

pt2pt wg

**FP16**

**Virtual MPI Forum Meeting**

**Jan 31, 2018**

**Atsushi Hori**

**RIKEN AICS**

# An Important Issue to be discussed

- **When shall we introduce the new datatypes?**
  - **As soon as,**
    - **ISO/IEEE/ANSI standardizes them, or**
    - **Major compilers accept them, or**
    - **An implementor decides to do so, ...**
- **It may take (more than ?) one year to put new datatypes into the standard, if we follow the normal procedure (passing two readings)**
  - **It is almost IMPOSSIBLE to have datatypes in the standard as soon as XYZ happens**

# New Ticket #74

- Replacement of ticket-65 and -66
- Stop introducing new types into the standard MPI standard only defines *fundamental* data types such as MPI\_INT, MPI\_FLOAT, and so on
  - decouple type definition from standard
  - give implementors the freedom of choice when to add and what to add (new basic datatypes)
- **Instead,**
  - Standard defines the *naming rule* of basic datatype names
  - Implementors can add new basic datatypes having the names following the naming rule

# Related Tickets

- **#65 16-bit floating-point support for C/C++**
  - ➔ MPI standard will not define FP16, but implementor can decide to support or not
- **#66 define language-agnostic, IEEE types**
  - ➔ Implementors can decide to implement new types
- **#69 Definition and Explanation of External32**
  - ➔ Left unchanged because of backward compatibility (wording might be updated)

# Definition of Terms

- *Basic datatypes*
  - primitive datatypes
    - can be appeared in the type maps
  - Two kinds
    - *Fundamental datatypes*
    - *Additional datatypes*
- *Fundamental datatypes*
  - Defined in the standard
- *Additional datatypes*
  - Added by an implementor

# Modification to the standard

- Adding one section into Chapter 4. Datatypes
- Fundamental data type definition
- Addition of fundamental data type names
  - Data type naming rule
  - Old basic data type names become synonyms (and deprecated ?)
  - ➔ If deprecated, then ALL deprecated datatypes in the current standard must be renamed to the new ones !!!!
- MPI\_type\_equivalent() function
  - to check if two data types are equivalent or not in an implementation

# **Impact to users**

- **Old datatype names are left unchanged as synonyms and no impact to current users**

# Details (1)

- **Fundamental Datatypes (in old names here)**
  - C MPI\_CHAR, MPI\_INT, MPI\_LONG, ...
  - Fortran MPI\_CHARACTER, MPI\_INTEGER, ...
  - MPI MPI\_BYTE, MPI\_PACKED, MPI\_AINT, ...
- **Implementation-defined Additional Datatypes**
  - *Data Type Naming Rule (later)*
  - An implementation may add basic datatype(s) if the implementation and associated compiler support it (them).
  - The current data type names not compliant with the rule are thought to be synonyms.



# Details (2)

- **Advice to users (draft)**
  - Excepting external32, MPI standard only defines the datatype names, regardless to the sizes of datatypes, data representation format, and if the data type is computed as is or any format conversion(s) is (are) accompanied with basics computation(s) [4 calculus, bit ops, etc.]. An implementation may add basic datatypes supported by the accompanied compiler. The datatype names added by an implementation follow the datatype naming rule in the standard so that users can identify the datatypes of the underlying programming language from the datatype names used in MPI. If any, documentation on these additional datatypes should be provided by the implementor.
  - Actual storage size of a datatype depends on an implementation, and accompanied compiler (see also `MPI_Type_size(x)`). For example, `MPI_LONG_INT` has the size of 4 bytes or 8 bytes depending on the compiler and/or processor. In the usage of MPI-IO with “external32,” however, datatype sizes and formats are explicitly defined in this standard for compatibility and portability (see also MPI-IO chapter XYZ).
- **Advice to implementors (draft)**
  - If some or all global reduction operations are offloaded to a device other than the host processor, e.g., network device, then the data sizes and binary formats of all datatypes must be exactly the same with the ones of the processor.

# New Data Type Function

- `int MPI_Type_equivalent( type0, type1, *flag )`
  - *flag* is set, if *type0* and *type1* are having the same binary format and size. Otherwise *flag* is reset.
  - If derived datatypes are given, then the type maps of both types must be equivalent. Here, “equivalent type maps” means all basic datatype and offset pairs in two type maps are the “equivalent” defined by `MPI_Type_equivalent()`.
  - ex)
    - `MPI_Type_equivalent( MPI_SHORT_FLOAT, MPI_FLOAT, flag )`
    - `MPI_Type_equivalent( MPI_REAL, MPI_FLOAT, flag )`

# Naming Rule Requirements

- Users can identify the underlying datatypes from the MPI datatype names
- Implementors can easily make new unique datatype names when they want to introduce new datatypes.
- How do we call the rule ? (straw vote)
  - **Datatype Naming Rule**, or
  - **Datatype Naming Convention**

# Type Naming Rule (1/3)

- MPI\_<CAT>\_<TYPENAME>[\_<EXTRA>], or  
MPI\_<CAT>\_\_<TYPENAME>[\_\_<EXTRA>]
  - examples
    - MPI\_C\_LONG\_LONG or MPI\_C\_\_LONG\_LONG
    - MPI\_CXX\_FLOAT or MPI\_CXX\_\_FLOAT
    - MPI\_F\_DOUBLE or MPI\_F\_\_DOUBLE
    - MPI\_F\_REAL\_2 or MPI\_F\_\_REAL\_\_2, ...
- MPI-defined types
  - MPI\_MPI\_BYTE or MPI\_MPI\_\_BYTE
  - MPI\_MPI\_PACKED, or MPI\_MPI\_\_PACKED, ...
- Language agnostic types
  - MPI\_IEEE\_BINARY32 or MPI\_IEEE\_\_BINARY32
  - MPI\_ISO\_DECIMAL64 or MPI\_ISO\_\_DECIMAL64, ...

# Type Naming Rule (2/3)

- `MPI_<KIND>_<TYPENAME>[_<EXTRA>]`  
`MPI_<KIND>__<TYPENAME>[__<EXTRA>]`
- **KIND**
  - Language “C”, “CXX”, “F”, ...  
(filename extension?)
  - MPI-defined “MPI”
  - Standard “ISO”, “IEEE”, “ANSI”, ...
- **Type Name** (defined by language, shortest one)
  - C/C++ “int”, “long long”, ...
  - Fortran “INTEGER”, “REAL”, ...
  - ISO/IEEE “binary16”, “decimal32”, ..
- **Extra**
  - Fortran “REAL(4)”, “COMPLEX(8)”, ...

# Type Naming Rule (3/3)

- MIN\_LOC, MAX\_LOC (if still needed)
  - `<TYPENAME0>__AND__<TYPENAME1>` or
  - `<TYPENAME0>__AT__<TYPENAME1>`
- MPI\_C\_FLOAT\_AND\_MPI\_C\_INT or  
MPI\_C\_\_FLOAT\_\_AND\_\_MPI\_C\_\_INT
- MPI\_C\_FLOAT\_AT\_MPI\_C\_INT or  
MPI\_C\_\_FLOAT\_\_AT\_\_MPI\_C\_\_INT
- Straw vote ?

# Old and New Type Names (1/2)

C Kind			Fortran Kind		
Current Name	Proposed New Name	Ext32 Size	Current Name	Proposed New Name	Ext32 Size
MPI_CHAR	<b>MPI_C_CHAR</b>	1	MPI_CHARACTER	<b>MPI_F_CHARACTER</b>	1
MPI_SHORT	<b>MPI_C_SHORT</b>	2	MPI_LOGICAL	<b>MPI_F_LOGICAL</b>	4
MPI_INT	<b>MPI_C_INT</b>	4	MPI_INTEGER	<b>MPI_F_INTEGER</b>	4
MPI_LONG	<b>MPI_C_LONG</b>	4	MPI_REAL	<b>MPI_F_REAL</b>	4
MPI_LONG_LONG_INT	<b>MPI_C_LONG_LONG</b>	8	MPI_DOUBLE_PRECISION	<b>MPI_F_DOUBLE_PRECISION</b>	8
MPI_LONG_LONG	<b>MPI_C_LONG_LONG</b>	8	MPI_COMPLEX	<b>MPI_F_COMPLEX</b>	2*4
MPI_SIGNED_CHAR	<b>MPI_C_SIGNED_CHAR</b>	1	MPI_DOUBLE_COMPLEX	<b>MPI_F_DOUBLE_COMPLEX</b>	2*8
MPI_UNSIGNED_CHAR	<b>MPI_C_UNSIGNED_CHAR</b>	1	(followings are optional datatypes in the current)		
MPI_UNSIGNED_SHORT	<b>MPI_C_UNSIGNED_SHORT</b>	2	MPI_INTEGER1	MPI_F_INTEGER_1	1
MPI_UNSIGNED	<b>MPI_C_UNSIGNED</b>	4	MPI_INTEGER2	MPI_F_INTEGER_2	2
MPI_UNSIGNED_LONG	<b>MPI_C_UNSIGNED_LONG</b>	4	MPI_INTEGER4	MPI_F_INTEGER_4	4
MPI_UNSIGNED_LONG_LONG	<b>MPI_C_UNSIGNED_LONG_LONG</b>	8	MPI_INTEGER8	MPI_F_INTEGER_8	8
MPI_FLOAT	<b>MPI_C_FLOAT</b>	4	MPI_INTEGER16	MPI_F_INTEGER_16	16
MPI_SHORT_FLOAT	<b>MPI_C_SHORT_FLOAT</b>	(2)	MPI_REAL2	MPI_F_REAL_2	2
MPI_DOUBLE	<b>MPI_C_DOUBLE</b>	8	MPI_REAL4	MPI_F_REAL_4	4
MPI_LONG_DOUBLE	<b>MPI_C_LONG_DOUBLE</b>	16	MPI_REAL8	MPI_F_REAL_8	8
MPI_WCHAR	<b>MPI_C_WCHAR</b>	2	MPI_REAL16	MPI_F_REAL_16	16
MPI_C_BOOL	<b>MPI_C_BOOL</b>	1	MPI_COMPLEX4	MPI_F_COMPLEX_4	2*2
MPI_INT8_T	<b>MPI_C_INT8_T</b>	1	MPI_COMPLEX8	MPI_F_COMPLEX_8	2*4
MPI_INT16_T	<b>MPI_C_INT16_T</b>	2	MPI_COMPLEX16	MPI_F_COMPLEX_16	2*8
MPI_INT32_T	<b>MPI_C_INT32_T</b>	4	MPI_COMPLEX32	MPI_F_COMPLEX_32	2*16
MPI_INT64_T	<b>MPI_C_INT64_T</b>	8			
MPI_UINT8_T	<b>MPI_C_UINT8_T</b>	1			
MPI_UINT16_T	<b>MPI_C_UINT16_T</b>	2			
MPI_UINT32_T	<b>MPI_C_UINT32_T</b>	4			
MPI_UINT64_T	<b>MPI_C_UINT64_T</b>	8			
MPI_C_COMPLEX	<b>MPI_C_COMPLEX</b>	2*4			
MPI_C_FLOAT_COMPLEX	<b>MPI_C_COMPLEX</b>	2*4			
MPI_C_SHORT_COMPLEX	<b>MPI_C_SHORT_COMPLEX</b>	(2*2)			
MPI_C_DOUBLE_COMPLEX	<b>MPI_C_DOUBLE_COMPLEX</b>	2*8			
MPI_C_LONG_DOUBLE_COMPLEX	<b>MPI_C_LONG_DOUBLE_COMPLEX</b>	2*16			

- The new type names in blue are *fundamental datatypes*
- Ext32 sizes in brackets are not defined in the current standard (and future)

# Old and New Type Names (2/2)

C++ Kind			MPI Kind		
Current Name	Proposed New Name	Ext32 Size	Current Name	Proposed New Name	Ext32 Size
MPI_CXX_BOOL	<b>MPI_CXX_BOOL</b>	1	MPI_BYTE	<b>MPI_MPI_BYTE</b>	1
MPI_CXX_FLOAT_COMPLEX	<b>MPI_CXX_FLOAT_COMPLEX</b>	2*4	MPI_PACKED	<b>MPI_MPI_PACKED</b>	(1)
MPI_CXX_DOUBLE_COMPLEX	<b>MPI_CXX_DOUBLE_COMPLEX</b>	2*8	MPI_AINT	<b>MPI_MPI_AINT</b>	8
MPI_CXX_LONG_DOUBLE_COMPLEX	<b>MPI_CXX_LONG_DOUBLE_COMPLEX</b>	2*16	MPI_COUNT	<b>MPI_MPI_COUNT</b>	8
			MPI_OFFSET	<b>MPI_MPI_OFFSET</b>	8