

UPC++ Specification  
v1.0 Draft 4

UPC++ Specification Working Group  
[upcxx-spec@googlegroups.com](mailto:upcxx-spec@googlegroups.com)

September 27, 2017

### **Abstract**

UPC++ is a C++11 library providing classes and functions that support Asynchronous Partitioned Global Address Space (APGAS) programming. We are revising the library under the auspices of the DOE's Exascale Computing Project, to meet the needs of applications requiring PGAS support. UPC++ is intended for implementing elaborate distributed data structures where communication is irregular or fine-grained. The UPC++ interfaces for moving non-contiguous data and handling memories with different optimal access methods are composable and similar to those used in conventional C++. The UPC++ programmer can expect communication to run at close to hardware speeds.

The key facilities in UPC++ are global pointers, that enable the programmer to express ownership information for improving locality, one sided communication, both put/get and RPC, futures and continuations. Futures to capture data readiness state, which is useful in making scheduling decisions, and continuations provide for completion handling via callbacks. Together, these enable the programmer to chain together a DAG of operations to execute asynchronously as high-latency dependencies become satisfied.

### **Acknowledgments**

This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration

Early development of UPC++ was supported by the Director, Office of Science, Office of Advanced Scientific Computing Research, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

### **Copyright**

This manuscript has been authored by an author at Lawrence Berkeley National Laboratory under Contract No. DE-AC02-05CH11231 with the U.S. Department of Energy. The U.S. Government retains, and the publisher, by accepting the article for publication, acknowledges, that the U.S. Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for U.S. Government purposes.

### **Legal Disclaimer**

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor the Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or the Regents of the University of California.

# Contents

<b>1</b>	<b>Overview and Scope</b>	<b>2</b>
1.1	Preliminaries . . . . .	2
1.2	Execution Model . . . . .	4
1.3	Memory Model . . . . .	5
1.4	Organization of this Document . . . . .	5
1.5	Document Conventions . . . . .	5
1.6	Glossary . . . . .	6
<b>2</b>	<b>Init and Finalize</b>	<b>9</b>
2.1	Overview . . . . .	9
2.2	Hello World . . . . .	9
2.3	API Reference . . . . .	10
<b>3</b>	<b>Global Pointers</b>	<b>12</b>
3.1	Overview . . . . .	12
3.2	API Reference . . . . .	13
<b>4</b>	<b>Storage Management</b>	<b>20</b>
4.1	Overview . . . . .	20
4.2	API Reference . . . . .	20
<b>5</b>	<b>Futures and Promises</b>	<b>24</b>
5.1	Overview . . . . .	24
5.2	The Basics of Asynchronous Communication . . . . .	24
5.3	Working with Promises . . . . .	26
5.4	Advanced Callbacks . . . . .	27
5.5	Execution Model . . . . .	29
5.6	Anonymous Dependencies . . . . .	31
5.7	Lifetime and Thread Safety . . . . .	32
5.8	API Reference . . . . .	33

---

5.8.1	future	33
5.8.2	promise	36
<b>6</b>	<b>Serialization</b>	<b>39</b>
6.1	Functions	40
<b>7</b>	<b>One-Sided Communication</b>	<b>41</b>
7.1	Overview	41
7.2	Completion	41
7.3	API Reference	42
7.3.1	Remote Puts	42
7.3.2	Remote Gets	43
<b>8</b>	<b>Remote Procedure Call</b>	<b>46</b>
8.1	Overview	46
8.2	Remote Hello World Example	47
8.3	API Reference	47
<b>9</b>	<b>Progress</b>	<b>52</b>
9.1	Overview	52
9.2	Restricted Context	53
9.3	Attentiveness	54
9.4	Thread Personas/Notification Affinity	54
9.5	API Reference	57
9.5.1	persona	58
9.5.2	persona_scope	61
9.5.3	Outgoing Progress	62
<b>10</b>	<b>Atomics</b>	<b>63</b>
10.1	Overview	63
10.2	API Reference	63
<b>11</b>	<b>Teams</b>	<b>66</b>
11.1	Overview	66
11.2	Local Teams	66
11.3	API Reference	67
11.3.1	team	67
11.3.2	team_id	69
11.3.3	Fundamental Teams	69
11.3.4	Collectives	70

<b>12 Distributed Objects</b>	<b>73</b>
12.1 Overview . . . . .	73
12.2 Building Distributed Objects . . . . .	74
12.3 Ensuring Distributed Existence . . . . .	74
12.4 API Reference . . . . .	75
<b>13 Non-Contiguous One-Sided Communication</b>	<b>78</b>
13.1 Overview . . . . .	78
13.2 API Reference . . . . .	79
13.2.1 Fragmented Put . . . . .	79
13.2.2 Fragmented Get . . . . .	81
13.2.3 Fragmented Put then RPC . . . . .	82
13.2.4 Fragmented Regular Put . . . . .	84
13.2.5 Fragmented Regular Get . . . . .	86
13.2.6 Fragmented Regular Put then RPC . . . . .	87
13.2.7 Strided Put . . . . .	89
13.2.8 Strided Get . . . . .	91
13.2.9 Strided Put then RPC . . . . .	92
<b>14 Memory Kinds</b>	<b>95</b>
<b>A Notes for Implementers</b>	<b>96</b>
A.1 <code>future_element_t</code> and <code>future_element_moved_t</code> . . . . .	96
A.2 <code>future&lt;T...&gt;::when_all</code> . . . . .	97
A.3 <code>to_future</code> . . . . .	98
A.4 <code>future_invoke_result_t</code> . . . . .	98
A.5 <code>promise_invoke_result_t</code> . . . . .	99
<b>Bibliography</b>	<b>100</b>

# 1 Chapter 1

## 2 Overview and Scope

### 3 1.1 Preliminaries

4 UPC++ is a C++11 library providing classes and functions that support Asynchronous  
5 Partitioned Global Address Space (APGAS) programming. The project began in 2012  
6 with a prototype AKA V0.1, described in the IPDPS14 paper by *Zheng et al.* [3]. This  
7 document describes a production version, V1.0, with the addition of several features and  
8 a new asynchronous API.

9 Under the PGAS model, a distributed memory parallel computer is viewed abstractly  
10 as a collection of *processing elements*, an individual computing resource, each with *local*  
11 *memory* (see Fig. 1.1). A processing element is called a *rank* in UPC++. The execution  
12 model of UPC++ is SPMD and the number of UPC++ ranks is fixed during program execution.

13 As with conventional C++ threads programming, ranks can access their respective  
14 local memory via a pointer. However, the PGAS abstraction supports a global address  
15 space, which is allocated in *shared segments* distributed over the ranks. A *global pointer*  
16 enables the programmer to move data in the shared segments between ranks as shown in  
17 Fig. 1.1. As with threads programming, references made via global pointers are subject to  
18 race conditions, and appropriate synchronization must be employed.

19 UPC++ global pointers are fundamentally different from conventional C-style pointers. A  
20 global pointer refers to a location in a shared segment. It cannot be dereferenced using the  
21  $\star$  operator, and it does not support conversions between pointers to base and derived types.  
22 It also cannot be constructed by the address-of operator. On the other hand, UPC++ global  
23 pointers *do* support some properties of a regular C pointer, such as pointer arithmetic and  
24 passing a pointer by value.

25 Notably, global pointers are used in *one-sided* communication: bulk copying operations  
26 (RMA) similar to *memcpy* but across ranks (Ch. 7), and in Remote Procedure Calls




Figure 1.1: Abstract Machine Model of a PGAS program memory

1 (RPC, Ch. 8). RPC enables the programmer to ship functions to other ranks, which is  
 2 useful in managing irregular distributed data structures. These ranks can push or pull  
 3 data via global pointers. *Futures* and *Promises* (Ch. 5) are used to determine completion  
 4 of communication or to provide handlers that respond to completion. Wherever possible,  
 5 UPC++ will engage low-level hardware support for communication and this capability is  
 6 crucial to UPC++’s support of *lightweight communication*.

7 UPC++’s design philosophy is to provide “close to the metal performance.” To meet this  
 8 requirement, UPC++ imposes certain restrictions. In particular, non-blocking communica-  
 9 tion is the default for nearly all operations defined in the API, and all communication is  
 10 explicit. These two restrictions encourage the programmer to write code that is perfor-  
 11 mant and make it more difficult to write code that is not. Conversely, UPC++ relaxes some  
 12 restrictions found in models such as MPI; in particular, it does not impose an in-order  
 13 delivery requirement between separate communication operations. The added flexibility  
 14 increases the possibility of overlapping communication and scheduling it appropriately.

15 UPC++ also avoids non-scalable constructs found in models such as UPC. For example,  
 16 it does not support shared distributed arrays or shared scalars. Instead, it provides dis-  
 17 tributed objects, which can be used to similar ends (Ch. 12). Distributed objects are useful  
 18 in solving the *bootstrapping problem*, whereby ranks need to distribute their local copies  
 19 of global pointers to other ranks. Though UPC++ does not directly provide multidimen-  
 20 sional arrays, applications that use UPC++ may define them. To this end, UPC++ supports  
 21 non-contiguous data transfers: vector, indexed, and strided data (Ch. 13).

22 Because UPC++ does not provide separate concurrent threads to manage progress, UPC++  
 23 must manage all progress inside active calls to the library. UPC++ has been designed with a  
 24 policy against the use of internal operating system threads. The strengths of this approach  
 25 are improved user-visibility into the resource requirements of UPC++ and better interoper-  
 26 ability with software packages and their possibly restrictive threading requirements. The  
 27 consequence, however, is that the user must be conscientious to balance the need for mak-  
 28 ing progress against the application’s need for CPU cycles. Chapter 9 discusses subtleties

1 of managing progress and how an application can arrange for UPC++ to advance the state  
2 of asynchronous communication.

3 Ranks may be grouped into teams (Ch. 11). A team can participate in collective  
4 operations. Teams are also the interface that UPC++ uses to propagate the shared memory  
5 capabilities of the underlying hardware and operating system and can let a programmer  
6 reason about hierarchical processor-memory organization, allowing an application to reduce  
7 its memory footprint. UPC++ supports atomic operations, currently on remote 32-bit and  
8 64-bit integers. Atomics are useful in managing distributed queues, hash tables, and so  
9 on. However, as explained in the discussion below on UPC++'s memory model, atomics are  
10 split phased and not handled the same way as they are in C++11 and other libraries.

11 UPC++ will support memory kinds (Ch. 14), whereby the programmer can identify re-  
12 gions of memory requiring different access methods or having different performance prop-  
13 erties, such as device memory. Since memory kinds will be implemented in Year 2, we will  
14 defer their detailed discussion until next year.

## 15 1.2 Execution Model

16 The UPC++ internal state contains, for each rank, internal unordered queues that are man-  
17 aged for the user. The UPC++ progress engine scans these queues for operations initiated by  
18 this rank, as well as externally generated operations that target this rank. The progress en-  
19 gine is active inside UPC++ calls only and is quiescent at other times, as there are no threads  
20 or background processes executing inside UPC++. This passive stance permits UPC++ to be  
21 driven by any other execution model a user might choose. This universality does place a  
22 small burden on the user: calling into the `progress` function. UPC++ relies on the user to  
23 make periodic calls into the `progress` function to ensure that UPC++ operations are com-  
24 pleted. `progress` is the mechanism by which the user loans UPC++ a thread of execution  
25 to perform operations that target the given rank. The user can determine that a specific  
26 operation completes by checking the status of its associated `future`, or by attaching a  
27 completion handler to the operation.

28 UPC++ presents a *thread-aware* programming model. It assumes that only one thread  
29 of execution is interacting with any UPC++ object. The abstraction for thread-awareness  
30 in UPC++ is the *persona*. A `future` produced by a thread of execution is associated with  
31 its persona, and transferring the `future` to another thread must be accompanied by trans-  
32 ferring the underlying persona. Each rank has a *master persona*, initially attached to the  
33 thread that calls `init`. Some UPC++ operations, such as `barrier`, require a thread to have  
34 exclusive access to the master persona to call them. Thus, the programmer is responsible  
35 for ensuring synchronized access to both personas and memory, and that access to shared  
36 data does not interfere with the internal operation of UPC++.



## 1.3 Memory Model

The UPC++ memory model differs from that of C++11 (and beyond) in that all updates are split-phased: every communication operation has a distinct initiate and wait step. Thus, atomic operations execute over a time interval, and the time intervals of successive operations that target the same datum must not overlap, or a data race will result.

UPC++ differs from MPI in that it doesn't guarantee in-order delivery. For example, if we overlap two successive RPC operations involving the same source and destination rank, we cannot say which one completes first.

## 1.4 Organization of this Document

This specification is intended to be a normative reference - a Programmer's Manual is forthcoming. For the purposes of understanding the key ideas in UPC++, we recommend that the novice reader skip Chapter 9 (Progress) and the advanced topics related to futures, personas, and continuation-based communication.

The organization for the rest of the document is as follows. Chapter 2 discusses the process of starting up and closing down UPC++. Global pointers (Ch. 3) are fundamental to the PGAS model, and Chapter 4 discusses storage allocation. Since UPC++ supports asynchronous communication only, UPC++ provides futures and promises (Ch. 5) to manage control flow and completion. Chapters 7 and 8 describe the two forms of asynchronous one-sided communication, `rput/rget` and RPC, respectively. Chapter 9 discusses progress. Chapter 10 discusses atomic operations. Chapter 11 discusses teams, which are a means of organizing UPC++ ranks. Chapter 12 discusses distributed objects. Chapter 13 discusses non-contiguous data transfers. Chapter 14 discusses memory kinds.

## 1.5 Document Conventions

1. C++ language keywords are in the color `mocha`.
2. UPC++ terms are set in the color `bright blue` except when they appear in a synopsis framebox.
3. All functions are declared `noexcept` unless specifically called out.
4. All entities are in the `upcxx` namespace unless otherwise qualified.

## 1.6 Glossary

**Affinity.** A binding of each location in a shared segment to a particular rank (generally the rank which allocated that shared object). Every byte of shared memory has affinity to exactly one rank (at least logically).

**C++ Concepts.** E.g. `TriviallyCopyable`. This document references C++ Concepts as defined in the C++14 standard [2] when specifying the semantics of types. However, compliant implementations are still possible within a compiler adhering to the earlier C++11 standard [1].

**Collective.** A constraint placed on some language operations which requires evaluation of such operations to be matched across all ranks. The behavior of collective operations is undefined unless all ranks execute the same sequence of collective operations.

A collective operation need not provide any actual synchronization between ranks, unless otherwise noted. The collective requirement simply states a relative ordering property of calls to collective operations that must be maintained in the parallel execution trace for all executions of any legal program. Some implementations may include unspecified synchronization between ranks within collective operations, but programs must not rely upon the presence or absence of such unspecified synchronization for correctness.

**Futures (and Promises)** (5) The primary mechanisms by which a UPC++ application interacts with non-blocking operations. The semantics of futures and promises in UPC++ differ from the those of standard C++. While futures in C++ facilitate communicating between threads, the intent of UPC++ futures is solely to provide an interface for managing and composing non-blocking operations, and they cannot be used directly to communicate between threads or ranks. A future is the interface through which the status of the operation can be queried and the results retrieved, and multiple future objects may be associated with the same promise. A future thus represents the consumer side of a non-blocking operation. A promise represents the producer side of the operation, and it is through the promise that the results of the operation are supplied and its dependencies fulfilled.

**Global pointer.** (3) The primary way to address memory in a shared memory segment of a UPC++ program. Global pointers can themselves be stored in shared memory or otherwise passed between ranks and retain their semantic meaning to any rank.

**Local.** Refers to an object or reference with affinity to a rank in the local team (11.2).

1 **Operation completion.** (7.2) The condition where a communication operation is  
2 complete with respect to the initiating rank, such that its effects are visible and  
3 that resources, such as source and destination memory regions, are no longer in  
4 use by UPC++.

5 **Persona.** (9.4) The abstraction for thread-awareness in UPC++. A UPC++ persona  
6 object represents a collection of UPC++-internal state usually attributed to a sin-  
7 gle thread. By making it a proper construct, UPC++ allows a single OS thread  
8 to switch between multiple application-defined roles for processing notifications.  
9 Personas act as the receivers for notifications generated by the UPC++ runtime.

10 **Private object.** An object outside the shared space that can be accessed only by  
11 the rank that owns it (e.g. an object on the program stack).

12 **Progress.** (9) The means by which the application allows the UPC++ runtime to  
13 advance the state of outstanding operations initiated by this or other ranks, to  
14 ensure they eventually complete.

15 **Rank.** An OS process that is a member of a UPC++ parallel job execution. UPC++  
16 uses a SPMD execution model, and the number of ranks is fixed during a given  
17 program execution. The placement of ranks across physical processors or NUMA  
18 domains is implementation-dependent.

19 **Referentially transparent.** A routine that is a pure function, where inputs alone  
20 determine the value returned by the function. For the same inputs, repeated  
21 calls to a referentially transparent function will always return the same result.

22 **Remote.** Refers to an object or reference whose affinity is not local to the current  
23 rank.

24 **Remote Procedure Call.** A communication operation that injects a function call  
25 invocation into the execution stream of another rank. These injections are  
26 one-sided, meaning the target rank need not explicitly expect the incoming op-  
27 eration or perform any specific action to receive it, aside from invoking UPC++  
28 progress.

29 **Serializable.** (6) A C++ object that is either TriviallyCopyable, or for which there  
30 is a user-supplied implementation of the visitor function `serialize`.

31 **Source completion.** The condition where a communication operation initiated by  
32 the current rank has advanced to a point where serialization of the local source  
33 memory regions for the operation has occurred, and the contents of those re-  
34 gions can be safely overwritten or reclaimed without affecting the behavior of the  
35 ongoing operation. Source completion does not generally imply operation com-  
36 pletion, and other effects of the operation (e.g., updating destination memory  
37 regions, or delivery to a remote rank) may still be in-progress.

1       **Shared segment.** A region of storage associated with a particular rank that is used  
2           to allocate shared objects that are accessible by any rank.

3       **Team.** A UPC++ object representing an ordered set of ranks.

4       **Thread (or OS thread).** An independent stream of executing instructions with  
5           private state. A rank process may contain many threads (created by the appli-  
6           cation), and each is associated with at least one persona.

# 1 Chapter 2

## 2 Init and Finalize

### 3 2.1 Overview

4 The `init` function must be called before any other UPC++ function can be invoked. This  
5 can happen anywhere in the program, so long as it appears before any UPC++ calls that  
6 require the library to be in an initialized state. The call is *collective*, meaning every process  
7 in the parallel job must enter this function if any are to participate in UPC++ operations.  
8 While `init` can be called more than once by each process in a program, only the first  
9 invocation will initialize UPC++, and the rest will merely increment the internal count of  
10 how many times `init` has been called. For each `init` call, a matching `finalize` call must  
11 eventually be made. `init` and `finalize` are not re-entrant and must be called by only  
12 a single thread of execution in each process. The thread that calls `init` has the *master*  
13 *persona* attached to it (see section 9.5.1 for more details of threading behavior). After the  
14 number of calls to `finalize` matches the number of calls to `init`, no UPC++ function that  
15 requires the library to be in an initialized state can be invoked until UPC++ is reinitialized  
16 by a subsequent call to `init`.

17 All UPC++ operations require the library to be in an initialized state unless otherwise  
18 specified, and violating this requirement results in undefined behavior. Member functions,  
19 constructors, and destructors are included in the set of operations that require UPC++ to  
20 be initialized, unless explicitly stated otherwise.

### 21 2.2 Hello World

22 A UPC++ installation should be able to compile and execute the simple *Hello World* program  
23 shown in Figure 2.1. The output of *Hello World*, however, is platform-dependent and may  
24 vary between different runs, since there is no synchronization to order the output between  
25 processes. Depending on the nature of the buffering protocol of `stdout`, output from

```
1 #include <upcxx/upcxx.hpp>
2 #include <iostream>
3 int main(int argc, char *argv[])
4 {
5     upcxx::init(); // initialize UPC++
6
7     std::cout << "Hello World"
8         << " ranks:" << upcxx::rank_n() // how many UPC++ ranks?
9         << " my rank: " << upcxx::rank_me() // which rank am I?
10        << std::endl;
11
12     upcxx::finalize(); // finalize UPC++
13     return 0;
14 }
```

Figure 2.1: *HelloWorld.cpp* program

1 different processes may even be interleaved.

## 2 2.3 API Reference

```
3 void init();
```

4 *Preconditions:* Called collectively by all processes in the parallel job. Calling  
5 thread must have the master persona (§9.5.1) if UPC++ is in an already-initialized  
6 state.

7 If there have been no previous calls to `init`, or if all previous calls to `init` have  
8 had matching calls to `finalize`, then this routine initializes the UPC++ library.  
9 Otherwise, leaves the library's state as is. Upon return, the calling thread will  
10 be attached to the master persona (§9.5.1).

11 *This function is legal to call when UPC++ is in the uninitialized state.*

```
12 void finalize();
```

13 *Preconditions:* Called collectively by all processes in the parallel job. Calling  
14 thread must have the master persona (§9.5.1), and UPC++ must be in an already-  
15 initialized state.

1 If this call matches the call to `init` that placed UPC++ in an initialized state,  
2 then this call uninitialized the UPC++ library. Otherwise, this function does not  
3 alter the library's state.

4 If this call uninitialized the UPC++ library while there are any asynchronous  
5 operations still in-flight, behavior is undefined. An operation is defined as in-  
6 flight if it was initiated but still requires internal-level or user-level progress  
7 from any persona on any rank in the job before it can complete. It is left  
8 to the application to define and implement their own specific approach to en-  
9 suring quiescence of in-flight operations. A potential quiescence API is being  
10 considered for future versions and feedback is encouraged.

# 1 Chapter 3

## 2 Global Pointers

### 3 3.1 Overview

4 The UPC++ `global_ptr` is the primary way to address memory in a remote shared memory  
5 segment of a UPC++ program. The next chapter discusses how memory in the shared  
6 segment is allocated to the user.

7 As mentioned in Chapter 1, a global pointer is a handle that may not be dereferenced.  
8 This restriction follows from the design decision to prohibit implicit communication. Logi-  
9 cally, a global pointer has two parts: a raw C++ pointer and an associated *affinity*, which  
10 is a binding of each location in a shared segment to a particular rank (generally the rank  
11 which allocated that shared object). In cases where the use of a `global_ptr` executes in  
12 a rank that has direct load/store access to the memory of the `global_ptr` (i.e. `is_local`  
13 is `true`), we may extract the raw pointer component, and benefit from the reduced cost  
14 of employing a local reference rather than a global one. To this end, UPC++ provides the  
15 `local()` function, which returns a raw C++ pointer. Calling `local()` on a `global_ptr`  
16 that references an address in a remote shared segment results in undefined behavior.

17 Global pointers have the following guarantees:

- 18 1. A `global_ptr<T>` is only valid if it is the null global pointer, it references a valid  
19 object, or it represents one element past the end of a valid array or non-array object.
- 20 2. Two global pointers compare equal if and only if they reference the same object, one  
21 past the end of the same array or non-array object, or are both null.
- 22 3. Equality of global pointers corresponds to observational equality, meaning that two  
23 global pointers which compare equal will produce equivalent behavior when inter-  
24 changed.

25 These facts become important given that UPC++ allows two ranks which are local to  
26 each other to map the same memory into their own virtual address spaces but possibly



1 with different virtual addresses. They also ensure that a global pointer can be viewed from  
 2 any rank to mean the same thing without need for translation.

## 3 3.2 API Reference

```
4 using intrank_t = /* implementation-defined */;
```

5 An implementation-defined signed integer type that represents a UPC++ rank  
 6 ID.

```
7 template<typename T>
8 struct global_ptr;
```

9 C++ Concepts: DefaultConstructible, TriviallyCopyable, TriviallyDestructible,  
 10 EqualityComparable, LessThanComparable, hashable

11 It is illegal for T to have any cv qualifiers: `std::is_const<T>::value` and  
 12 `std::is_volatile<T>::value` must both be false.

```
13 template<typename T>
14 struct global_ptr {
15     using element_type = T;
16     // ...
17 };
```

18 Member type that is an alias for the template parameter T.

```
19 template<typename T>
20 global_ptr<T>::global_ptr(T* ptr);
```

21 *Precondition:* `ptr` must be either null or an address in the shared segment (Ch.  
 22 4) of a rank in the local team (§11.2)

23 Constructs a global pointer corresponding to the given raw pointer. This con-  
 24 structor must be called explicitly.

25 *UPC++ progress level:* none

```
26 template<typename T>
27 global_ptr<T>::global_ptr(std::nullptr_t = nullptr);
```

1 Constructs a global pointer corresponding to a null pointer.

2 *This function is legal to call when UPC++ is in the uninitialized state.*

3 *UPC++ progress level: none*

```
4 template<typename T>  
5 global_ptr<T>::~global_ptr();
```

6 Trivial destructor. Does not delete or otherwise reclaim the raw pointer that  
7 this global pointer is referencing.

8 *This function is legal to call when UPC++ is in the uninitialized state.*

9 *UPC++ progress level: none*

```
10 template<typename T>  
11 bool global_ptr<T>::is_local() const;
```

12 Returns whether or not the calling rank has load/store access to the memory  
13 referenced by this pointer. Returns true if this is a null pointer, regardless of  
14 the context in which this query is called.

15 *UPC++ progress level: none*

```
16 template<typename T>  
17 bool global_ptr<T>::is_null() const;
```

18 Returns whether or not this global pointer corresponds to the null value, mean-  
19 ing that it references no memory. This query is purely a function of the global  
20 pointer instance, it is not affected by the context in which it is called.

21 *UPC++ progress level: none*

```
22 template<typename T>  
23 T* global_ptr<T>::local() const;
```

24 *Precondition: this->is\_local()*

25 Converts this global pointer into a raw pointer.

26 *UPC++ progress level: none*

```

1  template<typename T>
2  intrank_t global_ptr<T>::where() const;

```

3 Returns the rank in team `world()` with affinity to the T object pointed-to by  
4 this global pointer. The return value for `where()` on a null global pointer is  
5 an implementation-defined value. This query is purely a function of the global  
6 pointer instance, it is not affected by the context in which it is called.

7 *UPC++ progress level: none*

```

8  template<typename T>
9  global_ptr<T> global_ptr<T>::operator+(std::ptrdiff_t diff) const;
10 template<typename T>
11 global_ptr<T> operator+(std::ptrdiff_t diff, global_ptr<T> ptr);

```

12 *Precondition:* Either `diff == 0`, or the global pointer is pointing to the `i`th  
13 element of an array of `N` elements, where `i` may be equal to `N`, representing a  
14 one-past-the-end pointer. At least one of the indices `i+diff` or `i+diff-1` must  
15 be a valid element of the same array. A pointer to a non-array object is treated  
16 as a pointer to an array of size 1.

17 If `diff == 0`, returns a copy of the global pointer. Otherwise produces a  
18 pointer that references the element that is at `diff` positions greater than the  
19 current element, or a one-past-the-end pointer if the last element of the array  
20 is at `diff-1` positions greater than the current.

21 These routines are purely functions of their arguments, they are not affected  
22 by the context in which they are called.

23 *UPC++ progress level: none*

```

24 template<typename T>
25 global_ptr<T> global_ptr<T>::operator-(std::ptrdiff_t diff) const;

```

26 *Precondition:* Either `diff == 0`, or the global pointer is pointing to the `i`th  
27 element of an array of `N` elements, where `i` may be equal to `N`, representing a  
28 one-past-the-end pointer. At least one of the indices `i-diff` or `i-diff-1` must  
29 be a valid element of the same array. A pointer to a non-array object is treated  
30 as a pointer to an array of size 1.

31 If `diff == 0`, returns a copy of the global pointer. Otherwise produces a  
32 pointer that references the element that is at `diff` positions less than the

1 current element, or a one-past-the-end pointer if the last element of the array  
2 is at `diff+1` positions less than the current.

3 This routine is purely a function of its arguments, it is not affected by the  
4 context in which they are called.

5 *UPC++ progress level: none*

```
6 template<typename T>  
7 std::ptrdiff_t global_ptr<T>::operator-(global_ptr<T> rhs) const;
```

8 *Precondition:* Either `*this == rhs`, or this global pointer is pointing to the  
9 `i`th element of an array of `N` elements, and `rhs` is pointing at the `j`th element  
10 of the same array. Either pointer may also point one past the end of the array,  
11 so that `i` or `j` is equal to `N`. A pointer to a non-array object is treated as a  
12 pointer to an array of size 1.

13 If `*this == rhs`, results in 0. Otherwise, returns `i-j`.

14 This routine is purely a function of its arguments, it is not affected by the  
15 context in which it is called.

16 *UPC++ progress level: none*

```
17 template<typename T>  
18 global_ptr<T>& global_ptr<T>::operator++();
```

19 *Precondition:* the global pointer must be pointing to an element of an array or  
20 to a non-array object

21 Modifies this pointer to have the value `*this + 1` and returns a reference to  
22 this pointer.

23 This routine is purely a function of its instance, it is not affected by the context  
24 in which it is called.

25 *UPC++ progress level: none*

```
26 template<typename T>  
27 global_ptr<T> global_ptr<T>::operator++(int);
```

1     *Precondition:* the global pointer must be pointing to an element of an array or  
2     to a non-array object

3     Modifies this pointer to have the value `*this + 1` and returns a copy of the  
4     original pointer.

5     This routine is purely a function of its instance, it is not affected by the context  
6     in which it is called.

7     *UPC++ progress level:* none

```
8  template<typename T>  
9  global_ptr<T>& global_ptr<T>::operator--();
```

10     *Precondition:* the global pointer must either be pointing to the `i`th element of  
11     an array, where `i >= 1`, or one element past the end of an array or a non-array  
12     object

13     Modifies this pointer to have the value `*this - 1` and returns a reference to  
14     this pointer.

15     This routine is purely a function of its instance, it is not affected by the context  
16     in which it is called.

17     *UPC++ progress level:* none

```
18 template<typename T>  
19 global_ptr<T> global_ptr<T>::operator--(int);
```

20     *Precondition:* the global pointer must either be pointing to the `i`th element of  
21     an array, where `i >= 1`, or one element past the end of an array or a non-array  
22     object

23     Modifies this pointer to have the value `*this - 1` and returns a copy of the  
24     original pointer.

25     This routine is purely a function of its instance, it is not affected by the context  
26     in which it is called.

27     *UPC++ progress level:* none

```
28 template<typename T>  
29 bool global_ptr<T>::operator==(global_ptr<T> rhs) const;  
30 template<typename T>  
31 bool global_ptr<T>::operator!=(global_ptr<T> rhs) const;
```

```
1  template<typename T>
2  bool global_ptr<T>::operator<(global_ptr<T> rhs) const;
3  template<typename T>
4  bool global_ptr<T>::operator<=(global_ptr<T> rhs) const;
5  template<typename T>
6  bool global_ptr<T>::operator>(global_ptr<T> rhs) const;
7  template<typename T>
8  bool global_ptr<T>::operator>=(global_ptr<T> rhs) const;
```

9 Returns the result of comparing two global pointers. Two global pointers compare equal if they both represent null pointers, or if they represent the same memory address with affinity to the same rank. All other global pointers compare unequal.

13 A pointer to a non-array object is treated as a pointer to an array of size one. If two global pointers point to different elements of the same array, or to subobjects of two different elements of the same array, then the pointer to the element at the higher index compares greater than the pointer to the element at the lower index. If one pointer points to an element of an array or to a subobject of an element of an array, and the other pointer points one past the end of the array, then the latter compares greater than the former.

20 If global pointers  $p$  and  $q$  compare equal, then  $p == q$ ,  $p <= q$ , and  $p >= q$  all result in true while  $p != q$ ,  $p < q$ , and  $p > q$  all result in false. If  $p$  and  $q$  do not compare equal, then  $p != q$  is true while  $p == q$  is false.

23 If  $p$  compares greater than  $q$ , then  $p > q$ ,  $p >= q$ ,  $q < p$ , and  $q <= p$  all result in true while  $p < q$ ,  $p <= q$ ,  $q > p$ , and  $q >= p$  all result in false.

25 All other comparisons result in an unspecified value.

26 These routines are purely functions of their arguments, they are not affected by the context in which they are called.

28 *UPC++ progress level: none*

```
29 namespace std {
30     template<typename T>
31     struct less<global_ptr<T>>;
32     template<typename T>
33     struct less_equal<global_ptr<T>>;
34     template<typename T>
35     struct greater<global_ptr<T>>;
36     template<typename T>
```

```

1  struct greater_equal<global_ptr<T>>;
2  template<typename T>
3  struct hash<global_ptr<T>>;
4  }

```

Specializations of STL function objects for performing comparisons and computing hash values on global pointers. The specializations of `std::less`, `std::less_equal`, `std::greater`, and `std::greater_equal` all produce a strict total order over global pointers, even if the comparison operators do not. This strict total order is consistent with the partial order defined by the comparison operators.

*UPC++ progress level: none*

```

11 template<typename T>
12 std::ostream& operator<<(std::ostream &os, global_ptr<T> ptr);

```

Inserts an implementation-defined character representation of `ptr` into the output stream `os`. This function can be called on any valid global pointer, and the textual representation of two objects of type `global_ptr<T>` is identical if and only if the two global pointers compare equal.

```

17 template<typename T, typename U>
18 global_ptr<T> reinterpret_pointer_cast(global_ptr<U> ptr);

```

*Precondition:* the expression `reinterpret_cast<T*>((U*)nullptr)` must be well formed

Constructs a global pointer whose underlying raw pointer is obtained by using a cast expression on that of `ptr`. The affinity of the result is the same as that of `ptr`.

If `rp` is the raw pointer of `ptr`, then the raw pointer of the result is constructed by `reinterpret_cast<T*>(rp)`.

*UPC++ progress level: none*

# 1 Chapter 4

## 2 Storage Management

### 3 4.1 Overview

4 UPC++ provides two flavors of storage allocation involving the shared segment. The pair  
5 of functions `new_` and `delete_` will call the class constructors and destructors, respectively,  
6 as well as allocate and deallocate memory from the shared segment. The pair `allocate`  
7 and `deallocate` allocate and deallocate dynamic memory from the shared segment, but  
8 do not call C++ constructors or destructors. A user may call these functions directly, or  
9 use placement new, or other memory management practices.

### 10 4.2 API Reference

```
11 template<typename T, typename ...Args>  
12 global_ptr<T> new_(Args &&...args);
```

13 *Precondition:* `T(args...)` must be a valid call to a constructor for `T`.

14 Allocates space for an object of type `T` from the shared segment of the current  
15 rank. If the allocation succeeds, returns a pointer to the start of the allocated  
16 memory, and the object is initialized by invoking the constructor `T(args...)`.  
17 If the allocation fails, throws `std::bad_alloc`.

18 *Exceptions:* May throw `std::bad_alloc` or any exception thrown by the call  
19 `T(args...)`.

20 *UPC++ progress level:* none

```
21 template<typename T, typename ...Args>  
22 global_ptr<T> new_(const std::nothrow_t &tag, Args &&...args);
```



1     *Precondition:* `T(args...)` must be a valid call to a constructor for `T`.  
 2     Allocates space for an object of type `T` from the shared segment of the current  
 3     rank. If the allocation succeeds, returns a pointer to the start of the allocated  
 4     memory, and the object is initialized by invoking the constructor `T(args...)`.  
 5     If the allocation fails, returns a null pointer.  
 6     *Exceptions:* May throw any exception thrown by the call `T(args...)`.  
 7     *UPC++ progress level:* none

```
8  template<typename T>
9  global_ptr<T> new_array(size_t n);
```

10     *Precondition:* `T` must be `DefaultConstructible`.  
 11     Allocates space for an array of `n` objects of type `T` from the shared segment of  
 12     the current rank. If the allocation succeeds, returns a pointer to the start of  
 13     the allocated memory, and the objects are initialized by invoking their default  
 14     constructors. If the allocation fails, throws `std::bad_alloc`.  
 15     *Exceptions:* May throw `std::bad_alloc` or any exception thrown by the call  
 16     `T()`. If an exception is thrown by the constructor for `T`, then previously initial-  
 17     ized elements are destroyed in reverse order of construction.  
 18     *UPC++ progress level:* none

```
19 template<typename T>
20 global_ptr<T> new_array(size_t n, const std::nothrow_t &tag);
```

21     *Precondition:* `T` must be `DefaultConstructible`.  
 22     Allocates space for an array of `n` objects of type `T` from the shared segment of  
 23     the current rank. If the allocation succeeds, returns a pointer to the start of  
 24     the allocated memory, and the objects are initialized by invoking their default  
 25     constructors. If the allocation fails, returns a null pointer.  
 26     *Exceptions:* May throw any exception thrown by the call `T()`. If an exception  
 27     is thrown by the constructor for `T`, then previously initialized elements are  
 28     destroyed in reverse order of construction.  
 29     *UPC++ progress level:* none

```
30 template<typename T>
31 void delete_(global_ptr<T> g);
```

1     *Precondition:* T must be Destructible. g must be a non-deallocated pointer  
2     that resulted from a call to `new_<T, Args...>` on the current rank, for some  
3     value of `Args...`

4     Invokes the destructor on the given object and deallocates the storage allocated  
5     to it.

6     *Exceptions:* May throw any exception thrown by the the destructor for T.

7     *UPC++ progress level:* none

```
8  template<typename T>  
9  void delete_array(global_ptr<T> g);
```

10    *Precondition:* T must be Destructible. g must be a non-deallocated pointer  
11    that resulted from a call to `new_array<T>` on the current rank.

12    Invokes the destructor on each object in the given array and deallocates the  
13    storage allocated to it.

14    *Exceptions:* May throw any exception thrown by the the destructor for T.

15    *UPC++ progress level:* none

```
16 void* allocate(size_t size,  
17                size_t alignment = alignof(std::max_align_t));
```

18    *Precondition:* `alignment` is a valid alignment. `size` must be an integral mul-  
19    tiple of `alignment`.

20    Allocates `size` bytes of memory from the shared segment of the current rank,  
21    with alignment as specified by `alignment`. If the allocation succeeds, returns  
22    a pointer to the start of the allocated memory, and the allocated memory is  
23    uninitialized. If the allocation fails, returns a null pointer.

24    *UPC++ progress level:* none

```
25 template<typename T, size_t alignment = alignof(T)>  
26 global_ptr<T> allocate(size_t n=1);
```

27    *Precondition:* `alignment` is a valid alignment.

28    Allocates enough space for `n` objects of type T from the shared segment of  
29    the current rank, with the memory aligned as specified by `alignment`. If the  
30    allocation succeeds, returns a pointer to the start of the allocated memory, and

1 the allocated memory is uninitialized. If the allocation fails, returns a null  
2 pointer.

3 *UPC++ progress level: none*

4 `void deallocate(void* p);`

5 *Precondition:* `p` must be either a null pointer or a non-deallocated pointer that  
6 resulted from a call to the first form of `allocate` on the current rank.

7 Deallocates the storage previously allocated by a call to `allocate`. Does noth-  
8 ing if `p` is a null pointer.

9 *UPC++ progress level: none*

10 `template<typename T>`

11 `void deallocate(global_ptr<T> g);`

12 *Precondition:* `g` must be either a null pointer or a non-deallocated pointer that  
13 resulted from a call to `allocate<T, alignment>` on the current rank, for some  
14 value of `alignment`.

15 Deallocates the storage previously allocated by a call to `allocate`. Does noth-  
16 ing if `g` is a null pointer. Does not invoke the destructor for `T`.

17 *UPC++ progress level: none*

# Chapter 5

## Futures and Promises

### 5.1 Overview

In UPC++, the primary mechanisms by which a programmer interacts with non-blocking operations are futures and promises.<sup>1</sup> These two mechanisms, usually bound together under the umbrella concept of *futures*, are present in the C++11 standard. However, while we borrow some of the high-level concepts of C++'s futures, many of the semantics of `upcxx::future` and `upcxx::promise` differ from those of `std::future` and `std::promise`. In particular, while futures in C++ facilitate communicating between threads, the intent of UPC++ futures is solely to provide an interface for managing and composing non-blocking operations, and they cannot be used directly to communicate between threads or ranks.

A non-blocking operation is associated with a state that encapsulates both the status of the operation as well as any result values. Each such operation has an associated *promise* object, which can either be explicitly created by the user or implicitly by the runtime when a non-blocking operation is invoked. A promise represents the producer side of the operation, and it is through the promise that the results of the operation are supplied and its dependencies fulfilled. A *future* is the interface through which the status of the operation can be queried and the results retrieved, and multiple future objects may be associated with the same promise. A future thus represents the consumer side of a non-blocking operation.

### 5.2 The Basics of Asynchronous Communication

A programmer can invoke a non-blocking operation to be serviced by another rank, such as a one-sided get operation (Ch. 7) or a remote procedure call (Ch. 8). Such an operation

---

<sup>1</sup>Another mechanism, persona-targeted continuations, is discussed in §9.4.

1 creates an implicit promise and returns an associated future object to the user. When the  
 2 operation completes, the future becomes ready, and it can be used to access the results.  
 3 The following demonstrates an example using a remote get (see Ch. 9 on how to make  
 4 progress with UPC++):

```

5 global_ptr<double> ptr = /* obtain some remote pointer */;
6 future<double> fut = rget(ptr);           // initiate a remote get
7 // ...call into upcxx::progress() elided...
8 if (fut.ready()) {                       // check for readiness
9     double value = fut.result();         // retrieve result
10    std::cout << "got: " << value << '\n'; // use result
11 }

```

12 In general, a non-blocking operation will not complete immediately, so if a user needs  
 13 to wait on the readiness of a future, they must do so in a loop. To facilitate this, we  
 14 provide the `wait` member function, which waits on a future to complete while ensuring  
 15 that sufficient progress (Ch. 9) is made on internal and user-level state:

```

16 global_ptr<double> ptr = /* obtain some remote pointer */;
17 future<double> fut = rget(ptr);           // initiate a remote get
18 fut.wait();                               // wait for completion
19 double value = fut.result();             // retrieve result
20 std::cout << "got: " << value << '\n';   // use result

```

21 An alternative to waiting for completion of a future is to attach a *callback* or *completion*  
 22 *handler* to the future, to be executed when the future completes. This callback can be  
 23 any function object, including lambda (anonymous) functions, that can be called on the  
 24 results of the future, and is attached using `then`.

```

25 global_ptr<double> ptr = /* obtain some remote pointer */;
26 auto fut =
27 rget(ptr).then( // initiate a remote get and register a callback
28 // lambda callback function
29 [](double value) {
30     std::cout << "got: " << value << '\n'; // use result
31 }
32 );

```

33 The return value of `then` is another future representing the results of the callback, if  
 34 any. This permits the specification of a sequence of operations, each of which depends on  
 35 the results of the previous one.

36 A future can also represent the completion of a combination of several non-blocking  
 37 operations. Unlike the standard C++ future, `upcxx::future` is a variadic template, encapsulating  
 38 an arbitrary number of result values that can come from different operations. The  
 39 following example constructs a future that represents the results of two existing futures:

```
1 future<double> fut1 = /* one future */;  
2 future<int> fut2 = /* another future */;  
3 future<double, int> combined = when_all(fut1, fut2);
```

4 Here, `combined` represents the state and results of two futures, and it will be ready  
5 when both `fut1` and `fut2` are ready. The results of `combined` are a `std::tuple` whose  
6 components are the results of the source futures.

## 7 5.3 Working with Promises

8 In addition to the implicit promises created by non-blocking operations, a user may explic-  
9 itly create a promise object, obtain associated future objects, and then register non-blocking  
10 operations on the promise. This is useful in several cases, such as when a future is required  
11 before a non-blocking operation can be initiated, or where a single promise is used to count  
12 dependencies.

13 A promise can also be used to count *anonymous dependencies*, keeping track of opera-  
14 tions that complete without producing a value. Upon creation, a promise has a dependency  
15 count of one, representing the unfulfilled results or, if there are none, an anonymous de-  
16 pendency. Further anonymous dependencies can then be registered on the promise. When  
17 registration is complete, the original dependency can then be fulfilled to signal the end of  
18 registration. The following example keeps track of several remote put operations with a  
19 single promise:

```
20 global_ptr<int> ptrs[10] = /* some remote pointers */;  
21 // create a promise with no results  
22 // the dependency count starts at one  
23 promise<> prom;  
24  
25 // do 10 puts, registering each of them on the promise  
26 for (int i = 0; i < 10; i++) {  
27     // rput implicitly registers itself on the given promise  
28     rput(ptrs[i], prom);  
29 }  
30  
31 // fulfill initial anonymous dependency, since registration is done  
32 prom.finalize_anonymous();  
33  
34 // wait for the rput operations to complete  
35 prom.get_future().wait();
```

## 5.4 Advanced Callbacks

Polling for completion of a future allows simple overlap of communication and computation operations. However, it introduces the need for synchronization, and this requirement can diminish the benefits of overlap. To this end, many programs can benefit from the use of callbacks. Callbacks avoid the need for an explicit wait and enable reactive control flow: future completion triggers a callback. Callbacks allow operations to occur as soon as they are capable of executing, rather than artificially waiting for an unrelated operation to complete before being initiated.

Futures are the core abstraction for obtaining asynchronous results, and an API that supports asynchronous behavior can work with futures rather than values directly. Such an API can also work with immediately available values by having the caller wrap the values into a ready future using the `make_future` function template, as in this example that creates a future for an ordered pair of a `double` and an `int`:

```
14 void consume(future<int, double> fut);
15 consume(make_future(3, 4.1));
```

Given a future, we can attach a callback to be executed at some subsequent point when the future is ready using the `then` member function:

```
18 future<int, double> source = /* obtain a future */;
19 future<double> result = source.then(
20     [](int x, double y) {
21         return x + y;
22     }
23 );
```

In this example, `source` is a future representing an `int` and a `double` value. The argument of the call to `then` must be a function object that can be called on these values. Here, we use a lambda function that takes in an `int` and a `double`. The call to `then` returns a future that represents the result of calling the argument of `then` on the values contained in `source`. Since the lambda function above returns a `double`, the result of `then` is a `future<double>` that will hold the `double`'s value when it is ready.

There is also another case, when the callback returns a future, rather than some non-future type. In previous case, the result of `then()` is obtained by wrapping return type inside a future. In this case, this step is not needed, as we are already returning a future. Thus, the result of the call to `then` has the same type as the return type of the callback. However, there is an important difference: the result is a future, which may or may not be ready. In the first case, it is the returned non-future value that may or may not be ready. This subtle difference, allows the UPC++ programmer to chain the results of one asynchronous operation into the inputs of the next, to arbitrary degree of nesting.

```
38 future<int, double> source = /* obtain a future */;
```

```
1 future<double> result = source.then(  
2     [](int x, double y) {  
3         // return a future<double> that is ready  
4         return make_future(x + y);  
5     }  
6 );  
7 // result may not be ready, since the callback will not be executed  
8 // until source is ready
```

9 A callback may also initiate new asynchronous work and return a future representing  
10 the completion of that work:

```
11 global_ptr<int> remote_array = /* some remote array */;  
12  
13 // retrieve remote_array[0]  
14 future<int> elt0 = rget(remote_array);  
15  
16 // retrieve remote_array[remote_array[0]]  
17 future<int> elt_indirect = elt0.then(  
18     [=](int index) {  
19         return rget(remote_array + index);  
20     }  
21 );
```

22 The `then` member function is a combinator for constructing pipelines of transformations  
23 over futures. Given a future and a function that transforms that future's value into another  
24 value, `then` produces a future representing the post-transformation value. For example,  
25 we can future transform the value of `elt_indirect` above as follows:

```
26 future<int> elt_indirect_squared = elt_indirect.then(  
27     [](int value) {  
28         return value * value;  
29     }  
30 );
```

31 As the examples above demonstrate, the `then` member function allows a callback to  
32 depend on the result of another future. A more general pattern is for an operation to  
33 depend on the results of multiple futures. The `when_all` function template enables this  
34 by constructing a single future that combines the results of multiple futures:

```
35 future<int> value1 = /* ... */;  
36 future<double> value2 = /* ... */;  
37  
38 future<int, double> combined = when_all(value1, value2);  
39 future<double> result = combined.then(  
40     [](int i, double d) {  
41         return d * i;  
42     }  
43 );
```



```

1  [](int x, double y) {
2      return x + y;
3  }
4  );

```

A callback (made via `then`) can depend on multiple futures. We register the callback with a combined future, constructed with `when_all`. The `when_all` is restricted to combining lists of futures only. In the more general case, we may need to combine heterogeneous mixtures of future and non-future types. The `to_future` function template provides a further generalization, combining values from futures as well as raw (non-future) values themselves. While `when_all` can be used to meet this need (by wrapping raw values in calls to `make_future`), a call to `to_future` does this automatically:

```

12 future<int> value1 = /* ... */;
13 double value2 = /* ... */;
14
15 future<int, double> combined = to_future(value1, value2);
16 future<double> result = combined.then(
17     [](int x, double y) {
18         return x + y;
19     }
20 );

```

The results of a future can be obtained, if it is ready, as a `std::tuple` using the `result_tuple` member function of a future. Individual components can be retrieved by value with the `result` member function template or by r-value reference with `result_moved`. Unlike with `std::get`, it is not a compile-time error to use an invalid index with `result` or `result_moved`; instead, the return type is `void` for an invalid index. This simplifies writing generic functions on futures, such as the following C++14-compliant definition of `wait`:

```

28 template<typename ...T>
29 auto future<T...>::wait() {
30     while (!ready()) {
31         progress();
32     }
33     return result();
34 }

```

## 5.5 Execution Model

Futures have the capability to express dataflow/task-based programming, and other software frameworks provide thread-level parallelism by considering each callback to be a task

1 that can be run in an arbitrary worker thread. This is not the case in UPC++. In order  
2 to maximize performance, our approach to futures is purposefully ambivalent to issues of  
3 concurrency. A UPC++ implementation is allowed to take action as if the current thread is  
4 the only one that needs to be accounted for. This gives rise to a natural execution policy:  
5 callbacks registered against futures are always executed as soon as possible by the thread  
6 that discovers them. There are exactly two scenarios in which this may happen:

- 7 1. When a promise is fulfilled.
- 8 2. A callback is registered onto a ready future using the `then` member function.

9 Fulfilling a promise (via `fulfill_result`, `fulfill_anonymous` or `finalize_anonymous`)  
10 is the only operation that can take a future from a non-ready to a ready state, enabling  
11 callbacks that depend on it to execute. This makes promise fulfillment an obvious place  
12 for discovering and executing such callbacks. Thus, whenever a thread calls a fulfillment  
13 function on a promise, the user must anticipate that any newly available callbacks will be  
14 executed by the current thread before the fulfillment call returns.

15 The other place in which a callback will execute immediately is during the invocation  
16 of `then` on a future that is already in its ready state. In this case, the callback provided  
17 will fire immediately during the call to `then`.

18 There are some common programming contexts where it is not safe for a callback to  
19 execute during fulfillment of a promise. For example, it is generally unsafe to execute a  
20 callback that modifies a data structure while a thread is traversing the data structure. In  
21 such a situation, it is the user's responsibility to ensure that a conflicting callback will not  
22 execute. One solution is create a promise that represents a thread reaching its *safe-to-*  
23 *execute* context, and then adding it to the dependency list of any conflicting callback.

```
24 future<int> value = /* ... */;  
25 // create a promise representing a safe-to-execute state  
26 // dependency count is initially 1  
27 promise<> safe_state;  
28 // create a future that depends on both value and safe_state  
29 future<int> combined = when_all(value, safe_state.get_future());  
30 auto fut = // register a callback on the combined future  
31 combined.then(/* some callback that requires a safe state */);  
32 // do some work, potentially fulfilling value's promise...  
33 // signify a safe state  
34 safe_state.finalize_anonymous();  
35 // callback can now execute
```

36 As demonstrated above, the user can wait to fulfill the promise until it is safe to execute  
37 the callback, which will then allow it to execute.

## 5.6 Anonymous Dependencies

As demonstrated previously, promises can be used to both supply values as well as signal completion of events that do not produce a value. As such, a promise is a unified abstraction for tracking the completion of asynchronous operations, whether the operations produce a value or not. A promise represents at most one dependency that produces a value, but it can track any number of anonymous dependencies that do not result in a value.

When created, a promise starts with an initial dependency count of 1. For an empty promise (`promise<>`), this is necessarily an anonymous dependency, since an empty promise does not hold a value. For a non-empty promise, the initial count represents the sole dependency that produces a value. Further anonymous dependencies can be explicitly registered on a promise with the `require_anonymous` member function:

```
promise<int, double> pro; // initial dependency count is 1
pro.require_anonymous(10); // dependency count is now 11
```

The argument to `require_anonymous` must be strictly greater than the negation of the promise's dependency count, so that a call to `require_anonymous` never causes the dependency count to reach zero, putting the promise in the fulfilled state. In the example above, the argument must be greater than -1, and the given argument of 10 is valid.

Anonymous dependencies can be fulfilled by calling the `fulfill_anonymous` member function:

```
for (int i = 0; i < 5; i++) {
    pro.fulfill_anonymous(i);
} // dependency count is now 1
```

A non-anonymous dependency is fulfilled by calling `fulfill_result` with the produced values:

```
pro.fulfill_result(3, 4.1); // dependency count is now 0
assert(pro.get_future().ready());
```

While both empty and non-empty promises can be used to track anonymous dependencies, an empty promise is *only* able to track anonymous dependencies, so we expect that they will be the primary mechanism used to do so. As such, UPC++ operations that operate on promises make an important distinction between empty and non-empty promises: applying a UPC++ operation to an empty promise *does* increment its dependency count, calling a UPC++ operation on a non-empty promise *does not* increment its dependency count. The rationale for this is that an operation on a non-empty promise can only fulfill its initial, value-representing dependency, while an operation on an empty promise always fulfills an anonymous dependency. Rather than having the user manually increment the dependency count before calling an operation on an empty promise, UPC++ will implicitly perform this increment. This leads to the pattern, shown at the beginning of this chapter,

1 of registering operations on an empty promise and then finalizing the promise to take it  
2 out of registration mode:

```
3 global_ptr<int> ptrs[10] = /* some remote pointers */;  
4 promise<> prom; // dependency count is 1  
5  
6 for (int i = 0; i < 10; i++) {  
7     rput(ptrs[i], prom); // dependency count is incremented  
8 } // dependency count is now 11  
9  
10 prom.finalize_anonymous(); // decrement count, making it 10  
11  
12 // wait for the 10 rput operations to complete  
13 prom.get_future().wait();
```

14 A user familiar with UPC++ V0.1 will observe that empty promises subsume the ca-  
15 pabilities of events in UPC++ V0.1. In addition, they can take part in all the machinery  
16 of promises, futures, and callbacks, providing a much richer set of capabilities than were  
17 available in V0.1.

## 18 5.7 Lifetime and Thread Safety

19 Understanding the lifetime of objects in the presence of asynchronous control flow can be  
20 tricky. Objects must outlive the last callback that references them, which in general does  
21 not follow the scoped lifetimes of the call stack. For this reason, UPC++ automatically man-  
22 ages the state represented by futures and promises, and the state persists for as long as  
23 there is a future, promise, or dependent callback that references it. Thus, a user can con-  
24 struct intricate webs of callbacks over futures without worrying about explicitly managing  
25 the state representing the callbacks' dependencies or results.

26 Though UPC++ does not prescribe a specific management strategy, the semantics of  
27 futures and promises are analogous to those of standard C++11 smart pointers. As with  
28 `std::shared_ptr`, a future may be freely copied, and both the original and the copy  
29 represent the same state and are associated with the same promise. Thus, if one copy  
30 of a future becomes ready, then so will the other copies. On the other hand, a promise  
31 can be mutated by the user through its member functions, so allowing a promise to be  
32 copied would introduce the issue of aliasing. Instead, we adopt the same non-copyable, yet  
33 movable, semantics for a promise as `std::unique_ptr`.

34 Given that UPC++ futures and promises are already thread-unaware to allow the ex-  
35 ecution strategy to be straightforward and efficient, UPC++ also makes no thread safety  
36 guarantees about internal state management. This enables creation of copies of a future  
37 to be a very cheap operation. For example, a future can be captured by value by a lambda

1 function or passed by value without any performance penalties. On the other hand, the  
2 lack of thread safety means that sharing a future between threads must be handled with  
3 great caution. Even a simple operation such as making a copy of a future, as when passing  
4 it by value to a function, is unsafe if another thread is concurrently accessing an identical  
5 future, since the act of copying it can modify the internal management state. Thus, a  
6 mutex or other synchronization is required to ensure exclusive access to a future when  
7 performing any operation on it.

8 Fulfilling a promise gives rise to an even more stringent demand, since it can set off a  
9 cascade of callback execution. Before fulfilling a promise, the user must ensure that the  
10 thread has the exclusive right to mutate not just the future associated with the promise,  
11 but all other futures that are directly or indirectly dependent on fulfillment of the promise.  
12 Thus, when crafting their code, the user must properly manage exclusivity for *islands* of  
13 disjoint futures. We say that two futures are in *disjoint islands* if there is no dependency,  
14 direct or indirect, between them.

15 A reader having previous experience with futures will note that UPC++'s formulation is  
16 a significant departure from many other software packages. Futures are commonly used  
17 to pass data between threads, like a channel that a producing thread can supply a value  
18 into, notifying a consuming thread of its availability. UPC++, however, is intended for  
19 high-performance computing, and supporting concurrently shareable futures would require  
20 synchronization that would significantly degrade performance. As such, futures in UPC++  
21 are not intended to *directly* facilitate communication between threads. Rather, they are  
22 designed for a single thread to manage the non-determinism of reacting to the events  
23 delivered by concurrently executing agents, be they other threads or the network hardware.

## 24 5.8 API Reference

25 *UPC++ progress level for all functions in this chapter is: none*

### 26 5.8.1 future

```
27 template<typename ...T>  
28 class future;
```

29 C++ Concepts: DefaultConstructible, CopyConstructible, CopyAssignable,  
30 Destructible

31 It is illegal for any type in T to be void.

```
32 template<typename ...T>  
33 future<T...>::future();
```

1 Constructs a future that will never become ready.

2 *This function is legal to call when UPC++ is in the uninitialized state.*

```
3 template<typename ...T>  
4 future<T...>::~~future();
```

5 Destroys this future object.

6 *This function is legal to call when UPC++ is in the uninitialized state.*

```
7 template<typename ...T>  
8 future<T...> make_future(T ...results);
```

9 Constructs a trivially ready future from the given values.

```
10 template<typename ...T>  
11 bool future<T...>::ready() const;
```

12 Returns true if the future's result values have been supplied to it.

```
13 template<typename ...T>  
14 std::tuple<T...> const& future<T...>::result_tuple() const;
```

15 *Precondition:* `this->ready()`

16 Retrieves the tuple of result values for this future.

```
17 template<typename ...T>  
18 template<int I=0>  
19 future_element_t<I, future<T...>>  
20 future<T...>::result() const;
```

21 *Precondition:* `this->ready()`

22 Retrieves the  $I^{\text{th}}$  component (defaults to first) from the future's results tuple.  
23 The return type is `void` if  $I$  is an invalid index. Otherwise it is of type  $U$ , where  
24  $U$  is the  $I^{\text{th}}$  component of  $T$ .

```

1  template<typename ...T>
2  template<int I=0>
3  future_element_moved_t<I, future<T...>>
4  future<T...>::result_moved();

```

5     *Precondition:* `this->ready()`

6     Retrieves the  $I^{\text{th}}$  component (defaults to first) from the future's results tuple as  
7     an r-value reference, as if by calling `std::move` on the component. The return  
8     type is `void` if  $I$  is an invalid index. Otherwise it is of type `U&&`, where  $U$  is  
9     the  $I^{\text{th}}$  component of  $T$ . *Caution: this operation permits mutation of the value,*  
10     *via an rvalue reference which could be observed by further calls that return the*  
11     *result(s) of a future.*

```

12 template<typename ...T>
13 template<typename Func>
14 future_invoke_result_t<Func, T...>
15 future<T...>::then(Func func);

```

16     *Preconditions:* The call `func()` must not throw an exception.

17     Returns a new future representing the return value of the given function object  
18     `func` when invoked on the results of this future as its argument list. If `func`  
19     returns a future, then the result of `then` will be a semantically equivalent future,  
20     except that it will be in a non-ready state before `func` executes. If `func` does  
21     not return a future, then the return value of `then` is a future that encapsulates  
22     the result of `func`, and this future will also be in a non-ready state before `func`  
23     executes. If the return type of `func` is `void`, then the return type of `then` is  
24     `future<>`.

25     The function object will be invoked in one of two situations:

- 26     • Immediately before `then` returns if this future is in the ready state.
- 27     • During a promise fulfillment which would directly or indirectly make this  
28     future transition to the ready state.

```

29 template<typename ...T>
30 future_element_t<0, future <T...>> future<T...>::wait();

```

31     Waits for the future by repeatedly attempting UPC++ user-level progress and  
32     testing for readiness. See Ch. 9 for a discussion of progress. The return value  
33     is the same as that produced by calling `result()` on the future.

```
1 template<typename ...Futures>  
2 future<CTypes...> when_all(Futures ...fs);
```

3       Given a variadic list of futures as arguments, constructs a future representing  
4       the readiness of all arguments. The results tuple of this future will be the  
5       concatenated results tuples of the arguments. The type parameters of the  
6       returned object (`CTypes...`) is the ordered concatenation of the type parameter  
7       lists of the types in `Futures`.

```
8 template<typename ...T>  
9 future<CTypes...> to_future(T ...futures_or_results);
```

10       Given a variadic list of futures and/or non-futures as arguments, constructs a  
11       future representing the readiness of all the arguments that are futures. The  
12       results tuple of this future will be the concatenation of the result tuples of each  
13       future argument and the values of each non-future argument, in the order in  
14       which each argument occurs in `futures_or_results`. The type parameters of  
15       the returned object (`CTypes...`) is the concatenation of the type parameter  
16       lists of the future types in `T` and the non-future types themselves in `T`, in the  
17       order in which each type appears in `T`.

18       If none of the arguments are futures, then the resulting future object is trivially  
19       ready.

## 20 5.8.2 promise

```
21 template<typename ...T>  
22 class promise;
```

23       C++ Concepts: `DefaultConstructible`, `MoveConstructible`, `MoveAssignable`,  
24       `Destructible`

25       It is illegal for any type in `T` to be `void`.

```
26 template<typename ...T>  
27 promise<T...>::promise();
```

28       Constructs a promise with its results uninitialized and an initial dependency  
29       count of 1.

30       *This function is legal to call when UPC++ is in the uninitialized state.*



```
1  template<typename ...T>
2  promise<T...>::~~promise();
```

3 Destroys this promise object.

4 *This function is legal to call when UPC++ is in the uninitialized state.*

```
5  template<typename ...T>
6  void promise<T...>::require_anonymous(std::intptr_t count);
```

7 *Precondition:* The dependency count of this promise is greater than (-count)  
8 and greater than 0.

9 Adds `count` to this promise's dependency count.

```
10 template<typename ...T>
11 template<typename ...U>
12 void promise<T...>::fulfill_result(U &&...results);
```

13 *Precondition:* `fulfill_result` has not been called on this promise before, and  
14 the dependency count of this promise is greater than zero.

15 Initializes the promise's result tuple with the given values and decrements the  
16 dependency counter by 1. Requires that `T` and `U` have the same number of  
17 components, and that each component of `U` is implicitly convertible to the  
18 corresponding component of `T`. If the dependency counter reaches zero as a  
19 result of this call, the associated future is set to ready, and callbacks that are  
20 waiting on the future are executed on the calling thread before this function  
21 returns.

```
22 template<typename ...T>
23 void promise<T...>::fulfill_anonymous(std::intptr_t count);
```

24 *Precondition:* The dependency count of this promise is greater than or equal  
25 to `count`. If the dependency count is equal to `count` and `T` is not empty, then  
26 the results of this promise must have been previously supplied by a call to  
27 `fulfill_result`.

28 Subtracts `count` from the dependency counter. If this produces a zero counter  
29 value, the associated future is set to ready, and callbacks that are waiting on  
30 the future are executed on the calling thread before this function returns.

```
1 template<typename ...T>  
2 void promise<T...>::finalize_anonymous();
```

3       Equivalent to `this->fulfill_anonymous(1)`.

```
4 template<typename ...T>  
5 future<T...> promise<T...>::get_future() const;
```

6       Returns the future representing this promise being fulfilled. Repeated calls to  
7 `get_future` return equivalent futures with the guarantee that no additional  
8 memory allocation is performed.

# 1 Chapter 6

## 2 Serialization

3 As a communication library, UPC++ needs to send C++ types between ranks that might  
4 be separated by a network interface. The underlying GASNet networking interface sends  
5 and receives bytes, thus, UPC++ needs to be able to convert C++ types to and from bytes.

6 For standard TriviallyCopyable data types, UPC++ can serialize and deserialize these  
7 objects for the user without extra intervention on their part. For user data types that have  
8 more involved serialization requirements, the user needs to take two steps to inform UPC++  
9 about how to serialize the object.

- 10 1. Declare their type to be a friend of `access`
- 11 2. Implement the visitor function `serialize`

12 Figure 6.1 provides an example of this process. The definition of the `&` operator for the  
13 `Archive` class depends on whether UPC++ is serializing or deserializing an object instance.

14 UPC++ provides implementations of `operator&` for the C++ built-in types. UPC++ se-  
15 rialization is compatible with a subset of the Boost serialization interface. This does  
16 not imply that UPC++ includes or requires Boost as a dependency. The reference im-  
17 plementation of UPC++ does neither of these, it comes with its own implementation of  
18 serialization that simply adheres to the interface set by Boost. It is acceptable to have  
19 `friend boost::serialization::access` in place of `friend upcxx::access`. UPC++ will  
20 use your Boost serialization in that case.

21 There are restrictions on which actions serialization/deserialization routines may per-  
22 form. They are:

- 23 1. Serialization/deserialization may not call any UPC++ routine with a progress level  
24 other than `none`.
- 25 2. UPC++ must perceive these routines as referentially transparent. Loosely, this means  
26 that the routines should be “pure” functions between the native representation and  
27 a flat sequence of bytes.

```
15 class UserType {
16     // The user's fields and member declarations as usual.
17     int member1, member2;
18     // ...
19
20     // To enable the serializer to visit the member fields,
21     // the user provides this...
22     friend class upcxx::access;
23
24     // ...and this
25     template<typename Archive>
26     void serialize(Archive &ar, unsigned) {
27         ar & this->member1;
28         ar & this->member2;
29         // ...
30     }
31 };
```

Figure 6.1: An example of using `access` in a user-defined class

3. The routines must be thread-safe and permit concurrent invocation from multiple threads, even when serializing the same object.

## 6.1 Functions

In §7.2 (*Completions*) and Chapter 8 (*Remote Procedure Calls*) there are several cases where a C++ *FunctionObject* is expected to execute on a destination rank. In these cases the function arguments are serialized as described in this chapter. The *FunctionObject* itself is converted to a function pointer offset from a known *sentinel* in the source program's *code segment*. The details of the implementation are not described here but typical allowed *FunctionObjects* are

- C functions
- C++ global and file-scope functions
- Class static functions
- lambda functions

Calling member functions on remote objects requires additional steps described in Chapter 12 (*Distributed Objects*).

# 1 Chapter 7

## 2 One-Sided Communication

### 3 7.1 Overview

4 The main one-sided communication functions for UPC++ are `rput` and `rget`. Where possi-  
5 ble, the underlying transport layer will use RDMA techniques to provide the lowest-latency  
6 transport possible. The type `T` used by `rput` or `rget` needs to be **Serializable**, either in the  
7 sense of C++ `TriviallyCopyable` or by overriding the global `upcxx::serialize` function  
8 as described in Chapter 6 (*Serialization*).

### 9 7.2 Completion

10 Memory movement operations come with the concept of completion, meaning that the  
11 effect of the operation is now visible and that resources, such as memory on the source  
12 and destination sides, are no longer in use by UPC++. The user has choices in how they  
13 would like UPC++ to notify the application of completion: these are by future, promise, or  
14 continuation. Notification by future and promise was introduced in Ch. 5. Continuation  
15 style completion is explained in Ch. 9. An important aspect to clarify is that notification  
16 of completion only happens during user-level progress. Even if an operation completes  
17 early, including before the initiation operation returns, the application cannot learn this  
18 fact without entering user-progress. For futures and promises, only when the initiating  
19 thread (persona actually) enters user-level progress will the future or promise be eligible  
20 for taking on a readied or fulfilled state. Continuations will execute once a thread enters  
21 user-progress of the designated persona. See Ch. 9 for the full discussion on user-progress  
22 and personas.

## 7.3 API Reference

### 7.3.1 Remote Puts

```
3  template<typename T>
4  future<> rput(T value, global_ptr<T> dest);
5
6  template<typename T>
7  void rput(T value, global_ptr<T> dest, promise<> &completion);
8
9  template<typename T>
10 void rput(T value, global_ptr<T> dest,
11           persona &completion_recipient,
12           CompletionFunc completion_func);
```

*Precondition:* T must be Serializable. `dest` must reference a valid object of type T. In the second variant, `completion` must have a dependency count greater than zero. In the third variant, `CompletionFunc` must be a function-object type accepting no arguments, and the call `completion_func()` must not throw an exception.

Either serializes `value` immediately or copies it into an internal location for eventual serialization. After serialization, initiates a transfer of the data which will deserialize and store it in the memory referenced by `dest`.

Completion of the operation indicates that all aspects of the operation: serialization, deserialization, the remote store, and destruction of any internally managed T values are complete.

In the first variant, returns a future representing the completion of the operation.

In the second variant, the promise has its dependency count incremented immediately and fulfilled upon completion of the operation.

In the third variant, `completion_func` is enlisted in the given persona's user-progress upon completion of the operation (see §9.5.1 and Ch. 9).

*C++ memory ordering:* The writes to `dest` will have a *happens-before* relationship with the completion notification action (future readying, promise fulfillment, or persona continuation enlistment). In the third variant, all evaluations *sequenced-before* this call will have a *happens-before* relationship with the execution of `completion_func`.

*UPC++ progress level:* `internal`

```

1  template<typename T>
2  future<> rput(T const *src, global_ptr<T> dest, std::size_t count);
3
4  template<typename T>
5  void rput(T const *src, global_ptr<T> dest, std::size_t count
6           promise<> &completion);
7
8  template<typename T, typename Func>
9  void rput(T const *src, global_ptr<T> dest, std::size_t count,
10          persona &completion_recipient,
11          CompletionFunc completion_func);

```

12 *Precondition:* T must be Serializable. Addresses in the intervals [src,src+count)  
13 and [dest,dest+count) must all reference valid objects of type T. No object  
14 may be referenced by both intervals. In the second variant, the completion  
15 promise must have a dependency count greater than zero. In the third variant,  
16 CompletionFunc must be a function-object type accepting no arguments, and  
17 the call completion\_func() must not throw an exception.

18 Initiates an operation to serialize, transfer, deserialize, and store the count  
19 items of type T beginning at src to the memory beginning at dest.

20 Completion of this operation indicates that all source values have been serial-  
21 ized, deserialized, and the remote stores are complete. The values referenced  
22 in the [src,src+count) interval must not be modified until completion is in-  
23 dicated.

24 The first variant notifies completion via the readying of the returned future.

25 The second variant immediately increments the promise's dependency count  
26 and notifies completion by fulfilling that dependency.

27 The third variant notifies completion by enlisting completion\_func in the given  
28 persona's user-progress (see §9.5.1 and Ch. 9).

29 *C++ memory ordering:* The writes to dest will have a happens-before re-  
30 lationship with the completion notification action (future readying, promise  
31 fulfillment, or persona continuation enlistment). In the third variant, all eval-  
32 uations sequenced-before this call will have a happens-before relationship with  
33 the execution of completion\_func.

34 *UPC++ progress level:* internal

### 35 7.3.2 Remote Gets

```
1  template<typename T>
2  future<T> rget(global_ptr<T> src);
3
4  template<typename T>
5  void rget(global_ptr<T> src, promise<T> &completion);
6
7  template<typename T, typename CompletionFunc>
8  void rget(global_ptr<T> src,
9           persona &completion_recipient,
10          CompletionFunc completion_func);
```

11 *Precondition:* T must be Serializable. src must reference a valid object of type  
12 T. In the second variant, the completion promise must have a dependency  
13 count greater than zero and must not have had fulfill\_result called on it  
14 before. In the third variant, CompletionFunc must be a function-object type  
15 accepting a single argument of type T, and completion\_func must not throw  
16 an exception when invoked on its argument.

17 Initiates a transfer to this rank of a single value of type T located at src.  
18 Completion of the operation implies completion of serialization at the source  
19 side and deserialization at the initiator. Completion delivers the retrieved value  
20 directly in the notification.

21 The first variant notifies completion and the value by readying the future with  
22 that value.

23 The second variant notifies completion and the value by fulfilling the promise  
24 via fulfill\_result(value).

25 The third variant notifies completion and the value by enlisting the invocation  
26 of completion\_func(value) in the given persona's user-progress (see §9.5.1  
27 and Ch. 9).

28 *C++ memory ordering:* In the third variant, all evaluations *sequenced-before*  
29 this call will have a *happens-before* relationship with the subsequent invocation  
30 of completion\_func(value).

31 *UPC++ progress level:* internal

```
32  template<typename T>
33  future<> rget(global_ptr<T> src, T *dest, std::size_t count);
34
35  template<typename T>
36  void rget(global_ptr<T> src, T *dest, std::size_t count,
37          promise<> &completion);
```



```
1
2 template<typename T, typename CompletionFunc>
3 void rget(global_ptr<T> src, T *dest, std::size_t count,
4           persona &completion_recipient,
5           CompletionFunc completion_func);
```

6 *Precondition:* T must be Serializable. Addresses in the intervals  
7 [src, src+count) and [dest, dest+count) must all reference valid objects  
8 of type T. No object may be referenced by both intervals. In the second variant,  
9 completion must have a dependency count greater than zero. In the third vari-  
10 ant, CompletionFunc must be a function-object type accepting no arguments,  
11 and the call completion\_func() must not throw an exception.

12 Initiates a transfer of count values of type T beginning at src and assigns them  
13 to the locations beginning at dest.

14 Completion of the operation indicates completion of remote serialization, initiator-  
15 side deserialization, and all local assignments. The source values must not be  
16 modified until completion is notified.

17 The first variant notifies completion by readying the returned future.

18 The second variant immediately increments the promise's dependency count,  
19 and notifies completion by fulfilling that dependency.

20 The third variant notifies completion by enlisting completion\_func to be in-  
21 voked in the given persona's user-progress (see §9.5.1 and Ch. 9).

22 *C++ memory ordering:* In the third variant, all evaluations sequenced-before  
23 this call and the local assignments to dest will have a happens-before relation-  
24 ship with the invocation of completion\_func.

25 *UPC++ progress level:* internal

# 1 Chapter 8

## 2 Remote Procedure Call

### 3 8.1 Overview

4 UPC++ provides remote procedure calls (RPCs) for injecting function calls into other ranks.  
5 These injections are one-sided, meaning the recipient is not required to explicitly acknowl-  
6 edge which functions are expected. Concurrent with a rank's execution, incoming RPCs  
7 accumulate in an internal queue managed by UPC++. The only control a rank has over  
8 inbound RPCs is when it would like to check its inbox for arrived function calls and execute  
9 them. Draining the RPC inbox is one of the many responsibilities of the progress API (see  
10 Ch. 9, *Progress*).

11 There are two main flavors of RPC in UPC++: *fire-and-forget* (`rpc_ff`) and *round trip*  
12 (`rpc` without the promise argument). Each takes a function `Func` together with variadic  
13 arguments `Args`.

14 The `rpc_ff` call serializes the given function and arguments into a message destined  
15 for the recipient, and guarantees that this function call will be placed eventually in the  
16 recipient's inbox. The round-trip `rpc` call does the same, but also forces the recipient to  
17 reply to the sender of the RPC with a message containing the return value of the function,  
18 fulfilling the future returned by the sender's invocation of `rpc`. Thus, when the future is  
19 ready, the sender knows the recipient has executed the function call. Additionally, if the  
20 return value of `func` is a future, the recipient will wait for that future to become ready  
21 before sending its result back to the sender.

22 The call `rput_then_rpc` combines a remote put with an `rpc_ff`, and the RPC is invoked  
23 after the remote put completes.

24 There are important restrictions on what the permissible types for `func` and its bound  
25 arguments can be for RPC functions. First, the `Func` type must be a function object (has a  
26 publicly accessible overload of the function call operator, `operator()`). Second, both the  
27 `Func` and all `Args...` types must be Serializable (see Ch. 6, *Serialization*).

## 8.2 Remote Hello World Example

Figure 8.1 shows a simple alternative *Hello World* example where each rank issues an `rpc` to its neighbor, where the last rank wraps around to 0.

```

32 #include <upcxx/upcxx.hpp>
33 #include <iostream>
34 void hello_world(intrank_t num){
35     std::cout << "Rank " << num <<"  told rank " << upcxx::rank_me()
36         << " to say Hello World" << std::endl;
37 }
38 int main(int argc, char** argv){
39     upcxx::init();           // Start UPC++ state
40     intrank_t remote = (upcxx::rank_me()+1)%upcxx::rank_n();
41     auto f = upcxx::rpc(remote, hello_world, upcxx::rank_me());
42     f.wait();
43     upcxx::finalize();     // Close down UPC++ state
44     return 0;
45 }

```

Figure 8.1: HelloWorld with Remote Procedure Call

## 8.3 API Reference

```

5 template<typename Func, typename ...Args>
6 void rpc_ff(intrank_t recipient, Func &&func, Args &&...args);

```

*Precondition:* `Func` must be a Serializable type and a function-object type. Each of `Args...` must be a Serializable type, or `dist_object<T>&`, or `team&`. The call `func(args...)` must not throw an exception.

The `func` and `args...` are serialized immediately and retained internally until they are eventually sent. After their receipt on `recipient`, they are deserialized and `func(args...)` is enlisted for execution during user-level progress of the master persona. So long as the sending persona continues to make internal-level progress it is guaranteed that the message will eventually arrive at the recipient. See §9.5.3 `progress_required` for an understanding of how much internal-progress is necessary.

Special handling is applied to those members of `args` which are either a reference to `dist_object` type (see §12 Distributed Objects) or a `team` (see §11

1 Teams). These are serialized by their `dist_id` or `team_id` respectively. The  
2 recipient deserializes the id's and waits asynchronously until all of them have a  
3 corresponding instance constructed on the recipient. When that occurs, `func`  
4 is called with the recipient's instance references in place of those supplied at  
5 the send site.

6 *C++ memory ordering:* All evaluations *sequenced-before* this call will have a  
7 *happens-before* relationship with the recipient's invocation of `func`.

8 *UPC++ progress level:* `internal`

```
9 template<typename Func, typename ...Args>  
10 future_invoke_result_t<Func, Args...>  
11 rpc(inrank_t recipient, Func &&func, Args &&...args);
```

12 *Precondition:* `Func` must be a Serializable type and a function-object type.  
13 Each of `Args...` must be either a Serializable type, or `dist_object<T>&`, or  
14 `team&`. Additionally, `std::result_of<Func(Args...)>::type` must be a Se-  
15 rIALIZABLE type or `future<T...>`, where each type in `T...` must be Serializable.  
16 The call `func(args...)` must not throw an exception.

17 Similar to `rpc_ff`, this call sends `func` and `args...` to be executed remotely,  
18 but additionally returns a non-ready future which will be readied with the value  
19 returned from the remote invocation of `func(args...)`.

20 `func` and `args...` are either serialized immediately, or copy/moved (depending  
21 on the universal reference) to internal storage managed by UPC++ and serialized  
22 sometime later (the returned future's readying indicates that serialization is  
23 complete). The serialized values are then sent to `recipient`, and upon receipt  
24 are deserialized and `func(args...)` is enlisted for execution during user-level  
25 progress of the master persona.

26 If the result of `func(args...)` is a future, the return type of `rpc` is the same  
27 as that of the result, and the recipient will wait for the future to become ready  
28 before sending its results back to the sender. Otherwise, the return type of `rpc`  
29 is a future that encapsulates the result of `func(args...)`, unless the result  
30 of `func` is `void`, in which case it is `future<>`. Within user-progress of the  
31 recipient's master persona, the result from invoking `func(args...)` will be  
32 immediately serialized and eventually sent back to the initiating rank. Upon  
33 receipt, it will be deserialized and the action of readying the final future with  
34 that value will be enlisted into user-progress of the initiating persona.

35 The same special handling applied to `dist_object` and `team` arguments by  
36 `rpc_ff` is also done by `rpc`.

1 *C++ memory ordering:* All evaluations *sequenced-before* this call will have a  
 2 *happens-before* relationship with the invocation of `func`. The return from `func`,  
 3 and readying of that return value if it is a future, will have a *happens-before*  
 4 relationship with the readying of the final future.

5 *UPC++ progress level:* `internal`

```
6 template<typename Func, typename ...Args>
7 void rpc(intrank_t recipient,
8         promise_invoke_result_t<Func, Args...> &pro,
9         Func &&func, Args &&...args);
```

10 *Precondition:* `Func` must be a Serializable type and a function-object type.  
 11 Each of `Args...` must be a Serializable type, or `dist_object<T>&`, or `team&`.  
 12 Additionally, `std::result_of<Func(Args...)>::type` must be a Serializable  
 13 type or `future<T...>`, where each type in `T...` must be Serializable. The  
 14 call `func(args...)` must not throw an exception. The dependency count of  
 15 `pro` must be greater than zero, and if it is a non-empty promise, then its non-  
 16 anonymous dependency must not have been fulfilled.

17 Sends `func` and `args...` and sends back the result in the same way as the  
 18 future-returning variant of `rpc`, but instead fulfills the given promise with  
 19 the final value during user-progress of the initiating persona. If the result  
 20 of `func(args...)` is of the form `future<T...>`, then `pro` must have the type  
 21 `promise<T...>`. If the result is some other non-void type `T`, then `pro` must be  
 22 of type `promise<T>`. And if the result is `void`, `pro` must be of type `promise<>`.  
 23 In all cases where `pro` has type `promise<>`, the call to `rpc` increments the  
 24 anonymous dependency count of `pro`.

25 The same special handling applied to `dist_object` and `team` arguments by  
 26 `rpc_ff` is also done by `rpc`.

27 *C++ memory ordering:* All evaluations *sequenced-before* this call will have a  
 28 *happens-before* relationship with the invocation of `func`. The return from `func`,  
 29 and readying of that return value if it is a future, will have a *happens-before*  
 30 relationship with the fulfillment of the promise.

31 *UPC++ progress level:* `internal`

```
32 template<typename T,
33         typename RemoteCompletionFunc,
34         typename ...RemoteCompletionArgs>
35 future<> rput_then_rpc (
```

```
1   T const *src, global_ptr<T> dest,
2   std::size_t count,
3   RemoteCompletionFunc &&remote_completion_func,
4   RemoteCompletionArgs &&...remote_completion_args);
5
6   template<typename T,
7           typename RemoteCompletionFunc,
8           typename ...RemoteCompletionArgs>
9   void rput_then_rpc(
10  T const *src, global_ptr<T> dest,
11  std::size_t count,
12  promise<> &source_completion,
13  RemoteCompletionFunc &&remote_completion_func,
14  RemoteCompletionArgs &&...remote_completion_args);
15
16  template<typename T, typename SourceCompletionFunc,
17          typename RemoteCompletionFunc,
18          typename ...RemoteCompletionArgs>
19  void rput_then_rpc(
20  T const *src, global_ptr<T> dest,
21  std::size_t count,
22  persona &source_completion_persona,
23  SourceCompletionFunc source_completion_func,
24  RemoteCompletionFunc &&remote_completion_func,
25  RemoteCompletionArgs &&...remote_completion_args);
```

26 *Precondition:* RemoteCompletionFunc must be a Serializable type  
27 and a function-object type. Each of RemoteCompletionArgs...  
28 must either be a Serializable type, or dist\_object<U>&, or team&.  
29 SourceCompletionFunc must be a function-object type. The  
30 calls remote\_completion\_func(remote\_completion\_args...) and  
31 source\_completion\_func() must not throw an exception. Either dest  
32 or dest-1 must reference a valid object of type T. Addresses in the intervals  
33 [src, src+count) and [dest,dest+count) must all reference valid objects  
34 of type T. No object may be referenced by both intervals. For the second  
35 variant, the dependency count of source\_completion must be greater than  
36 zero.

37 Initiates a transfer of count items of type T from the local memory at src  
38 to the memory referenced by dest. Sends remote\_completion\_func and  
39 remote\_completion\_args... to the rank dest.where() (in the same manner  
40 as rpc\_ff) and enlists remote\_completion\_func(remote\_completion\_args...)

1 to be run in user-progress of the master persona after the transfer completes.  
2 Serialization of `remote_completion_func` and `remote_completion_args` hap-  
3 pens during this function call.

4 The initiating rank is notified of source completion, which only indicates that  
5 serialization of the source memory has occurred and the contents can be re-  
6 claimed. Source completion does not indicate the puts have become visible or  
7 that `remote_completion_func` has run on the target rank.

8 In the first variant, the resulting future represents source completion of the  
9 transfer. In the second variant, the dependency count of the given promise  
10 is incremented, and a dependency is fulfilled upon source completion of the  
11 transfer. In the third variant, `source_completion_func` is enlisted to the given  
12 persona's user-progress upon source completion of the transfer. The memory  
13 referenced by `src` must not be modified until notification of source completion.

14 *C++ memory ordering:* All evaluations *sequenced-before* this call and the puts  
15 from this call will have a *happens-before* relationship with the invocation of  
16 `remote_completion_func`. In the third variant, all evaluations *sequenced-*  
17 *before* this call will have a *happens-before* relationship with the invocation of  
18 `source_completion_func`.

19 *UPC++ progress level:* `internal`

# 1 Chapter 9

## 2 Progress

### 3 9.1 Overview

4 UPC++ presents a highly-asynchronous interface, but guarantees that user-provided call-  
5 backs will only ever run on user threads during calls to the library. This guarantees a good  
6 user-visibility of the resource requirements of UPC++, while providing a better interoper-  
7 ability with other software packages which may have restrictive threading requirements.  
8 However, such a design choice requires the application developer to be conscientious about  
9 providing UPC++ access to CPU cycles.

10 Progress in UPC++ refers to how the calling application allows the UPC++ internal run-  
11 time to advance the state of its outstanding asynchronous operations. Any asynchronous  
12 operation initiated by the user may require the application to give UPC++ access to the exe-  
13 cution thread periodically until the operation reports its completion. Such access is granted  
14 by simply making calls into UPC++. Each UPC++ function’s contract to the user contains its  
15 *progress guarantee* level. This is described by the members of the `upcxx::progress_level`  
16 enumerated type:

17 `progress_level::none` UPC++ will not attempt to advance the progress of asynchronous  
18 operations.

19 `progress_level::internal` UPC++ may advance its internal state, but no notifications  
20 will be delivered to the application. Thus, an application has very limited ways to  
21 “observe” the effects of such progress.

22 `progress_level::user` UPC++ may advance its internal state as well as signal completion  
23 of user-initiated operations. This may entail the firing of remotely injected procedure  
24 calls (RPCs), or readying/fulfillment of futures/promises and the ensuing callback  
25 cascade.



1 The most common progress guarantee made by UPC++ functions is *progress\_level::internal*.  
2 This ensures the delivery of notifications to remote ranks (or other threads) making *user*-  
3 level progress in a timely manner. In order to avoid having the user contend with the  
4 cost associated with callbacks and RPCs being run anytime a UPC++ function is entered,  
5 *progress\_level::user* is purposefully not the common case.

6 `progress` is the notable function enabling the application to make *user*-level progress.  
7 Its sole purpose is to look for ready operations involving this rank or thread and run the  
8 associated RPC/callback code.

```
9 upcxx::progress(progress_level lev = progress_level::user)
```

10 UPC++ execution phases which leverage asynchrony heavily tend to follow a particular  
11 program structure. First, initial communications are launched. Their completion callbacks  
12 might then perform a mixture of compute or further UPC++ communication with similar,  
13 cascading completion callbacks. Then, the application spins on `upcxx::progress()`,  
14 checking some designated application state which monitors the amount of pending outgoing/  
15 incoming/local work to be done. For the user, understanding which functions perform  
16 these progress spins becomes crucial, since any invocation of user-level progress may execute  
17 RPCs or callbacks.

## 18 9.2 Restricted Context

19 During user-level progress made by UPC++, callbacks may be executed. Such callbacks  
20 are subject to restrictions on how they may further invoke UPC++ themselves. We designate  
21 such restricted execution of callbacks as being in the *restricted context*. The general  
22 restriction is stated as:

23 *User code running in the restricted context must assume that for the duration*  
24 *of the context all other attempts at making user-level progress, from any thread*  
25 *on any rank, may result in a no-op every time.*

26 The immediate implication is that a thread which is already in the restricted context  
27 should assume no-op behavior from further attempts at making progress. This makes it  
28 pointless to try and wait for UPC++ notifications from within restricted context since there  
29 is no viable mechanism to make the notifications visible to the user. Thus, calling any  
30 routine which spins on user-level progress until some notification occurs will likely hang  
31 the thread.

## 1 9.3 Attentiveness

2 Many UPC++ operations have a mechanism to signal completion to the application. How-  
3 ever, a performance-oriented application will need to be aware of an additional asyn-  
4 chronous operation status indicator called *progress-required*. This status indicates that for  
5 a particular operation further advancements of the current rank or thread's *internal*-level  
6 progress are necessary so that completion regarding remote entities (e.g. notification of  
7 delivery) can be reached. Once an operation has left the progress-required state, UPC++  
8 guarantees that remote entities will see their side of the operations' completion without  
9 any further progress by the current compute resource. Applications will need to leverage  
10 this information for performance, as it is inadvisable for a compute resource to become  
11 inattentive to UPC++ progress (e.g. long bouts of arithmetic-heavy computation) while  
12 other entities depend on operations that require further servicing.

13 As said previously, nearly all UPC++ operations track their completion individually.  
14 However, it is not possible for the programmer to query UPC++ if individual operations  
15 no longer require further progress. Instead, the user may ask UPC++ when all operations  
16 initiated by this rank have reached a state at which they no longer require progress. This  
17 is achieved by using the following functions:

```
18 bool upcxx::progress_required();  
19 void upcxx::discharge();
```

20 The `progress_required` function reports whether this rank requires progress, allow-  
21 ing the application to know that there are still pending operations that will not achieve  
22 remote completion without further advancements to internal progress. This is of particular  
23 importance before an application enters a lapse of inattentiveness (for instance, performing  
24 expensive computations) in order to prevent slowing down remote entities.

25 The `discharge` function allows an application to ensure that UPC++ does not require  
26 progress anymore. It is equivalent to the following:

```
27 void upcxx::discharge() {  
28     while(upcxx::progress_required())  
29         upcxx::progress(upcxx::progress_level::internal);  
30 }
```

31 A well-behaved UPC++ application is encouraged to call `discharge` before any long lapse  
32 of attentiveness to progress.

## 33 9.4 Thread Personas/Notification Affinity

34 As explained in Chapter 5 *Futures and Promises*, futures require careful consideration  
35 when used in the presence of thread concurrency. It is crucial that UPC++ is very explicit

1 about how a multi-threaded application can safely use futures returned by UPC++ calls.

2 The most important thing an application has to be aware of is which thread UPC++  
 3 will use to signal completion of a given future. It is therefore extremely important to  
 4 know that UPC++ will use the same thread to which the future was returned by the UPC++  
 5 operation (i.e. the thread which invoked the operation in the first place). This means  
 6 that the thread which invoked a future-returning operation will be the only one able to  
 7 see that operation’s completion. As UPC++ triggers futures only during a call which makes  
 8 user-level progress, the invoking thread must continue to make such progress calls until  
 9 the future is satisfied. This requirement has the drawback of banning the application from  
 10 doing the following: initiating a future-returning operation on one thread, allowing that  
 11 thread to terminate or become permanently inattentive (e.g. sleeping in a thread pool),  
 12 and expecting a different thread to receive the future’s completion. This section will focus  
 13 on two ways the application can still attain this use-case.

14 The notion of “thread” has been used in a loose fashion throughout this document,  
 15 the natural interpretation being an operating system (OS) thread. More precisely, this  
 16 document uses the notion of “thread” to denote a UPC++ device referred to as *thread persona*  
 17 which generalizes the notion of operating system threads.

18 A UPC++ thread persona is a collection of UPC++-internal state usually attributed to a  
 19 single thread. By making it a proper construct, UPC++ allows a single OS thread to switch  
 20 between multiple application-defined roles for processing notifications. Personas act as the  
 21 receivers for notifications generated by the UPC++ runtime.

22 Values of type `upcxx::persona` are non-copyable, non-moveable objects which the  
 23 application can instantiate as desired. For each OS thread, UPC++ internally maintains  
 24 a stack of *active* persona references. The top of this stack is the *current* persona. All  
 25 asynchronous UPC++ operations will have their notification events (signaling of futures or  
 26 promises) sent to the current persona of the OS thread invoking the operation. Calls that  
 27 make user-level progress will process notifications destined to any of the active personas of  
 28 the invoking thread. The initial state of the persona stack consists of a single entry pointing  
 29 to a persona created by UPC++ which is dedicated to the current OS thread. Therefore,  
 30 if the application never makes any use of the persona API, notifications will be processed  
 31 solely by the OS thread that initiates the operation.

32 Pushing and popping personas from the persona stack (hence changing the current  
 33 persona) is done with the `upcxx::persona_scope` type.

```
34 namespace upcxx {
35
36     struct persona_scope {
37         // Make 'p' the new current persona for this OS thread.
38         persona_scope(persona &p);
39
40         // Acquire 'lock', then make 'p' the new current persona for
```

```
1     // this OS thread.
2     template<typename Lock>
3     persona_scope(Lock &lock, persona &p);
4
5     // Pop 'p' from persona stack, release 'lock' if any.
6     // Calling thread must be same for constructor and destructor.
7     ~persona_scope();
8 };
9
10 persona_scope& top_persona_scope();
11
12 persona_scope& default_persona_scope();
13
14 bool progress_required(persona_scope &ps = top_persona_scope());
15
16 void discharge(persona_scope &ps = top_persona_scope());
17
18 } // namespace upcxx
19
20 // Example demonstrating persona_scope.
21 upcxx::persona scheduler_persona;
22 std::mutex scheduler_lock;
23
24 { // Scope block delimits domain of persona_scope instance.
25     auto scope = upcxx::persona_scope(scheduler_lock, scheduler_persona);
26
27     // All following upcxx actions will use 'scheduler_persona'
28     // as current.
29
30     // ...
31
32     // 'scope' destructs:
33     // - 'scheduler_persona' dropped from active set if it
34     //   wasn't active before the scope's construction.
35     // - Previously current persona revived.
36     // - Lock released.
37 }
```

Since UPC++ will assume an OS thread has exclusive access to all of its active personas, it is the user's responsibility to ensure that no OS threads share an active persona concurrently. The use of the `persona_scope` constructor, which takes a lock-like synchronization primitive, is strongly encouraged to facilitate in enforcing this invariant.

1 There are two ways that asynchronous operations can be initiated by a given OS thread  
2 but retired in another. The first solution is simple:

- 3 1. The user defines a persona P.
- 4 2. Thread 1 activates P, initiates the asynchronous operation, and releases P.
- 5 3. Thread 1 synchronizes with Thread 2, indicating the operation has been initiated.
- 6 4. Thread 2 activates P, spins on `progress` until the operation completes.

7 Care must be taken that any futures created by phase 2 are never altered (uttered)  
8 concurrently. The same synchronization that was used to enforce exclusivity of persona  
9 acquisition can be leveraged to protect the future as well.

10 While this technique achieves our goal of different threads initiating and resolving  
11 asynchronous operations, it fails a different but also desirable property. It is often desirable  
12 to allow multiple threads to issue communication *concurrently* while delegating a separate  
13 thread to handle the notifications. To achieve this, it is clear that multiple personas are  
14 needed. Indeed, the exclusivity of a persona being current to only one OS thread prevents  
15 the application from concurrent initiation of communication.

16 In order to issue operations and concurrently retire them in a different thread, the user  
17 is strongly encouraged to use the callback-oriented API calls of UPC++ as opposed to the  
18 future or promise variants. An example of such a variant is:

```
19 template<typename T, typename CompletionFunc>
20 void upcxx::rput(T const *src, global_ptr<T> dest, std::size_t count,
21                persona &completion_recipient,
22                CompletionFunc completion_func);
```

23 In addition to the arguments necessary for the particular operation, the callback API  
24 takes a persona reference and a C++ function object (lambda, etc.) such that upon comple-  
25 tion of the operation, the designated persona shall execute the function object during its  
26 user-level progress. Using the callback API, it is simple to have multiple threads initiating  
27 communication concurrently with a designated thread receiving the completion notifica-  
28 tions. To achieve this, each operation is initiated by a thread using the agreed-upon  
29 persona of the receiver thread together with a callback that will incorporate knowledge of  
30 completion into the receiver's state.

## 31 9.5 API Reference

```
1 enum class progress_level {
2     /*none, -- not an actual member, conceptual only*/
3     internal,
4     user
5 };

6 void upcxx::progress(progress_level lev = progress_level::user);
```

7 This call will always attempt to advance internal progress.

8 If `lev == progress_level::user` then this thread is also used to execute any  
9 available user actions for the personas currently active. Actions include:

- 10 1. Either future-readying or promise-fulfilling completion notifications for  
11 asynchronous operations initiated by one of the active personas. By the  
12 execution model of futures and promises this can induce callback cascade.
- 13 2. Continuation-style completion notifications from operations initiated by  
14 any persona but designating one of the active personas as the completion  
15 recipient.
- 16 3. RPCs destined for this rank but only if the master persona is among the  
17 active set.
- 18 4. `lpc`'s destined for any of the active personas.

19 *UPC++ progress level: internal or user*

## 20 9.5.1 persona

```
21 class persona;
```

22 C++ Concepts: DefaultConstructible, Destructible

```
23 persona::persona();
```

24 Constructs a persona object with no enqueued operations.

25 *This function is legal to call when UPC++ is in the uninitialized state.*

```
26 persona::~~persona();
```

1 Destructs this persona object. If this persona is a member of any thread's  
 2 persona stack, the result of this call is undefined. If any operations are currently  
 3 enqueued on this persona, or if any operations initiated by this persona require  
 4 further progress, the result of this call is undefined.

5 *This function is legal to call when UPC++ is in the uninitialized state.*

```
6 template<typename Func>
7 void persona::lpc_ff(Func func);
```

8 *Precondition:* `Func` must be a function-object type that can be invoked on zero  
 9 arguments, and the call `func()` must not throw an exception.

10 `std::move`'s `func` into an unordered collection of type-erased function objects  
 11 to be executed during user-level progress of the targeted (this) persona. This  
 12 function is thread-safe, so it may be called from any thread to enqueue work  
 13 for this persona.

14 *C++ memory ordering:* All evaluations *sequenced-before* this call will have a  
 15 *happens-before* relationship with the invocation of `func`.

16 *UPC++ progress level:* none

```
17 template<typename Func>
18 future_invoke_result_t<Func> persona::lpc(Func func);
```

19 *Precondition:* `Func` must be a function-object type that can be invoked on zero  
 20 arguments, and the call `func()` must not throw an exception.

21 `std::move`'s `func` into an unordered collection of type-erased function objects  
 22 to be executed during user-level progress of the targeted (this) persona. The  
 23 return value of `func` is asynchronously returned to the currently active persona  
 24 in a future. If the return value of `func` is a future, then the targeted persona will  
 25 wait for that future before signaling the future returned by `lpc` with its value.  
 26 This function is thread-safe, so it may be called from any thread to enqueue  
 27 work for this persona. Note that the future returned by `lpc` is considered to  
 28 be owned by the currently active persona, the future returned by `func` (if any)  
 29 will be considered owned by the target (this) persona.

30 *C++ memory ordering:* All evaluations *sequenced-before* this call will have a  
 31 *happens-before* relationship with the invocation of `func`, and the invocation of  
 32 `func` will have a *happens-before* relationship with evaluations sequenced after  
 33 the signaling of the final future.

34 *UPC++ progress level:* none

```
1 persona& master_persona();
```

2 Returns a reference to the master persona automatically instantiated by the  
3 UPC++ runtime. The thread that executes `upcxx::init` implicitly acquires this  
4 persona as its current persona. The master persona is special in that it is the  
5 only one which will execute RPCs destined for this rank. Additionally, some  
6 UPC++ functions may only be called by a thread with the master persona in its  
7 active stack.

8 *UPC++ progress level: none*

```
9 persona& current_persona();
```

10 Returns a reference to the persona on the top of the thread's active persona  
11 stack.

12 *UPC++ progress level: none*

```
13 persona& default_persona();
```

14 Returns a reference to the persona instantiated automatically and uniquely for  
15 this OS thread. The default persona is always the bottom of and can never be  
16 removed from its designated OS thread's active stack.

17 *UPC++ progress level: none*

```
18 void liberate_master_persona();
```

19 *Precondition:* This thread must be the one which called `upcxx::init`, it must  
20 have not altered its persona stack since calling `init`, and it must not have  
21 called this function already since calling `init`.

22 The thread which invokes `upcxx::init` implicitly has the master persona at  
23 the top of its active stack, yet the user has no `persona_scope` to drop to allow  
24 other threads to acquire the persona. Thus, if the user intends for other threads  
25 to acquire the master persona, they should have the `init`-calling thread release  
26 the persona with this function so that it can be claimed by `persona_scope`'s.  
27 Generally, if this function is ever called, it is done soon after `init` and then the  
28 master persona should be reacquired by a `persona_scope`.

29 *UPC++ progress level: none*



## 1 9.5.2 `persona_scope`

2 `class persona_scope;`

3 C++ Concepts: Destructible, MoveConstructible

4 `persona_scope::persona_scope(persona &p);`

5 *Precondition:* Excluding this thread, `p` is not a member of any other thread's  
6 active stack.

7 Pushes `p` onto the top of the calling OS thread's active persona stack.

8 *UPC++ progress level: none*

9 `template<typename Mutex>`

10 `persona_scope::persona_scope(Mutex &mutex, persona &p);`

11 C++ Concepts of `Mutex`: `Mutex`

12 *Precondition:* `p` will only be a member of some thread's active stack if that  
13 thread holds `mutex` in a locked state.

14 Invokes `mutex.lock()`, then pushes `p` onto the OS thread's active persona  
15 stack.

16 *UPC++ progress level: none*

17 `persona_scope::~~persona_scope();`

18 *Precondition:* All `persona_scope`'s constructed on this thread since the con-  
19 struction of this instance have since destructed.

20 The persona supplied to this instance's constructor is popped from this thread's  
21 active stack. If this instance was constructed with the `mutex` constructor, then  
22 that `mutex` is unlocked.

23 *UPC++ progress level: none*

24 `persona_scope& top_persona_scope();`

25 Reference to the most recently constructed but not destructed `persona_scope`  
26 for this thread. Every thread begins with an implicitly instantiated scope point-  
27 ing to its default persona that survives for the duration of the thread's lifetime.

28 *UPC++ progress level: none*

1 `persona_scope& default_persona_scope();`

2 Every thread begins with an implicitly instantiated scope pointing to its default  
3 persona that survives for the duration of the thread's lifetime. This function  
4 returns a reference to that bottommost `persona_scope` for the calling thread,  
5 which points at the calling thread's `default_persona()`.

6 *UPC++ progress level: none*

### 7 9.5.3 Outgoing Progress

8 `bool progress_required(persona_scope &ps = top_persona_scope());`

9 *Precondition:* `ps` has been constructed by this thread.

10 For the set of personas included in this thread's active stack section bounded  
11 inclusively between `ps` and the current top, *nearly* answers if any UPC++ op-  
12 erations initiated by those personas require further advancement of internal-  
13 progress of their respective personas before their completion events will be  
14 eventually available to user-level progress on the destined ranks. The exact  
15 meaning of the return value depends on which personas are selected by `ps`:

- 16 • If `ps` *does not* include the master persona: A return value of `true` means  
17 that one or more of the personas indicated by `ps` requires further internal-  
18 progress to achieve completion of its outgoing operations. A value of `false`  
19 means that none of the personas indicated by `ps` require internal-progress,  
20 but internal-progress of the master persona might still be required.
- 21 • If `ps` *does* include the master persona: A return value of `true` means that  
22 one or more of the personas indicated by `ps` requires further internal-  
23 progress to achieve completion of its outgoing operations. A return value  
24 of `false` means that none of the non-master personas indicated by `ps`  
25 requires further internal-progress, but the master persona may or may not  
26 require further internal-progress.

27 *UPC++ progress level: none*

28 `void discharge(persona_scope &ps = top_persona_scope());`

29 Advances internal-progress enough to ensure that `progress_required(ps)` re-  
30 turns `false`.

31 *UPC++ progress level: internal*

# 1 Chapter 10

## 2 Atomics

### 3 10.1 Overview

4 UPC++ supports atomic operations on shared memory locations. Atomicity entails that a  
5 read-modify-write sequence on a memory location will happen without interference or inter-  
6 leaving with other concurrently executing atomic operations. Atomicity is not guaranteed  
7 if a memory location is concurrently targeted by both atomic and non-atomic operations.  
8 The order in which concurrent atomics update the same memory is not guaranteed, not  
9 even for successively issued operations by a single rank. Ordering of atomics with respect  
10 to other asynchronous operations is also not guaranteed. The only means to ensure such  
11 ordering is by waiting for one operation to complete before initiating its successor.

12 At this time, it is unclear how UPC++ will support mixing of atomic and non-atomic  
13 accesses to the same memory location. Until this is resolved, users must assume that  
14 for the duration of the program, once a memory location is accessed via a UPC++ atomic,  
15 only further atomic operations to that location will have meaningful results (note that even  
16 global barrier synchronization does not grant an exception to this rule). This unfortunately  
17 implies that deallocation of such memory is unsafe, as that would allow the memory to be  
18 reallocated to a context unaware of its constrained condition.

19 Each atomic operation works on a global pointer of an *approved atomic type*. Cur-  
20 rently, the approved atomic types are a subset of fundamental integer types, specifically:  
21 `std::int32_t`, `std::uint32_t`, `std::int64_t`, and `std::uint64_t`. All atomic opera-  
22 tions are non-blocking and return a future to indicate completion. UPC++ currently supports  
23 only a limited set of operations: `get`, `put`, and `fetch-and-add`.

### 24 10.2 API Reference

```
25 template<typename T>
```

```
1 future<T> atomic_get(global_ptr<T> p, std::memory_order order);
```

2 *Precondition:* T must be one of the approved atomic types. p must reference a  
3 valid object of type T. T must be the only type used by any atomic referencing  
4 any part of p's target memory for the entire lifetime of UPC++. order must be  
5 std::memory\_order\_relaxed or std::memory\_order\_acquire.

6 Initiates an atomic read of the object at location p and returns its value in a  
7 future.

8 *C++ memory ordering:* If order is std::memory\_order\_acquire then the  
9 read performed will have a *happens-before* relationship with the readying of the  
10 returned future and all evaluations *sequenced-after*.

11 *UPC++ progress level: internal*

```
12 template<typename T>  
13 future<> atomic_put(global_ptr<T> p, T val,  
14                    std::memory_order order);
```

15 *Precondition:* T must be one of the approved atomic types. p must reference a  
16 valid object of type T. T must be the only type used by any atomic referencing  
17 any part of p's target memory for the entire lifetime of UPC++. order must be  
18 std::memory\_order\_relaxed or std::memory\_order\_release.

19 Initiates an atomic write of val to the location p. Completion of the write is  
20 indicated in the returned future.

21 *C++ memory ordering:* If order is std::memory\_order\_release then all eval-  
22 uations *sequenced-before* this call will have a *happens-before* relationship with  
23 the write performed.

24 *UPC++ progress level: internal*

```
25 template<typename T>  
26 future<T> atomic_fetch_add(global_ptr<T> p, T val,  
27                          std::memory_order order);
```

28 *Precondition:* T must be one of the approved atomic types. p must refer-  
29 ence a valid object of type T. T must be the only type used by any atomic  
30 referencing any part of p's target memory for the entire lifetime of UPC+-.  
31 order must be std::memory\_order\_relaxed, std::memory\_order\_acquire,  
32 std::memory\_order\_release, or std::memory\_order\_acq\_rel.

1 Initiates the atomic read-modify-write operation consisting of: reading the  
2 value of the object located at `p`, adding `val` to it, and writing the new value  
3 back. The value returned in the future is the one initially read.

4 *C++ memory ordering:* If `order` is either `std::memory_order_release` or  
5 `std::memory_order_acq_rel` then all evaluations *sequenced-before* this call  
6 will have a *happens-before* relationship with the atomic action. If `order` is  
7 `std::memory_order_acquire` or `std::memory_order_acq_rel` then the atomic  
8 action will have a *happens-before* relationship with the readying of the returned  
9 future and all evaluations *sequenced-after*.

10 *UPC++ progress level:* **internal**

# 1 Chapter 11

## 2 Teams

### 3 11.1 Overview

4 UPC++ provides *teams* as a means of grouping ranks. UPC++ uses `teams` for collective op-  
5 erations. `team` construction is collective and should be considered moderately expensive  
6 and done as part of the set-up phase of a calculation. `teams` are similar to `MPI_Groups`  
7 and the default `team` is `world()`. `teams` are considered special when it comes to serial-  
8 ization. Each `team` has a unique `team_id` that is equal across the `team` and acts as an  
9 opaque handle. Any rank that is a member of the `team` can retrieve the `team` object with  
10 the `team_id::here()` function. Hence, coordinating ranks can reference specific `teams` by  
11 their `team_id`.

12 While a rank within a UPC++ SPMD program can have multiple `inrank_t` values that  
13 represent their relative placement in several `teams`, it is the `inrank_t` in the `world()`  
14 that is used in all UPC++ functions, unless otherwise specifically noted. For example,  
15 `broadcast_recv` uses the team-relative rank.

### 16 11.2 Local Teams

17 Each rank can obtain a reference to a special team by calling `local_team.global_ptr`'s to  
18 objects allocated by ranks within this `team` will report `is_local() == true` and `local()`  
19 will return a valid `T*` to that memory. The `global_ptr.where()` function will report the  
20 rank (in team `world()`) that originally acquired that memory using the functions in chapter  
21 4. It is not guaranteed that the `T*`'s obtained by different ranks to the same shared object  
22 will have bit-wise identical pointer values. In the general case, peers may have different  
23 virtual addresses for the same physical memory.

## 1 11.3 API Reference

### 2 11.3.1 team

```
3 class team;
```

4 C++ Concepts: MoveConstructible, Destructible

```
5 intrank_t team::rank_n() const;
```

6 Returns the number of ranks that are in the given team.

7 *UPC++ progress level: none*

```
8 intrank_t team::rank_me() const;
```

9 Returns the peer index of the caller in the given team.

10 *UPC++ progress level: none*

```
11 intrank_t team::operator[](intrank_t peer_index) const;
```

12 *Precondition: peer\_index >= 0 and peer\_index < rank\_n().*

13 Returns the index in the world() team for the rank associated with peer\_index  
14 in this team.

15 *UPC++ progress level: unspecified between none and internal*

```
16 intrank_t team::from_world(intrank_t world_index) const;
```

```
17 intrank_t team::from_world(intrank_t world_index,  
18                             intrank_t otherwise) const;
```

19 *Precondition: world\_index >= 0 and world\_index < world().rank\_n().* For  
20 the single argument overload, the rank associated with world\_index must be  
21 a member of this team.

22 Returns the peer index in this team of the rank associated with world\_index in  
23 the world() team. For the two argument overload, if the rank is not a member  
24 of this team then the value of otherwise is returned.

25 *UPC++ progress level: unspecified between none and internal*

```
1 team team::split(intrank_t color, intrank_t key);
```

2 *Precondition:* This function must be called collectively by all the ranks in this  
3 team, and it must be called by the thread that has the master persona (§9.5.1).  
4 No two ranks in the collective call may specify the same combination of `color`  
5 and `key`.

6 Splits the given team into subteams based on the `color` and `key` arguments.  
7 All ranks that call the function with the same `color` value will be separated  
8 into the same subteam. Ranks in the same subteam will be numbered according  
9 to their position in the sequence of sorted key values. The return value is the  
10 team representing the calling rank's new subteam. This call will invoke user-  
11 level progress, so the caller may expect incoming RPCs to fire before it returns.

12 *C++ memory ordering:* With respect to all threads participating in this col-  
13 lective, all evaluations which are *sequenced-before* their respective thread's in-  
14 vocation of this call will have a *happens-before* relationship with all evaluations  
15 sequenced after the call.

16 *UPC++ progress level:* `user`

```
17 team::team(team &&other);
```

18 *Precondition:* Calling thread must have the master persona.

19 Makes this instance the calling rank's representative of the team associated with  
20 `other`, transferring all state from `other`. Invalidates `other`, and any subsequent  
21 operations on `other`, except for destruction, produce undefined behavior.

22 *UPC++ progress level:* `none`

```
23 team::~~team();
```

24 *Precondition:* Calling thread must have the master persona.

25 If this instance has not been invalidated by being passed to the move construc-  
26 tor, then this will destroy the current rank's state associated with the team.  
27 Further lookups on this rank using the `team_id` corresponding to this team will  
28 have undefined behavior. If this instance has been invalidated by a move, then  
29 this call will have no effect.

30 *UPC++ progress level:* `none`

```
31 team_id team::id() const;
```

32 Returns the universal name associated with this team.

33 *UPC++ progress level:* `none`



### 1 11.3.2 team\_id

2 `class team_id;`

3 C++ Concepts: PODType, EqualityComparable, LessThanComparable, hash-  
4 able

5 A universal name representing a team.

6 `team& team_id::here() const;`

7 *Precondition:* The current rank must be a member of the `team` associated with  
8 this name, and it must have completed creation of the `team`.

9 Retrieves a reference to the `team` instance associated with this name.

10 *UPC++ progress level: none*

11 `future<team &> team_id::when_here() const;`

12 *Precondition:* The current rank must be a member of the `team` associated with  
13 this name. The calling thread must have the master persona.

14 Retrieves a future representing when the current rank constructs the `team` cor-  
15 responding to this name.

16 *UPC++ progress level: none*

### 17 11.3.3 Fundamental Teams

18 `team& world();`

19 Returns a reference to the team representing all the ranks in the program. It  
20 is illegal to perform a move on the returned team.

21 *UPC++ progress level: none*

22 `inrank_t rank_n();`

23 Returns the number of ranks that are in the world team. Equivalent to `world().rank_n()`.

24 *UPC++ progress level: none*

25 `inrank_t rank_me();`

1 Returns the peer index of the caller in the world team. Equivalent to `world().rank_me()`.

2 *UPC++ progress level: none*

3 `team& local_team();`

4 Returns a reference to the local team containing this rank. A local team represents a set of ranks which share physical memory (§11.2). It is illegal to perform a move on the returned team.

7 *UPC++ progress level: none*

8 `bool local_team_contains(intrank_t world_index);`

9 *Precondition:* `world_index >= 0` and `world_index < world().rank_n()`.

10 Determines if `world_index` is a member of the local team containing the this rank (§11.2). Equivalent to: `local_team().from_world(world_index,-1) >= 0`

12 *UPC++ progress level: none*

### 13 11.3.4 Collectives

14 `void barrier(team &team = world());`

15 *Precondition:* This function must be called collectively by all the ranks in the given team, and it must be called by the thread that has the master persona (§9.5.1).

18 Performs a barrier operation over the given team. The call will not return until all ranks in the team have entered the call. There is no implied relationship between this call and other in-flight operations. This call will invoke user-level progress, so the caller may expect incoming RPCs to fire before it returns.

22 *C++ memory ordering:* With respect to all threads participating in this collective, all evaluations which are *sequenced-before* their respective thread's invocation of this call will have a *happens-before* relationship with all evaluations sequenced after the call.

26 *UPC++ progress level: user*

27 `future<> barrier_async(team &team = world());`

1 *Precondition:* This function must be called collectively by all the ranks in the  
 2 given team, and it must be called by the thread that has the master persona  
 3 (§9.5.1).

4 Initiates an asynchronous barrier operation over the given team. The call will  
 5 return without waiting for other ranks to make the call. The returned future  
 6 will only become ready after all other ranks in the team have entered the call.

7 *C++ memory ordering:* With respect to all threads participating in this col-  
 8 lective, all evaluations which are *sequenced-before* their respective thread's in-  
 9 vocation of this call will have a *happens-before* relationship with all evaluations  
 10 sequenced after the signaling of the returned futures.

11 *UPC++ progress level:* **internal**

```
12 template<typename T, typename BinaryOp>
13 future<T> allreduce(T &&value, BinaryOp &&op, team &team = world());
```

14 *Precondition:* This function must be called collectively by all the ranks in the  
 15 given team, and it must be called by the thread that has the master persona  
 16 (§9.5.1). T must be Serializable. BinaryOp must be a function-object type  
 17 representing an associative and commutative mathematical operation taking  
 18 two values of type T and returning a value implicitly convertible to T. BinaryOp  
 19 must be referentially transparent and concurrently invocable. BinaryOp may  
 20 not invoke any UPC++ routine with a progress level other than none.

21 Performs a reduction operation over the ranks in the given team. If the team  
 22 contains only a single rank, then the resulting future will hold value. Oth-  
 23 erwise, initiates an asynchronous reduction over the values provided by each  
 24 rank. The reduction is performed in some non-deterministic order by applying  
 25 op to combine values and intermediate results. Each rank receives the result of  
 26 the reduction in the returned future.

27 *C++ memory ordering:* With respect to all threads participating in this col-  
 28 lective, all evaluations which are *sequenced-before* their respective thread's in-  
 29 vocation of this call will have a *happens-before* relationship with all evaluations  
 30 sequenced after the signaling of the returned futures.

31 *UPC++ progress level:* **internal**

```
32 template<typename T>
33 future<T> broadcast(T &&value, intrank_t sender,
34                   team &team = world());
```

```
1
2  template<typename T>
3  future<> broadcast(T *buffer, std::size_t count,
4                   intrank_t sender, team &team = world());
```

5     *Precondition:* The function must be called collectively by the ranks in the  
6     given team, and it must be called by the thread that has the master persona  
7     (§9.5.1). The value of `sender`, and `count` in the second variant, must be the  
8     same across all callers. In the second variant, the addresses in the interval  
9     [`buffer`,`buffer+count`) must all reference valid objects of type `T`. The type  
10    `T` must be `Serializable`.

11    Initiates an asynchronous broadcast (one-to-all) operation, with rank `sender`  
12    acting as the producer of the broadcast. In the first variant, `value` will be  
13    asynchronously sent to all ranks in the team, encapsulated in the returned  
14    future, which will be ready upon receipt of the value. In the second variant, the  
15    objects in [`buffer`,`buffer+count`) on rank `sender` are sent to the addresses  
16    [`buffer`,`buffer+count`) provided by the receiving ranks. The returned future  
17    signals completion of the operation with respect to the calling rank. For the  
18    sender, this indicates that the given buffer is available for reuse, and for a  
19    receiver, it indicates that the data have been received in its buffer.

20    *C++ memory ordering:* With respect to all threads participating in this col-  
21    lective, all evaluations which are *sequenced-before* the producing thread's invo-  
22    cation of this call will have a *happens-before* relationship with all evaluations  
23    sequenced after the signaling of the returned futures.

24    *UPC++ progress level:* `internal`

# 1 Chapter 12

## 2 Distributed Objects

### 3 12.1 Overview

4 In distributed-memory parallel programming, the concept of a single logical object parti-  
5 tioned over several ranks is a useful capability in many contexts: for example, geometric  
6 meshes, vectors, matrices, tensors, and associative maps. Since `UPC++` is a communication  
7 library, it strives to focus on the mechanisms of communication as opposed to the various  
8 programming idioms for managing distribution. However, a basic framework for users to  
9 implement their own distributed objects is useful and also enables `UPC++` to provide the  
10 user with the following valuable features:

- 11 1. Universal distributed object naming: per-object names that can be transmitted to  
12 other ranks while retaining their meaning.
- 13 2. Name-to-this mapping: Mapping between the universal name and the current rank's  
14 memory address holding that distributed object's state for the rank (the current  
15 rank's `this` pointer).

16 The need for universal distributed object naming stems primarily from RPC-based com-  
17 munication. If one rank needs to remotely invoke code on a peer's partition of a distributed  
18 object, there needs to be some mutually agreeable identifier for referring to that distributed  
19 object. For simplicity, this identifier value should be: identical across all ranks so that it  
20 may be freely communicated while maintaining its meaning. Moreover, the name should  
21 be `TriviallyCopyable` so that it may be serialized into `RPCs` efficiently (including with the  
22 auto-capture `[=]` lambda syntax), hashable, and comparable so that it works well with  
23 standard `C++` containers. `UPC++` provides distributed object names meeting these criteria  
24 as well as the registry for mapping names to and from the current rank's partition of the  
25 distributed object.

## 1 12.2 Building Distributed Objects

2 Distributed objects are built with the `upcxx::dist_object<T>` type. For all ranks in a  
3 given team, each rank constructs an instance of `dist_object<T>`, supplying a value of type  
4 `T` representing this rank's instance value. All ranks in the team must call this constructor  
5 collectively. Once construction completes, the distributed object has a universal name  
6 which can be used on any rank in the team to locate the resident instance. When the  
7 `dist_object<T>` is destructed the `T` value is also destructed. At this point the name  
8 will cease to carry meaning on this rank. Thus, the programmer should ensure that no  
9 rank destructs a distributed object until all name lookups destined for it complete and all  
10 hanging references of the form `T&` or `T*` to the value have expired.

11 The names of `dist_object<T>`'s are encoded by the `dist_id<T>` type. This type  
12 is `TriviallyCopyable`, `EqualityComparable`, `LessThanComparable`, `hashable`, and `trivially`  
13 `Serializable`. It has the members `.here()` and `.when_here()` for retrieving the resident  
14 `dist_object<T>` instance registered with the name.

## 15 12.3 Ensuring Distributed Existence

16 The `dist_object<T>` constructor requires it be called in a collective context, but it does  
17 not guarantee that, after the call, all other ranks in the team have exited or even reached  
18 the constructor. Thus users are required to guard against the possibility that when an RPC  
19 carrying an distributed object's name executes, the recipient rank may not yet have an  
20 entry for that name in its registry. Possible ways to deal with this include:

- 21 1. Barrier: Before issuing communication containing a `dist_id<T>` for a newly created  
22 distributed object, the relevant team completes a `barrier` to ensure global existence  
23 of the `dist_object<T>`.
- 24 2. Point to point: Before communicating a `dist_id<T>` with a given rank, the initiat-  
25 ing rank uses some two-party protocol to ensure that the peer has constructed the  
26 `dist_object<T>`.
- 27 3. Asynchronous point-to-point: The user performs no synchronization to ensure remote  
28 existence. Instead, an RPC is sent which, upon arrival, must wait asynchronously via  
29 a continuation for the peer to construct the distributed object.

30 UPC++ enables the asynchronous point-to-point approach implicitly when `dist_object<T>&`  
31 arguments are given to any of the RPC family of functions (see Ch. 8).

## 1 12.4 API Reference

```
2 template<typename T>
3 struct dist_object<T>;
```

4 C++ Concepts: MoveConstructible, Destructible

```
5 template<typename T>
6 dist_object<T>::dist_object(T value, team &team = world());
```

7 *Precondition:* Calling thread must have the master persona.

8 Constructs this rank's member of the distributed object identified by the col-  
 9 lective calling context across `team`. The initial value for this rank is given  
 10 in `value`. The future returned from `dist_id<T>::when_here` for the corre-  
 11 sponding `dist_id<T>` will be readied during this constructor. This implies  
 12 that continuations waiting for that future will execute before the constructor  
 13 returns.

14 *UPC++ progress level:* none

```
15 template<typename T>
16 template<typename ...Arg>
17 dist_object<T>::dist_object(team &team, Arg &&...arg);
```

18 *Precondition:* Calling thread must have the master persona.

19 Constructs this rank's member of the distributed object identified by the col-  
 20 lective calling context across `team`. The initial value for this rank is constructed  
 21 with `T(std::forward<Arg>(arg)...)...`. The result is undefined if this call  
 22 throws an exception. The future returned from `dist_id<T>::when_here` for  
 23 the corresponding `dist_id<T>` will be readied during this constructor. This  
 24 implies that continuations waiting for that future will execute before the con-  
 25 structor returns.

26 *UPC++ progress level:* none

```
27 template<typename T>
28 dist_object<T>::dist_object(dist_object<T> &&other);
```

1     *Precondition:* Calling thread must have the master persona.

2     Makes this instance the calling rank's representative of the distributed object  
3     associated with **other**, transferring all state from **other**. Invalidates **other**, and  
4     any subsequent operations on **other**, except for destruction, produce undefined  
5     behavior.

6     *UPC++ progress level:* none

```
7  template<typename T>  
8  dist_object<T>::~~dist_object();
```

9     *Precondition:* Calling thread must have the master persona.

10     If this instance has not been invalidated by being passed to the move construc-  
11     tor, then this will destroy the current rank's member of the distributed object.  
12     ~T() will be invoked on the resident instance, and further lookups on this rank  
13     using the `dist_id<T>` corresponding to this distributed object will have unde-  
14     fined behavior. If this instance has been invalidated by a move, then this call  
15     will have no effect.

16     *UPC++ progress level:* none

```
17 template<typename T>  
18 dist_id<T> dist_object<T>::id() const;
```

19     Returns the `dist_id<T>` representing the universal name of this distributed  
20     object.

21     *UPC++ progress level:* none

```
22 template<typename T>  
23 T* dist_object<T>::operator->() const;
```

24     Access to the current rank's value instance for this distributed object.

25     *UPC++ progress level:* none

```
26 template<typename T>  
27 T& dist_object<T>::operator*() const;
```

28     Access to the current rank's value instance for this distributed object.

29     *UPC++ progress level:* none



```
1 template<typename T>
2 struct dist_id<T>;
```

3 C++ Concepts: PODType, EqualityComparable, LessThanComparable, hash-  
4 able

```
5 template<typename T>
6 future<dist_object<T>&> dist_id<T>::when_here() const;
```

7 *Precondition:* The current rank's `dist_object<T>` instance associated with this  
8 name must not have been destroyed. The calling thread must have the master  
9 persona.

10 Retrieves a future representing when the current rank constructs the `dist_object<T>`  
11 corresponding to this name.

12 *UPC++ progress level: none*

```
13 template<typename T>
14 dist_object<T>& dist_id<T>::here() const;
```

15 *Precondition:* The current rank's `dist_object<T>` instance associated with  
16 this name must be alive. The calling thread must have the master persona.

17 Retrieves a reference to the current rank's `dist_object<T>` instance associated  
18 with this name.

19 *UPC++ progress level: none*

# Chapter 13

## Non-Contiguous One-Sided Communication

### 13.1 Overview

UPC++ provides functions to perform one-sided communications similar to `rget` and `rput` which are dedicated to handle data stored in non-contiguous buffers.

These functions are denoted with the `fragmented` keyword, and take two sequences of `std::pair` (or more generally `std::tuple`) describing how source and destination fragmented buffers should be accessed.

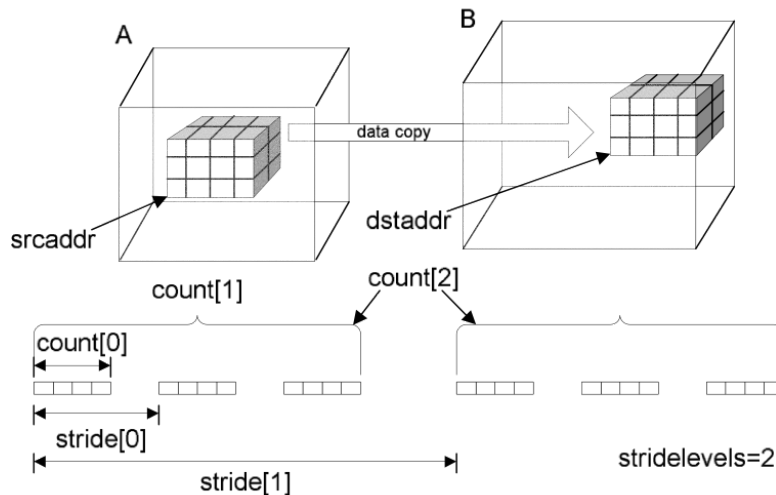


Figure 13.1: An example of a unit-stride  $i$  transfer between a `src` address and a `dst` address

The most general version of the API requires each `std::pair` to contain a local or

1 global pointer to a memory location in the first member while the second member contains  
2 the size of the contiguous chunk of memory to be transferred.

3 A second set of functions targets identical chunk sizes, thus requiring the user to provide  
4 pointers only. These functions are denoted by the `regular` keyword.

5 Finally, the third set of functions provide an API for strided accesses starting from  
6 two given source and destination addresses. An example of such a transfer is depicted in  
7 Figure 13.1. These are denoted by the `strided` keyword.

8 Each of the functions also has a `then_rpc` variant which executes a remote procedure  
9 call targeting the destination rank to signal completion of the transfer.

## 10 13.2 API Reference

### 11 13.2.1 Fragmented Put

```
12 // future variant
13 template<typename SrcIter, typename DestIter>
14 future<> rput_fragmented(
15     SrcIter src_runs_begin, SrcIter src_runs_end,
16     DestIter dest_runs_begin, DestIter dest_runs_end);
17
18 // promise variant
19 template<typename SrcIter, typename DestIter>
20 void rput_fragmented(
21     SrcIter src_runs_begin, SrcIter src_runs_end,
22     DestIter dest_runs_begin, DestIter dest_runs_end,
23     promise<> &completion);
24
25 // continuation variant
26 template<typename SrcIter, typename DestIter,
27         typename CompletionFunc>
28 void rput_fragmented(
29     SrcIter src_runs_begin, SrcIter src_runs_end,
30     DestIter dest_runs_begin, DestIter dest_runs_end,
31     persona &completion_recipient,
32     CompletionFunc completion_func);
```

33 *Preconditions:*

34 SrcIter and DestIter both satisfy the ForwardIterator C++ concept.  
35 std::get<0>(\*std::declval<SrcIter>()) has a return type convertible  
36 to T const\*, for some type T.

1        `std::get<1>(*std::declval<SrcIter>())` has a return type convertible  
2        to `std::size_t`.  
3        `std::get<0>(*std::declval<DestIter>())` has the return type `global_ptr<T>`,  
4        for the same type `T` as with `SrcIter`.  
5        `std::get<1>(*std::declval<DestIter>())` has a return type convert-  
6        ible to `std::size_t`.  
7        All destination addresses must be `global_ptr<T>`'s referencing memory  
8        with affinity to the same rank.  
9        The length of the expanded address sequence (the sum over the run  
10       lengths) must be the same for the source and destination sequences.  
11       `CompletionFunc` is a function-object type.

12       For some type `T`, takes a sequence of source addresses of `T const*` and a se-  
13       quence of destination addresses of `global_ptr<T>` and does the corresponding  
14       puts from each source address to the destination address of the same sequence  
15       position.

16       Address sequences are encoded in run-length form as sequences of runs, where  
17       each run is a pair consisting of a starting address plus the number of consecutive  
18       elements beginning at that address.

19       As an example of valid types for individual runs, `SrcIter` could be an iterator  
20       over elements of type `std::pair<T const*, std::size_t>`, and `DestIter` an  
21       iterator over `std::pair<global_ptr<T>, std::size_t>`. Variations replac-  
22       ing `std::pair` with `std::tuple` or `size_t` with other primitive integral types  
23       are also valid.

24       The sequence iterators must remain valid, and the underlying addresses and  
25       source memory contents must stay constant until completion is signaled. Only  
26       after completion is signaled can the address sequences and source memory be  
27       reclaimed by the application.

28       The destination memory regions must be completely disjoint and must not over-  
29       lap with any source memory regions, otherwise behavior is undefined. Source  
30       regions are permitted to overlap with each other.

31       In the future variant, completion is returned as a future.

32       In the promise variant, an anonymous dependency is added to the promise  
33       during the call and is fulfilled upon completion.

34       In the continuation variant, the `completion_func` function object is submitted  
35       to `completion_recipient`'s user-progress continuation queue upon comple-  
36       tion.

1 *C++ memory ordering:* In the continuation variant, all evaluations *sequenced-*  
2 *before* this call and the put operations from this call will have a *happens-before*  
3 relationship with the invocation of `completion_func`.

4 *UPC++ progress level:* `internal`

## 5 13.2.2 Fragmented Get

```
6 // future variant
7 template<typename SrcIter, typename DestIter>
8 future<> rget_fragmented(
9     SrcIter src_runs_begin, SrcIter src_runs_end,
10    DestIter dest_runs_begin, DestIter dest_runs_end);
11
12 // promise variant
13 template<typename SrcIter, typename DestIter>
14 void rget_fragmented(
15     SrcIter src_runs_begin, SrcIter src_runs_end,
16     DestIter dest_runs_begin, DestIter dest_runs_end,
17     promise<> &completion);
18
19 // continuation variant
20 template<typename SrcIter, typename DestIter,
21         typename CompletionFunc>
22 void rget_fragmented(
23     SrcIter src_runs_begin, SrcIter src_runs_end,
24     DestIter dest_runs_begin, DestIter dest_runs_end,
25     persona &completion_recipient,
26     CompletionFunc completion_func);
```

27 *Preconditions:*

28 SrcIter and DestIter both satisfy the ForwardIterator C++ concept.

29 `std::get<0>(*std::declval<SrcIter>())` has the type `global_ptr<T>`  
30 for some type T.

31 `std::get<1>(*std::declval<SrcIter>())` has a type convertible to `std::size_t`.

32 `std::get<0>(*std::declval<DestIter>())` has the type `T*`, for some  
33 type T.

34 `std::get<1>(*std::declval<DestIter>())` has a type convertible to  
35 `std::size_t`.

36 All source addresses must be `global_ptr<T>`'s referencing memory with  
37 affinity to the same rank.

1           The length of the expanded address sequence (the sum over the run  
2           lengths) must be the same for the source and destination sequences.

3           `CompletionFunc` is a function-object type.

4           For some type `T`, takes a sequence of source addresses of `global_ptr<T>` and a  
5           sequence of destination addresses of `T*` and does the corresponding gets from  
6           each source address to the destination address of the same sequence position.

7           Address sequences are encoded in run-length form as sequences of runs, where  
8           each run is a pair consisting of a starting address plus the number of consecutive  
9           elements beginning at that address.

10          As an example of valid types for individual runs, `DestIter` could be an it-  
11          erator over elements of type `std::pair<T*, std::size_t>`, and `SrcIter` an  
12          iterator over `std::pair<global_ptr<T>, std::size_t>`. Variations replac-  
13          ing `std::pair` with `std::tuple` or `size_t` with other primitive integral types  
14          are also valid.

15          The sequence iterators must remain valid, and the underlying addresses and  
16          source memory contents must stay constant until completion is signaled. Only  
17          after completion is signaled can the address sequences and source memory be  
18          reclaimed by the application.

19          The destination memory regions must be completely disjoint and must not over-  
20          lap with any source memory regions, otherwise behavior is undefined. Source  
21          regions are permitted to overlap with each other.

22          In the future variant, completion is returned as a future.

23          In the promise variant, an anonymous dependency is added to the promise  
24          during the call and is fulfilled upon completion.

25          In the continuation variant, the `completion_func` function object is submitted  
26          to `completion_recipient`'s user-progress continuation queue upon comple-  
27          tion.

28          *C++ memory ordering:* In the continuation variant, all evaluations *sequenced-*  
29          *before* this call and the gets from this call will have a *happens-before* relationship  
30          with the invocation of `completion_func`.

31          UPC++ progress level: `internal`

### 32 13.2.3 Fragmented Put then RPC

```
1 // future variant
2 template<typename SrcIter, typename DestIter,
3         typename RemoteCompletionFunc,
4         typename ...RemoteCompletionArgs>
5 future<> rput_fragmented_then_rpc(
6     SrcIter src_runs_begin, SrcIter src_runs_end,
7     DestIter dest_runs_begin, DestIter dest_runs_end,
8     intrank_t recipient,
9     RemoteCompletionFunc &&remote_completion_func,
10    RemoteCompletionArgs &&...remote_completion_args);
11
12 // promise variant
13 template<typename SrcIter, typename DestIter,
14         typename RemoteCompletionFunc,
15         typename ...RemoteCompletionArgs>
16 void rput_fragmented_then_rpc(
17     SrcIter src_runs_begin, SrcIter src_runs_end,
18     DestIter dest_runs_begin, DestIter dest_runs_end,
19     promise<> &source_completion,
20     intrank_t recipient,
21     RemoteCompletionFunc &&remote_completion_func,
22     RemoteCompletionArgs &&...remote_completion_args);
23
24 // continuation variant
25 template<typename SrcIter, typename DestIter,
26         typename SourceCompletionFunc,
27         typename RemoteCompletionFunc,
28         typename ...RemoteCompletionArgs>
29 void rput_fragmented_then_rpc(
30     SrcIter src_runs_begin, SrcIter src_runs_end,
31     DestIter dest_runs_begin, DestIter dest_runs_end,
32
33     persona &source_completion_recipient,
34     SourceCompletionFunc source_completion_func,
35
36     intrank_t recipient,
37     RemoteCompletionFunc &&remote_completion_func,
38     RemoteCompletionArgs &&...remote_completion_args);
```

```
39     Preconditions: Same as those in rput_fragmented with the addition that
40     RemoteCompletionFunc be Serializable and a function-object type and
41     that all RemoteCompletionArgs be Serializable, or dist_object<U>&, or
42     team&. The calls remote_completion_func(remote_completion_args...)
```

1 and `source_completion_func()` must not throw an exception. All remote  
2 memory referenced by the destination address sequence has affinity with rank  
3 `recipient`.

4 Performs the same series of puts as in `rput_fragmented`. Completion of the  
5 operation triggers a `rpc_ff` consisting of `remote_completion_func` invoked  
6 against `remote_completion_args` to the `recipient` rank. The current rank  
7 does not have access to this completion event. Instead, the current rank is  
8 notified of source completion. Source completion indicates only that the source  
9 and destination memory address sequences and source memory contents can be  
10 reclaimed. Source completion does not indicate the puts have become visible.

11 Serialization of `remote_completion_func` and `remote_completion_args` hap-  
12 pen during the function call.

13 In the future variant, source completion is returned as a future.

14 In the promise variant, an anonymous dependency is added to the promise  
15 during the call and is fulfilled upon source completion.

16 In the continuation variant, the `source_completion_func` function object is  
17 submitted to `source_completion_recipient`'s user-progress continuation queue  
18 upon source completion.

19 *C++ memory ordering:* All evaluations *sequenced-before* this call and the  
20 puts from this call will have a *happens-before* relationship with the invoca-  
21 tion of `remote_completion_func`. In the continuation variant, all evaluations  
22 *sequenced-before* this call will have a *happens-before* relationship with the invo-  
23 cation of `source_completion_func`.

24 *UPC++ progress level:* internal

## 25 13.2.4 Fragmented Regular Put

```
26 // future variant
27 template<typename SrcIter, typename DestIter>
28 future<> rput_fragmented_regular(
29     SrcIter src_runs_begin, SrcIter src_runs_end,
30     std::size_t src_run_length,
31     DestIter dest_runs_begin, DestIter dest_runs_end,
32     std::size_t dest_run_length);
33
34 // promise variant
35 template<typename SrcIter, typename DestIter>
36 void rput_fragmented_regular(
```



```
1  SrcIter src_runs_begin, SrcIter src_runs_end,
2  std::size_t src_run_length,
3  DestIter dest_runs_begin, DestIter dest_runs_end,
4  std::size_t dest_run_length,
5  promise<> &completion);
6
7  // continuation variant
8  template<typename SrcIter, typename DestIter,
9          typename CompletionFunc>
10 void rput_fragmented_regular(
11     SrcIter src_runs_begin, SrcIter src_runs_end,
12     std::size_t src_run_length,
13     DestIter dest_runs_begin, DestIter dest_runs_end,
14     std::size_t dest_run_length,
15     persona &completion_recipient,
16     CompletionFunc completion_func);
```

*Preconditions:*

SrcIter and DestIter both satisfy the ForwardIterator C++ concept.

\*std::declval<SrcIter>() has a type convertible to T const\*, for some type T.

\*std::declval<DestIter>() has the type global\_ptr<T>, for the same type T as with SrcIter.

All destination addresses must be global\_ptr<T>'s referencing memory with affinity to the same rank.

The length of the two sequences delimited by (src\_runs\_begin, src\_runs\_end) and (dest\_runs\_begin, dest\_runs\_end) multiplied by (src\_run\_length, dest\_run\_length) respectively must be the same.

CompletionFunc is a function-object type.

These calls have the same semantics as their rput\_fragmented counterparts with the difference that, for each sequence, all run lengths are the same and are factored out of the sequences into two extra parameters src\_run\_length and dest\_run\_length. Thus the iterated elements are no longer pairs, but just pointers (the first pair component).

The sequence iterators must remain valid, and the underlying addresses and source memory contents must stay constant until completion is signaled. Only after completion is signaled can the address sequences and source memory be reclaimed by the application.

In the future variant, completion is returned as a future.

1 In the promise variant, an anonymous dependency is added to the promise  
2 during the call and is fulfilled upon completion.

3 In the continuation variant, the `completion_func` function object is submitted  
4 to `completion_recipient`'s user-progress continuation queue upon comple-  
5 tion.

6 *C++ memory ordering:* In the continuation variant, all evaluations *sequenced-*  
7 *before* this call and the puts from this call will have a *happens-before* relationship  
8 with the invocation of `completion_func`.

9 *UPC++ progress level:* internal

### 10 13.2.5 Fragmented Regular Get

```
11 // future variant
12 template<typename SrcIter, typename DestIter>
13 future<> rget_fragmented_regular(
14     SrcIter src_runs_begin, SrcIter src_runs_end,
15     std::size_t src_run_length,
16     DestIter dest_runs_begin, DestIter dest_runs_end,
17     std::size_t dest_run_length);
18
19 // promise variant
20 template<typename SrcIter, typename DestIter>
21 void rget_fragmented_regular(
22     SrcIter src_runs_begin, SrcIter src_runs_end,
23     std::size_t src_run_length,
24     DestIter dest_runs_begin, DestIter dest_runs_end,
25     std::size_t dest_run_length,
26     promise<> &completion);
27
28 // continuation variant
29 template<typename SrcIter, typename DestIter,
30         typename CompletionFunc>
31 void rget_fragmented_regular(
32     SrcIter src_runs_begin, SrcIter src_runs_end,
33     std::size_t src_run_length,
34     DestIter dest_runs_begin, DestIter dest_runs_end,
35     std::size_t dest_run_length,
36     persona &completion_recipient,
37     CompletionFunc completion_func);
```

38 *Preconditions:*

1        `SrcIter` and `DestIter` both satisfy the `ForwardIterator` C++ concept.  
2        `*std::declval<DestIter>()` has a type convertible to `T*`, for some type  
3        `T`.  
4        `*std::declval<SrcIter>()` has the type `global_ptr<T>`, for the same  
5        type `T` as with `DestIter`.  
6        All source addresses must be `global_ptr<T>`'s referencing memory with  
7        affinity to the same rank.  
8        The length of the two sequences delimited by `(src_runs_begin, src_runs_end)`  
9        and `(dest_runs_begin, dest_runs_end)` multiplied by `(src_run_length,`  
10        `dest_run_length)` respectively must be the same.  
11        `CompletionFunc` is a function-object type.

12        These calls have the same semantics as their `rget_fragmented` counterparts  
13        with the difference that, for both sequences, all run lengths are the same and  
14        are factored out of the sequences into two extra parameters `src_run_length`  
15        and `dest_run_length`. Thus the iterated elements are no longer pairs, but just  
16        pointers (the first component).

17        The sequence iterators must remain valid, and the underlying addresses and  
18        source memory contents must stay constant until completion is signaled. Only  
19        after completion is signaled can the address sequences and source memory be  
20        reclaimed by the application.

21        In the future variant, completion is returned as a future.

22        In the promise variant, an anonymous dependency is added to the promise  
23        during the call and is fulfilled upon completion.

24        In the continuation variant, the `completion_func` function object is submitted  
25        to `completion_recipient`'s user-progress continuation queue upon comple-  
26        tion.

27        *C++ memory ordering:* In the continuation variant, all evaluations *sequenced-*  
28        *before* this call and the gets from this call will have a *happens-before* relationship  
29        with the invocation of `completion_func`.

30        *UPC++ progress level:* `internal`

## 31 13.2.6 Fragmented Regular Put then RPC

```
32 // future variant
33 template<typename SrcIter, typename DestIter>
34 future<> rput_fragmented_regular_then_rpc(
```

```
1   SrcIter src_runs_begin, SrcIter src_runs_end,
2   std::size_t src_run_length,
3   DestIter dest_runs_begin, DestIter dest_runs_end,
4   std::size_t dest_run_length,
5   intrank_t recipient,
6   RemoteCompletionFunc &&remote_completion_func,
7   RemoteCompletionArgs &&...remote_completion_args);
8
9 // promise variant
10 template<typename SrcIter, typename DestIter,
11         typename RemoteCompletionFunc,
12         typename ...RemoteCompletionArgs>
13 void rput_fragmented_regular_then_rpc(
14     SrcIter src_runs_begin, SrcIter src_runs_end,
15     std::size_t src_run_length,
16     DestIter dest_runs_begin, DestIter dest_runs_end,
17     std::size_t dest_run_length,
18     intrank_t recipient,
19     RemoteCompletionFunc &&remote_completion_func,
20     RemoteCompletionArgs &&...remote_completion_args);
21
22 // continuation variant
23 template<typename SrcIter, typename DestIter,
24         typename SourceCompletionFunc,
25         typename RemoteCompletionFunc,
26         typename ...RemoteCompletionArgs>
27 void rput_fragmented_regular_then_rpc(
28     SrcIter src_runs_begin, SrcIter src_runs_end,
29     std::size_t src_run_length,
30     DestIter dest_runs_begin, DestIter dest_runs_end,
31     std::size_t dest_run_length,
32
33     persona &source_completion_recipient,
34     SourceCompletionFunc source_completion_func,
35
36     intrank_t recipient,
37     RemoteCompletionFunc &&remote_completion_func,
38     RemoteCompletionArgs &&...remote_completion_args);
39
40     Preconditions: Same as those in rput_fragmented_regular with the addi-
41     tion that RemoteCompletionFunc be Serializable and a function-object type
42     and that all RemoteCompletionArgs be Serializable, or dist_object<U>&, or
43     team&. The calls remote_completion_func(remote_completion_args...)
```

1 and `source_completion_func()` must not throw an exception. All mem-  
2 ory referenced in the destination address sequence must have affinity with the  
3 **recipient** rank.

4 Performs the same series of puts as in `rput_fragmented_regular`. Completion  
5 of the operation triggers a `rpc_ff` consisting of `remote_completion_func` in-  
6 voked against `remote_completion_args` to the **recipient** rank. The current  
7 rank does not have access to this completion event. Instead, the current rank is  
8 notified of source completion. Source completion indicates only that the source  
9 and destination memory address sequences and source memory contents can be  
10 reclaimed. Source completion does not indicate the puts have become visible.

11 Serialization of `remote_completion_func` and `remote_completion_args` hap-  
12 pen during the function call.

13 In the future variant, source completion is returned as a future.

14 In the promise variant, an anonymous dependency is added to the promise  
15 during the call and is fulfilled upon source completion.

16 In the continuation variant, the `source_completion_func` function object is  
17 submitted to `source_completion_recipient`'s user-progress continuation queue  
18 upon source completion.

19 *C++ memory ordering:* All evaluations *sequenced-before* this call and the  
20 puts from this call will have a *happens-before* relationship with the invoca-  
21 tion of `remote_completion_func`. In the continuation variant, all evaluations  
22 *sequenced-before* this call will have a *happens-before* relationship with the invo-  
23 cation of `source_completion_func`.

24 *UPC++ progress level:* `internal`

## 25 13.2.7 Strided Put

```
26 // future variant
27 template<typename T, int Dim>
28 future<> rput_strided(
29     T const *src_base,
30     std::ptrdiff_t const *src_strides,
31     global_ptr<T> dest_base,
32     std::ptrdiff_t const *dest_strides,
33     std::size_t const *extents);
34
35 // promise variant
36 template<typename T, int Dim>
```

```
1 void rput_strided(  
2     T const *src_base,  
3     std::ptrdiff_t const *src_strides,  
4     global_ptr<T> dest_base,  
5     std::ptrdiff_t const *dest_strides,  
6     std::size_t const *extents,  
7     promise<> &completion);  
8  
9 // continuation variant  
10 template<typename T, int Dim, typename CompletionFunc>  
11 void rput_strided(  
12     T const *src_base,  
13     std::ptrdiff_t const *src_strides,  
14     global_ptr<T> dest_base,  
15     std::ptrdiff_t const *dest_strides,  
16     std::size_t const *extents,  
17     persona &completion_recipient,  
18     CompletionFunc completion_func);
```

19 *Precondition:* T must be a Serializable type. Dim must be non-negative. All  
20 source addresses and destination global pointers must reference valid objects  
21 of type T. Each of `src_strides[i]`, `dest_strides[i]`, and `extents[i]` must  
22 be valid objects of their respective pointed-to type for all  $0 \leq i < \text{Dim}$ .

23 If `Dim == 0`, `src_strides`, `dest_strides`, and `extents` are ignored, and the  
24 data movement performed is equivalent to `rput(src_base, dest_base, 1)`.

25 Otherwise, performs the semantic equivalent of many put's of type T. Let the  
26 *index space* be the set of integer vectors of dimension Dim in the bounding box  
27 with the inclusive lower bound at the all-zero origin, and the exclusive upper  
28 bound equal to `extents`. For each index vector `index` in the index space, there  
29 will be a put with source and destination addresses computed as:

```
30 // "dot" is the vector dot product.  
31 // Pointer arithmetic is done in bytes, not elements of T.  
32 // "dest_base" is a global_ptr, following syntax is  
33 // pseudo-code.  
34 src_address = src_base + dot(index, src_strides)  
35 dest_address = dest_base + dot(index, dest_strides)
```

36 The destination memory regions must be completely disjoint and must not over-  
37 lap with any source memory regions, otherwise behavior is undefined. Source  
38 regions are permitted to overlap with each other.

1 The contents of the source addresses must remain valid and constant until  
2 completion is signaled.

3 In the future variant, completion is returned as a future.

4 In the promise variant, an anonymous dependency is added to the promise  
5 during the call and is fulfilled upon completion.

6 In the continuation variant, the `completion_func` function object is submitted  
7 to `completion_recipient`'s user-progress continuation queue upon comple-  
8 tion.

9 *C++ memory ordering:* In the continuation variant, all evaluations *sequenced-*  
10 *before* this call and the puts from this call will have a *happens-before* relationship  
11 with the invocation of `completion_func`.

12 *UPC++ progress level:* `internal`

### 13 13.2.8 Strided Get

```
14 // future variant
15 template<typename T, int Dim>
16 future<> rget_strided(
17     global_ptr<T> src_base,
18     std::ptrdiff_t const *src_strides,
19     T *dest_base,
20     std::ptrdiff_t const *dest_strides,
21     std::size_t const *extents);
22
23 // promise variant
24 template<typename T, int Dim>
25 void rget_strided(
26     global_ptr<T> src_base,
27     std::ptrdiff_t const *src_strides,
28     T *dest_base,
29     std::ptrdiff_t const *dest_strides,
30     std::size_t const *extents,
31     promise<> &completion);
32
33 // continuation variant
34 template<typename T, int Dim, typename CompletionFunc>
35 void rget_strided(
36     global_ptr<T> src_base,
37     std::ptrdiff_t const *src_strides,
```

```
1  T *dest_base ,
2  std::ptrdiff_t const *dest_strides ,
3  std::size_t const *extents ,
4  persona &completion_recipient ,
5  CompletionFunc completion_func);
```

6 *Precondition:* T must be a Serializable type. Dim must be non-negative. All  
7 source global pointers and destination addresses must reference valid objects  
8 of type T. Each of `src_strides[i]`, `dest_strides[i]`, and `extents[i]` must  
9 be valid objects of their respective pointed-to type for all  $0 \leq i < \text{Dim}$ .

10 If `Dim == 0`, `src_strides`, `dest_strides`, and `extents` are ignored, and the  
11 data movement performed is equivalent to `rget(src_base, dest_base, 1)`.

12 Otherwise, performs the reverse direction of `rput_strided` where now the  
13 source memory is remote and the destination is local.

14 The destination memory regions must be completely disjoint and must not over-  
15 lap with any source memory regions, otherwise behavior is undefined. Source  
16 regions are permitted to overlap with each other.

17 The contents of the source addresses must remain valid and constant until  
18 completion is signaled.

19 In the future variant, completion is returned as a future.

20 In the promise variant, an anonymous dependency is added to the promise  
21 during the call and is fulfilled upon completion.

22 In the continuation variant, the `completion_func` function object is submitted  
23 to `completion_recipient`'s user-progress continuation queue upon comple-  
24 tion.

25 *C++ memory ordering:* In the continuation variant, all evaluations *sequenced-*  
26 *before* this call and the gets from this call will have a *happens-before* relationship  
27 with the invocation of `completion_func`.

28 *UPC++ progress level:* internal

### 29 13.2.9 Strided Put then RPC

```
30 // future variant
31 template<typename T, int Dim,
32         typename RemoteCompletionFunc ,
33         typename ...RemoteCompletionArgs >
34 future<> rput_strided_then_rpc (
35     T const *src_base ,
```



```
1  std::ptrdiff_t const *src_strides ,
2  global_ptr<T> dest_base ,
3  std::ptrdiff_t const *dest_strides ,
4  std::size_t const *extents ,
5  RemoteCompletionFunc &&remote_completion_func ,
6  RemoteCompletionArgs &&...remote_completion_args);
7
8  // promise variant
9  template<typename T, int Dim,
10         typename RemoteCompletionFunc ,
11         typename ...RemoteCompletionArgs>
12 void rput_strided_then_rpc(
13     T const *src_base ,
14     std::ptrdiff_t const *src_strides ,
15     global_ptr<T> dest_base ,
16     std::ptrdiff_t const *dest_strides ,
17     std::size_t const *extents ,
18     promise<> &source_completion ,
19     RemoteCompletionFunc &&remote_completion_func ,
20     RemoteCompletionArgs &&...remote_completion_args);
21
22 // continuation variant
23 template<typename T, int Dim,
24         typename SourceCompletionFunc ,
25         typename RemoteCompletionFunc ,
26         typename ...RemoteCompletionArgs>
27 void rput_strided_then_rpc(
28     T const *src_base ,
29     std::ptrdiff_t const *src_strides ,
30     global_ptr<T> dest_base ,
31     std::ptrdiff_t const *dest_strides ,
32     std::size_t const *extents ,
33
34     persona &source_completion_recipient ,
35     SourceCompletionFunc source_completion_func ,
36
37     RemoteCompletionFunc &&remote_completion_func ,
38     RemoteCompletionArgs &&...remote_completion_args);
```

```
39     Preconditions: Same as those in rput_strided with the addition that
40     RemoteCompletionFunc be Serializable and a function-object type and
41     that all RemoteCompletionArgs be Serializable, or dist_object<U>&, or
42     team&. The calls remote_completion_func(remote_completion_args...)
```

1 and `source_completion_func()` must not throw an exception. Either  
2 `dest_base` or `dest_base-1` must reference a valid object of type T.

3 Performs the same series of puts as `rput_strided` except for the completion  
4 semantics. Upon completion of the puts, the rank `dest_base.where()` is deliv-  
5 ered an `rpc_ff` of `remote_completion_func` invoked against `remote_completion_args`.  
6 Source completion is returned to the caller. Source completion signals that  
7 memory referenced by the source addresses may now be modified or reclaimed,  
8 it does not indicate completion of the puts.

9 Serialization of `remote_completion_func` and `remote_completion_args` oc-  
10 cur during this call.

11 In the future variant, source completion is returned as a future.

12 In the promise variant, an anonymous dependency is added to the promise  
13 during the call and is fulfilled upon source completion.

14 In the continuation variant, the `source_completion_func` function object is  
15 submitted to `source_completion_recipient`'s user-progress continuation queue  
16 upon source completion.

17 *C++ memory ordering:* All evaluations *sequenced-before* this call and the  
18 puts from this call will have a *happens-before* relationship with the invoca-  
19 tion of `remote_completion_func`. In the continuation variant, all evaluations  
20 *sequenced-before* this call will have a *happens-before* relationship with the invo-  
21 cation of `source_completion_func`.

22 *UPC++ progress level:* `internal`

# 1 Chapter 14

## 2 Memory Kinds

3 The memory kinds interface enables the programmer to identify regions of memory requir-  
4 ing different access methods or having different performance properties, and subsequently  
5 rely on the UPC++ communication services to perform transfers among such regions (both  
6 local and remote) in a manner transparent to the programmer. With GPU devices, HBM,  
7 scratch-pad memories, NVRAM and various types of storage-class and fabric-attached  
8 memory technologies featured in vendors' public road maps, UPC++ must be prepared to  
9 deal efficiently with data transfers among all the memory technologies in any given system.

10 Since memory kinds will be implemented in Year 2, we defer detailed discussion until  
11 next year.

# 1 Appendix A

## 2 Notes for Implementers

3 The following are possible implementations of template metaprogramming utilities for  
4 UPC++ features.

### 5 A.1 future\_element\_t and future\_element\_moved\_t

```
6 template<int I, typename T>
7 struct future_element; // undefined
8
9 template<int I, typename T, typename ...U>
10 struct future_element<I, future<T, U...>> {
11     typedef typename future_element<I-1, future<U...>>::type type;
12     typedef typename future_element<I-1, future<U...>>::moved_type
13         moved_type;
14 };
15
16 template<typename T, typename ...U>
17 struct future_element<0, future<T, U...>> {
18     typedef T type;
19     typedef T&& moved_type;
20 };
21
22 template<int I>
23 struct future_element<I, future<>> {
24     typedef void type;
25     typedef void moved_type;
26 };
27
```

```

1  template<int I, typename T>
2  using future_element_t = typename future_element<I, T>::type;
3
4  template<int I, typename T>
5  using future_element_moved_t =
6  typename future_element<I, T>::moved_type;

```

## 7 A.2 future<T...>::when\_all

8 Utility types:

```

9  template<template<typename ...Us> class T, typename A, typename B>
10 struct concat_type; // undefined
11
12 template<template<typename ...Us> class T,
13         typename ...As, typename... Bs>
14 struct concat_type<T, T<As...>, T<Bs...> > {
15     typedef T<As..., Bs...> type;
16 };
17
18 template<template<typename ...Us> class T,
19         typename A, typename... Bs>
20 struct concat_element_types {
21     typedef typename concat_element_types<T, Bs...>::type rest;
22     typedef typename concat_type<T, A, rest>::type type;
23 };
24
25 template<template<typename ...Us> class T, typename A>
26 struct concat_element_types<T, A> {
27     typedef A type;
28 };
29
30 template<template<typename ...Us> class T, typename ...U>
31 using concat_element_types_t =
32     typename concat_element_types<T, U...>::type;

```

33 Declaration of future<T...>::when\_all:

```

34 template<typename ...Futures>
35 concat_element_types_t<future, Futures...> when_all(Futures ...fs);

```

## 1 **A.3 to\_future**

2 Utility types:

```
3 template<typename T>  
4 struct future_type {  
5     typedef future<T> type;  
6 };  
7  
8 template<typename ...T>  
9 struct future_type<future<T...>> {  
10     typedef future<T...> type;  
11 };  
12  
13 template<>  
14 struct future_type<void> {  
15     typedef future<> type;  
16 };  
17  
18 template<typename T>  
19 using future_type_t = typename future_type<T>::type;  
20  
21 template<typename ...T>  
22 using future_types_t =  
23     concat_element_types_t<future, future_type_t<T>...>;
```

24 Declaration of to\_future:

```
25 template<typename ...U>  
26 future_types_t<U...> to_future(U ...futures_or_results);
```

## 27 **A.4 future\_invoke\_result\_t**

28 C++11-compliant implementation:

```
29 template<typename Func, typename... ArgTypes>  
30 using future_invoke_result_t =  
31     future_type_t<typename std::result_of<Func(ArgTypes...)>::type>;
```

32 C++17-compliant implementation:

```
33 template<typename Func, typename... ArgTypes>  
34 using future_invoke_result_t =  
35     future_type_t<std::invoke_result_t<Func, ArgTypes...>>;
```

## 1 **A.5** `promise_invoke_result_t`

2 Utility types:

```
3 template<typename T>  
4 struct promise_type {  
5   typedef promise<T> type;  
6 };  
7  
8 template<typename ...T>  
9 struct promise_type<future<T...>> {  
10   typedef promise<T...> type;  
11 };  
12  
13 template<>  
14 struct promise_type<void> {  
15   typedef promise<> type;  
16 };  
17  
18 template<typename T>  
19 using promise_type_t = typename promise_type<T>::type;
```

20 C++11-compliant implementation:

```
21 template<typename Func, typename... ArgTypes>  
22 using promise_invoke_result_t =  
23   promise_type_t<typename std::result_of<Func(ArgTypes...)>::type>;
```

24 C++17-compliant implementation:

```
25 template<typename Func, typename... ArgTypes>  
26 using promise_invoke_result_t =  
27   promise_type_t<std::invoke_result_t<Func, ArgTypes...>>;
```

# 1 Bibliography

- 2 [1] ISO. *ISO/IEC 14882:2011(E) Information technology - Programming Languages -*  
3 *C++*. Geneva, Switzerland, 2012.
- 4 [2] ISO. *ISO/IEC 14882:2014(E) Information technology - Programming Languages -*  
5 *C++, working draft*. Geneva, Switzerland, 2014.
- 6 [3] Y. Zheng, A. Kamil, M. B. Driscoll, H. Shan, and K. Yelick. UPC++: A PGAS ex-  
7 tension for C++. In *2014 IEEE 28th International Parallel and Distributed Processing*  
8 *Symposium*, pages 1105–1114, May 2014.