

Redesigning the Message Logging Model for High Performance

Aurelien Bouteiller, George Bosilca, Jack Dongarra
ICL, University of Tennessee Knoxville
Claxton 1122 Volunteer Boulevard
37996 Knoxville TN, USA

Abstract

Over the past decade the number of processors in the high performance facilities went up to hundreds of thousands. As a direct consequence, while the computational power follow the trend, the mean time between failures (MTBF) suffered, and it's now being counted in hours. In order to circumvent this limitation, a number of fault tolerant algorithms as well as execution environments have been developed using the message passing paradigm. Among them, message logging has been proved to achieve a better overall performance when the MTBF is low, mainly due to it's faster failure recovery. However, message logging suffers from a high overhead when no failure occurs. Therefore, in this paper we discuss a refinement of the message logging model intended to improve failure free message logging performance. The proposed approach simultaneously removes useless memory copies and reduces the number of logged events. We present the implementation of a pessimistic message logging protocol in Open MPI and compare it with the previous reference implementation MPICH-V2. Results outline a several order of magnitude improvement on performance and a zero overhead for most messages.

1 Introduction

The Top500 list, the list of the 500 most powerful supercomputers in the world, highlight a constant increase in the number of processors. An attentive reading shows that the top 10 of these supercomputers detain several thousands of processor each. Despite, a more careful design of the components in these supercomputers, the increase in the number of components directly affect the reliability of these systems. Moreover, the next computers in the Top500 list are build using commodity components, which greatly affect their reliability. While the Mean Time Between Failures (MTBF) on the BlueGene/L [18] is counted on days, the commodity clusters exhibit a usual MTBF of tens of hours.

With the expected increase in the number of processors in the next generation supercomputers, one might deduct that the probability of failures will continue to increase, leading to a drastic decrease in reliability.

Fault tolerant algorithm have a long history of research. Only recently, as the practical issue was raising, High Performance Computing (HPC) software have been adapted to deal with failures. As most HPC applications are using the Message Passing Interface (MPI) [19] to manage data transfers, introducing failure recovery features inside the MPI library automatically benefits to a large range of applications. One of the most popular automatic fault tolerant technique, coordinating checkpoint, builds a consistent recovery set [12, 15]. As today's HPC users are facing occasional failures, they have not suffered from its slow recovery procedure, involving restarting all the computing nodes even when only one failed. Considering future systems will endure higher fault frequency, recovery time could become another gap between the peak performance of the architecture and the effective performance users can actually harvest from the system.

Because message logging does not rely on such coordination, it is able to recover faster from failures. From previous results [12], it is expected to have a better typical application makespan than coordinated checkpoint when the MTBF is shorter than 9 hours while coordinated checkpoint would not be able to progress anymore for a MTBF shorter than 3 hours. Still, message logging suffers from a high overhead on communication performance. Moreover, the better the latency and bandwidth offered by newer high performance networks, the higher the relative overhead. Those drawbacks needs to be addressed to provide a resilient and fast fault tolerant MPI library to the HPC community. In this paper we propose a refinement of the classical model of message logging paradigm, closer to the reality of high performance network interface cards, where message receptions are decomposed in multiple dependent events. We better categorize message events allowing 1) to remove intermediate message copies on high performance networks and 2) to clearly recognize deterministic and non-deterministic

events, thus reducing the overall number of message requiring latency disturbing management. We demonstrate how this refinement can be used to reduce the fault free overhead of message logging protocols by implementing it in Open MPI [9]. Its performance is compared with the previous reference implementation of message logging MPICH-V2. Results outline a several orders of magnitude improvement of the fault free performance of pessimistic message logging and a drastic reduction in the overall number of logged events.

The rest of this paper is organized as follows: in the next section we recall the classical message logging and then depict the modifications we introduce to better fit HPC networks. In the third section we depict the implementation issues of the prototype in Open MPI. Fourth section presents experimental evaluation, followed by the related works and conclusion.

2 Message Logging

2.1 Classical Message Logging

Message logging is usually depicted using the more general model of message passing distributed systems. Communications between processes are considered explicit: processes explicitly request sending and receiving messages; and a message is considered as delivered only when the receive operation associated with the data movement complete. Additionally, from the perspective of the application each communication channel is FIFO, but there is no particular order on messages traveling along different channels. The execution model is pseudo-synchronous: there is no global shared clock among processes but there is some (potentially unknown) maximum propagation delay of messages in the network. An intuitive interpretation is to say the system is asynchronous and there is some *eventually reliable* failure detector.

Failures can affect both the processes and the network. Usually, network failures are managed by some CRC and message reemission provided by the hardware or low level software stack and do not need to be considered in the model. We consider that processes endure definitive crash failures, where a failed process stops sending any message.

Events Each computational or communication step of a process is an event. An execution is an alternate sequence of events and process states, the effect of an event on the preceding state leading the process to the new state. As the system is basically asynchronous, there is no direct time relationship between events occurring on different processes. However, Lamport defines a causal partial ordering between events with the *happened before* relationship [11]. It is noted $e < f$ when event f is causally influenced by e .

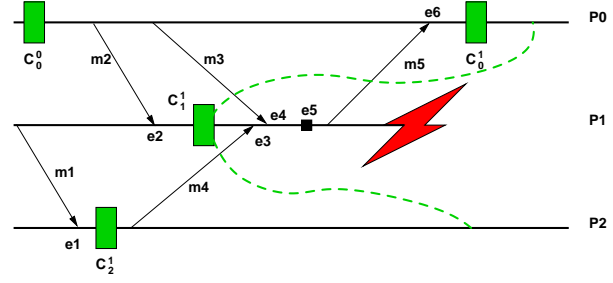


Figure 1. Example execution of a distributed system with checkpoints and inconsistent recovery line.

These events can be classified into two categories: deterministic and non-deterministic. An event is deterministic when from the current state there is only one outcome state for this event. On the contrary, if an event can result in several different states depending on its outcome, then it is non-deterministic. Examples of deterministic events are internal computations and message emissions, which follow the code-flow. Examples of non-deterministic events are message receptions, which depends on time constraints on message deliveries.

Checkpoints and Inconsistent States Checkpoints (i.e. complete images of the process memory space) are used to recover from failures. The recovery line is the configuration of the application after some processes have been reloaded from checkpoints. Unfortunately, checkpointing a distributed application is not as simple as storing each single process image without any coordination, as illustrated by the example execution of figure 1. When process P_1 fails, it rolls back to checkpoint C_1^1 . Messages from the past crossing the recovery line (m_3, m_4) are *in transit* messages: the restarted process will request their reception while the source process never sends them again, thus it is needed to save these messages. Messages from the future crossing the recovery line (m_5) are orphan: following the Lamport relationship current state of P_0 depends on reception of m_5 and by transitivity on any event that occurred on P_1 since C_1^1 (e_3, e_4, e_5). Since the channels between P_0 and P_1 and between P_2 and P_1 are asynchronous, the reception of m_3 and m_4 could occur in a different order during re-execution, leading during the recovery to a state of P_1 that diverge from the initial execution. As the current state of P_0 depends on states P_1 could not reach anymore, the overall state of the parallel application after the recovery is inconsistent. Checkpoints leading to an inconsistent state are useless and must be discarded, in the worst case all checkpoints are useless and the computation may have to be restarted

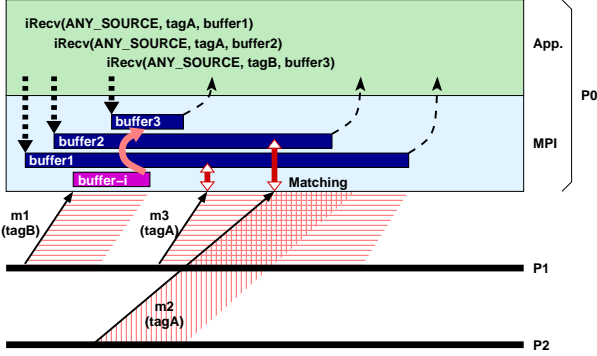


Figure 2. Steps in a zero copy MPI receive operation

from the beginning.

Recovery In event logging, processes are considered as *Piecewise deterministic*: only sparse non-deterministic events occur separating large parts of deterministic computation. Considering that non-deterministic events are committed during the initial execution into some *safe* repository, a recovering process is able to replay exactly the same order for all non-deterministic events. Therefore, it is able to reach exactly the same state as before the failure. Furthermore, message logging considers the network as the only source of non-determinism and only log the relative ordering of message from different senders (e_3, e_4 in fig. 1). The sufficient condition ensuring a successful recovery requires that a process must never depend on an unlogged event from another process. As the only way to create a dependency between processes is to send a message, all non-deterministic events occurring between two consecutive sends can be merged and committed together.

Event logging only save events in the remote repository, without storing the message payload. However, when a process is recovering, it needs to replay any reception that happened between last checkpoint and the failure and therefore requires those messages payload (m_3, m_4 in fig.1). During normal operation, every outgoing message is saved in the sender's volatile memory; a mechanism called sender-based message logging. This allows for the surviving processes to serve past messages to recovering processes, on demand without rolling back. Unlike events, sender-based data do not require stable or synchronous storage. Should a process holding useful sender-based data crash, the recovery procedure of this process replays every outgoing send and thus rebuilds the missing messages.

2.2 Non-blocking Communications

To reach top performance, the MPI standard defines a more sophisticated set of communication routines than simple blocking send and receive. One of the most important optimization for high throughput communication library is *zero copy*, the ability to send and receive directly in the application's user-space buffer without intermediary memory copies. Figure 2 shows the basic steps of non-blocking zero copy communications in MPI. First the application *requests* a message to be received, specifying the message source, tag and reception buffer. When a message arrives from the network, the source and the tag are compared to the pending requests. If the message does not match any pending request it is copied in a temporary buffer (m_1 is tagged *tagB* and all pending requests wants *tagA*) until a matching request is posted by the application. When the message match a pending request like m_3 and m_4 it is directly written in the receive buffer without intermediate copy. Because requests can be `ANY_SOURCE` the result of the matching may depend on the order of reception of the first fragment of the message (if m_2 have arrived earlier it would have been delivered in *buffer1*). Next step for the application might be to probe for message delivery readiness. The result of those probes function may depend on the message transfer termination time, but is not related to the matching order (m_3 is matched first but lasts longer than m_2).

Because classical message logging model assumes message reception is a single atomic event, it cannot catch the complexity of zero copy MPI communications involving distinct matching and delivery events. As an example in MPICH, only the lowest blocking point-to-point transport layer called the *device* matches the classical model, explaining why previous state of the art message logging implementation, such as MPICH-V, replaces the low level device by the `ch_v` fault tolerant one (see Figure 3(a)). This device has adequate properties regarding the hypothesis of message logging: 1) messages are delivered in one single atomic step to the application (though message interleave is allowed inside the `ch_v` device), 2) intermediate copies are made for every message to fulfill this atomic delivery requirement, the matching needs to be done at delivery time, 3) as the message logging mechanism replaces the regular low level device, it cannot easily benefit from zero copy and OS bypass features of modern network cards and 4) because it is not possible to distinguish the deterministic events at this software level, every message generates an event (which is obviously useless for deterministic receptions). We will show in the performance analysis section how this strong model requirements lead to dramatical performance overhead in MPI implementation when considering high performance interconnects.

2.3 Refinements for Zero Copy Messages

By interposing the event logging mechanism inside the MPI library, we can relax the strong model described previously. It is only necessary to log the communication events at the library level, and completely ignore the expensive events generated by the lower network layer. This requires to consider the particularity of the internal MPI library events, but allow us to use the optimized network layers provided by the implementation.

Network Events From the lower layer comes the packet related events: let m denote a message transferred in $length(m)$ network packets. We note r_m^i the i^{th} packet of message m , where $1 \leq i \leq length(m)$. Because the network is considered reliable and FIFO, we have $\forall 1 \leq i \leq length(m) - 1, r_m^i \prec r_m^{i+1}$. We denote $tag(m)$ the tag of message m and $src(m)$ its emitter. Packets are received atomically from the network layer.

Application Events From the upper layer, come the application related events. We note $Post(tag, source)$ a reception post, $Probe(tag, source)$ the event of checking the presence of a message, $Wait(n, \{R\})$ the event of waiting n completions of the set $\{R\}$ of request identifiers. Because the application is considered piecewise deterministic, we can assign a totally ordered sequence of identifiers to upper layer events. Let r_0 be a request identifier obtained by the $Post_0(tag_0, source_0)$ event. Since posting is the only way to obtain a request identifier, if $r_0 \in \{R\}$, $Post_0(tag_0, source_0) \prec Wait_0(n, \{R\})$. There is at most one event $Post$ per message and at least one $Wait$ event per message. If $r_{m_0}^1 \prec Probe_0(tag_0, source_0) \prec Post_0(tag_0, source_0)$, then $Probe_0(tag_0, source_0)$ must return true. Otherwise, it must return false. The main difference between $Probe$ and $Post$ is that in case $r_{m_0}^1$ precede one of these events, $Probe_0(tag_0, source_0)$ will not discard $r_{m_0}^1$, while $Post_0(tag_0, source_0)$ will always do.

Library Events The library events are the result of the combination of a network-layer event and an application-layer event. There are two categories of library events: 1) Matching (denoted by M) and 2) Completing (denoted by C). Matching binds a network communication with a message reception request; Completing checks the internal state of the communication library to determine the state of a message (completed or not).

- 1) To build a Matching event from a reception event and a Post event, we define a reception-matching pair of events: r_m^1 and $Post_0(tag_0, source_0)$ match for reception if and only if $(source_0 = src(m) \vee source_0 = ANY) \wedge (tag_0 = tag(m) \vee tag_0 =$

$ANY)$. The Matching event built from the reception-matching events is causally dependent from the two elements of the matching pair: $Post_0(tag_0, source_0) \prec M_0$ and $r_m^1 \prec M_0$. The reception-matching pair is deterministic if and only if $source_0 \neq ANY$. Additionally, based on the same rules, we can build a Matching from a Probe event and a reception event. In this case, the result of the Matching M_0 is successful if and only if $r_m^1 \prec Probe_0(tag_0, source_0)$. Otherwise, the Matching event takes a special value (undefined source). Because the order between r_m^1 and $Probe_0(tag_0, source_0)$ is non-deterministic, all probe-matching pair events are non-deterministic.

- 2) Similarly, to build a Completing event from a reception event and a Wait event, we define a completion-matching pair of events: $r_m^{length(m)}$ and $Wait(n, \{R\})$ match for completion if and only if there is a matching event M_0 built from r_m^1 containing the request identifier r_0 and $r_0 \in \{R\}$. The Completing event built from the completion-matching events is causally dependent from the two elements of the matching pair: $Wait(n, \{R\}) \prec C_0$ and $r_m^{length(m)} \prec C_0$. All the r_m^i events are non-deterministic per definition. Thus, every $Wait(n, \{R\})$ event is non-deterministic, because the result of these events depend upon the internal state of the library, which depends upon the $r_m^{length(m)}$ events. However, according to the matching and completion rules, if $r_m^{length(m)}$ and $Wait(n, \{R\})$ is a completion-matching pair, the Completing event built is deterministic if and only if $n_0 = |R_0|$ (case of Wait, WaitAll, Recv).

Although the refinement introduces many new events, most of them are not necessarily logged, since they are deterministic. Only non-deterministic events (non-deterministic Matching due to ANY sources; non-deterministic Matching due to probe-matching events; non-deterministic completion due to WaitSome, WaitAny, TestAll, Test, TestAny and TestSome) are logged and introduce a synchronization with the event logger.

3 Implementation in Open MPI

3.1 Generic Fault Tolerant Layer

Open MPI architecture is a typical example of the new generation MPI implementations. The figure 3(b) summarizes the Open MPI software stack dedicated to MPI communications. Regular components are summarized with plain lines, while the new fault tolerant components are dashed. At the lowest level, the BTL expose a set of

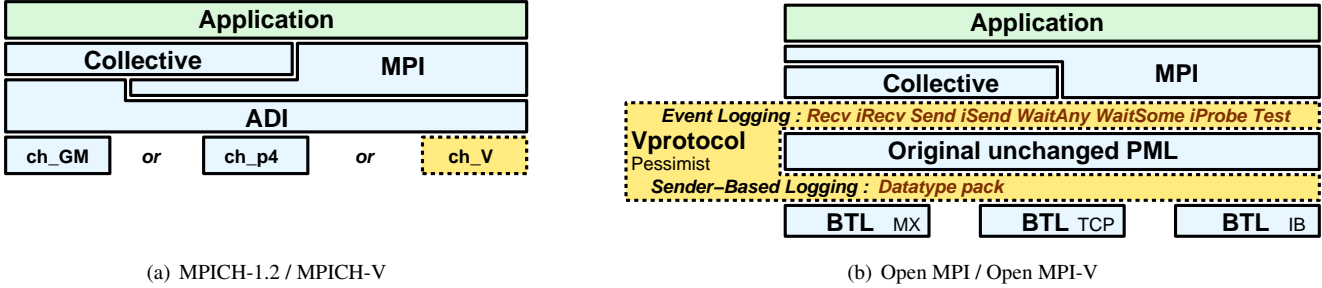


Figure 3. Comparison between the MPICH and the Open MPI architecture and the interposition level of fault tolerance (fault tolerant components are dashed)

communication primitives appropriate for both send/receive and RDMA interfaces. A BTL is MPI semantics agnostic; it simply moves a sequence of bytes (potentially non-contiguous) across the underlying transport. Multiple BTLs might be in use in the same time to strip data across multiple networks. The PML implements all logic for point-to-point MPI semantics including standard, buffered, ready, and synchronous communication modes. MPI message transfers are scheduled by the PML based on a specific policy according to short and long protocol, and using control messages (ACK/NACK/MATCH). Additionally, the PML is in charge of providing the MPI matching logic as well as reordering the out-of-order fragments. All remaining MPI functions, including some collective communications, are built on top of the PML interface. While in the current implementation of the fault tolerant components only point-to-point based collectives are supported, we plan to support in the near future other forms of collective communication implementations (such as hardware based collectives).

In order to integrate the fault tolerance capabilities in Open MPI, we added one new class of components, the `Vprotocol` (dashed in the figure 3(b)). A `Vprotocol` component is a parasite enveloping the default PML. Each such is an implementation of a particular fault tolerant algorithm; its goal is not to manage actual communications but to extent the PML with message logging features. As all of the Open MPI components, the `Vprotocol` module is loaded at runtime on user’s request and replaces some of the interface functions of the PML by its own. Once it has logged or modified the communication requests according to the needs of the fault tolerant algorithm, it calls the real PML to perform the actual communications. This modular design have several advantages compared to MPICH-V architecture: 1) it does not modify any core Open MPI component, regular PML message scheduling and device optimized BTL can be used, 2) expressing a particular fault tolerant protocol is easy, it is only focused on reacting to some events, not handling communications and 3) the best suited fault tolerant component can be selected at run time.

3.2 Pessimistic Message Logging Implementation

The `Vprotocol` pessimist is the implementation based on our refined model. It provides four main functionalities: sender-based message logging, remote event storage, any source reception event logging, non-deterministic delivery event logging. Each process have a local Lamport clock, used to mark events: during Send, iSend, Recv, iRecv and Start, every request receive the clock stamp as an unique identifier.

Sender-Based Logging The improvements we propose to the original model still rely on a sender-based message payload logging mechanism. We integrated the sender-based logging to the data-type engine of Open MPI. The data-type engine is in charge of packing (maybe non-contiguous) data into a flat format suitable for the receiver’s architecture. Each time a fragment of the message is packed, we copy the resulting data in a *mmaped* memory segment. Because the sender-based copy progress at the same speed as the network, it benefits from cache reuse and releases the send buffer at the same date. Data is then asynchronously written from memory to disk in background to decrease the memory footprint.

Event Logger Commits Non-deterministic events are sent to event loggers processes (EL). An EL is a special process added to the application outside the `MPI_COMM_WORLD`; several might be used simultaneously to improve scalability. Events are transmitted using non blocking MPI communications over an inter-communicator between the application process and the event logger. Though asynchronous, there is a transactional acknowledgement protocol to ensure that every event is safely logged before any MPI send can progress.

Any Source Receptions Any source logging is managed in the `iRecv`, `Recv` and `Start` functions. Each time an any

source receive is posted, the completion function of the corresponding request is modified. When the request is completed, the completion callback logs the event containing the request identifier and the matched source. During recovery, the first step is to retrieve the events related to the MPI process from the event logger. Then every promiscuous source is replaced by the well specified source of the event corresponding to the request identifier. Because channels are FIFO, enforcing the source is enough to replay the original matching order.

Non-Deterministic Deliveries Non-deterministic deliveries (NDD) are the `iProbe`, `WaitSome`, `WaitAny`, `Test`, `TestAll`, `TestSome` and `TestAny` functions. The Lamport clock is used to assign a unique identifier to every NDD operation. When a NDD ends, a new delivery event is created, containing the clock and the list of all the completed request identifiers. During replay, when the NDD clock is equal to the clock of the first event, the corresponding requests are completed by waiting for each of them.

It sometime happens that no request completed during a NDD. To avoid creating large number of events for consecutive unsuccessful NDD, we use lazy logging: only one event is created for all the consecutive failed NDD. If a send operation occurs, any pending NDD have to be flushed to the EL. If a NDD succeed, any pending lazy NDD is discarded. During recovery, NDD whose clock is lower than the first NDD event in the log have to return no request completed.

4 Performances

4.1 Experimental testbed

Experimental testbed includes two clusters. In the first cluster, each node is a dual Opteron 246 (2GHz) with 2GB DDR400 memory and a Myrinet 2000 PCI-E interconnect with a 16 ports Myrinet switch; this machine is used only for Myrinet 2000 experiments. In the second cluster, each node is a dual Xeon Woodcrest with 4GB DDR5300 memory. Two different networks are available on this machine, a Gigabit Ethernet and a Myrinet 10G interconnect with a 128 ports Myrinet switch. Software setup is Linux 2.6.18 using MX 1.2.0j. Benchmarks are compiled using gcc and gfortran 4.2.1 with the `-O3` flag. We used NetPIPE [16] to perform ping-pong tests, while the NAS Parallel Benchmarks 3.2.1 (NPB) and High Performance Linpack (HPL) are used to investigate application behavior. HPL uses threaded GotoBLAS. Because the proposed approach does not change the recovery strategy used in previous works, we only focus on failure free performance.

4.2 Benefits from Event Distinction

One of the main difference of the refined model is the split of message receptions into two distinct events. In the worst case that might lead to log twice as many events compared to the model used in other message logging implementations. However, better fitness between model and MPI internals allows for detecting (and discarding) deterministic events. Table 1 characterizes the amount of non-deterministic (actually logged in Open MPI-V) events compared to the overall number of exchanged messages. Though we investigated all the NPB kernels (BT, MG, SP, LU, CG, FT) to cover the widest spectrum of application patterns, we detected non-deterministic events in LU and MG only. In all other benchmarks, Open MPI-V does not log any event, thanks to the detection of deterministic messages. On both MG and LU, the only non-deterministic events are any-source messages ; there is no non-deterministic deliveries or probes. In MG, two third of the messages are deterministic, while in LU less than 1% are using the any source flag, outlining how the better fitting model allows to drastically decrease the overall number of logged events in most usual application patterns. As a comparison, MPICH-V2 logs at least one event for each message (and two for rendez-vous messages). According to our experiments, the same results hold for class A, C and D of the NAS. The ratio of logged events does not correlate with the number of computing processes in LU and decreases when more processes are used in MG, meaning that the fault tolerant version of the application is at least as scalable as the original one.

Avoiding logging of some events is expected to lower the latency cost of pessimistic protocol. Figure 4 presents the overhead on Myrinet round trip time of enabling the pessimistic fault tolerant algorithm. We compare Open MPI-Vpessimist with a similar non fault tolerant version of Open MPI, while we compare the reference message logging implementation MPICH-V2 with a similar version of MPICH-MX.

In this ping-pong test, all Recv operation are well specified sources and there is no `WaitAny`. As a consequence, Open MPI-Vpessimist does not create any event during the benchmark and reach exactly the same latency as Open MPI (3.79 μ s). To measure the specific cost of handling non-deterministic events in Open MPI-Vpessimist, we modified the NetPIPE benchmark code: every Recv have been replaced by the sequence of an any source `iRecv` and a `WaitAny`. This altered code generates two non-deterministic events for each messages. The impact on Open MPI-Vpessimist latency is a nearly three time increase in latency. The two events are merged into a single message to the event logger ; the next send is delayed until the acknowledgement comes back. This is the expected cost on latency

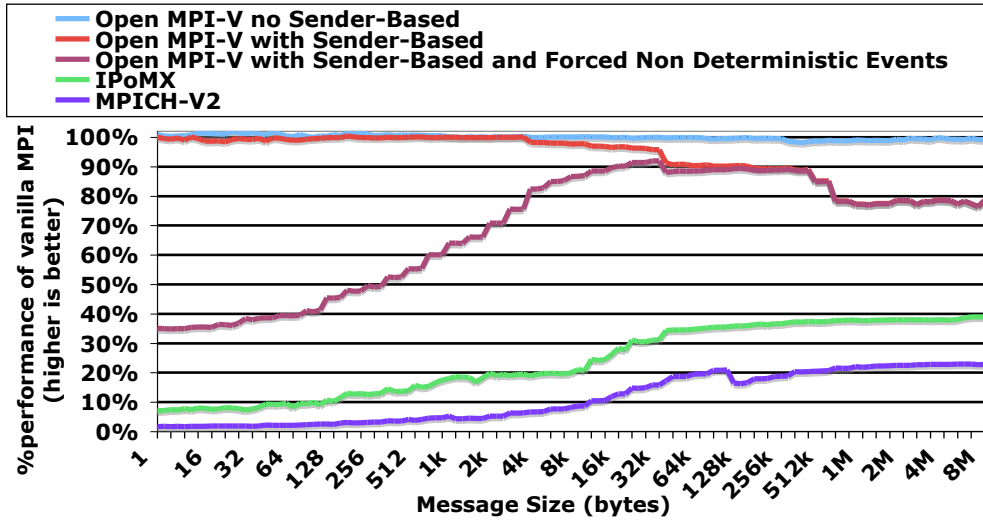


Figure 4. Myrinet 2000 ping-pong performance of pessimistic message logging in percentage of non fault tolerant MPI library

| | BT | SP | LU | | | | | |
|--------------------|-----|----|-------|-------|-------|-------|-------|-------|
| #processors | all | | 4 | 32 | 64 | 256 | 512 | 1024 |
| %non-deterministic | 0 | 0 | 1.13 | 0.66 | 0.80 | 0.80 | 0.75 | 0.57 |
| | FT | CG | MG | | | | | |
| #processors | all | | 4 | 32 | 64 | 256 | 512 | 1024 |
| %non-deterministic | 0 | 0 | 40.33 | 29.35 | 27.10 | 22.23 | 20.67 | 19.99 |

Table 1. Percentage of non-deterministic events to total number of exchanged messages on the NAS Parallel Benchmarks (Class B).

of pessimistic message logging : better detection of non-deterministic events removes the message logging cost for some common types of messages.

Because MPICH-V does not discard deterministic events from logging, there is a specific overhead ($40\mu s$ latency increase to reach $183\mu s$) for every message, even on the original deterministic benchmark. This specific overhead comes on top of those from memory copies.

4.3 Benefits from zero-copy receptions

Figure 4 shows the overhead of MPICH-V. With pessimistic protocol enabled, MPICH-V reaches only 22% of the MPICH-MX bandwidth. This bandwidth reduction roots in the number of memory copies on the critical path of messages. Because the message logging model used in MPICH-V assume delivery is atomic, it does not fit the intermediate layer of MPICH (similar to the PML layer of Open MPI) because it cannot accommodate the MPI matching and buffering logic. As a consequence event logging mechanism of MPICH-V replaces the low level *ch_mx* with a TCP/IP based device. The cost of memory copies introduced by this requirement is estimated by considering the performance of the NetPipe TCP benchmark on the IP emulation layer of MX: IPoMX. The cost of using TCP, with its internal copies and flow control protocol, is as high as 60% of the bandwidth and increases the latency from $3.16\mu s$ to $44.2\mu s$. In addition, the *ch_v* device itself needs to make an intermediate copy on the receiver to delay matching until the message is ready to be delivered. This is accountable for the 20% remaining overhead on bandwidth and increases the latency to $96.1\mu s$, even without enabling event logging.

On the contrary, in Open MPI-V the model fits tightly the behavior of MPI communications. The only memory copy come from the sender-based message payload logging ; there is no other memory copies. As a consequence Open MPI-V is able to reach a typical bandwidth as high as 1570Mbit/s (compared to 1870Mbit/s for base Open MPI and 1825Mbit/s for MPICH-MX). The difference between Open MPI-V with or without sender-based logging highlight the benefits of our cache reuse approach. While the sender-based copy fits in cache, the performance overhead of the extra copy is reduced to 11% and jumps to 28% for messages larger than 512kB.

4.4 Sender-based impact

While the overall number of memory copies have been greatly reduced, the sender-based message payload copy is mandatory and can't be avoided. Figure 5 helps understanding the source of this overhead by comparing the performance of Open MPI and Open MPI-V pessimist on different networks. As the sender-based copy is not on the critical

path of messages, there is no increase in latency, regardless of the network type. On Ethernet, bandwidth is unchanged as well, because the time to send the message on the wire is much larger than the time to perform the memory copy, thus a perfect overlap.

Counter-intuitively, Open MPI bandwidth for the non fault tolerant version is better on Myrinet 10G than on shared memory: the shared memory device uses a copy-in copy-out mechanism between processes, producing one extra memory access for each message (i.e. physically reducing the available bandwidth by two). Adding a third memory copy for handling sender-based logging to the two involved in regular shared memory transfer have up to 30% impact on bandwidth for large messages, even when this copy is asynchronous. This is the expected result considering that the performance bottleneck for shared memory network is the pressure on memory bus bandwidth.

As the sender-based message logging speed depends on memory bandwidth; the faster the network, the higher the relative copy time becomes. Myrinet 2G already reveal imperfect overlap between memory copies and network transmission, though when the message transfer fits in cache, the overhead is reduced by the memory reuse pattern of the sender-based mechanism. With the faster Myrinet 10G, the performance gap widen to 4.2Gbit/s (44% overhead). As the pressure on the memory subsystem is lower when using Myrinet 10G network than when using shared memory, one could expect sender-based copy to be less expensive in this context. However the comparison between Open MPI-V on Myrinet 10G and shared memory shows a similar maximum performance on both media, suggesting that some memory bandwidth is still available for improvements from better software engineering. Similarly, the presence of performance spikes for message sizes between 512kB and 2MB indicates that cache reuse strategy does not fit well the DMA mechanism used by this NIC.

4.5 Application performance and scalability

Figure 6(a) presents performance overhead of various numerical kernels on Myrinet 10G network with 64 nodes. Interestingly, the two benchmarks exhibiting non-deterministic events suffer from a mere 1% overhead compared to a non fault tolerant run. The more synchronous CG shows the highest performance degradation, topping to only 5% increase in execution time. Because there is no non-deterministic events in CG, overhead is solely due to sender-based payload logging.

Figure 6(b) compares performance of fault tolerant run of HPL with regular Open MPI on 90 quad core processors connected through Myrinet 10G, one thread per core. While the performance overhead is limited, it is independent of

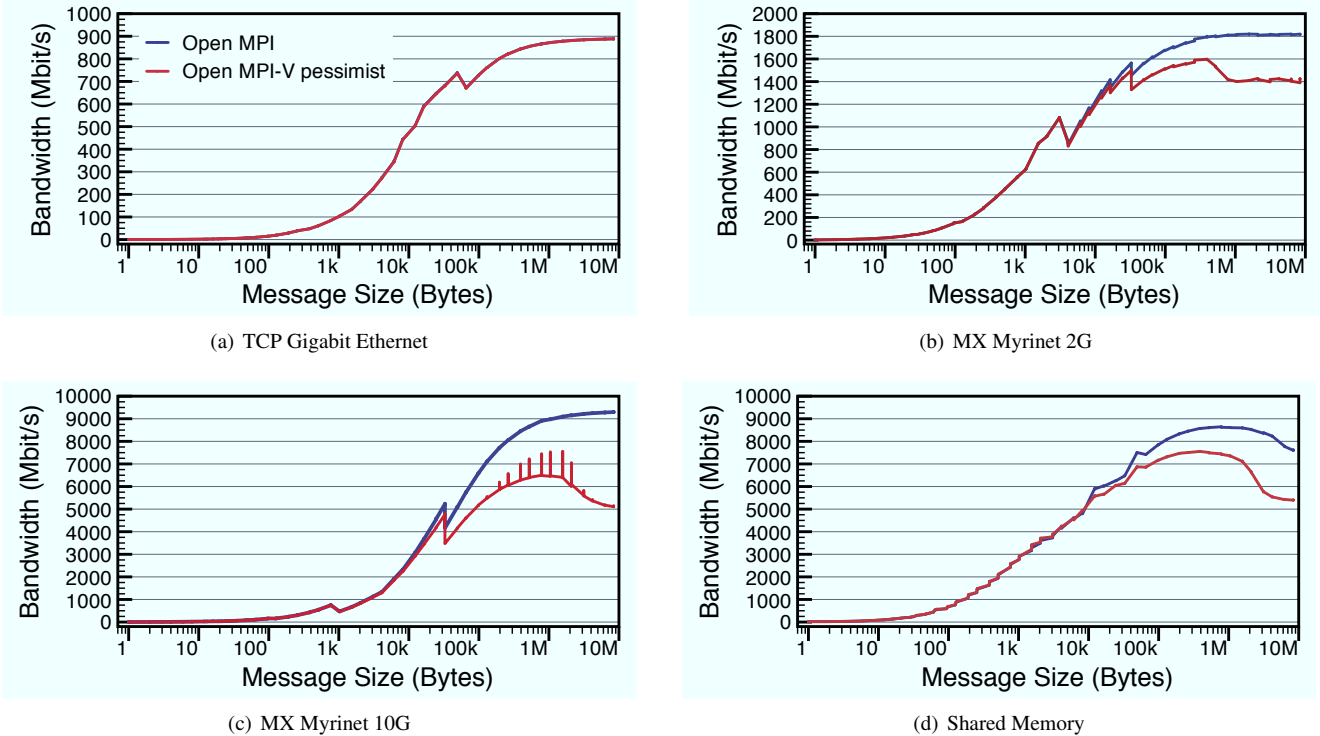


Figure 5. Ping-pong performance comparison between Open MPI and Open MPI-V pessimist on various networks.

the problem size. Similarly, for CG and LU (figures 6(c) and 6(d)), the scalability when the number of processes increase follows exactly the same trend for Open MPI and Open MPI-V. For up to 128 nodes, the scalability of the proposed message logging approach is excellent, regardless of the use of non-deterministic events by the application.

5 Related Works

Fault tolerance can be managed fully by the application [14, 4]; however software engineering cost can be greatly decreased by integrating fault tolerant mechanisms at the communication middleware level. FT-MPI [7, 8] aims at helping application to express its failure recovery policy by taking care of rebuilding internal MPI data structures (communicators, rank, etc.) and triggering user provided callbacks to restore a coherent application state when failure occurs. Though this approach is very efficient to minimize the cost of failure recovery technique, it still adds a significant level of complexity to the design of the application code.

Automatic fault tolerant MPI libraries are totally hiding failures from the application, thus avoiding any modification of the user’s code. A good review of the various techniques used to automatically ensure the successful recovery

of distributed applications from a checkpoint set is provided by [6].

Consistent recovery can be achieved automatically by building a coordinated checkpoint set where there exist no orphan message (with the Chandy & Lamport algorithm [3], CIC [10] or blocking the application until channels are empty). The blocking checkpointing approach have been used in LAM/MPI [15] while Chandy & Lamport algorithm have been used in CoCheck [17], MPICH-Vcl [2] and Déjà-vu. In all coordinated checkpoint techniques, the only consistent recovery line is when every process, including non failed ones, restart from checkpoint. The message logging model we propose does not have this requirement, which according to [12] allows for faster recovery.

Another way to ensure automatic consistent recovery is to use message logging. Manetho [5], Egida [13] and MPICH-V [1] are using several flavors of message logging (optimistic, pessimistic and causal). All are relying on the classical message logging model, as a consequence they are hooked into the lowest MPI level. Compared to our current work, they cannot distinguish between deterministic and non-deterministic events and introduce some extra memory copies leading to performance penalty on recent high throughput networks.

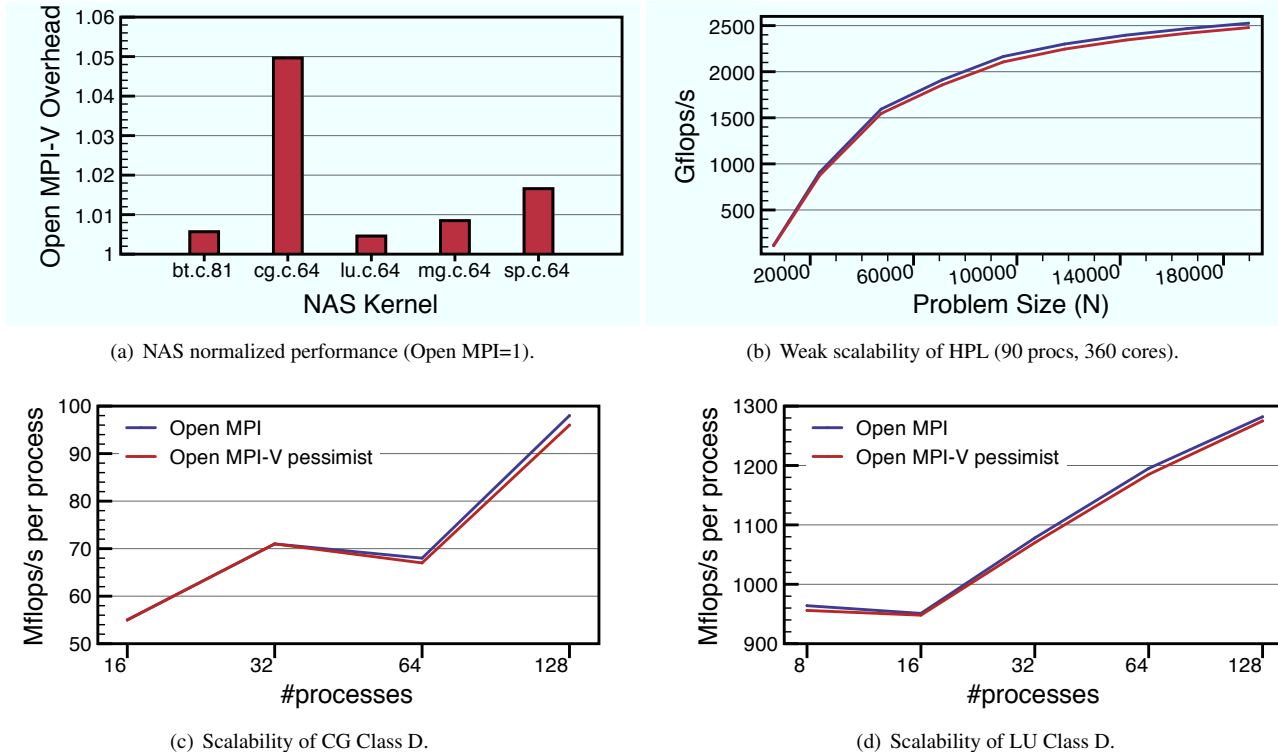


Figure 6. Application behavior comparison between Open MPI and Open MPI-V pessimist on Myrinet 10G.

6 Conclusion

In this paper we introduced a refinement of message logging model intended to reduce the raw overhead of this fault tolerant protocol. Unlike the classical model, it does not consider the message delivery as a single atomic step. Instead a message may generate two kinds of events: matching events at the beginning of any source receptions, and deliver events to count the number of time a message has been involved in a non-deterministic probe before delivery. Advantages of this model are 1) better fitting the actual MPI communication pattern 2) removing the need for an intermediate copy of message 3) allowing implementation of fault tolerant mechanisms at a higher level of the software hierarchy and then distinguish between non-deterministic and deterministic events.

We implemented a pessimistic message logging algorithm according to this model in Open MPI and compared its performance to the previous reference implementation of pessimistic message logging MPICH-V2. Results outlines a drastically lower cost of the fault tolerant framework. Thanks to the removal of intermediate message copies, Open MPI-V latency is 10.5 times better than MPICH-V2 in the worst case while bandwidth is multiplied by 4. Furthermore, because of the better detection of deterministic

events, most common types of messages do not have message logging overhead, leading to a 35 times better latency.

As a consequence, uncoordinated checkpointing results in less than 5% overhead on application performance, while scalability both in terms of number of nodes and data volume is close to perfect. This is a major improvement compared to previous uncoordinated fault tolerant approaches.

Future works

Direct consequence of our study is to try to eliminate the remaining cost from sender-based message logging. We plan to improve the pipelining between the fragment emission into the network and the sender-based copy to increase overlap and improve the cache reuse. Though it is impossible to reduce sender-based overhead on shared memory, the failure of a single core usually strongly correlates with the breakdown of the entire computing node. Therefore, coordinating checkpoint inside the same node and disabling sender-based could totally remove this cost while retaining the better scalability of the uncoordinated approach from a node perspective.

Next, we plan to better characterize the source of non-deterministic events. Those events may be generated whether by the numerical algorithm itself using any source

receptions or non-deterministic probes, or by the implementation of collective communications implementation over point-to-point. In this second case, some collaborative mechanisms may be involved to better reduce the cost of semantically determinist messages. With deeper modifications, we could even envision to disable completely logging during collective communication by adding some collective global success notification, allowing for replaying collectives as a whole instead of individual point-to-point messages.

One of the weaknesses of message logging, when compared to coordinated checkpoint, is a higher failure free overhead. Because it has been greatly improved by our work, the relative performance ordering of those two protocols could have changed. We plan next to make a comprehensive comparison between improved message logging and coordinated checkpoint, both in terms of failure free overhead, recovery speed and resiliency.

Last we could investigate the consequence of using application threads on the piecewise deterministic assumption, a scenario which will be more common with the dominance of multi-core processors. Because of a different ordering of MPI operations the unique request identifier and the probe clock may vary during recovery. Some mechanism may be designed to ensure the same identifiers are assigned during replay of threaded applications.

References

- [1] A. Bouteiller, T. Herault, G. Krawezik, P. Lemarinier, and F. Cappello. MPICH-V project: a multiprotocol automatic fault tolerant MPI. volume 20, pages 319–333. SAGE Publications, Summer 2006.
- [2] A. Bouteiller, P. Lemarinier, G. Krawezik, and F. Cappello. Coordinated checkpoint versus message log for fault tolerant MPI. In *IEEE International Conference on Cluster Computing (Cluster 2003)*. IEEE CS Press, December 2003.
- [3] K. M. Chandy and L. Lamport. Distributed snapshots : Determining global states of distributed systems. In *Transactions on Computer Systems*, volume 3(1), pages 63–75. ACM, February 1985.
- [4] Z. Chen, G. E. Fagg, E. Gabriel, J. Langou, T. Angskun, G. Bosilca, and J. Dongarra. Fault tolerant high performance computing by a coding approach. In *PPoPP '05: Proceedings of the tenth ACM SIGPLAN symposium on Principles and practice of parallel programming*, pages 213–223, New York, NY, USA, 2005. ACM Press.
- [5] Elnozahy, Elmootazbellah, and Zwaenepoel. Manetho: Transparent rollback-recovery with low overhead, limited rollback and fast output. *IEEE Transactions on Computing*, 41(5), May 1992.
- [6] M. Elnozahy, L. Alvisi, Y. M. Wang, and D. B. Johnson. A survey of rollback-recovery protocols in message-passing systems. *ACM Computing Surveys (CSUR)*, 34(3):375 – 408, september 2002.
- [7] G. Fagg and J. Dongarra. FT-MPI : Fault tolerant MPI, supporting dynamic applications in a dynamic world. In *7th Euro PVM/MPI User's Group Meeting2000*, volume 1908 / 2000, Balatonfüred, Hungary, september 2000. Springer-Verlag Heidelberg.
- [8] G. E. Fagg, A. Bukovsky, and J. J. Dongarra. HARNESS and fault tolerant MPI. *Parallel Computing*, 27(11):1479–1495, October 2001.
- [9] E. Gabriel, G. E. Fagg, G. Bosilca, T. Angskun, J. J. Dongarra, J. M. Squyres, V. Sahay, P. Kambadur, B. Barrett, A. Lumsdaine, R. H. Castain, D. J. Daniel, R. L. Graham, and T. S. Woodall. Open MPI: Goals, concept, and design of a next generation MPI implementation. In *Proceedings, 11th European PVM/MPI Users' Group Meeting*, pages 97–104, Budapest, Hungary, September 2004.
- [10] J.-M. H elary, A. Mostefaoui, and M. Raynal. Communication-induced determination of consistent snapshots. *IEEE Transactions on Parallel and Distributed Systems*, 10(9):865–877, 1999.
- [11] L. Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565, 1978.
- [12] P. Lemarinier, A. Bouteiller, T. Herault, G. Krawezik, and F. Cappello. Improved message logging versus improved coordinated checkpointing for fault tolerant MPI. In *IEEE International Conference on Cluster Computing (Cluster 2004)*. IEEE CS Press, 2004.
- [13] S. Rao, L. Alvisi, and H. M. Vin. Egida: An extensible toolkit for low-overhead fault-tolerance. In *29th Symposium on Fault-Tolerant Computing (FTCS'99)*, pages 48–55. IEEE CS Press, 1999.
- [14] A. Roy-Chowdhury and P. Banerjee. Algorithm-based fault location and recovery for matrix computations on multiprocessor systems. *IEEE Trans. Comput.*, 45(11):1239–1247, 1996.
- [15] S. Sankaran, J. M. Squyres, B. Barrett, A. Lumsdaine, J. Duell, P. Hargrove, and E. Roman. The LAM/MPI checkpoint/restart framework: System-initiated checkpointing. In *Proceedings, LACSI Symposium*, Sante Fe, New Mexico, USA, October 2003.
- [16] Q. Snell, A. Mikler, and J. Gustafson. Netpipe: A network protocol independent performance evaluator. In *IASTED International Conference on Intelligent Information Management and Systems.*, June 1996.
- [17] G. Stellner. CoCheck: Checkpointing and process migration for MPI. In *Proceedings of the 10th International Parallel Processing Symposium (IPPS '96)*, Honolulu, Hawaii, April 1996. IEEE CS Press.
- [18] The IBM LLNL BlueGene/L Team. An overview of the BlueGene/L supercomputer. In *Supercomputing '02: Proceedings of the 2002 ACM/IEEE conference on Supercomputing*, pages 1–22, Los Alamitos, CA, USA, 2002. IEEE Computer Society Press.
- [19] The MPI Forum. MPI: a message passing interface. In *Supercomputing '93: Proceedings of the 1993 ACM/IEEE conference on Supercomputing*, pages 878–883, New York, NY, USA, 1993. ACM Press.