The Journal of Open Source Software

chronovise: Measurement-Based Probabilistic Timing Analysis framework

Federico Reghenzani[1], **Giuseppe Massari**[1], and **William Fornaciari**[1]

1 Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano

DOI: 10.21105/joss.00711

Software

- Review
- Repository
- Archive

Submitted: 13 April 2018
Published: 28 August 2018

Summary

The rapid advance of computer architectures towards more powerful, but also more complex platforms, has the side effect of making the timing analysis of applications a challenging task (Cullmann et al., 2010). The increasing demand of computational power in cyber-physical systems (CPS) is getting hard to fulfill, if we consider typical real-time constrained applications. Time constraints in CPS are often mandatory requirements, i.e. they must be satisfied in any condition because of the mission-critical system purpose. The satisfaction of these constraints is traditionally demonstrated using well-established *static analyses*, providing the **Worst-Case Execution Time (WCET)** (Wilhelm et al., 2008). However, the increasing complexity of computing architectures – such as multi-core, multi-level caches, complex pipelines, etc. (Berg, Engblom, & Wilhelm, 2004) – makes these analyses computationally unaffordable or carrying out too pessimistic approximations. The problem grows when dealing with Commercial-Off-The-Shelf (COTS) hardware (Dasari, Akesson, Nélis, Awan, & Petters, 2013) and complex operating systems (Reghenzani, Massari, & Fornaciari, 2017).

Probabilistic approaches for hard real-time systems have been proposed as a possible solution to address this complexity increase (Bernat, Colin, & Petters, 2002). In particular, the **Measurement-Based Probabilistic Time Analysis (MBPTA)** (Cucu-Grosjean et al., 2012) is a probabilistic analysis branch for real-time systems to estimate the WCET directly from the observed execution times of real-time tasks. The time samples are collected across the application input domain and the WCET is provided in probablistic terms, the *probabilistic-WCET (pWCET)*, i.e. a WCET with a probability of observing higher execution times. The statistical theory at the basis of the WCET estimation is the **Extreme Value Theory (EVT)** (E. Castillo, Hadi, Balakrishnan, & Sarabia, 2005) (De Haan & Ferreira, 2007), typically used in natural disaster risk evaluation, However, to obtain a safe pWCET estimation, the execution time traces must fulfill certain requirements. In particular, MBPTA requires the time measurements to be (Kosmidis, 2017): (1) independent and identically distributed, (2) representative of all worst-case latencies. The first requirement comes from the EVT, it can be checked with suitable statistical tests and can be relaxed under some circumstances (Santinelli, Guet, & Morio, 2017), while the latter is relative to the input representativity and to the system (hardware/software) properties. Both requirements are necessary to obtain a safe, i.e. non-underestimated, pWCET.

The *chronovise* framework is an open-source software aiming at standardizing the flow of MBPTA process, integrating both estimation and testing phases. The few existing software presented in literature (Lu, Nolte, Kraft, & Norstrom, 2010) (Lesage, Griffin, Soboczenski, Bate, & Davis, 2015) lack of source code availability. Moreover, both works include a limited set of features, other than poor maturity level due to the missing integration of the most recent scientific contributions. Another software is available as open-source (Abella, 2017), but specialized for a variant of classical MBPTA analysis

called MBPTA-CV (Abella, Padilla, Castillo, & Cazorla, 2017). Our work aims at filling the absence of a stable software with a well-defined EVT execution flow. The proposed framework supports both Block-Maxima (BM), Peak-over-Threshold (PoT) and MBPTA-CV EVT approaches; the current available methods to estimate the extreme distribution. The output distribution respectively assumes the Generalized Extreme Value (GEV) and the Generalized Pareto Distribution (GPD) form. Three estimators, Maximum Likelihood Estimator (MLE) (Bücher & Segers, 2017), Generalized-MLE (GMLE) (Martins & Stedinger, 2000), Probability Weighted Moment (PWM, called also L-moments) (Hosking & Wallis, 1987), are already included, as well as some statistical tests: Kolmogorov-Smirnov (Massey, 1951) and (Modified) Anderson-Darling (Sinclair, Spurr, & Ahmad, 1990). Finally, the implementation of an overall results confidence estimation procedure is also available. The API provided allows users to specify or to implement new input generators and input representativity tests.

The software *chronovise* is in fact presented as a flexible and extensible framework, deployed as a static C++ library. The selection of C++ language enables the easy implementation of hardware-in-the-loop analyses. The underlying idea of *chronovise* is to provide a common framework for both researchers and users. Even if EVT is a well-known statistical theory, it is continuously evolving and it is still a hot topic in mathematical environment. The application of EVT in real-time computing is immature and it still requires several theoretical advances. This has led us to implement this software: enabling the exploitation of an already implemented EVT process, in order to perform experiments of new theories and methods, without the need to reimplement algorithms from scratch. With our framework we want to create a common *software-base*, that would increase both the replicability of the experiments and the reliability of the results, which are common issues in research. On the other hand, end-users – i.e. engineers that use the already available algorithms to estimate the pWCET – can just implement the measurement part and use the framework without introducing further changes.

References

Abella, J. (2017, November). MBPTA-CV. doi:10.5281/zenodo.1065776

Abella, J., Padilla, M., Castillo, J. D., & Cazorla, F. J. (2017). Measurement-based worst-case execution time estimation using the coefficient of variation. *ACM Transactions on Design Automation of Electronic Systems*, *22*(4), 72:1–72:29. doi:10.1145/3065924

Berg, C., Engblom, J., & Wilhelm, R. (2004). Requirements for and design of a processor with predictable timing. In L. Thiele & R. Wilhelm (Eds.), *Perspectives workshop: Design of systems with predictable behaviour*, Dagstuhl seminar proceedings. Dagstuhl, Germany: Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany. Retrieved from http://drops.dagstuhl.de/opus/volltexte/2004/5

Bernat, G., Colin, A., & Petters, S. M. (2002). WCET analysis of probabilistic hard real-time systems. In *23rd IEEE real-time systems symposium, 2002. RTSS 2002.* (pp. 279–288). doi:10.1109/REAL.2002.1181582

Bücher, A., & Segers, J. (2017). On the maximum likelihood estimator for the generalized extreme-value distribution. *Extremes*, *20*(4), 839–872. doi:10.1007/s10687-017-0292-6

Castillo, E., Hadi, A. S., Balakrishnan, N., & Sarabia, J.-M. (2005). Extreme value and related models with applications in engineering and science.

Cucu-Grosjean, L., Santinelli, L., Houston, M., Lo, C., Vardanega, T., Kosmidis, L., Abella, J., et al. (2012). Measurement-based probabilistic timing analysis for multi-path programs. In *2012 24th euromicro conference on real-time systems* (pp. 91–101).

doi:10.1109/ECRTS.2012.31

Cullmann, C., Ferdinand, C., Gebhard, G., Grund, D., Maiza, C., Reineke, J., Triquet, B., et al. (2010). Predictability considerations in the design of multi-core embedded systems. *Proceedings of Embedded Real Time Software and Systems*, 36–42.

Dasari, D., Akesson, B., Nélis, V., Awan, M. A., & Petters, S. M. (2013). Identifying the sources of unpredictability in cots-based multicore systems. In *2013 8th IEEE international symposium on industrial embedded systems (SIES)* (pp. 39–48). doi:10.1109/SIES.2013.6601469

De Haan, L., & Ferreira, A. (2007). *Extreme value theory: An introduction*. Springer Science & Business Media.

Hosking, J. R. M., & Wallis, J. R. (1987). Parameter and quantile estimation for the generalized pareto distribution. *Technometrics*, *29*(3), 339–349. doi:10.2307/1269343

Kosmidis, L. (2017). *Enabling caches in probabilistic timing analysis* (PhD thesis). Universitat Politècnica de Catalunya. Retrieved from https://www.tdx.cat/handle/10803/460819?locale-attribute=en

Lesage, B., Griffin, D., Soboczenski, F., Bate, I., & Davis, R. I. (2015). A framework for the evaluation of measurement-based timing analyses. In *Proceedings of the 23rd international conference on real time and networks systems*, RTNS '15 (pp. 35–44). New York, NY, USA: ACM. doi:10.1145/2834848.2834858

Lu, Y., Nolte, T., Kraft, J., & Norstrom, C. (2010). A statistical approach to response-time analysis of complex embedded real-time systems. In *2010 IEEE 16th international conference on embedded and real-time computing systems and applications* (pp. 153–160). doi:10.1109/rtcsa.2010.13

Martins, E. S., & Stedinger, J. R. (2000). Generalized maximum-likelihood generalized extreme-value quantile estimators for hydrologic data. *Water Resources Research*, *36*(3), 737–744. doi:10.1029/1999WR900330

Massey, F. J. (1951). The kolmogorov-smirnov test for goodness of fit. *Journal of the American Statistical Association*, *46*(253), 68–78. Retrieved from http://www.jstor.org/stable/2280095

Reghenzani, F., Massari, G., & Fornaciari, W. (2017). Mixed time-criticality process interferences characterization on a multicore linux system. In *2017 euromicro conference on digital system design (DSD)* (pp. 427–434). doi:10.1109/DSD.2017.18

Santinelli, L., Guet, F., & Morio, J. (2017). Revising measurement-based probabilistic timing analysis. In *2017 IEEE real-time and embedded technology and applications symposium (RTAS)* (pp. 199–208). doi:10.1109/RTAS.2017.16

Sinclair, C., Spurr, B., & Ahmad, M. (1990). Modified anderson darling test. *Communications in Statistics - Theory and Methods*, *19*(10), 3677–3686. doi:10.1080/03610929008830405

Wilhelm, R., Engblom, J., Ermedahl, A., Holsti, N., Thesing, S., Whalley, D., Bernat, G., et al. (2008). The worst-case execution-time problem. Overview of methods and survey of tools. *ACM Transaction of Embededded Computer Systems*, *7*(3), 36:1–36:53. doi:10.1145/1347375.1347389