

A Consensus-Based Fault-Tolerant Event Logger for High Performance Applications

Edson Tavares de Camargo^{1,2(✉)}, Elias P. Duarte Jr.², and Fernando Pedone³

¹ Department of Informatics, Federal University of Paraná (UFPR), Curitiba, Brazil

² Federal Technology University of Paraná (UTFPR), Toledo, Brazil

`edson@utfpr.edu.br`, `elias@inf.ufpr.br`

³ University of Lugano (USI), Lugano, Switzerland

`fernando.pedone@usi.ch`

Abstract. Most message logging protocols rely on a centralized event logger to store information (i.e., the determinants) to allow the recovery of an application process. This centralized approach, besides suffering from the single point of failure problem, represents a bottleneck for the efficiency of message logging protocols. In this work, we present a fault-tolerant distributed event logger based on consensus that outperforms the centralized approach. We implemented the event logger of MPI determinants using the Paxos algorithm. Our event logger inherits the Paxos properties: safety is guaranteed even if the system is asynchronous and liveness is guaranteed despite processes failures. Experimental results are reported for the performance of the distributed event logger based both on classic Paxos and parallel Paxos applied to AMG (Algebraic Multi-Grid) and NAS Parallel Benchmark applications.

1 Introduction

Most traditional strategies to deal with failures in HPC systems are based on rollback-recovery mechanisms [3, 6, 16]. These strategies allow applications to recover from failures without losing previously computed results. Message logging is a class of rollback-recovery technique that unlike coordinated checkpoint strategies does not require all processes to coordinate to save their state during normal execution and to restart after a single process failure.

Message logging relies on the piecewise deterministic assumption. This assumption states that all nondeterministic events that a process executes can be identified and the information necessary to replay each event during recovery can be logged in tuples called *determinants* [7]. By replaying the determinants in their exact original order, a process can deterministically recreate its pre-failure state. Most message logging protocols suppose that reception events (i.e., message receiving events) are the only possible nondeterministic events in the execution [3]. Consequently, a crucial task in message logging is to reliably save and restore the determinants without penalizing the performance.

The component responsible for reliably logging determinants is the *event logger*. The event logger receives the determinants from the application processes, stores them locally, and notifies the application processes. Previous works based on message logging typically assume that the event logger is a centralized entity (e.g., [1,4,17]), and thus it cannot tolerate failures. Indeed, the failure of the event logger would bring the execution to a halt as application processes would no longer be able to save the determinants.

The main goal of this paper is to propose a fault-tolerant event logger that has performance comparable to or better than a centralized event logger. In particular, our replicated event logger does not require extra system resources (i.e., physical nodes) in comparison with a centralized event logger and can tolerate a configurable number of failures. When configured to tolerate a single failure, our consensus-based event logger needs the same number of messages and communication steps (i.e., network delays) to log a determinant as a centralized event logger. We also show in the paper that the myth that fault tolerance introduces overheads is not completely unfounded since the indiscriminate use of existing fault-tolerance techniques can indeed lead to expensive solutions.

We implemented two fault-tolerant event loggers based on the Paxos algorithm [11]. One is based on classic Paxos and the other on a configuration we call parallel Paxos. We conducted a number of experiments comparing them to a centralized event logger. We evaluated the performance of our event logger implementations using the LU and MG kernels from the NAS Parallel Benchmark (NAS-PB) and the Algebraic MultiGrid (AMG) application. Our results show that the replicated event logger based on parallel Paxos consistently outperforms a centralized event logger while providing configurable fault tolerance.

The rest of the paper is organized as follows. Section 2 briefly overviews rollback-recovery, including message logging, and the event logger. Section 3 reviews the Paxos algorithm and presents our consensus-based event loggers. Section 4 presents our implementations of the event logger and experimental results and Sect. 5 concludes the paper.

2 Log-Based Rollback Recovery

Rollback-recovery techniques are often used to provide fault tolerance to HPC applications so that they can restart from a previously saved state [3,6]. Rollback-recovery assumes a distributed system that is a collection of application processes that communicate through a network and have access to stable storage that survives failures [7]. Processes save recovery information periodically on stable storage during their failure-free execution. After the occurrence of a failure, the process that failed uses the recovery information to restart its computation from a past state. The recovery information includes at least the state of the participating processes, called *checkpoints*. Some protocols may also include the logging of nondeterministic events, encoded in tuples called *determinants*.

Log-based approaches, or simply “message logging”, use both checkpoints and logging of nondeterministic events to avoid the drawbacks of both uncoordinated and coordinated checkpointing [14]. Message logging protocols assume the application is piece-wise deterministic [10]. This assumption asserts that all nondeterministic events executed by a process can be identified and the information necessary to replay each event during recovery can be logged in determinants. An *event* corresponds to a computational or communication step of a process. Most message logging protocols assume that message reception is the only nondeterministic event.

Depending on how determinants are logged, message logging protocols can be pessimistic, optimistic or causal [7]. In pessimistic logging, a process first stores the determinant of a nondeterministic event (e.g., in remote storage) before delivering the message. Despite the fact that pessimistic logging simplifies recovery and garbage collection, it presents an overhead on failure-free scenarios: the application has to wait for the determinant to be stored in order to proceed. In optimistic logging, processes log determinants asynchronously, thereby reducing the overhead. However, optimistic protocols allow orphan processes to be created due to failures and lead to more expensive recovery.

Usually, in message logging approaches the determinant of every received message is logged. However, it is possible to reduce the overall number of logged messages by identifying which events are deterministic and which are nondeterministic [2]. An event is deterministic when from the current state there is only one possible outcome state for the event. If an event can result in several different states, then it is nondeterministic. Message receptions with an explicitly identified sender are deterministic events and do not need to be logged; if the source is left unspecified then message receptions are nondeterministic. Message receptions with an explicitly identified sender are deterministic events and do not need to be logged; if the source is left unspecified then message receptions are nondeterministic. For example, a nondeterministic event occurs in MPI when the receiving process uses the tag `MPI_ANY_SOURCE` in `MPI_Recv`. As stated in [5], several MPI applications contain only deterministic communication events. However, many important MPI applications are nondeterministic, including all master-slave applications. Furthermore, programmers usually include nondeterministic communication events in the code to improve performance.

2.1 The Event Logger

The event logger plays an important role in message-logging protocols [1]. It receives the determinants from the application processes, stores them locally, and notifies the application processes after determinants are stored. The performance of the event logger has a major impact on the efficiency of message-logging protocols as showed in [1, 17, 18], and many protocols implement the event logger as a centralized (i.e., non-replicated) component [1, 3].

To the best of our knowledge, the only existing distributed event logger is proposed in [18] for O2P [17]. That protocol offers a distributed way to save determinants. Despite being related to our work, the entire solution fails if an

acknowledgement is not received by the sender; i.e., the fault tolerance of the solution is not guaranteed. Our proposed event logger based on consensus is both distributed and fault-tolerant. The Paxos protocol ensures progress with a majority of non-faulty processes and safety even if the system is asynchronous.

3 Consensus and Message Logging

Consensus is a fundamental abstraction in fault-tolerant distributed computing. In this section, we review the consensus problem, present Paxos, one of the most prominent consensus algorithms, and discuss how to efficiently implement a fault-tolerant event logger with Paxos.

3.1 Consensus and State Machine Replication

Consensus can be used to build a highly available event logging service using the state machine replication approach [19]. State machine replication regulates how commands must be propagated to and executed by the replicas in order for the service to be consistent. In our particular case, the commands are requests to save a determinant, propagated to and executed by replicas of the event logger. Command propagation has two requirements: (i) every non-faulty replica must receive every command and (ii) no two replicas can disagree on the order of received and executed commands. If command execution is deterministic, then replicas will reach the same state and produce the same output upon executing the same sequence of commands.

Intuitively, consensus captures the command propagation requirements of state machine replication. More precisely, consensus is defined by three abstract properties: (a) If a replica decides on a value, then the value was proposed by some process (*validity*). (b) No two replicas decide differently (*agreement*). (c) If a non-faulty process proposes a value, then eventually all non-faulty replicas decide some value (*termination*). From the requirements of state machine replication and the guarantees provided by consensus, it should be clear that state machine replication can be implemented as a series of consensus instances, where the i -th consensus instance decides on the i -th command (or batch of commands) to be executed by the replicas [11].

State machine replication and consensus provide a principled approach to ensuring that replicas are consistent despite failures. This approach should not be overlooked since ad hoc solutions to replication must face subtle impossibilities in the design of distributed systems subject to process failures [8].

3.2 The Paxos Protocol

Paxos is a fault-tolerant consensus algorithm designed for state machine replication [12]. Paxos has important characteristics: it is safe under asynchronous assumptions, live under weak synchronous assumptions, ensures progress with a majority of non-faulty processes, and assumes a crash-recovery failure model.




Fig. 1. Three implementations of an event logger. The centralized approach has a single event logger (R_0) and thus cannot tolerate any failures. The Paxos-based approaches can tolerate one failure (i.e., $f = 1$). Paxos coordinators already executed the first phase of the protocol and can proceed with the second phase upon receiving a command.

Paxos distinguishes the following roles that a process can play: *proposers*, *acceptors* and *learners*. Proposers propose a value, acceptors choose a value, and learners learn the decided value. A single process can assume any of those roles, and multiple roles simultaneously. Paxos is resilience-optimum [13]: to tolerate f failures it requires $2f + 1$ acceptors—that is, to ensure progress, a quorum of $f + 1$ acceptors must be non-faulty.

An instance of Paxos proceeds in two phases: during the first phase, a proposer selects a unique round number and sends a *prepare* request to a quorum of acceptors. Upon receiving a prepare request with a round number bigger than any round the acceptor previously received, the acceptor responds to the proposer promising that it will reject any future requests with smaller round numbers. If the acceptor already accepted a command for the current instance (explained next), it will return this command to the proposer, together with the round number received when the command was accepted. When the proposer receives answers from a quorum of acceptors, it proceeds to the second phase.

In the second phase, the proposer selects a command according to the following rule. If no acceptor in the quorum of responses accepted a command, the proposer can select a new command for the instance; however, if any of the acceptors returned a command in the first phase, the proposer chooses the command with the highest round number. The proposer then sends an *accept* request with the round number used in the first phase and the command chosen to a quorum of acceptors. When receiving such a request, the acceptors acknowledge it by sending a message to the coordinator and learners, unless the acceptors have already acknowledged another request with a higher round number. When a quorum of acceptors accepts a command consensus is reached.

If multiple proposers simultaneously execute the procedure above for the same instance, then no proposer may be able to execute the two phases of the protocol and reach consensus. To avoid scenarios in which proposers compete indefinitely, a *coordinator* process can be chosen. In this case, proposers submit commands to the coordinator, which executes the first and second phases of the protocol. If the coordinator fails, another process takes over its role. Paxos ensures consistency despite concurrent coordinators and termination in the presence of a single coordinator.

A coordinator can optimize performance by executing the first phase of the protocol for a batch of instances before it receives any commands [11]. This is possible because the coordinator only sends commands in the second phase of the protocol. With this optimization, a command can be chosen in three communication steps: the message from the proposer to the coordinator, the accept request from the coordinator to the acceptors, and the response to this request from the acceptors to the coordinator and learners.

3.3 Consensus-Based Message Logging

We now propose two protocols based on Paxos to render the event logger fault-tolerant: Classic Paxos and Parallel Paxos. Similarly to a centralized event logger, our protocols log nondeterministic events only; the message payload is saved by the sender- [9]. A determinant contains the sender of a message, the message identifier, and the message receiving order. Periodically, each process performs a checkpoint in order to save its state.

Our first protocol, Classic Paxos, is based on classic state machine replication. Application processes are proposers and the event logger replicas are acceptors and learners. Every application process submits commands to the coordinator, a process among the acceptors, to log determinants. The coordinator receives commands, executes Paxos to log the commands in a quorum of replicas, and sends replies to the application process (see Fig. 1). In “good executions” (i.e., in the absence of process failures) a determinant is logged after four communication steps and $2f + 2$ point-to-point message exchanges. By contrast, a centralized event logger can log events after two communication steps and two message exchanges.

In our second protocol, Parallel Paxos, we assign a separate sequence of Paxos executions to each application process. This means that each process has its set of replicas, which allows important optimizations. First, since each process has its own sequence of Paxos executions, the process does not compete with other processes in executions of Paxos and therefore, there is no need for a coordinator; in good executions, the process is the only proposer in its sequence of Paxos. Second, by using different sets of replicas, performance is no longer capped by what the coordinator and the acceptors can handle. Third, we can now co-locate the application process and the acceptor-coordinator in the same process. In good runs, this scheme can log a determinant after two communication steps and $2f$ messages.

An event logger implemented with Parallel Paxos presents the same number of communication steps as a centralized event logger, while tolerating a configurable number of failures and scaling performance. When configured to tolerate one failure ($f = 1$), it exchanges the same number of messages per logging operation as the centralized logger. Moreover, to save resources, “free acceptors” (i.e., acceptors not collocated with application processes) can be placed in the same physical node. For example, a single node can host all free acceptors, i.e. Parallel Paxos can use the same amount of nodes required by a centralized strategy.

Upon recovering from a failure, an application process must retrieve all its logged determinants. With the centralized approach, the application process contacts the event logger. With the replicated approaches, this is done by contacting a quorum of acceptors.

4 Evaluation

In this section we describe an implementation of the proposed consensus-based event logger and present experimental results, including a comparison with the traditional centralized alternative. Results are presented for the execution of three MPI applications: AMG 2013 (Algebraic Multigrid Solver) of the Lawrence Livermore National Laboratory, LU (Lower-Upper Gauss-Seidel solver) and MG (Multi-Grid on a sequence of meshes). Both LU and MG are on the NAS parallel benchmarks version 3.2. The applications were executed through Open MPI version 1.10. The experiments were conducted on a dedicated cluster that consists of 40 nodes each with two Intel(R) Quad-Core Xeon L5420 2.5 GHz processors and 8 Gbytes of RAM interconnected on a Gigabit Ethernet network.

We intercept MPI primitives using the MPI standard profiling interface (PMPI) [15]. If the interceptor detects a nondeterministic event, as defined in [2], it builds a determinant related to the event and makes a submission to the event logger. We implemented three event loggers: a traditional centralized logger, a distributed replicated logger based on Classic Paxos, and the distributed replicated logger based on Parallel Paxos (all described in Sect. 3.3). The Parallel Paxos event logger can be configured to log messages synchronously or asynchronously. In the synchronous mode, after submitting a determinant, an application process waits for an acknowledgement from the event logger before submitting the next determinant; in asynchronous mode the application process can submit multiple determinants before it receives acknowledgments from the event logger. Unless stated otherwise, our experiments use the synchronous mode. The interceptor and event loggers were implemented in C using the libevent version 2.0.22. We used the Paxos library libpaxos version 3.

The centralized event logger is hosted on a dedicated node. In Classic Paxos, a coordinator was deployed on a dedicated node while three acceptors (i.e., $f = 1$) were deployed each on a single node. There were also three learners, each one colocated with an acceptor. The learners are responsible for replying to the MPI processes as soon as a determinant is stored. In Parallel Paxos, each sequence of Paxos executions uses three acceptors (i.e., $f = 1$), each one deployed on a dedicated node. In Parallel Paxos each MPI process is both a proposer and a learner. Acceptors can be configured to store commands (i.e., determinants) on disk or in memory. In all experiments, the centralized event logger and the acceptors log values in main memory. We justify this choice by the fact that persistent memory technologies such as non-volatile RAM (NVRAM) and battery-backed memory are increasingly popular.

4.1 The Event Logger

To evaluate the performance of the event logger alone, we built a simple MPI application where a process only submits a new determinant to the logger after it receives a response acknowledging that the previously submitted determinant has been logged. Since application processes do not communicate among themselves in these experiments, determinants are fixed-content 50-byte messages.

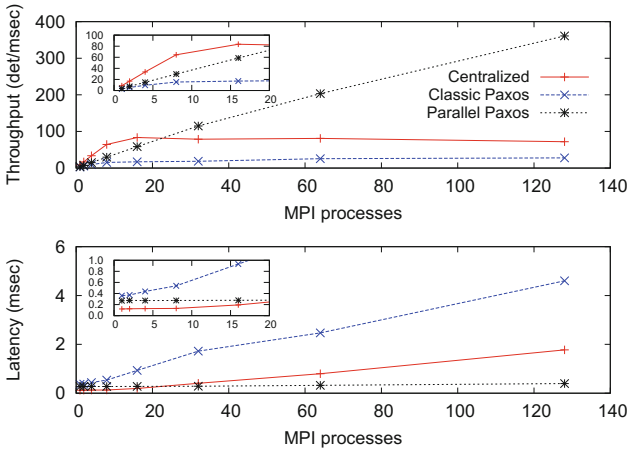


Fig. 2. Throughput and latency for the three event logger approaches.

Figure 2 shows the throughput in logged determinants per millisecond (det/msec) and latency (in msec) when we increased the number of MPI processes up to 128. The centralized event logger reaches the maximum throughput of about 83 det/msec with 16 processes. Classic Paxos reaches a maximum throughput of 28 det/msec with a latency of 4.6 msec with 128 MPI processes. Parallel Paxos never saturates in these experiments: throughput increases proportionally to the number of processes and the latency remains approximately constant, below 4 msec. With 128 processes, Parallel Paxos has 5 times the throughput of the centralized scheme and 13 times the throughput of Classic Paxos, with much lower latency.

4.2 AMG

AMG is a parallel algebraic multigrid solver for linear systems which can be classified as a nondeterministic application that employs “any-source” receptions and nondeterministic deliveries. All calls to `Iprobe` use the `any_source` tag and only one call to `Recv`, among many, uses the `any_source` tag. AMG also has calls to the `Test` and `Testall` primitives. Although during the execution there was a large number of `Iprobe`, `Test` and `Testall` invocations, a determinant for an

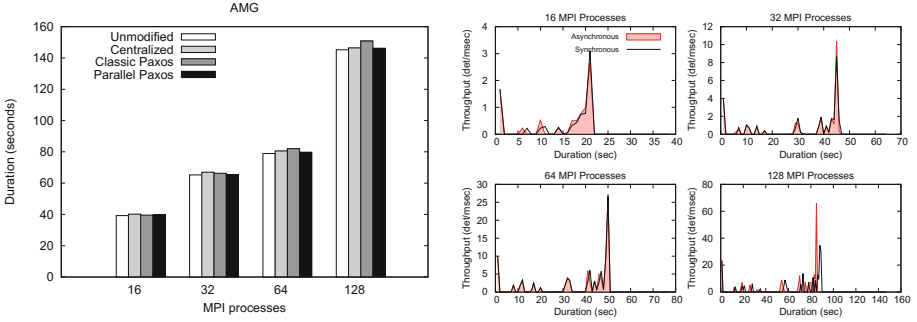


Fig. 3. AMG performance and throughput.




Fig. 4. LU class C performance and throughput.

`Iprobe` with the `any_source` tag is only created when the message is ready to be received. Similarly, for `Test` and `Testall` invocations, we count the number of invocations but only submit the determinant to the event logger when the MPI message related to `Test` or `Testall` is ready to be delivered.

Figure 3 (left) presents results for the AMG application, including the replicated event loggers. The “Unmodified” label refers to executions without any event logging. Although all strategies introduce an overhead, with 64 and 128 processes Classic Paxos increased the duration of the original application by approximately 3.8%, the worst-performing technique. The centralized scheme presented an overhead of approximately 2% for 16, 32 and 64 processes. Parallel Paxos presented the lowest overhead, below 1.3% for all configurations.

Figure 3 (right) shows the number of determinants logged per milliseconds considering both the synchronous and asynchronous Parallel Paxos modes, for 16, 32, 64 and 128 MPI processes. In the asynchronous mode, application processes never wait for the logger; thus, this case provides an upper bound for the performance of the event logger. As it can be seen, log requests are not uniformly distributed over time. For the case with 128 processes, between time instants 80s and 90s a peak can be distinguished, that reaches approximately 34 det/msec in the synchronous mode and 66 det/msec in the asynchro-

nous mode. These results help understand how the overhead of Parallel Paxos is distributed over time.

4.3 LU and MG

LU and MG are kernels of the NAS parallel benchmarks. As pointed out in [2], only MG and LU among the NAS-PB kernels generate nondeterministic events. We assess classes C and D of the kernels in deployments with 16, 32, 64 and 128 MPI processes. LU contains both `Recv` and `Irecv` primitives. The last one is used with the `any_source` tag. MG receives all its messages through `MPI_Irecv` with the `any_source` tag.

The total number of events logged in LU is less than 1% of the total of all its receptions in class C. The MG kernel however has almost 100% of nondeterministic events among its receptions. Although both AMG e MG solve similar problems, the reason for much more nondeterministic events in MG is its implementation. Unlike MG, AMG does not receive all its messages through `MPI_Irecv` with the `any_source` tag. This illustrates the fact that nondeterminism is often a programmer's choice (e.g., to boost performance), rather than a requirement coming from the problem being solved.

From our experimental evaluation, we concluded that logging determinants using any of the three event logging strategies presents nearly no overhead when logging events of classes C and D of the LU kernel. This is somewhat surprising since logging introduces some overhead in the AMG application and both classes C and D of the LU benchmark contain a higher percentage of nondeterministic events than AMG contains. Figure 4 (left) shows the results for class C of the LU kernel. By inspecting the number of determinants logged during the execution in Fig. 4 (right), we notice that determinants are more uniformly distributed in LU class C than in AMG and they happen at a rate that is within the limits the event logger can sustain (see Sect. 4.1). The difference between class C and D of the LU is that the last one has longer duration and lower throughput. As a consequence, the event logger never becomes an execution bottleneck in LU.

On the contrary, the logging of determinants introduces a considerable overhead to MG classes C and D (Figs. 5 and 6, respectively). In class C, while Classic Paxos presents an overhead of more than 125% and 200% for 64 and 128 processes, the overhead of the centralized event logger is below 31% and 55% for 64 and 128 processes, respectively. Parallel Paxos sports even lower overheads: 17,71% and 24,26% for 64 and 128 processes, respectively. The results for MG class D show a similar trend, with Parallel Paxos outperforming both the two other techniques. The MG kernel is highly communication-bound and all its receive events use the `any_source` tag. As the number of application processes increases, the event logger reaches its limits with the centralized and the Classic Paxos strategies. Parallel Paxos is able to scale performance by distributing the load among the various series of Paxos.

Figures 5 and 6 also show the rate of logged determinants per milliseconds for the synchronous and asynchronous Parallel Paxos-based event logger. The throughput of the synchronous mode is close to the asynchronous mode for MG

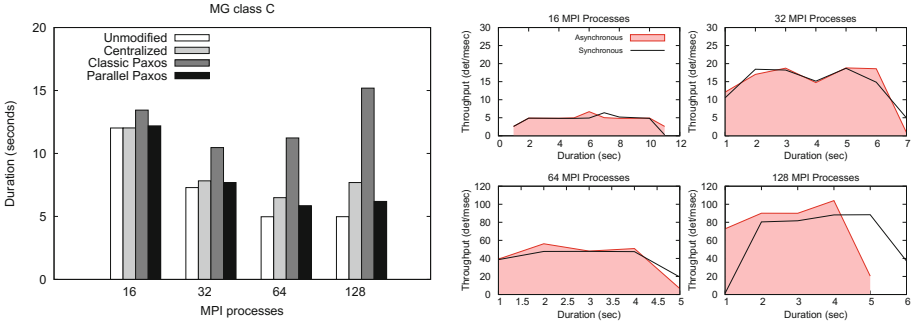


Fig. 5. MG class C performance and throughput.

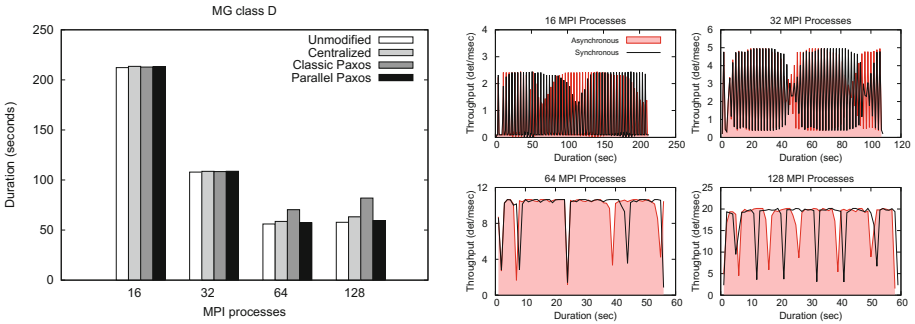


Fig. 6. MG class D performance and throughput.

class C with 16, 32 and 64 processes. For 128 processes, the asynchronous mode presents a throughput that is higher than that of the synchronous mode and finishes approximately 1 s earlier. MG class D has lower throughput than MG class C. The throughputs of both synchronous and asynchronous Parallel Paxos are very similar. In all configurations, both the synchronous and the asynchronous modes display a uniform rate over time.

5 Conclusion

In this work we presented a fault-tolerant and distributed event logger based on consensus for HPC applications. The event logger is the component responsible for reliably logging determinants and its performance can represent a significant impact on the efficiency of message logging protocols. We implemented two fault-tolerant event loggers based on the Paxos algorithm. By using Paxos, our event loggers guarantee safety even if the system is asynchronous and liveness despite processes failures. Our first protocol is based on classic state machine replication. In our second protocol, which we call Parallel Paxos, we assign a separate sequence of Paxos executions to each application process. We assessed

experimentally the performance of a centralized event logger and our two event loggers based on consensus. Besides evaluating the event loggers by themselves, we used three MPI applications to evaluate their performance: AMG, MG and LU. Results of all experiments show that the event logger based on Parallel Paxos always outperformed the centralized approach in terms of both the execution time and the throughput in terms of the number of determinants logged per millisecond.

The implementation of a recovery protocol using the nondeterministic events stored on the Parallel Paxos event logger is left as future work.

References

1. Bouteiller, A., Collin, B., Herault, T., Lemarinier, P., Cappello, F.: Impact of event logger on causal message logging protocols for fault tolerant MPI. In: IPDPS (2005)
2. Bouteiller, A., Bosilca, G., Dongarra, J.: Redesigning the message logging model for high performance. *Concurr. Comput.: Pract. Exp.* **22**(16), 2196–2211 (2010)
3. Bouteiller, A., Héroult, T., Krawezik, G., Lemarinier, P., Cappello, F.: MPICH-V project: a multiprotocol automatic fault-tolerant MPI. *Int. J. HPC Appl.* **20**(3), 319–333 (2006)
4. Bouteiller, A., Ropars, T., Bosilca, G., Morin, C., Dongarra, J.: Reasons for a pessimistic or optimistic message logging protocol in MPI uncoordinated failure, recovery. In: *Cluster* (2009)
5. Cappello, F., Guermouche, A., Snir, M.: On communication determinism in parallel HPC applications. In: *ICCCN* (2010)
6. Egwuotuoha, I.P., Levy, D., Selic, B., Chen, S.: A survey of fault tolerance mechanisms and checkpoint/restart implementations for HPC systems. *J. Supercomput.* **65**(3), 1302–1326 (2013)
7. Elnozahy, A., Wang, J.: A survey of rollback-recovery protocols in message-passing systems. *CSURV Comput. Surv.* **34**, 375–408 (2002)
8. Fischer, M.J., Lynch, N.A., Paterson, M.S.: Impossibility of distributed consensus with one faulty processor. *J. ACM* **32**(2), 374–382 (1985)
9. Johnson, D.B., Zwaenepoel, W.: Sender-based message logging. In: *FTCS* (1987)
10. Kshemkalyani, A.D., Singhal, M.: *Distributed Computing: Principles, Algorithms, and Systems*. Cambridge University Press, Cambridge (2011)
11. Lamport, L.: Paxos made simple. *SIGACTN: SIGACT News (ACM Spec. Interest Group Automata Comput. Theory)* **32**, 51–58 (2001)
12. Lamport, L.: The part-time parliament. *ACM Trans. Comput. Syst.* **16**(2), 133–169 (1998)
13. Lamport, L.: Lower bounds for asynchronous consensus. *Distrib. Comput.* **19**(2), 104–125 (2006)
14. Lemarinier, P., Bouteiller, A., Krawezik, G., Cappello, F.: Coordinated checkpoint versus message log for fault tolerant MPI. *Int. J. High Perform. Comput. Netw.* **2**, 146–155 (2006)
15. MPI Forum: document for a standard message-passing interface 3.1. Technical report, University of Tennessee (2015). <http://www.mpi-forum.org/docs/mpi-3.1>
16. Riesen, R., Ferreira, K., Silva, D.D., Lemarinier, P., Arnold, D., Bridges, P.G.: Alleviating scalability issues of checkpointing protocols. In: *SC* (2012)

17. Ropars, T., Morin, C.: Active optimistic message logging for reliable execution of MPI applications. In: Sips, H., Epema, D., Lin, H.-X. (eds.) Euro-Par 2009. LNCS, vol. 5704, pp. 615–626. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-03869-3_58](https://doi.org/10.1007/978-3-642-03869-3_58)
18. Ropars, T., Morin, C.: Improving message logging protocols scalability through distributed event logging. In: D’Ambra, P., Guarracino, M., Talia, D. (eds.) Euro-Par 2010. LNCS, vol. 6271, pp. 511–522. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-15277-1_49](https://doi.org/10.1007/978-3-642-15277-1_49)
19. Schneider, F.B.: Implementing fault-tolerant services using the state machine approach: a tutorial. *ACM Comput. Surv.* **22**(3), 299 (1990)