

D R A F T

Document for a Standard Message-Passing Interface

Message Passing Interface Forum

February 21, 2017

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.

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

Chapter 15

Process Fault Tolerance

15.1 Introduction

In distributed systems with numerous or complex components, a serious risk is that a component fault manifests as a process failure that disrupts the normal execution of a long running application. A process failure is a common 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: the affected process stops communicating permanently. This chapter introduces MPI features that support the development of applications, libraries, and programming languages that can tolerate MPI process failures. The primary goal is to specify error classes and interfaces that permit users to continue simple MPI communication (e.g., some point-to-point patterns) after failures have impacted the execution and rebuild MPI objects (communicators, files, etc.) as needed to restore the full capability of MPI to carry out elaborate communication operations (like collective communications). This specification does not include mechanisms to restore the data lost due to process failures. The literature is rich with diverse fault tolerance techniques that the users may employ at their discretion, including checkpoint-restart, algorithmic dataset recovery, and continuation ignoring failed MPI processes. All these fault tolerance approaches benefit from, and often require, the definitions and interfaces specified in this chapter in order to resume communicating after a failure.

The expected behavior of MPI in the case of an MPI 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 error (see Section 15.2); an MPI operation that does not involve a failed process will complete normally, unless interrupted by the user through provided functionality. Errors indicate only the local impact of the failure on an operation, and make no guarantee that other processes have also been notified of the same failure. Asynchronous failure propagation is not guaranteed or required, and users must exercise caution when determining the set of processes where a failure has been detected and raised an error. If an application needs global knowledge of failures, it can use the interfaces defined in Section 15.3 to explicitly propagate the notification of locally detected failures.

Some usage patterns on reliable machines do not require fault tolerance. An MPI implementation that does not tolerate process failures must never raise a process failure error (as listed in Section 15.4). Applications using the interfaces defined in this chapter must be portable across MPI implementations (including those which do not provide fault

1 tolerance, but in this case the interfaces may exhibit undefined behavior after a process
2 failure at any MPI process.) Fault tolerant applications may determine if the implementation
3 supports fault tolerance by querying the predefined attribute `MPI_FT` on `MPI_COMM_WORLD`
4 (see 8.1.2.)

5
6 *Advice to users.* Many of the operations and semantics described in this chapter
7 are applicable only when the MPI application has replaced the default error handler
8 `MPI_ERRORS_ARE_FATAL` on, at least, `MPI_COMM_WORLD`. (*End of advice to users.*)

10 15.2 Failure Notification

11
12 This section specifies the behavior of an MPI communication operation when failures occur
13 on MPI processes involved in the communication. An MPI process is considered involved in
14 a communication (for the purpose of this chapter) if any of the following is true:

- 15
16 • The process is in the group over which the operation is collective.
- 17
18 • The process is a destination or a specified or matched source in a point-to-point
19 communication.
- 20
21 • The operation is an `MPI_ANY_SOURCE` receive operation and the process belongs to
22 the source group.
- 23
24 • The process is a specified target in a remote memory operation.

25 An operation involving a failed MPI process must always complete in a finite amount of
26 time (possibly by raising one of the process failure error classes listed in Section 15.4). If an
27 operation does not involve a failed MPI process (such as a point-to-point message between
28 two non-failed MPI processes), it must not raise a process failure error.

29
30 *Advice to implementors.* An MPI implementation may provide failure detection only
31 for MPI processes involved in an ongoing operation and may postpone detection of
32 other failures until necessary. Moreover, as long as an implementation can complete
33 operations, it may choose to delay raising an error. Another valid implementation
34 might choose to raise an error as quickly as possible. (*End of advice to implementors.*)

35
36 When a communication operation raises a process failure error, it may not satisfy
37 its specification, (for example, a synchronizing operation may not have synchronized) and
38 the content of the output buffers, targeted memory, or output parameters is undefined.
39 Exceptions to this rule are explicitly stated in the remainder of this chapter. Error codes
40 returned from a function, output in arrays of error codes, or in status objects remain defined
41 after an operation raised a process failure error.

42 Nonblocking operations do not raise process failure errors during creation or initiation.
43 All process failure error raising is postponed until the corresponding completion function is
44 called.

15.2.1 Startup and Finalize

Initialization does not have any new semantics related to fault tolerance.

Advice to implementors. If an MPI process fails during MPI_INIT but its peers are able to complete the MPI_INIT successfully, then an implementation can return MPI_SUCCESS, produce an MPI_COMM_WORLD whose group contains failed MPI processes, and delay the reporting of the process failure to a subsequent MPI operation. (*End of advice to implementors.*)

MPI_FINALIZE will complete even in the presence of MPI process failures. If process 0 in MPI_COMM_WORLD has failed, it is possible that no MPI process returns from MPI_FINALIZE.

Advice to users. Fault tolerant applications are encouraged to implement all rank-specific code before the call to MPI_FINALIZE. In Example 8.10 in Section 8.7, the process with rank 0 in MPI_COMM_WORLD may have failed before, during, or after the call to MPI_FINALIZE, possibly leading to the code after MPI_FINALIZE never being executed. (*End of advice to users.*)

15.2.2 Point-to-Point and Collective Communication

An MPI implementation raises errors of the following classes in order to notify users that a point-to-point communication operation could not complete successfully because of the failure of at least one involved MPI process:

- MPI_ERR_PROC_FAILED_PENDING indicates, for a nonblocking communication, that the communication is a receive operation from MPI_ANY_SOURCE and no send operation has matched, yet a potential sending MPI process has failed. Neither the operation nor the request identifying the operation is completed.
- In all other cases, the operation raises an error 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 a point-to-point communication, it is completed. Communication involving the failed MPI process, initiated on this communicator after the error raised, must also raise an error of class MPI_ERR_PROC_FAILED.

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

Advice to users.

Depending on how the collective operation is implemented and when an MPI process failure occurs, some participating MPI processes may raise an error while other MPI processes return successfully from the same collective operation. For example, in MPI_BCAST, the root process may succeed before a failed MPI process disrupts the operation, resulting in some other processes raising an error.

(*End of advice to users.*)

1 *Advice to users.*

2 Note that communicator creation functions (e.g., `MPI_COMM_DUP` or
3 `MPI_COMM_SPLIT`) are collective operations. As such, if a failure happened during
4 the call, an error might be raised at some MPI processes while others succeed and
5 obtain a new communicator handle. Although it is valid to communicate between
6 MPI processes that succeeded in creating the new communicator handle, the user is
7 responsible for ensuring a consistent view of the communicator creation, if needed.
8 A conservative solution is to check the global outcome of the communicator creation
9 function with `MPI_COMM_AGREE` (defined in Section 15.3.1), as illustrated in Ex-
10 ample 15.1. (*End of advice to users.*)

11
12 After an MPI process failure, `MPI_COMM_FREE` (as with all other collective opera-
13 tions) may not complete successfully at all processes. For any MPI process that receives
14 the return code `MPI_SUCCESS`, the behavior is defined in Section 6.4.3. If an MPI process
15 raises a process failure error (classes `MPI_ERR_PROC_FAILED` or `MPI_ERR_REVOKED`), the
16 communicator handle `comm` is set to `MPI_COMM_NULL`; however, the implementation makes
17 no guarantee about the success or failure of the `MPI_COMM_FREE` operation, locally or
18 remotely.

19
20 *Advice to users.* Users are encouraged to call `MPI_COMM_FREE` on communicators
21 they do not wish to use anymore, even when they contain failed MPI processes. Al-
22 though the operation may raise a process failure error and not synchronize properly,
23 this gives a high quality implementation an opportunity to release local resources and
24 memory consumed by the object. (*End of advice to users.*)

25 26 15.2.3 Dynamic Process Management

27 *Rationale.* As with communicator creation functions, if a failure happens during a dy-
28 namic process management operation, an error might be raised at some MPI processes
29 while others succeed and obtain a new valid communicator. For most communicator
30 creation functions, users can validate the success of the operation by communicat-
31 ing on a pre-existing communicator spanning over the same group of processes (in
32 the worst case, from `MPI_COMM_WORLD`). This is however not always possible for
33 dynamic process management operations, since these operations can create a new
34 intercommunicator between previously disconnected MPI processes. The following
35 additional failure case semantics allow for users to validate, on the created intercom-
36 municator itself, the success of the dynamic process management operation. (*End of*
37 *rationale.*)

38
39 If the MPI implementation raises a process failure error at the root process in
40 `MPI_COMM_ACCEPT` or `MPI_COMM_CONNECT`, the corresponding operation must also
41 raise a process failure error at its root process.

42
43 *Advice to users.* The root process of an operation can succeed when a process failure
44 error is raised at some other non-root process. (*End of advice to users.*)

45
46 When using the intercommunicator returned from `MPI_COMM_SPAWN`,
47 `MPI_COMM_SPAWN_MULTIPLE`, or `MPI_COMM_GET_PARENT`, a communication for which
48 the root process of the spawn operation is the source or the destination must not deadlock.

When the root process raises a process failure error from a spawn operation, no MPI processes are spawned.

Advice to implementors. An implementation is allowed to abort a spawned MPI process during MPI_INIT when it cannot setup an intercommunicator with the root process of the spawn operation because of a process failure.

An implementation may report it spawned all the requested MPI processes even when a process created from MPI_COMM_SPAWN or MPI_COMM_SPAWN_MULTIPLE failed, and instead delay raising a process failure error to a later communication involving this process. (*End of advice to implementors.*)

Advice to users. To determine how many new MPI processes have effectively been spawned, the normal semantics for hard and soft spawn applies: if the requested number of processes is unavailable for a hard spawn, an error of class MPI_ERR_SPAWN is raised (possibly leaving MPI in an undefined state), and an appropriate error code is set in the array_of_errcodes parameter. Note however that an implementation may report that it has spawned the requested number of MPI processes even when some MPI processes have failed before exiting MPI_INIT. This condition can be detected by communicating over the created intercommunicator with these processes. (*End of advice to users.*)

Advice to implementors. MPI_COMM_JOIN does not require any supplementary semantics. When the remote MPI process on the fd socket has failed, the operation succeeds and sets intercomm to MPI_COMM_NULL. (*End of advice to implementors.*)

After an MPI process failure, MPI_COMM_DISCONNECT (as with all other collective operations) may not complete successfully at all MPI processes. For any process that receives the return code MPI_SUCCESS, the behavior is defined in 10.5.4. If an MPI process raises a process failure error (classes MPI_ERR_PROC_FAILED or MPI_ERR_REVOKED), the communicator handle comm is set to MPI_COMM_NULL; however, the implementation makes no guarantee about the success or failure of the MPI_COMM_DISCONNECT operation, locally or remotely.

Advice to users. Users are encouraged to call MPI_COMM_DISCONNECT on communicators they do not wish to use anymore, even when they contain failed MPI processes. Although the operation may raise a process failure error and not synchronize properly, this gives a high quality implementation an opportunity to release local resources and memory consumed by the object. (*End of advice to users.*)

15.2.4 One-Sided Communication

When an operation on a window raises a process failure error, the state of all data held in memory exposed by that window becomes undefined at all MPI processes for which a one-sided communication operation could have modified local data (a target in a remote write, or accumulate operation, or an origin in a remote read operation), and the operation completion has not been semantically guaranteed (as an example by a successful synchronization between the origin and the target, after the origin had issued an MPI_WIN_FLUSH).

1 *Advice to users.* Assessing if a particular portion of the exposed memory remains
2 correct is the responsibility of the user. Note that in passive target mode, when an
3 error is raised at the origin, the target memory data may become undefined before a
4 synchronization raises an error at the target.

5 The exposed memory data becomes undefined for all uses, not only the window in
6 which the error was raised. Any overlapping windows or uses involving shared memory
7 also read undefined data (even if they do not involve MPI calls). (*End of advice to*
8 *users.*)

9
10 *Advice to implementors.* A high quality implementation should limit the scope of the
11 exposed memory that becomes undefined (for example, only the memory addresses
12 and ranges that have been targeted by a remote write, or accumulate, or have been
13 an origin in a remote read). In that case, we encourage implementations to document
14 the provided behavior, and to expose the availability of this feature at runtime, as
15 an example by caching an implementation specific attribute on the window. (*End of*
16 *advice to implementors.*)

17
18 Non-synchronizing one-sided communication operations (as an example `MPI_GET`,
19 `MPI_PUT`) whose outputs are undefined, due to an MPI process failure, are not required to
20 raise a process failure error. However, if a communication cannot complete correctly due to
21 process failures, the synchronization operation must raise a process failure error at least at
22 the origin.

23
24 *Advice to implementors.* Non-synchronizing operations (`MPI_WIN_FLUSH_LOCAL`,
25 `MPI_WIN_FLUSH_LOCAL_ALL`) are not required to raise a process failure error. (*End*
26 *of advice to implementors.*)

27
28 *Advice to users.* As with collective operations over MPI communicators, active target
29 one-sided synchronization operations may raise a process failure error at some MPI
30 process while the corresponding operation returned `MPI_SUCCESS` at some other MPI
31 process. (*End of advice to users.*)

32 Passive target synchronization operations may raise a process failure error when any
33 MPI process in the window has failed (even when the target specified in the argument of
34 the passive target synchronization has not failed).

35
36 *Rationale.* An implementation of passive target synchronization may need to com-
37 municate with non-target MPI processes in the window, as an example, a previous
38 owner of an access epoch on the target window. (*End of rationale.*)

39
40 After an MPI process failure, `MPI_WIN_FREE` (as with all other collective operations)
41 may not complete successfully at all MPI processes. For any process that receives the return
42 code `MPI_SUCCESS`, the behavior is defined in Section 11.2.5. If a process raises a process
43 failure error (classes `MPI_ERR_PROC_FAILED` or `MPI_ERR_REVOKED`), the window handle
44 `win` is set to `MPI_WIN_NULL`; however, the implementation makes no guarantee about the
45 success or failure of the `MPI_WIN_FREE` operation, locally or remotely.

46
47 *Advice to users.* Users are encouraged to call `MPI_WIN_FREE` on windows they do
48 not wish to use anymore, even when they contain failed MPI processes. Although the

operation may raise a process failure error and not synchronize properly, this gives a high quality implementation an opportunity to release local resources and memory consumed by the object. Before calling `MPI_WIN_FREE`, it may be required to call `MPI_WIN_REVOKE` to close an epoch that couldn't be completed as a consequence of a process failure (see Section 15.3.2). (*End of advice to users.*)

15.2.5 I/O

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

If an MPI process failure prevents a file operation from completing, an MPI error of class `MPI_ERR_PROC_FAILED` is raised. Once an MPI implementation has raised an error of class `MPI_ERR_PROC_FAILED`, the state of the file pointers involved in the operation that raised the error is *undefined*.

Advice to users. Since collective I/O operations may not synchronize with other MPI 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 by using `MPI_FILE_SEEK` with the `MPI_SEEK_SET` update mode. (*End of advice to users.*)

After an MPI process failure, `MPI_FILE_CLOSE` (as with all other collective operations) may not complete successfully at all MPI processes. For any MPI process that receives the return code `MPI_SUCCESS`, the behavior is defined in Section 13.2.2. If an MPI process raises a process failure error (classes `MPI_ERR_PROC_FAILED` or `MPI_ERR_REVOKED`), the file handle `fh` is set to `MPI_FILE_NULL`; however, the implementation makes no guarantee about the success or failure of the `MPI_FILE_CLOSE` operation, locally or remotely.

Advice to users. Users are encouraged to call `MPI_FILE_CLOSE` on files they do not wish to use anymore, even when they contain failed MPI processes. Although the operation may raise a process failure error and not synchronize properly, this gives a high quality implementation an opportunity to release local resources and memory consumed by the object. (*End of advice to users.*)

15.3 Failure Mitigation Functions

15.3.1 Communicator Functions

Process failure notification is not global in MPI. MPI processes that do not call operations involving a failed MPI process are possibly never notified of its failure (see Section 15.2). If a notification must be propagated, MPI provides a function to revoke a communicator at all members.

```

1 MPI_COMM_REVOKE( comm )
2     IN      comm          communicator (handle)
3
4 int MPI_Comm_revoke(MPI_Comm comm)
5
6 MPI_Comm_revoke(comm, ierror)
7     TYPE(MPI_Comm), INTENT(IN) :: comm
8     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
9
10 MPI_COMM_REVOKE(COMM, IERROR)
11     INTEGER COMM, IERROR

```

12 This function notifies all MPI processes in the groups (local and remote) associated
13 with the communicator `comm` that this communicator is revoked. The revocation of a
14 communicator by any MPI process completes non-local MPI operations on `comm` at all MPI
15 processes by raising an error of class `MPI_ERR_REVOKED` (with the exception of
16 `MPI_COMM_SHRINK`, `MPI_COMM_AGREE`, and `MPI_COMM_IAGREE`). This function is
17 not collective and therefore does not have a matching call on remote MPI processes. All non-
18 failed MPI processes belonging to `comm` will be notified of the revocation despite failures.

19 A communicator is revoked at a given MPI process either when
20 `MPI_COMM_REVOKE` is locally called on it, or when any MPI operation on `comm` raises an
21 error of class `MPI_ERR_REVOKED` at that process. Once a communicator has been revoked
22 at an MPI process, all subsequent non-local operations on that communicator (with the
23 same exceptions as above), are considered local and must complete by raising an error of
24 class `MPI_ERR_REVOKED` at that MPI process.

```

25
26 MPI_COMM_SHRINK( comm, newcomm )
27
28     IN      comm          communicator (handle)
29     OUT    newcomm       communicator (handle)
30
31 int MPI_Comm_shrink(MPI_Comm comm, MPI_Comm* newcomm)
32
33 MPI_Comm_shrink(comm, newcomm, ierror)
34     TYPE(MPI_Comm), INTENT(IN) :: comm
35     TYPE(MPI_Comm), INTENT(OUT) :: newcomm
36     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
37
38 MPI_COMM_SHRINK(COMM, NEWCOMM, IERROR)
39     INTEGER COMM, NEWCOMM, IERROR

```

40 This collective operation creates a new intra- or intercommunicator
41 `newcomm` from the intra- or intercommunicator `comm`, respectively, by excluding the group
42 of failed MPI processes as agreed upon during the operation. The groups of `newcomm` must
43 include every MPI process that returns from `MPI_COMM_SHRINK`, and it must exclude
44 every MPI process whose failure caused an operation on `comm` to raise an MPI error of class
45 `MPI_ERR_PROC_FAILED` or `MPI_ERR_PROC_FAILED_PENDING` at a member of the groups of
46 `newcomm`, before that member initiated `MPI_COMM_SHRINK`. This call is semantically
47 equivalent to an `MPI_COMM_SPLIT` operation that would succeed despite failures, where
48

members of the groups of `newcomm` participate with the same color and a key equal to their rank in `comm`.

This function never raises an error of class `MPI_ERR_PROC_FAILED` or `MPI_ERR_REVOKED`. The defined semantics of `MPI_COMM_SHRINK` are maintained when `comm` is revoked, or when the group of `comm` contains failed MPI processes.

Advice to users. `MPI_COMM_SHRINK` is a collective operation, even when `comm` is revoked.

The group of `newcomm` may still contain failed MPI processes, whose failure will be detected in subsequent MPI operations. (*End of advice to users.*)

`MPI_COMM_FAILURE_ACK(comm)`

IN `comm` communicator (handle)

`int MPI_Comm_failure_ack(MPI_Comm comm)`

`MPI_Comm_failure_ack(comm, ierror)`

 TYPE(MPI_Comm), INTENT(IN) :: `comm`

 INTEGER, OPTIONAL, INTENT(OUT) :: `ierror`

`MPI_COMM_FAILURE_ACK(COMM, IERROR)`

 INTEGER COMM, IERROR

This local operation gives the users a way to *acknowledge* all locally notified failures on `comm`. After the call, unmatched `MPI_ANY_SOURCE` receive operations that would have raised an error of class `MPI_ERR_PROC_FAILED_PENDING` due to MPI process failure (see Section 15.2.2) proceed without further raising errors due to those acknowledged failures. Also after this call, `MPI_COMM_AGREE` will not raise an error of class `MPI_ERR_PROC_FAILED` due to those acknowledged failures (according to the specification found later in this section).

Advice to users.

Calling `MPI_COMM_FAILURE_ACK` on a communicator with failed MPI processes has no effect on collective operations (except for `MPI_COMM_AGREE`). If a collective operation would raise an error due to the communicator containing a failed process (as defined in Section 15.2.2), it can continue to raise an error even after the failure has been acknowledged. In order to use collective operations between MPI processes of a communicator that contains failed MPI processes, users should create a new communicator by calling `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)

```

1  int MPI_Comm_failure_get_acked(MPI_Comm comm, MPI_Group* failedgrp)
2
3  MPI_Comm_failure_get_acked(comm, failedgrp, ierror)
4      TYPE(MPI_Comm), INTENT(IN) :: comm
5      TYPE(MPI_Group), INTENT(OUT) :: failedgrp
6      INTEGER, OPTIONAL, INTENT(OUT) :: ierror
7
8  MPI_COMM_FAILURE_GET_ACKED(COMM, FAILEDGRP, IERROR)
9      INTEGER COMM, FAILEDGRP, IERROR

```

This local operation returns the group *failedgrp* of processes, from the communicator *comm*, that 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`.

Advice to users. When they are not separated by a call to `MPI_COMM_FAILURE_ACK`, multiple calls to `MPI_COMM_FAILURE_GET_ACKED` produce similar *failedgrp* groups; that is, the result when providing these groups to `MPI_GROUP_DIFFERENCE` is `MPI_EMPTY`. (*End of advice to users.*)

```

20 MPI_COMM_AGREE( comm, flag )
21
22     IN      comm      communicator (handle)
23     INOUT   flag      integer flag
24
25 int MPI_Comm_agree(MPI_Comm comm, int* flag)
26
27 MPI_Comm_agree(comm, flag, ierror)
28     TYPE(MPI_Comm), INTENT(IN) :: comm
29     INTEGER, INTENT(INOUT) :: flag
30     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
31
32 MPI_COMM_AGREE(COMM, FLAG, IERROR)
33     INTEGER COMM, FLAG, IERROR

```

The purpose of this collective communication is to agree on the integer value *flag* and on the group of failed processes in *comm*.

On completion, all non-failed MPI processes have agreed to set the output integer value of *flag* to the result of a bitwise ‘AND’ operation over the contributed input values of *flag*. If *comm* is an intercommunicator, the value of *flag* is a bitwise ‘AND’ operation over the values contributed by the remote group.

When an MPI process fails before contributing to the operation, the *flag* is computed ignoring its contribution, and `MPI_COMM_AGREE` raises an error of class `MPI_ERR_PROC_FAILED`. However, if all MPI processes have acknowledged this failure prior to the call to `MPI_COMM_AGREE`, using `MPI_COMM_FAILURE_ACK`, the error related to this failure is not raised. When an error of class `MPI_ERR_PROC_FAILED` is raised, it is consistently raised at all MPI processes, in both the local and remote groups (if applicable).

After `MPI_COMM_AGREE` raised an error of class `MPI_ERR_PROC_FAILED`, a subsequent call to `MPI_COMM_FAILURE_ACK` on *comm* acknowledges the failure of every MPI process that didn’t contribute to the computation of *flag*.

Advice to users. Using a combination of `MPI_COMM_FAILURE_ACK` and `MPI_COMM_AGREE` as illustrated in Example 15.3, users can propagate and synchronize the knowledge of failures across all MPI processes in `comm`. When `MPI_SUCCESS` is returned locally from `MPI_COMM_AGREE`, the operation has not raised an error of class `MPI_ERR_PROC_FAILED` at any MPI process and thereby returned `MPI_SUCCESS` at all other MPI processes. (*End of advice to users.*)

This function never raises an error of class `MPI_ERR_REVOKED`. The defined semantics of `MPI_COMM_AGREE` are maintained when `comm` is revoked, or when the group of `comm` contains failed MPI processes.

Advice to users. `MPI_COMM_AGREE` is a collective operation, even when `comm` is revoked. (*End of advice to users.*)

`MPI_COMM_IAGREE(comm, flag, req)`

IN	<code>comm</code>	communicator (handle)
INOUT	<code>flag</code>	integer flag
OUT	<code>req</code>	request (handle)

`int MPI_Comm_iagree(MPI_Comm comm, int* flag, MPI_Request* req)`

`MPI_Comm_iagree(comm, flag, req, ierror)`

<code>TYPE(MPI_Comm)</code> , <code>INTENT(IN)</code>	::	<code>comm</code>
<code>INTEGER</code> , <code>INTENT(INOUT)</code>	::	<code>flag</code>
<code>TYPE(MPI_Request)</code> , <code>INTENT(OUT)</code>	::	<code>req</code>
<code>INTEGER</code> , <code>OPTIONAL</code> , <code>INTENT(OUT)</code>	::	<code>ierror</code>

`MPI_COMM_IAGREE(COMM, FLAG, REQ, IERROR)`

`INTEGER COMM, FLAG, REQ, IERROR`

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

15.3.2 One-Sided Functions

`MPI_WIN_REVOKE(win)`

IN	<code>win</code>	window (handle)
----	------------------	-----------------

`int MPI_Win_revoke(MPI_Win win)`

`MPI_Win_revoke(win, ierror)`

<code>TYPE(MPI_Win)</code> , <code>INTENT(IN)</code>	::	<code>win</code>
<code>INTEGER</code> , <code>OPTIONAL</code> , <code>INTENT(OUT)</code>	::	<code>ierror</code>

`MPI_WIN_REVOKE(WIN, IERROR)`

`INTEGER WIN, IERROR`

This function notifies all MPI processes in the group associated with the window `win` that this window is revoked. The revocation of a window by any MPI process completes RMA operations on `win` at all MPI processes and RMA synchronizations on `win` raise an error of class `MPI_ERR_REVOKED`. This function is not collective and therefore does not have a matching call on remote MPI processes. All non-failed MPI processes belonging to `win` will be notified of the revocation despite failures.

A window is revoked at a given MPI process either when `MPI_WIN_REVOKE` is locally called on it, or when any MPI operation on `win` raises an error of class `MPI_ERR_REVOKED` at that process. Once a window has been revoked at an MPI process, all subsequent RMA operations on that window are considered local and RMA synchronizations must complete by raising an error of class `MPI_ERR_REVOKED` at that process. In addition, the current epoch is closed and RMA operations originating from this MPI process are interrupted and completed with undefined outputs.

```
MPI_WIN_GET_FAILED( win, failedgrp )
```

```
IN      win                window (handle)
OUT     failedgrp          group of failed processes (handle)
```

```
int MPI_Win_get_failed(MPI_Win win, MPI_Group* failedgrp)
```

```
MPI_Win_get_failed(win, failedgrp, ierror)
    TYPE(MPI_Win), INTENT(IN) :: win
    TYPE(MPI_Group), INTENT(OUT) :: failedgrp
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_WIN_GET_FAILED(WIN, FAILEDGRP, IERROR)
    INTEGER COMM, FAILEDGRP, IERROR
```

This local operation returns the group `failedgrp` of MPI processes from the window `win` that 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 update only the group of locally known failed MPI processes when it enters a synchronization function and must raise a process failure error. (*End of advice to users.*)

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

15.3.3 I/O Functions

```
MPI_FILE_REVOKE( fh )
```

```
IN      fh                file (handle)
```

```
int MPI_File_revoke(MPI_File fh)
```

```

MPI_File_revoke(fh, ierror)
    TYPE(MPI_File), INTENT(IN) :: fh
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_FILE_REVOKE(FH, IERROR)
    INTEGER FH, IERROR

```

This function notifies all MPI processes in the group associated with the file handle `fh` that this file handle is revoked. The revocation of a file handle by any MPI process completes non-local MPI operations on `fh` at all MPI processes by raising an error of class `MPI_ERR_REVOKED`. This function is not collective and therefore does not have a matching call on remote MPI processes. All non-failed MPI processes belonging to `fh` will be notified of the revocation despite failures.

A file handle is revoked at a given MPI process either when `MPI_FILE_REVOKE` is locally called on it, or when any MPI operation on `fh` raises an error of class `MPI_ERR_REVOKED` at that process. Once a file handle has been revoked at an MPI process, all subsequent non-local operations on that file handle are considered local and must complete by raising an error of class `MPI_ERR_REVOKED` at that process.

15.4 Process Failure Error Codes and Classes

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

<code>MPI_ERR_PROC_FAILED</code>	The operation could not complete because of an MPI process failure (a fail-stop failure).
<code>MPI_ERR_PROC_FAILED_PENDING</code>	The operation was interrupted by an MPI process failure (a fail-stop failure). The request is still pending and the operation may be completed later.
<code>MPI_ERR_REVOKED</code>	The communication object used in the operation has been revoked.

Table 15.1: Additional process fault tolerance error classes

15.5 Examples

15.5.1 Safe Communicator Creation

The example below illustrates how a new communicator can be safely created despite disruption by MPI 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 MPI `g` failed to create the child communicator handle, all MPI processes are notified by the value of the integer agreed on. MPI processes that had successfully created the child communicator handle destroy it, as it cannot be used consistently.

Example 15.1 Fault Tolerant Communicator Split Example

```

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

```

15.5.2 Obtaining the consistent group of failed processes

Users can invoke `MPI_COMM_FAILURE_ACK`, `MPI_COMM_FAILURE_GET_ACKED`, `MPI_WIN_GET_FAILED`, to obtain the group of failed MPI processes, as detected at the local MPI process. However, these operations are local, thereby the invocation of the same function at another MPI process can result in a different group of failed processes being returned.

In the following examples, we illustrate two different approaches that permit obtaining the consistent group of failed MPI processes across all MPI processes of a communicator. The first one employs `MPI_COMM_SHRINK` to create a temporary communicator excluding failed MPI processes. The second one employs `MPI_COMM_AGREE` to synchronize the set of acknowledged failures.

Example 15.2 Fault-Tolerant Consistent Group of Failures Example (Shrink variant)

```

38  Comm_failure_allget(MPI_Comm c, MPI_Group * g) {
39      MPI_Comm s; MPI_Group c_grp, s_grp;
40
41      /* Using shrink to create a new communicator, the underlying
42       * group is necessarily consistent across all processes, and excludes
43       * all processes detected to have failed before the call */
44      MPI_Comm_shrink(c, &s);
45      /* Extracting the groups from the communicators */
46      MPI_Comm_group(c, &c_grp);
47      MPI_Comm_group(s, &s_grp);
48      /* s_grp is the group of still alive processes, we want to

```



```

    * return the group of failed processes. */
MPI_Group_difference(c_grp, s_grp, g);

MPI_Group_free(&c_grp); MPI_Group_free(&s_grp);
MPI_Comm_free(&s);
}

```

Example 15.3 Fault-Tolerant Consistent Group of Failures Example (Agree variant)

```

Comm_failure_allget2(MPI_Comm c, MPI_Group * g) {
    int rc; int T=1;

    do {
        /* this routine is not pure: calling MPI_Comm_failure_ack
        * affects the state of the communicator c */
        MPI_Comm_failure_ack(comm);
        /* we simply ignore the value in this example */
        rc = MPI_Comm_agree(comm, &T);
    } while( rc != MPI_SUCCESS );
    /* after this loop, MPI_Comm_agree has returned MPI_SUCCESS at
    * all processes, so all processes have Acknowledged the same set of
    * failures. Let's get that set of failures in the g group. */
    MPI_Comm_failure_get_acked(comm, g);
}

```

15.5.3 Fault-Tolerant Master/Worker

The example below presents a master code that handles worker failures by discarding failed worker MPI processes and resubmitting the work to the remaining workers. 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 15.2.2.

Example 15.4 Fault-Tolerant Master Example

```

int master(void)
{
    MPI_Comm_set_errhandler(comm, MPI_ERRORS_RETURN);
    MPI_Comm_size(comm, &size);

    /* ... submit the initial work requests ... */

    /* Progress engine: Get answers, send new requests,
    and handle process failures */
    MPI_Irecv( buffer, 1, MPI_INT, MPI_ANY_SOURCE, tag, comm, &req );
    while( (active_workers > 0) && work_available ) {
        rc = MPI_Wait( &req, &status );
        if( MPI_SUCCESS == rc ) {
            /* ... process the answer and update work_available ... */
        }
    }
}

```

```

1     else {
2         MPI_Error_class(rc, &ec);
3         if( (MPI_ERR_PROC_FAILED == ec) ||
4             (MPI_ERR_PROC_FAILED_PENDING == ec) ) {
5             MPI_Comm_failure_ack(comm);
6             MPI_Comm_failure_get_acked(comm, &g);
7             MPI_Group_size(g, &gsize);
8
9             /* ... find the lost work and requeue it ... */
10
11            active_workers = size - gsize - 1;
12            MPI_Group_free(&g);
13
14            /* no need to repost when the request is still pending */
15            if( ec == MPI_ERR_PROC_FAILED_PENDING )
16                continue;
17        }
18    }
19    /* get ready to receive more notifications from workers */
20    MPI_Irecv( buffer, 1, MPI_INT, MPI_ANY_SOURCE, tag, comm, &req );
21 }
22 /* ... cancel request and cleanup ... */
23 }
24

```

15.5.4 Fault-Tolerant Iterative Refinement

The example below demonstrates a method of fault tolerance for detecting and handling 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 MPI process, the algorithm revokes the communicator, agrees on the presence of failures, and shrinks it to create a new communicator. By calling `MPI_COMM_REVOKE`, the algorithm ensures that all MPI processes will be notified of process failure and enter the `MPI_COMM_AGREE`. If an MPI process fails, the algorithm must complete at least one more iteration to ensure a correct answer.

Example 15.5 Fault-tolerant iterative refinement with shrink and agreement

```

35
36
37 while( gnorm > epsilon ) {
38     /* Add a computation iteration to converge and
39        compute local norm in lnorm */
40     rc = MPI_Allreduce(&lnorm, &gnorm, 1, MPI_DOUBLE, MPI_MAX, comm);
41     MPI_Error_class(rc, &ec);
42
43     if( (MPI_ERR_PROC_FAILED == ec) ||
44         (MPI_ERR_REVOKED == ec) ||
45         (gnorm <= epsilon) ) {
46
47         /* This process detected a failure, but other processes may have
48            * proceeded into the next MPI_Allreduce. Since this process

```

```
    * will not match that following MPI_Allreduce, these other      1
    * processes would be at risk of deadlocking. This process thus  2
    * calls MPI_Comm_revoke to interrupt other processes and notify  3
    * them that it has detected a failure and is leaving the        4
    * failure free execution path to go into recovery. */          5
if( MPI_ERR_PROC_FAILED == ec )                                   6
    MPI_Comm_revoke(comm);                                       7
                                                                    8
/* About to leave: let's be sure that everybody                   9
   received the same information */                               10
allsucceeded = (rc == MPI_SUCCESS);                             11
rc = MPI_Comm_agree(comm, &allsucceeded);                       12
MPI_Error_class(rc, &ec);                                       13
if( ec == MPI_ERR_PROC_FAILED || !allsucceeded ) {             14
    MPI_Comm_shrink(comm, &comm2);                               15
    MPI_Comm_free(comm); /* Release the revoked communicator */ 16
    comm = comm2;                                               17
    gnorm = epsilon + 1.0; /* Force one more iteration */      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](#), [9](#), [15](#)
CONST:MPI_Comm, [7–11](#)
CONST:MPI_COMM_NULL, [4](#), [5](#)
CONST:MPI_COMM_WORLD, [2–4](#)
CONST:MPI_EMPTY, [10](#)
CONST:MPI_ERR_PROC_FAILED, [3–10](#), [13](#)
CONST:MPI_ERR_PROC_FAILED_PENDING, [3](#), [8](#), [9](#), [13](#)
CONST:MPI_ERR_REVOKED, [4–8](#), [11–13](#)
CONST:MPI_ERR_SPAWN, [5](#)
CONST:MPI_ERRORS_ARE_FATAL, [2](#)
CONST:MPI_File, [12](#)
CONST:MPI_FILE_NULL, [7](#)
CONST:MPI_FT, [2](#)
CONST:MPI_Group, [9](#), [12](#)
CONST:MPI_GROUP_EMPTY, [10](#)
CONST:MPI_Request, [11](#)
CONST:MPI_SEEK_SET, [7](#)
CONST:MPI_SUCCESS, [3–7](#), [10](#)
CONST:MPI_Win, [11](#), [12](#)
CONST:MPI_WIN_NULL, [6](#)

EXAMPLES:Comm_failure_allget example, [14](#)
EXAMPLES:Comm_failure_allget2 example, [15](#)
EXAMPLES:Fault-tolerant iterative refinement with shrink and agreement, [16](#)
EXAMPLES:Master example, [15](#)
EXAMPLES:MPI_COMM_AGREE, [13](#), [15](#), [16](#)
EXAMPLES:MPI_COMM_FAILURE_ACK, [15](#)
EXAMPLES:MPI_COMM_FAILURE_GET_ACKED, [15](#)
EXAMPLES:MPI_COMM_FREE, [13](#), [14](#), [16](#)
EXAMPLES:MPI_COMM_GROUP, [14](#)
EXAMPLES:MPI_COMM_REVOKE, [16](#)
EXAMPLES:MPI_COMM_SHRINK, [14](#), [16](#)
EXAMPLES:MPI_COMM_SPLIT, [13](#)

EXAMPLES:MPI_GROUP_DIFFERENCE, [14](#)
EXAMPLES:MPI_GROUP_FREE, [14](#)
MPI_ALLREDUCE, [16](#)
MPI_BCAST, [3](#)
MPI_COMM_ACCEPT, [4](#)
MPI_COMM_AGREE, [4](#), [7–11](#), [13](#), [14](#), [16](#)
MPI_COMM_AGREE(comm, flag), [10](#)
MPI_COMM_CONNECT, [4](#)
MPI_COMM_DISCONNECT, [5](#)
MPI_COMM_DUP, [3](#)
MPI_COMM_FAILURE_ACK, [9](#), [10](#), [14](#)
MPI_COMM_FAILURE_ACK(comm), [9](#)
MPI_COMM_FAILURE_GET_ACKED, [10](#), [14](#)
MPI_COMM_FAILURE_GET_ACKED(comm, failedgrp), [9](#)
MPI_COMM_FREE, [4](#)
MPI_COMM_GET_PARENT, [4](#)
MPI_COMM_IAGREE, [8](#)
MPI_COMM_IAGREE(comm, flag, req), [11](#)
MPI_COMM_JOIN, [5](#)
MPI_COMM_REVOKE, [8](#), [16](#)
MPI_COMM_REVOKE(comm), [7](#)
MPI_COMM_SHRINK, [8](#), [9](#), [14](#)
MPI_COMM_SHRINK(comm, newcomm), [8](#)
MPI_COMM_SPAWN, [4](#), [5](#)
MPI_COMM_SPAWN_MULTIPLE, [4](#), [5](#)
MPI_COMM_SPLIT, [3](#), [8](#), [13](#)
MPI_FILE_CLOSE, [7](#)
MPI_FILE_REVOKE, [13](#)
MPI_FILE_REVOKE(fh), [12](#)
MPI_FILE_SEEK, [7](#)
MPI_FINALIZE, [3](#)
MPI_GET, [6](#)
MPI_GROUP_DIFFERENCE, [10](#)
MPI_INIT, [3](#), [5](#)

MPI_PUT, 6	1
MPI_WIN_FLUSH, 5	2
MPI_WIN_FLUSH_LOCAL, 6	3
MPI_WIN_FLUSH_LOCAL_ALL, 6	4
MPI_WIN_FREE, 6	5
MPI_WIN_GET_FAILED, 14	6
MPI_WIN_GET_FAILED(win, failedgrp), 12	7 8
MPI_WIN_REVOKE, 6 , 12	9
MPI_WIN_REVOKE(win), 11	10 11
TERM:fault tolerance, 1	12
ack, 9 , 10	13
agree, 10 , 11	14
communicator, 3 , 7	15
dynamic process, 4	16
error classes, 13	17
finalize, 2	18
I/O, 7 , 12	19
inquiry, 2	20
mitigation, 7	21
notification, 2 , 13	22
one-sided, 5 , 11	23
process failure, 1	24
revoke, 8 , 11 , 13	25
shrink, 8	26
startup, 2	27
TERM:process failure fault tolerance, 1	28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48