Re-Sampling for Statistical Timing Analysis of Real-Time Systems

Dorin Maxim, Mike Houston,

Luca Santinelli, Guillem Bernat,

Robert Davis and Liliana Cucu-Grosjean

PROARTIS · Inria · Université de Lorraine · The University of York · Rapita Systems

Probabilistic Real-Time Systems?

- Deterministic analysis can lead to significant overprovision in the system architecture.

- An alternative approach is to use probabilistic analysis. System reliability is typically expressed in terms of probabilities for hardware failures, memory failures, software faults, etc.

- For example, the reliability requirements placed on the timing behaviour of a system might indicate that the timing failure rate must be less than 10^{-9} per hour of operation.

Task Model

A set of n independent periodic tasks $\Gamma = \{\tau_1, \tau_2, \ldots, \tau_n\}$

Each task τ_i generates an infinite number of jobs

Jobs are independent of other jobs of the same task and those of other tasks

τ_i is characterized by:

$$\tau_i = (\mathcal{C}_i, T_i, D_i)$$

T_i being its period;

D_i being its relative deadline;

$\mathcal{C}_i$ being its execution time described by a ***random variable:***

$$\mathcal{C}_i = \begin{pmatrix} c_{i,k} \\ P(C_i = c_{i,k}) \end{pmatrix}$$

Response Time Computation

$$R_{i,j} = B_i(\lambda_{i,j}) \otimes I_i(\lambda_{i,j}) \otimes C_i$$

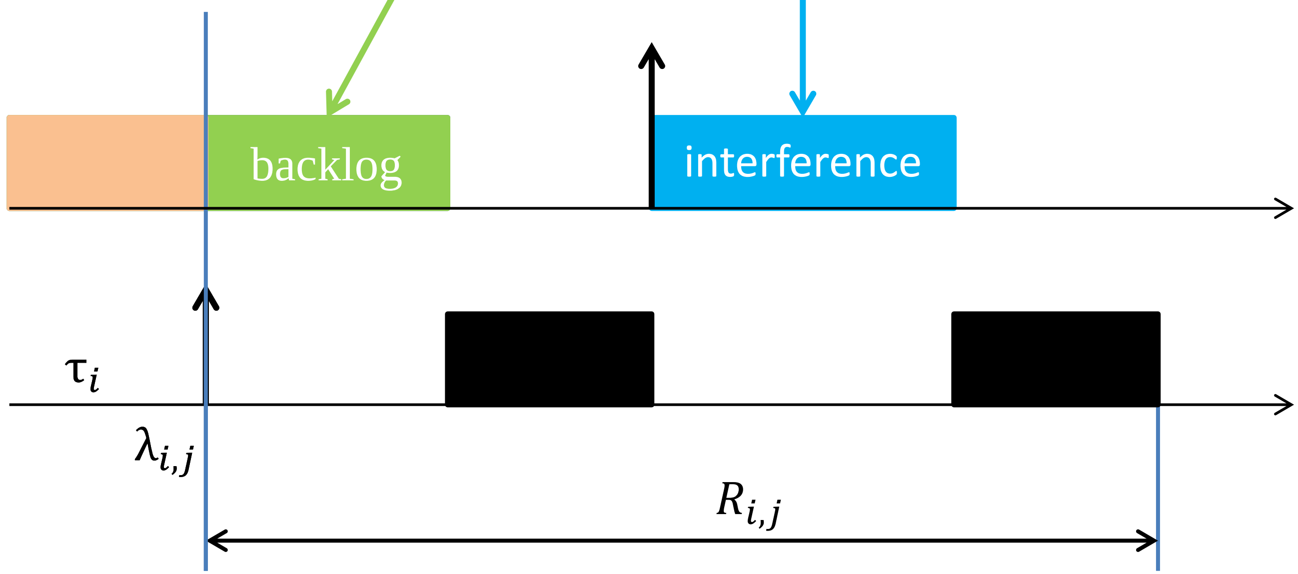

Response Time Computation

$$R_{i,j} = B_i(\lambda_{i,j}) \otimes I_i(\lambda_{i,j}) \otimes C_i$$

convolution

Example of Convolution

$$\begin{pmatrix} 2 & 7 \\ 0.6 & 0.4 \end{pmatrix} \otimes \begin{pmatrix} 11 & 18 \\ 0.9 & 0.1 \end{pmatrix} =$$

$$\begin{pmatrix} 13 & 18 & 20 & 25 \\ 0.54 & 0.36 & 0.06 & 0.04 \end{pmatrix}$$

Note: The result of a convolution between two random variables can have up to *m*n* values, where *m* and *n* are the cardinals of the two variables being convoluted, and for the result to be obtained *m*n* multiplications need to be performed.

Cost of Convolution

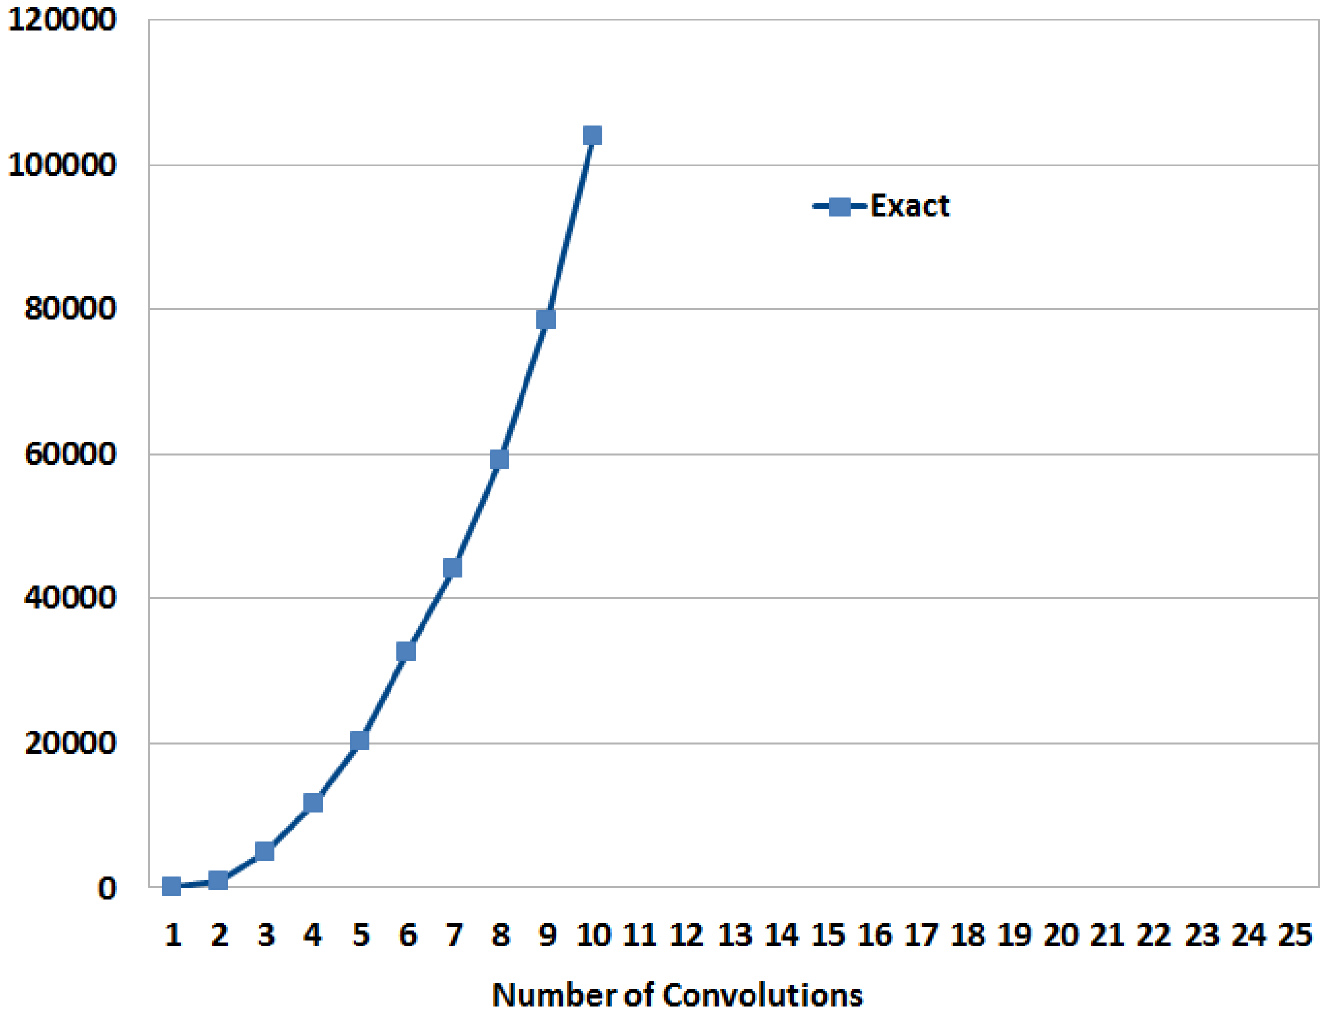

Reducing complexity: Re-sampling

The process of re-sampling a distribution $\mathcal{C}_i$ consists of reducing the initial distribution $\mathcal{C}_i$ from *n* values to *k* values, while not being optimistic.

Re-sampling

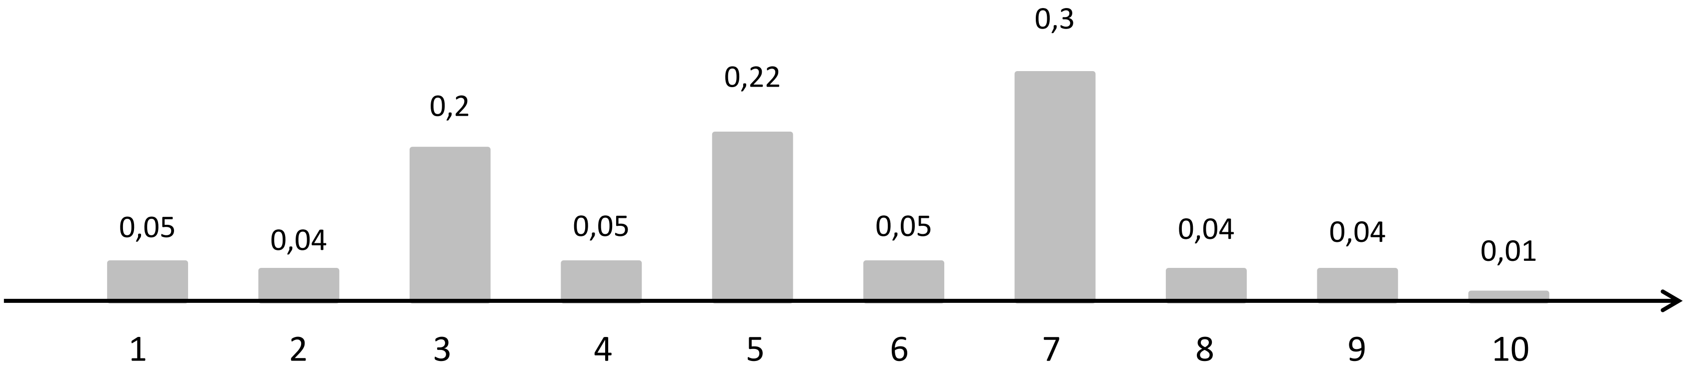

Re-sampling

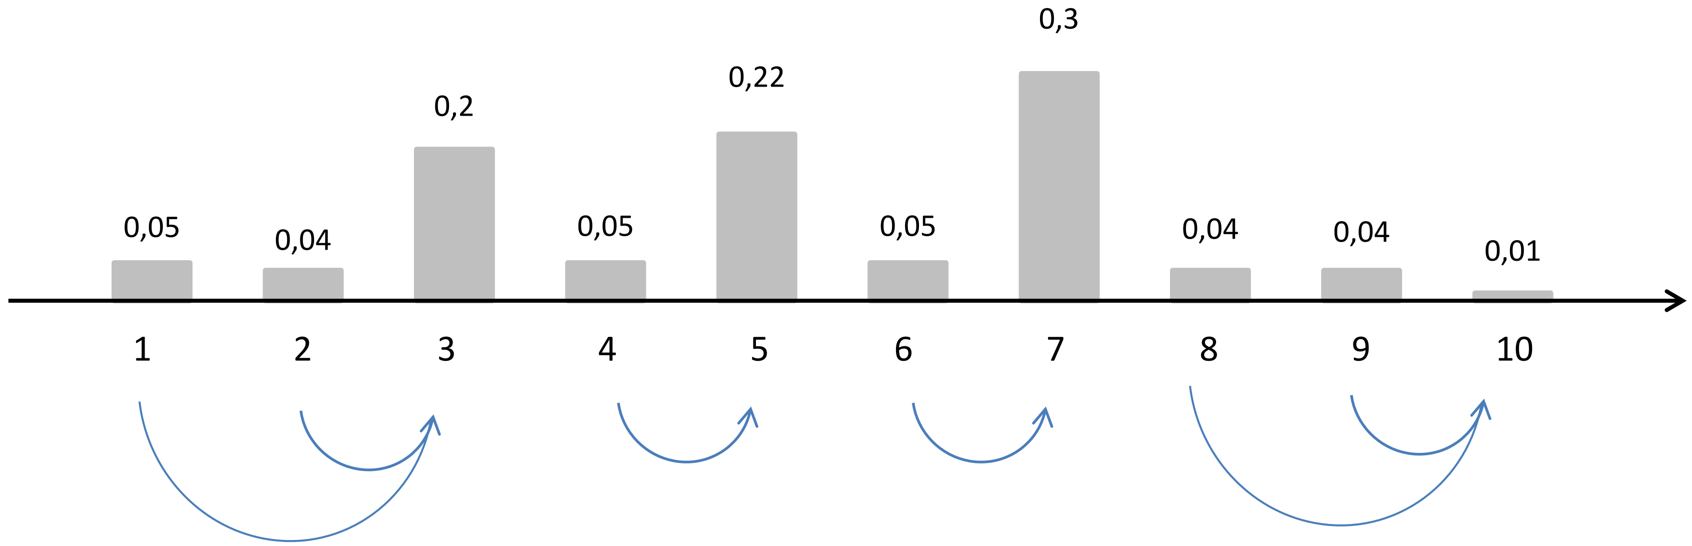

Re-sampling

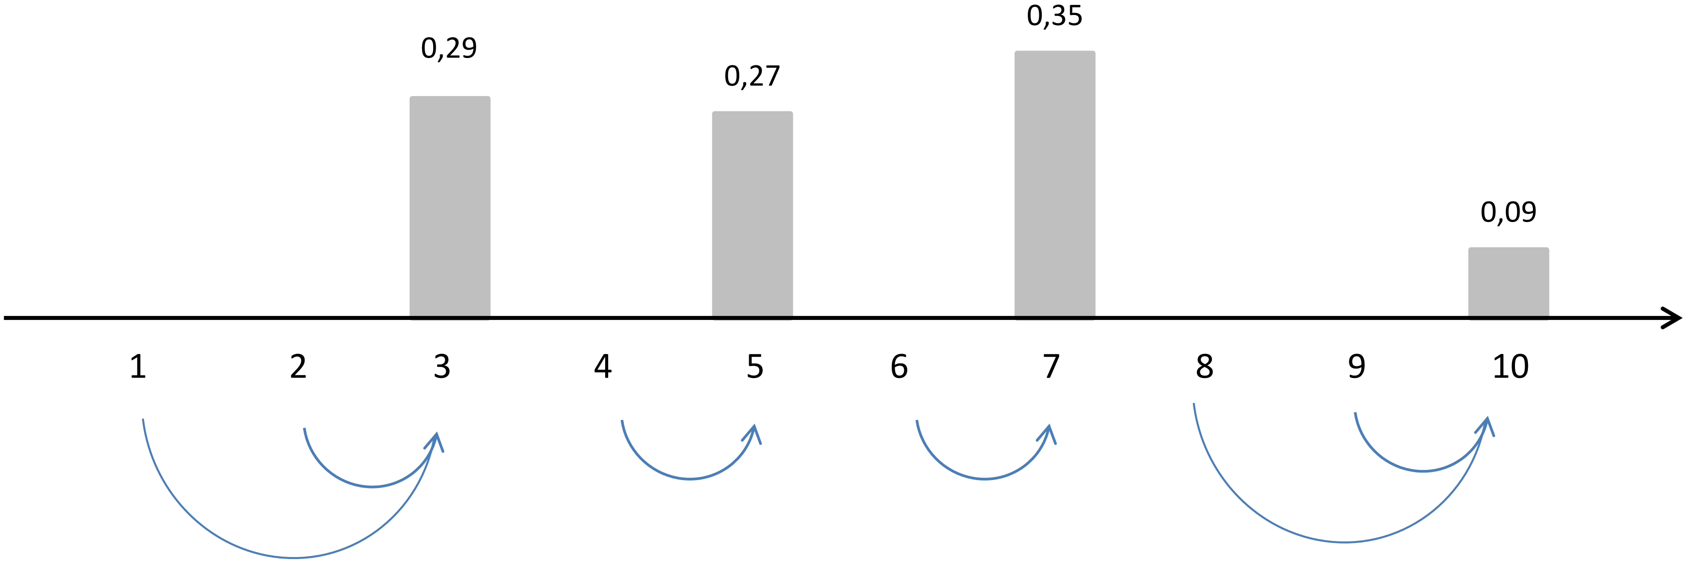

Pessimism

When re-sampling is performed, pessimism is introduced (due to the loss of information).

Response times obtained by using pessimistic random variables are guaranteed to be greater (i.e., worse) than the exact response times of the system. [Diaz 2004]

Quantifying the Pessimism

Metric for the relative pessimism:

$$W_i = \sum_{j=0}^{k_i} p_j * v_j$$

a.k.a. the *expectance* of the distribution.

Re-Sampling Techniques

Uniform Spacing Re-sampling

Domain Quantization

Reduced-Pessimism Re-sampling

Uniform Spacing Re-sampling

The re-sampling is done by choosing equally distanced values out of the initial distribution.

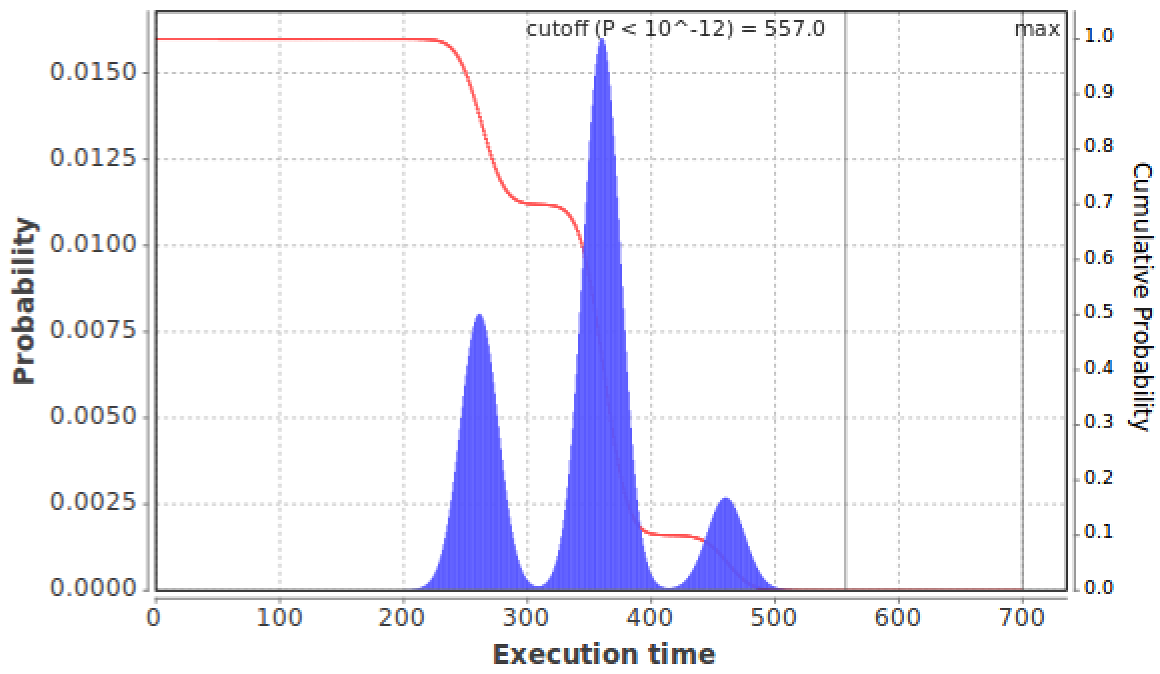

Uniform Spacing Re-sampling

The re-sampling is done by choosing equally distanced values out of the initial distribution.

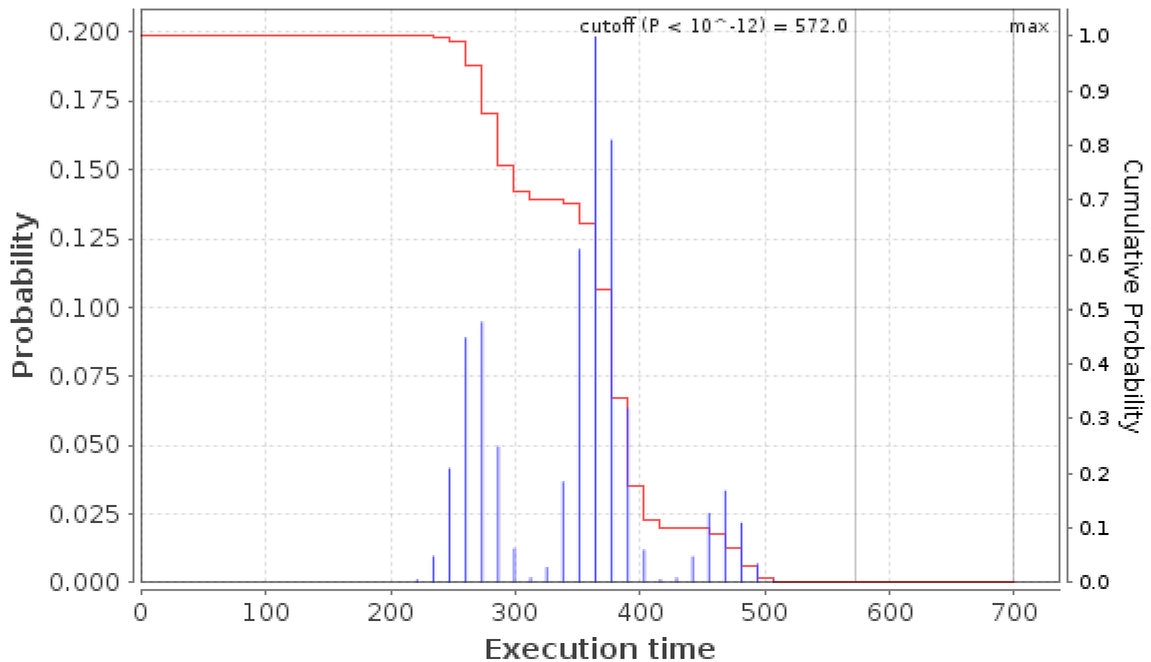

Domain Quantization

In the worst case, the distribution resulting from a convolution can have $m*n$ values.

In the best case, the distribution resulting from a convolution can have $m+n-1$ values.

For large distributions, having only $m+n-1$ values instead of $m*n$ values can make a big difference.

Domain Quantization

A way of decreasing the number of values in the resulting distribution is to quantize the values of the input distributions using a re-sampling strategy so that they have the same spacing between them.

$(1, 4, 7) \otimes (2, 6, 19) = (3, 6, 7, 9, 10, 13, 16, 19, 22)$

9 values

$(1, 4, 7) \otimes (9, 12, 15) = (10, 13, 16, 19, 22)$

only 5 values

Domain Quantization

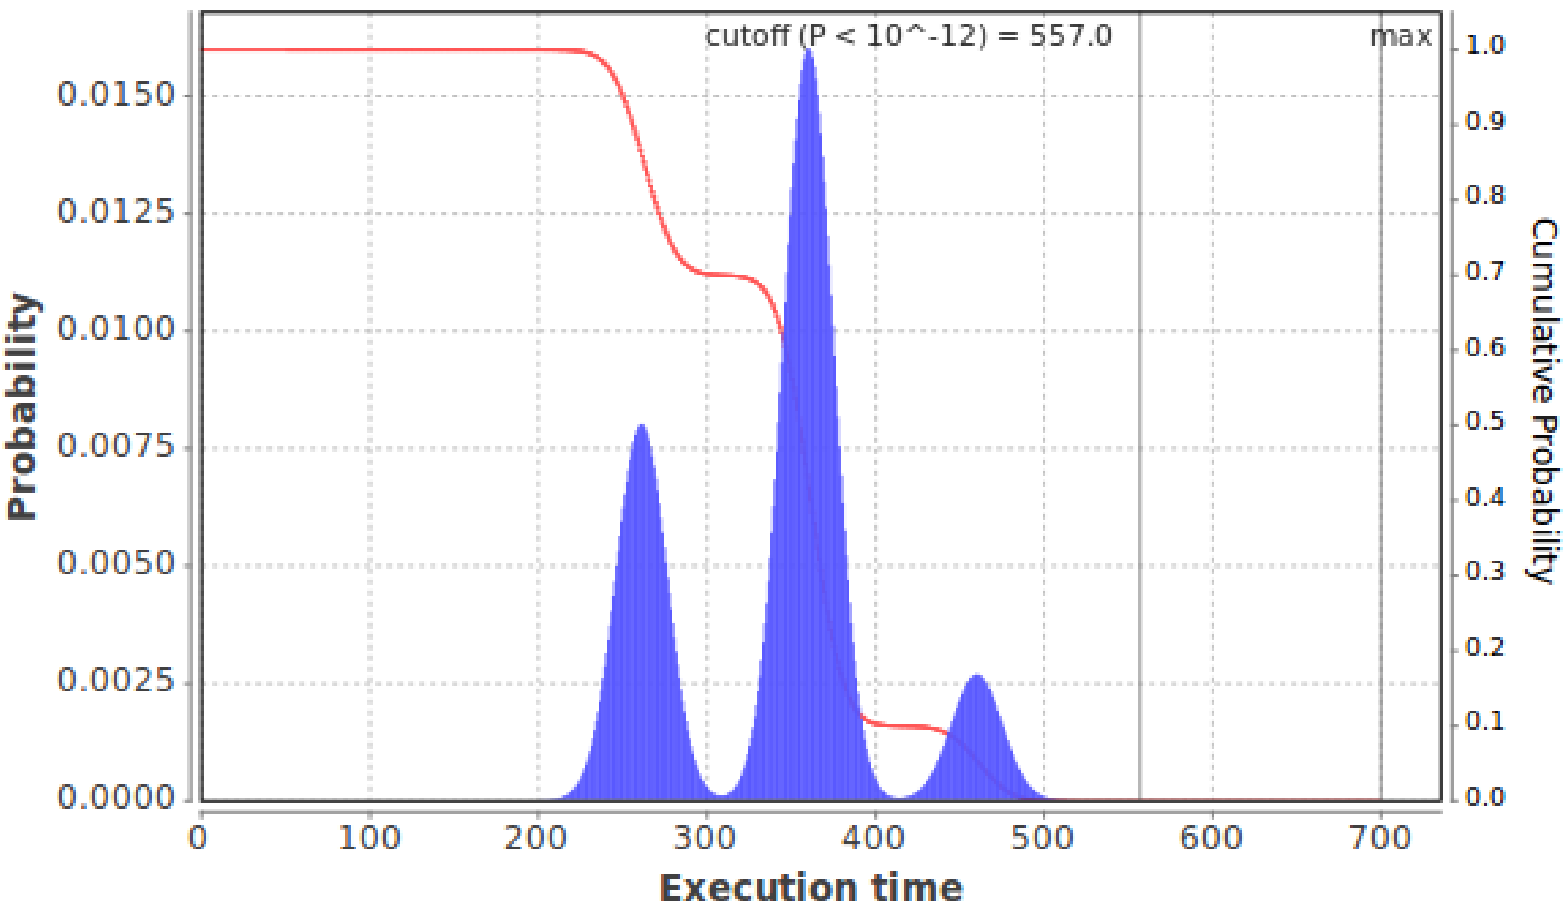

Domain Quantization

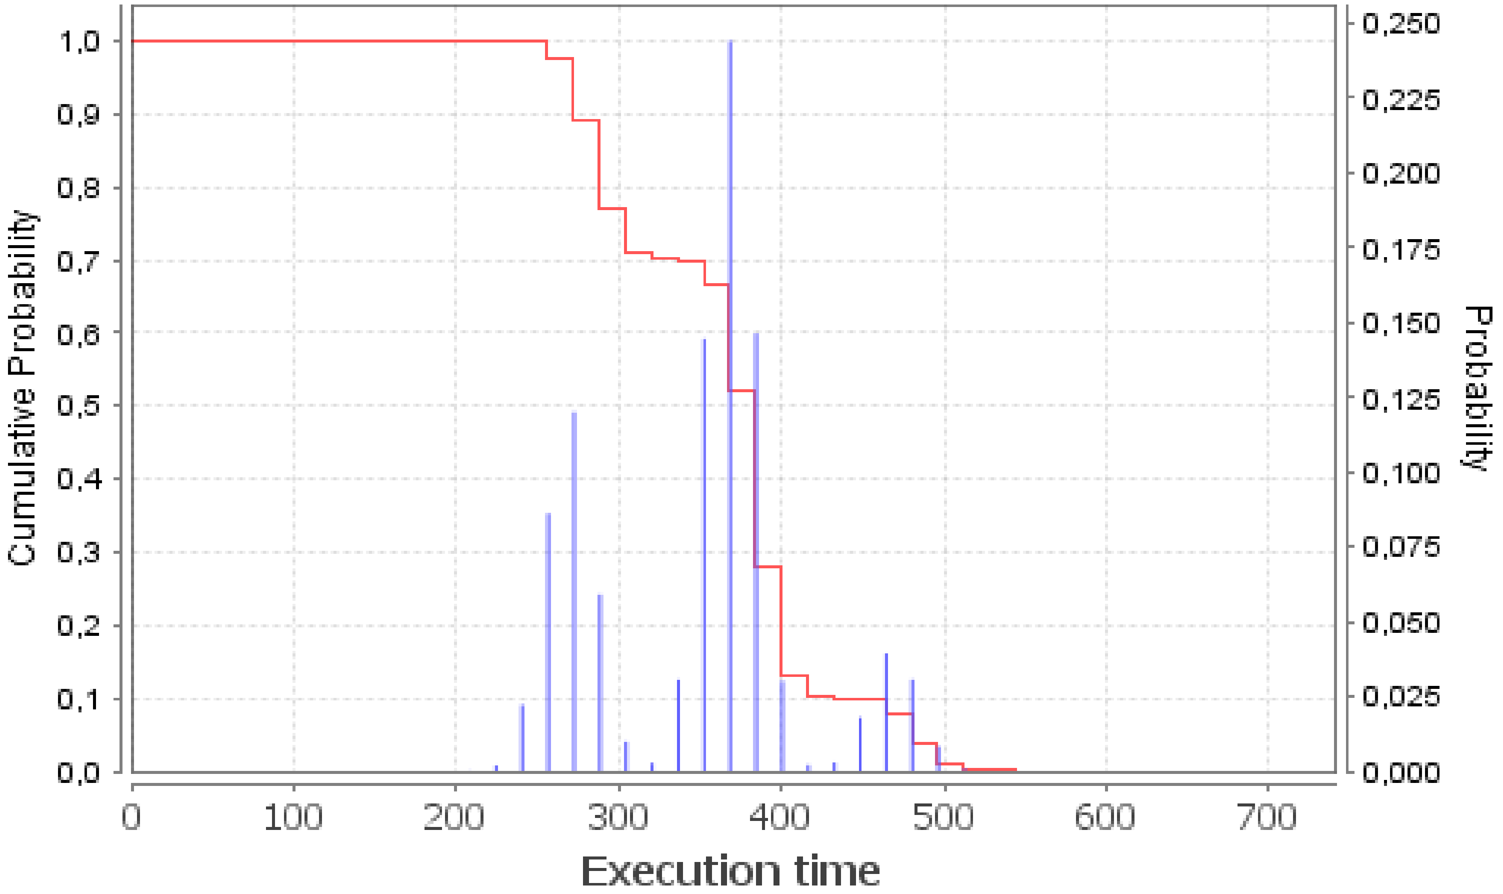

Reduced-Pessimism Re-sampling

A uniform selection of values for re-sampling will, most of the time, not result in a satisfactory reallocation of probability mass.

Reduced-Pessimism Re-sampling works by considering ranges of values and calculating the pessimism that would be introduced if the range of values were to be aggregated into a single entry with the highest value in the range taking all of the probability mass.

Reduced-Pessimism Re-sampling

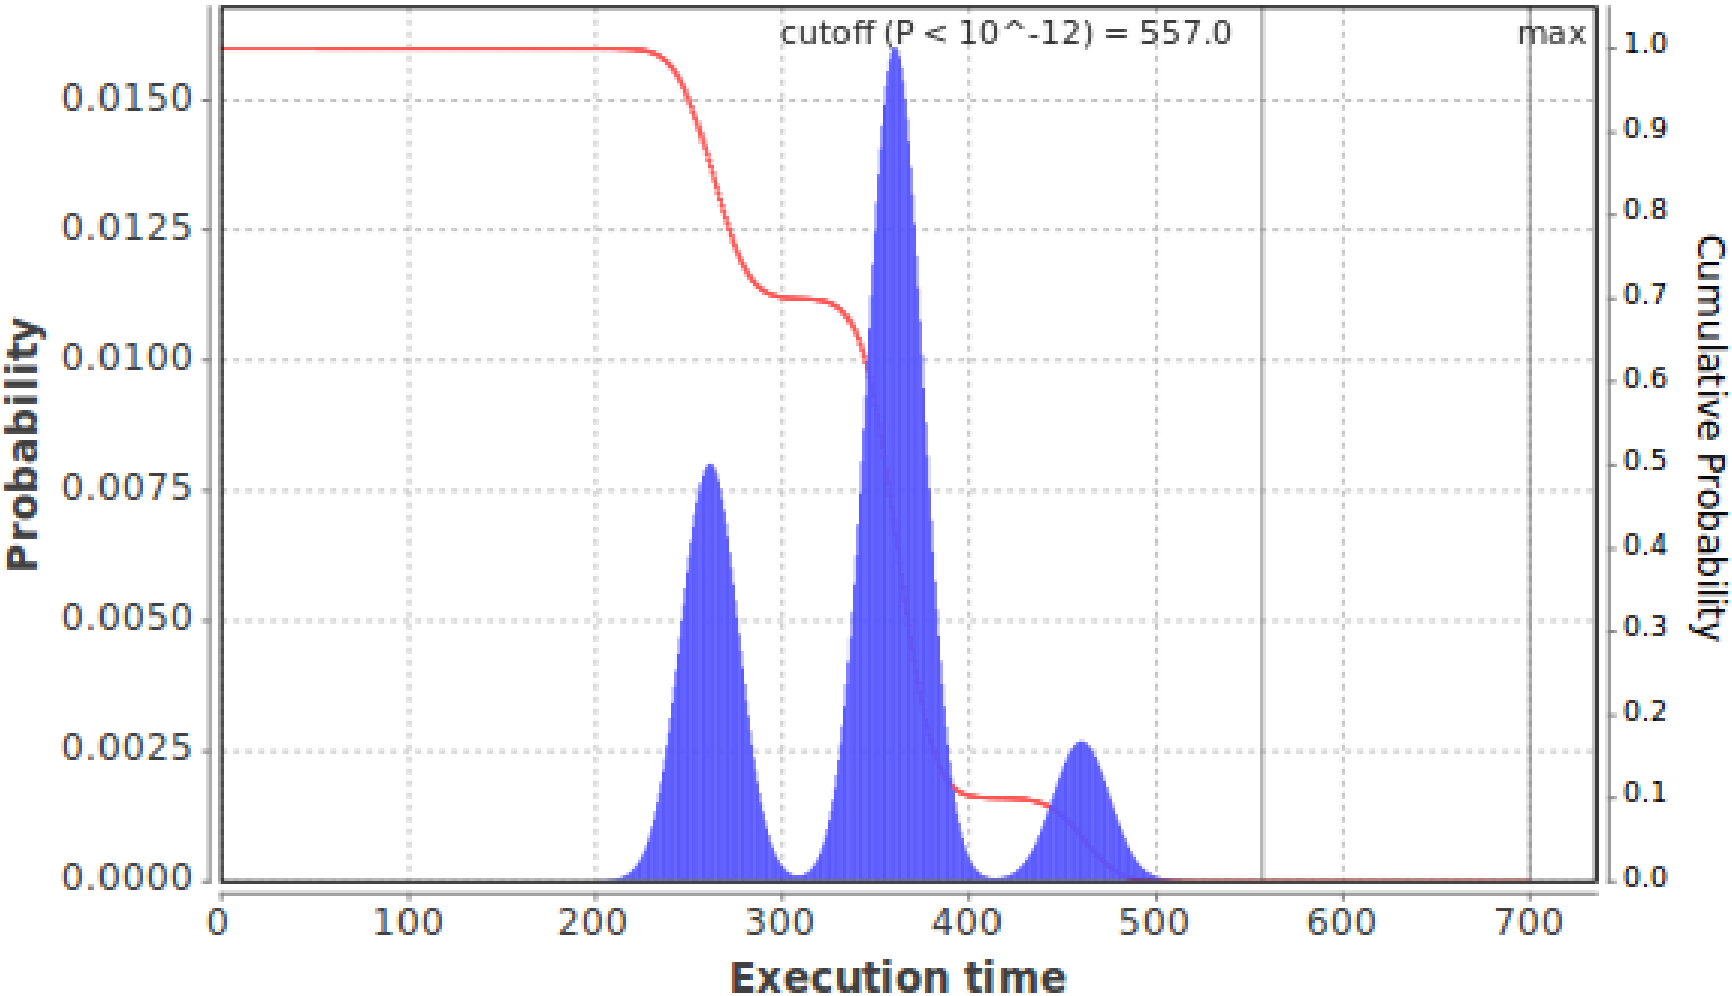

Reduced-Pessimism Re-sampling

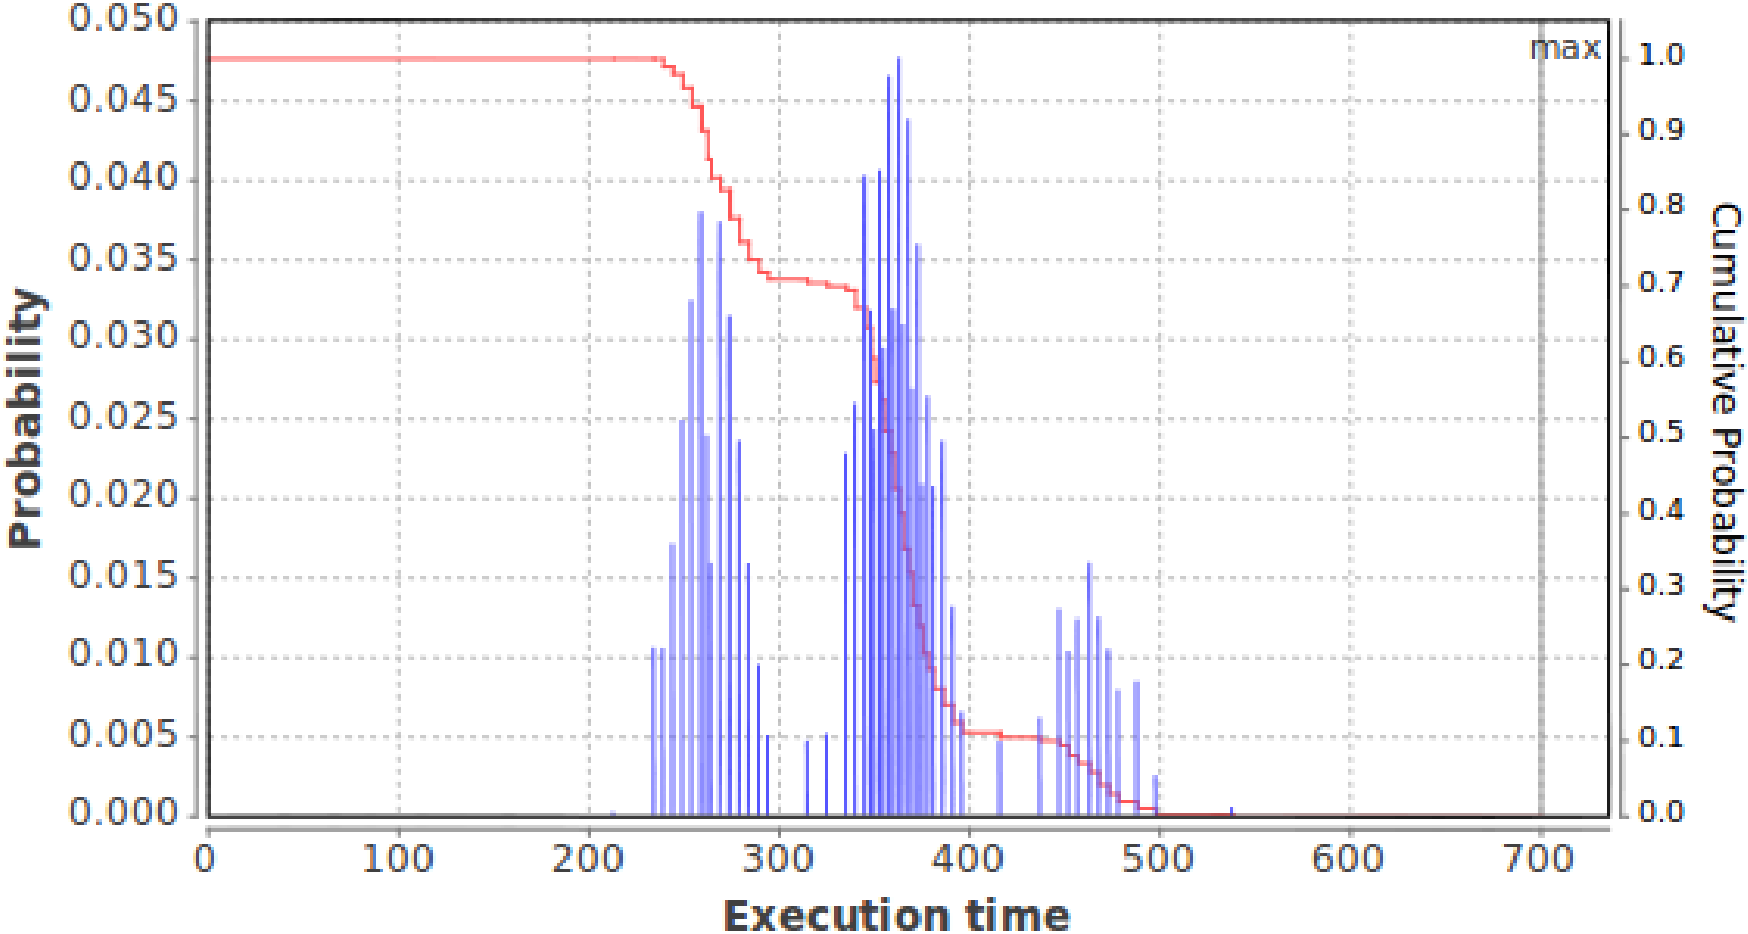

Experiments

Randomly generated sets of tasks with a PF characterizing the worst-case execution time of each task

Long periods in relation to the response time

=> no preemptions

Experiment 1

Computational cost: re-sampled vs. exact

- 10 task sets
- 25 tasks per task set
- 100 values per random variable in a range of approximately 10,000 possible values

Resulting distributions were re-sampled to 1000 values after each convolution

Computational cost

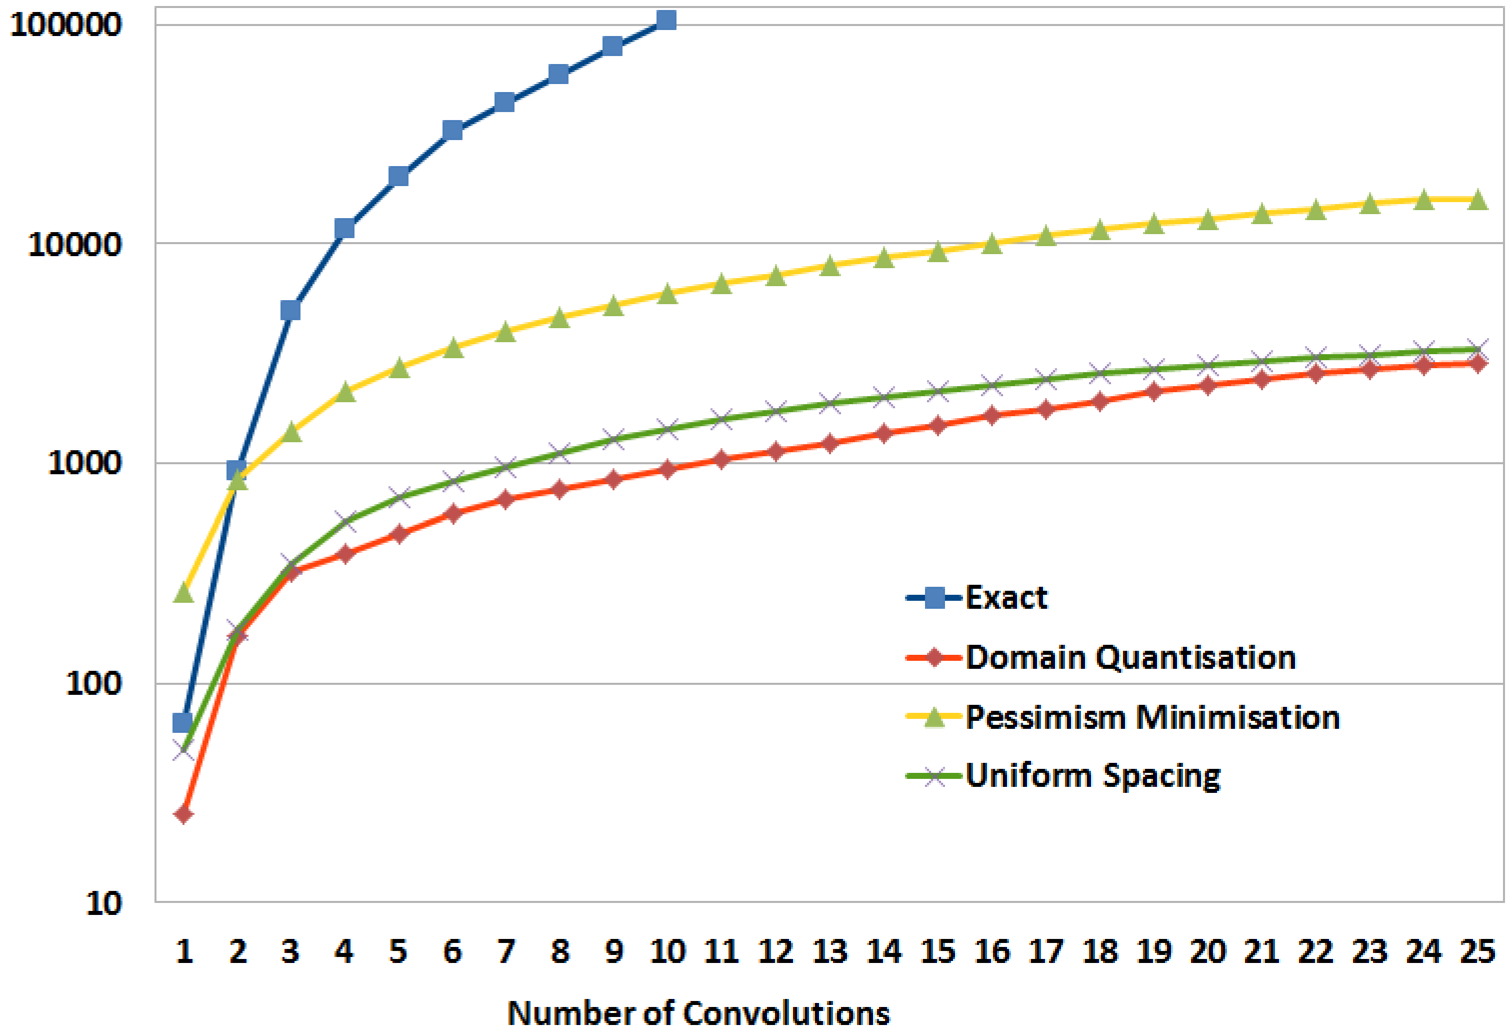

Experiment 2

Comparison of the exceedence function (1-CDF)

- Randomly generated task-set
- 25 tasks
- 100 values per random variable
- Re-sampling threshold = 100

25 convolutions are performed

=> the 25^{th} task is under analysis

Pessimism at response time level

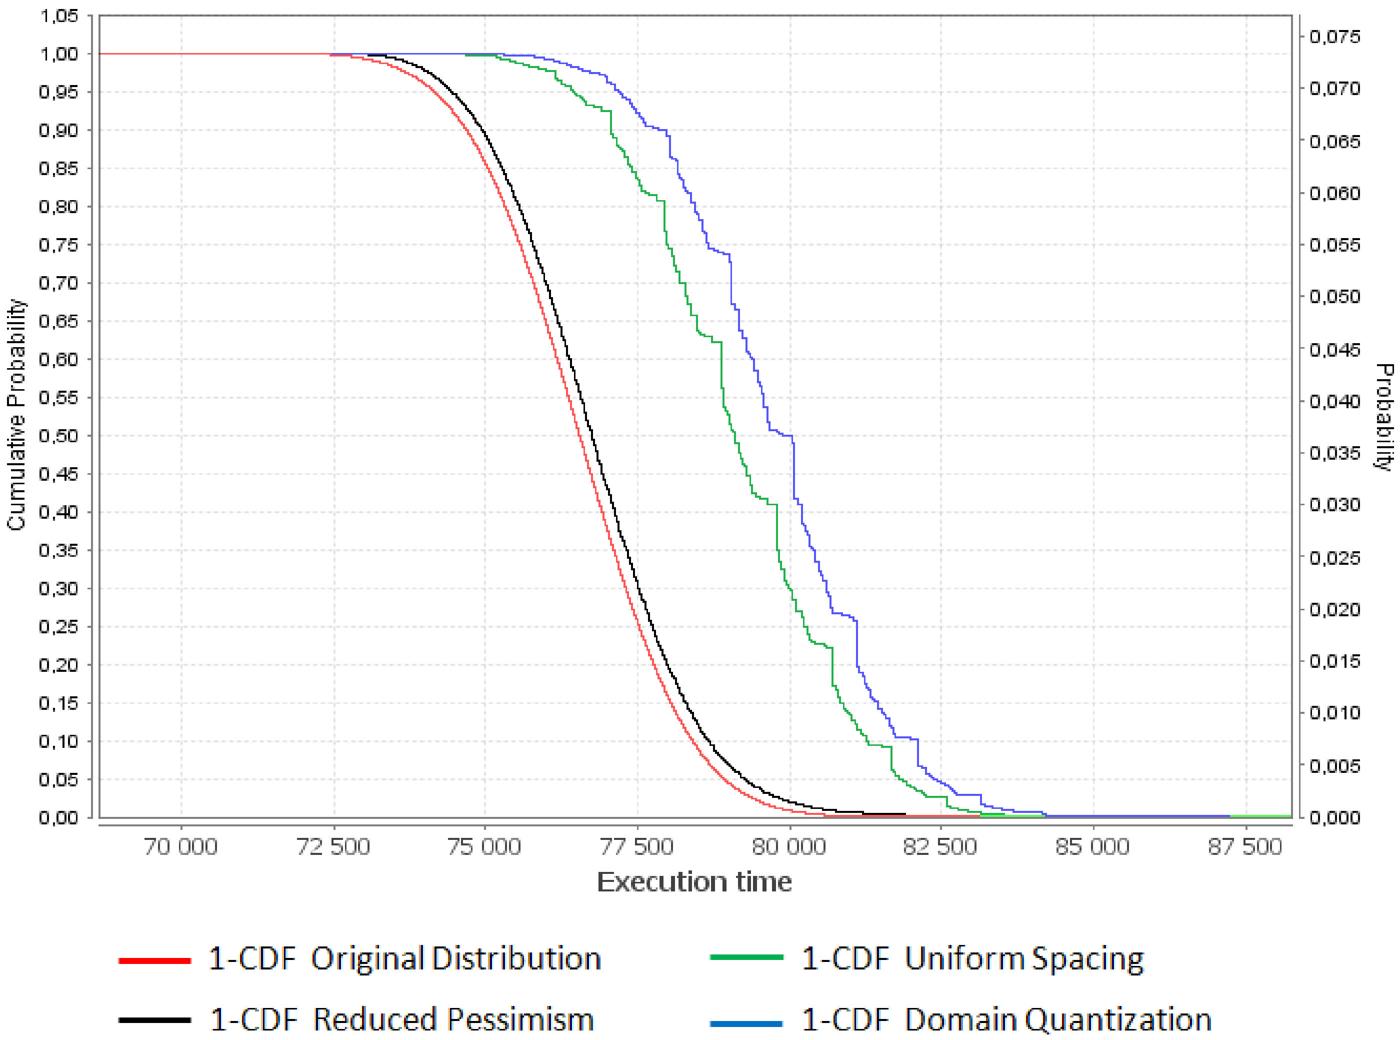

Pessimism at response time level

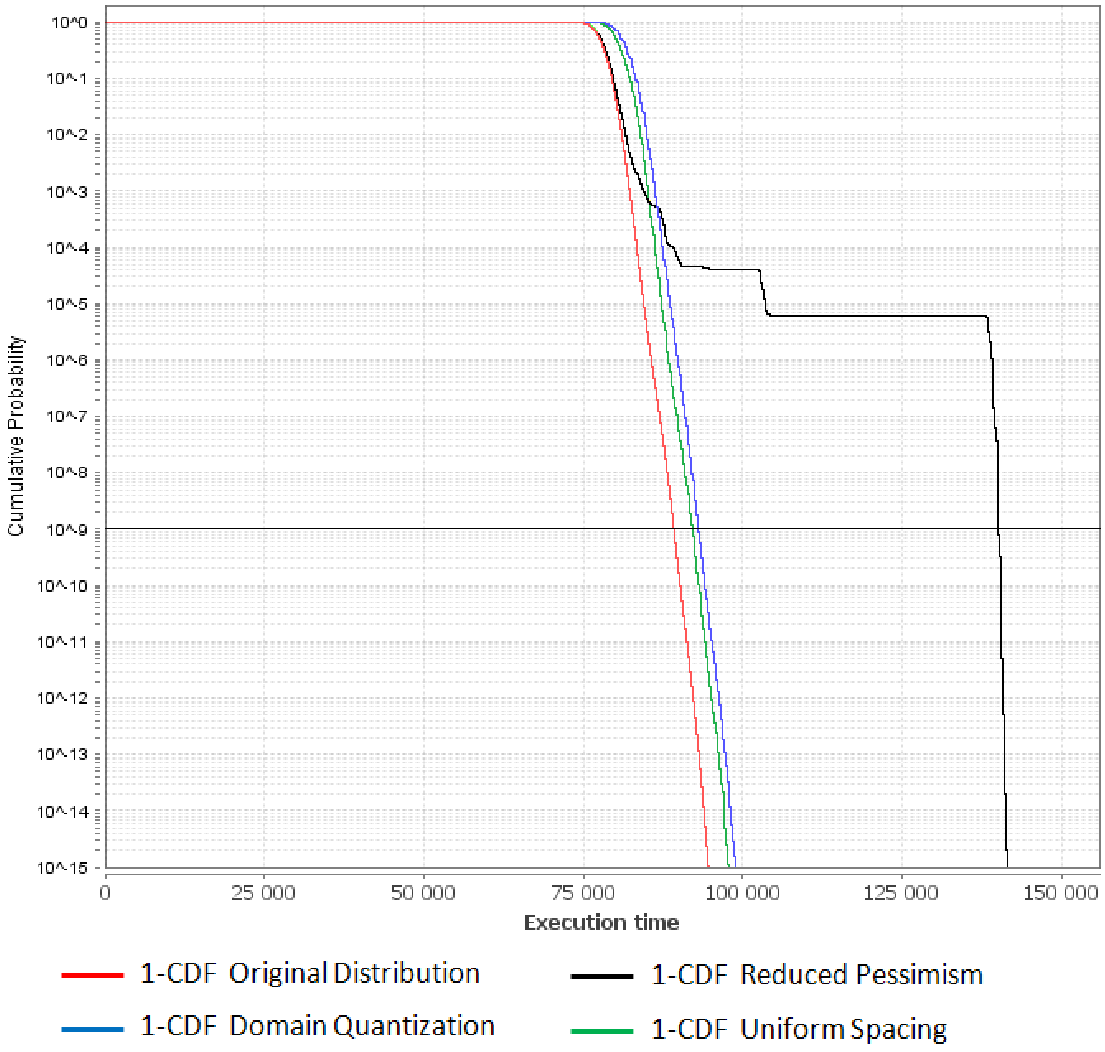

Experiment 3

Trade-off between performance and pessimism

- 10 task sets
- 25 tasks per task set
- 100 values per random variable

Varying the re-sampling threshold

Performance vs. pessimism

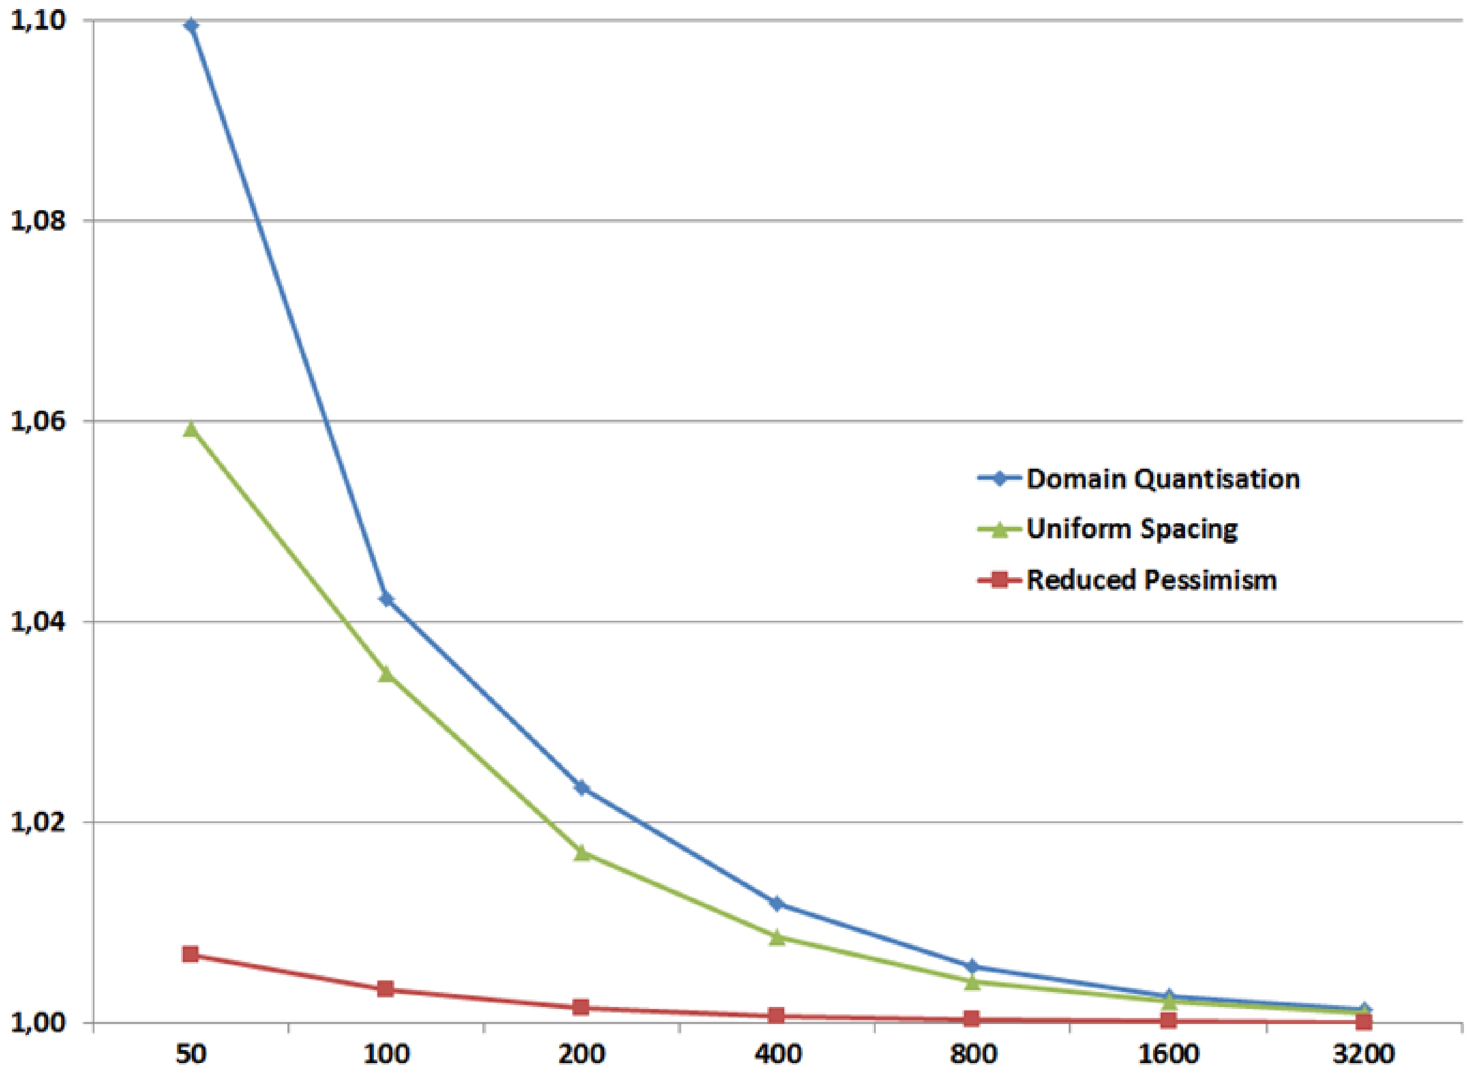

Conclusions

Re-sampling is a powerful way of reducing the complexity of the probabilistic analysis while introducing very little inaccuracy.

Thank You