

Reproducibility Report for the Paper: *Modeling of Request Cloning in Cloud Server Systems using Processor Sharing*

Alessandro Pellegrini
pellegrini@diag.uniroma1.it

February 7, 2020

Abstract

The authors have uploaded their artifact on Zenodo, which ensures a long-term retention of the artifact. The code is suitably documented, and some examples are given. A minimalistic overall description of the engine is provided. The artifact allows to setup the environment quite quickly, and the dependencies are well documented. The process to regenerate data for the figures in the paper completes, and all results are reproducible.

This paper can thus receive the *Artifacts Available* badge and the *Artifacts Evaluated—Functional*. Given the high quality of the artifact, also the *Artifacts Evaluated—Reusable* badge can be assigned.

1 Introduction

The paper *Modeling of Request Cloning in Cloud Server Systems using Processor Sharing* by Tommi Nylander, Johan Ruuskanen, Karl-Erik Årzén, and Martina Maggio [1] presents theoretical results related to the performance of client/server applications (in the context of cloud-based deploys) when request cloning is exploited to reduce the time to serve the requests.

The authors present a G/G/1 model which allows to represent a system of servers using cloning to serve the requests as if it were a single server, basing their approach on the idea that the service must be synchronized.

The authors validate their theoretical results by means of a discrete event simulation model, written in python, which is the focus of this report. The authors have also implemented the sequential runtime core for the model, thus making the artifact self-contained, which could be a benefit for some end users, provided they don't have strong performance requirements.

2 Replication of Computational Results

In this section we provide Figures taken from the original paper. The authors have given permission to use their original results in this report.

2.1 Software download and installation

The authors have provided a link to a permanent repository on Zenodo (<http://doi.org/10.5281/zenodo.3611398>), making it permanently available (**Artifacts Available badge**). This reviewer encourages the authors to place this link also in the main paper, to increase the dissemination of their results.

The amount of dependencies required to run the artifacts is extremely reduced. Simulations can be run by relying solely on python. There is a subsequent post-processing phase, required to analyze the results of the simulations to prepare the datasets to replot the figures which requires matlab. This last step, anyhow, is only related to the reproducibility phase, while the simulation model is self-contained in a reduced number of scripts.

The authors have provided a docker environment to run the simulations, which installs all the required python dependencies, which has been gratefully appreciated.

2.2 Quality of the artifact

The artifact is relevant to the associated paper, as it allows to re-generate all the data and all the figures, it is well documented (from the point of view of reproducibility), and it is easy to use thanks to the provided supporting environment. No component to re-run the simulations and to post-process the data is lacking (**Artifacts Evaluated—Functional**).

The authors have also provided the results of their simulation runs in the artifact. They propose to rely on this data to regenerate the plots, in case no sufficient computational power is available. I can confirm that by relying on this data, the post-processing and plotting modules do work, and the figures are 1:1 with respect to those in the paper. In the artifact, the authors also provide the post-processed data, thus making it possible to only replot the figures. Matlab scripts are written so that if these post-processed files are found, they are not overwritten. For this reproducibility assessment, this reviewer has rerun all simulations from scratch, and has cleared out all post-processed results before launching the post-processing activity.

The authors have provided a minimalistic documentation to the artifact, consisting in an html file with a high-level description of the simulation model. The authors have also provided a description of the timeline of a cloned request in their software. This documentation does not really allow to repurpose the artifact. Nevertheless, there are some examples which allow to get an insight on how to repurpose the artifact. The model/kernel documentation is good, making it easier to understand. Overall, this reviewer thinks that the quality of the artifact is good (**Artifacts Evaluated—Reusable**).

2.3 Replicating the experiments

Replication of the experiments has been carried out on multiple environments. Python simulations have been run on a cluster of 4 parallel high-end machines, all running Debian 9, in the provided docker container. Post processing and plotting have been run on an Arch Linux machine, using matlab R2019b rev 3.

The artifact is easy to run, and all data required to support the reproducibility is provided (**Artifacts Evaluated—Functional**). It took a more than one week to run all the simulations using ~ 150 cores. Initially, two of the post-processing matlab scripts were failing. This was preventing to generate the data which are later required by the tikz scripts in the main latex file to generate the plots.

The authors have been informed via email on Feb 02, 2020 of this issue, receiving all the data from simulations which were re-generated by this reviewer. The authors have replied on Feb 04, 2020, by acknowledging the issues and uploading a revised version of the artifact which has fixed all the aforementioned issues.

In this section, I provide a short description of my reproducibility results, and how they were obtained. There is a total of 11 Figures and 2 Tables in the original paper.

Tab. 1 in the paper shows probabilities from a distribution, which does not require to be replicated. Similarly, **Tab. 2** presents data coming from a theoretical analysis which does not require to be replicated.

Fig. 1 in the paper is a drawing and needs not to be reproduced. **Fig. 2** in the paper is a graphical representation of multiple equations, thus needs not to be reproduced.

Fig. 3 in the paper corresponds to Figure 1 in this report. The data for this figure is generated by running the `sim_gg1_3dist.py` wrapper script, which spawns a couple of simulation instances. Data is post processed by the `analyze_gg1_example.m` matlab script, which completes successfully. The results are visually and substantially comparable.

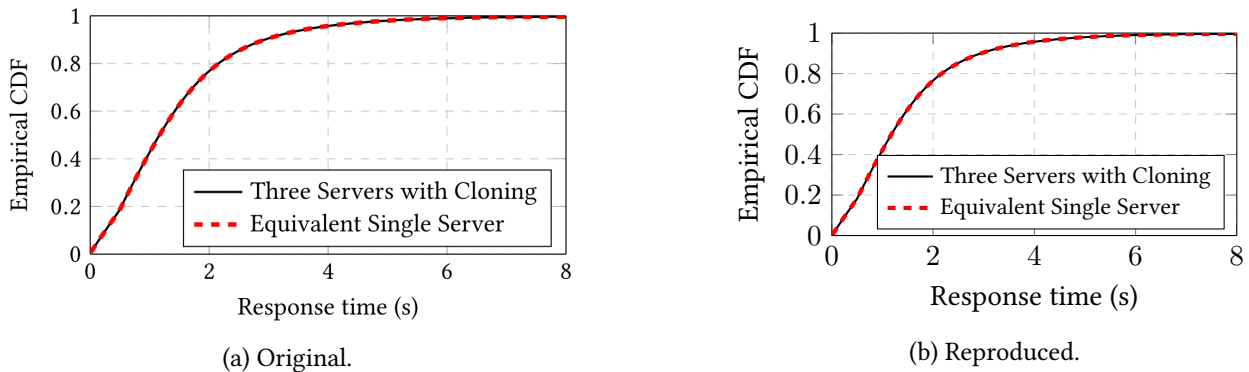


Figure 1: Empirical response time CDFs for the Example with Heterogeneous Servers. Data retrieved through 20 repeated simulations of 10^6 requests each. The 95% confidence intervals lie within the lines.

Fig. 4 in the paper corresponds is a visual representation of a theoretical analysis, which thus does not require reproducibility. **Fig. 5** in the paper is a drawing, which does not require reproducibility.

Fig. 6 in the paper corresponds to Figure 2 in this report. The data for this figure is generated by running the `sim_optimal_clone-ps.py` wrapper script, which spawns 3110 simulation instances. Data is post processed by the `analyze_opt_clone.m` matlab script, which completes successfully. The results are slightly different from the original ones in the paper, but the trends are comparable and do not diverge too much. I emphasize that this difference is also related to the reduced number of simulations which have been run for this particular figure, due to a clerical error in the first version of the artifact.

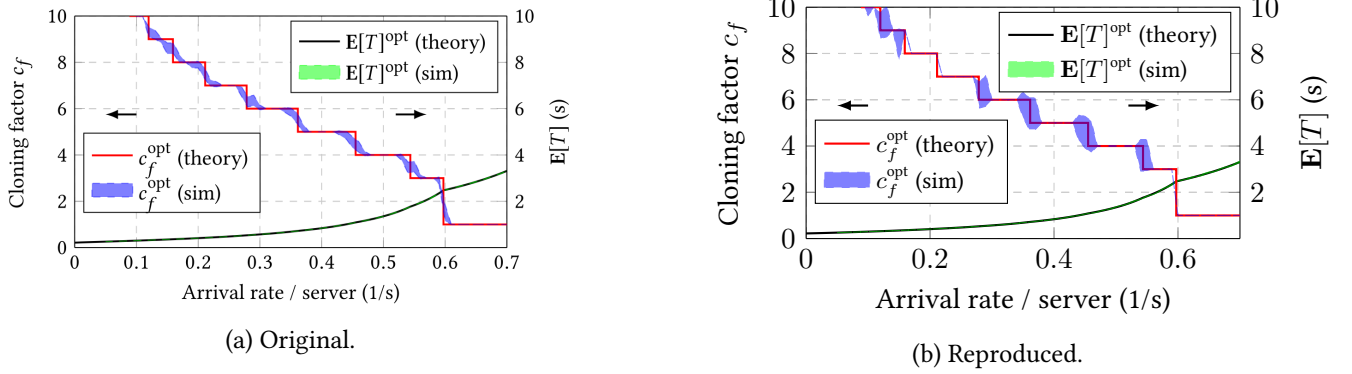


Figure 2: Clone-to-All: Comparison of theoretical values with 95% confidence intervals for the simulation results.

Fig. 7 in the paper corresponds to Figure 3 in this report. The data for this figure is generated by running the `sim_codesigns_icpe.py` wrapper script, which spawns 880 simulation instances. Data is post processed by the `analyze_codesigns.m` matlab script, which completes successfully. The results are visually and substantially comparable.

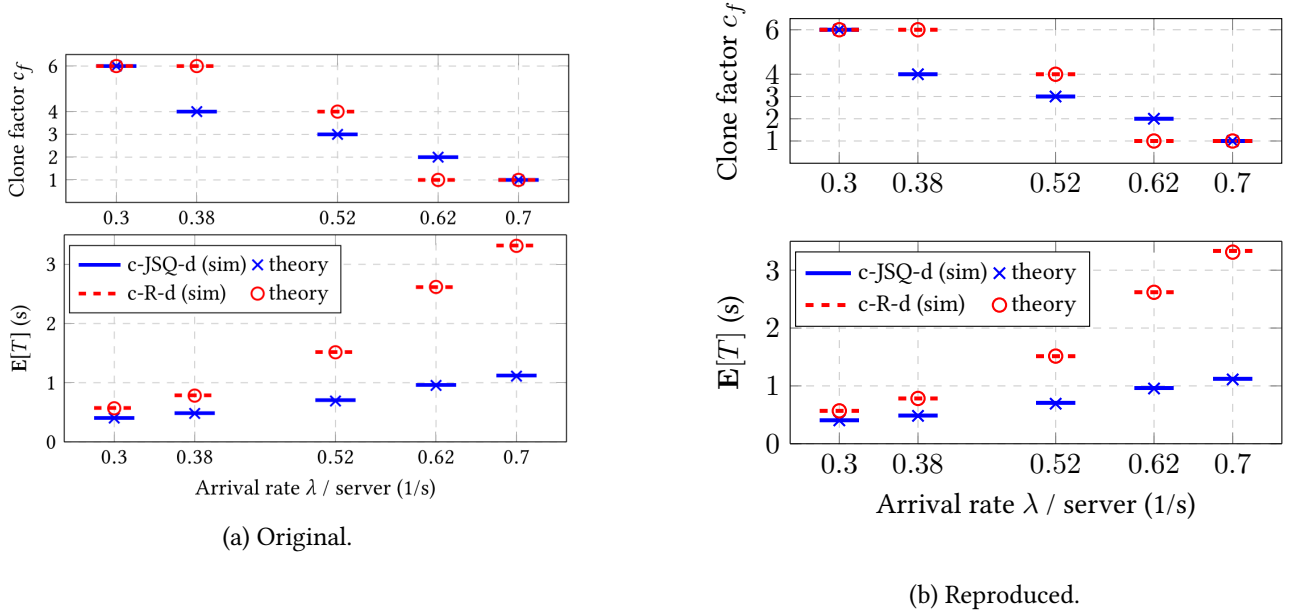
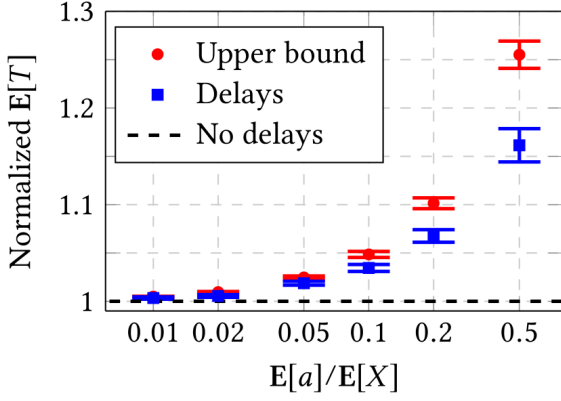



Figure 3: Co-designs: Comparison of theoretical values with 95% confidence intervals for the simulation results. The legend applies to both figures.

Fig. 8(a) in the paper corresponds to Figure 4 in this report. The data for this figure is generated by running the `sim_randomized_arrival_delays.py` wrapper script, which spawns 1000 simulation instances. Data is post processed by the `analyze_randomized_arrival_delays.m` matlab script, which completes successfully. The results are visually and substantially comparable, except for the point at $x = 0.5$, but the deviation can be deemed as negligible.

Fig. 8(b) in the paper corresponds to Figure 5 in this report. The data for this figure is generated by running the `sim_randomized_cancellation_delays.py` wrapper script, which spawns 1000 simulation instances. Data is post processed by the `analyze_randomized_cancellation_delays.m` matlab script, which completes successfully. The results are visually and substantially comparable, except for the




(a) Original.




(b) Reproduced.

Figure 4: Arrival delays only.

point at $x = 0.1$, but the deviation can be deemed as negligible.



(a) Original.



(b) Reproduced.

Figure 5: Cancellation delays only.


Fig. 8(c) in the paper corresponds to Figure 6 in this report. The data for this figure is generated by running the `sim_randomized_combined_delays.py` wrapper script, which spawns 1000 simulation instances. Data is post processed by the `analyze_randomized_combined_delays.m` matlab script, which completes successfully, except for the error bar at $x = 0.05$, but the deviation can be deemed as negligible and the trend might look overall better than the original plot in the paper.

Fig. 9 in the paper corresponds to Figure 7 in this report. The data for this figure is generated by running the `sim_randomized_sync_vs_nonsync_icpe.py` wrapper script, which spawns 4000 simulation instances. Data is post processed by the `analyze_randomized_scenarios.m` matlab script, which completes successfully, except for the points at $x = 0.9$, but the deviation can be deemed as negligible and the trend might look overall better than the original plot in the paper.

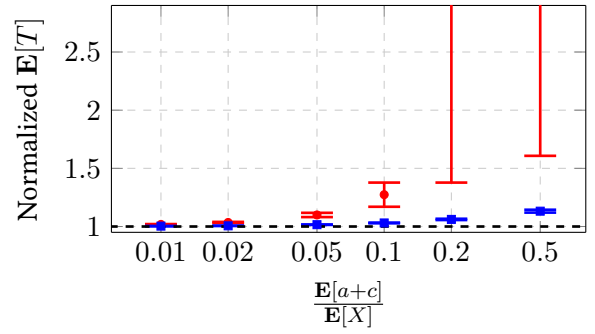
Overall, the plots which have been generated are consistent with the original publication.

References

- [1] Tommi Nylander, Johan Ruuskanen, Karl-Erik Årzén, and Martina Maggio, "Modeling of Request Cloning in Cloud Server Systems using Processor Sharing," In *Proceedings of the 11th ACM/SPEC International Conference on Performance Engineering*, April 2020.




(a) Original.




(b) Reproduced.

Figure 6: Both delays present.



(a) Original.



(b) Reproduced.

Figure 7: Simulations comparing $a\text{-}\ell\text{-d}$ to $c\text{-}\ell\text{-d}$ for random and JSQ. The normalization of $\mathbf{E}[T]$ is performed such that each value is divided by the value for the $c\text{-}\ell\text{-d}$ counterpart. The intervals represent 95% confidence intervals.