NC STATE UNIVERSITY

Safely Exploiting Multithreaded Processors to Tolerate Memory Latency in Real-Time Systems

Ali El-Haj-Mahmoud and Eric Rotenberg

Center for Embedded Systems Research (CESR)
Electrical & Computer Engineering
North Carolina State University

CESR

Embedded Processor Trends



- More demanding applications and user expectations
- Higher frequency
  - ARM-11 (0.13µ): 500 MHz
  - ARM-11 (0.10µ): 1 GHz
- Processor-memory speed gap

Memory Wall

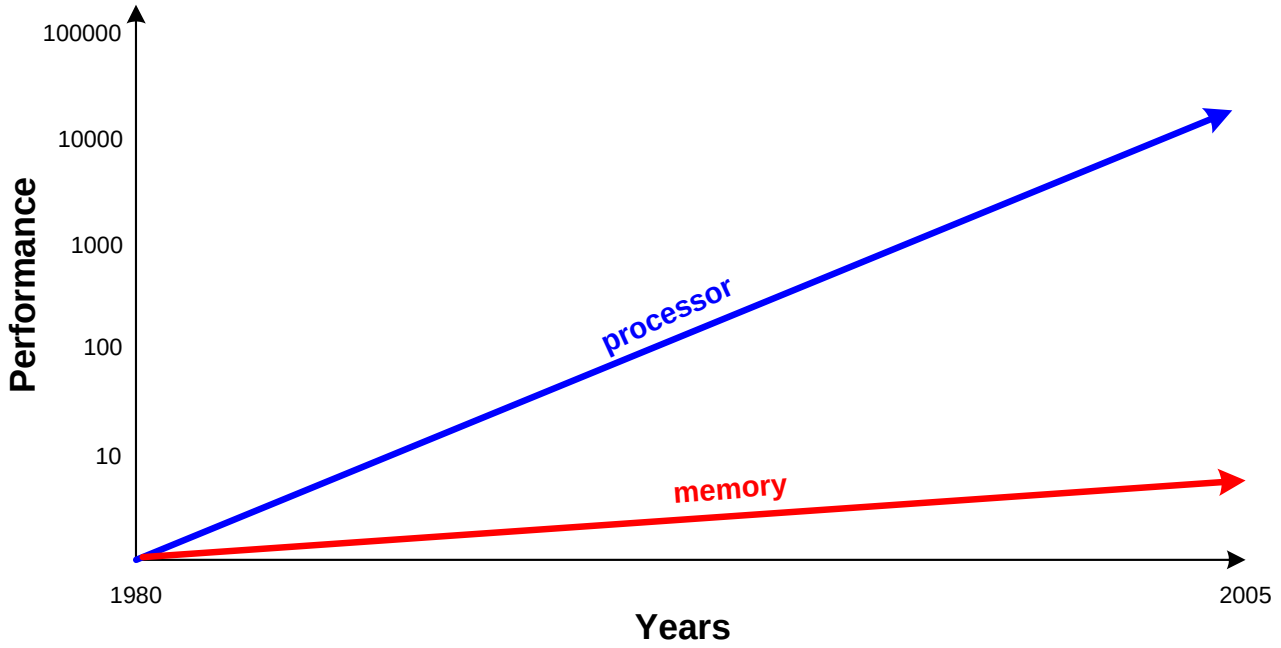

How to capitalize on higher frequency?

Multithreading

- *Switch-on-event coarse-grain multithreading*
  - Multiple register contexts for fast switching
  - Switch to alternate task when current task accesses memory
  - Overlap memory accesses with computation
- But what about multithreading in hard-real-time?
  - Require *analyzability*
  - Cannot rely on dynamic schemes

Exploiting Multithreading in Hard-Real-Time Systems

- Safety
  - Guarantee all tasks meet deadlines (worst-case)
  - Statically bound overlap under all scenarios
- Tractability
  - Confirm/disconfirm schedulability mathematically using closed-form tests
  - Consider each task individually

Classic Real-Time Scheduling

Example: Earliest Deadline First (EDF)

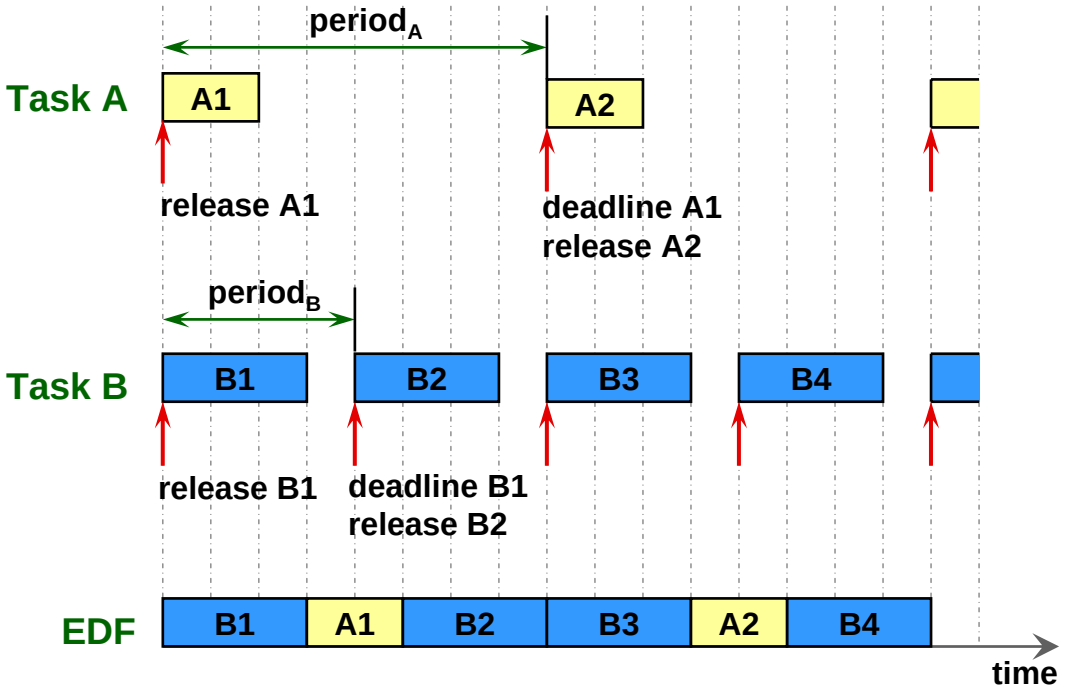

$$U_A = \frac{WCET_A}{period_A} = \frac{2}{8} = 0.25$$

$$U_B = \frac{WCET_B}{period_B} = \frac{3}{4} = 0.75$$

EDF schedulability test

$$U = \sum_i \frac{WCET_i}{period_i} \leq 1$$

- Utilization-based schedulability test
- No need to **construct** the schedule *a priori*

Performance vs. Tractability

- Performance: Memory overlap
- Tractability: Closed-form schedulability test

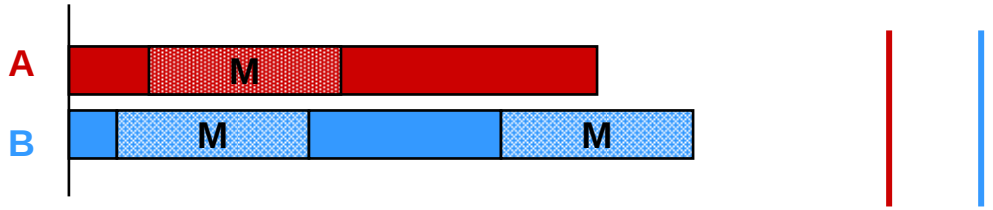

- Only basic parameters of tasks are known
  - WCET = C + M (from conventional WCET analysis)
  - Period = deadline

EDF (Classic Real-Time)

Low Performance but Tractable

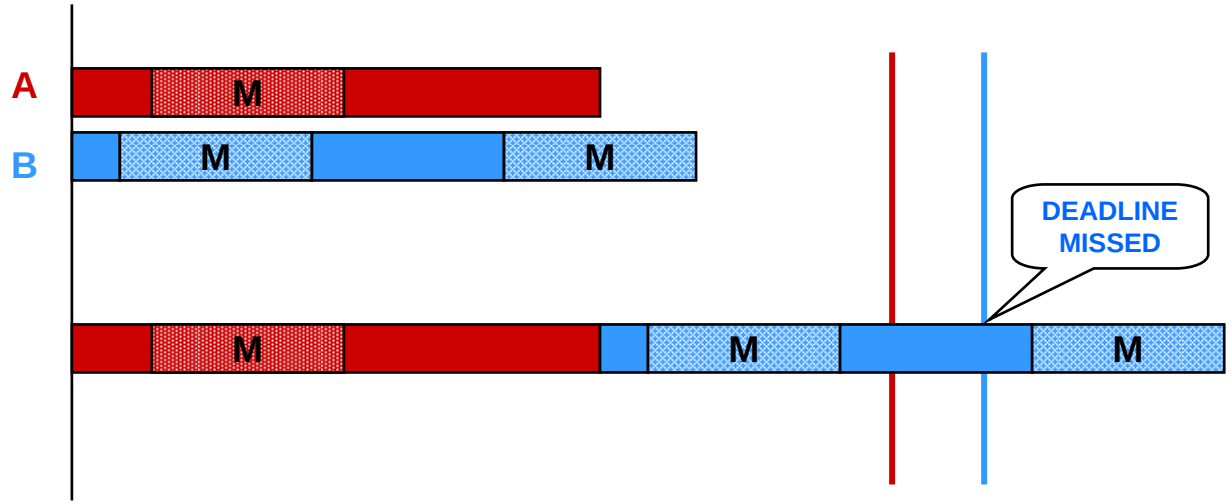

- Memory overlap: No
- Closed-form test: Yes

$$U = \frac{WCET_A}{period_A} + \frac{WCET_B}{period_B} > 1$$

EDF (Classic Real-Time)

Low Performance but Tractable

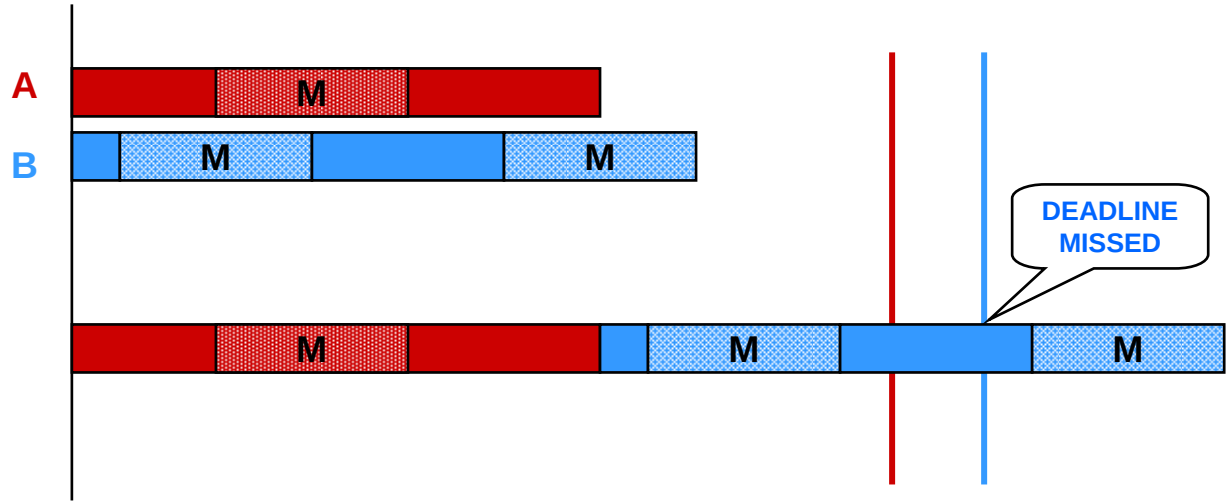

- Memory overlap: No
- Closed-form test: Yes

$$U = \frac{WCET_A}{period_A} + \frac{WCET_B}{period_B} > 1$$

EDF (Classic Real-Time)

Low Performance but Tractable

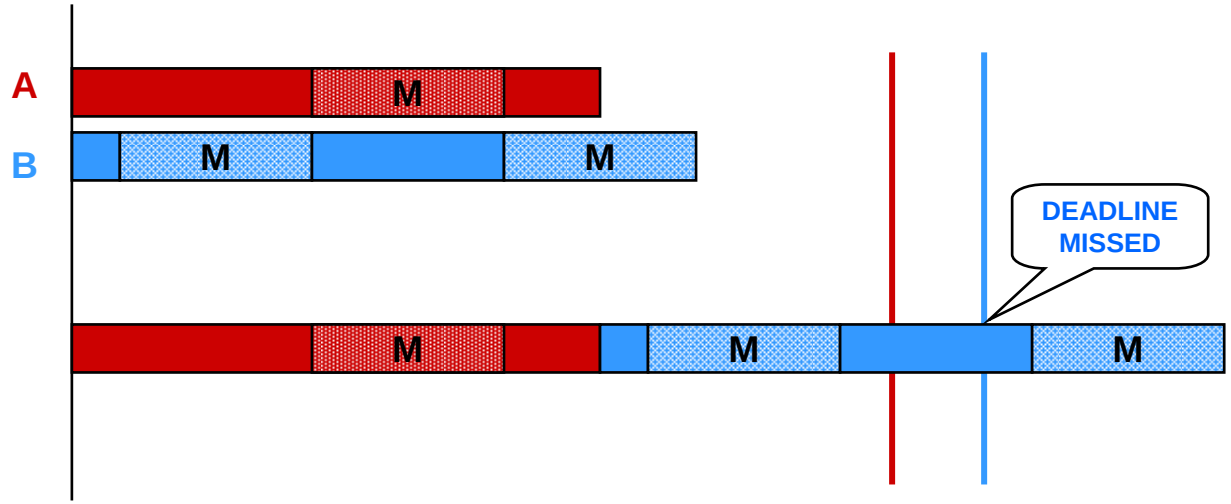

- Memory overlap: No
- Closed-form test: Yes

$$U = \frac{WCET_A}{period_A} + \frac{WCET_B}{period_B} > 1$$

EDF (Classic Real-Time)

Low Performance but Tractable

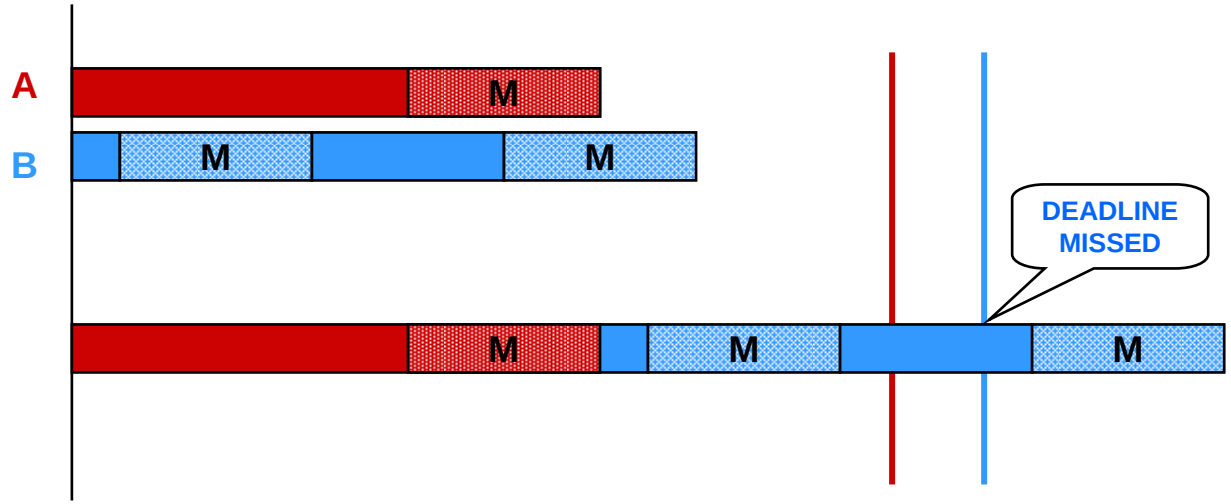

- Memory overlap: No
- Closed-form test: Yes

$$U = \frac{WCET_A}{period_A} + \frac{WCET_B}{period_B} > 1$$

Dynamic Switch on Memory Access
Possibly High Performance but Intractable

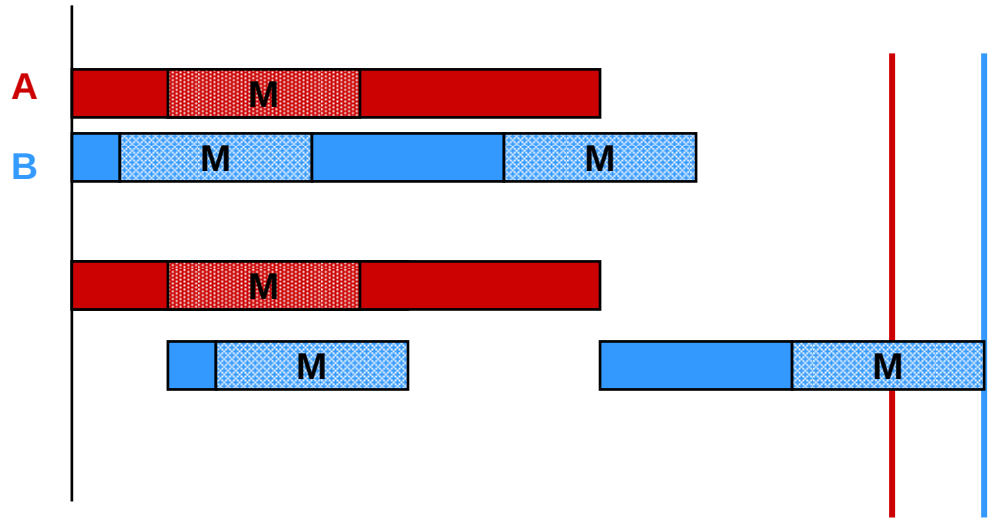

- Memory overlap: Possible, unfair for low priority
- Closed-form test: No
  - Must examine memory positioning!

Dynamic Switch on Memory Access
Possibly High Performance but Intractable

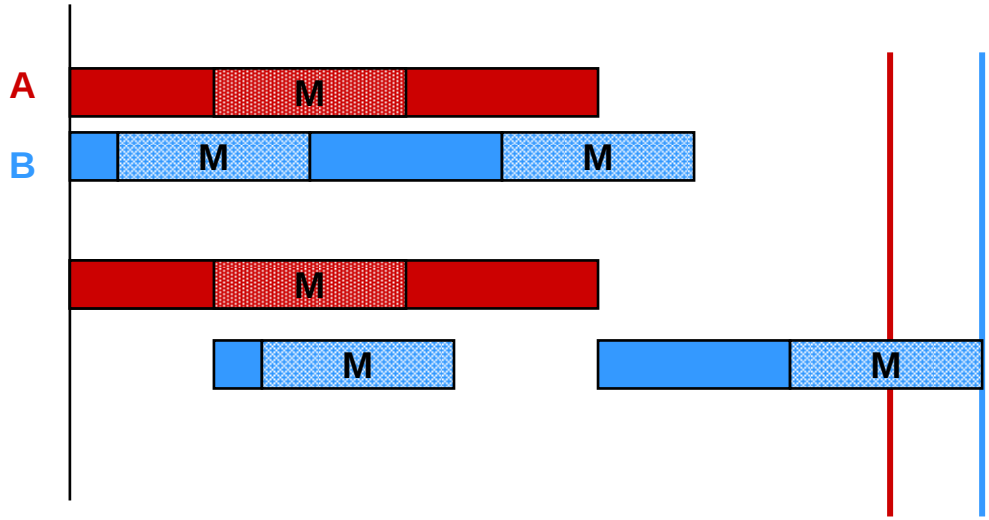

- Memory overlap: Possible, unfair for low priority
- Closed-form test: No
  - Must examine memory positioning!

Dynamic Switch on Memory Access
Possibly High Performance but Intractable

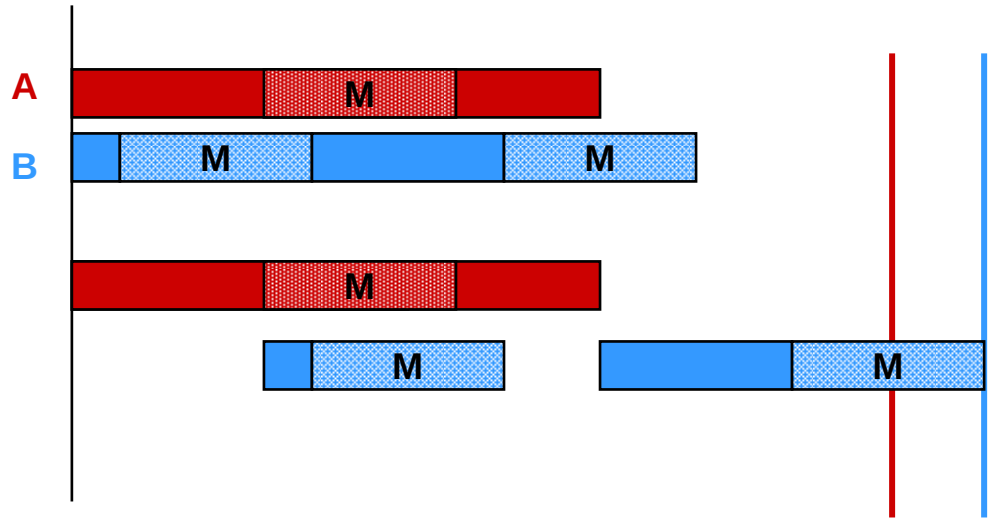

- Memory overlap: Possible, unfair for low priority
- Closed-form test: No
  - Must examine memory positioning!

Dynamic Switch on Memory Access
Possibly High Performance but Intractable

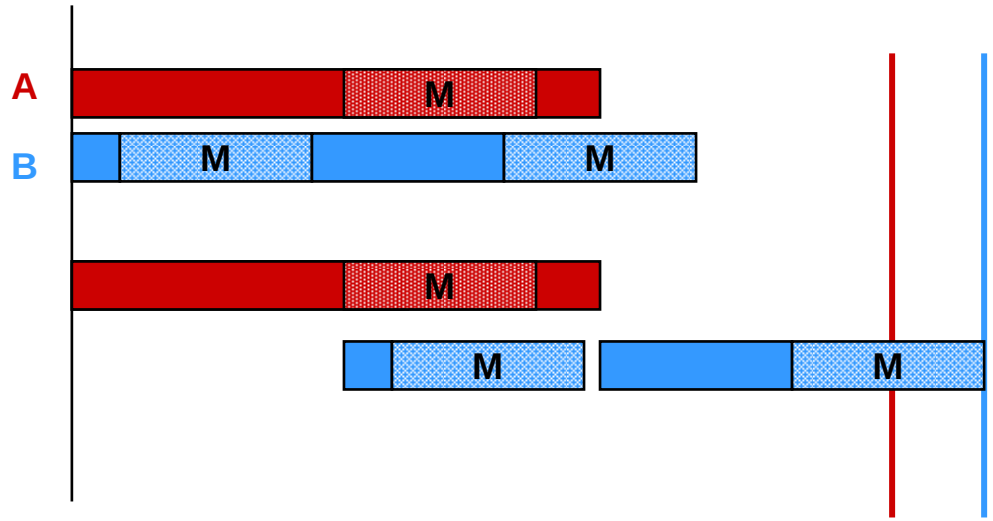

- Memory overlap: Possible, unfair for low priority
- Closed-form test: No
  - Must examine memory positioning!

Dynamic Switch on Memory Access
Possibly High Performance but Intractable

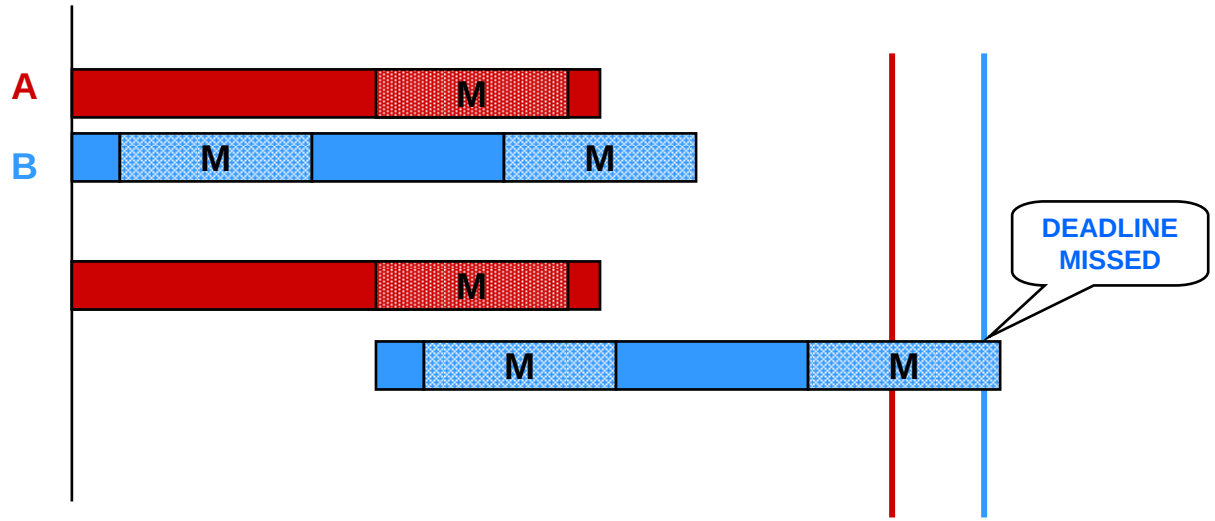

- Memory overlap: Possible, unfair for low priority
- Closed-form test: No
  - Must examine memory positioning!

Dynamic Switch on Memory Access
Possibly High Performance but Intractable

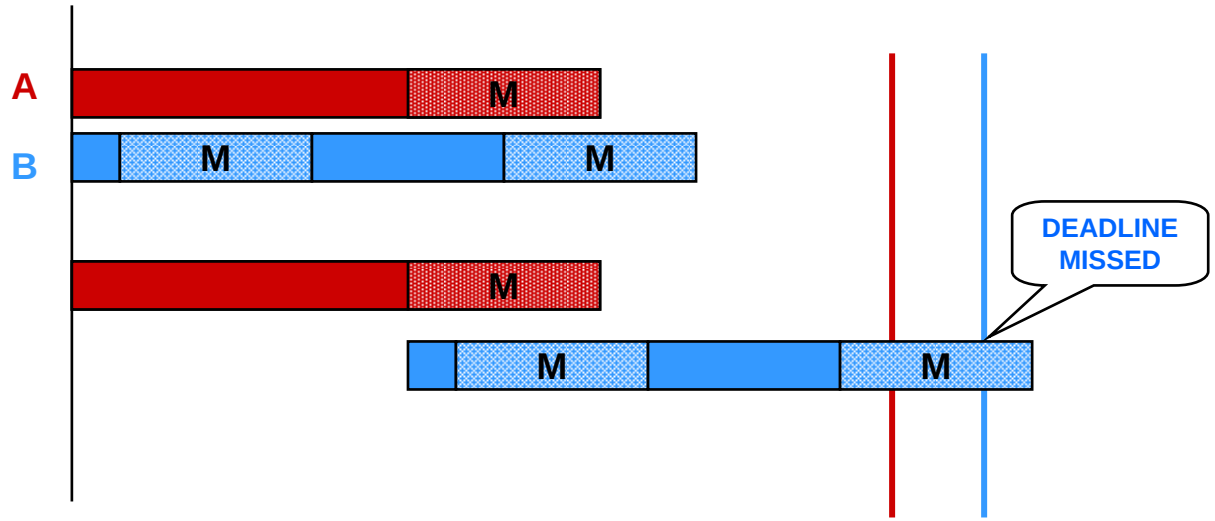

- Memory overlap: Possible, unfair for low priority
- Closed-form test: No
  - Must examine memory positioning!

Deterministic Switching

High Performance and Tractable

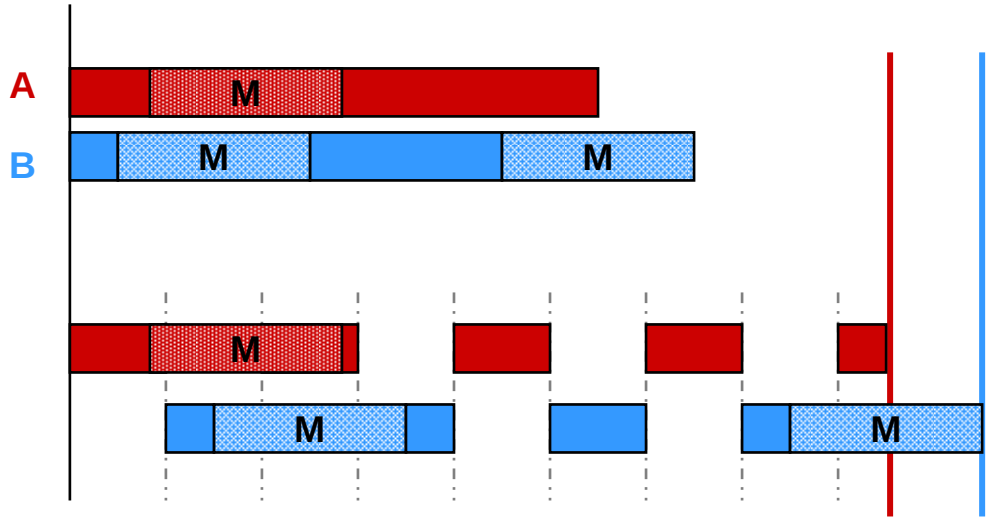

- Memory overlap: Yes (fair and bounded)
- Closed-form test: Yes

Deterministic Switching

High Performance and Tractable

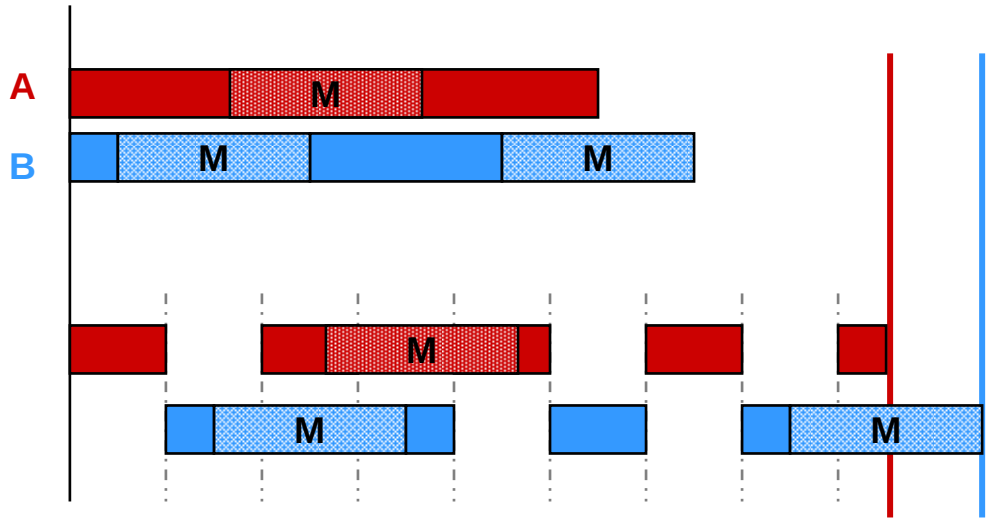

- Memory overlap: Yes (fair and bounded)
- Closed-form test: Yes

Deterministic Switching

High Performance and Tractable

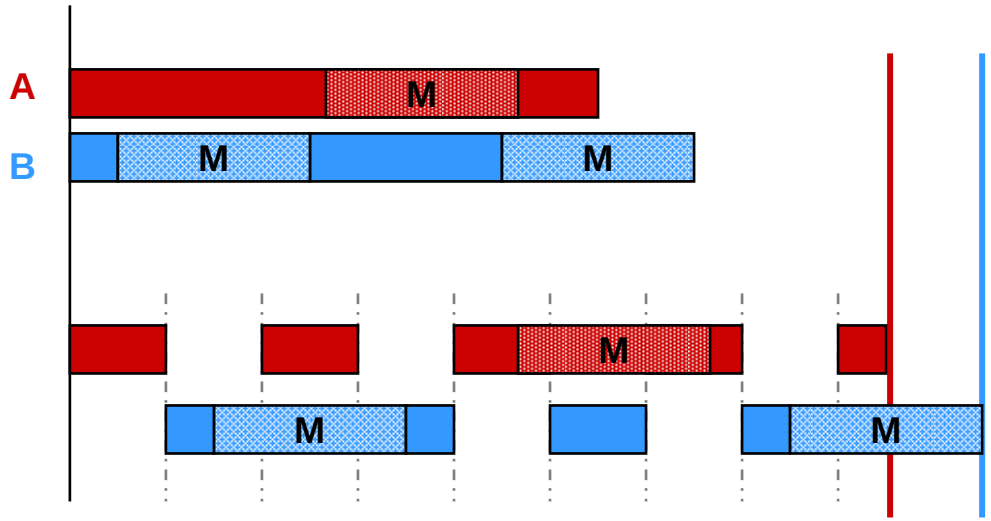

- Memory overlap: Yes (fair and bounded)
- Closed-form test: Yes

Deterministic Switching

High Performance and Tractable

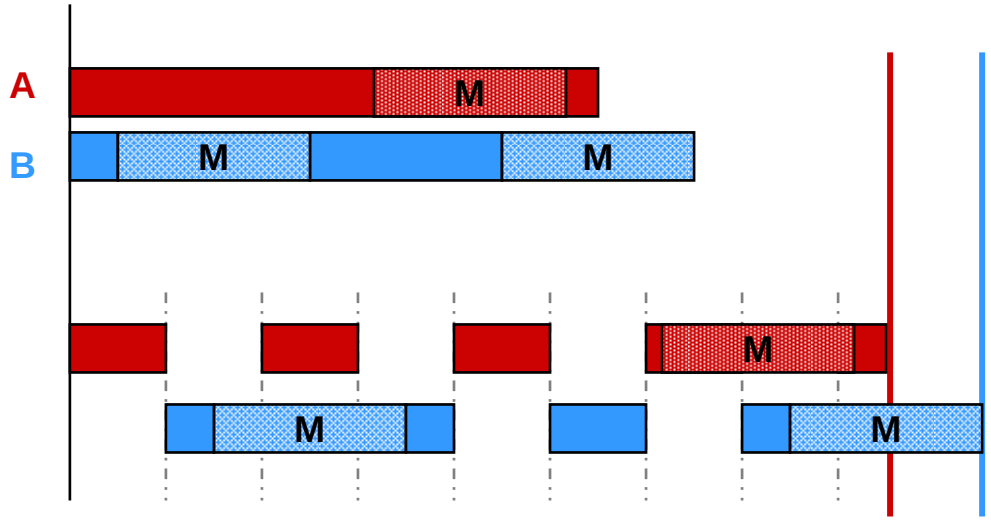

- Memory overlap: Yes (fair and bounded)
- Closed-form test: Yes

Deterministic Switching

High Performance and Tractable

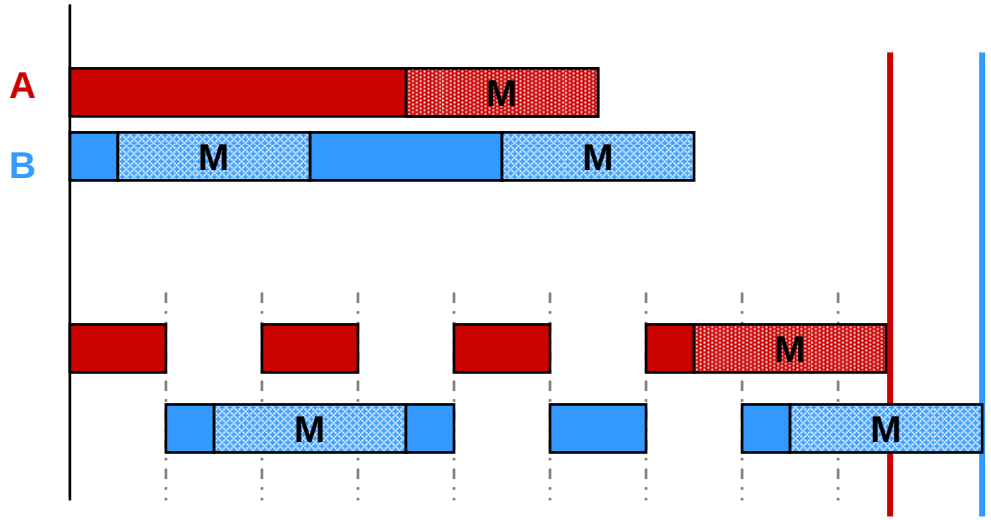

- Memory overlap: Yes (fair and bounded)
- Closed-form test: Yes

Tractability through Deterministic Switching

- Fully decouple independent tasks by forcing periodic switches
  - Every task gets a chance to initiate/overlap memory accesses
  - No scheduling dependences
  - No specificity regarding memory positioning

Weighted-Round-Robin (WRR)

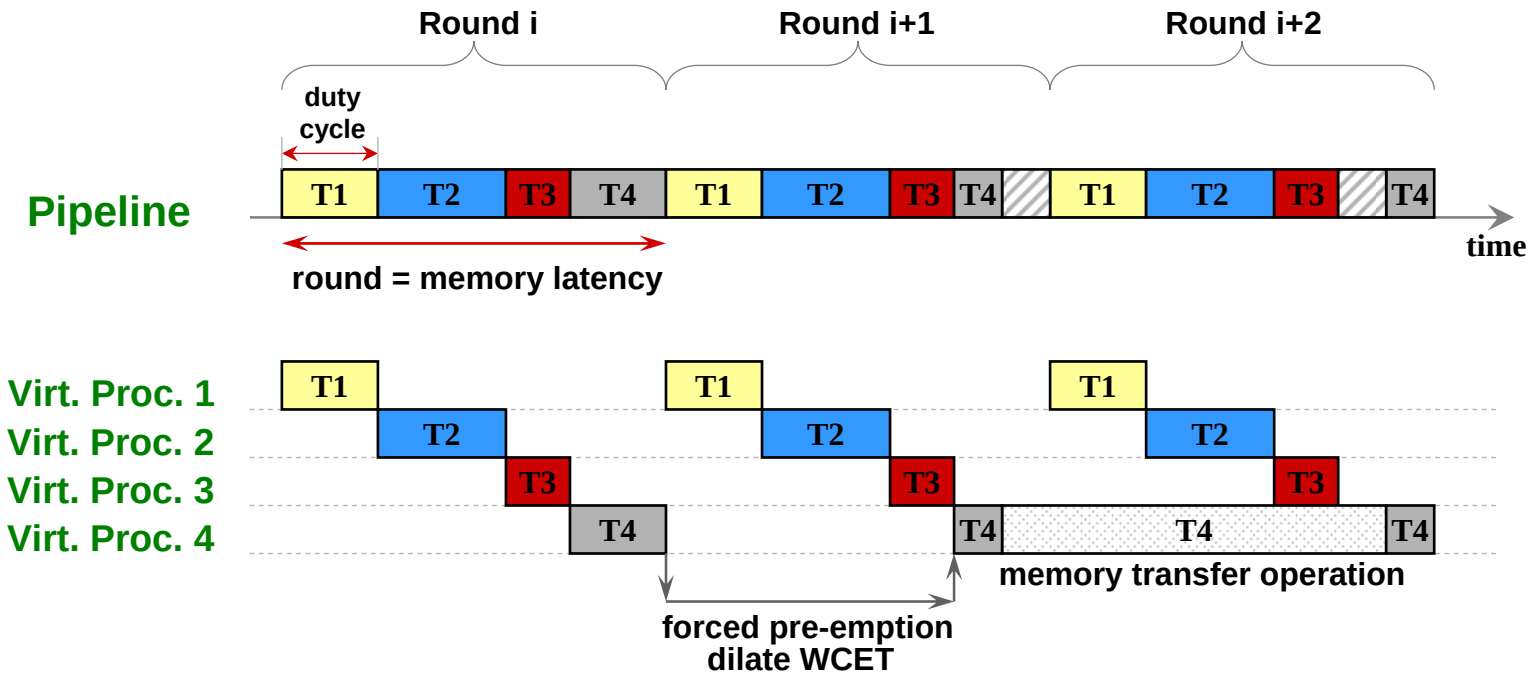

Analytical Framework for WRR

d: duty cycle $0 < d \leq 1$

P: period = deadline

WCET: Worst-Case Execution Time

WCET = C + M

C: aggregate computation time

M: aggregate memory time

WCET': dilated Worst-Case Execution Time

WCET' = (C/d) + M

Schedulability Test

1. Dilated task meets deadline on its virtual processor

Schedulability Test

1. Dilated task meets deadline on its virtual processor

$$WCET' \leq P \longrightarrow \frac{C}{d} + M \leq P$$

Schedulability Test

1. Dilated task meets deadline on its virtual processor

$$WCET' \leq P \longrightarrow \frac{C}{d} + M \leq P$$

$$d \geq \frac{C}{P - M}$$

Schedulability Test

1. Dilated task meets deadline on its virtual processor

$$WCET' \leq P \longrightarrow \frac{C}{d} + M \leq P$$

$$d \geq \frac{C}{P - M}$$

$$d = \frac{C}{P - M} = \frac{(C/P)}{1 - (M/P)}$$

Schedulability Test

1. Dilated task meets deadline on its virtual processor

$$WCET' \leq P \longrightarrow \frac{C}{d} + M \leq P$$

$$d \geq \frac{C}{P - M}$$

$$d = \frac{C}{P - M} = \frac{(C/P)}{1 - (M/P)}$$

2. Sum of all duty cycles less than or equal to 1

Schedulability Test

1. Dilated task meets deadline on its virtual processor

$$WCET' \leq P \longrightarrow \frac{C}{d} + M \leq P$$

$$d \geq \frac{C}{P - M}$$

$$d = \frac{C}{P - M} = \frac{(C/P)}{1 - (M/P)}$$

2. Sum of all duty cycles less than or equal to 1

$$\sum_{i=1}^{n} d_i \leq 1$$

Schedulability Test

1. Dilated task meets deadline on its virtual processor

$$WCET' \leq P \longrightarrow \frac{C}{d} + M \leq P$$

$$d \geq \frac{C}{P - M}$$

$$d = \frac{C}{P - M} = \frac{(C/P)}{1 - (M/P)}$$

2. Sum of all duty cycles less than or equal to 1

$$\sum_{i=1}^{n} d_i \leq 1 \longrightarrow \sum_{i=1}^{n} \frac{(C_i/P_i)}{1 - (M_i/P_i)} \leq 1$$

Generalized Analytical Framework

Multiple tasks per VP:

$$\frac{\left(\frac{C_1}{d_{vp}} + M_1\right)}{P_1} + \frac{\left(\frac{C_2}{d_{vp}} + M_2\right)}{P_2} + \cdots + \frac{\left(\frac{C_t}{d_{vp}} + M_t\right)}{P_t} \leq 1$$

$$d_{vp} = \frac{\sum_{j=1}^{t}\left(C_j / P_j\right)}{1 - \sum_{j=1}^{t}\left(M_j / P_j\right)}$$

Modeling Bus and Memory System

- Addressed in detail in paper
- Analysis accounts for:
  - Worst-case task serialization on memory bus
  - DRAM bank conflicts
  - Multiple VPs sharing single DRAM bank

Results

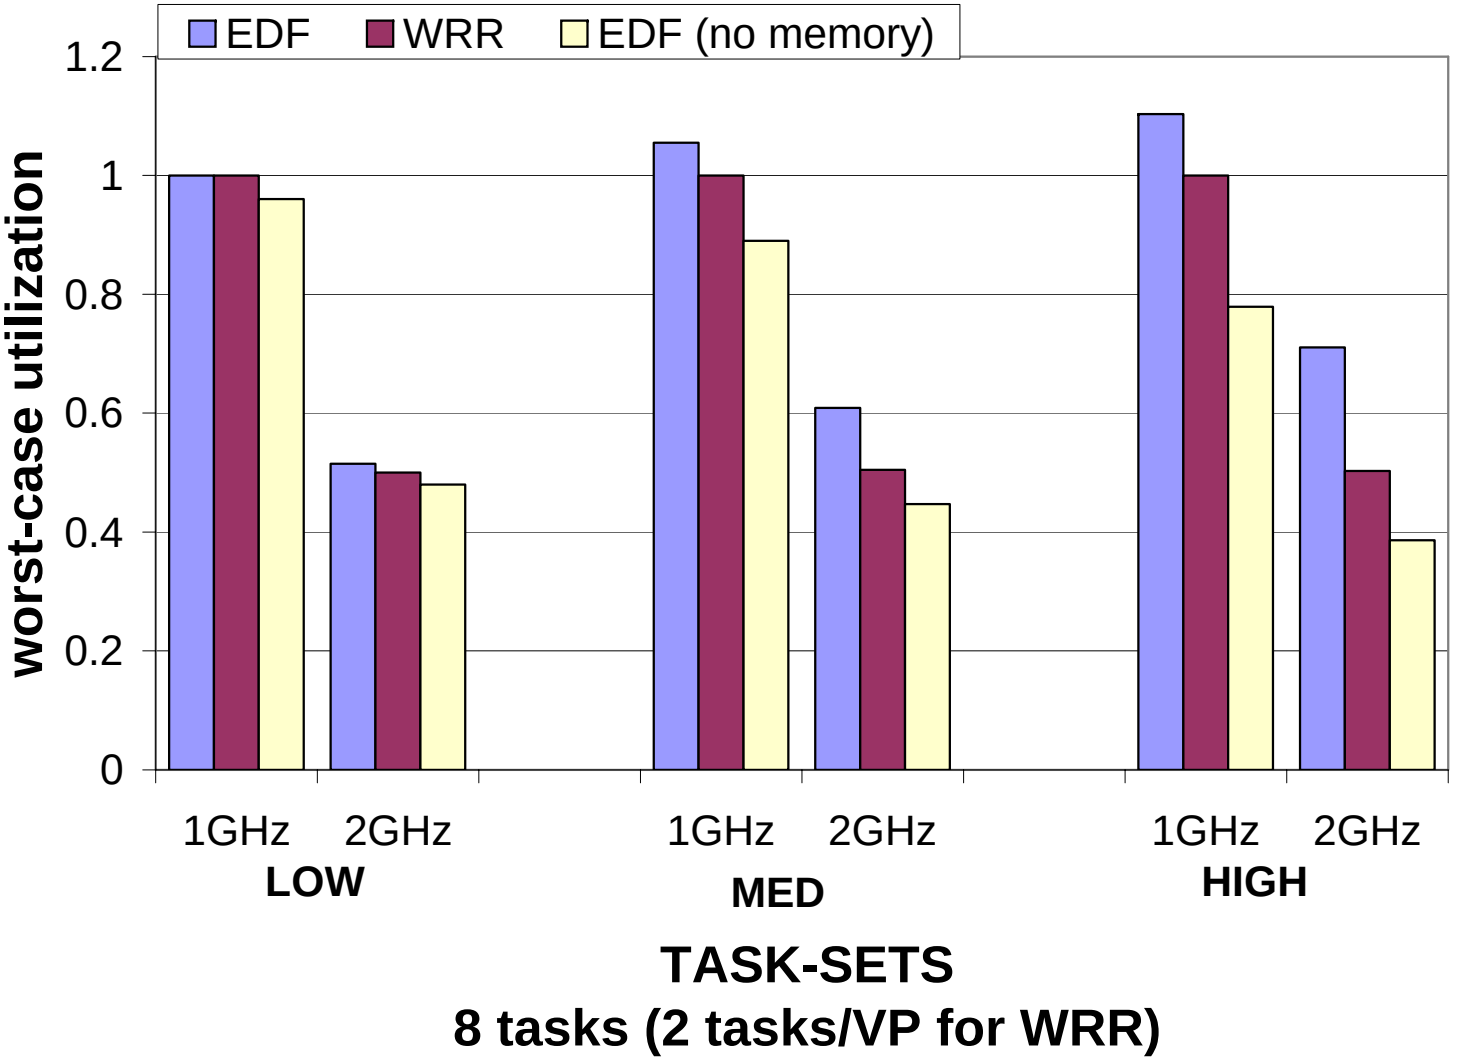
EDF
WRR
EDF (no memory)
worst-case utilization
1.2
1
0.8
0.6
0.4
0.2
0
1GHz
2GHz
LOW
1GHz
2GHz
MED
1GHz
2GHz
HIGH
TASK-SETS
8 tasks (2 tasks/VP for WRR)

Results

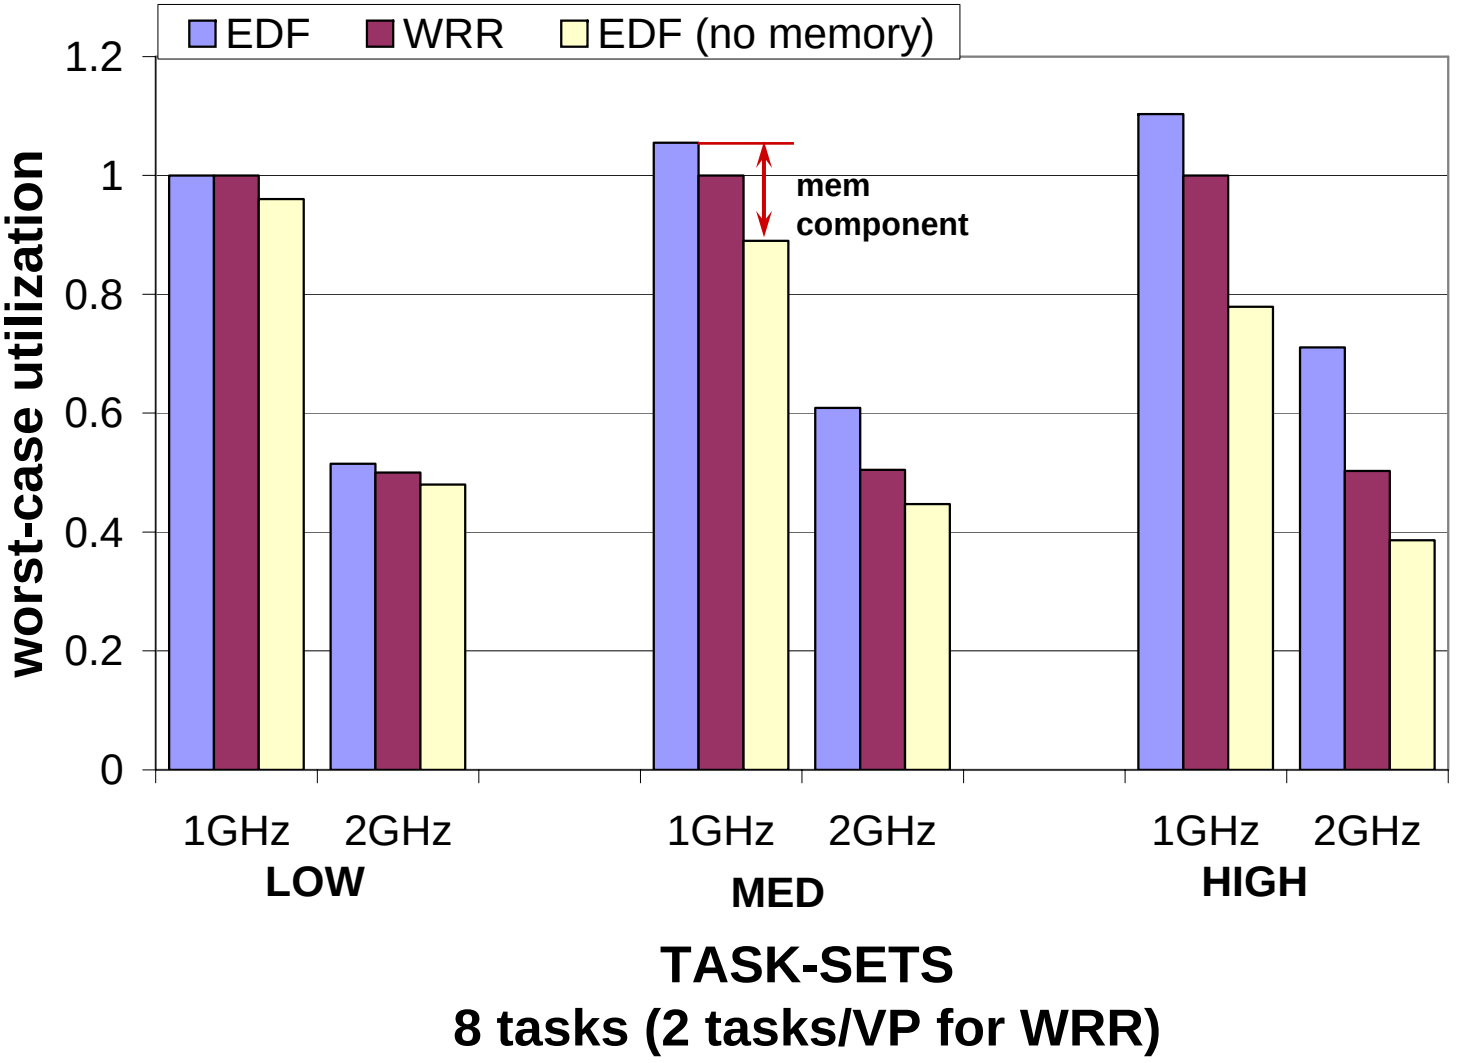
EDF
WRR
EDF (no memory)
1.2
1
0.8
0.6
0.4
0.2
0
worst-case utilization
mem component
1GHz
2GHz
LOW
1GHz
2GHz
MED
1GHz
2GHz
HIGH
TASK-SETS
8 tasks (2 tasks/VP for WRR)

Results

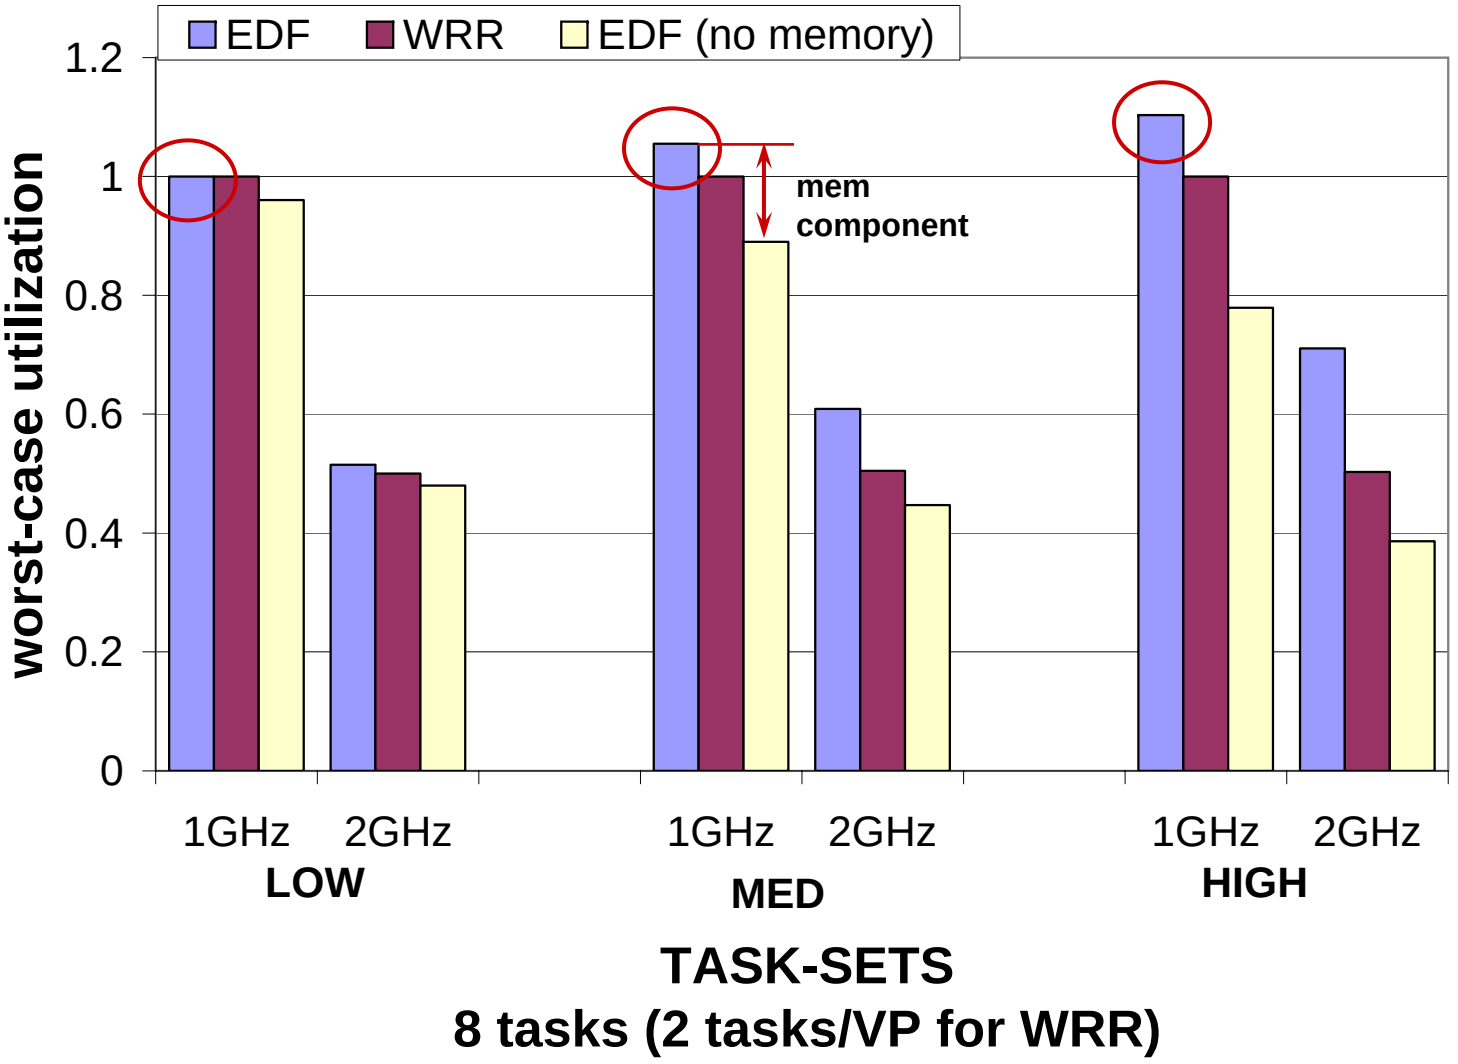
EDF
WRR
EDF (no memory)
1.2
1
0.8
0.6
0.4
0.2
0
worst-case utilization
mem component
1GHz
2GHz
LOW
1GHz
2GHz
MED
1GHz
2GHz
HIGH
TASK-SETS
8 tasks (2 tasks/VP for WRR)

Results

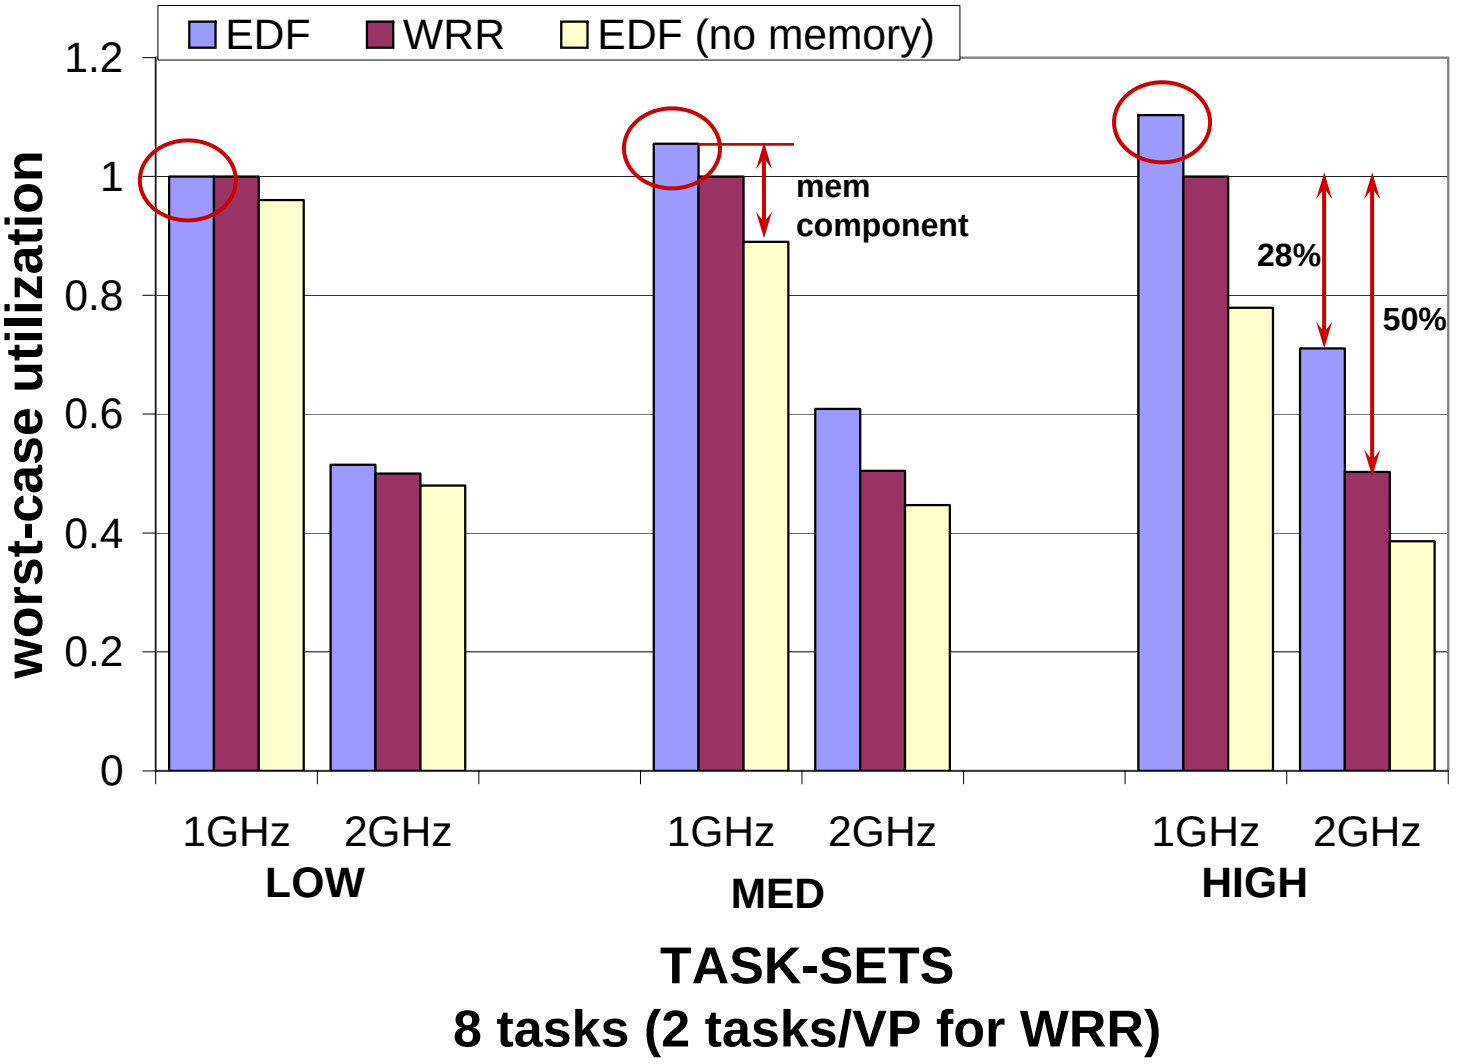
EDF
WRR
EDF (no memory)
worst-case utilization
1.2
1
0.8
0.6
0.4
0.2
0
mem component
28%
50%
1GHz
2GHz
LOW
MED
HIGH
TASK-SETS
8 tasks (2 tasks/VP for WRR)

Novel

	Memory overlap	Formalism (safety & tractability)
Classic real-time	No	Yes
Classic multithreading	Yes	No
Our real-time multithreading framework	Yes	Yes

Useful

- Fully capitalize on high-frequency embedded microprocessors
- Exceed schedulability limit of conventional real-time theory for uniprocessors by analytically bounding WCET overlap

Deployable

- Software-only solution
  - Use Ubicom IP3023 8-thread embedded microprocessor
  - Analytical framework + scheduling policy

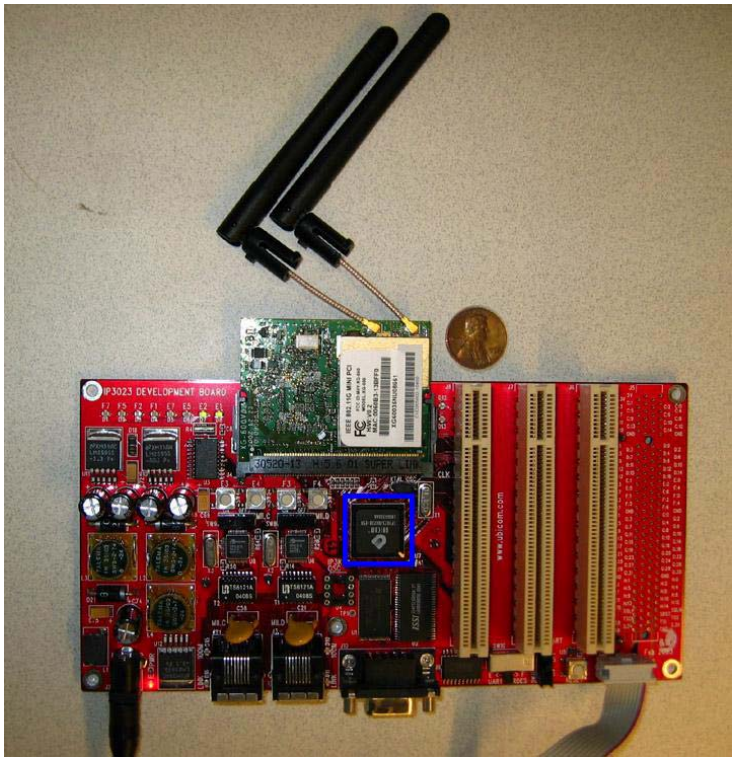

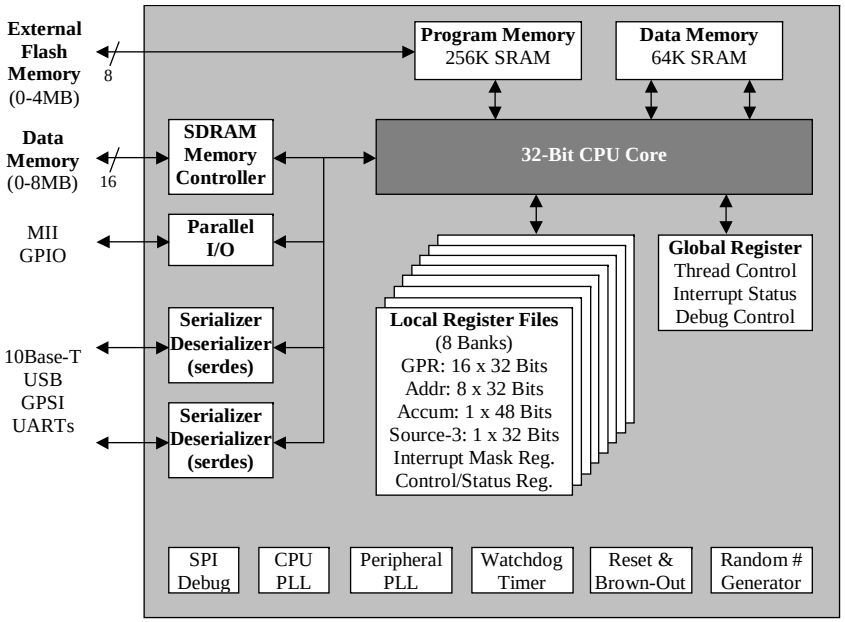

Summary

- Safely expose multithreading to hard-real-time schedulability analysis
- Bound computation / memory overlap
  - Offline closed-form schedulability test
  - Safe
  - Tractable
- Scale “Memory Wall” in embedded systems
  - Expose full benefits of high-frequency embedded processors

Questions?