

XcalableMP
<ex-scalable-em-p>
Language Specification

Version 1.4

XcalableMP Specification Working Group

November, 2018

Copyright ©2008-2018 XcalableMP Specification Working Group. Permission to copy without fee all or part of this material is granted, provided the XcalableMP Specification Working Group copyright notice and the title of this document are displayed. Notice is given that copying requires the explicit permission of the XcalableMP Specification Working Group.

History

Version 1.4: November, 2018

- 4.1.2 Combined Directive
- F DRAFT: Tasklet of upcoming XcalableMP 2.0

Version 1.3: August, 2017

- 2.8.2 Node Terminology
- 2.8.6 Local-view Terminology
- 3.1 Array Section Notation
- 3.2 Array assignment statements in C
- 3.5 Dynamic Allocation of Global Data in C
- 4.2 `nodes` Directive
- 4.3.1 `template` Directive
- 4.3.4 `align` Directive
- 4.4.3 `loop` Construct
- 4.5.1 `reflect` Construct
- 4.5.6 `wait_async` Construct
- 4.5.8 `reduce_shadow` Construct
- 5.7 Coarrays in XcalableMP C
- 5.8 Directives for the Local-view Programming
- 7.2.2 `xmpc_all_node_num`
- 7.2.5 `xmpc_node_num`
- 7.2.6 `xmpc_this_image`
- 7.2.8 `xmp_num_images`
- 7.3 Execution Control Functions
- 7.5.1 `xmp_malloc`
- 7.6.4 `xmp_nodes_attr`
- B.2.1 `xmp_array_gtol`
- B.2.2 `xmp_array_lsize`
- B.2.4 `xmp_array_lda`
- 7.9.1 `xmp_scatter`
- 7.9.3 `xmp_pack`
- 7.9.4 `xmp_unpack`
- C Memory-layout Model

Version 1.2.1: November, 2014 Corrections and clarifications to Version 1.2.

- 4.3.4 `align` Directive
- 4.4.3 `loop` Directive
- 5.7.1 [C] Declarations of Coarrays
- 5.8.1 [F] `local_alias` Directive
- 3.1 Array Section Notation
- 3.2 Array Assignment Statement

Version 1.2: November 20, 2013

Version 1.1: November 13, 2012

Version 1.0: November 14, 2011

Contents

1	Introduction	1
1.1	Features of XcalableMP	1
1.2	Scope	2
1.3	Organization of this Document	2
1.4	Changes to Version 1.4 from Version 1.3	2
1.5	Changes to Version 1.3 from Version 1.2.1	2
1.6	Changes to Version 1.1 from Version 1.2	3
2	Overview of the XcalableMP Model and Language	5
2.1	Hardware Model	5
2.2	Execution Model	5
2.3	Data Model	6
2.4	Global-view Programming Model	6
2.5	Local-view Programming Model	7
2.6	Interactions between Global View and Local View	8
2.7	Base Languages	8
2.8	Glossary	9
2.8.1	Language Terminology	9
2.8.2	Node Terminology	11
2.8.3	Data Terminology	12
2.8.4	Work Terminology	12
2.8.5	Communication and Synchronization Terminology	12
2.8.6	Local-view Terminology	13
3	Base Language Extensions in XcalableMP C	15
3.1	Array Section Notation	15
3.2	Array Assignment Statement	16
3.3	Built-in Functions for Array Section	17
3.4	Pointer to Global Data	17
3.4.1	Name of Global Array	17
3.4.2	Address-of Operator	17
3.5	Dynamic Allocation of Global Data	17
3.6	Descriptor-of Operator	18
4	Directives	19
4.1	Directive Format	19
4.1.1	General Rule	19
4.1.2	Combined Directive	21
4.2	<code>nodes</code> Directive	22
4.2.1	Node Reference	23

4.3	Template and Data Mapping Directives	24
4.3.1	<code>template</code> Directive	24
4.3.2	Template Reference	25
4.3.3	<code>distribute</code> Directive	26
4.3.4	<code>align</code> Directive	29
4.3.5	<code>shadow</code> Directive	31
4.3.6	<code>template.fix</code> Construct	33
4.4	Work Mapping Construct	35
4.4.1	<code>task</code> Construct	35
4.4.2	<code>tasks</code> Construct	36
4.4.3	<code>loop</code> Construct	38
4.4.4	<code>array</code> Construct	47
4.5	Global-view Communication and Synchronization Constructs	48
4.5.1	<code>reflect</code> Construct	48
4.5.2	<code>gmove</code> Construct	50
4.5.3	<code>barrier</code> Construct	52
4.5.4	<code>reduction</code> Construct	52
4.5.5	<code>bcast</code> Construct	55
4.5.6	<code>wait_async</code> Construct	56
4.5.7	<code>async</code> Clause	57
4.5.8	<code>reduce_shadow</code> Construct	57
5	Support for the Local-view Programming	61
5.1	Rules Determining Image Index	61
5.1.1	Primary Image Index	61
5.1.2	Image Index Determined by a <code>task</code> Directive	62
5.1.3	Current Image Index	62
5.1.4	Image Index Determined by a Non-primary Node Array	62
5.1.5	Image Index Determined by an Equivalenced Node Array	62
5.1.6	On-node Image Index	63
5.2	Basic Concepts	63
5.2.1	Examples	63
5.3	<code>coarray</code> Directive	64
5.3.1	Purpose and Form of the <code>coarray</code> Directive	64
5.3.2	An Example of the <code>coarray</code> Directive	65
5.4	<code>image</code> Directive	66
5.4.1	Purpose and Form of the <code>image</code> Directive	66
5.4.2	An Example of the <code>image</code> Directive	66
5.5	Image Index Translation Intrinsic Procedures	67
5.5.1	Translation to the Primary Image Index	67
5.5.2	Translation to the Current Image Index	68
5.6	Examples of Communication between Tasks	68
5.7	[C] Coarrays in XcalableMP C	71
5.7.1	[C] Declaration of Coarrays	71
5.7.2	[C] Reference of Coarrays	72
5.7.3	[C] Synchronization of Coarrays	72
5.8	Directives for the Local-view Programming	73
5.8.1	[F] <code>local_alias</code> Directive	73
5.8.2	<code>post</code> Construct	76
5.8.3	<code>wait</code> Construct	77

5.8.4	[C] lock/unlock Construct	78
6	Procedure Interfaces	81
6.1	General Rule	81
6.2	Argument Passing Mechanism in XcalableMP Fortran	81
6.2.1	Sequence Association of Global Data	82
6.2.2	Descriptor Association of Global Data	85
6.3	Argument-Passing Mechanism in XcalableMP C	88
7	Intrinsic and Library Procedures	93
7.1	Intrinsic Functions	93
7.1.1	xmp_desc_of	93
7.2	System Inquiry Functions	93
7.2.1	xmp_all_node_num	94
7.2.2	[C] xmpc_all_node_num	94
7.2.3	xmp_all_num_nodes	94
7.2.4	xmp_node_num	95
7.2.5	[C] xmpc_node_num	95
7.2.6	[C] xmpc_this_image	95
7.2.7	xmp_num_nodes	95
7.2.8	xmp_num_images	96
7.2.9	xmp_wtime	96
7.2.10	xmp_wtick	96
7.3	[C] Execution Control Functions	97
7.3.1	xmp_exit	97
7.4	Synchronization Functions	97
7.4.1	xmp_test_async	97
7.5	Memory Allocation Functions	97
7.5.1	[C] xmp_malloc	97
7.6	Mapping Inquiry Functions	98
7.6.1	xmp_nodes_ndims	98
7.6.2	xmp_nodes_index	98
7.6.3	xmp_nodes_size	99
7.6.4	xmp_nodes_attr	99
7.6.5	xmp_nodes_equiv	100
7.6.6	xmp_template_fixed	100
7.6.7	xmp_template_ndims	101
7.6.8	xmp_template_lbound	101
7.6.9	xmp_template_ubound	102
7.6.10	xmp_dist_format	102
7.6.11	xmp_dist_blocksize	103
7.6.12	xmp_dist_gblockmap	103
7.6.13	xmp_dist_nodes	104
7.6.14	xmp_dist_axis	104
7.6.15	xmp_align_axis	105
7.6.16	xmp_align_offset	105
7.6.17	xmp_align_replicated	106
7.6.18	xmp_align_template	106
7.6.19	xmp_array_ndims	106
7.6.20	xmp_array_lshadow	107

7.6.21	xmp_array_ushadow	107
7.6.22	xmp_array_lbound	108
7.6.23	xmp_array_ubound	108
7.7	[F] Array Intrinsic Functions of the Base Language	109
7.8	[C] Built-in Elemental Functions	109
7.9	Intrinsic/Built-in Transformational Procedures	109
7.9.1	xmp_scatter	110
7.9.2	xmp_gather	111
7.9.3	xmp_pack	111
7.9.4	xmp_unpack	112
7.9.5	xmp_transpose	112
7.9.6	xmp_matmul	112
7.9.7	xmp_sort_up	113
7.9.8	xmp_sort_down	113
8	OpenMP in XcalableMP Programs	115
	Bibliography	117
A	Programming Interface for MPI	119
A.1	Call MPI functions from an XcalableMP program	119
A.1.1	xmp_get_mpi_comm	119
A.1.2	xmp_init_mpi	119
A.1.3	xmp_finalize_mpi	120
A.2	Call XcalableMP functions from an MPI program	121
A.2.1	xmp_init	121
A.2.2	xmp_finalize	121
B	Interface to Numerical Libraries	125
B.1	Interface Design	125
B.2	Extended Mapping Inquiry Functions	125
B.2.1	xmp_array_gtol	126
B.2.2	xmp_array_lsize	126
B.2.3	xmp_array_laddr	127
B.2.4	xmp_array_lda	127
B.3	Example	127
C	Memory-layout Model	131
D	XcalableMP I/O	133
D.1	Categorization of I/O	133
D.1.1	Local I/O	133
D.1.2	Master I/O[F]	133
D.1.3	Global I/O	133
D.2	File Connection	134
D.2.1	File Connection in Local I/O	135
D.2.2	[F] File Connection in Master I/O	135
D.2.3	File Connection in Global I/O	135
D.3	Master I/O	135
D.3.1	master_io Construct	136
D.4	[F] Global I/O	137

D.4.1	Global I/O File Operation	138
D.4.2	Collective Global I/O Statement	140
D.4.3	Atomic Global I/O Statement	140
D.4.4	Direct Global I/O Statement	141
D.5	[C] Global I/O Library	141
D.5.1	Global I/O File Operation	144
D.5.2	Collective Global I/O Functions	146
D.5.3	Atomic Global I/O Functions	148
D.5.4	Direct Global I/O Functions	149
E	Memory Consistency Model	151
E.1	Execution Traces	151
E.1.1	Common Constraints	152
E.1.2	Constraints for Synchronous Communications	152
E.1.3	Constraints for Asynchronous Communications	152
F	DRAFT: Tasklet of upcoming XcalableMP 2.0	157
F.1	XcalableMP Extended Execution Model	157
F.2	Glossary	157
F.2.1	Node Terminology	157
F.2.2	Thread Terminology	157
F.2.3	Tasklet Terminology	157
F.3	Directives	158
F.3.1	Tasklet Constructs	158
F.3.2	Tasklet Scheduling	166
F.3.3	Communication Tasklet Constructs	166
G	Sample Programs	171

List of Figures

2.1	Hardware model.	5
2.2	Parallelization using the global-view programming model.	7
2.3	Local-view programming model.	8
2.4	Global view and local view.	9
4.1	Example showing shadow of a two-dimensional array.	33
4.2	Example of periodic shadow reflection.	49
6.1	Sequence association with a global dummy argument.	83
6.2	Sequence association with a local dummy argument.	84
6.3	Sequence association of a section of a global data object as an actual argument with a local dummy argument.	85
6.4	Sequence association of an element of a global data object as an actual argument with a local dummy argument.	86
6.5	Sequence association with a global dummy argument that has a full shadow.	86
6.6	Descriptor association with a global dummy argument.	88
6.7	Descriptor association with a local dummy argument.	89
6.8	Passing to a global dummy argument.	90
6.9	Passing to a local dummy argument.	91
6.10	Passing an element of a global data object as an actual argument to a local dummy argument.	91
B.1	Invocation of a library routine using an interface procedure.	125
C.1	Example of memory layout in the Omni XcalableMP compiler.	132
E.1	Constraints that are required by the XcalableMP memory consistency model.	152

List of Tables

7.1	Built-in elemental functions in XcalableMP C	110
D.1	Global I/O.	134
D.2	Operations for I/O.	137

Acknowledgment

The XcalableMP specification is designed by the XcalableMP Specification Working Group, which consists of the following members from academia, research laboratories, and industries.

- Tatsuya Abe Chiba Institute of Technology
- Tokuro Anzaki Hitachi
- Taisuke Boku University of Tsukuba
- Toshio Endo TITECH
- Yoshinari Fukui JAMSTEC
- Yasuharu Hayashi NEC
- Atsushi Hori RIKEN
- Kohichiro Hotta Fujitsu
- Hidetoshi Iwashita Fujitsu
- Susumu Komae AXE
- Atsushi Kubota Hiroshima City University
- Jinpil Lee RIKEN
- Toshiyuki Maeda Chiba Institute of Technology
- Motohiko Matsuda RIKEN
- Yuichi Matsuo JAXA
- Kazuo Minami RIKEN
- Shoji Morita AXE
- Hitoshi Murai RIKEN
- Kengo Nakajima University of Tokyo
- Takashi Nakamura JAXA
- Tomotake Nakamura Fujitsu
- Mamoru Nakano CRAY
- Masahiro Nakao RIKEN
- Takeshi Nanri Kyushu University
- Kiyoshi Negishi Hitachi
- Satoshi Ohshima Kyushu University
- Yasuo Okabe Kyoto University
- Hitoshi Sakagami NIFS
- Tomoko Sakari Fujitsu
- Shoich Sakon NEC
- Mitsuhisa Sato RIKEN
- Taizo Shimizu PC Cluster Consortium
- Takenori Shimosaka Hitachi
- Yoshihisa Shizawa RIST

- Shozo Takeoka AXE
- Hitoshi Uehara JAMSTEC
- Eiji Yamanaka Fujitsu
- Masahiro Yasugi Kyushu Institute of Technology
- Mitsuo Yokokawa Kobe University

This work was supported by “Seamless and Highly-productive Parallel Programming Environment for High-performance Computing” project funded by Ministry of Education, Culture, Sports, Science and Technology, Japan, and is supported by PC Cluster Consortium.

Chapter 1

Introduction

This document defines the XcalableMP specification, which is a directive-based language extension of Fortran and C for scalable and performance-aware parallel programming. The specification includes a collection of compiler directives and intrinsic and library procedures, and provides a model of parallel programming for distributed memory multiprocessor systems.

1.1 Features of XcalableMP

The features of XcalableMP are summarized as follows:

- XcalableMP supports typical parallelization based on the data-parallel paradigm and work mapping under the “global-view” programming model, and it enables the parallelization of the original sequential code using minimal modification with simple directives such as OpenMP [1]. Many ideas on “global-view” programming are inherited from High Performance Fortran (HPF) [2].
- The important design principle of XcalableMP is “performance-awareness.” All actions related to communication and synchronization are taken by directives (and coarray features), which is different from automatic parallelizing compilers. The user should be aware of the effect of the XcalableMP directives in the execution model for distributed-memory architecture.
- XcalableMP also includes features from Partitioned Global Address Space (PGAS) languages, such as coarray of the Fortran 2008 standard, for “local-view” programming.
- An extension of existing base languages with directives is useful to reduce code-rewriting and education costs. The XcalableMP language specification is defined as an extension to the Fortran and C base languages.
- For flexibility and extensibility, the execution model enables us to combine XcalableMP with explicit Message Passing Interface (MPI) [3] coding for more complicated and tuned parallel codes and libraries.
- For multi-core and SMP clusters, OpenMP directives can be combined into XcalableMP for thread programming inside each node as a hybrid programming model.

XcalableMP is being designed based on experiences gained during the development of HPF, HPF/JA [4], Fujitsu XPF (VPP FORTRAN) [5, 6], and OpenMPD [7].

1.2 Scope 1

The XcalableMP specification covers only user-directed parallelization, where the user explicitly specifies the behavior of the compiler and the runtime system in order to execute the program in parallel in a distributed-memory system. XcalableMP-compliant implementations are not required to automatically distribute data, detect parallelism, parallelize loops, or generate communications and synchronizations. 2
3
4
5
6

1.3 Organization of this Document 7

The remainder of this document is structured as follows: 8

- Chapter 2: Overview of the XcalableMP Model and Language 9
- Chapter 3: Base Language Extensions in XcalableMP C 10
- Chapter 4: Directives 11
- Chapter 5: Support for Local-view Programming 12
- Chapter 6: Procedure Interface 13
- Chapter 7: Intrinsic and Library Procedures 14
- Chapter 8: OpenMP in XcalableMP Programs 15

In addition, the following appendices are included in this document as proposals. 16

- Appendix A: Programming Interface for MPI 17
- Appendix B: Interface to Numerical Libraries 18
- Appendix C: Memory-layout Model 19
- Appendix D: XcalableMP I/O 20
- Appendix E: Memory Consistency Model 21
- Appendix F: DRAFT: Tasklet of upcoming XcalableMP 2.0 22

1.4 Changes to Version 1.4 from Version 1.3 23

- Combined directives in XcalableMP C are allowed. 24
- Add an appendix about the tasklet features of upcoming XcalableMP 2.0. 25

1.5 Changes to Version 1.3 from Version 1.2.1 26

- In XcalableMP C, a square bracket is available in *nodes-decl*, *nodes-ref*, *template-ref*, and *template-decl*. 27
28
- Add the `orthogonal` clause to the `reflect` directive in Section 4.5.1. 29
- Add `xmpc_all_node_num()` in Section 7.2.2. 30

- 1 • Add `xmp_node_num()` in Section 7.2.5.
- 2 • Add `xmp_this_image()` in Section 7.2.6.
- 3 • Add `xmp_num_images()` in Section 7.2.8.
- 4 • Modify `xmp_array_gtol()` in Section B.2.1.
- 5 • Change `xmp_array_lsize()` not to include shadow object in Section B.2.2.
- 6 • Create `xmp_array_lda()` from `xmp_array_lead_dim()` in Section B.2.4.
- 7 • In XcalableMP C, the dynamic allocataion of multi-dimensional global data is allowed.
- 8 • A restriction on the `align` directive is added.
- 9 • The `expand` and `margin` clauses of the `loop` construct are added.
- 10 • The meaning of a reduction-kind “-” in the `reduction` clause of the `loop` construct is
11 changed.
- 12 • The treatment for *async-id* not associated with any asynchronous communication is spec-
13 ified.
- 14 • The `reduce_shadow` construct is added.
- 15 • The description of the `local_alias` directive is modified.
- 16 • The `xmp_exit` library function is added.
- 17 • The specifications of `xmp_scatter`, `xmp_pack`, and `xmp_unpack` are modified.
- 18 • The memory consistency model of XcalableMP is discussed in the appendix.

19 1.6 Changes to Version 1.1 from Version 1.2

- 20 • The position of `align` directives for dummy arguments in XcalableMP C is specified.
- 21 • It is specified that aligned arrays cannot be initialized.
- 22 • The interpretation of a `reduction` clause of the `loop` directive is corrected.
- 23 • The syntax for declaring coarrays is changed.
- 24 • An assumed-shape array can be the target of the `local_alias` directive.
- 25 • The syntax and the semantics of the array section notation in XcalableMP C is modified.
- 26 • The syntax of the array assignment statement in XcalableMP C is extended.

Chapter 2

Overview of the XcalableMP Model and Language

2.1 Hardware Model

The target of XcalableMP is distributed-memory multicomputers (Figure 2.1). Each computation node, which may contain several cores, has its own local memory (shared by the cores, if any), and is connected with the others via an interconnection network. Each node can access its local memory directly and remote memory (the memory of another node) indirectly (i.e., through inter-node communication). However, it is assumed that accessing remote memory is much slower than accessing local memory.




Figure 2.1: Hardware model.

2.2 Execution Model

An XcalableMP program execution is based on the Single Program Multiple Data (SPMD) model, where each node starts execution from the same main routine, and continues to execute the same code independently (i.e., asynchronously), which is referred to as the *replicated execution*, until it encounters an XcalableMP construct.

A set of nodes that executes a procedure, statement, loop, a block, etc. is referred to as its *executing node set*, and is determined by the innermost `task`, `loop`, or `array` directive surrounding it dynamically, or at runtime. The *current executing node set* is an executing node set of the current context, which is managed by the XcalableMP runtime system on each node.

The current executing node set at the beginning of the program execution, or *entire node set*, is a node set that contains all the available nodes, which can be specified in an implementation-defined way (e.g., through a command-line option).

When a node encounters at runtime either a `loop`, `array`, or `task` construct, and is contained by the node set specified by the `on` clause of the directive, it updates the current executing node set with the specified one and executes the body of the construct, after which it resumes the last executing node set and proceeds to execute the subsequent statements.

In particular, when a node in the current executing node set encounters a `loop` or an `array` construct, it executes the loop or the array assignment in parallel with other nodes, so that each iteration of the loop or element of the assignment is independently executed by the node in which a specified data element resides.

When a node encounters a synchronization or a communication directive, synchronization or communication occurs between it and other nodes. That is, such *global constructs* are performed collectively by the current executing nodes. Note that neither synchronization nor communication occurs unless these constructs are being specified.

2.3 Data Model

There are two classes of data in XcalableMP: *global data* and *local data*. Data declared in an XcalableMP program are local by default.

Global data are distributed onto the executing node set by the `align` directive (see section 4.3.4). Each fragment of distributed global data is allocated in the local memory of a node in the executing node set.

Local data comprises all data that are not global. They are replicated within the local memory of each of the executing nodes.

A node can access directly only local data and sections of global data that reside in its local memory. To access data in remote memory, explicit communication must be specified in ways such as global communication constructs and coarray assignments.

In particular, in XcalableMP Fortran, for common blocks that include any global variables, it is implementation-defined what storage sequences they occupy and how storage association is defined between two of them.

2.4 Global-view Programming Model

The global-view programming model is useful when, starting from a sequential version of a program, the programmer parallelizes it in data-parallel style by adding directives with minimum modification. In the global-view programming model, the programmer describes the distribution of data among nodes using the data distribution directives. The `loop` construct assigns each iteration of a loop to the node at which the computed data is located. The global-view communication directives are used to synchronize nodes, maintain the consistency of shadow areas, and move sections of distributed data globally. Note that the programmer must specify explicitly communications to make all data references in the program local, and this is done using appropriate directives.

In many cases, the XcalableMP program according to the global-view programming model is based on a sequential program, and it can produce the same results, regardless of the number

1 of nodes (Figure 2.2).

2 There are three groups of directives for the global-view programming model. Because these
3 directives are ignored as a comment by the compilers of base languages (Fortran and C), an
4 XcalableMP program can be compiled by them to ensure that they run properly.

5 **Data Mapping**

6 Specifies the data distribution and mapping to nodes (partially inherited from HPF).

7 **Work Mapping (Parallelization)**

8 Assigns a work to a node set. The `loop` construct maps each iteration of a loop to nodes owning
9 a specific data elements. The `task` construct defines a set amount of work as a *task*, and assigns
10 it to a specific node set.

11 **Communication and Synchronization**

12 Specifies how to communicate and synchronize with the other compute nodes. In XcalableMP,
13 inter-node communication must be explicitly specified by the programmer. The compiler guar-
14 antees that no communication occurs unless it is explicitly specified by the programmer.

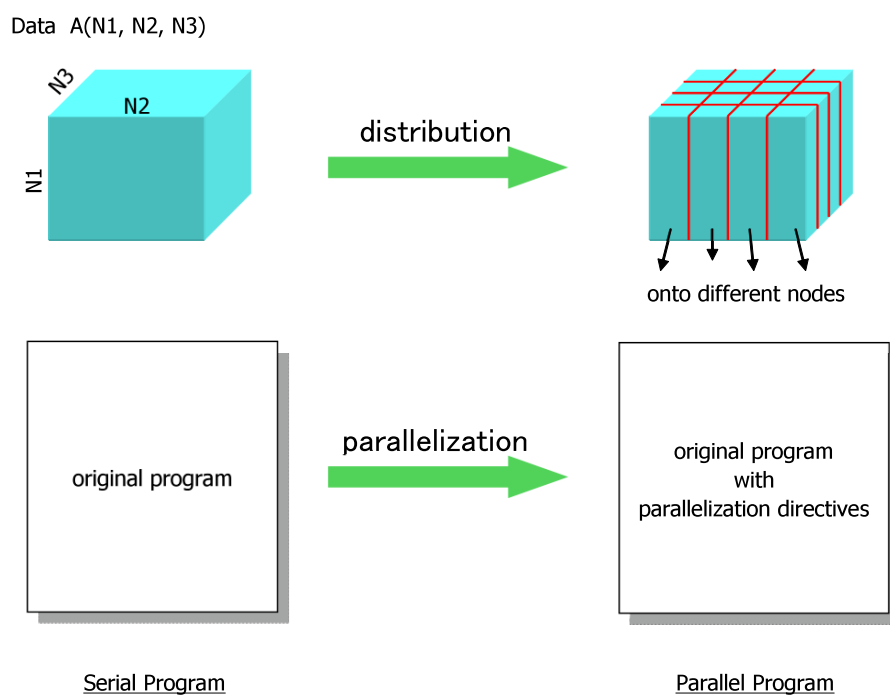


Figure 2.2: Parallelization using the global-view programming model.

15 **2.5 Local-view Programming Model**

16 The local-view programming model is suitable for programs that explicitly describe an algorithm
17 and a remote data reference that are to be executed by each node (Figure 2.3).

18 For the local-view programming model, some language extensions and directives are provided.
19 The coarray notation, which is imported from Fortran 2008, is one such extension, and can be

used to specify which replica of a local data is to be accessed. For example, the expression of $A(i)[N]$ is used to access an array element of $A(i)$ located on the node N . If the access is a reference, then a one-sided communication to get the value from the remote memory (i.e., the *get* operation) is issued by the executing node. If the access is a definition, then a one-sided communication to put a value to the remote memory (i.e., the *put* operation) is issued by the executing node.

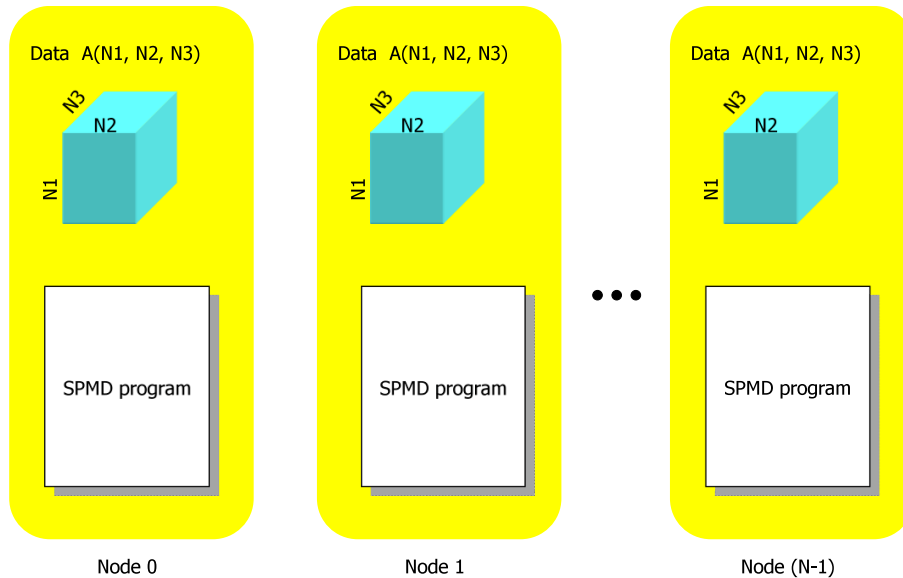


Figure 2.3: Local-view programming model.

2.6 Interactions between Global View and Local View

In the global view, nodes are used to distribute data and works. In the local view, nodes are used to address data in the coarray notation. In application programs, programmers should choose an appropriate data model according to the structure of the program. Figure 2.4 illustrates the global view and the local view of data.

Data may have both a global view and a local view, and can be accessed from either. XcalableMP provides some directives to give the local name (alias) to the global data declared in the global-view programming model to enable them to also be accessed in the local-view programming model. This feature is useful to optimize a certain part of the program by using explicit remote data access in the local-view programming model.

2.7 Base Languages

The XcalableMP language specification is defined based on Fortran and C as the base languages. More specifically, the base language of XcalableMP Fortran is Fortran 90 or later, and that of XcalableMP C is ISO C90 (ANSI C89) or later.

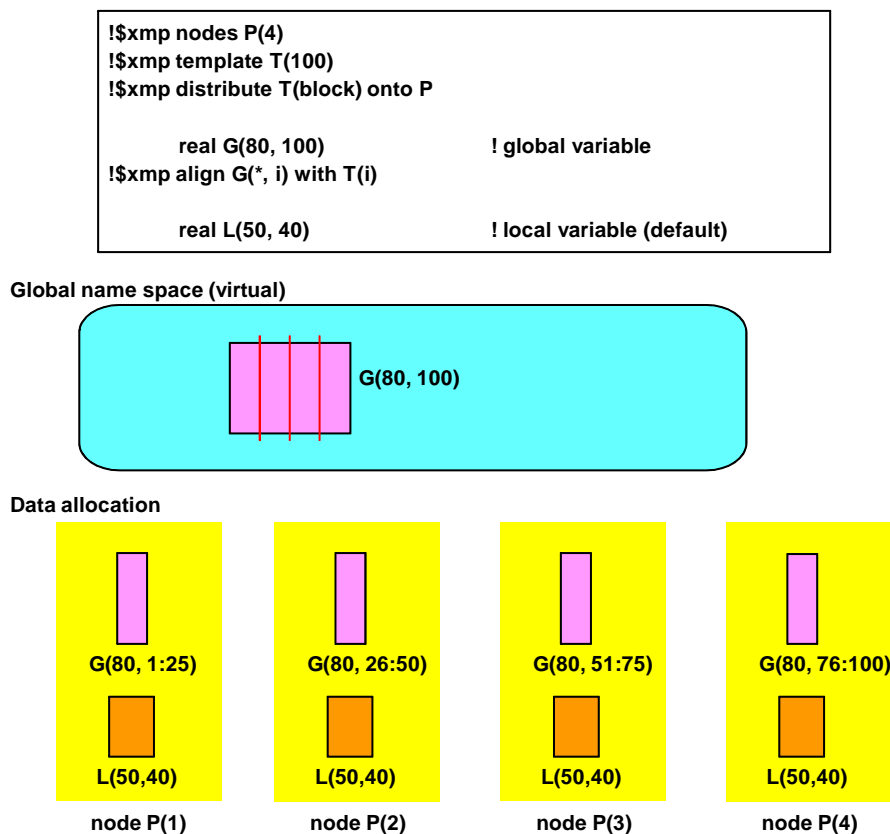


Figure 2.4: Global view and local view.

2.8 Glossary

2.8.1 Language Terminology

base language A programming language that serves as the foundation of the XcalableMP specification.

base program A program written in a base language.

XcalableMP

Fortran The XcalableMP specification for a base language Fortran, abbreviated as XMP/F.

XcalableMP C The XcalableMP specification for a base language C, abbreviated as XMP/C.

structured block For C, an executable statement, possibly compound, with a single entry at the top and a single exit at the bottom, or an XcalableMP construct. For Fortran, a block of executable statements with a single entry at the top and a single exit at the bottom, or an XcalableMP construct.

procedure A generic term used to refer to “procedure” (including subroutine and function) in XcalableMP Fortran and “function” in XcalableMP C.

directive In XcalableMP Fortran, a comment, and in XcalableMP C, a `#pragma`, that specifies XcalableMP program behavior.

declarative		
directive	An XcalableMP directive that may only be placed in a declarative context. A declarative directive has no associated executable user code; instead, it has one or more associated user declarations.	1 2 3
executable		
directive	An XcalableMP directive that is not declarative; it may be placed in an executable context.	4 5
construct	An XcalableMP executable directive (and for Fortran, the paired end directive, if any) and the associated statement, loop, or structured block, if any.	6 7 8
global construct	A construct that is executed collectively and synchronously by every node in the current executing node set. Global constructs are further classified into two groups of <i>global communication constructs</i> , such as gmove and barrier , which specify communication or synchronization, and <i>work mapping constructs</i> , such as loop , array , and tasks , which specify parallelization of loops, array assignments, or tasks.	9 10 11 12 13 14
template	A dummy array that represents an index space to be distributed onto a node set, which serves as the “template” of parallelization in XcalableMP and can be considered to represent a set of, for example, grid points in the grid method and particles in the particle method. A template is used in an XcalableMP program to specify the data and work mapping. Note that the lower bound of each dimension of a template is one in both XcalableMP Fortran and XcalableMP C.	15 16 17 18 19 20 21
data mapping	Allocating elements of an array to nodes in a node set by specifying with the align directive that the array is aligned with a distributed template.	22 23 24
work mapping	Assigning each of the iterations of a loop, the elements of an array assignment, or the tasks to nodes in a node set. Such work mapping is specified by aligning it with a template or distributing it onto a node set.	25 26 27 28
global	A data or a work is <i>global</i> if and only if there are one or more replicated instances of it, each of which is shared by the executing nodes.	29 30
local	A data or a work is <i>local</i> if and only if there is a replicated instance of it on each of the executing nodes.	31 32
global-view		
model	A programming or parallelization model on which parallel programs are written by specifying how to map global data and works onto nodes.	33 34 35
local-view model	A programming or parallelization model on which parallel programs are written by specifying how each node owns local data and performs local works.	36 37 38

1 2.8.2 Node Terminology

2 **node** An execution entity managed by the XcalableMP runtime system,
3 which has its own memory and can communicate with other nodes.
4 A node can execute one or more threads concurrently.

5 **node set** A totally ordered set of nodes.

6 **entire node set** A node set that contains all of the nodes participating in the exe-
7 cution of an XcalableMP program. It is the current executing node
8 set specified explicitly or implicitly at the beginning of the program
9 execution.

executing node

10 **set** A node set that contains all of the nodes participating in the execution
11 of a procedure, statement, construct, etc. of an XcalableMP program
12 is called its executing node set. In this document, this term is used to
13 represent the *current executing node set* unless it is ambiguous. The
14 executing node set at the beginning of the program execution is the
15 entire node set.

current executing node

16 **set** An executing node set of the current context, which is managed by the
17 XcalableMP runtime system. The current executing node set can be
18 modified by the **task**, **array**, or **loop** constructs.

19 **executing node** A node in the executing node set.

20 **node array** An XcalableMP entity of the same form as a Fortran array that rep-
21 resents a node set in XcalableMP programs. Each element of a node
22 array represents a node in the corresponding node set. A node array
23 is declared by the **nodes** directive. Note that the lower bound of each
24 dimension of a node array is one in both XcalableMP Fortran and
25 XcalableMP C.

entire node

26 **array** A node array corresponding to the entire node set. An entire node
27 array can be declared by a **NODES** directive without “=*node-ref*”.

executing node

28 **array** A node array corresponding to the executing node set. An execut-
29 ing node array corresponding to the procedure can be declared by a
30 **NODES** directive with the node reference “*”.

31 **parent node set** The parent node set of a node set is the last executing node set, which
32 encountered the innermost **task**, **loop**, or **array** construct that is be-
33 ing executed.

34 **node number** A unique number assigned to each node in a node set, which starts
35 from one and corresponds to its position within the node set that is
36 totally ordered.

2.8.3 Data Terminology

- variable** A named data storage block whose value can be defined and redefined during the execution of a program. Note that *variables* include array sections.
- global data** An array that is aligned with a template. Elements of a global data object are distributed onto nodes according to the distribution of the template. As a result, each node owns a part of the global data (called a *local section*), and can access it directly, but cannot access those on the other nodes.
- local data** Data that are not global. Each node owns a replica of a local data object, and can access it directly, but cannot access those on the other nodes. Note that the replicas of a local data object do not always have the same value.
- replicated data** Data whose storage is allocated on multiple nodes. A replicated data is either a local data object or a global data object replicated by an `align` directive.
- distribution** Assigning each element of a template to nodes in a node set in a specified manner. In a broad sense, it refers to assigning each element of an array, loop, etc.
- alignment** Associating each element of an array, loop, etc. with an element of the specified template. An element of the aligned array, loop, etc. is necessarily mapped to the same node as its associated element of the template.
- local section** A section of a global data object that is allocated as an array on each node at runtime. The local section of a global data object includes its shadow objects.
- shadow** An additional area of the local section of a distributed array, which is used to keep elements to be moved in from neighboring nodes.

2.8.4 Work Terminology

- task** A specific instance of executable codes that is defined by the `task` construct and executed by a node set specified by its `on` clause.

2.8.5 Communication and Synchronization Terminology

- communication** The movement of data between nodes. Communication in XcalableMP occurs only when the programmer instructs it explicitly using a global communication construct or a coarray reference.
- reduction** A procedure involving combining variables from each node in a specified manner and returning the result value. A reduction always involves communication. A reduction is specified by either the `on` clause of the loop construct or the `reduction` construct.

synchronization Synchronization is a mechanism to ensure that multiple nodes do not execute specific portions of a program at the same time. Synchronization between any number of nodes is specified by the `barrier` construct, and that between two nodes by the `post` and `wait` constructs.

asynchronous communication Communication that does not block, and which returns before it is complete. Thus, statements that follow it can overtake it. An asynchronous communication is specified by the `async` clause of global communication constructs or the `async` directive for a coarray reference.

2.8.6 Local-view Terminology

local alias An alias to the local section of a global data object, that is, a distributed array. A local alias can be used in XcalableMP programs in the same way as normal local data.

image An instance of an XcalableMP program corresponding to a respective node.

image set A totally ordered set of images.

image index An integer value that identifies an image in an image set, whose range is from one to the size of the image set.

entire image set The image set corresponding to the entire node set one to one in turn.
executing image

set An image set corresponding to the executing node set one to one in turn.

The executing image set at the beginning of the program execution comprises the entire image set.

allocation image

set An image set on which the coarray data object is allocated.

The allocation image set for a non-allocatable [F] or a static [C] coarray variable comprises the entire image set. Otherwise, the allocation image set for an allocatable [F] or an auto [C] coarray variable is the executing image set on which it is allocated unless it is specified by the `COARRAY` directive.

Chapter 3

Base Language Extensions in XcalableMP C

This chapter describes base language extensions in XcalableMP C that are not described in any other chapters.

3.1 Array Section Notation

Synopsis

The array section notation is a notation to describe a part of an array, which is adapted in Fortran.

Syntax

[C] *array-section* is *array-name*[{ *triplet* | *int-expr* }]...

where *triplet* is:

$$[base] : [length] [: step]$$

Description

In XcalableMP C, the base language C is extended so that a part of an array, i.e., an array section, can be put in an *array assignment statement*, which is described in 3.2, and some XcalableMP constructs. An array section is built from a subset of the elements of an array, which is specified by this notation including at least one *triplet*.

When *step* is positive, the *triplet* specifies a set of subscripts that is a regularly spaced integer sequence of length *length* beginning with *base* and proceeding in increments of *step* up to the largest. When *step* is negative, the *triplet* specifies a set of subscripts that is a regularly spaced integer sequence of length *length* beginning with *base* and proceeding in increments of *step* down to the smallest.

When *base* is omitted, it is assumed to be 0. When *length* is omitted, it is assumed to account for the remainder of the array dimension. When *step* is omitted, it is assumed to be 1.

An array section can be considered as a virtual array containing the set of elements from the original array, which is determined by all possible subscript lists that are specified by the sequence of *triplets* or *int-expr*'s in square brackets.

Restrictions

- [C] Each of *base*, *length*, and *step* must be an integer expression.
- [C] *length* must be greater than zero.
- [C] *step* must not be zero.

Example

Assuming that an array *A* is declared by the following statement,

```
int A[100];
```

some array sections can be specified as follows:

A[10:10]	array section of 10 elements from A[10] to A[19]
A[10:]	array section of 90 elements from A[10] to A[99]
A[:10]	array section of 10 elements from A[0] to A[9]
A[10:5:2]	array section of 5 elements from A[10] to A[18] by step 2
A[:]	the whole of A

3.2 Array Assignment Statement**Synopsis**

An array assignment statement copies a value into each element of an array section.

Syntax

```
[C] array-section [: [int-expr]...] = expression;
```

Description

The value of each element of the result of the right-hand side expression is assigned to the corresponding element of the array section on the left-hand side. When an operator or an elemental function (see section 7.8) is applied to array sections in the right-hand side expression, it is evaluated to an array section that has the same shape as that of the operands or arguments, and each element of which is the result of the operator or function applied to the corresponding element of the operands or arguments. A scalar object is assumed to be an array section that has the same shape as that of the array section(s), and where each element has its value.

Note that an array assignment is a statement, and therefore cannot appear as an expression in any other statements.

Restrictions

- [C] any array section appearing in the right-hand side expression and the left-hand side must have the same shape, i.e., the same number of dimensions and size of each dimension.
- [C] If *array-section* on the left-hand side is followed by “: [*int-expr*]...”, it must be a coarray.

1 Examples

2 An array assignment statement in the fourth line copies the elements B[0] through B[4] into
3 the elements A[5] through A[9].

4

```
----- XcalableMP C -----  
int A[10];  
int B[5];  
    ...  
A[5:5] = B[0:5];
```

5 3.3 Built-in Functions for Array Section

6 Some built-in functions are defined that can accept one or more array sections as arguments. In
7 addition, some of them are array-valued. Such array-valued functions can appear in the right-
8 hand side of an array assignment statement, and should be preceded by the `array` directive if
9 the array section is distributed.

10 All of the built-in functions for array sections are described in Sections 7.8 and 7.9.

11 3.4 Pointer to Global Data

12 3.4.1 Name of Global Array

13 The name of a global array is considered to represent an abstract entity in the XcalableMP
14 language. It is not interpreted as the pointer to the array, while the name of a local array is.

15 However, the name of a global array that appears in an expression is evaluated to the base
16 address of its local section on each node. The pointer can be operated on each node as if it were
17 a normal (local) pointer.

18 3.4.2 Address-of Operator

19 The result of the address-of operator (“&”) applied to an element of a global array is the pointer
20 to the corresponding element of its local section. Note that the value of the result pointer is
21 defined only on the node that owns the element. The pointer can be operated on the node as if
22 it were a normal (local) pointer.

23 As a result, for a global array `a`, `a` and `&a[0]` are not always evaluated to the same value.

24 3.5 Dynamic Allocation of Global Data

25 In XcalableMP C, it is possible to allocate global arrays at runtime. Such an allocation is done
26 by performing the following steps.

- 27 1. Declare a pointer to an object of the type of the global array to be allocated.
- 28 2. Align the pointer with a template as if it were an array.
- 29 3. Allocate a block of memory of the global size using the `xmp_malloc` library procedure, and
30 assign the return value to the pointer on each node.

```

XcalableMP C
#pragma nodes p(NP1, NP2)
#pragma xmp template t(:, :)
#pragma xmp distribute t(block, block) onto p
5 float (*pa)[N2];
#pragma xmp align pa[i][j] with t(i, j)
#pragma xmp template_fix t(0:N1-1, 0:N2-1)

pa = (float (*)[N2])xmp_malloc(xmp_desc_of(pa), N1, N2);
```

3.6 Descriptor-of Operator

The descriptor-of operator (“`xmp_desc_of`”) is introduced as a built-in operator in XcalableMP C.

The result of the descriptor-of operator applied to XcalableMP entities such as node arrays, templates, and global arrays is their *descriptor*, which can be used in various ways, including as an argument of some inquiry procedures. The type of the result, `xmp_desc_t`, is implementation-defined, and is defined in the `xmp.h` header file in XcalableMP C.

For the `xmp_desc_of` intrinsic function in XcalableMP Fortran, refer to section 7.1.1.

Chapter 4

Directives

This chapter describes the syntax and behavior of XcalableMP directives. In this document, the following notation is used to describe XcalableMP directives.

- xxx** **type-face** characters are used to indicate literal-type characters.
- xxx...* If the line is followed by "...", then xxx can be repeated.
- [xxx]* *xxx* is optional.
- The syntax rule continues.
- [F] The following lines are effective only in XcalableMP Fortran.
- [C] The following lines are effective only in XcalableMP C.

4.1 Directive Format

4.1.1 General Rule

In XcalableMP Fortran, XcalableMP directives are specified using special comments that are identified by unique sentinels `!$xmp`. An XcalableMP directive follows the rules for comment lines of either the Fortran free or fixed source form, depending on the source form of the surrounding program unit¹. XcalableMP Fortran directives are case insensitive.

[F] `!$xmp directive-name clause`

In XcalableMP C, XcalableMP directives are specified using the `#pragma` mechanism provided by the C standards. XcalableMP C directives are case-sensitive.

[C] `#pragma xmp directive-name clause`

Directives are classified as *declarative directives* and *executable directives*.

The declarative directive is a directive that may only be placed in a declarative context. A declarative directive has no associated executable user code. The scope rule of declarative directives obeys that of the declaration statements in the base language. For example, in XcalableMP Fortran, a node array declared by a `nodes` directive is visible only within either the program unit, the derived-type declaration, or the interface body that immediately surrounds the directives, unless it is overridden in the inner blocks or is use or host associated. Further, in XcalableMP C, a node array declared by a `nodes` directive is visible only in the range from the

¹Consequently, the rules of comment lines that an XcalableMP directive follows are the same as the ones followed by an OpenMP directive.

declaring point to the end of the block when placed within a block, or of the file when placed outside any blocks, unless overridden in the inner blocks.

Note that in XcalableMP Fortran, node arrays and templates in other scoping units are accessible by use or host association.

The following directives are declarative directives.

- `nodes`
- `template`
- `distribute`
- `align`
- `shadow`
- `coarray`

The executable directives are placed in an executable context. A stand-alone directive is an executable directive that has no associated user code, such as a `barrier` directive. An executable directive and its associated user code make up an XcalableMP construct, as in the following format:

```
[F]  !$xmp directive-name clause ...
      structured-block

[C]  #pragma xmp directive-name clause ...
      structured-block
```

Note that in XcalableMP Fortran, a corresponding `end` directive is required for some executable directives such as `task` and `tasks`, and in XcalableMP C, the associated statement can be a compound one.

The following directives are executable directives.

- `template_fix`
- `task`
- `tasks`
- `loop`
- `array`
- `reflect`
- `gmove`
- `barrier`
- `reduction`
- `bcast`
- `wait_async`

1 4.1.2 Combined Directive

2 Synopsis

3 Multiple attributes can be specified by one combined declarative directive, which is analogous
4 to type declaration statements using the “:.” punctuation.

5 Syntax

```
[F] !$xmp combined-directive      is combined-attribute [, combined-attribute ]... ::
                                     combined-decl [, combined-decl ]...
[C] #pragma xmp combined-directive is combined-attribute [, combined-attribute ]... ::
                                     combined-decl [, combined-decl ]...
```

7 *combined-attribute* is one of:

```
nodes
template
distribute ( dist-format [, dist-format]... ) onto nodes-name
8 align ( align-source [, align-source]... ) ■
                                     ■ with template-name ( align-subscript [, align-subscript]... )
shadow ( shadow-width [, shadow-width]... )
[F] dimension ( explicit-shape-spec [, explicit-shape-spec]... )
```

9 and *combined-decl* is one of:

```
nodes-decl
10 template-decl
array-name
```

11 Description

12 A combined directive is interpreted as if an object corresponding to each *combined-decl* is de-
13 clared in a directive corresponding to each *combined-attribute*, where all restrictions of each
14 directive, in addition to the following ones, are applied.

15 Restrictions

- 16 • The same kind of *combined-attribute* must not appear more than once in a given *combined-*
17 *directive*.
- 18 • If the `nodes` attribute appears in a *combined-directive*, each *combined-decl* must be a
19 *nodes-decl*.
- 20 • If the `template` or `distribute` attribute appears in a *combined-directive*, each *combined-*
21 *decl* must be a *template-decl*.
- 22 • If the `align` or `shadow` attribute appears in a *combined-directive*, each *combined-decl* must
23 be an *array-name*.
- 24 • [F] If the `dimension` attribute appears in a *combined-directive*, any object to which it
25 applies must be declared using either the `template` or the `nodes` attribute.

4.2 nodes Directive

Synopsis

The `nodes` directive declares a named node array.

Syntax

[F] `!$xmp nodes nodes-decl [, nodes-decl]...`

[C] `#pragma xmp nodes nodes-decl [, nodes-decl]...`

where *nodes-decl* is one of:

nodes-name (*nodes-spec* [, *nodes-spec*]...)

nodes-name (*nodes-spec* [, *nodes-spec*]...) = *nodes-ref*

[C] *nodes-name* [*nodes-spec*] [[*nodes-spec*]...]

[C] *nodes-name* [*nodes-spec*] [[*nodes-spec*]...] = *nodes-ref*

and *nodes-spec* must be one of:

int-expr

*

Description

The `nodes` directive declares a node array that corresponds to a node set.

The first and third forms of the `nodes` directive are used to declare a node array that corresponds to the entire node set. The second and fourth forms are used to declare a node array, each element of which is assigned to the node of the node set specified by *nodes-ref* at the corresponding position of its elements order. In the first and second forms, which use parentheses, the element order of the declared node array is based on Fortran 's. In the third and fourth forms, which use square brackets, the element order of the declared node array is based on C 's.

If *node-size* in the last dimension is "*" in the first and second forms, or if that in the first dimension is "*" in the third and fourth forms, then the size of the node array is automatically adjusted according to the total size of either the entire node set in the first and third forms or the referenced node set in the second and fourth forms.

Restrictions

- *nodes-name* must not conflict with any other local name in the same scoping unit.
- *nodes-spec* can be "*" only in the last dimension in the first and second forms, and *nodes-spec* can be "*" only in the first dimension in the third and fourth forms.
- *nodes-ref* must not reference *nodes-name* either directly or indirectly.
- If no *nodes-spec* is "*", then the product of all *nodes-spec* must be equal to the total size of the entire node set in the first and third forms, or the referenced node set in the second and fourth forms.
- *nodes-subscript* in *nodes-ref* must not be "*".

1 Examples

2 The following are examples of the first and the third forms that appears in the main program.
 3 Because the node array `p`, which corresponds to the entire node set, is declared to be of size 16,
 4 this program must be executed by 16 nodes.

XcalableMP Fortran	XcalableMP C
<pre> program main !\$xmp nodes p(16) !\$xmp nodes q(4,*) 5 !\$xmp nodes r(8)=p(3:10) 5 !\$xmp nodes z(2,3)=p(1:6) ... end program </pre>	<pre> int main() { #pragma xmp nodes p[16] #pragma xmp nodes q[*][4] #pragma xmp nodes r[8]=p[2:8] #pragma xmp nodes z[3][2]=p[0:6] ... } </pre>

6 The following are examples of a node declaration in a procedure. Because `p` is declared in
 7 the second and fourth forms to have a size of 16 and corresponds to the executing node set, the
 8 invocation of the `foo` function must be executed by 16 nodes. The node array `q` is declared in
 9 the first and third forms, and corresponds to the entire node set. The node array `r` is declared
 10 as a subset of `p`, and `x` as a subset of `q`.

XcalableMP Fortran	XcalableMP C
<pre> function foo() !\$xmp nodes p(16)=* !\$xmp nodes q(4,*) 11 !\$xmp nodes r(8)=p(3:10) 5 !\$xmp nodes x(2,3)=q(1:2,1:3) ... end function </pre>	<pre> void foo(){ #pragma xmp nodes p[16]=* #pragma xmp nodes q[*][4] #pragma xmp nodes r[8]=p[2:8] #pragma xmp nodes x[3][2]=q[0:3][0:2] ... } </pre>

12 4.2.1 Node Reference

13 Synopsis

14 The node reference is used to reference a node set.

15 Syntax

16 A node reference *nodes-ref* is specified by either the name of a node array or the “*” symbol.

17
$$\begin{array}{l} \textit{nodes-ref} \textbf{ is } \textit{nodes-name} \textit{ [(} \textit{nodes-subscript} \textit{ [, } \textit{nodes-subscript} \textit{] ...)]} \\ \textit{[C] nodes-ref} \textbf{ is } \textit{nodes-name} \textit{ [[} \textit{nodes-subscript} \textit{] [[} \textit{nodes-subscript} \textit{] ...]]} \\ \textbf{or} \textbf{ *} \end{array}$$

18 where *nodes-subscript* must be one of:

19 *int-expr*
 20 *triplet*
 *

21 Description

22 A node reference by *nodes-name* represents a node set corresponding to the node array specified
 23 by the name or its subarray. It is totally ordered in Fortran’s array element order in the first

form, and in C's array element order in the second form. A node reference by “*” represents the executing node set.

Specifically, the “*” symbol that appears as *nodes-subscript* in a dimension of *nodes-ref* is interpreted by each node at runtime as its position (coordinate) in the dimension of the referenced node array. Thus, a node reference $p(s_1, \dots, s_{k-1}, *, s_{k+1}, \dots, s_n)$ is interpreted as $p(s_1, \dots, s_{k-1}, j_k, s_{k+1}, \dots, s_n)$ on the node $p(j_1, \dots, j_{k-1}, j_k, j_{k+1}, \dots, j_n)$.

Note that “*” can be used only as the node reference in the on clause of some executable directives.

Examples

Assume that *p* is the name of a node array and that *m* is an integer variable.

- As a target node array in the `distribute` directive,

XcalableMP Fortran	XcalableMP C
<code>!\$xmp distribute a(block) onto p</code>	<code>#pragma xmp distribute a(block) onto p</code>

- To specify the node array to which the declared node array corresponds in the second and fourth forms of the `nodes` directive,

XcalableMP Fortran	XcalableMP C
<code>!\$xmp nodes r(2,2,4) = p(1:4,1:4)</code>	<code>#pragma xmp nodes r[4][2][2] = p[0:4][0:4]</code>
<code>!\$xmp nodes r(2,2,4) = p(1:16)</code>	<code>#pragma xmp nodes r[4][2][2] = p[0:16]</code>

- To specify the node array that corresponds to the executing node set of a task in the `task` directive,

XcalableMP Fortran	XcalableMP C
<code>!\$xmp task on p(1:4,1:4)</code>	<code>#pragma xmp task on p[0:4][0:4]</code>
<code>!\$xmp task on p(1:16)</code>	<code>#pragma xmp task on p[0:16]</code>
<code>!\$xmp task on p(:,*)</code>	<code>#pragma xmp task on p[*][:]</code>
<code>!\$xmp task on p(m)</code>	<code>#pragma xmp task on p[m]</code>

- To specify the node array that corresponds to the executing node set in the `barrier` and the `reduction` directive,

XcalableMP Fortran	XcalableMP C
<code>!\$xmp barrier on p(5:8)</code>	<code>#pragma xmp barrier on p[4:4]</code>
<code>!\$xmp reduction (+:a) on p(*,:)</code>	<code>#pragma xmp reduction (+:a) on p[:] [*]</code>

- To specify the source node and the node array that corresponds to the executing node set in the `bcast` directive,

XcalableMP Fortran	XcalableMP C
<code>!\$xmp bcast (b) from p(k) on p(:)</code>	<code>#pragma xmp (b) from p[k-1] on p[:]</code>

4.3 Template and Data Mapping Directives

4.3.1 template Directive

Synopsis

The `template` directive declares a template.

1 Syntax

[F] `!$xmp template template-decl [, template-decl]...`

2

[C] `#pragma xmp template template-decl [, template-decl]...`

3 where *template-decl* is:

template-name (*template-spec* [, *template-spec*]...)

4

[C] *template-name* [*template-spec-c*] [[*template-spec-c*]...]

5 and *template-spec* must be one of:

[*int-expr* :] *int-expr*

6

:

7 and *template-spec-c* must be one of:

int-expr

8

:

9 Description

10 The `template` directive declares a template with the shape specified by the sequence of *template-spec*'s or *template-spec-c*'s. If every *template-spec* or *template-spec-c* is “:”, then the shape of
 11 the template is initially undefined. This template must not be referenced until the shape is
 12 defined by a `template_fix` directive (see section 4.3.6) at runtime. If only *int-expr* is specified
 13 as *template-spec*, the default lower bound is one.
 14

15 Restrictions

- 16 • *template-name* must not conflict with any other local name in the same scoping unit.
- 17 • Every *template-spec* must be either [*int-expr* :] *int-expr* or “:”.
- 18 • Every *template-spec-c* must be either *int-expr* or “:”.

19 4.3.2 Template Reference**20 Synopsis**

21 The template reference expression specified in the `on` or the `from` clause of some directives is
 22 used to indirectly specify a node set.

23 Syntax

template-ref **is** *template-name* [(*template-subscript* [, *template-subscript*]...)]

24

[C] *template-ref* **is** *template-name* [[*template-subscript*] [[*template-subscript*]...]]

25 where *template-subscript* must be one of:

26

int-expr

27

triplet

*

Description

Being specified in the `on` or the `from` clause of some directives, the template reference refers to a subset of a node set in which the specified subset of the template resides.

Specifically, the “*” symbol that appears as *template-subscript* in a dimension of *template-ref* is interpreted by each node at runtime as the indices of the elements in the dimension that reside in the node. “*” in a template reference is similar to “*” in a node reference.

Examples

Assume that `t` is a template.

- In the `task` directive, the executing node set of the task can be indirectly specified using a template reference in the `on` clause.

XcalableMP Fortran	XcalableMP C
<code>!\$xmp task on t(1:m,1:n)</code>	<code>#pragma xmp task on t[0:n][0:m]</code>
<code>!\$xmp task on t</code>	<code>#pragma xmp task on t</code>

- In the `loop` directive, the executing node set of each iteration of the following loop is indirectly specified using a template reference in the `on` clause.

XcalableMP Fortran	XcalableMP C
<code>!\$xmp loop (i) on t(i-1)</code>	<code>#pragma xmp loop (i) on t[i-1]</code>

- In the `array` directive, the executing node set on which the associated array-assignment statement is performed in parallel is indirectly specified using a template reference in the `on` clause.

XcalableMP Fortran	XcalableMP C
<code>!\$xmp array on t(1:n)</code>	<code>#pragma xmp array on t[0:n]</code>

- In the `barrier`, `reduction`, and `bcast` directives, the node set that is to perform the operation collectively can be indirectly specified using a template reference in the `on` clause.

XcalableMP Fortran	XcalableMP C
<code>!\$xmp barrier on t(1:n)</code>	<code>#pragma xmp barrier on t[0:n]</code>
<code>!\$xmp reduction (+:a) on t(*,:)</code>	<code>#pragma xmp reduction (+:a) on t[:] [*]</code>
<code>!\$xmp bcast (b) on t(1:n)</code>	<code>#pragma xmp bcast (b) on t[0:n]</code>

4.3.3 distribute Directive

Synopsis

The `distribute` directive specifies the distribution of a template.

Syntax

[F] `!$xmp distribute template-name (dist-format [, dist-format]...) onto nodes-name`

[C] `#pragma xmp distribute template-name (dist-format [, dist-format]...)` ■
 ■ onto *nodes-name*

[C] `#pragma xmp distribute template-name [dist-format] [[dist-format] ...]` ■
 ■ onto *nodes-name*

1 where *dist-format* must be one of:

```

2     *
3     block [ ( int-expr ) ]
4     cyclic [ ( int-expr ) ]
5     gblock ( { * | int-array } )

```

3 Description

4 According to the specified distribution format, a template is distributed onto a specified node
5 array. The dimension of the node array that appears in the `onto` clause corresponds, in order
6 of left-to-right, to the dimension of the distributed template for which the corresponding *dist-*
7 *format* is not “*”.

8 Let *d* be the size of the dimension of the template, *p* be the size of the corresponding
9 dimension of the node array, `ceiling` and `mod` be Fortran’s intrinsic functions, and each of the
10 arithmetic operators be that of Fortran. The interpretation of *dist-format* is as follows:

11 “*” The dimension is not distributed.

12 `block` Equivalent to `block(ceiling(d/p))`.

13 `block(n)` The dimension of the template is divided into contiguous blocks of size *n*, which are
14 distributed onto the corresponding dimension of the node array. The dimension of the
15 template is divided into d/n blocks of size *n*, and one block of size `mod(d,n)` if any, and
16 each block is assigned sequentially to an index along the corresponding dimension of the
17 node array. Note that if $k = p - d/n - 1 > 0$, then there is no block assigned to the last *k*
18 indices.

19 `cyclic` Equivalent to `cyclic(1)`.

20 `cyclic(n)` The dimension of the template is divided into contiguous blocks of size *n*, and these
21 blocks are distributed onto the corresponding dimension of the node array in a round-robin
22 manner.

23 `gblock(m)` *m* is referred to as a mapping array. The dimension of the template is divided into
24 contiguous blocks so that the *i*’th block is of size `m(i)`, and these blocks are distributed
25 onto the corresponding dimension of the node array.

26 If at least one `gblock(*)` is specified in *dist-format*, then the template is initially undefined
27 and must not be referenced until the shape of the template is defined by `template_fix` directives
28 at runtime.

29 Restrictions

30 • [C] *template-name* must be declared by a `template` directive that lexically precedes the
31 directive.

32 • The number of *dist-format* that is not “*” must be equal to the rank of the node array
33 specified by *nodes-name*.

34 • The size of the dimension of the template specified by *template-name* that is distributed
35 by `block(n)` must be equal to or less than the product of the block size *n* and the size of
36 the corresponding dimension of the node array specified by *nodes-name*.

- The array *int-array* in parentheses following `gblock` must be an integer one-dimensional array, and its size must be equal to the size of the corresponding dimension of the node array specified by *nodes-name*.
- Every element of the array *int-array* in parentheses following `gblock` must have a value of a nonnegative integer.
- The sum of the elements of the array *int-array* in parentheses following `gblock` must be equal to the size of the corresponding dimension of the template specified by *template-name*.
- [C] A `distribute` directive for a template must precede any of its references in the executable code in the block.

Examples

Example 1

XcalableMP Fortran	XcalableMP C
<code>!\$xmp nodes p(4)</code>	<code>#pragma xmp nodes p[4]</code>
<code>!\$xmp template t(64)</code>	<code>#pragma xmp template t[64]</code>
<code>!\$xmp distribute t(block) onto p</code>	<code>#pragma xmp distribute t[block] onto p</code>

The template `t` is distributed in `block` format, as shown in the following table.

p(1)	t(1:16)	p[0]	t[0:16]
p(2)	t(17:32)	p[1]	t[16:16]
p(3)	t(33:48)	p[2]	t[32:16]
p(4)	t(49:64)	p[3]	t[48:16]

Example 2

XcalableMP Fortran	XcalableMP C
<code>!\$xmp nodes p(4)</code>	<code>#pragma xmp nodes p[4]</code>
<code>!\$xmp template t(64)</code>	<code>#pragma xmp template t[64]</code>
<code>!\$xmp distribute t(cyclic(8)) onto p</code>	<code>#pragma xmp distribute t[cyclic(8)] onto p</code>

The template `t` is distributed in `cyclic` format of size eight, as shown in the following table.

p(1)	t(1:8) t(33:40)	p[0]	t[0:8] t[32:8]
p(2)	t(9,16) t(41:48)	p[1]	t[8:8] t[40:8]
p(3)	t(17,24) t(49:56)	p[2]	t[16:8] t[48:8]
p(4)	t(25,32) t(57:64)	p[3]	t[24:8] t[56:8]

Example 3

XcalableMP Fortran	XcalableMP C
<code>!\$xmp nodes p(8,5)</code>	<code>#pragma xmp nodes p[5][8]</code>
<code>!\$xmp template t(64,64,64)</code>	<code>#pragma xmp template t[64][64][64]</code>
<code>!\$xmp distribute t(*,cyclic,block) onto p</code>	<code>#pragma xmp distribute t[block][cyclic][*] onto p</code>

The first dimension of the template `t` is not distributed. The second dimension is distributed onto the first dimension of the node array `p` in `cyclic` format. The third dimension is distributed onto the second dimension of `p` in `block` format. The results are as follows:

p(1,1)	t(1:64, 1:57:8, 1:13)	p[0][0]	t[0:13][0:8:8][0:64]
p(2,1)	t(1:64, 2:58:8, 1:13)	p[0][1]	t[0:13][1:8:8][0:64]
...
p(8,5)	t(1:64, 8:64:8, 53:64)	p[4][7]	t[52:12][7:8:8][0:64]

Note that the “64” in `template t` is not divisible by “5” in `node p`. Thus, the sizes of the blocks are different among nodes.

4.3.4 align Directive

Synopsis

The `align` directive specifies that an array is to be mapped in the same way as a specified template.

Syntax

```
[F] !$xmp align array-name ( align-source [, align-source]... )
      with template-name (align-subscript [, align-subscript]... )
```

```
[C] #pragma xmp align array-name [align-source] [[align-source]]...
      with template-name (align-subscript [, align-subscript]... )
      OR
      with template-name [align-subscript] [ [ align-subscript ]... ]
```

where *align-source* must be one of:

scalar-int-variable

*
:

and *align-subscript* must be one of:

scalar-int-variable [{ + | - } *int-expr*]

*
:

Note that the variable *scalar-int-variable* that appears in *align-source* is referred to as an “align dummy variable” and *int-expr* appearing in *align-subscript* as an “align offset.”

Description

The array specified by *array-name* is aligned with the template that is specified by *template-name* so that each element of the array indexed by the sequence of *align-sources* is aligned with the element of the template indexed by the sequence of *align-subscripts*, where *align-sources* and *align-subscripts* are interpreted as follows:

1. The first form of *align-source* and *align-subscript* represents an align dummy variable and an expression of it, respectively. The align dummy variable is considered to range over all valid index values in the corresponding dimension of the array.
2. The second form “*” of *align-source* and *align-subscript* represents a dummy variable (not an align dummy variable) that does not appear anywhere in the directive.

- The second form of *align-source* is said to “collapse” the corresponding dimension of the array. As a result, the index along the corresponding dimension does not affect the determination of the alignment. 1
2
3
 - The second form of *align-subscript* is said to “replicate” the array. Each element of the array is replicated, and is aligned to all index values in the corresponding dimension of the template. 4
5
6
3. The third form of *align-source* and the matching *align-subscript* represents the same align dummy variable whose range spans all valid index values in the corresponding dimension of the array. The matching of colons (“:”) in the sequence of *align-sources* and *align-subscripts* is determined as follows: 7
8
9
10
- [F] Colons in the sequence of *align-sources* and those in the sequence of *align-subscripts* are matched in corresponding left-to-right order, where any *align-source* and *align-subscript* that is not a colon is ignored. 11
12
13
 - [C] Colons in the sequence of *align-sources* in right-to-left order, and those in the sequence of (*align-subscript*)’s in left-to-right order are matched, or those in the sequence of [*align-subscript*]’s in right-to-left order are matched, where any *align-source* and *align-subscript* that is not a colon is ignored. 14
15
16
17

In XcalableMP C, an `align` directive for a dummy argument can be placed either outside the function body (as in the old style of C) or in it (as in the ANSI style). 18
19

Restrictions 20

- [C] *array-name* must be declared by a declaration statement that lexically precedes the directive. 21
22
- An align dummy variable may appear at most once in the sequence of *align-sources*. 23
- An align dummy variable may appear at most once in the sequence of *align-subscripts*. 24
- An *align-subscript* may contain at most one occurrence of an align dummy variable. 25
- The *int-expr* in an *align-subscript* may not contain any occurrence of an align dummy variable. 26
27
- The sequence of *align-sources* must contain exactly as many colons as contained by the sequence of *align-subscripts*. 28
29
- [F] The array specified by *array-name* must not appear as an *equivalence-object* in an `equivalence` statement. 30
31
- [C] An `align` directive for an array must precede any of its appearances in the executable code in the block. 32
33
- [F] The array specified by *array-name* shall not be initially defined. 34
- [C] The array specified by *array-name* shall not be initialized through an *initializer*. 35

1 **Examples**2 **Example 1**

	XcalableMP Fortran	XcalableMP C
3	<code>!\$xmp align a(i) with t(i)</code>	<code>#pragma xmp align a[i] with t[i]</code>

4 In XcalableMP Fortran, the array element `a(i)` is aligned with the template element `t(i)`. In XcalableMP C, the array element `a[i]` is aligned with the template element `t[i]`. These are equivalent to the following codes.

	XcalableMP Fortran	XcalableMP C
7	<code>!\$xmp align a(:) with t(:)</code>	<code>#pragma xmp align a[:] with t[:]</code>

8 **Example 2**

	XcalableMP Fortran	XcalableMP C
9	<code>!\$xmp align a(*,j) with t(j)</code>	<code>#pragma xmp align a[j][*] with t[j]</code>

10 In XcalableMP Fortran, the subarray `a(:,j)` is aligned with the template element `t(j)`. Note that the first dimension of `a` is collapsed. In XcalableMP C, the subarray `a[j][:]` is aligned with the template element `t[j]`. Note that the second dimension of `a` is collapsed.

13 **Example 3**

	XcalableMP Fortran	XcalableMP C
14	<code>!\$xmp align a(j) with t(*,j)</code>	<code>#pragma xmp align a[j] with t[j][*]</code>

15 In XcalableMP Fortran, the array element `a(j)` is replicated and aligned with each template element of `t(:,j)`. In XcalableMP C, the array element `a[j]` is replicated and aligned with each template element of `t[j][:]`.

18 **Example 4**

	XcalableMP Fortran	XcalableMP C
19	<code>!\$xmp template t(n1,n2)</code> <code>real a(m1,m2)</code> <code>!\$xmp align a(*,j) with t(*,j)</code>	<code>#pragma xmp template t[n2][n1]</code> <code>double a[m2][m1]</code> <code>#pragma xmp align a[j][*] with t[j][*]</code>

20 In XcalableMP Fortran, the subarray `a(:,j)` is aligned with each template element of `t(:,j)`. In XcalableMP C, the subarray `a[j][:]` is aligned with each template element of `t[j][:]`.

23 By replacing “*” of the array `a` and “*” of the template `t` with a dummy variable `i` and `k`, respectively, this alignment can be interpreted as the following mapping.

25 [F] $a(i,j) \rightarrow t(k,j) \mid (i,j,k) \in (1:n1, 1:n2, 1:m1)$
26 [C] $a[j][i] \rightarrow t[j][k] \mid (i,j,k) \in (0:n1, 0:n2, 0:m1)$

27 **4.3.5 shadow Directive**28 **Synopsis**

29 The `shadow` directive allocates the shadow area for a distributed array.

30 **Syntax**

[F] `!$xmp shadow array-name (shadow-width [, shadow-width]...)`

31

[C] `#pragma xmp shadow array-name [shadow-width] [shadow-width]...`

where *shadow-width* must be one of:

int-expr

int-expr : *int-expr*

*

Description

The **shadow** directive specifies the width of the shadow area of an array specified by *array-name*, which is used to communicate the neighbor element of the block of the array. When *shadow-width* is of the form “*int-expr* : *int-expr*,” the shadow area of the width specified by the first *int-expr* is added at the lower bound, and that specified by the second one is added at the upper bound in the dimension. When *shadow-width* is of the form *int-expr*, the shadow area of the same width specified is added at both the upper and lower bounds in the dimension. When *shadow-width* is of the form “*”, the entire area of the array is allocated on each node, and the area that it does not own is regarded as a shadow. This type of shadow is sometimes referred to as a “full shadow.”

Note that the shadow area of a multi-dimensional array includes “obliquely-neighboring” elements, which are owned by the node whose indices are different in more than one dimension, and that the shadow area can also be allocated at the global lower and upper bounds of an array.

The data stored in the storage area declared by the **shadow** directive is referred to as a *shadow object*. A shadow object represents an element of a distributed array, and corresponds to the data object that represents the same element as itself. The corresponding data object is referred to as the *reflection source* of the shadow object.

Restrictions

- [C] *array-name* must be declared by a declaration statement that lexically precedes the directive.
- The value specified by *shadow-width* must be a nonnegative integer.
- The number of *shadow-width* must be equal to the number of dimensions (or rank) of the array specified by *array-name*.
- [C] A **shadow** directive for an array must precede any of its appearances in the executable code in the block.

1 **Example**

```

XcalableMP Fortran
!$xmp nodes p(4,4)
!$xmp template t(64,64)
!$xmp distribute t(block,block) onto p
2
5   real a(64,64)
!$xmp align a(i,j) with t(i,j)
!$xmp shadow a(1,1)

```




Figure 4.1: Example showing shadow of a two-dimensional array.

3 The node $p(2,2)$ has $a(17:32,17:32)$ as a data object, and $a(16,16)$, $a(17:32,16)$,
4 $a(33,16)$, $a(16,17:32)$, $a(33,17:32)$, $a(16,33)$, $a(17:32,33)$, and $a(33,33)$ as shadow ob-
5 jects (Figure 4.1). Among them, $a(16,16)$, $a(33,16)$, $a(16,33)$, and $a(33,33)$ are “obliquely-
6 neighboring” elements of $p(2,2)$.

7 **4.3.6 template_fix Construct**8 **Synopsis**

9 This construct fixes the shape and/or the distribution of an undefined template.

10 **Syntax**

[F] `!$xmp template_fix` ■
 ■ [*dist-format* [, *dist-format*]...] *template-name* [(*template-spec* [, *template-spec*]...)]

[C] `#pragma xmp template_fix` ■
 ■ [(*dist-format* [, *dist-format*]...)] *template-name* [(*template-spec* [, *template-spec*]...)]
 or
 ■ [[*dist-format* [, *dist-format*]...] *template-name* [*template-spec*] [[*template-spec*]...]

13 where *template-spec* is:

14 [*int-expr* :] *int-expr*

15 and *dist-format* is one of:

*
 16 `block` [(*int-expr*)]
`cyclic` [(*int-expr*)]
`gblock` (*int-array*)

17 **Description**

18 The `template_fix` construct fixes the shape and/or the distribution of the template that is
 19 initially undefined, by specifying the sizes and/or the distribution format of each dimension
 20 at runtime. Arrays that are aligned with an initially undefined template must be allocatable
 21 arrays, in XcalableMP Fortran, or a pointer (see Section 3.5), in XcalableMP C, which cannot

be allocated until the template is fixed by the `template_fix` construct. No constructs that have such a template in their `on` clause should be encountered until the template is fixed by the `template_fix` construct. Any undefined template can be fixed only once by the `template_fix` construct in its scoping unit.

The meaning of the sequence of *dist-formats* is the same as that in the `distribute` directive.

Restrictions

- When a node encounters a `template_fix` construct at runtime, the template specified by *template-name* must be undefined.
- If the sequence of *dist-formats* exists in a `template_fix` construct, it must be identical to the sequence of *dist-formats* in the `distribute` directive for the template specified by *template-name*, except for *int-array* specified in the parenthesis immediately after `gblock`.
- Either the sequence of *dist-formats* or the sequence of *template-spec*'s should be given.

Example

XcalableMP Fortran	XcalableMP C
!\$xmp nodes p(*)	#pragma xmp nodes p[*]
!\$xmp template t(:)	#pragma xmp template t[:]
!\$xmp distribute t(gblock(*)) onto p	#pragma xmp distribute t[gblock(*)] onto p
real, allocatable :: a(:)	double *a;
!\$xmp align a(i) with t(i)	#pragma xmp align a[i] with t[i]
...	...
N = ...	N = ...;
M(...) = ...	M[] = {...};
...	...
!\$xmp template_fix(gblock(M)) t(N)	#pragma xmp template_fix[gblock(M)] t[N]
...	...
allocate (a(N))	a = xmp_malloc(xmp_desc_of(a), N);

Because the shape is `t(:)` or `t[:]` and the distribution format is `gblock(*)`, the template `t` is initially undefined. The allocatable array `a` is aligned with `t`. After the size `N` and the mapping array `M` is defined, `t` is fixed by the `template_fix` construct and `a` is allocated.

In XcalableMP C, it is possible to allocate global arrays at runtime only when they are one-dimensional. Such an allocation is done by performing the following steps.

1. Declare a pointer to an object of the type of the global array to be allocated.
2. Align the pointer with a template as if it were a one-dimensional array.
3. Allocate a storage of the global size with the function `xmp_malloc()` and assign the result value to the pointer on each node.

The functions `xmp_desc_of()` and `xmp_malloc()` are described in section 3.6 and 7.5.1, respectively.

1 4.4 Work Mapping Construct

2 4.4.1 task Construct

3 Synopsis

4 The **task** construct defines a task that is executed by a specified node set.

5 Syntax

```
[F] !$xmp task on {nodes-ref | template-ref}
    structured-block
    !$xmp end task
```

```
[C] #pragma xmp task on {nodes-ref | template-ref}
    structured-block
```

7 Description

8 When a node encounters a **task** construct at runtime, it executes the associated block (called a
9 *task*) if it is included by the node set specified by the **on** clause; otherwise, it skips the execution
10 of the block.

11 Unless a **task** construct is surrounded by a **tasks** construct, *nodes-ref* or *template-ref* in the
12 **on** clause is evaluated by the executing node set at the start of the task; otherwise, *nodes-ref*
13 and *template-ref* of the **task** construct are evaluated by the executing node set at the entry of
14 the **tasks** construct that immediately surrounds it. The current executing node set is set to be
15 that specified by the **on** clause at the entry of the **task** construct, and it is rewound to the last
16 one at the exit.

17 Restrictions

- 18 • The node set specified by *nodes-ref* or *template-ref* in the **on** clause must be a subset of
19 the parent node set.

20 Example

21 **Example 1** In XcalableMP Fortran, copies of variables **a** and **b** are replicated on nodes **nd(1)**
22 through **nd(8)**. A task defined by the **task** construct is executed only on **nd(1)**, and
23 defines the copies of **a** and **b** on a node **nd(1)**. The copies on nodes **nd(2)** through **nd(8)**
24 are not defined.

25 In XcalableMP C, copies of variables **a** and **b** are replicated on nodes **nd[0]** through **nd[7]**.
26 A task defined by the **task** construct is executed only on **nd[0]**, and defines the copies of
27 **a** and **b** on a node **nd[0]**. The copies on nodes **nd[1]** through **nd[7]** are not defined.

XcalableMP Fortran	XcalableMP C
<pre> !\$xmp nodes nd(8) !\$xmp template t(100) !\$xmp distribute t(block) onto nd real a, b; !\$xmp task on nd(1) read(*,*) a b = a*1.e-6 !\$xmp end task </pre>	<pre> #pragma xmp nodes nd[8] #pragma xmp template t[100] #pragma xmp distribute t[block] onto nd float a, b; #pragma xmp task on nd[0] { scanf ("%f", &a); b = a*1.e-6; } </pre>

Example 2 According to the `on` clause with a template reference, an assignment statement in the `task` construct is executed by the owner of the array element `a(:,j)` or `a[j][:]`.

XcalableMP Fortran	XcalableMP C
<pre> !\$xmp nodes nd(8) !\$xmp template t(100) !\$xmp distribute t(block) onto nd integer i,j real a(200,100) !\$xmp align a(*,j) with t(j) i = ... j = ... !\$xmp task on t(j) a(i,j) = 1.0 !\$xmp end task </pre>	<pre> #pragma xmp nodes nd[8] #pragma xmp template t[100] #pragma xmp distribute t(block) onto nd int i,j; float a[100][200]; #pragma align a[j][*] with t[j] i = ...; j = ...; #pragma xmp task on t[j] a[j][i] = 1.0; } </pre>

4.4.2 tasks Construct

Synopsis

The `tasks` construct is used to instruct the executing nodes to execute the multiple tasks that it surrounds in an arbitrary order.

Syntax

- [F] `!$xmp tasks`
task-construct
...
`!$xmp end tasks`
- [C] `#pragma xmp tasks`
{
task-construct
...
}

1 **Description**

2 **task** constructs surrounded by a **tasks** construct are executed in arbitrary order without implicit
 3 synchronization at the start of each task. As a result, if there are no overlaps between the
 4 executing node sets of the adjacent tasks, they can be executed in parallel.

5 *nodes-ref* or *template-ref* of each task immediately surrounded by a **tasks** construct is eval-
 6 uated by the executing node set at the entry of the **tasks** construct.

7 No implicit synchronization is performed at the start and end of the **tasks** construct.

8 **Example**

9 **Example 1** Three instances of subroutine **task1** are concurrently executed by node sets **p(1:500)**,
 10 **p(501:800)**, and **p(801:1000)**.

11

<pre> 11 XcalableMP Fortran 12 subroutine caller 13 !\$xmp nodes p(1000) 14 !\$xmp template tp(100) 15 !\$xmp distribute t(block) onto p 16 real a(100,100) 17 !\$xmp align a(*,k) with t(k) 18 ... 19 !\$xmp tasks 20 !\$xmp task on p(1:500) 21 call task1(a) 22 !\$xmp end task 23 !\$xmp task on p(501:800) 24 call task1(a) 25 !\$xmp end task 26 !\$xmp task on p(801:1000) 27 call task1(a) 28 !\$xmp end task 29 !\$xmp end tasks 30 ... 31 end subroutine </pre>	<pre> 32 XcalableMP Fortran 33 subroutine task1(a) 34 ... 35 !\$xmp nodes q(*)=* 36 ... 37 !\$xmp nodes p(1000) 38 !\$xmp distribute t(block) onto p 39 real a(100,100) 40 !\$xmp align a(*,k) with t(k) 41 ... 42 end subroutine </pre>
--	---

13 **Example 2** The first node **p(1)** executes the first and second tasks, the final node **p(8)** the
 14 second and the third tasks, and the other nodes **p(2)** through **p(7)** only the second task.

15

<pre> 16 XcalableMP Fortran 17 !\$xmp nodes p(8) 18 !\$xmp template t(100) 19 !\$xmp distribute t(block) onto p 20 real a(100) 21 !\$xmp align a(i) with t(i) 22 ... 23 !\$xmp tasks 24 ... 25 !\$xmp task on t(1) </pre>	
---	--

```

10      a(1) = 0.0
!$xmp end task

!$xmp task on t(2:99)
!$xmp loop on t(i)
15      do i=2,99
          a(i) = foo(i)
        enddo
!$xmp end task

20 !$xmp task on t(100)
      a(100) = 0.0
!$xmp end task

!$xmp end tasks

```

4.4.3 loop Construct

Synopsis

The loop construct specifies that each iteration of the following loop is executed by a node set that is specified by the on clause, so the iterations are distributed among nodes and executed in parallel.

Syntax

[F] !\$xmp loop [(*loop-index* [, *loop-index*]...)] on {*nodes-ref* | *template-ref*} ■
 ■ [*expand*(*expand-width* [, *expand-width*]...)] ■
 ■ [*margin*(*margin-width* [, *margin-width*]...)] ■
 ■ [*reduction-clause*]...
do-loops

[C] #pragma xmp loop [(*loop-index* [, *loop-index*]...)] on {*nodes-ref* | *template-ref*} ■
 ■ [*expand*(*expand-width* [, *expand-width*]...)] ■
 ■ [*margin*(*margin-width* [, *margin-width*]...)] ■
 ■ [*reduction-clause*]...
for-loops

where *expand-width* and *margin-width* must be one of:

```

[/unbound/] int-expr
[/unbound/] int-expr : int-expr

```

reduction-clause is:

```
reduction( reduction-kind : reduction-spec [, reduction-spec ]... )
```

reduction-kind is one of:

```

[F] +
    *
    -
    .and.
    .or.
    .eqv.
    .neqv.
    max
    min
    iand
    ior
    ieor
    firstmax
    firstmin
    lastmax
1    lastmin

```

```

[C] +
    *
    -
    &
    |
    ^

    &&
    ||
    max
    min
    firstmax
    firstmin
    lastmax
    lastmin

```

2 and *reduction-spec* is:

```

3     reduction-variable [ / location-variable [, location-variable ]... / ]

```

4 Description

5 A `loop` directive is associated with a loop nest consisting of one or more tightly nested loops
6 that follow the directive, and it distributes the execution of their iterations onto the node set
7 specified by the `on` clause.

8 The sequence of *loop-indices* in parenthesis denotes an index of an iteration of the loop nests.
9 If a control variable of a loop does not appear in the sequence, it is assumed that each of its
10 possible values is specified in the sequence. The sequence can be considered to denote a set of
11 indices of iterations. When the sequence is omitted, it is assumed that the control variables of
12 all the loops in the associated loop nests are specified.

13 When a *template-ref* is specified in the `on` clause, the associated loop is distributed so that
14 the iteration (set) indexed by the sequence of *loop-indices* is executed by the node onto which
15 a template element specified by the *template-ref* is distributed.

16 When a *nodes-ref* is specified in the `on` clause, the associated loop is distributed so that the
17 iteration (set) indexed by the sequence of *loop-indices* is executed by a node specified by the

nodes-ref.

In addition, the executing node set is updated to the node set specified by the `on` clause at the beginning of every iteration, and it is restored to the last one at the end of it.

When a *reduction-clause* is specified, a reduction operation of the kind specified by *reduction-kind* for a variable specified by *reduction-variable* is executed just after the execution of the loop nest.

When the `expand` clause is specified, and is of the form “*int-expr* : *int-expr*” in a dimension, the first *int-expr* is subtracted from the local lower bound in that dimension, and the second one is added to the local upper bound. When the `expand` clause is specified, and is of the form *int-expr*, the *int-expr* is subtracted from the local lower bound in that dimension, and is added to the local upper bounds. However, an “expanded” local iteration space does not spread out of the original global iteration space unless the `/unbound/` modifier is specified in *expand-width*.

When the `margin` clause is specified, the loop is transformed so that its local iteration space, *margin*, is:

$$\textit{margin} = \textit{expand} \triangle \textit{orig}$$

where *expand* is a local iteration space when an `expand` clause with the same argument(s) is specified, *orig* is a local iteration space when neither *expand* nor *margin*, and \triangle is the symmetric difference operator.

(Advice to programmers and implementers) Using the `expand` and `margin` clauses and asynchronous communication, programmers can overlap computation and communication as in the code left below. It is recommended for the implementation to support an extension that is a syntactic sugar for those sequence of constructs, such as the `peel_and_wait` clause in the code immediately following.

```

XcalableMP Fortran
!$xmp reflect (a) async(10)

!$xmp loop (i,j) on t(i,j)
!$xmp+      expand(-1,-1)
5   do j = 1, 16
      do i = 1, 16
          ...
      end do
    end do

!$xmp wait_async (10)

!$xmp loop (i,j) on t(i,j)
!$xmp+      margin(-1,-1)
15  do j = 1, 16
      do i = 1, 16
          ...
      end do
    end do

```

```

XcalableMP Fortran
!$xmp reflect (a) async(10)

!$xmp loop (i,j) on t(i,j)
!$xmp+ peel_and_wait(10, -1,-1)
5   do j = 1, 16
      do i = 1, 16
          ...
      end do
    end do

```

The reduction operation that is executed, except in cases with *reduction-kind* of `FIRSTMAX`, `FIRSTMIN`, `LASTMAX`, or `LASTMIN`, is equivalent to the `reduction` construct with *reduction-kind* of “+” for “-” in the clause and the same *reduction-kind* for the other kinds, the same *reduction-variable*, and an `on` clause obtained from that of the `loop` directive by replacing each *loop-index*

1 in the *nodes-ref* or the *template-ref* with a triplet representing the range of its value. As an
 2 example, the two codes below are therefore equivalent.

<pre> XcalableMP Fortran !\$xmp loop (j) on t(:,j) !\$xmp+ reduction(op:s) do j = js, je ... do i = 1, N s = s op a(i,j) end do ... end do </pre>	<pre> XcalableMP Fortran ! Initialize s_tmp to the identity ! element of the op operator s_tmp = ... !\$xmp loop (j) on t(:,j) do j = js, je ... do i = 1, N s_tmp = s_tmp op a(i,j) end do ... end do !\$xmp reduction(op:s_tmp) !\$xmp+ on t(*,js:je) s = s op s_tmp </pre>
--	---

4 In particular, for the reduction kinds of FIRSTMAX, FIRSTMIN, LASTMAX, and LASTMIN, in
 5 addition to a corresponding MAX or MIN reduction operation, the *location-variables* are set after
 6 executing the loop construct as follows:

- 7 • For FIRSTMAX and FIRSTMIN, they are set to their values at the end of the *first* iteration
 8 in which the *reduction-variable* takes the value of the reduction result, where *first* refers
 9 to the first position in the sequential order in which iterations of the associated loop nest
 10 were executed without parallelization.
- 11 • For LASTMAX and LASTMIN, they are set to their values at the end of the *last* iteration in
 12 which the *reduction-variable* takes the value of the reduction result, where *last* refers to the
 13 last position in the sequential order in which iterations of the associated loop nest were
 14 executed without parallelization.

15 Restrictions

- 16 • *loop-index* must be a control variable of a loop in the associated loop nest.
- 17 • A control variable of a loop can appear as *loop-index* at most once.
- 18 • The node set specified by *nodes-ref* or *template-ref* in the *on* clause must be a subset of
 19 the parent node set.
- 20 • The template specified by *template-ref* must be fixed before the loop construct is executed.
- 21 • The loop construct is global, which means that it must be executed by all of the executing
 22 nodes with the same values for each local variable referenced in the directive, and the lower
 23 bound, upper bound, and step of the associated loop.
- 24 • Either of the *expand* or *margin* clause, if any, can be specified.
- 25 • The number of *expand-width*, if any, must be equal to the number of dimensions (or rank)
 26 of the template specified by *template-ref* or of the node array specified by *node-ref*.

- The number of *margin-width*, if any, must be equal to the number of dimensions (or rank) of the template specified by *template-ref* or of the node array specified by *node-ref*.
- *reduction-spec* must have one or more *location-variable*'s if and only if *reduction-kind* is either FIRSTMAX, FIRSTMIN, LASTMAX, or LASTMIN.

Examples

Example 1

```

XcalableMP Fortran
!$xmp distribute t(block) onto p
!$xmp align (i) with t(i) :: a, b
...
!$xmp loop (i) on t(i)
5   do i = 1, N
      a(i) = 1.0
      b(i) = a(i)
   end do

```

The loop construct determines the node that executes each of the iterations, according to the distribution of template *t*, and distributes the execution. This example is syntactically equivalent to the one shown below, but will be faster because the iterations to be executed by each node can be determined before executing the loop.

```

XcalableMP Fortran
!$xmp distribute t(block) onto p
!$xmp align (i) with t(i) :: a, b
...
5   do i = 1, N
!$xmp task on t(i)
      a(i) = 1.0
      b(i) = a(i)
!$xmp end task
   end do

```

Example 2

```

XcalableMP Fortran
!$xmp distribute t(*,block) onto p
!$xmp align (i,j) with t(i,j) :: a, b
...
!$xmp loop (i,j) on t(i,j)
5   do j = 1, M
      do i = 1, N
          a(i,j) = 1.0
          b(i,j) = a(i,j)
      end do
10  end do

```

Because the first dimension of template *t* is not distributed, only the *j* loop, which is aligned with the second dimension of *t*, is distributed. This example is syntactically equivalent to the task construct shown below.


```

XcalableMP Fortran
!$xmp distribute t(*,block) onto p
!$xmp align (*,j) with t(*,j) :: a, b
...
do j = 1, M
5 !$xmp task on t(*,j)
do i = 1, N
a(i,j) = 1.0
b(i,j) = a(i,j)
end do
10 !$xmp end task
end do

```

1 Example 3

```

XcalableMP Fortran
!$xmp distribute t(block,block) onto p
!$xmp align (i,j) with t(i,j) :: a, b
...
!$xmp loop (i,j) on t(i,j)
5 do j = 1, M
do i = 1, N
a(i,j) = 1.0
b(i,j) = a(i,j)
end do
10 end do

```

2 The distribution of loops in the nested loop can be specified using the sequence of *loop-*
3 *indexes* in one loop construct. This example is equivalent to the loop shown below, but
4 will run faster because the iterations to be executed by each node can be determined
5 outside of the nested loop. Note that the node set specified by the inner `on` clause is a
6 subset of that specified by the outer one.

```

XcalableMP Fortran
!$xmp distribute t(block,block) onto p
!$xmp align (i,j) with t(i,j) :: a, b
...
!$xmp loop (j) on t(:,j)
5 do j = 1, M
!$xmp loop (i) on t(i,j)
do i = 1, N
a(i,j) = 1.0
b(i,j) = a(i,j)
10 end do
end do

```

7 Example 4

```

XcalableMP Fortran
!$xmp nodes p(10,3)
...
!$xmp loop on p(:,i)
do i = 1, 3

```

```

5      call subtask ( i )
      end do

```

Three node sets $p(:,1)$, $p(:,2)$, and $p(:,3)$ are created as the executing node sets, and each of them executes iterations 1, 2, and 3 of the associated loop, respectively. This example is equivalent to the loop containing task constructs (below left) or static tasks/task constructs (below right).

```

XcalableMP Fortran
!$xmp nodes p(10,3)
...
do i = 1, 3
!$xmp task on p(:,i)
5      call subtask ( i )
!$xmp end task
end do

```

```

XcalableMP Fortran
!$xmp nodes p(10,3)
...
!$xmp tasks
!$xmp task on p(:,1)
5      call subtask ( 1 )
!$xmp end task
!$xmp task on p(:,2)
5      call subtask ( 2 )
!$xmp end task
!$xmp task on p(:,3)
10     call subtask ( 3 )
!$xmp end task
!$xmp end tasks

```

Example 5

```

XcalableMP Fortran
...
lb(1) = 1
iub(1) = 10
5 lb(2) = 11
iub(2) = 25
lb(3) = 26
iub(3) = 50
!$xmp loop (i) on p(lb(i):iub(i))
10 do i = 1, 3
    call subtask ( i )
end do

```

The executing node sets of different sizes are created by $p(lb(i):iub(i))$ with different values of i for unbalanced workloads. This example is equivalent to the loop containing task constructs (below left) or static tasks/task constructs (below right).

<pre style="margin: 0;"> XcalableMP Fortran do i = 1, 3 !\$xmp task on p(lb(i):iub(i)) call subtask (i) !\$xmp end task end do ... </pre>	<pre style="margin: 0;"> XcalableMP Fortran !\$xmp tasks !\$xmp task on p(1:10) call subtask (1) !\$xmp end task !\$xmp task on p(11:25) call subtask (2) !\$xmp end task !\$xmp task on p(25:50) call subtask (3) !\$xmp end task !\$xmp end tasks </pre>
---	--

2 Example 6

```

XcalableMP Fortran
...
s = 0.0
!$xmp loop (i) on t(i) reduction(+:s)
do i = 1, N
    s = s + a(i)
end do

```

This loop computes the sum of $a(i)$ into the variable s on each node. Note that only the partial sum is computed on s without the reduction clause. This example is equivalent to the code given below.

```

XcalableMP Fortran
...
s = 0.0
!$xmp loop (i) on t(i)
do i = 1, N
    s = s + a(i)
end do
!$xmp reduction(+:s) on t(1:N)

```

6 Example 7

```

XcalableMP Fortran
...
amax = -1.0e30
ip = -1
jp = -1
!$xmp loop (i,j) on t(i,j) reduction(firstmax:amax/ip,jp/)
do j = 1, M
    do i = 1, N
        if( 1(i,j) .gt. amx ) then
            amx = a(i,j)
            ip = i
            jp = j
        end if
    end do
end do

```

```

        end do
    end do

```

This loop computes the maximum value of $a(i,j)$ and stores it into the variable `amax` in each node. In addition, the first indices for the maximum element of a are obtained in `ip` and `jp` after executing the loops. Note that this example cannot be written using the `reduction` construct.

Example 8

```

XcalableMP Fortran
!$xmp loop (i,j) on t(i,j) expand(/unbound/1,/unbound/1)
  do j = 1, 16
    do i = 1, 16
      ...
    end do
  end do

!$xmp loop (i,j) on t(i,j) margin(/unbound/1,/unbound/1)
  do j = 1, 16
    do i = 1, 16
      ...
    end do
  end do

```

Assuming that the template $t(100,100)$ is distributed in (block,block) onto a node array $p(4,4)$, the original local iteration space on $p(1,1)$, $orig_{1,1}$ is:

$$orig_{1,1} = \{ \begin{array}{cccc} (1,1), & (2,1), & (3,1), & (4,1), \\ (1,2), & (2,2), & (3,2), & (4,2), \\ (1,3), & (2,3), & (3,3), & (4,3), \\ (1,4), & (2,4), & (3,4), & (4,4) \end{array} \}$$

and it is expanded using the `expand` clause for the first loop, as follows:

$$expand(1,1)_{1,1} = \{ \begin{array}{cccccc} (0,0), & (0,1), & (0,2), & (0,3), & (0,4), & (0,5), \\ (1,0), & (1,1), & (1,2), & (1,3), & (1,4), & (1,5), \\ (2,0), & (2,1), & (2,2), & (2,3), & (2,4), & (2,5), \\ (3,0), & (3,1), & (3,2), & (3,3), & (3,4), & (3,5), \\ (4,0), & (4,1), & (4,2), & (4,3), & (4,4), & (4,5), \\ (5,0), & (5,1), & (5,2), & (5,3), & (5,4), & (5,5) \end{array} \}$$

Note that $expand(1,1)_{1,1}$ spreads out of the original global iteration space $\{(i,j) \mid 1 \leq i, j \leq 16\}$ because the `/unbound/` specifier is specified in the `expand` clause.

The local iteration space for the second loop with the `margin` clause is defined using the symmetric difference operator, as follows:

$$\begin{aligned} margin(1,1)_{1,1} &= expand(1,1)_{1,1} \Delta orig_{1,1} \\ &= \{ \begin{array}{cccccc} (0,0), & (0,1), & (0,2), & (0,3), & (0,4), & (0,5), \\ (1,0), & & & & & (1,5), \\ (2,0), & & & & & (2,5), \\ (3,0), & & & & & (3,5), \\ (4,0), & & & & & (4,5), \\ (5,0), & (5,1), & (5,2), & (5,3), & (5,4), & (5,5) \end{array} \} \end{aligned}$$

1 4.4.4 array Construct

2 Synopsis

3 The array construct divides the work of an array assignment between nodes.

4 Syntax

[F] !\$xmp array on *template-ref*
array-assignment-statement

5

[C] #pragma xmp array on *template-ref*
array-assignment-statement

6 Description

7 The array assignment is an alternative to a loop that performs an assignment to each element
 8 of an array. This directive specifies the parallel execution of an array assignment, where each
 9 sub-assignment and sub-operation of an element is executed by a node that is determined by
 10 the `on` clause.

11 Note that array assignments can also be used in XcalableMP C, which is one of the language
 12 extensions introduced by XcalableMP (see Section 3.2).

13 Restrictions

- 14 • The node set specified by *template-ref* in the `on` clause must be a subset of the parent node
 15 set.
- 16 • The template section specified by *template-ref* must have the same shape as the associated
 17 array assignment.
- 18 • The `array` construct is global and must be executed by all of the executing nodes with
 19 the same values for the variables that appear in the construct.

20 Examples

21 Example 1

```

_____ XcalableMP Fortran _____
!$xmp distribute t(block) onto p
!$xmp align (i) with t(i) :: a
...
!$xmp array on t(1:N)
5   a(1:N) = 1.0

```

22 This example is equivalent to the code shown below.

```

_____ XcalableMP Fortran _____
!$xmp distribute t(block) onto p
!$xmp align (i) with t(i) :: a
...
!$xmp loop on t(1:N)
5   do i = 1, N
      a(i) = 1.0
   end do

```

Example 2

```

XcalableMP Fortran
!$xmp template t(100,20)
!$xmp distribute t(block,block) onto p
    dimension a(100,20), b(100,20)
!$xmp align (i,j) with t(i,j) :: a, b
5     ...
!$xmp array on t
    a = b + 2.0

```

This example is equivalent to the code shown below.

```

XcalableMP Fortran
!$xmp template t(100,20)
!$xmp distribute t(block,block) onto p
    dimension a(100,20), b(100,20)
!$xmp align (i,j) with t(i,j) :: a, b
5     ...
!$xmp loop (i,j) on t(i,j)
    do j = 1, 20
        do i = 1, 100
            a(i,j) = b(i,j) + 2.0
10        end do
    end do

```

4.5 Global-view Communication and Synchronization Constructs**4.5.1 reflect Construct****Synopsis**

The `reflect` construct assigns the value of a reflection source to the corresponding shadow object.

Syntax

```

[F] !$xmp reflect ( array-name [, array-name...] ) ■
    ■ [width ( reflect-width [, reflect-width]... )] [orthogonal] [async ( async-id )]
[C] #pragma xmp reflect ( array-name [, array-name...] ) ■
    ■ [width ( reflect-width [, reflect-width]... )] [orthogonal] [async ( async-id )]

```

where *reflect-width* must be one of:

```

[/periodic/] int-expr
[/periodic/] int-expr : int-expr

```

Description

The `reflect` construct updates each of the shadow objects of the array specified by *array-name* with the value of its corresponding reflection source. Note that the shadow objects corresponding to elements at the non-orthogonal positions are also updated with this construct, unless the `orthogonal` clause is specified.

1 When the `width` clause is specified and takes the form “*int-expr* : *int-expr*” in a dimension,
 2 the shadow area having the width specified by the first *int-expr* at the lower bound and that
 3 specified by the second one at the upper bound in the dimension are updated. When the `width`
 4 clause is specified, and takes the form *int-expr*, the shadow areas having the same width specified
 5 at both the upper and lower bounds in the dimension are updated. When the `width` clause is
 6 omitted, the whole shadow area of the array is updated.

7 In particular, when the `/periodic/` modifier is specified in *reflect-width*, the update of the
 8 shadow object in the dimension is “periodic,” which means that the shadow object at the global
 9 lower (upper) bound is treated as if it corresponds to the data object of the global upper (lower)
 10 bound, and is updated with that value by the `reflect` construct.

11 When the `orthogonal` clause is specified, only the shadow objects corresponding to elements
 12 at the orthogonal positions are updated by the `reflect` construct.

13 When the `async` clause is specified, the statements following this construct may be executed
 14 before the operation is complete.

15 Restrictions

- 16 • The arrays specified by the sequence of *array-names* must be mapped onto the executing
 17 node set.
- 18 • The reflect width of each dimension specified by the *reflect-width* must not exceed the
 19 shadow width of the arrays.
- 20 • The `reflect` construct is global, which means that it must be executed by all nodes in
 21 the current executing node set, and each local variable referenced in the construct must
 22 have the same value among all of them.
- 23 • *async-id* must be an expression of type default integer in XcalableMP Fortran or type `int`
 24 in XcalableMP C.

25 Example

```

XcalableMP Fortran
!$xmp nodes p(4)
!$xmp template t(100)
!$xmp distribute t(block) onto p
5
   real a(100)
!$xmp align a(i) with t(i)
!$xmp shadow a(1)
   ...
10 !$xmp reflect (a) width (/periodic/1)
```




Figure 4.2: Example of periodic shadow reflection.

The `shadow` directive allocates “periodic” shadow areas of the array `a`. The `reflect` construct updates “periodically” the shadow area of `a` (Figure 4.2). A periodic shadow at the lower bound on the node `p(1)` is updated with the value of `a(100)` and that at the upper bound on `p(4)` with the value of `a(1)`.

4.5.2 `gmove` Construct

Synopsis

The `gmove` construct allows an assignment statement, which may cause communication, to be executed possibly in parallel by the executing nodes.

Syntax

```
[F] !$xmp gmove [in | out] [async ( async-id )]
[C] #pragma xmp gmove [in | out] [async ( async-id )]
```

Description

This construct copies the value of the right-hand side variable into the left-hand side of the associated assignment statement, which may cause communication between the executing nodes. Such communication is detected, scheduled, and performed by the XcalableMP runtime system.

There are three operating modes of the `gmove` construct:

- **collective mode**

When neither the `in` nor the `out` clause is specified, the copy operation is performed collectively, and results in implicit synchronization among the executing nodes.

If the `async` clause is not specified, then the construct is “synchronous,” and it is guaranteed that the left-hand side data can be read and overwritten, the right-hand side data can be overwritten, and all of the operations of the construct on the executing nodes are completed when returning from the construct; otherwise, the construct is “asynchronous,” and it is not guaranteed that the operations are completed, until the associating `wait_async` construct (Section 4.5.6) is completed.

- **in mode**

When the `in` clause is specified, the right-hand side data of the assignment, all or part of which may reside outside the executing node set, can be transferred from its owner nodes to the executing nodes by this construct.

If the `async` clause is not specified, then the construct is “synchronous,” and it is guaranteed that the left-hand side data can be read and overwritten, and that all of the operations of the construct on the owner nodes of the right-hand side and the executing nodes are completed when returning from the construct; otherwise, the construct is “asynchronous,” and it is not guaranteed that the operations are completed, until the associating `wait_async` construct (Section 4.5.6) is completed.

- **out mode**

When the `out` clause is specified, the left-hand side data of the assignment, all or part of which may reside outside the executing node set, can be transferred from the executing nodes to its owner nodes by this construct.

If the `async` clause is not specified, then the construct is “synchronous,” and it is guaranteed that the right-hand side data can be overwritten, and that all of the operations of

1 the construct on the owner nodes of the left-hand side and the executing nodes are com-
 2 pleted when returning from the construct; otherwise, the construct is “asynchronous,” and
 3 it is not guaranteed that the operations are completed, until the associating `wait_async`
 4 construct (Section 4.5.6) is completed.

5 When the `async` clause is specified, the statements following this construct may be executed
 6 before the operation is complete.

7 Restrictions

- 8 • The `gmove` construct must be followed by (i.e., associated with) a simple assignment state-
 9 ment that contains neither arithmetic operations nor function calls.
- 10 • The `gmove` construct is global, which means that it must be executed by all nodes in the
 11 current executing node set, and each local variable referenced in the construct must have
 12 the same value.
- 13 • If the `gmove` construct is in the *collective* mode, then all elements of the distributed arrays
 14 appearing on both the left-hand side and the right-hand side of the associated assignment
 15 statement must reside in the executing node set.
- 16 • If the `gmove` construct is in the *in* mode, then all elements of the distributed array appearing
 17 on the left-hand side of the associated assignment statement must reside in the executing
 18 node set.
- 19 • If the `gmove` construct is in the *out* mode, then all elements of the distributed array
 20 appearing on the right-hand side of the associated assignment statement must reside in
 21 the executing node set.
- 22 • *async-id* must be an expression of type default integer in XcalableMP Fortran or type `int`
 23 in XcalableMP C.

24 Examples

25 **Example 1: Array assignment** If the arrays on both the left-hand side and the right-hand
 26 side are distributed, then the copy operation is performed using all-to-all communication.
 27 If the left-hand side is a replicated array, this copy is performed using multi-cast commu-
 28 nication. If the right-hand side is a replicated array, then no communication is required.

	XcalableMP Fortran		XcalableMP C
29	<pre>!\$xmp gmove a(:,1:N) = b(:,3,0:N-1)</pre>		<pre>#pragma xmp gmove a[1:N][:] = b[0:N][3][:];</pre>

30 **Example 2: Scalar assignment to an array** When the right-hand side is an element of a
 31 distributed array, the copy operation is performed by broadcast communication from the
 32 owner of the element. If the right-hand side is a replicated array, then no communication
 33 is required.

	XcalableMP Fortran		XcalableMP C
34	<pre>!\$xmp gmove a(:,1:N) = c(k)</pre>		<pre>#pragma xmp gmove a[1:N][:] = c[k]</pre>

Example 3: in mode assignment Because `b(3)` referenced on the right-hand side of the `gmove` construct does not reside in the executing node set (`p(1:2)`), the construct is executed in the *in* mode. Thus, `b(3)` is transferred from its owner node `p(3)` to the executing node set.

Until `p(1:2)` returns from the construct, there is no gurantee that any node can read and overwrite `a(1:2)`, and that any relevant operations on `p(1:2)` and `p(3)` are completed.

```

XcalableMP Fortran
!$xmp nodes p(4)
!$xmp template t(4)
!$xmp distribute t(block) onto p

5      real a(4), b(4)
!$xmp align (i) with t(i) : a, b
      ...
!$xmp task on p(1:2)
      ...
10    !$xmp gmove in
      a(1:2) = b(2:3)
      ...
!$xmp end task

```

4.5.3 barrier Construct

Synopsis

The `barrier` construct specifies an explicit barrier at the point at which the construct appears.

Syntax

```
[F]  !$xmp barrier [on nodes-ref | template-ref]
[C]  #pragma xmp barrier [on nodes-ref | template-ref]
```

Description

The barrier operation is performed among the node set specified by the `on` clause. If no `on` clause is specified, then it is assumed that the current executing node set is specified in it.

Note that an `on` clause may represent multiple node sets. In such a case, a barrier operation is performed in each node set.

Restriction

- The node set specified by the `on` clause must be a subset of the executing node set.

4.5.4 reduction Construct

Synopsis

The `reduction` construct performs a reduction operation among nodes.

1 **Syntax**

2 [F] `!$xmp reduction (reduction-kind : variable [, variable]...)` ■
 ■ [*on node-ref* | *template-ref*] [*async* (*async-id*)]

3 where *reduction-kind* is one of:

4 +
 *
 .and.
 .or.
 .eqv.
 .neqv.
 max
 min
 iand
 ior
 ieor

5 [C] `#pragma xmp reduction (reduction-kind : variable [, variable]...)` ■
 ■ [*on node-ref* | *template-ref*] [*async* (*async-id*)]

6 where *reduction-kind* is one of:

7 +
 *
 &
 |
 ~
 &&
 ||
 max
 min

8 **Description**

9 The `reduction` construct performs a type of reduction operation specified by *reduction-kind*
 10 for the specified local variables among the node set specified by the `on` clause, and it sets
 11 the reduction results to the variables on each of the nodes. Note that some of the reduction
 12 operations, namely, `FIRSTMAX`, `FIRSTMIN`, `LASTMAX`, and `LASTMIN`, which can be specified in the
 13 `reduction` clause of the `loop` directive, cannot be specified in the `reduction` construct because
 14 their semantics are not defined for it. The variable specified by *variable*, which is the target of
 15 the reduction operation, is referred to as the “reduction variable.” After the reduction operation,
 16 the value of a reduction variable becomes the same in every node that performs the operation.

17 The reduction result is computed by combining the reduction variables on all of the nodes
 18 using the reduction operator. The ordering of this reduction is unspecified.

19 When the `async` clause is specified, the statements following this construct may be executed
 20 before the operation is complete.

21 When *template-ref* is specified in the `on` clause, the operation is performed in a node set
 22 that consists of nodes onto which the specified template section is distributed. Therefore, before
 23 the `reduction` construct is executed, the referenced template must be fixed. When *nodes-ref*
 24 is specified in the `on` clause, the operation is performed in the specified node set. When the `on`
 25 clause is omitted, the operation is performed in the executing node set.

Note that an `on` clause may represent multiple node sets. In such a case, a reduction operation is performed in each node set.

Restrictions

- The variables specified by the sequence of *variables* must either not be aligned or must be replicated among nodes of the node set specified by the `on` clause.
- The `reduction` construct is global, which means that it must be executed by all nodes in the current executing node set, and each local variable referenced in the construct must have the same value.
- *async-id* must be an expression of type default integer in XcalableMP Fortran or type `int` in XcalableMP C.
- The node set specified by the `on` clause must be a subset of the executing node set.

Examples

Example 1

```

XcalableMP Fortran
!$xmp reduction(+:s)
!$xmp reduction(max:aa) on t(*,:)
!$xmp reduction(min:bb) on p(10:30)

```

In the first line, the reduction operation calculates the sum of the scalar variable `s` in the executing node set, and the result is stored in the variable in each node.

The reduction operation in the second line computes the maximum value of the variable `aa` in each node set onto which each of the template sections specified by `t(*, :)` is distributed.

In the third line, the minimum value of the variable `bb` in the node set specified by `p(10:30)` is calculated. This example is equivalent to the following code using the `task` construct.

```

XcalableMP Fortran
!$xmp task on p(10:30)
!$xmp reduction(min:bb)
!$xmp end task

```

Example 2

```

XcalableMP Fortran
dimension a(n,n), p(n), w(n)
!$xmp align a(i,j) with t(i,j)
!$xmp align p(i) with t(i,*)
!$xmp align w(j) with t(*,j)
...
!$xmp loop (j) on t(:,j)
do j = 1, n
sum = 0
!$xmp loop (i) on t(i,j) reduction(+:sum)
do i = 1, n
sum = sum + a(i,j) * p(i)
end do

```

```

        w(j) = sum
    end do

```

1 This code computes the matrix vector product, where a **reduction** clause is specified for
2 the loop construct of the inner loop. This is equivalent to the following code snippet.

```

----- XcalableMP Fortran -----
!$xmp loop (j) on t(:,j)
    do j = 1, n
        sum = 0
!$xmp loop (i) on t(i,j)
5        do i = 1, n
            sum = sum + a(i,j) * p(i)
        end do
!$xmp reduction(+:sum) on t(1:n,j)
10        w(j) = sum
    end do

```

3 In these cases, the reduction operation on the scalar variable **sum** is performed for every
4 iteration in the outer loop, which may cause a large overhead. To reduce this overhead,
5 the **reduction** clause should be specified in the **loop** construct for the outer loop. This is
6 because the node set in which the reduction operation is performed is determined on the
7 basis of its **on** clause (see 4.4.3), and the **on** clause of the outer **loop** construct is different
8 from that of the inner one. However, this code can be modified using the **reduction**
9 construct as follows:

```

----- XcalableMP Fortran -----
    dimension a(n,n), p(n), w(n)
!$xmp align a(i,j) with t(i,j)
!$xmp align p(i) with t(i,*)
!$xmp align w(j) with t(*,j)
5    ...
!$xmp loop (j) on t(:,j)
    do j = 1, n
        sum = 0
!$xmp loop (i) on t(i,j)
10        do i = 1, n
            sum = sum + a(i,j) * p(i)
        end do
        w(j) = sum
    end do
15 !$xmp reduction(+:w) on t(1:n,*)

```

10 This code performs a reduction operation on the array **w** only once, which may result in
11 faster operation.

12 4.5.5 bcast Construct

13 Synopsis

14 The **bcast** construct performs broadcast communication from a specified node.

Syntax

```
[F] !$xmp bcast ( variable [, variable]... ) [from nodes-ref | template-ref]
      [on nodes-ref] | template-ref [async ( async-id )]
[C] #pragma xmp bcast ( variable [, variable]... ) [from nodes-ref | template-ref]
      [on nodes-ref] | template-ref [async ( async-id )]
```

Description

The values of the variables specified by the sequence of *variables* (called *broadcast variables*) are broadcasted from the node specified by the **from** clause (called the *source node*) to each of the nodes in the node set specified by the **on** clause. After executing this construct, the values of the broadcast variables become the same as those in the source node. If the **from** clause is omitted, then the *first* node, that is, the leading one in Fortran's array element order, of the node set specified by the **on** clause is assumed to be a source node. If the **on** clause is omitted, then it is assumed that the current executing node set is specified in it.

When the **async** clause is specified, the statements following this construct may be executed before the operation is complete.

Restrictions

- The variables specified by the sequence of *variables* must either not be aligned or must be replicated among nodes of the node set specified by the **on** clause.
- The **bcast** construct is global, which means that it must be executed by all nodes in the current executing node set, and each local variable referenced in the construct must have the same value among all of them.
- *async-id* must be an expression of type default integer in XcalableMP Fortran or type **int** in XcalableMP C.
- The node set specified by the **on** clause must be a subset of the executing node set.
- The source node specified by the **from** clause must belong to the node set specified by the **on** clause.
- The source node specified by the **from** clause must be one node.

4.5.6 wait_async Construct**Synopsis**

The **wait_async** construct guarantees that asynchronous communications specified by *async-id* are complete.

Syntax

```
[F] !$xmp wait_async ( async-id [, async-id ]... ) [on nodes-ref | template-ref]
[C] #pragma xmp wait_async ( async-id [, async-id ]... ) [on nodes-ref | template-ref]
```

Description

The **wait_async** construct will block, and therefore statements following it will not be executed, until the completion of all of the asynchronous communications that are specified by *async-id*'s and issued on the node set specified by the **on** clause. If an *async-id* is not associated with any asynchronous communication, the **wait_async** construct ignores it.

1 Restrictions

- 2 • *async-id* must be an expression of type default integer in XcalableMP Fortran or type `int`
3 in XcalableMP C.
- 4 • ~~*async-id* must be associated with an asynchronous communication using the `async` clause
5 of a communication construct.~~
- 6 • The `wait_async` construct is global, which means that it must be executed by all nodes in
7 the current executing node set, and each local variable referenced in the construct must
8 have the same value among all of them.
- 9 • The node set specified by the `on` clause must be the same as those of the global constructs
10 that initiate the asynchronous communications specified by *async-id*.

11 4.5.7 `async` Clause

12 Synopsis

13 The `async` clause of the `reflect`, `gmove`, `reduction`, and `bcast` constructs enables the corre-
14 sponding communication to be performed asynchronously.

15 Description

16 Communication corresponding to the construct with an `async` clause is performed asynchronously,
17 that is, it is initiated but not completed, and therefore, statements following it may be executed
18 before the communication is complete.

19 Example

XcalableMP Fortran

```
!$xmp reflect (a) async(1)
  S1
!$xmp wait_async(1)
  S2
```

20 The `reflect` construct on the first line matches the `wait` construct on the third line because
21 both of their *async-id* evaluate to one. These constructs ensure that statements in `S1` can be
22 executed before the `reflect` communication is complete, and no statement in `S2` is executed
23 until the `reflect` communication is complete.

24 4.5.8 `reduce_shadow` Construct

25 Synopsis

26 The `reduce_shadow` construct adds values of shadow objects to their reflection source.

27 Syntax

[F] `!$xmp reduce_shadow (array-name [, array-name]...)` ■
 ■ `[width (reflect-width [, reflect-width]...)] [orthogonal] [async (async-id)]`
 28 [C] `#pragma xmp reduce_shadow (array-name [, array-name]...)` ■
 ■ `[width (reflect-width [, reflect-width]...)] [orthogonal] [async (async-id)]`

Description

The `reduce_shadow` construct adds values of shadow objects of the array specified by *array-name* to their reflection source. Note that the shadow objects corresponding to elements at the non-orthogonal positions are also added as the default behavior.

When the `width` clause is specified and has the form “*int-expr* : *int-expr*” in a dimension, the shadow areas having the width specified by the first *int-expr* at the lower bound, and that specified by the second one at the upper bound in the dimension are added. When the `width` clause is specified and has the form *int-expr*, the shadow areas having the same width specified at both the upper and lower bounds in the dimension are added. When the `width` clause is omitted, the whole shadow area of the array is added.

In particular, when the `/periodic/` modifier is specified in *reflect-width*, the addition of the shadow object in the dimension is “periodic,” which means that the shadow object at the global lower (upper) bound is treated as if it corresponds to the data object of the global upper (lower) bound and is added by the `reduce_shadow` construct.

When the `orthogonal` clause is specified, the shadow object is added only by orthogonal nodes.

When the `async` clause is specified, the statements following this construct may be executed before the operation is complete.

Restrictions

- The arrays specified by the sequence of *array-names* must be mapped onto the executing node set.
- The width of each dimension specified by *reflect-width* must not exceed the shadow width of the arrays.
- The `reduce_shadow` construct is global, which means that it must be executed by all nodes in the current executing node set, and each local variable referenced in the construct must have the same value among all of them.
- *async-id* must be an expression of type default integer in XcalableMP Fortran or type `int` in XcalableMP C.

Examples

XcalableMP Fortran

```

    real rho(n,n)
!$xmp align rho(i,j) with t1(i,j)
!$xmp shadow rho(1:1)

5     real f(m)
       integer x(m), y(m)
!$xmp align (k) with t2(k) : f, x, y

!$xmp loop on t2(k)
10    do i = 1, no
       ix = x(i)
       iy = y(i)
       dx = x(i) - ix
       dy = y(i) - iy
15    rho(ix ,iy ) = rho(ix ,iy ) + (1.0-dx)*(1.0-dy)*f(i)

```



```

        rho(ix+1,iy ) = rho(ix+1,iy ) +      dx *(1.0-dy)*f(i)
        rho(ix  ,iy+1) = rho(ix  ,iy+1) + (1.0-dx)*      dy *f(i)
        rho(ix+1,iy+1) = rho(ix+1,iy+1) +      dx *      dy *f(i)
    end do
20 !$xmp reduce_shadow (rho)
    !$xmp reflect (rho)

```

- 1 Assume that a two-dimensional field `rho` and `m` particles are both distributed onto nodes.
- 2 On each node, a contribution of a particle `f(k)` is added to the nearest grid point of the field
- 3 and its neighbors, which may be in the shadow area on the node. In the last two lines, the
- 4 values of the shadow area from neighboring nodes are added to the corresponding data object,
- 5 and the results are then copied back to the shadow area on the neighboring nodes.

Chapter 5

Support for the Local-view Programming

This chapter describes the coarray features in XcalableMP, which are based on that of Fortran 2008. Note that they are also available in XcalableMP C. In addition, this chapter describes some directives for the local-view programming.

The coarray features in Fortran 2008 are extended and integrated into XcalableMP. The specifications in this chapter are designed to achieve the following purposes:

- Upward compatibility to the Fortran 2008 coarray features
If an XcalableMP Fortran program does not contain any XMP directives, any standard-conforming Fortran 2008 program remains standard conforming under XcalableMP. In this sense, the interpretations and extensions defined in this chapter are upward compatible with the Fortran International Standard, ISO/IEC 1539-1:2010 (Fortran 2008).
- Support for task parallelism
XcalableMP makes it possible to construct a task parallel program by combining multiple Fortran 2008 codes, which may be developed independently, with minimum modifications.
- Integration of global-view style programming and local-view style programming
In XcalableMP, users can appropriately use global-view style programming of XcalableMP or local-view style programming, which is typically used in MPI or Fortran 2008 programs, according to the characteristics of the code in a program.
- Possibility of support for multiple topologies of a computing system
An XcalableMP processor may allow users to specify the correspondence between node arrays and the topologies of a computing system, and to exploit the full potential of a particular system.

5.1 Rules Determining Image Index

This section defines how the image index of an image in a set of images is determined in association with a node array and a `task` construct.

5.1.1 Primary Image Index

Every image has a default image index in all of the images at the invocation of a program. In XcalableMP, the default image index is the primary image index, and is an integer value that ranges from one to the number of images at the invocation of a program.

A primary node array corresponds to all of the images at the invocation of a program, and it also corresponds to all of the nodes at the invocation of a program. The primary image index of an image is the (Fortran) subscript order value of the uniquely corresponding element of a primary node array.

5.1.2 Image Index Determined by a task Directive

The execution of a `task` directive determines that a set of nodes (and the corresponding set of images) forms an executing node set. If a name of a node array or a subobject of a node array appears in the `task` directive, the nodes and the corresponding images in the executing node set are ordered in (Fortran) array element order in the node array or the subobject of the node array. If a name of a template array or a subobject of a template array appears in the `task` directive, the nodes and the corresponding images in the executing node set are ordered in (Fortran) array element order in the corresponding subobject of the node array. The image index of an image in the determined set of images is the integer order value in the range one to the cardinality of the set of images.

5.1.3 Current Image Index

The image index of an image in the current set of images is the current image index.

A current executing node array corresponds to the current set of images and also the current executing node set at the evaluation of the declaration of the node array. Each image in the current set of images corresponds to the element of an executing node array whose subscript order value is the same as the current image index of the image when the evaluation of the declaration of the executing node array is being evaluated. In particular, when all `task` directive constructs are completed, the current image index of an image is the same as the primary image index.

5.1.4 Image Index Determined by a Non-primary Node Array

A non-primary node array corresponds to all of the images at the invocation of a program, and it also corresponds to all the nodes at the invocation of a program. The correspondence between each image and each element of a non-primary node array is processor dependent. A processor may support any means to specify the correspondence.

The image index of an image in all of the images at the invocation of a program is the subscript order value of the corresponding element of a non-primary node array. This is the case if and only if the current set of images corresponds to the non-primary node whole array in which the nodes in the executing node set are ordered in (Fortran) array-element order in the non-primary node whole array. The image index is a non-primary image index.

The correspondence between the primary image index and a non-primary image index of the same image is processor dependent. Between any two distinct non-primary node arrays, the correspondence between a non-primary image index and the other non-primary image index of the same image is processor dependent unless they have the same shape. If two non-primary node arrays have the same shape, the corresponding elements of the node arrays correspond to the same image.

5.1.5 Image Index Determined by an Equivalenced Node Array

A `nodes` directive with “=*node-ref*” that is not “=*” or “=**” specifies that each element of the declared node array corresponds in (Fortran) array-element order to that of the *node-ref*, which is the name of a node array or a subobject of a node array. The nodes in the declared node array and the corresponding images are ordered in (Fortran) array-element order in the

1 *node-ref*. The image index of an image in the set of images corresponding to the declared node
 2 array is the integer order value ranging from one to the cardinality of the set of images.

3 5.1.6 On-node Image Index

4 XcalableMP supports the `coarray` directive and the `image` directive to specify that an image
 5 index indicates the image corresponding to the element of a particular node array whose subscript
 6 order value is the same as the image index. The image index is an on-node image index for the
 7 specified node array. Because an evaluation of the declaration of a node array determines a set
 8 of images corresponding to the node array, the directives specify that the set of images is the
 9 “all images” for the image indices affected by the directives. In particular, the on-node image
 10 index for a primary node array is the primary image index.

11 5.2 Basic Concepts

12 In XcalableMP, “all images” in Fortran 2008 changes coupled with the execution of `task` con-
 13 structs, and refers to the current set of images. In particular, when an `allocate` statement is
 14 executed for which an *allocate-object* is a coarray, there is an implicit synchronization of all the
 15 images in the current set of images. On each image in the current set of images, execution of the
 16 segment following the statement is delayed until all other images in the set have executed the
 17 same statement the same number of times. When a `deallocate` statement is executed for which
 18 an *allocate-object* is a coarray, there is an implicit synchronization of all the images in the current
 19 set of images. On each image in the current set of images, execution of the segment following
 20 the statement is delayed until all other images in the set have executed the same statement the
 21 same number of times.

- 22 • When an allocatable coarray is allocated during the execution of `task` constructs, the
 23 coarray shall be subsequently deallocated before the completion of the `task` construct
 24 whose `task` directive is the most recently executed one in the `task` constructs that are not
 25 completed at the allocation.

26 The image index determined by an image selector indicates the current image index by
 27 default. Coarrays are visible within the range of the “all images,” and are accessed using the
 28 current image index by default. The image index that appears in an executable statement
 29 indicates the current image index by default.

30 5.2.1 Examples

- 31 • In the following code fragment, the value of a coarray `b` on the images 1, 2, 3, and 4,
 32 which constitute the executing node set and correspond to `node(5)`, `node(6)`, `node(7)`,
 33 and `node(8)` respectively, is defined with the value of the coarray `a` on `node(5)`.

```

                                XcalableMP Fortran
program xmpcoarray
!$xmp nodes node(8)=**      ! A primary node array.
!$xmp task on node(5:8)    ! The executing node set
    call sub                ! corresponds to node(5:8).
5 !$xmp end task
    end

subroutine sub
    real, save :: a[*], b[*] ! The images 1, 2, 3,
```

```

10      :                               ! and 4 correspond to node(5:8),
      b = a[1]                          ! respectively.

```

- In the following code fragment, an allocatable coarray `a` is allocated on the images 1, 2, 3, and 4, which constitute the executing node set and correspond to `node(5)`, `node(6)`, `node(7)`, and `node(8)`, respectively.

```

                                     XscalableMP Fortran
      program xmpcoarray
!$xmp nodes node(8)=**
!$xmp task on node(5:8) ! The executing node set
      call sub2          ! corresponds to node(5:8).
5 !$xmp end task
      end

      subroutine sub2
      real, allocatable :: a(:)[: ]
10      :
      allocate(a(0:99)[*])

```

Note

- The result of `xmp_num_nodes()` is always the same as that of `NUM_IMAGES()`.
- The result of `xmp_node_num()` is always the same as that of `THIS_IMAGE()`.
- In a `read` statement, an io-unit that is an asterisk identifies an external unit that is preconnected for a sequential formatted input only on the image whose primary image index is one.

5.3 coarray Directive

5.3.1 Purpose and Form of the coarray Directive

The `coarray` directive maps coarrays onto a node array and the set of images that corresponds to the node array. An image index determined by an image selector for a coarray that appears in a `coarray` directive always indicates the on-node image index for the node array; that is, the specified image corresponds to the node whose subscript order value in the node array is the same as the image index.

A coarray appearing in a `coarray` directive is an on-node coarray of the node array that is specified in the `coarray` directive.

```
[F] !$xmp coarray on node-name :: object-name-list
```

```
[C] #pragma xmp coarray on node-name :: object-name-list
```

- An *object-name* shall be a name of a coarray declared in the same scoping unit.
- The same *object-name* shall not appear more than once in `coarray` directives in a scoping unit.

- 1 • If an *object-name* is a name of an allocatable object, the current set of images at the
2 allocation and the deallocation of the object shall correspond to the node array specified
3 as the *node-name*, and the current image index of each image shall be the same as the
4 subscript order value of the corresponding element of the node array.
- 5 • If an *object-name* is the name of an allocated allocatable dummy argument, the set of
6 images onto which it is mapped shall be a subset of the set of images that has most
7 recently allocated the corresponding argument in the chain of argument associations.
- 8 • If an *object-name* is the name of a nonallocatable dummy argument whose ultimate argu-
9 ment has an allocatable attribute, the set of images onto which the *object-name* is mapped
10 shall be a subset of the set of images that has most recently allocated the corresponding
11 argument in the chain of argument associations.
- 12 • The image index determined by an image selector for an on-node coarray shall be within
13 the range of one to the size of the node array onto which the on-node coarray is mapped.
- 14 • THIS_IMAGE(COARRAY[,DIM]) shall be invoked by the image contained in the set of
15 images onto which the COARRAY argument is mapped if the COARRAY argument ap-
16 pears in a `coarray` directive.

17 Note

- 18 • The result value of THIS_IMAGE(COARRAY) is the sequence of cosubscript values for the
19 COARRAY argument that would specify the current image index of the invoking image,
20 if the COARRAY argument does not appear in a `coarray` directive. The result value
21 of THIS_IMAGE(COARRAY) is the sequence of cosubscript values for the COARRAY
22 argument that would specify the on-node image index of the invoking image for the node
23 array onto which the COARRAY argument is mapped if the COARRAY argument appears
24 in a `coarray` directive.
- 25 • The result value of THIS_IMAGE(COARRAY,DIM) is the value of cosubscript DIM in
26 the sequence of cosubscript values for the COARRAY argument that would specify the
27 current image index of the invoking image if the COARRAY argument does not appear in
28 a `coarray` directive. The result value of THIS_IMAGE(COARRAY,DIM) is the value of
29 cosubscript DIM in the sequence of cosubscript values for the COARRAY argument that
30 would specify the on-node image index of the invoking image for the node array onto which
31 the COARRAY argument is mapped if the COARRAY argument appears in a `coarray`
32 directive.

33 5.3.2 An Example of the `coarray` Directive

XcalableMP Fortran

```

module global
!$xmp nodes node(8)**
  real s[*]           ! The coarray s is always
!$xmp coarray on node :: s ! visible on node(1:8).
end global

program coarray
  use global
!$xmp task on node(5:8) ! The executing node set
  call sub             ! consists of node(5:8).

```

```

!$xmp end task
end

subroutine sub
use global
15 real, save :: a[*]    ! The images 1, 2, 3, and 4
    :                ! correspond to node(5:8), respectively.
if(this_image().eq.1)then ! The value of the coarray a on node(5)
    s[1] = a          ! defines that of the coarray s on node(1)
20 endif

```

5.4 image Directive 1

5.4.1 Purpose and Form of the image Directive 2

The `image` directive specifies that an image index in the following executable statement indicates the on-node image index of the node array specified in the `image` directive unless the image index is determined by an image selector. 3

The `image` directive also specifies that the execution of a `sync all` statement performs a synchronization of all of the images corresponding to the node array specified in the `image` directive. 4

[F] `!$xmp image (node-name)` 5

[C] `#pragma xmp image (node-name)` 6

- An `image` directive shall be followed by a `sync all` statement, an image control statement that contains *image-set*, or a reference to an intrinsic procedure that has `IMAGES` argument. 7

5.4.2 An Example of the image Directive 8

```

XcalableMP Fortran
module global
!$xmp nodes node(8)**
real s[*]          ! The coarray s is always visible
!$xmp coarray on node :: s ! on node(1:8).
5 end global

program image
use global
!$xmp tasks
10 !$xmp task on node(1:4)
    call subA ! The executing node set consists of node(1:4).
!$xmp end task
!$xmp task on node(5:8)
    call subB ! The executing node set consists of node(5:8).
15 !$xmp end task
!$xmp end tasks
end

```



```

subroutine subA
20 use global
   real, save :: a[*] ! The images 1, 2, 3, and 4
      :                ! correspond to node(1:4), respectively.
!$xmp image(node)      ! Synchronization between node(1:4) and
   sync images(5)      ! node(5).
25 a = s[1]             ! a on node(1:4) is defined using
      :                ! the value of s on node(1).
end subroutine

subroutine subB
30 use global
   real, save :: b[*] ! The images 1, 2, 3, and 4
      :                ! correspond to node(5:8), respectively.
   if(this_image() .eq. 1)then ! The image 1 indicates node(5).
      s[1] = b             ! s on node(1) is defined using the value of
35                          ! b on node(5).
!$xmp image(node)      ! Synchronization between
      sync images((/1,2,3,4/)) ! node(5) and node(1:4).
   endif
      :
40 end subroutine

```

1 5.5 Image Index Translation Intrinsic Procedures

2 XcalableMP supports intrinsic procedures to translate image indices between different sets of
3 images.

4 5.5.1 Translation to the Primary Image Index

5 `xmp_get_primary_image_index(NUMBER,INDEX,PRI_INDEX,NODE_DESC)`

6 **Description.** Translate image indices to the primary image indices.

7 **Class.** Subroutine.

8 **Arguments.**

9 **NUMBER** shall be a scalar of type default integer. It is an INTENT(IN) argument.

10 **INDEX** shall be a rank-one array of type default integer. The size of **INDEX** shall be
11 greater than or equal to the value of **NUMBER**. It is an INTENT(IN) argument.
12 The value of each element of **INDEX** shall be within the range one to the size of the
13 node array specified in **NODE_DESC** if **NODE_DESC** appears. The value of each
14 element of **INDEX** shall be within the range one to the cardinality of the current
15 set of images if **NODE_DESC** does not appear.

16 **PRI_INDEX** shall be a rank-one array of type default integer. The size of **PRI_INDEX**
17 shall be greater than or equal to the value of **NUMBER**. It is an INTENT(OUT)
18 argument. If **NODE_DESC** appears, **PRI_INDEX(i)** is assigned the primary im-
19 age index corresponding to the element of the node array specified in **NODE_DESC**
20 whose subscript order value is **INDEX(i)**; otherwise, **PRI_INDEX(i)** is assigned

the primary image index corresponding to the image whose current image index is **INDEX(i)**.

NODE_DESC (optional) shall be a descriptor of a node array. It is an **INTENT(IN)** argument. **NODE_DESC** shall appear in XcalableMP C.

Example. In the following code fragment, the value of `index(1:4)` is `(/5,6,7,8/)`.

```

XcalableMP Fortran
!$xmp nodes node(1:8)=**      ! A primary node array
!$xmp nodes subnode(4)=node(5:8)
  integer index(4)
  call xmp_get_primary_image_index&
5      &(4, (/1,2,3,4/), index, xmp_desc_of(subnode))

```

5.5.2 Translation to the Current Image Index

`xmp_get_image_index(NUMBER,INDEX,CUR_INDEX,NODE_DESC)`

Description. Translate image indices to the current image indices.

Class. Subroutine.

Arguments.

NUMBER shall be a scalar of type default integer. It is an **INTENT(IN)** argument.

INDEX shall be a rank-one array of type default integer. The size of **INDEX** shall be greater than or equal to the value of **NUMBER**. It is an **INTENT(IN)** argument. The value of each element of **INDEX** shall be within the range one to the size of the node array specified in **NODE_DESC**.

CUR_INDEX shall be a rank-one array of type default integer. The size of **CUR_INDEX** shall be greater than or equal to the value of **NUMBER**. It is an **INTENT(OUT)** argument. If the current image index corresponding to the element of the node-array specified in **NODE_DESC** whose subscript order value is **INDEX(i)** exists, **CUR_INDEX(i)** is assigned the current image index; otherwise, **CUR_INDEX(i)** is assigned zero.

NODE_DESC shall be a descriptor of a node array. It is an **INTENT(IN)** argument.

Example. In the following code fragment, the value of `index(1:4)` is `(/1,2,3,4/)`.

```

XcalableMP Fortran
!$xmp nodes node(1:8)=**
  integer index(4)
!$xmp task on node(5:8)
  call xmp_get_image_index&
5      &(4, (/5,6,7,8/), index, xmp_desc_of(node))
!$xmp end task

```

5.6 Examples of Communication between Tasks

- In the following program fragment, two tasks communicate with each other with synchronization.

```

                                XcalableMP Fortran
module nodes
!$xmp nodes node(8)=**           ! A primary node array
integer, parameter :: n=2
!$xmp nodes subnodeA(n)=node(1:n) ! subnodeA is for taskA.
5 !$xmp nodes subnodeB(8-n)=node(n+1:8) ! subnodeB is for taskB.
endmodule

module intertask
10 use nodes
real,save :: dA[*],dB[*]
endmodule

use nodes
!$xmp tasks
15 !$xmp task on subnodeA ! The taskA is invoked on subnodeA.
call taskA
!$xmp end task
!$xmp task on subnodeB ! The taskB is invoked on subnodeB.
call taskB
20 !$xmp end task
!$xmp end tasks
end

subroutine taskA
25 use intertask
:
me = this_image() ! The value of me is i on subnodeA(i).
if(me.eq.1)then
call xmp_get_primary_image_index& ! The value of iyouabs
30 &(1,(/1/),iyouabs,subnodeB) ! is n+1.
!$xmp image(node) ! Synchronization between
sync images(iyouabs) ! node(1) and node(n+1).
call exchange(dA,dB,iyouabs)
!$xmp image(node) ! Synchronization between
35 sync images(iyouabs) ! node(1) and node(n+1).
endif
sync all ! Synchronization within subnodeA.
if(me.ne.1)dA = dA[1]
sync all ! Synchronization within subnodeA.
40 :
end

subroutine taskB
45 use intertask
:
me = this_image() ! The value of me is i on subnodeB(i).
if(me.eq.1)then
call xmp_get_primary_image_index& ! The value of iyouabs
&(1,(/1/),iyouabs,subnodeA) ! is 1.

```

```

50 !$xmp  image(node)                ! Synchronization between
      sync images(iyouabs)          ! node(n+1) and node(1).
      call exchange(dB,dA,iyouabs)
!$xmp  image(node)                ! Synchronization between
      sync images(iyouabs)          ! node(n+1) and node(1).
55      endif
      sync all                      ! Synchronization within subnodeB.
      if(me.ne.1)dB = dB[1]
      sync all                      ! Synchronization within subnodeB.
60      end

      subroutine exchange(mine,yours,iput)
      use nodes
      real :: mine[*],yours[*]      ! mine and yours are always
65 !$xmp coarray on node :: mine,yours ! visible on node(1:8).

      yours[iput] = mine ! node(1) puts mine to yours[n+1] and
                          ! node(n+1) puts mine to yours[1].
      end

```

- In the following program fragment, two tasks communicate with each other without one-to-one synchronization.

1
2

```

XcalableMP Fortran
!$xmp nodes node(8)=**          ! A primary node array
:
!$xmp tasks
!$xmp  task on(node(1:n))
5      call taskA(n)            ! The taskA is invoked on node(1:n)
!$xmp  end task
!$xmp  task on(node(n+1:8))
      call taskB(8-n)          ! The taskB is invoked on node(n+1:8)
!$xmp  end task
10 !$xmp end tasks
    end

    subroutine taskA(n)
    real,save :: yours[*],mine[*]
15 !$xmp nodes subnode(n)=*      ! An executing node array

    me = this_image()
    if(me.eq. NUM_IMAGES())then
        call xmp_get_primary_image_index(1,me,meabs) ! meabs=n.
20        call exchange(yours,mine,meabs,meabs+1,NUM_IMAGES())
    endif
    sync all                      ! Synchronization within node(1:n).
    if(me.ne.NUM_IMAGES())mine = mine[NUM_IMAGES()]
    sync all                      ! Synchronization within node(1:n).
25    end

```

```

    subroutine taskB(m)
    real,save :: yours[*],mine[*]
!$xmp nodes subnode(m)=*          ! An executing node array
30
    me = this_image()
    if(me.eq.1)then
        call xmp_get_abs_image_index(1,me,meabs) ! meabs=n+1.
        call exchange(yours,mine,meabs,meabs-1,NUM_IMAGES())
35
    endif
    sync all                        ! Synchronization within node(n+1:8).
    if(me.ne.1)mine = mine[1]
    sync all                        ! Synchronization within node(n+1:8).
    end

40
    subroutine exchange(yours,mine,meabs,iyouabs,nnodes)
    USE, INTRINSIC :: ISO_FORTRAN_ENV
    real :: yours[*],mine[*]
    real, save :: s[*]              ! only for exchange.
45    TYPE(LOCK_TYPE),save :: lock[*] ! for lock.
!$xmp nodes subnode(nnodes)=*     ! An executing node array.
!$xmp nodes node(8)=**           ! The coarrays s and lock are
!$xmp coarray on node :: s,lock   ! always visible on node(1:8).

50    LOCK(lock[meabs]) ! node(n) puts yours[n] to s[n] and
    s[meabs] = yours   ! node(n+1) puts yours[n+1] to s[n+1].
    UNLOCK(lock[meabs])

    LOCK(lock[iyouabs]) ! node(n) gets s[n+1] into mine[n] and
55    mine = s[iyouabs] ! node(n+1) gets s[n] into mine[n+1].
    UNLOCK(lock[iyouabs])
    end

```

1 5.7 [C] Coarrays in XcalableMP C

2 This section describes the coarray features for XcalableMP C.

3 5.7.1 [C] Declaration of Coarrays

4 Synopsis

5 Coarrays are declared in XcalableMP C.

6 Syntax

7 [C] *data-type variable* : *codimensions*

8 where *codimensions* is:

9 *[[int-expr]...][*]*

Description

For XcalableMP C, coarrays are declared with a colon and square bracket, where *codimensions* specify the shape of a variable.

Note that the `coarray` directive for defining a coarray in the XcalableMP specification 1.0 (page 49) is obsolete.

Restrictions

- A coarray *variable* must have a global scope.

Examples

```

XcalableMP C
#pragma xmp nodes p[16]
float x:[*];

```

A variable *x* that has a global scope is declared as a coarray.

5.7.2 [C] Reference of Coarrays**Synopsis**

A coarray can be directly referenced or defined by any node. The target node is specified using an extended notation in XcalableMP C.

Syntax

[C] *variable* : [*int-expr*]...

Description

A sequence of [*int-expr*]’s preceded by a colon in XcalableMP C determines the image index for a coarray to be accessed. Note that the image index in XcalableMP C is 0-origin while the image index in XcalableMP Fortran is 1-origin.

A reference of coarrays can appear in the same place as an reference of normal variables in the base languages.

Examples

In the following codes, the second image ([C] image index 1/[F] image index 2) gets all values of array B on the first image ([C] image index 0/[F] image index 1) to array A on the second image.

```

XcalableMP C
int A[100]:[*], B[100]:[*];

if(xmpc_this_image() == 1){
    A[:] = B[:]:[0];
}

```

```

XcalableMP Fortran
integer :: A(100)[*], B(100)[*]

if (this_image() == 2) then
    A(:) = B(:)[1]
end if

```

5.7.3 [C] Synchronization of Coarrays**Synopsis**

XcalableMP C provides synchronization functions for coarrays.

1 Format

```

[C] void xmp_sync_all(int* status)
[C] void xmp_sync_memory(int* status)
2 [C] void xmp_sync_image(int image, int* status)
[C] void xmp_sync_images(int num, int* image_set, int* status)
[C] void xmp_sync_images_all(int* status)

```

3 Description

- 4 • `xmp_sync_all` is equivalent to the `sync all` statement in Fortran 2008.
- 5 • `xmp_sync_memory` is equivalent to the `sync memory` statement in Fortran 2008.
- 6 • A combination of `xmp_sync_image`, `xmp_sync_images`, and `xmp_sync_images_all` is equivalent to the `sync images` statement in Fortran 2008.
- 7
- 8 – `xmp_sync_image` is to synchronize one image.
- 9 – `xmp_sync_images` is to synchronize some images.
- 10 – `xmp_sync_images_all` is to synchronize all images.

11 Arguments

- 12 • The argument *status* is defined with one of the follow symbolic constants.
- 13 – `XMP_STAT_SUCCESS`
- 14 – `XMP_STAT_STOPPED_IMAGE`

15 If an execution of the function is successful, the *status* is defined using `XMP_STAT_SUCCESS`.
 16 The condition where the *status* is defined using `XMP_STAT_STOPPED_IMAGE` is the same
 17 as that where the *status* is defined using `STAT_STOPPED_IMAGE` in Fortran 2008. These
 18 symbolic constants are defined in “xmp.h.” If any other error condition occurs during the
 19 execution of these functions, the *status* is defined with a value that is different from the
 20 value of `XMP_STAT_SUCCESS` and `XMP_STAT_STOPPED_IMAGE`.

- 21 • In `xmp_sync_image`, the variable *image* determines a target image index.
- 22 • In `xmp_sync_images`, the variable *num* is a number of target images, and the variable
 23 *image_set* is an array in which the target image set is defined.

24 5.8 Directives for the Local-view Programming**25 5.8.1 [F] local_alias Directive****26 Synopsis**

27 In XcalableMP Fortran, the `local_alias` directive declares a local data object as an alias to
 28 the local section of a mapped array.

29 Syntax

```
30 [F] !$xmp local_alias local-array-name => global-array-name
```

Description

The LOCAL_ALIAS directive associates a non-mapped array with an explicitly mapped array. The non-mapped array is an associating local array and the explicitly mapped array is an associated global array. The shape of the associating local array is the same as that of the node-local portion of the associated global array including the shadow area. Each element of the associating local array shares the same storage unit in array-element order with that of the node-local portion of the associated global array. An associating local array and the corresponding global array always have the same allocation status. An associating local array always has the dynamic type and type parameter values of the corresponding associated global array.

An associating local array may be a coarray. An associating local array that is a coarray is an on-node coarray of the node array onto which the corresponding associated global array is mapped. All specifications and restrictions on coarrays are also applied to an associating local array that is a coarray, with the exception that an associating local array is always declared with *assumed-shape-spec-list* of the same rank as the associated global array. In particular, a processor shall ensure that an associating local array that is a coarray has the same bounds on all the images corresponding to the node array onto which the corresponding associated global array is mapped. The mapping attributes that are allowed for an associated global array are processor dependent.

Note that the base language Fortran is extended so that a deferred-shape array that is neither an allocatable array nor an array pointer is declared if it is specified as a *local-array-name* in the `local_alias` directive.

In XcalableMP C, the `address-of` operator is applied to global data substitutes for the `local_alias` directive (see 3.4).

Restrictions

- A *global-array-name* shall be the name of an explicitly mapped array declared in the same scoping unit.
- A *local-array-name* shall be the name of a non-mapped array declared in the same scoping unit.
- A *local-array-name* shall not be a dummy argument.
- An associating local array shall have the declared type and type parameters of the corresponding associated global array.
- An associating local array shall be declared with *assumed-shape-spec-list* of the same rank as the corresponding associated global array.
- A *local-array-name* shall appear in a `coarray` directive in the same scoping unit and the *node-name* in the `coarray` directive shall be the name of the node array onto which the associated global array is mapped.
- If an associated global array is a dummy argument and corresponds to an associating local array that is a coarray, the corresponding effective argument shall be an explicitly mapped array or a subobject of an explicitly mapped array whose name appears in a LOCAL_ALIAS directive, and the corresponding associating local array shall be a coarray.
- If a dummy argument is a coarray and the corresponding ultimate argument is a coarray appearing in a LOCAL_ALIAS directive, the dummy argument shall appear in a COAR-

1 RAY directive with a node array corresponding to a subset of the set of images that
2 corresponds to the node array onto which the ultimate argument is mapped.

3 Examples

4 Example 1

```

XcalableMP Fortran
!$xmp nodes n(4)
!$xmp template :: t(100)
!$xmp distribute (block) onto n :: t

5     real :: a(100)
!$xmp align (i) with t(i) :: a
!$xmp shadow (1) :: a

     real :: b(:)
10 !$xmp local_alias b => a

```

5 The array **a** is distributed by block onto four nodes. The node **n(2)** has its local section of
6 25 elements (**a(25:50)**) with shadow areas of size one on both the upper and lower bounds.
7 The local alias **b** is an array of 27 elements (**b(1:27)**) on **n(2)**. The table below shows
8 the correspondence of each element of **a** to that of **b** on **n(2)**.

a	b
lower shadow	1
26	2
27	3
28	4
...	...
50	26
upper shadow	27

10 Example 2

```

XcalableMP Fortran
!$xmp nodes n(4)
!$xmp template :: t(100)
!$xmp distribute (cyclic) onto n :: t

5     real :: a(100)
!$xmp align (i) with t(i) :: a

     real :: b(0:)
!$xmp local_alias b => a

```

11 An array **a** is distributed cyclically onto four nodes. Node **n(2)** has its local section of
12 25 elements (**a(2:100:4)**). The lower bound of local alias **b** is declared to be zero. As a
13 result, **b** is an array of size 25 whose lower bound is zero (**b(0:24)**) on **n(2)**. The table
14 below shows the correspondence of each element of **a** to that of **b** on **n(2)**.

a	b
2	0
6	1
10	2
...	...
98	24

1

Example 3

2

```

XcalableMP Fortran
!$xmp nodes n(4)
!$xmp template :: t(:)
!$xmp distribute (block) onto n :: t

5      real, allocatable :: a(:)
!$xmp align (i) with t(i) :: a

      real :: b(:)[*]
!$xmp local_alias b => a
10     ...

!$xmp template_fix :: t(128)

15     allocate (a(128))

      if (me < 4) b(4) = b(4)[me +1]

```

Because the global array `a` is an allocatable array, its local alias `b` is not defined when the subroutine starts execution. `b` is defined when `a` is allocated at the `allocate` statement. Note that `b` is declared as a coarray, and can therefore be accessed in the same manner as a normal coarray.

3

4

5

6

5.8.2 post Construct

7

Synopsis

8

The `post` construct, in combination with the `wait` construct, specifies the point-to-point synchronization.

9

10

Syntax

11

```

[F]  !$xmp post ( nodes-ref, tag )
[C]  #pragma xmp post ( nodes-ref, tag )

```

12

Description

13

This construct ensures that the execution of statements that precede it is completed before statements that follow the matching `wait` construct start are executed.

14

15

A `post` construct issued with the arguments of `nodes-ref` and `tag` on a node (called a *posting node*) dynamically matches at most one `wait` construct issued with the arguments of the posting node (unless omitted) and the same value as `tag` (unless omitted) by the node specified by `nodes-ref`.

16

17

18

19

1 **Restriction**

- 2 • *nodes-ref* must represent one node.
- 3 • *tag* must be an expression of type default integer, in XcalableMP Fortran, or type `int`, in
- 4 XcalableMP C.

5 **Example**6 **Example 1**

XcalableMP Fortran S1 !\$xmp post (p(2), 1)	XcalableMP Fortran !\$xmp wait (p(1), 1) S2
---	---

8 It is assumed that the code of the left is executed by the node `p(1)`, while that on the

9 right is executed by node `p(2)`.

10 The `post` construct on the left matches the `wait` construct on the right because their *nodes-*

11 *refs* represent each other and both *tags*'s have the same value of one. These constructs

12 ensure that no statement in `S2` is executed by `p(2)` until the execution of all statements

13 in `S1` is completed by `p(1)`.

14 **Example 2**

XcalableMP Fortran !\$xmp wait S3

15 It is assumed that this code is executed by node `p(2)`.

16 The `post` construct in the left code in Example 1 may match this `wait` construct because

17 both *nodes-ref* and *tag* are omitted in this `wait` construct.

18 **5.8.3 wait Construct**19 **Synopsis**

20 The `wait` construct, in combination with the `post` construct, specifies a point-to-point synchro-

21 nization.

22 **Syntax**

23 [F] !\$xmp wait [(*nodes-ref* [, *tag*])]

[C] #pragma xmp wait [(*nodes-ref* [, *tag*])]

24 **Description**

25 This construct prohibits statements that follow from being executed until the execution of all

26 statements preceding a matching `post` construct is completed on the node specified by *node-ref*.

27 A `wait` construct that is issued with the arguments of *nodes-ref* and *tag* on a node (called a

28 *waiting node*) dynamically matches a `post` construct issued with the arguments of the waiting

29 node and the same value as *tag* by the node specified by *nodes-ref*.

30 If *tag* is omitted, then the `wait` construct can match a `post` construct that is issued with

31 the arguments of the waiting node and any tag by the node specified by *nodes-ref*. If both *tag*

32 and *nodes-ref* are omitted, then the `wait` construct can match a `post` construct that is issued

33 with the arguments of the waiting node and any tag on any node.

Restriction

- *nodes-ref* must represent one node.
- *tag* must be an expression of type default integer, in XcalableMP Fortran, or type `int`, in XcalableMP C.

5.8.4 [C] lock/unlock Construct**Synopsis**

The `lock/unlock` constructs are equivalent to the `lock/unlock` statements in Fortran 2008.

Syntax

```
[C] #include <xmp.h>
[C] xmp_lock_t lock-object [, lock-object ]...
[C] #pragma xmp lock (lock-object) [ acquired_lock (success) ] [ stat (status) ]
[C] #pragma xmp unlock (lock-object) [ stat (status) ]
```

Please note the following points:

- The type `xmp_lock_t` is defined in “`xmp.h`”.
- The variable *lock-object* is a coarray.
- The variable *success* is an expression of type `int`.
- The variable *status* is an expression of type `int`.

Description

The `lock` construct, in combination with the `unlock` construct, is used to control a *lock-object*. The *lock-object* must be defined as a coarray to control it on a target node. The *lock-object* must be an expression of type `xmp_lock_t`, which is an opaque object defined in “`xmp.h`”.

If the `acquired_lock` clause is not used in the `lock` construct and the *lock-object* is locked, the node stops at the `lock` construct until the *lock-object* is unlocked by a different node. If the `acquired_lock` clause is used in the `lock` construct and the *lock-object* is locked by a different node, the node does not stop at the `lock` construct and the variable *success* is defined with the value `false`; `lock` construct leaves the *lock-object* unchanged. If the `acquired_lock` clause is used in the `lock` construct and the *lock-object* is unlocked, the variable *success* is defined with the value `true`.

The *status* is defined with one of the follow symbolic constants when executing the `lock/unlock` construct.

- `XMP_STAT_SUCCESS`
- `XMP_STAT_LOCKED`
- `XMP_STAT_UNLOCKED`
- `XMP_STAT_LOCKED_OTHER_IMAGE`

If the execution of the `lock/unlock` construct is successful, the *status* is defined with `XMP_STAT_SUCCESS`. The condition where the *status* is defined with `XMP_STAT_LOCKED`, `XMP_STAT_UNLOCKED`, or `XMP_STAT_LOCKED_OTHER_IMAGE` is the same as that where the *status* is defined with `STAT_LOCKED`, `STAT_UNLOCKED`, or `STAT_LOCKED_OTHER_IMAGE` in Fortran 2008. These symbolic constants are defined in “`xmp.h`”. If any other error condition occurs during the execution of these constructs, the *status* is defined with a value that is different from the value of `XMP_STAT_SUCCESS`, `XMP_STAT_LOCKED`, `XMP_STAT_UNLOCKED`, and `XMP_STAT_LOCKED_OTHER_IMAGE`.

1 **Example**

```
_____ XcalableMP C _____  
#include "xmp.h"  
  
xmp_lock_t lock_obj:[*];  
int A:[*], B;  
5 #pragma xmp nodes p[2]  
...  
#pragma xmp lock(lock_obj:[2])  
    if(xmp_node_num() == 1){  
        A:[2] = B;  
10    }  
#pragma xmp unlock(lock_obj:[2])
```


Chapter 6

Procedure Interfaces

This chapter describes the procedure interfaces, that is, how procedures are invoked and arguments are passed, in XcalableMP.

In order to achieve high composability of XcalableMP programs, it is one of the most important requirement that XcalableMP procedures can invoke procedures written in the base language with as few restrictions as possible.

6.1 General Rule

In XcalableMP, a procedure invocation is itself a local operation, and does not cause any communication or synchronization at runtime. Thus, a node can invoke any procedure, whether written in XcalableMP or in the base language, at any point during the execution. There is no restriction on the characteristics of procedures invoked by an XcalableMP procedure, except for a few ones on its argument, which are explained below.

Local data in the actual or dummy argument list (referred to as a *local actual argument* and a *local dummy argument*, respectively) are treated by the XcalableMP compiler in the same manner as the compiler of the base language. This rule makes it possible for a local actual argument in a procedure written in XcalableMP to be associated with a dummy argument of a procedure written in the base language.

If both an actual argument and its associated dummy argument are coarrays, they must be declared on the same node set.

Implementation The XcalableMP compiler does not transform either local actual or dummy arguments, so the backend compiler of the base language can treat them in its usual way.

The rest of this chapter specifies how global data appearing in an actual and/or dummy argument list (referred to as a *global actual argument* and a *global dummy argument*, respectively) are processed by the XcalableMP compiler.

6.2 Argument Passing Mechanism in XcalableMP Fortran

Either of the following global data can be put in the actual argument list:

- an array name;
- an array element; or
- an array section that satisfies both of the following conditions:

- its subscript list is a list of zero or more colons (“:”) followed by zero or more *int-expr*'s;
- the subscript of the dimension having a shadow is *int-expr* unless it is the last dimension.

There are two kinds of argument association for global data in XcalableMP Fortran: one is *sequence association*, which is for global dummies that are an explicit-shape or assumed-size array, and the other is *descriptor association*, which is for all other.

6.2.1 Sequence Association of Global Data

The concept of sequence association in Fortran is extended for global actual and dummy arguments in XcalableMP as follows.

If the actual argument is an array name or an array section that satisfies the above conditions, it represents an element sequence consisting of the elements of its local section in Fortran's array element order on each node. In addition, if the actual argument is an element of a global data object, it represents an element sequence consisting of the corresponding element in the local section and each element that follows it in array element order on each node.

An global actual argument that represents an element sequence and corresponds to a global dummy argument is sequence associated with the dummy argument if the dummy argument is an explicit-shape or assumed-size array. According to this (extended) rule of sequence association, each element of the element sequence represented by the global actual argument is associated with the element of the local section of the global dummy argument that has the same position in array element order.

Sequence association is the default rule of association for global actual arguments, and it is therefore applied unless it is obvious from the interface of the invoked procedure that the corresponding dummy argument is neither an explicit-shape nor assumed-size array.

Implementation In order to implement sequence association, the name, a section, or an element of global data appearing as an actual argument is treated by the XcalableMP compiler as the base address of its local section on each node, and the global data appearing as the corresponding dummy argument is initialized at runtime so that it is composed of the local sections, each of which starts from the address received as the argument. On a node that does not have the local section corresponding to the actual argument, an unspecified value (e.g. null) is received.

Such an implementation implies that in many cases, in order to associate properly a global actual argument with the global dummy argument, their mappings (including their shadow attributes) must be identical.

Examples

Example 1 Both the actual argument **a** and the dummy argument **x** are global explicit-shape arrays, and therefore, **a** is sequence associated with **x**.

The base address of the local section of **a** is passed between these subroutines on each node. Each of the local sections of **x** starts from the received address (Figure 6.1).

XcalableMP Fortran

```

subroutine xmp_sub1
!$xmp nodes p(4)
!$xmp template t(100)
!$xmp distribute t(block) onto p

```



```

5      real a(100)
!$xmp align a(i) with t(i)
!$xmp shadow a(1:1)
      call xmp_sub2(a)
      end subroutine

10     subroutine xmp_sub2(x)
!$xmp nodes p(4)
!$xmp template t(100)
!$xmp distribute t(block) onto p
15     real x(100)
!$xmp align x(i) with t(i)
!$xmp shadow x(1:1)
      ...

```




Figure 6.1: Sequence association with a global dummy argument.

1 **Example 2** The actual argument `a` is a global explicit-shape array, and the dummy argument
2 `x` is a local explicit-shape. Sequence association is also applied in this case.

3 The caller subroutine `xmp_sub1` passes the base address of the local section of `a` on each
4 node, and the callee `f_sub2` receives it and initializes `x` with the storage starting from it
5 (Figure 6.2).

```

----- XcalableMP Fortran -----
      subroutine xmp_sub1
!$xmp nodes p(4)
!$xmp template t(100)
!$xmp distribute t(block) onto p
5      real a(100)
!$xmp align a(i) with t(i)
!$xmp shadow a(1:1)
      n = 1 + 100/4 + 1
      call f_sub2(a,n)
10     end subroutine

```

```

----- Fortran -----
      subroutine f_sub2(x,n)
      real x(n)
      ...

```




Figure 6.2: Sequence association with a local dummy argument.

Example 3 The actual argument $a(:,1)$ is a contiguous section of a global data object, and the dummy argument x is a local explicit-shape array. Sequence association is applied in this case, but only the node $p(1)$ owns the section. Hence, f_sub2 is invoked only by $p(1)$ (Figure 6.3).

1
2
3
4

```

XcalableMP Fortran
subroutine xmp_sub1
!$xmp nodes p(4)
!$xmp template t(100,100)
!$xmp distribute t(*,block) onto p
5  real a(100,100)
!$xmp align a(i,j) with t(i,j)
!$xmp shadow a(0,1:1)
n = 100
!$xmp task on p(1)
10 call f_sub2(a(:,1),n)
!$xmp end task
end subroutine

Fortran
subroutine f_sub2(x,n)
real x(n)
...
```




Figure 6.3: Sequence association of a section of a global data object as an actual argument with a local dummy argument.

- 1 **Example 4** The actual argument $a(1)$ is an element of the global data, and the dummy argument x is a local explicit-shape array. Sequence association is applied in this case, but
 2 only the node $p(1)$ owns the element. Hence, f_sub2 is invoked only by $p(1)$ (Figure 6.4).
 3

XcalableMP Fortran	
5	<pre> subroutine xmp_sub1 !\$xmp nodes p(4) !\$xmp template t(100) !\$xmp distribute t(block) onto p real a(100) !\$xmp align a(i) with t(i) !\$xmp shadow a(1:1) n = 100/4 !\$xmp task on p(1) call f_sub2(a(1),n) !\$xmp end task end subroutine </pre>
10	<pre> subroutine f_sub2(x,n) real x(n) ...</pre>

- 4 **Example 5** Even if either the global actual or dummy argument has a full shadow, the rule of
 5 sequence association is the same in principle. Hence, the base address of the local section
 6 of a is passed between these subroutines on each node, and each local section of x starts
 7 from the received address (Figure 6.5).

8 6.2.2 Descriptor Association of Global Data

- 9 When the actual argument is a global data object, and it is obvious from the interface of
 10 the invoked procedure that the corresponding dummy argument is neither an explicit-shape
 11 nor assumed-size array, the actual argument is *descriptor associated* with the dummy argument.
 12 According to the descriptor association rule, the dummy argument inherits its shape and storage
 13 from the actual argument.




Figure 6.4: Sequence association of an element of a global data object as an actual argument with a local dummy argument.




Figure 6.5: Sequence association with a global dummy argument that has a full shadow.

Implementation In order to implement the descriptor association, a global actual argument 1
is treated by the XcalableMP compiler: 2

- as if it were the *global-data descriptor* of the actual array, which is an internal data structure 3
managed by the XcalableMP runtime system to store information on a global data object 4
(see 7.1.1) if the dummy is a global data object; or 5
- as it is an array representing the local section of the actual array, which is to be processed 6
by the backend Fortran compiler in the same manner as usual data if the dummy is a local 7
data object. 8

For the first case, a global dummy is initialized at runtime with a copy of the global-data 9
descriptor received. 10

When an actual argument is descriptor associated with the dummy argument and their 11
mappings are not identical, the XcalableMP runtime system may detect and report the error. 12

1 **Examples**

2 **Example 1** There is an explicit interface of the subroutine `xmp_sub2` specified by an interface
 3 block in the subroutine `xmp_sub1`, from which it is found that the dummy argument `x`
 4 is a global assumed-shape array. Therefore, the global actual argument `a` is descriptor
 5 associated with the global dummy argument `x`.

6 It is the global-data descriptor of `a` that is passed between these subroutines. The dummy
 7 argument `x` is initialized by the XcalableMP runtime system on the basis of the information
 8 extracted from the descriptor received (Figure 6.6).

```

                                XcalableMP Fortran
subroutine xmp_sub1

!$xmp nodes p(4)
!$xmp template t(100)
5 !$xmp distribute t(block) onto p
  real a(100)
!$xmp align a(i) with t(i)
!$xmp shadow a(1:1)

10  interface
    subroutine xmp_sub2(x)
!$xmp nodes p(4)
!$xmp template t(100)
!$xmp distribute t(block) onto p
15  real x(:)
!$xmp align x(i) with t(i)
!$xmp shadow a(1:1)
    end subroutine xmp_sub2
    end interface

20  call xmp_sub2(a)

    end subroutine

25  subroutine xmp_sub2(x)
!$xmp nodes p(4)
!$xmp template t(100)
!$xmp distribute t(block) onto p
    real x(:)
30 !$xmp align x(i) with t(i)
!$xmp shadow a(1:1)
    ...

```

9 **Example 2** There is the explicit interface of the subroutine `f_sub2`, which is written in Fortran,
 10 specified by an interface block in the subroutine `xmp_sub1`, and the dummy argument `x` is
 11 a local (i.e., non-mapped) assumed-shape array. Therefore, the global actual argument `a`
 12 is descriptor associated with the local dummy argument `x`.

13 The global actual argument is replaced with its local section by the XcalableMP compiler,
 14 and the association of the local section with the dummy argument is to be processed by
 15 the backend Fortran compiler in the same manner as usual data (Figure 6.7).

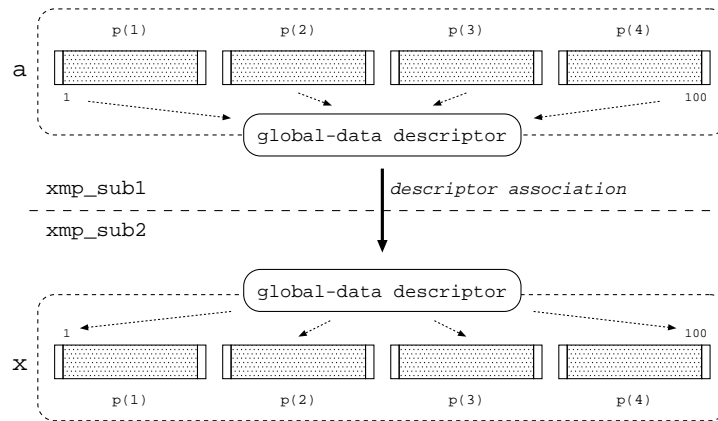


Figure 6.6: Descriptor association with a global dummy argument.

```

XcalableMP Fortran
subroutine xmp_sub1
!$xmp nodes p(4)
!$xmp template t(100)
5 !$xmp distribute t(block) onto p
  real a(100)
!$xmp align a(i) with t(i)
!$xmp shadow a(1:1)

10  interface
    subroutine f_sub2(x)
    real x(:)
    end subroutine f_sub2
  end interface

15  call f_sub2(a)

  end subroutine

Fortran
subroutine f_sub2(x)
  real x(:)
  ...

```

6.3 Argument-Passing Mechanism in XcalableMP C

When an actual argument is a global data object, it is passed by the address of its local section. When a dummy argument is a global data object, an address is received and used as the base address of each of its local sections.

Implementation The name of a global data object appearing as an actual argument is treated by the XcalableMP compiler as the pointer to the first element of its local section on each




Figure 6.7: Descriptor association with a local dummy argument.

1 node. On a node onto which no part of the global data object is mapped, the pointer is set
 2 to an unspecified value (e.g., null). Note that an element of a global data object in the actual
 3 argument list is treated in the same manner as those in other usual statements because an array
 4 element is passed by value, as in C.

5 The name of a global data object appearing as a dummy argument is treated by the Xcal-
 6 ableMP compiler as the pointer to the first element of its local section on each node. As a result,
 7 it is initialized at runtime so that it is composed of the local sections on the executing nodes.

8 Such an implementation implies that in many cases, in order to pass properly a global actual
 9 argument to the corresponding global dummy argument, their mappings (including their shadow
 10 attributes) must be identical.

11 Examples

12 **Example 1** The global actual argument `a` is treated by the XcalableMP compiler as the pointer
 13 to the first element of its local section, which is passed to the callee, on each node.

14 The global dummy argument `x` is initialized so that each of its local sections starts from
 15 the address held by the received pointer (Figure 6.8).

```

XcalableMP C
void xmp_func1()
{
  #pragma xmp nodes p[4]
  #pragma xmp template t[100]
5  #pragma xmp distribute t[block] onto p
  float a[100];
  #pragma xmp align a[i] with t[i]
  #pragma xmp shadow a[1:1]

10  xmp_func2(a);
}

void xmp_func2(float x[100])
{
15 #pragma xmp nodes p[4]
  #pragma xmp template t[100]
  #pragma xmp distribute t[block] onto p

```

```

#pragma xmp align x[i] with t[i]
#pragma xmp shadow a[1:1]
20  ...

```

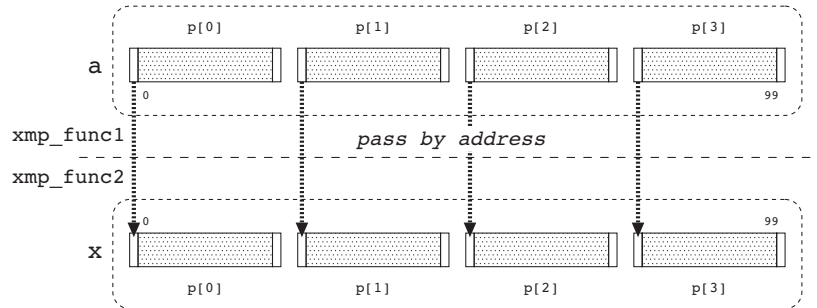


Figure 6.8: Passing to a global dummy argument.

Example 2 The global actual argument `a` is treated by the XcalableMP compiler as the pointer to the first element of its local section, which is passed to the callee, on each node. 1

The local dummy argument `x` on each node starts from the address held by the received pointer (Figure 6.9). 2

The local dummy argument `x` on each node starts from the address held by the received pointer (Figure 6.9). 3

```

----- XcalableMP C -----
void xmp_func1()
{
  #pragma xmp nodes p[4]
  #pragma xmp template t[100]
5 #pragma xmp distribute t[block] onto p
  float a[100];
  #pragma xmp align a[i] with t[i]
  #pragma xmp shadow a[1:1]
10  c_func2(a);
}

```

```

----- C -----
void c_func2(float x[27])
{
  ...
}

```

Example 3 The actual argument `a[0]` is an element of a global data object, and the dummy argument `x` is a scalar, where the normal argument-passing rule of C for variables of a basic type (i.e., “pass-by-value”) is applied. However, only the node `p[0]` owns the element. Hence, `c_func2` is invoked only by `p[0]` (Figure 6.10). 5

```

----- XcalableMP C -----
void xmp_func1()
{
  #pragma xmp nodes p[4]
  #pragma xmp template t[100]
5 #pragma xmp distribute t[block] onto p

```

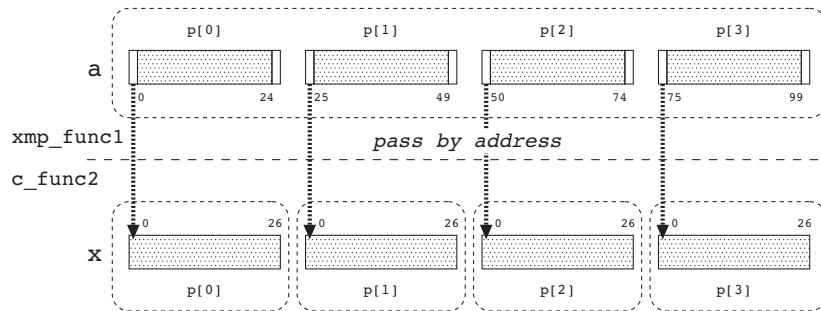



Figure 6.9: Passing to a local dummy argument.

```

float a[100];
#pragma xmp align a[i] with t[i]
#pragma xmp shadow a[1:1]
10 #pragma xmp task on p[0]
    c_func2(a[0]);
}

```

C

```

void c_func2(float x)
{
    ...
}

```




Figure 6.10: Passing an element of a global data object as an actual argument to a local dummy argument.

Chapter 7

Intrinsic and Library Procedures

This specification defines various procedures that perform a system inquiry, synchronization, computation, etc. The procedures are provided as intrinsic procedures in XcalableMP Fortran, and as library procedures in XcalableMP C.

7.1 Intrinsic Functions

7.1.1 `xmp_desc_of`

Format

```
[F] type(xmp_desc) xmp_desc_of(xmp_entity)
```

Note that `xmp_desc_of` is an intrinsic function in XcalableMP Fortran or a built-in operator in XcalableMP C. For the `xmp_desc_of` operator, refer to section 3.6.

Synopsis

`xmp_desc_of` returns a descriptor to retrieve information of the specified global array, template, or node array. The resulting descriptor can be used as an input argument of mapping inquiry functions.

The type of descriptors, `type(xmp_desc)`, in XcalableMP Fortran, and `xmp_desc_t`, in XcalableMP C, is implementation-defined, and it is defined in a Fortran module named `xmp_lib` or a Fortran include file named `xmp_lib.h`.

Arguments

The argument or operand `xmp_entity` is the name of either a global array, a template, or a node array.

7.2 System Inquiry Functions

- `xmp_all_node_num`
- [C] `xmpc_all_node_num`
- `xmp_all_num_nodes`
- `xmp_node_num`
- [C] `xmpc_node_num`

- [C] `xmpc_this_image` 1
- `xmp_num_nodes` 2
- `xmp_num_images` 3
- `xmp_wtime` 4
- `xmp_wtick` 5

7.2.1 `xmp_all_node_num` 6

Format 7

[F] integer function `xmp_all_node_num()` 8
 [C] int `xmp_all_node_num(void)`

Synopsis 9

The `xmp_all_node_num` routine returns the node number, within the entire node set, of the node that calls `xmp_all_node_num`. 10
 11

Arguments 12

none. 13

7.2.2 [C] `xmpc_all_node_num` 14

Format 15

[C] int `xmpc_all_node_num(void)` 16

Synopsis 17

The `xmpc_all_node_num` routine returns the node number -1 , within the entire node set, of the node that calls `xmpc_all_node_num`. 18
 19

Arguments 20

none. 21

7.2.3 `xmp_all_num_nodes` 22

Format 23

[F] integer function `xmp_all_num_nodes()` 24
 [C] int `xmp_all_num_nodes(void)`

Synopsis 25

The `xmp_all_num_nodes` routine returns the number of nodes in the entire node set. 26

Arguments 27

none. 28

1 7.2.4 xmp_node_num

2 Format

3 [F] integer function xmp_node_num()
3 [C] int xmp_node_num(void)

4 Synopsis

5 The `xmp_node_num` routine returns the node number, within the current executing node set, of
6 the node that calls `xmp_node_num`.

7 Arguments

8 none.

9 7.2.5 [C] xmpc_node_num

10 Format

11 [C] int xmpc_node_num(void)

12 Synopsis

13 The `xmpc_node_num` routine returns the node number -1 , within the current executing node set,
14 of the node that calls `xmpc_node_num`.

15 Arguments

16 none.

17 7.2.6 [C] xmpc_this_image

18 Format

19 [C] int xmpc_this_image(void)

20 Synopsis

21 The `xmpc_this_image` routine is identical to the `xmpc_node_num` routine.

22 Arguments

23 none.

24 7.2.7 xmp_num_nodes

25 Format

26 [F] integer function xmp_num_nodes()
26 [C] int xmp_num_nodes(void)

27 Synopsis

28 The `xmp_num_nodes` routine returns the number of the executing nodes.

Arguments	1
none.	2
7.2.8 xmp_num_images	3
Format	4
[F] integer function xmp_num_images()	5
[C] int xmp_num_images(void)	5
Synopsis	6
The xmp_num_images routine is identical to the xmp_num_nodes routine.	7
Arguments	8
none.	9
7.2.9 xmp_wtime	10
Format	11
[F] double precision function xmp_wtime()	12
[C] double xmp_wtime(void)	12
Synopsis	13
The xmp_wtime routine returns elapsed wall-clock time in seconds since some time in the past.	14
The “time in the past” is guaranteed not to change during the life of the process. There is no	15
requirement that different nodes return “the same time.”	16
Arguments	17
none.	18
7.2.10 xmp_wtick	19
Format	20
[F] double precision function xmp_wtick()	21
[C] double xmp_wtick(void)	21
Synopsis	22
The xmp_wtick routine returns the resolution of the timer used by xmp_wtime. It returns a	23
double-precision value that is equal to the number of seconds between successive clock ticks.	24
Arguments	25
none.	26

1 7.3 [C] Execution Control Functions

2 7.3.1 xmp_exit

3 Format

4 [C] void xmp_exit(int status)

5 Synopsis

6 xmp_exit terminates an XcalableMP program normally. The value of the argument `status`
7 returned to the host environment is the same as that by the `exit` standard library function of
8 the base language.

9 xmp_exit must be collectively invoked by every node in the entire node set; otherwise, the
10 behavior is undefined.

11 Arguments

12 The argument `status` is a status code to be returned to the host environment.

13 7.4 Synchronization Functions

14 7.4.1 xmp_test_async

[F] logical function xmp_test_async(async_id)
integer async_id

[C] int xmp_test_async(int async_id)

16 Synopsis

17 The xmp_test_async routine returns `.true.` in XcalableMP Fortran, or 1 in XcalableMP C, if
18 an asynchronous communication specified by the argument `async_id` is complete; otherwise, it
19 returns `.false.` or 0.

20 Arguments

21 The argument `async_id` is an integer expression that specifies an asynchronous communication
22 initiated by a global communication construct with the `async` clause.

23 7.5 Memory Allocation Functions

24 7.5.1 [C] xmp_malloc

25 void* xmp_malloc(xmp_desc_t d, size_t size0, size_t size1, ...)

26 Synopsis

27 The xmp_malloc routine allocates storage for the local section of a global array of size `size0`×`size1`×...
28 that is associated with the descriptor specified by `d`, and returns the pointer to it on each node.
29 For an example of xmp_malloc, refer to section 3.5.

Arguments

- `d` is the descriptor associated with the pointer to a global array to be allocated.
- `size0`, `size1`, ... are the sizes of the dimensions of the global array to be allocated.

7.6 Mapping Inquiry Functions

All mapping inquiry functions are specified as integer functions. These functions return zero upon success and an implementation-defined negative integer value upon failure.

7.6.1 xmp_nodes_ndims**Format**

```
[F] integer function xmp_nodes_ndims(d, ndims)
      type(xmp_desc) d
      integer        ndims
[C]  int            xmp_nodes_ndims(xmp_desc_t d, int *ndims)
```

Synopsis

The `xmp_nodes_ndims` function provides the rank of the target node array.

Input Arguments

- `d` is a descriptor of a node array.

Output Arguments

- `ndims` is the rank of the node array specified by `d`.

7.6.2 xmp_nodes_index**Format**

```
[F] integer function xmp_nodes_index(d, dim, index)
      type(xmp_desc) d
      integer        dim
      integer        index
[C]  int            xmp_nodes_index(xmp_desc_t d, int dim, int *index)
```

Synopsis

The `xmp_nodes_index` function provides the indices of the executing node in the target node array.

Input Arguments

- `d` is a descriptor of a node array.
- `dim` is the target dimension of the node array.

Output Arguments

- `index` is an index of the target dimension of the node array specified by `d`.

1 **7.6.3 xmp_nodes_size**2 **Format**

```

[F] integer function xmp_nodes_size(d, dim, size)
      type(xmp_desc)  d
3      integer        dim
      integer        size
[C] int              xmp_nodes_size(xmp_desc_t d, int dim, int *size)

```

4 **Synopsis**

5 The `xmp_nodes_size` function provides the size of each dimension of the target node array.

6 **Input Arguments**

- 7 • `d` is a descriptor of a node array.
- 8 • `dim` is the target dimension of the node array.

9 **Output Arguments**

- 10 • `size` is the extent of the target dimension of the node array specified by
- 11 `t d`.

12 **7.6.4 xmp_nodes_attr**13 **Format**

```

[F] integer function xmp_nodes_attr(d, attr)
      type(xmp_desc)  d
14      integer        attr
[C] int              xmp_nodes_attr(xmp_desc_t d, int *attr)

```

15 **Synopsis**

16 The `xmp_nodes_attr` function provides the attribute of the target node array. The output value

17 of the argument `attr` is one of:

```

      XMP_ENTIRE_NODES      (Entire nodes)
18      XMP_EXECUTING_NODES (Executing nodes)
      XMP_EQUIVALENCE_NODES (Equivalence nodes)

```

19 These are named constants that are defined in module `xmp_lib` and in the include file

20 `xmp_lib.h` in XcalableMP Fortran, and symbolic constants that are defined in the header file

21 `xmp.h` in XcalableMP C.

22 **Input Arguments**

- 23 • `d` is a descriptor of a node array.

24 **Output Arguments**

- 25 • `attr` is an attribute of the target node array specified by `d`.

7.6.5 xmp_nodes_equiv 1

Format 2

```
[F] integer function xmp_nodes_equiv(d, dn, lb, ub, st)
      type(xmp_desc) d
      type(xmp_desc) dn
      integer lb(*)
      integer ub(*)
      integer st(*)
[C] int xmp_nodes_equiv(xmp_desc_t d, xmp_desc_t *dn,
                        int lb[], int ub[], int st[])

```

Synopsis 4

The `xmp_nodes_equiv` function provides the descriptor of a node array as well as a subscript list that represents a node set that is assigned to the target node array in the `nodes` directive. This function returns with a failure when the target node array is not declared as equivalenced.

Input Arguments 8

- `d` is a descriptor of a node array.

Output Arguments 10

- `dn` is the descriptor of the referenced node array if the target node array is declared as equivalenced; otherwise, `dn` is set to undefined.
- `lb` is a one-dimensional integer array the extent of which must be more than or equal to the rank of the referenced node array. The *i*-th element of `lb` is set to the lower bound of the *i*-th subscript of the node reference unless it is “*”, or to undefined otherwise.
- `ub` is a one-dimensional integer array the extent of which must be more than or equal to the rank of the referenced node array. The *i*-th element of `ub` is set to the upper bound of the *i*-th subscript of the node reference unless it is “*”, or to undefined otherwise.
- `st` is a one-dimensional integer array the extent of which must be more than or equal to the rank of the referenced node array. The *i*-th element of `st` is set to the stride of the *i*-th subscript of the node reference unless it is “*”, or to zero otherwise.

7.6.6 xmp_template_fixed 22

Format 23

```
[F] integer function xmp_template_fixed(d, fixed)
      type(xmp_desc) d
      logical fixed
[C] int xmp_template_fixed(xmp_desc_t d, int *fixed)

```

Synopsis 25

The `xmp_template_fixed` function provides the logical value that shows whether the template is fixed or not.

1 **Input Arguments**

- 2
- `d` is a descriptor of a template.

3 **Output Arguments**

- 4
- `fixed` is set to true in XcalableMP Fortran and an implementation-defined non-zero integer value in XcalableMP C if the template specified by `d` is fixed; otherwise, it is set to false in XcalableMP Fortran and zero in XcalableMP C.

7 **7.6.7 xmp_template_ndims**8 **Format**

[F]	integer function	<code>xmp_template_ndims(d, ndims)</code>
	<code>type(xmp_desc)</code>	<code>d</code>
9	integer	<code>ndims</code>
[C]	int	<code>xmp_template_ndims(xmp_desc_t d, int *ndims)</code>

10 **Synopsis**11 The `xmp_template_ndims` function provides the rank of the target template.12 **Input Arguments**

- 13
- `d` is a descriptor of a template.

14 **Output Arguments**

- 15
- `ndims` is the rank of the template specified by `d`.

16 **7.6.8 xmp_template_lbound**17 **Format**

[F]	integer function	<code>xmp_template_lbound(d, dim, lbound)</code>
	<code>type(xmp_desc)</code>	<code>d</code>
18	integer	<code>dim</code>
	integer	<code>lbound</code>
[C]	int	<code>xmp_template_lbound(xmp_desc_t d, int dim, int *lbound)</code>

19 **Synopsis**20 The `xmp_template_lbound` function provides the lower bound of each dimension of the template.
21 This function returns with a failure when the lower bound is not fixed.22 **Input Arguments**

- 23
- `d` is a descriptor of a template.
 - `dim` is the target dimension of the template.

25 **Output Arguments**

- 26
- `lbound` is the lower bound of the target dimension of the template specified by `d`. When the lower bound is not fixed, it is set to undefined.

7.6.9 xmp_template_ubound

1

Format

2

```
[F] integer function xmp_template_ubound(d, dim, ubound)
      type(xmp_desc) d
      integer        dim
      integer        ubound
```

3

```
[C] int xmp_template_ubound(xmp_desc_t d, int dim, int *ubound)
```

Synopsis

4

The `xmp_template_ubound` function provides the upper bound of each dimension of the template. This function returns with a failure when the upper bound is not fixed.

5

6

Input Arguments

7

- `d` is a descriptor of a template.
- `dim` is the target dimension of the template.

8

9

Output Arguments

10

- `ubound` is an upper bound of the target dimension of the template specified by `d`. When the upper bound is not fixed, it is set to undefined.

11

12

7.6.10 xmp_dist_format

13

Format

14

```
[F] integer function xmp_dist_format(d, dim, format)
      type(xmp_desc) d
      integer        dim
      integer        format
```

15

```
[C] int xmp_dist_format(xmp_desc_t d, int dim, int *format)
```

Synopsis

16

The `xmp_dist_format` function provides the distribution format of a dimension of a template. The output value of the argument `format` is one of:

17

18

```
XMP_NOT_DISTRIBUTED (not distributed)
XMP_BLOCK            (block distribution)
XMP_CYCLIC           (cyclic distribution)
XMP_GBLOCK           (gblock distribution)
```

19

These symbolic constants are defined in “`xmp.h`”.

20

Input Arguments

21

- `d` is a descriptor of a template.
- `dim` is the target dimension of the template.

22

23

Output Arguments

24

- `format` is a distribution format of the target dimension of the template specified by `d`.

25

1 **7.6.11 xmp_dist_blocksize**2 **Format**

	[F]	integer function	xmp_dist_blocksize(d, dim, blocksize)
		type(xmp_desc)	d
3		integer	dim
		integer	blocksize
	[C]	int	xmp_dist_blocksize(xmp_desc_t d, int dim, int *blocksize)

4 **Synopsis**

5 The `xmp_dist_blocksize` function provides the block width of a dimension of a template.

6 **Input Arguments**

- 7 • `d` is a descriptor of a template.
- 8 • `dim` is the target dimension of the template.

9 **Output Arguments**

- 10 • `blocksize` is the block width of the target dimension of the template specified by `d`.

11 **7.6.12 xmp_dist_gblockmap**12 **Format**

	[F]	integer function	xmp_dist_gblockmap(d, dim, map)
		type(xmp_desc)	d
13		integer	dim
		integer	map(N)
	[C]	int	xmp_dist_gblockmap(xmp_desc_t d, int dim, int map[])

14 **Synopsis**

15 The `xmp_dist_gblockmap` function provides the mapping array of the `gblock` distribution.

16 When the `dim`-th dimension of the global array is distributed by `gblock` and its mapping
 17 array is fixed, this function returns zero; otherwise, it returns an implementation-defined negative
 18 integer value.

19 **Input Arguments**

- 20 • `d` is a descriptor of a template.
- 21 • `dim` is the target dimension of the template.

22 **Output Arguments**

- 23 • `map` is a one-dimensional integer array the extent of which is more than the size of the
 24 corresponding dimension of the node array onto which the template is distributed.

25 The `i`-th element of `map` is set to the value of the `i`-th element of the target mapping array.

7.6.13 xmp_dist_nodes 1

Format 2

[F]	integer function	xmp_dist_nodes(d, dn)	
	type(xmp_desc)	d	
	type(xmp_desc)	dn	3
[C]	int	xmp_dist_nodes(xmp_desc_t d, xmp_desc_t *dn)	

Synopsis 4

The `xmp_dist_nodes` function provides the descriptor of the node array onto which a template is distributed. 5
6

Input Arguments 7

- `d` is a descriptor of a template. 8

Output Arguments 9

- `dn` is the descriptor of the node array. 10

7.6.14 xmp_dist_axis 11

Format 12

[F]	integer function	xmp_dist_axis(d, dim, axis)	
	type(xmp_desc)	d	
	integer	dim	13
	integer	axis	
[C]	int	xmp_dist_axis(xmp_desc_t d, int dim, int *axis)	

Synopsis 14

The `xmp_dist_axis` function provides the dimension of the node array onto which a dimension of a template is distributed. This function returns with a failure when the dimension of the template is not distributed. 15
16
17

Input Arguments 18

- `d` is a descriptor of a template. 19
- `dim` is the target dimension of the template. 20

Output Arguments 21

- `axis` is a dimension of the node array onto which the target dimension of the template specified by `d` is distributed. When the dimension of the template is not distributed, it is set to undefined. 22
23
24

1 7.6.15 xmp_align_axis

2 Format

```
[F] integer function xmp_align_axis(d, dim, axis)
      type(xmp_desc) d
3      integer      dim
      integer      axis
[C] int            xmp_align_axis(xmp_desc_t d, int dim, int *axis)
```

4 Synopsis

5 The `xmp_align_axis` function provides the dimension of the template with which a dimension of
6 a global array is aligned. This function returns with a failure when the dimension of the global
7 array is not aligned.

8 Input Arguments

- 9 • `d` is a descriptor of a global array.
- 10 • `dim` is the target dimension of the global array.

11 Output Arguments

- 12 • `axis` is the dimension of the template with which the target dimension of the global array
13 specified by `d` is aligned. When the dimension of the global array is not aligned, or is
14 collapsed, it is set to undefined.

15 7.6.16 xmp_align_offset

16 Format

```
[F] integer function xmp_align_offset(d, dim, offset)
      type(xmp_desc) d
17      integer      dim
      integer      offset
[C] int            xmp_align_offset(xmp_desc_t d, int dim, int *offset)
```

18 Synopsis

19 The `xmp_align_offset` function provides the align offset for a dimension of a global array. This
20 function returns with a failure when there is no offset.

21 Input Arguments

- 22 • `d` is a descriptor of a global array.
- 23 • `dim` is the target dimension of the global array.

24 Output Arguments

- 25 • `offset` is the align offset for the target dimension of the global array specified by `d`. When
26 there is no offset, it is set to undefined.

7.6.17 xmp_align_replicated 1

Format 2

[F]	integer function	xmp_align_replicated(d, dim, replicated)	
	type(xmp_desc)	d	
	integer	dim	3
	logical	replicated	
[C]	int	xmp_align_replicated(xmp_desc_t d, int dim, int *replicated)	

Synopsis 4

The `xmp_align_replicated` function provides the logical value that shows whether or not the dimension of the template with which a global array is aligned is replicated. 5
6

Input Arguments 7

- `d` is a descriptor of a global array. 8
- `dim` is the target dimension of the template with which the global array is aligned. 9

Output Arguments 10

- `replicated` is a logical scalar, which is set to true if the dimension of the template is replicated. 11
12

7.6.18 xmp_align_template 13

Format 14

[F]	integer function	xmp_align_template(d, dt)	
	type(xmp_desc)	d	
	type(xmp_desc)	dt	15
[C]	int	xmp_align_template(xmp_desc_t d, xmp_desc_t *dn)	

Synopsis 16

The `xmp_align_template` function provides the descriptor of the template with which a global array is aligned. 17
18

Input Arguments 19

- `d` is a descriptor of a global array. 20

Output Arguments 21

- `dt` is the descriptor of the template. 22

7.6.19 xmp_array_ndims 23

Format 24

[F]	integer function	xmp_array_ndims(d, ndims)	
	type(xmp_desc)	d	
	integer	ndims	25
[C]	int	xmp_array_ndims(xmp_desc_t d, int *ndims)	

1 Synopsis

2 The `xmp_array_ndims` function provides the rank of a global array.

3 Input Arguments

- 4 • `d` is a descriptor of a global array.

5 Output Arguments

- 6 • `ndims` is the rank of the global array specified by `d`.

7 7.6.20 xmp_array_lshadow**8 Format**

[F]	integer function	<code>xmp_array_lshadow(d, dim, lshadow)</code>
	<code>type(xmp_desc)</code>	<code>d</code>
9	integer	<code>dim</code>
	integer	<code>lshadow</code>
[C]	int	<code>xmp_array_lshadow(xmp_desc_t d, int dim, int *lshadow)</code>

10 Synopsis

11 The `xmp_array_lshadow` function provides the size of the lower shadow of a dimension of a global
12 array.

13 Input Arguments

- 14 • `d` is a descriptor of a global array.
- 15 • `dim` is the target dimension of the global array.

16 Output Arguments

- 17 • `lshadow` is the size of the lower shadow of the target dimension of the global array specified
18 by `d`.

19 7.6.21 xmp_array_ushadow**20 Format**

[F]	integer function	<code>xmp_array_ushadow(d, dim, ushadow)</code>
	<code>type(xmp_desc)</code>	<code>d</code>
21	integer	<code>dim</code>
	integer	<code>ushadow</code>
[C]	int	<code>xmp_array_ushadow(xmp_desc_t d, int dim, int *ushadow)</code>

22 Synopsis

23 The `xmp_array_ushadow` function provides the size of the upper shadow of a dimension of a
24 global array.

Input Arguments

- `d` is a descriptor of a global array.
- `dim` is the target dimension of the global array.

Output Arguments

- `ushadow` is the size of the upper shadow of the target dimension of the global array specified by `d`.

7.6.22 xmp_array_lbound**Format**

```
[F] integer function xmp_array_lbound(d, dim, lbound)
      type(xmp_desc) d
      integer        dim
      integer        lbound
[C] int             xmp_array_lbound(xmp_desc_t d, int dim, int *lbound)
```

Synopsis

The `xmp_array_lbound` function provides the lower bound of a dimension of a global array. This function returns with a failure when the lower bound is not fixed.

Input Arguments

- `d` is a descriptor of a global array.
- `dim` is the target dimension of the global array.

Output Arguments

- `lbound` is the lower bound of the target dimension of the global array specified by `d`. When the lower bound is not fixed, it is set to undefined.

7.6.23 xmp_array_ubound**Format**

```
[F] integer function xmp_array_ubound(d, dim, ubound)
      type(xmp_desc) d
      integer        dim
      integer        ubound
[C] int             xmp_array_ubound(xmp_desc_t d, int dim, int *ubound)
```

Synopsis

The `xmp_array_ubound` function provides the upper bound of a dimension of a global array. This function returns with a failure when the upper bound is not fixed.

Input Arguments

- `d` is a descriptor of a global array.
- `dim` is the target dimension of the global array.

1 **Output Arguments**

- 2 • `ubound` is the upper bound of the target dimension of the global array specified by `d`.
3 When the upper bound is not fixed, it is set to undefined.

4 **7.7 [F] Array Intrinsic Functions of the Base Language**

5 The array intrinsic functions of the base language Fortran are classified into three classes: *in-*
6 *quiry*, *elemental*, and *transformational*.

7 This section specifies how these functions work in the XMP/F programs when a global array
8 appears as an argument.

- 9 • **Inquiry functions**

10 The inquiry functions with a global array or its subobject being an argument are regarded
11 as inquiries about the global array, and return its “global” properties as if it were not
12 distributed.

- 13 • **Elemental functions**

14 The result of the elemental functions with a global array or its subobject being an argument
15 has the same shape and mapping as the argument. Note that such a reference of these
16 elemental functions is in effect limited to be in the `array` construct.

- 17 • **Transformational functions**

18 It is unspecified how the transformational functions work when a global array or its subob-
19 ject appears as an argument. A processor shall detect such a reference of these functions
20 and issue a warning message for it. Some intrinsic transformational subroutines are defined
21 in section 7.9 as alternatives to these transformational functions.

22 **7.8 [C] Built-in Elemental Functions**

23 Some built-in elemental functions that can operate each element of array arguments are defined
24 in XcalableMP C. Such a built-in function accepts one or more array sections as its arguments
25 and returns an array-valued result having the same shape and mapping as the argument. The
26 values of the elements of the result are the same as what would have been obtained if the scalar
27 function of the C standard library had been applied separately to the corresponding elements
28 of each array argument.

29 These functions may appear on the right-hand side of an array assignment statement, and
30 it should be preceded by the `array` directive if the array section is distributed.

31 Table 7.1 shows the list of built-in elemental functions in XcalableMP C. Their elementwise
32 behavior is the same as those of the corresponding functions in the C standard library.

33 **7.9 Intrinsic/Built-in Transformational Procedures**

34 Some intrinsic/built-in transformational procedures are defined for the non-elemental operations
35 of arrays.

36 Note that each “array argument” of the following procedures must be an array name or an
37 array section, in XcalableMP Fortran, or an array section, in XcalableMP C, that represents
38 the whole array.

Table 7.1: Built-in elemental functions in XcalableMP C. (The first line refers to the element type of their argument(s) and return value.)

double	float	long double
acos	acosf	acosl
asin	asinf	asinl
atan	atanf	atanl
atan2	atan2f	atan2l
cos	cosf	cosl
sin	sinf	sinl
tan	tanf	tanl
cosh	coshf	coshl
sinh	sinhf	sinhl
tanh	tanhf	tanhl
exp	expf	expl
frexp	frexpf	frexpl
ldexp	ldexpf	ldexpl
log	logf	logl
log10	log10f	log10l
fabs	fabsf	fabsl
pow	powf	powl
sqrt	sqrtf	sqrtl
ceil	ceilf	ceill
floor	floorf	floorl
fmod	fmodf	fmodl

7.9.1 xmp_scatter

1

Format

2

[F] `xmp_scatter(x, a, idx1, ..., idxn)`

[C] `void xmp_scatter(x[:], ..., a[:], ..., idx1[:], ..., ..., idxn[:], ...)`

3

Synopsis

4

The `xmp_scatter` procedure copies the value of each element of an array `a` to the corresponding element of an array `x` that is determined by vectors `idx1`, ..., `idxn`.

5

6

This procedure produces the same result as the following Fortran assignment statement when `x`, `a`, and `idx1`, ..., `idxn` are not mapped.

7

8

```
x(idx1(:, :, ...), ..., idxn(:, :, ...)) = a(:, :, ...)
```

9

If any of the vectors `idx1`, ..., `idxn` have two or more elements with the same value, the behavior and the result of `xmp_scatter` is unspecified.

10

11

Output Arguments

12

- `x` is an array of any type, shape, and mapping.

13

1 **Input Arguments**

- 2 • **a** is an array of the same type as **x** and any shape and mapping.
- 3 • **idx1**, ..., **idxn** are integer arrays of the same shape and mapping as **a**. The number of
- 4 **idx**'s is equal to the rank of **x**.

5 **7.9.2 xmp_gather**

6 **Format**

```
7 [F] xmp_gather(x, a, idx1, ..., idxn)
  [C] void xmp_gather(x[:], ..., a[:], ..., idx1[:], ..., ..., idxn[:], ...)
```

8 **Synopsis**

9 The `xmp_gather` procedure copies the value of each element of an array **a** determined by vectors
10 **idx1**, ..., **idxn** to the corresponding element of an array **x**.

11 This procedure produces the same result as the following Fortran assignment statement when
12 **x**, **a**, and **idx1**, ..., **idxn** are not mapped.

```
13 x(:, :, ...) = a(idx1(:, :, ...), ..., idxn(:, :, ...))
```

14 **Output Arguments**

- 15 • **x** is an array of any type, shape, and mapping.

16 **Input Arguments**

- 17 • **a** is an array of the same type as **x** and any shape and mapping.
- 18 • **idx1**, ..., **idxn** are integer arrays of the same shape and mapping as **x**. The number of
- 19 **idx**'s is equal to the rank of **a**.

20 **7.9.3 xmp_pack**

21 **Format**

```
22 [F] xmp_pack(v, a, [mask])
  [C] void xmp_pack(v[:], a[:], ..., [mask[:], ...])
```

23 **Synopsis**

24 The `xmp_pack` procedure packs all of the elements of an array **a**, if **mask** is not specified, or
25 the elements selected by **mask**, to a vector **v** according to the array element order of the base
26 language.

27 **Output Arguments**

- 28 • **v** is a one-dimensional array of any type, size, and mapping.

29 **Input Arguments**

- 30 • **a** is an array of the same type as **v** and any shape and mapping.
- 31 • (optional) **mask** is an array of default logical, in XcalableMP Fortran, or of type `_Bool`, in
32 XcalableMP C, that has the same shape and mapping as **a**.

7.9.4 xmp_unpack 1

Format 2

```
[F]      xmp_unpack(a, v, [mask])
[C] void xmp_unpack(a[:]...., v[:], [mask[:]....])
```

Synopsis 4

The `xmp_unpack` procedure unpacks a vector `v` to all the elements of an array `a`, if `mask` is not specified, or the elements selected by a mask `mask` according to the array element order of the base language. 5
6
7

Output Arguments 8

- `a` is an array of any type, shape, and mapping. 9

Input Arguments 10

- `v` is a one-dimensional array of the same type of `a` and any shape and mapping. 11
- (optional) `mask` is an array of default logical, in XcalableMP Fortran, or of type `_Bool`, in XcalableMP C, that has the same shape and mapping as `a`. 12
13

7.9.5 xmp_transpose 14

Format 15

```
[F]      xmp_transpose(x, a, opt)
[C] void xmp_transpose(x[:][:], a[:][:], int opt)
```

Synopsis 17

The `xmp_transpose` procedure sets the result obtained by transposing a matrix `a` to a matrix `x`. 18

Output Arguments 19

- `x` is a two-dimensional array of any type, shape, and mapping. 20

Input Arguments 21

- `a` is a two-dimensional array of the same type as `x` and any mapping. The extent of the first dimension is equal to that of the second dimension of `x`, and the extent of the second dimension is equal to that of the first dimension of `x`. 22
23
24
- `opt` is an integer scalar. If `opt` is 0, the value of `a` remains unchanged after calling this procedure. If `opt` is 1, the value may be changed. 25
26

7.9.6 xmp_matmul 27

Format 28

```
[F]      xmp_matmul(x, a, b)
[C] void xmp_matmul(x[:][:], a[:][:], b[:][:])
```

29

1 Synopsis

2 The `xmp_matmul` procedure computes the product of matrices `a` and `b`, and it sets the result to
3 a matrix `x`.

4 Output Arguments

- 5 • `x` is a two-dimensional array of any numerical type, shape and mapping.

6 Input Arguments

- 7 • `a` is a two-dimensional array of the same type of `x` and any mapping. The extent of the
8 first dimension is equal to that of `x`.
- 9 • `b` is a two-dimensional array of the same type of `x` and any mapping. The extent of the
10 first dimension is equal to that of the second dimension of `a`, and the extent of the second
11 dimension is equal to that of `x`.

12 7.9.7 `xmp_sort_up`

13 Format

```
14 [F]      xmp_sort_up(v1, v2)  
14 [C] void xmp_sort_up(v1[:], v2[:])
```

15 Synopsis

16 The `xmp_sort_up` procedure sets the result obtained by sorting elements of a vector `v2` in as-
17 cending order to a vector `v1`.

18 Output Arguments

- 19 • `v1` is a one-dimensional array of any numerical type, shape, and mapping.

20 Input Arguments

- 21 • `v2` is a one-dimensional array of the same type, shape, and mapping as `v1`.

22 7.9.8 `xmp_sort_down`

23 Format

```
24 [F]      xmp_sort_down(v1, v2)  
24 [C] void xmp_sort_down(v1[:], v2[:])
```

25 Synopsis

26 The `xmp_sort_down` procedure sets the result obtained by sorting elements of a vector `v2` in
27 descending order to a vector `v1`.

28 Output Arguments

- 29 • `v1` is a one-dimensional array of any numerical type, shape and mapping.

Input Arguments

1

- `v2` is a one-dimensional array of the same type, shape, and mapping as `v1`.

2

Chapter 8

OpenMP in XcalableMP Programs

The usage of OpenMP directives in XcalableMP programs is subjected to the following basic rule.

- XcalableMP directives and the invocation of an XcalableMP intrinsic/built-in procedure should be single-threaded, and they may therefore be placed in the sequential part, or one of the `single`, `master`, or `critical` regions that are closely nested inside a `parallel` region whose parent thread is the initial thread;
- with the exception that the XcalableMP's `loop` directive that controls a loop can be placed immediately inside the OpenMP's parallel loop directive (`parallel do` for Fortran and `parallel for` for C), which controls the identical loop.

The behavior of coarray references in a `parallel` region is implementation-defined.

Examples

Assume that the following codes are placed in the sequential part of the program.

```
_____ XcalableMP C _____  
#pragma omp parallel for  
for (...){  
    #pragma xmp barrier // NG because not single-threaded  
}
```

```
_____ XcalableMP C _____  
#pragma omp parallel for  
for (...){  
    #pragma omp single  
    {  
5      #pragma xmp barrier // OK because single-threaded  
                                     // (inside a single region)  
    }  
}
```

```
_____ XcalableMP C _____  
#pragma omp parallel for  
#pragma xmp loop // OK because immediately nested  
for (...){  
    ...  
5 }
```

```
                    XcalableMP C
#pragma xmp loop // OK because single-threaded (not nested)
#pragma omp parallel for
for (...){
    ...
5 }

```

```
                    XcalableMP C
#pragma xmp loop // OK because single threaded (not nested)
for (...){
    #pragma omp parallel for
    for (...) { ... }
5 }

```

```
                    XcalableMP C
#pragma omp parallel for
for (...){
    #pragma xmp loop // NG because not immediately nested
    for (...) { ... }
5 }

```

Bibliography

- 2 [1] OpenMP Architecture Review Board, “OpenMP Application Program Interface Version
3 3.1”, <http://www.openmp.org/mp-documents/OpenMP3.1.pdf> (2011).
- 4 [2] High Performance Fortran Forum, “High Performance Fortran Language Specification Ver-
5 sion 2.0”, <http://hpf.rice.edu/versions/hpf2/hpf-v20.pdf> (1997).
- 6 [3] Message Passing Interface Forum, “MPI: A Message-Passing Interface Standard Version
7 2.2”, <http://www.mpi-forum.org/docs/mpi-2.2/mpi22-report.pdf> (2009).
- 8 [4] Japan Association of High Performance Fortran, “HPF/JA Language Specification”, [http:
9 //www.hpfpc.org/jahpf/spec/hpfja-v10-eng.pdf](http://www.hpfpc.org/jahpf/spec/hpfja-v10-eng.pdf) (1999).
- 10 [5] Yuanyuan Zhang, Hidetoshi Iwashita, Kuninori Ishii, Masanori Kaneko, Tomotake Naka-
11 mura, and Kohichiro Hotta, “Hybrid Parallel Programming on SMP Clusters Using XP-
12 Fortran and OpenMP”, Proceedings of the International Workshop on OpenMP (IWOMP
13 2010), Vol. 6132 of Lecture Notes in Computer Science, pp. 133–148, Springer (2010).
- 14 [6] Hidetoshi Iwashita, Naoki Sueyasu, Sachio Kamiya, and Matthijs van Waveren, “VPP
15 Fortran and the design of HPF/JA extensions”, Concurrency and Computation — Practice
16 & Experience, Vol. 14, No. 8–9, pp. 575–588, Wiley (2002).
- 17 [7] Jinpil Lee, Mitsuhsa Sato, and Taisuke Boku, “OpenMPD: A Directive-Based Data Parallel
18 Language Extension for Distributed Memory Systems”, Proceedings of the 2008 Interna-
19 tional Conference on Parallel Processing, pp. 121-128 (2008).

1 Appendix A

2 Programming Interface for MPI

3 This chapter describes the programming interface for MPI, which is widely used for parallel
4 programming in cluster computing. Users can introduce MPI functions to XcalableMP using
5 the interface.

6 A.1 Call MPI functions from an XcalableMP program

7 XcalableMP provides the following user API functions to call MPI functions from an XcalableMP
8 program.

- 9 • `xmp_get_mpi_comm`
- 10 • `xmp_init_mpi`
- 11 • `xmp_finalize_mpi`

12 A.1.1 `xmp_get_mpi_comm`

13 Format

14 [F] integer function `xmp_get_mpi_comm()`
[C] MPI_Comm `xmp_get_mpi_comm(void)`

15 Synopsis

16 `xmp_get_mpi_comm` returns the handle of the communicator associated with the executing node
17 set.

18 Arguments

19 none.

20 A.1.2 `xmp_init_mpi`

21 Format

22 [F] `xmp_init_mpi()`
[C] void `xmp_init_mpi(int *argc, char ***argv)`

23 Synopsis

24 `xmp_init_mpi` initializes the MPI execution environment.

Arguments

In XcalableMP C, the command-line arguments `argc` and `argv` should be given to `xmp_init_mpi`.

A.1.3 `xmp_finalize_mpi`

Format

[F] `xmp_finalize_mpi()`

[C] `void xmp_finalize_mpi(void)`

Synopsis

`xmp_finalize_mpi` terminates the MPI execution environment.

Arguments

none.

Example

XcalableMP C

```
#include <stdio.h>
#include "mpi.h"
#include "xmp.h"
5 #pragma xmp nodes p[4]

int main(int argc, char *argv[]) {
    xmp_init_mpi(&argc, &argv)

10     int rank, size;
    MPI_Comm_rank(MPI_COMM_WORLD, &rank);
    MPI_Comm_size(MPI_COMM_WORLD, &size);

    #pragma xmp task on p[1:2]
15     {
        MPI_Comm comm = xmp_get_mpi_comm(); // get the MPI communicator of p[1:2]

        int rank, size;
        MPI_Comm_rank(comm, &rank);
20     MPI_Comm_size(comm, &size);
    }

    xmp_finalize_mpi();

25     return 0;
}
```

1 A.2 Call XcalableMP functions from an MPI program

2 XcalableMP provides the following user API functions to call XcalableMP functions from an
3 MPI program.

- 4 • `xmp_init`
- 5 • `xmp_finalize`

6 The XcalableMP functions should appear between `xmp_init` and `xmp_finalize`. Please refer
7 to chapter 6 and examples in this section about how arguments are passed to the XcalableMP
8 functions.

9 A.2.1 `xmp_init`

10 Format

```

11 [F]          xmp_init(comm)
           integer comm
12 [C] void     xmp_init(MPI_Comm comm)
```

12 Synopsis

13 `xmp_init` initializes the XcalableMP execution environment. The MPI communicator indicated
14 in `xmp_init` is used as an executing node set in the XcalableMP functions. `xmp_init` should
15 appear after `MPI_Init`.

16 Arguments

17 MPI Communicator `comm` should be given to `xmp_init`.

18 A.2.2 `xmp_finalize`

19 Format

```

20 [F]          xmp_finalize()
21 [C] void     xmp_finalize(void)
```

21 Synopsis

22 `xmp_finalize` finalizes the XcalableMP execution environment. `xmp_finalize` should appear
23 before `MPI_Finalize`.

24 Arguments

25 none.

26 Example (C language)

```

MPI C
-----
#include <mpi.h>
#include <xmp.h>
#define N 5

int main(int argc, char **argv)
{
```

```

int a[N], comm_size;

MPI_Init(&argc, &argv);
MPI_Comm_size(MPI_COMM_WORLD, &comm_size);

xmp_init(MPI_COMM_WORLD);
foo(N*comm_size, a); // foo() is an XMP function
xmp_finalize();

MPI_Finalize();
return 0;
}

```

10

15

XcalableMP C

```

void foo(int total_elements, int a[total_elements])
{
#pragma xmp nodes p[*]
#pragma xmp template t[total_elements]
5 #pragma xmp distribute t[block] onto p
#pragma xmp align a[i] with t[i]

#pragma xmp loop on t[i]
10   for(int i=0;i<total_elements;i++)
      a[i] = i;
}

```

10

Example (Fortran)

MPI Fortran

```

program test
  include 'mpif.h'
  integer, parameter :: N = 5
  integer :: a(N), ierr, comm_size

  call MPI_INIT(ierr)
  call MPI_COMM_SIZE(MPI_COMM_WORLD, comm_size, ierr)

  call xmp_init(MPI_COMM_WORLD)
  call foo(comm_size*5, a)
  call xmp_finalize()

  call MPI_FINALIZE(ierr)
end program test

```

5

10

XcalableMP Fortran

```

subroutine foo(total_elements, a)
  integer total_elements, a(total_elements)
!$xmp nodes p(*)
!$xmp template t(total_elements)
5 !$xmp distribute t(block) onto p
!$xmp align a(i) with t(i)

```

5


```
10 | !$xmp loop on t(i)
    |   do i=1, total_elements
    |     a(i) = i
    |   end do
    | end subroutine foo
```


1 Appendix B

2 Interface to Numerical Libraries

3 This chapter describes the XcalableMP interfaces to existing MPI parallel libraries, which is
4 effective to achieve high productivity and performance of XcalableMP programs.

5 B.1 Interface Design

6 The recommended design of the interface is as follows:

- 7 • Numerical library routines can be invoked by an XcalableMP procedure through an inter-
8 face procedure (Figure B.1).




Figure B.1: Invocation of a library routine using an interface procedure.

- 9 • When the numerical library routine requires information regarding a global array, the
10 interface extracts it from the descriptor using query routines provided by XcalableMP,
11 and passes it to the numerical library routine as an argument.
- 12 • The interface does not affect the behavior of numerical library routines except for restric-
13 tions concerning the XcalableMP specification.

14 B.2 Extended Mapping Inquiry Functions

15 In this section, the extended mapping inquiry functions, which are implementation-defined, are
16 shown. Specifications of the functions below are obtained from the Omni XcalableMP compiler
17 (<http://www.xcalablemp.org/download.html>).

B.2.1 xmp_array_gtol 1

[F]	integer function	xmp_array_gtol(d, dim, g_idx, l_idx)	
	type(xmp_desc)	d	
	integer	dim	
	integer	g_idx	2
	integer	l_idx	

[C]	void	xmp_array_gtol(xmp_desc_t d, int dim, int g_idx, int* l_idx)	
-----	------	--	--

Synopsis 3

The `xmp_array_gtol` function translates a global index specified by `g_idx` of a global array specified by `d` into the corresponding index of its local section, and sets it to an array specified by `l_idx`. If the element of the specified index does not reside in the caller of the function, the resulting array is set to an unspecified value. 4
5
6
7

Input Arguments 8

- `d` is a descriptor of the global array. 9
- `dim` is the target dimension of the global array. 10
- `g_idx` is an index of the global array. 11

Output Argument 12

- `l_idx` is an index of the local array. 13

B.2.2 xmp_array_lsize 14

Format 15

[F]	integer function	xmp_array_lsize(d, dim, lsize)	
	type(xmp_desc)	d	
	integer	dim	16
	integer	lsize	

[C]	int	xmp_array_lsize(xmp_desc_t d, int dim, int *lsize)	
-----	-----	--	--

Synopsis 17

The `xmp_array_lsize` function provides the local size of each dimension of the target global array. Note that the local size does not include the size of the shadow. 18
19

Input Arguments 20

- `d` is a descriptor of a global array. 21
- `dim` is the target dimension of the global array. 22

Output Argument 23

- `lsize` is the local size of the target dimension of the global array. 24

1 B.2.3 xmp_array_laddr

2 Format

3 [C] int xmp_array_laddr(xmp_desc_t d, void **laddr)

4 Synopsis

5 The `xmp_array_laddr` function provides the local address of the target global array.

6 Input Arguments

- 7 • `d` is a descriptor of a global array.

8 Output Arguments

- 9 • `laddr` is the local address of the target global array.

10 B.2.4 xmp_array_lda

11 Format

12	[F] integer function	xmp_array_lda(d, lda)
	type(xmp_desc)	d
	integer	lda
	[C] int	xmp_array_lda(xmp_desc_t d, int* lda)

13 Synopsis

14 The `xmp_array_lda` function provides the leading dimension of the two-dimensional global array.
 15 This function is used to call numerical libraries, such as BLAS.

16 Input Argument

- 17 • `d` is a descriptor of a global array, which must be a two-dimensional array.

18 Output Argument

- 19 • `lda` is a leading dimension of the target global array.

20 B.3 Example

21 This section shows the interface to ScaLAPACK as an example of the XscalableMP interface to
 22 numerical libraries.

23 ScaLAPACK is a linear algebra library for distributed-memory. Communication processes
 24 in the ScaLAPACK routines depend on BLACS (Basic Linear Algebraic Communication Sub-
 25 programs). ScaLAPACK library routines invoked from XscalableMP procedures also depend on
 26 BLACS.

27 **Example 1** This example shows an implementation of the interface for the ScaLAPACK driver
 28 routine `pdgesv`.

```

                                XcalableMP Fortran
subroutine ixmp_pdgesv(n,nrhs,a,ia,ja,da,ipiv,b,ib,jb,db,ictxt,info)

use xmp_lib

5  integer n,nrhs,ia,ja,ib,jb,ictxt,info,desca(9),descb(9),ierr
double precision a,b
type(xmp_desc) da,db,dta,dtb
integer lbound_a1,ubound_a1,lbound_a2,ubound_a2
integer blocksize_a1,blocksize_a2,lead_dim_a
10 integer lbound_b1,ubound_b1,lbound_b2,ubound_b2
integer blocksize_b1,blocksize_b2,lead_dim_b

ierr=xmp_array_lbound(da,1,lbound_a1)
ierr=xmp_array_ubound(da,1,ubound_a1)
15 ierr=xmp_array_lbound(da,2,lbound_a2)
ierr=xmp_array_ubound(da,2,ubound_a2)
ierr=xmp_align_template(da,dta)
ierr=xmp_dist_blocksize(dta,1,blocksize_a1)
ierr=xmp_dist_blocksize(dta,2,blocksize_a2)
20 ierr=xmp_array_lead_dim(da,1,lead_dim_a)

ierr=xmp_array_lbound(db,1,lbound_b1)
ierr=xmp_array_ubound(db,1,ubound_b1)
ierr=xmp_array_lbound(db,2,lbound_b2)
25 ierr=xmp_array_ubound(db,2,ubound_b2)
ierr=xmp_align_template(db,dtb)
ierr=xmp_dist_blocksize(dtb,1,blocksize_b1)
ierr=xmp_dist_blocksize(dtb,2,blocksize_b2)
ierr=xmp_array_lead_dim(db,1,lead_dim_b)

30
desca(1)=1
desca(2)=ictxt
desca(3)=ubound_a1-lbound_a1+1
desca(4)=ubound_a2-lbound_a2+1
35 desca(5)=blocksize_a1
desca(6)=blocksize_a2
desca(7)=0
desca(8)=0
desca(9)=lead_dim_a

40
descb(1)=1
descb(2)=ictxt
descb(3)=ubound_b1-lbound_b1+1
descb(4)=ubound_b2-lbound_b2+1
45 descb(5)=blocksize_b1
descb(6)=blocksize_b2
descb(7)=0
descb(8)=0
descb(9)=lead_dim_b

```

```

50      call pdgesv(n,nhrs,a,ia,ja,desca,ipiv,b,ib,jb,descb,info)

      return
      end
55

```

1 **Example 2** This example shows an XcalableMP procedure using the interface of Example 1.

```

----- XcalableMP Fortran -----
program xmptdgesv

use xmp_lib

5      double precision a(1000,1000)
      double precision b(1000)
      integer ipiv(2*1000,2)
!$xmp nodes p(2,2)
!$xmp template t(1000,1000)
10 !$xmp template t1(2*1000,2)
!$xmp distribute t(block,block) onto p
!$xmp distribute t1(block,block) onto p
!$xmp align a(i,j) with t(i,j)
!$xmp align ipiv(i,j) with t1(i,j)
15 !$xmp align b(i) with t(i,*)
      ...
      integer i,j,ictxt
      integer m=1000,n=1000,nprow=2,npcol=2
      integer icontxt=-1,iwhat=0
20      integer nrhs=1,ia=1,ja=1,ib=1,jb=1,info
      character*1 order
      ...
      order="C"
      ...
25      call blacs_get(icontxt,iwhat,ictxt)
      call blacs_gridinit(ictxt,order,nprow,npcol)
      ...
!$xmp loop (i,j) on t(i,j)
      do j=1,n
30          do i=1,m
              a(i,j) = ...
          end do
      end do
      ...
35 !$xmp loop on t(i,*)
      do i=1,m
          b(i)= ...
      end do
      ...
40      call ixmp_pdgesv(n,nrhs,a,ia,ja,xmp_desc_of(a),ipiv,
*          b,ib,jb,xmp_desc_of(b),ictxt,info)

```

45

```
...  
call blacs_gridexit(ictxt)  
...  
stop  
end
```


1 Appendix C

2 Memory-layout Model

3 In this chapter, the memory-layout model of global data in the Omni XcalableMP compiler
4 (<http://www.xcalablemp.org/download.html>) is presented for reference.

5 According to the XcalableMP specification, a global array is distributed onto a node array
6 according to the data-mapping directives, and as a result, a node owns a set of elements.

7 On each node, all and only the elements of the global array that it owns are gathered to form
8 the local array having the same rank as the global. For each axis of the global data, all and only
9 the indices that the node owns are packed to the axis of the local array so that the sequence
10 can be maintained, with the shadow area, if any, added at the lower and/or upper bound of the
11 axis.

12 Eventually, the local array is stored in memory on each node according to the rule for storing
13 arrays in the base language, that is, in row-major order in XMP/Fortran and in column-major
14 order in XMP/C.

15 Note that owing to the model above, the memory usage may be non-uniform among the
16 nodes.

17 Example

```
----- XcalableMP Fortran -----  
!$xmp nodes p(4,4)  
!$xmp template t(64,64)  
!$xmp distribute t(block,block) onto p  
5      real a(64,64)  
!$xmp align a(i,j) with t(i,j)  
!$xmp shadow a(1,1)
```

18 The array `a` is distributed by a format of `(block,block)` onto a two-dimensional node array
19 `p`, and each node owns a local array including a shadow area. Then, the local array is stored in
20 memory on each node, as shown in Figure C.1.




Figure C.1: Example of memory layout in the Omni XcalableMP compiler.

1 Appendix D

2 XcalableMP I/O

3 D.1 Categorization of I/O

4 XcalableMP has three kinds of I/O.

5 D.1.1 Local I/O

6 Local I/O is a method that is employed to use I/O statements and standard I/O functions in
7 the base languages, where I/O statements and functions are used without any directives.

8 I/O statements (in XcalableMP Fortran) and I/O functions (in XcalableMP C) are executed
9 locally similar to other execution statements. It depends on the system which nodes can handle
10 the I/O statements and functions.

11 Local I/O can read a file written by the base language, and vice versa.

12 [F] The name of a global array in the I/O list describes the entire area of the array located
13 in each node.

14 An array element of a global array can be referred to as an I/O item only in the node where
15 it is located.

16 [F] No array section of a global array can be referred to as an I/O item.

17 D.1.2 Master I/O[F]

18 Master I/O is input and output for the file that corresponds to an executing node set. Master
19 I/O is a collective execution.

20 In master I/O, a global data object is input and output as if it was executed only by a master
21 node, which represents the executing node set, through its local copy of the data.

22 The master node is chosen by the system arbitrarily from among the executing node set,
23 and is unique to the executing node set during execution of the program.

24 Master I/O is provided in the form of directives of XcalableMP Fortran.

25 A global array as an I/O item is accessed in the sequential order of array elements. When
26 a local variable is read from a file, the value is copied to all nodes of the executing node set.
27 When a local variable or an expression is written to a file, only the value of the data on the
28 master node is written.

29 Master I/O can read a file written in the base language, and vice versa.

30 D.1.3 Global I/O

31 Global I/O is input and output for a file that corresponds to an executing node set. Some
32 executions of the global I/O are collective and the others are independent. In a large system with

Table D.1: Global I/O.

	independent/collective	access method
Collective I/O	collective	sequential access
Atomic I/O	independent	sequential access
Direct I/O	independent	direct access

many nodes, the global I/O can be expected to have higher speed and less memory consumption execution than master I/O.

[F] It is provided in the form of directives for some of I/O statements, such as OPEN, CLOSE, READ, and WRITE statements.

[C] It is provided in the form of service functions and an include file.

Global I/O can handle only unformatted (binary) files. In XcalableMP Fortran, implied DO loops and some specifiers cannot be used. In XcalableMP C, formatted I/O libraries, including `fprintf()` and `fscanf()`, are not provided.

Global I/O can read a file written in MPI-IO, and vice versa.

[F] File formats are not compatible between XcalableMP Fortran and the base language because global I/O does not generate or access file headers and footers that are specific to the base language.

There are three kinds of global I/O, as shown in Table D.1. **Collective** global I/O is for collective execution and sequential file access. It handles global data in a sequential order, similar to master I/O. **Atomic** global I/O is for independent execution and sequential file access. Executing nodes share file positioning of the global I/O file and execute each I/O statement and library call mutually. **Direct** global I/O is for independent execution and direct file access. Each executing node has its own file positioning and accesses a shared file independently.

Restriction

- The name of a global array may not be declared in a namelist group. That is, NAMELIST I/O is not allowed for global arrays.

Advice to programmers

Local I/O is useful for debugging that is focused on a node because local I/O is executed on each node individually.

Master I/O is a directive extension, where the execution result matches that of the base language, ignoring directive lines.

Global I/O aims for highly parallel I/O using thousands of nodes. It is limited to binary files, and it avoids the extreme concentration of computational load and memory consumption to specific nodes using MPI-IO or other parallel I/O techniques.

D.2 File Connection

A file is connected to a unit in XcalableMP Fortran and to a file handler in XcalableMP C. This operation is called **file connection**. Local I/O connects a file to each node independently. Master I/O and global I/O connect a file to an executing node set collectively.

There are two ways of specifying file connections, dynamic connection and preconnection. Dynamic connection connects a file during the execution of the program. Preconnection connects a file at the beginning of execution of the program, and it can therefore execute I/O statements

1 and functions without the prior execution of an OPEN statement or a function call to open the
2 file.

3 **D.2.1 File Connection in Local I/O**

4 The language processor of the base language connects the file to each node. It is implementation-
5 defined which nodes can access the standard input, output, and error files. The behavior of the
6 accesses to files having the same name on multiple nodes is also implementation-defined. The
7 only primary node can access the standard input, output, and error files.

8 **D.2.2 [F] File Connection in Master I/O**

9 An OPEN statement that is specified with a master I/O directive connects a file to the executing
10 node set. When a master I/O file is connected by a READ statement or a WRITE statement
11 without encountering any OPEN statement, the name and attributes of the file depend on the
12 language system of the base language. Disconnection from a master I/O file is executed by a
13 CLOSE statement or by the termination of the program.

14 The dynamic connection must be executed collectively by all nodes sharing the file with the
15 same unit number. Two executing node sets may employ the same unit number only if they
16 have no common node.

17 The standard input, output, and error files are preconnected to the entire node set. There-
18 fore, master I/O executed on the entire node set is always allowed without OPEN or CLOSE
19 statements.

20 **D.2.3 File Connection in Global I/O**

21 The dynamic connection of global I/O is a collective execution, and is valid for the executing
22 node set. Global I/O files cannot be preconnected.

23 **[F]**

24 An OPEN statement that is specified with a global I/O directive connects a file to the executing
25 node set. Disconnection from a global I/O file is executed by a CLOSE statement or by the
26 termination of the program.

27 The dynamic connection must be executed collectively by all nodes sharing the file with the
28 same unit number. Two executing node sets may employ the same unit number only if they
29 have no common node.

30 **[C]**

31 A library function to open a global I/O file connects the file to the executing node set. Discon-
32 nection from a global I/O file is executed by a library function to close the file or terminate the
33 program.

34 **D.3 Master I/O**

35 A master I/O construct executes data transfer between a file and an executing node set via a
36 master node of the executing node set. For a global array, the virtual sequential order of the
37 array elements is visible.

D.3.1 master_io Construct 1

Syntax 2

[F] !\$xmp master_io 3
 io-statement

[F] !\$xmp master_io begin 3
 io-statement

...

!\$xmp master_io end 4

where *io-statement* is one of: 4

- OPEN statement 5
- CLOSE statement 6
- READ statement 7
- WRITE statement 8
- PRINT statement 9
- BACKSPACE statement 10
- ENDFILE statement 11
- REWIND statement 12
- INQUIRE statement 13

Restriction 14

- The following items, including a global array or a subobject of a global array, must not appear in an input item or output item. 15
 - A substring-range 17
 - A section-subscript 18
 - An expression including operators 19
 - An *io-implied-do-control* 20
- An I/O statement specified with a master I/O directive must be executed collectively on the node set that is connected to the file. 21
- Internal file I/O is not permitted to be a master I/O. 23

Description 24

An I/O statement that is specified with a master I/O directive accesses a file whose format is the same as that of the base language. The access, including connection, disconnection, input and output, file positioning, and inquiry, is collective, and must be executed on the same node set as the one on which the file was connected. 25

A master node, which is a unique node to an executing node set, is chosen by the language system. Master I/O works as if all file accesses were executed only on the master node. 26

The operations for I/O items are summarized in Table D.2. 27

28

29

30

31

Table D.2: Operations for I/O.

	I/O item	operation
input item	name of global array	The data elements that are read from the file in the sequential order of array elements are distributed onto the global array on the node set. The file positioning increases according to the size of data.
	array element of global array	The data element that is read from the file is copied to the element of the global array on the node to which it is mapped. The file positioning increases according to the size of data.
	local variable	The data element that is read from the file is replicated to the local variables on all nodes of the executing node set. The file positioning increases according to the size of data.
	implied DO loop	For each input item, repeat the above operation.
output item	name of global array	The data elements of the global array are collected and are written to the file in the sequential order of array elements. The file positioning increases according to the size of data.
	array element of global array	The element of the global array is written to the file. A file position increases according to the size of data.
	local variable and expression	The value evaluated on the master node is written to the file. The file positioning increases according to the size of data.
	implied DO loop	For each output item, repeat the above operation.

1 Namelist input and output statements cannot treat global arrays. A namelist output state-
 2 ment writes the values on the master node to the file. In the namelist input, each item of the
 3 namelist is read from the file to the master node if it is recorded in the file. Then, all items of
 4 the namelist are replicated onto all nodes of the executing node set from the master node even
 5 if some items are not read from the file.

6 IOSTAT and SIZE specifiers and specifiers of the INQUIRE statement that can return values
 7 always return the same value among the executing node set.

8 When a condition that is specified by the ERR, END, or EOR specifier is satisfied, all nodes
 9 of the executing node set are branched together to the same statement.

10 **Advice to implementers**

11 It is recommended to provide such a compiler option that local I/O statements (specified without
 12 directives) are regarded as master I/O statements (specified with `master_io` directives).

13 **D.4 [F] Global I/O**

14 Global I/O performs unformatted data transfer, and can be expected to have a higher perfor-
 15 mance and lower memory consumption than master I/O. The file format is compatible with the

one in MPI-IO. 1

There are three kinds of Global I/O, namely collective I/O, atomic I/O, and direct I/O. 2

D.4.1 Global I/O File Operation 3

`global_io` construct is defined as follows. 4

Syntax 5

```
[F] !$xmp global_io [atomic / direct]
      io-statement
```

```
[F] !$xmp global_io [atomic / direct] begin
      io-statement
```

```
...
!$xmp end global_io
```

The first syntax is just a shorthand of the second syntax. 7

Restriction 8

I/O statements and specifiers that are available for an *io-statement* are shown in the following table. The definition of each specifier is described in the specification of the base language. 9

Case of `global_io` construct without a direct clause: 10

I/O statement	available specifiers
OPEN	UNIT, IOSTAT, FILE, STATUS, POSITION, ACTION, ACCESS, FORM
CLOSE	UNIT, IOSTAT, STATUS
READ	UNIT, IOSTAT
WRITE	UNIT, IOSTAT

Case of `global_io` construct with a direct clause: 11

I/O statement	available specifiers
OPEN	UNIT, IOSTAT, FILE, STATUS, RECL, ACTION, ACCESS, FORM
CLOSE	UNIT, IOSTAT, STATUS
READ	UNIT, REC, IOSTAT
WRITE	UNIT, REC, IOSTAT

The input item and output item of a data transfer statement with a `global_io` directive must be the name of a variable. 12

Description 13

Global I/O construct connects, disconnects, inputs, and outputs the global I/O file, which is compatible with MPI-IO. 14

The standard input, output, and error files cannot be a Global I/O file. A Global I/O file cannot preconnect to any unit or any file handler, and must be explicitly connected by the OPEN statement that is specified with a `global_io` directive. 15

The OPEN statement that is specified with a `global_io` directive is collective execution, and the file is shared among the executing node set. A file that has already been opened by another 16

1 OPEN statement with a `global_io` directive cannot be reopened by an OPEN statement with
 2 or without a `global_io` directive before closing it.

3 A global I/O file must be disconnected explicitly by a CLOSE statement that is specified with
 4 a `global_io` directive; otherwise, the result of I/O is not guaranteed. The CLOSE statement
 5 that is specified with a `global_io` directive is a collective execution, and must be executed by
 6 the same executing node set as the one where the OPEN statement is executed.

7 Utilizable values of the specifiers in I/O statements are shown in the following table. Defi-
 8 nitions of the specifiers are described in the specification of the base language.

9 • OPEN statement

specifiers	value	default
UNIT	external file unit (scalar constant expression)	not omissible
FILE	file name (scalar CHARACTER expression)	not omissible
STATUS	'OLD', 'NEW', 'REPLACE' or 'UNKNOWN'	'UNKNOWN'
POSITION	'ASIS', 'REWIND' or 'APPEND'	'ASIS'
ACTION	'READ', 'WRITE' or 'READ-WRITE'	implementation-defined
RECL	the value of the record length (scalar constant expression)	not omissible
ACCESS	'SEQUENTIAL' or 'DIRECT'	'SEQUENTIAL'
FORM	'FORMATTED' or 'UNFORMATTED'	For direct access, UNFORMATTED. For sequential access, this specifier shall not be omitted.

10 POSITION is available only if the directive has no direct clause. RECL is available only
 11 if the directive has a direct clause. For direct I/O, the ACCESS specifier shall appear and
 12 the value shall be evaluated to DIRECT. For collective I/O and atomic I/O, the value of
 13 the ACCESS specifier shall be evaluated to SEQUENTIAL if this specifier appears. For
 14 collective I/O and atomic I/O, the FORM specifier shall appear and the value shall be
 15 evaluated to UNFORMATTED. For direct I/O, the value of the FORM specifier shall be
 16 evaluated to UNFORMATTED if this specifier appears.

17 • CLOSE statement

specifiers	value	default
UNIT	external file unit (scalar constant expression)	not omissible.
STATUS	'KEEP' or 'DELETE'	'KEEP'

18 • READ/WRITE statement

19 REC is available only if the directive has a direct clause.

20 • When a scalar variable of default INTEGER is specified to the IOSTAT specifier, it be-
 21 comes defined with an error code after execution.

specifiers	value	default
UNIT	external file unit (scalar constant expression)	not omissible
REC	the value of the number of record (scalar constant expression)	not omissible

OPEN, CLOSE, READ, and WRITE statements that are specified with `global_io` directives without atomic or direct clauses are called collective OPEN, collective CLOSE, collective READ, and collective WRITE statements, respectively. All of these statements are called collective I/O statements.

OPEN, CLOSE, READ, and WRITE statements specified with `global_io` directives having atomic clauses are called atomic OPEN, atomic CLOSE, atomic READ, and atomic WRITE statements, respectively. All of these statements are called atomic I/O statements.

OPEN, CLOSE, READ, and WRITE statements specified with `global_io` directives with direct clauses are called direct OPEN, direct CLOSE, direct READ, and direct WRITE statements, respectively. All of these statements are called direct I/O statements.

The file connected by a collective, atomic, or direct OPEN statement can be read/written only by the same type of READ/WRITE statements. The file can be disconnected by the same type of CLOSE statement. Different types of global I/O cannot be executed together for the same file or the same unit. For example, atomic I/O statements cannot be executed for the unit connected by a collective OPEN statement.

D.4.1.1 `file_sync_all` Directive

Two data accesses cause a conflict if they access the same absolute byte displacements of the same file, and at least one is a write access. When two accesses to the same file conflict in direct or collective I/O, the following `file_sync_all` directive to the file must be executed.

Syntax

```
!$xmp file_sync_all([UNIT=]file-unit-number)
```

The `file_sync_all` directive is an execution directive and collective to the nodes connected to the specified file-unit-number. The execution of a `file_sync_all` directive first synchronizes all the nodes connected to the specified file-unit-number, and then causes all previous writes to the file by the nodes to be transferred to the storage device. If some nodes have made updates to the file, then all such updates become visible to subsequent reads of the file by the nodes.

D.4.2 Collective Global I/O Statement

Collective I/O statements read/write shared files and can handle global arrays.

All collective I/O statements execute collectively. In collective I/O, all accesses to a file, such as connection, disconnection, input, and output, must be executed on the same executing node set.

The operations for I/O items are summarized in the following table.

D.4.3 Atomic Global I/O Statement

Atomic I/O statements read/write shared files exclusively among executing nodes in arbitrary order. Because it is a nondeterministic parallel execution, the results may differ every time it is executed, even for the same program.

I/O item		operation
input item	name of global array	The values read from a file are assigned to the elements of the global array. The file positioning increases according to the size of the data.
	local variable	The values read from the file are replicated into the local array on all executing nodes. The file positioning increases according to the length of the data.
output item	name of global array	The values of a global array are written to the file in the sequential order of the array elements. The file positioning increases according to the size of the data.
	local variable, expression	The values evaluated on a node are arbitrarily selected by the language processor from the executing node set. The file positioning increases according to the size of the data.

1 Atomic OPEN and CLOSE statements are executed collectively, while atomic READ and
2 WRITE statements are executed independently. A file connected by an atomic OPEN statement
3 can be disconnected only by an atomic CLOSE statement executed on the same executing node
4 set. Atomic READ and WRITE statements can be executed on any single node of the same
5 executing node set.

6 Atomic READ and WRITE statements are exclusively executed. The unit of exclusive
7 operation is a single READ statement or a single WRITE statement.

8 The initial file positioning is determined by the POSITION specifier of the atomic OPEN
9 statement. Then, the file positioning seeks in every READ and WRITE statement according to
10 the length of the input/output data.

11 D.4.4 Direct Global I/O Statement

12 Direct I/O statements read/write shared files by specifying the file positioning for each node.

13 Direct OPEN and CLOSE statements are executed collectively, while direct READ and
14 WRITE statements are executed independently. A file connected by a direct OPEN statement
15 can be disconnected only by a direct CLOSE statement executed on the same executing node
16 set. Direct READ and WRITE statements can be executed on any single node of the same
17 executing node set.

18 Direct READ and WRITE statements read/write local data at the file positioning specified
19 by the REC specifier independently. The file positioning is shifted from the top of the file on
20 the basis of the product of the specifiers RECL (of OPEN statement) and REC (of READ and
21 WRITE statement).

22 In order to guarantee the order of direct I/O statements to the same file position, the file
23 should be closed or the file_sync_all directive should be executed between these statements.
24 Otherwise, the outcome of multiple accesses to the same file position, in which at least one is a
25 write access, is implementation-defined.

26 D.5 [C] Global I/O Library

27 XcalableMP C provides some data types defined in the include file “xmp.h”, a set of library
28 functions with arguments of the data types, and built-in operators to get values of the data
29 types from names of a variable, a template, etc.

30 The following types are provided.

- `xmp_file_t` : file handle 1
 - `xmp_rang_t` : descriptor of array section 2
- The following library functions are provided. Collective function names end with `_all`. 3
- global I/O file operation 4
 - `xmp_fopen_all` : file open 5
 - `xmp_fclose_all` : file close 6
 - `xmp_fseek` : setting (individual) file pointer 7
 - `xmp_fseek_shared_all` : setting shared file pointer 8
 - `xmp_ftell` : displacement of (individual) file pointer 9
 - `xmp_ftell_shared` : displacement of shared file pointer 10
 - `xmp_file_sync_all` : file synchronization 11
 - collective I/O 12
 - `xmp_file_set_view_all` : setting file view 13
 - `xmp_file_clear_view_all` : initializing file view 14
 - `xmp_fread_all` : collective read of local data 15
 - `xmp_fwrite_all` : collective write of local data 16
 - `xmp_fread_darray_all` : collective read of global data 17
 - `xmp_fwrite_darray_all` : collective write of global data 18
 - atomic I/O 19
 - `xmp_fread_shared` : atomic read 20
 - `xmp_fwrite_shared` : atomic write 21
 - direct I/O 22
 - `xmp_fread` : direct read 23
 - `xmp_fwrite` : direct write 24

Data type 25

The following data types are defined in include file `xmp_io.h`. 26

`xmp_file_t` A file handler. It is connected to a file when the file is opened. It has a shared file pointer and an individual file pointer to point where data should be read/written in the file. 27-29

A shared file pointer is a shared resource among all nodes of the node set that has opened the file. Atomic I/O uses a shared file pointer. An (individual) file pointer is an individual resource on each node. Collective I/O and direct I/O use individual file pointers. 30-32

These two file pointers are managed in the structure `xmp_file_t`, and can be controlled and referenced only through the provided library functions. 33-34

`xmp_range_t` Descriptor of array section, including lower bound, upper bound, and stride for each dimension. Functions for operating the descriptor are shown in the following table. The `xmp_allocate_range()` function is used to allocate memory. The `xmp_set_range()` function is used to set ranges of an array section. The `xmp_free_range()` function releases the memory for the descriptor. 35-39

function name	xmp_range_t *xmp_allocate_range(n_dim)	
argument	int n_dim	the number of dimensions
return value	xmp_range_t*	descriptor of array section. NULL is returned when a program abends.

function name	void xmp_set_range(rp, i_dim, lb, length, step)	
argument	xmp_range_t *rp	descriptor
	int i_dim	target dimension
	int lb	lower bound of array section in the dimension i_dim
	int length	length of array section in the dimension i_dim
	int step	stride of array section in the dimension i_dim

function name	void xmp_free_range(rp)	
argument	xmp_range_t *rp	descriptor of array section.

D.5.1 Global I/O File Operation

D.5.1.1 xmp_fopen_all

xmp_fopen_all opens a global I/O file. Collective execution.

function name	xmp_file_t *xmp_fopen_all(fname, amode)	
argument	const char *fname	file name
	const char *amode	equivalent to fopen of POSIX. combination of “rwa+”
return value	xmp_file_t*	file structure. NULL is returned when a program abend.

File view is initialized, where file view is based on the MPI-IO file view mechanism. The value of shared and individual file pointers depends on the value of amode.

amode	intended purpose
r	Open for reading only. File pointer points to the beginning of the file.
r+	Open an existing file for update (reading and writing). File pointer points to the beginning of the file.
w	Create for writing. If a file having that name already exists, it will be overwritten. File pointer points to the beginning of the file.
w+	Create a new file for update (reading and writing). If a file having that name already exists, it will be overwritten. File pointer points to the beginning of the file.
a	Append; open for writing at end-of-file or create for writing if the file does not exist. File pointer points to the end of the file.
a+	Open for append; open (or create if the file does not exist) for update at the end of the file. File pointer points to the beginning of the file.

D.5.1.2 xmp_fclose_all

xmp_fclose_all closes a global I/O file. Collective execution.

function name	int *xmp_fclose_all(fh)	
argument	xmp_file_t *fh	file structure
return value	int	0: normal termination 1: abnormal termination. fh is NULL. 2: abnormal termination. error in MPI_File_close.

D.5.1.3 xmp_fseek

xmp_fseek sets the individual file pointer in the file structure. Independent execution.

function name	int xmp_fseek(fh, offset, whence)	
argument	xmp_file_t *fh	file structure
	long long offset	displacement of current file view from position of whence
	int whence	choose file position SEEK_SET: the beginning of the file SEEK_CUR: current position SEEK_END: the end of the file
return value	int	0: normal termination an integer other than 0: abnormal termination

1 D.5.1.4 xmp_fseek_shared

2 xmp_fseek_shared sets the shared file pointer in the file structure. Independent execution.

function name	int xmp_fseek_shared(fh, offset, whence)	
argument	xmp_file_t *fh	file structure
	long long offset	displacement of current file view from position of whence
	int whence	choose file position SEEK_SET: the beginning of the file SEEK_CUR: current position SEEK_END: the end of the file
return value	int	0: normal termination an integer other than 0: abnormal termination

3 D.5.1.5 xmp_ftell

4 xmp_ftell returns the position of the individual file pointer in the file structure. Independent
5 execution.

function name	long long xmp_ftell(fh)	
argument	xmp_file_t *fh	file structure
return value	long long	Upon successful completion, the function shall open the file and return a non-negative integer representing the lowest-numbered unused file descriptor. Otherwise, a negative number shall be returned.

6 D.5.1.6 xmp_ftell_shared

7 xmp_ftell_shared returns the position of the shared file pointer in the file structure. Independent
8 execution.

function name	long long xmp_ftell_shared(fh)	
argument	xmp_file_t *fh	file structure
return value	long long	Upon successful completion, the function shall open the file and return a non-negative integer representing the lowest numbered unused file descriptor. Otherwise, negative number shall be returned.

D.5.1.7 xmp_file_sync_all

xmp_file_sync_all guarantees completion of access to the file from nodes sharing the file. Two data accesses conflict if they access the same absolute byte displacements of the same file, and at least one is a write access. When two accesses A1 and A2 to the same file conflict in direct or collective I/O, an xmp_file_sync_all to the file must be invoked between A1 and A2; otherwise, the outcome of the accesses is undefined. Collective execution.

function name	int xmp_file_sync_all(fh)	
argument	xmp_file_t *fh	file structure
return value	int	0: normal termination an integer other than 0: abnormal termination

D.5.2 Collective Global I/O Functions

Collective I/O is executed collectively, but using the individual pointer. It reads/writes data from the position of the individual file pointer and moves the position forward by the length of the data.

Before the file access, a file view is often specified. A file view, like a window to the file, spans the positions corresponding to the array elements that are owned by each node. For more details of file view, refer to the MPI 2.0 specification.

D.5.2.1 xmp_file_set_view_all

xmp_file_set_view_all sets a file view to the file. Collective execution.

function name	int xmp_file_set_view_all(fh, disp, desc, rp)	
argument	xmp_file_t *fh	file structure
	long long disp	displacement from the beginning of the file.
	xmp_desc_t desc	descriptor
	xmp_range_t *rp	range descriptor
return value	int	0: normal termination an integer other than 0: abnormal termination

The file view of distributed *desc* limited to range *rp* is set into file structure *fh*.

1 **D.5.2.2 xmp_file_clear_view_all**

2 xmp_file_clear_view_all clears the file view. Collective execution.

3 The positions of the shared and individual file pointers are set to disp, and the elemental
4 data type and the file type are set to MPI_BYTE.

function name	int xmp_file_clear_view_all(fh, disp)	
argument	xmp_file_t *fh	file structure
	long long disp	displacement from the beginning of the file.
return value	int	0: normal termination an integer other than 0: abnormal termination

5 **D.5.2.3 xmp_fread_all**6 xmp_fread_all reads the same data from the position of the shared file pointer onto all of the
7 executing nodes. Collective execution.

function name	ssize_t xmp_fread_all(fh, buffer, size, count)	
argument	xmp_file_t *fh	file structure
	void *buffer	beginning address of read variables
	size_t size	the size of a read data element
	size_t count	the number of read data elements
return value	ssize_t	Upon successful completion, return the size of read data. Otherwise, negative number shall be returned.

8 **D.5.2.4 xmp_fwrite_all**9 xmp_fwrite_all writes individual data on all of the executing nodes to the position of the shared
10 file pointer. Collective execution.11 It is assumed that the file view is set in advance. Each node writes its data into its own file
12 view.

function name	ssize_t xmp_fwrite_all(fh, buffer, size, count)	
argument	xmp_file_t *fh	file structure
	void *buffer	beginning address of written variables
	size_t size	the size of a written data element
	size_t count	the number of written data elements
return value	ssize_t	Upon successful completion, return the size of written data. Otherwise, negative number shall be returned.

13 **D.5.2.5 xmp_fread_darray_all**14 xmp_fread_darray_all reads data cooperatively to the global array from the position of the shared
15 file pointer.16 Data is read from the file to distributed *desc* limited to range *rp*.

function name	ssize_t xmp_fread_darray_all(fh, desc, rp)	
argument	xmp_file_t *fh	file structure
	xmp_desc_t desc	descriptor
	xmp_range_t *rp	range descriptor
return value	ssize_t	Upon successful completion, return the size of read data. Otherwise, negative number shall be returned.

D.5.2.6 xmp_fwrite_darray_all

xmp_fwrite_darray_all writes data cooperatively from the global array to the position of the shared file pointer.

function name	ssize_t xmp_fwrite_darray_all(fh, desc, rp)	
argument	xmp_file_t *fh	file structure
	xmp_desc_t desc	descriptor
	xmp_range_t *rp	range descriptor
return value	ssize_t	Upon successful completion, return the size of read data. Otherwise, negative number shall be returned.

Data is written from distributed *desc* limited to range *rp* to the file.

D.5.3 Atomic Global I/O Functions

Atomic I/O is executed independently, but using the shared pointer. It exclusively reads/writes local data from the position of the shared file pointer, and moves the position forward by the length of the data.

Before atomic I/O is executed, the file view must be cleared.

[Rationale]

Although the file views must be the same on all processes in order to use the shared file pointer, the xmp_file_set_view_all function may set different file views for all nodes. Thus, before atomic I/O is used, the file view must be cleared.

D.5.3.1 xmp_fread_shared

xmp_fread_shared exclusively reads local data from the position of the shared file pointer, and moves the position forward by the length of the data. Independent execution.

function name	ssize_t xmp_fread_shared(fh, buffer, size, count)	
argument	xmp_file_t *fh	file structure
	void *buffer	beginning address of read variables
	size_t size	the size of a read data element
	size_t count	the number of read data elements
return value	ssize_t	Upon successful completion, return the size of read data. Otherwise, negative number shall be returned.

1 D.5.3.2 xmp_fwrite_shared

2 xmp_fwrite_shared exclusively writes local data to the position of the shared file pointer and
3 moves the position forward by the length of the data. Independent execution.

function name	ssize_t xmp_fwrite_shared(fh, buffer, size, count)	
argument	xmp_file_t *fh	file structure
	void *buffer	beginning address of written variables
	size_t size	the size of a written data element
	size_t count	the number of written data elements
return value	ssize_t	Upon successful completion, return the size of written data. Otherwise, negative number shall be returned.

4 D.5.4 Direct Global I/O Functions

5 Direct I/O is executed independently and uses the individual pointer. It individually reads/writes
6 local data from the position of the individual file pointer, and moves the position forward by
7 the length of the data, considering the file view.

8 In order to guarantee the order by xmp_fread and xmp_fwrite functions to the same file
9 position, the file should be closed or the xmp_file_sync_all function should be executed between
10 these functions. Otherwise, the outcome of multiple accesses to the same file position, in which
11 at least one is the xmp_fwrite function, is implementation dependent.

12 Advice to programmers

13 Function xmp_fseek is useful for setting the individual file pointer. It is not recommended to
14 use it together with the file view because of its complexity.

15 D.5.4.1 xmp_fread

16 xmp_fread reads data from the position of the individual file pointer and moves the position
17 forward by the length of the data. Independent execution.

function name	ssize_t xmp_fread(fh, buffer, size, count)	
argument	xmp_file_t *fh	file structure
	void *buffer	beginning address of read variables
	size_t size	the size of a read data element
	size_t count	the number of read data elements
return value	ssize_t	Upon successful completion, return the size of read data. Otherwise, negative number shall be returned.

18 D.5.4.2 xmp_fwrite

19 xmp_fwrite writes data to the position of the individual file pointer and moves the position
20 forward by the length of the data. Independent execution.

function name	ssize_t xmp_fwrite(fh, buffer, size, count)	
argument	xmp_file_t *fh	file structure
	void *buffer	beginning address of written variables
	size_t size	the size of a written data element
	size_t count	the number of written data elements
return value	ssize_t	Upon successful completion, return the size of written data. Otherwise, negative number shall be returned.

1 Appendix E

2 Memory Consistency Model

3 This chapter explains the memory consistency model that is adopted by XcalableMP.

4 Memory consistency models have specified rules regarding multiple data accesses to memory.
5 Because XcalableMP is an extension of the base languages, and its memory consistency model
6 is defined as an extension to them, that is, XcalableMP follows all of the rules that are adopted
7 by base languages.

8 In addition, XcalableMP introduces some rules about global view. In global view, global
9 communication constructs are used to access distributed data. Furthermore, distributed data
10 can be accessed by designating data in local view. Conversely, non-distributed data can be
11 accessed by designating distributed data using global communication constructs in global view.
12 These are not considered under the memory consistency models of the base language because
13 global view is a new concept that was introduced by XcalableMP.

14 Please recall that global communication constructs are collective as described in Section 2.8.

15 E.1 Execution Traces

16 This section explains execution traces that are enabled by the Xcalable memory consistency
17 model.

18 First, instructions are defined as

$$i := \text{xmp_syn} \mid \text{xmp_asyn}(\text{async-id}) \mid \text{wait_asyn}(\text{async-id}) \mid \text{f_stmt}$$

19 where `xmp_syn` denotes a global communication construct with no `async` clause, `xmp_asyn(async-id)`
20 denotes a global communication construct with the clause `async(async-id)`, and `f_stmt` is a
21 statement.

22 Next, operations are defined as

$$o := \text{Fetch}^j i \mid \text{Execute}^j i \mid \text{Reflect}^j i$$

23 where j is a positive integer.

24 Operation `Fetchj i` denotes that instruction i is fetched j times. The integer j is incremented
25 whenever a loop is exited. The instructions that are called multiple times in loops are identified
26 by js . Operation `Executej i` denotes that instruction i is executed, while operation `Reflectj i`
27 denotes that the effect of instruction i is saved to physical memories.

28 Finally, the memory consistency model defines constraints written by a partial order \leq
29 on operations as described below. Execution traces are defined as sequences of operations that
30 follow the order. In the following, $o_1 < o_2$ denotes $o_1 \leq o_2$ and $o_1 \not\leq o_2$. In addition, $o_1 < o_2 < o_3$
31 denotes $o_1 < o_2$ and $o_2 < o_3$.

$$\begin{aligned} \text{Fetch}^{j_1} i_1 < \text{Fetch}^{j_2} i_2 &\text{ implies } \text{Execute}^{j_1} i_1 < \text{Execute}^{j_2} i_2 && \text{(i)} \\ \text{Execute}^{j_1} \text{xmp_syn} < \text{Execute}^{j_2} i_2 &\text{ implies } \text{Reflect}^{j_1} \text{xmp_syn} < \text{Execute}^{j_2} i_2 && \text{(ii)} \\ \text{Execute}^{j_1} \text{xmp_asyn}(\text{async-id}) < \text{Execute}^{j_3} \text{wait_asyn}(\text{async-id}) < \text{Execute}^{j_2} i_2 &\text{ implies } \\ \text{Reflect}^{j_1} \text{xmp_asyn}(\text{async-id}) < \text{Execute}^{j_2} i_2 &&& \text{(iii)} \end{aligned}$$

Figure E.1: Constraints that are required by the XcalableMP memory consistency model.

E.1.1 Common Constraints 1

In this subsection, we explain some constraints that are common to both synchronous and asynchronous communications. 2
3

In the XcalableMP memory consistency model, instructions are executed in the order in which they are fetched. This is represented by i in Figure E.1. 4
5

E.1.2 Constraints for Synchronous Communications 6

The constructs `reflect`, `gmove` (and its subsequent assignment statement), `reduction`, and `bcast` are synchronous if `async` is not specified. This means that executions of these constructs guarantee the completion of data synchronization. That is, global communication constructs read data that are written by previously executed statements, and their subsequent statements and global communication constructs read data that are written by global communication constructs. This is given by ii in Figure E.1 7
8
9
10
11
12

For example, in the following code, the assignment statement `g(:)=h(:)` is guaranteed to be completed before the second `gmove` construct is executed. Therefore, the value of `g(i)` must be `i` when the assignment statement `x(:)=g(6:10)` is executed. 13
14
15

Finally, the value of `x(i)` on `p(1)` should be `i+5`. 16

```

XcalableMP Fortran
!$xmp nodes p(2)
!$xmp template t(10)
!$xmp distribute (block) onto p :: t
integer :: g(10), h(10)
5 !$xmp align (i) with t(i) :: g, h
integer x(5)

!$xmp loop on t(i)
do i=1,10
10 h(i)=i
end do

!$xmp gmove
g(:)=h(:)
15 !$xmp gmove
x(:)=g(6:10)
```

E.1.3 Constraints for Asynchronous Communications 17

The constructs `reflect`, `gmove` (and its following assignment statement), `reduction`, and `bcast` are asynchronous if `asyncs` are specified. Completions of data read and written by these global 18
19

1 communication constructs are not guaranteed until `wait_asyncs` are executed. This is repre-
 2 sented by iii in Figure E.1.

3 For example, in the following code, the assignment statement `g(:)=h(:)` may not be com-
 4 pleted before the second `gmove` construct is executed as the first `gmove` construct has `async`
 5 clause. Therefore, the value of `g(i)` is not guaranteed to be `i+5`. Of course, the value of `x(i)`
 6 on `p(1)` is not guaranteed to be `i+5`.

```

XcalableMP Fortran
!$xmp nodes p(2)
!$xmp template t(10)
!$xmp distribute (block) onto p :: t
integer :: g(10), h(10)
5 !$xmp align (i) with t(i) :: g, h
integer x(5)

!$xmp loop on t(i)
do i=1,10
10 h(i)=i
end do

!$xmp gmove async(1)
g(:)=h(:)
15 !$xmp gmove
x(:)=g(6:10)
!$xmp wait_async(1)

```

7 The `wait_async(async-id)` guarantees the completion of a global communication construct
 8 that has *async-id*. Therefore, the value of `x(i)` is not guaranteed to be `i+5` in the following
 9 program:

```

XcalableMP Fortran
!$xmp nodes p(2)
!$xmp template t(10)
!$xmp distribute (block) onto p :: t
integer :: g(10), h(10)
5 !$xmp align (i) with t(i) :: g, h
integer x(5)

!$xmp loop on t(i)
do i=1,10
10 h(i)=i
end do

!$xmp gmove async(1)
g(:)=h(:)
15 !$xmp wait_async(1)
!$xmp gmove
x(:)=g(6:10)

```

10 Assignment statements in local view and `gmove` constructs in global view may race. The
 11 value of `x(5)` is not guaranteed to be 6, and may be 10 in the following program:

```

XcalableMP Fortran
!$xmp nodes p(2)
!$xmp template t(10)
!$xmp distribute (block) onto p :: t
integer :: g(10), h(10)
5 !$xmp align (i) with t(i) :: g, h
integer x(5)

integer l(5), m(5)
!$xmp local_alias l => g
10 !$xmp local_alias m => h

!$xmp loop on t(i)
do i=1,10
h(i)=i
15 end do

!$xmp gmove async(1)
g(:)=h(:)
l(5)=6
20 !$xmp wait_async(1)
x(5)=l(5)

```

By avoiding the race, the value of `x(5)` is guaranteed to be 6 as follows:

1

```

XcalableMP Fortran
!$xmp nodes p(2)
!$xmp template t(10)
!$xmp distribute (block) onto p :: t
integer :: g(10), h(10)
5 !$xmp align (i) with t(i) :: g, h
integer x(5)
integer l(5), m(5)
!$xmp local_alias l => g
!$xmp local_alias m => h
10
!$xmp loop on t(i)
do i=1,10
h(i)=i
end do
15
!$xmp gmove async(1)
g(:)=h(:)
!$xmp wait_async(1)
l(5)=6
20 x(5)=l(5)

```

Please note that function calls have no synchronization at its entrance/exit. In the following program, the value of `x(5)` is not guaranteed to be 6:

2

3

```

XcalableMP Fortran
!$xmp nodes p(2)
!$xmp template t(10)

```



```
!$xmp distribute (block) onto p :: t
integer :: g(10), h(10)
5 !$xmp align (i) with t(i) :: g, h
integer x(5)
integer l(5), m(5)
!$xmp local_alias l => g
!$xmp local_alias m => h
10
!$xmp loop on t(i)
do i=1,10
h(i)=i
end do
15
!$xmp gmove async(1)
call sub(g,h)
l(5)=6
!$xmp wait_async(1)
20 x(5)=l(5)
```


1 Appendix F

2 DRAFT: Tasklet of upcoming 3 XcalableMP 2.0

4 This chapter shows a draft of the specification of the tasklet features from upcoming XcalableMP
5 2.0. The models and any other rules of XcalableMP 2.0 follows those of XcalableMP unless
6 specified.

7 F.1 XcalableMP Extended Execution Model

8 When a node encounters a `tasklets` construct at runtime, a set of threads are created, an
9 implicit tasklet for *structured-block* is generated on the node, and one of the threads begins
10 execution of it.

11 When a thread encounters a `tasklet` construct at runtime, a new tasklet is generated on the
12 node. Execution of generated tasklets is assigned to one of the threads on the node, subject to
13 the thread's availability to execute work. Thus, execution of the new tasklet could be immediate,
14 or deferred until later according to the tasklet scheduling constraint (Section F.3.2) and thread
15 availability.

16 At the end of the `tasklets` construct, there is an implicit `taskletwait` construct to complete
17 all of the tasklets generated in the construct, after which the threads created at the beginning
18 of the construct are terminated, and the node resumes the execution.

19 F.2 Glossary

20 F.2.1 Node Terminology

21 **node** A logical entity managed by the XcalableMP runtime system, which
22 has its own local memory and can communicate with each other, and
23 on which one or more threads can execute inside the `tasklets` region.

24 F.2.2 Thread Terminology

25 **thread** An execution entity of tasklets, which execute on a node inside a
26 `tasklets` region.

27 F.2.3 Tasklet Terminology

28 **tasklet** A specific instance of executable code and its data environment, gener-
29 ated when a thread encounters a `tasklet` construct.

tasklet		
dependence	An ordering relation between two sibling tasklets: the dependent tasklet and a previously generated predecessor tasklet. The tasklet dependence is fulfilled when the predecessor tasklet has completed.	1 2 3
dependent		
tasklet	A tasklet that because of a tasklet dependence cannot be executed until its predecessor tasklets have completed.	4 5
predecessor		
tasklet	A tasklet that must complete before its dependent tasklets can be executed.	6 7

F.3 Directives 8

F.3.1 Tasklet Constructs 9

F.3.1.1 tasklets Construct 10

Synopsis 11

The `tasklets` construct starts a region for tasklet execution. See Section ?? for a general description of the XMP's tasklet execution model. 12
13

Syntax 14

```
[F] !$xmp tasklets [on {nodes-ref | template-ref}]
    structured-block
    !$xmp end tasklets 15
```

```
[C] #pragma xmp tasklets [on {nodes-ref | template-ref}]
    structured-block
```

Description 16

When a node encounters a `tasklets` construct at runtime, an implicit tasklet for *structured-block* is created and a thread on the node begins execution of it. There is an implicit `taskletwait` construct at the end of the `tasklets` construct. 17
18
19

In a `tasklets` region, only directives listed below can be specified. 20

- `tasklet` 21
- `taskletyield` 22
- `taskletwait` 23
- `tasklet reflect` 24
- `tasklet gmove` 25
- `tasklet bcast` 26
- `tasklet reduction` 27
- `tasklet reduce_shadow` 28

1 Note that a `tasklets` region may not contain another `tasklets` construct, that is, `tasklets`
 2 cannot be nested.

3 In addition, no OpenMP directives can be specified in a `tasklets` region.

4 **Restrictions**

5 • The `tasklets` construct is global, which means that it must be executed by all nodes in
 6 the current executing node set, and each local variable referenced in the construct must
 7 have the same value.

8 • The node set specified by the `on` clause must be a subset of the executing node set.

9 **F.3.1.2 tasklet Construct**

10 **Synopsis**

11 The `tasklet` construct defines a tasklet that will be executed by a specified node set.

12 **Syntax**

```
[F] !$xmp tasklet [on {nodes-ref | template-ref}] [depend-clause ]...
    structured-block
    !$xmp end tasklet
```

```
[C] #pragma xmp tasklet [on {nodes-ref | template-ref}] [depend-clause ]...
    structured-block
```

14 where *depend-clause* is one of:

```
in ( variable [, variable]... )
out ( variable [, variable]... )
inout ( variable [, variable]... )
```

```
pro_post ( {nodes-ref | template-ref} [, tag] )
pro_wait [( {nodes-ref | template-ref} [, tag] )]
epi_post ( {nodes-ref | template-ref} [, tag] )
epi_wait [( {nodes-ref | template-ref} [, tag] )]
```

```
remote_in [( variable [, {nodes-ref | template-ref}] [, tag] )]
remote_out [( variable [, {nodes-ref | template-ref}] [, tag] )]
accept_remote_in ( variable, {nodes-ref | template-ref} [, tag] )
accept_remote_out ( variable, {nodes-ref | template-ref} [, tag] )
```

18 **Description**

19 When a thread encounters a `tasklet` construct at runtime, it generates a tasklet from the code
 20 for the associated block and put the tasklet into the tasklet pool if it is included by the node
 21 set specified by the `on` clause; otherwise, it skips the block.

22 If the `on` clause is omitted, it is assumed that the nodes that owns a variable selected by the
 23 implementation from the *variables* specified in the `out` clause, if any, is specified in it; otherwise
 24 it is assumed that the current executing node set is specified in it.

25 The tasklet may have some scheduling constraints defined by *depend-clause*. The specifica-
 26 tion of each *depend-clause* is described in the following sections.

Restrictions

- The node set specified by *nodes-ref* or *template-ref* in the **on** clause must be a subset of the parent node set.

F.3.1.3 in/out/inout Clauses**Synopsis**

The **in**, **out**, and **inout** clauses specify dependences of tasklets within a node.

Syntax

```

in ( variable [, variable]... )
out ( variable [, variable]... )
inout ( variable [, variable]... )

```

Description

The **in**, **out**, and **inout** clauses work in the same way as the **depend** clause of the **task** directive in OpenMP 4.0 or later, with regard to dependences of tasklets within a node.

The **in** clause. The generated tasklet will be a dependent tasklet of all previously generated sibling tasklets that reference at least one of the list items in an **out** or **inout** clause.

The **out** and **inout** clauses. The generated tasklet will be a dependent tasklet of all previously generated sibling tasklets that reference at least one of the list items in an **in**, **out**, or **inout** clause.

Restrictions

- *variables* used in these clauses of the same tasklet or sibling tasklets must indicate identical storage locations or disjoint storage locations.
- *variables* cannot be zero-length array sections.
- A variable that is part of another variable (such as an element of a structure) but is not an array element or an array section cannot appear in these clause.

F.3.1.4 pro_post Clause**Synopsis**

The **pro_post** clause, in combination with the **pro_wait** or **epi_wait** construct, specifies a point-to-point synchronization between tasklets.

Syntax

```

pro_post ( { nodes-ref | template-ref } [, tag] )

```

Description

This clause ensures that the predecessor tasklets have completed on the local node before tasklets that have a matching **pro_wait** clause are scheduled, or tasklets that have a matching **epi_wait** clause are completed on remote nodes.

A **pro_post** clause having the arguments of *nodes-ref/template-ref* and *tag*, of a tasklet on a node (called a *posting node*) dynamically matches at most one **pro_wait** or **epi_wait** clause

1 having the arguments of the posting node (unless omitted) and the same value as *tag* (unless
2 omitted), of a tasklet on the node specified by *nodes-ref/template-ref*.

3 Restrictions

- 4 • *tag* must be an expression of type default integer, in XcalableMP Fortran, or type `int`, in
5 XcalableMP C.

6 F.3.1.5 `pro_wait` Clause

7 Synopsis

8 The `pro_wait` clause, in combination with the `pro_post` or `epi_post` clause, specifies a point-
9 to-point synchronization between tasklets.

10 Syntax

11 `pro_wait` [({*nodes-ref* | *template-ref*} [, *tag*])]

12 Description

13 This clause prohibits the tasklet from being scheduled on the local node until tasklets that have
14 a matching `pro_post` clause are scheduled, or tasklets that have a matching `epi_post` clause
15 are completed on remote nodes.

16 A `pro_wait` clause having the arguments of *nodes-ref/template-ref* and *tag*, of a tasklet on
17 a node (called a *waiting node*) dynamically matches a `pro_post` clause having the arguments
18 of the waiting node and the same value as *tag*, of a tasklet on the node specified by *nodes-*
19 *ref/template-ref*.

20 If *tag* is omitted, then the `pro_wait` construct can match a `pro_post` or `epi_post` clause
21 having the arguments of the waiting node and any tag, of tasklet on the node specified by *nodes-*
22 *ref/template-ref*. If both *nodes-ref/template-ref* and *tag* are omitted, then the `pro_wait` clause
23 can match a `pro_post` or `epi_post` clause having the arguments of any node and any tag.

24 Restrictions

- 25 • *tag* must be an expression of type default integer, in XcalableMP Fortran, or type `int`, in
26 XcalableMP C.

27 F.3.1.6 `epi_post` Clause

28 Synopsis

29 The `epi_post` clause, in combination with the `pro_wait` or `epi_wait` construct, specifies a
30 point-to-point synchronization between tasklets.

31 Syntax

32 `epi_post` ({*nodes-ref* | *template-ref*} [, *tag*])

Description

This clause ensures that the tasklet have completed on the local node before tasklets that have a matching `pro_wait` clause are scheduled, or tasklets that have a matching `epi_wait` clause are completed on remote nodes.

A `pro_post` clause having the arguments of *nodes-ref/template-ref* and *tag*, of a tasklet on a node (called a *posting node*) dynamically matches at most one `pro_wait` or `epi_wait` clause having the arguments of the posting node (unless omitted) and the same value as *tag* (unless omitted), of a tasklet on the node specified by *nodes-ref/template-ref*.

Restrictions

- *tag* must be an expression of type default integer, in XcalableMP Fortran, or type `int`, in XcalableMP C.

F.3.1.7 epi_wait Clause**Synopsis**

The `epi_wait` clause, in combination with the `pro_post` or `epi_post` clause, specifies a point-to-point synchronization between tasklets.

Syntax

```
epi_wait [( {nodes-ref | template-ref} [, tag] )]
```

Description

This clause prohibits the tasklet from being completed on the local node until tasklets that have a matching `pro_post` clause are scheduled, or tasklets that have a matching `epi_post` clause are completed on remote nodes.

A `pro_wait` clause having the arguments of *nodes-ref/template-ref* and *tag*, of a tasklet on a node (called a *waiting node*) dynamically matches a `pro_post` clause having the arguments of the waiting node and the same value as *tag*, of a tasklet on the node specified by *nodes-ref/template-ref*.

If *tag* is omitted, then the `pro_wait` construct can match a `pro_post` or `epi_post` clause having the arguments of the waiting node and any tag, of tasklet on the node specified by *nodes-ref/template-ref*. If both *nodes-ref/template-ref* and *tag* are omitted, then the `pro_wait` clause can match a `pro_post` or `epi_post` clause having the arguments of any node and any tag.

Restrictions

- *tag* must be an expression of type default integer, in XcalableMP Fortran, or type `int`, in XcalableMP C.

F.3.1.8 remote_in Clauses**Synopsis**

The `remote_in` clause specifies a predecessor tasklet with regard to a data dependency derived from a remote-read (get) operation.

1 **Syntax**2 `remote_in [({variable | *} [, {nodes-ref | template-ref}] [, tag])]`3 **Description**4 A `remote_in` clause can be regarded as a syntactic sugar for the combination of `in`, `pro_wait`,
5 and `epi_post` clauses.6 If `*` is specified as the first argument or all of the arguments are omitted, no `in` clause is
7 implied.8 If the second argument is omitted, the implied `pro_wait` clause will match the first incoming
9 `pro_post` or `epi_post` clause from a node and the implied `epi_post` will work as if the node
10 was specified in it.11 **Restrictions**12 • *variables* used in these clauses of the same tasklet or sibling tasklets must indicate identical
13 storage locations or disjoint storage locations.14 • *variables* cannot be zero-length array sections.15 • A variable that is part of another variable (such as an element of a structure) but is not
16 an array element or an array section cannot appear in these clause.17 • *tag* must be an expression of type default integer, in XcalableMP Fortran, or type `int`, in
18 XcalableMP C.19 **F.3.1.9 accept_remote_in Clause**20 **Synopsis**21 The `accept_remote_in` clause specifies a dependent tasklet with regard to a data dependency
22 derived from a remote-read (get) operation.23 **Syntax**24 `accept_remote_in (variable, {nodes-ref | template-ref} [, tag])`25 **Description**26 A `accept_remote_in` clause can be regarded as a syntactic sugar for the combination of `in`,
27 `pro_post`, and `epi_wait` clauses.28 (Advice to implementers) The action for the implied `epi_wait` clause may be
29 deferred until a dependent tasklet of the tasklet with respect to the the implied `in`
30 clause is scheduled.31 **Restrictions**32 • *variables* used in these clauses of the same tasklet or sibling tasklets must indicate identical
33 storage locations or disjoint storage locations.34 • *variables* cannot be zero-length array sections.

- A variable that is part of another variable (such as an element of a structure) but is not an array element or an array section cannot appear in these clause. 1
2
- *tag* must be an expression of type default integer, in XcalableMP Fortran, or type `int`, in XcalableMP C. 3
4

F.3.1.10 `remote_out` Clause 5

Synopsis 6

The `remote_out` clause specifies a predecessor tasklet with regard to a data dependency derived from a remote-write (put) operation. 7
8

Syntax 9

```
remote_out [( {variable | *} [, {nodes-ref | template-ref}] [, tag] )] 10
```

Description 11

A `remote_out` clause can be regarded as a syntactic sugar for the combination of `out`, `pro_wait`, and `epi_post` clauses. 12
13

If `*` is specified as the first argument or all of the arguments are omitted, no `in` clause is implied. 14
15

If the second argument is omitted, the implied `pro_wait` clause will match the first incoming `pro_post` or `epi_post` clause from a node and the implied `epi_post` will work as if the node was specified in it. 16
17
18

Restrictions 19

- *variables* used in these clauses of the same tasklet or sibling tasklets must indicate identical storage locations or disjoint storage locations. 20
21
- *variables* cannot be zero-length array sections. 22
- A variable that is part of another variable (such as an element of a structure) but is not an array element or an array section cannot appear in these clause. 23
24
- *tag* must be an expression of type default integer, in XcalableMP Fortran, or type `int`, in XcalableMP C. 25
26

F.3.1.11 `accept_remote_out` Clause 27

Synopsis 28

The `accept_remote_out` clause specifies a dependent tasklet with regard to a data dependency derived from a remote-write (put) operation. 29
30

Syntax 31

```
accept_remote_out ( variable, {nodes-ref | template-ref} [, tag] ) 32
```

1 Description

2 A `accept_remote_out` clause can be regarded as a syntactic sugar for the combination of `out`,
3 `pro_post`, and `epi_wait` clauses.

4 (Advice to implementers) The action for the implied `pro_post` clause may be
5 advanced after the scheduling constraint enforced by the implied `out` clause is met.

6 Restrictions

- 7 • *variables* used in these clauses of the same tasklet or sibling tasklets must indicate identical
8 storage locations or disjoint storage locations.
- 9 • *variables* cannot be zero-length array sections.
- 10 • A variable that is part of another variable (such as an element of a structure) but is not
11 an array element or an array section cannot appear in these clause.
- 12 • *tag* must be an expression of type default integer, in XcalableMP Fortran, or type `int`, in
13 XcalableMP C.

14 F.3.1.12 `taskletyield` Construct

15 Synopsis

16 The `taskletyield` construct specifies that the current tasklet can be suspended in favor of
17 execution of a different tasklet.

18 Syntax

[F] `!$xmp taskletyield`

19

[C] `#pragma xmp taskletyield`

20 Description

21 The `taskletyield` construct includes an explicit task scheduling point in the current tasklet.

22 F.3.1.13 `taskletwait` Construct

23 Synopsis

24 The `taskletwait` construct specifies a wait on the completion of child tasklets of the current
25 tasklet.

26 Syntax

[F] `!$xmp taskletwait`

27

[C] `#pragma xmp taskletwait`

28 Description

29 The `taskletwait` construct includes an implicit task scheduling point in the current tasklet.
30 The current tasklet is suspended at the tasklet scheduling point until all child tasks that it
31 generated before the `taskletwait` construct complete execution.

F.3.2 Tasklet Scheduling

Whenever a thread reaches a tasklet scheduling point, the implementation may cause it to perform a tasklet switch, beginning or resuming execution of a different tasklet. Tasklet scheduling points are implied at the following locations:

- the point immediately following the generation of a tasklet
- the point of completion of a tasklet that has one or more implicit or explicit `epi_wait` clauses
- after the point of completion of a tasklet
- at a `taskletyield` construct
- at an implicit and explicit `taskletwait` construct

When a thread encounters a tasklet scheduling point, it may do either of the following, subject to the Tasklet Scheduling Constraints (below):

- begin execution of a tasklet
- resume any suspended tasklet

If more than one of the above choices is available, it is unspecified as to which will be chosen. Tasklet Scheduling Constraints are as follows:

1. A dependent tasklet shall not be scheduled until its tasklet dependences are fulfilled.
2. A tasklet that has one or more implicit or explicit `pro_wait` clauses shall not be scheduled until all of the `pro_wait` clauses are matched with `pro_post` or `epi_post` clauses.

F.3.3 Communication Tasklet Constructs

F.3.3.1 Overview

The communication tasklet constructs define inter-node interactions between tasklets on the basis of XMP's global-view communication constructs.

They are executed as a set of tasklets, each of which produces the same effect on a node as the corresponding global-view communication construct does.

(Advice to implementers) To implement the above feature of the communication tasklet constructs, they should be based on remote read (`get`) operations and point-to-point synchronizations.

F.3.3.2 `tasklet reflect` Construct

Synopsis

The `tasklet reflect` construct defines tasklets that update the shadow area of a global array like the `reflect` directive.

1 **Syntax**

```

[F]  !$xmp tasklet reflect ( array-name [, array-name]... ) █
      █ [width ( reflect-width [, reflect-width]... )] [orthogonal] █
      █ [on {nodes-ref | template-ref}]
[C]  #pragma xmp tasklet reflect ( array-name [, array-name]... ) █
      █ [width ( reflect-width [, reflect-width]... )] [orthogonal] █
      █ [on {nodes-ref | template-ref}]

```

3 **Description**

4 The `tasklet reflect` construct generates a tasklet on each of the nodes specified by the `on`
5 clause, which will produce the same effect as the `reflect` construct having the same clauses
6 does. In addition, for each of the arrays specified by the sequence of *array-names*, an `inout`
7 clause for it implicitly added to the generated tasklets.

8 Note that `tasklet reflect` is a local construct, unlike `reflect`, and therefore the node set
9 specified by the `on` clause need not include all of the nodes onto which the target arrays are
10 mapped.

11 **Restrictions**

- 12 • The reflect width of each dimension specified by the *reflect-width* must not exceed the
13 shadow width of the arrays.
- 14 • The node set specified by the `on` clause must be a subset of the executing node set.

15 **Example**16 **F.3.3.3 tasklet gmove Construct**17 **Synopsis**

18 The `tasklet gmove` construct defines tasklets that copy the variable from the right-hand side
19 (rhs) into the left-hand side (lhs) of the associated assignment statement like the `gmove` construct.

20 **Syntax**

```

[F]  !$xmp tasklet gmove [on {nodes-ref | template-ref}]
[C]  #pragma xmp tasklet gmove [on {nodes-ref | template-ref}]

```

22 **Description**

23 The `tasklet gmove` construct generates a tasklet on each of the nodes specified by the `on`
24 clause, which will produce the same effect as the `gmove` construct having the same clauses does.
25 In addition, for the variables on the right-hand and left-hand side of the associated assignment
26 statement, `in` and `out` clauses are implicitly added to the generated tasklets, respectively,

27 Note that `tasklet gmove` is a local construct, unlike `gmove`, and therefore the node set
28 specified by the `on` clause need not include all of the nodes onto which global arrays appearing
29 in the associated statement, if any, are mapped.

30 **Restrictions**

- 31 • The `tasklet gmove` construct must be followed by (i.e., associated with) a simple assign-
32 ment statement that contains neither arithmetic operations nor function calls.

- The node set specified by the `on` clause must be a subset of the executing node set.

F.3.3.4 `tasklet bcast` Construct

Synopsis

The `tasklet bcast` construct defines tasklets that perform broadcast communication from a specified node like the `bcast` construct.

Syntax

```
[F] !$xmp tasklet bcast ( variable [, variable]... ) [from nodes-ref | template-ref] ■
                               ■ [on nodes-ref] | template-ref
[C] #pragma xmp tasklet bcast ( variable [, variable]... ) [from nodes-ref | template-ref] ■
                               ■ [on nodes-ref] | template-ref
```

Description

The `tasklet bcast` construct generates a tasklet on each of the nodes specified by the `on` clause, which will produce the same effect as the `bcast` construct having the same clauses does. In addition, for each of the variables specified by the sequence of *variables*, an `inout` clause for it implicitly added to the generated tasklets.

Note that `tasklet bcast` is a local construct, unlike `bcast`.

Restrictions

- The variables specified by the sequence of *variables* must either not be aligned or must be replicated among nodes of the node set specified by the `on` clause.
- The node set specified by the `on` clause must be a subset of the executing node set.
- The source node specified by the `from` clause must belong to the node set specified by the `on` clause.
- The source node specified by the `from` clause must be one node.

F.3.3.5 `tasklet reduction` Construct

Synopsis

The `tasklet reduction` construct defines tasklets that perform a reduction operation like the `reduction` construct.

Syntax

```
[F] !$xmp tasklet reduction ( reduction-kind : variable [, variable]... ) ■
                               ■ [on node-ref] | template-ref
[C] #pragma xmp tasklet reduction ( reduction-kind : variable [, variable]... ) ■
                               ■ [on node-ref] | template-ref
```

1 Description

2 The `tasklet reduction` construct generates a tasklet on each of the nodes specified by the `on`
 3 clause, which will produce the same effect as the `reduction` construct having the same clauses
 4 does. In addition, for each of the variables specified by the sequence of *variables*, an `inout`
 5 clause for it implicitly added to the generated tasklets.

6 Note that `tasklet reduction` is a local construct, unlike `reduction`.

7 Restrictions

- 8 • The variables specified by the sequence of *variables* must either not be aligned or must be
 9 replicated among nodes of the node set specified by the `on` clause.
- 10 • The node set specified by the `on` clause must be a subset of the executing node set.

11 F.3.3.6 `tasklet reduce_shadow` Construct

12 Synopsis

13 The `tasklet reduce_shadow` construct defines tasklets that add values of shadow objects to
 14 their reflection source like the `reduce_shadow` construct.

15 Syntax

```

16 [F] !$xmp tasklet reduce_shadow ( array-name [, array-name]... ) ■
      ■ [width ( reflect-width [, reflect-width]... )] [orthogonal] ■
      ■ [on {nodes-ref | template-ref}]
[C] #pragma xmp tasklet reduce_shadow ( array-name [, array-name]... ) ■
      ■ [width ( reflect-width [, reflect-width]... )] [orthogonal] ■
      ■ [on {nodes-ref | template-ref}]
```

17 Description

18 The `tasklet reduce_shadow` construct generates a tasklet on each of the nodes specified by
 19 the `on` clause, which will produce the same effect as the `reduce_shadow` construct having the
 20 same clauses does. In addition, for each of the arrays specified by the sequence of *array-names*,
 21 an `inout` clause for it implicitly added to the generated tasklets.

22 Note that `tasklet reduce_shadow` is a local construct, unlike `reduce_shadow`, and therefore
 23 the node set specified by the `on` clause need not include all of the nodes onto which the target
 24 arrays are mapped.

25 Restrictions

- 26 • The width of each dimension specified by *reflect-width* must not exceed the shadow width
 27 of the arrays.
- 28 • The node set specified by the `on` clause must be a subset of the executing node set.

1 Appendix G

2 Sample Programs

3 Example 1

```

XcalableMP C
/*
 * A parallel explicit solver of Laplace equation in \XMP
 */
#pragma xmp nodes p(NPROCS)
5 #pragma xmp template t(1:N)
#pragma xmp distribute t(block) onto p

double u[XSIZE+2][YSIZE+2],
      uu[XSIZE+2][YSIZE+2];
10 #pragma xmp align u[i][*] to t(i)
#pragma xmp align uu[i][*] to t(i)
#pragma xmp shadow uu[1:1][0:0]

lap_main()
15 {
  int x,y,k;
  double sum;
  for(k = 0; k < NITER; k++){
    /* old <- new */
20 #pragma xmp loop on t(x)
    for(x = 1; x <= XSIZE; x++)
      for(y = 1; y <= YSIZE; y++)
        uu[x][y] = u[x][y];
    #pragma xmp reflect (uu)
25 #pragma xmp loop on t(x)
    for(x = 1; x <= XSIZE; x++)
      for(y = 1; y <= YSIZE; y++)
        u[x][y] = (uu[x-1][y] + uu[x+1][y] +
30          uu[x][y-1] + uu[x][y+1])/4.0;

    sum = 0.0;
    #pragma xmp loop on t[x] reduction(+:sum)
    for(x = 1; x <= XSIZE; x++)
35      for(y = 1; y <= YSIZE; y++)
```

```

        sum += (uu[x][y]-u[x][y]);
#pragma xmp task on p(1)
        printf("sum = %g\n",sum);
}

```

Example 2

1

```

                                XcalableMP C
/*
 * Linpack in XcalableMP (Gaussian elimination with partial pivoting)
 * 1D distribution version
 */
5 #pragma xmp nodes p(*)
#pragma xmp template t(0:LDA-1)
#pragma xmp distribute t(cyclic) onto p

double pvt_v[N]; // local
10
/* gaussian elimination with partial pivoting */
dgefa(double a[n][LDA],int lda, int n,int ipvt,int *info)
#pragma xmp align a[:] [i] with t(i)
{
15     REAL t;
     int idamax(),j,k,kp1,l,nm1,i;
     REAL x_pvt;

     nm1 = n - 1;
20     for (k = 0; k < nm1; k++) {
         kp1 = k + 1;
         /* find l = pivot index */
         l = A_idamax(k,n-k,a[k]);
         ipvt[k] = l;
25
         /* if (a[k][l] != ZERO) */
#ifdef XMP
#pragma xmp gmove
         pvt_v[k:n-k] = a[l][k:n-k];
30 #else
         for(i = k; i < n; i++) pvt_v[i] = a[i][l];
#endif

         /* interchange if necessary */
35         if (l != k){
#ifdef XMP
#pragma xmp gmove
         a[l][:] = a[k][:];
#pragma xmp gmove
40         a[k][:] = pvt_v[:];
#else
         for(i = k; i < n; i++) a[i][l] = a[i][k];
         for(i = k; i < n; i++) a[i][k] = pvt_v[i];

```

```

45 #endif
    }
    /* compute multipliers */
    t = -ONE/pvt_v[k];
    A_dscal(k+1, n-(k+1),t,a[k]);

50    /* row elimination with column indexing */
    for (j = kp1; j < n; j++) {
        t = pvt_v[j];
        A_daxpy(k+1,n-(k+1),t,a[k],a[j]);
    }
55 }
    ipvt[n-1] = n-1;
}

dgesl(double a[n][LDA],int lda,int n,int pvt[n],double b,int job)
60 #pragma xmp align a[:] [i] with t(i)
#pragma xmp align b[i] with t(i)
{
    REAL t;
    int k, kb, l, nm1;

65    nm1 = n - 1;
    /* job = 0 , solve a * x = b, first solve l*y = b */
    for (k = 0; k < nm1; k++) {
        l = ipvt[k];
70 #pragma xmp gmove
        t = b[l];
        if (l != k){
            #pragma xmp gmove
                b[l] = b[k];
75 #pragma xmp gmove
                b[k] = t;
        }
        A_daxpy(k+1,n-(k+1),t,a[k],b);
    }

80    /* now solve u*x = y */
    for (kb = 0; kb < n; kb++) {
        k = n - (kb + 1);
        #pragma xmp task on t(k)
85 {
            b[k] = b[k]/a[k][k];
            t = -b[k];
        }
        #pragma xmp bcast (t) from t(k)
90    A_daxpy(0,k,t,a[k],b);
    }
}

```

```

/*
95 * distributed array based routine
*/
A_daxpy(int b,int n,double da,double dx[n],double dy[n])
#pragma xmp align dx[i] with t(i)
#pragma xmp align dy[i] with t(i)
100 {
    int i,ix,iy,m,mpl;
    if(n <= 0) return;
    if(da == ZERO) return;
    /* code for both increments equal to 1 */
105 #pragma xmp loop on t(b+i)
    for (i = 0;i < n; i++) {
        dy[b+i] = dy[b+i] + da*dx[b+i];
    }
}
110
int A_idamax(int b,int n,double dx[n])
#pragma xmp align dx[i] with t(i)
{
    double dmax, g_dmax;
115     int i, ix, itemp;
    if(n == 1) return(0);

    /* code for increment equal to 1 */
    itemp = 0;
120     dmax = 0.0;
#pragma xmp loop on t(i) reduction(lastmax:dmax/itemp/)
    for (i = b; i < n; i++) {
        if(fabs((double)dx[i]) > dmax) {
125             itemp = i;
            dmax = fabs((double)dx[i]);
        }
    }
    return (itemp);
}
130
A_dscal(int b,int n,double da,double dx[n])
#pragma xmp align dx[i] with t(i)
#pragma xmp align dy[i] with t(i)
{
135     int i;
    if(n <= 0)return;

    /* code for increment equal to 1 */
#pragma xmp loop on t(i)
140     for (i = b; i < n; i++)
        dx[i] = da*dx[i];
}

```

Index

- `/periodic/` modifier, 49, 58
- `/unbound/` modifier, 40
- address-of operator, 17
- `align`, 29
- align dummy variable, 29
- align offset, 29
- alignment, 12
- allocation image set, 13
- `array`, 47
- array assignment in XMP/C, 16
- array intrinsic functions, 109
- array section in XMP/C, 15
- `async` clause, 57
- asynchronous communication, 13
- `barrier`, 52
- base language, 9
- base program, 9
- `bcast`, 55
- `block`, 27
- broadcast variables, 56
- built-in elemental functions, 109
- built-in functions of XMP/C, 17
- built-in transformational procedures, 109
- coarray reference, 72
- collapse, 30
- collective mode (of `gmove`), 50
- combined directive, 21
- communication, 12
- construct, 10
- current executing node set, 6, 11
- `cyclic`, 27
- data mapping, 10
- declarative directive, 10
- declarative directives, 19
- dependent tasklet, 158
- descriptor, 18
- descriptor association, 82, 85
- descriptor-of operator, 18, 93
- Directive
 - `align`, 29
 - `array`, 47
 - `async` clause, 57
 - `barrier`, 52
 - `bcast`, 55
 - `distribute`, 26
 - `gmove`, 50
 - `local_alias`, 73
 - `lock`, 78
 - `loop`, 38, 55
 - `nodes`, 22
 - `post`, 76, 77
 - `reduce_shadow`, 57
 - `reduction`, 52
 - `reflect`, 48
 - `shadow`, 31
 - `task`, 35, 37
 - `tasklet`, 159
 - `tasklets`, 158
 - `tasks`, 35, 36
 - `template`, 24
 - `unlock`, 78
 - `wait`, 76, 77
 - `wait_async`, 56
- directive, 9, 19, 158
- `distribute`, 26
- distribution, 12
- distribution format
 - `*`, 27
 - `block`, 27
 - `cyclic`, 27
 - `gblock`, 27
- entire image set, 13
- entire node array, 11
- entire node set, 6, 11
- Example
 - `align`, 31, 129
 - `array`, 47
 - array assignment in XMP/C, 17
 - array section in XMP/C, 16
 - `async`, 57

- coarray, 72
- coarray reference, 72
- distribute, 129
- dynamic allocation in XMP/C, 17
- end task, 35, 37
- end tasks, 37
- gmove, 51
- library interface, 127
- local_alias, 75
- loop, 41, 42, 129
- memory-layout, 131
- MPI interface, 120
- node reference, 24
- nodes, 23, 129
- OpenMP in XcalableMP programs, 115
- post, 77
- procedure interface, 82, 87
- reduce_shadow, 58
- reduction, 54
- reflect, 49
- shadow, 33, 49
- task, 35, 37, 42
- tasks, 37
- template, 26, 129
- template_fix, 17, 34, 76
- wait, 77
- wait_async, 57
- xmp_desc_of, 17
- xmp_malloc, 17
- executable directive, 10
- executable directives, 19
- executing image set, 13
- executing node, 11
- executing node array, 11
- executing node set, 6, 11
- full shadow, 32
- gblock, 27
- global, 10
- global actual argument, 81
- global communication constructs, 10
- global construct, 10
- global constructs, 6
- global data, 6, 12
- global dummy argument, 81
- global-view model, 10
- gmove, 50
- image, 13
- image index, 13
- image set, 13
- in mode (of gmove), 50
- Intrinsic and Library Procedures
 - xmp_all_node_num, 94
 - xmp_all_num_nodes, 94
 - xmp_array_gtol, 126
 - xmp_exit, 97
 - xmp_malloc, 17, 97
 - xmp_node_num, 95
 - xmp_num_nodes, 95
 - xmp_test_async, 97
 - xmp_wtick, 96
 - xmp_wtime, 96
 - xmpc_all_node_num, 94
- intrinsic transformational procedures, 109
- Laplace, 171
- library interface, 125
- Linpac, 172
- local, 10
- local actual argument, 81
- local alias, 13
- local data, 6, 12
- local dummy argument, 81
- local section, 12
- local-view model, 10
- local_alias, 73
- location-variable, 41
- lock, 78
- loop, 38, 55
- node, 11, 157
- node array, 11
- node number, 11
- node reference, 23, 24
- node set, 11
- nodes, 22
- out mode (of gmove), 50
- parent node set, 11
- post, 76, 77
- posting node, 76
- predecessor tasklet, 158
- procedure, 9
- procedure interface, 81
- reduce_shadow, 57
- reduction, 12
- reduction, 52
- reduction variable, 53
- reflect, 48

- reflection source, 32
- replicate, 30
- replicated data, 12
- replicated execution, 5
- Sample Program
 - Laplace, 171
 - Linpack, 172
- sequence association, 82
- shadow, 12
- shadow, 31
- shadow object, 32
- source node, 56
- structured block, 9
- synchronization, 13
- Syntax
 - accept_remote_in, 163
 - accept_remote_out, 164
 - align, 29
 - array, 47
 - array assignment in XMP/C, 16
 - array section in XMP/C, 15
 - barrier, 52
 - bcast, 56
 - coarray, 71
 - coarray reference, 72
 - directive, 19
 - distribute, 26
 - epi_wait, 161, 162
 - gmove, 50
 - in, 160
 - inout, 160
 - local_alias, 73
 - lock, 78
 - loop, 38
 - node reference, 23
 - nodes, 22
 - out, 160
 - post, 76
 - pro_post, 160
 - pro_wait, 161
 - reduce_shadow, 57
 - reduction, 53
 - reflect, 48
 - remote_in, 163
 - remote_out, 164
 - shadow, 31
 - task, 35
 - tasklet, 158, 159
 - tasklet bcast, 168
 - tasklet gmove, 167
 - tasklet reduce_shadow, 169
 - tasklet reduction, 168
 - tasklet reflect, 167
 - taskletwait, 165
 - taskletyield, 165
 - tasks, 36
 - template, 25
 - template reference, 25
 - template_fix, 33
 - wait, 77
 - wait_async, 56
- task, 7, 12
- task, 35, 37
- tasklet, 157
- tasklet, 159
- tasklet dependence, 158
- tasklets, 158
- tasks, 35, 36
- template, 10
- template, 24
- template reference, 25
- thread, 157
- unlock, 78
- variable, 12
- wait, 76, 77
- wait_async, 56
- waiting node, 77
- work mapping, 10
- work mapping constructs, 10
- XcalableMP C, 9
- XcalableMP Fortran, 9
- xmp_align_axis, 105
- xmp_align_offset, 105
- xmp_align_replicated, 106
- xmp_align_template, 106
- xmp_all_node_num, 94
- xmp_all_num_nodes, 94
- xmp_array_gtol, 126
- xmp_array_laddr, 127
- xmp_array_lbound, 108
- xmp_array_lshadow, 107
- xmp_array_lsize, 126
- xmp_array_ndims, 106
- xmp_array_ubound, 108
- xmp_array_ushadow, 107
- xmp_desc_of, 18, 93

xmp_dist_axis, 104
xmp_dist_blocksize, 103
xmp_dist_gblockmap, 103
xmp_dist_format, 102
xmp_dist_nodes, 104
xmp_exit, 97
xmp_finalize, 121
xmp_finalize_mpi, 120
xmp_gather, 111
xmp_get_mpi_comm, 119
xmp_init, 121
xmp_init_mpi, 119
xmp_malloc, 17, 97
xmp_matmul, 112
xmp_node_num, 95
xmp_nodes_size, 99
xmp_nodes_attr, 99
xmp_nodes_equiv, 100
xmp_nodes_index, 98
xmp_nodes_ndims, 98
xmp_num_nodes, 95
xmp_pack, 111
xmp_scatter, 110
xmp_sort_down, 113
xmp_sort_up, 113
xmp_template_fixed, 100
xmp_template_lbound, 101
xmp_template_ndims, 101
xmp_template_ubound, 102
xmp_test_async, 97
xmp_transpose, 112
xmp_unpack, 112
xmp_wtick, 96
xmp_wtime, 96
xmpc_all_node_num, 94