Unikernel-Based Real-Time Virtualization Under Deferrable Servers: Analysis and Realization (Artifact)

Kuan-Hsun Chen
University of Twente, The Netherlands

Mario Günzel
TU Dortmund University, Germany

Boguslaw Jablkowski
EMVICORE GmbH, Dortmund, Germany

Markus Buschhoff
EMVICORE GmbH, Dortmund, Germany

Jian-Jia Chen
TU Dortmund University, Germany

Abstract

This artifact provides the source code to validate and reproduce the numerical results of the associated paper "Unikernel-Based Real-Time Virtualization under Deferrable Servers: Analysis and Realization". Due to the nature of a close-source project with the company, i.e., EMVICORE GmbH, the source code of the case study in Section 6.2 is not included in this artifact.

2012 ACM Subject Classification Computer systems organization → Embedded and cyber-physical systems; Software and its engineering → Real-time systems software

Keywords and phrases Unikernel, Virtualization, Reservation Servers, Deferrable Servers, Cyber-Physical Systems, Real-Time Systems

Digital Object Identifier 10.4230/DARTS.8.1.2

Funding This work has been supported by Deutsche Forschungsgemeinschaft (DFG), as part of Sus-Aware (Project No. 398602212). This result is part of a project (PropRT) that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 865170)

Related Article Kuan-Hsun Chen, Mario Günzel, Boguslaw Jablkowski, Markus Buschhoff, and Jian-Jia Chen, "Unikernel-Based Real-Time Virtualization Under Deferrable Servers: Analysis and Realization", in 34th Euromicro Conference on Real-Time Systems (ECRTS 2022), LIPIcs, Vol. 231, pp. 6:1–6:22, 2022. https://doi.org/10.4230/LIPIcs.ECRTS.2022.6

Related Conference 34th Euromicro Conference on Real-Time Systems (ECRTS 2022), July 5–8, 2022, Modena, Italy

1 Scope

This artifact presents the source code of the response-time analysis used in the numerical simulation reported in Section 6.1. Particularly, two figures, i.e., Figure 6 and 7 in the paper, can be reproduced by executing `auto.sh`. It thereby supports our claim, that the proposed analysis outperforms the state of the art, i.e., the converted Real-Time Calculus-based approach.

Due to the nature of a close-source project with the company, i.e., EMVICORE GmbH, the source code of the case study in Section 6.2 is not included in this artifact.

Dagstuhl Artifacts Series
Schloss Dagstuhl – Leibniz-Zentrum für Informatik,
Dagstuhl Publishing, Germany

2 Content

The artifact package includes the analysis implementation and the setup of numerical simulation. In particular, the `res` directory contains:

- `benchmark.py`: Server and task creation
- `our_analysis.py`: Our analysis
- `plot.py`: Plotting functionality
- `rtc_cb.py`: RTC-based analysis

In addition, a bash-script `auto.sh` is provided to automize the evaluation. The detailed document can be found as a `README.md` with installation and usage instructions.

3 Getting the artifact

The artifact endorsed by the Artifact Evaluation Committee is available free of charge on the Dagstuhl Research Online Publication Server (DROPS). In addition, the artifact is also available at: `https://github.com/tu-dortmund-ls12-rt/unikernel-based_deferrable_server_analysis`.

4 Tested platforms

The artifact was tested on a laptop computer using 64-bit Arch Linux 5.17.3 with i7-10610U CPU and 16GB main memory; it does not assume or require any particular hardware configuration. It took about 170 seconds with this machine to obtain Figure 6 and 7, when set `num_processors = 5` in `main.py`. The artifact should work on any system that supports Python 3.9 and Python 3.10.

5 License

The artifact is available under the MIT License.

6 MD5 sum of the artifact

e3c69a027defea668a42c6abe4019d00

7 Size of the artifact

12.6 KiB