Activities and Results of the Recent Meeting of the International Exascale Software Project (IESP), San Francisco, CA, USA, April 2011

Jack Dongarra, Pete Beckman, Terry Moore

1. Introduction

At the IESP's recent meeting in San Francisco, the ideas of "hardware/software co-design" and the "co-design process" took center stage as organizing concepts for the community's exascale effort. Although these ideas were pioneered for use in the development of mobile and embedded processors and microcontrollers, the IESP community, facing the formidable challenges of computing at extreme scale, recognized early on that they could and should be generalized for use in this more complex domain. The first version of the Roadmap, set in motion at the Tskuba meeting, proposed the use of "Application Co-Design Vehicles" to provide targets for, and feedback to, the research and development efforts of the HPC software community. As the meeting in San Francisco progressed, it was clear that this generalized co-design approach had undergone significant development over the last year, and is now being used to structure the inter-community collaboration on the climb to exascale.

2. Overview of Plenary Sessions

Judging from the plenary presentations by different international groups on the first day of the meeting, the trek upward from petascale is proceeding apace in other countries and regions. In the morning, a series of reports from the *European Exascale Software Initiative* (*EESI*) showed that the EU community is not only making significant progress on domain applications, enabling technologies, and system software, but that they are also taking an increasingly larger role in the IESP effort as a whole. Especially noteworthy in this latter respect is the summary report of Bernd Mohr regarding the survey he conducted into the compositions of the software stacks currently in use at HPC centers around the world. This survey, probably the first one of its kind, puts down an important marker to monitor global progress from the current stack to the one that the extreme scale systems of the future will require. In the late afternoon, reports from the Asian wing of the IESP —Japan and China— showed early versions of this new stack will need to arrive soon, since these countries have already dramatically increased the size of the systems their scientists are using, with plans for even more substantial increases over the next one to two years.

The update on US exascale efforts, presented in the early afternoon, focused on the *Exascale Software Center* (*ESC*) now under development by the DOE. The ESC's aggressive plan for creating and deploying an exascale-capable system software stack is built around the co-design process, with science-driven, DOE *Co-Design Centers*, focused on different areas of the DOE's research portfolio, working to inform the evolution of the projected stack. The most tangible expression of this effort is the ESC's "Application Inventory" (or "Co-Design Code Team Survey"), which is essentially an interviewing tool intended to help ascertain what the science application teams are using now, what they believe they will need to achieve their goals for exascale science, and what their most critical design tradeoff decisions will be, relative to the exascale stack. During the months preceding the San Francisco meeting, different members of ESC development teams used the Inventory to systematically query the application groups at the co-design centers. Not only were the results of these interviews summarized and presented at the meeting, but as the progress reports on the layers of the emerging ESC stack were made, it was evident that the teams developing those layers had benefitted from, and were responding to, insights gleaned from the Inventory interviews.

In a similar way, the reports from leaders of several application Co-Design Centers, given in their plenary, also showed the positive effects of engaging in the co-design process. In the past, domain science researchers describing their applications to a computer science and engineering audience would often focus purely on scientific aspects of the application and be relatively thin on facts critical to the design tradeoffs that the hardware and software infrastructure builders cared about. By contrast, the reports from these DOE Co-Design Centers in mission critical application areas —Nuclear Energy, Fusion, Materials, High-Energy Density Physics, and Combustion— were richer in such detail. They show that involvement in the co-design process can improve the ability of scientists to translate their projected research plans into design requirements and tradeoff rationales for future computational infrastructure. These presentations set the stage well for the breakout sessions to follow.

3. Overview of Breakout Sessions

The breakout sessions on the second day divided into two stages: In the first part, the meeting participants from the different constituent groups engaged in the co-design process —hardware, software, and science applications— met with the goal of articulating what they either need from, or can to provide to, the **other** two groups. The slides presenting the output of these sessions, all of which are available on line, show that each of the co-design areas is developing a much clearer, and more specific, conception of how to work with the other two groups.

The hardware group, for example, is exploring the possibility of providing an abstract machine model (or simulator) of a future exascale platform sufficient to allow the other groups to evaluate tradeoffs and bottlenecks. But perhaps the most significant new idea to emerge out of this initial breakout came from the application group. The idea is that each application area should produce a "*skeleton*" or "*compact application*," i.e. a much smaller version of the full science application which, although it might not produce the right physical answers or fully represent the science involved, would represent the underlying algorithms, and be much easier for hardware and software infrastructure developers to install and test. Skeleton applications would complement the experience-based projections gathered from the application inventories with miniature models of the full applications, which could be probed for data on how the applications would behave on real systems. The creation of these compact apps would provide the co-design process with a code base that could put it on a much more solid empirical foundation.

The three groups in the final breakout session were charged with setting the agenda of action items for the next six months and with contributing to the next update of the IESP roadmap. Each one, however, attended to a different aspect or part of the overall effort. The largest group focused on the co-design process. Examining the execution plans of the Co-Design Centers represented at the meeting, this group endeavored to identify common elements, problems, and opportunities for collaboration in co-design across the different centers. A second group took up the questions surrounding IESP governance; building on the work of previous meetings, this group created a straw-man international governance model to be incorporated into the Roadmap. Finally, the third group tried to identify and explore specific areas where, given the extant extreme scale computing programs in different countries, international collaboration in the development of a common exascale stack would be both feasible and mutually beneficial. Written results from all three groups will be integrated into the next version of the IESP Roadmap, which will be published this summer.

4. Summary of Main Outcomes and Action Items

In summary, the main outcomes and action items from the San Francisco meeting include the following:

- *Organize CDC efforts around compact/skeleton applications*: The development of compact or skeleton applications offers a natural way to focus the efforts of the Co-Design Centers because both the hardware and software communities can use them to explore potential tradeoffs and bottlenecks, quantitatively. The community needs to act soon to create a common repository for these compact apps so that as large a community as possible can benefit from them and complementary efforts to create and maintain them can be synchronized.

- *Refine and normalize the Application Inventory*: There was general agreement at the meeting that the Application Inventory developed by the ESC is an excellent information gathering tool for groups engaged in co-design. But if the results of such interviews are going to be shared and compared across national boundaries and over time, it is equally clear that the survey should be standardized as much as possible, at least informally. The application survey is already on-line and further work to accomplish this goal is planned before the next meeting.
- *Establish a common terminology*: The highly interdisciplinary nature of the co-design process, as well as diversity and distribution of potential co-design centers/groups, makes it essential to choose a common terminology so that different groups can coordinate and leverage each other's experience. The IESP will explore the possibility of using its wiki as the basis for creating a dictionary of terms and concepts for extreme scale computing.
- *Develop a process for maintaining the IESP Stack Survey*: The on-line publication, in early June, of the IESP survey of current software stacks at HPC centers worldwide, represents a new level of shared awareness and cooperation within the global Computational Science community. Meeting participants generally agreed that establishing and maintaining this baseline of comparison for the deployment of software infrastructure was highly valuable, and its publication after the meeting was delayed only because there were outstanding results from a few centers needed to make it relatively complete.
- *Share lessons learned on NDA process*: Since close collaboration with hardware vendors is an integral part of successful co-design, one obvious, non-technical impediment to a successful outcome is the need for non-disclosure agreements (NDAs) to protect the companies' intellectual property. To help address this problem, Bill Kramer of NCSA volunteered at the San Francisco meeting to write an analysis and lessons learned for the NDA process used with IBM in the development and purchase of Blue Waters. That document is now in hand. We plan to make it available to the community and examine ways to build on it before the next meeting.
- *Update and publish a new iteration of the IESP Roadmap*: Several of the results presented above, as well as other material from the breakout sessions and contributions from individuals, need to be incorporated into a new iteration of the IESP Roadmap. Although there have been contributions and suggested changes to the technical sections, the most substantial changes will be to the sections on IESP organization and co-design. The co-design section will need to be completely rewritten.

5. Conclusion and future meetings

The IESP meeting in San Francisco was remarkably productive, especially given the fact that it was originally scheduled for Kyoto, Japan, and had to be moved at the last minute in the aftermath of the earthquake, tidal wave, and Fukushima nuclear disaster that struck Japan in early March. Arrangements for a suitable new venue had to be made in haste, travel arrangements of the attendees had to be quickly reshuffled, and the agenda had to be changed to account for scheduled participants who had to cancel, and new participants who could now attend because of the new location. Final versions of the meeting agenda and attendee list are attached to this report.

Attendees at the San Francisco meeting agreed that the next IESP meeting would be held in the fall of 2011 in Cologne, Germany. If the established pattern were followed for the subsequent meeting in the spring of 2012, it would be held in the USA. But in view of the fact that, because of the cancellation of the Kyoto meeting, Asian locations have effectively missed an entire cycle, the idea of following the normal order may need to be reconsidered for that meeting.