

the Exascale Report™



The voice of the emerging exascale community

<http://TheExascaleReport.com>
December 2012

2013: Shaping Up To Be a Pivotal Year

A letter to the emerging exascale community

By: Professor Thomas Sterling, Indiana University

As we bring 2012 to a close, and celebrate the winter solstice and other events of perceived personal importance, it is incumbent upon us to pause and consider progress towards selected goals and lessons acquired through the processes of their pursuit.

For those of us committed to the advancement of high performance computing, the realization of Exaflops capability is one such selected goal; the means of achievement of which is uncertain and a source of contention among experts.

This past year has seen significant progress towards the discovery of those means in the US through the initiative of the DOE in their on-going Co-design Program and their new X-stack Program, providing a framework for investigation in application and system development towards practical exploitation of usable exaflops computing systems. The terms “practical” and “usable” are key to the long-term goal. There is little doubt that a collection of a few hundred million ALUs can be integrated to form a machine with a raw peak performance of a quintillion floating-point operations per second before the end of this decade. Making it work for some highly visible self-selected applications of suitable regular and static properties is also likely, even using conventional practices that are anticipated to be extended in to that time frame.

However, devising systems able to support a broad diversity of essential applications that do not conform to these favorable properties and that are representable in some schema at sufficiently low power and high reliability is far more challenging. But this is so far merely an assertion and has not been adequately demonstrated.

This next year demands the delivery by research scientists of a demonstration of the requirements that have to be satisfied for effective exaflops computing and a determination of and if conventional methods can or cannot achieve them. Such a result will permit the entire community to work together towards a commonly recognized goal rather than to continue to engage at cross purposes in diverse strategies.

Perhaps it is too early to make this determination, but at some point it will be essential to guide the development of future hardware designs, software system development, and programming model specification. Towards this end are needed new performance models and metrics that might inform future directions and establish shared basis for consideration and comparison of alternatives. These combined with detailed assessment of real-world applications and perhaps the derivation of new parallel algorithms may yield the necessary and probably new methodology for exaflops computing systems and applications.

At the risk of appearing naïve, I am sanguine that the foundation established by the DOE programs will lead to the deep understanding required to move forward with confidence and alacrity towards a mutually defensible shared goal. At the outset, 2013 holds the promise of being the pivotal period in this quest.

Happy New Year to all of my colleagues and partners in this joint enterprise.

Thomas Sterling

Professor, School of Informatics and Computing

Chief Scientist and Associate Director, CREST

Indiana University



© 2012, The Exascale Report™