



Argonne
NATIONAL
LABORATORY

... for a brighter future

Hybrid Programming in MPI-3

Pavan Balaji, Argonne National Laboratory



U.S. Department
of Energy

UChicago ►
Argonne_{LLC}



A U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC

Status Update

- We have two proposals from the working group so far
 - Matched probe (Torsten Hoefler)
 - Presented at earlier Forum meetings
 - MPI Endpoints (Marc Snir)
 - Under discussion within the working group for a few months
 - Several drafts of the proposal reviewed and discussed
 - First presentation of the concepts today
 - Other proposals have been merged into the MPI Endpoints proposal (e.g., threads as processes proposal, interoperating with PGAS models)
 - Some other proposals (or subsets of proposals) have been dropped or held back for the time being

Relevant issues yet to be covered

- Interoperating with other models
 - MPI + CUDA/OpenCL
 - Being studied by Purushotham Bangalore
 - MPI + PGAS (as processes)
 - Being studied by Abhinav Vishnu
 - MPI + TBB/Ct
 - Being studied by Alexander Supalov
 - Interoperating with higher-level models above MPI
 - We have not yet decided if this will be covered by the working group
 - Jeff Squyres and Torsten Hoefler will present details at one of the future working group telecons

Other discussion items

- Who will manage resources in hybrid models?
 - E.g., how many processes and how many threads?
 - Discussed: the working group mostly considers this to be out-of-scope for MPI (or at least MPI-3)
- Differentiating thread package flavors
 - MPI does not distinguish thread packages (e.g., pthreads vs. Windows threads vs. Solaris threads)
 - MPI implementations have to be given this information out-of-band (e.g., through configure options)
 - This is needed to access thread-local storage, locks, etc.
 - Should MPI-3 continue to ignore this aspect?
 - Should the user specify how to access it as a part of MPI-3?