**SC13 BOF Summary: Coprocessors, GPUs, MICs: Enticing Users to Jump onto Accelerators**

The BOF presented two Centers' experiences, the NASA Center for Climate Simulation and the Application Acceleration Center of Excellence, University of Tennessee, with supporting users as they migrated their code to the Xeon Phi architecture. Both Centers described the need for a stable environment (i.e. OS, compilers), the importance of support tools such as VTune, DDT, Map, etc, and the overarching need for training support staff and the users in the new architecture, its impact on coding requirements and how to use the tools to get the most performance possible from the Xeon Phis. All agreed that while most users' code ports quickly, a lot of code rewriting is required, particularly if the code is not well vectorized to start. There was additional agreement that the current efforts are a necessary exercise in preparation for the future generations of many core processors which are the future of supercomputing. Both Centers stressed the need for getting the right tools and training in place early in the process and providing thorough documentation.

Questions from the audience and the discussion that followed addressed the similar need to provide staff who can support the code modifications whether conducting a GPU rollout or a Phi rollout, the need to express and exploit parallelism, the growth of disparity between legacy and newly written, optimized code and the desire for increased support from compiler developers to provide mechanisms for code that is architecture agnostic. There was discussion on how to improve communication between the users, possibly by holding seminars where code owners are make 10-15 minute presentations on their code. There was also some lively discussion on the need for an application catalog listing code that had been ported and detailing the level of performance improvement, whether successful or not, and the effort involved. An observation was made that even if the ported code ran at the same speed on the Xeon Phi, the lower acquisition and operating cost of the Xeon Phi made that a positive for the Center and its users. It must, however, be noted that this will require effort on the part of both the user and the support staff.

Ideally, by next year, the community will have made advances in the identification of a set of best practices for migrating to the Xeon Phi and created a list of ported applications and their success rates.