

Example: single application, no library

BSP application

At barrier, take global checkpoint

Rank j notified of node failure when communicates with node

Node fails -> subset of rank that communicate with node are notified

At next barrier, all will be notified

If rank j notified of failure, initiate global protocol

If rank 0 informed of failure, tells all to roll back to prior checkpoint

Lessons learned:

Synchronous notification is sufficient

Example: Charm++ overdecomposition on top of MPI

Charm++ runs multiple “message passing objects”(chares) on each core

If communication between set of ranks requires frequent retransmits this maps to a “slow-down in communication” event within this set.

Ranks in set get notified

Notified ranks collaborate in redistributing chares to new locations to avoid slow communicator

Rank j notified of node failure when communicates with node

Need to define healing of MPI state (MPI_COMM_WORLD)