



RMA Working Group Status

Brian Barrett, Pavan Balaji,
Torsten Hoefler, Rolf Rabenseifner,
Rajeev Thakur, Keith D. Underwood

1

New Plan: Merge and then Split the Proposals

- Some ideas were killed
- Other ideas were merged
- Merged ideas will be brought forth as a “small” set of individually motivated proposals
 - *Proposal 1: Basic semantics to improve application usage scenarios*
 - *Proposal 2: Extensions/Improvements to the memory model*
 - *Proposal 3: User defined RMW/Accumulate ops*
 - *Proposal 4: Query for cache coherency support*
 - *Proposal 5: Individual requests for completion*
 - *Proposal 6: Remote copy*

2



Typed compare and swap

- Are we willing to add a typed compare-and-swap?
 - Thought is that compare and swap would be for one element only
 - Type would match on source and target
- One version of compare and swap:
 - `MPI_Comp_and_swap(void *src, void *comp, MPI_Datatype type, int rank, MPI_Aint displacement, MPI_Win win);`
- Alternative would be one function per type
 - `MPI_Comp_and_swap_int(int src, int comp, int rank, MPI_Aint displacement, MPI_Win win);`
 - `MPI_Comp_and_swap_float(float src, float comp, int rank, MPI_Aint displacement, MPI_Win win);`

3



User defined Operations

- Should we allow them?
 - Requires a collective function call to register it across all processes
 - All restrictions as MPI_Reduce apply
- Pros:
 - Turing complete (users can implement their own handlers)
 - E.g., Reset of target buffer
- Cons:
 - Locking issues (concurrent RMWs or ACCs from multiple origins on the same target)
 - Progression issues

4



What are the goals of RMA?

- Three basic options:
 - Support direct usage by users
 - Support 3rd party communication libraries
 - Which ones?
 - GASNet?
 - ARMCI?
 - Support high level PGAS capabilities
 - Which ones?
 - UPC?
 - CAF?
 - Global Array?

5



MPI_ALLOC_MEM

- Lock/Unlock semantics in MPI-2.2 require only memory allocated with MPI_ALLOC_MEM to be used for portability
 - Reason: this allows some platforms to implement passive target without an asynchronous agent
- Vote:
 - Does this still need to be a case for “portability” (it’ll continue to be required for “performance”)?

6

