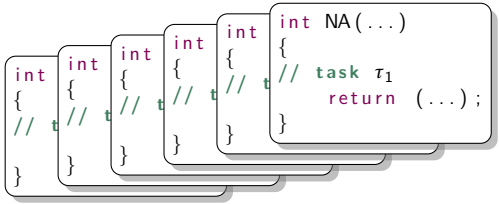

Static, time triggered, non-preemptive scheduling

✓ Respect the dependency constraints

✓ account for precise upper bounds on the interference

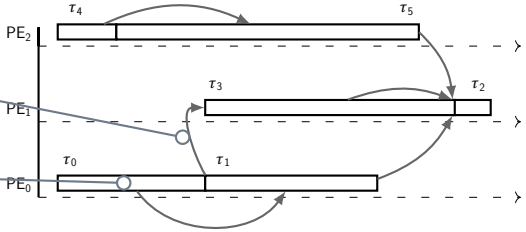

Model of the Application

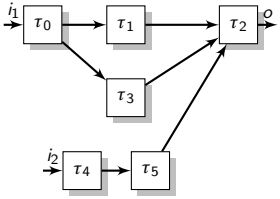

An execution instance is:

- Direct Acyclic Task Graph
- Mono-rate (or at least harmonic rates)
- Fixed mapping and execution order

Model of the Application

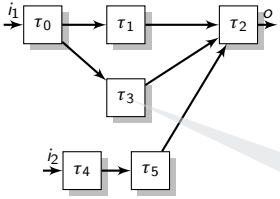

An execution instance is:

- Direct Acyclic Task Graph
- Mono-rate (or at least harmonic rates)
- Fixed mapping and execution order

Each task τ_i:

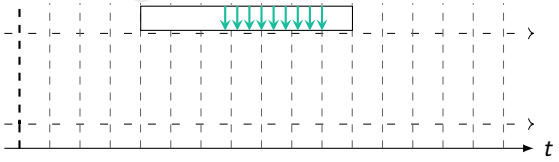

Model of the Application

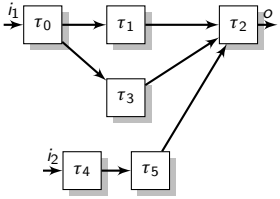

An execution instance is:

- Direct Acyclic Task Graph
- Mono-rate (or at least harmonic rates)
- Fixed mapping and execution order

Each task τ_i:

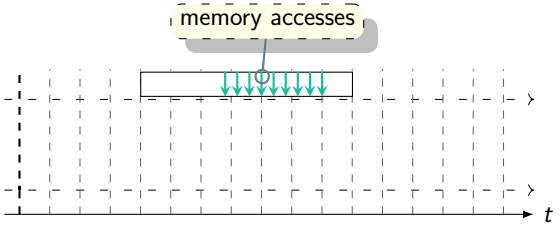

Model of the Application

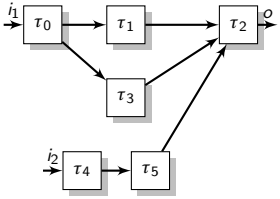

An execution instance is:

- Direct Acyclic Task Graph
- Mono-rate (or at least harmonic rates)
- Fixed mapping and execution order

Each task τ_i:

- **Release date** (rel_i). **Response time** (R_i)

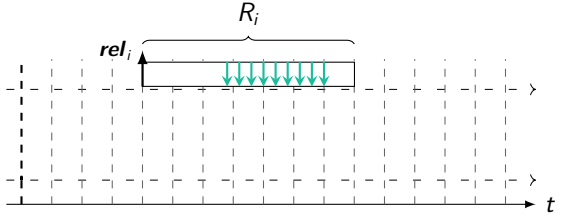

Model of the Application

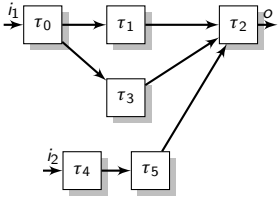

An execution instance is:

- Direct Acyclic Task Graph
- Mono-rate (or at least harmonic rates)
- Fixed mapping and execution order

Each task τ_i:

- **Release date** (rel_i). **Response time** (R_i)

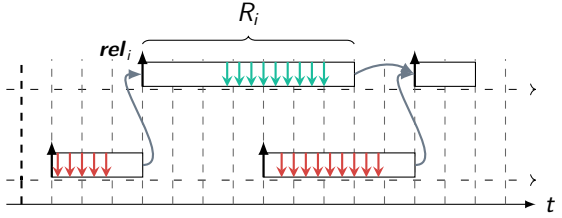

Model of the Application

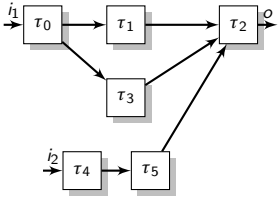

An execution instance is:

- Direct Acyclic Task Graph
- Mono-rate (or at least harmonic rates)
- Fixed mapping and execution order

Each task τ_i:

- **Release date** (rel_i). **Response time** (R_i)

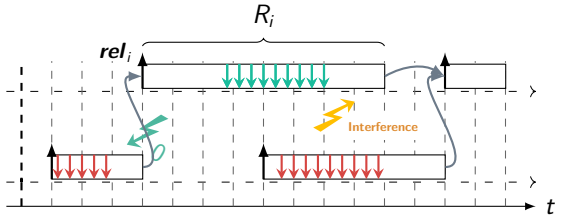

Model of the Application

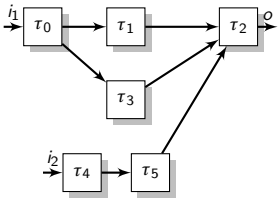

An execution instance is:

- Direct Acyclic Task Graph
- Mono-rate (or at least harmonic rates)
- Fixed mapping and execution order

Each task τ_i:

- **Release date** (rel_i). **Response time** (R_i)

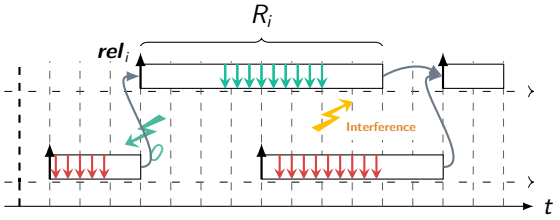

Static Non-Preemptive Scheduling

- Find R_i including interference
- Find rel_i respecting dependencies

Outline: Many-Core Response Time Analysis

Response Time Analysis

Existing work [Altmeyer et al., 2015]

$$\forall i: \; R_i \; = \; WCET_i \; + \; I^{BUS}(R_i) \; + \ldots$$

For each task i:

- Response Time
- WCET in isolation
- Bus Interference

Independent tasks

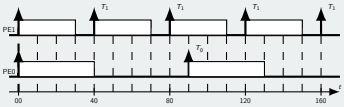

Response Time Analysis

Existing work [Altmeyer et al., 2015]

$$\forall i:\ R_i = WCET_i + I^{BUS}(R_i) + \ldots$$

For each task i:

- Response Time
- WCET in isolation
- Bus Interference

Independent tasks

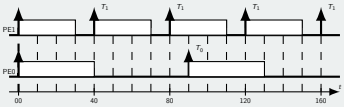

Contribution (in RTNS 2016)

$$\forall i:\ R_i = \mathrm{WCET}_i + I^{BUS}(R_i,\ \Theta) + \ldots$$

- WCET in isolation
- Set of release dates of all tasks
- Bounded interference

Dependent tasks

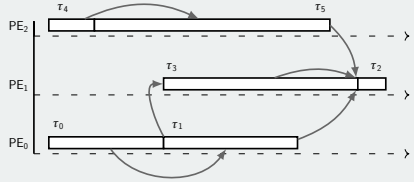

Response Time Analysis

Existing work [Altmeyer et al., 2015]

$$\forall i:\ R_i = WCET_i + I^{BUS}(R_i) + \ldots$$

For each task i:

- Response Time
- WCET in isolation
- Bus Interference

Independent tasks

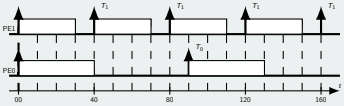

Contribution (in RTNS 2016)

$$\forall i:\ R_i = \mathrm{WCET}_i + I^{BUS}(R_i, \Theta) + \ldots$$

- WCET in isolation
- Set of release dates of all tasks
- Bounded interference

Dependent tasks

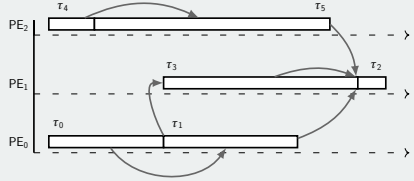

Recursive formula ⇒ iterative algorithm.

Response Time Analysis with Dependencies

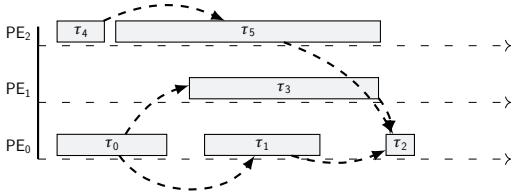

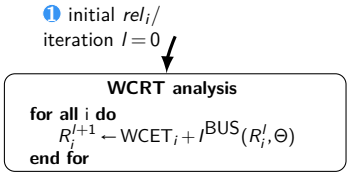

1. Start with initial release dates

Response Time Analysis with Dependencies

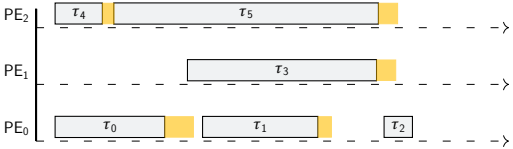

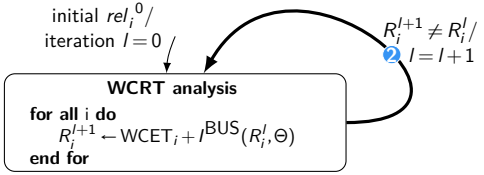

1. Start with initial release dates
2. Compute response times

...

Response Time Analysis with Dependencies

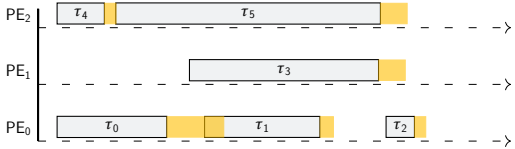

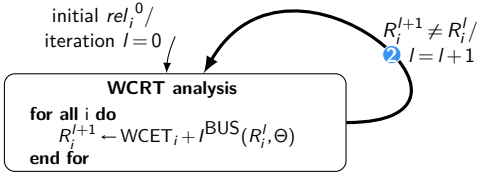

1. Start with initial release dates
2. Compute response times

... ...

Response Time Analysis with Dependencies

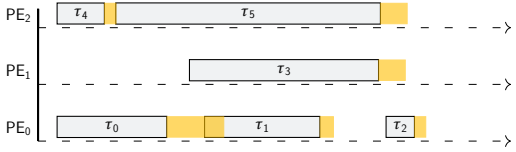

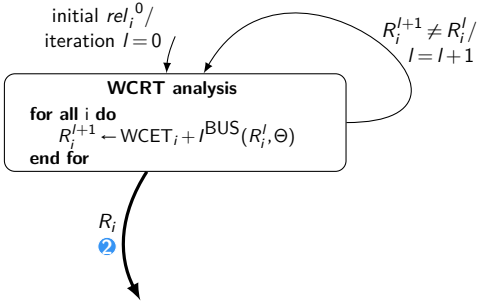

1. Start with initial release dates
2. Compute response times

   ... ... a fixed-point is reached!

Response Time Analysis with Dependencies

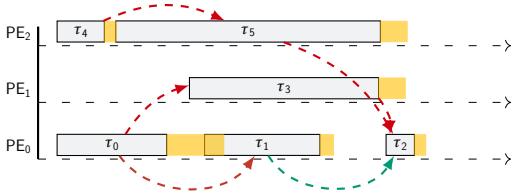

1. Start with initial release dates
2. Compute response times

   ... ... a fixed-point is reached!
3. Update the release dates

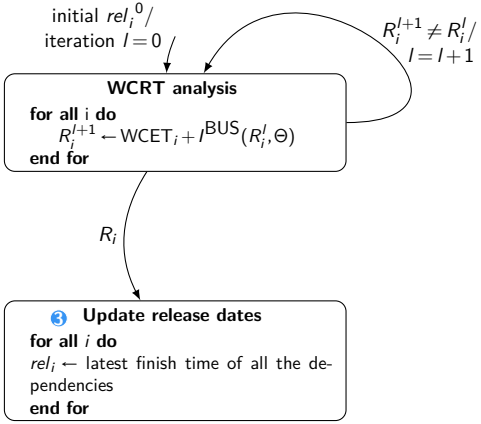

Response Time Analysis with Dependencies

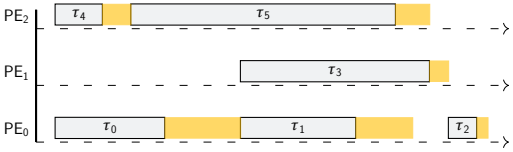

1. Start with initial release dates
2. Compute response times

   ... ... a fixed-point is reached!
3. Update the release dates
4. Repeat

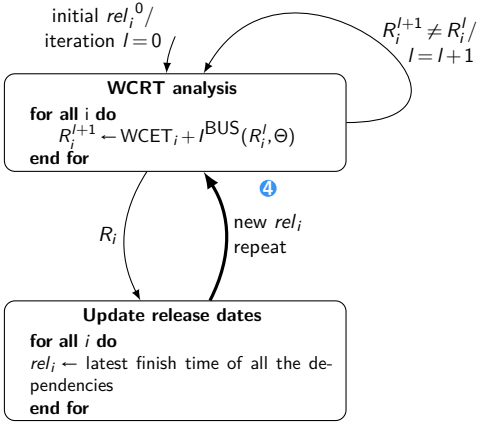

Response Time Analysis with Dependencies

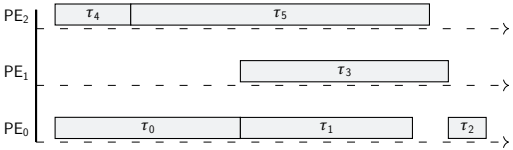

1. Start with initial release dates
2. Compute response times
   ... ... a fixed-point is reached!
3. Update the release dates
4. Repeat until no release date changes
   (another fixed-point iteration).

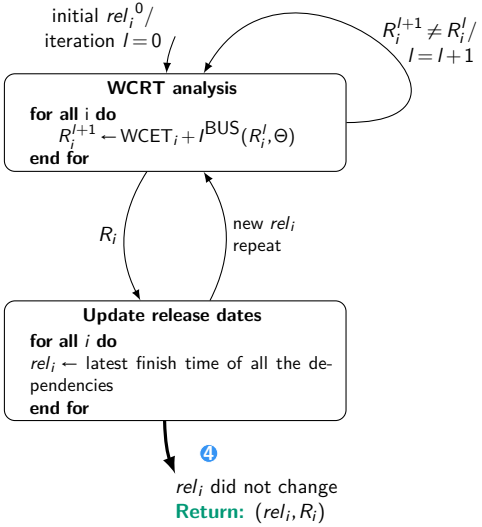

Convergence Toward a Fixed-point

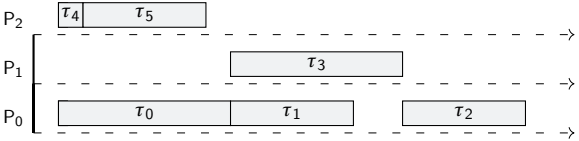

- Convergence of the 1^{st} fixed-point iteration:

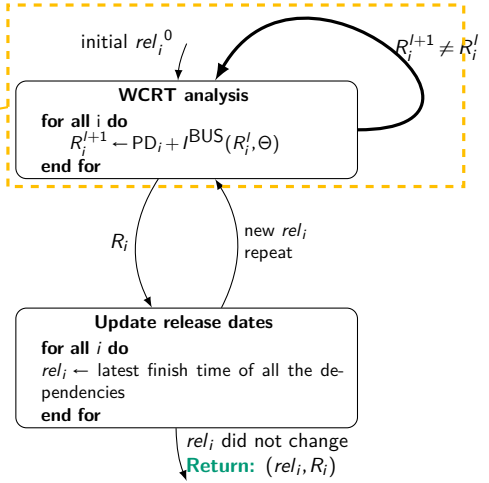

Convergence Toward a Fixed-point

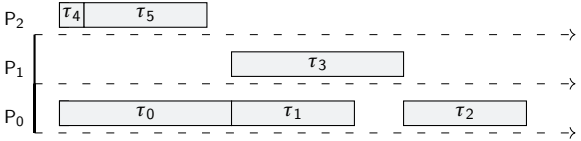

- Convergence of the 1^{st} fixed-point iteration:
  - Monotonic and bounded ✓

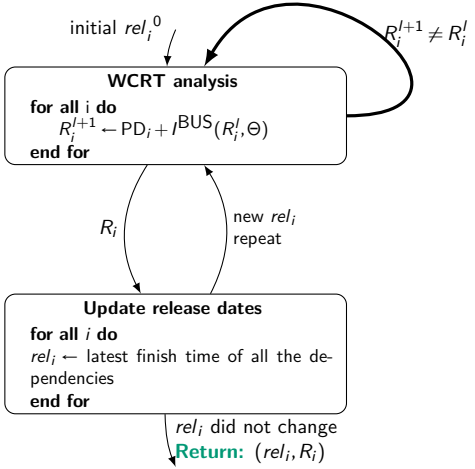

Convergence Toward a Fixed-point

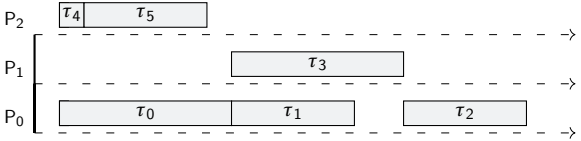

- Convergence of the 1^{st} fixed-point iteration:
  - Monotonic and bounded ✓
- Convergence of the 2^{nd} fixed-point iteration:

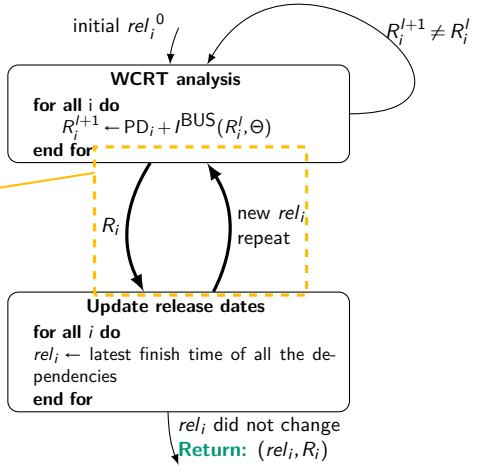

Convergence Toward a Fixed-point

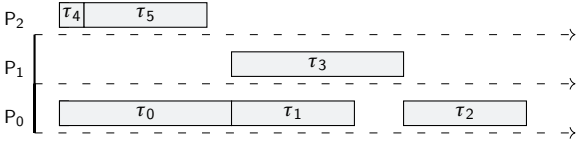

- Convergence of the 1^{st} fixed-point iteration:
  - Monotonic and bounded ✓
- Convergence of the 2^{nd} fixed-point iteration:
  - No monotonicity: R_i and rel_i may grow or shrink at each iteration. ?

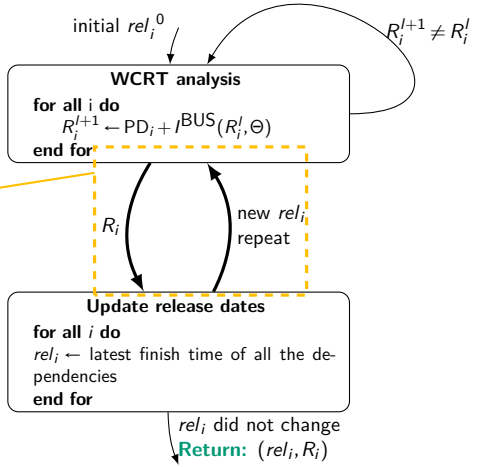

Convergence Toward a Fixed-point

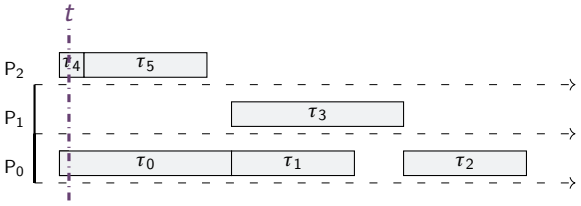

- Convergence of the 1^{st} fixed-point iteration:
  - Monotonic and bounded ✓
- Convergence of the 2^{nd} fixed-point iteration:
  - No monotonicity: R_i and rel_i may grow or shrink at each iteration. ?

Theorem

At each iteration, at least one task finds its final release date.

Full proof in the manuscript.

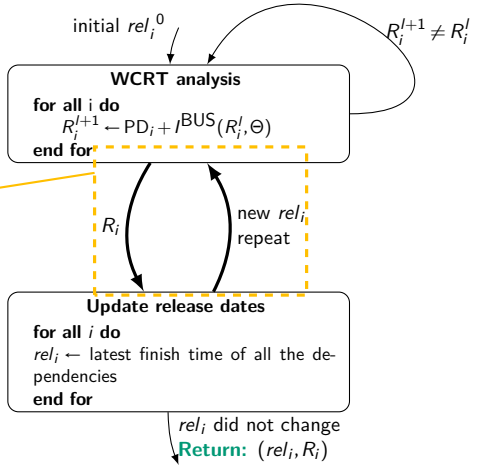

Convergence Toward a Fixed-point

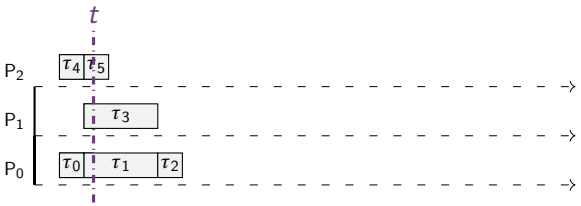

- Convergence of the 1^{st} fixed-point iteration:
  - Monotonic and bounded ✓
- Convergence of the 2^{nd} fixed-point iteration:
  - No monotonicity: R_i and rel_i may grow or shrink at each iteration. ?

Theorem

At each iteration, at least one task finds its final release date.

Full proof in the manuscript.

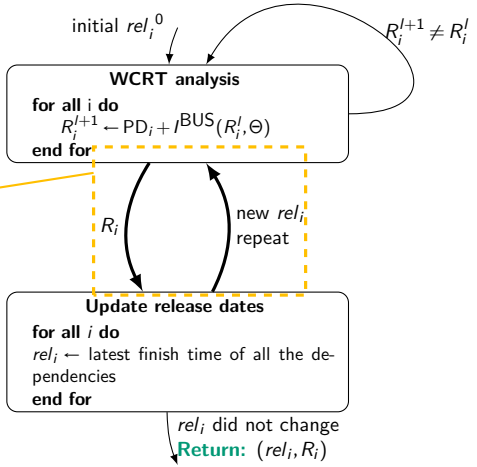

Convergence Toward a Fixed-point

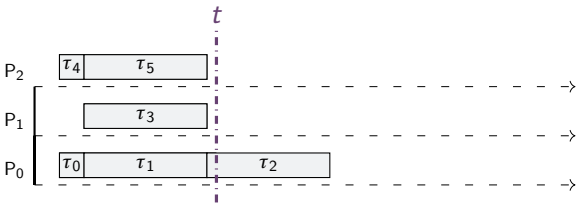

- Convergence of the 1^{st} fixed-point iteration:
  - Monotonic and bounded ✓
- Convergence of the 2^{nd} fixed-point iteration:
  - No monotonicity: R_i and rel_i may grow or shrink at each iteration. ?

Theorem

At each iteration, at least one task finds its final release date.

Full proof in the manuscript.

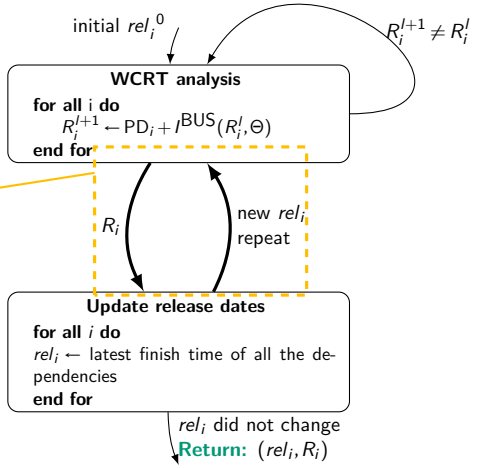

Convergence Toward a Fixed-point

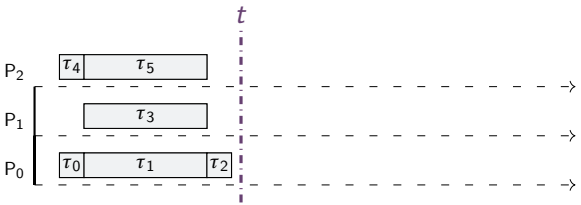

- Convergence of the 1^{st} fixed-point iteration:
  - Monotonic and bounded ✓
- Convergence of the 2^{nd} fixed-point iteration:
  - No monotonicity: R_i and rel_i may grow or shrink at each iteration. ?

Theorem

At each iteration, at least one task finds its final release date.

Full proof in the manuscript.

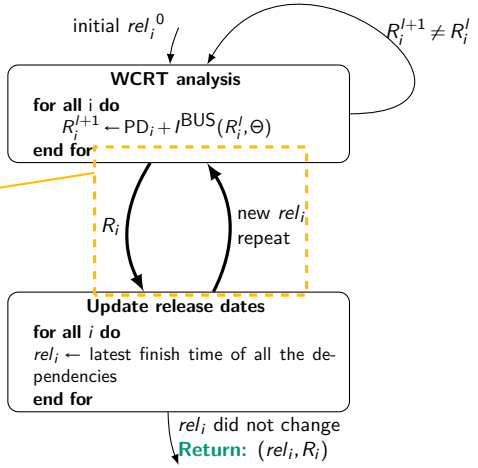

Outline: Many-Core Response Time Analysis

Target Many-Core: Kalray MPPA2

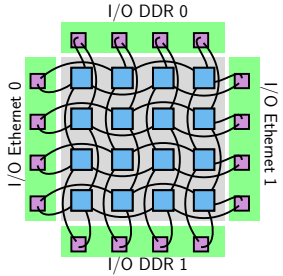

- Kalray MPPA2 (codenamed Bostan)
- 16 compute clusters + 4 I/O clusters
- Dual NoC

Target Many-Core: Kalray MPPA2

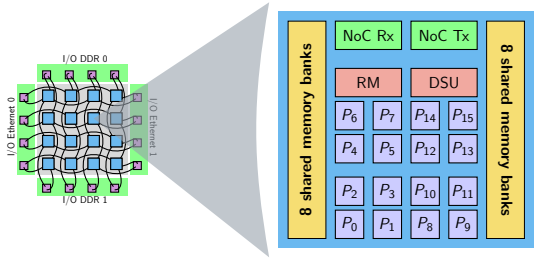

Per cluster:

- 16 cores + 1 Resource Manager
- NoC Tx, NoC Rx, Debug Unit
- 16 shared memory banks (total: 2 MB)

Target Many-Core: Kalray MPPA2

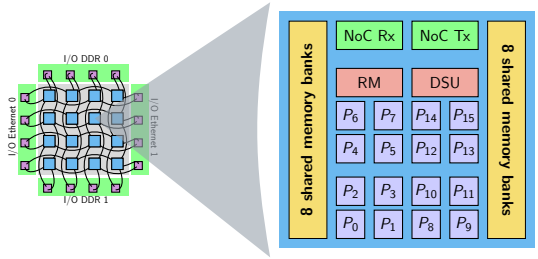

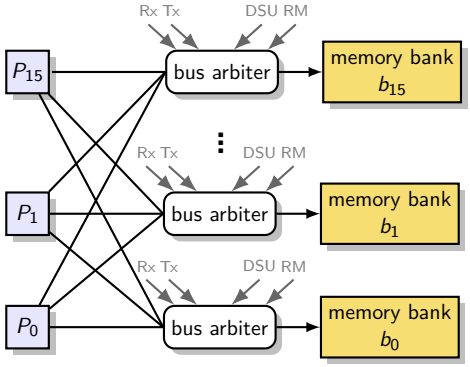

Per cluster:

- 16 cores + 1 Resource Manager
- NoC Tx, NoC Rx, Debug Unit
- 16 shared memory banks (total: 2 MB)
- 1 bus arbiter per memory bank

Target Many-Core: Kalray MPPA2

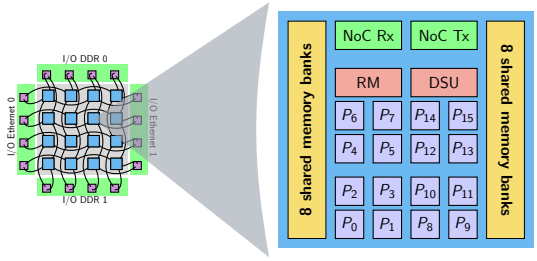

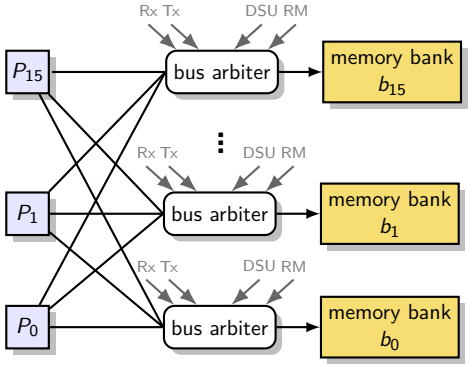

Per cluster:

- 16 cores + 1 Resource Manager
- NoC Tx, NoC Rx, Debug Unit
- 16 shared memory banks (total: 2 MB)
- 1 bus arbiter per memory bank

- Possible spatial isolation
  assigning memory banks to cores
- Task execution model:
  - execute in a "local" bank
  - write to a "remote" bank
- Interference from communications

Model of the MPPA2 Bus

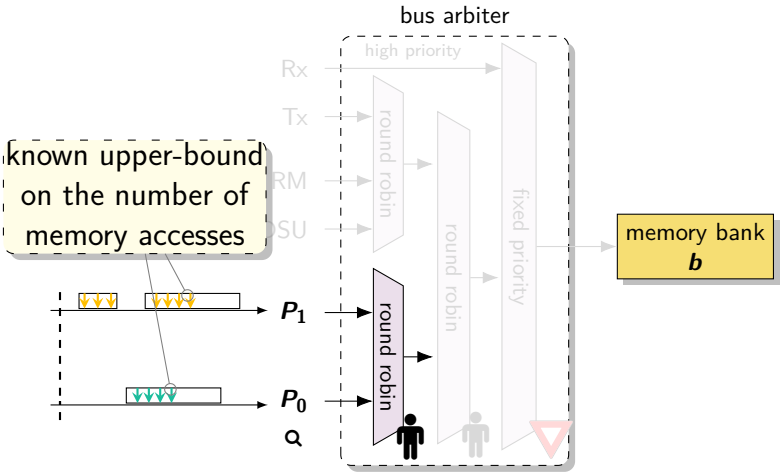

Shared Bus Model

$$I^{BUS}(R_i,\Theta) = \sum_{b\in\mathscr{B}} I_b^{BUS}(R_i,\Theta)$$

where $\mathscr{B}$: a set of memory banks

Bus interference $I_b^{\mathrm{BUS}}(R_i,\Theta)$:

Model of the MPPA2 Bus

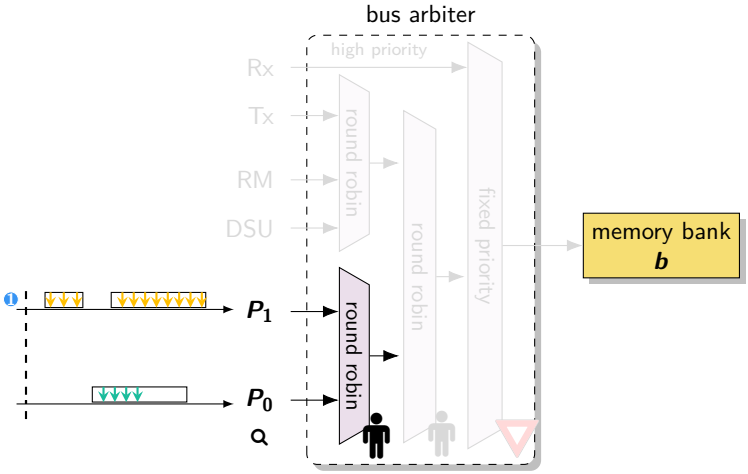

Shared Bus Model

$$I^{BUS}(R_i,\Theta) = \sum_{b\in\mathscr{B}} I_b^{BUS}(R_i,\Theta)$$

where $\mathscr{B}$: a set of memory banks

Bus interference $I_b^{\mathrm{BUS}}(R_i,\Theta)$:

1. $\lightning(P_0)$=↓↓↓...↓

Model of the MPPA2 Bus

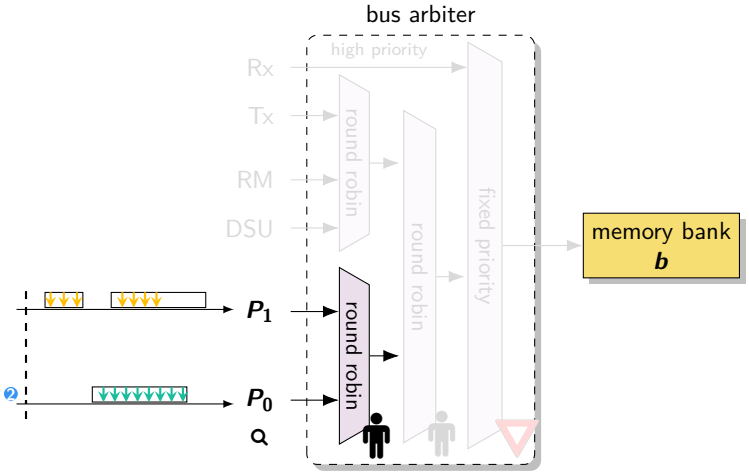

Shared Bus Model

$$I^{BUS}(R_i,\Theta) = \sum_{b\in\mathscr{B}} I_b^{BUS}(R_i,\Theta)$$

where $\mathscr{B}$: a set of memory banks

Bus interference $I_b^{\mathrm{BUS}}(R_i,\Theta)$:

1. $\lightning(P_0)$=↓↓↓...↓
2. $\lightning(P_1)$=↓↓↓...↓

Model of the MPPA2 Bus

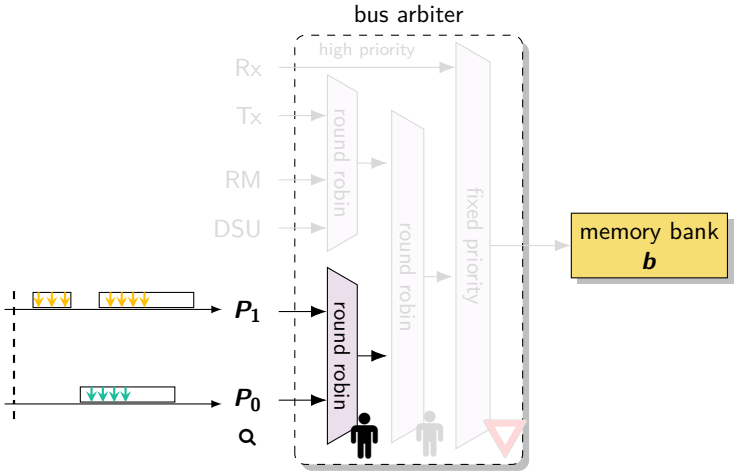

Shared Bus Model

$$I^{BUS}(R_i,\Theta) = \sum_{b\in\mathscr{B}} I_b^{BUS}(R_i,\Theta)$$

where $\mathscr{B}$: a set of memory banks

Bus interference $I_b^{\mathrm{BUS}}(R_i,\Theta)$:

1. $\lightning(P_0)$=↓↓↓...↓
2. $\lightning(P_1)$=↓↓↓...↓

Model of the MPPA2 Bus

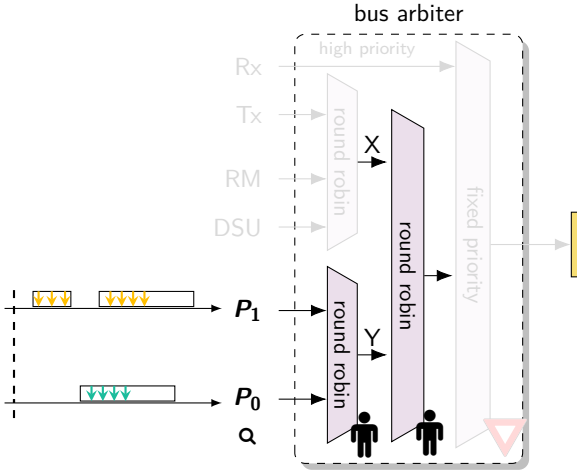

Shared Bus Model

$$I^{BUS}(R_i,\Theta) = \sum_{b\in\mathscr{B}} I_b^{BUS}(R_i,\Theta)$$

where $\mathscr{B}$: a set of memory banks

Bus interference $I_b^{\mathrm{BUS}}(R_i,\Theta)$:

1. $\lightning(P_0)$=↓↓↓...↓
2. $\lightning(P_1)$=↓↓↓...↓

$$\min(\lightning(P_0), \lightning(P_1)) \; + \min(\lightning(Y), \lightning(\mathrm{X}))$$

Model of the MPPA2 Bus

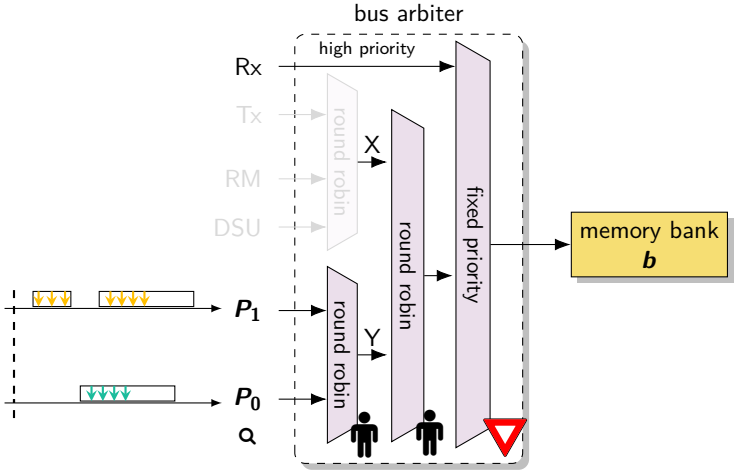

Shared Bus Model

$$I^{BUS}(R_i,\Theta) = \sum_{b \in \mathscr{B}} I_b^{BUS}(R_i,\Theta)$$

where $\mathscr{B}$: a set of memory banks

Bus interference $I_b^{\mathrm{BUS}}(R_i,\Theta)$:

1. $\lightning(P_0)$=↓↓↓…↓
2. $\lightning(P_1)$=↓↓↓…↓

$$\min(\lightning(P_0), \lightning(P_1)) \; + \min(\lightning(Y), \lightning(\mathrm{X})\;) + \lightning(\mathrm{Rx})$$

Model of the MPPA2 Bus

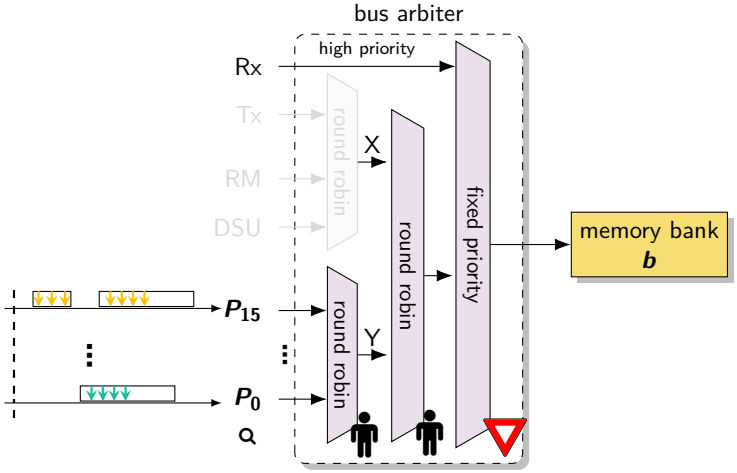

Shared Bus Model

$$I^{BUS}(R_i,\Theta) = \sum_{b\in\mathscr{B}} I_b^{BUS}(R_i,\Theta)$$

where $\mathscr{B}$: a set of memory banks

Bus interference $I_b^{\mathrm{BUS}}(R_i,\Theta)$:

1. $\lightning(P_0)$=↓↓↓…↓
2. $\lightning(P_1)$=↓↓↓…↓

$$\sum_{1\le i\le 15} \min(\lightning(P_0), \lightning(P_i)) \; + \min(\lightning(Y), \lightning(\mathrm{X})\;) + \lightning(\mathrm{Rx})$$

Model of the MPPA2 Bus

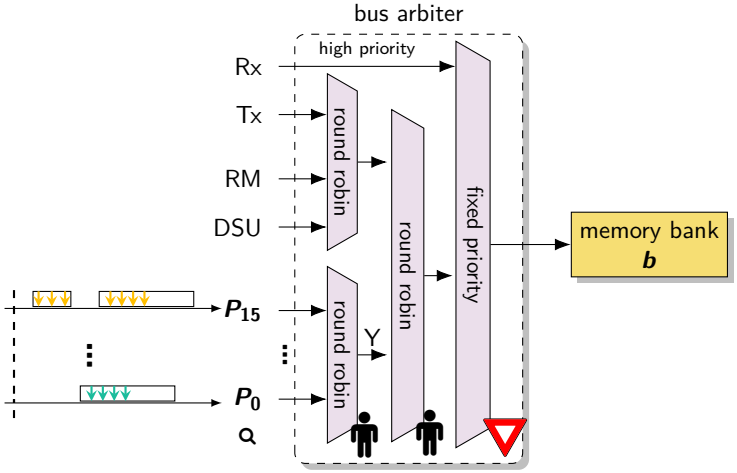

Shared Bus Model

$$I^{BUS}(R_i,\Theta) = \sum_{b\in\mathscr{B}} I_b^{BUS}(R_i,\Theta)$$

where $\mathscr{B}$: a set of memory banks

Bus interference $I_b^{\mathrm{BUS}}(R_i,\Theta)$:

1. $\lightning(P_0)$=↓↓↓…↓
2. $\lightning(P_1)$=↓↓↓…↓

$$\sum_{1\le i\le 15} \min(\lightning(P_0), \lightning(P_i)) \;+\; \min(\lightning(Y), \lightning(\mathrm{X})\;) + \lightning(\mathrm{Rx})$$

$X = \sum\{Tx, RM, DSU\}$

Outline: Many-Core Response Time Analysis

Case Study: ROSACE, a Flight Management System Controller [2]

High level representation

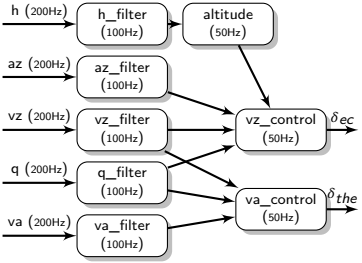

[2] Pagetti et al., 2014

Case Study: ROSACE, a Flight Management System Controller [2]

High level representation

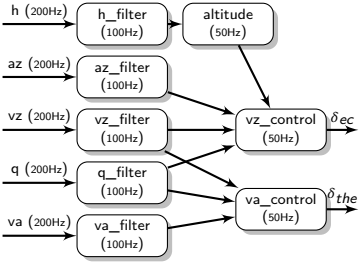

Unrolled execution

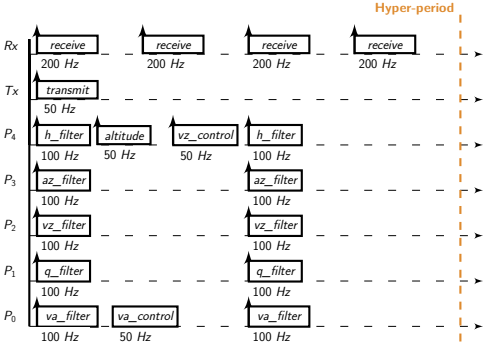

[2] Pagetti et al., 2014

Evaluation: ROSACE Case Study

Function	WCET (cycles)	Memory accesses
altitude	275	22
az_filter	274	22
h_filter	326	24
q_filter	338	24
va_control	303	24
va_filter	301	23
vz_control	320	25
vz_filter	334	25

- Values obtained from measurements
- Memory accesses from data and instruction cache misses + communications
- Moreover:
  - *NoC Rx*: writes 5 words
  - *NoC Tx*: reads 2 words

Evaluation: ROSACE Case Study

Function	WCET (cycles)	Memory accesses
altitude	275	22
az_filter	274	22
h_filter	326	24
q_filter	338	24
va_control	303	24
va_filter	301	23
vz_control	320	25
vz_filter	334	25

- Values obtained from measurements
- Memory accesses from data and instruction cache misses + communications
- Moreover:
  - *NoC Rx*: writes 5 words
  - *NoC Tx*: reads 2 words

Experiments: Find the smallest schedulable hyper-period

Evaluation: Experiments

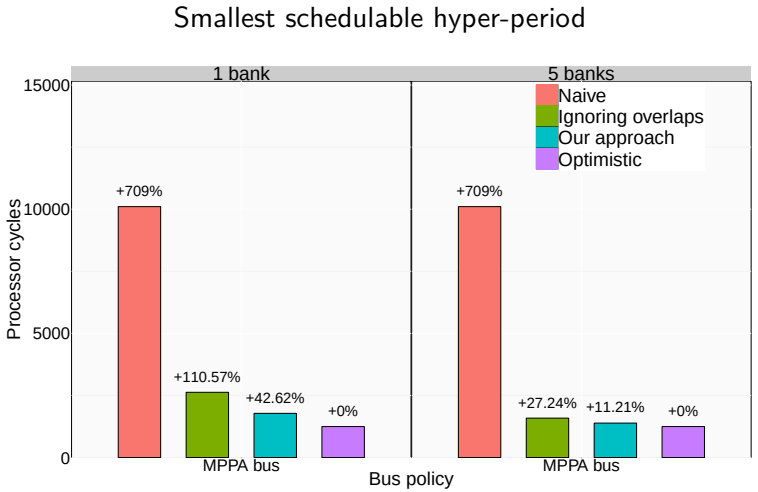

Evaluation: Experiments

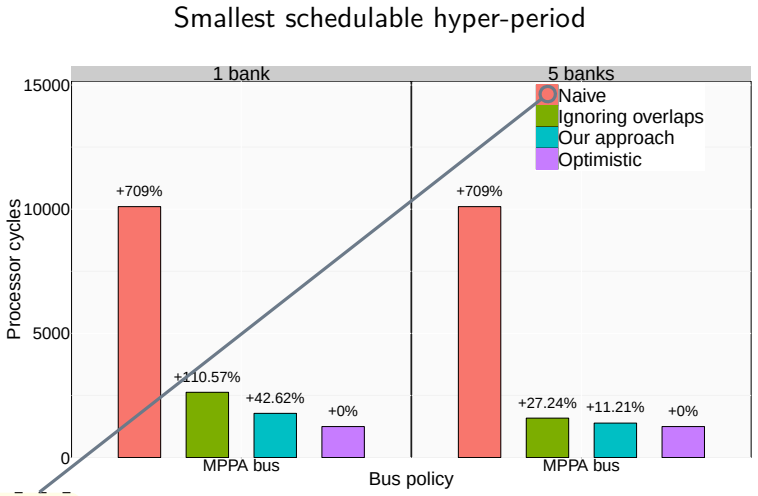

Naive:
- All accesses interfere
- Rx bounded by 1 access per bank

Evaluation: Experiments

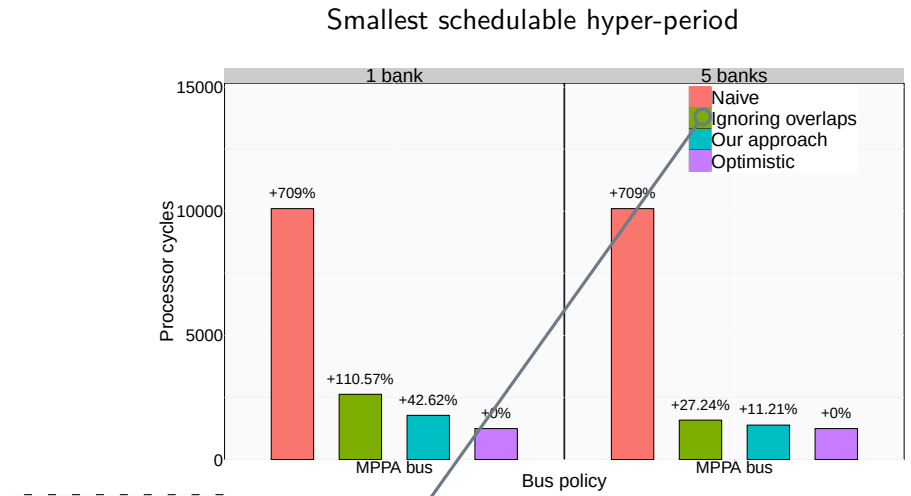

Evaluation: Experiments

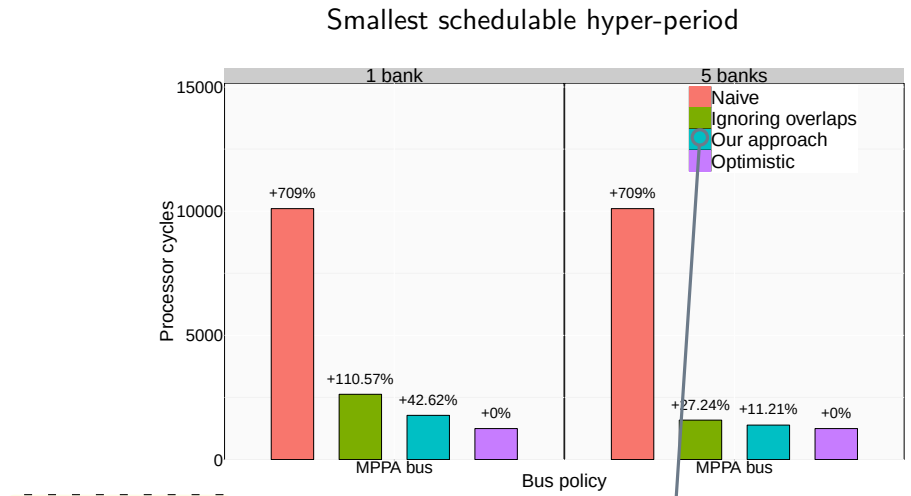

Evaluation: Experiments

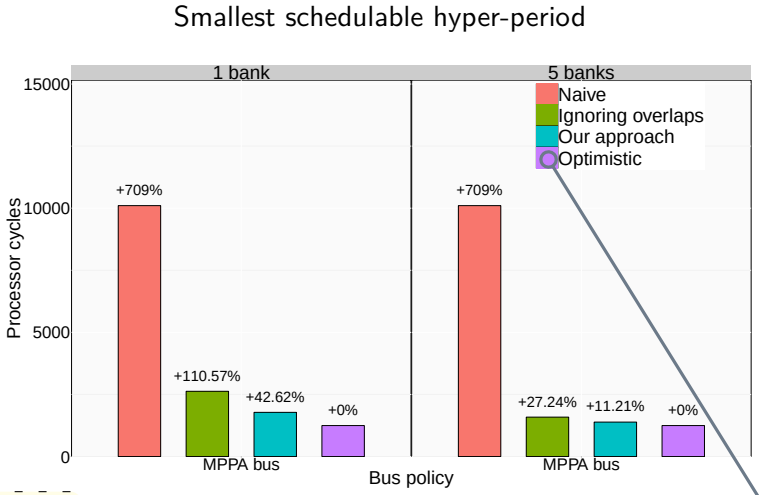

Naive:
- All accesses interfere
- Rx bounded by 1 access per bank

Ignoring overlap:
We don't use the release dates

Our approach:
We use the release dates

Optimistic:
No bus interference

Evaluation: Experiments

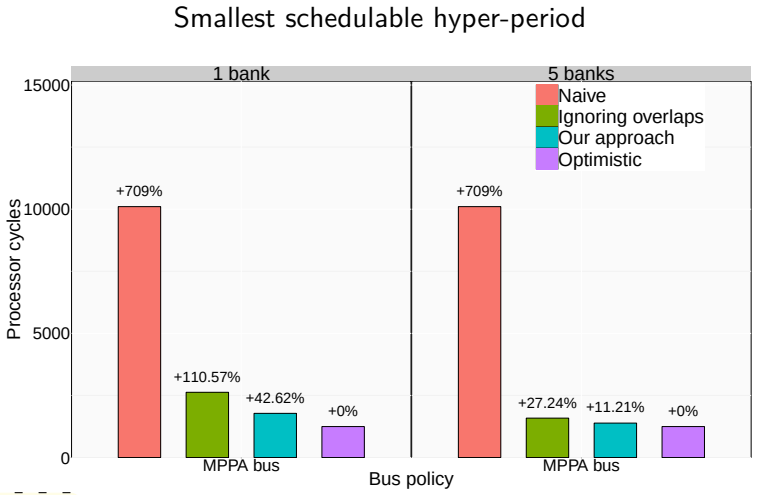

Naive:
- All accesses interfere
- Rx bounded by 1 access per bank

Ignoring overlap:
We don't use the release dates

Our approach:
We use the release dates

Optimistic:
No bus interference

Evaluation: CAPACITES Project

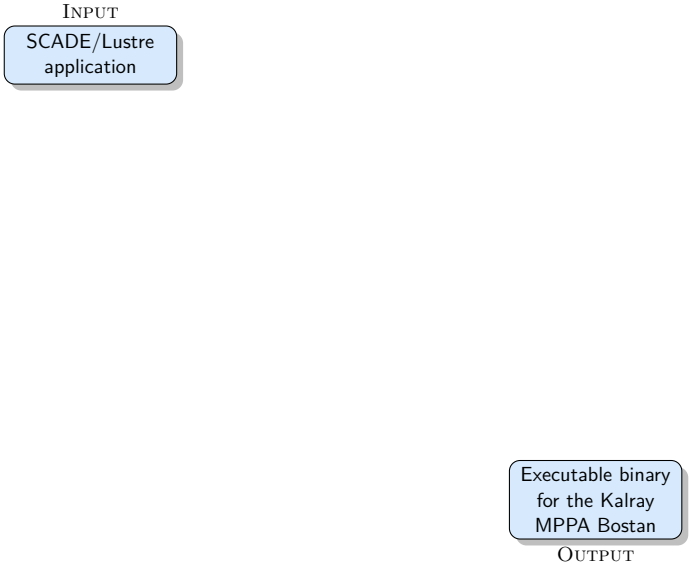

Evaluation: CAPACITES Project

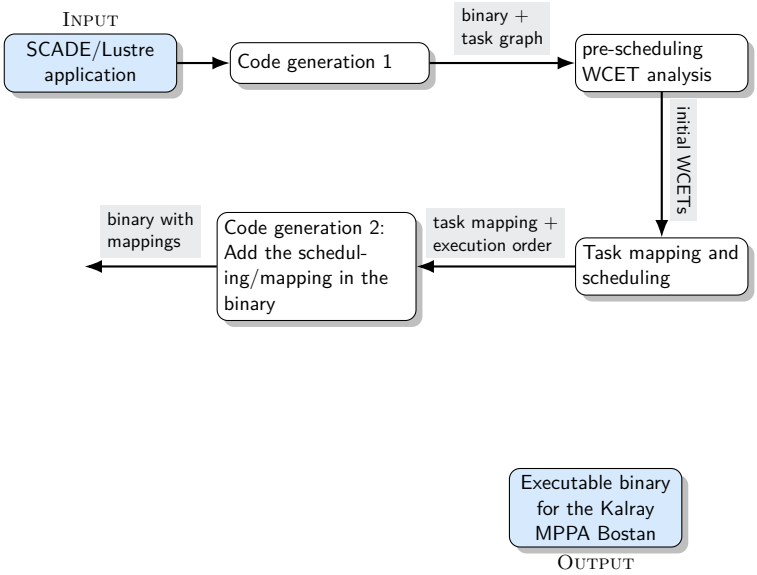

Evaluation: CAPACITES Project

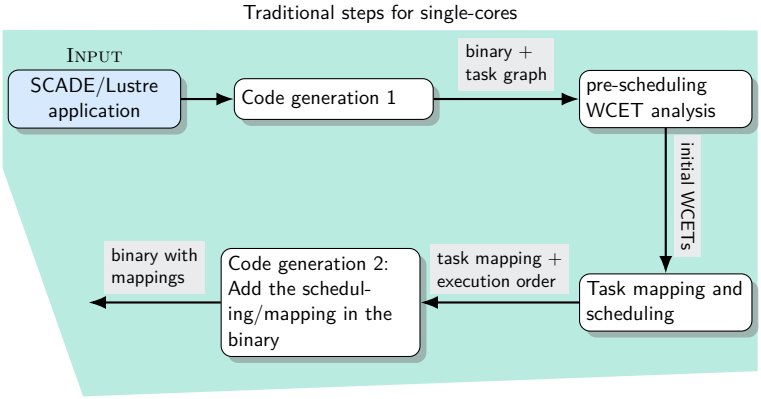

Evaluation: CAPACITES Project

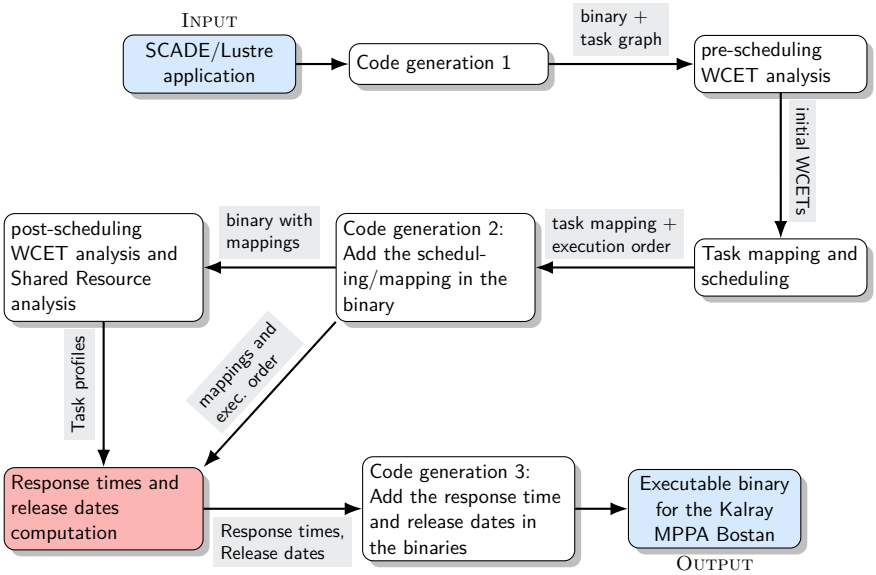

Outline: Many-Core Response Time Analysis

Summary of Part II

- A response time analysis of synchronous data flow programs on the Kalray MPPA2

Summary of Part II

- A response time analysis of synchronous data flow programs on the Kalray MPPA2
- Given:
  - Task profiles: WCET in isolation and number of accesses
  - Mapping of Tasks
  - Execution Order

Summary of Part II

- A response time analysis of synchronous data flow programs on the Kalray MPPA2
- Given:
  - Task profiles: WCET in isolation and number of accesses
  - Mapping of Tasks
  - Execution Order
- We compute:
  - Tight response times taking into account the interference
  - Release dates respecting the dependency constraints

Summary of Part II

- A response time analysis of synchronous data flow programs on the Kalray MPPA2
- Given:
  - Task profiles: WCET in isolation and number of accesses
  - Mapping of Tasks
  - Execution Order
- We compute:
  - Tight response times taking into account the interference
  - Release dates respecting the dependency constraints

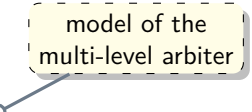

Summary of Part II

- A response time analysis of synchronous data flow programs on the Kalray MPPA2
- Given:
  - Task profiles: WCET in isolation and number of accesses
  - Mapping of Tasks
  - Execution Order
- We compute:
  - Tight response times taking into account the interference
  - Release dates respecting the dependency constraints

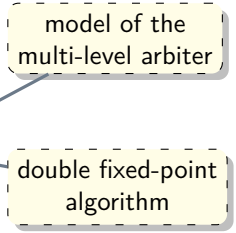

Summary of Part II

- A response time analysis of synchronous data flow programs on the Kalray MPPA2
- Given:
  - Task profiles: WCET in isolation and number of accesses
  - Mapping of Tasks
  - Execution Order
- We compute:
  - Tight response times taking into account the interference
  - Release dates respecting the dependency constraints

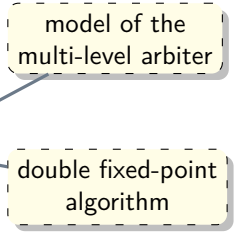

- **Find in the manuscript:**
  - Execution phases: execution phase + communication phase
  - Support of: accesses pipelining, blocking and non-blocking accesses, bursts of accesses
  - More experiments with randomly generated task graphs

Future Work of Part II

- Model of the Resource Manager
- Analysis with a Real-Time Operating System

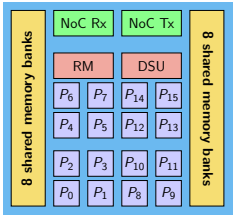

Future Work of Part II

- Model of the Resource Manager
- Analysis with a Real-Time Operating System

tighter estimation of context switches and other interrupts

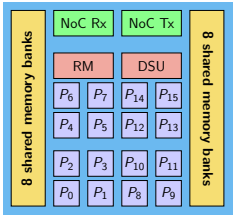

Future Work of Part II

tighter estimation of context switches and other interrupts

- Model of the Resource Manager
- Analysis with a Real-Time Operating System

- Model of the NoC accesses.

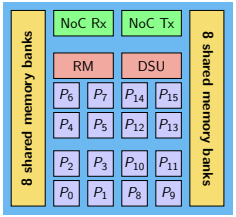

Future Work of Part II

- Model of the Resource Manager
- Analysis with a Real-Time Operating System

tighter estimation of context switches and other interrupts

- Model of the NoC accesses.
- Comparison between estimated and measured response times

use the output of any NoC analysis

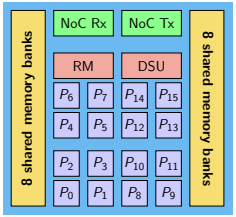

Conclusion

Conclusion

- Multi/Many-cores in Real-Time Systems
  - Full isolation: for example TDMA
  - Bounded interference

Conclusion

- Multi/Many-cores in Real-Time Systems
  - Full isolation: for example TDMA
  - Bounded interference
- Precise modeling of hardware components

Conclusion

- Multi/Many-cores in Real-Time Systems
  - Full isolation: for example TDMA
  - Bounded interference
- Precise modeling of hardware components
- Research directions for multi-core analysis:

Multi-core scheduling and timing analysis framework (in RTSOPS 2016)

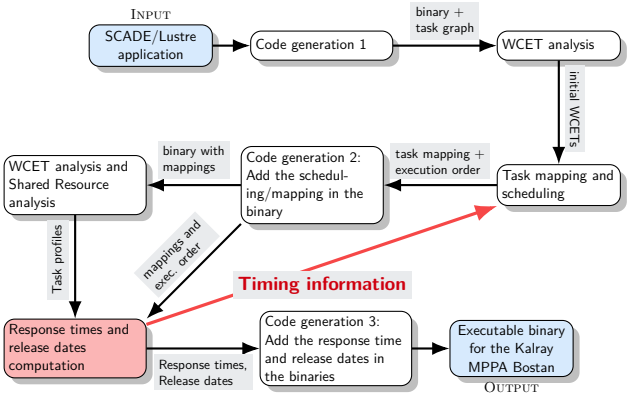

Many-Core Timing Analysis of Real-Time Systems
and its application to an industrial processor

Hamza Rihani

Université Grenoble Alpes / Verimag

Publications:

Rihani, Hamza et al. (2015). "WCET analysis in shared resources real-time systems with TDMA buses". In: *Proceedings of the 23rd International Conference on Real-Time Networks and Systems*.

Rihani, Hamza, Claire Maiza, and Matthieu Moy (2016a). "Efficient Execution of Dependent Tasks on Many-Core Processors". In: *RTSOPS 2016*. 7th International Real-Time Scheduling Open Problems Seminar. Toulouse, France.

Rihani, Hamza et al. (2016b). "Response Time Analysis of Synchronous Data Flow Programs on a Many-Core Processor". In: *Proceedings of the 24th International Conference on Real-Time Networks and Systems (RTNS)*.

This work is funded by grant CAPACITES (PIA-FSN2 n°P3425-146798) from the French *Ministère de l'économie, des finances et de l'industrie*.

References I


Altmeyer, Sebastian et al. (2015). "A Generic and Compositional Framework for Multicore Response Time Analysis". In: *Proceedings of the 23rd International Conference on Real Time and Networks Systems (RTNS)*, pp. 129–138.


Chattopadhyay, Sudipta, Abhik Roychoudhury, and Tulika Mitra (2010). "Modeling Shared Cache and Bus in Multi-cores for Timing Analysis". In: *Proceedings of the 13th International Workshop on Software and Compilers for Embedded Systems*. SCOPES '10. St. Goar, Germany: ACM, 6:1–6:10.


Henry, Julien et al. (2014). "How to Compute Worst-case Execution Time by Optimization Modulo Theory and a Clever Encoding of Program Semantics". In: *Proceedings of the 2014 SIGPLAN/SIGBED Conference on Languages, Compilers and Tools for Embedded Systems (LCTES)*, pp. 43–52.


Kelter, Timon et al. (2014). "Static Analysis of Multi-core TDMA Resource Arbitration Delays". In: *Real-Time Syst.* 50.2, pp. 185–229.


Rosèn, Jacob et al. (2007). "Bus Access Optimization for Predictable Implementation of Real-Time Applications on Multiprocessor Systems-on-Chip". In: *RTSS 2007*.

BACKUP

Multicore Response Time Analysis

Example: Fixed Priority bus arbiter, PE1 $>$ PE0
Bus access delay $= 10$

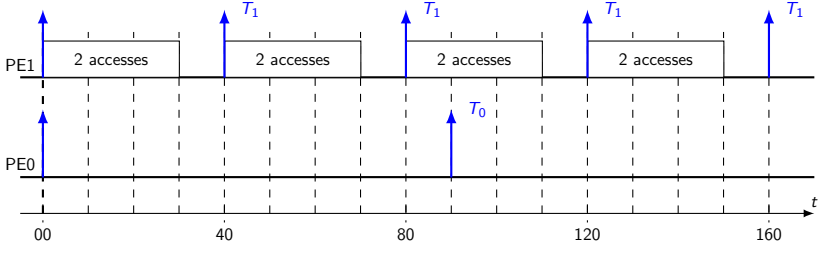

[1] Altmeyer et al., RTNS 2015

Multicore Response Time Analysis

Example: Fixed Priority bus arbiter, PE1 $>$ PE0
Bus access delay $= 10$

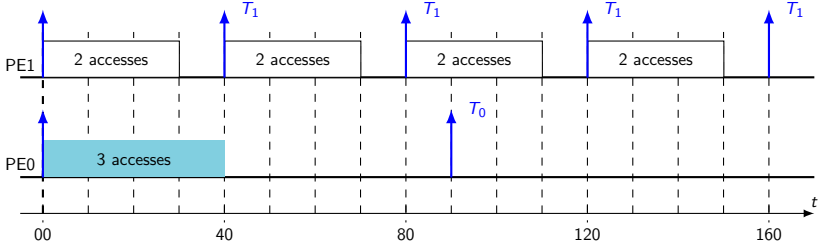

- Task of interest running on $PE0$:
  $R_0 = 10 + 3 \times 10$ (response time in isolation)

[1] Altmeyer et al., RTNS 2015

Multicore Response Time Analysis

Example: Fixed Priority bus arbiter, PE1 > PE0
Bus access delay = 10

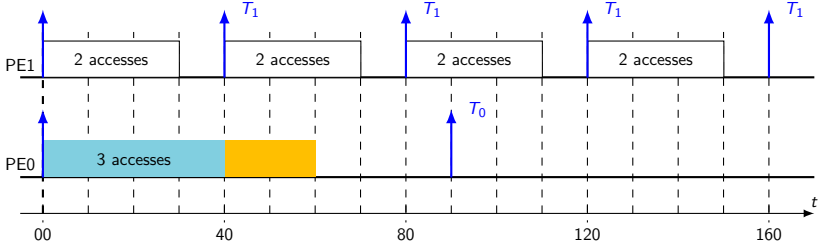

- Task of interest running on $PE0$:

  $R_0 = 10 + 3 \times 10$ (response time in isolation)

  $R_1 = 10 + 3 \times 10 + 2 \times 10 = 60$

[1] Altmeyer et al., RTNS 2015

Multicore Response Time Analysis

Example: Fixed Priority bus arbiter, PE1 > PE0
Bus access delay = 10

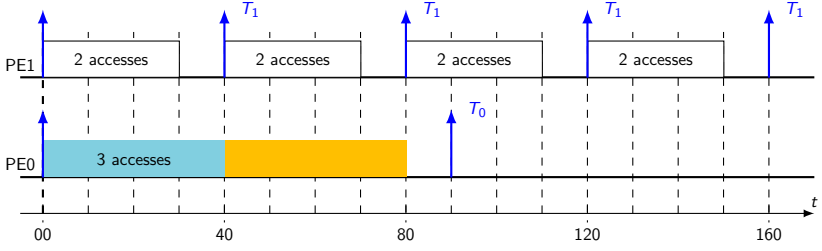

- Task of interest running on $PE0$:

  $R_0 = 10 + 3 \times 10$ (response time in isolation)

  $R_1 = 10 + 3 \times 10 + 2 \times 10 = 60$

  $R_2 = 10 + 3 \times 10 + 2 \times 10 + 2 \times 10 = 80$

[1] Altmeyer et al., RTNS 2015

Multicore Response Time Analysis

Example: Fixed Priority bus arbiter, PE1 > PE0
Bus access delay = 10

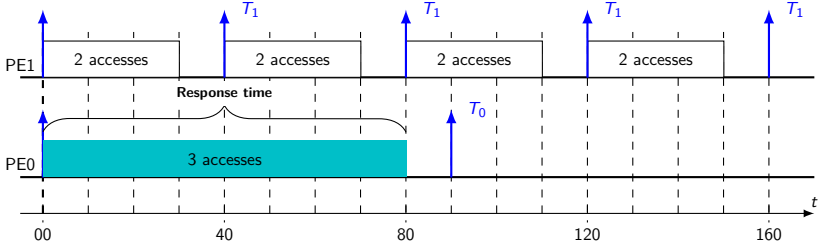

- Task of interest running on $PE0$:

  $R_0 = 10 + 3 \times 10$ (response time in isolation)

  $R_1 = 10 + 3 \times 10 + 2 \times 10 = 60$

  $R_2 = 10 + 3 \times 10 + 2 \times 10 + 2 \times 10 = 80$

  $R_3 = 10 + 3 \times 10 + 2 \times 10 + 2 \times 10 + 0 = 80$ (fixed-point)

[1] Altmeyer et al., RTNS 2015

Evaluation: Runtime Perfomance

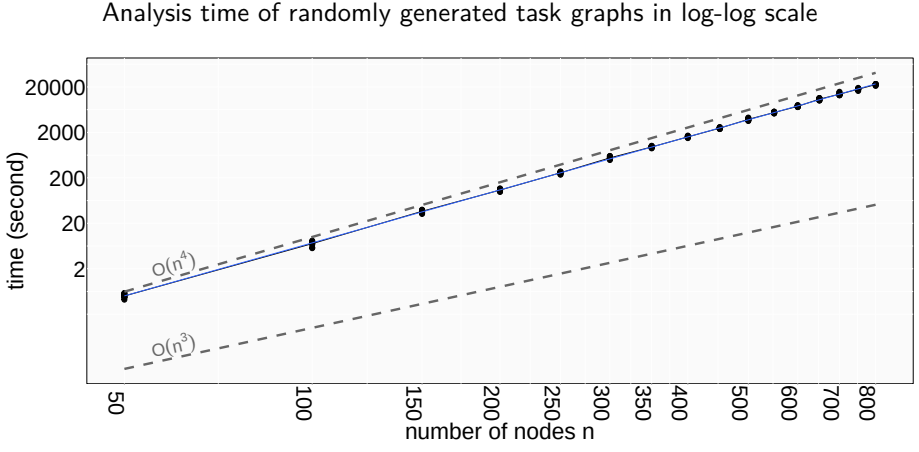

Theoretical complexity $\mathcal{O}(n^4)$. Experimental complexity $\mathcal{O}(n^{3.87})$.

Randomly Generated Task Graphs

Fat Task Graph

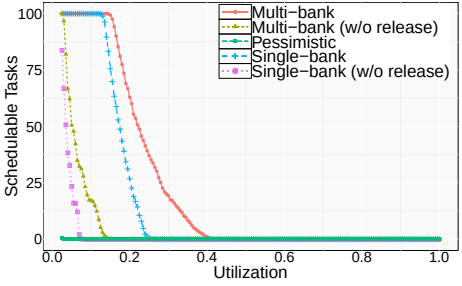

Long Task Graph

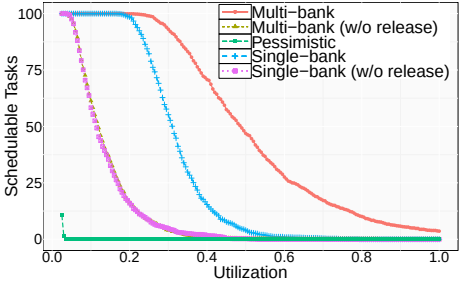

Cached Memory Operations

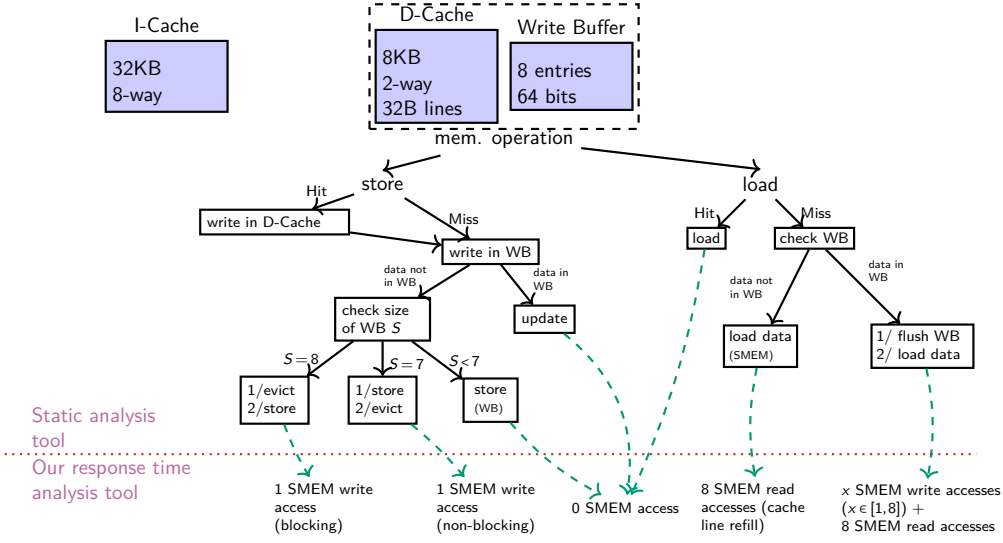

In isolation:

1 access = 10 cycles / 1 cache line refill = 17 cycles

Offset-based SMT Encoding

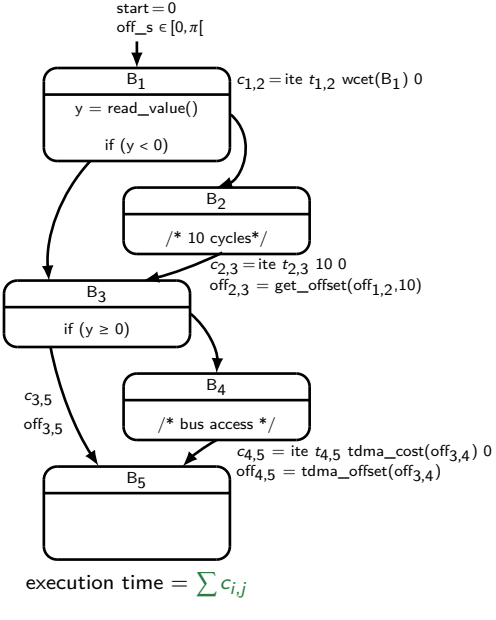

- $\text{off}_{i,j} = e_{i,j} \mod \pi$

  Encode the costs of the basic blocks

  $e_{i,j}$ (absolute time) $\dashrightarrow$ $c_{i,j}$ (cost)

$c_{i,j} =$ ite $t_{i,j}$ *cost* 0

"ite C A B" $\Leftrightarrow$ "if C then A else B"

Offset-based SMT Encoding

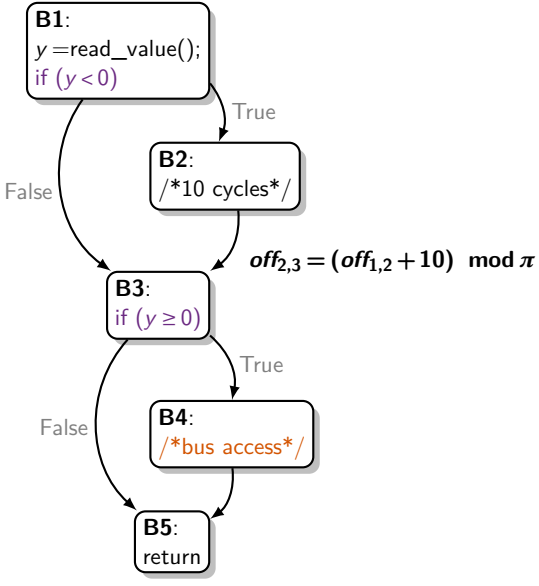

Offset-based SMT Encoding

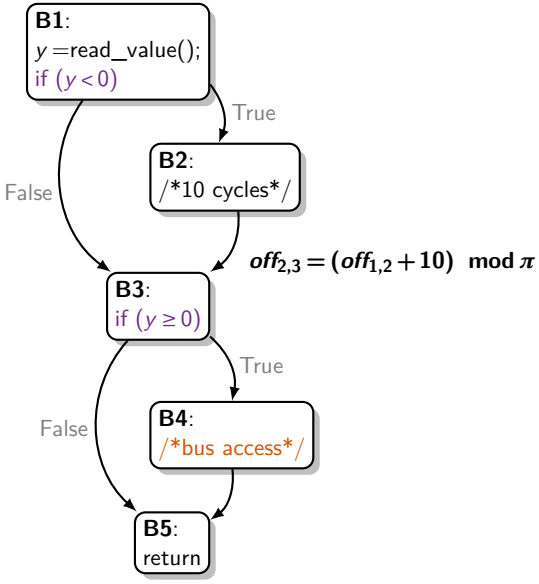

$$\overbrace{(\text{off}_{i,j}+c)}^{<\pi} \mod \pi$$

$$\Updownarrow$$

$$\overbrace{(\text{off}_{i,j}+c \mod \pi)}^{\stackrel{\text{def}}{=}\ \alpha\ <\ 2\pi} \underline{\mod}\ \pi$$

$$\downarrow$$

$$\text{if } \alpha < \pi \text{ then } \alpha \text{ else } \alpha - \pi$$

Offset-based SMT Encoding

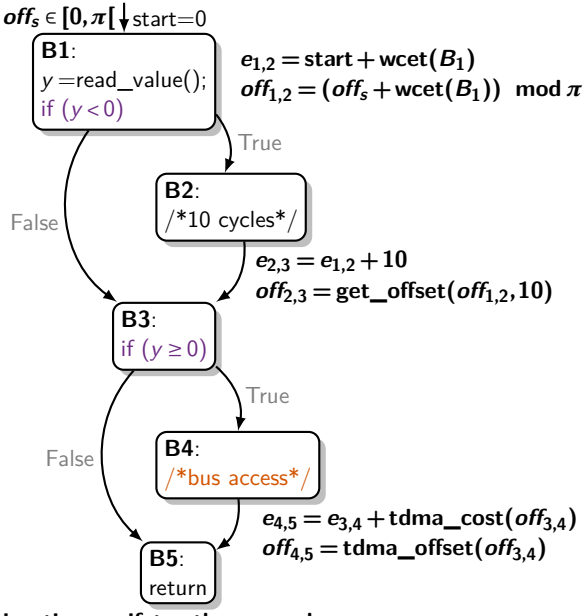

execution time $=$ **if** $t_{3,5}$ **then** $e_{3,5}$ **else** $e_{4,5}$

Using TDMA functions

- *if..then..else* encoding

  $\text{off} = \mathit{ite}\ t_{13}\ \text{off}_{13}\ \text{off}_{23}$

  $\text{off}_{35} = \text{off}$

  $\text{off}_{34} = \text{off}$

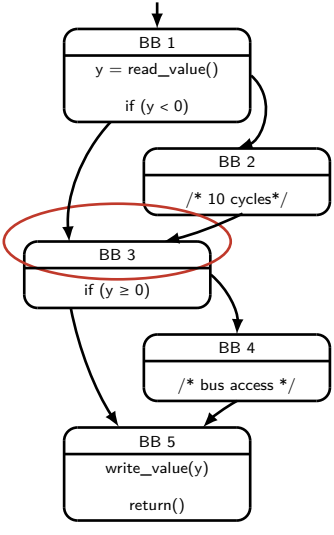

- *sum* encoding

  $\text{off} = \text{off}_{13} + \text{off}_{23}$

  $\text{off}_{35} = \mathit{ite}\ t_{35}\ \text{off}\ 0$

  $\text{off}_{34} = \mathit{ite}\ t_{34}\ \text{off}\ 0$

Using TDMA functions

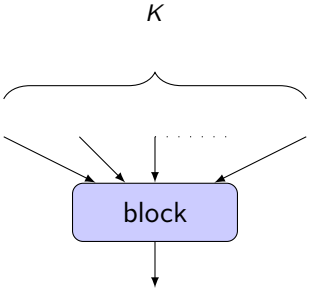

if..then..else (ite) encoding:

$$off_{i,j} = (\textbf{if } t_{1,i} \textbf{ then } off_{1,i} \textbf{ else if } t_{2,i} \textbf{ then } off_{2,i} \textbf{ else } \ldots \textbf{ else if } t_{K,i} \textbf{ then } off_{K,i} \textbf{ else } 0)$$

sum encoding:

$$off_i = \sum_{k=1}^{k=K} off_{k,i}$$

$$off_{i,j} = \textbf{if } t_{i,j} \textbf{ then } off_i \textbf{ else } 0$$

Performance 3

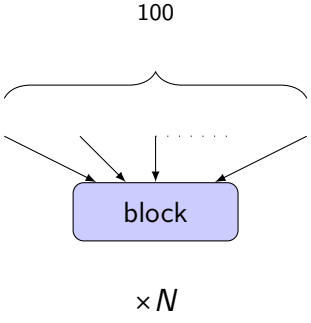

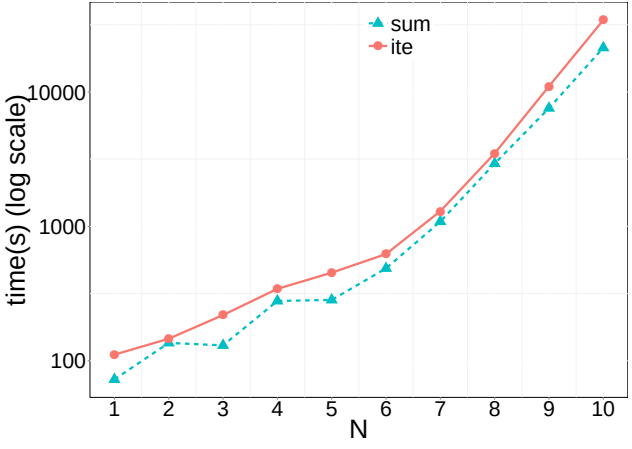

How it works?

- Example with binary search:

```
Testing wcet >= 0... SAT (value found = 18).
                              New interval = [18, 73].
Testing wcet >= 46... UNSAT. New interval = [18, 45].
Testing wcet >= 32... UNSAT. New interval = [18, 31].
Testing wcet >= 25... UNSAT. New interval = [18, 24].
Testing wcet >= 21... UNSAT. New interval = [18, 20].
Testing wcet >= 19... UNSAT. New interval = [18, 18].
The maximum value of wcet is 18 .
Computation time is 0.010000s
```

Evaluation: Analysis Time

Analysis time

Name	$\pi = 40,\ \sigma = 20,\ acc = 10$	$\pi = 400,\ \sigma = 200,\ acc = 40$
bs	**0.45s**	0.80s
insertsort	1.37s	7.19s
jfdctint	44.10s	55.47s
fdct	41.36s	34.42s
compressdata	4.66s	3.44s
fly-by-wire	28.78s	**109.37s**