

Endpoints Revisions Discussion Items

MPI Forum, Hybrid WG

March 24, 2014

Design Goals for the Interface

- Endpoints should be a generalization of the existing communicators concept
 - Not a new, orthogonal type of communicator
 - Query a communicator's type?
- Endpoints should be usable when initialized with SINGLE, FUNNELED, SERIALIZED, and MULTIPLE
 - If we remove the relaxed collective semantic for `MPI_Comm_free()`, we can only support MULTIPLE
 - Required because of collective attribute callbacks?
- *Next time: Create full discussion around this issue so WG understands implications.*

Issue: MPI_COMM_COMPARE

- Current approach:
 - Return MPI_IDENT
 - User must check rank of identical handles to see if they are different endpoints in the same communicator
- Proposal: Add MPI_ALIASED or MPI_ENDPOINTS
 - Indicates that communicator handles correspond to different ranks in the same communicator
- *Next time: Define MPI_ALIASED and review*

Issue: Query number of endpoints

- Challenge: This is not straightforward, because the user could call `MPI_Comm_create_endpoints()` using an endpoints communicator
 - I.e. there are multiple ancestors within my `MPI_COMM_WORLD` process
- Query needs a point of reference, must be done with respect to a parent communicator
 - What if query is done with an ancestor communicator, rather than the immediate parent?
 - What if the parent communicator from which the endpoints communicator has been freed?
 - Answering these queries is hard, and could require MPI to do a lot of internal bookkeeping
- These are the reasons why we chose to leave a query function for a future proposal...

Straw Man Query Function

`MPI_Comm_query_ep(MPI_Comm ref_comm, MPI_Comm ep_comm, Output?)`

- Local-only query
 - Cannot query endpoints at a remote process
 - Intended to return the number of ranks in `ep_comm` that share a common parent rank in `ref_comm`
- What is the output?
 - Number of ranks in `ep_comm` that share a rank in `ref_comm`?
 - What if `ref_comm` is an endpoints communicator?
 - Return number of endpoints that share the same ancestor rank as the `ep_comm` handle?
 - Can you query the `ref_comm` or `ep_comm` ranks?
 - Can you query the `MPI_Comm` handles for these ranks?
 - **This is what is required in order to solve the problem of deciding how to call collectives.**
- *Decision: WG decided not to add a query function.*
- *TODO: Prepare notes on why we made this decision and gather feedback from Forum members that requested the query function.*