

# 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 ~~return~~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 ~~completes~~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

~~In this section, we specify~~

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 ~~either:~~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; 1  
2
3. the operation is an MPI\_ANY\_SOURCE ~~reception and the~~ receive operation and the failed process belongs to the source group. 3  
4

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. 5  
6  
7

*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.) 8  
9  
10  
11  
12  
13

*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. 14  
15  
16  
17  
18

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. 19  
20  
21  
22  
23

In case of a failure of involved processes, the completion of an unmatched MPI reception from MPI\_ANY\_SOURCE is undecidable. Such communication is marked with an error of class MPI\_ERR\_PENDING and the completion operation returns. If the operation worked on a request, and the request was allocated by a nonblocking communication call, then the request is still valid and pending. To acknowledge this failure and discover which processes failed, the user should call MPI\_COMM\_FAILURE\_ACK. 24  
25  
26  
27  
28  
29  
30

*Advice to implementers.* When an operation on MPI\_ANY\_SOURCE has matched internally, a subsequent process failure on this operation must return an error of class MPI\_ERR\_PROC\_FAILED (like a named receive). (End of advice to implementers.) 31  
32  
33  
34

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. 35  
36  
37

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. 38  
39  
40

*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 intra-communicator, processes failing before entering the operation provoke all surviving ranks to return MPI\_ERR\_PROC\_FAILED. Similarly, on an inter-communicator, 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. A conservative method 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.)

### 17.2.1 Dynamic Process Management

Dynamic process management function 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.

### 17.2.2 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. 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 Chapter 11. Similar to collective operations over MPI Communicators, it is possible that some processes could have detected the failure and returned with MPI\_ERR\_PROC\_FAILED, while others could have returned successfully. The status of the operations occurring during the epoch completed with an error related to process failure is detailed below.

1. For MPI\_WIN\_FENCE operations following a failure, MPI makes no guarantees about the state of the destination memory.
2. If a failure is to be reported during active target communication functions MPI\_WIN\_COMPLETE and 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. MPI\_WIN\_LOCK and MPI\_WIN\_UNLOCK operations return an error of class MPI\_ERR\_PROC\_FAILED if the target rank has 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.

### 17.2.3 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, failures may not be reported during an MPI\_FILE\_WRITE\_XXX operation, however they 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 Handling Mitigation 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. If global knowledge is required, MPI provides a function to globally invalidate a communicator.

### 17.3.1 Communicator Functions

**MPI\_COMM\_INVALIDATE( comm )**

IN comm communicator (handle)

This function eventually notifies all processes ~~of all groups within 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 ~~and MPI.COMM.FREE~~. 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 (Page ??? line ???)~~, are considered local and must return with an error of class MPI\_ERR\_INVALIDATED. If an implementation chooses to implement MPI.COMM.FREE, ~~must fail with an error of class~~ as a local operation (see Page 209 Line 1), it is allowed to succeed.

*Note for the Forum* The text of Page 208 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.

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

~~Because MPI.CommOMM.fFree is also allowed to complete after the communicator has been invalidated. It is the user's responsibility to ensure that there are no pending messages remaining on the invalidated communicator. If the communicator is not invalid but there are known failed processes inside, it is up to the user to clean up its own requests~~ REE 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 partitions the group associated with *comm* into two disjoint subgroups: the group of failed processes and the group of alive processes. A new communicator is created for the group of alive processes and returned as *newcomm*. This function is illegal on a communicator which has not been invalidated and will return an error of class MPI\_ERR\_ARG. This call returns the same value on all ranks, even if failures happen during the call. If the return

is MPI\_SUCCESS, 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 an agreement is made among living processes to determine the group of failed processes whose implicit contribution is MPI\_UNDEFINED.

*Advice to users.* This call does not guarantee that all processes in *newcomm* are alive, but that all processes in *newcomm* agreed on a consistent set that includes at least the union of processes locally known to have failed before the call. 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](#) proceed without further reporting acknowledged failures.

*Advice to users.* It is an incorrect MPI code to call a collective communication on a communicator with acknowledged failures. 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 were 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*. It returns the same return code on all ranks, even if failures happen during the call. Upon successful completion, all living processes uniformly set the value of *flag* to the result of a logical 'AND' operation over the value contributed from all ranks. Any failed process is assumed to have participated with *flag* = *false*.

If *comm* is an inter-communicator, 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 contribute with *flag = false*).

### 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 17.2.3). 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

1

`MPI_ERR_PROC_FAILED`

A process in the operation has failed (a fail-stop failure).

`MPI_ERR_INVALIDATED`

The communicator [object](#) used in the operation was invalidated.

2