VisiSoft Meets These Needs!

Horst Simon on the HPC Slowdown (February 13, 2015)

At an HPC meetup event in San Francisco on Feb 10, Berkeley Lab Deputy Director Horst Simon makes the case that Moore's law and parallelism can no longer be counted on to provide the exponential growth that has been driving high-performance computing for six decades.

Horst gives measurable evidence of his claim that Moore's law is running out of steam ...

If indeed Moore's law is coming to an end, there will be a need for new architectures and new technologies, ... in June 2008 something happened to cause a leveling out of the slope of this extrapolation. A similar break point also appears in June 2013.

... this five-year span marks a turning point where the growth attributed to Moore's law and parallelism are no longer there. It's a case that is supported further by the lack of turnover in the top ten machines, with this grouping remaining virtually unchanged for two years.

There are ... steep technical challenges ... such as overcoming data bottlenecks and power constraints. ... "the only thing standing between us and an exascale machine is a lot of money – billion of dollars of investment and maybe a power bill of $50-100 million a year."

Is the U.S. Falling Behind in Supercomputing and Exascale? (January 29, 2015)

By: John Russell

Few dispute the importance of supercomputing to U.S. competitiveness. The argument is around whether current government efforts – primarily through the Advanced Scientific Computing Research (ASCR) program within the U.S. Department of Energy (DOE) – are effective and sufficient or wasteful and excessive.

... a House of Representatives hearing (Subcommittee on Energy – Supercomputing and American Technology Leadership) argued the decline in U.S. supercomputer research ... the United States has recently dropped from first to seventh place in the world. ... within about five years China will surpass the U.S. in both research funding as a fraction of GDP and absolute funding ... ..

At least as interesting ... was discussion around key technology challenges including the importance of co-design principles (simultaneously algorithm, software and hardware development) in supercomputing

Technology transfer was ... cited. Moving applications and technology out from national supercomputing centers into the mainstream can be challenging. ...

Here are a few key points of the bill:

... partner with universities, National Laboratories, and industry

... include ... any computer technologies that show promise of substantial reductions in power requirements and substantial gains in parallelism of multicore processors, concurrency, memory and storage, bandwidth, and reliability.

New Report Explores the Value of Government Leadership in Supercomputing for Industrial Competitiveness (October 02, 2014)

(Washington, DC) The **Council on Competitiveness** today released *... the Benefits of Supercomputing Investment for U.S. Industry*.. The report ... explores what actions are likely to unleash greater industrial competitiveness.

Continued below ...

VisiSoft Meets These Needs!

Key Findings:

1. More than one-third of U.S. industry representatives surveyed claim their most demanding high performance computing (HPC) applications could utilize 1,000-fold increases in computing capability over the next five years;

2. Software scalability is the #1 most significant limiting factor to achieve the next 10-fold improvement in performance, and it is the #2 most significant limiting factor to reach a 1,000-fold improvement;

3. ... An overwhelming majority of respondents believe HPC is a matter of competitive survival and is critical to the future direction of their businesses. ... many respondents note that links between government and industry need to be strengthened.

"... The challenge to lead in HPC is complex and continuous, but essential to America's economy, security and innovative capacity."

... **Steven E. Koonin,** ... former Undersecretary of Science at the U.S. Department of Energy, added, "Through modeling, simulation and analysis, HPC paves a path for many entities - companies, cities, or individuals to do business in new ways, develop revolutionary products, offer services, and interact in ways that improve everything from health and safety to productivity and entertainment."

NSF PROGRAM SOLICITATION NSF 145-16 (February 2014)

Synopsis of Program:

> ... semiconductor technology is facing fundamental physical limits and single processor performance has plateaued ... parallelism has become critically important.
>
> The Exploiting Parallelism and Scalability (XPS) program aims to support groundbreaking research leading to a new era of parallel computing. Achieving the needed breakthroughs will require a collaborative effort among researchers ... and will be built on new concepts, theories, and foundational principles. New approaches to achieve scalable performance and usability need new abstract models and algorithms, new programming models and languages, new hardware architectures, compilers, operating systems and run-time systems, and must exploit domain and application-specific knowledge. Research is also needed on energy efficiency, communication efficiency, ...

INTRODUCTION

... Parallelism has become critically important at many levels, ... Multi- and many-core processors, ... and the software environments harnessing these resources comprise a new cyber infrastructure enabling a new set of global applications with tremendous economic and social impact. ... At the same time, a main driver of continued performance improvement is ending: semiconductor technology is facing fundamental physical limits and single processor performance has plateaued.

To continue improving performance, we need a new era of parallel computing, driven by novel, groundbreaking research in all areas impacting parallel performance and scalability.

PROGRAM DESCRIPTION

... New approaches should encompass both software and hardware to achieve scalable performance and usability through new abstract models and algorithms, programming models and languages, data models and declarative query languages, hardware architectures, compilers and runtime systems.

Research on foundational principles should engender a paradigm shift in the ways in which one conceives, develops, analyzes, and uses parallel algorithms, languages, and concurrency. Foundational research should be guided by crucial design principles and constraints impacting these principles.

VisiSoft Meets These Needs!

The Right Direction For Change

As stated below, top computer engineers describe the difficulty encountered when using current software approaches, and the direction needed for change.

Justin Rattner (former Chief Technical Officer, Intel),
"As hardware technology approaches the terascale level on the desktop, software has fallen further behind. -- One result has been a lack of parallel programming applications to leverage dual-and multi-core processing technology. Intel is looking for new languages for programming in parallel."

Chuck Moore (formerly Chief Technology Officer, AMD)
"The industry is in a little bit of a panic about how to program multi-core processors, especially heterogeneous ones. To make effective use of multi-core hardware today you need a PhD in computer science. That can't continue if we want to enable heterogeneous CPUs."

Craig Mundie (Chief Research & Strategy Officer, Microsoft)
"To maximize computing horsepower, software makers will need to change how software programmers work. Only a handful of programmers in the world know how to write software code to divide computing tasks into chunks that can be processed at the same time instead of a traditional, linear, one-job-at-a-time approach. -- A new programming language will be required, and could affect how almost every piece of software is written. --- This problem will be hard."

It is apparent from the above statements that a new software technology is required. Trying to see how many CPUs can be crammed into a computer is not the problem. The real economic requirement is for: ***a software approach that minimizes the number of processors needed to meet the runtime constraints for a given application.***

This problem has been described by many of the most experienced people in the field. They clearly state that the current approach to building software - using C-based languages and OOP - is a huge step backward. In the original designers' own words, C was never intended to be a real programming language, and the band-aids that have been added do not solve the problem.

Why can't large complex software systems be decomposed into independent parts, to run concurrently on relatively simple parallel processor chips? The answer is "They can!" But, as stated above by top engineers, this requires getting out of the current software rut in which we are stuck. It implies the acceptance of a totally new approach to designing software as provided with VisiSoft, a CAD system for software, one whose economic measures are obvious and overwhelming.

VisiSoft Economic Multipliers

Many High Performance Computer (HPC) users, e.g., weather forecasters, are complaining that they cannot add more detail to improve the accuracy of their models. They are at their limit of running times because of the need for more processing power.

Current HPC Processor Utilization Efficiencies (PUEs) are typically 7% to 10%, taking 100 processors to achieve a speed multiplier of 7 to 10. This coupled with their requirements for high power, air conditioning, and floor space contribute to the high cost of running HPC centers.

VisiSoft solves these problems.

When software is written in VisiSoft - versus a C-based language or FORTRAN, one can achieve speed increases of 10 to 100 - just on a single processor - in a PC.

Using VisiSoft on a parallel processor, one can get Processor Utilization Efficiencies (PUEs) that range from 75% to 95% versus the current HPC PUEs of 7% to 10%. This is another factor of 10 faster using the same number of processors.

Using the lower end of the 100 to 1000 times speed range, this implies that a 32 processor PC could potentially run faster than a 3000 processor HPC.

Now consider the footprint of a 3000 processor HPC - with 30 boards in a rack compared to a single PC board. This presents another factor of 10 speed difference due to memory boundary crossing delays.

This presents a total potential speed multiplier that is 10,000 times faster.

Private organizations are concerned about economics. To keep their costs down, they likely want to minimize the number of processors required to meet a speed constraint.

Let's assume we only get a factor of 100 instead of 10,000. This still implies that a 32 processor PC will beat a 3000 processor HPC in speed - a huge competitive economic position.

Speaking of costs, subject area experts can use the VisiSoft CAD system directly, without help from programmers. Software can be built 2 to 5 times faster, and upgraded easily compared to all other approaches.

Various experiments have already been run to prove all this. It's time to repeat them and prove the substantial value hidden in this disruptive technology. The economics are obvious.

We must get people to:

Repeat the Experiments & Compare the Results!

Visit our website at: www.VisiSoft.com