A Proposal for User-Level Failure Mitigation in the MPI-3 Standard

February 22, 2012

Abstract

This chapter describes a flexible approach providing process fault tolerance by allowing the application to react to failures while maintaining a minimal execution path in failure-free executions. The focus is on returning control to the application by avoiding deadlocks due to failures within the MPI library. No implicit, asynchronous error notification is required. Instead, functions are provided to allow processes to invalidate any communication object, thus preventing any process from waiting indefinitely on calls involving the invalidated objects. We consider the proposed set of functions to constitute a minimal basis which allows libraries and applications to increase the fault tolerance capabilities by supporting additional types of failures, and to build other desired strategies and consistency models to tolerate faults.

Chapter 17

Process Fault Tolerance

17.1 Introduction

MPI processes may fail at any time during execution. Long running and large scale applications are at increased risk of encountering process failures during normal execution. This chapter introduces the MPI features that support the development of applications and libraries that can tolerate process failures. The approach described in this chapter is intended to prevent the deadlock of processes while avoiding any impact on the failure-free execution of an application.

The expected behavior of MPI in case of a process failure is defined by the following statements: any MPI call that involves a failed process must not block indefinitely, but either succeed or raise an MPI error (see Section 17.2); asynchronous failure propagation is not required by the MPI standard, an MPI call that does not involve the failed process will complete normally. If an application needs global knowledge of failures, it can use the interfaces defined in Section 17.3 to explicitly propagate locally detected failures.

Advice to users. Many of the operations and semantics described in this chapter are only applicable when the MPI application has replaced the default error handler MPI_ERRORS_ARE_FATAL on, at least, MPI_COMM_WORLD. (End of advice to users.)

17.2 Failure Notification

This section specifies the behavior of an MPI communication call when failures happened on processes involved in the communication. A process is considered as involved in a communication if any of the following is true:

1. the operation is a collective call and the process appears in one of the groups on which the operation is applied;

2. the process is a named or matched destination or source in a point-to-point communication;

3. the operation is an MPI_ANY_SOURCE receive operation and the failed process belongs to the source group.

Therefore, if an operation does not involve a failed process (such as a point to point message between two non-failed processes), it must not return a process failure error.

Advice to implementers. It is a legitimate implementation to provide failure detection only for processes involved in an ongoing operation and postpone detection of other failures until necessary. Moreover, as long as an implementation can complete operations, it may choose to delay returning an error. Another valid implementation might choose to return an error to the user as quickly as possible. (End of advice to implementers.)

Note for the Forum The text of Page 65, lines 28-33 must be changed to allow MPI_IPROBE to set flag=true and return the appropriate status, if an error is detected during an MPI_IPROBE. MPI_PROBE is defined as behaving as MPI_IPROBE so it should be sufficient. Similarly, the same effort should be done for MPI_MPROBE and MPI_MRECV.

Non-blocking operations must not return an error about process failures during initialization. All process failure errors are postponed until the corresponding completion function is called.

17.2.1 Point-to-Point and Collective Communication

When a failure prevents the MPI implementation from completing a point-to-point communication, the communication is marked as completed with an error of class MPI_ERR_PROC_FAILED. Further point-to-point communication with the same process on this communicator must also return MPI_ERR_PROC_FAILED.

MPI libraries can not determine if the completion of an unmatched reception operation of type MPI_ANY_SOURCE can succeed when one of the potential senders has failed. If the reception has matched internally, a subsequent process failure on this operation must return an error of class MPI_ERR_PROC_FAILED (as if it were a named receive). Otherwise, the communication is marked with an error of class MPI_ERR_PENDING and the completion operation returns. If the operation worked on a request allocated by a nonblocking communication call, then the request is still valid and pending. To acknowledge a failure and discover which processes failed, the user should call MPI_COMM_FAILURE_ACK.

When a collective operation cannot be completed because of the failure of an involved process, the collective operation eventually returns an error of class MPI_ERR_PROC_FAILED. The content of the output buffers is undefined.

Advice to users. Depending on how the collective operation is implemented and when a process failure occurs, some participating alive processes may raise an error while other processes return successfully from the same collective operation. For example, in MPI_Bcast, the root process is likely to succeed before a failed process disrupts the operation, resulting in some other processes returning an error. However, it is noteworthy that for non-rooted collective operations on an intracommunicator, processes failing before entering the operation provoke all surviving ranks to return MPI_ERR_PROC_FAILED. Similarly, on an intercommunicator, processes of the remote group failing before entering the operation have the same effect on all surviving ranks of the local group. (End of advice to users.)

Advice to users. Note that communicator creation functions (like MPI_COMM_DUP or MPI_COMM_SPLIT) are collective operations. As such, if a failure happened during the call, an error might be returned to some processes while others succeed and obtain a new communicator. It is the responsibility of the user to ensure that all involved processes have a consistent view of the communicator creation, if needed. A conservative solution is to invalidate the parent communicator if the operation fails, otherwise call an MPI_Barrier on the parent communicator and invalidate the new communicator if the MPI_Barrier fails. (End of advice to users.)

17.2.2 Dynamic Process Management

Dynamic process management functions require some additional semantics from the MPI implementation as detailed below.

1. If the MPI implementation decides to return an error related to process failure at the root process of MPI_COMM_CONNECT or MPI_COMM_ACCEPT, the root processes of both intracommunicators must return an error of class MPI_ERR_PROC_FAILED (unless required to return MPI_ERR_INVALIDATED as defined by 17.3.1).

2. If the MPI implementation decides to return an error related to process failure at the root process of MPI_COMM_SPAWN, no spawned processes should be able to communicate on the created intercommunicator.

Advice to users. As with communicator creation functions, it is possible that if a failure happens during dynamic process management calls, an error might be returned to some processes while others succeed and obtain a new communicator. (End of advice to users.

17.2.3 One-Sided Communication

As with all non-blocking operations, one-sided communication operations should delay all failure notification to their synchronization calls and return MPI_ERR_PROC_FAILED (see Section 17.2). If the implementation decides to return an

error related to process failure from the synchronization function, the epoch behavior is unchanged from the definitions in Section 11.4. Similar to collective operations over MPI communicators, it is possible that some processes could have detected the failure and returned MPI_ERR_PROC_FAILED, while others could have returned MPI_SUCCESS.

The status of the operations occurring during the epoch which completed with an error related to process failure are detailed below.

1. For MPI_WIN_FENCE operations which return an error class related to process failure, MPI makes no guarantee about the state of the destination memory.

2. If a failure is to be reported during active target communication functions MPI_WIN_COMPLETE or MPI_WIN_WAIT (or the non-blocking equivalent MPI_WIN_TEST), the epoch is considered completed and all operations not involving the failed processes are completed successfully.

3. If the target rank has failed, MPI_WIN_LOCK and MPI_WIN_UNLOCK operations return an error of class MPI_ERR_PROC_FAILED. If the owner of a lock has failed, the lock can not be acquired again and all subsequent operations on the lock must fail with an error of class MPI_ERR_PROC_FAILED.

17.2.4 I/O

Due to the fact that MPI I/O writing operations can choose to buffer data to improve performance, for the purposes of process fault tolerance, all I/O data writing operations are treated as operations which synchronize on MPI_FILE_SYNC. Therefore (as described for non-blocking operations in Section 17.2), failures may not be reported during an MPI_FILE_WRITE_XXX operation but must be reported by the next MPI_FILE_SYNC. In this case, all alive processes must uniformly return either success or a failure of class MPI_ERR_PROC_FAILED.

Once MPI has returned an error of class MPI_ERR_PROC_FAILED, it makes no guarantees about the position of the file pointer following any previous operations. The only way to know the current location by calling the local functions MPI_FILE_GET_POSITION or MPI_FILE_GET_POSITION_SHARED.

17.3 Failure Mitigation Functions

17.3.1 Communicator Functions

MPI provides no guarantee of global knowledge of a process failure. Only processes involved in a communication with the failed process are guaranteed to eventually detect its failure (see Section 17.2). If global knowledge is required, MPI provides a function to globally invalidate a communicator.

MPI_COMM_INVALIDATE(comm)

IN	comm	**communicator (handle)**

This function eventually notifies all processes in the groups (local and remote) associated with the communicator *comm* that this communicator is now considered invalid. An invalid communicator preempts any non-local MPI calls on *comm*, with the exception of MPI_COMM_SHRINK. A communicator becomes invalid as soon as:

1. MPI_COMM_INVALIDATE is locally called on it;
2. Or any MPI function returned MPI_ERR_INVALIDATED (or such error field was set in the status pertaining to a request on this communicator).

Once a communicator has been invalidated, all subsequent non-local calls on that communicator, with the exception of MPI_COMM_SHRINK and MPI_AGREEMENT, are considered local and must return with an error of class MPI_ERR_INVALIDATED. If an implementation chooses to implement MPI_COMM_FREE as a local operation (see Page 209 Line 1), it is allowed to succeed on an invalidated communicator.

Note for the Forum The text of Page 208 lines 39-43 must be amended to provide the following advice to implementers.

The implementation should make a best effort to free an invalidated communicator locally and return MPI_SUCCESS. Otherwise, it must return MPI_ERR_INVALIDATED.

Note for the Forum The text of Page 208 lines 39-48 must be amended to provide the following advice to users.

Because MPI_COMM_FREE resets the MPI_Errhandler of a communicator to MPI_ERRORS_ARE_FATAL, fault tolerant applications should complete all pending communications before calling MPI_COMM_FREE.

MPI_COMM_SHRINK(comm, newcomm)

IN	comm	**communicator (handle)**
OUT	newcomm	**communicator (handle)**

This function creates a new intra or inter communicator *newcomm* from the invalidated intra or inter communicator *comm* respectively by excluding its failed processes as detailed below. It is erroneous MPI code to call MPI_COMM_SHRINK on a communicator which has not been invalidated (as defined above) and will return an error of class MPI_ERR_ARG.

This function must not return an error due to process failure (error classes MPI_ERR_PROC_FAILED and MPI_ERR_INVALIDATED). Upon successful completion, an agreement is made among living processes to determine the group of failed processes. This group includes at least all processes whose failure has been notified to the user. The call is semantically equivalent to MPI_COMM_SPLIT where living processes participate with the same color and a key

equal to their rank in *comm* and failed processes implicitly contribute MPI_UNDEFINED.

Advice to users. This call does not guarantee that all processes in *newcomm* are alive. Any new failure will be detected in subsequent MPI calls. (End of advice to users.)

MPI_COMM_FAILURE_ACK(comm)

IN	comm	**communicator (handle)**

This local function gives the users a way to acknowledge all locally notified failures on *comm*. After the call, operations that would have returned MPI_ERR_PENDING due to process failure (see Section 17.2.1) proceed without further reporting acknowledged failures.

Advice to users. It is erroneous MPI code to call a collective communication on a communicator with acknowledged failures. Such calls will continue to return an error of class MPI_ERR_PROC_FAILED as defined in Section 17.2.1. To reliably use collective operations on a communicator with failed processes, the communicator should first be invalidated using MPI_COMM_INVALIDATE and then a new communicator should be created using MPI_COMM_SHRINK. (End advice to users.)

MPI_COMM_FAILURE_GET_ACKED(comm, failedgroup)

IN	comm	**communicator (handle)**
OUT	failedgroup	**group (handle)**

This local function returns the group *failedgroup* of processes from the communicator *comm* which have been locally acknowledged as failed by preceding calls to MPI_COMM_FAILURE_ACK.

MPI_AGREEMENT(comm, flag)

IN	comm	**communicator (handle)**
INOUT	flag	**boolean flag**

This function performs a collective operation among all living processes in *comm*. On completion, all living processes must agree to set the value of *flag* to the result of a logical *'AND'* operation over the contributed values. This function must not return an error due to process failure (error classes MPI_ERR_PROC_FAILED and MPI_ERR_INVALIDATED), and failed processes do not contribute to the operation.

If *comm* is an intercommunicator, the return value is uniform over both groups and (if applicable) the value of *flag* is a logical *'AND'* operation over the values contributed by the remote group (where failed processes do not contribute to the operation).

Advice to users. MPI_AGREEMENT maintains its collective meaning even if the *comm* is invalidated.

17.3.2 One-Sided Functions

MPI_WIN_INVALIDATE (win)

IN	win	**window (handle)**

This function eventually notifies all ranks within the window *win* that this window is now considered invalid. An invalid window preempts any non-local MPI calls on *win*. Once a window has been invalidated, all subsequent non-local calls on that window are considered local and must fail with an error of class MPI_ERR_INVALIDATED.

MPI_WIN_GET_FAILED(win, failedgroup)

IN	win	**window (handle)**
OUT	failedgroup	**group (handle)**

This local function returns the group *failedgroup* of processes from the window *win* which are locally known to have failed.

Advice to users. MPI makes no assumption about asynchronous progress of the failure detection. A valid MPI implementation may choose to only update the group of locally known failed processes when it enters a synchronization function. (End advice to users.)

Advice to users. It is possible that only the calling process has detected the reported failure. If global knowledge is necessary, processes detecting failures should use the call MPI_WIN_INVALIDATE. (End advice to users.)

17.3.3 I/O Functions

MPI_FILE_INVALIDATE (fh)

IN	fh	**file (handle)**

This function eventually notifies all ranks within file *fh* that this file is now considered invalid. An invalid file preempts any non-local completion calls MPI calls on *file* (see Section 17.2.4). Once a file has been invalidated, all subsequent non-local calls on the file must fail with an error of class MPI_ERR_INVALIDATED.

17.4 Error Codes and Classes

MPI_ERR_PROC_FAILED	A process in the operation has failed (a fail-stop failure).
MPI_ERR_INVALIDATED	The communication object used in the operation was invalidated.