

D R A F T

Document for a Standard Message-Passing Interface

Message Passing Interface Forum

November 20, 2013

This work was supported in part by NSF and ARPA under NSF contract CDA-9115428 and Esprit under project HPC Standards (21111).

This is the result of a LaTeX run of a draft of a single chapter of the MPIF Final Report document.

Chapter 17

Process Fault Tolerance

17.1 Introduction

~~Long running and large scale applications are at increased risk of encountering process failures during normal execution. We consider a process failure~~ In distributed systems with numerous or complex components, the risk that the fault of a component manifests as a process failure that disrupts the normal execution of a long running application is serious. A process failure is a common ultimate outcome for many hardware, network or software faults that cause a process to crash; It can be more formally defined as a fail-stop failure; failed processes become: the failed process becomes permanently unresponsive to communications. This chapter introduces ~~the~~ MPI features that support the development of applications ~~and libraries, libraries, and programming languages~~ that can tolerate process failures. The ~~approach described in this chapter is intended to prevent the deadlock of processes while avoiding impact on the failure-free execution of an application.~~ primary goal is to specify error classes and interfaces that permit users to continue simple MPI communication operations after failures have impacted the execution, and rebuild MPI objects (communicators, files, etc.) as needed to restore the full capability of MPI to carry elaborate communication operations (like collective communications.) This specification does not include mechanisms to restore the lost data from failed processes; the literature is rich with wildly varied fault tolerance techniques that the users may employ at their discretion, including checkpoint-restart, algorithmic dataset recovery, or continuation ignoring failed processes. All these fault tolerance approaches benefit from, and often require, the definitions and interfaces specified in this chapter to resume communicating after a failure.

The expected behavior of MPI in the case of a process failure is defined by the following statements: any MPI operation that involves a failed process must not block indefinitely, but either succeed or raise an MPI exception (see Section 17.2); an MPI operation that does not involve ~~the a~~ failed process will complete normally, unless interrupted by the user through provided functionality. Exceptions only indicate the local impact of the failure on an operation. Asynchronous failure propagation is not ~~required~~ guaranteed or required and users must exercise caution when reasoning on the set of ranks where a failure has been detected and raised an exception. If an application needs global knowledge of failures, it can use the interfaces defined in Section 17.3 to explicitly propagate the notification of locally detected failures.

~~An-~~

The typical usage pattern on some reliable machines may not command the need for fault tolerance. An MPI implementation that does not tolerate process failures must provide the interfaces and semantics defined in this chapter as long as no failure occurred. It must never raise an exception of class MPI_ERR_PROC_FAILED or MPI_ERR_PENDING because of to report a process failure.

~~This chapter does not define process failure semantics for the operations specified in Chapters , therefore they remain undefined by the MPI standard~~Fault tolerant applications using the interfaces defined in this chapter must compile, link and run successfully in failure free executions.

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 operation when failures occur on processes involved in the communication. A process is considered involved in a communication if any of the following is true:

1. the operation is collective and the process appears in one of the groups of the associated communication object;
2. the process is a specified 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 involving a failed process must always complete in a finite amount of time (possibly by raising a process failure exception). If an operation does not involve a failed process (such as a point-to-point message between two non-failed processes), it must not raise a process failure exception.

Advice to implementors. A correct MPI implementation may 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 raising an ~~error~~exception. Another valid implementation might choose to raise an ~~error~~exception as quickly as possible. (*End of advice to implementors.*)

When a communication operation raises an exception related to process failure, it may not satisfy its specification. In particular, a synchronizing operations may not have synchronized, and the content of the output buffers, targeted memory or output parameters is undefined. Operations excepting this rule are explicitly stated in the remainder of this chapter.

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

17.2.1 Startup and Finalize

Advice to implementors. If a process fails during MPI_INIT but its peers are able to complete the MPI_INIT successfully, then a high quality implementation will return MPI_SUCCESS and delay the reporting of the process failure to a subsequent MPI operation. (*End of advice to implementors.*)

MPI_FINALIZE will complete successfully even in the presence of process failures.

Advice to users. ~~Considering Example 8.10 in Section 8.7, the process with rank 0 in may have failed before, during, or after the call to MPI raises exceptions only before MPI_FINALIZE. MPI only provides failure detection capabilities up to when~~ is invoked and ~~thereby~~ provides no support for fault tolerance during or after MPI_FINALIZE. Applications are encouraged to implement all rank-specific code before the call to MPI_FINALIZE ~~to handle the case where process~~. Considering Example 8.10 in Section 8.7, the process with rank 0 in MPI_COMM_WORLD fails may have failed before, during, or after the call to MPI_FINALIZE, possibly leading to this code never being executed. (*End of advice to users.*)

17.2.2 Point-to-Point and Collective Communication

An MPI implementation raises exceptions of the following error classes to notify users that a point-to-point communication operation could not complete successfully because of the failure of involved processes:

- MPI_ERR_PENDING indicates, for a non-blocking communication, that the communication is a receive operation from MPI_ANY_SOURCE and no matching send has been posted, yet a potential sender process has failed. Neither the operation nor the request identifying the operation are completed. Note that the same error class is also used in status when another communication raises an exception during the same operation (as defined in Section 3.7.5).
- In all other cases, the operation raises an exception of class MPI_ERR_PROC_FAILED to indicate that the failure prevents the operation from following its failure-free specification. If there is a request identifying the point-to-point communication, it is completed. Future point-to-point communication with the same process on this communicator must also raise MPI_ERR_PROC_FAILED.

Advice to users.

To acknowledge a failure and discover which processes failed, the user should call MPI_COMM_FAILURE_ACK (as defined in Section 17.3.1).

(*End of advice to users.*)

When a collective operation cannot be completed because of the failure of an involved process, the collective operation raises an ~~error-exception~~ of class MPI_ERR_PROC_FAILED.

Advice to users.

Depending on how the collective operation is implemented and when a process failure occurs, some participating alive processes may raise an exception while other processes

return successfully from the same collective operation. For example, in MPI_BCAST, the root process may succeed before a failed process disrupts the operation, resulting in some other processes raising an ~~error-exception~~.

However, it is noteworthy that ~~for collective operations on an intracommunicator in which all processes contribute to the result and all processes receive the result, processes which do not enter the operation due to process failure provoke all surviving ranks to raise . Similarly, for the same collective operations on an intercommunicator, a process in the remote group which failed before entering the operation has the same effect on all surviving ranks of the local group~~ some operation semantic may force raising an exception at all ranks, if a process fails before entering the operation. As an example, on an intracommunicator, if a process raises an exception from an operation matching a particular process, the process receiving that exception can then assume that in a following MPI_BARRIER on this communicator, all ranks will raise an exception MPI_ERR_PROC_FAILED, because the participating process is known to have failed before entering the barrier.

(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-exception~~ might be raised at some processes while others succeed and obtain a new communicator. While it is valid to communicate between processes which succeeded to create the 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~~ the possibility to match posted operations. A conservative solution is ~~to have each process either revoke (see check the global outcome of the communicator creation function with MPI_COMM_AGREE (defined in Section 17.3.1) the parent communicator if the operation fails, or call an on the parent communicator and then revoke the new communicator if the fails.~~

, as illustrated in Example 17.1. *(End of advice to users.)*

~~When a communication operation raises an exception related to process failure , the content of the output buffers is undefined.~~

After a process failure, MPI_COMM_FREE, as with all other collective operations, may not complete successfully at all ranks. For any rank which receives the return code MPI_SUCCESS, the behavior is defined as in Section 6.4.3. If a rank raises a process failure exception (MPI_ERR_PROC_FAILED or MPI_ERR_REVOKED), the implementation makes no guarantee about the success or failure of the MPI_COMM_FREE operation remotely; however it still attempts to clean up any local data used by the Window object. This will be signified by returning MPI_COMM_NULL only when the object has successfully been freed locally.

17.2.3 Dynamic Process Management

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

1. If the MPI implementation raises an ~~error-exception~~ related to process failure to the root process of MPI_COMM_CONNECT or MPI_COMM_ACCEPT, at least the

root processes of both intracommunicators must raise the same ~~error-exception~~ of class MPI_ERR_PROC_FAILED (unless required to raise MPI_ERR_REVOKED as defined by 17.3.1). The same is true if the implementation ~~returns an error to raises an exception at~~ any process in MPI_COMM_JOIN.

2. If the MPI implementation raises an ~~error-exception~~ related to process failure to the root process of MPI_COMM_SPAWN or MPI_COMM_SPAWN_MULTIPLE, 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 operations, an ~~error-exception~~ might be raised at some processes while others succeed and obtain a new communicator. (*End of advice to users.*)

17.2.4 One-Sided Communication

One-Sided communication operations must provide failure notification in their synchronization operations which may raise an ~~error-exception~~ due to process failure (see Section 17.2). If the implementation does not raise an ~~error-exception~~ related to process failure in the synchronization function, the epoch behavior is unchanged from the definitions in Section 11.5. As with collective operations over MPI communicators, it is possible that some processes have detected a failure and raised MPI_ERR_PROC_FAILED, while others returned MPI_SUCCESS. Once the implementation ~~returns an error raises an exception~~ related to process failure on a specific window in a synchronization function, all subsequent operations on that window ~~much also return an error code must also raise an exception~~ related to process failure.

Unless specified below, the state of memory targeted by any process in an epoch in which operations raised an ~~error-exception~~ related to process failure is undefined, with the exception of memory targeted by remote read operations (and operations which are semantically equivalent to read operations, such as an MPI_ACCUMULATE with MPI_NO_OP as the operation). All other window locations are valid.

~~If a failure is to be reported during active target communication functions-~~

~~If an exception is raised from an active target synchronization operation~~

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 must complete successfully.

~~If the-~~

MPI_WIN_LOCK ~~and~~ MPI_WIN_UNLOCK ~~may raise~~ MPI_ERR_PROC_FAILED ~~when any process in the window has failed. An implementation cannot block indefinitely in a correct program waiting for a lock to be acquired; If the owner of the lock has failed, some other process trying to acquire the lock either succeeds or raises an exception of class MPI_ERR_PROC_FAILED. If the~~ target rank has failed, MPI_WIN_LOCK and MPI_WIN_UNLOCK operations ~~raise an error must raise an exception~~ of class MPI_ERR_PROC_FAILED. The lock cannot be acquired again at any target in the window, and all subsequent operations on the lock must raise MPI_ERR_PROC_FAILED. ~~As with communicator-based operations, an implementation cannot block indefinitely in a correct program waiting for a lock to be acquired. If the owner of the lock has failed, some other process should be notified via the return code --~~

Advice to implementors. If a non-target rank in the window fails, ~~it is possible that the implementation will~~ a high quality implementation may be able to mask such ~~an error~~ a fault inside the locking algorithm and continue to allow the remaining ranks to acquire the lock without raising errors. (*End of advice to implementors.*)

It is possible that request-based RMA operations complete successfully (via operations such as MPI_TEST or MPI_WAIT) while the enclosing epoch completes by raising an exception due to a process failure. In this scenario, the local buffer is valid but the remote targeted memory is undefined.

After a process failure ~~has been detected~~, MPI_WIN_FREE, as with all other collective operations, may not complete successfully ~~on at~~ all ranks. For any rank which receives the return code MPI_SUCCESS, the behavior is defined as in Section 11.2.5. If a rank ~~receives a return code related to process failure~~ raises a process failure exception (MPI_ERR_PROC_FAILED or MPI_ERR_REVOKED), the implementation makes no guarantee about the success or failure of the MPI_WIN_FREE operation remotely, ~~though it should still attempt~~; however it still attempts to clean up any local data used by the Window object. This will be signified by returning MPI_FILE_NULL only when the object has successfully been freed locally.

~~It is possible that request-based RMA operations complete successfully (via operations such as or) while the enclosing epoch completes by raising error due to process failure. In this scenario, the local buffer is valid but the remote targeted memory is undefined.~~

17.2.5 I/O

I/O backend failure error classes and their consequences are defined in Section 13.7. The following section defines the behavior of I/O operations when MPI process failures prevent their successful completion.

~~Since collective I/O operations may not synchronize with other processes, process failures may not be reported during a collective I/O operation.~~ If a process failure prevents a file operation from completing, an MPI exception of class MPI_ERR_PROC_FAILED is raised. Once an MPI implementation has raised an error exception of class MPI_ERR_PROC_FAILED, the state of the file pointer involved in the operation which ~~returned the error code raised the exception~~ is *undefined*.

Advice to users.

Since collective I/O operations may not synchronize with other processes, process failures may not be reported during a collective I/O operation. Users are encouraged to use MPI_COMM_AGREE on a communicator containing the same group as the file handle, when they need to deduce the completion status of collective operations on file handles and maintain a consistent view of file pointers. The file pointer can be reset using ~~with the~~ MPI_FILE_SEEK with the MPI_SEEK_SET update mode.

(*End of advice to users.*)

After a process failure ~~has been detected~~, MPI_FILE_CLOSE, as with all other collective operations, may not complete successfully ~~on at~~ all ranks. For any rank which receives the return code MPI_SUCCESS, the behavior is defined as in Section ~~11.2.5~~ 13.2.2. If a rank ~~receives a return code related to process failure~~ raises a process failure exception

(MPI_ERR_PROC_FAILED or MPI_ERR_REVOKED), the implementation makes no guarantee about the success or failure of the MPI_FILE_CLOSE operation remotely, ~~though it should still attempt~~; however it still attempts to clean up any local data used by the File objecthandle. This will be signified by returning MPI_FILE_NULL only when the object has successfully been freed locally.

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 operation 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 revoke a communicator at all members.

MPI_COMM_REVOKE(comm)

IN comm communicator (handle)

int MPI_Comm_revoke(MPI_Comm comm)

MPI_COMM_REVOKE(COMM, IERROR)

INTEGER COMM, IERROR

This function notifies all processes in the groups (local and remote) associated with the communicator `comm` that this communicator is now considered revoked. This function is not collective and therefore does not have a matching call on remote processes. ~~It is erroneous to call on a communicator for which no operation raised an MPI exception related to process failure.~~ All alive processes belonging to `comm` will be notified of the revocation despite failures. The revocation of a communicator completes any non-local MPI operations on `comm` by raising an ~~error-exception~~ of class MPI_ERR_REVOKED, with the exception of MPI_COMM_SHRINK and MPI_COMM_AGREE (and its nonblocking equivalent). A communicator becomes revoked as soon as:

1. MPI_COMM_REVOKE is locally called on it;
2. Any MPI operation raised an ~~error-exception~~ of class MPI_ERR_REVOKED because another process in `comm` has called MPI_COMM_REVOKE.

Once a communicator has been revoked, all subsequent non-local operations on that communicator, with the exception of MPI_COMM_SHRINK and MPI_COMM_AGREE (and its nonblocking equivalent), are considered local and must complete by raising an ~~error-exception~~ of class MPI_ERR_REVOKED.

Advice to users. High quality implementations are encouraged to do their best to free resources locally when the user calls free operations on revoked communication objects, or communication objects containing failed processes. (*End of advice to users.*)

```

1 MPI_COMM_SHRINK( comm, newcomm )
2     IN      comm      communicator (handle)
3     OUT     newcomm   communicator (handle)
4
5
6 int MPI_Comm_shrink(MPI_Comm comm, MPI_Comm* newcomm)
7 MPI_COMM_SHRINK(COMM, NEWCOMM, IERROR)
8     INTEGER COMM, NEWCOMM, IERROR
9

```

This collective operation creates a new intra or inter communicator `newcomm` from the intra or inter communicator `comm` respectively by excluding its failed processes as detailed below. It is valid MPI code to call `MPI_COMM_SHRINK` on a communicator which has been revoked (as defined above).

This function must not raise an [error-exception](#) due to process failures (error classes `MPI_ERR_PROC_FAILED` and `MPI_ERR_REVOKED`). All processes agree on the content of the group of processes that failed. This group includes at least every process failure that has raised an MPI exception of class `MPI_ERR_PROC_FAILED` or `MPI_ERR_PENDING`. The call is semantically equivalent to an `MPI_COMM_SPLIT` operation that would succeed despite failures, and 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 operations. (*End of advice to users.*)

```

27 MPI_COMM_FAILURE_ACK( comm )
28     IN      comm      communicator (handle)
29
30
31 int MPI_Comm_failure_ack(MPI_Comm comm)
32 MPI_COMM_FAILURE_ACK(COMM, IERROR)
33     INTEGER COMM, IERROR
34

```

This local operation gives the users a way to *acknowledge* all locally notified failures on `comm`. After the call, unmatched `MPI_ANY_SOURCE` receptions that would have raised an [error-code-exception](#) `MPI_ERR_PENDING` due to process failure (see Section 17.2.2) proceed without further [reporting-of-errors-raising-exceptions](#) due to those acknowledged failures.

Advice to users. Calling `MPI_COMM_FAILURE_ACK` on a communicator with failed processes does not allow that communicator to be used successfully for collective operations. Collective communication on a communicator with acknowledged failures will continue to raise an [error-exception](#) of class `MPI_ERR_PROC_FAILED` as defined in Section 17.2.2. To [reliably-use-collective-operations-on-a-communicator-with-resume-using-collective-operations-when-a-communicator-contains](#) failed processes, [the-communicator-should-first-be-revoked-using-and-then](#) a new communicator should be created using `MPI_COMM_SHRINK`. (*End of advice to users.*)

MPI_COMM_FAILURE_GET_ACKED(comm, failedgrp)

IN comm communicator (handle)
 OUT failedgrp group of failed processes (handle)

int MPI_Comm_failure_get_acked(MPI_Comm comm, MPI_Group* failedgrp)

MPI_COMM_FAILURE_GET_ACKED(COMM, FAILEDGRP, IERROR)
 INTEGER COMM, FAILEDGRP, IERROR

This local operation returns the group failedgrp of processes, from the communicator comm, which have been locally acknowledged as failed by preceding calls to MPI_COMM_FAILURE_ACK. The failedgrp can be empty, that is, equal to MPI_GROUP_EMPTY.

MPI_COMM_AGREE(comm, flag)

IN comm communicator (handle)
 INOUT flag boolean flag

int MPI_Comm_agree(MPI_Comm comm, int* flag)

MPI_COMM_AGREE(COMM, FLAG, IERROR)
 LOGICAL FLAG
 INTEGER COMM, IERROR

This function performs a collective operation on the group of living processes in comm. On completion, all living processes ~~must~~ agree to set the output value of flag to the result of a logical 'AND' operation over the contributed input values of flag. ~~This function must not raise an error due to process failure (error classes and), and processes~~ Processes that failed before entering the call do not contribute to the operation.

If comm is an intercommunicator, 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).

When MPI_COMM_AGREE completes, the group of living processes is consistent, therefore all failures in comm to date have been detected. If comm contains failures that have not been acknowledged by MPI_COMM_FAILURES_ACK, MPI_COMM_AGREE raises an exception of class MPI_ERR_PROC_FAILED; such an exception is also raised at all surviving ranks. The value of flag is correct when this exception is raised. This function never raise an exception of class MPI_ERR_REVOKED.

Advice to users. MPI_COMM_AGREE maintains its collective behavior even if the comm is revoked. *(End of advice to users.)*

```

1 MPI_COMM_IAGREE( comm, flag, req )
2     IN      comm      communicator (handle)
3
4     INOUT   flag      boolean flag
5
6     OUT     req       request (handle)
7
8 int MPI_Comm_iagree(MPI_Comm comm, int* flag, MPI_Request* req)
9
10 MPI_COMM_IAGREE(COMM, FLAG, REQ, IERROR)
11     LOGICAL FLAG
12     INTEGER COMM, REQ, IERROR

```

12 This function has the same semantics as MPI_COMM_AGREE except that it is non-
13 blocking.

15 17.3.2 One-Sided Functions

```

18 MPI_WIN_REVOKE( win )
19
20     IN      win      window (handle)
21
22 int MPI_Win_revoke(MPI_Win win)
23
24 MPI_WIN_REVOKE(WIN, IERROR)
25     INTEGER WIN, IERROR

```

26 This function notifies all processes within the window win that this window is now
27 considered revoked. ~~A revoked~~ This function is not collective and therefore does not have
28 a matching call on remote processes. All alive processes belonging to win will be notified of
29 the revocation despite failures. The revocation of a window completes any non-local MPI
30 operations on win ~~with error and causes any new operations to complete with error~~ by raising
31 an exception of class MPI_ERR_REVOKED. Once a window has been revoked, all subsequent
32 non-local operations on that window are considered local and must ~~fail with an error~~ raise
33 an exception of class MPI_ERR_REVOKED.

```

34
35
36 MPI_WIN_GET_FAILED( win, failedgrp )
37     IN      win      window (handle)
38
39     OUT     failedgrp group of failed processes (handle)
40
41 int MPI_Win_get_failed(MPI_Win win, MPI_Group* failedgrp)
42
43 MPI_WIN_GET_FAILED(WIN, FAILEDGRP, IERROR)
44     INTEGER COMM, FAILEDGRP, IERROR

```

45 This local operation returns the group failedgrp of processes from the window win which
46 are locally known to have failed.

47
48

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 and must raise a process failure exception. (*End of 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_REVOKED. (*End of advice to users.*)

17.3.3 I/O Functions

MPI_FILE_REVOKE(fh)

IN fh file (handle)

int MPI_File_revoke(MPI_File fh)

MPI_FILE_REVOKE(FH, IERROR)

INTEGER FH, IERROR

This function notifies all ~~ranks within file~~ processes within the file handle fh that this file handle is now considered revoked.

~~Ongoing~~ This function is not collective and therefore does not have a matching call on remote processes. All alive processes belonging to the file handle fh will be notified of the revocation despite failures. The revocation of a file handle completes any non-local completion operations on a revoked file handle raise MPI operations on win by raising an exception of class MPI_ERR_REVOKED. Once a file handle has been revoked, all subsequent non-local operations on the file handle that file handle are considered local and must raise an MPI exception of class MPI_ERR_REVOKED.

17.4 Error Codes and Classes

The following error classes are added to those defined in Section 8.4:

MPI_ERR_PROC_FAILED	The operation could not complete because of a process failure (a fail-stop failure).
MPI_ERR_REVOKED	The communication object used in the operation has been revoked.

Table 17.1: Additional process fault tolerance error classes

17.5 Examples

The example below illustrates how a new communicator can be safely created despite disruption by process failures. A child communicator is created with MPI_COMM_SPLIT, then the global success of the operation is verified with MPI_COMM_AGREE. If any process

failed to create the child communicator, all processes are notified by the value of the boolean flag agreed on. Processes that had successfully created the child communicator destroy it, as it cannot be safely used.

Example 17.1 Fault Tolerant Communicator Split Example

```

1  int Comm_split_consistent(MPI_Comm parent, int color, int key, MPI_Comm* child)
2  {
3
4
5
6      rc = MPI_Comm_split(parent, color, key, child);
7
8      split_ok = (MPI_SUCCESS == rc);
9
10     MPI_Comm_agree(parent, &split_ok);
11     if(split_ok) {
12         /* All surviving processes have created the "child" comm
13          * It may contain supplementary failures and the first
14          * operation on it may raise an exception, but it is a
15          * workable object that will yield well specified outcomes */
16         return MPI_SUCCESS;
17     }
18     else {
19         /* At least one process did not create the child comm properly
20          * if the local rank did succeed in creating it, it disposes
21          * of it, as it is a broken, unuseable object */
22         if(MPI_SUCCESS == rc) {
23             MPI_Comm_free(child);
24         }
25         return MPI_ERR_PROC_FAILED;
26     }
27 }
28
29

```

17.5.1 Master/Worker

The example below presents a master code that handles failures by ignoring failed processes and resubmitting requests. It demonstrates the different failure cases that may occur when posting receptions from `MPI_ANY_SOURCE` as discussed in the advice to users in Section 17.2.2.

Example 17.2 Fault-Tolerant Master Example

```

37 int master(void)
38 {
39     MPI_Comm_set_errhandler(comm, MPI_ERRORS_RETURN);
40     MPI_Comm_size(comm, &size);
41
42     /* ... submit the initial work requests ... */
43
44     MPI_Irecv( buffer, 1, MPI_INT, MPI_ANY_SOURCE, tag, comm, &req );
45
46     /* Progress engine: Get answers, send new requests,
47      and handle process failures */
48

```

```

while( (active_workers > 0) && work_available ) {
    rc = MPI_Wait( &req, &status );

    if( (MPI_ERR_PROC_FAILED == rc) || (MPI_ERR_PENDING == rc) ) {
        MPI_Comm_failure_ack(comm);
        MPI_Comm_failure_get_acked(comm, &g);
        MPI_Group_size(g, &gsize);

        /* ... find the lost work and requeue it ... */

        active_workers = size - gsize - 1;
        MPI_Group_free(&g);

        /* repost the request if it matched the failed process */
        if( rc == MPI_ERR_PROC_FAILED )
            MPI_Irecv( buffer, 1, MPI_INT, MPI_ANY_SOURCE,
                      tag, comm, &req );
    }

    continue;
}

/* ... process the answer and update work_available ... */
MPI_Irecv( buffer, 1, MPI_INT, MPI_ANY_SOURCE, tag, comm, &req );
}

/* ... cancel request and cleanup ... */
}

```

17.5.2 Iterative Refinement

The example below demonstrates a method of fault-tolerance to detect and handle failures. At each iteration, the algorithm checks the return code of the `MPI_ALLREDUCE`. If the return code indicates a process failure for at least one process, the algorithm revokes the communicator, agrees on the presence of failures, and later shrinks it to create a new communicator. By calling `MPI_COMM_REVOKE`, the algorithm ensures that all processes will be notified of process failure and enter the `MPI_COMM_AGREE`. If a process fails, the algorithm must complete at least one more iteration to ensure a correct answer.

Example 17.3 Fault-tolerant iterative refinement with shrink and agreement

```

while( gnorm > epsilon ) {
    /* Add a computation iteration to converge and
       compute local norm in lnorm */
    rc = MPI_Allreduce( &lnorm, &gnorm, 1, MPI_DOUBLE, MPI_MAX, comm);

    if( (MPI_ERR_PROC_FAILED == rc) ||
        (MPI_ERR_COMM_REVOKE == rc) ||
        (gnorm <= epsilon) ) {

```

```

1
2  \DIFaddbegin \DIFadd{/* This rank detected a failure, but other ranks may have
3  * proceeded into the next MPI_Allreduce. Since this rank
4  * will not match that following MPI_Allreduce, these other
5  * ranks would be at risk of deadlocking, so this process
6  * calls MPI_Comm_revoke to interrupt other ranks and notify
7  * them that it has detected a failure and is leaving the
8  * failure free execution path to go into recovery. */
9  }\DIFaddend if( MPI_ERR_PROC_FAILED == rc )
10     MPI_Comm_revoke(comm);
11
12     /* About to leave: let's be sure that everybody
13     received the same information */
14     allsucceeded = (rc == MPI_SUCCESS);
15     \DIFaddbegin \DIFadd{rc = }\DIFaddend MPI_Comm_agree(comm, &allsucceeded);
16     if( \DIFaddbegin \DIFadd{rc == MPI_ERR_PROC_FAILED || }\DIFaddend !allsucceeded ) {
17         \DIFdelbegin \DIFdel{/* We plan to join the shrink, thus the communicator
18         should be marked as revoked */
19         MPI_Comm_revoke(comm);
20         }\DIFdelend MPI_Comm_shrink(comm, &comm2);
21         MPI_Comm_free(comm); /* Release the revoked communicator */
22         comm = comm2; gnorm = epsilon + 1.0; /* Force one more iteration */
23     }
24 }
25 }
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

```


Bibliography

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

Index

- CONST:MPI_ANY_SOURCE, [2](#), [3](#), [8](#), [12](#)
- CONST:MPI_Comm, [7–10](#)
- CONST:MPI_COMM_NULL, [4](#)
- CONST:MPI_COMM_WORLD, [2](#), [3](#)
- CONST:MPI_ERR_PENDING, [2](#), [3](#), [8](#)
- CONST:MPI_ERR_PROC_FAILED, [2–6](#), [8](#),
[9](#), [11](#)
- CONST:MPI_ERR_REVOKED, [4–11](#)
- CONST:MPI_ERRORS_ARE_FATAL, [2](#)
- CONST:MPI_File, [11](#)
- CONST:MPI_FILE_NULL, [6](#), [7](#)
- CONST:MPI_Group, [9](#), [10](#)
- CONST:MPI_GROUP_EMPTY, [9](#)
- CONST:MPI_NO_OP, [5](#)
- CONST:MPI_Request, [10](#)
- CONST:MPI_SEEK_SET, [6](#)
- CONST:MPI_SUCCESS, [3–6](#)
- CONST:MPI_UNDEFINED, [8](#)
- CONST:MPI_Win, [10](#)

- EXAMPLES:Fault-tolerant iterative refinement with shrink and agreement, [13](#)
- EXAMPLES:Master example, [12](#)
- EXAMPLES:MPI_COMM_AGREE, [12](#), [13](#)
- EXAMPLES:MPI_COMM_FAILURE_ACK, [12](#)
- EXAMPLES:MPI_COMM_FAILURE_GET_ACKED, [12](#)
- EXAMPLES:MPI_COMM_FREE, [12](#), [13](#)
- EXAMPLES:MPI_COMM_REVOKE, [13](#)
- EXAMPLES:MPI_COMM_SHRINK, [13](#)
- EXAMPLES:MPI_COMM_SPLIT, [12](#)

- MPI_ACCUMULATE, [5](#)
- MPI_ALLREDUCE, [13](#)
- MPI_BARRIER, [4](#)
- MPI_BCAST, [4](#)
- MPI_COMM_ACCEPT, [4](#)
- MPI_COMM_AGREE, [4](#), [6](#), [7](#), [9–11](#), [13](#)
- MPI_COMM_AGREE(comm, flag), [9](#)
- MPI_COMM_CONNECT, [4](#)
- MPI_COMM_DUP, [4](#)
- MPI_COMM_FAILURE_ACK, [3](#), [8](#), [9](#)
- MPI_COMM_FAILURE_ACK(comm), [8](#)
- MPI_COMM_FAILURE_GET_ACKED(comm, failedgrp), [9](#)
- MPI_COMM_FAILURES_ACK, [9](#)
- MPI_COMM_FREE, [4](#)
- MPI_COMM_IAGREE(comm, flag, req), [10](#)
- MPI_COMM_JOIN, [5](#)
- MPI_COMM_REVOKE, [7](#), [13](#)
- MPI_COMM_REVOKE(comm), [7](#)
- MPI_COMM_SHRINK, [7](#), [8](#)
- MPI_COMM_SHRINK(comm, newcomm), [8](#)
- MPI_COMM_SPAWN, [5](#)
- MPI_COMM_SPAWN_MULTIPLE, [5](#)
- MPI_COMM_SPLIT, [4](#), [8](#), [11](#)
- MPI_FILE_CLOSE, [6](#)
- MPI_FILE_REVOKE(fh), [11](#)
- MPI_FILE_SEEK, [6](#)
- MPI_FINALIZE, [3](#)
- MPI_INIT, [3](#)
- MPI_TEST, [6](#)
- MPI_WAIT, [6](#)
- MPI_WIN_COMPLETE, [5](#)
- MPI_WIN_FREE, [6](#)
- MPI_WIN_GET_FAILED(win, failedgrp), [10](#)
- MPI_WIN_LOCK, [5](#)
- MPI_WIN_REVOKE(win), [10](#)
- MPI_WIN_REVOKED, [11](#)
- MPI_WIN_TEST, [5](#)
- MPI_WIN_UNLOCK, [5](#)
- MPI_WIN_WAIT, [5](#)