

Analyzing ARM's MPAM From the Perspective of Time Predictability

M. Zini, D. Casini, A. Biondi,

"Analyzing ARM's MPAM From the Perspective of Time Predictability",
in *Transactions on Computers*, 2022

ReTiS Lab, Scuola Superiore Sant'Anna, Pisa, Italy




Sant'Anna
Scuola Universitaria Superiore Pisa



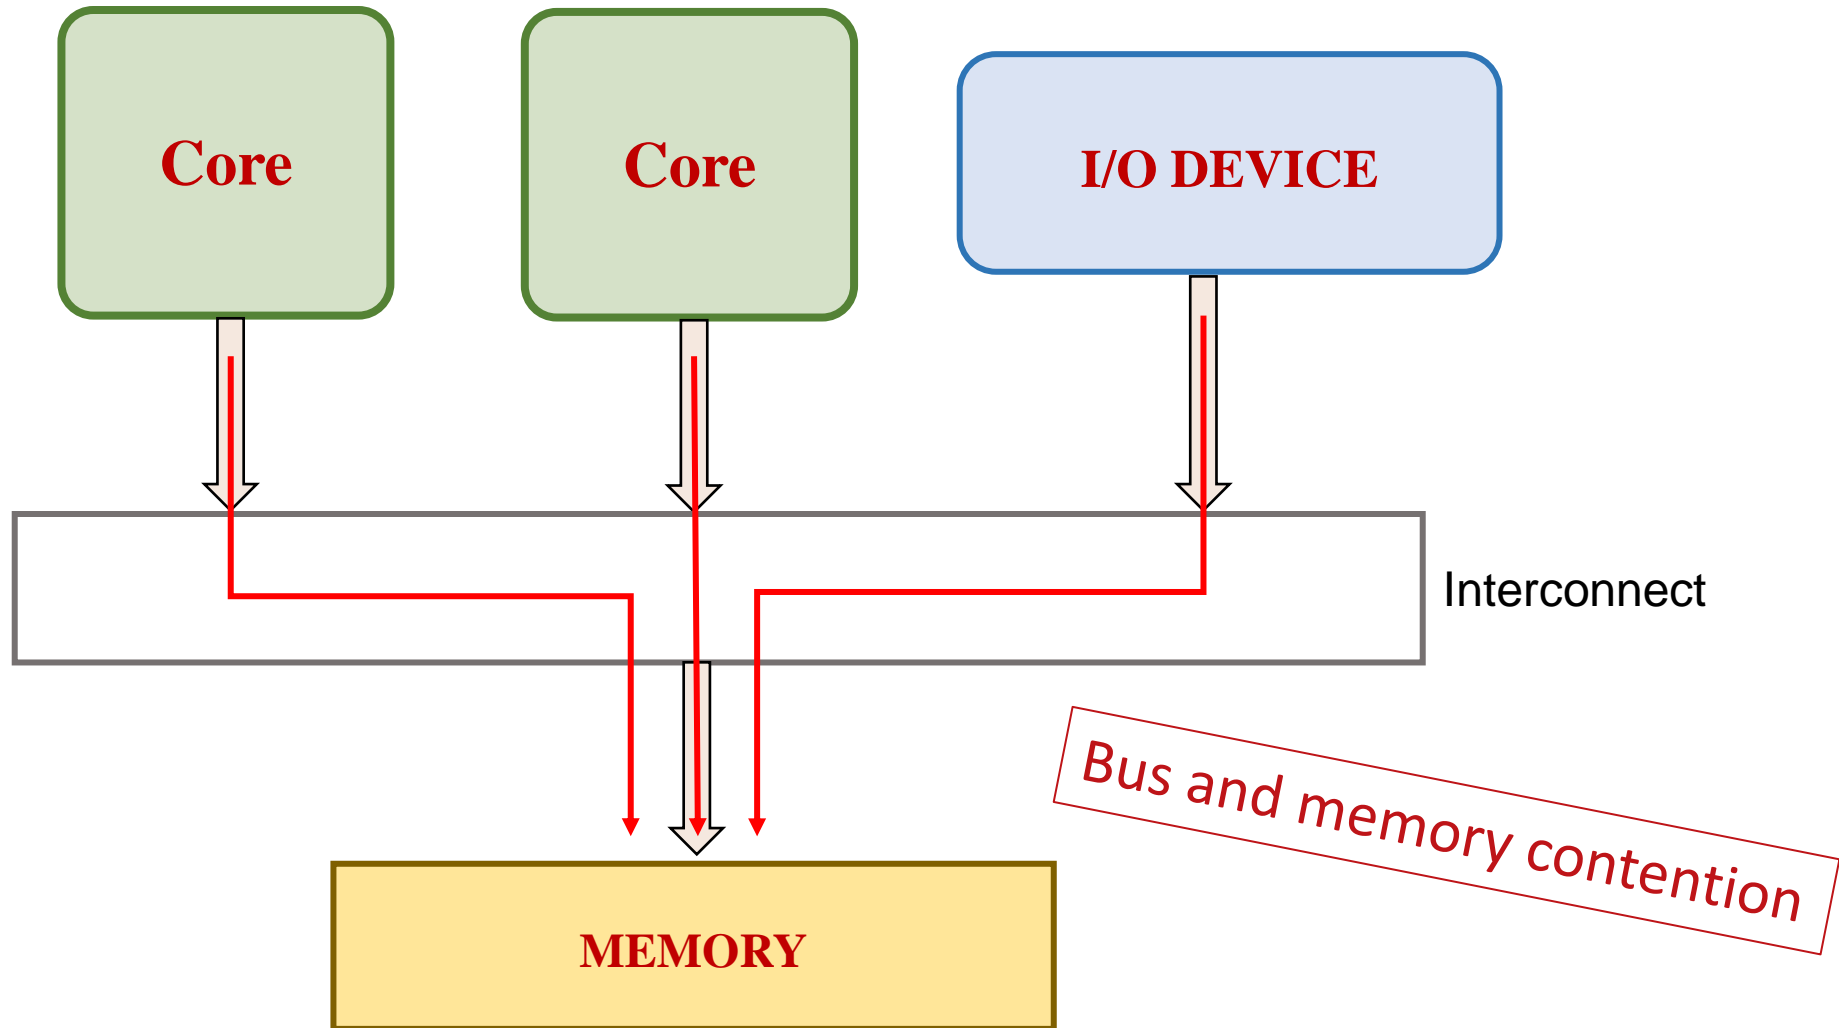
Sources of delay

- In modern computing platform, not only the processor is a **shared resource** among tasks
- Systems that require predictability need to take into account other factors: two of them are **cache and memory contention**

Cache contention



Bus and Memory Contention



Memory system resource Partitioning and Monitoring

Extension of ARMv8-A architecture

- The MPAM specifications describe **hardware mechanisms** to limit cache and memory access
- Such mechanism can be integrated inside CPU cores and Memory System Components

- **State of the art**
- **MPAM architecture**
- **Modelling the system**
- **Results**

- **State of the art**
- MPAM architecture
- Modelling the system
- Results

Software-based techniques:

- Cache coloring
- Bandwidth partitioning with usage monitors

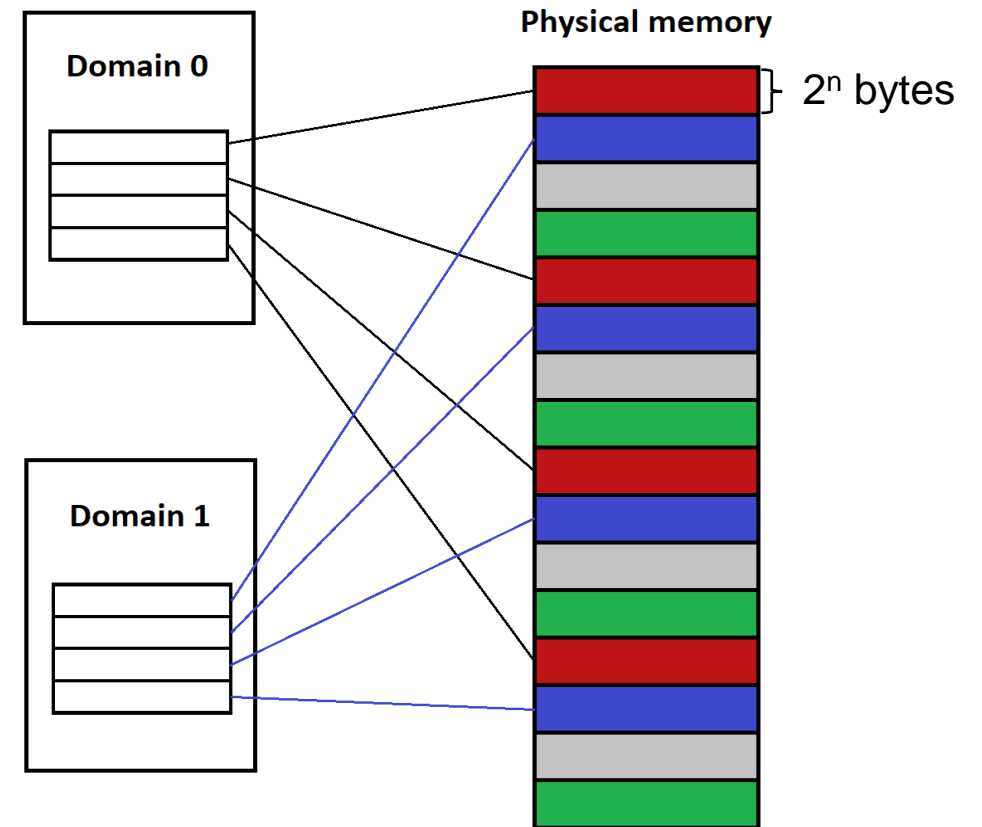

Hardware-based techniques:

- Intel Resource Director Technology (RDT)

State of the art: cache coloring

Exploit relation between:


- Physical page number
- Cache index




State of the art: bandwidth partitioning

- Memory bandwidth monitors
- Explicit preemption of the tasks

Bandwidth monitor

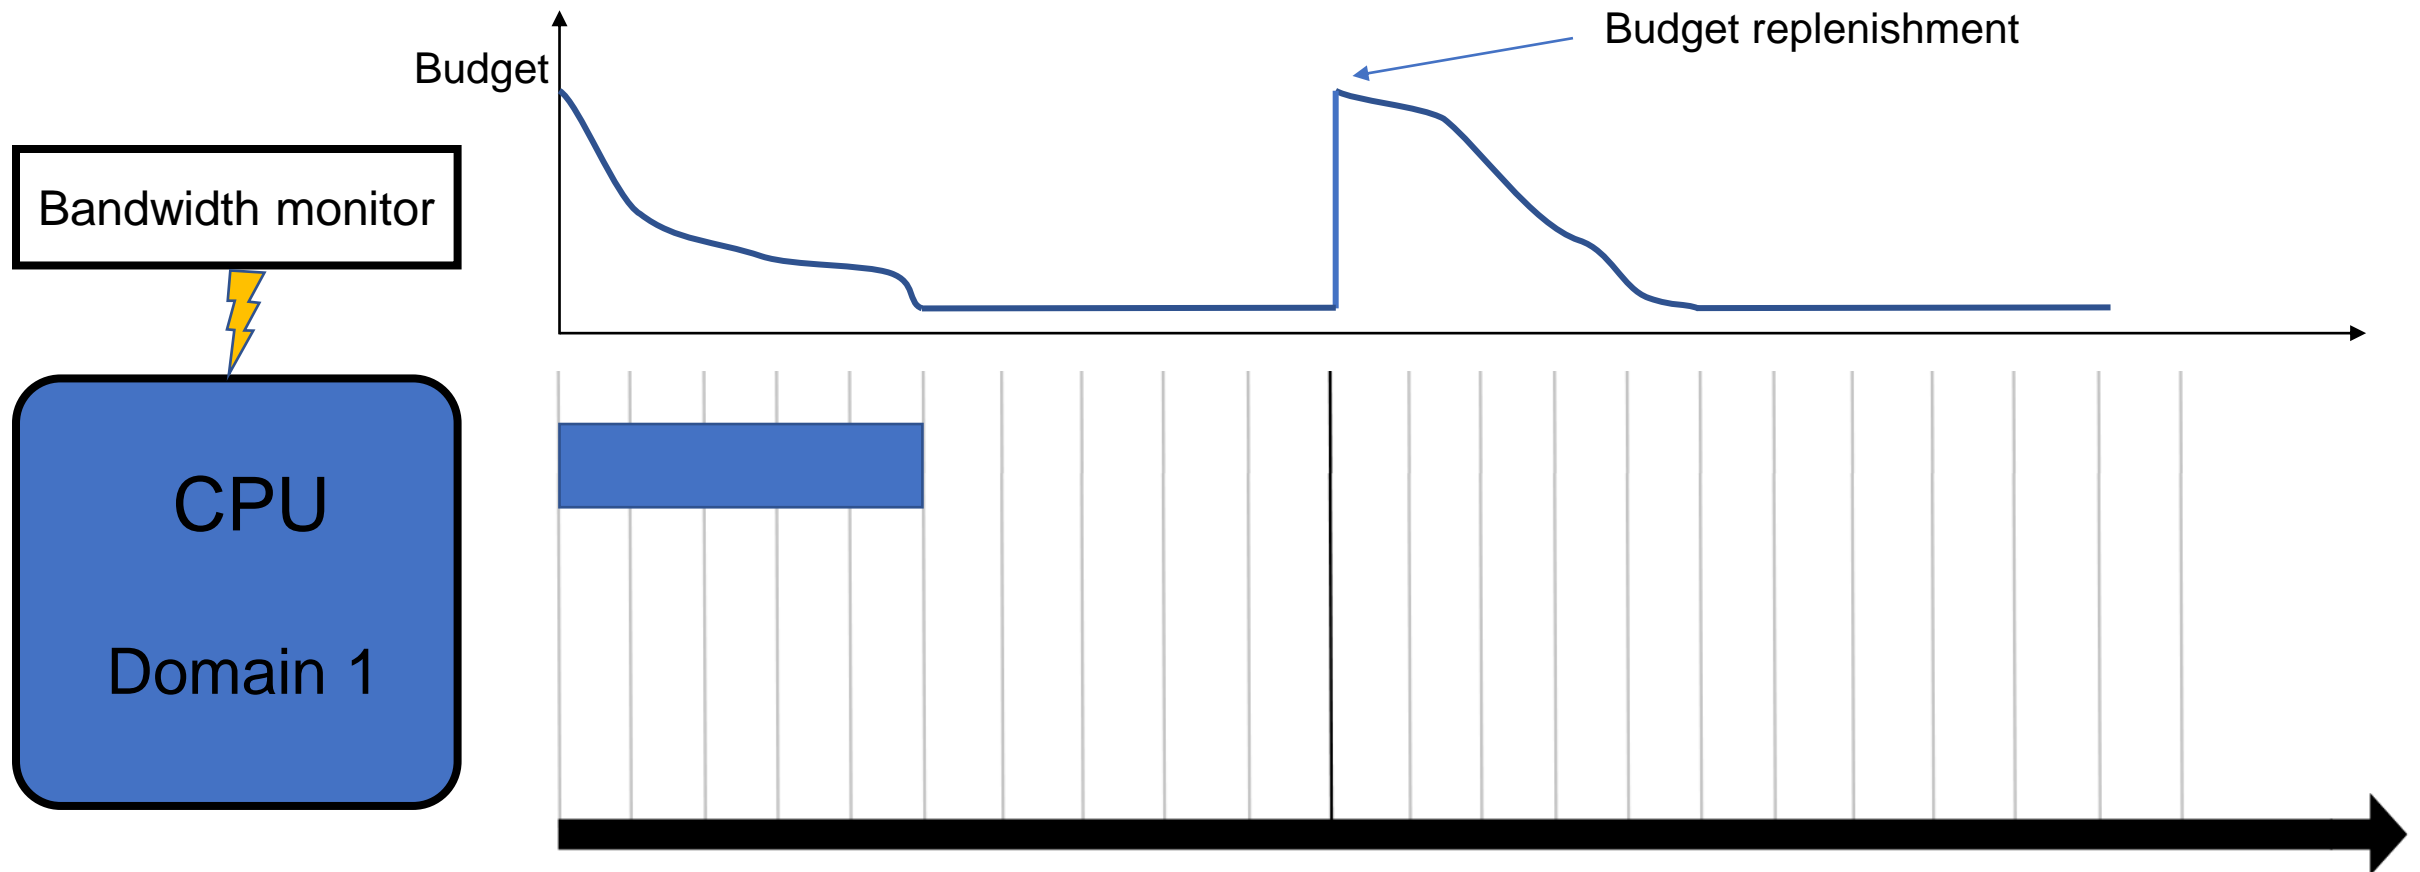


CPU
Domain 1



State of the art: bandwidth partitioning


- Memory bandwidth monitors
- Explicit preemption of the tasks




State of the art: bandwidth partitioning

- Memory bandwidth monitors
- Explicit preemption of the tasks

Bandwidth monitor



CPU
Domain 2



Intel Resource Director Technology functionalities

- Cache Monitoring Technology
- Memory Bandwidth Monitoring
- Cache Allocation Technology
- Code and Data Prioritization
- Memory Bandwidth Allocation

Example of CAT-Only Usage - 16 bit Capacity Masks

COS0	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	Traditional CAT
COS1	0	0	0	0	1	1	1	1	0	0	0	0	0	0	0	
COS2	0	0	0	0	0	0	0	0	1	1	0	0	0	0	0	
COS3	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	

Example of Code/Data Prioritization Usage - 16 bit Capacity Masks

COS0.Data	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	CAT with CDP
COS0.Code	0	0	0	0	1	1	0	0	0	0	0	0	0	0	0	
COS1.Data	0	0	0	0	0	0	1	1	1	1	0	0	0	0	0	
COS1.Code	0	0	0	0	0	0	0	0	1	1	1	0	0	0	0	
Other COS.Data	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	
Other COS.Code	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	

- State of the art
- **MPAM architecture**
- Modelling the system
- Results

Memory system resource Partitioning and Monitoring

Extension of ARMv8-A architecture

- **Resource Monitoring**
- **Resource Partitioning**

Cache and memory bandwidth

Resource monitoring

- **Cache storage usage monitors**
- **Memory-bandwidth usage monitors**

Resource partitioning

- **Cache partitioning**
 - Portion partitioning
 - Minimum-Maximum capacity partitioning
 - Associativity partitioning
- **Memory bandwidth partitioning**
 - Portion partitioning
 - Proportional stride partitioning
 - Minimum-maximum partitioning
 - Priority partitioning

**The MPAM specifications are
vague**

**Many details are
implementation defined**

**This is needed because ARM
has to take into account many
different needs**

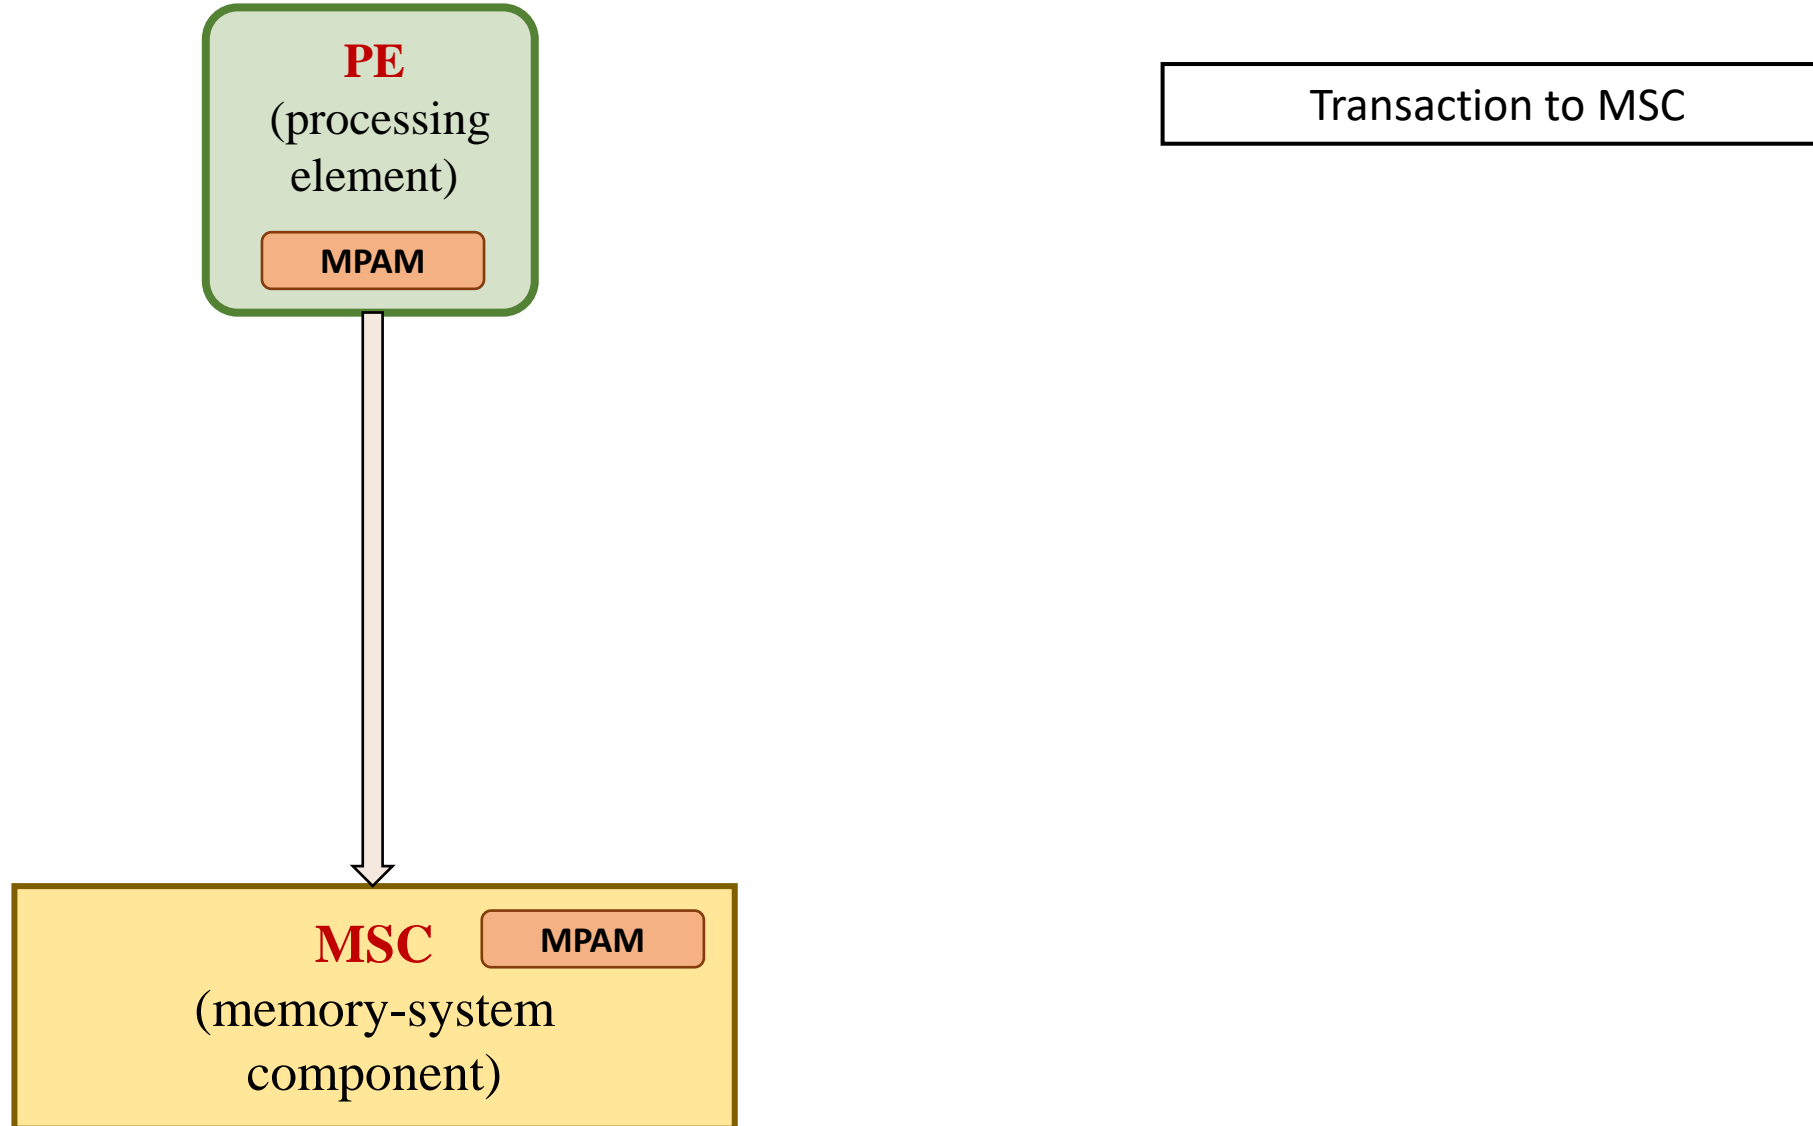
MPAM implementations

No **practical implementation** to date for
real-time systems

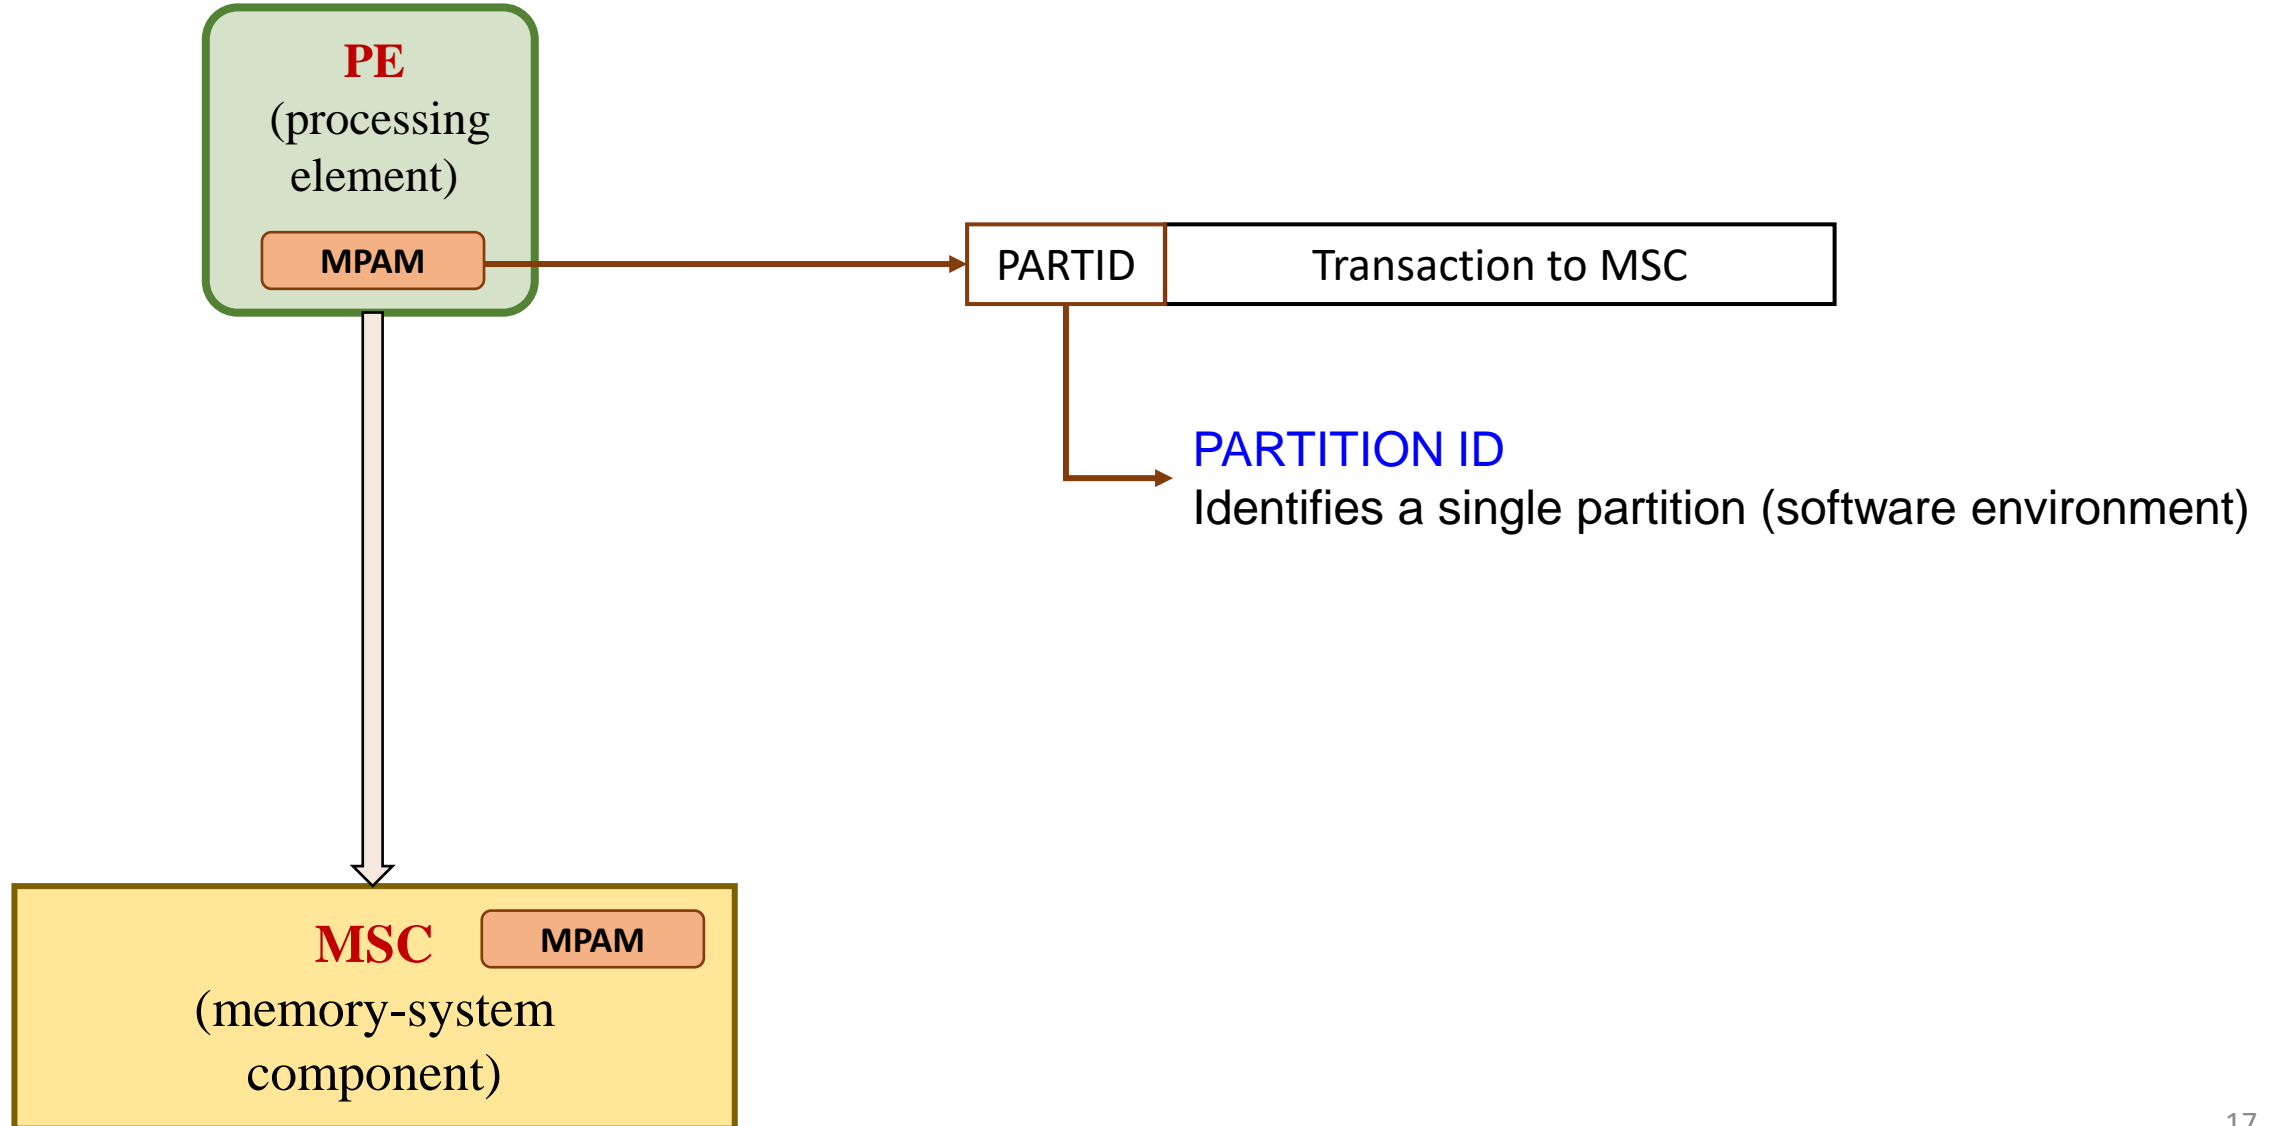


Possibility to influence the future
design choices of hardware vendors

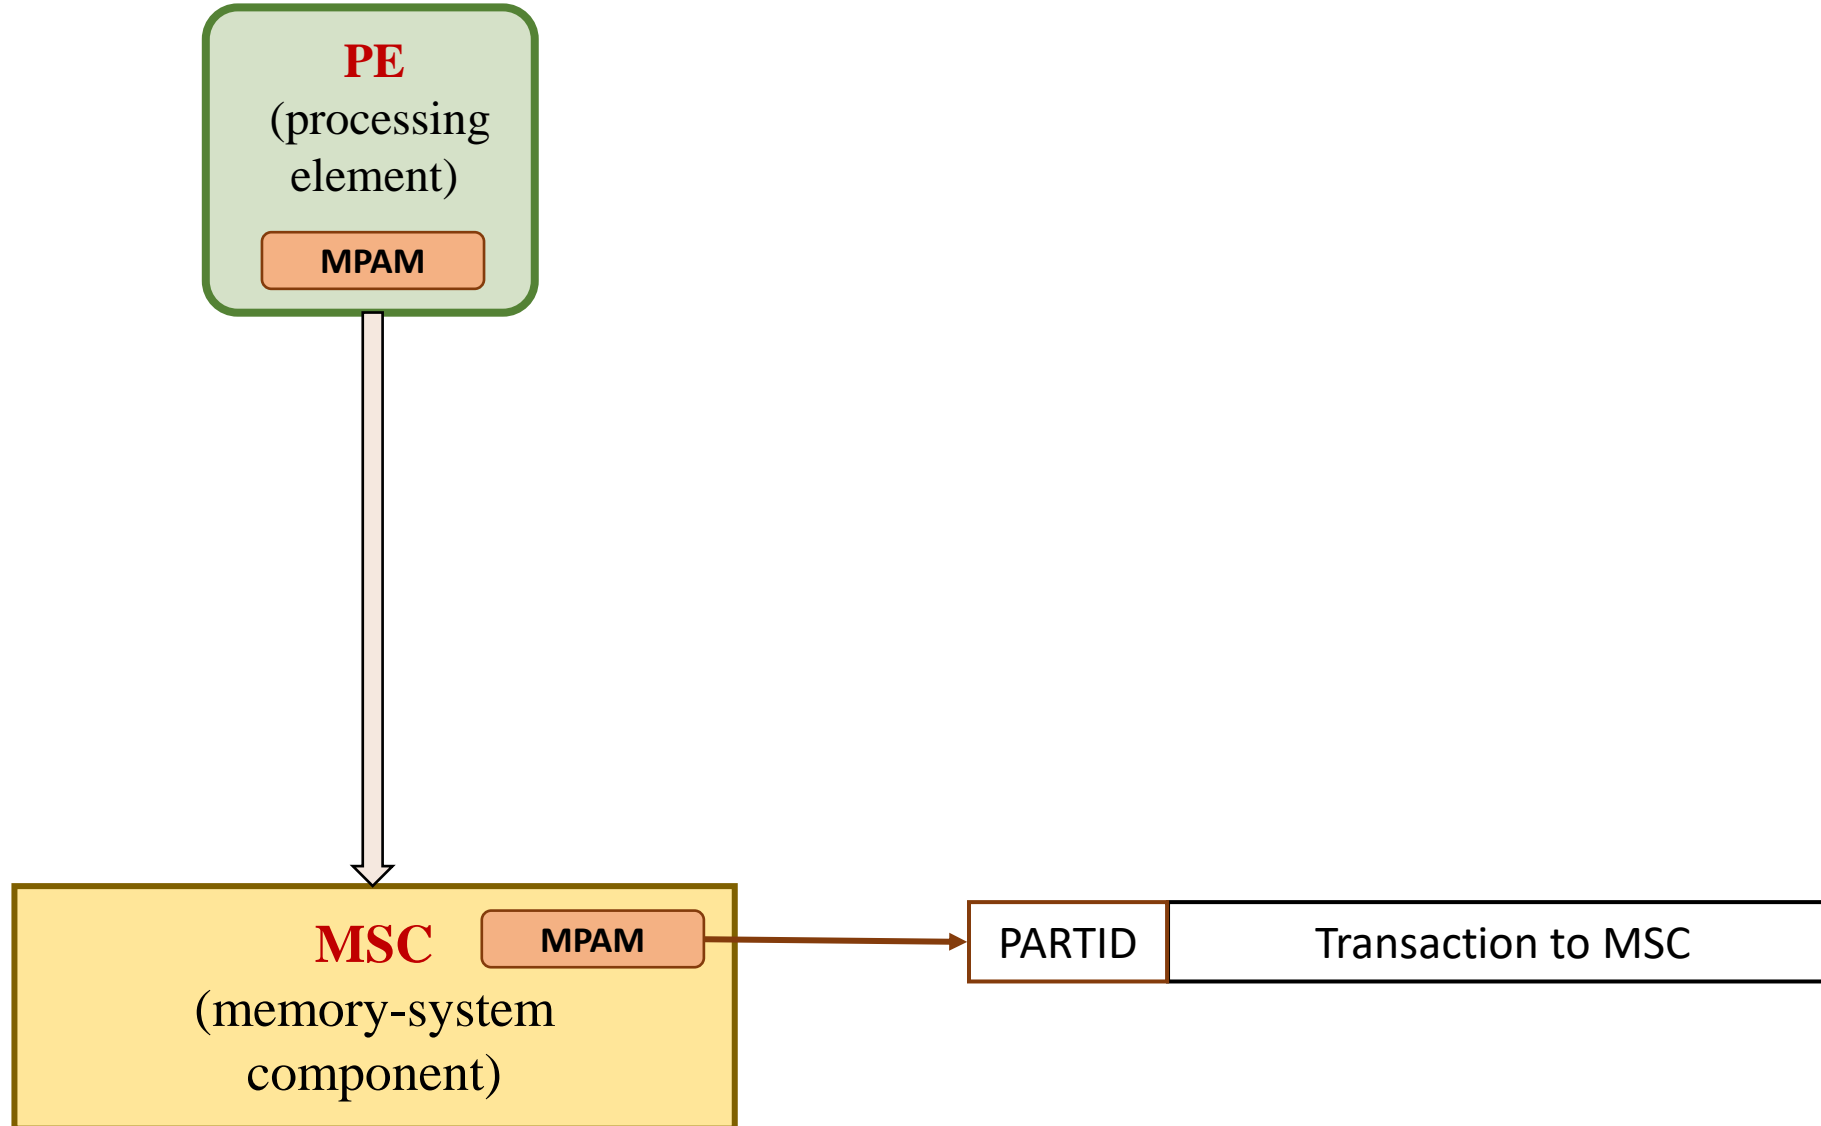
MPAM System Architecture



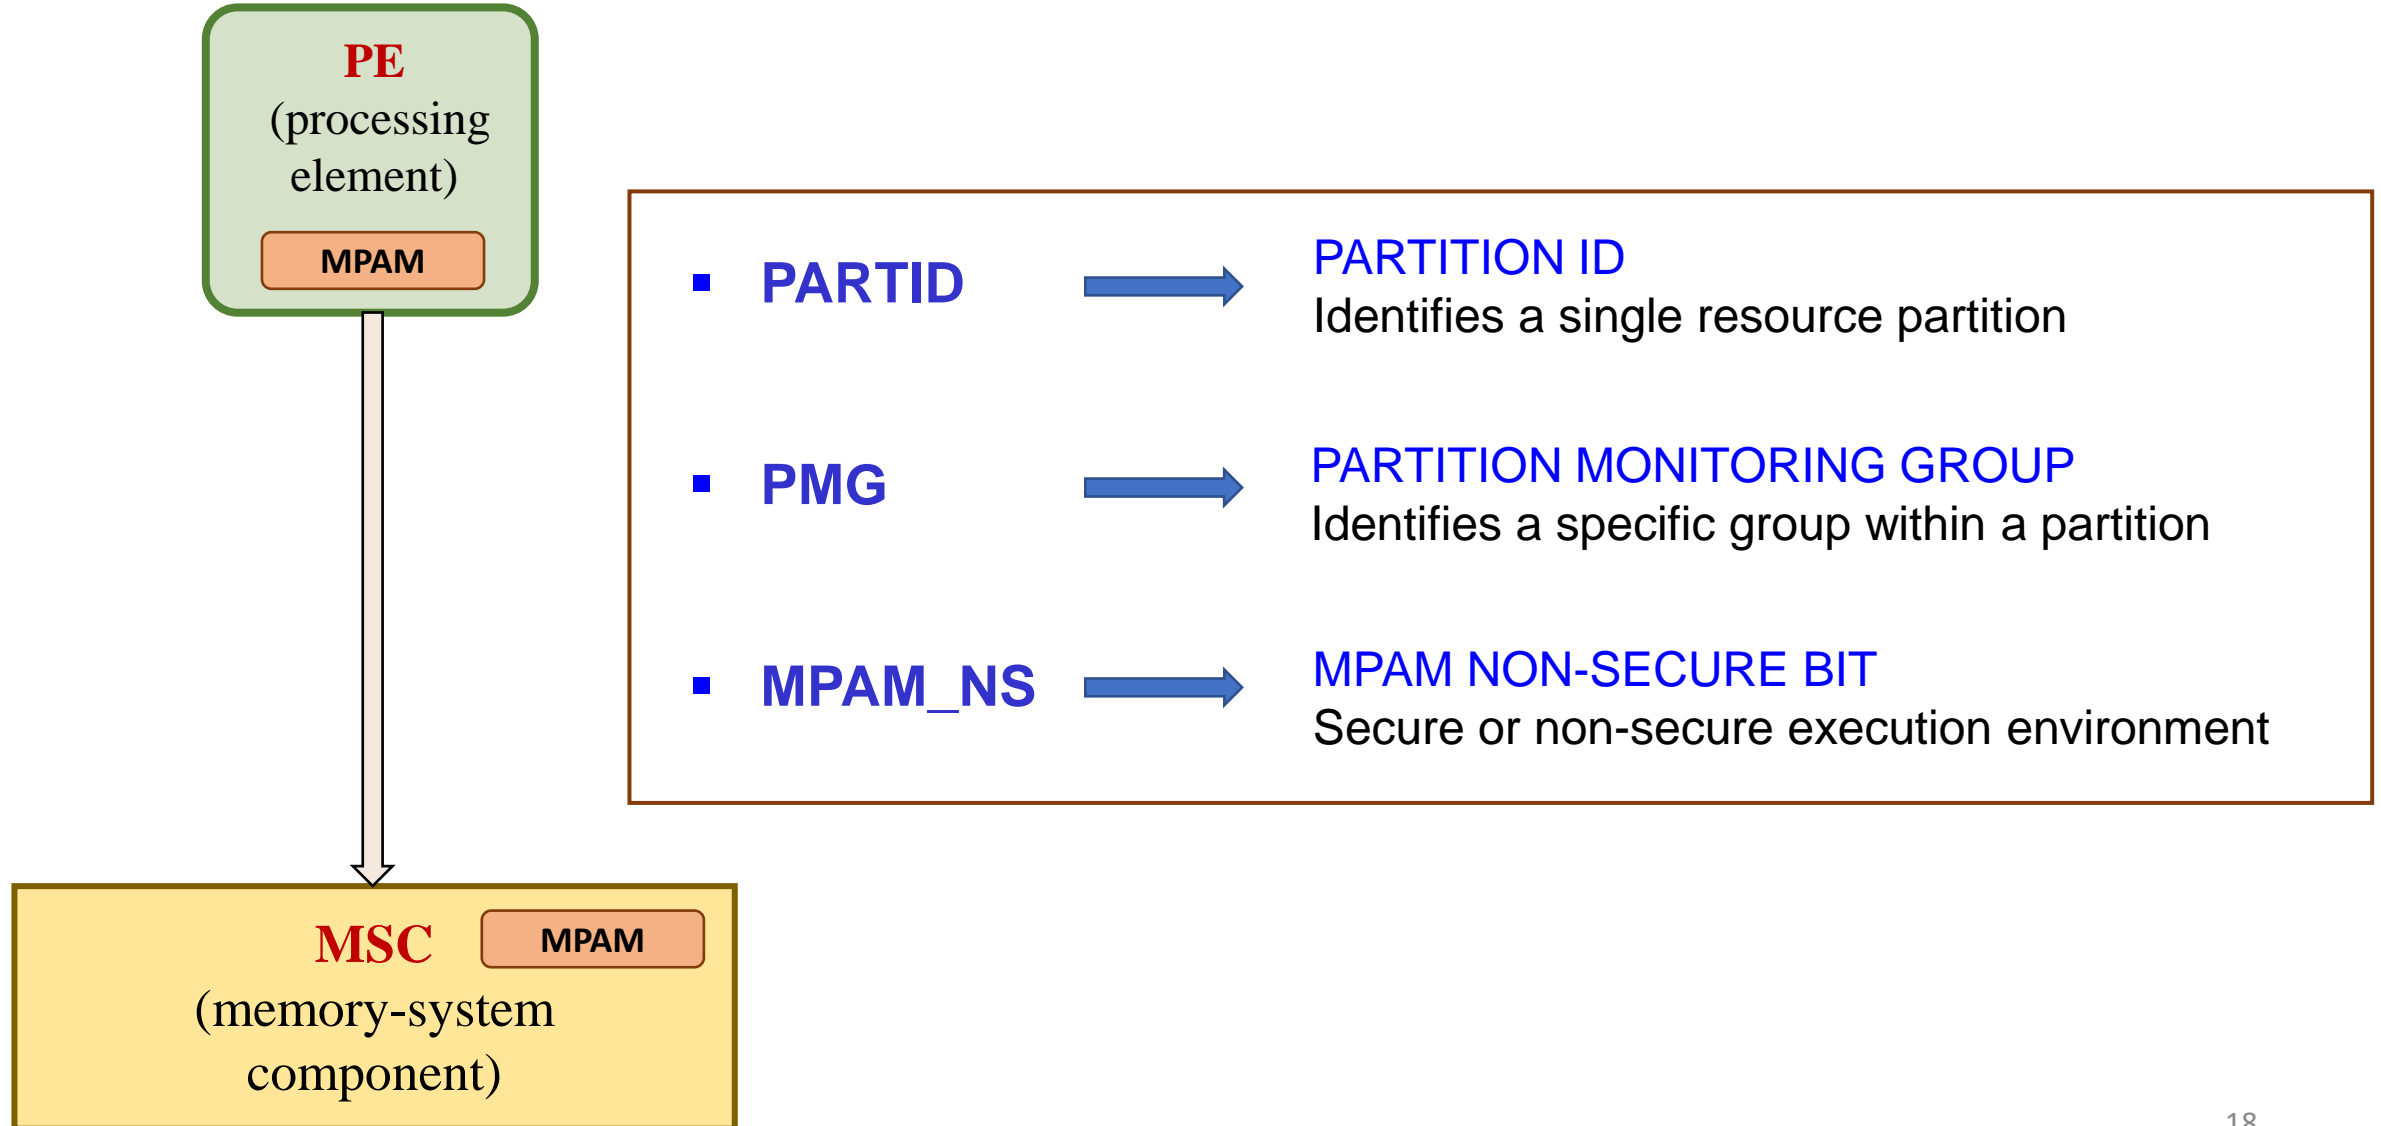
MPAM System Architecture



MPAM System Architecture



MPAM System Architecture



- State of the art
- MPAM architecture
- **Modelling the system**
- Results

Memory system resource Partitioning and Monitoring

Extension of ARMv8-A architecture

- **Resource Monitoring**


- **Resource Partitioning**

Cache and memory bandwidth


- Priority partitioning
- Minimum-maximum partitioning

MSC Bandwidth Control

- **Priority partitioning:** priority-based requests scheduling

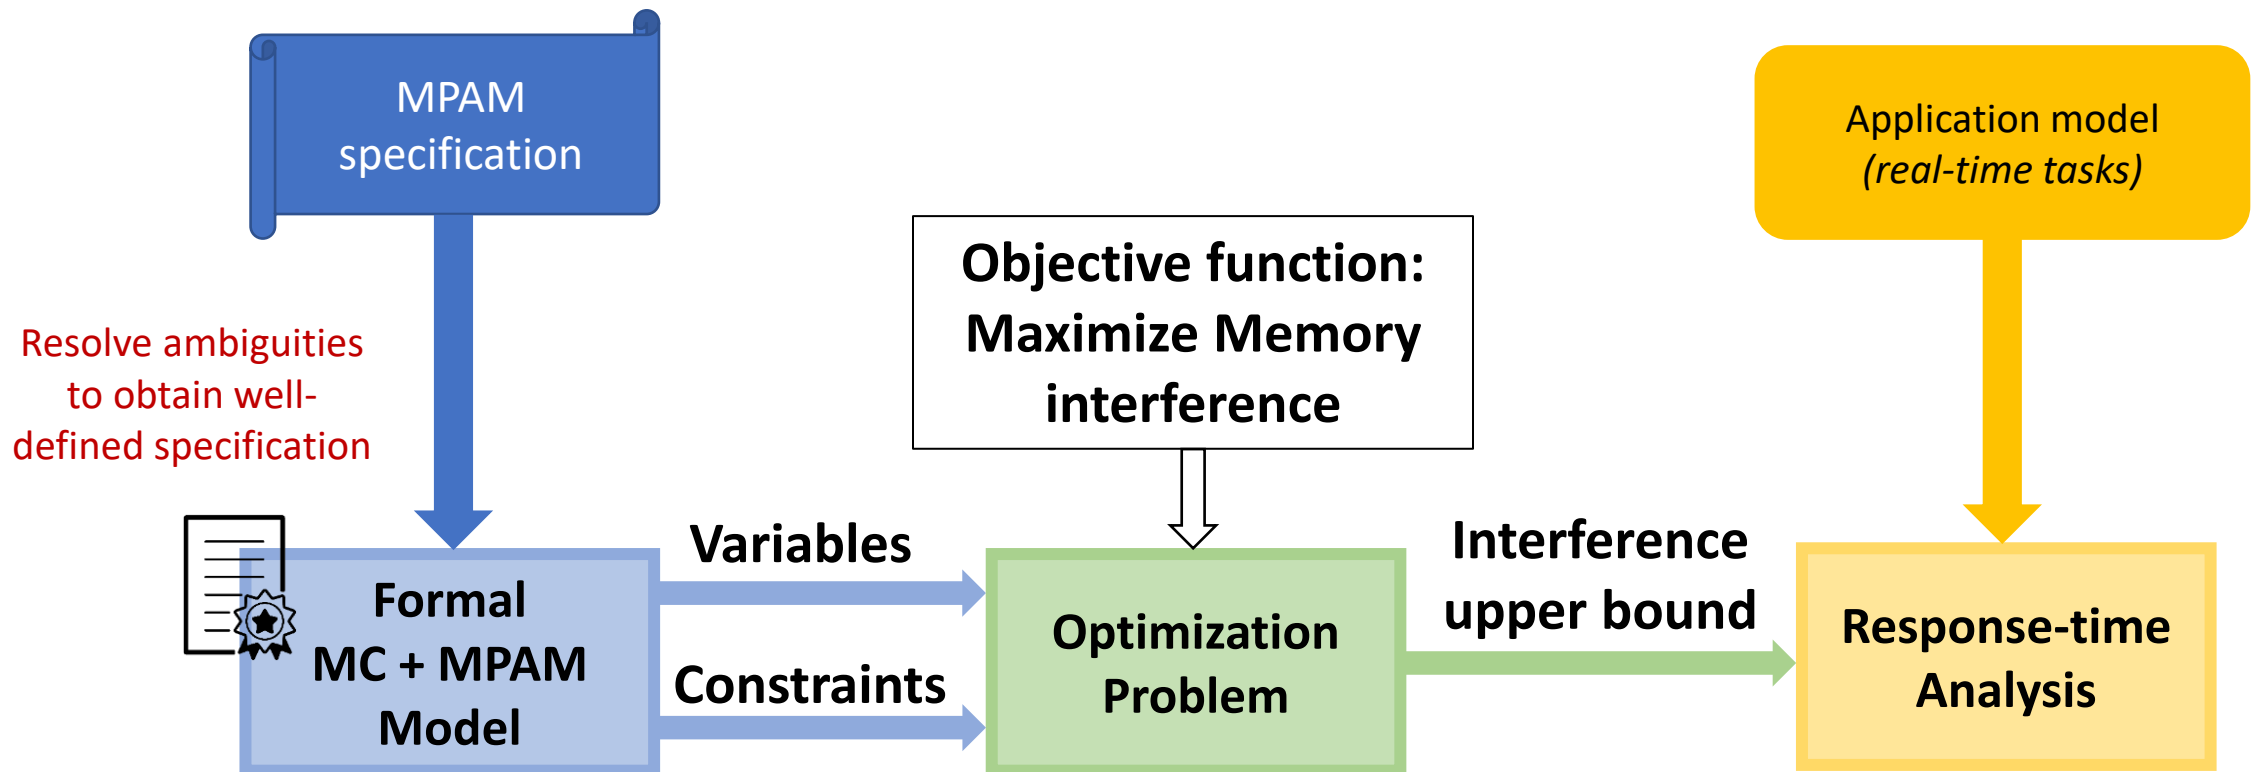


- **Memory-bandwidth minimum and maximum partitioning:** a minimum and maximum budget can be specified for every partition



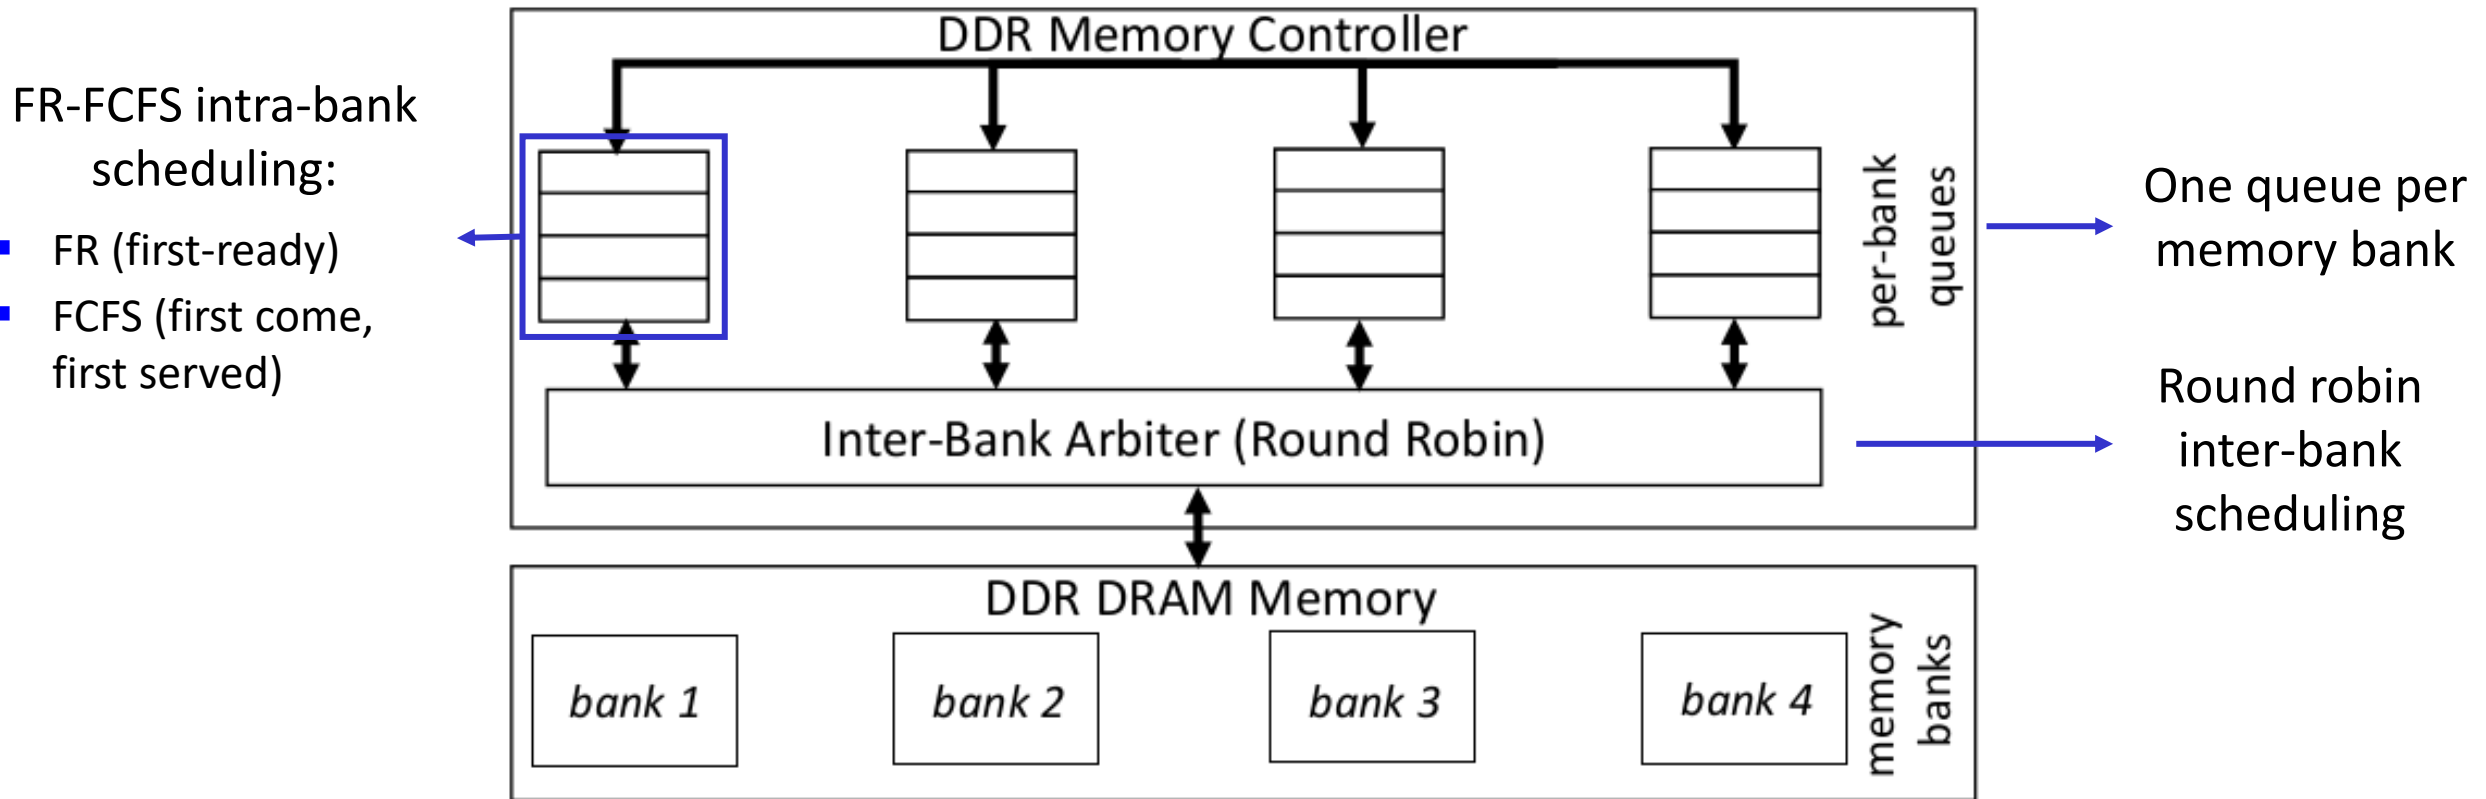
Modelling the system

The system has been modelled with an optimization problem, in order to compute the worst-case memory interference with and without the MPAM's functionalities.



Memory controller model

The memory controller model is based on a generalization of models already used in literature. The assumptions of the existing models have been relaxed.



First-ready strategy

The intra-bank queues are scheduled according to the FR-FCFS policy (with thresholding)

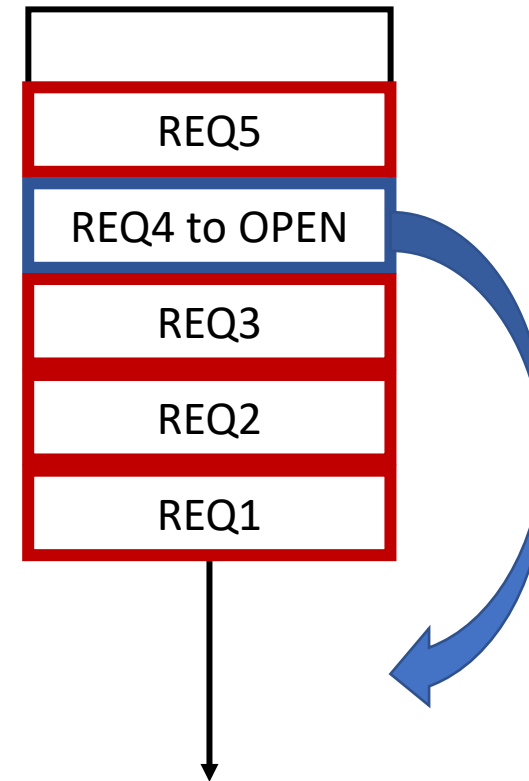
- FCFS (first come, first served)

- FR (first-ready):

If a request targets an *open* memory row, it is prioritized

The previously served request targeted the same memory row

Intra-bank queue



Served first

MC + MPAM

How do we **combine** the rules of the MC
and the rules of MPAM?



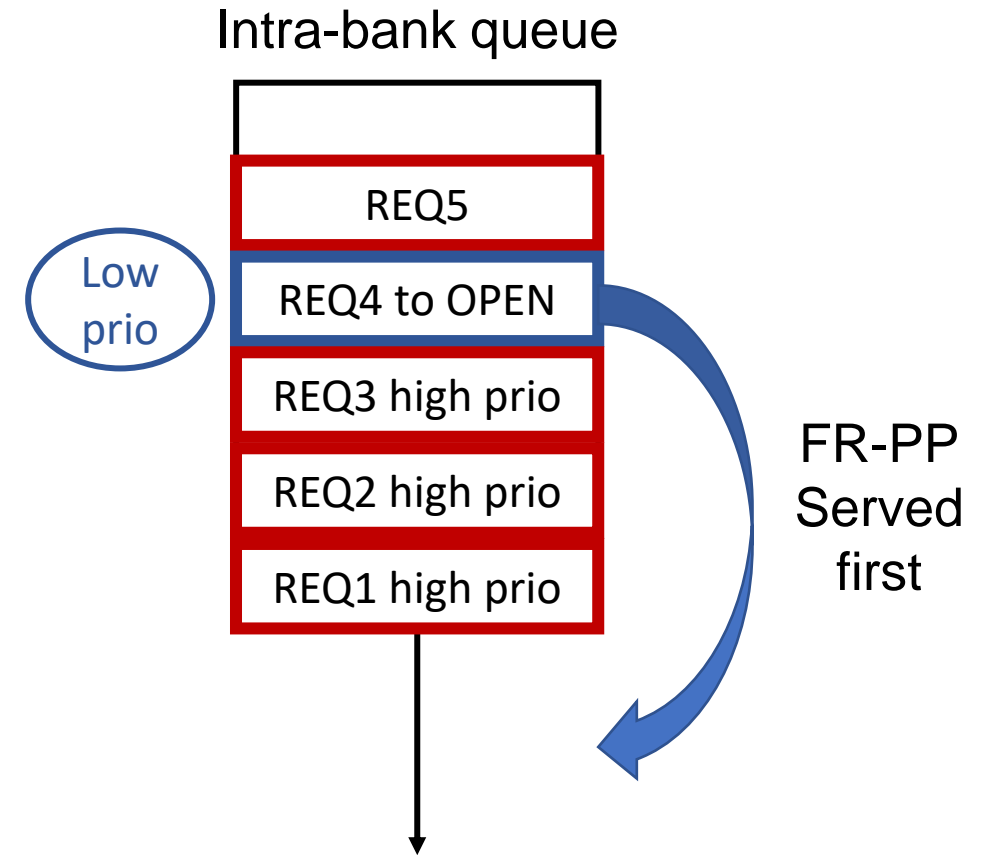
We need to find a **global solution** that
preserves the properties of the two
components

Modelling the system

Two possible implementations of priority partitioning were compared:

➤ **FR-PP (first-ready - priority partitioning):**

In case a request targeting an open row is pending, it has precedence over other requests with any priority



Modelling the system

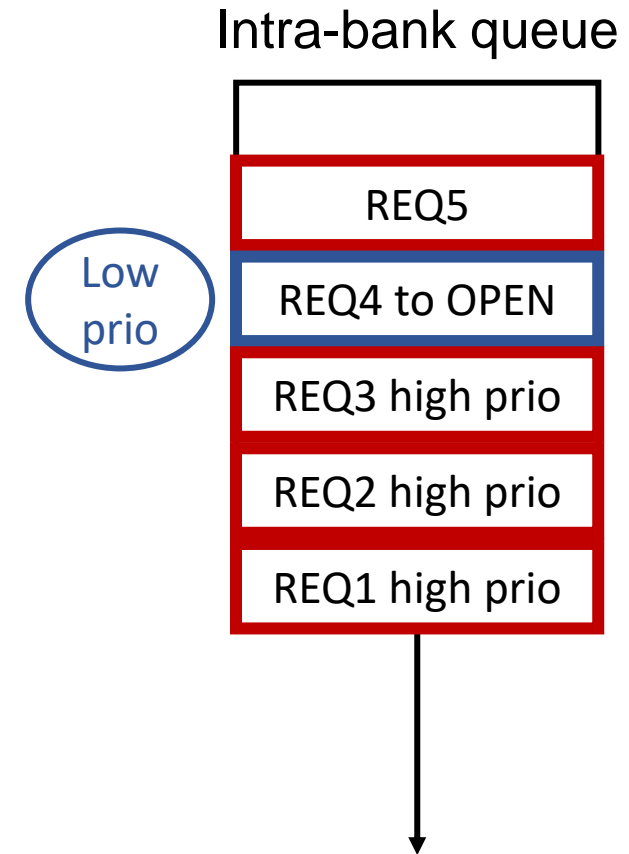
Two possible implementations of priority partitioning were compared:

➤ **FR-PP (first-ready - priority partitioning):**

In case a request targeting an open row is pending, it has precedence over other requests with any priority

➤ **PP-FR (priority partitioning - first-ready):**

In case a request targeting an open row is pending, it DOES NOT have precedence over other requests with higher priority



Modelling the system

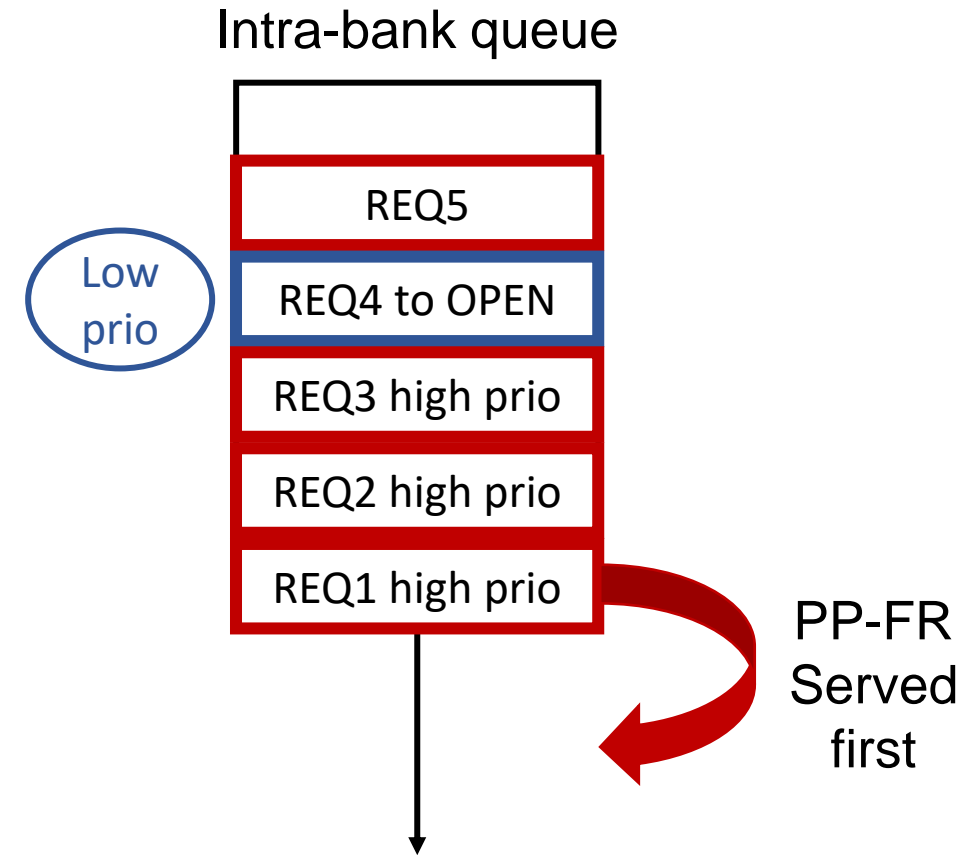
Two possible implementations of priority partitioning were compared:

➤ **FR-PP (first-ready - priority partitioning):**

In case a request targeting an open row is pending, it has precedence over other requests with any priority

➤ **PP-FR (priority partitioning - first-ready):**

In case a request targeting an open row is pending, it DOES NOT have precedence over other requests with higher priority




- State of the art
- MPAM architecture
- Modelling the system
- **Results**

Priority Partitioning

The model has been tested with a task set derived from the WATERS 2019 industrial challenge by Bosch

(a) max improvement



Priority Partitioning

The model has been tested with a task set derived from the WATERS 2019 industrial challenge by Bosch

(a) max improvement




For each task, the plot shows the improvement when that task is assigned the highest priority

Priority Partitioning

(b) rate monotonic

% of memory interference, w.r.t. the case w/o MPAM

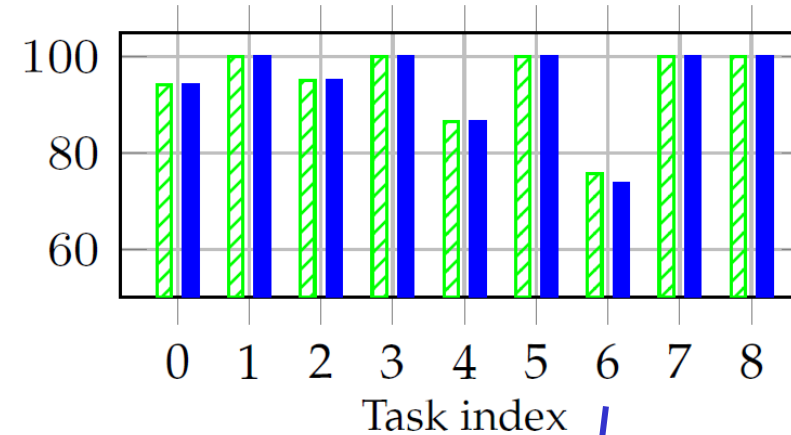


Task 1 has the smallest period



Task 1 has the biggest improvement


(c) request monotonic



Task 6 has the highest number of requests



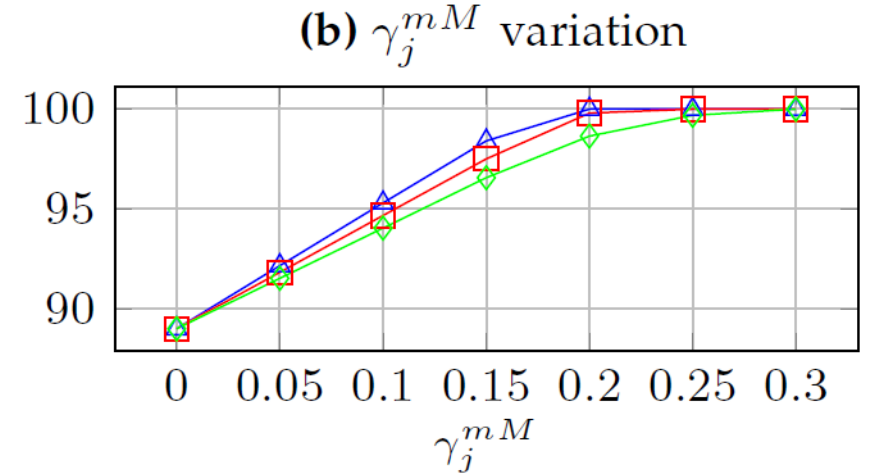



Task 6 has the biggest improvement



Min-max Partitioning

% of memory interference, w.r.t. the case w/o MPAM



$$\gamma = \frac{\#req \text{ high prio}}{\#req \text{ medium prio}}$$

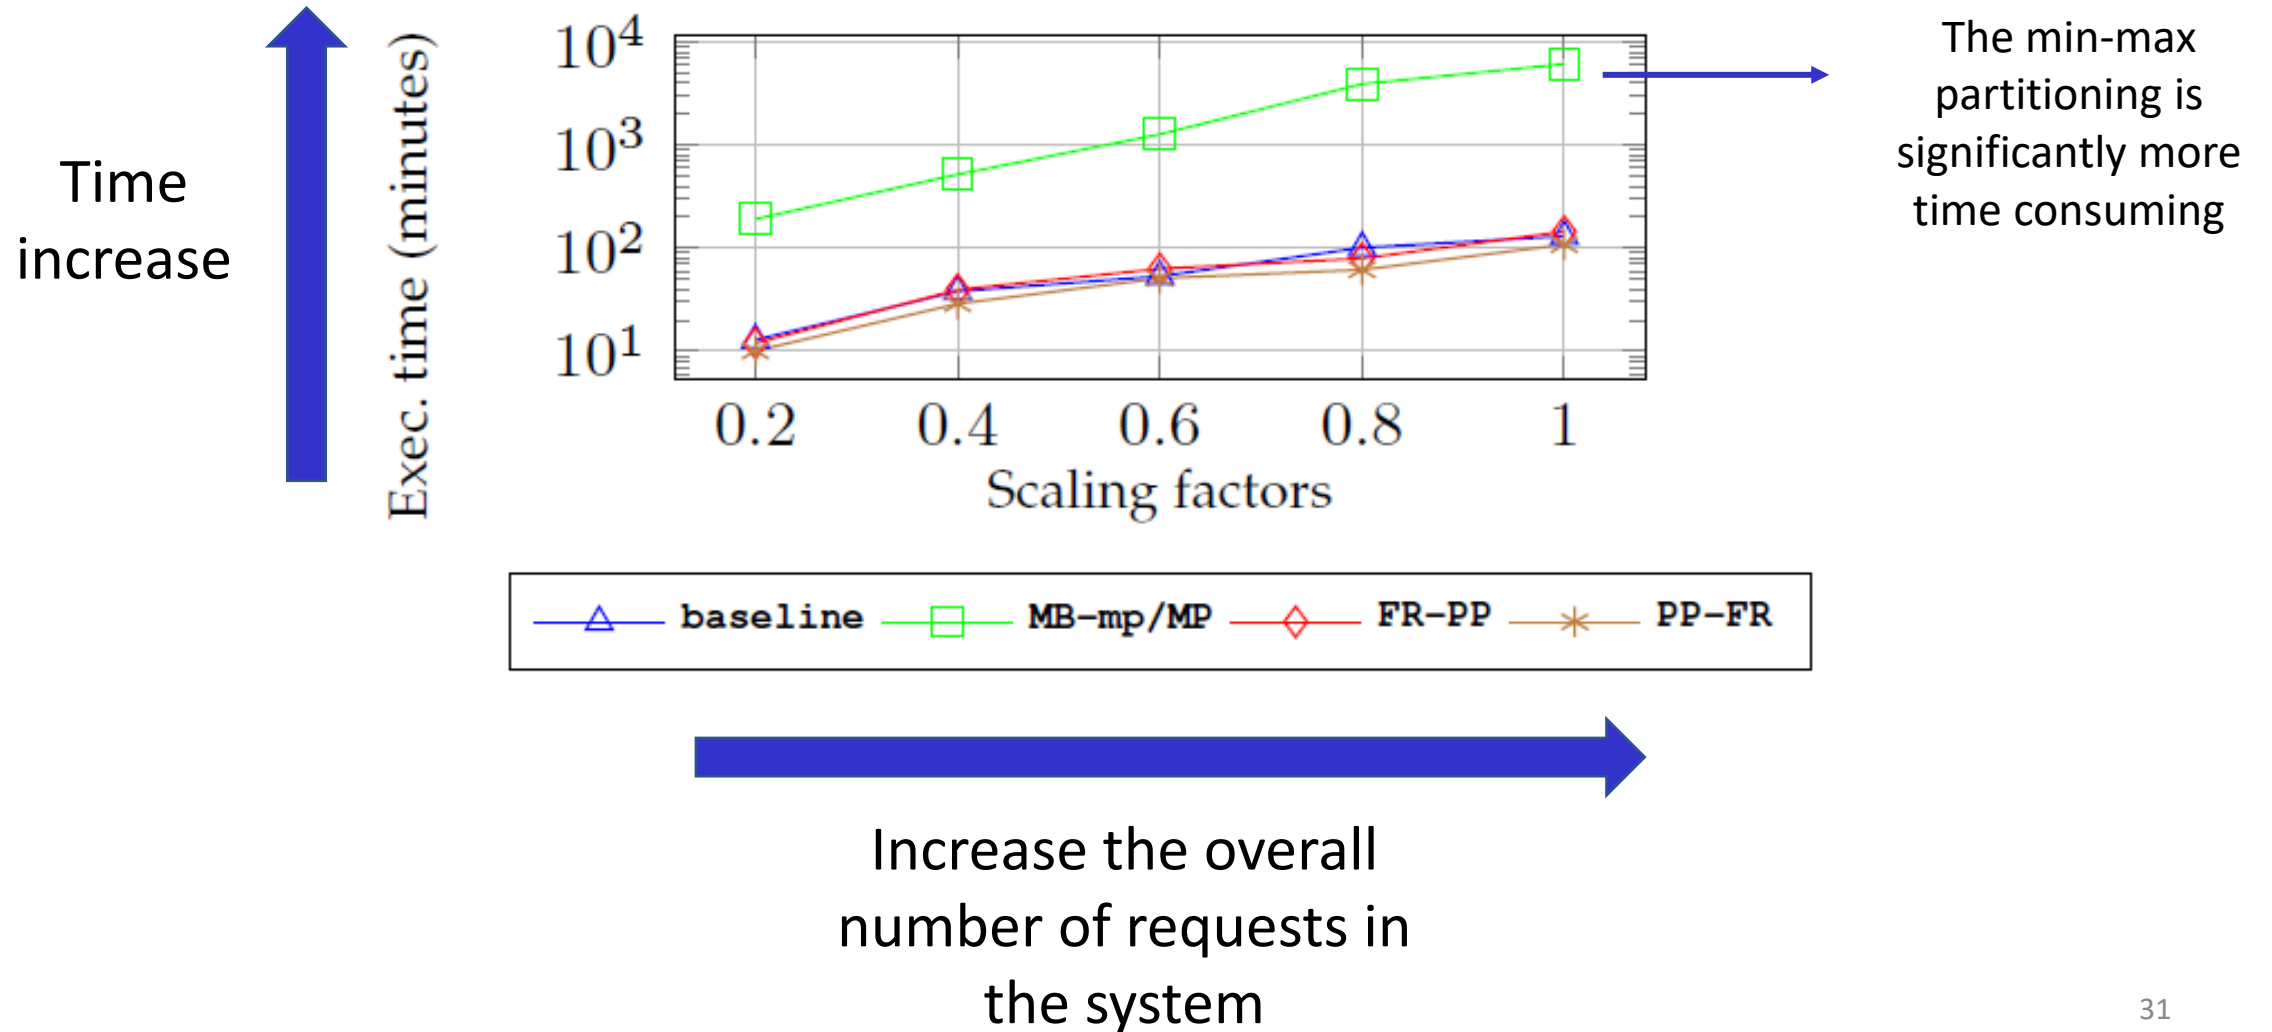
Increase the number of medium priority requests of interfering tasks

$$\beta = \frac{\#req \text{ medium prio}}{\#req \text{ total}}$$

Increase the number of high priority requests of interfering tasks

Scalability

We also assessed how the number of requests affects the execution time of the optimization problem



Conclusions

- **MPAM offers interesting tools to reduce memory contention**
- **The priority partitioning strategy can bring good results**
- **The min-max strategy offers very reduced benefits in the worst case**

Often, simpler is better for predictability

- **Analysis of other MPAM strategies**
- **Use more aggressive configuration of min-max strategy**
- **Evaluation of the strategies' effectiveness in the average-case scenario (in a simulated environment)**

Thank you!

Matteo Zini

matteo.zini@santannapisa.it