

FINJ: A Fault Injection Tool for HPC Systems

Alessio Netti¹, Zeynep Kiziltan¹, Ozalp Babaoglu¹
Alina Sirbu², Andrea Bartolini³, Andrea Borghesi³

¹ Department of Computer Science and Engineering
University of Bologna, Italy

{alessio.netti, zeynep.kiziltan, ozalp.babaoglu}@unibo.it

² Department of Computer Science
University of Pisa, Italy
alina.sirbu@unipi.it

³ Department of Electrical, Electronic and Information Engineering
University of Bologna, Italy
{andrea.bartolini, andrea.borghesi}@unibo.it

Abstract. We present FINJ, a high-level fault injection tool for High-Performance Computing (HPC) systems. FINJ provides support for custom workloads and allows generation of anomalous conditions through the use of fault-triggering executable programs. FINJ can be integrated seamlessly with any other lower-level fault injection tool, allowing users to create and monitor a variety of highly-complex and diverse fault conditions in HPC systems that would be difficult to recreate in practice. FINJ is suitable for large-scale experiments involving many, potentially interacting nodes in HPC systems, making it a very versatile design and evaluation tool.

Keywords: exascale systems, resiliency, fault detection, monitoring, benchmarking, open-source

1 Introduction

Motivation. High-Performance Computing (HPC) systems have become indispensable for economic growth and scientific progress in our modern society. As the performance of HPC systems increases, the value of the results they produce increases through higher-fidelity simulations, better predictive models and analysis of greater quantities of data. The resulting techniques, policy decisions and vastly-improved manufacturing processes in areas such as agriculture, engineering, transportation, materials, energy, health care, security and the environment are bound to impact most aspects of our lives. Today, HPC systems are also being used as fundamental “instruments” to achieve groundbreaking results in basic sciences ranging from particle physics to cosmology. Yet, many important problems in various fields remain unsolvable with current computational resources. *Exascale* HPC systems, capable of 10^{18} operations per second and rivaling the computational power of the human brain at the neuron level, are believed to

be essential for solving such problems [2]. Reaching exascale performance is the *moonshot* for modern HPC systems with many nations and companies engaged in an *arms race* towards achieving it.

Exascale systems, when they arrive, will come at a significant cost: scaling current technologies to exascale performance through massive parallelism will result in systems that have prohibitively-high levels of power consumption [17] and excessively-high failure rates [4]. Thus, to be usable in production environments with acceptable *Quality of Service* levels, exascale systems need to improve their power efficiency and resiliency by several orders of magnitude.

In our terminology, a *fault* is defined as an anomalous behavior at the software or hardware level that can lead to illegal system states (*errors*) and, in the worst case, to service interruptions (*failures*) [7]. In this paper, we limit our attention to improving the resiliency of HPC systems through the use of mechanisms for predicting, detecting and preventing errors and failures. An important technique in this endeavor is *fault injection*: the deliberate triggering of faults in a system so as to observe their behavior in a controlled environment, enable development of new prediction and response techniques and testing of existing ones [11]. For fault injection to be effective, dedicated tools are necessary, allowing users to trigger complex and realistic fault scenarios in a reproducible manner.

Related Work. Fault injection for prediction and detection purposes has been a topic of great interest in recent years. In [16, 8, 9, 6], the authors employed software-based fault injection techniques to observe the behavior and performance variations of HPC systems in anomalous conditions, and to detect such faults using system performance metrics. However, while characterizing the fault-simulating programs that were used, these works do not focus on the tools used to inject and coordinate the faults themselves in the system.

Several studies have proposed fault injection tools with varying levels of abstraction. Calhoun et al. [3] devised a compiler-level fault injection tool focused on memory bit-flip errors, targeting HPC applications. De Bardeleben et al. [5] proposed a logic error-oriented fault injection tool. This tool is designed to inject faults in virtual machines, by exploiting emulated machine instructions through the open-source virtual machine and processor emulator (QEMU). Both works focus on low-level fault-specific tools and do not provide functionality for the injection of complex workloads, and for the collection of produced data, if any.

Stott et al. [15] proposed NFTAPE, a high-level and generic tool for fault injection. This tool is designed to be integrated with other fault injection tools and triggers at various levels, allowing for the automation of long and complex experiments. The tool however has aged considerably, and is not publicly available. A similar fault injection tool was proposed by Naughton et al. [14], however, to the best of our knowledge, it has never progressed past the prototype stage and is also not publicly available.

Joshi et al. [12] introduced the PREFAIL tool, which allows for the injection of failures at any code entry point in the underlying operating system. This tool, like NFTAPE, employs a coordinator process for the execution of complex experiments. It is targeted at a specific type of fault (code-level errors) and does

not permit performing experiments focused on performance degradation and interference, among other fault types. Similarly, the tool proposed by Gunawi et al. [10], named FATE, allows the execution of long experiments; furthermore, it is focused on reproducing specific fault sequences, simulating real scenarios. Like PREFAIL, it is limited to a specific fault type, namely I/O errors, thus greatly limiting its scope.

Contributions. The main contribution of this paper is the design and implementation of FINJ, an open-source Python tool for fault injection targeted at HPC systems. An innovative feature of FINJ is the possibility of seamless integration with other injection tools targeted at specific fault types, thus enabling users to coordinate faults from different sources and different system levels. FINJ also provides *Workload* support, thus permitting users to specify lists of applications to be executed and faults to be triggered on multiple nodes at specific times with specific durations. FINJ thus represents a high-level, flexible tool, enabling users to perform complex and reproducible experiments, aimed at revealing the complex relations that may exist between faults, application behavior and the system itself. To the best of our knowledge, FINJ is the first portable, open-source tool that allows users to perform and control complex injection experiments, that can be easily integrated with heterogeneous faults types, and that includes workload support.

Organization. The rest of the paper is structured as follows. In Section 2, we describe the FINJ architecture (Section 2.1), its components (Section 2.2) and their implementation (Section 2.3). In Section 3, we present a simple use case to show how FINJ can be deployed, while Section 4 concludes the paper.

2 FINJ Architecture

In this Section we discuss how fault injection is achieved in FINJ. We then present its architecture, together with some implementation details. Due to its portable and modular nature, customizing FINJ for different purposes is easy.

2.1 Architecture Overview

Fault injection in FINJ is achieved through *tasks* that are executed on target nodes: each task corresponds to a particular application, which can either be a benchmark program or a fault-triggering program. This approach allows for great flexibility, as FINJ can be integrated with any type of low-level fault injection tool. A task is defined by the following attributes:

- *args*: the full shell command required to run the selected task. The command must refer to an executable file that can be accessed from the target hosts;
- *timestamp*: the time in seconds at which the task must be started, relative to the starting time of the injection session;

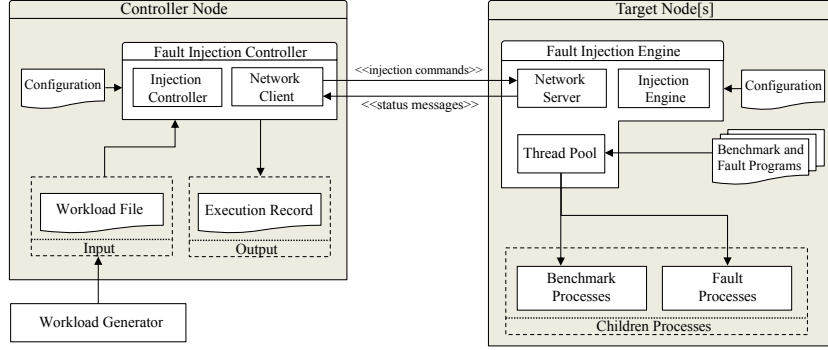


Fig. 1. Architecture of the FINJ tool showing the division between a controller node (left) and a target node (right).

- *duration*: the task’s *maximum allowed* duration, expressed in seconds, after which it will be abruptly terminated. This duration can serve as an *exact* duration as well, with FINJ restarting the task if it finishes earlier, and terminating it if it lasts more. This behavior depends on the FINJ configuration (see Section 2.2). A duration of 0 implies that the task is always allowed to run until its termination;
- *isFault*: defines whether the task corresponds to a fault-triggering program, or to a benchmark application;
- *seqNum*: a sequence number used to uniquely identify the task;
- *cores*: the list of CPU cores that the task is allowed to use on target nodes, enforced through a *NUMA Control* policy [13]; this attribute is optional.

A set of tasks defines a *workload*, which is a succession of scheduled fault and benchmark executions at specific times, reproducing a realistic working environment for the fault injection process. A particular execution of a given workload then constitutes an *injection session*.

FINJ consists of two basic components: a fault injection *controller*, and a fault injection *engine*. The two components correspond to processes that must be run on the nodes subject to injection experiments. Communication between them is achieved through TCP *sockets* using a simple message-based protocol. The high-level structure of the FINJ architecture is illustrated in Figure 1.

FINJ Controller. The controller is the orchestrator of the injection process, and should be run on an external node that is not affected by the faults. The controller maintains connections to all nodes involved in the injection session, which run fault injection *engine* instances and whose addresses are specified by users when launching the program. Therefore, injection sessions can be performed on multiple nodes at the same time. The controller reads task entries from the selected workload: the reading process is incremental, and tasks are read in real-time a few minutes before their expected execution, according to their relative

time-stamp. For each task the controller sends a command to all target hosts, instructing them to start the new task at the specified time. Finally, the controller collects all status messages produced by the target hosts, and stores them in a separate file for each host. These status messages are related to the start and termination of each single task, besides status changes in the host (for example, when connection is lost and re-established).

FINJ Engine. The engine is structured as a daemon, perpetually running on nodes that are expected to be subject to injection sessions. The engine waits for task commands to be received from remote controller instances. Engines can be connected to multiple controllers at the same time, and status messages will be sent to all of them. However, task commands are accepted from one controller at a time, which is defined as the *master* of the injection session. The engine manages received task commands by assigning them to a dedicated thread from a *pool*. The thread manages all aspects related to the execution of the task, such as spawning the necessary subprocesses and sending status messages to controllers when relevant events (such as the start or termination of the task) occur.

2.2 Components

FINJ is based on a highly modular architecture, and therefore it is very easy to customize single components in order to add or tune features.

Network. Engine and controller instances communicate through a network layer in the FINJ tool. Communication is achieved through TCP sockets that are used to implement a simple message-based protocol. This design choice is motivated by the fact that the volume of data sent during injection sessions is extremely low, while high reliability is a desirable quality. Specifically, a message *client* and *server* were implemented: clients are used by FINJ controllers in order to connect to servers hosted on FINJ engine instances. Messages can then be either *commands*, related to single tasks and imposed by controllers, or *status* messages, which are sent by engines and are related to status changes in their system. All messages are in the form of *dictionaries*. This component also handles resiliency features such as automatic re-connection from clients to servers, since temporary connection losses are to be expected in a fault injection context.

Thread Pool. Task commands in FINJ engines are assigned to a thread in a pool as they are received: each thread manages all aspects of a task assigned to it. Specifically, the thread sleeps until the scheduled starting time of the task (according to its time-stamp); then, it spawns a subprocess running the specified task, and sends a message to all connected controller hosts to inform them of the event. At this point, the thread waits for the task's termination, depending on its duration and on the current configuration. Finally, the thread sends a new status message to all connected hosts informing them of the task's termination, and returns to sleep. The amount of threads in the pool, which is a configurable parameter, determines the maximum number of concurrent tasks that can be executed in a given injection session.

Input and Output. In FINJ, input and output of all data related to injection sessions are performed by controller instances, and are handled by *reader* and *writer* entities. By default, these employ the CSV format, which was chosen due to its extreme simplicity and generality, but they can be easily customized by users for other formats. *Input* in FINJ is constituted by *workload* files: as mentioned in Section 2.1, these files include one entry for each task that must be executed in the injection session. Using the CSV format makes workload files extremely readable, and manually writing workloads corresponding to highly specific test cases can be easily achieved as well. FINJ *output*, instead, is made up of two parts. The first is the *execution log*, which contains entries corresponding to status changes in the target node, namely the start and termination of tasks, errors that are encountered if any, and connection loss or recovery events. The second part of FINJ output is related to tasks: all output text written to the *stdout* or *stderr* channels during their execution, if any, is reported to controllers, and is stored in separate plain-text files in a directory alongside the main output file, each named according to the task’s name and sequence number.

Configuration. The FINJ tool’s runtime behavior is customizable by means of a configuration file. This file is in JSON format and includes several options that alter the behavior of either controller or engine instances. Among the basic options, it is possible to specify the listening TCP port for engine instances, and the list of addresses of target hosts, to which controller instances should connect at launch time. The latter is useful when injection sessions must be performed on large sets of nodes, whose addresses can be conveniently stored in a file. More complex options are also available: for instance, it is possible to define a series of commands corresponding to tasks that must be launched together with FINJ, and must be terminated with it. This option proves especially useful when users wish to set up monitoring frameworks, such as the *Lightweight Distributed Metric Service* (LDMS) [1], to be launched together with FINJ in order to collect system performance metrics during injection sessions.

Workload Generation. While writing workload files manually is possible, this is time-consuming and not desirable for long injection sessions. Therefore, we implemented in FINJ a *workload generation* tool, which can be used to automatically generate workload files with certain statistical features, while trying to combine flexibility and ease of use. The workload generation process is controlled by three parameters: a maximum *time span* for the total duration of the workload expressed in seconds, a statistical distribution for the *duration* of tasks, and another one for their *inter-arrival* times. These distributions are separated in two sets, for fault and benchmark tasks, thus amounting to a total of four. They can be either specified analytically by the user or can be fitted from real data, thus reproducing realistic behavior.

A workload is composed as a series of fault and benchmark tasks that are selected from a list of possible shell commands. To control the composition of workloads, users can optionally associate to each command a probability for its selection during the generation process, and a list of CPU cores for its execution,

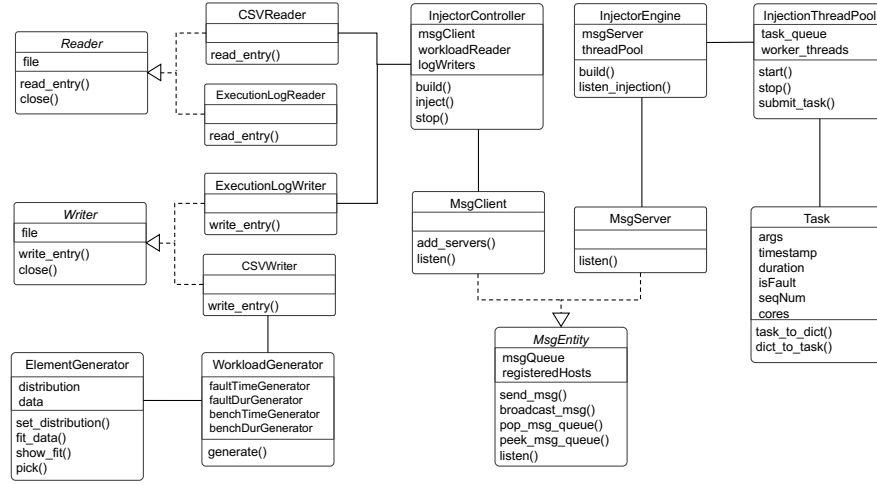


Fig. 2. Class diagram of the FINJ tool.

as explained in Section 2.1. By default, commands are picked uniformly. Having defined its parameters, the workload generation process is then fairly simple: tasks are randomly generated in order to achieve statistical features close to those specified as input, and are written to an output CSV file, until the maximum imposed time span is reached. Alongside the full workload, a *probe* file is also produced: this workload file contains one entry for each task type, all with a short fixed duration, and represents a lightweight workload version. This file can be used during the setup phase to test the correct configuration of the system, making sure that all tasks are correctly found and executed on the target hosts, without having to run the entire heavy workload.

2.3 Implementation

FINJ is implemented in Python, an object-oriented, high-level interpreted programming language⁴, and can be used on all major operating systems. All FINJ dependencies are included in the Python distribution, and the only optional external dependency is the *scipy* package, which is needed for the workload generation functionality. The source code is publicly available on GitHub⁵ under the MIT license, together with its documentation, usage examples and several fault-triggering programs. FINJ works on Python versions 3.4 and above.

In Figure 2 we illustrate the class diagram for the FINJ tool. The *engine* and *controller* entities are respectively represented by the *InjectorEngine* and *InjectorController* classes. Users can instantiate these classes and start injection sessions directly, by using the *listen* method to put the engine in listening mode,

⁴ <https://www.python.org/events/python-events/>

⁵ <https://github.com/AlessioNetti/fault.injector>

and the *inject* method of the controller, which allows to start the injection session itself. However, scripts are supplied with FINJ to create controller and engine instances from a command-line interface, simplifying the process. This method will be discussed in Section 3. The *InjectionThreadPool* class, instead, supplies the thread pool implementation used to execute and manage tasks.

The network layer of the tool is represented by the *MsgClient* and *MsgServer* classes, which implement the message and queue-based client and server used for communication. Both classes are implementations of the *MsgEntity* abstract class, which provides the interface for sending and receiving messages, and implements the basic mechanisms that regulate the access to the underlying queue.

Input and output are instead handled by the *Reader* and *Writer* abstract classes and their implementations: *CSVReader* and *CSVWriter* handle the reading and writing of workload files, while *ExecutionLogReader* and *ExecutionLogWriter* handle execution logs generated by injection sessions. Since these classes are all implementations of abstract interfaces, it is easy for users to customize them for different formats. Tasks are modeled by the *Task* class that contains all attributes specified in Section 2.1.

Lastly, access to the workload generator is provided through the *WorkloadGenerator* class, which is the interface used to set up and start the generation process. This class is backed by the *ElementGenerator* class, which offers basic functionality for fitting data and generating random values. This class acts as a wrapper on scipy’s *rv_continuous* class, which generates random variables.

3 Using FINJ

In this section we sketch how FINJ can be deployed and used through its command-line interface. We also demonstrate the flow of execution through a concrete example carried out on a real HPC node.

3.1 Instantiation

Two Python scripts named *finj_engine* and *finj_controller* are supplied to start *engine* and *controller* instances, respectively. In the following we will describe the use of these scripts. Due to space reasons, we will not present all available configuration options for FINJ engines and controllers, but these are thoroughly explained on the GitHub repository for the tool.

The first script, *finj_engine*, should be run on target nodes for injection, and will start the FINJ engine daemon listening for injection requests by controllers:

```
python finj_engine [ -c CONFIGFILE ] [ -p PORT ]
```

The script can take two arguments. The first is *-c*, which is the path to the JSON configuration file for the engine. The second option, *-p*, is instead the TCP port to be used by the engine for listening to external requests. Both arguments are optional, and omitting them will result in the engine using a default configuration. Once started, the engine will idle until a controller connects to it.

A controller instance can be started with the *finj_controller* script:

```
timestamp;duration;seqNum;isFault;cores;args
0;1719;1;False;0;./hpl_lininput
587;291;2;True;0;./leak 291 1
1171;244;3;True;0;sudo ./cpufreq 244
```

Fig. 3. A sample CSV workload that can be used with FINJ.

```
python finj_controller [ -s ] [ -c CONFIGFILE ] [ -w
    ↪ WORKLOADFILE ] [ -a ADDRESSLIST ] [ -m MAXTASKS ]
```

The `-c` option works in the same manner as seen for the `finj_engine` script. The `-w` option is instead the path to a CSV workload file to be injected in the target hosts. If none is supplied, the controller will connect to the target hosts in *listening* mode, thus not injecting any workload, but only receiving and storing all status messages sent from them. This feature is useful when users wish to remotely monitor an ongoing injection session. The `-a` option accepts a list of addresses for target hosts running engine instances, specified in `hostname:port` format and separated by commas. If no address is supplied, the controller will search for valid addresses in the current configuration file, if any, and abort if none is found. Finally, the `-m` option specifies a maximum number of tasks to be injected from the selected workload, and the `-s` option enables the *silent* mode, in which all log messages from the controller are suppressed, except for errors.

3.2 Sample Execution

In this subsection we will consider a sample fault injection session carried out using FINJ. The employed workload file, named *sample.csv*, is illustrated in Figure 3. The test was carried out on one node of an HPC system equipped with two Intel Xeon E5-2630 v3 CPUs, 128GB of RAM, and running CentOS 7.3.

In this workload, the first task is the Intel Distribution⁶ for the well-known *High-Performance Linpack* (HPL) benchmark, optimized for Intel Xeon CPUs. This task starts at time 0 in the workload, and has a maximum allowed duration of 30 minutes. The following two tasks are fault-triggering programs: *leak*, a program that creates a low-intensity memory leak in the system, eventually using all available RAM, and *cpufreq*, which uses the Intel P-State driver in the Linux kernel⁷ to dynamically reduce the maximum allowed CPU frequency, emulating performance degradation. This last script requires appropriate permissions, so that users can access the files controlling Linux CPU governors. All of the tasks are forced to run on core 0 of the target host’s CPU, and the HPL benchmark was configured to use only one thread. Also note that the tasks must be available at the specified path on the systems running the FINJ engine, which in this case is relative to the location of the launching script.

⁶ <https://software.intel.com/en-us/mkl-windows-developer-guide-overview-of-the-intel-distribution-for-linpack-benchmark>

⁷ <https://www.kernel.org/doc/Documentation/cpu-freq/intel-pstate.txt>

```

timestamp;type;args;seqNum;duration;isFault;cores;error
1522524987;command_session_s;None;None;None;None;None;None
1522525007;status_start;./hpl_lininput;1;1719;False;0;None
1522525594;status_start;./leak 291 1;2;291;True;0;None
1522525885;status_end;./leak 291 1;2;291;True;0;None
1522526178;status_start;sudo ./cpufreq 244;3;244;True;0;None
1522526422;status_end;sudo ./cpufreq 244;3;244;True;0;None
1522526726;status_end;./hpl_lininput;1;1719;False;0;None
1522526727;command_session_e;None;None;None;None;None;None

```

Fig. 4. A sample output file produced by FINJ after an injection session for the workload specified in Figure 3.

Having defined the workload, the injection engine and controller must be started. Using the default configuration, and supposing that the test must be performed locally, this can be accomplished with the two following commands:

```

python finj_engine -p 30000 &
python finj_controller -w sample.csv -a localhost:30000

```

The controller instance will then connect to the engine and start executing the workload, storing all output in a unique CSV file for each target host. When this process is finished, the controller terminates. The output CSV files for our example have the format shown in Figure 4: each entry represents a status change event, which in this case is the start or termination of tasks, and is flagged with its absolute time-stamp on the target host. In addition, we also find an *error* field, detailing possible errors that were encountered. Note that the file is opened and closed by session *start* and *end* entries: the presence of these ensures that the injection process did not encounter errors and that the entire workload was processed successfully.

At this point, the data generated by FINJ can be easily compared with other data, for example performance metrics collected through a monitoring framework, in order to perform fault detection studies or simply to better understand the system’s behavior under faults. For this test, we used the LDMS framework [1] to collect resource usage and performance metrics on the target host at each second, for the duration of the injection session. In Figure 5 we show the total RAM usage and the CPU frequency of core 0. The benchmark’s profile is simple, showing a constant CPU frequency while RAM usage slowly increases as the application performs tests on increasing matrix sizes. On the other hand, the effect of our fault programs, marked in gray, can be clearly observed in the system: the *cpufreq* fault causes a sudden drop in CPU frequency, resulting in reduced performance and longer computation times, while the *leak* fault causes a steady, linear increase in RAM usage. Even though saturation of the available RAM is not reached, this peculiar behavior of the leak fault can be used for prediction purposes, opening up the possibility of preventing failure events.




Fig. 5. CPU Frequency and RAM Usage, as monitored on the target system during the sample injection session.

4 Conclusions

We have presented FINJ, a high-level tool for fault injection and monitoring in HPC systems. FINJ allows for automation of complex experiments, and for reproducing anomalous behaviors in a deterministic, simple way. FINJ is open-source and implemented in Python, an object-oriented interpreted programming language available on all major operating systems, and has no dependencies for its core operation. This, together with the simplicity of its command-line interface, makes the deployment of FINJ on large-scale systems trivial. Since FINJ is based on the use of tasks, which are external fault-triggering executable programs, the tool can be seamlessly integrated with lower-level fault injection tools of any kind, that range from the application level to the kernel, or even hardware level. The use of workloads in FINJ also allows to reproduce complex, specific fault conditions in HPC systems, and to reliably perform large-scale experiments involving multiple nodes of a system at the same time.

As future work, we plan to perform scalability studies on the FINJ tool, especially for the *controller* component, by deploying it on a large-scale HPC environment. We have already performed extensive testing on the system presented in Section 3.2 with excellent preliminary results. Also, we plan to develop additional features, such as the ability to build workloads in which the order of tasks is defined by *causal* relationships rather than time-stamps, which might simplify the triggering of extremely specific anomalous states in a given system.

Acknowledgements. A. Netti has been supported by a research fellowship from the *Oprecomp-Open Transprecision Computing* project. A. Sîrbu has been partially funded by the EU project *SoBigData Research Infrastructure — Big Data and Social Mining Ecosystem* (grant agreement 654024).

References

1. Agelastos, A., Allan, B., Brandt, J., Cassella, P., Enos, J., Fullop, J., Gentile, A., Monk, S., Naksinehaboon, N., Ogden, J., et al.: The lightweight distributed metric service: a scalable infrastructure for continuous monitoring of large scale computing systems and applications. In: Proc. of SC 2014. pp. 154–165. IEEE (2014)
2. Ashby, S., Beckman, P., Chen, J., Colella, P., Collins, B., Crawford, D., et al.: The opportunities and challenges of exascale computing. Summary Report of the Advanced Scientific Computing Advisory Committee (ASCAC) Subcommittee pp. 1–77 (2010)
3. Calhoun, J., Olson, L., Snir, M.: Flipit: An llvm based fault injector for hpc. In: Proc. of Euro-Par 2014. pp. 547–558. Springer (2014)
4. Cappello, F., Geist, A., Gropp, W., Kale, S., Kramer, B., Snir, M.: Toward exascale resilience: 2014 update. Supercomputing frontiers and innovations 1(1), 5–28 (2014)
5. DeBardeleben, N., Blanchard, S., Guan, Q., Zhang, Z., Fu, S.: Experimental framework for injecting logic errors in a virtual machine to profile applications for soft error resilience. In: Proc. of Euro-Par 2011. pp. 282–291. Springer (2011)
6. Ferreira, K.B., Bridges, P., Brightwell, R.: Characterizing application sensitivity to os interference using kernel-level noise injection. In: Proc. of SC 2008. p. 19. IEEE Press (2008)
7. Gainaru, A., Cappello, F.: Errors and faults. In: Fault-Tolerance Techniques for High-Performance Computing, pp. 89–144. Springer (2015)
8. Guan, Q., Chiu, C.C., Fu, S.: Cda: A cloud dependability analysis framework for characterizing system dependability in cloud computing infrastructures. In: Proc. of PRDC 2012. pp. 11–20. IEEE (2012)
9. Guan, Q., Fu, S.: Adaptive anomaly identification by exploring metric subspace in cloud computing infrastructures. In: Proc. of SRDS 2013. pp. 205–214. IEEE (2013)
10. Gunawi, H.S., Do, T., Joshi, P., Alvaro, P., Hellerstein, J.M., Arpaci-Dusseau, A.C., Arpaci-Dusseau, R.H., Sen, K., Borthakur, D.: Fate and destini: A framework for cloud recovery testing. In: Proc. of NSDI 2011. p. 239 (2011)
11. Hsueh, M.C., Tsai, T.K., Iyer, R.K.: Fault injection techniques and tools. Computer 30(4), 75–82 (1997)
12. Joshi, P., Gunawi, H.S., Sen, K.: Prefail: A programmable tool for multiple-failure injection. In: ACM SIGPLAN Notices. vol. 46, pp. 171–188. ACM (2011)
13. Lameter, C.: Numa (non-uniform memory access): An overview. Queue 11(7), 40 (2013)
14. Naughton, T., Bland, W., Vallee, G., Engelmann, C., Scott, S.L.: Fault injection framework for system resilience evaluation: fake faults for finding future failures. In: Proc. of Resilience 2009. pp. 23–28. ACM (2009)
15. Stott, D.T., Floering, B., Burke, D., Kalbarczpk, Z., Iyer, R.K.: Nftape: a framework for assessing dependability in distributed systems with lightweight fault injectors. In: Proc. of IPDS 2000. pp. 91–100. IEEE (2000)
16. Tuncer, O., Ates, E., Zhang, Y., Turk, A., Brandt, J., Leung, V.J., Egele, M., Coskun, A.K.: Diagnosing performance variations in hpc applications using machine learning. In: Proc. of ISC 2017. pp. 355–373. Springer (2017)
17. Villa, O., Johnson, D.R., Oconnor, M., Bolotin, E., Nellans, D., Luitjens, J., et al.: Scaling the power wall: a path to exascale. In: Proc. of SC 2014. pp. 830–841. IEEE (2014)