

## ICS Workshop Position Paper

### Opportunities for engagement between XSEDE, CDI practitioners and tool developers

Phil Blood<sup>1</sup>, Sergiu Sanielevici<sup>1</sup> and Robert Sinkovits<sup>2</sup>

<sup>1</sup> Pittsburgh Supercomputing Center, <sup>2</sup> San Diego Supercomputer Center

The mission of the XSEDE Extended Collaborative Support Services (ECSS) program is to improve the productivity of the XSEDE user community through successful, meaningful collaborations which optimize applications, improve work and data flows, increase effective use of the XSEDE digital infrastructure and broadly expand the XSEDE user base by engaging members of underrepresented communities and domain areas.

Resources allocated through XSEDE include high-performance computing (HPC), high-throughput computing (HTC), visualization, cloud and storage systems. The ECSS staff bring expertise in performance tuning, software parallelization, the efficient use of accelerators, I/O optimization, data analytics, workflows, science gateway development, machine learning and other areas. The ultimate goal is not only to help users make best use of XSEDE resources, but also to transfer knowledge of best practices in computationally intensive and data intensive research.

XSEDE is looking to collaborate with other organizations that encounter the same challenges that we do. For example, we need to deal with issues of scale. Since the collaborations require the dedicated involvement of ECSS staff, our impact is directly proportional to the size of the ECSS workforce and hence our budget. We confront hurdles related to the availability and usability of tools for profiling, tracing and debugging applications and to the limitations of optimizing, vectorizing and auto-parallelizing compilers. And finally, we face challenges in the dissemination of knowledge and information. Ideally each ECSS project will result not only in better code, but the transfer of knowledge to the XSEDE researcher and her team. Even in the best-case scenario, this only impacts a limited number of users.

We are interested in tools that will help our staff become more productive and therefore increase the number of projects that they can complete. Tools with shallower learning curves will also make it easier for us to educate our XSEDE PIs and their teams in the use of the tools. Finally, better tools can also help us to have better outcomes for individual projects, for example, by identifying opportunities for optimization or parallelization that may have otherwise been missed.

In addition to improved tools, we would also like to explore better ways to disseminate the information that was gained through the ECSS projects. We are starting to take advantage of the XSEDE Technical Report series, which gives our staff an opportunity to publish findings and results that are not necessarily a good fit for journal publications or conference proceedings. For example, these reports may go into the low-level details illustrating how code was modified to achieve higher performance or improved scalability. These technical reports can capture valuable information that might otherwise have been lost. We are open to the idea of working with other organizations to expand the scope beyond ECSS and create a broader repository that can be used as a starting point for others who need to carry out similar work.