

ADCL: Abstract Data and Communication Library
User level API functions

January 22, 2014

(c) Parallel Software Technologies Laboratory,
Department of Computer Science,
University of Houston, 2007, 2014

Contents

1	Introduction and Terminology	2
2	Environmental control functions	4
2.1	Initializing ADCL	4
2.2	Shutting down ADCL	4
2.3	ADCL program skeletons	5
2.4	ADCL error codes	6
3	High level API	8
3.1	Ibcast	8
3.2	Ialltoall	9
3.3	Reduce	9
3.4	Allreduce	10
3.5	Alltoall	10
3.6	Alltoally	11
3.7	Allgatherv	11
3.8	Example	12
4	Timer object	14
5	Attributes and Attribute-sets	17
6	Functions and Function-sets	19
6.1	Predefined Function-sets	24
6.2	Examples	24
7	Topology	27
7.1	Example	28
8	Vectors and Vector-maps	30
9	Requests	36
9.1	Constructors and Destructors	36
9.2	Initiating communication operations	38
9.3	Retrieving internal components	39
9.4	Save/Restore request status	40
9.5	Examples	41

1 Introduction and Terminology

ADCL is an application level communication library aiming at providing the highest possible performance for application level communication patterns. The library provides for each communication pattern a large number of implementations and incorporates a runtime selection logic in order to choose the implementation leading to the highest performance of the application on the current platform. The library provides multiple different runtime selection algorithms, including a brute force search strategy which tests all available implementations of a given communication pattern; a heuristic relying on orthogonal tuning of attributes characterizing an implementation; and a 2k factorial design search strategy designed for large search spaces and correlated attributes values.

ADCL is not supposed to be a replacement for MPI, but an add-on library. In fact, ADCL relies and exploits many of the features of MPI, and does not offer for example point-to-point transfer primitives. The ADCL API offers high level interfaces of application level collective operations. The high level interfaces are required in order to be able to switch within the library the implementation of the according collective operation without modifying the application itself. Thus, ADCL complements functionality available within MPI. The main objects within the ADCL API are:

- `ADCL_Attributes`, which are an abstraction for particular characteristic of a function/implementation. Each attribute is represented by the set of possible values for this characteristic.
- `ADCL_Attrsets`, which represent a collection of ADCL attributes
- `ADCL_Functions`, each of them being equivalent to an actual implementation of a particular communication pattern. An ADCL Function can have an attribute-set attached to it, in which case the values for each of the attributes in the attribute-set for this particular function have to be defined.
- `ADCL_Fnctsets`, which represent a collection of ADCL functions providing the same functionality. All Functions in a function-set have to have the same attribute-set. ADCL provides pre-defined function sets, such as for neighborhood communication (`ADCL_FNCTSET_NEIGHBORHOOD`) or for shift operations (`ADCL_FNCTSET_SHIFT`). The user can however also register its own functions in order to utilize the ADCL runtime selection logic.

- 1 • **ADCL_Topology** objects, which provides a description of the process
2 topology and neighborhood relations within the application.

- 3 • **ADCL_Vector** which specify the data structures to be used during the
4 communication and the actual data. The user can for example *register*
5 a data structure such as a vector or a matrix with the ADCL library,
6 detailing how many dimensions the object has, the extent of each
7 dimension, which parts of the matrix shall be used for communication,
8 the basic datatype of the object, and the pointer to the data array of
9 the object.

- 10 • **ADCL_Vmap**: which defines portion of a vector to be used in a particular
11 communication operation/

- 12 • **ADCL_Request** objects, which combines a process topology, a function-
13 set and a vector object. The application can initiate a communication
14 by starting a particular ADCL request.

- 15 • **ADCL_Timer** object which allows to register a request with a timer
16 object in order to base the tuning of a functionset on a larger code
17 section instead of the function execution time itself. In addition, the
18 timer object allows to co-tune multiple functionsets within the same
19 code sequence.

20 This document discusses each of the objects in details and explains the
21 user level API functions. Starting from the Fall 2013 release of ADCL, a
22 higher-level API has also been released with provides MPI like interfaces for
23 many of the predefined operations supported by ADCL. In addition, ADCL
24 supports the notion of historic learning for some predefined function sets,
25 i.e. learning across multiple executions of the same operation.

1 **2 Environmental control functions**

2 This section discusses the general functions required to establish the ADCL
3 environment and to shut it down. All ADCL functions return error codes.
4 ADCL leaves it up to the application to take the appropriate actions in case
5 an error occurs. The only exception to that rule is if an error occurs within
6 an MPI function called by ADCL, since MPI's default error behavior is to
7 abort in case of an error. However, the user can change the default behavior
8 of the MPI library by setting the default error handler of `MPI_COMM_WORLD`
9 to `MPI_ERRORS_RETURN` (see also section 7.2 in the MPI-1 [?] specification).

10 ADCL provides C and F90 interfaces for most functions. The Fortran
11 interface of a routines contains an additional argument compared to its
12 C counterpart, namely the error code. Furthermore, all Fortran ADCL
13 objects are defined as integers, leaning on the approach chosen by MPI. A
14 C application has to include the ADCL header file called `ADCL.h`, a Fortran
15 application has to include the file `ADCL.inc` in any routine utilizing ADCL
16 functions.

17 **2.1 Initializing ADCL**

```
18  
19 int ADCL_Init ( void );  
20  
21 subroutine ADCL_Init ( ierror )  
22 integer ierror  
23
```

24 `ADCL_Init` initializes the ADCL execution environment. The function
25 allocates internal data structures required for ADCL, and has to be called
26 therefore before any other ADCL function. Upon success, ADCL returns
27 `ADCL_SUCCESS`. It is recommended to call `ADCL_Init` right after `MPI_Init`.
28 It is erroneous to call `ADCL_Init` multiple times.

29 **2.2 Shutting down ADCL**

```
30  
31 int ADCL_Finalize ( void );  
32  
33 subroutine ADCL_Finalize ( ierror )  
34 integer ierror
```

1 **ADCL_Finalize** finalizes the ADCL environment. Since the function
2 deallocates internal data structures, it should be called at the very end of the
3 application, but before **MPI_Finalize**. It is erroneous to call **ADCL_Finalize**
4 multiple times.

5 **2.3 ADCL program skeletons**

6 Using the two functions described above, the following presents the skeleton
7 for any ADCL application.

```
8 #include <stdio.h>
9 #include "mpi.h"
10 #include "ADCL.h"
11
12 int main ( int argc, char **argv )
13 {
14     MPI_Init ( &argc, &argv );
15     ADCL_Init ();
16
17     ...
18     ADCL_Finalize ();
19     MPI_Finalize ();
20     return 0;
21 }
```

22 Accordingly, the fortran skeleton looks as follows:

```
23 program ADCLskeleton
24     include 'mpif.h'
25     include 'adcl.inc'
26
27     integer ierror
28
29     call MPI_Init ( ierror )
30     call ADCL_Init (ierror )
31
32     ...
33     call ADCL_Finalize ( ierror )
34     call MPI_Finalize ( ierror )
35 end program ADCLskeleton
```

1 **2.4 ADCL error codes**

2 The following is a list of error codes as defined by ADCL.

- 3 ● ADCL_SUCCESS : no error
- 4 ● ADCL_NO_MEMORY: internal memory allocation failed
- 5 ● ADCL_ERROR_INTERNAL : internal ADCL error
- 6 ● ADCL_USER_ERROR: generic user error
- 7 ● ADCL_UNDEFINED: undefined behavior
- 8 ● ADCL_NOT_FOUND : object not found
- 9 ● ADCL_INVALID_ARG : invalid argument passed by user to an ADCL
10 function. Generic error code, only used if one of the codes below do
11 not match.
- 12 ● ADCL_INVALID_NDIMS : invalid number of dimension passed by user to
13 an ADCL function.
- 14 ● ADCL_INVALID_DIMS : invalid dimension passed by user to an ADCL
15 function.
- 16 ● ADCL_INVALID_HWIDTH : invalid number of halo-cells passed by user to
17 an ADCL function.
- 18 ● ADCL_INVALID_DAT: invalid MPI datatype passed by user to an ADCL
19 function.
- 20 ● ADCL_INVALID_DATA: invalid buffer pointer passed by user to an ADCL
21 function.
- 22 ● ADCL_INVALID_COMTYPE: invalid communication type passed by user to
23 an ADCL function.
- 24 ● ADCL_INVALID_COMM: invalid MPI communicator passed by user to an
25 ADCL function.
- 26 ● ADCL_INVALID_REQUEST: invalid ADCL request passed by user to an
27 ADCL function.
- 28 ● ADCL_INVALID_NC : invalid NC argument passed by user to an ADCL
29 function.

- 1 ● ADCL_INVALID_TYPE: ?
- 2 ● ADCL_INVALID_TOPOLOGY: invalid ADCL topology passed by user to an
3 ADCL function.
- 4 ● ADCL_INVALID_ATTRIBUTE: invalid ADCL attribute passed by user to
5 an ADCL function.
- 6 ● ADCL_INVALID_ATTRSET: invalid ADCL attribute-set passed by user to
7 an ADCL function.
- 8 ● ADCL_INVALID_FUNCTION: invalid ADCL function passed by user to an
9 ADCL function.
- 10 ● ADCL_INVALID_WORK_FUNCTION_PTR: invalid ADCL function pointer passed
11 by user to an ADCL function.
- 12 ● ADCL_INVALID_FNCTSET: invalid ADCL function-set passed by user to
13 an ADCL function.
- 14 ● ADCL_INVALID_VECTOR: invalid ADCL vector passed by user to an ADCL
15 function.
- 16 ● ADCL_INVALID_VECTORSET: invalid ADCL vector set passed by user to
17 an ADCL function.
- 18 ● ADCL_INVALID_DIRECTION: invalid direction argument passed by user
19 to an ADCL function.

1 **3 High level API**

2 The high level API in ADCL imitates the classic MPI collective commu-
3 nications API. It allows the user to bypass the creation of `ADCL_Topology`,
4 `ADCL_Vmap` and `ADCL_Vector` when using the ADCL predefined function-sets.
5 So far, it has been implemented for `ADCL_FNCTSET_IBCAST`, `ADCL_FNCTSET_IALLTOALL`,
6 `ADCL_FNCTSET_ALLTOALL`, `ADCL_FNCTSET_ALLTOALLV`, `ADCL_FNCTSET_REDUCE`,
7 `ADCL_FNCTSET_ALLREDUCE` and `ADCL_FNCTSET_ALLGATHERV` function-sets. When
8 the user calls a high level function, all these steps are handled internally, and
9 a persistent ADCL request is created in the same way as `MPI_Send_init`.
10 Cartesian neighborhood communication is as of today the only operation that
11 still has to use the original and slightly more complex API outlined later in
12 the document.

13 Once a persistent communication handle has been created, an actual
14 communication step can be started using either `ADCL_Request_start`) for
15 blocking operations or using a `ADCL_Request_init`) and `ADCL_Request_wait`
16 for non-blocking operations. For more details on the operations
17 permitted on the request object, please refer to section 9. A request
18 object can also be used in combination with the timer object as
19 explained in section 4.

20 **3.1 Ibcast**

21 (Requires a functional implementation of the LibNBC library)

```
22 int ADCL_Ibcast_init ( void *buffer, int count,  
23                      MPI_Datatype datatype, int root, MPI_Comm comm,  
24                      ADCL_Request* req);
```

25 with

- 26 • `buffer(IN/OUT)`: starting address of buffer.
- 27 • `count(IN)`: number of entries in buffer.
- 28 • `datatype(IN)`: data type of buffer.
- 29 • `root(IN)`: rank of broadcast root.
- 30 • `comm(IN)`: MPI communicator.
- 31 • `req(OUT)`: handle to the newly created ADCL request object.

1 **3.2 Ialltoall**

2 (Requires a functional implementation of the LibNBC library)

```
3 int ADCL_Ialltoall_init ( void *sendbuf, int sendcount,  
4     MPI_Datatype sendtype, void *recvbuf, int recvcount,  
5     MPI_Datatype recvtpe, MPI_Comm comm, ADCL_Request* req)
```

6 with

- 7 • sendbuf(IN): starting address of send buffer.
- 8 • sendcount(IN): number of elements to send to each process.
- 9 • sendtype(IN): data type of elements of send buffer.
- 10 • recvbuf(OUT): address of receive buffer.
- 11 • recvcount(IN): number of elements received from any process.
- 12 • recvtpe(IN): data type of elements of receive buffer.
- 13 • comm(IN): MPI communicator.
- 14 • req(OUT): handle to the newly created ADCL request object.

15 **3.3 Reduce**

```
16 int ADCL_Reduce_init ( void *sendbuf, void *recvbuf, int count,  
17     MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm,  
18     ADCL_Request* req)
```

19 with

- 20 • sendbuf(IN): starting address of send buffer.
- 21 • recvbuf(OUT): address of receive buffer.
- 22 • count(IN): number of elements in send buffer.
- 23 • datatype(IN): data type of elements.
- 24 • op(IN): MPI operation to perform (eg. MPI_SUM).
- 25 • root(IN): rank of root process.
- 26 • comm(IN): MPI communicator.
- 27 • req(OUT): handle to the newly created ADCL request object.

1 **3.4 Allreduce**

```
2 int ADCL_Allreduce_init ( void *sendbuf, void *recvbuf, int count,  
3     MPI_Datatype datatype, MPI_Op op, MPI_Comm comm,  
4     ADCL_Request* req)
```

5 with

- 6 • sendbuf(IN): starting address of send buffer.
- 7 • recvbuf(OUT): address of receive buffer.
- 8 • count(IN): number of elements send buffer.
- 9 • datatype(IN): data type of elements.
- 10 • op(IN): MPI operation to perform (eg. MPI_SUM).
- 11 • comm(IN): MPI communicator.
- 12 • req(OUT): handle to the newly created ADCL request object.

13 **3.5 Alltoall**

```
14 int ADCL_Alltoall_init ( void *sendbuf, int sendcount,  
15     MPI_Datatype sendtype, void *recvbuf, int recvcount,  
16     MPI_Datatype recvttype, MPI_Comm comm, ADCL_Request* req)
```

17 with

- 18 • sendbuf(IN): starting address of send buffer.
- 19 • sendcount(IN): number of elements to send to each process.
- 20 • sendtype(IN): data type of elements of send buffer.
- 21 • recvbuf(OUT): address of receive buffer.
- 22 • recvcount(IN): number of elements received from any process.
- 23 • recvttype(IN): data type of receive buffer elements.
- 24 • comm(IN): MPI communicator.
- 25 • req(OUT): handle to the newly created ADCL request object.

1 3.6 Alltoallv

```
2 int ADCL_Alltoallv_init ( void *sendbuf, int* sendcnts, int *sdispls,  
3     MPI_Datatype sendtype, void *recvbuf, int* recvcnts,  
4     int *rdispls, MPI_Datatype recvtype, MPI_Comm comm,  
5     ADCL_Request* req)
```

6 with

- 7 • sendbuf(IN): starting address of send buffer.
- 8 • sendcnts(IN): integer array equal to the group size specifying
9 the number of elements to send to each process.
- 10 • sdispls(IN): integer array (of length group size). Entry j
11 specifies the displacement (relative to sendbuf from which
12 to take the outgoing data destined for process j.
- 13 • sendtype(IN): data type of elements of send buffer.
- 14 • recvbuf(OUT): address of receive buffer.
- 15 • recvcnts(IN): integer array equal to the group size specifying
16 the maximum number of elements that can be received from each
17 process.
- 18 • rdispls(IN): integer array (of length group size). Entry i
19 specifies the displacement (relative to recvbuf) at which to
20 place the incoming data from process i
- 21 • recvtype(IN): data type of receive buffer elements.
- 22 • comm(IN): MPI communicator.
- 23 • req(OUT): handle to the newly created ADCL request object.

24 3.7 Allgatherv

```
25 int ADCL_Allgatherv_init ( void *sendbuf, int sendcount,  
26     MPI_Datatype sendtype, void *recvbuf, int* recvcnts,  
27     int *displs, MPI_Datatype recvtype, MPI_Comm comm,  
28     ADCL_Request* req)
```

29 with

- 1 • `sendbuf(IN)`: starting address of send buffer.
- 2 • `sendcount(IN)`: number of elements in send buffer.
- 3 • `sendtype(IN)`: data type of elements of send buffer.
- 4 • `recvbuf(OUT)`: address of receive buffer.
- 5 • `recvcnts(IN)`: integer array (of length group size) containing
- 6 the number of elements that are to be received from each process.
- 7 • `rdispls(IN)`: integer array (of length group size). Entry `i`
- 8 specifies the displacement (relative to `recvbuf`) at which
- 9 to place the incoming data from process `i`.
- 10 • `recvtype(IN)`: data type of receive buffer elements.
- 11 • `comm(IN)`: MPI communicator.
- 12 • `req(OUT)`: handle to the newly created ADCL request object.

13 **3.8 Example**

14 The following example shows the usage of the ADCL high level API
15 to broadcast an integer.

```
16 #include <stdio.h>
17 #include "mpi.h"
18 #include "ADCL.h"
19 #include "nbc.h"
20
21 int main ( int argc, char ** argv )
22 {
23
24     ADCL_Request request;
25
26     MPI_Init ( &argc, &argv );
27     ADCL_Init ();
28
29     int data = 10;
30
31     /*****/
32
```

```
1   ADCL_Ibcast_init(&data, 1, MPI_INT, 0, MPI_COMM_WORLD, &request);
2
3   /* start the communication. call the following function an
4      arbitrary number of times */
5   ADCL_Request_start ( request );
6
7
8   /* if communication is done, free all handles */
9   ADCL_Request_free ( &request );
10
11  ADCL_Finalize ();
12  MPI_Finalize ();
13  return 0;
14 }
```

1 4 Timer object

2 ADCL tunes a functionset by default based on the time spent in the
3 corresponding function. Users might desire to tune a particular
4 communication operation based on a large codesequence, e.g. an
5 entire iteration of their iterative algorithm or similar, well defined
6 sections in the code. There are three specific instance where it
7 makes sense to tune an operation based on a larger code sequence
8 to potentially includes also computational operations:

- 9 1. for very short running operations, where clock resolution and
10 synchronization of processes might be an issue.
- 11 2. for non-blocking operations, in which (theoratically) the bulk
12 of the communication operation would happen in the background,
13 and thus timing the actual function call is nearly irrelevant.
- 14 3. for code sections which contain multiple operations that shall
15 be tuned simultaniously. Tuning each operation separatly can
16 lead in such a scenario to optimal execution time of each operation
17 individually, but not globally from the application perspective.

18 The timer object allows to register an ADCL_Request with a timer,
19 and add function calls to start and stop the timer outside of the
20 ADCL_Request_start / _init() functions.

```
21 int ADCL_Timer_create ( int nreq, ADCL_Request *reqs, ADCL_Timer *timer);
```

22 with

- 23 • nreqs(IN): number of requests to be associated with the timer
24 object.
- 25 • reqs(IN): array of dimension nreqs containing the array of
26 ADCL_Requests to be associated with the timer object.
- 27 • timer(OUT): handle to ADCL timer object.

```
28 int ADCL_Timer_free ( ADCL_Timer *timer );
```

29 with

1 • timer(INOUT): handle to the timer object allocated with ADCL_Timer-
2 _create. Upon successful completion, the handle will be set
3 to ADCL_TIMER_NULL.

4 Important The timer object should be freed before freeing the
5 requests attached to it, otherwise you can get an unexpected behavior
6 in some cases.

```
7 int ADCL_Timer_start ( ADCL_Timer timer );  
8 int ADCL_Timer_stop ( ADCL_Timer timer );
```

9 start and stop the timer. The execution time between the start
10 and the stop function will be stored with each ADCL Request that
11 has been associated with the provided timer object.

12 The following shows an example on how to associate a request
13 with a timer object and how to utilize it in the code.

```
14 int main ( int argc, char ** argv )  
15 /*****  
16 {  
17     /* General variables */  
18     int rank, size, it;  
19     double datain[10], dataout[10];  
20     ADCL_Request req;  
21     ADCL_Timer timer;  
22  
23     /* Initiate the MPI environment */  
24     MPI_Init ( &argc, &argv );  
25     MPI_Comm_rank ( MPI_COMM_WORLD, &rank );  
26     MPI_Comm_size ( MPI_COMM_WORLD, &size );  
27  
28     /* Initiate the ADCL library */  
29     ADCL_Init ();  
30  
31     /* Initialize datain */  
32     for (it=0; it <10; it++ ) {  
33         datain[it] = ...;  
34         dataout[it] = 0.0;  
35     }  
36  
37     /* Initialize a persistent Allreduce operation */
```



```

1  ADCL_Allreduce_init ( datain, dataout, 10, MPI_DOUBLE, MPI_SUM,
2                          MPI_COMM_WORLD, &req);
3
4  /* define timer object */
5  ADCL_Timer_create ( 1, &req, &timer );
6
7  for (it=0; it<MAXIT; it++){
8      ADCL_Timer_start( timer );
9
10     /* perform some computation */
11     ...
12     /* Start the communication */
13     ADCL_Request_start ( req1 );
14
15     /* perform some more computation */
16     ....
17     ADCL_Timer_stop( timer );
18 }
19
20 ADCL_Timer_free   ( &timer );
21 ADCL_Request_free ( &req1 );
22
23 ADCL_Finalize ();
24 MPI_Finalize ();
25 return 0;
26 }
27

```

1 5 Attributes and Attribute-sets

2 An ADCL_Attribute is an abstraction for a particular characteristic
3 of a function/implementation. Each attribute is represented by
4 a set of possible values. An ADCL_Attrset is a group of ADCL attributes.
5 In the following, we present the constructors and destructors for
6 both ADCL objects.

```
7 int ADCL_Attribute_create ( int maxnvalues, int *array_of_values,  
8     char ** array_of_value_names, char *attr_name,  
9     ADCL_Attribute *attr);
```

10

```
11 subroutine ADCL_Attribute_create ( maxnvalues, array_of_values,  
12     attr, ierror )  
13 integer maxnvalues, attr, ierr  
14 integer array_of_values (*)
```

15 with

- 16 • maxnvalues(IN): number of possible values for this attribute
- 17 • array_of_values(IN): integer array of size maxnvalues containing
18 the possible values for this attribute. The values in this
19 array have to be monotony increasing. However, they do not
20 have to be contiguous.
- 21 • array_of_value_names(IN): character array specifying a name
22 for each attribute value. Not available in the fortran interface.
- 23 • attr_name(IN): name of the attribute. Not available in the
24 fortran interface.
- 25 • attr(OUT): handle to the attribute object

26

```
27 int ADCL_Attribute_free ( ADCL_Attribute *attr );
```

28

```
29 subroutine ADCL_Attribute_free ( attr, ierror )  
30 integer attr, ierror
```

31 with

```

1      • attr(INOUT): handle to the ADCL attribute to be freed. The
2        handle has to be a valid attribute. Upon successful completion,
3        the handle will be set to ADCL_ATTRIBUTE_NULL.

4  int ADCL_Attrset_create ( int maxnum,
5                          ADCL_Attribute *array_of_attrs, ADCL_Attrset *attrset );
6
7  subroutine ADCL_Attrset_create ( maxnum, array_of_attrs,
8                                attrset, ierror )
9  integer maxnum, attrset, ierror
10 integer array_of_attrs (*)

11 with

12      • maxnum(IN): number of attributes to be grouped

13      • array_of_attrs(IN): array of size maxnum containing the handles
14        to attributes. Each entry of the array has to be a valid ADCL
15        attribute created with ADCL_Attribute_create.

16      • attrset(OUT): handle to the attribute-set object

17

18  int ADCL_Attrset_free ( ADCL_Attrset *attrset );
19
20  subroutine ADCL_Attrset_free ( attrset, ierror )
21  integer attrset, ierror

22  with

23      • attrset(INOUT): handle to the ADCL attribute-set to be freed.
24        The handle has to be a valid attribute-set. Upon successful
25        completion, the handle will be set to ADCL_ATTRSET_NULL.

```

1 6 Functions and Function-sets

2 An ADCL_Function is the equivalent to an actual implementation of
3 a particular communication pattern. An ADCL Function can have an
4 attribute-set attached to it, in which case the values for each
5 of the attributes in the attribute-set for this particular function
6 have to be defined. A user can however also decide not to attach
7 an attribute-set to a function by passing in ADCL_ATTRSET_NULL at
8 the particular argument.

9 An ADCL_Fnctset is a collection of ADCL functions providing the
10 same functionality. All Functions in a function-set have to have
11 the same attribute-set. ADCL provides pre-defined function sets,
12 see section 6.1 for a list of predefined function-sets. The user
13 can also register its own functions in order to utilize the ADCL
14 runtime selection logic.

```
15 typedef void ADCL_work_fnct_ptr ( ADCL_Request req );
```

```
16  
17 int ADCL_Function_create ( ADCL_work_fnct_ptr *fnctp,  
18     ADCL_Attrset attrset, int *array_of_attrvalues,  
19     char *name, ADCL_Function *fnct);
```

```
20  
21 subroutine ADCL_Function_create ( fnctp, attrset,  
22     array_of_attrvalues, name, fnct, ierror )
```

```
23 external fnctp  
24 integer attrset, fnct, ierror  
25 integer array_of_attrvalues(*)  
26 char name (*)
```

```
27
```

```
28 with
```

- 29 • fnctp(IN): function pointer to the actual implementation. The
30 prototype has to be of type ADCL_work_fnct_ptr.
- 31 • attrset(IN): valid ADCL attribute-set handle, or ADCL_ATTRSET_NULL.
32 Passing the NULL attribute-set object in forces ADCL to use
33 the brute-force runtime search algorithm.
- 34 • array_of_attrvalues(IN): if an attribute-set has been specified,
35 this array of integers has to provide the values for each attribute
36 in the attribute-set.

- 1 • name(IN): name for the function. The length of the character
- 2 string can not exceed ADCL_MAX_NAMELEN. It is allowed to pass
- 3 in a NULL pointer instead of a string.
- 4 • fnct(OUT): handle to the ADCL function object.

5

```

6 int ADCL_Function_create_async ( ADCL_work_fnct_ptr *init_fnct,
7     ADCL_work_fnct_ptr *wait_fnct,
8     ADCL_Attrset attrset, int *array_of_attrvalues,
9     char *name, ADCL_Function *fnct);
10
11 subroutine ADCL_Function_create_async ( init_fnct, wait_fnct,
12     attrset, array_of_attrvalues, name, fnct, ierror )
13 external init_fnct, wait_fnct
14 integer attrset, fnct, ierror
15 integer array_of_attrvalues(*)
16 char name (*)

```

17 with

- 18 • init_fnct(IN): function pointer to the actual implementation
- 19 of the initiation function. The prototype has to be of type
- 20 ADCL_work_fnct_ptr.
- 21 • wait_fnct(IN): function pointer to the actual implementation
- 22 of the completion function. The prototype has to be of type
- 23 ADCL_work_fnct_ptr.
- 24 • attrset(IN): valid ADCL attribute-set handle, or ADCL_ATTRSET_NULL.
- 25 Passing the NULL attribute-set object in forces ADCL to use
- 26 the brute-force runtime search algorithm.
- 27 • array_of_attrvalues(IN): if an attribute-set has been specified,
- 28 this array of integers has to provide the values for each attribute
- 29 in the attribute-set.
- 30 • name(IN): name for the function. The length of the character
- 31 string can not exceed ADCL_MAX_NAMELEN. It is allowed to pass
- 32 in a NULL pointer instead of a string.
- 33 • fnct(OUT): handle to the ADCL function object.

```

1
2 int ADCL_Function_free ( ADCL_Function *fnct );
3
4 subroutine ADCL_Function_free ( fnct, ierror )
5 integer fnct, ierror
6 with
7     • fnct(INOUT): valid handle to an ADCL function. Upon return,
8       the handle is set to ADCL_FUNCTION_NULL
9
10 int ADCL_Fnctset_create ( int maxnum, ADCL_Function *fncts,
11                          char *name, ADCL_Fnctset *fnctset );
12
13 subroutine ADCL_Fnctset_create ( maxnum, fncts, name,
14                                fnctset, ierror )
15 integer maxnum, fnctset, ierror
16 integer fncts(*)
17 char name (*)
18 with
19     • maxnum(IN): number of ADCL functions to be bundled to a function
20       set
21     • fncts(IN): array of size maxnum containing the handles to the
22       ADCL functions. All functions have to provide the same attribute-set.
23     • name(IN): name of the function set. It is allowed to pass
24       in a NULL pointer instead of a name. The length of the character
25       string can not exceed ADCL_MAX_NAMELEN.
26     • fnctset(OUT): handle for the ADCL function-set.
27
28 The following function provides a short cut for scenarios, where
29 a single function shall be executed with different attribute values.
30 Thus, instead of the user having to create all the individual functions
31 and register them with ADCL, this interface provides the opportunity
32 to specify the function pointer, an attribute set, and get an Functionset
33 handle back. ADCL will create internally a function with the same
34 function pointer for each possible combination of attribute values.
35 Some combinations of attribute values can be also excluded.

```

```

1 int ADCL_Fnctset_create_single ( ADCL_work_fnct_ptr *init_fnct,
2     ADCL_work_fnct_ptr *wait_fnct, ADCL_Attrset attrset,
3     char *name, int **without_attr_combinations,
4     int num_without_attr_combinations,
5     ADCL_Fnctset *fnctset );
6
7 subroutine ADCL_Fnctset_create_single ( void *init_fnct,
8     void *wait_fnct, int *attrset, char *name,
9     int *without_attr_vals,
10    int *num_without_attr_vals,
11    int *fnctset, int *ierror, int name_len );
12 external init_fnct, wait_fnct
13 integer attrset, without_attr_vals, num_without_attr_vals
14 integer fnctset, ierror
15 char name (*)
16
17 with
18
19     • init_fnct(IN): function pointer to the actual implementation
20       of the init function. The prototype has to be of type ADCL_work_fnct_ptr.
21
22     • wait_fnct(IN): function pointer to the actual implementation
23       of the completion function. The prototype has to be of type
24       ADCL_work_fnct_ptr. It is allowed to pass in a NULL pointer
25       if there is no need to a completion function.
26
27     • attrset(IN): valid ADCL attribute-set handle, or ADCL_ATTRSET_NULL.
28       It is illegal for this function to pass in the NULL attribute-set
29       handle.
30
31     • name(IN): name for the function. The length of the character
32       string can not exceed ADCL_MAX_NAMELEN. It is allowed to pass
33       in a NULL pointer instead of a string.
34
35     • without_attr_vals(IN): array containing the attribute values
36       that the user want to exclude.
37
38     • num_without_attr_vals(IN): number of entries in the array containing
39       the attribute values that the user want to exclude.
40
41     • fnctset(OUT): handle to the ADCL function-set object.
42
43
44
45
46
47
48
49
50
51
52
53
54

```

```
1 int ADCL_Functset_free ( ADCL_Functset *fctset );
2
3 subroutine ADCL_Functset_free ( fctset, ierror )
4 integer fctset, ierror
5 with
6     • fctset(INOUT): valid ADCL function-set handle to be freed.
7     After successful completion, the handle is set to ADCL_FNCTSET_NULL.
```


1 **6.1 Predefined Function-sets**

2 ADCL currently supports the following predefined function-sets:

- 3 • ADCL_FNCTSET_NEIGHBORHOOD: a function-set supporting n-dimensional
4 neighborhood communication. The dimension of the neighborhood
5 communication is determined by the dimension of the vector
6 object used during the request creation time.
- 7 • ADCL_FNCTSET_ALLGATHERV: a function-set supporting an MPI_Allgather
8 style communication.
- 9 • ADCL_FNCTSET_ALLREDUCE: a function-set supporting an MPI_Allreduce
10 style communication.
- 11 • ADCL_FNCTSET_REDUCE: a function-set supporting an MPI_Reduce
12 style communication.
- 13 • ADCL_FNCTSET_ALLTOALL: a function-set supporting an MPI_Alltoall
14 style communication.
- 15 • ADCL_FNCTSET_ALLTOALLV: a function-set supporting an MPI_Alltoallv
16 style communication.
- 17 • ADCL_FNCTSET_IBCAST: a function-set supporting a non-blocking
18 broadcast operation similarly to MPI_Ibcast. Requires a functional
19 implementation of the LibNBC library.
- 20 • ADCL_FNCTSET_IALLTOALL: a function-set supporting a non-blocking
21 all-to-all operation similarly to MPI_IAlltoall. Requires a
22 functional implementation of the LibNBC library.

23 Other predefined function-sets might be added in following releases
24 of ADCL.

25 **6.2 Examples**

26 The following is an example for a code registering three functions
27 without an attribute-set, and combining them to a function-set.
28 Since there is no attribute-set attached to the functions, the array
29 of attribute-values argument of the ADCL function constructors can
30 be a NULL pointer. Please note, that not setting an attribute-set
31 forces ADCL to fall back to the brute force runtime selection logic.

```

1  #include <stdio.h>
2  #include "ADCL.h"
3  #include "mpi.h"
4
5  void test_func_1 ( ADCL_Request req );
6  void test_func_2 ( ADCL_Request req );
7  void test_func_3 ( ADCL_Request req );
8
9  int main ( int argc, char ** argv )
10 {
11     ADCL_Function funcs[3];
12     ADCL_Fnctset fctset;
13     int i;
14
15     MPI_Init ( &argc, &argv );
16     ADCL_Init ();
17
18     ADCL_Function_create ( (ADCL_work_fnct_ptr *)test_func_1,
19         ADCL_ATTRSET_NULL, NULL, "test_func_1", &(funcs[0]));
20     ADCL_Function_create ( (ADCL_work_fnct_ptr *)test_func_2,
21         ADCL_ATTRSET_NULL, NULL, "test_func_2", &(funcs[1]));
22     ADCL_Function_create ( (ADCL_work_fnct_ptr *)test_func_3,
23         ADCL_ATTRSET_NULL, NULL, "test_func_3", &(funcs[2]));
24
25     ADCL_Fnctset_create ( 3, funcs, "trivial functions",
26         &fctset );
27
28     /* Do something with the fctset */
29
30     ADCL_Fnctset_free ( &fctset );
31     for ( i=0; i<3; i++ ) {
32         ADCL_Function_free ( &funcs[i] );
33     }
34     ADCL_Finalize ();
35     MPI_Finalize ();
36     return 0;
37 }
38
39 void test_func_1 ( ADCL_Request req ) {

```

```
1     int rank;
2
3     MPI_Comm_rank ( MPI_COMM_WORLD, &rank );
4     printf("%d: In test_func_1 \n", rank);
5     return;
6 }
7
8 void test_func_2 ( ADCL_Request req ) {
9     int rank;
10
11     MPI_Comm_rank ( MPI_COMM_WORLD, &rank );
12     printf("%d: In test_func_2\n", rank);
13     return;
14 }
15
16 void test_func_3 ( ADCL_Request req ) {
17     int rank;
18
19     MPI_Comm_rank ( MPI_COMM_WORLD, &rank );
20     printf("%d: In test_func_3\n", rank);
21     return;
22 }
23
```

1 7 Topology

2 An ADCL_Topology objects contains the description of the process
3 topology and neighborhood relations within the application. The
4 reason ADCL introduces this abstraction and does not rely entirely
5 on the MPI cartesian communicators is, that many codes have a multi-dimensional
6 process distribution, where the process dimensions are running in
7 a different order compared to a cartesian MPI communicator. Thus,
8 ADCL provides a topology constructor using MPI cartesian communicators,
9 but also a generic constructor, where the user can define for each
10 process who its neighbors are.

```
11 int ADCL_Topology_create ( MPI_Comm cart_comm,  
12     ADCL_Topology *topo);
```

13

```
14 subroutine ADCL_Topology_create ( cart_comm, topo, ierror )  
15 integer cart_comm, topo, ierror
```

16 with

17 • cart_comm(IN): MPI cartesian communicator. A call to MPI_Topo_test
18 has to return MPI_CART for cart_comm to be valid in this function
19 call.

20 • topo(OUT): handle of an ADCL topology object.

21

```
22 int ADCL_Topology_create_generic ( int ndims, int *lneighbors,  
23     int *rneighbors, int *coords, int direction,  
24     MPI_Comm comm, ADCL_Topology *topo);
```

25

```
26 subroutine ADCL_Topology_create_generic ( ndims, lneighbors,  
27     rneighbors, coords, direction, comm, topo, ierror )  
28 integer ndims, direction, comm, topo, ierror  
29 integer lneighbors(*), rneighbors(*), coords(*)
```

30 with

31 • ndims(IN): number of dimensions of the process topology.

32 • lneighbors(IN): integer array of dimension ndims containing
33 the ranks of the left neighbors for each dimension. In case

```

1     a left neighbor does not exist (e.g. the process is at the
2     boundary of the process topology), the according entry in the
3     array has to be set to MPI_PROC_NULL.

4     • rneighbors(IN): integer array of dimension ndims containing
5     the ranks of the right neighbors for each dimension. In case
6     a right neighbor does not exist (e.g. the process is at the
7     boundary of the process topology), the according entry in the
8     array has to be set to MPI_PROC_NULL.

9     • coords(IN): integer array of dimension ndims containing the
10    coordinates of the process in the process topology.

11    • direction(IN): ?

12    • comm(IN): valid MPI communicator. In contrary to the previous
13    function, this does not have to be an MPI cartesian communicator.

14    • topo(OUT): handle of an ADCL topology object.

15

16 int ADCL_Topology_free ( ADCL_Topology *topo );

17 with

18    • topo(INOUT): ADCL topology object to be released. Upon successful
19    completion, the handle will be set to ADCL_TOPOLOGY_NULL.

```

20 7.1 Example

21 The following example shows how to register a 2-D process topology
22 with ADCL for an arbitrary number of processes using MPI cartesian
23 communicators.

```

24 #include <stdio.h>
25 #include "ADCL.h"
26 #include "mpi.h"
27
28 int main ( int argc, char ** argv )
29 {
30     MPI_Comm cart_comm;
31     ADCL_Topology topo;

```

```
1     int cdims[]={0,0}, periods[]={0,0};
2
3     MPI_Init ( &argc, &argv );
4     ADCL_Init ();
5
6     MPI_Dims_create ( size, 2, cdims );
7     MPI_Cart_create ( MPI_COMM_WORLD, 2, cdims, periods, 0,
8                     &cart_comm);
9     ADCL_Topology_create ( cart_comm, &topo );
10
11    /* do something useful with the topology object */
12
13    ADCL_Topology_free ( &topo );
14    MPI_Comm_free ( &cart_comm );
15
16    ADCL_Finalize ();
17    MPI_Finalize ();
18    return 0;
19 }
```

1 8 Vectors and Vector-maps

2 An ADCL_Vector combined with the ADCL_Vmap specifies the data structures
3 to be used during the communication and the actual data.

4 For the collectives defined in the MPI standard it is noticed
5 that the arguments of the parameters can be separated into three
6 groups: information concerning the data (buffer, data type), concerning
7 the process topology (communicator, root) and related to the communication
8 pattern (element counts, reduction operation, array of element counts
9 or displacements). As an example, for the MPI_Bcast interface,
10 the information about the data consists of the buffer and the data
11 type, root and the MPI communicator comm give information about
12 the process topology and count is related to the communication pattern.
13 As each of the parameters of the MPI collectives falls into one
14 of these three groups, they form the basis of the ADCL objects:
15 information concerning the data is stored in the ADCL vector object,
16 information concerning the process topology in the ADCL topology
17 object and information related to the communication pattern becomes
18 part of the new vmap object. As of today, there are five different
19 Vmap types defined in ADCL:'

- 20 • HALO: specifies the number of halo-cells in each direction
21 (hwidth).
- 22 • ALL: a value specifying the number of elements used in the
23 communication (count).
- 24 • REDUCE: a value specifying the number of elements used in the
25 communication (count) along with an operation (op) to be performed
26 on the data.
- 27 • LIST: an array of values specifying the number of elements
28 communicated to each process, with the element at position
29 i in the array used to communicate to rank i . In addition,
30 an array of displacements has to be provided following the
31 same rules as for the array of counts arguments.
- 32 • INPLACE: an abstract vmap object that can be used similarly
33 to MPI_INPLACE as an argument to some function sets.

34 The following are the (brief) interfaces to allocate each Vmap
35 object. For the sake of brevity we do not detail all arguments
36 at this point.

```

1 int ADCL_Vmap_halo_allocate ( int hwidth, ADCL_Vmap *vec );
2 int ADCL_Vmap_list_allocate ( int size, int* rcnts, int* displ,
3                               ADCL_Vmap *vec );
4 int ADCL_Vmap_allreduce_allocate ( MPI_Op op, ADCL_Vmap *vec );
5 int ADCL_Vmap_reduce_allocate ( MPI_Op op, ADCL_Vmap *vec );
6 int ADCL_Vmap_alltoall_allocate ( int scnt, int rcnt, ADCL_Vmap *vec );
7 int ADCL_Vmap_all_allocate ( ADCL_Vmap *vec );
8 int ADCL_Vmap_inplace_allocate ( ADCL_Vmap *vec );
9 int ADCL_Vmap_free ( ADCL_Vmap *vec );

```

10 Similarly to all other objects, the Vmap object can be freed
11 after utilization using

```

12 int ADCL_Vmap_free ( ADCL_Vmap *vmap )

```

13 ADCL distinguishes between allocating a vector and registering
14 a vector. The difference is, that in the first case, the library
15 allocates the memory for the object as specified by the user, while
16 in the second case the memory has to be allocated by the application.

```

17 int ADCL_Vector_allocate_generic ( int ndims, int *dims, int nc,
18                                   ADCL_Vmap vmap, MPI_Datatype dat,
19                                   void *data, ADCL_Vector *vec )
20

```

21 with

- 22 ● `ndims(IN)`: number of dimension of the data structure excluding
23 the `nc` argument detailed below.
- 24 ● `dims(IN)`: array of dimension `ndims` containing the extent of
25 each dimension including the halo-cells.
- 26 ● `nc(IN)`: number of elements per entry. Since many scientific
27 codes have matrices where each entry of the matrix is a submatrix
28 itself, this argument gives the possibility to specify the
29 dimension of the submatrix.
- 30 ● `vmap(IN)`: vector map to be used for this vector.
- 31 ● `dat(IN)`: basic MPI datatype describing the data type of the
32 matrix

1 • data(OUT): pointer to the buffer allocated. This pointer will
2 be required for calculations within the user code, since the
3 buffer has been allocated by ADCL. ADCL guarantees, that a
4 contiguous memory location has been allocated for multi-dimensional
5 arrays. Please note, that the buffer pointer is **not** the beginning
6 of the data array, but the pointer to the multi-dimensional
7 matrix itself. As an example if ndims=2 and the temporary
8 variable used within ADCL to allocate the 2-D array is called
9 tmp_matrix, then the buffer pointer returned in this argument
10 is data = tmp_matrix, which is not equal to data = &(tmp_matrix[0][0])
11 in C!

12 • vec(OUT): handle to ADCL vector object.

13 There is no Fortran interface defined for this routine. For a discussion,
14 why no Fortran interface is provided for this routine, please refer
15 to the discussion of MPI Alloc_mem in the MPI-2 [?] specification
16 section 4.1.1.

17

```
18 int ADCL_Vector_register_generic ( int ndims, int *dims, int nc,  
19                                 ADCL_Vmap vmap, MPI_Datatype dat,  
20                                 void *data, ADCL_Vector *vec )
```

21

```
22 subroutine ADCL_Vector_register_generic ( ndims, dims, nc,  
23                                           vmap, dat, data, vec, ierror )
```

```
24 integer ndims, nc, vmap, dat, vec, ierror
```

```
25 integer dims(*)
```

```
26 TYPE data(*)
```

27 with

28 • ndims(IN): number of dimension of the data structure excluding
29 the nc argument detailed below.

30 • dims(IN): array of dimension ndims containing the extent of
31 each dimension including the halo-cells.

32 • nc(IN): number of elements per entry. Since many scientific
33 codes have matrices where each entry of the matrix is a submatrix
34 itself, this argument gives the possibility to specify the
35 dimension of the submatrix.

- 1 • vmap(IN): variable describing which parts of the matrix can
2 be used for communication. The Vmap object is typically the
3 result of an ADCL_Vmap_allocate* function as shown above.
- 4 • dat(IN): basic MPI datatype describing the data type of the
5 matrix
- 6 • data(IN): pointer to the data array. Please note, that the
7 buffer pointer has to be the pointer to the multi-dimensional
8 matrix itself, and not the beginning of the data array. As
9 an example if ndims=2 and the variable used within application
10 is defined as double tmp_matrix[10][10], the buffer pointer
11 passed to ADCL has to be tmp_matrix, which is not equal to &(tmp_matrix[0][0])
12 in C!
- 13 • vec(OUT): handle to ADCL vector object.

14

```
15 int ADCL_Vector_free ( ADCL_Vector *vec );
```

16 with

- 17 • vec(INOUT): handle to the vector object allocated with ADCL_Vector-
18 _allocate. It is illegal to call this function with a vector
19 object registered with ADCL_Vector_register. Upon successful
20 completion, the handle will be set to ADCL_VECTOR_NULL.

21 There is no Fortran interface defined for this routine. For a discussion,
22 why no Fortran interface is provided for this routine, please refer
23 to the discussion of MPI_Alloc_mem in the MPI-2 [?] specification
24 section 4.1.1.

```
25 int ADCL_Vector_deregister ( ADCL_Vector *vec );
```

26

```
27 subroutine ADCL_Vector_deregister ( vec, ierror )  
28 integer vec, ierror
```

29 with

- 30 • vec(INOUT): handle to the vector object registered with ADCL_Vector-
31 _register. It is illegal to call this function with a vector
32 object allocated with ADCL_Vector_allocate. Upon successful
33 completion, the handle will be set to ADCL_VECTOR_NULL.

```

1     The following example assumes, that each entry of this matrix
2     contains of a vector with five elements.

3     #include <stdio.h>
4     #include "ADCL.h"
5     #include "mpi.h"
6
7     /* Dimensions of the data matrix per process */
8     #define DIM0  8
9     #define DIM1  4
10
11    int main ( int argc, char ** argv )
12    {
13        int dims[2];
14        int nc=5, hwidth=1;
15        double ***data;
16        ADCL_Vector vec;
17        ADCL_Vmap vmap;
18
19        MPI_Init ( &argc, &argv );
20        ADCL_Init ();
21
22        dims[0] = DIM0 + 2*hwidth;
23        dims[1] = DIM1 + 2*hwidth;
24        ADCL_Vmap_halo_allocate ( hwidth, &vmap);
25        ADCL_Vector_allocate_generic ( 2,  dims, nc, vmap,
26            MPI_DOUBLE, &data, &vec );
27
28        /* now you can access the elements of the vector in order to perform
29           computations, e.g */
30        for ( i=1; i<DIM0; i++ ) {
31            for ( j=1; j<DIM1; j++ ) {
32                for ( k=0; k < nc; k++ ) {
33                    data[i][j][k] = ...
34                }
35            }
36
37            ADCL_Vector_free ( &vec );
38            ADCL_Vmap_free ( &vmap );
39            ADCL_Finalize ();

```

```
1     MPI_Finalize ();  
2     return 0;  
3 }
```

1 9 Requests

2 An ADCL_Request object combines a process topology, a function-set
3 and a vector object to a single user-level handle. Using this handle,
4 the application can initiate a communication by starting a particular
5 ADCL request. In the following we will detail request constructors
6 and destructors, the available functions for initiating a communication,
7 and reflection functions in order to retrieve certain components
8 of an ADCL request from a user level function.

9 9.1 Constructors and Destructors

```
10 int ADCL_Request_create ( ADCL_Vector vec, ADCL_Topology topo,  
11                          ADCL_Fnctset fnctset, ADCL_Request *req );
```

```
12  
13 subroutine ADCL_Request_create ( vec, topo, fnctset, req, ierror)  
14 integer vec, topo, fnctset, req, ierror
```

15 with

- 16 • `vec(IN)`: vector object identifying the data arrays used for
17 the communication. The vector object can be `ADCL_VECTOR_NULL`
18 for a user defined function-set following the requirements
19 and restrictions of the according function-set. However, for
20 the predefined function-sets, the vector object has to be provided
21 following the dimension matching rule explained below.
- 22 • `topo(IN)`: topology object identifying the neighboring processes.
23 The topology object can be `ADCL_TOPOLOGY_NULL` for user defined
24 function-set following the requirements and restrictions of
25 the according function-set. However, for the predefined function-sets,
26 the topology object has to be provided following the dimension
27 matching rule explained below.
- 28 • `fnctset(IN)`: ADCL function-set to be evaluated
- 29 • `req(OUT)`: handle to the newly created ADCL request object

30
31 *Dimension matching rule:* The dimension of the vector object has
32 to be equal to the dimension of the process topology. The user
33 has one additional degree of freedom by using the `nc` argument in

1 vector creation, which allows to match a vector of dimension $n+$
2 1 to a n dimensional process topology, by not distributing the last
3 dimension of the vector.

4 *Remark on the dimension matching rule:* It is fairly straight forward
5 to define a more generic vector object, where the user could define
6 a multi-dimensional vector and assign to each of the dimensions
7 a flag which indicates whether the data is distributed accordingly
8 in this dimension or not. The solution applied so far restricts
9 this approach by reducing the dimension of a matrix which is not
10 distributed across the processes to the last dimension. This is
11 due to our findings, that it is a common approach in many codes.
12 However, if an end-user faces a problem requiring the generic solution,
13 please contact the authors of ADCL. In any case, the number of dimension
14 of the matrix distributed across the process has to match the number
15 of dimensions of the process topology.

```
16 int ADCL_Request_create_generic ( ADCL_Vectset vecset,  
17     ADCL_Topology topo, ADCL_Fnctset fnctset,  
18     ADCL_Request *req );
```

```
19  
20 subroutine ADCL_Request_create_generic ( vecset, topo,  
21     fnctset, req, ierror )  
22 integer vecset, topo, fnctset, req, ierror
```

23 with

- 24 • vecset(IN): vector set to be used in the request. The extent
25 of the array of send vectors and receive vectors has to match
26 the number of dimensions of the topology object.
- 27 • topo(IN): topology object identifying the neighboring processes.
28 The topology object can be ADCL_TOPOLOGY_NULL for user defined
29 function-set following the requirements and restrictions of
30 the according function-set. However, for the predefined function-sets,
31 the topology object has to be provided following the dimension
32 matching rule explained below.
- 33 • fnctset(IN): ADCL function-set to be evaluated
- 34 • req(OUT): handle to the newly created ADCL request object

35

```
1 int ADCL_Request_free ( ADCL_Request *req );
2
3 subroutine ADCL_Request_free ( req, ierror )
4 integer req, ierror
5
6 with
7     • req(INOUT): valid ADCL request handle. After successful completion,
8       the handle will be set to ADCL_REQUEST_NULL.
```

9 9.2 Initiating communication operations

10 The following function can be used to start a blocking instance
11 of the communication pattern.

```
12 int ADCL_Request_start ( ADCL_Request req );
13
14 subroutine ADCL_Request_start ( req, ierror )
15 integer req, ierror
16
17 with
18     • req(IN): valid ADCL request
```

19 The following two functions represent a non-blocking version
20 of ADCL_Request_start. There are no restrictions in ADCL on how
21 many ADCL requests can be ongoing simultaneously. However, for
22 a single request a call to ADCL_Request_init has to be followed by
23 a call to ADCL_Request_wait.

```
24 int ADCL_Request_init ( ADCL_Request req );
25
26 subroutine ADCL_Request_init ( req, ierror )
27 integer req, ierror
28
29 with
30     • req(IN): valid ADCL request
```



```

1 with
2     • req(IN): valid ADCL request handle
3     • function_name(OUT): Name of the function currently being called.
4       This string has to be freed by the user level function.
5     • attribute_names(OUT):
6     • num_of_attributes(OUT): number of attributes. This value is
7       equal to the extent of the attribute_names, attribute_values_names
8       and attribute_values arrays.
9     • attribute_values_names(OUT): array of names for the attribute
10      values of the current function.
11     • attribute_values(OUT): array of integers containing the attribute
12      values of the current function.
13

```

14 9.4 Save/Restore request status

15 The following function helps the user to save the progress in the
16 search process of the fastest implementation when using the brute
17 force strategy. Then, the user can simply resume the search from
18 where he stopped by restoring the request status using the function
19 ADCL_Request_restore_status.

```

20 int ADCL_Request_save_status ( ADCL_Request req, int *tested_num,
21                               double **unfiltered_avg, double **filtered_avg,
22                               double **outliers, int *winner_so_far );

```

```

23 with
24     • req(IN): valid ADCL request handle
25     • tested_num(OUT): number of already evaluated ADCL functions.
26     • unfiltered_avg(OUT): unfiltered average of the evaluated ADCL
27       functions.
28     • filtered_avg(OUT): filtered average of the evaluated ADCL functions.

```

```

1      • outliers(OUT): number of outliers detected in the evaluated
2      ADCL functions.

3      • winner_so_far(OUT): the winning implementation so far.

4  int ADCL_Request_restore_status ( ADCL_Request req, int tested_num,
5      double *unfiltered_avg, double *filtered_avg,
6      double *outliers );

7  with

8      • req(IN): valid ADCL request handle

9      • tested_num(IN): number of already evaluated ADCL functions.

10     • unfiltered_avg(IN): unfiltered average of each evaluated ADCL
11     functions.

12     • filtered_avg(IN): filtered average of each evaluated ADCL functions.

13     • outliers(IN): number of outliers detected in the evaluated
14     ADCL functions.

```

15 9.5 Examples

16 The following example demonstrates the usage of the predefined function
17 set ADCL_FNCTSET_NEIGHBORHOOD for a three dimensional matrix and
18 process topology.

```

19 #include <stdio.h>
20 #include "ADCL.h"
21 #include "mpi.h"
22
23 /* Dimensions of the data matrix per process */
24 #define DIM0 64
25 #define DIM1 32
26 #define DIM2 32
27
28 int main ( int argc, char ** argv )
29 {
30     int hwidth, dims[3];
31     double ***data;
32     int cdims[]={0,0,0};

```

```

1   int periods[]={0,0,0};
2
3   MPI_Comm cart_comm;
4   ADCL_Vector vec;
5   ADCL_Topology topo;
6   ADCL_Request request;
7
8   MPI_Init ( &argc, &argv );
9   ADCL_Init ();
10
11  MPI_Dims_create ( size, 3, cdims );
12  MPI_Cart_create ( MPI_COMM_WORLD, 3, cdims, periods, 0, &cart_comm);
13  ADCL_Topology_create ( cart_comm, &topo );
14
15  /*****
16  /* Test 1: hwidth=1, nc=0 */
17  hwidth=1;
18  dims[0] = DIM0 + 2*hwidth;
19  dims[1] = DIM1 + 2*hwidth;
20  dims[2] = DIM2 + 2*hwidth;
21  ADCL_Vector_allocate ( 3, dims, 0, hwidth, MPI_DOUBLE, &data, &vec );
22  ADCL_Request_create ( vec, topo, ADCL_FNCTSET_NEIGHBORHOOD, &request );
23
24  /* start the computation and communication. call the following function
25     an arbitrary number of times */
26  ADCL_Request_start ( request );
27
28
29  /* if computation/communication is done, free all handles */
30  ADCL_Request_free ( &request );
31  ADCL_Vector_free ( &vec );
32  ADCL_Topology_free ( &topo );
33  MPI_Comm_free ( &cart_comm );
34
35  ADCL_Finalize ();
36  MPI_Finalize ();
37  return 0;
38  }

```