

# Overview of UNH EXS 1.4.0 for Programmers

Robert Russell <rdr@iol.unh.edu>

Patrick MacArthur <pmacarth@iol.unh.edu>

University of New Hampshire InterOperability Laboratory

Durham, New Hampshire 03824

2016-12-07

## Contents

<b>1</b>	<b>Introduction</b>	<b>4</b>
<b>2</b>	<b>Overview of the EXS style of programming</b>	<b>5</b>
2.1	Asynchronous I/O . . . . .	5
2.1.1	Creating an event queue . . . . .	6
2.1.2	Deleting an event queue . . . . .	8
2.1.3	Using an event queue . . . . .	8
2.2	Memory registration . . . . .	9
<b>3</b>	<b>Establishing the EXS environment</b>	<b>9</b>
<b>4</b>	<b>Managing an EXS connection</b>	<b>10</b>
4.1	Creating an EXS socket . . . . .	10
4.2	Establishing an EXS client connection . . . . .	11
4.3	Establishing a server EXS connection . . . . .	11
4.3.1	Establishing a server listening post . . . . .	11
4.3.2	Accepting connections on the server . . . . .	12
4.4	Closing an EXS connection . . . . .	15
4.5	Half-closing an EXS connection . . . . .	16
<b>5</b>	<b>Basic data transfer over an EXS connection</b>	<b>17</b>
5.1	Sending data asynchronously . . . . .	17
5.2	Receiving data asynchronously . . . . .	18
5.3	Waiting for asynchronous I/O completion events . . . . .	18
5.4	Matching sends with receives . . . . .	19
5.4.1	Receiver's buffer is greater than or equal to amount of data in sender's packet . . . . .	19
5.4.2	Receiver's buffer is less than amount of data in sender's packet . . . . .	20

<b>6</b>	<b>Basic flow control within the EXS interface</b>	<b>20</b>
6.1	Send and Receive Credits . . . . .	20
6.2	Negotiations at connection establishment . . . . .	21
6.3	Matching advertisements and receives . . . . .	21
6.4	Number of internal buffers allocated . . . . .	22
<b>7</b>	<b>Tuning the UNH EXS interface</b>	<b>23</b>
7.1	Credit negotiation . . . . .	23
7.2	Small unregistered packets . . . . .	24
7.3	Using the hardware inline feature on <code>exs_send()</code> . . . . .	24
7.4	Pinning the EXS completion thread to a CPU . . . . .	25
7.5	Busy polling for completions . . . . .	26
7.6	Stream receive buffer . . . . .	27
<b>8</b>	<b>Registered and unregistered memory</b>	<b>27</b>
8.1	Explicit memory registration and deregistration . . . . .	28
8.2	Implicit memory registration and deregistration . . . . .	28
8.3	Memory alignment . . . . .	29
<b>9</b>	<b>Synchronous I/O</b>	<b>29</b>
9.1	Sending data synchronously . . . . .	29
9.2	Receiving data synchronously . . . . .	30
9.3	Establishing an EXS client connection synchronously . . . . .	30
9.4	Accepting connections on the server synchronously . . . . .	31
9.5	Closing EXS connections synchronously . . . . .	31
<b>10</b>	<b>Converting programs from using normal sockets to using EXS sockets</b>	<b>32</b>
10.1	Client using normal sockets . . . . .	32
10.2	Client using EXS sockets in synchronous mode with implicit memory registration	32
10.3	Client using EXS sockets in synchronous mode with explicit memory registration	33
10.4	Client using EXS sockets in asynchronous mode with implicit memory registration . . . . .	33
10.5	Client using EXS sockets in asynchronous mode with explicit memory registration . . . . .	34
<b>11</b>	<b>Status of UNH EXS 1.4.0</b>	<b>35</b>
11.1	Comparison with the Extended Sockets API (ES-API) Issue 1.0 Specification	35
11.2	Modifications to the ES-API standard in the UNH EXS implementation . . .	35
11.2.1	<code>exs.h</code> header file . . . . .	35
11.2.2	<code>exs_accept()</code> . . . . .	36
11.2.3	<code>exs_close()</code> . . . . .	36
11.2.4	<code>exs_connect()</code> . . . . .	36
11.2.5	<code>exs_mregister()</code> . . . . .	36
11.2.6	<code>exs_qcreate()</code> . . . . .	36
11.2.7	<code>exs_qmodify()</code> . . . . .	37

11.2.8	exs_qstatus()	37
11.2.9	exs_recv()	37
11.2.10	exs_send()	38
11.2.11	exs_sendfile()	38
11.3	Known deficiencies	38
11.3.1	thread cancellation	38
11.3.2	fork()	39


# 1 Introduction

The **Extended Sockets API** (ES-API) is a specification published by the **OpenGroup** that defines extensions to the traditional socket API in order to provide asynchronous I/O and also memory registration for Remote Direct Memory Access (RDMA). These two major new features enable programmers to take advantage of today’s multi-core processors and RDMA network hardware, such as InfiniBand, iWARP and RoCE interfaces, in a convenient yet efficient manner.

This document describes version 1.4.0 of the UNH EXS interface, an extended implementation of the ES-API. Section 11 outlines the differences between version 1.4.0 and the ES-API standard.

The UNH EXS interface provides most of the features specified in the ES-API, and provides a few additional features that give the programmer more flexibility. For example, the programmer can choose to program with synchronous rather than asynchronous I/O, and/or to program with or without memory registration. The programmer can also conveniently “tune” certain aspects of the EXS interface to take advantage of application requirements in order to provide better performance.

In addition, the UNH EXS interface is designed to be implemented entirely in user space on the Linux operating system. This latter feature provides easy porting, modification and adoption of EXS, since it requires no changes to existing Linux kernels. The current implementation of the EXS interface is based on **OpenFabrics Software** (OFS), a set of free software packages provided by the **OpenFabrics Alliance** (OFA). OFS is incorporated into many Linux distributions and is also released as the **OpenFabrics Enterprise Distribution** (OFED), which includes the entire set of user space OFS along with backports of the kernel space OFS to several enterprise Linux distributions. OFS runs in both user and kernel space, and provides efficient, asynchronous access to various types of RDMA networking hardware, currently InfiniBand, iWARP and RoCE. The layering structure of this system is shown below. The term **RNIC** stands for “RDMA Network Interface Card”, which is the term used by iWARP for the hardware interface card that enables a computer to use zero-copy transfers over Ethernet fibers and cables. **HCA** stands for “Host Channel Adapter”, which is the term used by InfiniBand for the hardware interface card that enables a computer to use zero-copy transfers over InfiniBand fabric. RoCE does not seem to have a special name for it’s interface card. We will use the term **Channel Adapter (CA)** to refer generically to any of the 3 hardware interface cards.



The principal difference between the UNH EXS interface and the OpenGroup’s ES-API is

that the UNH EXS interface runs entirely in user space. Therefore, it is not integrated with “normal” kernel sockets. This means EXS functions cannot be used with “normal” sockets, and EXS sockets cannot be used with “normal” socket functions. This, in turn, means the UNH EXS interface had to add numerous functions, such as `exs_socket()`, `exs_bind()`, etc., that the ES-API does not define because the ES-API expects to be implemented in the kernel and integrated with “normal” sockets. The ES-API is intended to be an extension to existing “normal” sockets and therefore does not need to redefine “normal” functions.

Another difference is that the UNH EXS interface is designed as a “thin”, efficient layer between the user and the Channel Adapter via the OFED stack. It therefore retains much of the packet orientation of the underlying protocols for both `SOCK_STREAM` and `SOCK_SEQPACKET` sockets.

This document gives an overview of the general concepts of EXS, and how the UNH EXS interface can generally be used by a programmer. It is not a reference manual. Section 2 starts by giving an overview of the style of programming for which the features offered by EXS are most suitable. Subsequent topics include EXS asynchronous I/O, dealing with event queues using the `exs_qcreate()`, `exs_qdelete()` and `exs_qdequeue()` functions, and the need for memory registration when using RDMA. Section 3 describes how a user program initializes the EXS environment with the `exs_init()` function. Section 4 describes how users create and manage EXS client and server connections by using the `exs_socket()`, `exs_bind()`, `exs_connect()`, `exs_listen()`, `exs_accept()`, and `exs_close()` functions. Section 5 describes how users transfer data over EXS connections by using the `exs_send()` and `exs_recv()` functions. Section 6 describes the internal flow control mechanisms utilized by the UNH EXS interface and their effects on performance. Section 7 describes some ways EXS performance can be tuned through the use of the `exs_fcntl()` function. Section 8 describes ways to utilize registered and unregistered memory by using the `exs_mregister()` and `exs_mderegister()` functions, and the `EXS_MHANDLE_UNREGISTERED` value. Section 9 describes ways to utilize synchronous and asynchronous I/O and the various flags that control this. Section 10 gives a series of sample client programs showing how to convert from a client using normal sockets (with no access to RDMA hardware) to a client using EXS sockets (with full access to RDMA hardware) in asynchronous mode and explicit memory registration. Section 11 describes the current status of UNH EXS 1.4.0 and the differences between it and the ES-API standard.

## 2 Overview of the EXS style of programming

The general paradigm for programming with EXS is “threads” and “events” programming. Although it is not necessary to use threads in order to use EXS, it is the simplest way to utilize the parallelism provided by multi-core processors and to take advantage of the asynchronous I/O facilities in EXS. Threads are also utilized internally by the UNH EXS interface.

### 2.1 Asynchronous I/O

EXS accomplishes asynchronous I/O by partitioning the user’s interaction with I/O via the EXS interface into two distinct phases: the “start” phase, and the “completion” phase.

Step	User thread	EXS interface
1	call asynchronous function	
2		check parameters
3		return error
4	deal with error	

Table 1: Start phase results in an error, I/O not started.

Each asynchronous operation starts with a start phase. If that phase fails, then there is no subsequent completion phase—instead, the user must deal with the error. But if the start phase succeeds, the asynchronous EXS operation proceeds in parallel with independent processing by the user thread that started it. The completion phase begins when the user calls a separate function that waits for the asynchronous operation to finish. This function returns detailed results about the asynchronous operation that can indicate either success or failure of the operation.

In EXS, the start phase of an asynchronous operation is accomplished by new EXS functions with names similar to the standard socket and UNIX I/O function calls. These include: `exs_accept()`, `exs_close()`, `exs_connect()`, `exs_recv()`, and `exs_send()`. These functions first verify that their user-supplied parameter values make such an operation possible. They then cause the EXS interface to “start” an operation but do not wait for the operation to actually take place. Instead they immediately return to the user after giving the OFED stack all the information necessary to proceed with the operation in parallel with (i.e., asynchronously to) the user’s threads.

As a parameter to these new asynchronous EXS functions the user must specify a pointer to a “queue” object previously created by the user (as discussed next in this section). When the asynchronous operation completes, either successfully or not, the EXS interface will “post” a completion event to that queue. This “event” is a structure which contains detailed information about the success or failure of the operation. The user obtains this information by “dequeuing” the event structure from the queue. If the EXS interface posts the event before the user tries to dequeue it, the queue object stores the information until the user performs the dequeue operation. If the user tries to dequeue an event before the asynchronous operation is finished, the dequeue operation will block until either the EXS interface posts the event or an amount of time specified by the user as a parameter to the dequeue operation elapses. This can be an indefinite amount of time. There are 3 possible scenarios, as illustrated in the following diagrams.

### 2.1.1 Creating an event queue

A user creates an event queue by using:

```
qhandle = exs_qcreate(depth);
```

where **depth** specifies the guaranteed minimum number of events that the user wants to be able to store in this queue. This function dynamically allocates memory to hold at least that number of event structures. A user can create numerous different event queues, but these are valid only in the context of the calling process (and its threads)—these event queues are

Step	User thread	EXS interface
1	call asynchronous function	
2		check parameters
3		start I/O operation
4		return ok
5	proceed with other work	proceed with I/O operation
6	proceed with other work	post I/O completion event
7	proceed with other work	
8	dequeue I/O completion event	
9	process I/O completion event	

Table 2: Start phase ok, user thread calls `exs_qdequeue` after I/O is complete.

Step	User thread	EXS interface
1	call asynchronous function	
2		check parameters
3		start I/O operation
4		return ok
5	proceed with other work	proceed with I/O operation
6	dequeue I/O completion event	proceed with I/O operation
7	wait for I/O completion event	proceed with I/O operation
8	wait for I/O completion event	post I/O completion event
9	process I/O completion event	

Table 3: Start phase ok, user thread calls `exs_qdequeue` before I/O is complete.

not valid in any child processes forked by this parent. The value returned in **qhandle** will be NULL if there was an error of some sort, in which case the error code is stored in the global **errno**. Otherwise, the value returned in **qhandle** is a “handle” (i.e., a pointer) of type `exs_qhandle_t`, which must be passed as a parameter in subsequent EXS function calls that start asynchronous operations. This handle identifies the queue to which an asynchronous operation will post its completion event.

### 2.1.2 Deleting an event queue

A user deletes an event queue by using:

```
result = exs_qdelete(qhandle);
```

where **qhandle** identifies a previously created event queue. The value returned in **result** will be 0 if the call to `exs_qdelete()` was successful, or -1 if there was an error of some sort, in which case the error code is stored in the global **errno**.

Note that any completion events posted to, but not yet dequeued from, this queue will be lost when this function is called. Also note that any outstanding asynchronous operations that reference this queue and that have been started but not yet completed will cause the call to `exs_qdelete()` to fail.

### 2.1.3 Using an event queue

Following a successful call to one or more EXS asynchronous functions, the user must use:

```
nevents = exs_qdequeue(qhandle, event_vector, count, timeout);
```

to wait for those EXS asynchronous functions to complete. In the call to `exs_qdequeue()`, **qhandle** must be the same value used previously in the EXS asynchronous function call that started an operation, **event\_vector** must be a user-provided array big enough to hold **count** events of type `exs_event_t`, and **timeout** is a pointer to a structure of type **struct timeval** containing the maximum amount of time the user wants to wait for an event to happen. The **timeout** parameter can be NULL if the user wants to wait indefinitely. This call to `exs_qdequeue()` will block until either the **timeout** elapses or at least one event is posted to the queue and is copied into **event\_vector**. The value returned in **nevents** will be 0 if a timeout occurred, a positive value if **nevents** events were removed from the queue and copied into **event\_vector**, or -1 if there was an error of some sort, in which case the error code is stored in the global **errno**.

If the **nevents** value returned by `exs_qdequeue()` is positive, then each of the first **nevents** structures of the **event\_vector** array will have been filled in by the EXS interface with information to identify the operation that posted the event and to convey back to the user the final results of that operation. For each of these structures in the **event\_vector** array, the following fields will be filled in as follows:

- `exs_evt_errno` is 0 if the asynchronous I/O operation completed successfully, otherwise it contains an error code (i.e., a Linux **errno** value) to indicate why it failed.
- `exs_evt_socket` is a copy of the **fd** parameter value used in the EXS asynchronous function to identify the connection.

- `exs_evt_ahandle` is a copy of the **ahandle** parameter value used in the EXS asynchronous function. This value is chosen by the user for identification purposes only, and is opaque to (i.e., not used by) the EXS interface.
- `exs_evt_type` is a constant value indicating which type of EXS asynchronous function caused the event. These values will be discussed below with the individual EXS asynchronous functions. Examples include `EXS_EVT_ACCEPT`, `EXS_EVT_CLOSE`, `EXS_EVT_CONNECT`, `EXS_EVT_RECV` and `EXS_EVT_SEND`.
- `exs_evt_union` contains a structure whose type depends on the `exs_evt_type` value.

The fields in these structures will be discussed below with the individual EXS asynchronous functions.

## 2.2 Memory registration

In order to utilize the direct memory-to-memory transfer feature of the RDMA interface hardware, the CA requires that the virtual memory involved on both ends of a transfer be “registered”. This registration accomplishes several necessary functions:

1. It establishes the location and size of a memory area to be utilized in RDMA transfers;
2. It establishes the read/write/execute permissions granted to both the local and remote CAs for accessing that memory area;
3. It “pins” the user’s virtual memory area into real (i.e., physical) memory so that a CA can access the memory without going through the CPU’s paging hardware.

Note that any type of virtual memory can be registered—stack memory, heap memory, global memory, etc. Common usage will most likely be to register dynamically allocated memory, but since this is not necessary, EXS does NOT allocate the memory it is registering.

The mechanism by which a user explicitly registers and deregisters memory is described in section 8 below. That section also describes a means by which the user can request dynamic registration and deregistration of memory as part of an individual EXS function call. A word of warning, however: dynamic registration and deregistration is an expensive procedure, and should be used only as a temporary “bridge” when converting an existing application from “normal” sockets to EXS. The benefits of EXS require explicit memory registration by the user program.

## 3 Establishing the EXS environment

Unlike with normal sockets, a program that wants to utilize EXS must explicitly initialize the EXS environment by using:

```
result = exs_init(version);
```

where **version** is the version of the EXS interface the user wishes to use. The initial (and currently only) version is given in the constant `EXS_VERSION1`, which is defined in the file `exs.h` which, in turn, must be included by all programs wishing to utilize EXS. The current version of the library is given by the global constant `EXS_VERSION`, and can be referenced by the application but should **not** be used as the argument to `exs_init()` since this defeats the ABI compatibility checking within `exs_init()`. The **result** value is 0 if the `exs_init()` was successful, or -1 if there was an error of some sort, in which case the error code is stored in the global **errno**.

This function will initialize the EXS interface for the calling process. It must be performed exactly once before any other EXS function is called, and is usually placed near the beginning of program execution.

## 4 Managing an EXS connection

As with classic sockets, there are three different types of connections a user can create:

1. a client connection
2. a server listening post
3. a server agent connection

### 4.1 Creating an EXS socket

In order to establish either a client connection or a server listening post, the user must first set up the local endpoint for the connection by using:

```
fd = exs_socket(domain, socket_type, protocol);
```

where the value of **domain** must be one of `PF_INET` or `PF_INET6` (or equivalently, `AF_INET` and `AF_INET6`), the value of **socket\_type** must be either `SOCK_STREAM` or `SOCK_SEQPACKET`, and **protocol** should be **0** (some other values are accepted for compatibility reasons). This is a synchronous operation, so there is no event associated with its completion. The value returned in **fd** will be -1 if there was an error of some sort, in which case the error code is stored in the global **errno**; otherwise, the value returned in **fd** is a non-negative integer **file descriptor** which must be used to identify this connection in all subsequent calls to EXS functions.

As of version 1.3.0, a large intermediate receive buffer will be allocated for sockets of type **SOCK\_STREAM** unless this feature is turned off via the interfaces described in section 7. The size of this buffer defaults to 3 Gigabytes, and may also be modified via the interfaces described in section 7.

Which EXS functions to call after a successful return from `exs_socket()` depends on whether the user wants to establish a client connection or a server listening post.

## 4.2 Establishing an EXS client connection

The client end of a connection is the simplest to establish, although this operation will not succeed unless the remote server's listening post has been previously established (discussed in section 4.3.1). Creation of a client connection is started by:

```
result = exs_connect(fd, server_address, server_addrlen, connect_flags,  
                    timeout, qhandle, ahandle);
```

This function is normally the next EXS function called by a client after a successful call to `exs_socket()`, and `fd` identifies the endpoint established by the `exs_socket()`. The `server_address` argument is a pointer to a structure of type `struct sockaddr` into which the user has stored the IPv4/IPv6 address and port number of the server's listening post. `server_addrlen` is the size in bytes of the structure pointed to by `server_address`. The value of the `connect_flags` parameter is usually 0, but see below for the other possibilities. `timeout` is a structure of type `struct timeval` which specifies the maximum amount of time the client thread is willing to wait for the `exs_connect()` to complete successfully. If the `timeout` pointer is NULL, `exs_connect()` will wait indefinitely. `qhandle` identifies an event queue that will be used to wait for the completion of this `exs_connect()`, and `ahandle` is an arbitrary pointer value chosen by the user which will be returned in the `exs_evt_ahandle` field of the event notification (discussed previously in section 2.1.3).

The value returned in `result` will be 0 if the `exs_connect()` was successfully started, or -1 if there was an error of some sort, in which case the error code is stored in the global `errno` and no asynchronous activity is started.

If the `exs_connect()` started successfully, it will operate asynchronously with the user thread that started it. During this time the user should not call any additional EXS functions for this connection, because the internal state of the connection will be undefined until the `exs_connect()` has completed and the EXS interface has posted to the user's `qhandle` an event whose `exs_evt_type` field contains the value `EXS_EVT_CONNECT`. Once this event has been received, the user is able to use the connection to transmit data to and from the remote server.

The only `connect_flags` value currently defined for `exs_connect()` is `EXS_BLOCK`. When this flag value is present, the value of the `qhandle` and `ahandle` parameters are ignored by the EXS interface, and can be NULL. The `EXS_BLOCK` flag value should be supplied when the user wants the `exs_connect()` to operate synchronously rather than asynchronously.

## 4.3 Establishing a server EXS connection

Establishing a server is a bit more complex than establishing a client. As with regular sockets it involves two distinct sequences.

### 4.3.1 Establishing a server listening post

A **listening post** is the socket endpoint set up by a server in order to allow clients to be able to contact the server. To enable such contact, a listening post must be bound to an IPv4/IPv6 address and port number pair that is known to clients. This IPv4/IPv6 address

and port number pair is analogous to a “1-800” number set up by a business—it must be known to customers who will use it to contact the business by telephone.

The first step in establishing a listening post is to create the socket endpoint using the `exs_socket()` function, as already discussed in section 4.1. The next step is to bind that socket endpoint to the IPv4/IPv6 address and port number pair which will be utilized by clients in the `exs_connect()` function, as just discussed in section 4.2. This binding is done by using:

```
result = exs_bind(fd, server_address, server_addrlen);
```

The `fd` parameter identifies the endpoint established by a previously successful call to `exs_socket()`. `server_address` is a pointer to a structure of type `struct sockaddr` into which the user has stored the IPv4/IPv6 address and port number to be assigned by the EXS interface to the server, and `server_addrlen` is the size in bytes of the structure pointed to by `server_address`. This is a synchronous operation, so there is no event associated with its completion. The value returned in `result` will be 0 if the call to `exs_bind()` was successful, or -1 if there was an error of some sort, in which case the error code is stored in the global `errno`.

After a successful call to `exs_bind()`, the next step in establishing a listening post is to identify the socket to the EXS interface as a listening post and to establish a **backlog** for it. This **backlog** is analogous to establishing the maximum number of calls which can be kept waiting on a 1-800 line, and is accomplished using:

```
result = exs_listen(fd, backlog);
```

`fd` has the same value as that used in the previous call to `exs_bind()`, and **backlog** is the maximum number of client connections that can be “kept” on hold until a server connection dedicated to a new client can be set up (i.e., until a customer can be switched to a free agent in the 1-800 call center). This is a synchronous operation, so there is no event associated with its completion. The value returned in `result` will be 0 if the `exs_listen()` was successful, or -1 if there was an error of some sort, in which case the error code is stored in the global `errno`.

At this point the server process is ready to somehow convey the IPv4/IPv6 address (or DNS name) and port number of its listening post to potential client processes all over the world. Well established servers, such as the World Wide **WHTTP** service, have been assigned “Well Known Ports” by IANA (the Internet Assigned Numbers Authority) (e.g., port 80 has been assigned to the HTTP service), so this knowledge is available to every browser in the world, and programmers only have to know the IPv4/IPv6 address (or DNS name) of their destination in order to connect their client process to an HTTP server. Most programmers do not have the luxury of working with Well Known Ports. Therefore, the means by which the server’s IPv4/IPv6 address and port number are made known to clients is outside the scope of EXS.

### 4.3.2 Accepting connections on the server

Once a server has established a listening post, it needs to set up to accept connections from clients. This is analogous to hiring agents to answer the phones in the 1-800 call center. When a business answers a call to the 1-800 number, the call is switched to a separate line

into a call center where a single agent deals exclusively with that individual customer—the 1-800 number remains ready to accept new calls and to switch them to other agents. Similarly, when the listening post accepts a connection from a client, it creates a new socket that will deal exclusively with that client—the listening post itself remains ready to accept new connections, but never actually transfers any data with any clients. The listening post is set up to do this using:

```
result = exs_accept(fd, address_vector, count, accept_flags, qhandle);
```

The **fd** parameter is the same as the one used in a previously successful call to `exs_listen()`, **address\_vector** is a pointer to an array of structures of type `exs_acceptaddr`, and **count** is the number of elements in that array. The value of the **accept\_flags** parameter will usually be 0; other flag values currently defined for `exs_accept()` are discussed below. **qhandle** identifies an event queue on which a user can wait for the completion of this `exs_accept()`.

The value returned in **result** will be 0 if the `exs_accept()` was successfully started, or -1 if there was an error of some sort, in which case the error code is stored in the global **errno** and no asynchronous activity is started.

The array of structures pointed to by **address\_vector** must have been allocated by the user before calling `exs_accept()`. Continuing the 1-800 analogy, there will be one element in this array for each agent available in the call center. Each element of the array contains the following three fields, which must be initialized by the user prior to calling `exs_accept()` as follows:

**exs\_addr** is a pointer to a structure of type `struct sockaddr` into which the IPv4/IPv6 address and port number of a new remote client will be stored by the EXS interface (this is how the agent is able to identify the customer who is calling). If the user does not wish to get this information, this pointer can be NULL or the value of the `exs_addrlen` parameter can be 0.

**exs\_addrlen** is the number of bytes allocated by the user to the structure pointed to by `exs_addr`. If that pointer is NULL, or if the user does not wish to get the remote client's IPv4/IPV6 address in the `exs_addr` parameter, the value of `exs_addrlen` can be 0.

**exs\_ahandle** is an arbitrary pointer value chosen by the user for identification purposes only, and is never looked at by the EXS interface. It will be returned in the `exs_evt_ahandle` field of the event notification (as previously discussed in section 2.1.3).

Following the call to `exs_accept()`, the user must use:

```
nevents = exs_qdequeue(qhandle, event_vector, count, timeout);
```

to wait for clients to perform `exs_connect()` operations to this server's listening post. The `exs_qdequeue()` function has already been explained in section 2.1.3. To repeat from that section, **qhandle** must point to the same event queue as that used in the call to `exs_accept()`, **event\_vector** must be a user provided array big enough to hold **count** events of type `exs_event_t`, and **timeout** is a pointer to a structure of type `struct timeval` containing the maximum amount of time the user wants to wait for an event to happen. The **timeout** parameter can be NULL if the user wants to wait indefinitely.

If the **nevents** value returned by `exs_qdequeue()` is positive, then each of the first **nevents** structures of the **event\_vector** array were filled in by the EXS interface with information to identify each operation that posted an event to this event queue and to convey the final results of each operation back to the user. For an event associated with an `exs_accept()`, the following fields of the `exs_event_t` structure will be filled in (as generically described in section 2.1.3):

**exs\_evt\_errno** is 0 if the `exs_accept()` operation completed successfully, otherwise it contains an error code (i.e., a Linux **errno** value) to indicate why it failed.

**exs\_evt\_socket** is a copy of the listening post's **fd** value that was used as a parameter in the call to `exs_accept()`,

**exs\_evt\_ahandle** is a copy of the **ahandle** value stored by the user in the `exs_ahandle` field of an element in the array pointed to by the **address\_vector** parameter to the `exs_accept()`. This value is chosen by the user for identification purposes only, and is not used by the EXS interface.

**exs\_evt\_type** is `EXS_EVT_ACCEPT`.

**exs\_evt\_union** contains a structure of type `exs_evt_accept` that has been filled in by the EXS interface with values in the following fields:

**exs\_evt\_new\_socket** is the file descriptor identifying the new connection to the client—this value should be used by the server from this point on to transmit data to/from that client. This is analogous in the 1-800 call center to the line to an exclusive agent to which the customer is switched. However, the analogy is not perfect because the call center line must already exist (as extension 123, for example), whereas this file descriptor represents a new socket created internally by the `exs_accept()` function. It is as if the `exs_accept()` internally calls `exs_socket()` for each new client, and the socket created by this internal call to `exs_socket()` has a file descriptor that is different from any existing file descriptors. Note that this scheme is not new to EXS—the **accept()** function for normal sockets works in exactly the same manner.

**exs\_evt\_addr** is a copy of the `exs_addr` pointer from an element in the array pointed to by the **address\_vector** parameter to the `exs_accept()`. The structure pointed to by the `exs_evt_addr` field is of type **struct sockaddr**, and will now contain the IPv4/IPv6 address and port number of the remote client at the other end of the new connection.

**exs\_evt\_addrlen** contains the number of bytes used to store the client's IPv4/IPv6 address and port number in the structure pointed to by `exs_evt_addr`.

Once a new connection to a client has been indicated by a successful completion event, it is common for a server process to spawn a new “agent” thread to deal exclusively with that client. This is analogous to switching a customer's call to an agent in the 1-800 call center. The file descriptor from the `exs_evt_new_socket` field in the `exs_evt_accept` structure

should be used by this agent thread in all the `exs_send()` and `exs_recv()` functions for transactions with that client. When the agent thread's dealings with this client are finished, that thread should call `exs_close()` with this **fd** and then terminate. Just the agent thread needs to terminate, not the server listening post thread. Here the 1-800 analogy breaks down, because human agents will just hang up with one customer and wait for a call from another customer. With EXS, as with normal sockets, there is no way to reuse a socket—a new one must be created by `exs_accept()` as each new client connects.

The only **accept\_flags** value currently defined for `exs_accept()` is `EXS_BLOCK`. When this flag value is present, the value of the **count** parameter MUST be exactly 1, since the **fd** for only 1 remote connection can be returned by 1 call. The value of the **qhandle** parameter is ignored by the EXS interface, and can be `NULL`. The `EXS_BLOCK` flag value should be supplied when the user wants the `exs_accept()` to operate synchronously rather than asynchronously. A synchronous `exs_accept()` blocks until a remote client connects, at which time the **result** returned by `exs_accept()` will be the **fd** for the new connection to the remote client. Note that in this case no `exs_event_t` structure is generated for the user, so the other fields in that structure are not available to the user. This means that the `exs_ahandle` field in the first (and only) **address\_vector** is also ignored by the EXS interface, and can be `NULL`.

## 4.4 Closing an EXS connection

When the user has finished using a connection of any of the three types discussed above, the user calls the following function:

```
result = exs_close(fd, close_flags, qhandle, ahandle);
```

where **fd** identifies the connection to be closed, **flags** contains flags that modify the normal behavior of this call, **qhandle** identifies an event queue that will be used to wait for the completion of this call to `exs_close()`, and **ahandle** is an arbitrary pointer value chosen by the user which will be returned in the `exs_evt_ahandle` field of the event notification (discussed previously in section 2.1.3).

The value returned in **result** will be 0 if `exs_close()` was successfully started, or -1 if there was an error of some sort, in which case the error code is stored in the global **errno** and no asynchronous activity is started.

If the `exs_close()` started successfully, it will operate asynchronously to the user thread that called it. During this time the user must not call any additional EXS functions using this **fd** to identify the connection, because the **fd** will become invalid and the internal state of the connection will become undefined once `exs_close()` is called (the **fd** is invalidated synchronously, even if the subsequent cleanup happens asynchronously). When the `exs_close()` completes, the EXS interface will post to the user's **qhandle** an event whose `exs_evt_type` field contains the value `EXS_EVT_CLOSE`. The `exs_evt_errno`, `exs_evt_socket`, and `exs_evt_ahandle` fields of the event will be filled in as previously described in section 2.1.3. No additional information is stored in the `exs_evt_union` field for events of this type.

It is highly recommended that, prior to exiting a program or thread, a user always call `exs_close()` for each connection controlled by that program or thread. The main reason for this is that EXS is inherently asynchronous, so many EXS functions called by a user

simply start an asynchronous operation—the real work takes place asynchronously to the thread that called the EXS function. Therefore, when writing code, a user might not be aware of all that is being accomplished asynchronously, so that what looks to be finished to a user might not be finished in the EXS interface. Calling the `exs_close()` function ensures that this asynchronous activity is finished, whereas exiting the program or thread would not ensure this, and thus all data might not be transmitted. Calling `exs_close()` also helps to cleanly shutdown the other end of the connection. Finally, calling `exs_close()` allows the **fd** associated with the connection to be reused.

The `exs_close()` operation supports two flags. When the `EXS_BLOCK` flag value is present, the value of the **qhandle** and **ahandle** parameters are ignored by the EXS interface, and can be `NULL`. The `EXS_BLOCK` flag value should be supplied when the user wants the `exs_close()` to operate synchronously rather than asynchronously. When the `EXS_DONTLINGER` flag value is present, the `exs_close()` function will immediately terminate the connection instead of first waiting for pending sends to terminate. This will also cause any outstanding and future send and receive operations at the remote endpoint to immediately fail, setting `exs_evt_errno` or **errno** (as appropriate) to **ECONNRESET**.

## 4.5 Half-closing an EXS connection

In some situations, it is useful to shut down only half of a connection. In particular, for some request/response protocols, the client may send a single request and not need to send any further messages. In this case, shutting the connection down only for writing will indicate EOF to the receiver but still allow the receiver to send its response on the socket.

When the user wishes to half-close a connection, the user calls the following function:

```
result = exs_shutdown(fd, how, shutdown_flags, qhandle, ahandle);
```

where **fd** identifies the connection to be closed, **how** identifies in which direction to shut down communication, **shutdown\_flags** contains flags that modify the normal behavior of this call, **qhandle** identifies an event queue that will be used to wait for the completion of this call to `exs_shutdown()`, and **ahandle** is an arbitrary pointer value chosen by the user which will be returned in the `exs_evt_ahandle` field of the event notification (discussed previously in section 2.1.3).

The value returned in **result** will be 0 if `exs_shutdown()` was successfully started, or -1 if there was an error of some sort, in which case the error code is stored in the global **errno** and no asynchronous activity is started. When the `exs_shutdown()` completes, the EXS interface will post to the user's **qhandle** an event whose `exs_evt_type` field contains the value `EXS_EVT_SHUTDOWN`. The `exs_evt_errno`, `exs_evt_socket`, and `exs_evt_ahandle` fields of the event will be filled in as previously described in section 2.1.3. No additional information is currently stored in the `exs_evt_union` field for events of this type.

The valid values for **how** are: **SHUT\_WR**, which shuts down the connection for sending (writing) only; **SHUT\_RD**, which shuts down the connection for receiving (reading) only; and **SHUT\_RDWR**, which shuts down the connection in both directions.

The `exs_shutdown()` function will fail (setting **errno** to **EBUSY**) if an `exs_shutdown()` operation is outstanding in the same direction (locally or remotely); this is to avoid posting an event too early or needing complex internal logic to ensure that we post the correct

number of completion events in order. If an `exs_shutdown()` operation in a given direction has already completed, future `exs_shutdown()` operations in that direction are treated as a no-op and will immediately generate an asynchronous event on the supplied event queue.

Unlike `exs_close()`, `exs_shutdown()` does *not* invalidate the `fd` of the connection, even if the call results in communication being shut down in both directions. This is because it is generally expected that communication will continue in the opposite direction of the direction that communication is being shut down. This capability can also be exploited to ensure that there is no activity on a socket prior to invalidating its `fd` by shutting down communication in both directions, which may be useful for a program which opens many sockets and tracks them in some data structure. However, the user is still responsible for eventually calling `exs_close()` (discussed in section 4.4) on the socket to invalidate its `fd` (freeing it for reuse) and completely releasing the resources used by the connection.

The only `shutdown_flags` value currently defined for `exs_shutdown()` is `EXS_BLOCK`. When this flag value is present, the value of the `qhandle` and `ahandle` parameters are ignored by the EXS interface, and can be `NULL`. The `EXS_BLOCK` flag value should be supplied when the user wants the `exs_shutdown()` to operate synchronously rather than asynchronously.

## 5 Basic data transfer over an EXS connection

There are two basic, complementary operations that transfer data over an EXS connection: `exs_send()` and `exs_recv()`. These two operations are used regardless of how the connection was established, since both clients and servers need to be able to both send and receive data. Both of these functions are asynchronous, so they only start the data transfer—the user must explicitly call `exs_qdequeue()` to know when the transfer finishes.

### 5.1 Sending data asynchronously

A user sends data asynchronously by using:

```
result = exs_send(fd, send_buffer, send_length, send_flags,
                 qhandle, ahandle, mhandle);
```

where `fd` identifies the connection, and the user has filled `send_buffer` with `send_length` bytes of data prior to calling `exs_send()`. This `send_buffer` must be completely within an area of memory that was assigned the memory registration key `mhandle` in a previous call to `exs_mregister()`. `send_flags` will usually have the value 0; other flags currently defined for `exs_send()` are discussed in sections 9 and 10 below. `qhandle` identifies an event queue previously created by the user to wait for completion events. `ahandle` is an arbitrary pointer value chosen by the user which will be returned in the `exs_evt_ahandle` field of the completion event notification (discussed previously in section 2.1.3).

If `exs_send()` returns 0, then the operation has started and an asynchronous event will be posted on the user's event queue once the operation has completed. If the connection has been gracefully shut down for writing, `exs_send()` will return -1 and set `errno` to `EPIPE`, and no asynchronous completion event will be posted to the event queue. If the connection has been abruptly terminated, `exs_send()` will return -1 and set `errno` to `ECONNRESET`,

and no asynchronous completion event will be posted to the event queue. If the socket was never connected before calling `exs_send()`, `exs_send()` will return -1 and set `errno` to **ENOTCONN**, and no asynchronous completion event will be posted to the event queue.

## 5.2 Receiving data asynchronously

A user receives data by using:

```
result = exs_recv(fd, recv_buffer, max_length, recv_flags,
                 qhandle, ahandle, mhandle);
```

where **fd** identifies the connection, and the user has reserved a **recv\_buffer** capable of holding **max\_length** bytes of data. This **recv\_buffer** must be completely within an area of memory that was assigned the memory registration key **mhandle** in a previous call to `exs_mregister()`. **recv\_flags** will usually have the value 0, unless other flags are desired. As of version 1.3.0, for **SOCK\_STREAM** sockets, the **MSG\_WAITALL** flag will make the `exs_recv()` operation wait until it can fill all **max\_length** bytes; otherwise, the operation will complete as soon as any data is received. For **SOCK\_SEQPACKET** sockets, the **MSG\_WAITALL** flag has no effect. Other effects of this flag and additional flags currently defined for `exs_recv()` are discussed in sections 9 and 10 below. **qhandle** identifies an event queue previously created by the user to wait for completion events. **ahandle** is an arbitrary pointer value chosen by the user which will be returned in the `exs_evt_ahandle` field of the completion event notification (discussed previously in section 2.1.3).

If `exs_recv()` returns 0, then the operation has started and an asynchronous event will be posted on the user's event queue once the operation has completed. If the connection has been gracefully shut down for reading (i.e., via the `exs_shutdown()` or `exs_close()` functions on the local or remote side of the connection), `exs_recv()` will complete asynchronously with both an `exs_evt_errno` and `exs_evt_length` of 0 (this is done because a return value of 0 from `exs_recv()` would otherwise be ambiguous, since it is used to mean that an asynchronous event will follow). If the connection was terminated abruptly, `exs_recv()` will return -1 and set `errno` to **ECONNRESET**, and no asynchronous completion event will be posted to the event queue. If the socket was never connected before calling `exs_recv()`, `exs_recv()` will return -1 and set `errno` to **ENOTCONN**, and no asynchronous completion event will be posted to the event queue.

## 5.3 Waiting for asynchronous I/O completion events

The value returned by `exs_send()` or `exs_recv()` in **result** will be 0 if the operation was successfully started, or -1 if there was an error of some sort, in which case the error code is stored in the global **errno** and no asynchronous activity is started.

Once an `exs_send()` or `exs_recv()` is successfully started, it will operate asynchronously to the user program. During this time the user must not in any way modify the area of memory pointed to by the buffer parameter, because the data in that buffer has still not been transferred across the connection. Eventually the user must use:

```
nevents = exs_qdequeue(qhandle, event_vector, count, timeout);
```

to wait for the operation to finish so that the buffer can be safely accessed again. The `exs_qdequeue()` call was explained in section 2.1.3. `qhandle` must be the value used in the `exs_send()` or `exs_recv()`, `event_vector` must be an array big enough to hold `count` events of type `exs_event_t`, and `timeout` is the maximum amount of time the user wants to wait for an event to happen (NULL for an indefinite wait).

If the value returned in `nevents` is positive, then each of the first `nevents` structures of the `event_vector` array were filled in by the EXS interface with information to identify the operation that caused this event and to convey the final results of that operation back to the user. For each of these structures, the fields `exs_evt_errno`, `exs_evt_socket`, and `exs_evt_ahandle` will be filled in by the EXS interface with the values explained previously (in section 2.1.3). In addition:

`exs_evt_type` will indicate which type of operation caused the event: `EXS_EVT_SEND` if the operation was `exs_send()`, or `EXS_EVT_RECV` if the operation was `exs_recv()`.

`exs_evt_union` will contain a structure of type `exs_evt_xfer`

- `exs_evt_buffer` is the buffer address specified in the `exs_send()` or `exs_recv()`
- `exs_evt_mhandle` is the `mhandle` specified in the `exs_send()` or `exs_recv()`
- `exs_evt_amount_lost` is described below in section 5.4
- `exs_evt_length` is the number of bytes successfully transferred by this operation.

When the value of `exs_evt_type` is `EXS_EVT_SEND` and the operation was successful, the value of `exs_evt_length` will always be equal to the value of the `send_length` parameter in the original `exs_send()`. When the value of `exs_evt_type` is `EXS_EVT_RECV` and the operation was successful, the value of `exs_evt_length` will always be less than or equal to the value of the `max_length` parameter in the original `exs_recv()`.

## 5.4 Matching sends with receives

Since the amount of data sent in one packet by an `exs_send()` may differ from the amount of data requested in the remote side's corresponding `exs_recv()`, the EXS interface layer must rationalize any difference. For efficiency, the way it does this is very packet-oriented, and hence differs somewhat from the way this is done in traditional sockets.

### 5.4.1 Receiver's buffer is greater than or equal to amount of data in sender's packet

If the receiver provides a buffer whose size is greater than or equal to the size of the data in the matching send packet, there is no problem—the EXS interface delivers all the sent data into the beginning of the receiver's buffer, any remaining space at the end of the receiver's buffer is left undefined, the value returned in `exs_evt_length` is the exact number of bytes delivered into the buffer, and `exs_evt_amount_lost` is always zero.

### 5.4.2 Receiver’s buffer is less than amount of data in sender’s packet

If the receiver provides a buffer whose size is less than the size of the data in the matching send packet, the resolution depends on the type of socket in use. For all socket types, the EXS interface completely fills the receiver’s buffer with the first part of the sent data, and returns in `exs_evt_length` the exact number of bytes delivered into the buffer.

If the socket type is **SOCK\_STREAM**, the EXS interface ignores packet boundaries, so it saves the remainder of the data that is in the matching send packet and uses it to match with subsequent `exs_recv()` calls. No data is ever lost, so the value returned to the receiver in `exs_evt_amount_lost` is always zero.

If the socket type is **SOCK\_SEQPACKET**, the EXS interface maintains packet boundaries, so it discards the remainder of the data in the matching send packet, and returns to the receiver in `exs_evt_amount_lost` the number of bytes discarded.

## 6 Basic flow control within the EXS interface

Because data sent by `exs_send()` operations on one end of a connection must be delivered into buffers specified by `exs_recv()` operations on the other end, the EXS interfaces on both ends must coordinate the flow of this data so that neither side runs out of buffers. This is done in a manner that is largely transparent to the user.

Since a user must allocate and fill a memory buffer before specifying it as the **send\_buffer** parameter to `exs_send()`, the EXS interface does not have to provide any additional storage for data on the sending side. However, the interface cannot actually send data until it knows for sure that the user on the receiving side has provided a corresponding buffer into which the sender’s data can be delivered without any additional copying or buffering (because EXS uses direct memory to memory transfers). The EXS interface uses a credit mechanism to accomplish this.

### 6.1 Send and Receive Credits

The interface on each end of an established connection internally maintains two dynamically varying local credit values. At any time “`send_credits`” is the maximum number of packets this interface is allowed to start sending to the other side using `exs_send()`; and “`recv_credits`” is the maximum number of packets this interface is allowed to start receiving from the other side using `exs_recv()`. At any time the value of **send\_credits** on one side must equal the value of **recv\_credits** on the other side, and vice versa. As explained in the next section (6.2), when an EXS connection is first established, the EXS interfaces on both sides negotiate these numbers so that they are initialized to the same values.

Each time a user issues an `exs_recv()`, the receiving EXS interface reduces its local **recv\_credits** by one. If the balance would become negative, the receiving interface sets **errno** to **EBUSY** and returns -1 (the caller is expected to then wait for previous `exs_recv()` operations to complete by calling `exs_qdequeue()`). If the caller wishes to wait until a credit is available, the `EXS_CREDIT_WAIT` flag can be provided, but then it is up to the caller to ensure that this will not cause a deadlock. If the balance would not become negative, the receiving interface sends an “advertisement” to the sending side and then returns a value of

0 to the caller of the `exs_rcv()` to indicate that the `exs_rcv()` has started successfully. This advertisement contains no data, but rather information (or “metadata”) describing the memory on the receiver that is now ready to receive data from the sender (i.e., its length, location, and memory registration key).

The sending interface must keep track locally of the `send_credits` it has negotiated with the receiver. Each time a sending user issues an `exs_send()`, the sending EXS interface reduces its local `send_credits` by one. If the balance would become negative, the sending interface sets `errno` to **EBUSY** and returns -1 (the caller is expected to then wait for previous `exs_send()` operations to complete by calling `exs_qdequeue()`). If the caller wishes to wait until a credit is available, the `EXS_CREDIT_WAIT` flag can be provided, but then it is up to the caller to ensure that this will not cause a deadlock. If the balance would not become negative, the sending interface adds the information (i.e., “metadata”) from this `exs_send()` to an internal queue and returns a value of 0 to the caller of the `exs_send()` to indicate that the `exs_send()` has been started successfully.

## 6.2 Negotiations at connection establishment

At the time an EXS connection is first established, the client side automatically sends the server a short “setup request” message that contains the client’s EXS version number (currently 1), and the initial values of the client’s **send\_credits**, and **rcv\_credits**. Upon receiving this “setup request” message, the server side of the connection compares the values contained in this message with its own corresponding values. The **minimum** of each corresponding value is used to reset the server’s own value and to build a “setup response” message that it sends back to the client. Once the client receives this “setup response” message, it uses those values to set its own corresponding values. Consequently, from that point in time on, both ends of the newly established connection have the same value for each of the corresponding parameters, and the flow control mechanism will now function properly. These negotiations are transparent to the user.

Following these negotiations, the EXS interface on each side of a newly established connection sets up one internal receive buffer for each local **send\_credit** that it has negotiated with the corresponding receiving side. These buffers are used to store EXS advertisements as they arrive from the remote end. When the receiving interface sends an advertisement to the sender, the advertisement is delivered directly into one of these interface buffers and the sending interface is notified of this arrival by the OFED stack. Since the receiving interface should never send more advertisements than it has local **rcv\_credits**, and since the sending interface has posted one receive buffer for each of its local **send\_credits**, no advertisements should ever be lost for lack of a buffer when they arrive on the sending side.

## 6.3 Matching advertisements and receives

Each EXS interface keeps track of advertisements it has received and `exs_send()` operations that its sending user has started. Whenever the user starts an `exs_send()` operation, the EXS interface looks for an already received advertisement to match it with. Similarly, whenever an advertisement is received from the remote receiver, the receiving interface looks for

an already started `exs_send()` to match it with. Matching occurs in the manner already described in section 5.4. The sending interface then issues an `RDMA_WRITE_WITH_IMM` operation to its local CA (via the OFED stack) in order to actually transfer data into the user's buffer on the receiving side (as indicated in the advertisement) directly from the user's buffer on the sending side (as indicated in the `exs_send()`) without any extra copying or CPU intervention on either side. This `RDMA_WRITE_WITH_IMM` operation is transparent to the user on both sides of a connection.

When the receiving interface is notified by its receiving CA (via the OFED stack) that a remotely issued `RDMA_WRITE_WITH_IMM` operation has completed, it will do two things.

1. It uses the “immediate” value supplied in the local `RDMA_WRITE_WITH_IMM` completion to locate the corresponding advertisement previously issued by that receiver and posts the receiver's completion event to tell the receiving user that its `exs_recv()` has completed, and to convey the results to the receiving user via the `exs_event_t` structure.
2. It increments its local **recv\_credits**, and a new `exs_recv()` operation may commence using this newly available credit.

When the sending interface is notified by its sending CA (via the OFED stack) that the locally issued `RDMA_WRITE_WITH_IMM` operation has completed, it also does two things.

1. It posts the sender's completion event to tell the sending user that its `exs_send()` has completed and to convey the results to the sending user via the `exs_event_t` structure.
2. It increments its local **send\_credits** for this connection, and a new `exs_send()` operation may commence using this newly available credit.

Clearly this credit mechanism requires each EXS interface to allocate some hidden, internal buffers for exchanging advertisements. However, these buffers and advertisements are small, since they are used only to send and receive a limited amount of control information, not an unlimited amount of user data. Indeed, these buffers are small enough to be embedded in internal control blocks that contain additional information needed locally by the EXS interface.

## 6.4 Number of internal buffers allocated

Using the mechanism discussed in the next section, a user can set the local values to use when negotiating the initial **send\_credits** and **recv\_credits** values. Therefore, it is important for the user to understand how the EXS interface uses these values to allocate its internal buffers. The receiving interface needs two control blocks for each local **recv\_credit** it has initially negotiated with the sending interface: one for sending an advertisement, and one for receiving the corresponding remotely issued `RDMA_WRITE_WITH_IMM` completion. The sending interface also needs two control blocks for each local **send\_credit** that it has negotiated with the receiver: one for keeping track of waiting `exs_send()` operations that have

been started but for which no advertisement has been received yet, and one for receiving an advertisement. However, one extra control block is allocated for each local **send\_credit** to ensure that receive buffers are always posted regardless of the inevitable time delays between sending or receiving a message and being notified of the completion of that operation (at which time a control block can be reused).

## 7 Tuning the UNH EXS interface

One of the differences between the UNH EXS interface and the ES-API is an additional mechanism by which a user can “tune” some aspects of the EXS interface in order to increase performance for a particular application. The user does this through the use of:

```
result = exs_fcntl(fd, command, argument);
```

The value of the **fd** parameter identifies the socket to be tuned, the value of the **command** parameter indicates what the user wants to do to the socket, and the value of the **argument** parameter depends on the **command**. This is a synchronous operation, so there is no event associated with its completion. The value returned in **result** will also depend on the command, although for all commands it will be -1 if there was an error of some sort, in which case the error code is stored in the global **errno**.

### 7.1 Credit negotiation

As previously discussed in section 6, the credit value used to control the flow of data in EXS is negotiated at the time a new connection is established. The value used in this negotiation can be set by the user through the use of the `exs_fcntl()` function prior to the `exs_connect()` and `exs_accept()` calls on the **fd** parameter. When the value of the **command** parameter to the `exs_fcntl()` function is `EXS_F_SETFLOWCONTROLCREDITS`, the EXS interface will use the value of the **argument** parameter as the local value to be used in the negotiation of the local **send\_credits** and **recv\_credits**. The default value is 32. In previous versions, 32 was also the maximum acceptable value, but this maximum has been removed as of version 1.3.0. Negotiation occurs as part of connection establishment, and the result is the minimum value supplied by either side. This **command** parameter value cannot be used in an `exs_fcntl()` call on an established connection. The result returned by a successful `exs_fcntl()` call is the old value of the corresponding local credit value.

When the **command** parameter to `exs_fcntl()` is `EXS_F_GETFLOWCONTROLCREDITS`, the result returned by the `exs_fcntl()` depends on whether or not the connection has been established. If `exs_fcntl()` is called with this **command** parameter value prior to a successful call to `exs_connect()` or `exs_accept()` on the **fd** parameter, the value returned will be the local value used in a future negotiation on that connection. If it is called after a connection was successfully established, the value returned will be the local value that resulted from the negotiation at the time the connection was established. Note that this return value is not the dynamically varying credit value used to control the flow on the connection, but the negotiated limit on that credit value.

## 7.2 Small unregistered packets

Because of the overhead involved in dynamically registering and unregistering memory (see section 8), it may be faster to send and receive small amounts of data by having the EXS library simply copy the data into/out of preregistered library buffers. Thus, in `SOCK_SEQPACKET` mode, UNH EXS supports copying unregistered small packets. The definition of a “small” packet can be controlled through the use of the `EXS_F_SETSPMAXSIZE` value of the **command** parameter to the `exs_fcntl()` function, in which case the value of the **argument** parameter should be the value to be used in the negotiation of the definition of “small”. The default value is 0. Negotiation occurs as part of connection establishment, and the result is the minimum value supplied by either side. This **command** parameter value cannot be used in an `exs_fcntl()` call on an established connection. The result returned by a successful `exs_fcntl()` call is the old value of the corresponding local small packet max size.

Setting a positive value for `EXS_F_SETSPMAXSIZE` effects the operation of `exs_send()` and `exs_recv()` called with a value of `EXS_MHANDLE_UNREGISTERED` for the required **mhandle** parameter. If the value of the required **length** parameter is less than or equal to the small packet max size, then rather than dynamically registering and unregistering the data supplied in the function call, the data is copied into a preregistered library buffer on `exs_send()` or copied out of a preregistered library buffer on `exs_recv()`.

When the **command** parameter to the `exs_fcntl()` function is `EXS_F_GETSPMAXSIZE`, the result returned by the `exs_fcntl()` depends on whether or not the connection has been established. If `exs_fcntl()` is called with this **command** parameter value prior to a successful call to `exs_connect()` or `exs_accept()` on the **fd** parameter, the value returned will be the local value used in a future negotiation on that connection. If it is called after a connection was successfully established, the value returned will be the local value that resulted from the negotiation at the time the connection was established.

In **SOCK\_STREAM** mode, this is not supported since the stream mode does its own buffering by default.

## 7.3 Using the hardware inline feature on `exs_send()`

Many RDMA interface cards support the optional “inline” feature, which allows the card to enqueue a copy of small amounts of data as part of the metadata it enqueues on the send queue at the time it starts an `exs_send()` operation. The effect of this is to make the transfer slightly faster (i.e., to exhibit lower latency), because the data is already on the interface card at the time it actually moves data onto the wire, so there is no need to involve the memory bus in the transfer itself (the memory bus was involved at the time the `exs_send()` was enqueued).

The definition of “small amounts” of data can be controlled to some extent through the use of the `EXS_F_SETINLINEMAXSIZE` value of the **command** parameter to the `exs_fcntl()` function, in which case the value of the **argument** parameter should be the definition of “small amounts”. The default value is the largest value acceptable to the interface card. Negotiation occurs as part of connection establishment, and is explained next. This **command** parameter value cannot be used in an `exs_fcntl()` call on an established connection. The

result returned by a successful `exs_fcntl()` call is the old value of the corresponding local inline max size.

The user does not have complete control over this value, because this is an optional feature of RDMA interface cards, and the maximum size possible depends on the particular card used. Therefore, if the user uses the `EXS_F_SETINLINEMAXSIZE` value in the **command** parameter to the `exs_fcntl()` function, then the user's value of the **argument** parameter is only a starting point for the EXS library to negotiate with the local interface card to set the maximum inline size to use. If the user's value is acceptable to the local interface card, then that is the value used. If the user's value is not acceptable, then the EXS library will silently find and use the largest value acceptable to the interface card but smaller than the user's value.

Note that the scope of the maximum inline size is the local interface card—it is not negotiated with the remote side when a connection is established by an `exs_connect()` or `exs_accept()`, because the interface card on the remote side could be different from the local interface card, and might support a different maximum value. In any case, maximum inline size is only utilized by the EXS library on an `exs_send()`, since it only effects the queuing performed locally when this function is called—it has no effect on an `exs_recv()` or on data transferred on the wire.

Note also that the maximum inline size is independent of the small packet max size. The maximum inline size is a hardware option that improves the latency of `exs_send()` operations, regardless of whether or not the **mhandle** parameter is given as `EXS_MHANDLE_UNREGISTERED` in the `exs_send()`. Therefore the EXS library will use it by default on all calls to `exs_send()` in which the value of the **send\_length** parameter is less than or equal to the maximum inline size. (Obviously if the user sets this maximum to 0, this feature will not be used.) The small packet max size is a software option that improves the latency of `exs_send()`, but only when the **mhandle** parameter is given as `EXS_MHANDLE_UNREGISTERED`. The EXS library uses will use it only on these calls, and only if the value of the **send\_length** parameter is less than or equal to the small packet max size.

When the **command** parameter to the `exs_fcntl()` function is `EXS_F_GETINLINEMAXSIZE`, the result returned by the `exs_fcntl()` depends on whether or not the connection has been established. If `exs_fcntl()` is called with this **command** parameter value prior to a successful call to `exs_connect()` or `exs_accept()` on the **fd** parameter, the value returned will be the local value used in a future negotiation with the local interface when a connection is established. If it is called after a connection was successfully established, the value returned will be the local value that resulted from the negotiation at the time the connection was established.

## 7.4 Pinning the EXS completion thread to a CPU

Because of the asynchronous nature of its operation, the EXS library utilizes a “completion thread”. One of the fine-tuning options available to users is the ability to pin this completion thread to a particular CPU. If the user pins his/her threads to different CPUs, the execution load will be distributed across a multi-core platform, which should give better performance. If thread pinning is not performed, then the completion thread can be dynamically assigned to various CPUs by the kernel scheduler.

Completion thread pinning is controlled through the use of the `EXS_F_SETCOMPTHREADCPU` value of the **command** parameter to the `exs_fcntl()` function, in which case the value of the **argument** parameter should be the number (starting at 0) of the CPU to which the completion thread for the connection indicated by the **fd** parameter should be pinned. The default value is `INT_MAX`, which means that the completion thread is NOT pinned to any CPU. Pinning is done at the time a connection is established, and may be changed dynamically after connection establishment. The result returned by a successful `exs_fcntl()` call is the old CPU number. If this return value is `INT_MAX`, it means the completion thread was not previously pinned to any CPU, and can be dynamically assigned to CPUs by the kernel scheduler.

## 7.5 Busy polling for completions

By default, the UNH EXS completion thread releases the CPU whenever it has no work to do in order to not consume CPU cycles unnecessarily. However, this requires kernel intervention, both to release the CPU and again when the thread needs to be reawakened, and this can add several microseconds overhead to the performance of an EXS transaction. To avoid this, the user can select the “busy polling” option for a completion thread. This option means that the completion thread will react faster to the completion of transactions, which usually produces lower latency. But it also means the completion thread will never give up the CPU, so that it will consume 100

Busy polling in the completion thread of the connection indicated by the **fd** is controlled through the use of the `EXS_F_SETFD` value of the **command** parameter to the `exs_fcntl()` function, in which case the value of the **argument** parameter should be a bit-mask containing the `EXS_FD_BUSYPOLL` flag. The choice used by a completion thread is determined when that thread is created as part of connection establishment. By default, this flag is not set, which means that busy polling is not employed by the completion thread for an **fd**. The **argument** bit-mask containing the `EXS_FD_BUSYPOLL` flag cannot be used in an `exs_fcntl()` call on an established connection. The result returned by a successful `exs_fcntl()` call is the old flags bit-mask.

When the value of the **command** parameter to the `exs_fcntl()` function is `EXS_F_GETFD`, the result returned by the `exs_fcntl()` is the current flags value at that time. The value of the **argument** parameter is ignored.

The proper way to set the `EXS_FD_BUSYPOLL` flag is to:

1. call `exs_fcntl()` with the `EXS_F_GETFD` value for the **command** parameter,
2. OR the result returned by that `exs_fcntl()` call with the `EXS_FD_BUSYPOLL` flag,
3. use the result of that OR operation as the **argument** parameter in a call to `exs_fcntl()` with the `EXS_F_SETFD` value for the **command** parameter.

Doing it this way ensures that only a single flag value (in this case, the `EXS_FD_BUSYPOLL` flag) gets changed.

Note that it is also possible to separately configure event queues to exclusively busy poll or wait for notifications (the default is busy poll for a few cycles before waiting for notification). This is described in section 11.2.6.

## 7.6 Stream receive buffer

As mentioned earlier in section 4.1, version 1.3.0 introduces an intermediate receive buffer for **SOCK\_STREAM** sockets. This buffer is meant to decrease latency for long-distance communications that is introduced by the advertisement mechanism. When no advertisements are pending, the sender will write data into this intermediate buffer instead of waiting for an advertisement. `exs_recv()` will return data from this buffer whenever it contains data, and will send advertisements only if there is no data in the buffer. As of version 1.3.2, the receiver will send advertisements even if the **MSG\_WAITALL** flag is not present. The sender will make a best effort to write directly to the user-supplied memory area whenever it gets an advertisement. However, it is possible that the sender had already written the data to the intermediate buffer before it received the advertisement, in which case it will ignore the advertisement. The actual algorithm by which this occurs is much more complex than this simple overview implies, but a full discussion of the algorithm is outside the scope of this document.

The intermediate receive buffer defaults to 3 Gigabytes in size, and this is also its maximum size due to limitations of the underlying OFED stack. There is a fallback mechanism to decrease the size if the requested size cannot be allocated. However, due to the way that memory allocation works in Linux and the fact that this intermediate buffer must be pinned in virtual memory, this fallback mechanism is unlikely to work in practice if the available physical memory is less than the desired buffer size.

To decrease the size of the intermediate receive buffer, call `exsfcntl()` with the **command** parameter set to `EXS_F_SETSTREAMBUFSIZE` and the **argument** parameter set to the desired size in bytes. This parameter must be set on the client side of the connection before calling `exs_connect()` and on the server side before calling `exs_accept()` for the first time. Note that as of UNH EXS 1.3.6, this will change the size of the remote receive buffer on each side independently. That is, setting the size on the server side will change the size of the receive buffer on the server side only, and setting the size on the client side will change the size of the receive buffer on the client side only. The similar `EXS_F_GETSTREAMBUFSIZE` may be used at any time to determine the actual size of the intermediate receive buffer, which may be less than the requested size due to memory constraints.

In many cases, the intermediate receive buffer will increase performance dramatically; however, there may be some use cases where this extra buffer is unwanted. To turn this feature off completely, set the desired stream buffer size to 0 with the `EXS_F_SETSTREAMBUFSIZE` parameter. Note that currently, this also sets the `EXS_FD_NODELAY` flag for compatibility with previous releases.

In the future, we plan to additionally support a send-side Nagle buffer for small packets. The size of this buffer would then be controllable via the `EXS_F_SETSPMAXSIZE` command to the `exsfcntl()` function mentioned in the previous section.

## 8 Registered and unregistered memory

The EXS interface is designed to transfer data using direct memory-to-memory transfers with no extra copying. This requires that memory buffers involved at each end of the

RDMA transmission be “registered” with the CA on each end prior to one side issuing the `exs_send()` and the other side issuing the `exs_recv()`. This normally requires the user to explicitly register and deregister the memory used in EXS transfers, on both the sending and receiving sides (see section 8.1). However, it is possible for a user to implicitly register and deregister the memory used on either the sending or receiving side of an EXS transfer, or both (see section 8.2).

## 8.1 Explicit memory registration and deregistration

A user explicitly registers memory using:

```
mhandle = exs_mregister(address, length, flags);
```

where **address** points to an area of memory containing **length** bytes to be registered. The value of **flags** is usually 0 to allow both send and receive operations on the memory area. The flag `EXS_MRF_RECV_DISABLE` can be specified to only allow send operations on the memory area. Note that prior to calling `exs_mregister()`, the memory at **address** must already be allocated by the user, either statically or dynamically.

If the `exs_mregister()` is successful, the value returned in **mhandle** will be an opaque memory handle that can be used as a parameter to an `exs_send()` when the send buffer is located anywhere within this area of memory, and/or as a parameter to an `exs_recv()` when the receive buffer is located anywhere within this area of memory. If the `exs_mregister()` fails for any reason, the value returned in **mhandle** will be the constant `EXS_MHANDLE_INVALID` and an error code will be stored in the global **errno**.

Note that once an **mhandle** has been successfully registered, it can be used repeatedly in subsequent calls to `exs_send()` and/or `exs_recv()` that have buffers in that memory area. Only when a user is completely finished using an explicitly registered area of memory for I/O does he or she deregister it using:

```
result = exs_mderegister(mhandle, flags);
```

where **mhandle** must be the value returned by a prior successful call to `exs_mregister()`, and the value of **flags** must be 0 since no flags are currently supported for `exs_mderegister()`.

## 8.2 Implicit memory registration and deregistration

The EXS standard provides a simple mechanism to implicitly register and deregister memory buffers used in `exs_send()` and `exs_recv()` operations. To do this, wherever an **mhandle** parameter is required in an EXS function call, the user simply supplies the constant `EXS_MHANDLE_UNREGISTERED` to indicate that the user has not explicitly registered the buffer parameter specified in that function call. Given this value, the EXS interface will dynamically register and deregister the buffer as necessary. Of course this adds considerable overhead in interface to dynamically register and deregister memory as part of an `exs_send()` or `exs_recv()` call, but if an area of memory is used for I/O only once or twice, rather than repeatedly, the user may find it more convenient to let the EXS interface perform the registration required by the CA rather than coding out explicit calls to `exs_mregister()` and `exs_mderegister()`.

Note that memory registration applies only to the process calling the `exs_mregister()` function (and to threads attached to that process). If a process forks a child process, that child does not inherit any of the parent's memory registrations.

### 8.3 Memory alignment

For best performance, large messages should be aligned at a cache-line boundary, since the channel adapter copies a cache line at a time. If a large user data buffer is unaligned, each cache line must be accessed twice by the channel adapter, which will have a noticeable effect on throughput. On most modern systems, cache lines are usually 64, 128, or 256 bytes. For the channel adapter, this can be queried using the `lspci` utility. The CPU's cache line size can be queried by examining the `/proc/cpuinfo` file on Linux. To allocate an aligned memory area, one can use the `posix_memalign()` function defined in POSIX.1-2001 or the `aligned_alloc()` function defined in ISO/IEC 9899:2011 (C11).

## 9 Synchronous I/O

The EXS interface is designed to transfer data asynchronously, as already described in section 5. However, some applications have no need for asynchronous I/O, although they still want to use RDMA. To accommodate this type of application, the user can simply supply the value `EXS_BLOCK` in the **flags** parameter to `exs_accept()`, `exs_connect()`, `exs_close()`, `exs_send()`, or `exs_recv()`. When this flag is present, **qhandle** and **ahandle** parameters required in these EXS function calls are ignored by the EXS interface and can be `NULL`. The `EXS_BLOCK` flag indicates that the user wants this function to both start an operation and wait for its completion. If the result returned by this function is -1, the error code may apply either to the start phase or the completion phase—the user has to somehow determine which. Otherwise, the result is a value taken from one of the fields in the successful completion event (which is hidden from the user by the EXS interface). For an `exs_send()` or `exs_recv()`, the result will be the number of bytes actually transmitted. For an `exs_accept()` it will be the **fd** of the new connection. For an `exs_connect()` or `exs_close()` it will be 0.

For convenience when the user wishes to both block and use unregistered memory in a send or receive operation, two additional functions have been provided: `exs_write()` and `exs_read()`. These are described below in terms of their `exs_send()` and `exs_recv()` equivalents.

### 9.1 Sending data synchronously

A user sends registered data synchronously by using:

```
result = exs_send(fd, write_buffer, write_length, EXS_BLOCK, NULL, NULL,
                 mhandle);
```

or the more convenient:

```
result = exs_blocking_send(fd, write_buffer, write_length, send_flags,
                           mhandle);
```

in which the user does not have to supply the value `EXS_BLOCK` in the `send_flags` parameter.

A user sends unregistered data synchronously by using either:

```
result = exs_send(fd, write_buffer, write_length, EXS_BLOCK, NULL, NULL,
                 EXS_MHANDLE_UNREGISTERED);
```

or the more convenient:

```
result = exs_write(fd, write_buffer, write_length);
```

In all these situations, `fd` identifies the connection, and the user has filled `write_buffer` with `write_length` bytes of data prior to the call. When using these functions, the EXS interface will automatically wait for completion of the data transfer before returning. If the transfer is successful, the value returned in `result` will be the number of bytes actually written. Otherwise, the value returned in `result` will be -1 to indicate that there was an error of some sort, in which case the error code is stored in the global `errno`.

## 9.2 Receiving data synchronously

A user receives registered data synchronously by using:

```
result = exs_recv(fd, read_buffer, max_length, EXS_BLOCK, NULL, NULL,
                 mhandle);
```

or the more convenient:

```
result = exs_blocking_recv(fd, read_buffer, max_length, recv_flags,
                          mhandle);
```

in which the user does not have to supply the value `EXS_BLOCK` in the `recv_flags` parameter.

A user receives unregistered data synchronously by using either:

```
result = exs_recv(fd, read_buffer, max_length, EXS_BLOCK, NULL, NULL,
                 EXS_MHANDLE_UNREGISTERED);
```

or the more convenient:

```
result = exs_read(fd, read_buffer, max_length);
```

In all these situations, `fd` identifies the EXS connection, and the user has reserved a `read_buffer` capable of holding `max_length` bytes of data. When using these functions, the EXS interface will automatically wait for completion of the data transfer before returning. If the transfer is successful, the value returned in `result` will be the number of bytes actually read. Otherwise, the value returned in `result` will be -1 to indicate that there was an error of some sort, in which case the error code is stored in the global `errno`.

## 9.3 Establishing an EXS client connection synchronously

The `EXS_BLOCK` constant can also be used as the value of the `flags` parameter in the `exs_connect()` function so that it will operate synchronously rather than asynchronously. When this flag value is present, the value of the `qhandle` and `ahandle` parameters to this function are ignored by the EXS interface, and can be `NULL`. For convenience and similarity to the equivalent functions for normal sockets, the following synchronous function is also provided:

```
result = exs_blocking_connect(fd, server_address, server_addrlen);
```

where the parameters are identical to those for the corresponding “normal” TCP/IP socket `connect()` function.

## 9.4 Accepting connections on the server synchronously

The `EXS_BLOCK` constant can also be used as the value of the `flags` parameter in the `exs_accept()` function so that it will operate synchronously rather than asynchronously. When this flag value is present, the value of the `qhandle` parameter to this function is ignored by the EXS interface, and can be `NULL`. For convenience and similarity to the equivalent function for normal sockets, the following synchronous function is also provided:

```
result = exs_blocking_accept(fd, peer_address, &peer_addrlen);
```

`exs_blocking_accept()` blocks until a remote client connects, at which time the **result** it returns will be the **fd** for the new connection to the remote client. No `exs_event_t` structure is generated for the user, so the other fields in that structure are not available to the user.

It is important to note that the second and third parameters to `exs_blocking_accept()` differ from the corresponding parameters to `exs_accept()` (see section 4.3.2). Instead of giving an **address\_vector** array of structures of type `exs_acceptaddr` and a **count** parameter indicating the number of elements in the array, `exs_blocking_accept()` takes as its **peer\_address** parameter a pointer to a structure of type `struct sockaddr` into which the IPv4/IPv6 address and port number of a new remote client will be stored by the EXS interface. The **peer\_addrlen** parameter is the number of bytes allocated by the user to the structure pointed to by **peer\_address**. These two parameters to `exs_blocking_accept()` are identical to the first two fields in an element of the **address\_vector** parameter to `exs_accept()`, and make the 3 parameters to `exs_blocking_accept()` identical to the 3 parameters of the “normal” TCP/IP socket function `accept()`.

Note also that the `EXS_BLOCK` flag does NOT have to be included in the **accept\_flags** parameter to `exs_blocking_accept()`.

## 9.5 Closing EXS connections synchronously

An EXS connection of any type can be closed synchronously by using either:

```
result = exs_close(fd, EXS_BLOCK, NULL, NULL);
```

or the more convenient:

```
result = exs_blocking_close(fd);
```

where **fd** indicates the connection to be closed.

## 10 Converting programs from using normal sockets to using EXS sockets

The availability in EXS of both implicit memory registration and synchronous I/O means that users who wish to convert existing programs using “normal” sockets (see section 10.1) to full use of EXS sockets have a choice of how to proceed. They can start by using EXS sockets in synchronous mode and use only implicit memory registration (see section 10.2). Such usage is almost identical with the use of “normal” sockets (essentially just the function names change), but it does give the user access to the RDMA hardware. Once that is working, the user can choose between converting first to using explicit memory registration while continuing to use synchronous mode (see section 10.3), or converting first to using asynchronous mode while continuing to use implicit memory registration (see section 10.4). Once that step has been finished and is working properly, the other step can be taken to give fully asynchronous operation with explicitly registered memory (see section 10.5).

### 10.1 Client using normal sockets

The following gives an example of the complete conversion just mentioned, starting with the general outline for a client that uses “normal” sockets (and therefore cannot use RDMA hardware):

```
fd = socket (PF_INET, SOCK_STREAM, 0);
connect (fd, server_address, server_addrlen);
loop
    write (fd, out_buffer, out_bytes);
    in_bytes = read (fd, in_buffer, maxbytes);
endloop;
close (fd);
```

### 10.2 Client using EXS sockets in synchronous mode with implicit memory registration

The first conversion step mentioned above gives an almost identical program that uses synchronous mode and implicit memory registration so that it can therefore use RDMA hardware. Note that this program is identical to the normal program except for the introduction of the one-time call to `exs_init()` at the start, and the name changes of the various socket functions.

```
exs_init (EXS_VERSION);
fd = exs_socket (PF_INET, SOCK_STREAM, 0);
exs_blocking_connect (fd, server_address, server_addrlen);
loop
    exs_write (fd, out_buffer, out_bytes);
    in_bytes = exs_read (fd, in_buffer, maxbytes);
endloop;
exs_blocking_close (fd);
```

### 10.3 Client using EXS sockets in synchronous mode with explicit memory registration

We now have a choice of which feature of EXS to apply first. Let's choose to register memory before we go to asynchronous operation (so we will continue to use only EXS functions that are "blocking"). The general outline for this version of the client would be:

```
exs_init(EXS_VERSION);
fd = exs_socket(PF_INET, SOCK_STREAM, 0);
exs_blocking_connect(fd, server_address, server_addrlen);
in_mhandle = exs_mregister(in_buffer, max_in_bytes, flags);
out_mhandle = exs_mregister(out_buffer, max_out_bytes, flags);
loop
    exs_blocking_send(fd, out_buffer, out_bytes, 0, out_mhandle);
    in_bytes = exs_blocking_recv(fd, in_buffer, in_bytes, 0, in_mhandle);
endloop;
exs_mderegister(out_mhandle);
exs_mderegister(in_mhandle);
exs_blocking_close(fd);
```

### 10.4 Client using EXS sockets in asynchronous mode with implicit memory registration

Alternatively, we could choose to utilize EXS asynchronous operations before registering memory. The general outline for this version of the client would be:

```
exs_init(EXS_VERSION);
fd = exs_socket(PF_INET, SOCK_STREAM, 0);
management_qhandle = exs_qcreate(1);
exs_connect(fd, server_address, server_addrlen, 0, NULL, management_qhandle,
NULL);
/*---- perform computation in parallel with EXS activity ----*/
exs_qdequeue(management_qhandle, &management_event, 1, NULL);
in_qhandle = exs_qcreate(3);
out_qhandle = exs_qcreate(3);
loop
    exs_send(fd, out_buffer, out_bytes, 0, out_qhandle,
NULL, EXS_MHANDLE_UNREGISTERED);
    /*---- perform computation in parallel with data transfer ----*/
    exs_qdequeue(out_qhandle, &out_event, 1, NULL);
    exs_recv(fd, in_buffer, max_in_bytes, 0, in_qhandle,
NULL, EXS_MHANDLE_UNREGISTERED);
    /*---- perform computation in parallel with data transfer ----*/
    exs_qdequeue(in_qhandle, &in_event, 1, NULL);
    in_bytes = in_event.exs_evt_union.exs_evt_xfer.exs_evt_length;
endloop;
exs_qdelete(out_qhandle);
exs_qdelete(in_qhandle);
exs_close(fd, 0, management_qhandle, NULL);
/*---- perform computation in parallel with EXS activity ----*/
exs_qdequeue(management_qhandle, &management_event, 1, NULL);
exs_qdelete(management_qhandle);
```

## 10.5 Client using EXS sockets in asynchronous mode with explicit memory registration

Our final step is to combine the changes made independently in the previous two steps, giving us a program that uses both EXS registered memory and EXS asynchronous I/O for RDMA transfers:

```
exs_init(EXS_VERSION);
fd = exs_socket(PF_INET, SOCK_STREAM, 0);
management_qhandle = exs_qcreate(1);
exs_connect(fd, server_address, server_addrlen, 0, NULL, management_qhandle,
            NULL);
/*---- perform computation in parallel with EXS activity ----*/
exs_qdequeue(management_qhandle, &management_event, 1, NULL);
in_mhandle = exs_mregister(in_buffer, max_in_bytes, flags);
out_mhandle = exs_mregister(out_buffer, max_out_bytes, flags);
in_qhandle = exs_qcreate(3);
out_qhandle = exs_qcreate(3);
loop
    exs_send(fd, out_buffer, out_bytes, 0, out_qhandle, NULL, out_mhandle);
    /*---- perform computation in parallel with data transfer ----*/
    exs_qdequeue(out_qhandle, &out_event, 1, NULL);
    exs_rcv(fd, in_buffer, max_in_bytes, 0, in_qhandle, NULL, in_mhandle);
    /*---- perform computation in parallel with data transfer ----*/
    exs_qdequeue(in_qhandle, &in_event, 1, NULL);
    in_bytes = in_event.exs_evt_union.exs_evt_xfer.exs_evt_length;
endloop;
exs_qdelete(out_qhandle);
exs_qdelete(in_qhandle);
exs_mderegister(out_mhandle);
exs_mderegister(in_mhandle);
exs_close(fd, 0, management_qhandle, NULL);
/*---- perform computation in parallel with EXS activity ----*/
exs_qdequeue(management_qhandle, &management_event, 1, NULL);
exs_qdelete(management_qhandle);
```

This program runs, but as it stands there isn't much parallel activity between the asynchronous EXS activity and the user thread, which should perform parallel computation in the places now marked in the code only by appropriate comments. To take advantage of the potential parallelism, this program needs to be modified to perform useful computation between the call of an `exs_close()` or an `exs_connect()` or an `exs_send()` or an `exs_rcv()` that starts an EXS operation and the corresponding `exs_qdequeue()` that waits for the completion of the EXS operation.

## 11 Status of UNH EXS 1.4.0

### 11.1 Comparison with the Extended Sockets API (ES-API) Issue 1.0 Specification

UNH EXS function	origin	UNH EXS status	section discussed
<code>exs_accept()</code>	ES-API standard	implemented	4.3.2
<code>exs_bind()</code>	POSIX	implemented	4.3.1
<code>exs_blocking_accept()</code>	POSIX	implemented	9.4
<code>exs_blocking_close()</code>	POSIX	implemented	9.5
<code>exs_blocking_connect()</code>	POSIX	implemented	9.3
<code>exs_blocking_recv()</code>	non-standard	implemented	9.1
<code>exs_blocking_send()</code>	non-standard	implemented	9.1
<code>exs_cancel()</code>	ES-API standard	not implemented	
<code>exs_close()</code>	non-standard	implemented	4.4
<code>exs_connect()</code>	ES-API standard	implemented	4.2
<code>exs_init()</code>	ES-API standard	implemented	3.1
<code>exs_fcntl()</code>	non-standard	implemented	7
<code>exs_listen()</code>	POSIX	implemented	4.3.1
<code>exs_mderegister()</code>	ES-API standard	implemented	8.1
<code>exs_mmodify()</code>	ES-API standard	not implemented	
<code>exs_mregister()</code>	ES-API standard	implemented	8.1
<code>exs_poll()</code>	ES-API standard	not implemented	
<code>exs_qcreate()</code>	ES-API standard	implemented	2.1.1
<code>exs_qdelete()</code>	ES-API standard	implemented	2.1.2
<code>exs_qdequeue()</code>	ES-API standard	implemented	2.1.3
<code>exs_qmodify()</code>	ES-API standard	implemented	
<code>exs_qstatus()</code>	ES-API standard	implemented	
<code>exs_read()</code>	POSIX	implemented	9.2
<code>exs_recv()</code>	ES-API standard	implemented	5.2
<code>exs_recvmsg()</code>	ES-API standard	not implemented	
<code>exs_send()</code>	ES-API standard	implemented	5.1
<code>exs_sendfile()</code>	ES-API standard	not implemented	
<code>exs_sendmsg()</code>	ES-API standard	not implemented	
<code>exs_shutdown()</code>	non-standard	implemented	4.5
<code>exs_socket()</code>	POSIX	implemented	4.1
<code>exs_write()</code>	POSIX	implemented	9.1

### 11.2 Modifications to the ES-API standard in the UNH EXS implementation

#### 11.2.1 `exs.h` header file

The definitions of all symbols, structures, and function prototypes introduced by UNH EXS are found in the header file “`exs.h`”, not “`sys/exs.h`” as stated in the ES-API standard. There-

fore, each “.c” file using UNH EXS should have the following line after all other “include” directives at the beginning of the compilation unit:

```
#include <exs.h>
```

The `EXS_DISABLE_DEPRECATED` feature test macro, if defined before “exs.h” is included, will cause deprecated symbols to not be made visible to the compiler. This allows users to verify that their programs do not make use of deprecated symbols.

### 11.2.2 `exs_accept()`

The `EXS_BLOCK` flag has been added in UNH-EXS to indicate that the user wants this function to both start an operation and wait for its completion. When this flag is present, the value of the **qhandle** parameter required in the `exs_accept()` function call is ignored by the EXS interface and can be NULL. The result returned by a successful `exs_accept()` with the `EXS_BLOCK` flag is the **fd** of the new connection to a remote client.

### 11.2.3 `exs_close()`

The `EXS_BLOCK` flag has been added in UNH-EXS to indicate that the user wants this function to both start an operation and wait for its completion. When this flag is present, the values of the **qhandle** and **ahandle** parameters that are required in the `exs_close()` function call are ignored by the EXS interface and can be NULL.

### 11.2.4 `exs_connect()`

A non-NULL value for the **timeout** parameter, allowed in the ES-API standard, is not yet supported in UNH-EXS.

The `EXS_BLOCK` flag has been added in UNH-EXS to indicate that the user wants this function to both start an operation and wait for its completion. When this flag is present, the values of the **qhandle** and **ahandle** parameters that are required in the `exs_connect()` function call are ignored by the EXS interface and can be NULL.

### 11.2.5 `exs_mregister()`

The nonstandard `EXS_LOCAL_READ`, `EXS_LOCAL_WRITE`, `EXS_REMOTE_READ`, `EXS_REMOTE_WRITE`, and `EXS_ACCESS_ALL` flags introduced in prior UNH EXS versions are deprecated as of UNH EXS 1.3.6. Prior to UNH EXS 1.3.6, memory registration handles created without specifying any flags were usable only with `exs_send()`. In UNH EXS 1.3.6, this default is changed and memory registration handles created without specifying any flags are now usable with both `exs_send()` and `exs_recv()`, and a new flag `EXS_MRF_RECV_DISABLE` was added to create a memory registration handle usable for sending but not receiving.

### 11.2.6 `exs_qcreate()`

The `exs_qcreate()` function takes an integer parameter specifying the depth of the queue. In UNH EXS, this is only used to determine the number of event structures to pre-allocate; the queue can hold a theoretically unlimited number of elements. That said, limiting the

usage of the queue to the specified depth will improve performance by not requiring calls to `malloc()` in the data transfer path.

### 11.2.7 `exs_qmodify()`

The `exs_qmodify()` function, in addition to accepting the standard `EXS_QATTR_SIGNAL` and `EXS_QATTR_EVENTS` attributes, takes an attribute specific to UNH EXS, `EXS_QATTR_WAIT`, which was introduced in UNH EXS 1.3.6. This attribute specifies how `exs_qdequeue()` should wait if there are no events posted when it is called. Note that this attribute may only be modified prior to the first time that the queue is used, since it may alter the way that the queue is synchronized internally for access by multiple threads.

The default value of this attribute is `EXS_WAIT_ADAPTIVE`, meaning `exs_qdequeue()` will busy poll the queue for a given number of cycles before giving up the CPU and waiting for notification of a new event. The number of cycles to spin is adjusted on-the-fly based on whether or not events are retrieved via the initial busy polling loop. The exact number of cycles and how it is adjusted are subject to change in future versions of UNH EXS.

The other values of this attribute are `EXS_WAIT_BUSY_POLL` and `EXS_WAIT_NOTIFY`, which cause `exs_qdequeue()` to busy poll or give up the CPU and wait for notification, respectively.

Note that the wait type for event queues is completely independent of the wait type for the RDMA completion thread, as discussed in Section 7.5.

### 11.2.8 `exs_qstatus()`

The `exs_qstatus()` function accepts the `EXS_QATTR_WAIT` attribute, which returns the wait type of the queue.

### 11.2.9 `exs_recv()`

The ES-API standard `MSG_PEEK` and `MSG_OOB` flag values are not supported in UNH-EXS.

The `EXS_DONTWAIT` flag is ignored as of UNH-EXS 1.3.1.

The `EXS_CREDIT_WAIT` flag has been added in UNH-EXS 1.3.2 to indicate that the user wants this function to wait until a credit is available, but otherwise run asynchronously.

The `EXS_BLOCK` flag has been added in UNH-EXS to indicate that the user wants this function to both start an operation and wait for its completion. When this flag is present, the values of the `qhandle` and `ahandle` parameters that are required in the `exs_recv()` function call are ignored by the EXS interface and can be NULL.

The `EXS_UNSIGNALLED` flag has been added in UNH-EXS to indicate that the user does not want the completion of an asynchronous `exs_recv()` to generate an event. It is ignored when the `EXS_BLOCK` flag is present. When `EXS_UNSIGNALLED` is present, the value of the `qhandle` parameter that is required in the `exs_recv()` function call may be NULL, in which case no event is generated upon the completion of the `exs_recv()`. However, if the value of the `qhandle` parameter is not NULL, an event will be generated in the event queue ONLY if the operation did NOT complete successfully—no event is generated if the operation completed successfully.

### 11.2.10 `exs_send()`

The ES-API standard `MSG_EOR` and `MSG_OOB` flags are not supported in UNH-EXS.

The `EXS_DONTWAIT` flag is ignored as of UNH-EXS 1.3.1.

The `EXS_CREDIT_WAIT` flag has been added in UNH-EXS 1.3.2 to indicate that the user wants this function to wait until a credit is available, but otherwise run asynchronously.

The `EXS_BLOCK` flag has been added in UNH-EXS to indicate that the user wants this function to both start an operation and wait for its completion. When this flag is present, the values of the **qhandle** and **ahandle** parameters that are required in the `exs_send()` function call are ignored by the EXS interface and can be `NULL`.

The `EXS_UNSIGNALED` flag has been added in UNH-EXS to indicate that the user does not want the completion of an asynchronous `exs_send()` to generate an event. It is ignored when the `EXS_BLOCK` flag is present. When `EXS_UNSIGNALED` is present, the value of the **qhandle** parameter that is required in the `exs_send()` function call may be `NULL`, in which case no event is generated upon the completion of the `exs_send()`. However, if the value of the **qhandle** parameter is not `NULL`, an event will be generated in the event queue **ONLY** if the operation did **NOT** complete successfully—no event is generated if the operation completed successfully.

### 11.2.11 `exs_sendfile()`

`exs_sendfile()` currently only supports file sizes of under 1GB, and only supports a single `sendvec` element per call (`sendvec_cnt == 1`). Additionally, the `exs_iov` field of `union exs_xfvec` is not supported.

UNH EXS supports an additional member of `union exs_xfvec` not specified in the ES-API standard, `exs_pathv`, which takes a pathname string as opposed to a file descriptor. This will cause UNH EXS to implicitly open the file at the start of the operation and close it at the end of the operation.

The `EXS_UNSIGNALED` flag has been added in UNH-EXS to indicate that the user does not want the completion of an asynchronous `exs_send()` to generate an event. It is ignored when the `EXS_BLOCK` flag is present. When `EXS_UNSIGNALED` is present, the value of the **qhandle** parameter that is required in the `exs_send()` function call may be `NULL`, in which case no event is generated upon the completion of the `exs_send()`. However, if the value of the **qhandle** parameter is not `NULL`, an event will be generated in the event queue **ONLY** if the operation did **NOT** complete successfully—no event is generated if the operation completed successfully.

## 11.3 Known deficiencies

### 11.3.1 thread cancellation

At the present time, the UNH EXS library functions are **NOT** cancellation safe, because there are **NO** cancellation cleanup handlers implemented for any of them. Users are therefore advised **NOT** to call `pthread_cancel()` for any of their threads when they might be executing in code using the UNH EXS library.

### 11.3.2 fork()

The UNH EXS library has several limitations regarding the forking of child processes, due to limitations of the OpenFabrics software stack that it is built on top of:

- Calling `exs_init()` in a child process is not supported, unless the process has first called `exec()`. This means that EXS is completely unusable in a child process after a `fork()`, unless the child process calls `exec()`.
- Registered memory is unavailable to child processes. That is, the memory area that is registered will not be mapped in the child process. See the Linux `madvise(2)` manual page for more information.

If you do not need fork support, you may set the **EXS\_FORK\_UNSAFE** environment variable. This means that your application *must not* call `fork()` once your process has called `exs_init()`, whether implicitly or explicitly, or silent data corruption may result.