**QUEEN'S
UNIVERSITY
BELFAST**

Implementing Efficient Message Logging Protocols as MPI Application Extensions

Dichev, K., & Nikolopoulos, D. S. (2019). Implementing Efficient Message Logging Protocols as MPI Application Extensions. In *EuroMPI 2019 10/09/2019 → 13/09/2019 Zurich, Switzerland* [8] ACM.
<https://doi.org/10.1145/3343211.3343219>

Published in:

EuroMPI 2019 10/09/2019 → 13/09/2019 Zurich, Switzerland

Document Version:

Peer reviewed version

Queen's University Belfast - Research Portal:

[Link to publication record in Queen's University Belfast Research Portal](#)

Publisher rights

Copyright 2019 ACM. This work is made available online in accordance with the publisher's policies. Please refer to any applicable terms of use of the publisher.

General rights

Copyright for the publications made accessible via the Queen's University Belfast Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The Research Portal is Queen's institutional repository that provides access to Queen's research output. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact openaccess@qub.ac.uk.

Implementing Efficient Message Logging Protocols as MPI Application Extensions

Kiril Dichev
Queen's University Belfast
Belfast, United Kingdom
K.Dichev@qub.ac.uk

Dimitrios S. Nikolopoulos
Queen's University Belfast
Belfast, United Kingdom
D.Nikolopoulos@qub.ac.uk

ABSTRACT

Message logging protocols are enablers of local rollback, a more efficient alternative to global rollback, for fault tolerant MPI applications. Until now, message logging MPI implementations have incurred the overheads of a redesign and redeployment of an MPI library, as well as continued performance penalties across various kernels. Successful research efforts for message logging implementations do exist, but not a single one of them can be easily deployed today by more than a few experts. In contrast, in this work we build efficient message logging capabilities on top of an MPI library with no message logging capabilities; we do so for two different HPC kernels, one with a global exchange pattern (CG), and one with a neighbourhood exchange pattern (LULESH). While our library of choice ULFM detects failure and recovers MPI communicators, we build on that to then restore the intra- and inter-process data consistency of both applications. This task turns out to be challenging, and we present the methodology for doing so in this work. In the end, we achieve message logging capabilities for each kernel, without the need for an actual message logging runtime underneath. On the performance side, we match state-of-the-art solutions and (a) eliminate event logging and the event logger component altogether, and (b) design a hybrid protocol, which gracefully shifts between global and local rollback, depending on the available payload logging memory. Such a hybrid protocol between local and global rollback has not been previously proposed to our knowledge. Our extensions span a few hundred lines of code for each kernel, are open-sourced, and enable local and global rollback after process failure.

1 INTRODUCTION

It is widely accepted that compute clusters and supercomputers are transitioning towards systems of millions of compute units to satisfy the requirements of compute-intensive parallel scientific applications. With this increase in compute components, a proportional decrease in the Mean-Time-Between-Failure (MTBF) across parallel executions will follow [21, 24], which would make highly scalable parallel application runs infeasible without integrating resilience.

In this manuscript, we focus on recovery from fail-stop errors, i.e. any failures leading to the unexpected termination of an MPI process and the loss of its data; a node crash is among the possible causes of fail-stop errors. For such failures, checkpoint/restart (C/R) strategies are commonly used; they introduce time redundancy due to the rollback of execution but require fewer additional resources than resource replication techniques. The recent advances in fault-tolerant MPI library implementations, such as ULFM [6],

have integrated efficient and scalable detection and recovery primitives to allow MPI applications to deal with failures. C/R with global rollback is the most widely used fault tolerance technique in HPC applications today; in global rollback, all processes roll back to the last globally consistent checkpoint, and continue execution. Global rollback can be universally applied to all types of application kernels, and its significant advantages in failure-prone executions are widely known and accepted.

In order to objectively argue which aspect of recovery should be optimised, quantitative data should be analysed. A large-scale study of the highly optimised S3D application on the Titan supercomputer from Gamell et al. [14] shows that the rollback phase is the main culprit. Therefore, optimisations reducing the amount of rollback need to be studied in the area of fault tolerance research.

There are multiple ways to optimise rollback after failure. One possibility, which has seen considerable attention in recent years, is the classic direction of fault tolerance studies – checkpoint/restart. Optimised checkpoint/restarts naturally lead to reduced rollback, even if all processes need to roll back.

An entirely orthogonal idea is to allow surviving processes to perform as little recompute as possible, by only rolling back failed processes after failure. This is the realm of local rollback recovery, and of message logging protocols, the enablers of local rollback. Upon failure, message logs can be replayed to a failed process, and if correctly designed, message logging can allow survivors to not roll back, a potential advantage over all traditional global rollback strategies. Various message logging protocols [4], exist, such as pessimistic, optimistic, and causal message logging. A large body of research exists, particularly with the MPICH-V and MPICH-V2 projects, which studied various runtime designs and protocols for message logging. Unfortunately, various kernels consistently show large runtime overheads even for well-designed message logging runtimes. In addition, both of these projects have disappeared as research efforts. More recent MPI implementations also struggle to provide easy to deploy message logging support; e.g. Charm++ has abandoned such support in their latest stable releases. The most up-to-date message logging protocol can be found in the pessimistic VProtocol within the Open MPI codebase; however, the support is only partial (payload logging), and just a stepping stone to a complete message logging runtime, and needs additional features to be made functional, not all of which are freely available and easy to use [18]. Today, not a single easy-to-use MPI implementation with message logging capabilities exists.

In this work, we challenge this landscape and implement a message logging protocol as an efficient and configurable extension of HPC application kernels. We describe the various issues we faced and how we resolved them. We rely on ULFM, the de-facto

fault-tolerant MPI implementation, for capabilities such as failure detection and communicator recovery. No message logging MPI library, or other components, such as external payload or event loggers, are required. Our HPC kernels of choice are the MPI version of the NAS CG benchmark, and the larger proxy app LULESH from LLNL. The former has always been among the most challenging kernels due to its global and intensive communication, whereas the latter represents a different class of stencil-like codes with neighbourhood exchange.

After implementing message logging as an extension to applications, we also achieve some interesting optimisations, such as (a) no event logging or event loggers for these kernels (b) configurable payload logging, since payloads can be conceived as fine-grained application checkpoints.

The novel contributions of this work are:

- A methodology for implementing message logging protocols as compact extensions of HPC kernels, exemplified for two HPC kernels. The methodology includes:
 - The concept of iterations as transactions, which helps us resolve intra-process data consistency issues, as well as eliminate event logging
 - The snapshot of the computational wave after failure, which can be used to replay messages and restore inter-process data consistency, on top of the intra-process consistency.
- Due to the scope of payload logging within the application, a hybrid scheme with graceful degradation between local and global rollback is introduced, depending on available payload capacity.

The implemented codes are open-sourced, and rely on other open-sourced codes, such as NPB, LULESH, and ULFM.

The paper is organised as follows: the related work in message logging is presented in Sect. 2. In Sect. 3, we introduce the concept of transactional MPI iterations, which allows us to restore intra-process data consistency. We then focus on inter-process data consistency in Sect. 4. We continue with the technical aspects, such as the send wrapper (Sect. 5) and replay routine (Sect. 6) implementations. We briefly summarise the extensions required for CG and LULESH in Sect. 7. Our experiments are shown in Sect. 8, after which we conclude the paper (Sect. 9).

2 RELATED WORK

The idea of implementing message logging protocols in application kernels is not entirely new. Gamell et al. [13] outlined the use of sender-based logs for a stencil code called S3D, integrating the rollback on top of ULFM; while the authors never open-source their work, we clearly outline the design we use for two kernels. We suspect that LULESH has some similarities in its neighbourhood exchange patterns to S3D; However, the NAS CG kernel is more challenging with its global communication patterns, such as the allreduce, and possibly better placed to stress message logging implementations.

Within MPI runtimes, message logging protocols are a widely studied area of fault tolerance. All of them are generally applicable to various kernels, whereas our contributions are applied per HPC application as kernel extensions. The seminal paper of Strom [22]

first introduced optimistic message logging protocols; it also formalised orphans and determinants. Researchers have proposed an optimised version of optimistic message logging [11, 20], which reduces the extent of piggybacked determinants, due to the deterministic nature of communication events. Implementations of causal message logging protocols in MPI libraries include MPICH-V [7], MPICH-V2 [10], and Charm++ [19]. The work of Losada et al. [18] implements the most recent variation on pessimistic message logging protocols; it employs the VProtocol in Open MPI¹, and a source-to-source compiler called CPPC. Our work, in contrast to pessimistic, optimistic, and causal message logging protocols, entirely disables event logging; we statically analyse the underlying applications, and guarantee the absence of non-deterministic events that need to be replayed. In this respect, the only contribution that performs a similar optimisation is of Bouteiller et al. [9]. Their pessimistic protocol is highly optimised to reduce event logging to challenging non-deterministic events; for example, any-sender non-blocking receives are such events. The authors log no events for the NAS CG benchmark, similarly to our work, even though using a runtime mechanism, rather than a kernel extension.

In terms of exploring determinism of HPC codes, interesting contributions exist, which similarly to our work look for optimisations based on the absence of non-determinism. Guermouche et al. [15] also carefully examines the communication properties of various kernels, and argue that most codes are send-deterministic, which allows for more efficient protocols to be implemented; the focus is specifically on causal message logging protocols. In a rather complex runtime extension, the authors manage to reduce the number of logged messages, while still preserving the causality of their message logging algorithm, for send-deterministic applications. However, to reduce the message logs, the protocol introduces some overhead to messages, such as piggybacked data (e.g. epoch numbers); often, a large set of processes needs to roll back, while in our proposal only the failed process rolls back. There are some similarities to our optimisations in these related works; we similarly perform “static analysis of applications consisting of looking at each communication pattern and studying its deterministic nature”. However, there are numerous differences in the underlying techniques: among others, we work in the context of the application, and not the MPI implementation.

In terms of payload logging, and optimisations for that, Losada et al. cap the message logs by observing the checkpoint intervals, and discard older logs. They enable this optimisation via application instrumentation with a source-to-source compiler (CPPC). We also cap the payload logs to be contained within a checkpoint interval. However, we do not stop there – the capping of payload logging may continue gradually, allowing for a graceful degradation between global and local rollback. To the best of our knowledge, no message logging protocol has ever implemented such a hybrid solution.

We also give some background to the use of the NAS CG benchmark as a challenging problem for message logging protocols here. In a study for MPICH-V, Bouteiller et al. [8] examine the NAS benchmarks, CG, BT, LU, and FT; the authors find that LU and CG (in that order) cause the most memory and compute overheads for an implementation of a causal message logging protocol. The

¹not freely available at the time of writing

	[9]	[19]	[11]	this work
protocol	pessimistic	causal	optimistic	event-less
implemented in	Open MPI	Charm++	Open MPI	CG and LULESH (extended)
event logging required	not for CG	yes	yes	no
message logging required	yes	yes	yes	yes
requires tracking determinants	no	yes	yes	no
optimised determinant protocol?	–	yes	yes	–
applicability	general	general	general	per kernel
overheads	≈ 5% for CG	≈ 1% for CG	≈ 5% for CG	fully configurable
main culprit	payloads	not detailed	payloads	payloads

Table 1: A table of related work exploring message logging. We also compare the related work to ours in various important aspects, including NAS CG results where available.

same researchers [9] later introduce an optimised pessimistic message logging protocol, which among others reduces overhead for kernels with non-deterministic patterns, such as LU. After optimisation, they observe the highest runtime overhead with NAS CG, at 5%, for 64 node runs. More recently, Meneses et al. [19] study the benchmark CG, MG, BT, and DT from the same suite, with implementation of pessimistic and causal message logging protocols, and find that pessimistic message logging CG shows 17% runtime overheads, but that a causal message logging shows significant improvements, with around 1% runtime overheads for the CG kernel.

As a summary and comparison with closely related work, we provide Table 1; whenever possible, we provide some reference numbers for CG; such numbers do not exist in these contributions for LULESH.

In terms of availability, MPICH-V and MPICH-V2 are deprecated and not available, and the message logging capabilities in Charm++ are not functional in recent stable releases. The only currently usable message logging protocol, the pessimist Vprotocol in Open MPI, is in an experimental stage, and not ready for use out-of-the-box; recently proposed optimised protocols are not currently available for testing [18]. We believe this landscape clearly motivates our work to implement message logging capabilities as application extensions on top of popular and widely-used MPI implementations. In this way, compact implementations of message logging as kernel extensions can be tested or even ported from our open-sourced work.

Figure 1: Illustration of ensuring the transactional nature of an iteration i : until the last MPI op in iteration i , a rollback is possible, therefore operations need to be idempotent. After the last MPI op call within an iteration, a commit is guaranteed, and operations may change global data irreversibly.

3 INTRA-PROCESS DATA CONSISTENCY

In this section, we guarantee data consistency within a process, even after an iteration is rolled back; we deal with inter-process data consistency in a later section.

3.1 Transactional MPI Iterations

To guarantee data consistency of each process after failure, we take the view that an iteration can be designed as an atomic transaction. The fundamental property of a transaction is that it can either be committed or aborted; if aborted, it leaves the application in its pre-transaction state. Like any MPI code, each iteration of an HPC kernel consists of two types of operation – local operations, or non-local operations (MPI calls). For an iteration to have transactional nature, we must be able to roll it back whenever a failure is detected. In fault-tolerant MPI implementations, only in MPI functions can report the failure of another process. In effect, this reduces the problem to the illustration of Fig. 1. All phases including the very last MPI operations in an iteration (in yellow) may be repeated. Therefore, all operations need to be idempotent. If, in contrast, some operation is non-idempotent, the rollback cannot happen without further action. For example, if one of the local operations (in blue) preceding the last MPI operation is an increment on global data, the iteration is not transactional in nature. A repetition of iteration i will lead to incorrect global data, since the increment of global data would be performed twice. A kernel iteration is transactional in nature if it performs all non-idempotent operations only after the very last MPI call in an iteration.

We hold that all non-transactional iterations can be made transactional, with some memory duplication, and in some cases added copy operations. For example, the illustrated code in Alg. 1 is non-transactional: an increment on global data is a non-idempotent operation, and it happens before an MPI call. If we perform a rollback of this iteration, not reading the last checkpoint, this may result in x being incremented twice for the same iteration. The solution to this problem is outlined in Alg. 2, and consists of guaranteeing that no non-idempotent operations are called *until after the last MPI call* in the iteration. If a failure is detected during any MPI call, the transaction (or iteration) can be rolled back with no further action, and without undesired side effects on global data, since all operations are idempotent. After this call, the occurrence of a failure cannot be detected in the same iteration anymore, i.e. the iteration would be “committed”.

Algorithm 1 Illustration of a non-transactional MPI iteration

```
// non-transactional
for (int i=0; i<NITER; i++) {
    x++; // non-idempotent, re-run bad!
    MPI_Sendrecv(&x,...,&y,...);
    // further local operations
    ...
}
```

Algorithm 2 Illustration of a transactional MPI iteration

```
// transactional
for (int i=0; i<NITER; i++) {
    x_tmp = x;
    x_tmp++;
    MPI_Sendrecv(&x_tmp,...,&y_tmp,...);
    // non-idempotent, but will never be reversed
    x = x_tmp;
    y = y_tmp;
    // further local operations
    ...
}
```

Figure 2: Different phases of NAS CG and LULESH iterations. Phases in blue and brown are local updates. Phases in yellow are MPI communication events, and detectors of failure. We implement the copy to/from temp buffers (brown) as enablers of transactional iterations. Phases in red provide recovery of communicators (ULFM), as well as local rollback via our replay routine.

3.2 CG and LULESH iterations as transactions

Consider the local operations and non-local operations illustrated in Fig. 2 for the NAS CG and LULESH benchmarks. For each kernel, local operations are outlined in blue and brown, and non-local MPI calls, are outlined in yellow. In their original implementations, both CG and LULESH are not transactional in nature; they write into global data before their last MPI operations in each iteration, that is both kernels have non-idempotent operations which means that any rollback will break the data consistency of the application, unless code modifications similar to these shown in Alg. 2 are made. We did this for both kernels, and these modifications are highlighted in brown in the figure. The recovery of failed processes includes a sequence of stages, highlighted as red boxes in the figure; some of them must be performed by a fault-tolerant MPI implementation (such as ULFM). Others include our extensions to restore data consistency, and to provide local rollback.

- the MPI implementation must have a failure detection mechanism. The failure detection can only be triggered during MPI calls, when a surviving MPI process communicates with either a failed process, or a process which has already revoked the communicator.
- the developer can register an error handler routine to handle failures. The routine generally has 2 stages: (a) recovery of the MPI communicator (MPI-specific) and (b) recovery of the application state (application-specific).

Our work is concerned with restoring data consistency within a process, followed by global data consistency between MPI processes as well. We note here that we strictly require an MPI implementation to provide primitives for failure detection and communicator recovery; without these primitives, our contribution would fail. This requirement is rather generic, and is true also for automated global rollback; MPI communicator recovery is essential if any rollback is dynamic rather than triggered externally. Our solution is embedded in the application, and no message logging capabilities are required of the MPI runtime, while local rollback is still achieved.

It is worth noting that transactional logic of iterations was initially conceived for practical reasons; upon process failure, an error handler interrupts the control flow of the surviving processes. After the successful MPI communicator recovery, we need to pass control back to the application kernel – but where? If we interpret an iteration as a transaction, our task is easy – we roll back to the start of the current iteration. Note that this control flow after failure differs from its counterpart with message logging protocols in runtimes. For example, pessimistic protocols synchronously log events before a message is passed, so no rollback of survivors to earlier program stages is required, or even possible.

A failed process is always restarted from the last global checkpoint.

3.3 Why we need no event logging/logger

Event logging has always been an integral part of any message logging protocol within MPI. In the context of MPI applications, every send and receive is usually logged, since the runtime has no concept of the application state, and needs to recover each communication in case of a failure. Event logging usually incurs an overhead proportional to the latency on the underlying network, since event loggers

are quite often placed on different nodes from the application processes nodes, and the popular family of pessimistic message logging protocols requires every event to be synchronously logged before progress. There have been some optimisations in event logging in the past. For example, optimistic message logging has previously removed the requirement of synchronously logging events, masking some of the incurred overheads. Another interesting optimisation [9] is the ability of an optimised implementation to only log challenging events (such as any-sender receives), which in effect result in similar disabling of event logging as our work for some kernels, e.g. the NAS CG benchmark.

In this work, we take a substantially different view to events. Since we implement our protocol as extension of application kernels, it is not MPI communication events that are key, but application events that irreversibly modify the data of a process. As a consequence of this viewpoint, we can try to reduce the number of these irreversible events. This view in itself does not eliminate event logging. Consider, for example, the outline of the different phases of the LULESH mini-app of Fig. 2. Every yellow phase entails numerous non-blocking send/rcv calls (exchanges with neighbours in three dimensions), which normally would need to be replayed in the same order in a log-enabled recovery. However, by conceiving each iteration as a *transaction*, we can abort entire iterations, and eliminate event logging altogether. During recovery each application only needs to provide its current iteration. The simplification is significant; since each process can safely ignore the exact phase within its iteration where a failure was detected.

As a consequence of disabling event logging or event loggers, we do not carry the runtime overheads incurred in MPI implementations. In contrast, runtime-based logging employs event loggers as remote processes, including long-existing efforts such as MPICH-V, or newer implementations such as Open MPI's VProtocol. We still need payload logging, as an extension of each MPI process.

We note that our design is currently applied with some assumptions on the underlying kernel in mind:

- a kernel is deterministic (but not necessarily send-deterministic), which includes every HPC kernel we are aware of (see [15] for overview)
- a coordinated checkpointing scheme, possibly asynchronous, is in place, for the given kernel.

The fact that we require an explicitly or implicitly coordinated checkpointing scheme carries important simplifications for us, e.g. the fact that orphan or in-transit messages can never occur.

4 INTER-PROCESS DATA CONSISTENCY AFTER FAILURE

Despite the fact that the transactional nature of iterations now guarantees that the data of each process remains consistent after failure, inter-process data consistency is an unresolved issue. Unfortunately, it is a very likely outcome that MPI processes detect failure in different iterations, and therefore the application will have an inconsistent global data state.

Following scenarios are all likely, even in tightly-coupled HPC codes:

- Survivors are very likely to be in different iterations after failure detection.

Figure 3: Illustration of partitioning a matrix for NAS CG with 16 MPI processes; each process has 2 reduce partners (in black) and 1 transpose partner (red).

Figure 4: Illustration of a snapshot of a real failure scenario for cg.C.16 (16 processes). Process 2 fails in iteration 13, and is restarted at the last checkpoint. Not all survivors detect the failure in the same iteration.

- Restarted process and survivors are almost certainly in different iterations upon failure detection.

We will detail the issues and their solutions in this section.

4.1 Survivors – different iterations

An important issue for most MPI kernels, even when they have desirable properties such as determinism, is the non-transactional nature of MPI point-to-point communication calls. The completion of an MPI point-to-point call, for example, is not a guarantee of a delivered message; the most obvious example is the MPI_Send call, which can be buffered locally (for example, small-size message and the internal use of eager vs rendezvous protocols). A matching receive may still be unsuccessful due to a failure of the receiving process, leading to inconsistencies in the global state. This non-transactional nature of MPI point-to-point communication requires us to carefully consider how to recover (see also e.g. [12, 17, 18]).

Consider a 16 process run with CG (cg.C.16), where a square matrix is partitioned among processes as illustrated in Fig. 3. There are 2 reduce partners (black edges), and one transpose partner (red edges), for each process.

A snapshot of an actual failure scenario for this setting is illustrated in Fig. 4. All ranks write their part of a global checkpoint

every 10 iterations, the most recent checkpoint being at the end of iteration 10. Rank 2 crashes unexpectedly in the middle of iteration 13, just after the exchange with the transpose partner 8. Each process can only notice the failure in an MPI call (phases highlighted in yellow for CG and LULESH in Fig. 2). Rank 0 and/or 3 will immediately notice the failure of rank 2, since they cannot possibly receive data they need. The communicator will be revoked. Yet it is not necessary that all MPI processes notice the revoked communicator and trigger recovery in the same iteration. In a typical scenario, processes {0,1,3,12,13,14,15} are notified in the same iteration 13, being direct or “close” communication partners with 2 after its failure. However, processes {4,...,11} are notified while already in the following iteration 14 – their completion of iteration 13 was successful, having received their required messages from their communication partners.

Due to the outlined non-transactional nature of MPI P2P calls, a number of surviving processes are in iteration 14 and others in iteration 13. All processes abort their current iteration, and restart it. The iteration-14 survivors are now responsible for replaying the message contents of iteration 13 to all survivors left behind, in order to restore consistency.

4.2 Survivors and restarted process – different iterations

As illustrated in Fig. 4, process 2 fails at iteration 13, reloads from the last checkpoint it has written, and starts at iteration 11. The survivors have obviously progressed further, being in iterations 13 and 14.

To restore the consistency, process 2 would read the checkpoint of iteration 10, and iterate through til iteration 13, receiving replayed messages. However, process 2 must not send any messages to any peers til iteration 13, since these messages have already been successfully sent (otherwise survivors would not have progressed). Of the surviving processes, 0 and 3 are reduction peers of 2, and 8 is the transpose peer of 2. Therefore, peers 0, 3, and 8, all need to replay messages from iterations 11, 12, and 13 to rank 2.

5 THE SEND WRAPPER

The send wrapper needs to intercept each send call of the application; a receive equivalent is not required. Unfortunately, the send wrapper is not simply overwriting the MPI calls, and the use of e.g. the MPI profiling interface is not possible. The wrapper for each application needs to know the context and phase at which the sends were intercepted.

All of the send calls are intercepted by the wrapper, and the implementation is illustrated in Alg. 3. First, a check is done if the send is into the future ($peer_iters[me] < peer_iters[dest]$). In this case, a send is cancelled, since the communication partner has already received the message. This formulation generically applies to both restarted process, and some of the survivor processes, as illustrated in Fig. 4:

- a restarted process in previous iterations in regard to survivors does not need to send them a message.

- a surviving process in previous iterations in regard to other survivors, due to the inconsistencies caused by MPI’s non-transactional semantics, does not need to send them a message.

If the send is not into the future, the send operation may be a normal send ($peer_iters[me] == peer_iters[dest]$) or a replay ($peer_iters[me] > peer_iters[dest]$). For a normal send operation, we log the payload, unless we apply payload capping (discussed later in this section). If the send is a replay, we read the stored payload log, and send it as a replayed MPI message to a communication peer (either a restarted process, or a survivor left behind during recovery).

Algorithm 3 Outline of send wrapper routine. The send wrapper replaces all MPI send calls of the application, and is also called by the replay routine.

```
int send_wrapper(void * buf..., int current_iter) {
    // match send with the right iteration
    if (current_iter == peer_iters[dest]) {
        // regular send
        if (peer_iters[me] == peer_iters[dest])
            // sender-based payload logging
            // with potential capping capabilities
            append_log(buf, current_iter);
        MPI_Send(buf, ...);
        // replay message from buffer
    } else if (peer_iters[me] > peer_iters[dest]) {
        log_buf = get_log(current_iter, ...);
        MPI_Send(log_buf, ...);
        ...
    }
    else {return 0;} // filter out sends into future
}
else {return 0;} // send irrelevant to sender
}
```

The high-level logic of the send wrapper is quite generic; unfortunately, the payload logging, which depend on the application data, is application-dependent in each of our extensions.

Because all payload logging is managed by the `append_log` routine of Alg. 3, and within the application instead of the runtime, there are interesting optimisations which can be made to enable capping of logging message payloads. The payloads of each application are conceptually not different from mini-checkpoints, in contrast to a runtime which only sees them as raw data. These mini-checkpoints can be capped to save runtime and memory overheads; this, however, can limit the extent to which we can replay messages. Yet it never affects global rollback – we never compromise on checkpointing for global rollbacks. The capping is easy to implement; details can be found in the source code.

6 A CONDITIONAL REPLAY ROUTINE FOR A HYBRID ROLLBACK PROTOCOL

Upon failure detection, all processes synchronise to find everyone’s current iteration, which represents the front line of computation.

Fig. 4, with survivors and restarted process in different iterations, represents such a front line. The restarted process follows the normal application logic, having started from the last checkpoint iteration, with one notable exception – all sends of the restarted process into the future are filtered out by the send wrapper of Alg. 3.

The recovery routine, which enables local rollback for all survivors, is at the core of our work, and illustrated in Alg. 4. The routine enables a hybrid protocol between local and global rollback, which is the task of survivor processes.

Algorithm 4 Outline of replay routine, called by all processes after a failure. All processes broadcast the front line of computation. If the required payloads are available, only the restarted process rolls back; otherwise all processes roll back.

```

void replay() {
    // gather the current iteration of each peer
    MPI_Allgather(peer_iters, ...);
    if (me == failed) { // restarted
        read_checkpoint();
    }
    else { // survivor
        maxit = max(peer_iters);
        minit = min(peer_iters);
        // do we have the required payloads
        // stored for local rollback?
        bool local = (maxit % CP_INT < LOG_SIZE);
        // survivors don't roll back, but replay messages
        if (local) {
            for (int i=minit; i++; i<maxit) {
                // replicate send sequence of kernel
                // note: wrapper will use stored payloads
                send_wrapper(NULL, ..., i);
                ...
            }
        }
        else { // global rollback necessary
            read_checkpoint();
        }
    }
}

```

If a log buffer is large enough to provide replay messages for the latest failure, a local rollback is enabled. The local rollback is enabled as a sequence of replayed messages by all survivors. Note that replayed messages are needed for restarted processes as well as survivors stuck in past iterations. Since the purpose of a replay is for each process to send messages into the past, and since each process gets the global front line of computation at the start, no distinction between failed and surviving processes is required. Once a process completes this process of supporting everyone left behind, it can safely resume with its normal iterations. Note that the process supporting others may be in the past in relation to other processes after the replay routine, but it has successfully replayed all its messages at the end of the routine.

Otherwise, if the log buffer is not large enough to provide replay messages reaching into the most advanced current iteration, all

Figure 5: Illustration of capping the payload buffer, and the resulting hybrid protocol between local and global rollback: failures where all survivors are within the green zone allow local rollback; failures where any survivors are in the red zone require global rollback.

survivors roll back together with the restarted process, resorting to global rollback.

This flexibility is illustrated in Fig. 5, which simply uses a checkpoint interval of 4 ($CP_INT = 4$), and a payload buffer limited to 2 iterations ($LOG_SIZE = 2$). The limited payload buffer reduced all memory and runtime overheads due to payloads to 50% of what an efficient message logging protocol would require. In our hybrid protocol, if after failure a survivor has progressed past these green iterations, it has stopped logging message payloads (red iterations). Therefore, it cannot replay messages and local rollback is impossible; However, since we never compromise on global rollback, we can still perform global rollback. If all survivors are within the first 2 green iterations when a failure is detected, local rollback is guaranteed.

We are not aware of any work in the MPI community implementing a hybrid mechanism between local and global rollback, which can gradually shift between either, depending on the user requirements. While this modification is easy for message logging application kernels, whether this is possible with a message logging MPI runtime, acting on specific application hooks, is a more challenging research question.

The high-level logic of the replay routine is generic, but the sequence and arguments of the MPI sends is application-dependent. It must accurately mimic the underlying application iterations. Perhaps an advanced source-to-source compiler may be able to support this process in the future.

7 EXTENSIONS TO CG AND LULESH

7.1 CG from NAS Parallel Benchmarks

The NAS parallel benchmarks are documented by Bailey et al. [5] and widely used in the HPC domain. We have chosen the NAS benchmark suite to demonstrate our approach mainly due to its popularity in the HPC community. We hope that by providing the extended version of the benchmark, researchers and developers will find it suitable for reproducible experiments and comparison. In terms of development, NAS benchmarks provide pre-defined correctness checks, which are extremely useful to us when integrating fault tolerance mechanisms.

The focus on CG as iterative kernel is due to the fact that its communication pattern generates larger overheads than most other

kernels, posing a significant challenge for message logging protocols.

The MPI version of the NAS CG benchmark is available in Fortran. While the preferred option for us would be to only make incremental changes to the original Fortran code, various challenges have been reported when using ULFM with Fortran applications (see Weeks et al. [23]). Therefore, despite the effort, we decided to port the MPI version from Fortran to C.

7.2 LULESH

LULESH is a hydrodynamics code [16] provided by LLNL, and a challenging mini-app for scientists to study in the HPC programming efforts towards exascale. Our reference code is the provided checkpointing variant of the 1.0 LULESH code.

LULESH is interesting to us since it is a much larger, and more realistic HPC code, than the very compute-intensive but often benchmark-oriented CG kernel. Despite the complexities of the code, and the significant larger number of communication calls across the code, we do not perceive LULESH as more challenging in its communication patterns than CG. Its neighbourhood exchanges in 3D are reminiscent of stencil codes, which do not require collective communication.

7.3 Problem Sizes

For the NAS benchmarks, the problem sizes across classes A-E are well documented for CG; our tests are mostly limited to classes C-E. For LULESH, we use the predefined 45^3 elements per domain, and experiment with different process sizes, each with this fixed domain size (weak scaling).

7.4 Summary of extensions

Our extensions for either CG or LULESH consist of following steps:

- Introduce global rollback via traditional checkpointing, where unavailable (up to 220 lines for LULESH)
- Introduce ULFM routines for revoking and recovering MPI communicators after failure. This code (≈ 200 lines) is a copy/paste from existing tutorial versions [3].
- Introduce message logging, which consists of:
 - buffer copying for transactional iterations, allowing transparent rollback within an iteration (10 lines for CG, 120 lines for LULESH)
 - modified sender calls, and replay routines (around 200 lines for both)

We summarise the lines of code (LOC) for each of these in Tab. 2.

The most lines specific to message logging are introduced for the replay routine of Fig. 4, which mimics the communication pattern of each application kernel. We have outlined our extensions previously; our message logging versions of CG and LULESH are open-sourced and freely available [2].

–	Message logging CG	Message logging LULESH
C/R	50	220
ULFM-specific	200	200
Message logging (replay, sender wrapper, payload logging)	200	200
Temp buffers (for transactions)	10	120
Total	1800	6500

Table 2: LOC for message logging CG and message logging LULESH.

8 EXPERIMENTS

8.1 Verification

Given the fact that we introduced message logging for the two kernels, the most significant outcome of our work is not performance, but the validation that the local rollback works correctly. Our validation was as follows:

- For each checkpoint interval CP_INT for either of the two kernels, crash a fixed MPI process in consecutive runs in iterations i , with $0 \leq i < CP_INT$.
- Use an application metric to compare the final result with the fault-free execution. NAS already provides invaluable verification tests (ζ and r values). For LULESH, we read the final origin energy.

For either of the two kernels, and local rollback, we verified that errors in the replay routine, or not enabling transactional iterations, are among the main causes of incorrect results. Global rollback and failure-free runs produce correct results regardless of transactional iterations or replay routine.

8.2 Performance measures

Here, we give an overview of the runtime and memory overheads we measure for payload logging for either kernel.

8.2.1 Experimental Platform. We use the Cirrus cluster in EPCC (University of Edinburgh) for our experiments; each of 280 nodes consisting of 2 18-core Broadwell processors with 128 GB RAM per processor. The interconnect is FDR infiniband. Details are provided online [1]. We measured the bandwidth and latency across two nodes via NetPIPE; the maximum bandwidth is close to 50 Gbps, and the latency across nodes ≈ 2 microseconds.

We use ULFM 2.0 rc1 from November 2017 for all our experiments, configured with `--with-ft`.

8.2.2 Payloads – memory overheads. Here, we quantify the memory overheads for message payload logging. We can measure the per-log and the total heap space usage during an execution (we have modified static to dynamic memory allocation in the code) by using the popular Valgrind Massif tool for heap profiling.

The payload overheads are given in Table 3 for 16-proc CG runs and 27-proc LULESH runs. Every process logs in the range of 450KB – 9MB for each single iteration, which represents a memory overhead of up to 0.1% for CG, and up to 0.7% for LULESH, for each single iteration.

<kernel> – <#procs>	heap use without logging	payload size per process and iteration	percentage increase (heap)
cg.B.16	≈ 434 MB	≈ 450 KB	0.1%
cg.C.16	≈ 1.08 GB	≈ 900 KB	0.09%
cg.D.16	≈ 20 GB	≈ 9 MB	0.045%
LULESH.27	≈ 70 MB	≈ 500 KB	0.7%

Table 3: Memory overheads of payload logging for 16-process CG and 27-process LULESH.

<kernel> – <#procs>	I/O	MPI time	payload copy	buffer copy	total
lulesh.216	0.6 (2%)	4.2 (12.7%)	0.006 (0.02%)	2.7 (8%)	33.2
cg.C.256	1.2 (21%)	1.4 (25%)	0.35 (6.2%)	0.053 (1%)	5.6

Table 4: Runtimes for a 75 iteration LULESH run (216 processes on 6 nodes), and a 75 iteration cg.C.256 run (256 procs on 8 nodes). Results are aggregations after manually instrumenting the phases of each kernel. Readings taken from rank 0; all units are seconds.

The accumulated overheads then depend on the checkpoint interval if we log all payloads, or if enabled, the capped payload logs of the hybrid protocol. For example, if we checkpoint only once each 26 CG iterations, we would have payload overheads of up to 234 MB per MPI process for cg.D.16, which is over 1% memory overhead per process. If we checkpoint each 20 iterations for LULESH, we would have payload overheads of up to 20 MB per MPI process, which is ≈ 28% memory overhead per process. Payload capping can reduce these memory overheads.

8.2.3 Payloads and temporary buffers – runtime overheads. For completeness, we show a more detailed breakdown of execution time in Table 4. It includes I/O overheads, time spent in MPI calls, payload copy overhead, and buffer copy overhead to enable transactional iterations; the wall clock times for LULESH and CG are also included, and used to calculate the relative contributions of the corresponding phases.

For all measured times, we have manually instrumented both CG and LULESH across different phases. We also show the percentage of these times to the overall wall clock times. Since CG works with larger data volumes than LULESH, the MPI and I/O times represent a higher percentage of the overall runtime. We are more interested in the overheads for payloads, and also for the temporary buffers enabling transactional iterations, both parts of our extensions. CG incurs significant payload overheads of up to 6.2% with payload logging, compared to the almost non-existent payload overheads for LULESH; again, this relates to the significant communication exchange for CG. These results are consistent with the results of related work with message logging protocols. In contrast, the buffer duplication to enable transactions comes at a high price for LULESH (8%), compared to the low overhead for CG (1%). Since we choose to create deep copies of the entire LULESH domain each iteration,

Figure 6: Total summed iterations for 16-node runs with cg.D.64 and LULESH.64, depending on where a failure is triggered, and type of rollback.

which represents all application data, this is also not surprising. In contrast, for CG we only copy a few persistent vectors each iteration.

Again, payload capping can reduce the runtime overheads, which are significant for CG. The copy into temporary buffers cannot be reduced in the same way, since transactional iterations are a requirement for the proposed approach; however, since we copy the entire application data for LULESH, various optimisations seem feasible.

8.2.4 Wall clock times. In this section, we measure what are the consequences of local and global rollback to recompute, as well as overall wall clock times.

We outlined our non-trivial setting earlier in this section – we crash a fixed MPI process in consecutive runs for iterations i , with $0 \leq i < CP_INT$. We set the CP_INT for CG to 25 iterations, and for LULESH to 20 iterations. Within this setting, we first determine the total number of compute iterations, summed across all processes, as a function of when a crash happens (see Fig. 6). For the same runs, we measure the wall clock time to completion after failure for all runs; the times are shown in Fig. 7. As we can see, local rollback as expected decreases consistently the number of recomputed iterations, since only restarted processes roll back. However, this does not lead to significantly faster overall execution times. Since the restarted process is on the critical path of recovery both for local and global rollback, the results are not surprising. We believe in general faster execution time should not be expected of local rollback. Rather, local rollback needs to be combined with other techniques to yield immediate end-to-end benefits; we have previously demonstrated such techniques [12] by applying frequency scaling to the idle processes. Another potential optimisation is to collocate with other processes during recovery on idle processors.

Finally, we have validated the hybrid local-global rollback, by setting the capacity of the payload logs to $LOG_SIZE = \frac{CP_INT}{2}$. As

Figure 7: Wall clock time for 16-node runs with cg.D.64 and LULESH.64, depending on where a failure is triggered, and type of rollback. The plot indicates that local rollback does not offer significant time reduction in overall execution time, compared to global rollback.

Figure 8: Demonstration of the hybrid protocol which automatically switch between local and global rollback. In this example, the payload capacity is $\frac{CP_INT}{2}$ for each kernel.

expected, the protocol gracefully switches between local rollback, if a failure strikes early during the CP_INT window, and global rollback, if the failure strikes later. We show the runtime results for the hybrid approach, again testing different iterations of failure for a fixed process, in Fig. 8. While the hybrid protocol dynamically switches between local and global rollback, always providing correct results, the overall wall clock time linearly increases either way, since when a failure strikes later, more iterations need to be repeated by one process, or all processes.

9 CONCLUSION

In this work, we introduced message logging capabilities into two popular HPC kernels, the NAS CG benchmark, and the mini-app LULESH. One of our main motivators for doing so is the fact that message logging MPI runtimes get published today, but are not easy to deploy by anyone but a few experts in these protocols. We introduced a methodology for implementing message logging into HPC kernels; this included using additional temporary buffers to enable what we called transactional MPI iterations, so that the rollback of a process does not break its data consistency. We then introduced sender-based payload logging and a replay mechanism, which enabled us to restore the inter-process data consistency. All of the above techniques were presented in a generic way, but aspects such as the payload logging or sequence of replayed messages remained kernel-specific. As a performance optimisation which has never been proposed before, we were also able to introduce a hybrid rollback protocol, which gracefully shifts from local to global rollback, if we choose to decrease the available memory for payload logging.

While our focus is on validation rather than performance, we presented experimental results, which show that local rollback is fully functional, and therefore a complete replacement for message logging MPI runtimes, for these kernels, and carries less overheads than most implementations (e.g. no event logger required). However, the idling processes due to reduced rollback did not immediately reduce overall runtime; we have pointed to techniques which exploit the reduced rollback to bring end-to-end runtime or energy benefits. We hope our methodology and codebase may be useful to the community for experimenting with message logging protocols as compact HPC kernels extensions, as a replacement of complex and difficult to set up message logging runtimes.

ACKNOWLEDGMENT

This work used the Cirrus UK National Tier-2 HPC Service at EPCC (<http://www.cirrus.ac.uk>) funded by the University of Edinburgh and EPSRC (EP/P020267/1).

REFERENCES

- [1] Cirrus hardware. <http://www.cirrus.ac.uk/about/hardware.html>.
- [2] Message logging kernels. <https://github.com/KADichev/message-logging-kernels>.
- [3] ULFM Tutorial for SC'17. <http://fault-tolerance.org/downloads/SC17-handson.pdf>.
- [4] Lorenzo Alvisi and Keith Marzullo. Message logging: Pessimistic, optimistic, causal, and optimal. *IEEE Transactions on Software Engineering*, 24(2):149–159, 1998.
- [5] David H Bailey, Eric Barszcz, John T Barton, David S Browning, Robert L Carter, Leonardo Dagum, Rod A Fatoohi, Paul O Frederickson, Thomas A Lasinski, Rob S Schreiber, et al. The nas parallel benchmarks. *The International Journal of Supercomputing Applications*, 5(3):63–73, 1991.
- [6] Wesley Bland, Aurelien Bouteiller, Thomas Herault, George Bosilca, and Jack Dongarra. Post-failure recovery of mpi communication capability: Design and rationale. *The International Journal of High Performance Computing Applications*, 27(3):244–254, 2013.
- [7] George Bosilca, Aurelien Bouteiller, Franck Cappello, Samir Djilali, Gilles Fedak, Cecile Germain, Thomas Herault, Pierre Lemarinier, Oleg Lodygensky, Frederic Magniette, et al. Mpich-v: Toward a scalable fault tolerant mpi for volatile nodes. In *Supercomputing, ACM/IEEE 2002 Conference*, pages 29–29. IEEE, 2002.
- [8] A. Bouteiller, B. Collin, T. Herault, P. Lemarinier, and F. Cappello. Impact of event logger on causal message logging protocols for fault tolerant mpi. In *19th IEEE International Parallel and Distributed Processing Symposium*, pages 10 pp.–, April 2005.

- [9] Aurelien Bouteiller, George Bosilca, and Jack Dongarra. Redesigning the message logging model for high performance. *Concurrency and Computation: Practice and Experience*, 22(16):2196–2211, 2010.
- [10] Aurélien Bouteiller, Franck Cappello, Thomas Herault, Gérard Krawezik, Pierre Lemarini, and Frédéric Magniette. Mpich-v2: a fault tolerant mpi for volatile nodes based on pessimistic sender based message logging. In *Proceedings of the 2003 ACM/IEEE conference on Supercomputing*, page 25. ACM, 2003.
- [11] Aurelien Bouteiller, Thomas Ropars, George Bosilca, Christine Morin, and Jack Dongarra. Reasons for a pessimistic or optimistic message logging protocol in mpi uncoordinated failure, recovery. In *2009 IEEE International Conference on Cluster Computing and Workshops*, pages 1–9. IEEE, 2009.
- [12] Kiril Dichev, Kirk Cameron, and Dimitrios S Nikolopoulos. Energy-efficient localised rollback via data flow analysis and frequency scaling. In *Proceedings of the 25th European MPI Users' Group Meeting*, page 11. ACM, 2018.
- [13] M. Gamell, K. Teranishi, J. Mayo, H. Kolla, M. A. Heroux, J. Chen, and M. Parashar. Modeling and simulating multiple failure masking enabled by local recovery for stencil-based applications at extreme scales. *IEEE Transactions on Parallel and Distributed Systems*, 28(10):2881–2895, Oct 2017.
- [14] Marc Gamell, Daniel S. Katz, Hemanth Kolla, Jacqueline Chen, Scott Klasky, and Manish Parashar. Exploring automatic, online failure recovery for scientific applications at extreme scales. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC '14*, pages 895–906, Piscataway, NJ, USA, 2014. IEEE Press.
- [15] Amina Guermouche, Thomas Ropars, Elisabeth Brunet, Marc Snir, and Franck Cappello. Uncoordinated checkpointing without domino effect for send-deterministic mpi applications. In *Parallel & Distributed Processing Symposium (IPDPS), 2011 IEEE International*, pages 989–1000. IEEE, 2011.
- [16] Ian Karlin, A. Bhatele, B. L. Chamberlain, J. Cohen, Z. Devito, M. Gokhale, R. Haque, R. Hornung, J. Keasler, D. Laney, E. Luke, S. Lloyd, J. McGraw, R. Neely, D. Richards, M. Schulz, C. H. Still, F. Wang, and D. Wong. Lulesh programming model and performance ports overview. Technical Report LLNL-TR-608824, December 2012.
- [17] Ignacio Laguna, David F. Richards, Todd Gamblin, Martin Schulz, and Bronis R. de Supinski. Evaluating user-level fault tolerance for mpi applications. In *Proceedings of the 21st European MPI Users' Group Meeting, EuroMPI/ASIA '14*, pages 57:57–57:62, New York, NY, USA, 2014. ACM.
- [18] Nuria Losada, George Bosilca, Aurelien Bouteiller, Patricia González, and María J. Martín. Local rollback for resilient mpi applications with application-level checkpointing and message logging. *Future Generation Computer Systems*, 91:450–464, 02-2019 2019.
- [19] Esteban Meneses, Greg Bronevetsky, and Laxmikant V Kale. Evaluation of simple causal message logging for large-scale fault tolerant hpc systems. In *Parallel and Distributed Processing Workshops and Phd Forum (IPDPSW), 2011 IEEE International Symposium on*, pages 1533–1540. IEEE, 2011.
- [20] Thomas Ropars and Christine Morin. *O2P: An Extremely Optimistic Message Logging Protocol*. PhD thesis, INRIA, 2006.
- [21] B. Schroeder and G. Gibson. A large-scale study of failures in high-performance computing systems. *IEEE Transactions on Dependable and Secure Computing*, 7(4):337–350, Oct 2010.
- [22] Rob Strom and Shaula Yemini. Optimistic recovery in distributed systems. *ACM Transactions on Computer Systems (TOCS)*, 3(3):204–226, 1985.
- [23] Nathan Weeks, Glenn Luecke, Pieter Maris, and James Vary. Challenges in developing mpi fault-tolerant fortran applications. In *2018 IEEE International Conference on Cluster Computing (CLUSTER)*, pages 524–531. IEEE, 2018.
- [24] Gengbin Zheng, Xiang Ni, and Laxmikant V Kalé. A scalable double in-memory checkpoint and restart scheme towards exascale. In *Dependable Systems and Networks Workshops (DSN-W), 2012 IEEE/IFIP 42nd International Conference on*, pages 1–6. IEEE, 2012.