

第IV部

附属書

ここでは、High Performance Fortran 言語 バージョン 2.0 (第 I 部と第 II 部で記述) および、公認拡張仕様 (第 III 部で記述) の機能の文法と意味について記述を、参照利用のためにまとめる。ここは、厳密な意味での HPF 言語仕様書の一部ではない。

附属書 A 構文規則

この附属書は、本 High Performance Fortran 言語仕様書で定義される構文規則の一覧である。

A.2 記述法と構文

A.2.2 指示文の構文

H201 *hpf-directive-line* is *directive-origin hpf-directive*

H202 *directive-origin* is !HPF\$
or CHPF\$
or *HPF\$

H203 *hpf-directive* is *specification-directive*
or *executable-directive*

H204 *specification-directive* is *processors-directive*
or *align-directive*
or *distribute-directive*
or *inherit-directive*
or *template-directive*
or *combined-directive*
or *sequence-directive*

H205 *executable-directive* is *independent-directive*

制約: *hpf-directive-line* には、他の文があってはならない。

制約: *specification-directive* は、宣言構文 (*declaration-construct*) が現れてもよい場所のみ現れてもよい。

制約: *executable-directive* は、実行構文 (*executable-construct*) が現れてもよい場所のみ現れてもよい。

制約: *hpf-directive-line* は、プログラム単位内の周囲の形式に従って、Fortran の自由形式 (F95:3.3.1.1) または固定形式 (F95:3.3.2.1) の注釈行の規則のどちらかに従う。(F95:3.3)

H206	<i>specification-directive-extended</i>	is	<i>processors-directive</i>	1
		or	<i>subset-directive</i>	2
		or	<i>align-directive</i>	3
		or	<i>distribute-directive</i>	4
		or	<i>inherit-directive</i>	5
		or	<i>template-directive</i>	6
		or	<i>combined-directive</i>	7
		or	<i>sequence-directive</i>	8
		or	<i>dynamic-directive</i>	9
		or	<i>range-directive</i>	10
		or	<i>shadow-directive</i>	11
				12
				13
				14
H207	<i>executable-directive-extended</i>	is	<i>independent-directive</i>	15
		or	<i>realign-directive</i>	16
		or	<i>redistribute-directive</i>	17
		or	<i>on-directive</i>	18
		or	<i>resident-directive</i>	19
				20
H208	<i>executable-construct-extended</i>	is	<i>action-stmt</i>	21
		or	<i>case-construct</i>	22
		or	<i>do-construct</i>	23
		or	<i>if-construct</i>	24
		or	<i>where-construct</i>	25
		or	<i>on-construct</i>	26
		or	<i>resident-construct</i>	27
		or	<i>task-region-construct</i>	28
				29
				30
				31
A.3	データマッピング			32
				33
A.3.2	データ整列とデータ分散指示文の構文			34
H301	<i>combined-directive</i>	is	<i>combined-attribute-list</i> :: <i>combined-decl-list</i>	35
H302	<i>combined-attribute</i>	is	ALIGN <i>align-attribute-stuff</i>	36
		or	DISTRIBUTE <i>dist-attribute-stuff</i>	37
		or	INHERIT	38
		or	TEMPLATE	39
		or	PROCESSORS	40
		or	DIMENSION (<i>explicit-shape-spec-list</i>)	41
				42
H303	<i>combined-decl</i>	is	<i>hpf-entity</i> [(<i>explicit-shape-spec-list</i>)]	43
		or	<i>object-name</i>	44
				45
				46
				47
				48

1 H304 *hpf-entity* is *processors-name*
2 or *template-name*
3

4 制約: 同じ種類の *combined-attribute* を、一つの *combined-directive* 中で 2 回以上指定しては
5 ならない。
6

7 制約: DIMENSION 属性が *combined-directive* に現れた場合、それが適用されるすべての言語
8 要素は、HPF の TEMPLATE や PROCESSORS 型指定子で宣言されなければならない。
9

10 A.3.3 DISTRIBUTE 指示文

12 H305 *distribute-directive* is DISTRIBUTE *distributee dist-directive-stuff*

13 H306 *dist-directive-stuff* is *dist-format-clause* [*dist-onto-clause*]

14 H307 *dist-attribute-stuff* is *dist-directive-stuff*
15 or *dist-onto-clause*
16

17 H308 *distributee* is *object-name*
18 or *template-name*
19

20 H309 *dist-format-clause* is (*dist-format-list*)
21 or * (*dist-format-list*)
22 or *
23

24 H310 *dist-format* is BLOCK [(*scalar-int-expr*)]
25 or CYCLIC [(*scalar-int-expr*)]
26 or *
27

28 H311 *dist-onto-clause* is ONTO *dist-target*
29

30 H312 *dist-target* is *processors-name*
31 or * *processors-name*
32 or *
33

34 制約: *distributee* である実体名 (*object-name*) は、単純な名前でなければならず、部分実体特
35 定子や成分名 (*component-name*) であってはならない。
36

37 制約: *distributee* である実体名 (*object-name*) は、*alignee* であってはならない。
38

39 制約: *distributee* である実体名 (*object-name*) は、POINTER 属性をもつことはできない。
40

41 制約: *distributee* である実体名 (*object-name*) は、TARGET 属性をもつことはできない。
42

43 制約: *distributee* がスカラであるとき、*dist-format-list*(およびそれを囲む丸括弧) は現れては
44 ならない。この場合、文形式の指示文は、*dist-format-clause* が “*” である場合だけが
45 許される。
46

47 制約: *dist-format-list* が指定されたら、その長さはそれぞれの *distributee* の次元数と等しく
48 なければならない。

制約: *dist-format-list* と *dist-target* の両方が指定されたとき、“*” でない *dist-format-list* の要素の数は、そのプロセッサ構成の次元数と等しくなければならない。

制約: *dist-format-list* がなく *dist-target* が指定されたとき、それぞれの *distributee* の次元数は、そのプロセッサ構成の次元数と等しくなければならない。

制約: DISTRIBUTE 指示文で、*dist-format-clause* または *dist-target* のどちらかが “*” で始まるときは、すべての *distributee* は仮引数でなければならない。

制約: DISTRIBUTE 指示文の *dist-format* 中の整数式 (*scalar-int-expr*) は、すべて宣言式 (*specification-expr*) でなくてはならない。

A.3.4 ALIGN 指示文

H313 *align-directive* is ALIGN *alignnee align-directive-stuff*

H314 *align-directive-stuff* is (*align-source-list*) *align-with-clause*

H315 *align-attribute-stuff* is [(*align-source-list*)] *align-with-clause*

H316 *alignnee* is *object-name*

H317 *align-source* is :

or *

or *align-dummy*

H318 *align-dummy* is *scalar-int-variable*

制約: *alignnee* である実体名 (*object-name*) は、単純な名前であらなければならない、部分実体特定子や成分名 (*component-name*) であってはならない。

制約: *alignnee* である実体名 (*object-name*) は、*distributee* であってはならない。

制約: *alignnee* である実体名 (*object-name*) は、POINTER 属性をもつことはできない。

制約: *alignnee* である実体名 (*object-name*) は、TARGET 属性をもつことはできない。

制約: *alignnee* がスカラであるとき、*align-source-list* (およびそれを囲む丸括弧) は現れてはならない。この場合、文形式の指示文は許されない。

制約: *align-source-list* が指定されたとき、その長さは *alignnee* の次元数と等しくなければならない。

制約: *align-dummy* は、名前付き変数であらなければならない。

制約: 実体は、INHERIT 属性と ALIGN 属性の両方をもつことはできない。

H319 *align-with-clause* is WITH *align-spec*

H320 *align-spec* is *align-target* [(*align-subscript-list*)]

or * *align-target* [(*align-subscript-list*)]

1	H321	<i>align-target</i>	is	<i>object-name</i>
2			or	<i>template-name</i>
3				
4	H322	<i>align-subscript</i>	is	<i>int-expr</i>
5			or	<i>align-subscript-use</i>
6			or	<i>subscript-triplet</i>
7			or	*
8				
9	H323	<i>align-subscript-use</i>	is	[[<i>int-level-two-expr</i>] <i>add-op</i>]
10				<i>align-add-operand</i>
11			or	<i>align-subscript-use add-op int-add-operand</i>
12	H324	<i>align-add-operand</i>	is	[<i>int-add-operand</i> *] <i>align-primary</i>
13			or	<i>align-add-operand</i> * <i>int-mult-operand</i>
14				
15	H325	<i>align-primary</i>	is	<i>align-dummy</i>
16			or	(<i>align-subscript-use</i>)
17				
18	H326	<i>int-add-operand</i>	is	<i>add-operand</i>
19	H327	<i>int-mult-operand</i>	is	<i>mult-operand</i>
20				
21	H328	<i>int-level-two-expr</i>	is	<i>level-2-expr</i>
22				

23 制約: *align-target* である実体名 (*object-name*) は、単純な名前でなければならず、部分実体
24 特定子や成分名 (*component-name*) であってはならない。

25 制約: *align-target* は、OPTIONAL 属性をもつことはできない。

26 制約: ALIGN 指示文の *align-spec* が “*” で始まるときは、すべての *alignee* は仮引数でな
27 ければならない。

28 制約: *align-directive* 中の整数式 (*int-expr*)、*int-level-two-expr*、*int-add-operand* および *int-*
29 *mult-operand* は、宣言式でなくてはならない。

30 制約: *align-directive* 中の *align-subscript* である添字三つ組 (*subscript-triplet*) の添字 (*sub-*
31 *script*) や刻み幅 (*stride*) は、宣言式でなくてはならない。

32 制約: *align-subscript-list* では、同一の *align-dummy* は高々1回だけ現れることができる。

33 制約: *align-subscript-use* 式には、*align-dummy* は高々1回だけ現れることができる。

34 制約: *align-dummy* として使われた *scalar-int-variable* は、上記の文法により明確に許され
35 た場所以外では *align-spec* の中には現れてはならない。言い代えれば、一つの *align-*
36 *dummy* に、*align-dummy* を含まない整数宣言式を加数や乗数として加えることでの
37 み、*align-subscript-use* を構成することができる。

38 制約: *align-subscript* 中の添字 (*subscript*) には、いかなる *align-dummy* も現れてはならない。

39 制約: *int-add-operand*、*int-mult-operand* および *int-level-two-expr* は、整数型でなくてはな
40 らない。

A.3.6 PROCESSORS 指示文

H329 *processors-directive* is PROCESSORS *processors-decl-list*
H330 *processors-decl* is *processors-name*
[(*explicit-shape-spec-list*)]

A.3.7 TEMPLATE 指示文

H331 *template-directive* is TEMPLATE *template-decl-list*
H332 *template-decl* is *template-name* [(*explicit-shape-spec-list*)]

A.3.8 記憶列結合と順序結合

H333 *sequence-directive* is SEQUENCE [[::] *association-name-list*]
or NO SEQUENCE [[::] *association-name-list*]
H334 *association-name* is *object-name*
or / [*common-block-name*] /

制約: いかなる有効域においても、データ実体の名前と共通ブロック名は、*sequence-directive* の中に高々一度しか現れてはならない。

制約: *association-name-list* のない *sequence-directive* は、同じ有効域の中で一つしか許されない。

A.4 副プログラム境界でのデータマッピング

A.4.4 整列

H401 *inherit-directive* is INHERIT *inheritee-list*
H402 *inheritee* is *object-name*

制約: *inheritee* は仮引数でなければならない。

制約: *inheritee* は *alignee* であってはならない。

制約: *inheritee* は *distributee* であってはならない。

A.5 INDEPENDENT 指示文及び関連の指示文

A.5.1 INDEPENDENT 指示文

H501 *independent-directive* is INDEPENDENT [, *new-clause*]
[, *reduction-clause*]

1 H502 *new-clause* is NEW (*variable-name-list*)
2
3 H503 *reduction-clause* is REDUCTION (*reduction-variable-list*)
4
5 H504 *reduction-variable* is *array-variable-name*
6 or *scalar-variable-name*
7 or *structure-component*

8 制約: *independent-directive* の後の最初の非注釈行は、*do-stmt*、*forall-stmt*、または *forall-*
9 *construct* でなければならない。

11 制約: *independent-directive* の後の最初の非注釈行が *do-stmt* である場合、その文は *do-vari-*
12 *able* を含む *loop-control* オプションを持たなければならない。

14 制約: NEW 節または REDUCTION 節がある場合、指示文の後の最初の非注釈行は *do-stmt* でな
15 なければならない。

17 制約: NEW 節または REDUCTION 節内に指定する *variable* 及びそれらの成分及び要素は、次の
18 ものであってはならない。

- 20 • 仮引数
- 21
- 22 • SAVE 属性または TARGET 属性を持つもの
- 23
- 24 • COMMON ブロックに現われるもの
- 25
- 26 • EQUIVALENCE 文により他の実体と記憶列結合するもの
- 27
- 28 • 参照結合したもの
- 29
- 30 • 親子結合により他の有効域でアクセスされるもの

31 制約: *reduction-variable* として現われる変数は、同じ *independent-directive* の *new-clause* 中
32 に現われてはならず、*independent-directive* が適用される後続の *do-stmt*、*forall-stmt*
33 及び *forall-construct* の範囲内 (すなわち、ソース上でのループ本体部) の *new-clause*
34 及び *reduction-clause* に現われてはならない。

36 制約: *reduction-variable* 中の *structure-component* は、*subscript-section-list* を含んではなら
37 ない。

39 制約: *reduction-variable*¹として現われる変数は、組込み型でなければならない。また、
40 CHARACTER 型であってはならない。

48 ¹原文は *reduction-var*

H505	<i>reduction-stmt</i>	is <i>variable</i> = <i>variable mult-op mult-operand</i>	1
		or <i>variable</i> = <i>add-operand * variable</i>	2
		or <i>variable</i> = <i>variable add-op add-operand</i>	3
		or <i>variable</i> = <i>level-2-expr + variable</i>	4
		or <i>variable</i> = <i>variable and-op and-operand</i>	5
		or <i>variable</i> = <i>and-operand and-op variable</i>	6
		or <i>variable</i> = <i>variable or-op or-operand</i>	7
		or <i>variable</i> = <i>or-operand or-op variable</i>	8
		or <i>variable</i> = <i>variable equiv-op equiv-operand</i>	9
		or <i>variable</i> = <i>equiv-operand equiv-op variable</i>	10
		or <i>variable</i> = <i>reduction-function (variable , expr)</i> ¹²	11
		or <i>variable</i> = <i>reduction-function (expr , variable)</i> ¹³	12
			13
			14
H506	<i>reduction-function</i>	is MAX	15
		or MIN	16
		or IAND	17
		or IOR	18
		or IEOR	19
			20
			21
制約:	<i>reduction-stmt</i> 中に <i>variable</i> が 2 つ現われる場合、その 2 つは字面的に等しくなければ		22
	ならない。		23
			24
			25
A.6	外来プログラム単位		26
			27
A.6.2	外来プログラム単位の宣言		28
			29
H601	<i>function-stmt</i>	is [<i>prefix</i>] FUNCTION <i>function-name</i>	30
		([<i>dummy-arg-name-list</i>])	31
		[RESULT (<i>result-name</i>)]	32
			33
H602	<i>subroutine-stmt</i>	is [<i>prefix</i>] SUBROUTINE <i>subroutine-name</i>	34
		[([<i>dummy-arg-list</i>])]	35
			36
H603	<i>prefix</i>	is <i>prefix-spec</i> [<i>prefix-spec</i>] ...	37
			38
H604	<i>prefix-spec</i>	is <i>type-spec</i>	39
		or RECURSIVE	40
		or PURE	41
		or ELEMENTAL	42
		or <i>extrinsic-prefix</i>	43
			44
制約:	任意の HPF の外部副プログラム内で、どの内部副プログラムも、その親と同じ外来種		45
	別でなければならず、また、明示的に外来種別が与えられていない内部副プログラム		46
	に対しては、その親と同じ外来種別が仮定される。		47
			48
H605	<i>program-stmt</i>	is [<i>extrinsic-prefix</i>] PROGRAM <i>program-name</i>	49

1 H606 *module-stmt* is [*extrinsic-prefix*] MODULE *module-name*
2
3 H607 *block-data-stmt* is [*extrinsic-prefix*] BLOCK DATA
4 [*block-data-name*]

5
6 制約: 任意の HPF のモジュール内で、どのモジュール副プログラムも、その親と同じ外来種
7 別でなければならず、また明示的に外来種別が与えられていないモジュール副プログ
8 ラムに対しては、その親と同じ外来種別が仮定される。

9
10 制約: HPF の任意の主プログラムあるいはモジュール副プログラム内で、どの内部副プログ
11 ラムも、その親と同じ外来種別でなければならず、また明示的に外来種別が与えられ
12 ていない内部副プログラムに対しては、その親と同じ外来種別が仮定される。

13
14 H608 *extrinsic-prefix* is EXTRINSIC (*extrinsic-spec*)

15
16 H609 *extrinsic-spec* is *extrinsic-spec-arg-list*
17 or *extrinsic-kind-keyword*

18
19 H610 *extrinsic-spec-arg* is *language*
20 or *model*
21 or *external-name*

22
23 H611 *language* is [LANGUAGE =]
24 *scalar-char-initialization-expr*

25
26 H612 *model* is [MODEL =]
27 *scalar-char-initialization-expr*

28
29 H613 *external-name* is [EXTERNAL_NAME =]
30 *scalar-char-initialization-expr*

31 制約: *extrinsic-spec-arg-list* の中には、*language*、*model*、あるいは *external-name* の少なく
32 とも 1 つが指定されなければならない、またどれも 2 回以上指定することはできない。

33
34 制約: もし、LANGUAGE=を使わずに *language* を指定する場合、*language* は *extrinsic-spec-arg-*
35 *list* 中の最初の要素でなければならない。もし、MODEL=を使わずに *model* を指定する場
36 合、LANGUAGE=のない *language* が *extrinsic-spec-arg-list* 中の最初の要素であり、*model*
37 が 2 番目の要素でなければならない。もし、EXTERNAL_NAME=を使わずに *external-name*
38 を指定する場合、LANGUAGE=のない *language* が *extrinsic-spec-arg-list* 中の最初の要素
39 であり、MODEL=のない *model* が 2 番目の要素でなければならない。

40
41 制約: LANGUAGE=、MODEL=、EXTERNAL_NAME=を伴う形は、上で禁止された場合を除いて、い
42 かなる順番で書いても構わない。

43
44 これらの *extrinsic-spec-arg-list* に関する規則は、あたかも EXTRINSIC が、LANGUAGE、
45 MODEL、EXTERNAL_NAME という、それぞれが OPTIONAL であるような *dummy-arg-list*
46 を使用した明示的引用仕様を伴った手続であるかのようなものであることに注意され
47 たい。

48

制約: *language* の中では、*char-initialization-expr* の値は、以下のものが許される。

- 'HPF' HPF 言語を指す。もし *model* が明示的に指定されていない場合、*model* には 'GLOBAL' が暗黙に仮定される。
- 'FORTRAN' ANSI または ISO 規格の Fortran 言語を指す。もし *model* が明示的に指定されていない場合、*model* には 'SERIAL' が暗黙に仮定される。
- 'F77' 以前の ANSI または ISO 規格である FORTRAN 77 言語を指す。もし *model* が明示的に指定されていない場合、*model* には 'SERIAL' が暗黙に仮定される。
- 'C' ANSI 規格の C 言語を指す。もし *model* が明示的に指定されていない場合、*model* には 'SERIAL' が暗黙に仮定される。
- 実装依存の値。暗黙の *model* は実装に依存する。

ほとんどの実装にとって、'C' は、引用仕様本体 (*interface-body*) 中に記述された FUNCTION 文か SUBROUTINE 文でしか許されないことに注意されたい。

制約: *language* が指定されていない場合、親有効域と同じものが仮定される。

制約: *model* において、*char-initialization-expr* の値は、以下のものが許される。

- 'GLOBAL' グローバルモデルを指す。
- 'LOCAL' ローカルモデルを指す。
- 'SERIAL' シリアルモデルを指す。
- 実装依存の値。

制約: *model* が指定されず、*language* の指定から暗黙に仮定されない場合、親有効域と同じものが仮定される。

制約: 名称が HPF の 3 文字から始まる全ての *language* 及び *model* は、本仕様及びその後継仕様の現在あるいは将来における定義のために予約されている。

制約: *external-name* において、*scalar-char-initialization-expr* の値は、その用途が外来種別によって決定される文字列である。例えば、ある外来種別は、*external-name* を、その手順が C の手順から参照された場合の名前を指定するために使用するかもしれない。そのような実装では、ユーザはコンパイラに、その名前が C コンパイラに理解できるように変換を行なうことを期待するであろう。もし *external-name* が指定されていない場合、その値は実装依存となる。

H614 *extrinsic-kind-keyword* is HPF
 or HPF_LOCAL
 or HPF_SERIAL

制約: EXTRINSIC(HPF) は EXTRINSIC('HPF', 'GLOBAL') と同値である。*extrinsic-prefix* が存在しないとき、HPF コンパイラはコンパイル単位を外来種別 HPF に属するかのよう解釈する。従って、HPF コンパイラにとって、EXTRINSIC(HPF) あるいは

1 EXTRINSIC('HPF', GLOBAL') と指定するのは冗長である。しかし、このような明示
2 的な指定は、複数の外来種別をサポートしているコンパイラを使用する場合に必要と
3 なるかもしれない。

4
5 制約: EXTRINSIC(HPF_LOCAL) は EXTRINSIC('HPF', 'LOCAL') と同値である。外来種別が
6 HPF_LOCAL であるような主プログラムは、外来種別が HPF で、実行部がそのサブルー
7 チンの呼出しだけから構成された主プログラムから引数なしで呼び出される、外来種
8 別 HPF_LOCAL のサブルーチンであるかのように振る舞う。

9
10 制約: EXTRINSIC(HPF_SERIAL) は EXTRINSIC('HPF', 'SERIAL') と同値である。外来種別が
11 HPF_SERIAL であるような主プログラムは、外来種別が HPF で、実行部分はそのサブ
12 ルーチンの呼出しだけから構成された主プログラムから引数なしで呼び出される、外
13 来種別 HPF_SERIAL のサブルーチンであるかのように振る舞う。

14
15 制約: 名称が HPF の 3 文字から始まる全ての *extrinsic-kind-keyword* は、本仕様及びその後継
16 仕様の現在あるいは将来における定義のために予約されている。

18 A.8 データマッピングの公認拡張仕様

20 A.8.2 拡張データマッピング指示文の属性形式の構文

21
22 H801 *combined-attribute-extended* is ALIGN *align-attribute-stuff*
23 or DISTRIBUTE *dist-attribute-stuff*
24 or INHERIT
25 or TEMPLATE
26 or PROCESSORS
27 or DIMENSION (*explicit-shape-spec-list*)
28 or DYNAMIC
29 or RANGE *range-attr-stuff*
30 or SHADOW *shadow-attr-stuff*
31 or SUBSET
32
33
34
35

36 制約: SUBSET 属性は、プロセッサ構成にだけ適用できる。

38 A.8.3 REDISTRIBUTE 指示文

39
40 H802 *redistribute-directive* is REDISTRIBUTE *distributee dist-directive-stuff*
41 or REDISTRIBUTE *dist-attribute-stuff* ::
42 *distributee-list*
43

44 制約: REDISTRIBUTE 指示文に現れる *distributee* は、DYNAMIC 属性をもたなければならない
45 (8.5 節 参照)。

46
47 制約: REDISTRIBUTE 指示文の *distributee* は、ALIGN または REALIGN 指示文の *alignee* とし
48 て現れてはならない。

制約: REDISTRIBUTE 指示文の *dist-format-clause* や *dist-target* のどちらも、“*” で始まってはならない。

A.8.4 REALIGN 指示文

H803 *realign-directive* **is** REALIGN *alignee align-directive-stuff*
 or REALIGN *align-attribute-stuff :: alignee-list*

制約: REALIGN 指示文に現れる *alignee* は、DYNAMIC 属性をもたなければならない (8.5 節参照)。

制約: *align-with-clause* に指定された *align-target* が DYNAMIC 属性をもつ場合、*alignee* も同様に DYNAMIC 属性をもたなければならない。

制約: REALIGN 指示文の *alignee* は、DISTRIBUTE や REDISTRIBUTE 指示文の *distributee* であってはならない。

A.8.5 DYNAMIC 指示文

H804 *dynamic-directive* **is** DYNAMIC *alignee-or-distributee-list*

H805 *alignee-or-distributee* **is** *alignee*
 or *distributee*

制約: 共通ブロックの実体は、DYNAMIC と宣言することはできず、また、DYNAMIC である実体 (またはテンプレート) に整列することはできない。(ここで制約されているようなことをしたいならば、共通ブロックの代わりにモジュールを使わなければならない。)

制約: 構造型の成分は、POINTER 属性をもつ場合だけ DYNAMIC 属性をもつことができる。(詳しくは、8.9 節を参照されたい。)

制約: SAVE 属性をもつ実体は、DYNAMIC と宣言することはできず、また、DYNAMIC である実体 (またはテンプレート) に整列することはできない。

A.8.7 部分プロセッサへのマッピング

H806 *extended-dist-target* **is** *processors-name* [(*section-subscript-list*)]
 or * *processors-name* [(*section-subscript-list*)]
 or *

制約: 部分配列添字並び (*section-subscript-list*) の部分配列添字 (*section-subscript*) は、ベクトル添字 (*vector-subscript*) であってはならず、添字 (*subscript*) または添字三つ組 (*subscript-triplet*) でなければならない。

制約: 部分配列添字並び (*section-subscript-list*) では、部分配列添字 (*section-subscript*) の数は、*processor-name* の次元数と等しくなければならない。

H808	<i>alignee-extended</i>	is <i>object-name</i>	1
		or <i>component-name</i>	2
		or <i>structure-component</i>	3
			4
制約:	構造型の成分は、その型が明示的にマップされていない場合に限って明示的に整列することができる。		5
			6
			7
制約:	構造型のデータ実体は、その構造型が明示的にマップされた型でない場合に限って明示的に整列することができる。		8
			9
			10
制約:	ALIGN 指示文の <i>alignee</i> は、構造体成分 (<i>structure-component</i>) であってはならない。		11
			12
制約:	構造型定義の中に現れる ALIGN 指示文の <i>alignee</i> は、構造型の成分の成分名 (<i>component-name</i>) でなければならない。		13
			14
			15
制約:	構造型定義の中に現れる ALIGN 指示文に限って、 <i>alignee</i> を成分名 (<i>component-name</i>) とすることができる。		16
			17
			18
制約:	REALIGN 指示文の中でだけ <i>alignee</i> を構造体成分 (<i>structure-component</i>) にすることができる。このとき、右端を除くすべての部分参照はスカラ (0 次元) でなければならない。構造体成分の右端の部分参照は、DYNAMIC 属性をもっていなければならない。		19
			20
			21
			22
H809	<i>align-target-extended</i>	is <i>object-name</i>	23
		or <i>template-name</i>	24
		or <i>component-name</i>	25
		or <i>structure-component</i>	26
			27
			28
制約:	構造型定義の中に現れる ALIGN 指示文に限って、整列先を成分名 (<i>component-name</i>) とすることができる。		29
			30
			31
制約:	<i>align-target</i> が構造体成分 (<i>structure-component</i>) であるとき、その右端を除くすべての部分参照はスカラ (0 次元) でなければならない。		32
			33
			34
A.8.10 新しい分散形式			
			35
H810	<i>extended-dist-format</i>	is BLOCK [(<i>int-expr</i>)]	36
		or CYCLIC [(<i>int-expr</i>)]	37
		or GEN_BLOCK (<i>int-array</i>)	38
		or INDIRECT (<i>int-array</i>)	39
		or *	40
			41
			42
制約:	DISTRIBUTE 指示文または REDISTRIBUTE 指示文の <i>extended-dist-format</i> の中に現れる整数型配列 (<i>int-array</i>) は、整数型の一次元配列でなければならない。		43
			44
			45
制約:	DISTRIBUTE 指示文の <i>extended-dist-format</i> の中に現れる整数型配列 (<i>int-array</i>) は、制限式 (<i>restricted-expr</i>) でなければならない。		46
			47
			48

1 制約: GEN_BLOCK 分散に現れる整数型配列 (*int-array*) の大きさは、分散先のプロセッサ構成
2 の対応する次元の寸法と等しくなければならない。

3 制約: INDIRECT 分散に現れる整数型配列 (*int-array*) の大きさは、その分散が適用される *dis-*
4 *tributee* の対応する次元の寸法と等しくなければならない。
5

7 A.8.11 RANGE 指示文

8
9 H811 *range-directive* is RANGE *ranger range-attr-stuff*

10 H812 *ranger* is *object-name*
11 or *template-name*

12
13 H813 *range-attr-stuff* is *range-distribution-list*

14 H814 *range-distribution* is (*range-attr-list*)

15 H815 *range-attr* is *range-dist-format*
16 or ALL

17
18 H816 *range-dist-format* is BLOCK [()]
19 or CYCLIC [()]
20 or GEN_BLOCK
21 or INDIRECT
22 or *

23
24
25
26
27 制約: 少なくとも以下のいずれかが成り立たなければならない。

- 28 • *ranger* は DYNAMIC 属性をもつ。
- 29 • *ranger* は INHERIT 属性をもつ。
- 30 • *ranger* は DISTRIBUTE 指示文が *combined-directive* で指定され、その *dist-format-*
31 *clause* は*である。

32
33
34 制約: *range-attr-list* の長さはそれぞれ、*ranger* の次元数と等しくなければならない。

35
36 制約: *ranger* は ALIGN 指示文または REALIGN 指示文の *alignee* であってはならない。

37 38 A.8.12 SHADOW 指示文

39
40 H817 *shadow-directive* is SHADOW *shadow-target shadow-attr-stuff*

41 H818 *shadow-target* is *object-name*
42 or *component-name*

43
44 H819 *shadow-attr-stuff* is (*shadow-spec-list*)

45 H820 *shadow-spec* is *width*
46 or *low-width* : *high-width*
47
48

H821	<i>width</i>	is	<i>int-expr</i>	1
H822	<i>low-width</i>	is	<i>int-expr</i>	2
H823	<i>high-width</i>	is	<i>int-expr</i>	3
				4
				5
制約:	<i>width</i> 、 <i>low-width</i> または <i>high-width</i> として現れる整数式 (<i>int-expr</i>) は、0 以上の値をもつ宣言式 (<i>specification-expr</i>) でなければならない。			6
				7
				8

A.9 データとタスク並列に対する公認拡張

A.9.1 活動プロセッサ集合

H901	<i>subset-directive</i>	is	SUBSET <i>processors-name</i>	13
				14
				15

A.9.2 ON 指示文

H902	<i>on-directive</i>	is	ON <i>on-stuff</i>	17
H903	<i>on-stuff</i>	is	<i>home</i> [, <i>resident-clause</i>] [, <i>new-clause</i>]	18
H904	<i>on-construct</i>	is	<i>directive-origin block-on-directive</i> <i>block</i> <i>directive-origin end-on-directive</i>	19
				20
				21
				22
				23
				24
H905	<i>block-on-directive</i>	is	ON <i>on-stuff</i> BEGIN	25
H906	<i>end-on-directive</i>	is	END ON	26
H907	<i>home</i>	is	HOME (<i>variable</i>) or HOME (<i>template-elmt</i>) or (<i>processors-elmt</i>)	27
				28
				29
				30
				31
H908	<i>template-elmt</i>	is	<i>template-name</i> [(<i>section-subscript-list</i>)]	32
H909	<i>processors-elmt</i>	is	<i>processors-name</i> [(<i>section-subscript-list</i>)]	33
				34
				35

A.9.3 RESIDENT 節、指示文、構文

H910	<i>resident-clause</i>	is	RESIDENT <i>resident-stuff</i>	36
H911	<i>resident-stuff</i>	is	[(<i>res-object-list</i>)]	37
H912	<i>resident-directive</i>	is	RESIDENT <i>resident-stuff</i>	38
H913	<i>resident-construct</i>	is	<i>directive-origin block-resident-directive</i> <i>block</i> <i>directive-origin end-resident-directive</i>	39
				40
				41
				42
				43
				44
				45
				46
				47
H914	<i>block-resident-directive</i>	is	RESIDENT <i>resident-stuff</i> BEGIN	48

1 H915 *end-resident-directive* is END RESIDENT

2 H916 *res-object* is *object*

4 A.9.4 TASK_REGION 構文

6 H917 *task-region-construct* is

8 *directive-origin block-task-region-directive*
9 *block*

10 *directive-origin end-task-region-directive*

11 H918 *block-task-region-directive* is TASK_REGION

12 H919 *end-task-region-directive* is END TASK_REGION

15 A.10 非同期入出力に関する公認拡張

17 or ASYNCHRONOUS

18 or ID = *scalar-default-int-variable*

19 or ASYNCHRONOUS

21 制約: ASYNCHRONOUS 指定子又は ID=指定子のどちらか一方を指定した場合、その両方を指定
22 しなければならない。

23 制約: ASYNCHRONOUS 指定子を指定した場合、REC=指定子を指定しなければならず、かつ
24 「書式」を指定してはならず、かつ「変数群名」を指定してはならない。

25 制約: ASYNCHRONOUS 指定子を指定した場合、データ転送文中のどの式においても、関数を
26 引用してはならない。

27 or ID = *scalar-default-int-variable*

28 or PENDING = *scalar-default-logical-variable*

29 制約: INQUIRE 文中に FILE=指定子を指定した場合、ID=指定子や PENDING=指定子を指定し
30 てはならない。

31 制約: ID=指定子又は PENDING=指定子のどちらか一方を指定する場合、その両方を指定しな
32 ければならない。

39 A.10.1 WAIT 文

41 H1001 *wait-stmt* is WAIT (*wait-spec-list*)

42 H1002 *wait-spec* is UNIT = *io-unit*

43 or ID = *scalar-default-int-expr*

44 or ERR = *label*

45 or IOSTAT = *label*

制約: *wait-spec-list* には、UNIT=指定子をちょうど 1 つ、ID=指定子をちょうど 1 つと、他の各指定子を高々1 つ、指定しなければならない。

A.11 HPF 外来機能に関する公認拡張

A.11.2 外来言語との結合仕様

H1101 *type-declaration-stmt-extended* is *type-spec* [[, *attr-spec-extended*] ... ::]
entity-decl-list

H1102 *attr-spec-extended* is PARAMETER
or *access-spec*
or ALLOCATABLE
or DIMENSION (*array-spec*)
or EXTERNAL
or INTENT (*intent-spec*)
or INTRINSIC
or OPTIONAL
or POINTER
or SAVE
or TARGET
or MAP_TO (*map-to-spec*)
or LAYOUT (*layout-spec*)
or PASS_BY (*pass-by-spec*)

H1103 *map-to-spec* is *scalar-char-initialization-expr*

H1104 *layout-spec* is *scalar-char-initialization-expr*

H1105 *pass-by-spec* is *scalar-char-initialization-expr*

制約: 同じ *attr-spec-extended* を、一つの *type-declaration-stmt* 中で 2 回以上指定してはならない。

制約: 一つのデータ要素には、一つの有効域内で、どの属性も 2 回以上明示的に指定してはならない。

制約: MAP_TO 属性、LAYOUT 属性、及び PASS_BY 属性は、これらの属性が明示的に定義されている外来種別の有効域内で、仮引数に対してだけ指定できる。

附属書 B 構文記号の索引

この附属書は、構文規則で用いられる記号の索引を示す。H で始まる識別番号は、本 High Performance Fortran 言語仕様書の構文規則を表し、その規則の全体は附属書 A に示されている。R で始まる識別番号は、Fortran 言語規格 (“Fortran 95”) の構文規則を表す。

B.1 構文規則の左辺に現れる非終端記号

記号	定義箇所	参照箇所
<i>action-stmt</i>	R216	H208
<i>add-op</i>	R710	H323 H505
<i>add-operand</i>	R706	H326 H505
<i>align-add-operand</i>	H324	H323 H324
<i>align-attribute-stuff</i>	H315	H302 H801 H803
<i>align-directive</i>	H313	H204 H206
<i>align-directive-stuff</i>	H314	H313 H803
<i>align-dummy</i>	H318	H317 H325
<i>align-primary</i>	H325	H324
<i>align-source</i>	H317	H314 H315
<i>align-spec</i>	H320	H319
<i>align-subscript</i>	H322	H320
<i>align-subscript-use</i>	H323	H322 H323 H325
<i>align-target</i>	H321	H320
<i>align-target-extended</i>	H809	
<i>align-with-clause</i>	H319	H314 H315
<i>alignee</i>	H316	H313 H803 H805
<i>alignee-extended</i>	H808	
<i>alignee-or-distributec</i>	H805	H804
<i>allocate-object</i>	R625	
<i>allocate-stmt</i>	R622	
<i>and-op</i>	R720	H505
<i>and-operand</i>	R715	H505
<i>array-constructor</i>	R432	
<i>array-spec</i>	R513	H1102
<i>assignment-stmt</i>	R735	
<i>association-name</i>	H334	H333
<i>attr-spec</i>	R503	

<i>attr-spec-extended</i>	H1102	H1101				1	
<i>block</i>	R801	H904	H913	H917		2	
<i>block-data-stmt</i>	H607					3	
<i>block-on-directive</i>	H905	H904				4	
<i>block-resident-directive</i>	H914	H913				5	
<i>block-task-region-directive</i>	H918	H917				6	
<i>call-stmt</i>	R1211					7	
<i>case-construct</i>	R808	H208				8	
<i>combined-attribute</i>	H302	H301				9	
<i>combined-attribute-extended</i>	H801					10	
<i>combined-decl</i>	H303	H301				11	
<i>combined-directive</i>	H301	H204	H206			12	
<i>data-stmt</i>	R532					13	
<i>deallocate-stmt</i>	R631					14	
<i>directive-origin</i>	H202	H201	H904	H913	H917	15	
<i>dist-attribute-stuff</i>	H307	H302	H801	H802		16	
<i>dist-directive-stuff</i>	H306	H305	H307	H802		17	
<i>dist-format</i>	H310	H309				18	
<i>dist-format-clause</i>	H309	H306				19	
<i>dist-onto-clause</i>	H311	H306	H307			20	
<i>dist-target</i>	H312	H311				21	
<i>distribute-directive</i>	H305	H204	H206			22	
<i>distributtee</i>	H308	H305	H802	H805		23	
<i>distributtee-extended</i>	H807					24	
<i>do-construct</i>	R816	H208				25	
<i>dummy-arg</i>	R1223	H602				26	
<i>dynamic-directive</i>	H804	H206				27	
<i>end-function-stmt</i>	R1220					28	
<i>end-on-directive</i>	H906	H904				29	
<i>end-resident-directive</i>	H915	H913				30	
<i>end-subroutine-stmt</i>	R1224					31	
<i>end-task-region-directive</i>	H919	H917				32	
<i>entity-decl</i>	R504	H1101				33	
<i>equiv-op</i>	R722	H505				34	
<i>equiv-operand</i>	R717	H505				35	
<i>executable-construct</i>	R215					36	
<i>executable-construct-extended</i>	H208					37	
<i>executable-directive</i>	H205	H203				38	
<i>executable-directive-extended</i>	H207					39	
<i>execution-part</i>	R208					40	
<i>explicit-shape-spec</i>	R514	H302	H303	H330	H332	H801	41
<i>expr</i>	R723	H505					42

1	<i>extended-dist-format</i>	H810					
2	<i>extended-dist-target</i>	H806					
3	<i>external-name</i>	H613	H610				
4	<i>extrinsic-kind-keyword</i>	H614	H609				
5	<i>extrinsic-prefix</i>	H608	H604	H605	H606	H607	
6	<i>extrinsic-spec</i>	H609	H608				
7	<i>extrinsic-spec-arg</i>	H610	H609				
8	<i>function-reference</i>	R1210					
9	<i>function-stmt</i>	H601					
10	<i>function-subprogram</i>	R1216					
11	<i>high-width</i>	H823	H820				
12	<i>home</i>	H907	H903				
13	<i>hpf-directive</i>	H203	H201				
14	<i>hpf-directive-line</i>	H201					
15	<i>hpf-entity</i>	H304	H303				
16	<i>if-construct</i>	R802	H208				
17	<i>independent-directive</i>	H501	H205	H207			
18	<i>inherit-directive</i>	H401	H204	H206			
19	<i>inheritee</i>	H402	H401				
20	<i>input-item</i>	R914					
21	<i>int-add-operand</i>	H326	H323	H324			
22	<i>int-expr</i>	R728	H310	H322	H810	H821	H822
23	<i>int-level-two-expr</i>	H328	H323				
24	<i>int-mult-operand</i>	H327	H324				
25	<i>int-variable</i>	R607	H318				
26	<i>interface-body</i>	R1205					
27	<i>internal-subprogram-part</i>	R210					
28	<i>io-unit</i>	R901	H1002				
29	<i>kind-selector</i>	R506					
30	<i>label</i>	R313	H1002				
31	<i>language</i>	H611	H610				
32	<i>layout-spec</i>	H1104	H1102				
33	<i>level-2-expr</i>	R707	H328	H505			
34	<i>low-width</i>	H822	H820				
35	<i>map-to-spec</i>	H1103	H1102				
36	<i>mask-expr</i>	R743					
37	<i>model</i>	H612	H610				
38	<i>module-stmt</i>	H606					
39	<i>mult-op</i>	R709	H505				
40	<i>mult-operand</i>	R705	H327	H505			
41	<i>namelist-stmt</i>	R544					
42	<i>new-clause</i>	H502	H501	H903			

<i>nullify-stmt</i>	R629			1
<i>on-construct</i>	H904	H208		2
<i>on-directive</i>	H902	H207		3
<i>on-stuff</i>	H903	H902	H905	4
<i>or-op</i>	R721	H505		5
<i>or-operand</i>	R716	H505		6
<i>output-item</i>	R915			7
<i>pass-by-spec</i>	H1105	H1102		8
<i>pointer-assignment-stmt</i>	R736			9
<i>pointer-object</i>	R630			10
<i>prefix</i>	H603	H601	H602	11
<i>prefix-spec</i>	H604	H603		12
<i>processors-decl</i>	H330	H329		13
<i>processors-directive</i>	H329	H204	H206	14
<i>processors-elmnt</i>	H909	H907		15
<i>program-stmt</i>	H605			16
<i>range-attr</i>	H815	H814		17
<i>range-attr-stuff</i>	H813	H801	H811	18
<i>range-directive</i>	H811	H206		19
<i>range-dist-format</i>	H816	H815		20
<i>range-distribution</i>	H814	H813		21
<i>ranger</i>	H812	H811		22
<i>read-stmt</i>	R909			23
<i>realign-directive</i>	H803	H207		24
<i>redistribute-directive</i>	H802	H207		25
<i>reduction-clause</i>	H503	H501		26
<i>reduction-function</i>	H506	H505		27
<i>reduction-stmt</i>	H505			28
<i>reduction-variable</i>	H504	H503		29
<i>res-object</i>	H916	H911		30
<i>resident-clause</i>	H910	H903		31
<i>resident-construct</i>	H913	H208		32
<i>resident-directive</i>	H912	H207		33
<i>resident-stuff</i>	H911	H910	H912 H914	34
<i>section-subscript</i>	R618	H806	H908 H909	35
<i>sequence-directive</i>	H333	H204	H206	36
<i>shadow-attr-stuff</i>	H819	H801	H817	37
<i>shadow-directive</i>	H817	H206		38
<i>shadow-spec</i>	H820	H819		39
<i>shadow-target</i>	H818	H817		40
<i>specification-directive</i>	H204	H203		41
<i>specification-directive-extended</i>	H206			42

<i>object</i>	H916	1
<i>object-name</i>	H303 H308 H316 H321 H334 H402	2
	H807 H808 H809 H812 H818	3
<i>processors-name</i>	H304 H312 H330 H806 H901 H909	4
<i>program-name</i>	H605	5
<i>result-name</i>	H601	6
<i>subroutine-name</i>	H602	7
<i>template-name</i>	H304 H308 H321 H332 H807 H809	8
	H812 H908	9
<i>variable-name</i>	H502 H504	10

B.3 終端記号

記号	参照箇所	
!HPF\$	H202	11
(H302 H303 H309 H310 H314 H315	12
	H320 H325 H330 H332 H502 H503	13
	H505 H601 H602 H608 H801 H806	14
	H810 H814 H816 H819 H907 H908	15
	H909 H911 H1001 H1102	16
)	H302 H303 H309 H310 H314 H315	17
	H320 H325 H330 H332 H502 H503	18
	H505 H601 H602 H608 H801 H806	19
	H810 H814 H816 H819 H907 H908	20
	H909 H911 H1001 H1102	21
*	H309 H310 H312 H317 H320 H322	22
	H324 H505 H806 H810 H816	23
*HPF\$	H202	24
+	H505	25
,	H501 H505 H903 H1101	26
/	H334	27
:	H317 H820	28
::	H301 H333 H802 H803 H1101	29
=	H505 H611 H612 H613 H1002	30
ALIGN	H302 H313 H801	31
ALL	H815	32
ALLOCATABLE	H1102	33
BEGIN	H905 H914	34
BLOCK	H310 H607 H810 H816	35
CHPF\$	H202	36
CYCLIC	H310 H810 H816	37

1	DATA	H607
2	DIMENSION	H302 H801 H1102
3	DISTRIBUTE	H302 H305 H801
4	DYNAMIC	H801 H804
5	ELEMENTAL	H604
6	END	H906 H915 H919
7	ERR	H1002
8	EXTERNAL	H1102
9	EXTERNAL_NAME	H613
10	EXTRINSIC	H608
11	FUNCTION	H601
12	GEN_BLOCK	H810 H816
13	HOME	H907
14	HPF	H614
15	HPF_LOCAL	H614
16	HPF_SERIAL	H614
17	IAND	H506
18	ID	H1002
19	IEOR	H506
20	INDEPENDENT	H501
21	INDIRECT	H810 H816
22	INHERIT	H302 H401 H801
23	INTENT	H1102
24	INTRINSIC	H1102
25	IOR	H506
26	IOSTAT	H1002
27	LANGUAGE	H611
28	LAYOUT	H1102
29	MAP_TO	H1102
30	MAX	H506
31	MIN	H506
32	MODEL	H612
33	MODULE	H606
34	NEW	H502
35	NO	H333
36	ON	H902 H905 H906
37	ONTO	H311
38	OPTIONAL	H1102
39	PARAMETER	H1102
40	PASS_BY	H1102
41	POINTER	H1102
42	PROCESSORS	H302 H329 H801

PROGRAM	H605	1
PURE	H604	2
RANGE	H801 H811	3
REALIGN	H803	4
RECURSIVE	H604	5
REDISTRIBUTE	H802	6
REDUCTION	H503	7
RESIDENT	H910 H912 H914 H915	8
RESULT	H601	9
SAVE	H1102	10
SEQUENCE	H333	11
SHADOW	H801 H817	12
SUBROUTINE	H602	13
SUBSET	H801 H901	14
TARGET	H1102	15
TASK_REGION	H918 H919	16
TEMPLATE	H302 H331 H801	17
UNIT	H1002	18
WAIT	H1001	19
WITH	H319	20
		21
		22
		23
		24
		25
		26
		27
		28
		29
		30
		31
		32
		33
		34
		35
		36
		37
		38
		39
		40
		41
		42
		43
		44
		45
		46
		47
		48

附属書 C HPF 1.1 Subset

As part of the definition of the previous version of the High Performance Fortran language, HPF 1.1, a subset language was formally defined, based on the Fortran 77 language. The goal was to permit more rapid implementations of a useful subset of HPF that did not require full implementation of the new ANSI/ISO standard Fortran (“Fortran 90”).

No subset language is defined as part of the current version, HPF 2.0. This Annex is included in the HPF 2.0 language document as a convenient summary of the HPF 1.1 Subset, which has served as a minimum requirement for HPF implementations.

C.1 Fortran 90 Features in the HPF 1.1 Subset

The features of the HPF 1.1 subset languages are listed below. For reference, the section numbers from the Fortran 90 standard are given along with the related syntax rule numbers:

- All FORTRAN 77 standard conforming features, except for storage and sequence association.
- The Fortran 90 definitions of MIL-STD-1753 features:
 - DO WHILE statement (8.1.4.1.1 / R821)
 - END DO statement (8.1.4.1.1 / R825)
 - IMPLICIT NONE statement (5.3 / R540)
 - INCLUDE line (3.4)
 - scalar bit manipulation intrinsic procedures: IOR, IAND, NOT, IEXOR, ISHFT, ISHFTC, BTEST, IBSET, IBCLR, IBITS, MVBITS (13.13)
 - binary, octal and hexadecimal constants for use in DATA statements (4.3.1.1 / R407 and 5.2.9 / R533)
- Arithmetic and logical array features:
 - array sections (6.2.2.3 / R618–621)
 - * subscript triplet notation (6.2.2.3.1)
 - * vector-valued subscripts (6.2.2.3.2)
 - array constructors limited to one level of implied DO (4.5 / R431)
 - arithmetic and logical operations on whole arrays and array sections (2.4.3, 2.4.5, and 7.1)

- array assignment (2.4.5, 7.5, 7.5.1.4, and 7.5.1.5) 1
 - masked array assignment (7.5.3) 2
 - * WHERE statement (7.5.3 / R738) 3
 - * block WHERE . . . ELSEWHERE construct (7.5.3 / R739) 4
 - array-valued external functions (12.5.2.2) 5
 - automatic arrays (5.1.2.4.1) 6
 - ALLOCATABLE arrays and the ALLOCATE and DEALLOCATE statements (5.1.2.4.3, 6.3.1 / R622, and 6.3.3 / R631) 7
 - assumed-shape arrays (5.1.2.4.2 / R516) 8
- Intrinsic procedures: 9
- The list of intrinsic functions and subroutines below is a combination of (a) routines which are entirely new to Fortran and (b) routines that have always been part of Fortran, but have been extended here to new argument and result types. The new or extended definitions of these routines are part of the subset. If a FORTRAN 77 routine is not included in this list, then only the original FORTRAN 77 definition is part of the subset. 10
- For all of the intrinsics that have an optional argument DIM, only actual argument expressions for DIM that are initialization expressions are part of the subset. The intrinsics with this constraint are marked with † in the list below. 11
- the argument presence inquiry function: PRESENT (13.10.1) 12
 - all the numeric elemental functions: ABS, AIMAG, AINT, ANINT, CEILING, CMPLX, CONJG, DBLE, DIM, DPROD, FLOOR, INT, MAX, MIN, MOD, MODULO, NINT, REAL, SIGN (13.10.2) 13
 - all mathematical elemental functions: ACOS, ASIN, ATAN, ATAN2, COS, COSH, EXP, LOG, LOG10, SIN, SINH, SQRT, TAN, TANH (13.10.3) 14
 - all the bit manipulation elemental functions : BTEST, IAND, IBCLR, IBITS, IBSET, IEOR, IOR, ISHFT, ISHFTC, NOT (13.10.10) 15
 - all the vector and matrix multiply functions: DOT_PRODUCT, MATMUL (13.10.13) 16
 - all the array reduction functions: ALL†, ANY†, COUNT†, MAXVAL†, MINVAL†, PRODUCT†, SUM†(13.10.14) 17
 - all the array inquiry functions: ALLOCATED, LBOUND†, SHAPE, SIZE†, UBOUND†(13.10.15) 18
 - all the array construction functions: MERGE, PACK, SPREAD†, UNPACK (13.10.16) 19
 - the array reshape function: RESHAPE (13.10.17) 20
 - all the array manipulation functions: CSHIFT†, EOSHIFT†, TRANSPOSE (13.10.18) 21
 - all array location functions: MAXLOC†, MINLOC†(13.10.19) 22

- 1 – all intrinsic subroutines: `DATE_AND_TIME`, `MVBITS`, `RANDOM_NUMBER`, `RANDOM_SEED`,
- 2 `SYSTEM_CLOCK` (3.11)
- 3
- 4 • Declarations:
- 5 – Type declaration statements, with all forms of *type-spec* except *kind-selector*
- 6 and `TYPE`(type-name), and all forms of *attr-spec* except *access-spec*, `TARGET`, and
- 7 `POINTER`. (5.1 / R501-503, R510)
- 8
- 9 – attribute specification statements: `ALLOCATABLE`, `INTENT`, `OPTIONAL`, `PARAMETER`,
- 10 `SAVE` (5.2)
- 11
- 12 • Procedure features:
- 13
- 14 – `INTERFACE` blocks with no *generic-spec* or *module-procedure-stmt* (12.3.2.1)
- 15 – optional arguments (5.2.2)
- 16 – keyword argument passing (12.4.1 /R1212)
- 17
- 18 • Syntax improvements:
- 19
- 20 – long (31 character) names (3.2.2)
- 21 – lower case letters (3.1.7)
- 22 – use of “_” in names (3.1.3)
- 23 – “!” initiated comments, both full line and trailing (3.3.2.1)
- 24
- 25
- 26

27 C.2 HPF 1.1 Directives and Language Extensions in the HPF 1.1 Subset

28
29 The following HPF 1.1 directives and language extensions to Fortran 90 were included in
30 the HPF 1.1 Subset:

- 31
- 32 • The basic data distribution and alignment directives: `ALIGN`, `DISTRIBUTE`,
- 33 `PROCESSORS`. and `TEMPLATE`.
- 34
- 35 • The *forall-statement* (but not the *forall-construct*).
- 36
- 37 • The `INDEPENDENT` directive.
- 38
- 39 • The `SEQUENCE` and `NO SEQUENCE` directives.
- 40
- 41 • The system inquiry intrinsic functions `NUMBER_OF_PROCESSORS` and
- 42 `PROCESSORS_SHAPE`.
- 43 • The computational intrinsic functions `ILEN`, and the HPF extended Fortran intrin-
- 44 sics `MAXLOC` and `MINLOC`, with the restriction that any actual argument expression
- 45 corresponding to an optional `DIM` argument must be an initialization expression.
- 46

47 For a discussion of the rationale by which features were chosen for the HPF 1.1 Subset,
48 please consult HPF Language Specification Version 1.1.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

附属書D Previous HPFF Acknowledgments

The current HPF 2.0 document would not have been possible without the contributions of the previous series of HPFF meetings. Following are the acknowledgements for those efforts.

D.1 HPFF Acknowledgments

Technical development for HPF 1.0 was carried out by subgroups, and was reviewed by the full committee. Many people served in positions of responsibility:

- Ken Kennedy, Convener and Meeting Chair;
- Charles Koelbel, Executive Director and Head of the FORALL Subgroup;
- Mary Zosel, Head of the Fortran 90 and Storage Association Subgroup;
- Guy Steele, Head of the Data Distribution Subgroup;
- Rob Schreiber, Head of the Intrinsic Subgroup;
- Bob Knighten, Head of the Parallel I/O Subgroup;
- Marc Snir, Head of the Extrinsic Subgroup;
- Joel Williamson and Marina Chen, Heads of the Subroutine Interface Subgroup; and
- David Loveman, Editor.

Geoffrey Fox convened the first HPFF meeting with Ken Kennedy and later led a group to develop benchmarks for HPF. Clemens-August Thole organized a group in Europe and was instrumental in making this an international effort. Charles Koelbel produced detailed meeting minutes that were invaluable to subgroup heads in preparing successive revisions to the draft proposal. Guy Steele developed L^AT_EX macros for a variety of tasks, including formatting BNF grammar, Fortran code and pseudocode, and commentary material; the document would have been much less aesthetically pleasing without his efforts.

Many companies, universities, and other entities supported their employees' attendance at the HPFF meetings, both directly and indirectly. The following organizations were represented at two or more meetings by the following individuals (not including those present at the first HPFF meeting in January of 1992, for which there is no accurate attendee list): Alliant Computer Systems CorporationDavid Reese

Amoco Production Company	Jerrold Wagener, Rex Page	1
Applied Parallel Research	John Levesque, Rony Sawdayi, Gene Wagenbreth	2
Archipel	Jean-Laurent Philippe	3
CONVEX Computer Corporation	Joel Williamson	4
Cornell Theory Center	David Presberg	5
Cray Research, Inc.	Tom MacDonald, Andy Meltzer	6
Digital Equipment Corporation	David Loveman	7
Fujitsu America	Siamak Hassanzadeh, Ken Muira	8
Fujitsu Laboratories	Hidetoshi Iwashita	9
GMD-I1.T, Sankt Augustin	Clemens-August Thole	10
Hewlett Packard	Maureen Hoffert, Tin-Fook Ngai, Richard Schooler	11
IBM	Alan Adamson, Randy Scarborough, Marc Snir, Kate Stewart	12
Institute for Computer Applications in Science & Engineering ...	Piyush Mehrotra	13
Intel Supercomputer Systems Division	Bob Knighten	14
Lahey Computer	Lev Dyadkin, Richard Fuhler, Thomas Lahey, Matt Snyder	15
Lawrence Livermore National Laboratory	Mary Zosel	16
Los Alamos National Laboratory	Ralph Brickner, Margaret Simmons	17
Louisiana State University	J. Ramanujam	18
MasPar Computer Corporation	Richard Swift	19
Meiko, Inc.	James Cownie	20
nCUBE, Inc.	Barry Keane, Venkata Konda	21
Ohio State University	P. Sadayappan	22
Oregon Graduate Institute of Science and Technology	Robert Babb II	23
The Portland Group, Inc.	Vince Schuster	24
Research Institute for Advanced Computer Science	Robert Schreiber	25
Rice University	Ken Kennedy, Charles Koelbel	26
Schlumberger	Peter Highnam	27
Shell	Don Heller	28
State University of New York at Buffalo	Min-You Wu	29
SunPro and Sun Microsystems	Prakash Narayan, Douglas Walls	30
Syracuse University	Alok Choudhary, Tom Haupt	31
TNO-TU Delft	Edwin Paalvast, Henk Sips	32
Thinking Machines Corporation	Jim Bailey, Richard Shapiro, Guy Steele	33
United Technologies Corporation	Richard Shapiro	34
University of Stuttgart	Uwe Geuder, Bernhard Woerner, Roland Zink	35
University of Southampton	John Merlin	36
University of Vienna	Barbara Chapman, Hans Zima	37
Yale University	Marina Chen, Alope Majumdar	38

Many people contributed sections to the final language specification and HPF Journal of Development, including Alok Choudhary, Geoffrey Fox, Tom Haupt, Maureen Hoffert, Ken Kennedy, Robert Knighten, Charles Koelbel, David Loveman, Piyush Mehrotra, John

1 Merlin, Tin-Fook Ngai, Rex Page, Sanjay Ranka, Robert Schreiber, Richard Shapiro, Marc
2 Snir, Matt Snyder, Guy Steele, Richard Swift, Min-You Wu, and Mary Zosel. Many others
3 contributed shorter passages and examples and corrected errors.

4 Because public input was encouraged on electronic mailing lists, it is impossible to
5 identify all who contributed to discussions; the entire mailing list was over 500 names long.
6 Following are some of the active participants in the HPFF process not mentioned above:

7	N. Arunasalam	Werner Assmann	Marc Baber
8	Babak Bagheri	Vasanth Bala	Jason Behm
9	Peter Belmont	Mike Bernhardt	Keith Bierman
10	Christian Bishof	John Bolstad	William Camp
11	Duane Carbon	Richard Carpenter	Brice Cassenti
12	Doreen Cheng	Mark Christon	Fabien Coelho
13	Robert Corbett	Bill Crutchfield	J. C. Diaz
14	James Demmel	Alan Egolf	Bo Einarsson
15	Pablo Elustondo	Robert Ferrell	Rhys Francis
16	Hans-Hermann Frese	Steve Goldhaber	Brent Gorda
17	Rick Gorton	Robert Halstead	Reinhard von Hanxleden
18	Hiroki Honda	Carol Hoover	Steven Huss-Lederman
19	Ken Jacobsen	Elaine Jacobson	Behm Jason
20	Alan Karp	Ronan Keryell	Anthony Kimball
21	Ross Knippe	Bruce Knobe	David Kotz
22	Ed Krall	Tom Lake	Peter Lawrence
23	Bryan Lawver	Bruce Leasure	Stewart Levin
24	David Levine	Theodore Lewis	Woody Lichtenstein
25	Ruth Lovely	Doug MacDonald	Raymond Man
26	Stephen Mark	Philippe Marquet	Jeanne Martin
27	Oliver McBryan	Charlie McDowell	Michael Metcalf
28	Charles Mosher	Len Moss	Lenore Mullin
29	Yoichi Muraoka	Bernie Murray	Vicki Newton
30	Dale Nielsen	Kayutov Nikolay	Steve O'Neale
31	Jeff Painter	Cherri Pancake	Harvey Richardson
32	Bob Riley	Kevin Robert	Ron Schmucker
33	J.L. Schonfelder	Doug Scofield	David Serafini
34	G.M. Sigut	Anthony Skjellum	Niraj Srivastava
35	Paul St.Pierre	Nick Stanford	Mia Stephens
36	Jaspal Subhlok	Xiaobai Sun	Hanna Szoke
37	Bernard Tourancheau	Anna Tsao	Alex Vasilevsky
38	Stephen Vavasis	Arthur Veen	Brian Wake
39	Ji Wang	Karen Warren	D.C.B. Watson
40	Matthijs van Waveren	Robert Weaver	Fred Webb
41	Stephen Whitley	Michael Wolfe	Fujio Yamamoto
42	Marco Zagha		

The following organizations made the language draft available by anonymous FTP access and/or mail servers: AT&T Bell Laboratories, Cornell Theory Center, GMD-I1.T (Sankt Augustin), Oak Ridge National Laboratory, Rice University, Syracuse University, and Thinking Machines Corporation. These outlets were instrumental in distributing the document.

The High Performance Fortran Forum also received a great deal of volunteer effort in nontechnical areas. Theresa Chatman and Ann Redelfs were responsible for most of the meeting planning and organization, including the first HPFF meeting, which drew over 125 people. Shaun Bonton, Rachele Harless, Rhonda Perales, Seryu Patel, and Daniel Swint helped with many logistical details. Danny Powell spent a great deal of time handling the financial details of the project. Without these people, it is unlikely that HPF would have been completed.

HPFF operated on a very tight budget (in reality, it had no budget when the first meeting was announced). The first meeting in Houston was entirely financed from the conferences budget of the Center for Research on Parallel Computation, an NSF Science and Technology Center. DARPA and NSF have supported research at various institutions that have made a significant contribution towards the development of High Performance Fortran. Their sponsored projects at Rice, Syracuse, and Yale Universities were particularly influential in the HPFF process. Support for several European participants was provided by ESPRIT through projects P6643 (PPPE) and P6516 (PREPARE).

D.2 HPFF94 Acknowledgments

The HPF 1.1 version of the document was prepared during the HPFF94 series of meetings. A number of people shared technical responsibilities for the activities of the HPFF94 meetings:

- Ken Kennedy, Convener and Meeting Chair;
- Mary Zosel, Executive Director and head of CCI Group 2;
- Richard Shapiro, Head of CCI Group 1;
- Ian Foster, Head of Tasking Subgroup;
- Alok Choudhary, Head of Parallel I/O Subgroup;
- Chuck Koelbel, Head of Irregular Distributions Subgroup;
- Rob Schreiber, Head of Implementation Subgroup;
- Joel Saltz, Head of Benchmarks Subgroup;
- David Loveman, Editor, assisted by Chuck Koelbel, Rob Schreiber, Guy Steele, and Mary Zosel, section editors.

1 Attendance at the HPFF94 meetings included the following people from organizations
2 that were represented two or more times.

3 Don Heller Ames Laboratory
4 Jerrold Wagener Amoco Production Company
5 John Levesque Applied Parallel Research
6 Ian Foster Argonne National Laboratory
7 Terry Pratt CESDIS/NASA Goddard
8 Jim Cowie Cooperating Systems
9 Andy Meltzer, Jon Steidel Cray Research, Inc.
10 David Loveman Digital Equipment Corporation
11 Bruce Olsen Hewlett Packard
12 E. Nunohiro, Satoshi Itoh Hitachi
13 Henry Zongaro IBM
14 Piyush Mehrotra ... Institute for Computer Applications in Science & Engineering
15 Bob Knighten, Roy Touzeau Intel SSD
16 Mary Zosel, Bor Chan, Karen Warren ...Lawrence Livermore National Laboratory
17 Ralph Brickner Los Alamos National Laboratory
18 J. Ramanujam Louisiana State University
19 Paula Vaughan, Donna Reese Mississippi State University and NSF ERC
20 Shoichi Sakon, Yoshiki Seo..... NEC
21 P. Sadayappan, Chua-Huang Huang Ohio State University
22 Andrew Johnson OSF Research Institute
23 Chip Rodden, Jeff Vanderlip Pacific Sierra Research
24 Larry Meadows, Doug Miles The Portland Group, Inc.
25 Robert Schreiber Research Institute for Advanced Computer Science
26 Ken Kennedy, Charles Koelbel Rice University
27 Ira Baxter Schlumberger
28 Alok Choudhary Syracuse University
29 Guy Steele Thinking Machines Corporation, Sun Microsystems
30 Richard Shapiro Thinking Machines Corp., Silicon Graphics Inc.
31 Scott Baden, Val Donaldson University of California, San Diego
32 Robert Babb University of Denver
33 Joel Saltz, Paul Havlak University of Maryland
34 Nicole Nemer-Preece University of Missouri-Rolla
35 Hans Zima, Siegfried Benkner, Thomas Fahringer University of Vienna

36 An important activity of HPFF94 was the processing of the many items submitted for
37 comment and interpretation which led to the HPF 1.1 update of the language document.
38 A special acknowledgement goes to Henry Zongaro, IBM, for many thoughtful questions
39 exposing dark corners of language design that were previously overlooked, and to Guy
40 Steele, Thinking Machines/Sun Microsystems for his analysis of, and solutions for some of
41 the thornier issues discussed. And general thanks to the people who submitted comments
42 and interpretation requests, including:
43
44
45
46
47
48

David Loveman, Michael Hennecke, James Cownie, Adam Marshall, Stephen Ehrlich,
Mary Zosel, Matt Snyder, Larry Meadows, Dick Hendrickson, Dave Watson, John Merlin,
Vasanth Bala, Paul.Wesson, Denis.Hugli, Stanly Steinberg, Henk Sips, Henry Zongaro,
Eiji_Nunohiro, Jens Bloch Helmers, Rob Schreiber, David B. Serafini, and Allan Knies.

Other special mention goes to Chuck Koelbel at Rice University for continued maintenance of the HPFF mailing lists, to Donna Reese and staff at Mississippi State University for establishing and maintaining a WWW home-page for HPFF, and to the University of Maryland for establishing a benchmark FTP site.

Theresa Chatman and staff at Rice University were responsible for meeting planning and organization and Danny Powell continued to handle financial details of the project.

HPFF94 received direct support for research and administrative activities from grants from ARPA, DOE, and NSF.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

関連図書

- [1] American National Standards Institute, Inc., 1430 Broadway, New York, NY. *American National Standard Programming Language FORTRAN*, ANSI X3.9-1978, approved April 3, 1978.
- [2] American National Standards Institute, Inc., 1430 Broadway, New York, NY. *American National Standard for Information Systems Programming Language FORTRAN*, S8 (X3.9-198x) Revision of X3.9-1978, Draft S8, Version 104, April 1987.
- [3] High Performance Fortran Forum. *High Performance Fortran Language Specification* Scientific Programming, 2,1, 1993. Also published as: CRPC-TR92225, Center for Research on Parallel Computation, Rice University, Houston, TX, 1992 (revised May 1993). Also published as: Fortran Forum, 12,4, Dec. 1993 and 13,2, June 1994.
- [4] High Performance Fortran Forum. *High Performance Fortran Language Specification, version 1.0* CRPC-TR92225, Center for Research on Parallel Computation, Rice University, Houston, TX, 1992 (revised May 1993).
- [5] US Department of Defense. *Military Standard, MIL-STD-1753: FORTRAN, DoD Supplement to American National Standard X3.9-1978*, November 9, 1978.

附属書 E Policy and Mechanism for Recognized Extrinsic Interfaces

HPF defines certain extrinsics such as `HPF_LOCAL`, and `HPF_SERIAL` as interfaces that HPFF believes are useful to the HPF community. But there are many more such extrinsic interfaces beyond those maintained by HPFF. HPFF has adopted a policy of formally recognizing certain extrinsic interface definitions, where the interface, and its addition to the HPF document is considered to be a service to the HPF community. Examples are language bindings to HPF or library packages.

E.1 Extrinsic Policy

To be considered for HPFF recognition, a proposed extrinsic must demonstrate the following things. It should be noted, however, that meeting these criteria does not guarantee acceptance of a proposed interface by HPFF.

- conformance to HPF rules for calling extrinsics,
- significant new functionality,
- existing practice such as users, implementations, etc.,
- institutional backing with evidence of ongoing support,
- coherent documentation,
- non-proprietary interface definition, and
- copyright goes to HPFF for interface, with permission to use (royalty free).

If a proposed extrinsic is accepted by HPFF, then:

- HPFF will recognize the interface and reference it in documentation, but HPFF does not assume responsibility for the extrinsic or its interface.
- The sponsor of the extrinsic must continue to conform to the HPF interface rules for extrinsics. The interface HPFF approves must not change without HPFF approval.
- The sponsor must assume responsibility for any CCI requests concerning the extrinsic.

A list of recognized extrinsic interfaces will be included in HPF documentation, with the following guidelines:

- There should be a single page introduction to the extrinsic which contains:
 - the name of the extrinsic,
 - a brief abstract of functionality,
 - a brief and informal description of the interface,
 - information about platform and system availability, and
 - reference and contacts for formal documentation, continued responsibility, and additional information (e.g. compiler availability).
- There should be about two pages with short examples of usage.
- A short paper with the formal definition of the interface and an informal description of the functionality of the extrinsic.

E.2 Extrinsic Interface Mechanism

The HPF www-home page will have instructions for submission of an extrinsic interface. For HPFF consideration, the sponsor prepares a proposal that includes:

- a statement of what significant new functionality is provided,
- a description of existing practice,
- a statement of institutional backing with evidence of ongoing support,
- a copy of the complete documentation or a reference to an online version of the documentation,
- a draft of the text (described above) that would be included in the HPFF documentation, and
- a statement justifying the claim that the interface follows HPF conventions for calling extrinsics.

If the proposed extrinsic interface is approved by HPFF, the sponsor then submits:

- a formal statement for HPFF records that the interface definition is non-proprietary and that the copyright of the interface belongs HPFF,
- the formal contact for CCI and continued maintenance of the interface, and
- a copy of the interface documentation formatted for HPFF use, including a copy in the current document and web mark-up languages.

附属書 F HPF_CRAFT

HPF_CRAFT is a hybrid language, combining an SPMD execution model with high performing HPF features. The model combines the multi-threaded execution of HPF_LOCAL and the HPF syntax. The goal of HPF_CRAFT is to attain the potential performance of an SPMD programming model with access to HPF features and a well-defined extrinsic interface to HPF.

F.1 Introduction

HPF_CRAFT is a hybrid language, combining an SPMD execution model with high performing and portable HPF features. The model combines the multi-threaded execution of HPF_LOCAL and the HPF syntax and features. The goal of HPF_CRAFT is to attain the potential performance of an SPMD programming model with access to HPF features and a well-defined extrinsic interface to HPF. It is built on top of the HPF_LOCAL extrinsic environment.

SPMD features and a multi-threaded model allow the user to take advantage of the performance and opportunity for low level access of a more general purpose programming model. Including HPF data distribution features gives the programmer access to high performing aspects of both models, but with the added responsibility of working with a more low-level execution model. HPF_CRAFT is best suited for platforms that support one way communication features, but is consistent with HPF and easily targeted for platforms that have HPF and can support SPMD programming styles.

The HPF features included in HPF_CRAFT are a subset of the full HPF language chosen for their performance and their broad portability and ease of use. HPF_CRAFT contains additional features to support SPMD programming styles. There are some differences from HPF, however. For example, I/O causes differences; in HPF_CRAFT different processors are allowed to read from different files at the same time, in HPF the processors must all read from the same file. The differences in the models are principally caused by the multi-threaded execution model and the introduction of HPF_LOCAL data rules.

HPF_CRAFT allows for the notion of *private data*. Data defaults to a mapping in which data items are allocated so that each processor has a unique copy. The values of the individual data items and the flow of control may vary from processor to processor within HPF_CRAFT. This behavior is consistent with the behavior of HPF_LOCAL. In HPF_CRAFT a processor may be individually named and code executed based upon which processor it is executing on. HPF_CRAFT also allows for the notion of *private loops*. A private loop is executed in entirety by each processor.

The rules governing the interface to HPF_CRAFT subprograms are similar to those for the HPF_LOCAL interface. Dummy arguments use a hybrid of the interfaces between HPF and itself and that of HPF and HPF_LOCAL. Explicitly mapped dummy arguments behave just as they do in HPF, while default (private) dummy arguments use the HPF_LOCAL calling convention.

HPF_CRAFT will be initially made available on Cray MPP systems and may also be available on Cray vector architectures. Future versions of HPF_CRAFT are possible on other vendor's architectures as well.

HPF_CRAFT is being implemented for Cray Research by The Portland Group, Inc. For Cray systems, HPF_CRAFT may be obtained through the Cray Research Inc. Orderdesk,

Cray Research Inc.
orderdesk@cray.com
(612) 683-5907

Additional formal documentation, requests, and suggestions can be made to

The Portland Group
9150 SW Pioneer Ct., Suite H
Wilsonville, OR 97070
(503) 682-2806
trs@pgroup.com

F.2 Examples of Use

HPF_CRAFT is intended for use in circumstances where greater control and performance are desired for MIMD style architectures. Since data may be declared to be private, local control is made more available and since processor information is available message passing and direct memory access programming styles can be seamlessly integrated with explicitly mapped data.

The following examples show some of the capabilities of HPF_CRAFT that are different from those of HPF. Others, such as integrated message passing and synchronization primitives are not shown. Much of HPF can also be used within HPF_CRAFT.

Example 1 illustrates the difference between the default distribution for data and the distribution of mapped data.

! Example 1

```
INTEGER PRIVATE_A(100, 20), PRIVATE_B(12, 256), PRIVATE_C
INTEGER MAPPED_A(100, 20), MAPPED_B(12, 256), MAPPED_C
!HPF$ DISTRIBUTE MAPPED_A(BLOCK, BLOCK), MAPPED_B(BLOCK, *), MAPPED_C
```

1 In the above example, given 8 processors, there would be $8 * 100 * 20$ (or 16,000) elements
2 in the array `PRIVATE_A`. Each processor contains an entire array named `PRIVATE_A`. The
3 elements of `PRIVATE_A` on processor 1 cannot be referenced using implicit syntax by any
4 other processor. There are only $100 * 20$ (or 2000) elements of array `MAPPED_A`, however,
5 and these elements are distributed about the machine in a (`BLOCK, BLOCK`) fashion.

6 The difference between the `PRIVATE_A` declaration in `HPF_CRAFT` and that in `HPF`
7 is the most instructive. In `HPF_CRAFT` each processor contains one copy of the array,
8 and the values of the elements of the array may vary from processor to processor. `HPF`
9 implementations are permitted to make one copy of the array per processor the default, but
10 the values of these copies must remain coherent across all processors. In `HPF` there is no
11 way to write a conforming program in which different processors have different values for
12 the same array.
13

14 Example 2 shows the usefulness of the `ON` clause for the `INDEPENDENT` loop as well as
15 giving an example of how private data may be used.
16

17 ! Example 2
18

```
19 PRIVATE_C = 0  
20 !HPF$ INDEPENDENT (I, J) ON MAPPED_B(I, J)  
21 DO J=1,256  
22 DO I=1,12  
23 MAPPED_B(I, J) = MAPPED_B(I, J) + 5  
24 PRIVATE_C = PRIVATE_C + MAPPED_B(I, J)  
25 ENDDO  
26 ENDDO  
27  
28  
29  
30
```

31 In this example, each iteration is executed on the processor containing the data that is
32 mapped to it. The user was allowed to specify this.

33 In addition, the private variable `PRIVATE_C` is used to compute a total for each processor.
34 At the end of execution of the loop, the values of `PRIVATE_C` may be different on each
35 processor depending upon the values in the elements of the array on each processor. This
36 data may be used as is, or it can be quickly summed using a barrier or an `ATOMIC UPDATE`.
37

38 Example 3 shows the final total value being combined into the variable `MAPPED_C` whose
39 value is available to all processors.
40

41 ! Example 3
42

```
43 MAPPED_C = 0  
44 !HPF$ ATOMIC UPDATE  
45 MAPPED_C = MAPPED_C + PRIVATE_C  
46  
47  
48
```

Example 4 shows how the language allows private data to vary from processor to processor.

! Example 4

```
IF (MY_PE() .EQ. 5) THEN
  PRIVATE_C = some-big-expression
ENDIF
```

In this example, PRIVATE_C on processor 5 will have the result of *some-big-expression*. Each processor can do distinctly different work and communicate through mapped data.

The code fragment in Example 5 is from an application and shows a few features of the language.

! Example 5

```
!HPF$ GEOMETRY G(*, CYCLIC)
  REAL FX(100,100), FY(100,100), FZ(100,100)
!HPF$ DISTRIBUTE (G) :: FX,FY,FZ
  REAL FXP(100,16,100), FYP(100,16,100)
!HPF$ DISTRIBUTE FXP(*,*, BLOCK) FYP(*,*, BLOCK)
  INTEGER CELL, ATOM, MAP(1000), NACELL(1000)

!HPF$ INDEPENDENT (CELL) ON FX(1,CELL)
  DO CELL=1,100
    JCELLO = 16*(CELL-1)
    DO NABOR = 1, 13
      JCELL = MAP(JCELLO+NABOR)
      DO ATOM=1, NACELL(CELL)
        FX(ATOM, CELL) = FX(ATOM, CELL) + FXP(ATOM, NABOR, JCELL)
        FY(ATOM, CELL) = FY(ATOM, CELL) + FYP(ATOM, NABOR, JCELL)
      ENDDO
    ENDDO
  ENDDO
```

The GEOMETRY directive allows the user to generically specify a mapping and use it to apply to many arrays (they need not have the same extents.)

Example 5 has a single INDEPENDENT loop which is the outer loop. It executes 100 iterations total. Within this loop the private value of JCELLO is set for each processor (ensuring that it is a local computation everywhere.) Nested inside the INDEPENDENT loop is a private loop; this loop executes 13 times *per* processor. Inside this loop JCELL is

1 computed locally on each processor, minimizing unnecessary communication. Finally the
2 innermost loop is also private.

3 4 **F.3 External Interface** 5

6 This section describes the behavior when an HPF_CRAFT routine is called from HPF.

7 The calling convention and argument passing rules for HPF_CRAFT are a hybrid of
8 those for HPF calling HPF_LOCAL and HPF calling HPF. Explicit interfaces are required.
9 Where dummy arguments are private (default) storage, the HPF calling HPF_LOCAL con-
10 ventions are used. Where dummy arguments are explicitly mapped, the calling convention
11 matches HPF calling HPF.

12
13 There are a number of constraints on HPF_CRAFT routines that are called from HPF.
14 The following is a list of restrictions placed on HPF_CRAFT routines called from HPF:
15

- 16 • Recursive HPF_CRAFT routines cannot be called from HPF.
- 17
18 • HPF_CRAFT routines called from HPF may only enter the routine at a single place
19 (no alternate entries).
- 20
21 • An HPF_CRAFT supprogram may not be invoked directly or indirectly from within
22 the body of a FORALL construct or within the body of an INDEPENDENT DO loop that
23 is inside an HPF program.
- 24
25 • The attributes (type, kind, rank, optional, intent) of the dummy arguments in a
26 supprogram called by HPF must match the attributes of the corresponding dummy
27 arguments in the explicit interface.
- 28
29 • A dummy argument of an HPF_CRAFT supprogram called by HPF
 - 30 – must not be a procedure name.
 - 31 – must not have the POINTER attribute.
 - 32 – must not be sequential, unless it is also PE_PRIVATE.
 - 33 – must have assumed shape even when it is explicit shape in the interface.
 - 34 – must have assumed shape even when it is explicit shape in the interface.
 - 35 – must have assumed shape even when it is explicit shape in the interface.
 - 36 – if scalar, it must be mapped so that each processor has a copy of the argument.
 - 37
- 38 • The default mapping of scalar dummy arguments and of scalar function results when
39 an HPF program calls an HPF_CRAFT routine is that it is replicated on each pro-
40 cessor.
41

42 If a dummy argument of an EXTRINSIC('HPF_CRAFT') routine interface block is an
43 array and the dummy argument of the HPF_CRAFT supprogram has the default private
44 mapping, then the corresponding dummy argument in the specification of the HPF_CRAFT
45 procedure must be an array of the same rank, type, and type parameters. When the extrinsic
46 procedure is invoked, the dummy argument is associated with the local array that consists
47 of the subgrid of the global array that is stored locally.
48

If the dummy argument of the HPF_CRAFT supprogram is explicitly mapped, it must have the same mapping as the dummy argument of the EXTRINSIC('HPF_CRAFT') supprogram. Note that this restriction does not require actual and dummy arguments to match and is no more stringent than saying that mappings of dummy arguments in interface blocks must match those in the actual routine.

F.4 Execution Model

HPF_CRAFT is built upon the fundamental execution model of HPF_LOCAL, augmented with data mapping and work distribution features from HPF. It is also augmented with explicit low-level control features, many taken from Cray Research's CRAFT language.

In HPF_CRAFT there is a single task on each processor and all tasks begin executing in parallel, with data defaulting to a private distribution, the same default distribution used in HPF_LOCAL. Each processor gets a copy of the data storage unless specified otherwise by the user. Consequently I/O works identically to I/O in HPF_LOCAL and message passing libraries are easily integrated.

Simply stated, the execution model is that of HPF_LOCAL.

To provide correct behavior when explicitly mapped data is involved, this model defines implicit barrier points at which the execution model requires that all processors must stop and wait for the execution of all other processors before continuing. These barriers add additional semantics to the HPF_LOCAL behavior. An implementation may remove any of these barriers that are deemed unnecessary, but *every* processor must participate in the barriers at each one of these points.

The points where there are implicit barriers are conceptually after those instances in which the processors in the HPF_CRAFT program are executing cooperatively, as if in an HPF program (e.g., after an INDEPENDENT loop). An HPF_CRAFT program treats operations on explicitly mapped objects as if they were operations in an HPF program and it treats operations on private data as if they were executed within the HPF_LOCAL framework. It is occasionally useful for an advanced programmer to indicate to the compilation system where barriers are not needed; HPF_CRAFT has syntax to allow this capability.

F.5 HPF_CRAFT Functional Summary

HPF_CRAFT contains a number of features not available in HPF, and restricts the usage of many of the features currently available. The following is a concise list of the differences.

- INDEPENDENT has been extended to better support an ON clause.
- There are new rules defining the interaction of explicitly mapped and private data.
- Parallel inquiry intrinsics IN_PARALLEL() and IN_INDEPENDENT() have been added.
- Serial regions (MASTER / END MASTER) have been added.
- Explicit synchronization primitives are provided.

- The `ATOMIC UPDATE`, `SYMMETRIC`, and `GEOMETRY` directives have been added.
- Many other compiler information directives have been added to assist the compiler in producing good quality code.

F.5.1 Data Mapping Features

Data mapping features provided are those that have been found useful most often. When data is explicitly mapped, only one copy of the data storage is created unless the explicit mapping directs otherwise. The value of explicitly mapped replicated data items must be consistent between processors as is the case in HPF. Storage and sequence association for explicitly mapped arrays is not guaranteed in HPF_CRAFT. For private data, storage and sequence association follows the Fortran 90 rules.

A new directive is included for completeness: `PE_PRIVATE`, which specifies that the data should conform to the default behavior. The values of private variables may vary on different processors.

F.5.2 Subprogram Interfaces

The behavior and requirements of an HPF_CRAFT program at subprogram interfaces may be divided into three cases. Each case is also available using some combination of HPF and HPF_LOCAL. For dummy arguments that are explicitly mapped, the behavior is identical to that of HPF. All processors must cooperate in a subprogram invocation that remaps or explicitly maps data. In other words, if an explicit interface is required (by the HPF rules) or the subprogram declares explicitly mapped data, the subprogram must be called on all processors. Processors need not cooperate if there are only reads to non-local data. The `INHERIT` attribute may only be applied to explicitly mapped data.

Data that has the default private mapping (case two) the behavior of an HPF_CRAFT subprogram at subprogram interfaces is identical to that of HPF_LOCAL. Data is passed individually on every processor and the processors need not interact in any way.

When a subprogram is passed actual arguments that are a combination of both explicitly mapped data and private data, the explicitly mapped data follows the HPF rules and the private data follows the HPF_LOCAL rules.

In case three, the user has the option of passing data with explicitly mapped actual arguments to dummy arguments that are not explicitly mapped (i.e., private.) The mapping rules for this data are identical to the mapping rules when HPF calls an HPF_LOCAL subprogram. The data remains “in-place.” All HPF arrays are logically carved up into pieces; the HPF_CRAFT procedure executing on a particular physical processor sees an array containing just those elements of the global array that are mapped to that physical processor. There is implicit barrier synchronization after an `INDEPENDENT` loop. Transfer of control into or out of an `INDEPENDENT` loop is prohibited.

Finally, it is undefined behavior when an actual argument is private and the dummy argument is explicitly mapped. A definition could be supplied for this interaction, but

it is the same solution that one might propose for a calling sequence when HPF_LOCAL subprograms call HPF subprograms.

F.5.3 The INDEPENDENT Directive

The INDEPENDENT directive is part of HPF_CRAFT with the same semantics as in HPF. However, within INDEPENDENT loops the values of private data may vary from processor to processor. INDEPENDENT applied to FORALL has identical syntax and semantics as in HPF.

An HPF independent loop optionally may have a NEW clause. The NEW clause is not required by HPF_CRAFT for default (not explicitly mapped) data. In HPF_CRAFT data defaults to private so values may differ from processor to processor.

Private data has slightly different behavior than data specified in the NEW clause. The value of a private datum on each processor can be used beyond a single iteration of the loop. Private data may be used to compute local sums, for example. The values of data items named in a NEW clause may not be used beyond a single iteration. The NEW clause asserts that the INDEPENDENT directive would be valid if new objects were created for the variables named in the clause for each iteration of the loop. The semantics of the NEW clause are identical in HPF_CRAFT and HPF.

The semantics of an INDEPENDENT applied to loops containing private data references changes with respect to the private data. The change can be summarized to say that instead of indicating that iterations have no dependencies upon one-another, with respect to the private data, iterations on different processors have no dependencies upon one-another.

F.5.4 The ON Clause

In addition to the version of INDEPENDENT available from HPF, a new version of INDEPENDENT is included that incorporates the ON clause. There are a number of differences between the versions of INDEPENDENT with and without the ON clause.

The new version of the INDEPENDENT directive may be applied to the first of a group of tightly nested loops and may apply to more than one of them. This more easily facilitates the use of the ON clause. The current INDEPENDENT directive applies only to a single loop nest. The INDEPENDENT directive is extended so that multiple loop nests can be named. The general syntax for these new independent loops is as follows:

```
!HPF$ INDEPENDENT (I1,I2,... ,In) ON array-name(h1(I1),h2(I2),... ,hn(In))
  DO I1 = L1, U1, S1
    DO I2 = L2, U2, S2
      ...
      DO In = Ln, Un, Sn
        ...
      END DO
    ...
  END DO
END DO
```

1 The syntax and semantics of `INDEPENDENT` with the `ON` clause are different from its
2 syntax and semantics without the `ON` clause. With the `ON` clause the directive states that
3 there are no cross-processor dependencies, but there may be dependencies between iterations
4 on a processor. There is an implicit barrier synchronization after an `INDEPENDENT` loop.
5 Transfer of control into or out of an `INDEPENDENT` loop is prohibited.

6 The iteration space of an `INDEPENDENT` nest must be rectangular. That is, the lower
7 loop bound, the upper loop bound, and the step expression for each loop indicated by the
8 `INDEPENDENT` induction list must be invariant with regard to the `INDEPENDENT` nest. Each
9 index expression of *array-name* in the `ON` clause (the functions h_i above,) must be one of
10 the following two forms:
11

12 $[a * \text{loop_control_variable} +] b$
13 $[a * \text{loop_control_variable} -] b$
14

15 where a and b must be integer values; they can be expressions, constants, or variables. The
16 values of a and b must be invariant with regard to the `INDEPENDENT` loop nest. For example,
17 specifying `A(I,J,K)` is valid. Specifying `A(3,I+J,K)` is not valid. Specifying `A(I,I,K)` is
18 not valid because `I` appears twice. Division is prohibited in any index expression of the `ON`
19 clause.
20

21 F.5.5 Array Syntax

22 Array syntax is treated identically in `HPF_CRAFT` as in `HPF` for explicitly mapped objects.
23 For private objects the behavior is identical to that of `HPF_LOCAL`. When private objects
24 and explicitly mapped objects are combined the rules are as follows:
25

26 $result = rhs_1 \text{ op}_1 rhs_2 \text{ op}_2 \dots \text{ op}_m rhs_n$
27

- 28 • If *result* is explicitly mapped and all *rhs* arrays are explicitly mapped, the work is
29 distributed as in `HPF`.
- 30 • If *result* is private and all *rhs* arrays are private the computation is done on all pro-
31 cessors as an `HPF_LOCAL` program would do it.
- 32 • If *result* is private and all *rhs* arrays are explicitly mapped, the work is distributed as
33 in `HPF` and the values of the results are broadcast to the *result* on each processor.
- 34 • If *result* is explicitly mapped and *not* all *rhs* arrays are explicitly mapped, the results
35 of the operation are undefined, unless all corresponding elements of all private *rhs*
36 arrays have the same values.
- 37 • If *result* is private and some, but not all *rhs* arrays are explicitly mapped, the value
38 is computed on each processor and saved to the local *result*.

39 All processors must participate in any array syntax statement in which the value of an
40 explicitly mapped array is modified, and there is implicit barrier synchronization after the
41 statement executes.
42
43
44
45
46
47
48

F.5.6 Treatment of FORALL and WHERE Statements

The FORALL and WHERE statements are treated exactly as in HPF when data is explicitly mapped. When private data is modified, the statement is executed separately on each processor. Finally, when data in a FORALL or WHERE are mixed, the rules for array syntax apply. If any explicitly mapped data item is modified in a *forall-stmt* or *where-stmt* then arrays in the *forall-header* or *where-header* must be explicitly mapped. In a FORALL construct, if any explicitly mapped array is modified, all modified arrays must be explicitly mapped. There is an implicit barrier synchronization after FORALL and WHERE statements if any arrays in the *forall-header* or *where-header* are explicitly mapped.

F.5.7 Synchronization Primitives

A number of synchronization primitives are provided. These primitives include:

- Barriers (test, set, wait)
- Locks (test, set, clear)
- Critical Sections
- Events (test, set, wait, clear)

Barriers provides an explicit mechanism for a task to indicate its arrival at a program point and to wait there until all other tasks arrive. A task may test and optionally wait at an explicit barrier point. In the following example, a barrier is used to make sure that *block3* is not entered by any task until all tasks have completed execution of *block1*.

```
block1
CALL SET_BARRIER()
block2
CALL WAIT_BARRIER()
block3
```

The following example performs a similar function as above. However, while waiting for all tasks to arrive at the barrier, the early tasks perform work within a loop.

```
block1
CALL SET_BARRIER()
DO WHILE (.NOT. TEST_BARRIER())
    block2
END DO
block3
```

Locks are used to prevent the simultaneous access of data by multiple tasks.

The SET_LOCK(*lock*) intrinsic sets the mapped integer variable *lock* atomically. If the lock is already set, the task that called SET_LOCK is suspended until the lock is cleared by another task and then sets it. Individual locks may be tested or cleared using *result* = TEST_LOCK(*lock*) and CLEAR_LOCK(*lock*) respectively.

1 A *critical section* protects access to a section of code rather than to a data object. The
2 **CRITICAL** directive marks the beginning of a code region in which only one task can enter
3 at a time. The **END CRITICAL** directive marks the end of the critical section. Transfer of
4 control into or out of a critical section is prohibited.
5

```
6   !HPF$ CRITICAL  
7       GLOBAL_SUM = GLOBAL_SUM + LOCAL_SUM  
8   !HPF$ END CRITICAL  
9
```

10
11
12 Events are typically used to record the state of a program's execution and to commu-
13 nicate that state to another task. Because they do not set locks, as do the lock routines
14 described earlier, they cannot easily be used to enforce serial access of data. They are suited
15 to work such as signalling other tasks when a certain value has been located in a search
16 procedure. There are four routines needed to perform the event functions, and each requires
17 a mapped argument.

18 The **SET_EVENT**(*event*) routine sets or *posts* an event; it declares that an action has
19 been accomplished or a certain point in the program has been reached. A task can post
20 an event at any time, whether the state of the event is cleared or already posted. The
21 **CLEAR_EVENT**(*event*) routine clears an event, the **WAIT_EVENT**(*event*) routine waits until a
22 particular event is posted, and the *result* = **TEST_EVENT**(*event*) function returns a logical
23 value indicating whether a particular event has been posted.
24
25

26 **F.5.8 Barrier Removal**

27
28 You can explicitly remove an implicit barrier after any **INDEPENDENT** loop, or after any
29 array syntax statement that modifies explicitly mapped arrays, by using the **NO BARRIER**
30 directive.
31

```
32   !HPF$ NO BARRIER  
33
```

34 **F.5.9 Serial Regions**

35
36 It is often useful to enter a region where only one task is executing. This is particularly
37 useful for certain types of I/O. To facilitate this, two directives are provided. In addition,
38 one may optionally attach a **COPY** clause to the **END MASTER** directive which specifies the
39 private data items whose values should be broadcast to all processors. The syntax of this
40 directive is:
41

```
42   !HPF$ MASTER  
43       sequential region  
44       ...  
45   !HPF$ END MASTER [, COPY( var1 [, var2, ..., varn ])]  
46  
47  
48
```

where *var* is **SYMMETRIC** private data to be copied to the same named private data on other processors.

If a routine is called within a serial region, the routine executes serially; there is no way to get back to parallel execution within the routine. All explicitly mapped data is accessible from within routines called in a serial region, but a routine called from within a serial region cannot allocate explicitly mapped data or remap data. All processors must participate in the invocation of the serial region. Transfer of control into or out of a serial region is not permitted.

F.5.10 Libraries

The HPF Local Routine Library is available in **HPF_CRAFT**. The **HPF_LOCAL** extrinsic environment contains a number of libraries that are useful for local SPMD programming and a number of libraries that allow the user to determine global (rather than local) state information. These library procedures take as input the name of a dummy argument and return information on the corresponding global HPF actual argument. They may only be invoked by an **HPF_CRAFT** procedure that was directly invoked by global HPF code. They may be called only for private data. The libraries reside in a module called **HPF_LOCAL_LIBRARY**.

The HPF Library is available to **HPF_CRAFT** when called with data that is explicitly mapped and all processors are participating in the call. In addition, as in **HPF_LOCAL**, the entire HPF Library is available for use with private data. Mixing private and explicitly mapped data in calls to the HPF library produces undefined behavior.

F.5.11 Parallel Inquiry Intrinsics

These intrinsic functions are provided as an extension to HPF. They return a logical value that provides information to the programmer about the state of execution in a program.

IN_PARALLEL()
IN_INDEPENDENT()

F.5.12 Task Identity

MY_PE() may be used to return the local processor number. The physical processors are identified by an integer in the range of 0 to $n-1$ where n is the value returned by the global **HPF_LIBRARY** function **NUMBER_OF_PROCESSORS**. Processor identifiers are returned by **ABSTRACT_TO_PHYSICAL**, which establishes the one-to-one correspondence between the abstract processors of an HPF processors arrangement and the physical processors. Also, the local library function **MY_PROCESSOR** returns the identifier of the task executing the call.

F.5.13 Parallelism Specification Directives

These directives allow a user to assert that a routine will only be called from within a parallel region, a serial region, or from within both regions. Without these directives an

1 implementation might be required to generate two versions of code for each routine, de-
2 pending upon implementation strategies. The directives simply make the generated code
3 size smaller and remove a test.

```
4 !HPF$ PARALLEL_ONLY  
5 !HPF$ SERIAL_ONLY  
6 !HPF$ PARALLEL_AND_SERIAL  
7
```

8
9
10 The default is PARALLEL_ONLY.

11 12 **F.5.14 The SYMMETRIC Directive**

13 SYMMETRIC variables are private data that are guaranteed to be at the same storage location
14 on every processor. The feature is beneficial to implementations that provide one-way com-
15 munication functionality. One task can either *get* or *put* data into another task's symmetric
16 data location, without involving the other task. There is an implicit barrier synchronization
17 after SYMMETRIC data is allocated.

```
18  
19  
20 REAL PRIV1(100), PRIV2  
21 !HPF$ SYMMETRIC PRIV1, PRIV2  
22
```

23 24 **F.5.15 The RESIDENT Directive**

25 The RESIDENT directive can be specified at the loop level and at the routine level. It is
26 an assertion that the references to particular variables in the routine (or loop) are only
27 references to data that are local to the task making the assertion. In the following loop, all
28 references to arrays A, B, and C are local to the task executing each iteration.

```
29  
30 REAL A(100), B(100), C(100)  
31 INTEGER IX(100)  
32 !HPF$ DISTRIBUTE A(BLOCK), B(BLOCK), C(BLOCK)  
33 !HPF$ RESIDENT A  
34  
35  
36 ...  
37 !HPF$ INDEPENDENT (I) ON B(I) RESIDENT(C)  
38 DO I = 1, 100  
39 A(IX(I)) = B(I) + C(IX(I))  
40 END DO  
41
```

42 43 **F.5.16 The ATOMIC UPDATE Directive**

44 In HPF_CRAFT, the ATOMIC UPDATE directive tells the compiler that a particular data item
45 or the elements of a particular array for a specified operation must be updated atomically.
46 This can be used within loops or in array syntax and applies to both the elements of an
47 array with an assignment of a permutation and the elements of an array within a loop.
48

In the following example, all references to `R(IX(I))` occur atomically, thus eliminating the possibility that different iterations might try to modify the same element concurrently.

```
REAL R(200), S(1000)
INTEGER IX(1000)
!HPF$ DISTRIBUTE R(BLOCK), S(BLOCK), IX(BLOCK)

!HPF$ INDEPENDENT (I) ON S(I)
DO I = 1, 1000
!HPF$     ATOMIC UPDATE
           R(IX(I)) = R(IX(I)) + S(I)
END DO
```

F.5.17 The GEOMETRY Directive

The `GEOMETRY` directive is similar to a `typedef` in C, only it is for data mapping. It allows the user to conveniently change the mappings of many arrays at the same time. It is similar in many ways to the `TEMPLATE` directive, but since it is bound to no particular extent it is sometimes easier to apply.

```
!HPF$ GEOMETRY geom( $d_1$  [,  $d_2$ , ...,  $d_n$ ])
!HPF$ DISTRIBUTE ( geom ) [::] var1[, var2, ..., varm]
```

Where d_i indicates one of the allowable distribution formats.

```
!HPF$ GEOMETRY GBB(BLOCK, CYCLIC)
REAL A(300,300), B(400,400)
!HPF$ DISTRIBUTE (GBB) :: A, B
!   if GBB changes then both A and B change
```


附属書 G The FORTRAN 77 Local Library

The HPF standard now describes an `EXTRINSIC(LANGUAGE='F77',MODEL='LOCAL')` interface, or `EXTRINSIC(F77_LOCAL)` to use the keyword identification (see Section 11.6 for its description), similar in characteristics to the `EXTRINSIC(LANGUAGE='HPF',MODEL='LOCAL')` and `EXTRINSIC(LANGUAGE='FORTRAN',MODEL='LOCAL')` interfaces. This section describes a set of library routines to make it easier to make use of the `F77_LOCAL` interface when passing distributed array data. These library routines can facilitate, for example, a portable blend of global data parallel code with preexisting FORTRAN 77-based code using explicit message passing calls for interprocessor communication. The FORTRAN 77 Local Library interface described in this section was originally developed as part of Thinking Machines TMHPF and is now supported by Sun Microsystems Inc. For suggestions, requests, or corrections concerning this interface, please contact

Sun Microsystems Inc.
High Performance Computing
M/S UCHLO5-104
5 Omni Way
Chelmsford, MA 01824
f77-local-library@sun.com

G.1 Introduction

The basic constraints for the local model (Section 11.1) together with the `F77_LOCAL`-specific argument passing options (Section 11.6) define the nature of the `F77_LOCAL` interface: how control is to be transferred from a global HPF procedure to a set of local procedures described by an `EXTRINSIC(F77_LOCAL)` procedure interface and how data can be passed between these two types of procedures: by reference or by descriptor, and with or without temporary local reordering of data to satisfy FORTRAN 77 provisions for sequential, contiguous storage of array data in Fortran array element order. These alternative methods of argument passing can be obtained by use of the two special-purpose attributes for extrinsic dummy arguments defined for `LANGUAGE='F77'` routines: `LAYOUT('F77_ARRAY')` (the default) vs. `LAYOUT('HPF_ARRAY')`, and `PASS_BY('*')` (the default) vs. `PASS_BY('HPF_HANDLE')`. However, to take advantage of the option allowing one to pass global HPF array “handles” to local FORTRAN 77 procedures and then obtain information locally about how the local

portion of a given parallel array is actually distributed requires special inquiry routines comparable to the HPF Local Library of functions. Since this library is not only described as a module, but uses many features such as array-valued functions and optional arguments not available in FORTRAN 77 code, it is recommended that a modified FORTRAN 77 interface to this library be provided in the manner described below. Furthermore, there is the problem of describing local portions of parallel arrays in the FORTRAN 77 code used in each local routine called from a global HPF one. Since assumed-shape syntax may not be used, explicit shape arrays are required. But it is common for global distribution of arbitrary sized arrays to result in local portions of arrays that do not have constant shapes on all processors, and the actual extents in each processor cannot necessarily be predicted in advance. In order to allow programmers to obtain axis extent information at run time from the HPF global caller, a special HPF-callable subgrid inquiry subroutine is provided. A FORTRAN 77 callable version of the same routine is also described below, for flexibility in programming.

G.2 Summary

- One HPF-callable subgrid inquiry subroutine

HPF_SUBGRID_INFO

- A set of FORTRAN 77-callable inquiry subroutines

F77_SUBGRID_INFO

F77_GLOBAL_ALIGNMENT

F77_GLOBAL_DISTRIBUTION

F77_GLOBAL_TEMPLATE

F77_ABSTRACT_TO_PHYSICAL

F77_PHYSICAL_TO_ABSTRACT

F77_LOCAL_TO_GLOBAL

F77_GLOBAL_TO_LOCAL

F77_LOCAL_BLKCNT

F77_LOCAL_LINDEX

F77_LOCAL_UINDEX

F77_GLOBAL_SHAPE

F77_GLOBAL_SIZE

F77_SHAPE

F77_SIZE

F77_MY_PROCESSOR

G.3 Global HPF Subgrid Inquiry Routine

The `F77_LOCAL` library interface includes only one global HPF subroutine, `HPF_SUBGRID_INFO`, whose implementation should be added as an extension to the standard HPF Library module. Its purpose is to provide per-processor information about the local subgrids of distributed arrays. This information is often critical when passing such arrays to local procedures written in FORTRAN 77, where array argument shapes must be stated explicitly in the local procedure (except in the last dimension; there are “assumed size” but no “assumed shape” arrays), but may be expressed in terms of arguments passed at run time (“adjustable shape arrays”). Thus the subgrid parameters obtained from this subgrid inquiry routine can be passed as arguments to the local routines and used there to describe the extents of the locally visible portions of global HPF arrays, as the example in Section G.5 will demonstrate.

**HPF_SUBGRID_INFO (ARRAY, IERR, DIM, LB, UB, STRIDE,
LB_EMBED, UB_EMBED, AXIS_MAP)**

Description. Gives local information about local subgrid allocation onto each processor of a distributed array; callable from a global HPF routine.

Class. Inquiry subroutine.

Arguments.

ARRAY is a nonsequential array of any type, size, shape, or mapping. It is an `INTENT (IN)` argument.

IERR is a scalar integer of default kind. It is an `INTENT (OUT)` argument. Its return value is zero upon successful return and nonzero otherwise. Errors result if local subgrids cannot be expressed as array sections of `ARRAY`.

If any of the optional arguments `LB_EMBED`, `UB_EMBED`, or `AXIS_MAP` is present, then a nonzero value is also returned if the compiler does not organize the local data in serial memory by sequence associating a larger “embedding” array (see Section G.3.1 below for more explanation).

DIM (optional) is a scalar integer of default kind. It is an `INTENT (IN)` argument. `DIM` indicates the axis along which return values are desired. If `DIM` is not present, values are returned for all axes.

LB (optional) is an `INTENT (OUT)`, default integer array. If this argument is present, and if the value returned in `IERR` is zero, the values returned in array `LB` are the lower bounds in

	global coordinates of each processor's subgrid, along one (if DIM is present) or each dimension of ARRAY.	1 2
UB (optional)	is an INTENT (OUT), default integer array. If this argu- ment is present, and if the value returned in IERR is zero, the values returned in array UB are the upper bounds in global coordinates of each processor's subgrid, along one (if DIM is present) or each dimension of ARRAY.	3 4 5 6 7 8
STRIDE (optional)	is an INTENT (OUT), default integer array. If this argu- ment is present, and if the value returned in IERR is zero, the values returned in array STRIDE are the strides in lo- cal memory between elements of each processor's subgrid, along one (if DIM is present) or each dimension of ARRAY.	9 10 11 12 13 14
LB_EMBED (optional)	is an INTENT (OUT), default integer array. If this argu- ment is present, and if the value returned in IERR is zero, the values returned in array LB_EMBED are the lower bounds in global coordinates of the actual global array elements allocated on each processor, possibly a superset of the user-visible subgrid, along one (if DIM is present) or each dimension of ARRAY.	15 16 17 18 19 20 21 22 23
UB_EMBED (optional)	is an INTENT (OUT), default integer array. If this argu- ment is present, and if the value returned in IERR is zero, the values returned in array UB_EMBED are the upper bounds in global coordinates of the actual global array el- ements allocated on each processor, possibly a superset of the user-visible subgrid, along one (if DIM is present) or each dimension of ARRAY.	24 25 26 27 28 29 30 31
AXIS_MAP (optional)	is a rank 2, INTENT (OUT), default integer array. If this argument is present, its shape must be at least $[n, r]$, where n is the number of processors and r is the rank of ARRAY. If the value returned in IERR is zero, the values returned in $AXIS_MAP(i, 1:r)$ represent the numbers of the axes of the subgrid on processor i from fastest varying to slowest varying, and form a permutation of the sequence $1, 2, \dots, r$.	32 33 34 35 36 37 38 39 40 41 42

For the last six arguments, LB, UB, STRIDE, LB_EMBED, UB_EMBED, and AXIS_MAP, each array has a first axis of extent at least n , where n is the number of processors, and the first n indices of that axis of each array must be distributed (perhaps via an explicit CYCLIC or BLOCK distribution) one index per processor. If a second dimension is needed, it should be a collapsed axis of extent at least equal to the rank of ARRAY.

43
44
45
46
47
48

1 If `HPF_SUBGRID_INFO` is called, and the elements of `ARRAY` that are local to any particular
2 processor are not representable as an array section of the global user array, then a nonzero
3 value is returned for `IERR`. Otherwise, if any of the optional arguments `LB`, `UB`, or `STRIDE`
4 is present, then the lower bounds, upper bounds, or strides, respectively, that describe the
5 local array sections are returned in terms of one-based, global coordinates.
6

7 **G.3.1 Subgrid Inquiries Involving Embedding Arrays**

9 In the common case in which the elements of each local subgrid of the global array argument
10 are distributed across processors, with no overlap, and allocated in local memory like a local
11 FORTRAN 77 array, as a contiguous sequence of elements in Fortran array element order,
12 these three last optional arguments would not be required.
13

14 However, some implementations may choose less common layouts in local memory,
15 that involve “embedding” these elements in a larger array section of equal rank that *is*
16 sequence-associated in serial memory. For example, alignment of axes of arrays in different
17 orders may result in a permuting embedding of the subgrid. Or axes of subgrids may be
18 padded with ghost cells, either for stencil optimizations or to achieve same-size subgrids on
19 all nodes.
20

21 In variations such as these, we may still view the subgrid as being “embedded” in a
22 sequence associated array which may be accessible in `F77_LOCAL` operations, if the permu-
23 tation of axes, shape of any embedding array, and offsets into that array can be obtained at
24 runtime. The last three arguments of `HPF_SUBGRID_INFO` are provided to allow programmers
25 to obtain this information when it is appropriate, with the help of the `IERR` flag to signal
26 when this is not the case.
27

28 In this mapping, local memory has been allocated for a larger array section, with co-
29 ordinates (`LB_EMBED` : `UB_EMBED` : `STRIDE`). The coordinates of the *actual* computational
30 elements are limited to the subset (`LB` : `UB` : `STRIDE`). The sequence association is gen-
31 eralized to an arbitrary mapping of axes. Here, `AXIS_MAP` numbers the axes from fastest
32 varying to slowest varying. If `LB_EMBED`, `UB_EMBED`, or `AXIS_MAP` is specified in a call to
33 `HPF_SUBGRID_INFO` but `ARRAY` does not satisfy the assumptions of this mapping model, then
34 a nonzero value is returned for `IERR`.
35

36 **G.4 Local FORTRAN 77 Inquiry Routines**

37 Here the F77-callable inquiry subroutines are described briefly. These provide essentially
38 the same capability as the combination of the HPF intrinsic array inquiry functions such
39 as `SHAPE` and `SIZE`, together with the `HPF_LOCAL_LIBRARY` inquiry routines. The subrou-
40 tine `F77_SUBGRID_INFO` serves as a local counterpart to the globally callable subroutine
41 `HPF_SUBGRID_INFO` described above. In all of the following:
42
43
44

- 45 • `ARRAY` is a dummy argument passed in from a global HPF caller using the `LAYOUT`
46 (`'HPF_ARRAY'`) attribute and declared within the FORTRAN 77 local subroutine as
47 a scalar integer variable. It is an `INTENT (IN)` argument.
48

- DIM is a scalar integer of default kind. It is an `INTENT (IN)` argument. This argument specifies a particular axis of the global array associated with `ARRAY` or, if `DIM = -1`, inquiry is for all axes.
- An “inquiry result” is an `INTENT (OUT)` argument. If `DIM = -1`, it is a rank-one array of size equal to at least the rank of the global array associated with `ARRAY`, returning information associated with all axes. If `DIM` is positive, the “inquiry result” is a scalar, returning information only for the axis indicated by `DIM`.
- The arguments are defined in the same way as for the corresponding `HPF` or `HPF_LOCAL` routines unless otherwise noted. See the description of `HPF_SUBGRID_INFO` above and Section 11.7.1 for full specifications of the similarly-named `HPF_LOCAL_LIBRARY` procedures.

F77_SUBGRID_INFO (ARRAY, IERR1, IERR2, DIM, LB, UB, STRIDE, LB_EMBED, UB_EMBED, AXIS_MAP)

Description. This is a FORTRAN 77-callable version of the HPF subroutine `HPF_SUBGRID_INFO`.

Arguments.

`IERR1` is a scalar integer of default kind. It is an `INTENT (OUT)` argument. Its return value is zero if `LB`, `UB`, and `STRIDE` were determined successfully and nonzero otherwise.

`IERR2` is a scalar integer of default kind. It is an `INTENT (OUT)` argument. Its return value is zero if `LB_EMBED` and `UB_EMBED` were determined successfully and nonzero otherwise.

`LB`, `UB`, `STRIDE`, `LB_EMBED`, `UB_EMBED`, `AXIS_MAP` are “inquiry results” of default integer type. They are the lower and upper bounds and strides of the array sections describing the local data (in terms of global indices), the lower and upper bounds of the embedding arrays (again, in terms of global indices), and the axes of the embedding arrays to which the axes of `ARRAY` are mapped.

F77_GLOBAL_ALIGNMENT (ALIGNEE, LB, UB, STRIDE, AXIS_MAP, IDENTITY_MAP, DYNAMIC, NCOPIES)

Description. This is a FORTRAN 77-callable version of the `HPF_LOCAL` subroutine `GLOBAL_ALIGNMENT`. All but the first are `INTENT (OUT)` arguments whose return values are as specified by the corresponding HPF routine.

TEMPLATE_RANK	is a scalar integer of default kind.	1
LB, UB, AXIS_INFO	are integer arrays of rank one. Their size must be at least equal to the rank of the align-target to which the global HPF array associated with ALIGNEE is ultimately aligned.	2 3 4 5
AXIS_TYPE	is a CHARACTER*10 array of rank one. Its size must be at least equal to the rank of the align-target to which the global HPF array associated with ALIGNEE is ultimately aligned.	6 7 8 9 10
NUMBER_ALIGNED	is a scalar integer of default kind.	11
DYNAMIC	is a scalar logical.	12 13 14

F77_ABSTRACT_TO_PHYSICAL(ARRAY, INDEX, PROC)

Description. This is a FORTRAN 77-callable version of the HPF_LOCAL subroutine ABSTRACT_TO_PHYSICAL.

Arguments.

INDEX	is a rank-one, INTENT (IN), integer array.	15 16 17 18 19 20 21 22
PROC	is a scalar, INTENT (OUT), integer.	23 24

F77_PHYSICAL_TO_ABSTRACT(ARRAY, PROC, INDEX)

Description. This is a FORTRAN 77-callable version of the HPF_LOCAL subroutine PHYSICAL_TO_ABSTRACT.

Arguments.

PROC	is a scalar, INTENT (IN), integer.	25 26 27 28 29 30 31 32 33
INDEX	is a rank-one, INTENT (OUT), integer array.	34 35

F77_LOCAL_TO_GLOBAL(ARRAY, L_INDEX, G_INDEX)

Description. This is a FORTRAN 77-callable version of the HPF_LOCAL subroutine LOCAL_TO_GLOBAL.

Arguments.

L_INDEX	is a rank-one, INTENT (IN), integer array.	36 37 38 39 40 41 42 43 44
G_INDEX	is a rank-one, INTENT (OUT), integer array.	45 46 47 48

1 **F77_GLOBAL_TO_LOCAL**(ARRAY, G_INDEX, L_INDEX, LOCAL,
2 NCOPIES, PROCS)

3 **Description.** This is a FORTRAN 77-callable version of the HPF_LOCAL subroutine
4 GLOBAL_TO_LOCAL.
5

6 **Arguments.**

7
8 G_INDEX is a rank-one, INTENT (IN), integer array.
9
10 L_INDEX is a rank-one, INTENT (OUT), integer array.
11
12 LOCAL is a scalar, INTENT (OUT), logical.
13
14 NCOPIES is a scalar, INTENT (OUT), integer.
15
16 PROCS is a rank-one, integer array whose size is at least the
17 number of processors that hold copies of the identified
18 element.

19 **F77_LOCAL_BLKCNT**(L_BLKCNT, ARRAY, DIM, PROC)
20

21 **Description.** This is a FORTRAN 77-callable version of the HPF_LOCAL function
22 LOCAL_BLKCNT.
23

24 **Arguments.**

25
26 L_BLKCNT is an “inquiry result” of type integer.
27
28 PROC is a scalar integer of default kind. It must be a valid
29 processor number or, if PROC = -1, the value returned
30 by F77_MY_PROCESSOR() is implied.
31

32 **F77_LOCAL_LINDEX**(L_LINDEX, ARRAY, DIM, PROC)
33

34 **Description.** This is a FORTRAN 77-callable version of the HPF_LOCAL function
35 LOCAL_LINDEX.
36

37 **Arguments.**

38
39 L_LINDEX is a rank-one, integer array of size equal to at least the
40 value returned by F77_LOCAL_BLKCNT.
41
42 DIM may not be -1.
43
44 PROC is a scalar integer of default kind. It must be a valid
45 processor number or, if PROC = -1, the value returned
46 by F77_MY_PROCESSOR() is implied.
47
48

F77_LOCAL_UIINDEX(L_UIINDEX, ARRAY, DIM, PROC)

Description. This is a FORTRAN 77-callable version of the HPF_LOCAL function LOCAL_UIINDEX.

Arguments.

L_UIINDEX is a rank-one, integer array of size equal to at least the value returned by F77_LOCAL_BLKCNT.

DIM may not be -1.

PROC is a scalar integer of default kind. It must be a valid processor number or, if PROC = -1, the value returned by F77_MY_PROCESSOR() is implied.

F77_GLOBAL_SHAPE(SHAPE, ARRAY)

Description. This is a FORTRAN 77-callable version of the HPF_LOCAL function GLOBAL_SHAPE.

Arguments.

SHAPE is a rank-one, integer array of size equal to at least the rank of the global array associated with ARRAY. Its return value is the shape of that global array.

F77_GLOBAL_SIZE(SIZE, ARRAY, DIM)

Description. This is a FORTRAN 77-callable version of the HPF_LOCAL function GLOBAL_SIZE.

Arguments.

SIZE is a scalar integer equal to the extent of axis DIM of the global array associated with ARRAY or, if DIM = -1, the total number of elements in that global array.

F77_SHAPE(SHAPE, ARRAY)

Description. This is a FORTRAN 77-callable version of the HPF intrinsic SHAPE, as it would behave as called from HPF_LOCAL.

Arguments.

SHAPE is a rank-one, integer array of size equal to at least the rank of the subgrid associated with ARRAY. Its return value is the shape of that subgrid.

F77_SIZE(SIZE, ARRAY, DIM)

Description. This is a FORTRAN 77-callable version of the HPF intrinsic SIZE, as it would behave as called from HPF_LOCAL.

Arguments.

SIZE is a scalar integer equal to the extent of axis DIM of the subgrid associated with ARRAY or, if DIM = -1, the total number of elements in that subgrid.

F77_MY_PROCESSOR(MY_PROC)

Description. This is a FORTRAN 77-callable version of the HPF_LOCAL function MY_PROCESSOR.

Arguments.

MY_PROC is a scalar, INTENT (OUT), integer. Its value is the identifying number of the physical processor from which this call is made.

G.5 Programming Example Using HPF_SUBGRID_INFO

G.5.1 HPF Caller

```
PROGRAM EXAMPLE
! Declare the data array and a verification copy
  INTEGER, PARAMETER :: NX = 100, NY = 100
  REAL, DIMENSION(NX,NY) :: X, Y
!HPF$ DISTRIBUTE(BLOCK,BLOCK) :: X, Y
! The global sum will be computed
! by forming partial sums on the processors
  REAL PARTIAL_SUM(NUMBER_OF_PROCESSORS())
!HPF$ DISTRIBUTE PARTIAL_SUM(BLOCK)
! Local subgrid parameters are declared per processor
! for a rank-two array
  INTEGER, DIMENSION(NUMBER_OF_PROCESSORS(),2) ::
    & LB, UB, NUMBER
!HPF$ DISTRIBUTE(BLOCK,*) :: LB, UB, NUMBER
! Define interfaces
  INTERFACE
    EXTRINSIC(F77_LOCAL) SUBROUTINE LOCAL1
      & ( LB1, UB1, LB2, UB2, NX, X )
! Arrays LB1, UB1, LB2, UB2, and X are passed by default
```

```

! as LAYOUT('F77_ARRAY') and PASS_BY('*')
      INTEGER, DIMENSION(:) :: LB1, UB1, LB2, UB2
      INTEGER NX
      REAL X(:, :)
!HPF$ DISTRIBUTE(BLOCK) :: LB1, UB1, LB2, UB2
!HPF$ DISTRIBUTE(BLOCK,BLOCK) :: X
      END
      EXTRINSIC(F77_LOCAL) SUBROUTINE LOCAL2(N,X,R)
! Arrays N, X, and R are passed by default
! as LAYOUT('F77_ARRAY') and PASS_BY('*')
      INTEGER N(:)
      REAL X(:, :), R(:)
!HPF$ DISTRIBUTE N(BLOCK)
!HPF$ DISTRIBUTE X(BLOCK,BLOCK)
!HPF$ DISTRIBUTE R(BLOCK)
      END
      END INTERFACE

! Determine result using only global HPF
      ! Initialize values
      FORALL (I=1:NX,J=1:NY) X(I,J) = I + (J-1) * NX
      ! Determine and report global sum
      PRINT *, 'GLOBAL HPF RESULT: ',SUM(X)
! Determine result using local subroutines
      ! Initialize values ( assume stride = 1 )
      CALL HPF_SUBGRID_INFO( Y, IERR, LB=LB, UB=UB )
      IF (IERR.NE.0) STOP 'ERROR!'
      CALL LOCAL1( LB(:,1), UB(:,1), LB(:,2), UB(:,2), NX, Y )
      ! Determine and report global sum
      NUMBER = UB - LB + 1
      CALL LOCAL2 ( NUMBER(:,1) * NUMBER(:,2) , Y , PARTIAL_SUM )
      PRINT *, 'F77_LOCAL RESULT #1 : ',SUM(PARTIAL_SUM)
      END

```

G.5.2 FORTRAN 77 Callee

```

      SUBROUTINE LOCAL1( LB1, UB1, LB2, UB2, NX, X )
! The global actual arguments passed to LB1, UB1, LB2, and UB2
! have only one element apiece and so can be treated as scalars
! in the local Fortran 77 procedures
      INTEGER LB1, UB1, LB2, UB2
! NX contains the global extent of the first dimension

```

```

1      ! of the global array associated with local array X
2      INTEGER NX
3      ! Note that X may have no local elements.
4      REAL X ( LB1 : UB1 , LB2 : UB2 )
5      ! Initialize the elements of the array, if any
6      DO J = LB2, UB2
7          DO I = LB2, UB2
8              X(I,J) = I + (J-1) * NX
9          END DO
10     END DO
11     END DO
12     END
13
14     SUBROUTINE LOCAL2(N,X,R)
15     ! Here, the rank of the original array is unimportant
16     ! Only the total number of local elements is needed
17     !
18     INTEGER N
19     REAL X(N), R
20     ! If N is zero, local array X has no elements, but R
21     ! still computes the correct local sum
22     R = 0.
23     DO I = 1, N
24         R = R + X(I)
25     END DO
26     END
27
28
29
30

```

G.6 Programming Example Using F77-Callable Inquiry Subroutines

This example performs only the initialization of the above example. It illustrates use of the F77-callable inquiry routines on descriptors passed from HPF, as well as the addressing of uncompressed local subgrid data in terms of “embedding arrays.”

G.6.1 HPF Caller

```

37     PROGRAM EXAMPLE
38     INTEGER, PARAMETER :: NX = 100, NY = 100
39     REAL, DIMENSION(NX,NY) :: X
40     !HPF$ DISTRIBUTE(BLOCK,BLOCK) :: X
41     ! Local subgrid parameters are declared per processor
42     ! for a rank-two array
43     INTEGER, DIMENSION(NUMBER_OF_PROCESSORS(),2) ::
44     & LB, UB, LB_EMBED, UB_EMBED
45     !HPF$ DISTRIBUTE(BLOCK,*) :: LB, UB, LB_EMBED, UB_EMBED
46     ! Define interfaces
47
48

```

```

INTERFACE
    EXTRINSIC(F77_LOCAL) SUBROUTINE LOCAL1(
& LB1, UB1, LB_EMBED1, UB_EMBED1,
& LB2, UB2, LB_EMBED2, UB_EMBED2, X, X_DESC )
    INTEGER, DIMENSION(:) ::
& LB1, UB1, LB_EMBED1, UB_EMBED1,
& LB2, UB2, LB_EMBED2, UB_EMBED2
! X is passed twice, both times without local reordering.
! First, it is passed by reference for accessing array elements.
    REAL, DIMENSION(:,,:), LAYOUT('HPF_ARRAY'),
& PASS_BY('*') :: X
! It is also passed by descriptor for use in F77 LOCAL
! LIBRARY subroutines only.
    REAL, DIMENSION(:,,:), LAYOUT('HPF_ARRAY'),
& PASS_BY('HPF_HANDLE') :: X_DESC
!HPF$ DISTRIBUTE(BLOCK) :: LB1, UB1, LB_EMBED1, UB_EMBED1
!HPF$ DISTRIBUTE(BLOCK) :: LB2, UB2, LB_EMBED2, UB_EMBED2
!HPF$ DISTRIBUTE(BLOCK,BLOCK) :: X
    END
END INTERFACE
! Initialize values
! ( Assume stride = 1 and no axis permutation )
    CALL HPF_SUBGRID_INFO( X, IERR,
& LB=LB, LB_EMBED=LB_EMBED,
& UB=UB, UB_EMBED=UB_EMBED)
    IF (IERR.NE.0) STOP 'ERROR!'
    CALL LOCAL1(
& LB(:,1), UB(:,1), LB_EMBED(:,1), UB_EMBED(:,1),
& LB(:,2), UB(:,2), LB_EMBED(:,2), UB_EMBED(:,2), X, X )
    END

```

G.6.2 FORTRAN 77 Callee

```

SUBROUTINE LOCAL1(
& LB1, UB1, LB_EMBED1, UB_EMBED1,
& LB2, UB2, LB_EMBED2, UB_EMBED2, X, X_DESC )
    INTEGER LB1, UB1, LB_EMBED1, UB_EMBED1
    INTEGER LB2, UB2, LB_EMBED2, UB_EMBED2
! The subgrid has been passed in its 'embedded' form
    REAL X ( LB_EMBED1 : UB_EMBED1 , LB_EMBED2 : UB_EMBED2 )
! Locally X_DESC is declared as an INTEGER
    INTEGER X_DESC

```

```
1      ! Get the global extent of the first axis
2      ! This is an HPF_LOCAL type of inquiry routine with an
3      ! 'F77_' prefix
4          CALL F77_GLOBAL_SIZE(NX,X,1)
5      ! Otherwise, initialize elements of the array
6      ! Loop only over actual array elements
7          DO J = LB2, UB2
8              DO I = LB2, UB2
9                  X(I,J) = I + (J-1) * NX
10             END DO
11         END DO
12     END DO
13 END
```