

Towards Enhancing OpenMP Expressiveness and Performance

Haoqiang Jin

NASA Ames Research Center
hjin@nas.nasa.gov

Abstract. Since its introduction in 1997, OpenMP has become the de facto standard for shared memory parallel programming. The notable advantages of the model are its global view of memory space that simplifies programming development and its incremental approach toward parallelization. However, it is very challenge to scale OpenMP codes to tens or hundreds of processors. This problem becomes even more profound with the recent introduction of multi-core, multi-chip architectures. Several extensions have been introduced to enhance OpenMP expressiveness and performance, including thread subteams and workqueuing. In this talk, we describe applications that expose the limitation of the current OpenMP and examine the impact of these extensions on application performance. We focus on exploiting multi-level parallelism and dealing with unbalanced workload in applications with these extensions and compare with other programming approaches, such as hybrid. Our experience has demonstrated the importance of the new language features for OpenMP applications to scale well on large shared memory parallel systems.