UPC-IO: A Parallel I/O API for UPC

V1.02

Tarek El-Ghazawi
François Cantonnet
Proshanta Saha
Yiyi Yao
The George Washington University
801 22nd Street NW • Suite 607
Washington,DC 20052, USA
{tarek, fcantonn, sahap, yyy}@gwu.edu

Rajeev Thakur
Rob Ross
Mathematics and Computer Science Division
Argonne National Laboratory
9700 S. Cass Avenue
Argonne, IL 60439, USA
{thakur, rross}@mcs.anl.gov

Dan Bonachea
Dept. of Computer Science
University of California, Berkeley
Berkeley, CA 94720, USA
bonachea@cs.berkeley.edu

Sept 25, 2006

1

Contents

A   Proposed Additions and Extensions

1   This section contains proposed additions and extensions to the UPC specification. Such proposals are included when stable enough for developers to implement and for users to study and experiment with them. However, their presence does not suggest long term support. When fully stable and tested, they will be moved to the main body of the specification.

2   This section also describes the process used to add new items to the specification, which starts with inclusion in this section. Requirements for inclusion are:[1]

1. A documented API which shall use the format and conventions of this specification and [ISO/IEC00].

2. Either a complete, publicly available, implementation of the API or a set of publicly available example programs which demonstrate the interface.

3. The concurrence of the UPC consortium that its inclusion would be in the best interest of the language.

3   If all implementations drop support for an extension and/or all interested parties no longer believe the extension is worth pursuing, then it may simply be dropped. Otherwise, the requirements for inclusion of an extension in the main body of the specification are:

1. Six months residence in this section.

2. The existence of either one (or more) publicly available "reference" implementation written in standard UPC OR at least two independent implementations (possibly specific to a given UPC implementation).

---

[1]These requirements ensure that most of the semantic issues that arise during initial implementation have been addressed and prevents the accumulation of interfaces that no one commits to implement. Nothing prevents the circulation of more informal *what if* interface proposals from circulating in the community before an extension reaches this point.

3. The existence of a significant base of experimental user experience which demonstrates positive results with a substantial portion of the proposed API.

4. The concurrence of the UPC consortium that its inclusion would be in the best interest of the language.

4   For each extension, there shall be a predefined *feature macro* beginning with `__UPC` which will be defined by an implementation to be the interface version of the extension if it is supported, otherwise undefined.

A.1 UPC Parallel I/O <upc_io.h>

1 This subsection provides the UPC parallel extensions of [ISO/IEC00 Sec 7.19]. All the characteristics of library functions described in [ISO/IEC00 Sec 7.1.4] apply to these as well. Implementations that support this interface shall predefine the feature macro `__UPC_IO__` to the value 1.

Common Constraints

2 All UPC-IO functions are collective and must be called by all threads collectively.[2]

3 If a program calls `exit`, `upc_global_exit`, or returns from `main` with a UPC file still open, the file will automatically be closed at program termination, and the effect will be equivalent to `upc_all_fclose` being implicitly called on the file.

4 If a program attempts to read past the end of a file, the read function will read data up to the end of file and return the number of bytes actually read, which may be less than the amount requested.

5 Writing past the end of a file increases the file size.

6 If a program seeks to a location past the end of a file and writes starting from that location, the data in the intermediate (unwritten) portion of the file is undefined. For example, if a program opens a new file (of size 0 bytes), seeks to offset 1024 and writes some data beginning from that offset, the data at offsets 0–1023 is undefined. Seeking past the end of file and performing a write causes the current file size to immediately be extended up to the end of the write. However, just seeking past the end of file or attempting to read past the end of file, without a write, does not extend the file size.

7 All generic pointers-to-shared passed to the I/O functions (as function arguments or indirectly through the list I/O arguments) are treated as if they had a phase field of zero (that is, the input phase is ignored).

8 All UPC-IO read/write functions take an argument `flags` of type `upc_flag_t`. The semantics of this argument is defined in Section **??**. These semantics apply only to memory locations in user-provided buffers, not to the read/write

---

[2]Note that collective does not necessarily imply barrier synchronization. The synchronization behavior of the UPC-IO data movement library functions is explicitly controlled by using the `flags` flag argument. See Section **??** for details.

operations on the storage medium or any buffer memory internal to the library implementation.

9   The `flags` flag is included even on the fread/fwrite_local functions (which take a pointer-to-local as the buffer argument) in order to provide well-defined semantics for the case where one or more of the pointer-to-local arguments references a shared object (with local affinity). In the case where all of the pointer-to-local arguments in a given call reference only private objects, the `flags` flag provides no useful additional guarantees and is recommended to be passed as `UPC_IN_NOSYNC|UPC_OUT_NOSYNC` to maximize performance.

10   The UPC-IO read/write functions can be either a synchronous or an asynchronous operation. The default is set to a synchronous operation. If the `UPC_ASYNC` flag is specified as a parameter to the `upc_flag_t` argument, the function will be an asynchronous operation.

11   The arguments to all UPC-IO functions are single-valued except where explicitly noted in the function description.

12   UPC-IO, by default, supports weak consistency and atomicity semantics. The default (weak) semantics are as follows. The data written to a file by a thread is only guaranteed to be visible to another thread after all threads have collectively closed or synchronized the file.

13   Writes to a file from a given thread are always guaranteed to be visible to subsequent file reads by the *same* thread, even without an intervening call to collectively close or synchronize the file.

14   Byte-level data consistency is supported.

15   If concurrent writes from multiple threads overlap in the file, the resulting data in the overlapping region is undefined with the weak consistency and atomicity semantics

16   When reading data being concurrently written by another thread, the data that gets read is undefined with the weak consistency and atomicity semantics.

17   File reads into overlapping locations in a shared buffer in memory using individual file pointers or list I/O functions leads to undefined data in the target buffer under the weak consistency and atomicity semantics.

18   A given file must not be opened at same time by the POSIX I/O and UPC-IO

libraries.

19 Except where otherwise noted, all UPC-IO functions return NON-single-valued errors; that is, the occurrence of an error need only be reported to at least one thread, and the `errno` value reported to each such thread may differ. When an error is reported to ANY thread, the position of ALL file pointers for the relevant file handle becomes undefined.

20 The error values that UPC-IO functions may set in `errno` are implementation-defined, however the `perror()` and `strerror()` functions are still guaranteed to work properly with `errno` values generated by UPC-IO.

21 UPC-IO functions can not be called between `upc_notify` and corresponding `upc_wait` statements.

A.1.1   Background

A.1.1.1   File Accessing and File Pointers

1 Collective UPC-IO accesses can be done in and out of shared and private buffers, thus local and shared reads and writes are generally supported. In each of these cases, file pointers could be either common or individual. Note that in UPC-IO, common file pointers cannot be used in conjunction with pointer-to-local buffers. File pointer modes are specified by passing a flag to the collective `upc_all_fopen` function and can be changed using `upc_all_fcntl`. When a file is opened with the common file pointer flag, all threads share a common file pointer. When a file is opened with the individual file pointer flag, each thread gets its own file pointer.

2 UPC-IO also provides file-pointer-independent list file accesses by specifying explicit offsets and sizes of data that is to be accessed. List IO can also be used with either pointer-to-local buffers or pointer-to-shared buffers.

A.1.1.2   Synchronous and Asynchronous I/O

1 I/O operations can be synchronous (blocking) or asynchronous (non-blocking). While synchronous calls are quite simple and easy to use from a programming point of view, asynchronous operations allow the overlapping of computation

and I/O to achieve improved performance. Synchronous calls block and wait until the corresponding I/O operation is completed. On the other hand, an asynchronous call starts an I/O operation and returns immediately. Thus, the executing process can turn its attention to other processing needs while the I/O is progressing.

2   UPC-IO supports both synchronous and asynchronous I/O functionality. The asynchronous I/O operations have the restriction that only one (collective) asynchronous operation can be active at a time on a given file handle. That is, an asynchronous I/O operation must be completed by calling `upc_all_ftest_async` or `upc_all_fwait_async` before another asynchronous I/O operation can be called on the same file handle. This restriction is similar to the restriction MPI-IO [MPI2] has on split-collective I/O functions: only one split collective operation can be outstanding on an MPI-IO file handle at any time.

A.1.1.3   Consistency and Atomicity Semantics

1   The consistency semantics define when the data written to a file by a thread is visible to other threads. The atomicity semantics define the outcome of operations in which multiple threads write concurrently to a file or shared buffer and some of the writes overlap each other. For performance reasons, UPC-IO uses weak consistency and atomicity semantics by default. The user can select stronger semantics either by opening the file with the flag `UPC_STRONG_CA` or by calling `upc_all_fcntl` with the command `UPC_SET_STRONG_CA_SEMANTICS`.

2   The default (weak) semantics are as follows. The data written by a thread is only guaranteed to be visible to another thread after all threads have called `upc_all_fclose` or `upc_all_fsync`. (Note that the data *may* be visible to other threads before the call to `upc_all_fclose` or `upc_all_fsync` and that the data may become visible to different threads at different times.) Writes from a given thread are always guaranteed to be visible to subsequent reads by the *same* thread, even without an intervening call to `upc_all_fclose` or `upc_all_fsync`. Byte-level data consistency is supported. So for example, if thread 0 writes one byte at offset 0 in the file and thread 1 writes one byte at offset 1 in the file, the data from both threads will get written to the file. If concurrent writes from multiple threads overlap in the file, the resulting data in the overlapping region is undefined. Similarly, if one thread tries to

8

read the data being concurrently written by another thread, the data that gets read is undefined. Concurrent in this context means any two read/write operations to the same file handle with no intervening calls to `upc_all_fsync` or `upc_all_fclose`.

3   For the functions that read into or write from a shared buffer using a common file pointer, the weak consistency semantics are defined as follows. Each call to `upc_all_{fread,fwrite}_shared[_async]` with a common file pointer behaves as if the read/write operations were performed by a single, distinct, anonymous thread which is different from any compute thread (and different for each operation). In other words, NO file reads are guaranteed to see the results of file writes using the common file pointer until after a close or sync under the default weak consistency semantics.

4   By passing the `UPC_STRONG_CA` flag to `upc_all_fopen` or by calling `upc_all_fcntl` with the command `UPC_SET_STRONG_CA_SEMANTICS`, the user selects strong consistency and atomicity semantics. In this case, the data written by a thread is visible to other threads as soon as the file write on the calling thread returns. In the case of writes from multiple threads to overlapping regions in the file, the result would be as if the individual write function from each thread occurred atomically in some (unspecified) order. Overlapping writes to a file in a single (list I/O) write function on a single thread are not permitted (see Section A.2). While strong consistency and atomicity semantics are selected on a given file handle, the `flags` argument to all fread/fwrite functions on that handle is ignored and always treated as `UPC_IN_ALLSYNC | UPC_OUT_ALLSYNC`.

5   The consistency semantics also define the outcome in the case of overlapping reads into a shared buffer in memory when using individual file pointers or list I/O functions. By default, the data in the overlapping space is undefined. If the user selects strong consistency and atomicity, the result would be as if the individual read functions from each thread occurred atomically in some (unspecified) order. Overlapping reads into memory buffers in a single (list I/O) read function on a single thread are not permitted (see Section A.2).

6   Note that in strong consistency and atomicity mode, atomicity is guaranteed at the UPC-IO function level. The entire operation specified by a single function is performed atomically, regardless of whether it represents a single, contiguous read/write or multiple noncontiguous reads or writes as in a list I/O function.

9

A.1.1.4   File Interoperability

1   UPC-IO does not specify how an implementation may store the data in a file
on the storage device. Accordingly, it is implementation-defined whether or
not a file created by UPC-IO can be directly accessed by using C/POSIX I/O
functions. However, the UPC-IO implementation must specify how the user
can retrieve the file from the storage system as a linear sequence of bytes
and vice versa. Similarly, the implementation must specify how familiar
operations, such as the equivalent of POSIX `ls`, `cp`, `rm`, and `mv` can be
performed on the file.

A.1.2   Predefined Types

1   The following types are defined in `<upc_io.h>`

2   `upc_off_t` is a signed integral type that is capable of representing the size of
the largest file supported by the implementation.

3   `upc_file_t` is an opaque shared data type of incomplete type (as defined in
[ISO/IEC00 Sec 6.2.5]) that represents an open file handle.

4   `upc_file_t` objects are always manipulated via a pointer (that is, `upc_file_t`
`*`).

5   `upc_file_t` is a shared data type. It is allowed to pass a (`upc_file_t *`)
across threads, and two pointers to `upc_file_t` that reference the same logical
file handle will always compare equal.

A.1.3   UPC File Operations

Common Constraints

1   When a file is opened with an individual file pointer, each thread will get its
own file pointer and advance through the file at its own pace.

2   When a common file pointer is used, all threads positioned in the file will be
aligned with that common file pointer.

3   Common file pointers cannot be used in conjunction with pointers-to-local
(and hence cannot operate on private objects).

4   No function in this section may be called while an asynchronous operation is pending on the file handle, except where otherwise noted.

A.1.3.1   The `upc_all_fopen` function

Synopsis

1
```
#include <upc.h>
#include <upc_io.h>
upc_file_t *upc_all_fopen(const char *fname, int flags,
        size_t numhints, struct upc_hint const *hints);
```

Description

2   `upc_all_fopen` opens the file identified by the filename `fname` for input/output operations.

3   The flags parameter specifies the access mode. The valid flags and their meanings are listed below. Of these flags, exactly one of `UPC_RDONLY`, `UPC_WRONLY`, or `UPC_RDWR`, and one of `UPC_COMMON_FP` or `UPC_INDIVIDUAL_FP`, must be used. Other flags are optional. Multiple flags can be combined by using the bitwise OR operator (|), and each flag has a unique bitwise representation that can be unambiguously tested using the bitwise AND operator(&).

    `UPC_RDONLY` Open the file in read-only mode

    `UPC_WRONLY` Open the file in write-only mode

    `UPC_RDWR` Open the file in read/write mode

    `UPC_INDIVIDUAL_FP` Use an individual file pointer for all file accesses (other than list I/O)

    `UPC_COMMON_FP` Use the common file pointer for all file accesses (other than list I/O)

    `UPC_APPEND` Set the *initial* position of the file pointer to end of file. (The file pointer is not moved to the end of file after each read/write)

UPC_CREATE Create the file if it does not already exist. If the named file does not exist and this flag is not passed, the function fails with an error.

UPC_EXCL Must be used in conjunction with UPC_CREATE. The open will fail if the file already exists.

UPC_STRONG_CA Set strong consistency and atomicity semantics

UPC_TRUNC Open the file and truncate it to zero length. The file must be opened before writing.

UPC_DELETE_ON_CLOSE Delete the file automatically on close

4   The UPC_COMMON_FP flag specifies that all accesses (except for the list I/O operations) will use the common file pointer. The UPC_INDIVIDUAL_FP flag specifies that all accesses will use individual file pointers (except for the list I/O operations). Either UPC_COMMON_FP or UPC_INDIVIDUAL_FP must be specified or upc_all_fopen will return an error.

5   The UPC_STRONG_CA flag specifies strong consistency and atomicity semantics. The data written by a thread is visible to other threads as soon as the write on the calling thread returns. In the case of writes from multiple threads to overlapping regions in the file, the result would be as if the individual write function from each thread occurred atomically in some (unspecified) order. In the case of overlapping reads into a shared buffer in memory when using individual file pointers or list I/O functions, the result would be as if the individual read functions from each thread occurred atomically in some (unspecified) order.

6   If the flag UPC_STRONG_CA is not specified, weak semantics are provided. The data written by a thread is only guaranteed to be visible to another thread after all threads have called upc_all_fclose or upc_all_fsync. (Note that the data *may* be visible to other threads before the call to upc_all_fclose or upc_all_fsync and that the data may become visible to different threads at different times.) Writes from a given thread are always guaranteed to be visible to subsequent reads by the *same* thread, even without an intervening call to upc_all_fclose or upc_all_fsync. Byte-level data consistency is supported. For the purposes of atomicity and consistency semantics, each call to upc_all_{fread,fwrite}_shared[_async] with a common file pointer

12

behaves as if the read/write operations were performed by a single, distinct, anonymous thread which is different from any compute thread (and different for each operation)."[3]

7   Hints can be passed to the UPC-IO library as an array of key-value pairs[4] of strings. `numhints` specifies the number of hints in the `hints` array; if `numhints` is zero, the `hints` pointer is ignored. The user can free the `hints` array and associated character strings as soon as the open call returns. The following type is defined in `<upc_io.h>`:

```
struct upc_hint
```

holds each element of the `hints` array and contain at least the following initial members, in this order.

```
const char *key;
const char *value;
```

8   `upc_all_fopen` defines a number hints. An implementation is free to support additional hints. An implementation is free to ignore any hint provided by the user. Implementations should *silently* ignore any hints they do not support or recognize. The predefined hints and their meanings are defined in A.3. An implementation is not required to interpret these hint keys, but if it does interpret the hint key, it must provide the functionality described. All hints are single-valued character strings (the content is single-valued, not the location).

9   A file on the storage device is in the *open* state from the beginning of a successful open call to the end of the matching successful close call on the file handle. It is erroneous to have the same file *open* simultaneously with two `upc_all_fopen` calls, or with a `upc_all_fopen` call and a POSIX/C `open` or `fopen` call.

10   The user is responsible for ensuring that the file referenced by the `fname` argument refers to a single UPC-IO file. The actual argument passed on

---

[3]In other words, NO reads are guaranteed to see the results of writes using the common file pointer until after a close or sync when `UPC_STRONG_CA` is not specified.

[4]The contents of the key/value pairs passed by all the threads must be single-valued.

each thread may be different because the file name spaces may be different on different threads, but they must all refer to the same logical UPC-IO file.

11 On success, the function returns a pointer to a file handle that can be used to perform other operations on the file.

12 `upc_all_fopen` provides single-valued errors - if an error occurs, the function returns `NULL` on ALL threads, and sets `errno` appropriately to the same value on all threads.

A.1.3.2 The `upc_all_fclose` function

Synopsis

1
```
#include <upc.h>
#include <upc_io.h>
int upc_all_fclose (upc_file_t *fd);
```

Description

2 `upc_all_fclose` executes an implicit `upc_all_fsync` on `fd` and then closes the file associated with `fd`.

3 The function returns 0 on success. If `fd` is not valid or if an outstanding asynchronous operation on `fd` has not been completed, the function will return an error.

4 `upc_all_fclose` provides single-valued errors - if an error occurs, the function returns –1 on ALL threads, and sets `errno` appropriately to the same value on all threads.

5 After a file has been closed with `upc_all_fclose`, the file is allowed to be opened and the data in it can be accessed by using regular C/POSIX I/O calls.

A.1.3.3 The `upc_all_fsync` function

Synopsis

1

```
#include <upc.h>
#include <upc_io.h>
int upc_all_fsync(upc_file_t *fd);
```

Description

2   `upc_all_fsync` ensures that any data that has been written to the file asso-
ciated with `fd` but not yet transferred to the storage device is transferred to
the storage device. It also ensures that subsequent file reads from any thread
will see any previously written values (that have not yet been overwritten).

3   There is an implied barrier immediately before `upc_all_fsync` returns.

4   The function returns 0 on success. On error, it returns –1 and sets `errno`
appropriately.

A.1.3.4 The `upc_all_fseek` function

Synopsis

1
```
#include <upc.h>
#include <upc_io.h>
upc_off_t upc_all_fseek(upc_file_t *fd, upc_off_t offset,
        int origin);
```

Description

2   `upc_all_fseek` sets the current position of the file pointer associated with
`fd`.

3   This offset can be relative to the current position of the file pointer, to the
beginning of the file, or to the end of the file. The offset can be negative,
which allows seeking backwards.

4   The origin parameter can be specified as UPC_SEEK_SET, UPC_SEEK_CUR, or
UPC_SEEK_END, respectively, to indicate that the offset must be computed
from the beginning of the file, the current location of the file pointer, or the
end of the file.

5   In the case of a common file pointer, all threads must specify the same offset

and origin. In the case of an individual file pointer, each thread may specify a different offset and origin.

6   It is allowed to seek past the end of file. It is erroneous to seek to a negative position in the file. See the Common Constraints number 5 at the beginning of Section A.1.3 for more details.

7   The current position of the file pointer can be determined by calling `upc_all_fseek(fd, 0, UPC_SEEK_CUR)`.

8   On success, the function returns the current location of the file pointer in bytes. If there is an error, it returns –1 and sets `errno` appropriately.

A.1.3.5   The `upc_all_fset_size` function

Synopsis

1

```
#include <upc.h>
#include <upc_io.h>
int upc_all_fset_size(upc_file_t *fd, upc_off_t size);
```

Description

2   `upc_all_fset_size` executes an implicit `upc_all_fsync` on `fd` and resizes the file associated with `fd`. The file handle must be open for writing.

3   `size` is measured in bytes from the beginning of the file.

4   If `size` is less than the current file size, the file is truncated at the position defined by `size`. The implementation is free to deallocate file blocks located beyond this position.

5   If `size` is greater than the current file size, the file size increases to `size`. Regions of the file that have been previously written are unaffected. The values of data in the new regions in the file (between the old size and size) are undefined.

6   If this function truncates a file, it is possible that the individual and common file pointers may point beyond the end of file. This is allowed and is equivalent to seeking past the end of file (see the Common Constraints in Section A.1.3 for the semantics of seeking past the end of file).

7   It is unspecified whether and under what conditions this function actually allocates file space on the storage device. Use `upc_all_fpreallocate` to force file space to be reserved on the storage device.

8   Calling this function does not affect the individual or common file pointers.

9   The function returns 0 on success. On error, it returns –1 and sets `errno` appropriately.

A.1.3.6   The `upc_all_fget_size` function

Synopsis

1
```
#include <upc.h>
#include <upc_io.h>
upc_off_t upc_all_fget_size(upc_file_t *fd);
```

Description

2   `upc_all_fget_size` returns the current size in bytes of the file associated with `fd` on success. On error, it returns –1 and sets `errno` appropriately.

A.1.3.7   The `upc_all_fpreallocate` function

Synopsis

1
```
#include <upc.h>
#include <upc_io.h>
int upc_all_fpreallocate(upc_file_t *fd, upc_off_t size);
```

Description

2   `upc_all_fpreallocate` ensures that storage space is allocated for the first `size` bytes of the file associated with `fd`. The file handle must be open for writing.

3   Regions of the file that have previously been written are unaffected. For newly allocated regions of the file, `upc_all_fpreallocate` has the same effect as writing undefined data.

4   If `size` is greater than the current file size, the file size increases to `size`. If `size` is less than or equal to the current file size, the file size is unchanged.

5   Calling this function does not affect the individual or common file pointers.

6   The function returns 0 on success. On error, it returns –1 and sets `errno` appropriately.

A.1.3.8   The `upc_all_fcntl` function

Synopsis

1
```
#include <upc.h>
#include <upc_io.h>
int upc_all_fcntl(upc_file_t *fd,  int cmd, void *arg);
```

Description

2   `upc_all_fcntl` performs one of various miscellaneous operations related to the file specified by `fd`, as determined by `cmd`. The valid commands `cmd` and their associated argument `arg` are explained below.

>   UPC_GET_CA_SEMANTICS Get the current consistency and atomicity semantics for `fd`. The argument `arg` is ignored. The return value is `UPC_STRONG_CA` for strong consistency and atomicity semantics and 0 for the default weak consistency and atomicity semantics.

>   UPC_SET_WEAK_CA_SEMANTICS Executes an implicit `upc_all_fsync` on `fd` and sets `fd` to use the weak consistency and atomicity semantics (or leaves the mode unchanged if that mode is already selected). The argument `arg` is ignored. The return value is 0 on success. On error, this function returns -1 and sets `errno` appropriately.

>   UPC_SET_STRONG_CA_SEMANTICS Executes an implicit `upc_all_fsync` on `fd` and sets `fd` to use the strong consistency and atomicity semantics (or leaves the mode unchanged if that mode is already selected). The argument `arg` is ignored. The return value is 0 on success. On error, this function returns -1 and sets `errno` appropriately.

UPC_GET_FP Get the type of the current file pointer for `fd`. The argument `arg` is ignored. The return value is either `UPC_COMMON_FP` in case of a common file pointer, or `UPC_INDIVIDUAL_FP` for individual file pointers.

UPC_SET_COMMON_FP Executes an implicit `upc_all_fsync` on `fd`, sets the current file access pointer mechanism for `fd` to a common file pointer (or leaves it unchanged if that mode is already selected), and seeks to the beginning of the file. The argument `arg` is ignored. The return value is 0 on success. On error, this function returns -1 and sets `errno` appropriately.

UPC_SET_INDIVIDUAL_FP Executes an implicit `upc_all_fsync` on `fd`, sets the current file access pointer mechanism for `fd` to an individual file pointer (or leaves the mode unchanged if that mode is already selected), and seeks to the beginning of the file. The argument `arg` is ignored. The return value is 0 on success. On error, this function returns -1 and sets `errno` appropriately.

UPC_GET_FL Get all the flags specified during the `upc_all_fopen` call for `fd`, as modified by any subsequent mode changes using the `upc_all_fcntl(UPC_SET_*)` commands. The argument `arg` is ignored. The return value has same format as the `flags` parameter in `upc_all_fopen`.

UPC_GET_FN Get the file name provided by each thread in the `upc_all_fopen` call that created `fd`. The argument `arg` is a valid (`const char**`) pointing to a (`const char*`) location in which a pointer to file name will be written. Writes a (`const char*`) into `*arg` pointing to the filename in implementation-maintained read-only memory, which will remain valid until the file handle is closed or until the next `upc_all_fcntl` call on that file handle.

UPC_GET_HINTS Retrieve the hints applicable to `fd`. The argument `arg` is a valid (`const upc_hint_t**`) pointing to a (`const upc_hint_t*`) location in which a pointer to the hints array will be written. Writes a (`const upc_hint_t*`) into `*arg` pointing to an array of `upc_hint_t`'s in implementation-maintained read-only memory, which will remain valid until the file handle is closed or until the next `upc_all_fnctl` call on that file handle. The number of hints in the array is returned by the

call. The hints in the array may be a subset of those specified at file open time, if the implementation ignored some unrecognized or unsupported hints.

UPC_SET_HINT Executes an implicit upc_all_fsync on fd and sets a new hint to fd. The argument arg points to one single-valued upc_hint_t hint to be applied. If the given hint key has already been applied to fd, the current value for that hint is replaced with the provided value. The return value is 0 on success. On error, this function returns -1 and sets errno appropriately.

UPC_ASYNC_OUTSTANDING Returns 1 if there is an asynchronous operation outstanding on fd, or 0 otherwise.

3   In case of a non valid fd, upc_all_fcntl returns -1 and sets errno appropriately.

4   It *is* allowed to call upc_all_fcntl(UPC_ASYNC_OUTSTANDING) when an asynchronous operation is outstanding (but it is still disallowed to call upc_all_fcntl with any other argument when an asynchronous operation is outstanding).

A.1.4   Reading Data

Common Constraints

1   No function in this section A.1.4 may be called while an asynchronous operation is pending on the file handle.

A.1.4.1   The upc_all_fread_local function

Synopsis

1
```
#include <upc.h>
#include <upc_io.h>
upc_off_t upc_all_fread_local(upc_file_t *fd, void *buffer,
    size_t size, size_t nmemb, upc_flag_t flags);
```

Description

2   `upc_all_fread_local` reads data from a file into a local buffer on each thread.

3   This function can be called only if the current file pointer type is an individual file pointer, and the file handle is open for reading.

4   `buffer` is a pointer to an array into which data will be read, and each thread may pass a different value for `buffer`.

5   Each thread reads `(size*nmemb)` bytes of data from the file at the position indicated by its individual file pointer into the buffer. Each thread may pass a different value for `size` and `nmemb`. If `size` or `nmemb` is zero, the `buffer` argument is ignored and that thread performs no I/O.

6   On success, the function returns the number of bytes read into the local buffer of the calling thread, and the individual file pointer of the thread is incremented by that amount. On error, it returns –1 and sets `errno` appropriately.

A.1.4.2   The `upc_all_fread_shared` function

Synopsis

1
```
#include <upc.h>
#include <upc_io.h>
upc_off_t upc_all_fread_shared(upc_file_t *fd,
    shared void *buffer, size_t blocksize, size_t size,
    size_t nmemb, upc_flag_t flags);
```

Description

2   `upc_all_fread_shared` reads data from a file into a shared buffer in memory.

3   The function can be called when the current file pointer type is either a common file pointer or an individual file pointer. The file handle must be open for reading.

4   `buffer` is a pointer to an array into which data will be read. It must be a pointer to shared data and may have affinity to any thread. `blocksize` is the block size of the shared buffer in elements (of `size` bytes each). A `blocksize` of 0 indicates an indefinite blocking factor.

5   In the case of individual file pointers, the following rules apply: Each thread may pass a different address for the `buffer` parameter. Each thread reads (`size*nmemb`) bytes of data from the file at the position indicated by its individual file pointer into its buffer. Each thread may pass a different value for `blocksize`, `size` and `nmemb`. If `size` or `nmemb` is zero, the `buffer` argument is ignored and that thread performs no I/O. On success, the function returns the number of bytes read by the calling thread, and the individual file pointer of the thread is incremented by that amount.

6   In the case of a common file pointer, the following rules apply: All threads must pass the same address for the `buffer` parameter, and the same value for `blocksize`, `size` and `nmemb`. The effect is that (`size*nmemb`) bytes of data are read from the file at the position indicated by the common file pointer into the buffer. If `size` or `nmemb` is zero, the `buffer` argument is ignored and the operation has no effect. On success, the function returns the total number of bytes read by all threads, and the common file pointer is incremented by that amount.

7   If reading with individual file pointers results in overlapping reads into the shared buffer, the result is determined by whether the file was opened with the `UPC_STRONG_CA` flag or not (see Section A.1.3.1).

8   The function returns –1 on error and sets `errno` appropriately.

A.1.5   Writing Data

Common Constraints

1   No function in this section A.1.5 may be called while an asynchronous operation is pending on the file handle.

A.1.5.1   The `upc_all_fwrite_local` function

Synopsis

1
```
#include <upc.h>
#include <upc_io.h>
upc_off_t upc_all_fwrite_local(upc_file_t *fd, void *buffer,
    size_t size, size_t nmemb, upc_flag_t flags);
```

Description

2   `upc_all_fwrite_local` writes data from a local buffer on each thread into a file.

3   This function can be called only if the current file pointer type is an individual file pointer, and the file handle is open for writing.

4   `buffer` is a pointer to an array from which data will be written, and each thread may pass a different value for `buffer`.

5   Each thread writes `(size*nmemb)` bytes of data from the buffer to the file at the position indicated by its individual file pointer. Each thread may pass a different value for `size` and `nmemb`. If `size` or `nmemb` is zero, the `buffer` argument is ignored and that thread performs no I/O.

6   If any of the writes result in overlapping accesses in the file, the result is determined by the current consistency and atomicity semantics mode in effect for `fd` (see A.1.3.1).

7   On success, the function returns the number of bytes written by the calling thread, and the individual file pointer of the thread is incremented by that amount. On error, it returns –1 and sets `errno` appropriately.

A.1.5.2   The `upc_all_fwrite_shared` function

Synopsis

1
```
    #include <upc.h>
    #include <upc_io.h>
    upc_off_t upc_all_fwrite_shared(upc_file_t *fd,
        shared void *buffer, size_t blocksize,  size_t size,
        size_t nmemb, upc_flag_t flags);
```

Description

2   `upc_all_fwrite_shared` writes data from a shared buffer in memory to a file.

3   The function can be called if the current file pointer type is either a common file pointer or an individual file pointer. The file handle must be open for writing.

4   `buffer` is a pointer to an array from which data will be written. It must be a pointer to shared data and may have affinity to any thread. `blocksize` is the block size of the shared buffer in elements (of `size` bytes each). A `blocksize` of 0 indicates an indefinite blocking factor.

5   In the case of individual file pointers, the following rules apply: Each thread may pass a different address for the `buffer` parameter. Each thread writes (`size*nmemb`) bytes of data from its buffer to the file at the position indicated by its individual file pointer. Each thread may pass a different value for `blocksize`, `size` and `nmemb`. If `size` or `nmemb` is zero, the `buffer` argument is ignored and that thread performs no I/O. On success, the function returns the number of bytes written by the calling thread, and the individual file pointer of the thread is incremented by that amount.

6   In the case of a common file pointer, the following rules apply: All threads must pass the same address for the `buffer` parameter, and the same value for `blocksize`, `size` and `nmemb`. The effect is that (`size*nmemb`) bytes of data are written from the buffer to the file at the position indicated by the common file pointer. If `size` or `nmemb` is zero, the `buffer` argument is ignored and the operation has no effect. On success, the function returns the total number of bytes written by all threads, and the common file pointer is incremented by that amount.

7   If writing with individual file pointers results in overlapping accesses in the file, the result is determined by the current consistency and atomicity semantics mode in effect for `fd` (see Section A.1.3.1).

8   The function returns –1 on error and sets `errno` appropriately.

A.1.6   Asynchronous I/O

Common Constraints

1   Only one asynchronous I/O operation can be outstanding on a UPC-IO file handle at any time. If an application attempts to initiate a second asynchronous I/O operation while one is still outstanding on the same file handle the behavior is undefined – however, high-quality implementations will issue a fatal error.

2   For asynchronous read operations, the contents of the destination memory are

undefined until after a successful `upc_all_fwait_async` or `upc_all_ftest_async` on the file handle. For asynchronous write operations, the source memory may not be safely modified until after a successful `upc_all_fwait_async` or `upc_all_ftest_async` on the file handle.

3 An implementation is free to block for completion of an operation in the asynchronous initiation call or in the `upc_all_ftest_async` call (or both). High-quality implementations are recommended to minimize the amount of time spent within the asynchronous initiation or `upc_all_ftest_async` call.

4 In the case of list I/O functions, the user may modify or free the lists after the asynchronous I/O operation has been initiated.

5 The semantics of the flags of type `upc_flag_t` when applied to the async variants of the fread/fwrite functions should be interpreted as follows: constraints that reference entry to a function call correspond to entering the fread_async-/fwrite_async call that initiates the asynchronous operation, and constraints that reference returning from a function call correspond to returning from the `upc_all_fwait_async()` or successful `upc_all_ftest_async()` call that completes the asynchronous operation. Also, note that the flags which govern an asynchronous operation are passed to the library during the asynchronous initiation call.

A.1.6.1 The `upc_all_fwait_async` function

Synopsis

1
```
#include <upc.h>
#include <upc_io.h>
upc_off_t upc_all_fwait_async(upc_file_t *fd)
```

Description

2 `upc_all_fwait_async` completes the previously issued asynchronous I/O operation on the file handle `fd`, blocking if necessary.

3 It is erroneous to call this function if there is no outstanding asynchronous I/O operation associated with `fd`.

4 On success, the function returns the number of bytes read or written by the asynchronous I/O operation as specified by the blocking variant of the

function used to initiate the asynchronous operation. On error, it returns –1 and sets `errno` appropriately, and the outstanding asynchronous operation (if any) becomes no longer outstanding.

A.1.6.2   The `upc_all_ftest_async` function

Synopsis

1
```
#include <upc.h>
#include <upc_io.h>
upc_off_t upc_all_ftest_async(upc_file_t *fd, int *flag)
```

Description

2   `upc_all_ftest_async` tests whether the outstanding asynchronous I/O operation associated with `fd` has completed.

3   If the operation has completed, the function sets `flag=1` and the asynchronous operation becomes no longer outstanding;[5] otherwise it sets `flag=0`. The same value of `flag` is set on all threads.

4   If the operation was completed, the function returns the number of bytes that were read or written as specified by the blocking variant of the function used to initiate the asynchronous operation. On error, it returns –1 and sets `errno` appropriately, and sets the `flag=1`, and the outstanding asynchronous operation (if any) becomes no longer outstanding.

5   It is erroneous to call this function if there is no outstanding asynchronous I/O operation associated with `fd`.

---

[5]This implies it is disallowed to call `upc_all_fwait_async` or `upc_all_ftest_async` immediately after a successful `upc_all_ftest_async` on that file handle.

A.2 List I/O

This subsection describes optional extensions to the UPC-IO interface defined in section . High-quality implementations of UPC-IO are encouraged to provide this interface to support advanced users and library writers. Implementations that support this interface shall predefine the feature macro tt __UPC_LIST_IO__ to the value 1.

Common Constraints

1   List I/O functions take a list of addresses/offsets and corresponding lengths in memory and file to read from or write to.

2   List I/O functions can be called regardless of whether the current file pointer type is individual or common.

3   File pointers are not updated as a result of a list I/O read/write operation.

4   Types declared in `<upc_io.h>`are

```
struct upc_local_memvec
```

which contains at least the initial members, in this order:

```
void *baseaddr;
size_t len;
```

and is a memory vector element pointing to a contiguous region of local memory.

5   `struct upc_shared_memvec`

which contains at least the initial members, in this order:

```
shared void *baseaddr;
size_t blocksize;
size_t len;
```

and is a memory vector element pointing to a blocked region of shared memory.

6   27

```
struct upc_filevec
```

which contains at least the initial members, in this order:

```
upc_off_t offset;
size_t len;
```

and is a file vector element pointing to a contiguous region of a file.

For all cases these vector element types specify regions which are `len` bytes long. If `len` is zero, the entry is ignored. `blocksize` is the block size of the shared buffer in bytes. A `blocksize` of 0 indicates an indefinite blocking factor.

7   The `memvec` argument passed to any list I/O *read* function by a single thread must not specify overlapping regions in memory.

8   The base addresses passed to `memvec` can be in any order.

9   The `filevec` argument passed to any list I/O *write* function by a single thread must not specify overlapping regions in the file.

10 The offsets passed in `filevec` must be in monotonically non-decreasing order.

11 No function in this section (A.2) may be called while an asynchronous operation is pending on the file handle.

12 No function in this section (A.2) implies the presence of barriers at entry or exit. However, the programmer is advised to use a barrier after calling `upc_all_fread_list_shared` to ensure that the entire shared buffer has been filled up, and similarly, use a barrier before calling `upc_all_fwrite_list_shared` to ensure that the entire shared buffer is up-to-date before being written to the file.

13 For all the list I/O functions, each thread passes an independent set of memory and file vectors. Passing the same vectors on two or more threads specifies redundant work. The file pointer is a single-valued argument, all other arguments to the list I/O functions are NOT single-valued.

A.2.1 The `upc_all_fread_list_local` function

Synopsis

1
```
#include <upc.h>
#include <upc_io.h>
upc_off_t upc_all_fread_list_local(upc_file_t *fd,
    size_t memvec_entries, struct upc_local_memvec const *memvec,
    size_t filevec_entries, struct upc_filevec const *filevec,
    upc_flag_t flags);
```

Description

2   `upc_all_fread_list_local` reads data from a file into local buffers in memory. The file handle must be open for reading.

3   `memvec_entries` indicates the number of entries in the array `memvec` and `filevec_entries` indicates the number of entries in the array `filevec`. The values may be 0, in which case the `memvec` or `filevec` argument is ignored and no locations are specified for I/O.

4   The result is as if data were read in order from the list of locations specified by `filevec` and placed in memory in the order specified by the list of locations in `memvec`. The total amount of data specified by `memvec` must equal the total amount of data specified by `filevec`.

5   On success, the function returns the number of bytes read by the calling thread. On error, it returns –1 and sets `errno` appropriately.

A.2.2 The `upc_all_fread_list_shared` function

Synopsis

1
```
#include <upc.h>
#include <upc_io.h>
```

```
upc_off_t upc_all_fread_list_shared(upc_file_t *fd,
    size_t memvec_entries, struct upc_shared_memvec const *memvec,
    size_t filevec_entries, struct upc_filevec const *filevec,
    upc_flag_t flags);
```

Description

2   `upc_all_fread_list_shared` reads data from a file into various locations of a shared buffer in memory. The file handle must be open for reading.

3   `memvec_entries` indicates the number of entries in the array `memvec` and `filevec_entries` indicates the number of entries in the array `filevec`. The values may be 0, in which case the `memvec` or `filevec` argument is ignored and no locations are specified for I/O.

4   The result is as if data were read in order from the list of locations specified by `filevec` and placed in memory in the order specified by the list of locations in `memvec`. The total amount of data specified by `memvec` must equal the total amount of data specified by `filevec`.

5   If any of the reads from different threads result in overlapping regions in memory, the result is determined by the current consistency and atomicity semantics mode in effect for `fd` (see Section A.1.3.1).

6   On success, the function returns the number of bytes read by the calling thread. On error, it returns –1 and sets `errno` appropriately.

A.2.3   The upc_all_fwrite_list_local function

Synopsis

1
```
#include <upc.h>
#include <upc_io.h>
upc_off_t upc_all_fwrite_list_local(upc_file_t *fd,
    size_t memvec_entries, struct upc_local_memvec const *memvec,
    size_t filevec_entries, struct upc_filevec const *filevec,
    upc_flag_t flags);
```

Description

2 `upc_all_fwrite_list_local` writes data from local buffers in memory to a file. The file handle must be open for writing.

3 `memvec_entries` indicates the number of entries in the array `memvec` and `filevec_entries` indicates the number of entries in the array `filevec`. The values may be 0, in which case the `memvec` or `filevec` argument is ignored and no locations are specified for I/O.

4 The result is as if data were written from memory locations in the order specified by the list of locations in `memvec` to locations in the file in the order specified by the list in `filevec`. The total amount of data specified by `memvec` must equal the total amount of data specified by `filevec`.

5 If any of the writes from different threads result in overlapping accesses in the file, the result is determined by the current consistency and atomicity semantics mode in effect for `fd` (see Section A.1.3.1).

6 On success, the function returns the number of bytes written by the calling thread. On error, it returns –1 and sets `errno` appropriately.

A.2.4   The `upc_all_fwrite_list_shared` function

Synopsis

1
```
#include <upc.h>
#include <upc_io.h>
upc_off_t upc_all_fwrite_list_shared(upc_file_t *fd,
    size_t memvec_entries, struct upc_shared_memvec const *memvec,
    size_t filevec_entries, struct upc_filevec const *filevec,
    upc_flag_t flags);
```

Description

2 `upc_all_fwrite_list_shared` writes data from various locations of a shared buffer in memory to a file. The file handle must be open for writing.

3 `memvec_entries` indicates the number of entries in the array `memvec` and

filevec_entries indicates the number of entries in the array filevec. The values may be 0, in which case the memvec or filevec argument is ignored and no locations are specified for I/O.

4   The result is as if data were written from memory locations in the order specified by the list of locations in memvec to locations in the file in the order specified by the list in filevec. The total amount of data specified by memvec must equal the total amount of data specified by filevec.

5   If any of the writes from different threads result in overlapping accesses in the file, the result is determined by the current consistency and atomicity semantics mode in effect for fd (see Section A.1.3.1).

6   On success, the function returns the number of bytes written by the calling thread. On error, it returns –1 and sets errno appropriately.

A.3   UPC-IO HINTS

This section provides a list of predefined hints and their meanings. An implementation is not required to interpret these hint keys, but if it does interpret the hint key, it must provide the functionality described. All hints are single-valued character strings (the content is single-valued, not the location). Please refer to A.1.3.1 for syntax and usage of hints.

`access_style` (comma-separated list of strings): indicates the manner in which the file is expected to be accessed. The hint value is a comma-separated list of any the following: "read_once", "write_once", "read_mostly", "write_mostly", "sequential", and "random". Passing such a hint does not place any constraints on how the file may actually be accessed by the program, although accessing the file in a way that is different from the specified hint may result in lower performance.

`collective_buffering` (boolean): specifies whether the application may benefit from collective buffering optimizations. Allowed values for this key are "true" and "false". Collective buffering parameters can be further directed via additional hints: cb_buffer_size, and cb_nodes.

`cb_buffer_size` (decimal integer): specifies the total buffer space that the implementation can use on each thread for collective buffering.

`cb_nodes` (decimal integer): specifies the number of target threads or I/O nodes to be used for collective buffering.

`file_perm` (string): specifies the file permissions to use for file creation. The set of allowed values for this key is implementation defined.

`io_node_list` (comma separated list of strings): specifies the list of I/O devices that should be used to store the file and is only relevant when the file is created.

`nb_proc` (decimal integer): specifies the number of threads that will typically be used to run programs that access this file and is only relevant when the file is created.

`striping_factor` (decimal integer): specifies the number of I/O devices that the file should be striped across and is relevant only when the file is created.

start_io_device (decimal integer): specifies the number of the first I/O device from which to start striping the file and is relevant only when the file is created.

striping_unit (decimal integer): specifies the striping unit to be used for the file. The striping unit is the amount of consecutive data assigned to one I/O device before progressing to the next device, when striping across a number of devices. It is expressed in bytes. This hint is relevant only when the file is created.