Embedded Electronics and Computing Systems
DEPARTMENT OF INFORMATION ENGINEERING AND COMPUTER SCIENCE

UNIVERSITÀ DEGLI STUDI DI TRENTO

Probabilistic Real-Time Guarantees: There is life beyond the i.i.d. assumption

Bernardo Villalba Frías[1], Luigi Palopoli[1], Luca Abeni[2] and Daniele Fontanelli[3]

[1] Department of Information Engineering and Computer Science – University of Trento
[2] ReTiS Laboratory – Scuola Superiore Sant'Anna
[3] Department of Industrial Engineering – University of Trento

Outline

- **Introduction**.
- **Limitations** in the current **Probabilistic Guarantees** analysis.
- The **Markov Computation Time Model (MCTM)**.
- **Stochastic Analysis** of the MCTM.
- **Estimating** the MCTM **parameters**.
- Experimental **Validation**.

Soft Real-Time Systems

- Hard real-time systems:
  - Mature and effective methods.
  - However, many recent real-time systems **are not** hard real-time.
- Soft real-time systems:
  - Resilient to **occasional** and **controlled** timing failures.

Examples

Web TV

Infotainment

Examples

Web TV

Infotainment

Visual Control

Probabilistic Guarantees

- **Performance specifications** stated in **probabilistic terms**:
  - Associate each task with a probability of respecting the deadline.
  - Timing requirement to enforce: respect the deadline with a given probability.
- State of the art: **Numerical methods** and **analytic bounds**.
- Stochastic **analysis** based on the **i.i.d. assumption**.

Reservation-based scheduling

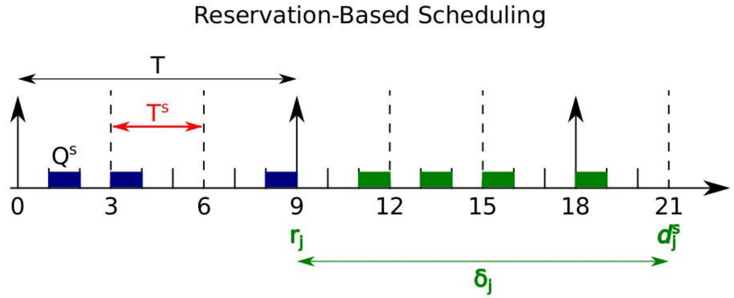

- Resource Reservations:
  - Scheduling parameters: Q^s and T^s.
  - Executes Q^s time units in every reservation period T^s.
  - Guarantees temporal isolation.
- CBS scheduler: SCHED_DEADLINE.
- The probability of **respecting the deadline** is:

$$\mathbf{Pr}\left\{\delta_j \leq T\right\}$$

CBS as a Markov chain

- Model for the evolution of δ_j :

$$v_1 = c_1$$

$$v_j = \max\{0, v_{j-1} - \frac{T}{T^s} Q^s\} + c_j$$

$$\delta_j = \left\lceil \frac{v_j}{Q^s} \right\rceil T^s$$

with $C(c) = \mathbf{Pr}\{c_j = c\}$

CBS as a Markov chain

- Model for the evolution of δ_j :

$$v_1 = c_1$$
$$v_j = \max\{0, v_{j-1} - \frac{T}{T^s} Q^s\} + c_j$$
$$\delta_j = \left\lceil \frac{v_j}{Q^s} \right\rceil T^s$$

with $C(c) = \mathbf{Pr}\{c_j = c\}$

- State of the system:

$$\mathcal{S}_h(j) = \{v_j = c^{\min} + h\}$$

CBS as a Markov chain

- Model for the evolution of δ_j :

$$v_1 = c_1$$
$$v_j = \max\{0, v_{j-1} - \frac{T}{T^s} Q^s\} + c_j$$
$$\delta_j = \left\lceil \frac{v_j}{Q^s} \right\rceil T^s$$

with $C(c) = \mathbf{Pr}\{c_j = c\}$

- State of the system:

$$\mathcal{S}_h(j) = \{v_j = c^{\min} + h\}$$

- Define the probability:

$$\pi_h(j) = \mathbf{Pr}\{\mathcal{S}_h(j)\}$$

CBS as a Markov chain

- Dynamic of the workload modeled as a queue[1]:

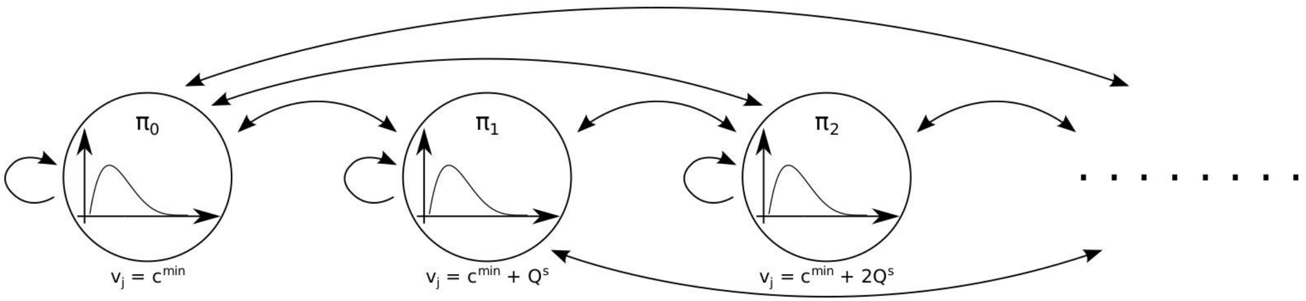

[1] Palopoli, L., Fontanelli, D., Abeni, L. and Villalba Frías, B., "An Analytical Solution for Probabilistic Guarantees of Reservation Based Soft Real-Time Systems", IEEE Transactions on Parallel and Distributed Systems, vol 27, no. 3, pp. 640 – 653, March 2016.

CBS as a Markov chain

- Dynamic of the workload modeled as a queue[1]:

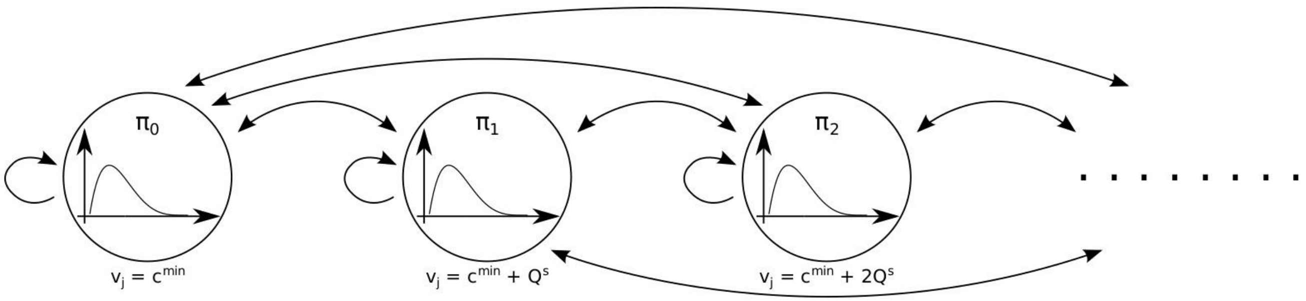

- Introduce the vector:

$$\Pi(j) = [\pi_0(j)\ \ \pi_1(j)\ \ \pi_2(j)\ \ \ldots]$$

[1] Palopoli, L., Fontanelli, D., Abeni, L. and Villalba Frías, B., "An Analytical Solution for Probabilistic Guarantees of Reservation Based Soft Real-Time Systems", IEEE Transactions on Parallel and Distributed Systems, vol 27, no. 3, pp. 640 – 653, March 2016.

CBS as a Markov chain

- Dynamic of the workload modeled as a queue[1]:

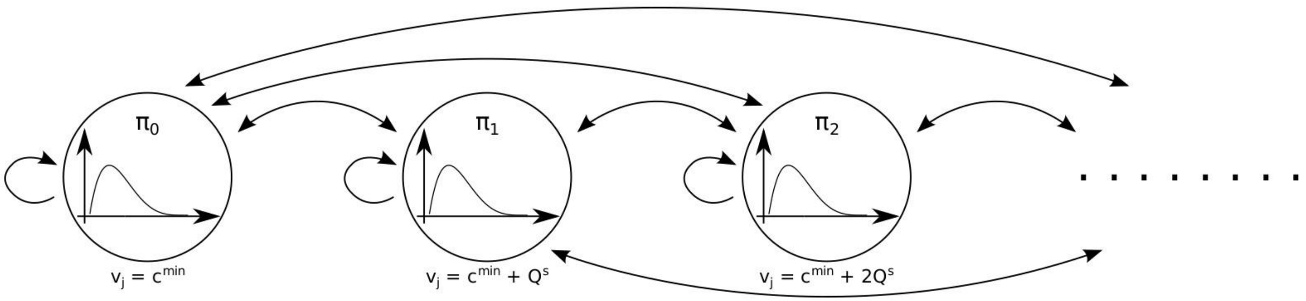

- Introduce the vector:

$$\Pi(j) = [\pi_0(j) \ \ \pi_1(j) \ \ \pi_2(j) \ \ \ldots]$$

- Probability evolution:

$$\Pi(j) = \Pi(j-1)P$$

$$P = \begin{bmatrix} a_0 & a_1 & a_2 & \ldots & a_n & 0 & 0 & 0 & \ldots \\ a_0 & a_1 & a_2 & \ldots & a_n & 0 & 0 & 0 & \ldots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ldots \\ a_0 & a_1 & a_2 & \ldots & a_n & 0 & 0 & 0 & \ldots \\ 0 & a_0 & a_1 & a_2 & \ldots & a_n & 0 & 0 & \ldots \\ 0 & 0 & a_0 & a_1 & a_2 & \ldots & a_n & 0 & \ldots \\ \vdots & \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots \end{bmatrix}$$

[1] Palopoli, L., Fontanelli, D., Abeni, L. and Villalba Frías, B., "An Analytical Solution for Probabilistic Guarantees of Reservation Based Soft Real-Time Systems", IEEE Transactions on Parallel and Distributed Systems, vol 27, no. 3, pp. 640 – 653, March 2016.

CBS as a Markov chain

- Dynamic of the workload modeled as a queue[1]:

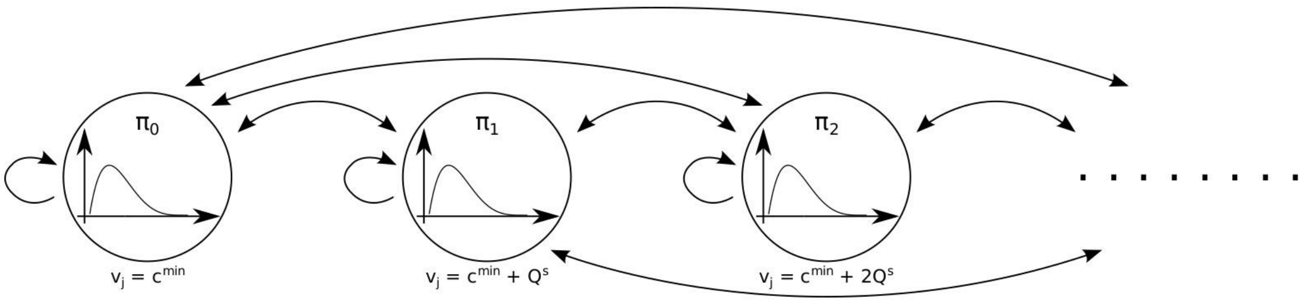

- Introduce the vector:

$$\Pi(j) = [\pi_0(j) \;\; \pi_1(j) \;\; \pi_2(j) \;\; \ldots]$$

- Probability evolution:

$$\Pi(j) = \Pi(j-1)P$$

- Steady state:

$$\overline{\Pi} = \lim_{j \to \infty} \Pi(j)$$

$$P = \begin{bmatrix} a_0 & a_1 & a_2 & \ldots & a_n & 0 & 0 & 0 & \ldots \\ a_0 & a_1 & a_2 & \ldots & a_n & 0 & 0 & 0 & \ldots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ldots \\ a_0 & a_1 & a_2 & \ldots & a_n & 0 & 0 & 0 & \ldots \\ 0 & a_0 & a_1 & a_2 & \ldots & a_n & 0 & 0 & \ldots \\ 0 & 0 & a_0 & a_1 & a_2 & \ldots & a_n & 0 & \ldots \\ \vdots & \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots \end{bmatrix}$$

[1] Palopoli, L., Fontanelli, D., Abeni, L. and Villalba Frías, B., "An Analytical Solution for Probabilistic Guarantees of Reservation Based Soft Real-Time Systems", IEEE Transactions on Parallel and Distributed Systems, vol 27, no. 3, pp. 640 – 653, March 2016.

Results on i.i.d. execution times

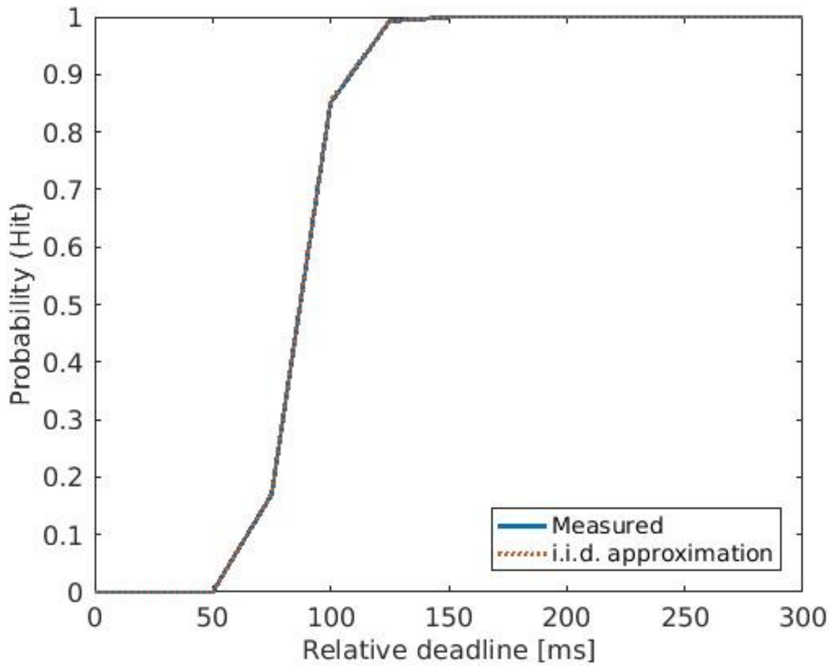

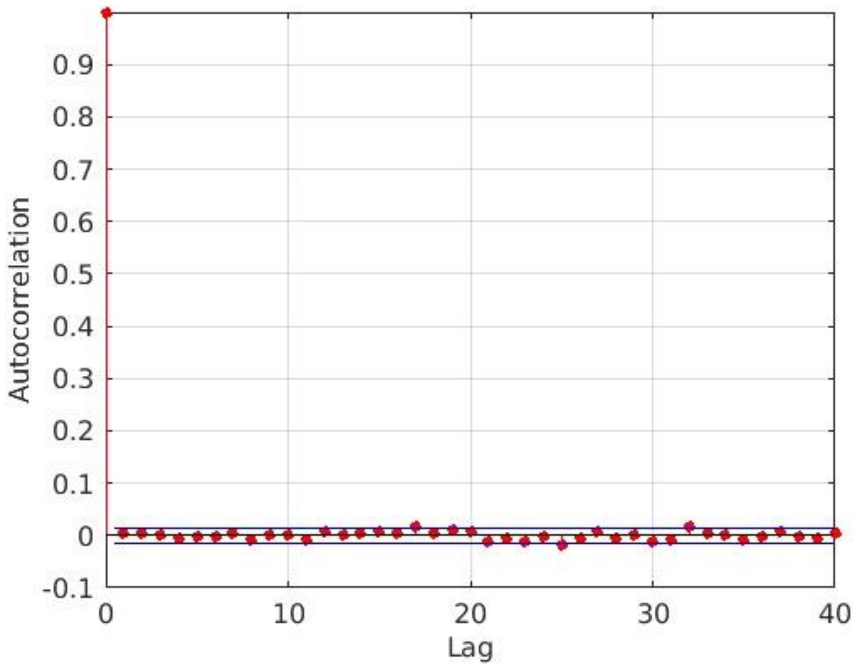

- Test for independence[1]: Runs of "above and below the mean" with 0.05 significance level:

z-statistic	p-value (>0.05)	hypothesis
-0.5295	0.2938	Accepted

[1] Liu, R., Mills, A. and Anderson, J., "Independence Thresholds: Balancing Tractability and Practicality in Soft Real-Time Stochastic Analysis", Proceedings of the IEEE Real-Time Systems Symposium, Rome, Italy, December 2014.

Results on non i.i.d. execution times

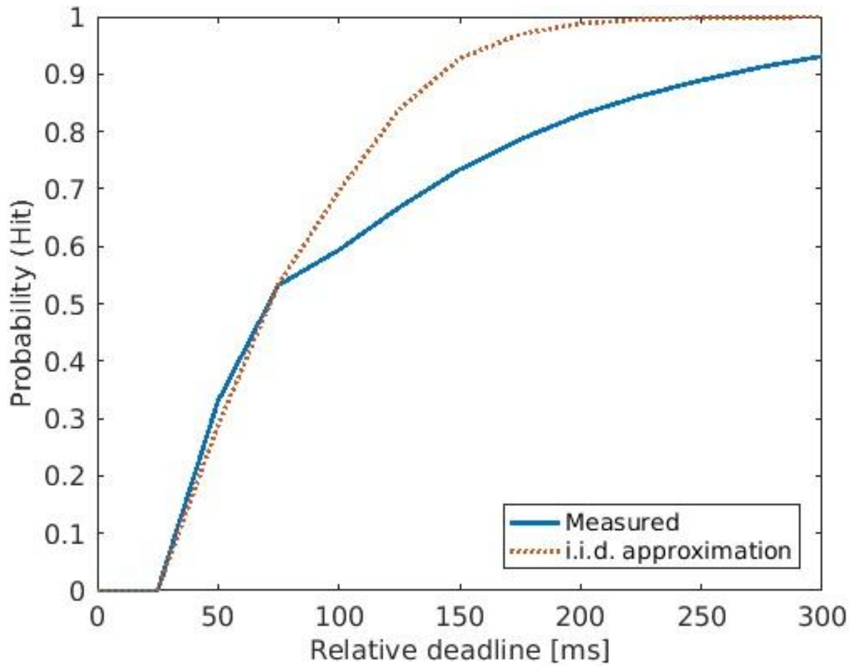

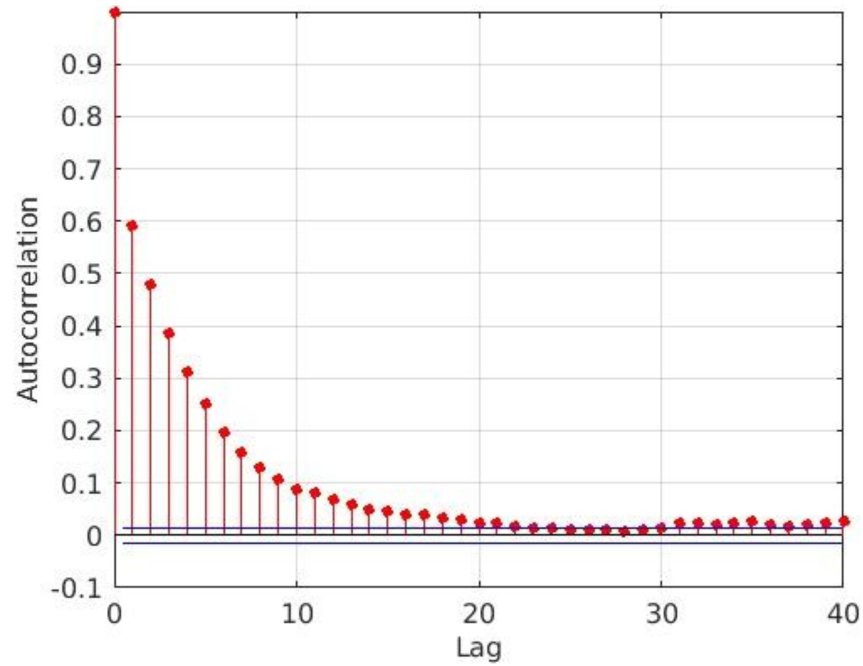

- Test for independence: Runs of "above and below the mean" with 0.05 significance level:

z-statistic	p-value (>0.05)	hypothesis
-100.5715	0.0000	Rejected

A motivating example

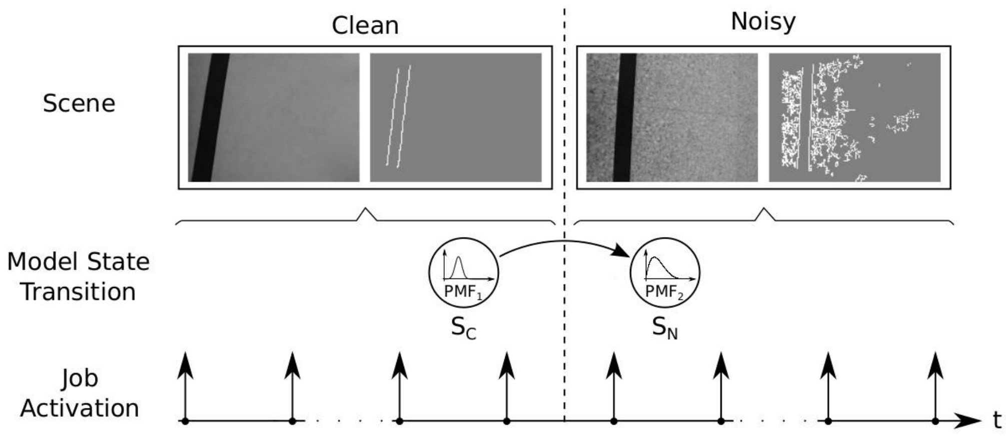

- Extended use of **randomized algorithms** in robotics.
- Path-following robot with image processing.
- Possibly, different types of environment.
- The robot will remain in **one mode** for a while.
- Then, it will **switch mode**… continuously.

The modeling idea

- A **Markov chain** to model the "**mode change**":

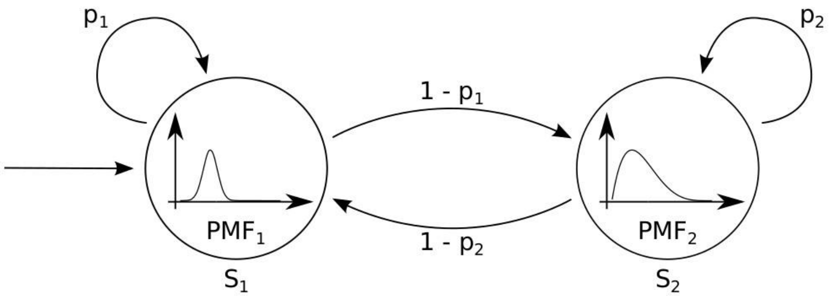

The modeling idea

- A **Markov chain** to model the "**mode change**":

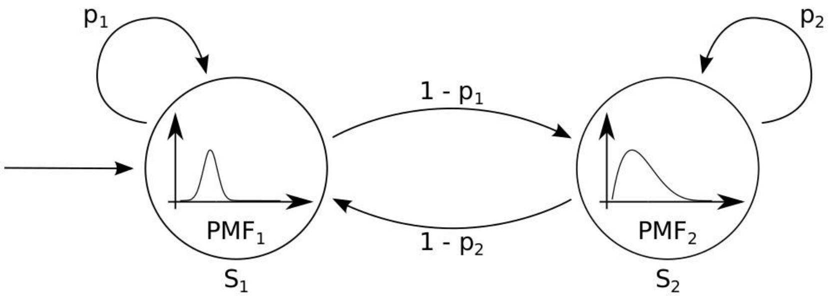

- The switching behavior introduces dependencies.
- Finite number of modes.
- If transitions are defined as a Markov Process.
- Model corresponds to a Markov Modulated Process.
- More precisely: a hidden Markov model.

The Markov Computation Time Model

- A Markov Computation Time Model (MCTM) is defined as the triple $\{\mathcal{M}, \mathcal{P}, \mathcal{C}\}$

$$\mathcal{M} = \{m_1, \ldots, m_N\}$$

$$\mathcal{P} = (p_{a,b}), \ \forall a, b \in \mathcal{M}$$

$$p_{a,b} = \mathbf{Pr}\{m_j = b \,|m_{j-1} = a\}$$

$$\mathcal{C} = \left\{C_{m_j} : m_j \in \mathcal{M}\right\}$$

The Markov Computation Time Model

- A Markov Computation Time Model (MCTM) is defined as the triple $\{\mathcal{M}, \mathcal{P}, \mathcal{C}\}$

$$\mathcal{M} = \{m_1, \ldots, m_N\}$$

$$\mathcal{P} = (p_{a,b}), \;\; \forall a, b \in \mathcal{M}$$

$$p_{a,b} = \mathbf{Pr}\{m_j = b \,|\, m_{j-1} = a\}$$

$$\mathcal{C} = \left\{C_{m_j} : m_j \in \mathcal{M}\right\}$$

- Assumptions:
  - Job's execution time **only depends** on the current mode.
  - The "mode change" event is **independent** both from:
    - The current computation workload.
    - The execution time required by the previous job.

The Markov Computation Time Model

1

By modeling the **dependencies** as a Markov system with a **switching behavior**, it is possible to **describe**, more precisely, the **non i.i.d. execution times**.

Stochastic Analysis

- The new model of the system:

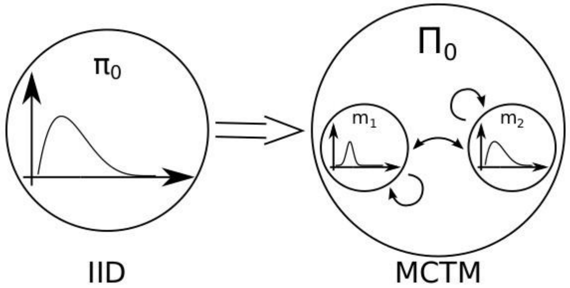

Stochastic Analysis

- The new model of the system:

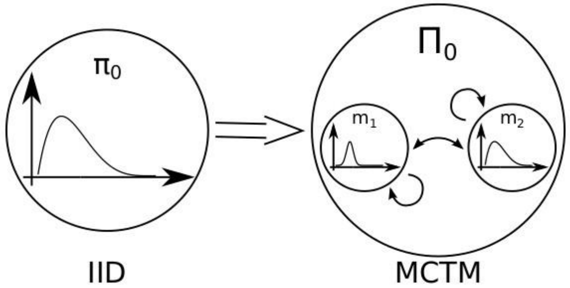

$$\mathcal{S}_h(j) = \{v_j = c^{\min} + h\}$$
$$\pi_h(j) = \mathbf{Pr}\{\mathcal{S}_h(j)\}$$
$$\Pi(j) = [\pi_0(j)\ \ \pi_1(j)\ \ \pi_2(j)\ \ \ldots]$$

$$\Longrightarrow$$

$$\mathcal{S}_{g,\,h}(j) = \{m_j = g\} \wedge \{v_j = c^{\min} + h\}$$
$$\pi_{g,\,h}(j) = \mathbf{Pr}\{\mathcal{S}_{g,\,h}(j)\}$$
$$\Pi_h(j) = [\pi_{1,h}(j)\ \ \pi_{2,h}(j)\ \ \ldots\ \ \pi_{N,h}(j)]$$
$$\Pi(j) = [\Pi_0(j)\ \ \Pi_1(j)\ \ \Pi_2(j)\ \ \ldots]$$

Stochastic Analysis

- The new model of the system:

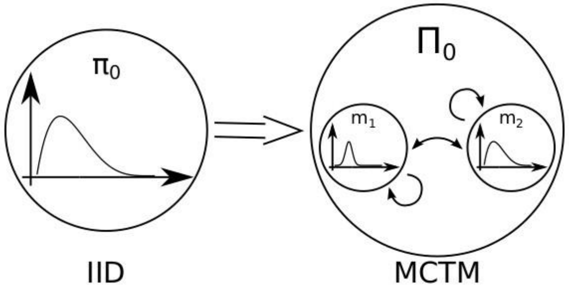

$$\mathcal{S}_h(j) = \{v_j = c^{\min} + h\}$$
$$\pi_h(j) = \mathbf{Pr}\{\mathcal{S}_h(j)\}$$
$$\Pi(j) = [\pi_0(j) \;\; \pi_1(j) \;\; \pi_2(j) \;\; \ldots]$$

$$\Longrightarrow$$

$$\mathcal{S}_{g,\,h}(j) = \{m_j = g\} \wedge \{v_j = c^{\min} + h\}$$
$$\pi_{g,\,h}(j) = \mathbf{Pr}\{\mathcal{S}_{g,\,h}(j)\}$$
$$\Pi_h(j) = [\pi_{1,h}(j) \;\; \pi_{2,h}(j) \;\; \ldots \;\; \pi_{N,h}(j)]$$
$$\Pi(j) = [\Pi_0(j) \;\; \Pi_1(j) \;\; \Pi_2(j) \;\; \ldots]$$

- Probability evolution:

$$\Pi(j) = \Pi(j-1)P$$
$$\overline{\Pi} = \lim_{j \to \infty} \Pi(j)$$

A Markov chain

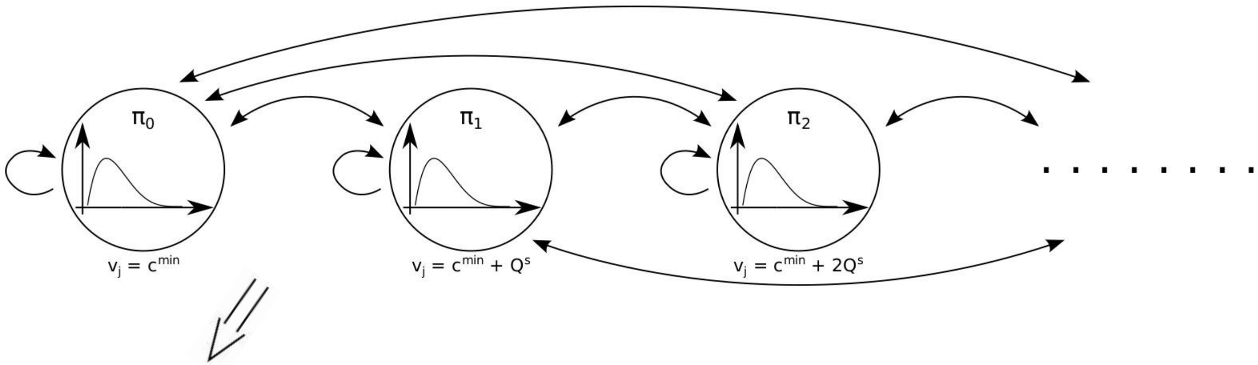

$$P = \begin{bmatrix} a_0 & a_1 & a_2 & \dots & a_n & 0 & 0 & 0 & \dots \\ a_0 & a_1 & a_2 & \dots & a_n & 0 & 0 & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \dots \\ a_0 & a_1 & a_2 & \dots & a_n & 0 & 0 & 0 & \dots \\ 0 & a_0 & a_1 & a_2 & \dots & a_n & 0 & 0 & \dots \\ 0 & 0 & a_0 & a_1 & a_2 & \dots & a_n & 0 & \dots \\ \vdots & \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots \end{bmatrix}$$

A Markov chain

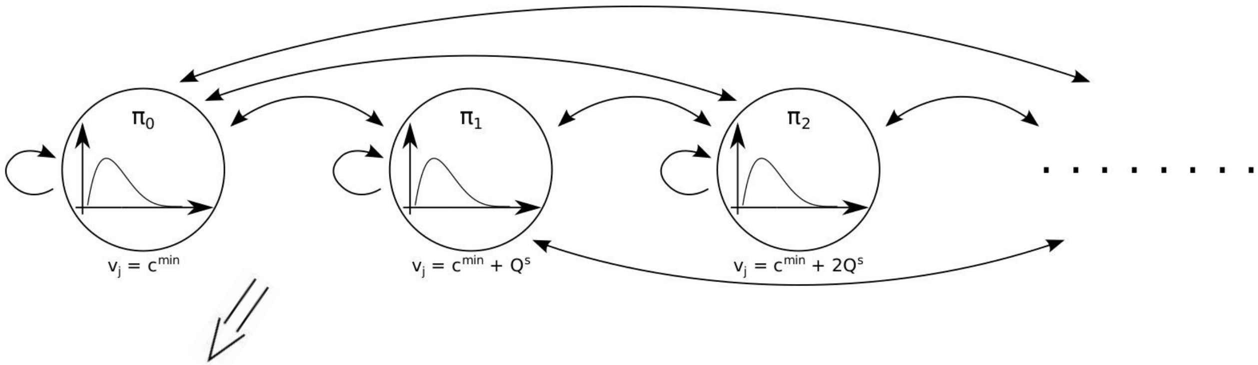

$$P = \begin{bmatrix} a_0 & a_1 & a_2 & \dots & a_n & 0 & 0 & 0 & \dots \\ a_0 & a_1 & a_2 & \dots & a_n & 0 & 0 & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \dots \\ a_0 & a_1 & a_2 & \dots & a_n & 0 & 0 & 0 & \dots \\ 0 & a_0 & a_1 & a_2 & \dots & a_n & 0 & 0 & \dots \\ 0 & 0 & a_0 & a_1 & a_2 & \dots & a_n & 0 & \dots \\ \vdots & \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots \end{bmatrix}$$

$$P = \begin{bmatrix} P_0 & P_1 & P_2 & \dots & P_R & \mathbf{0} & \mathbf{0} & \mathbf{0} & \dots \\ P_0 & P_1 & P_2 & \dots & P_R & \mathbf{0} & \mathbf{0} & \mathbf{0} & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \dots \\ P_0 & P_1 & P_2 & \dots & P_R & \mathbf{0} & \mathbf{0} & \mathbf{0} & \dots \\ \mathbf{0} & P_0 & P_1 & P_2 & \dots & P_R & \mathbf{0} & \mathbf{0} & \dots \\ \mathbf{0} & \mathbf{0} & P_0 & P_1 & P_2 & \dots & P_R & \mathbf{0} & \dots \\ \vdots & \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots \end{bmatrix}$$

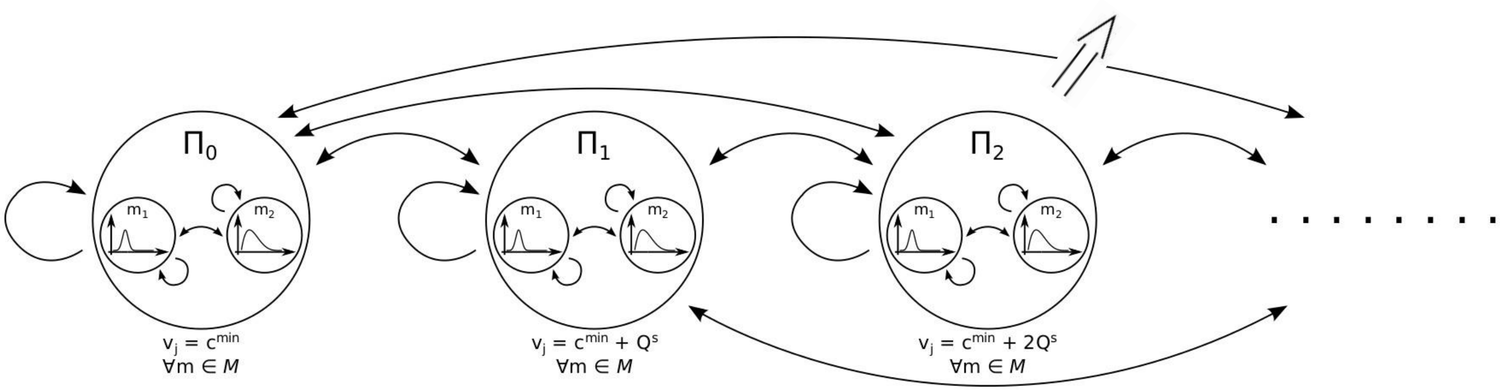

Stochastic Analysis

- Each individual block (P_e) is given by:

$$P_e = \begin{bmatrix} p_{1,1} \cdot \alpha_{1,e} & p_{1,2} \cdot \alpha_{2,e} & \cdots & p_{1,N} \cdot \alpha_{N,e} \\ p_{2,1} \cdot \alpha_{1,e} & p_{2,2} \cdot \alpha_{2,e} & \cdots & p_{2,N} \cdot \alpha_{N,e} \\ \cdots & \cdots & \cdots & \cdots \\ p_{N,1} \cdot \alpha_{1,e} & p_{N,2} \cdot \alpha_{2,e} & \cdots & p_{N,N} \cdot \alpha_{N,e} \end{bmatrix}$$

with:

$$p_{a,\,b} = \mathbf{Pr}\left\{m_j = b \,|\, m_{j-1} = a\right\}$$

$$\alpha_{b,\,h} = C_b(c^{\min} + h)$$

Stochastic Analysis

- Each individual block (P_e) is given by:

$$P_e = \begin{bmatrix} p_{1,1} \cdot \alpha_{1,e} & p_{1,2} \cdot \alpha_{2,e} & \dots & p_{1,N} \cdot \alpha_{N,e} \\ p_{2,1} \cdot \alpha_{1,e} & p_{2,2} \cdot \alpha_{2,e} & \dots & p_{2,N} \cdot \alpha_{N,e} \\ \dots & \dots & \dots & \dots \\ p_{N,1} \cdot \alpha_{1,e} & p_{N,2} \cdot \alpha_{2,e} & \dots & p_{N,N} \cdot \alpha_{N,e} \end{bmatrix}$$

with:

$$p_{a,\,b} = \mathbf{Pr}\left\{m_j = b \,|m_{j-1} = a\right\}$$

$$\alpha_{b,\,h} = C_b(c^{\min} + h)$$

- Numerically solved:
  - Cyclic or Logarithmic Reduction.

Stochastic Analysis

- Each individual block (P_e) is given by:

$$P_e = \begin{bmatrix} p_{1,1} \cdot \alpha_{1,e} & p_{1,2} \cdot \alpha_{2,e} & \cdots & p_{1,N} \cdot \alpha_{N,e} \\ p_{2,1} \cdot \alpha_{1,e} & p_{2,2} \cdot \alpha_{2,e} & \cdots & p_{2,N} \cdot \alpha_{N,e} \\ \cdots & \cdots & \cdots & \cdots \\ p_{N,1} \cdot \alpha_{1,e} & p_{N,2} \cdot \alpha_{2,e} & \cdots & p_{N,N} \cdot \alpha_{N,e} \end{bmatrix}$$

with:

$$p_{a,\,b} = \mathbf{Pr}\left\{m_j = b \,|m_{j-1} = a\right\}$$

$$\alpha_{b,\,h} = C_b(c^{\min} + h)$$

- Numerically solved:
  - Cyclic or Logarithmic Reduction.
- Steady state distribution of the response time:

$$\lim_{j\to\infty} \mathbf{Pr}\left\{\delta_j \leq D\right\}$$

Stochastic Analysis

2

By presenting **the MCTM** as a QBDP and using the **available numeric solutions**, it is possible to obtain the **probability of respecting the deadline** for the proposed model.

Parameter Identification

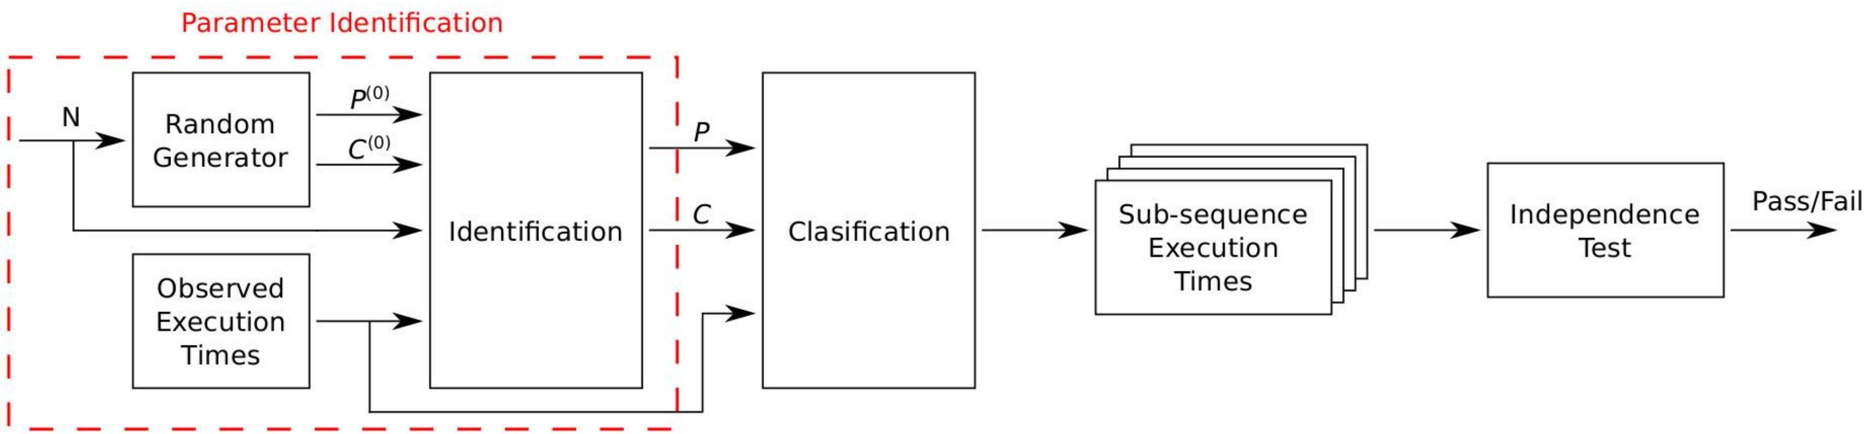

- **Identify** the values of $\{\mathcal{M}, \mathcal{P}, \mathcal{C}\}$
- Solved by the **Baum-Welch algorithm**:
  - Iterative estimation of the parameters.
  - Convergence to the maximum likelihood matrices.

Parameter Validation

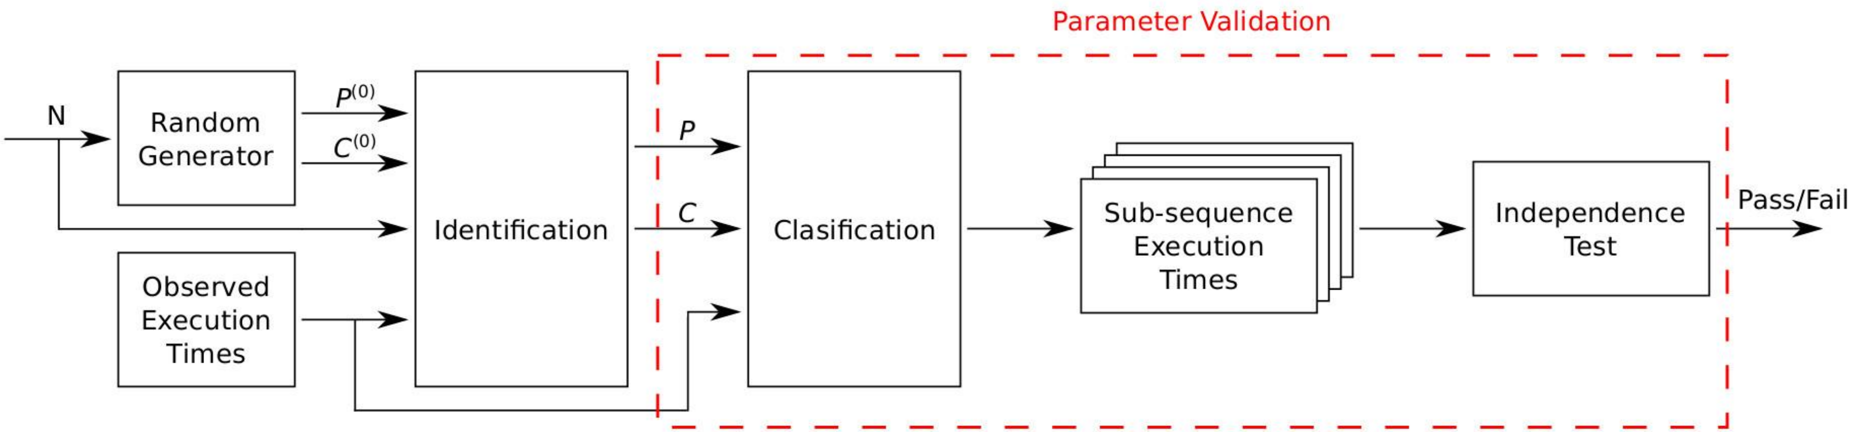

- **Validate** the estimated values of $\{\mathcal{M}, \mathcal{P}, \mathcal{C}\}$
- Solved by the **Viterbi algorithm**:
  - Generates a sequence of hidden states.
  - Obtains N sub-sequences of execution times.
- **Perform** a numerical **test for independence**.

Results on non i.i.d. execution times

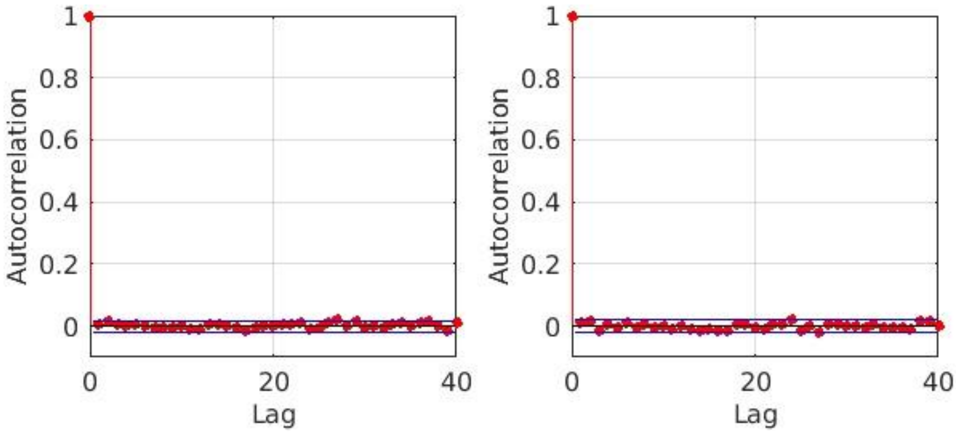

Mode	z-statistic	p-value (>0.05)	hypothesis
1	0.6585	0.7449	Accepted
2	-0.3214	0.3739	Accepted

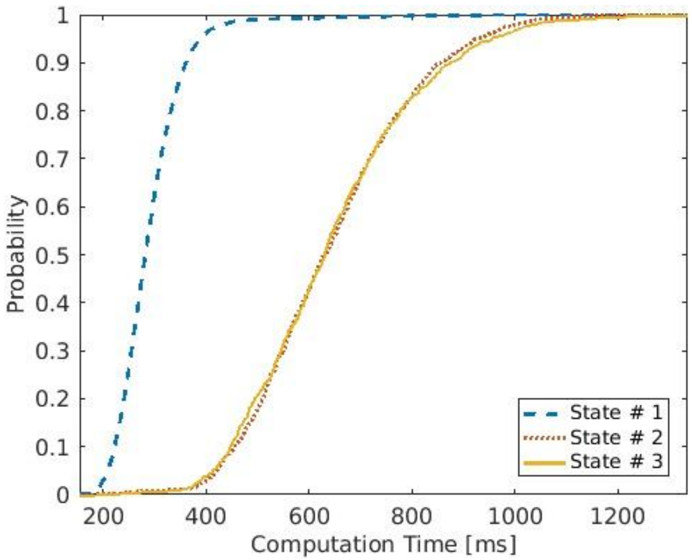

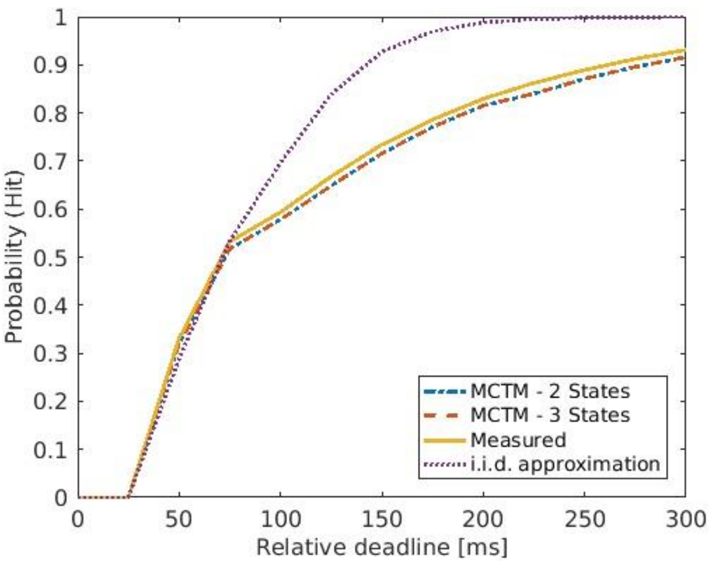

Parameter Estimation

3

By applying **standard identification techniques** on hidden Markov models, it is possible to **estimate the parameters** describing the **Markov Computation Time Model**.

Experimental Setup

- Robotic vision application.
- Two operating conditions.
- 100 sequences of observations.
- 18400 execution times per sequence.
- WandBoard running Ubuntu.
- Linux Kernel 4.8.1

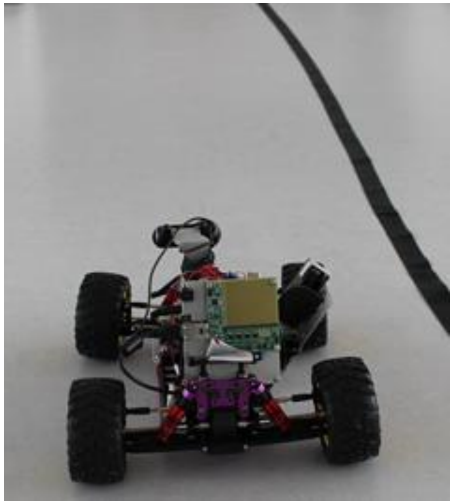

Parameter Estimation

- Execution times obtained from image processing algorithm.
- Strong autocorrelation.
- Test for independence: Runs of “above and below the mean” with 0.05 significance level.

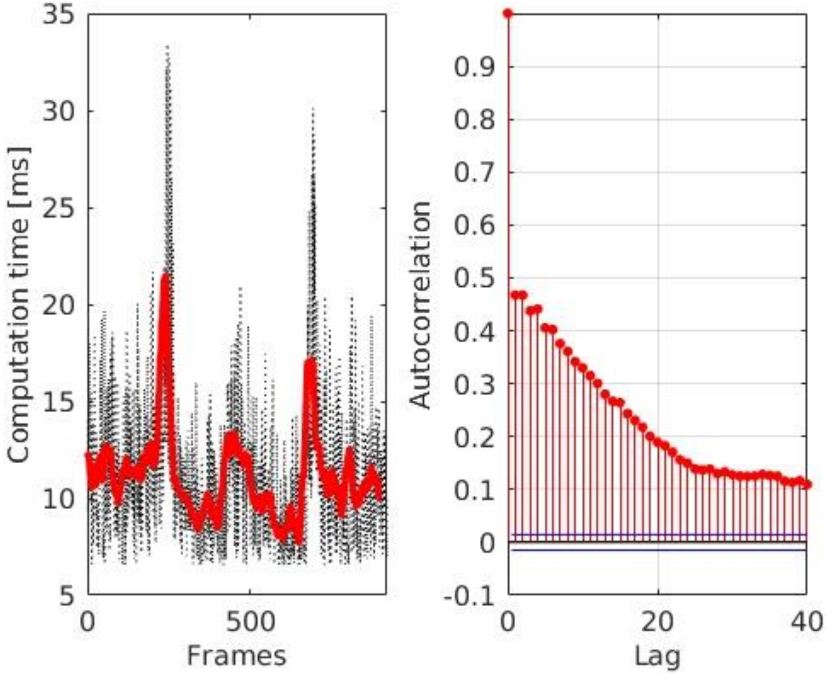

z-statistic	p-value (>0.05)	hypothesis
-37.3271	0.0000	Rejected

Parameter Estimation

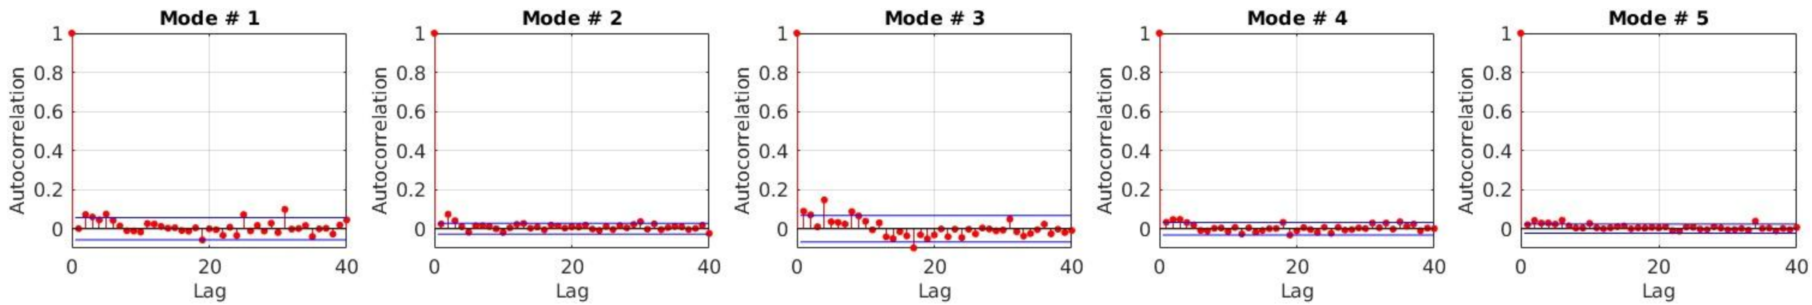

- Applying the estimation technique for different input data.
- Consistently identified a 5-modes MCTM.
- Test for independence: Runs of “above and below the mean” with 0.05 significance level.

Mode	z-statistic	p-value (>0.05)	hypothesis
1	-1.3929	0.0818	Accepted
2	-1.1932	0.1164	Accepted
3	-1.1088	0.1338	Accepted
4	-1.5830	0.0567	Accepted
5	-0.6522	0.2571	Accepted

Fixed bandwidth (16%)

- Scheduling parameters:
  - T = 100 ms.
  - (T^s, Q^s) = (25 ms, 4 ms).
- Probability of **respecting the deadline**:

$$\frac{\#\text{ of jobs respecting the deadline}}{\#\text{ of jobs}}$$

- **Overestimation** when considering the i.i.d. model.
- **Good match**: MCTM vs. Real application.

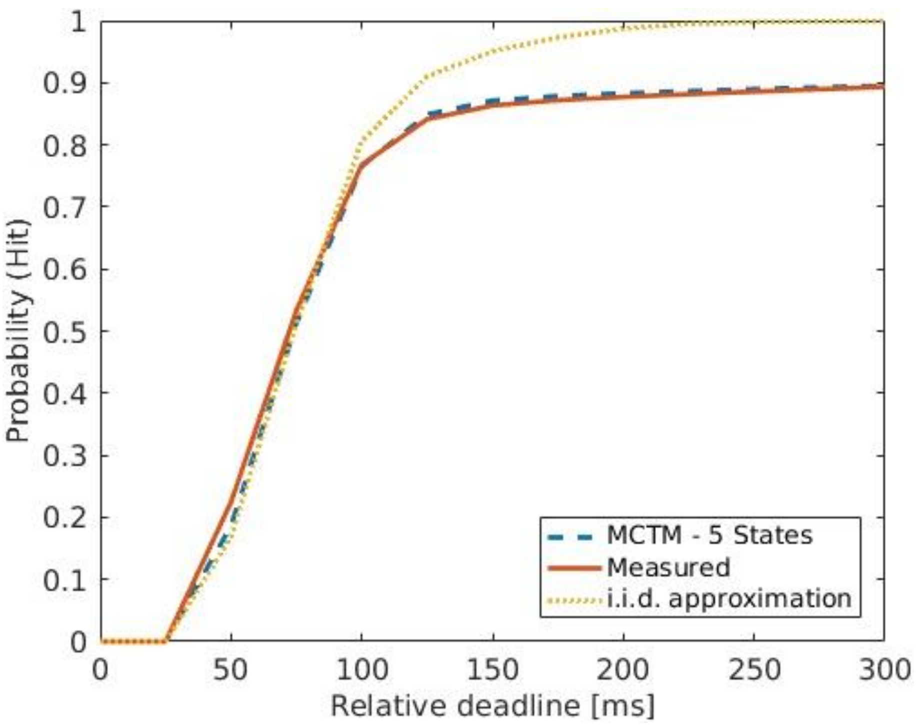

Fixed bandwidth (70%)

- Scheduling parameters:
  - T = 200 ms.
  - (T^s, Q^s) = (50 ms, 35 ms).
- Probability of **respecting the deadline**:

$$\frac{\#\text{ of jobs respecting the deadline}}{\#\text{ of jobs}}$$

- **Overestimation** when considering the i.i.d. model
- **Good match**: MCTM vs. Real application.

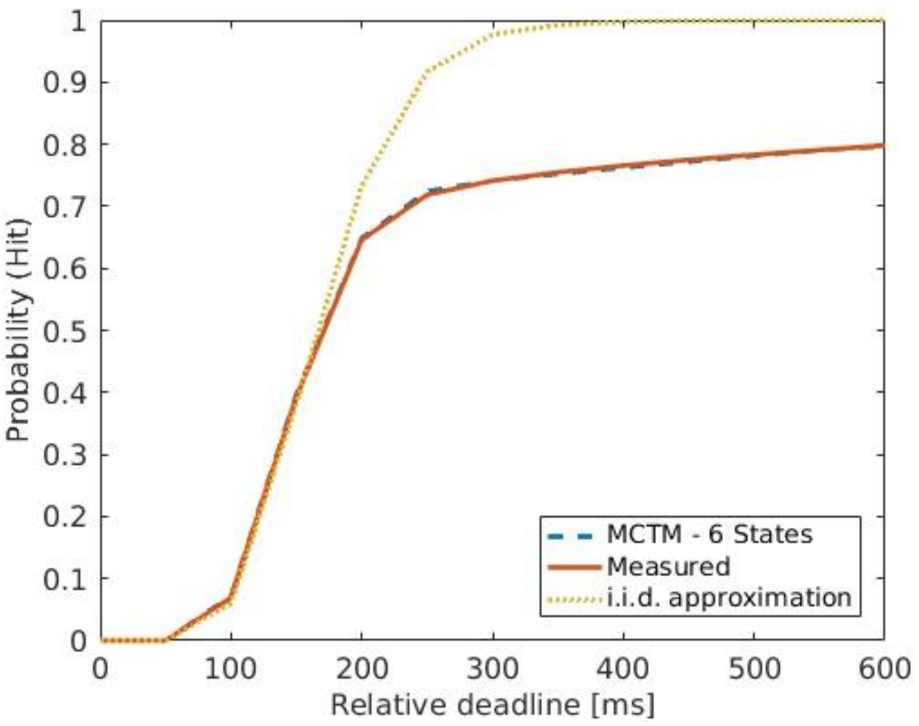

Fixed deadline (D = T)

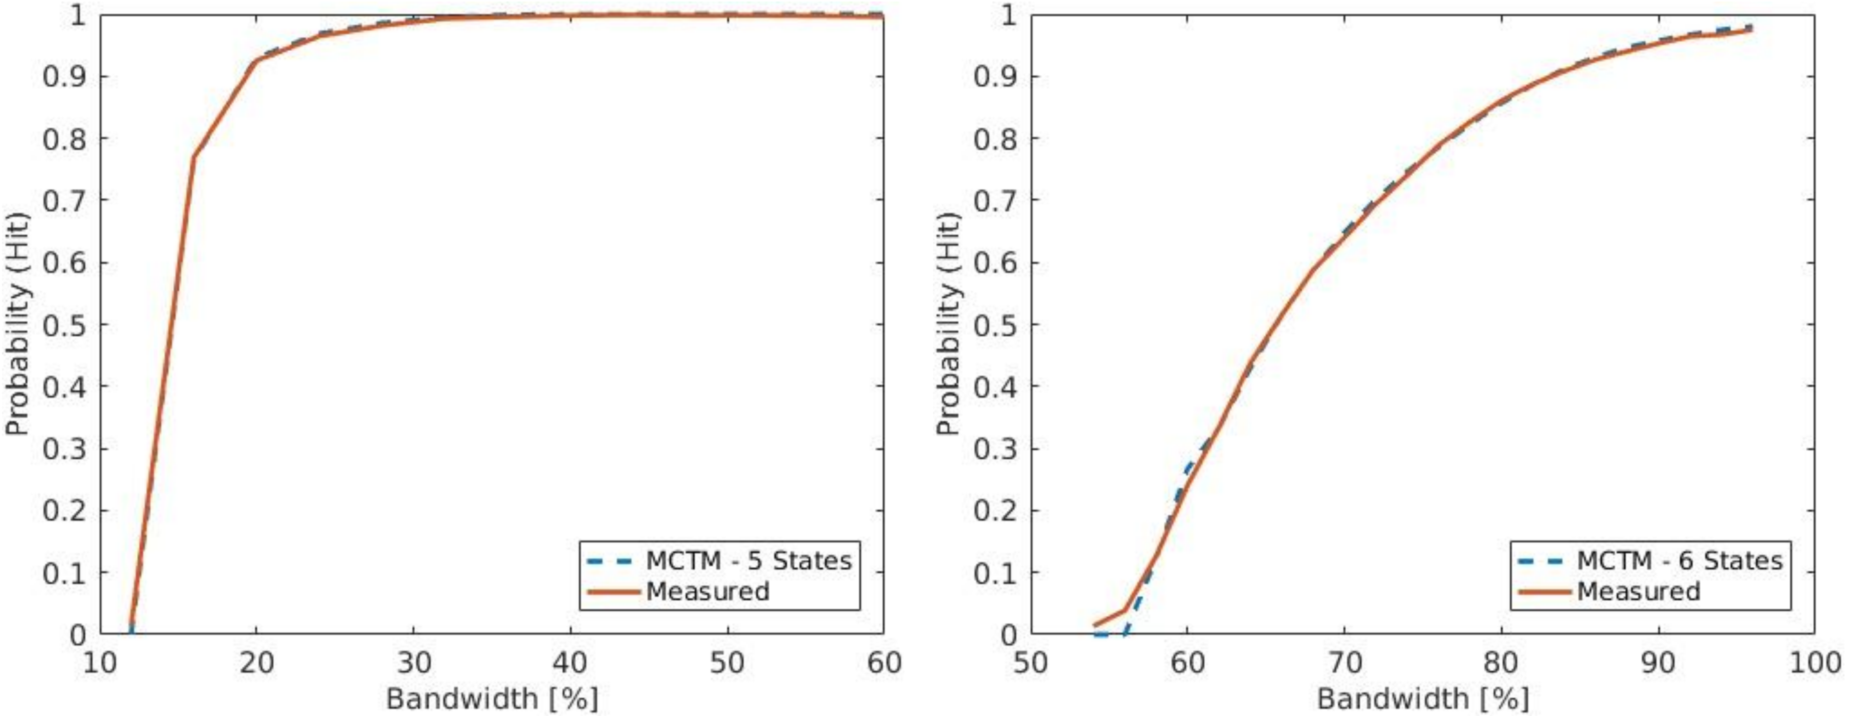

- Different values for the bandwidth are explored.
- **Good performance** of the approach compared with the real application.

Conclusions

- Provided **probabilistic guarantees** for soft real-time systems **characterized by dependencies** in the **execution times**.
- Introduced a **Markovian representation** of the system to **model** these **dependencies**.
- Adapted the techniques for **probabilistic guarantees** to the case of **MCTM**.
- Shown a technique for the **estimation** of the **MCTM parameters**.

Probabilistic Real-Time Guarantees: There is life beyond the i.i.d. assumption

Thanks.

Bernardo Villalba Frías
br.villalbafrias@unitn.it