Probabilistic Real-Time Scheduling and its Possible Link to Mixed-Criticality Systems

Georg von der Brüggen[1], Sergey Bozhko[2], Mario Günzel[1],
Kuan-Hsun Chen[3], Jian-Jia Chen[1], and Björn B. Brandenburg[2]

[1]TU Dortmund, Germany
[2]Max Planck Institute for Software Systems, Germany
[3]University of Twente, The Netherlands

05 December 2022

Uncertain Execution Behaviour in Mixed-Criticality

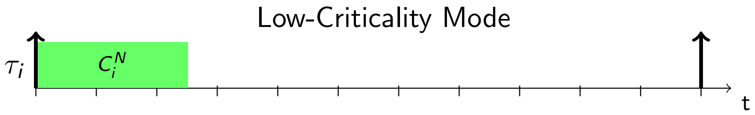

- $D_i = T_i$
- $\tau_i = (C_i^N,\ T_i)$

Uncertain Execution Behaviour in Mixed-Criticality

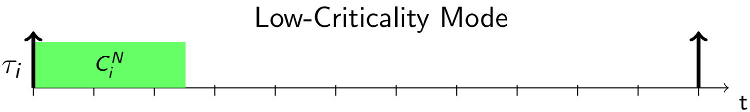

- System-wide switch to high-criticality mode

- $D_i = T_i$
- $\tau_i = (C_i^N,\ T_i)$

Uncertain Execution Behaviour in Mixed-Criticality

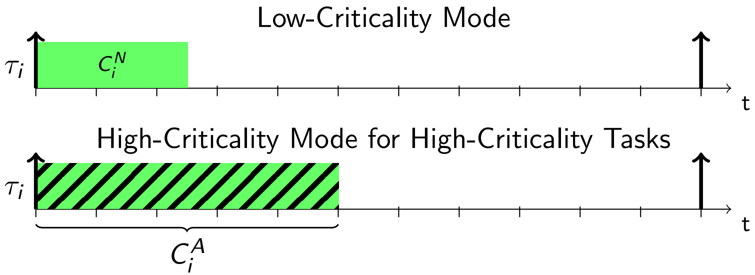

- System-wide switch to high-criticality mode

- $D_i = T_i$
- $\tau_i = (C_i^N,\ T_i)$, $C_i^A = 2 \cdot C_i^N$

Uncertain Execution Behaviour in Mixed-Criticality

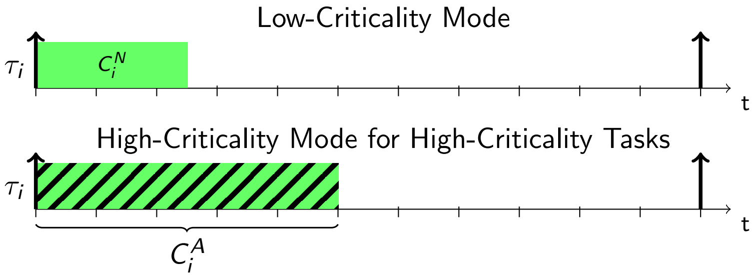

- System-wide switch to high-criticality mode
- Low-criticality tasks: No guarantees, best effort

- $D_i = T_i$
- $\tau_i = (C_i^N,\ T_i)$, $C_i^A = 2 \cdot C_i^N$

Worst-Case Response-Time Analysis

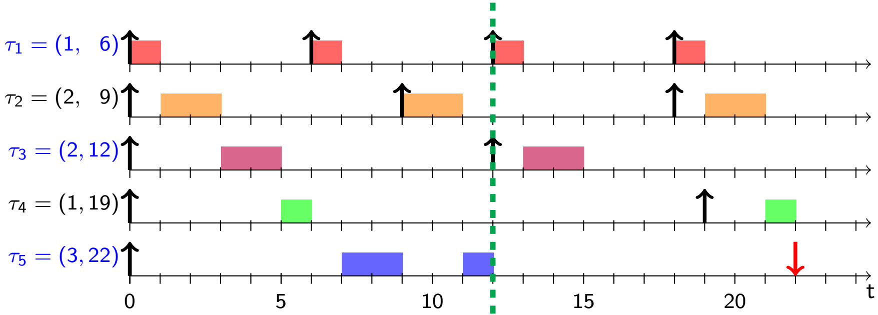

1. Low-criticality mode

Worst-Case Response-Time Analysis

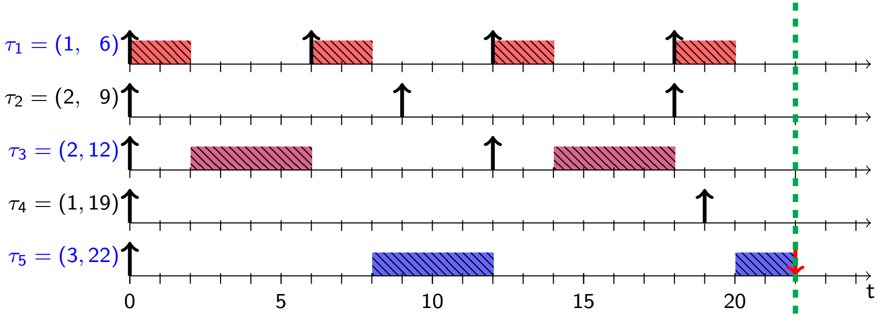

1. Low-criticality mode
2. High-criticality mode

Worst-Case Response-Time Analysis

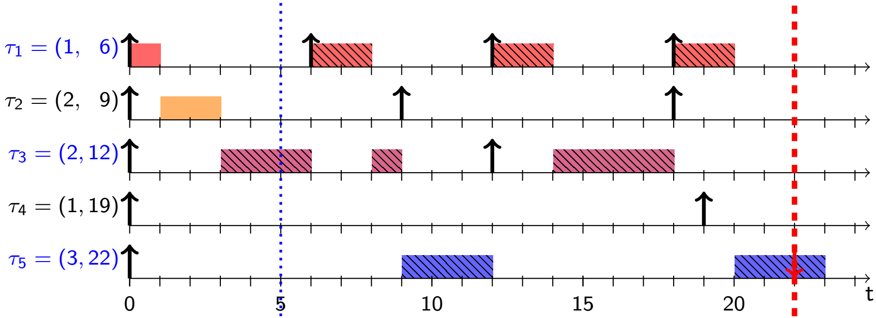

1. Low-criticality mode
2. High-criticality mode
3. Mode switch (usually the bottleneck)

Uncertain Execution Behaviour in Probabilistic View

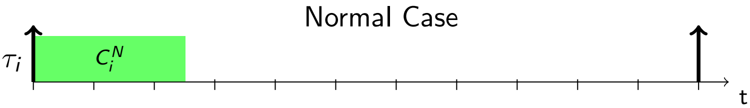

- $D_i = T_i$
- $\tau_i = (C_i^N,\ T_i)$

Uncertain Execution Behaviour in Probabilistic View

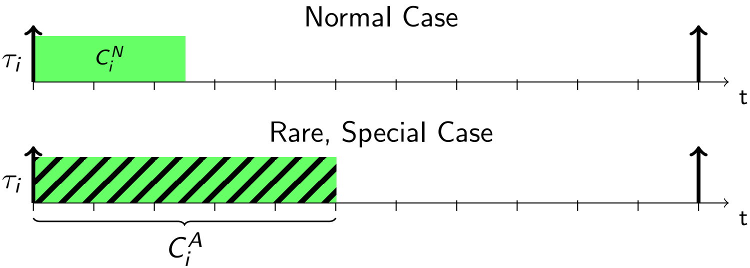

- $D_i = T_i$
- $\tau_i = (C_i^N,\ T_i)$, $C_i^A = 2 \cdot C_i^N$

Uncertain Execution Behaviour in Probabilistic View

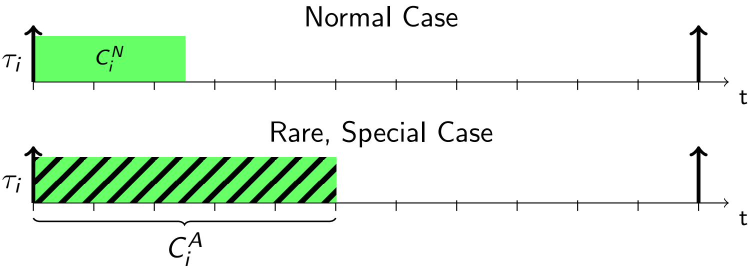

- $D_i = T_i$
- $\tau_i = (C_i^N,\ T_i)$, $C_i^A = 2 \cdot C_i^N$
- $\mathbb{P}(C_i^A) + \mathbb{P}(C_i^N) = 1$
- $\mathbb{P}(C_i^A) \ll \mathbb{P}(C_i^N)$

Probabilistic Worst-Case Response-Time Analysis

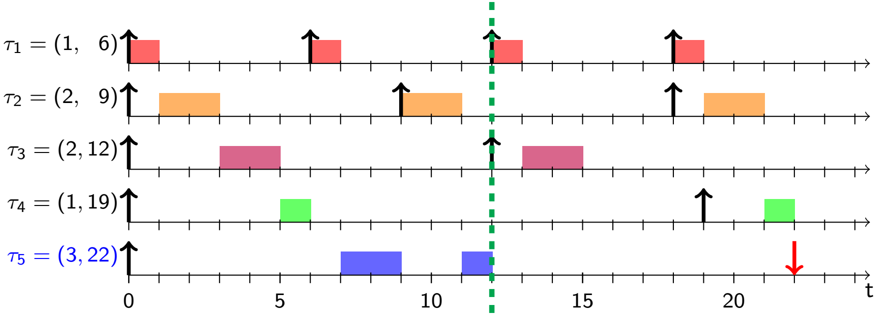

- Do all tasks meet their deadline under all circumstances?

Probabilistic Worst-Case Response-Time Analysis

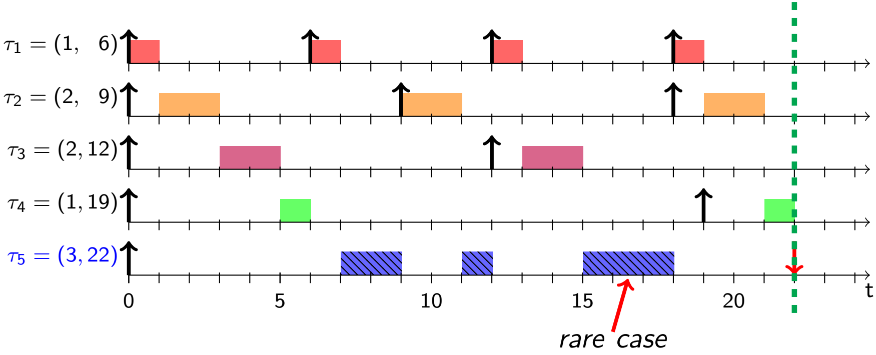

- Do all tasks meet their deadline under all circumstances?

Probabilistic Worst-Case Response-Time Analysis

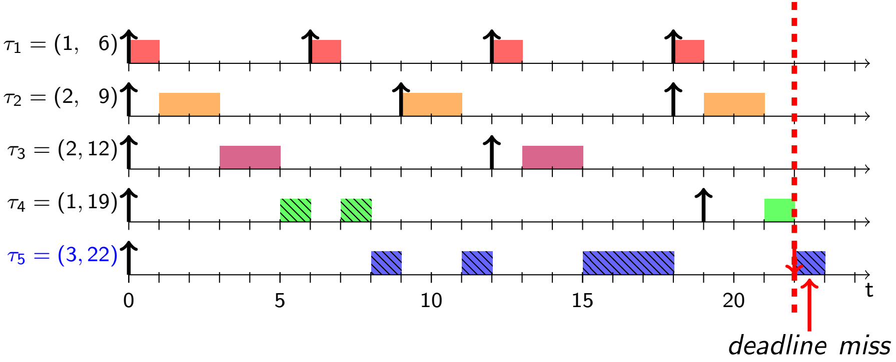

- Do all tasks meet their deadline under all circumstances?

Probabilistic Worst-Case Response-Time Analysis

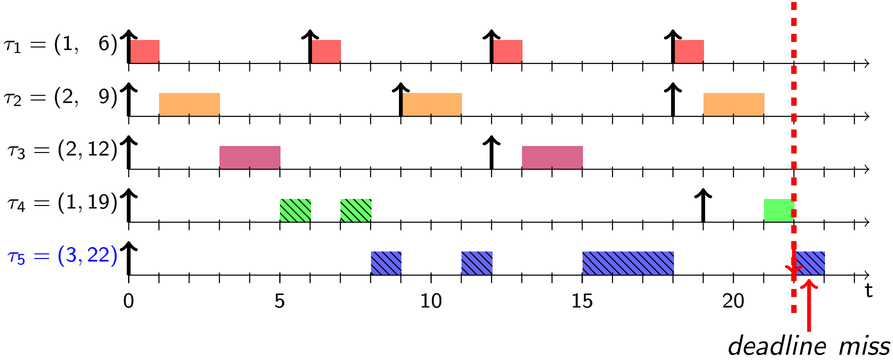

- ~~Do all tasks meet their deadline under all circumstances?~~
- What is the probability that a job misses its deadline?

Links Between Mixed-Criticality and Probabilistic Scheduling

- Recent Mixed-Criticality research often considers return to low-criticality mode

Links Between Mixed-Criticality and Probabilistic Scheduling

- Recent Mixed-Criticality research often considers return to low-criticality mode
- That only a subset of tasks exhibits abnormal behaviour seems natural

Links Between Mixed-Criticality and Probabilistic Scheduling

- Recent Mixed-Criticality research often considers return to low-criticality mode
- That only a subset of tasks exhibits abnormal behaviour seems natural

$\Rightarrow$ Frequent system-wide mode switches which may be both costly and unnecessary

Links Between Mixed-Criticality and Probabilistic Scheduling

- Recent Mixed-Criticality research often considers return to low-criticality mode
- That only a subset of tasks exhibits abnormal behaviour seems natural

$\Rightarrow$ Frequent system-wide mode switches which may be both costly and unnecessary

Idea: Mode switch based on deadline miss probability

Soft and Firm Real-Time Systems

Important problem in industry (Akesson et al., RTSS 2020)

- 62%: system includes soft or firm real-time components
- 45%: the most critical function can miss some deadlines
- Only 15%: deadlines can never be missed

B. Akesson, M. Nasri, G. Nelissen, S. Altmeyer, and R. I. Davis
"A comprehensive survey of industry practice in real-time systems", RTSS 2020

Soft and Firm Real-Time Systems

Important problem in industry (Akesson et al., RTSS 2020)

- 62%: system includes soft or firm real-time components
- 45%: the most critical function can miss some deadlines
- Only 15%: deadlines can never be missed

The probability that deadlines are missed must be quantified

B. Akesson, M. Nasri, G. Nelissen, S. Altmeyer, and R. I. Davis
“A comprehensive survey of industry practice in real-time systems”, RTSS 2020

Quantifying Deadline Misses

Deadline Miss Rate: Percentage of deadline misses (in the long run)

Quantifying Deadline Misses

Deadline Miss Rate: Percentage of deadline misses (in the long run)

Worst-Case Deadline Failure Probability:
Maximum probability of any job to be the first to miss its deadline

Quantifying Deadline Misses

Deadline Miss Rate: Percentage of deadline misses (in the long run)

Worst-Case Deadline Failure Probability:
Maximum probability of any job to be the first to miss its deadline

- Mean time to failure due to a deadline miss
- Bounds the deadline-miss rate (no backlog)

Complexity Issues of the Naive Solution

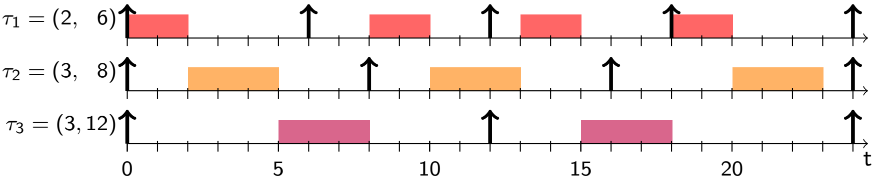

- Consider all jobs in the hyperperiod individually

Complexity Issues of the Naive Solution

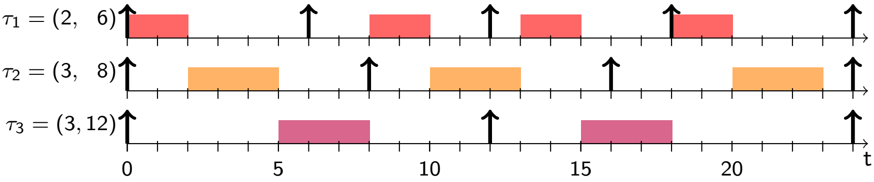

- Consider all jobs in the hyperperiod individually
- For each job, calculate the probability to miss its deadline

Complexity Issues of the Naive Solution

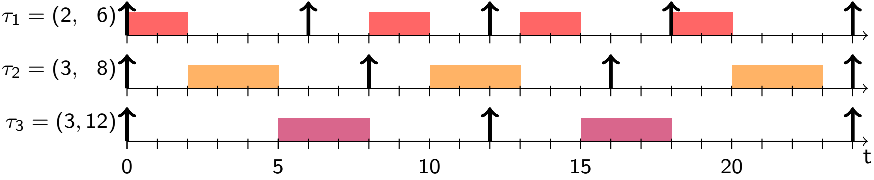

- Consider all jobs in the hyperperiod individually
- For each job, calculate the probability to miss its deadline

$\Rightarrow$ Computational complexity is too high to be feasible in practice

How much can we gain?

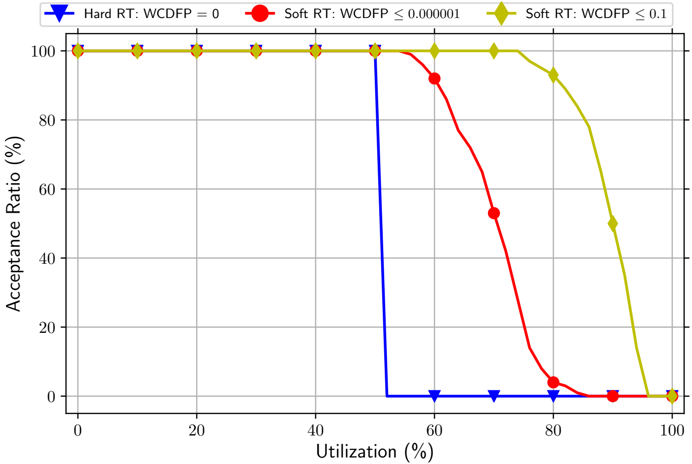

Setting for rare mode:
$P_i(A) = 0.01$
$C_i^A = 2 \cdot C_i^N$

G. v. d. Brüggen, N. Piatkowski, K.-H. Chen, J.-J. Chen, K. Morik, and B. Brandenburg
“Efficiently Approximating the Worst-Case Deadline Failure Probability under EDF”, RTSS 2021

Major Research Questions

1. Which scenarios must be considered?

Major Research Questions

1. Which scenarios must be considered?
   - depends on inter-arrival constraints
   - depends on scheduling policy

Major Research Questions

1. Which scenarios must be considered?
   - depends on inter-arrival constraints
   - depends on scheduling policy
   - depends on considered metric

Major Research Questions

1. Which scenarios must be considered?
   - depends on inter-arrival constraints
   - depends on scheduling policy
   - depends on considered metric
   - reduce number of scenarios without under-approximation

Major Research Questions

1. Which scenarios must be considered?
   - depends on inter-arrival constraints
   - depends on scheduling policy
   - depends on considered metric
   - reduce number of scenarios without under-approximation
   - reduce number of scenarios without being too pessimistic

Major Research Questions

1. Which scenarios must be considered?
   - depends on inter-arrival constraints
   - depends on scheduling policy
   - depends on considered metric
   - reduce number of scenarios without under-approximation
   - reduce number of scenarios without being too pessimistic
2. How to calculate the miss probability for a job efficiently?

Major Research Questions

1. Which scenarios must be considered?
   - depends on inter-arrival constraints
   - depends on scheduling policy
   - depends on considered metric
   - reduce number of scenarios without under-approximation
   - reduce number of scenarios without being too pessimistic
2. How to calculate the miss probability for a job efficiently?
   - determine probability for too much workload in an interval

Major Research Questions

1. Which scenarios must be considered?
   - depends on inter-arrival constraints
   - depends on scheduling policy
   - depends on considered metric
   - reduce number of scenarios without under-approximation
   - reduce number of scenarios without being too pessimistic
2. How to calculate the miss probability for a job efficiently?
   - determine probability for too much workload in an interval
3. How can dependencies be handled?

State of the Art - Calculation

- Job-based approaches

Job-Based Approaches

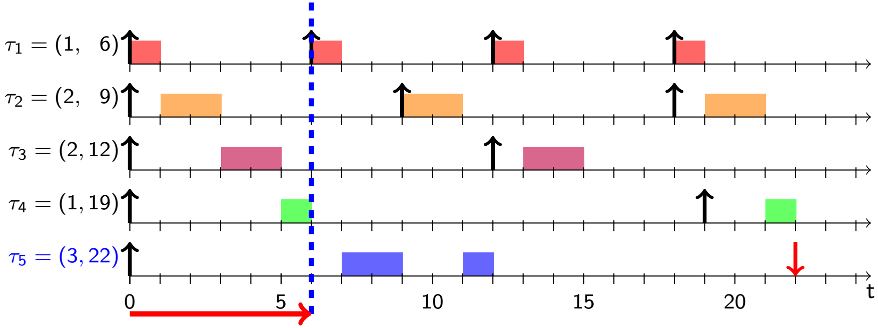

Job-Based Approaches

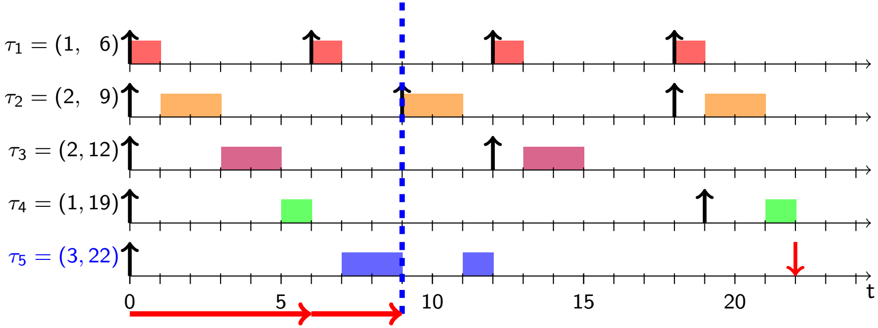

Job-Based Approaches

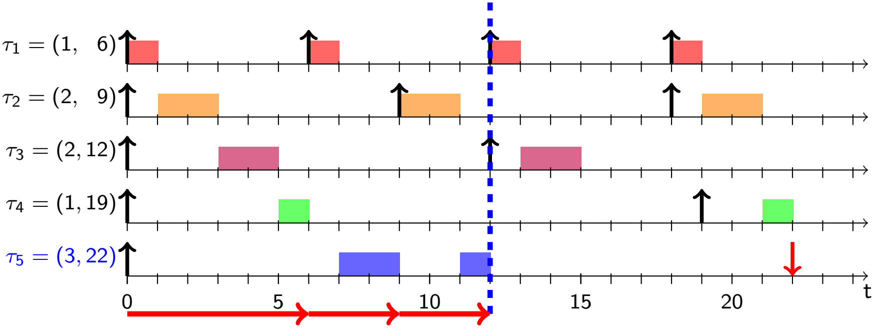

Job-Based Approaches

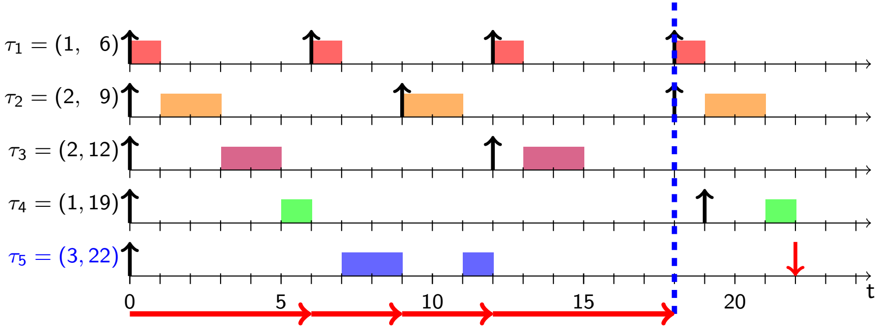

Job-Based Approaches

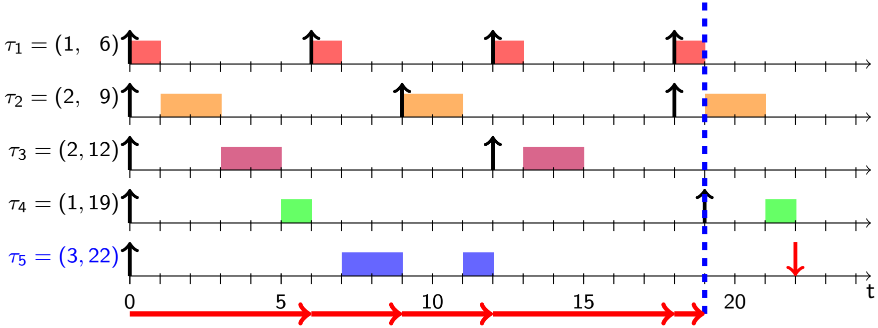

Job-Based Approaches

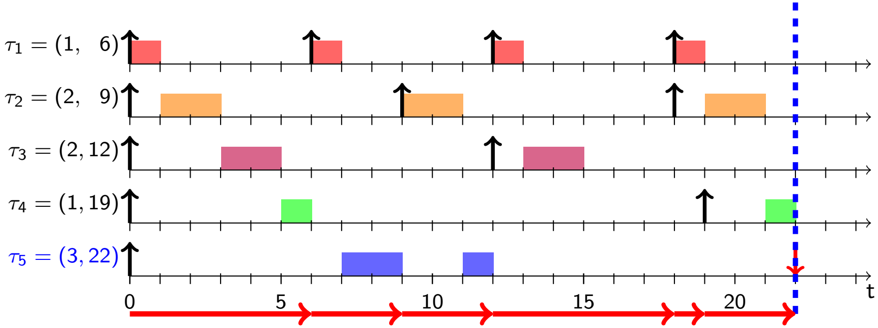

Direct Convolution

$$\frac{C_1}{\mathbb{P}_1} = \begin{pmatrix} 3 & 5 \\ 0.9 & 0.1 \end{pmatrix} \otimes \begin{pmatrix} 5 & 6 \\ 0.8 & 0.2 \end{pmatrix} = \frac{C_2}{\mathbb{P}_2}$$

Direct Convolution

$$\begin{matrix} C_1 \\ \mathbb{P}_1 \end{matrix} = \begin{pmatrix} 3 & 5 \\ 0.9 & 0.1 \end{pmatrix} \otimes \begin{pmatrix} 5 & 6 \\ 0.8 & 0.2 \end{pmatrix} = \begin{matrix} C_2 \\ \mathbb{P}_2 \end{matrix}$$

$$\Big\downarrow$$

$$\begin{pmatrix} \quad \end{pmatrix}$$

Direct Convolution

$$\begin{matrix} C_1 \\ \mathbb{P}_1 \end{matrix} = \begin{pmatrix} \color{red}{3} & 5 \\ \color{red}{0.9} & 0.1 \end{pmatrix} \otimes \begin{pmatrix} \color{red}{5} & 6 \\ \color{red}{0.8} & 0.2 \end{pmatrix} = \begin{matrix} C_2 \\ \mathbb{P}_2 \end{matrix}$$

$$\Big\downarrow$$

$$\begin{pmatrix} \color{red}{8} & & & \\ \color{red}{0.72} & & & \end{pmatrix}$$

Direct Convolution

$$\begin{matrix} C_1 \\ \mathbb{P}_1 \end{matrix} = \begin{pmatrix} \color{red}{3} & 5 \\ \color{red}{0.9} & 0.1 \end{pmatrix} \otimes \begin{pmatrix} 5 & \color{red}{6} \\ 0.8 & \color{red}{0.2} \end{pmatrix} = \begin{matrix} C_2 \\ \mathbb{P}_2 \end{matrix}$$

$$\Downarrow$$

$$\begin{pmatrix} 8 & \color{red}{9} & & \\ 0.72 & \color{red}{0.18} & & \end{pmatrix}$$

Direct Convolution

$$\begin{matrix} C_1 \\ \mathbb{P}_1 \end{matrix} = \begin{pmatrix} 3 & \color{red}{5} \\ 0.9 & \color{red}{0.1} \end{pmatrix} \otimes \begin{pmatrix} \color{red}{5} & 6 \\ \color{red}{0.8} & 0.2 \end{pmatrix} = \begin{matrix} C_2 \\ \mathbb{P}_2 \end{matrix}$$

$$\Big\downarrow$$

$$\begin{pmatrix} 8 & 9 & \color{red}{10} & \quad \\ 0.72 & 0.18 & \color{red}{0.08} & \quad \end{pmatrix}$$

Direct Convolution

$$\begin{matrix} C_1 \\ \mathbb{P}_1 \end{matrix} = \begin{pmatrix} 3 & \color{red}{5} \\ 0.9 & \color{red}{0.1} \end{pmatrix} \otimes \begin{pmatrix} 5 & \color{red}{6} \\ 0.8 & \color{red}{0.2} \end{pmatrix} = \begin{matrix} C_2 \\ \mathbb{P}_2 \end{matrix}$$

$$\Big\downarrow$$

$$\begin{pmatrix} 8 & 9 & 10 & \color{red}{11} \\ 0.72 & 0.18 & 0.08 & \color{red}{0.02} \end{pmatrix}$$

Job-Level Convolution

τ_1

$\begin{matrix} C_1 \\ \mathbb{P}_1 \end{matrix} = \begin{pmatrix} 3 & 5 \\ 0.9 & 0.1 \end{pmatrix}$

$D_1 = T_1 = 8$

τ_2

$\begin{matrix} C_2 \\ \mathbb{P}_2 \end{matrix} = \begin{pmatrix} 5 & 6 \\ 0.8 & 0.2 \end{pmatrix}$

$D_2 = T_2 = 14$

Job-Level Convolution

τ_1

$$\begin{matrix} C_1 \\ \mathbb{P}_1 \end{matrix} = \begin{pmatrix} 3 & 5 \\ 0.9 & 0.1 \end{pmatrix}$$

$D_1 = T_1 = 8$

τ_2

$$\begin{matrix} C_2 \\ \mathbb{P}_2 \end{matrix} = \begin{pmatrix} 5 & 6 \\ 0.8 & 0.2 \end{pmatrix}$$

$D_2 = T_2 = 14$

$t = 0$ $t = 8$ $t = 14$

Job-Level Convolution

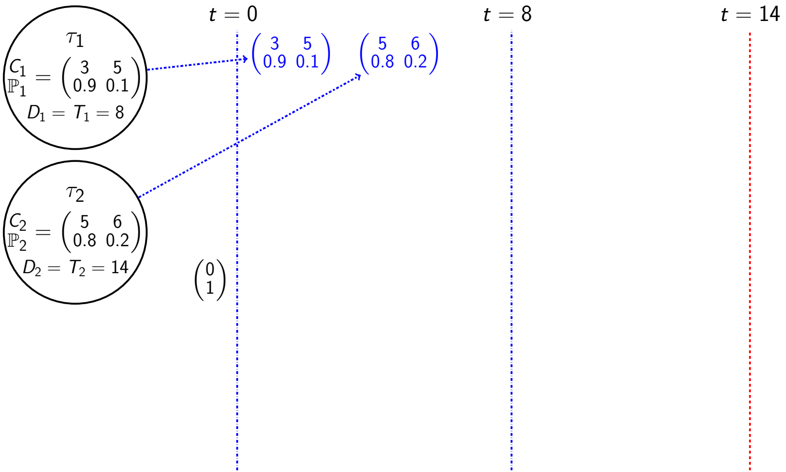

Job-Level Convolution

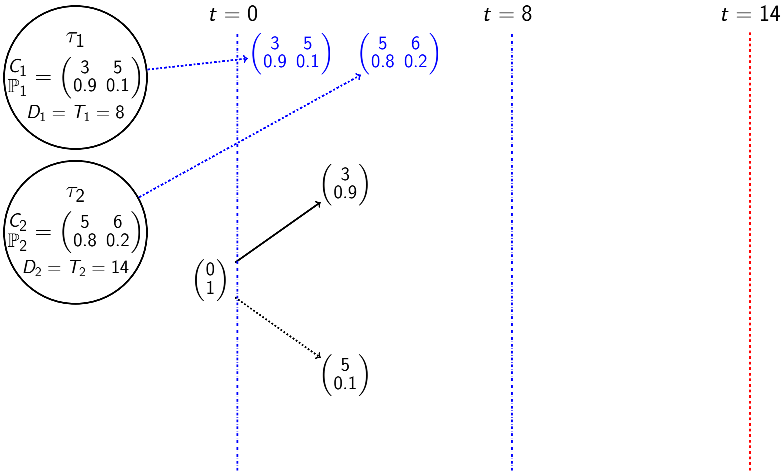

Job-Level Convolution

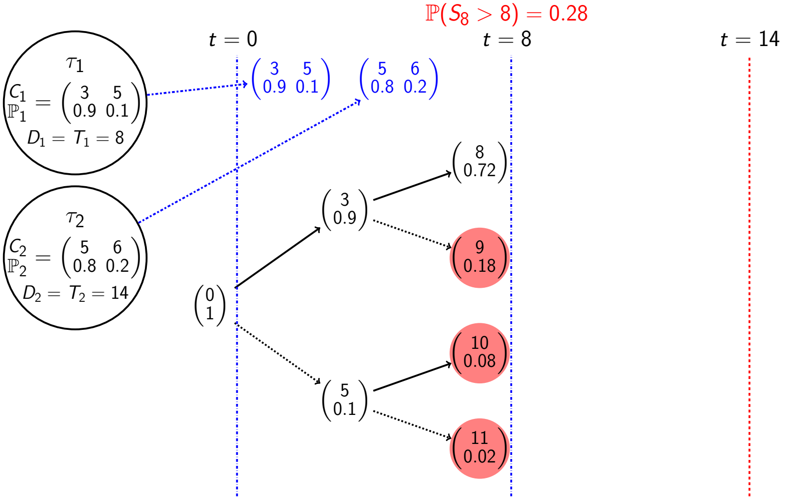

Job-Level Convolution

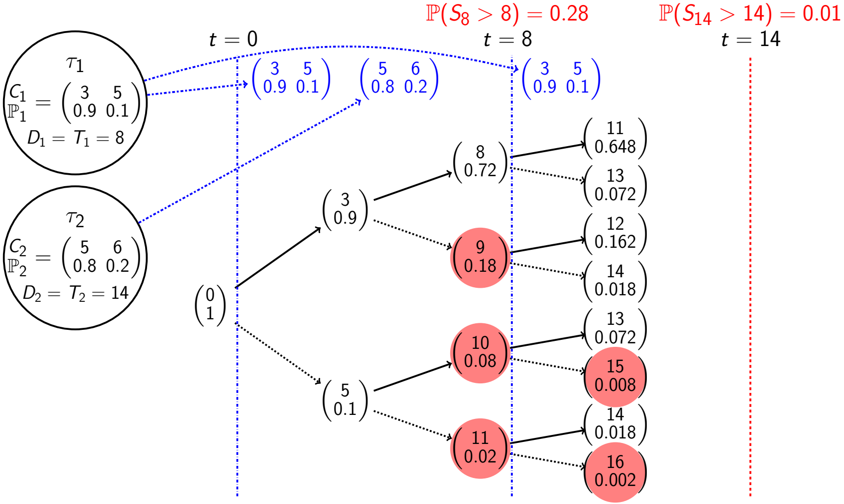

Job-Level Convolution

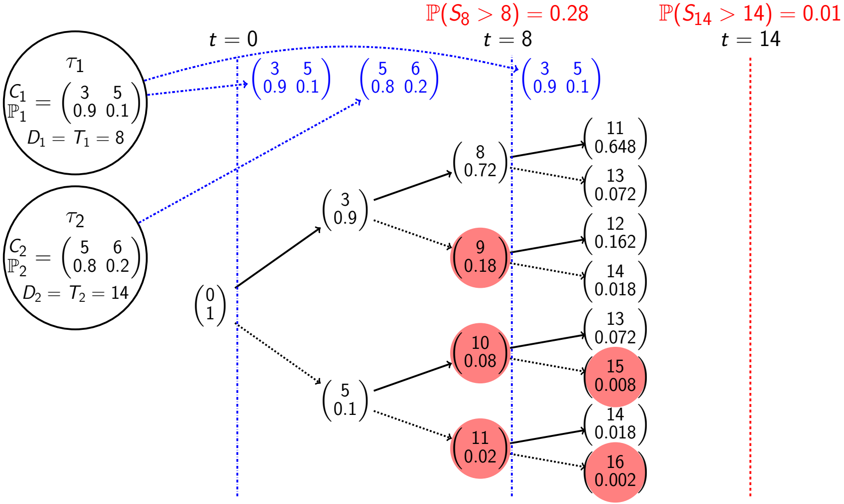

Exponential in the number of jobs

Job-Level Convolution

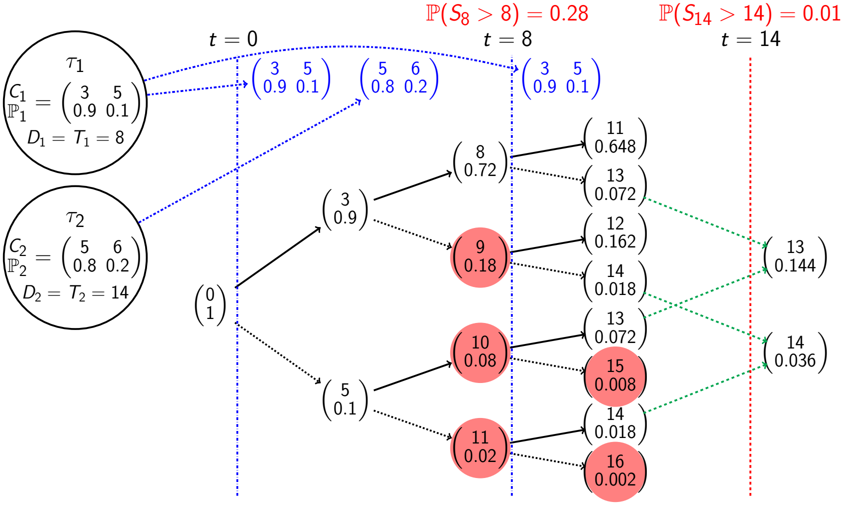

Exponential in the number of jobs

State of the Art - Calculation

- Job-based + theoretically exact
  - Direct job-level convolution - exponential # jobs
  - Re-sampling (Maxim and Cucu-Grosjean, RTSS 13)
    Improved re-sampling (Markovic et al., ECRTS 21)

Re-Sampling

$$\binom{7}{0.84} \quad \binom{8}{0.15} \quad \binom{9}{0.012} \quad \binom{10}{4.9{\cdot}10^{-04}} \quad \binom{11}{1.3{\cdot}10^{-05}} \quad \binom{12}{1.9{\cdot}10^{-07}} \quad \binom{13}{1.6{\cdot}10^{-09}} \quad \binom{14}{6.1{\cdot}10^{-12}}$$

Total C_i	7	8	9	10	11	12	13	14
Probability	0.84	0.15	0.012	$4.9 \cdot 10^{-04}$	$1.3 \cdot 10^{-05}$	$1.9 \cdot 10^{-07}$	$1.6 \cdot 10^{-09}$	$6.1 \cdot 10^{-12}$

Total C_i	7	8	9	10	11	
Probability	0.84	0.15	0.012	$4.9 \cdot 10^{-04}$	$1.3 \cdot 10^{-05}$	

Re-Sampling

$\binom{7}{0.84}$ $\binom{8}{0.15}$ $\binom{9}{0.012}$ $\binom{10}{4.9 \cdot 10^{-04}}$ $\binom{11}{1.3 \cdot 10^{-05}}$ $\binom{12}{1.9 \cdot 10^{-07}}$ $\binom{13}{1.6 \cdot 10^{-09}}$ $\binom{14}{6.1 \cdot 10^{-12}}$

Total C_i	7	8	9	10	11	12	13	14
Probability	0.84	0.15	0.012	$4.9 \cdot 10^{-04}$	$1.3 \cdot 10^{-05}$	$1.9 \cdot 10^{-07}$	$1.6 \cdot 10^{-09}$	$6.1 \cdot 10^{-12}$

Total C_i	7	8	9	10	11	
Probability	0.84	0.15	0.012	$4.9 \cdot 10^{-04}$	$1.3 \cdot 10^{-05}$	

Re-Sampling

$\binom{7}{0.84}$ $\binom{8}{0.15}$ $\binom{9}{0.012}$ $\binom{10}{4.9 \cdot 10^{-04}}$ $\binom{11}{1.3 \cdot 10^{-05}}$

Total C_i	7	8	9	10	11	12	13	14
Probability	0.84	0.15	0.012	$4.9 \cdot 10^{-04}$	$1.3 \cdot 10^{-05}$	$1.9 \cdot 10^{-07}$	$1.6 \cdot 10^{-09}$	$6.1 \cdot 10^{-12}$
Total C_i	7	8	9	10	11			
Probability	0.84	0.15	0.012	$4.9 \cdot 10^{-04}$	$1.3 \cdot 10^{-05}$			

Re-Sampling

$\begin{pmatrix} 7 \\ 0.84 \end{pmatrix}$ $\begin{pmatrix} 8 \\ 0.15 \end{pmatrix}$ $\begin{pmatrix} 9 \\ 0.012 \end{pmatrix}$ $\begin{pmatrix} 10 \\ 4.9 \cdot 10^{-04} \end{pmatrix}$ $\begin{pmatrix} 11 \\ 1.3 \cdot 10^{-05} \end{pmatrix}$

Total C_i	7	8	9	10	11	12	13	14
Probability	0.84	0.15	0.012	$4.9 \cdot 10^{-04}$	$1.3 \cdot 10^{-05}$	$1.9 \cdot 10^{-07}$	$1.6 \cdot 10^{-09}$	$6.1 \cdot 10^{-12}$

Total C_i	7	8	9	10	11	
Probability	0.84	0.15	0.012	$4.9 \cdot 10^{-04}$	$1.3 \cdot 10^{-05}$	

Re-Sampling

$\binom{7}{0.84}$ $\binom{8}{0.15}$ $\binom{9}{0.012}$ $\binom{10}{4.9 \cdot 10^{-04}}$ $\binom{11}{1.3 \cdot 10^{-05}}$ $\binom{14}{\ }$

Total C_i	7	8	9	10	11	12	13	14
Probability	0.84	0.15	0.012	$4.9 \cdot 10^{-04}$	$1.3 \cdot 10^{-05}$	$1.9 \cdot 10^{-07}$	$1.6 \cdot 10^{-09}$	$6.1 \cdot 10^{-12}$

Total C_i	7	8	9	10	11	14
Probability	0.84	0.15	0.012	$4.9 \cdot 10^{-04}$	$1.3 \cdot 10^{-05}$	

Re-Sampling

$\binom{7}{0.84}$ $\binom{8}{0.15}$ $\binom{9}{0.012}$ $\binom{10}{4.9 \cdot 10^{-04}}$ $\binom{11}{1.3 \cdot 10^{-05}}$ $\binom{14}{1.916061 \cdot 10^{-07}}$

Total C_i	7	8	9	10	11	12	13	14
Probability	0.84	0.15	0.012	$4.9 \cdot 10^{-04}$	$1.3 \cdot 10^{-05}$	$1.9 \cdot 10^{-07}$	$1.6 \cdot 10^{-09}$	$6.1 \cdot 10^{-12}$

Total C_i	7	8	9	10	11	14
Probability	0.84	0.15	0.012	$4.9 \cdot 10^{-04}$	$1.3 \cdot 10^{-05}$	$1.916061 \cdot 10^{-07}$

State of the Art - Calculation

- Job-based + theoretically exact
  - Direct job-level convolution - exponential # jobs
  - Re-sampling (Maxim and Cucu-Grosjean, RTSS 13)
    Improved re-sampling (Markovic et al., ECRTS 21)
    - precision loss potentially large
  - Cyclic convolution (Markovic et al., ECRTS 21)
    + faster than direct convolution

State of the Art - Calculation

- Job-based + theoretically exact
  - Direct job-level convolution - exponential # jobs
  - Re-sampling (Maxim and Cucu-Grosjean, RTSS 13)
    Improved re-sampling (Markovic et al., ECRTS 21)
    - precision loss potentially large
  - Cyclic convolution (Markovic et al., ECRTS 21)
    + faster than direct convolution
  - Monte-Carlo RTA (Bozhko et al., RTSS 21)

Monte-Carlo Response Time Analysis

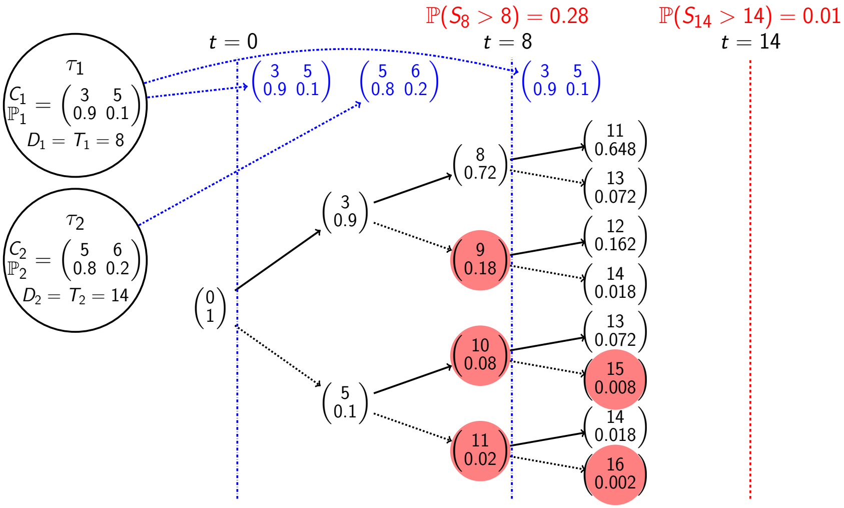

Monte-Carlo Response Time Analysis

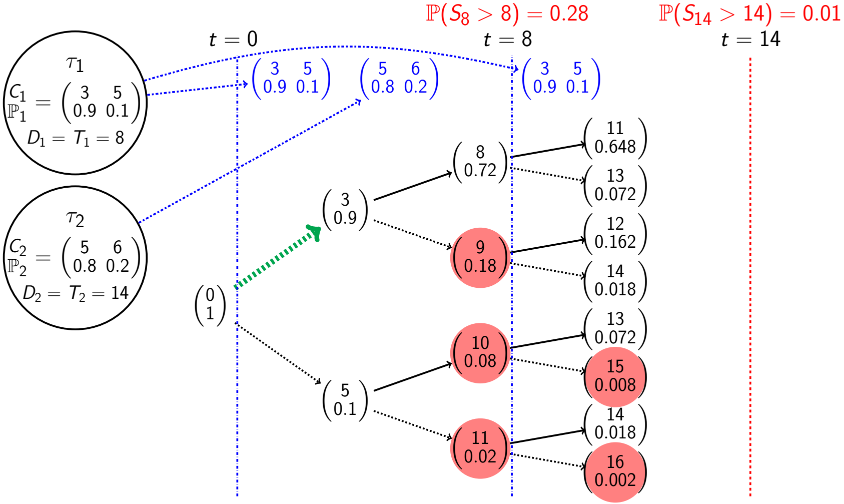

Monte-Carlo Response Time Analysis

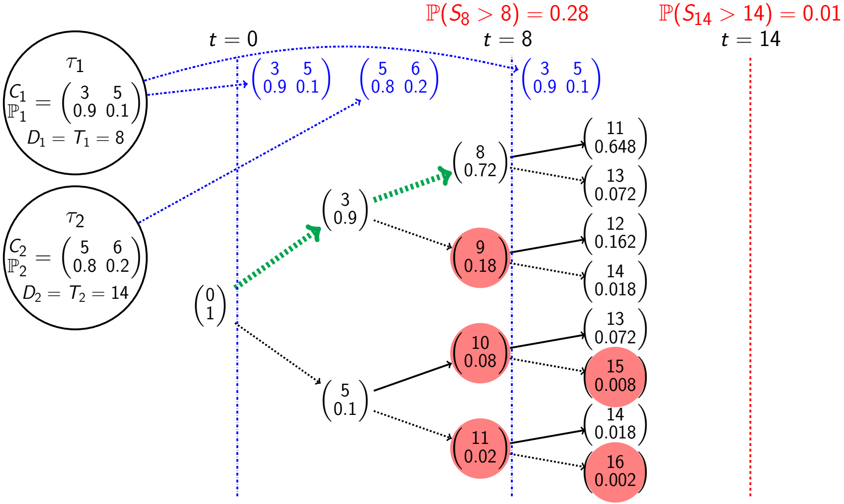

Monte-Carlo Response Time Analysis

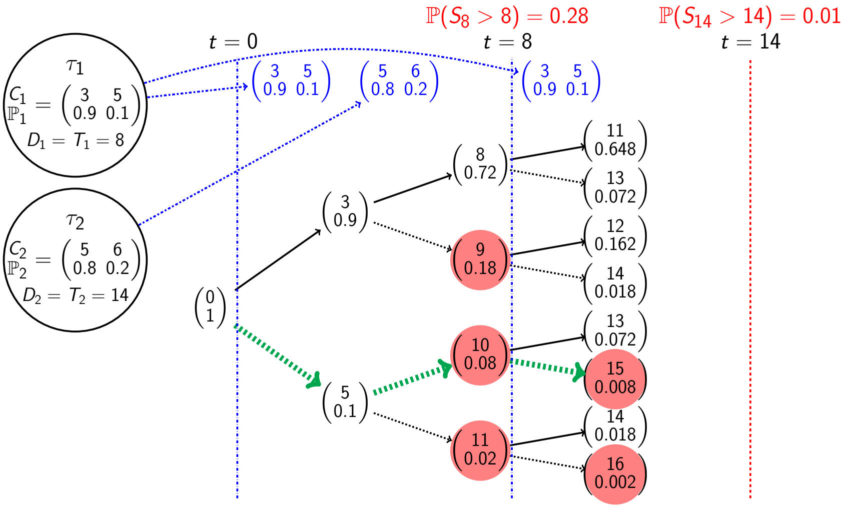

Monte-Carlo Response Time Analysis

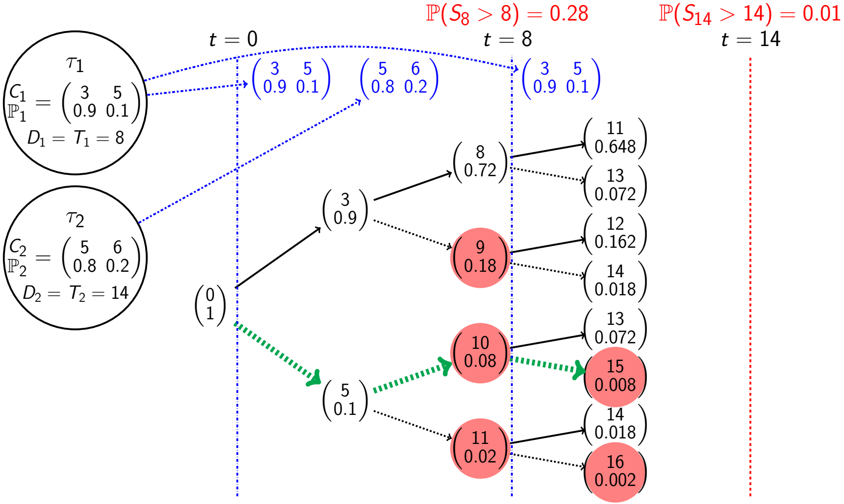

Count: #samples and #misses $\Rightarrow$ Bernoulli trial, bounded error

State of the Art - Calculation

- Job-based + theoretically exact
  - Direct job-level convolution - exponential # jobs
  - Re-sampling (Maxim and Cucu-Grosjean, RTSS 13)
    Improved re-sampling (Markovic et al., ECRTS 21)
    - precision loss potentially large
  - Cyclic convolution (Markovic et al., ECRTS 21)
    + faster than direct convolution
  - Monte-Carlo RTA (Bozhko et al., RTSS 21)
    + scales well with # jobs - precision depends on # samples

State of the Art - Calculation

- Job-based + theoretically exact
  - Direct job-level convolution - exponential # jobs
  - Re-sampling (Maxim and Cucu-Grosjean, RTSS 13)
    Improved re-sampling (Markovic et al., ECRTS 21)
    - precision loss potentially large
  - Cyclic convolution (Markovic et al., ECRTS 21)
    + faster than direct convolution
  - Monte-Carlo RTA (Bozhko et al., RTSS 21)
    + scales well with # jobs - precision depends on # samples
- Interval-based

Interval-Based Approaches

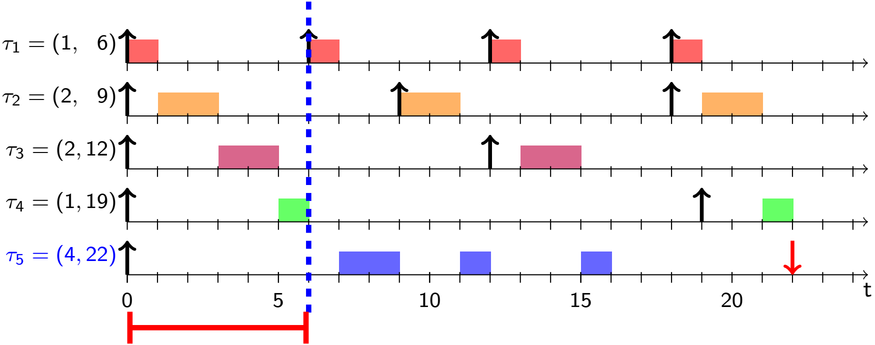

Interval-Based Approaches

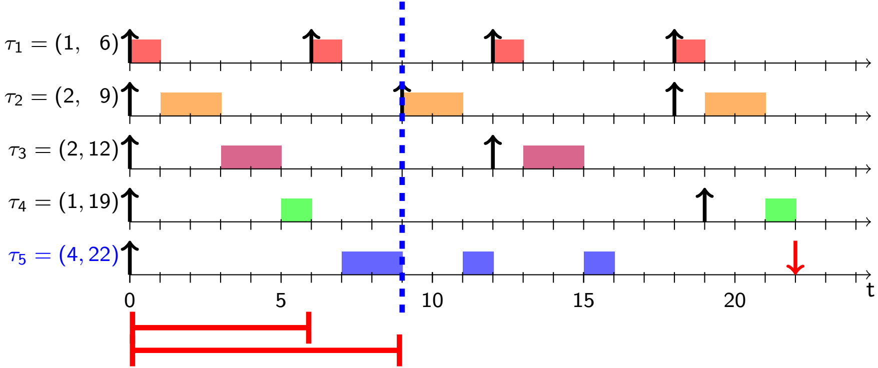

Interval-Based Approaches

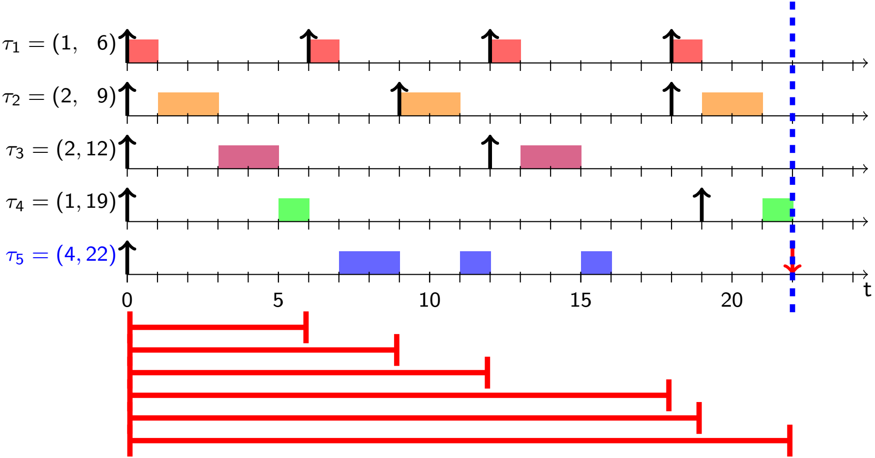

State of the Art - Calculation

- Job-based + theoretically exact
  - Direct job-level convolution - exponential # jobs
  - Re-sampling (Maxim and Cucu-Grosjean, RTSS 13)
    Improved re-sampling (Markovic et al., ECRTS 21)
    - precision loss potentially large
  - Cyclic convolution (Markovic et al., ECRTS 21)
    + faster than direct convolution
  - Monte-Carlo RTA (Bozhko et al., RTSS 21)
    + scales well with # jobs - precision depends on # samples
- Interval-based + scalable - iid assumption - unknown loss
  - Task-level convolution (von der Brüggen et al., ECRTS 18)
    + no additional loss - slow

State of the Art - Calculation

- Job-based + theoretically exact
  - Direct job-level convolution - exponential # jobs
  - Re-sampling (Maxim and Cucu-Grosjean, RTSS 13)
    Improved re-sampling (Markovic et al., ECRTS 21)
    - precision loss potentially large
  - Cyclic convolution (Markovic et al., ECRTS 21)
    + faster than direct convolution
  - Monte-Carlo RTA (Bozhko et al., RTSS 21)
    + scales well with # jobs - precision depends on # samples
- Interval-based + scalable - iid assumption - unknown loss
  - Task-level convolution (von der Brüggen et al., ECRTS 18)
    + no additional loss - slow
  - Bernstein's and Hoeffding's Bound (von der Brüggen et al., ECRTS 18)
    + very fast - very imprecise - unknown additional loss

State of the Art - Calculation

- Job-based + theoretically exact
  - Direct job-level convolution - exponential # jobs
  - Re-sampling (Maxim and Cucu-Grosjean, RTSS 13)
    Improved re-sampling (Markovic et al., ECRTS 21)
    - precision loss potentially large
  - Cyclic convolution (Markovic et al., ECRTS 21)
    + faster than direct convolution
  - Monte-Carlo RTA (Bozhko et al., RTSS 21)
    + scales well with # jobs - precision depends on # samples
- Interval-based + scalable - iid assumption - unknown loss
  - Task-level convolution (von der Brüggen et al., ECRTS 18)
    + no additional loss - slow
  - Bernstein's and Hoeffding's Bound (von der Brüggen et al., ECRTS 18)
    + very fast - very imprecise - unknown additional loss
  - Improved Chernoff Bounds (Chen and Chen, SIES 17 and Chen et al., DATE 19)
    + fast - imprecise - unknown additional loss

State of the Art - Scenarios

- Worst-Case Deadline Failure Probability
- Deadline Miss Rate

State of the Art - Scenarios

- Worst-Case Deadline Failure Probability
  - FP: Related to critical instant (Chen et al., RTSS 22)
- Deadline Miss Rate

State of the Art - Scenarios

- Worst-Case Deadline Failure Probability
  - FP: Related to critical instant (Chen et al., RTSS 22)
  - EDF: Align all deadlines (von der Brüggen et al., RTSS 21)
- Deadline Miss Rate

State of the Art - Scenarios

- Worst-Case Deadline Failure Probability
  - FP: Related to critical instant (Chen et al., RTSS 22)
  - EDF: Align all deadlines (von der Brüggen et al., RTSS 21)
- Deadline Miss Rate
  - FP: —

State of the Art - Scenarios

- Worst-Case Deadline Failure Probability
  - FP: Related to critical instant (Chen et al., RTSS 22)
  - EDF: Align all deadlines (von der Brüggen et al., RTSS 21)
- Deadline Miss Rate
  - FP: —
  - EDF: —

State of the Art - Dependencies

- Either iid (i.e., no dependencies) or acyclic, triggered mode dependencies

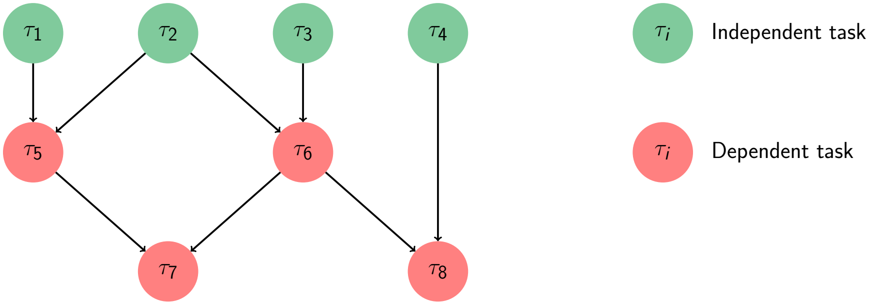

State of the Art - Dependencies

- Either iid (i.e., no dependencies) or acyclic, triggered mode dependencies
- Bounded length of triggered interval

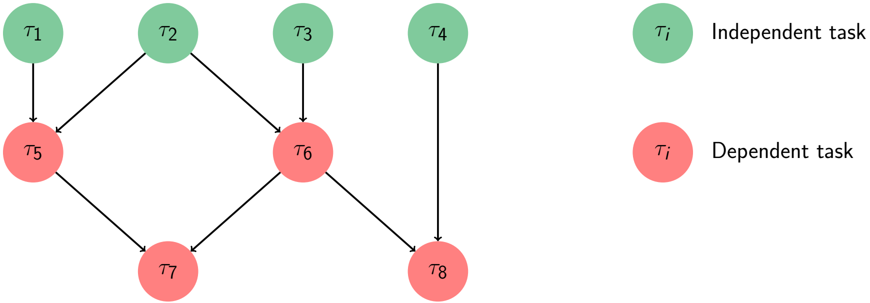

State of the Art - Dependencies

- Either iid (i.e., no dependencies) or acyclic, triggered mode dependencies
- Bounded length of triggered interval
- Over-approximation using task-level convolution

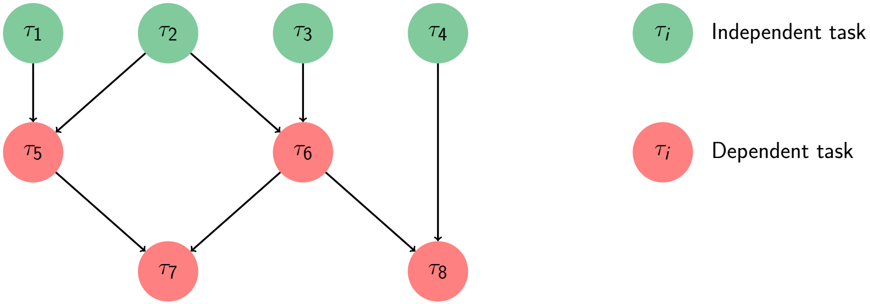

Conclusion - State of the Art for Probabilistic Scheduling

- Scenarios
  - WCDFP - scenario or over-approximation known
  - DMR - scenarios unknown
  - Scenarios when backlog is considered unknown
- Calculation methods with different precision/runtime tradeoffs
- Dependencies increase the complexity
  - Only limited scenarios examined
  - Many calculation methods require independent probabilities

Link - Mode Switch Probability

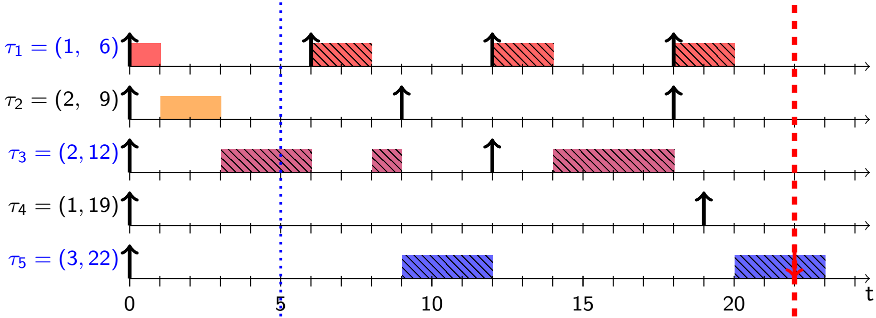

- Mode switch: bottleneck due pre-mode-change low-criticality interference

Link - Mode Switch Probability

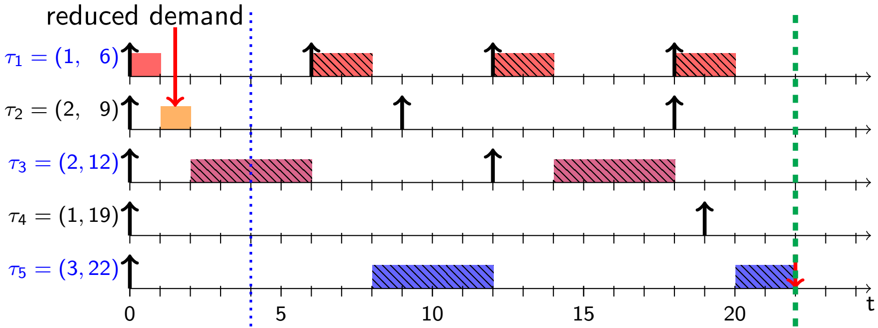

- Mode switch: bottleneck due pre-mode-change low-criticality interference
- Low-criticality resource demand and mode switch often not happen simultaneously

Link - Mode Switch Probability

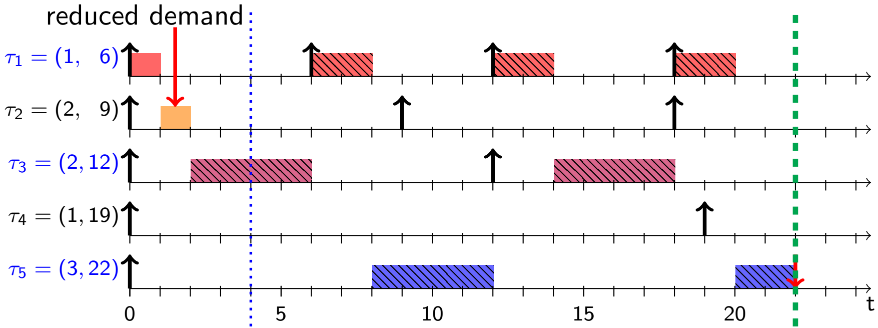

- Mode switch: bottleneck due pre-mode-change low-criticality interference
- Low-criticality resource demand and mode switch often not happen simultaneously

$\Rightarrow$ Exploit that probability of simultaneous occurrence is low during analysis

Link - Probabilistic Guarantees for Low-Criticality Tasks

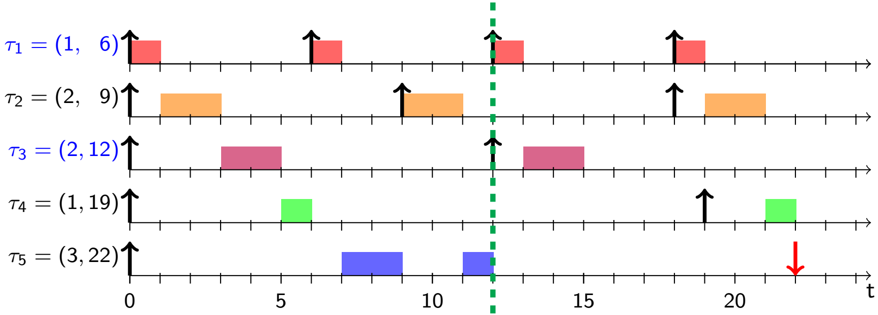

- Idea: Sample multiple schedules (e.g., probabilistic time of the mode switch)
- Determine miss probability for τ_5: Count # samples and # misses

Link - Probabilistic Guarantees for Low-Criticality Tasks

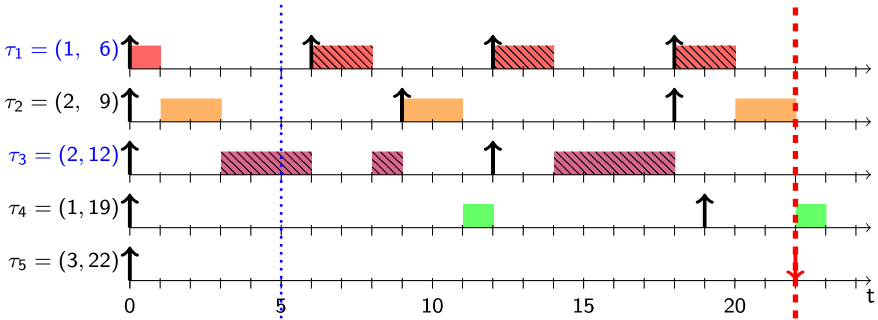

- Idea: Sample multiple schedules (e.g., probabilistic time of the mode switch)
- Determine miss probability for τ_5: Count # samples and # misses

Link - Probabilistic Guarantees for Low-Criticality Tasks

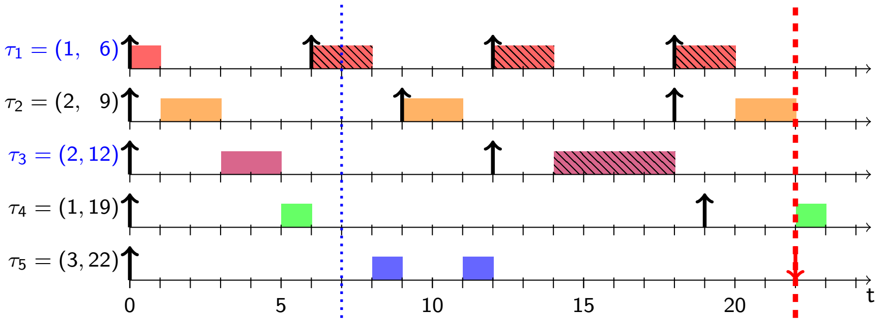

- Idea: Sample multiple schedules (e.g., probabilistic time of the mode switch)
- Determine miss probability for τ_5: Count # samples and # misses

Link - Probabilistic Guarantees for Low-Criticality Tasks

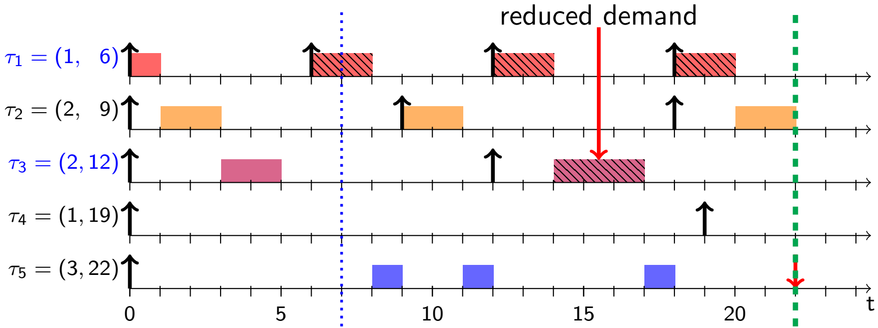

- Idea: Sample multiple schedules (e.g., probabilistic time of the mode switch)
- Determine miss probability for τ_5: Count # samples and # misses
- Consider probabilistic mode switch and probabilistic execution time

Conclusion

- Start discussion which links are potentially interesting
- Three main research questions in probabilistic scheduling
  - Scenario to consider: In many cases unknown
  - Calculation methods: Different precision/runtime tradeoffs
  - Dependent probabilities: Not many solutions due to complexity issues
- Additional possible links between probabilistic scheduling and mixed-criticality
  - Analysing mode switch probability
  - Probabilistic guarantees for low-criticality tasks
  - Interpret high- and low-criticality parameters (different levels of assurance) as different levels of confidence in correspondence to mode probabilities
- Additional challenges
  - Detecting a budged overrun vs. refuting an execution time distribution
  - Complex interplay of execution time dependencies in mixed-criticality systems

Conclusion

- Start discussion which links are potentially interesting
- Three main research questions in probabilistic scheduling
    - Scenario to consider: In many cases unknown
    - Calculation methods: Different precision/runtime tradeoffs
    - Dependent probabilities: Not many solutions due to complexity issues
- Additional possible links between probabilistic scheduling and mixed-criticality
    - Analysing mode switch probability
    - Probabilistic guarantees for low-criticality tasks
    - Interpret high- and low-criticality parameters (different levels of assurance) as different levels of confidence in correspondence to mode probabilities
- Additional challenges
    - Detecting a budged overrun vs. refuting an execution time distribution
    - Complex interplay of execution time dependencies in mixed-criticality systems

Thank You!