

CISTER SEMINAR SERIES

Analyzing the Contention on the shared memory bus for COTS-Based Multicores

Presented by

Dakshina Dasari

29 April 2011

Agenda



Motivation

Problem of
shared low
level
resources

Proposed
Method

Future work

COTS-based Multicores

Increasingly used in embedded systems

- Low power, high computing capabilities
- Faster to design and market

Finding WCET in multicores difficult

- COTS: Undocumented parameters
- COTS: Not predictable
- Shared resource contention (low-level)
- Uniprocessor theories developed not applicable

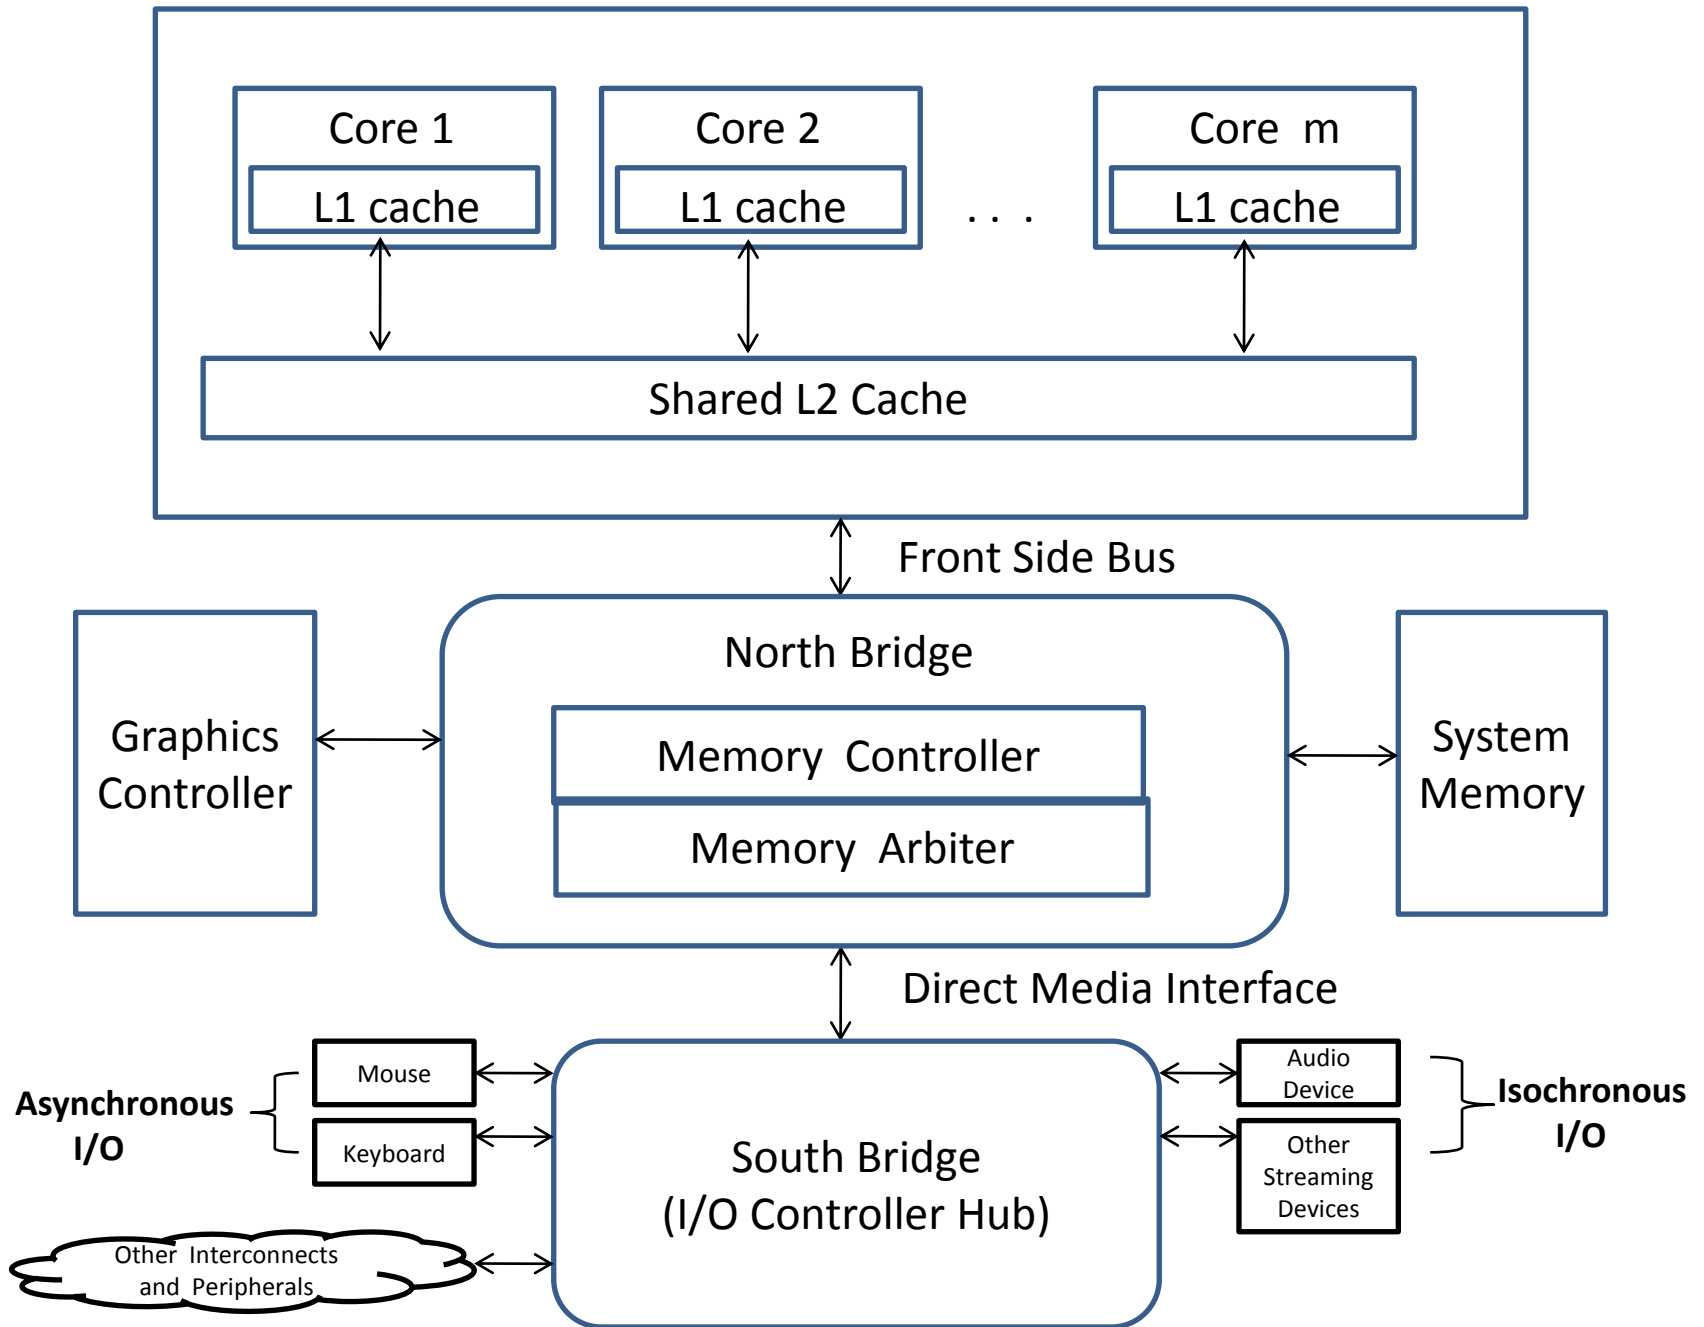
COTS: Commercial-off -the -shelf

Implications

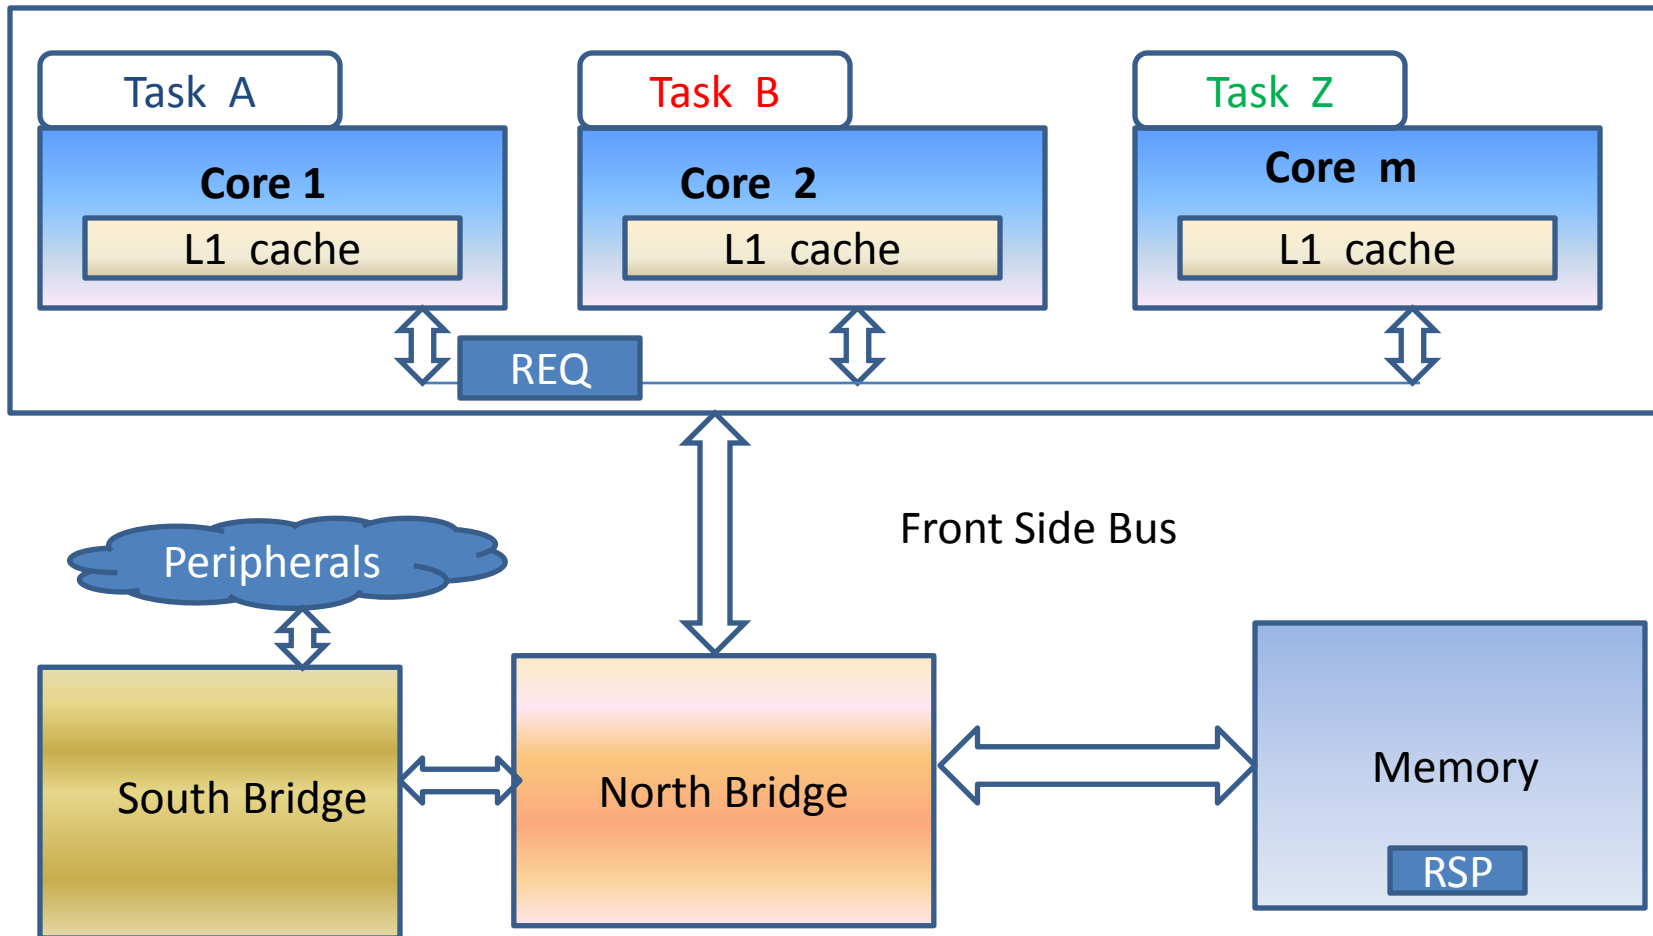
- Usage of very simple models in research
 - Do not reflect underlying hardware
- Generalized assumptions
 - Non tight WCET estimates

The industry trend does not seem to be towards building predictable systems ☹️

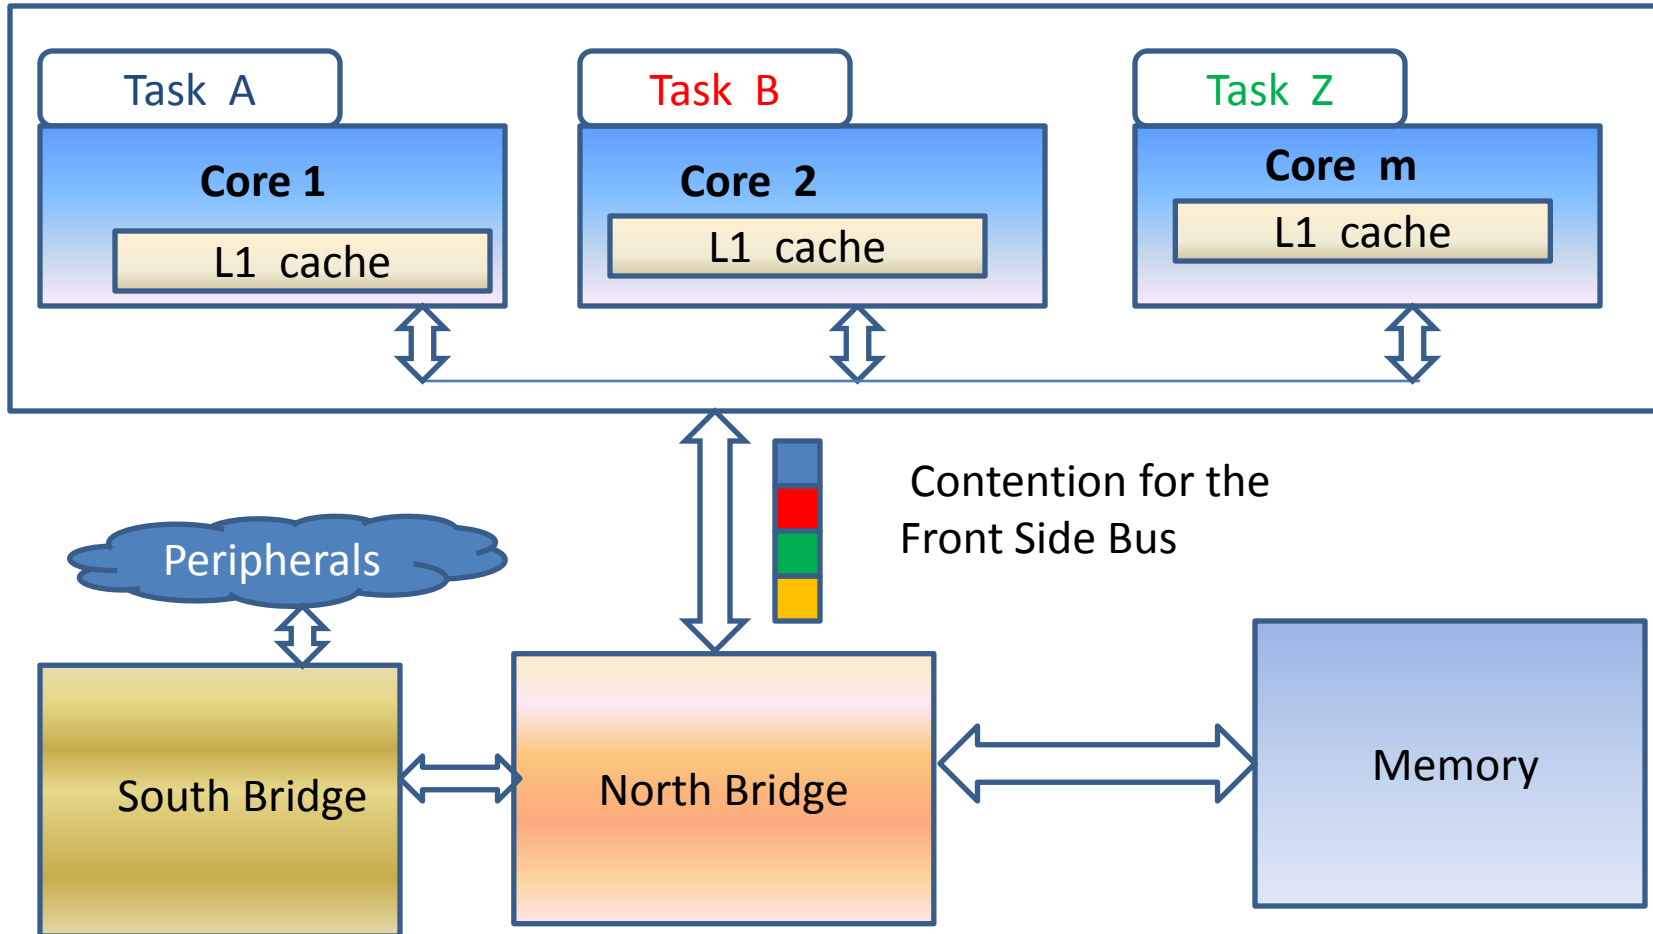
But performance oriented systems



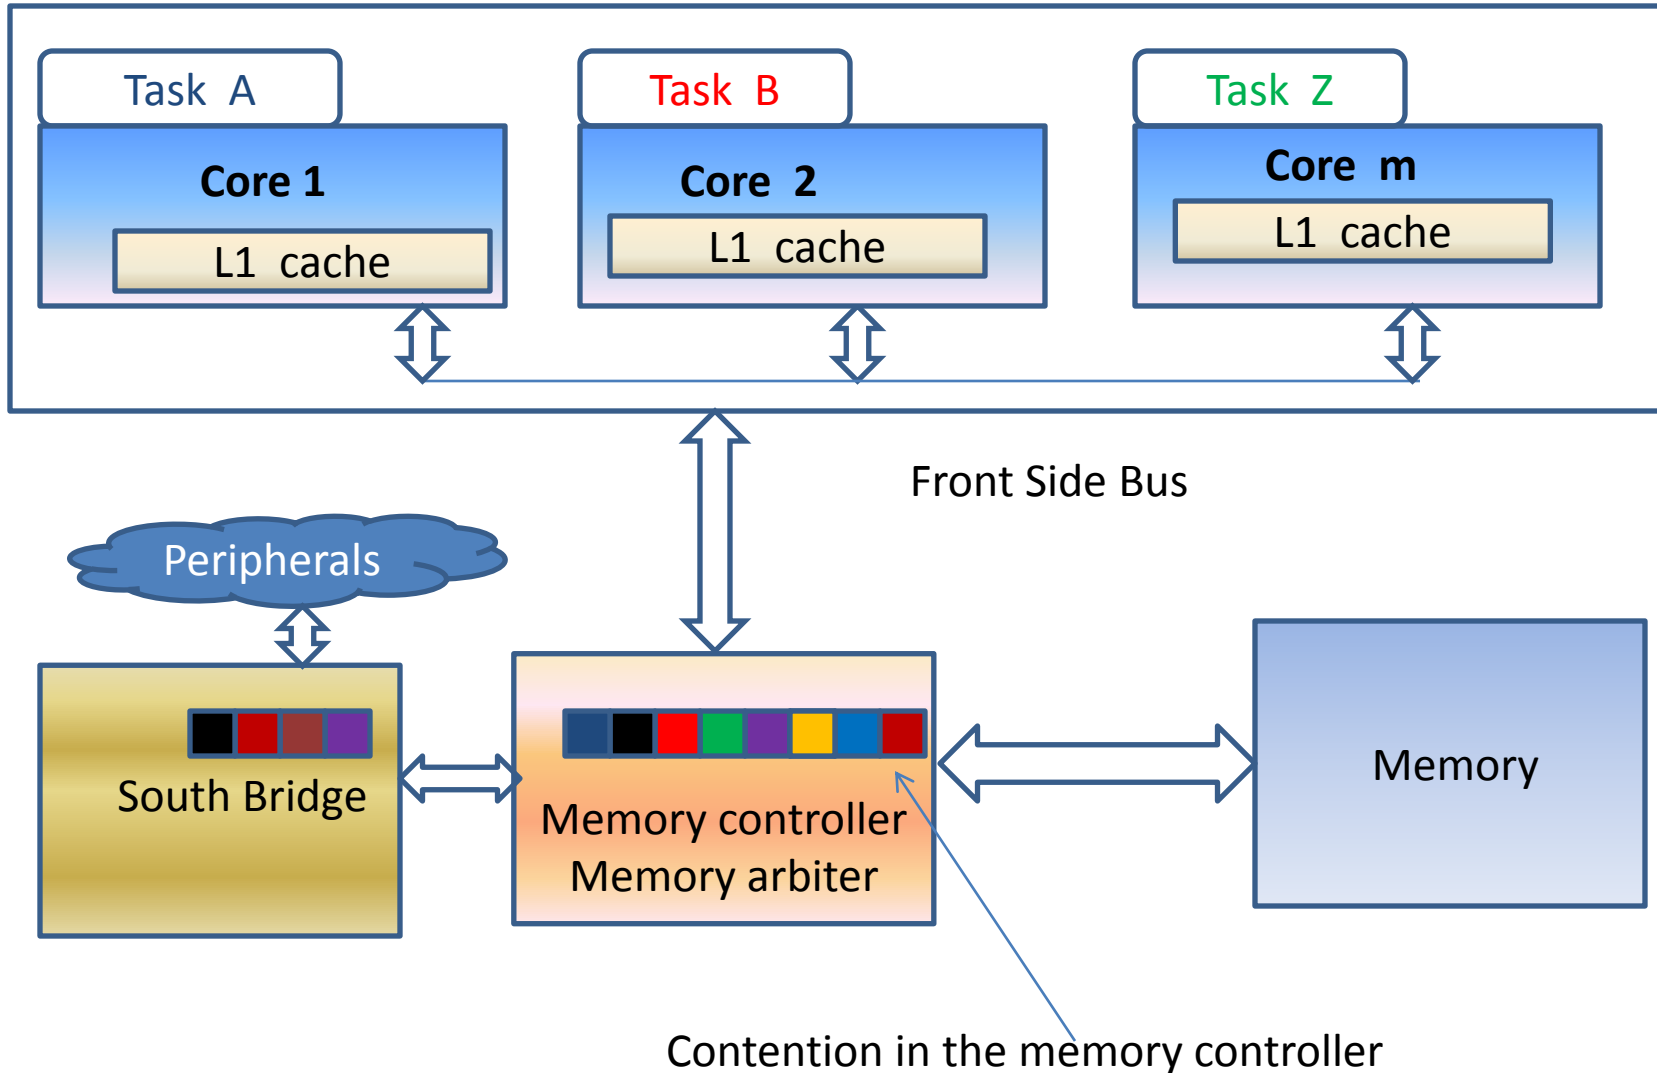
Shared resource contention



Shared resource contention



Shared resource contention



Nondeterminism in computing accurate WCET

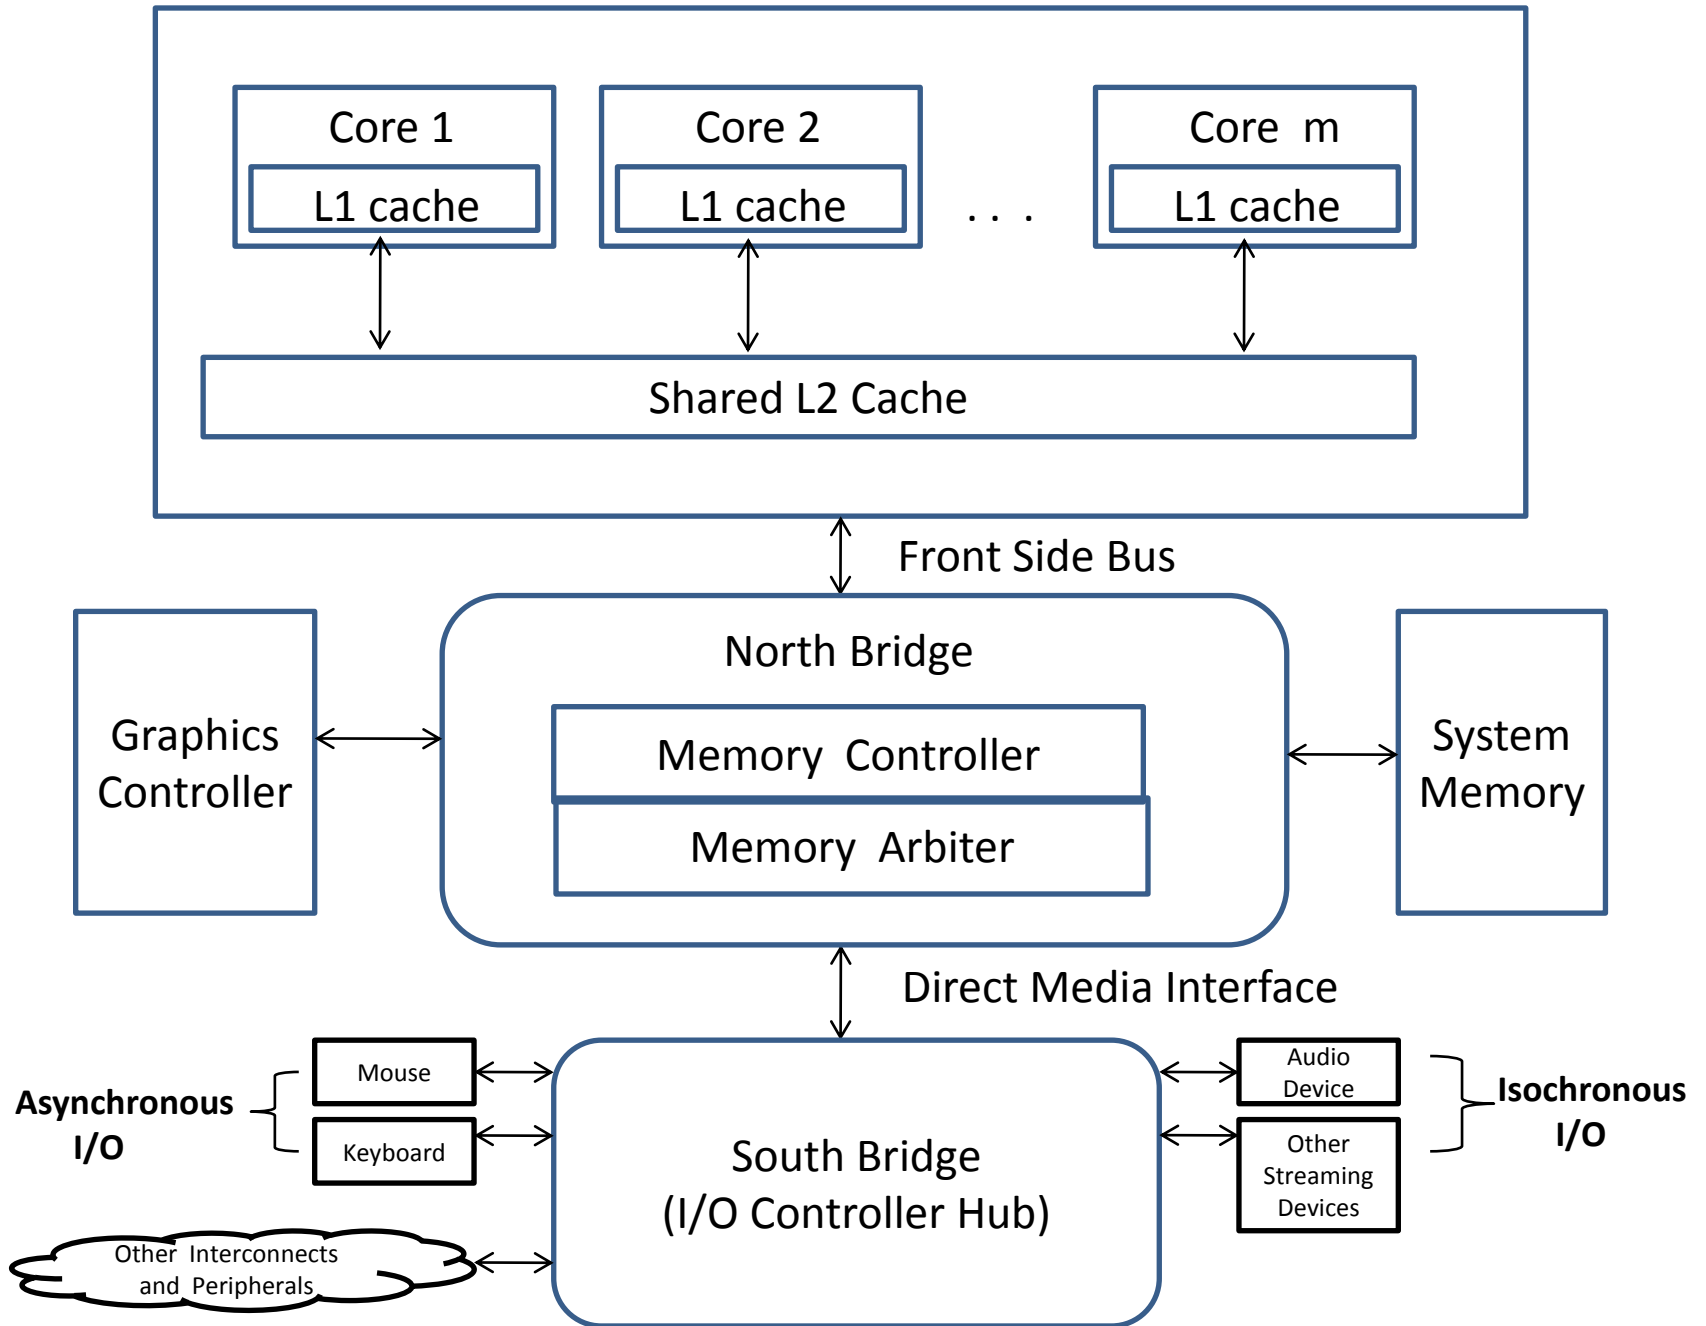
- Total Time for a request =
 T_FSB + // FSB contention
 T_FSB_NB + // transmission over FSB
 T_NB // NB contention
 T_NB_MEM + // tx time
 T_MEM + //memory access time
 T_MEM_NB + // tx time
 T_NB_FSB // tx time

Nontrivial to accurately determine : T_FSB, T_NB and T_MEM

Issues


Non-accuracy due to some undocumented parameters :

- Size of buffers in NB not stated
- Arbitration algorithm in the NB is vendor proprietary
- Memory access time is variable for each request and dependent on memory access scheduling techniques



Contention in the FSB

- Resolved using a Round Round Algorithm
 - (Disclaimer : wrt to intel processors)
 - Fairness : Order of transmission is fixed apriori (1-2-3-4-1)
 - Bus owner parks onto the bus until other owners assert the bus-request line
 - To Reduce switching overhead
 - Non-idling: A bus owner can keep transmitting when other cores do not transmit



Contention in the North Bridge

Request Type	Service slots (system cycles)
DRAM Maintenance Requests (Refresh)	X (High Priority = 1)
Display (Isochronous)	Y
Streaming (Isochronous)	Z
CPU (Asynchronous)	W
	Total = N system cycles

- Schedule period repeats after every N cycles
- Flexible, Slot based mechanism
- Tries to meet QoS requirements of Isochronous (Periodic requests) with low-latency requirements of Asynchronous requests

Contention in the North Bridge

Request Type	Service slots (system cycles)
DRAM Maintenance Requests (Refresh)	X (High Priority = 1)
Display (Isochronous)	Y
Streaming (Isochronous)	Z
CPU (Asynchronous)	W
	Total = N system cycles

- Flexible but non-predictable
- Weights assigned to request types not specified
- Difficult to accurately compute an upper-bound

Nondeterminism in computing accurate WCET

- Total Time for a request =
 T_FSB + // FSB contention
 T_FSB_NB + // transmission over FSB
 T_NB // NB contention
 T_NB_MEM + // tx time
 T_MEM + //memory access time
 T_MEM_NB + // tx time
 T_NB_FSB // tx time

Nontrivial to accurately determine : T_FSB, T_NB and T_MEM


Summary of the discussion

- Method to obtain maximum time to service a request (TR) by adding individual factors difficult
- Workaround :
 - Measure end to end latency for a large number of requests
 - Record the maximum value
 - Use this value for WCET estimation


Problem Definition

- Compute the WCET of a task, considering contention on the bus on , given the following :
- WCET in isolation
- A multicore system with
 - Private caches : Cores do not share the cache
 - Shared front side bus with **Round Robin Bus Arbitration Algorithm**
- Task model
 - Non pre-emptive (Tasks run uninterrupted)
 - No Cache Related Pre-emption Delay and context switch overhead
 - Constrained deadline ($D_i \leq T_i$) Periodic tasks
 - Partitioned scheduling (Tasks do not migrate)

Round Robin algorithm



Round Robin algorithm



$$C_i^{\text{mix}} = C_i^{\text{iso}}$$


where

TR: Time to serve a request


C_i^{iso} :WCET in isolation

C_i^{mix} :WCET when run with other tasks

Round Robin algorithm



Round Robin algorithm



Blocked by 7 requests

$$C_i^{\text{mix}} = C_i^{\text{iso}} + 7 * \text{TR}$$


where

TR: Time to serve a request

C_i^{iso} :WCET in isolation

C_i^{mix} :WCET when run with other tasks

Round Robin algorithm



No blocking from other tasks




Blocked by 7 requests



Worst case Blocked by 9 requests

Round Robin algorithm



Worst case Blocked by 9 requests

$$C_i^{\text{mix}} = C_i^{\text{iso}} + 9 * TR$$


where

TR: Time to serve a request

C_i^{iso} : WCET in isolation

C_i^{mix} : WCET when run with other tasks

Round Robin algorithm



Worst case Blocked by 9 requests

$$C_i^{\text{mix}} = C_i^{\text{iso}} + 9 * TR$$

where

TR: Time to serve a request

C_i^{iso} : WCET in isolation


C_i^{mix} : WCET when run with other tasks

m : number of cores

$RQST_i(t)$: Requests generated by task i in time t

$$C_i^{\text{mix}} = C_i^{\text{iso}} + RQST_i(C_i^{\text{iso}}) * (m-1) * TR$$

Round Robin algorithm




Worst case Blocked by #Max = $(m-1) * RQST_i(C_i^{iso})$ requests

$$C_i^{mix} = C_i^{iso} + RQST_i(C_i^{iso}) * (m-1) * TR$$

Very pessimistic !!

- Tasks on other cores may not generate #Max requests
- There may be no tasks scheduled on the other cores

Round Robin algorithm



Worst case Blocked by Max = $(m-1) * RQST_i(C_i^{iso})$
requests

$$C_i^{mix} = C_i^{iso} + RQST_i(C_i^{iso}) * (m-1) * TR$$

Very pessimistic !!

- Ex : Task i generates 2000 requests $RQST_i(C_i^{iso}) = 2000$
- Co-scheduled tasks on other cores generate 20 requests
- By the bound $C_i^{mix} = C_i^{iso} + 2000 * (3) * TR$
- **Actual** : $C_i^{mix} = C_i^{iso} + 20 * TR$

Round Robin algorithm

Pessimistic bound:

$$C_i^{\text{mix}} = C_i^{\text{iso}} + \text{RQST}_i(C_i^{\text{iso}}) * (m-1) * \text{TR}$$

For tighter WCET bounds we need:

$$C_i^{\text{mix}} = C_i^{\text{iso}} + \text{Requests_from_other_cores} * \text{TR}$$

(during execution of task i)

We need a **Per-Core Request Estimator Function**

Round Robin algorithm

$$C_i^{\text{mix}} = C_i^{\text{iso}} + \text{Requests_from_other_cores} * \text{TR}$$

(during execution of task i)

We need a **Per-Core Request Estimator Function**

PCRE_j(t) : Returns maximum number of requests generated by tasks scheduled on core 'j' during time interval 't'

The Method

Task A



ADVERSARY



Core 1

Core 2

Core 3

Core 4

Interference queue of length = $(m-1) * RQST_A(C_A^{iso}) = 3 * 3 = 9$ slots



At time 0

$$PCRE_2(0) = PCRE_3(0) = PCRE_4(0) = 0$$

$m = 4$ cores

$$C_A^{iso} = 4 \quad TR = 0.05$$

$$C_A^{mix} = 4$$

The Method

Task A



ADVERSARY



Core 1

Core 2

Core 3

Core 4

Interference queue of length = $(m-1) * RQST_A(C_A^{iso}) = 3 * 3 = 9$ slots



At time 0

$$PCRE_2(0) = PCRE_3(0) = PCRE_4(0) = 0$$

$m = 4$ cores


$$C_A^{iso} = 4 \quad TR = 0.05$$

$$C_A^{mix} = 4$$

The Method



Core 1



Core 2

Core 3

Core 4



At time = 4

$$PCRE_2(4) = 1 \quad PCRE_3(4) = 1 \quad PCRE_4(4) = 2$$

$m = 4$ cores

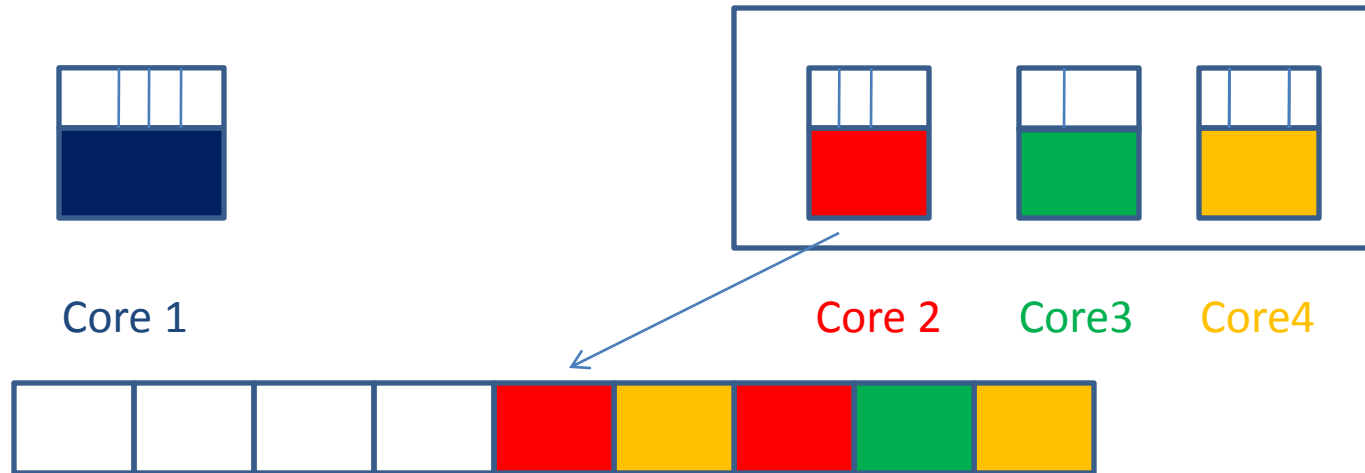
$$C_A^{iso} = 4 \quad TR = 0.05$$

$$C_A^{mix} = 4 + 4 * 0.05 \\ = 4.20$$



Execution time of task A increases

The Method



At time = 4.20

Δ : Increased execution time

$$PCRE_2(\Delta) = 0 \quad PCRE_3(\Delta) = 1 \quad PCRE_4(\Delta) = 0$$

$$C_A^{iso} = 4 \quad TR = 0.05$$

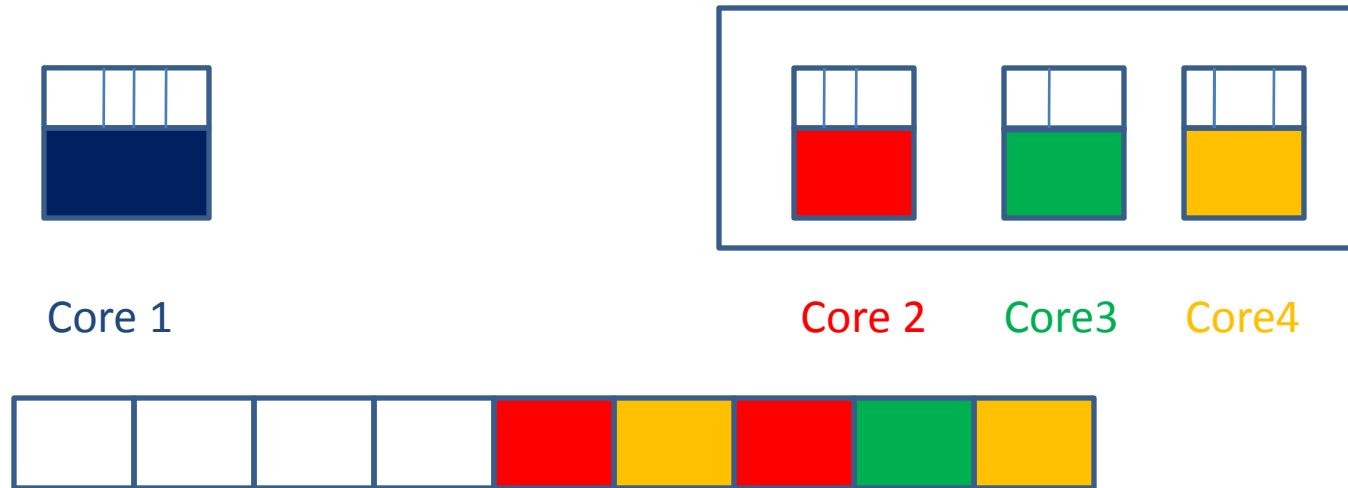
$$C_A^{mix} = 4.20 + 1 * 0.05$$

$$= 4.25$$

Execution time of task A further increases



The Method



At time = 4.25

Δ : Increased execution time

$PCRE_2(\Delta) = 0$ $PCRE_3(\Delta) = 0$ $PCRE_4(\Delta) = 0$

No more requests from other cores !!!

$C_A^{iso} = 4$ $TR = 0.05$

$C_A^{mix} = 4.25 + 0$
 $= 4.25$

Final wcet = 4.25

The algorithm

Initialization Step

$$C_i^0 = C_i^{\text{iso}}$$

$$iqlen_i^0 = RQST_i(C_i^0) * (m-1)$$

$$\text{external_rqst}_i^0 = \sum_{j \neq \pi(i)} PCRE_j(C_i^0)$$

$$\text{blocking_rqst}_i^0 = \min(iqlen_i^0, \text{external_rqst}_i^0)$$

Iteration Step

$$C_i^k = C_i^{k-1} * \text{blocking_rqst}_i^{k-1} * TR$$

$$iqlen_i^k = iqlen_i^{k-1} - \text{blocking_rqst}_i^{k-1}$$

$$\text{external_rqst}_i^k = \sum_{j \neq \pi(i)} (PCRE_j(C_i^k) - PCRE_j(C_i^{k-1}))$$

$$\text{blocking_rqst}_i^k = \min(iqlen_i^{k-1}, \text{external_rqst}_i^{k-1})$$

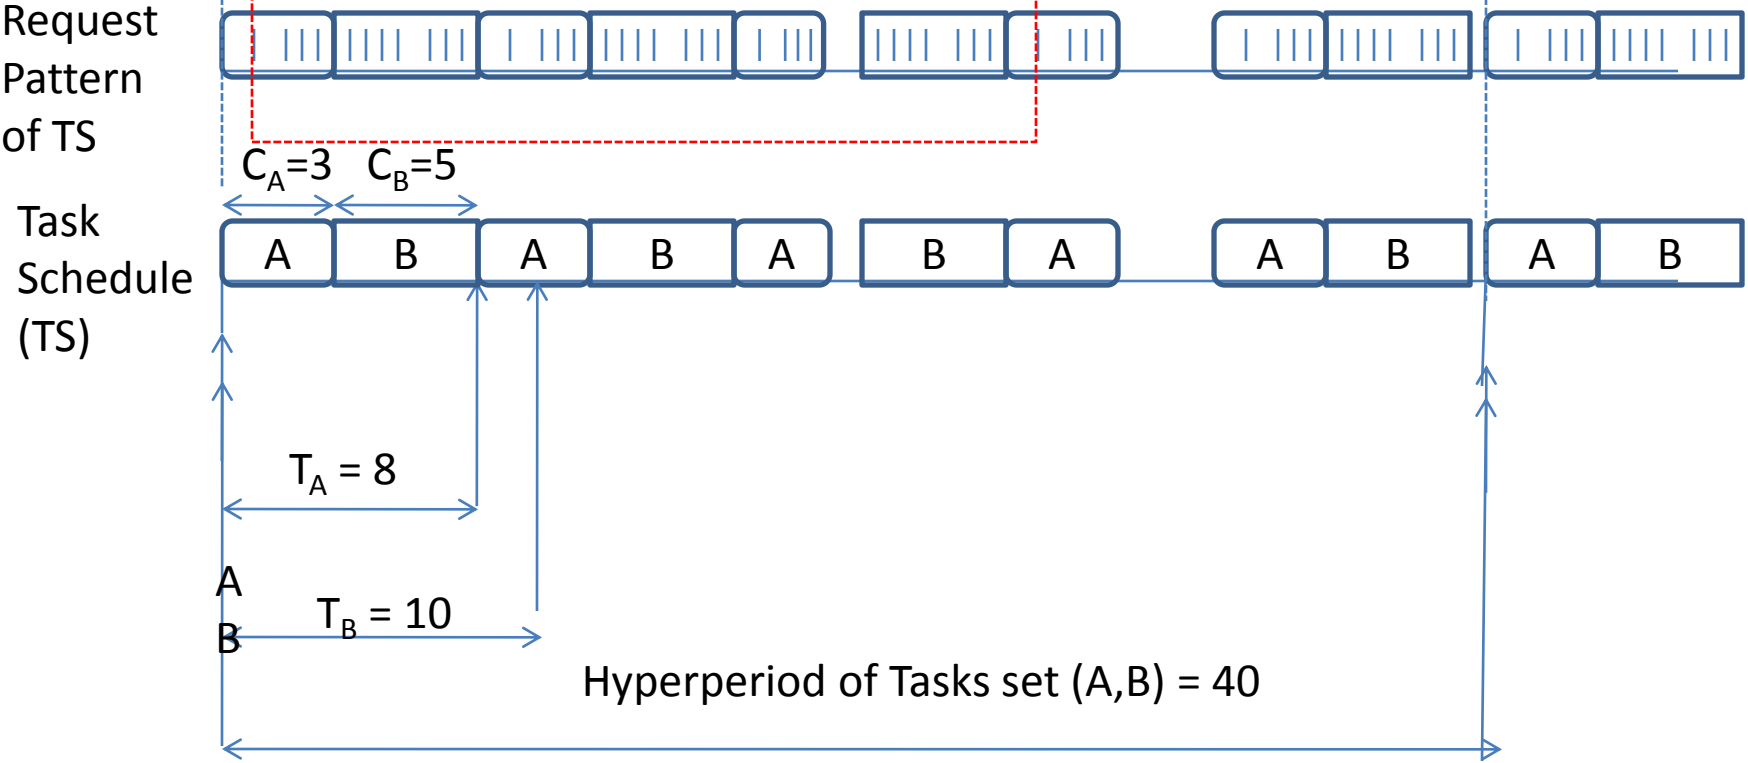
Stopping Conditions :

$$iqlen_i^k = 0 \quad \text{RR Upper bound reached}$$

$$\text{blocking_rqst}_i^k = 0 \quad \text{No more requests from other cores}$$

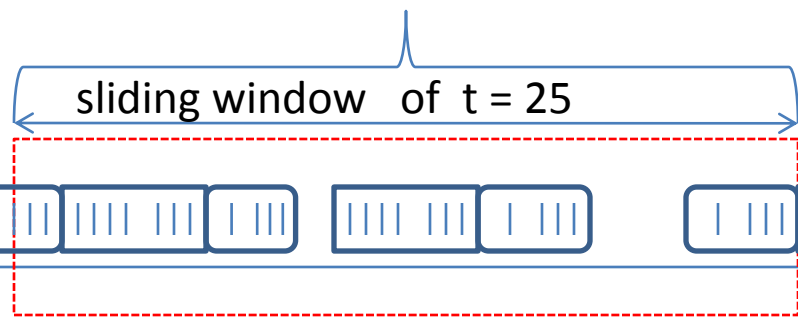
$$R_A=4 \quad R_B=7$$

Schedule repeats hereon



Per Core Request Estimator

$R_A=4$ $R_B=7$



Schedule repeats hereon

Request Pattern of TS

$C_A=3$ $C_B=5$

Task Schedule (TS)



$T_A = 8$

$T_B = 10$

A
B

Hyperperiod of Tasks set (A,B) = 40

Per Core Request Estimator

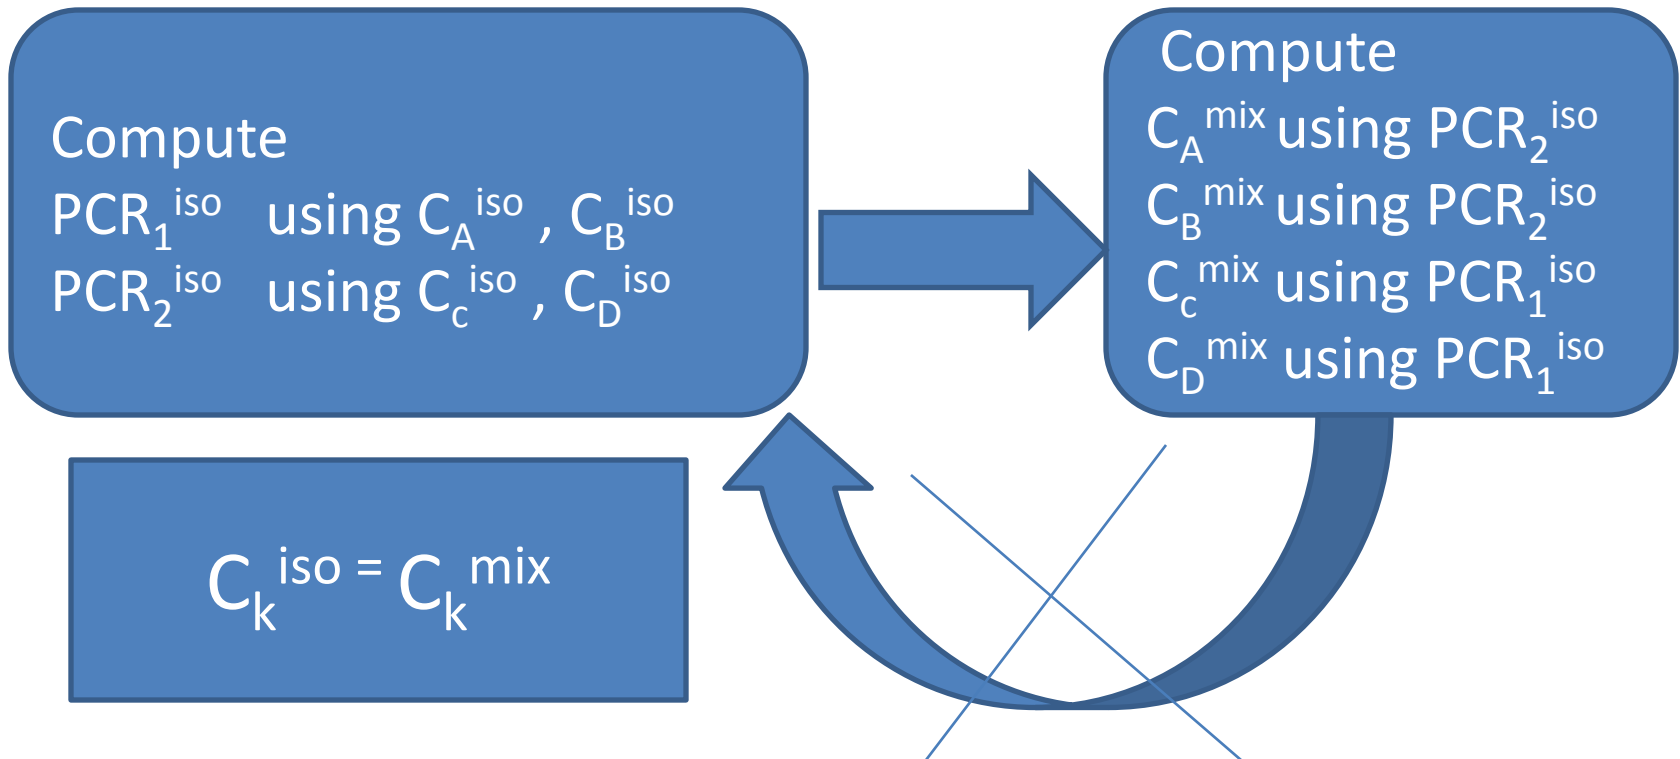
Obtaining request patterns

- PCRE(t) depends on the exact request pattern of the tasks
 - **Measurements**
 - Performance monitoring counters
 - Special purpose registers in microprocessors
 - Reset counter, select event (like L1 cache misses)
 - Code block to be monitored
 - Stop counter, Read values
 - Static analysis

System wide analysis

Tasks A, B assigned to core 1

Tasks C, D assigned to core 2



System wide analysis

Why:

Task A



Private caches
No extra cache misses

Non preemptive tasks
No extra cache misses

Increase only due to bus
contention

Increased execution does not increase
number of requests generated


Request density decreases .

Therefore $PCR_i^{iso}(t) \geq PCR_i^{mix}(t)$

System wide analysis

Tasks A, B assigned to core 1

Tasks C, D assigned to core 2



Other Related work carried out

- Response time analysis for an arbitration-agnostic bus contention algorithm for COTS-based systems
- Measurement based framework for generating a request profile for a task
 - Uses performance monitoring counters
- Preliminary shared cache analysis

Future Work

- Reducing Bus Contention with resource-aware schedulers
- Addressing contention considering shared caches
- Addressing Pre-emptive task models