
On the Road to DiPOSH: Adventures in High-Performance OpenSHMEM

Camille Coti¹, Allen D. Malony²

¹LIPN, Université Paris 13, Sorbonne Paris Cité, France

²University of Oregon, USA

Camille.Coti@lipn.univ-paris13.fr

malony@cs.uoregon.edu

Future HPC programming systems must address the challenge of how to integrate shared and distributed memory parallelism. The growing number of server cores argues in favor of shared memory multithreading at the node level, but makes interfacing with distributed communication libraries more problematic. Alternatively, implementing rich message passing libraries to run across codes can be cumbersome and inefficient. The paper describes an attempt to address the challenge with OpenSHMEM, where a lean API makes for a high-performance shared memory operation and communication semantics maps directly to fast networking hardware. DiPOSH is our initial attempt to implement OpenSHMEM with these objectives. Starting with our node-level POSH design, we leveraged MPI one-sided support to get initial internode functionality. The paper reports our progress. To our pleasant surprise, we discovered a natural and compatible integration of OpenSHMEM and MPI, in contrast to what is found in MPI+X hybrids today.

Keywords: Run-time environment, Communication library, Distributed OpenSHMEM.