

Reproducibility Report for ACM SIGMOD 2021 Paper: “Clonos: Consistent Causal Recovery for Highly-Available Streaming Dataflows”

DONGHE KANG and SPYROS BLANAS, The Ohio State University, USA

This report tries to reproduce the Clonos system on an AWS instance. We can reproduce most figures and major findings in the paper. The experiments to evaluate overhead and Nexmark Q3 can always succeed but experiments of Nexmark Q8, concurrent failures and multiple failures may fail.

1 INTRODUCTION

This report describes the reproducibility result of the Clonos system [1], which is built on top of Apache Flink. Clonos provides a fault-tolerance approach to recover from failure fast. The goal of this report is to replay experiments in the paper to compare Clonos and Flink in terms of overhead under normal operations and recovery performance under failures. Experiments of the overhead and Nexmark Q3 reproduced results and findings while experiments of Nexmark Q8, concurrent failures and multiple failures failed in some configurations.

2 SUBMISSION

The submission includes the source code of Clonos, the scripts to compile, deploy and run experiments and the script to generate figures. The submission uses Docker to install dependent and distributed components such as Hadoop, Kafka, etc. It also has a pre-built Docker image of Clonos. Data and queries are automatically populated in experiments.

- **GitHub repository:** <https://github.com/PSilvestre/ClonosReproducibility>
- **Instructions to install dependencies:** `aws-instructions`
- **Parameters to configure experiments:** `experimental_parameters.sh`
- **Script to compile and run experiments:** `0_workflow.sh`

3 HARDWARE AND SOFTWARE ENVIRONMENT

We tried to replay experiments on an AWS **m5a.8xlarge** instance with the image **Ubuntu Server 18.04 LTS (HVM)**. Table 1 summarizes the configuration of the experimental environment.

Table 1. Hardware & Software environment

	Paper	Repro Review
CPU	Intel Xeon Gold 6130	AMD EPYC 7571
cores	16	16
GHz	2.1	2.1
RAM	96GB	123GB
Storage	SSD	EBS gp2 (SSD)

4 REPRODUCIBILITY EVALUATION

4.1 Process

After downloading the code from the GitHub repository, the first step is to install Docker following the `aws-instructions` file. The installation is straightforward if the user has the root privilege. A challenge is that most clusters disallow users to get the privilege.

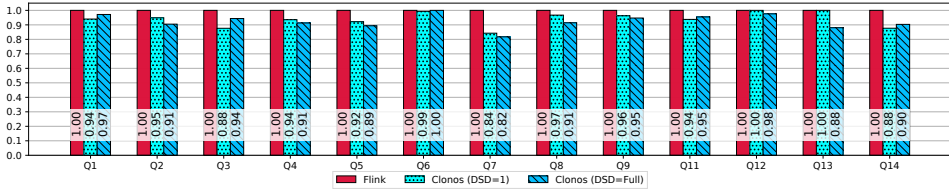
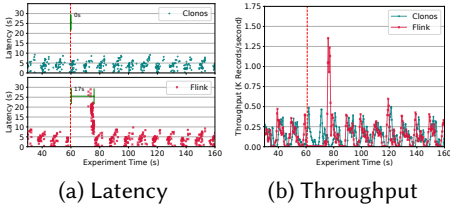
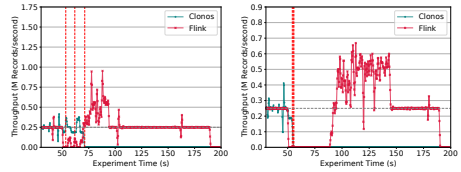


Fig. 1. Relative throughput of Clonos compared to Flink under normal operations.



(a) Latency (b) Throughput

Fig. 2. Performance of Nexmark Q3 under single-operator failures.



(a) Multiple (b) Concurrent

Fig. 3. Multiple and concurrent failures.

The second step is to configure experimental parameters in the files *experimental_parameters.sh*, *2.1_nexmark_experiments.sh* and *2.2_synthetic_experiments.sh*. The configuration is easy as parameters have detailed explanations.

The last step is to run all experiments by script *0_workflow.sh*. Experiments may be blocked randomly when we run on AWS. The authors added one option in the script to run a part of the experiments so that we can restart from the blocked experiment. The option is also helpful to tune parameters in different experiments.

4.2 Results

We can reproduce most experiments and major findings in the paper. The first experimental question in the paper is the overhead of Clonos under normal operations. Figure 1 shows the experimental result we produced to measure the overhead. Clonos and Flink have comparable throughput, which is the same as the result in the paper.

The second major finding in the paper is that Clonos can recover from single-operator failure faster than Flink. Figure 2 shows the reproduced result to evaluate Nexmark Q3. The result has the same pattern as the result in the paper. We cannot reproduce the experiment of Nexmark Q8 in some configurations as the throughput of Clonos is 0 after the failure. Clonos has the same problem in the experiments of concurrent and multiple failures (Figure 3).

The last finding in the paper is that Clonos recovers faster when there are multiple or concurrent failures. The experiment of multiple failures reproduced the finding but the results of the two experiments sometimes have different patterns as in the paper. In Figure 3a, Clonos has higher throughput than Flink when failures happen from 50th to 75th seconds. However, the throughput of Clonos is 0 after the 75th second which is different from the paper. Figure 3b shows a similar pattern. Authors explained that this may be because Clonos failed to recover and halted.

REFERENCES

- [1] Pedro F. Silvestre, Marios Fragkoulis, Diomidis Spinellis, and Asterios Katsifodimos. 2021. *Clonos: Consistent Causal Recovery for Highly-Available Streaming Dataflows*. Association for Computing Machinery, New York, NY, USA, 1637–1650. <https://doi.org/10.1145/3448016.3457320>