Probabilistic real-time: restart from scratch

3rd Italian Workshop on Embedded Systems

Federico Reghenzani <federico.reghenzani@polimi.it>
William Fornaciari <william.fornaciari@polimi.it>

HEAP Lab group

HEAP Lab is a cross-disciplinary research team of ~20 people in Politecnico di Milano (Italy)

- Check our webpage:
  - http://www.heaplab.deib.polimi.it/

Federico Reghenzani

2nd year PhD student

<federico.reghenzani@polimi.it>

Probabilistic real-time computing

- In recent years, architectures **increased in complexity** to respond to the high computational power needs (e.g. multi-core, multi-level caches, complex pipelines, speculation, etc.)
- The **Worst-Case Execution Time (WCET)** is hard to be computed in modern architectures
  - ➔ Interferences between tasks (both intra- and inter-cores)
  - ➔ Unexpected latencies (e.g. SMI, memory controller)
  - ➔ Timing anomalies
- Traditional WCET analyses require too computational power or produce too pessimistic results
  → a proposed solution is **probabilistic real-time**

Probabilistic approaches

It is possible to distinguish two main approaches in probabilistic real-time:

- Static Probabilistic Timing Analyses (SPTA)
  - An extension of traditional WCET analyses with probabilities
  - Currently it has not arisen much interest
- Measurement-Based Probabilistic Timing Analysis (MBPTA)

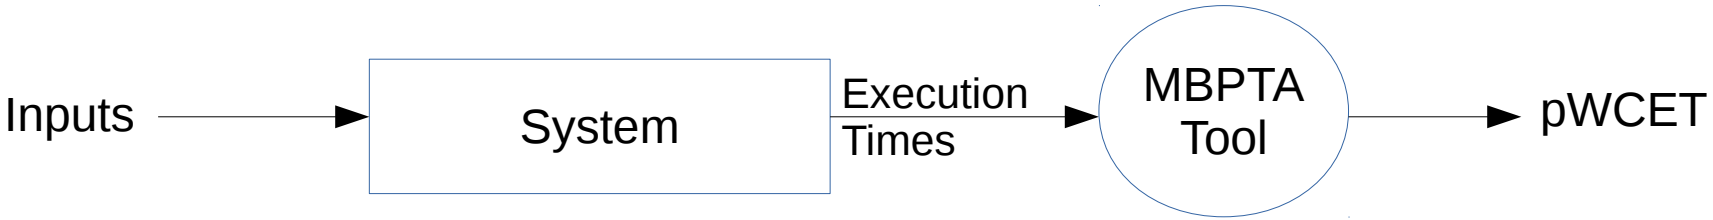

The probabilistic-WCET

The probabilistic-WCET (pWCET) is a statistical distribution with the following cumulative distribution function:

$$p = P(X > \overline{WCET})$$

- Given a violation probability, we can compute the WCET
- Given a WCET, we can compute the violation probability

Provided that the result is **safe** it is possible to use it for hard real-time system: this probability will be just another term in the failure analysis of critical systems.

Extreme Value Theory (1/2)

The Extreme Value Theory (EVT) is a statistical theory dealing with the probability of *extreme events*

- It is traditionally used to predict natural disasters

The basic idea is that the distribution tail of a set of measurements converges to a well-known form, the **Generalized Extreme Value** distribution *:

$$G(x) = \begin{cases} e^{-e^{\frac{x-\mu}{\sigma}}} & \xi = 0 \\ e^{-[1+\xi\frac{x-\mu}{\sigma}]^{-1/\xi}} & \xi \neq 0 \end{cases}$$

Extreme Value Theory (2/2)

To estimate this distribution a couple of approaches exists:

- **Block-Maxima (BM)**:
  - Group the time measurements in block of fixed size and takes the maximum value

- **Peak-over-Threshold (PoT)**:
  - Define a threshold and discard all the time measurements below this value

Conditions for EVT results safety

However, to obtain valid results, the input measurements must satisfy a couple of hypotheses:

- The time measurements must be i.i.d.
  - Strongly depends on the processor and system architecture.
- The distribution of maximas must be in the *domain of attraction* of exteme distributions
  - It may depend on several factors, including how the distribution is estimated.

What we DON’T do

- We are not working on architecture for probabilistic real-time support
  - e.g. randomized caches
- Currently, we are not working on a specific use-case or application
  - In the future we hope to extend the theoretical results to real world
  - Proposal of integration with non-critical software

What we do

Current ongoing research:

- Analysis of statistical tests for EVT hypotheses checking
- Probabilistic scheduling in Mixed-Criticality scenarios
- Using probabilistic results in soft real-time context
- Using probabilistic approach in HPC world
- Developing the tool *chronovise* for MBPTA analyses

Verifying hypotheses

There are several statistical tests to verify the hypotheses, however, there is a lack of systematic approaches.

- We are working on how to apply statistical tests for probabilistic real-time scenario
- Verifying both i.i.d. and *domain of attraction* hypotheses is extremely important to ensure the result soundness

Another important condition, as for traditional approaches, is the input representativity.

The statistical power problem

Performing the tests and obtaining positive results are sufficient to ensure the safety of pWCET?

NO

Statistical tests suffer of false positive and false negative results. In particular, we are studying the impact of the test statistical power in the reliability of the pWCET result.

→ A necessary step for certifiability

Probabilistic mixed-criticality

Questions we have in plan to answer:

- How we can schedule mixed-criticality workloads when we have a probabilistic bound on the WCET?
- Can the pWCET computation for low-criticality tasks be used to improve the mixed-criticality scheduling performance?
  - In other words, can we probabilistically increase the utilization factor?

MBPTA if hypotheses violated

It has been shown that the result of EVT is not “too much dependent” on i.i.d. hypothesis

- If i.i.d. violated, we cannot provide any safe guarantee
  - No certifiability
- But it can be used for relaxed constrained of soft real-time task

Beyond Embedded Systems

Probabilistic real-time aapproaches may be used in different contexts not necessarily embedded systems, e.g. HPC.

- Probabilistic real-time may enable real-time HPC with some sort of guarantees more interesting than average case
- We can deal with performance/predictability tradeoff of HPC software using heterogeneous platforms
- It may be possible to help resource allocation algorithms

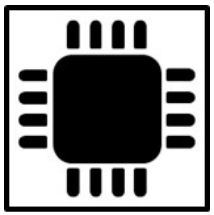

embedded

chronovise

We are developing a software called *chronovise* to perform MBPTA analysis and the full EVT stack

- Open Source and available at https://github.com/federeghe/chronovise (doi:10.21105/joss.00711)
- The *chronovise* aims at standardizing the EVT flow, including the generation of representative inputs, the verification of i.i.d. hypothesis, the estimation of final pWCET
- It's written in C++ and can be used for online analyses
  - No need to know a priori the sample size

Conclusions

- Real-Time is not ready for modern architectures
  - Lack of well-established WCET analyses for multi-core
  - Too complex analyses for complex mechanisms
- Probabilistic real-time may be a solution
  - The result correctness and reliability depends on statistical hypotheses
  - Not-easy to prove the hypotheses validity
- We are trying to fill the (big) theoretical gaps still affecting the probabilistic theory and we propose the use of probabilistic approach in novel scenarios, like HPC or MCS.

Questions?

THANKS FOR YOUR ATTENTION