


精彩故事：稀疏矩阵乘法在 Xeon Phi 和其他架构下的性能比较； HPC 能效的综述； OpenMP、MPI 和混合科学应用中的弱可扩展性能模型；随时随地的集群监控的探索；一个数据密集型云存储的框架

March 21, 2013

The Week in HPC Research

Tiffany Trader



The top research stories of the week have been hand-selected from leading scientific centers, prominent journals and relevant conference proceedings. Here's another diverse set of items, including an evaluation of sparse matrix multiplication performance on Xeon Phi versus four other architectures; a survey of HPC energy efficiency; performance modeling of OpenMP, MPI and hybrid scientific applications using weak scaling; an exploration of anywhere, anytime cluster monitoring; and a framework for data-intensive cloud storage.

Evaluating Sparse Matrix Multiplication Kernels on Intel Xeon Phi

The Intel Xeon Phi made a big splash at SC12, and computer scientists are eager to put the coprocessor through its paces. Such is the case with a team of researchers from the Ohio State University, who authored a recent [paper](#), describing their work evaluating sparse matrix multiplication kernels on the Intel Xeon Phi.



As the team notes, the Phi sports 61 cores, each supporting 4 hardware threads with 512-bit wide SIMD registers for a theoretical peak performance of 1 teraflops double precision.

Their paper is meant to serve as an introduction to the Phi architecture and to analyze its peak performance using the sparse matrix as a test application. It's a good choice to test the Phi's capabilities because it is representative of many large-scale applications and because it is a difficult problem for coprocessor architectures.

As the team writes: "Many scientific applications involve operations on large sparse matrices such as linear solvers, eigensolver, and graph mining algorithms. The core of most of these applications involves the multiplication of a large, sparse matrix with a dense vector (SpMV)."

They also note that "the irregularity and sparsity of SpMV-like kernels create several problems for these architectures [i.e. accelerators/coprocessors]."

The researchers compared the sparse matrix multiplication performance of Xeon Phi with four other architectures: two dual Intel Xeon processors, X5680 (Westmere) and E5-2670 (Sandy Bridge), as well as two NVIDIA Tesla GPUs C2050 and K20. Their results of their experiment show that the Phi offered superior performance.

They write that "although the design of a Xeon Phi core is not much different than those of the cores in modern processors, its large number of cores and hyperthreading capability allow many application to saturate the available memory bandwidth, which is not the case for many cutting-edge processors. Yet, our performance studies show that it is the memory latency not the bandwidth which creates a bottleneck for SpMV on this architecture. Finally, our experiments show that Xeon Phi's sparse kernel performance is very promising and even better than that of cutting-edge general purpose processors and GPUs."

Energy Awareness in HPC: A Survey

A group of researchers from the Walchand College of Engineering, in the city of Sangli, Maharashtra, India, have published a [paper](#) addressing one of the most pressing problems in high-performance computing: energy-efficiency.



Figure 1. Taxonomy of Research Methodology in HPC

The team sets out by acknowledging the increased awareness of energy and costs associated with power management for high performance computing. They write that “power control is becoming a key challenge for effectively operating a modern high end computing infrastructures such as server, clusters, data centers and grids,” although the scope of the paper is primarily concerned with cluster systems.

The researchers argue that developing energy efficient computer designs is the next major goal of the high performance computing. The paper presents a survey and classification of energy efficient techniques for cluster computing. The research outlines both hardware and software related variables and sub-classes thereof. An important point made in the paper is that performance itself does not become a secondary objective but it is understood that power is a constraint to increasing performance.

Performance Modeling of Hybrid MPI/OpenMP Applications at Scale

Texas A&M University computer scientists Xingfu Wu and Valerie Taylor are exploring a performance modeling framework based on memory bandwidth contention time and a parameterized communication model. They have co-authored a [paper](#) describing their work with modeling and predicting the performance of OpenMP, MPI and hybrid scientific applications using weak scaling on large-scale multicore supercomputers.

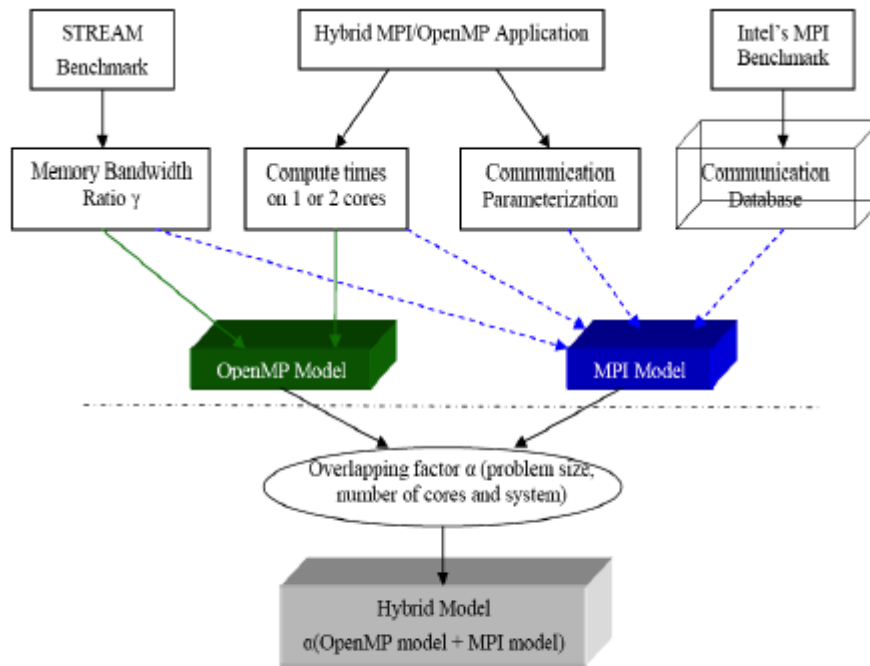


Figure 4. Framework for modeling hybrid MPI/OpenMP programs

The research team employed STREAM memory benchmarks to identify initial performance and model validation of MPI and OpenMP applications. They also used the hybrid large-scale scientific application Gyrokinetic Toroidal Code in magnetic fusion to validate the performance model.

The experiment used three different supercomputers: an IBM POWER4, POWER5+ and BlueGene/P. Study results showed an error rate of less than 7.77% for predicting the performance of hybrid MPI/OpenMP GTC on up to 512 cores on these multicore systems.

Anywhere Anytime Cluster Monitoring

A trio of computer scientists from Shandong University in Jinan, China, are exploring the feasibility of anywhere, anytime cluster monitoring. More specifically, they are working to design and implement a cluster monitoring system based on Android.




Fig. 1. The system structure

The team starts with the view that high performance computing (HPC) has been democratized to the point that HPC clusters have become an important resource for many scientific fields, including graphics, biology, physics, climate research, and many others. Still, depending on local funding realities, the availability of such machines is almost universally constrained. In light of this, monitoring becomes an essential task necessary for the efficient utilization and management of limited resources. However, as the researchers observe, traditional cluster monitoring systems demonstrate poor mobility, which stymies proper management.

The authors are seeking to improve the flexibility of monitoring systems and improve the communication between administrators. They assert that the mobile cluster monitoring system outlined in their paper "will make it possible to monitor the whole cluster anywhere and anytime to allow administrators to manage, diagnose, and troubleshoot cluster issues more accurately and promptly."

The system they developed is based on the Android platform, the brainchild of Google, and built on open source monitoring tools, Gaglia and Nagios. The design uses a client-server model, where the server probes the data via monitoring tools and produces a global view of the data. The mobile client gets the monitoring packages by Socket. Then, the cluster's status is displayed on the Android application.

Their work was published as a chapter in the book, [Pervasive Computing and the Networked World](#).

A Framework for Cloud Storage

UK computer scientists Victor Chang, Robert John Walters and Gary Wills set out to explore the topic of cloud storage and bioinformatics in a private cloud deployment. They've written a [paper](#) about their experience to serve as a resource for other researchers with data-intensive compute needs who are interested in analyzing the benefits of a cloud model.




Figure 3: Selected brain imaging

Among the many benefits of the cloud model are its cost-savings potential, agility, efficiency, resource consolidation, business opportunities and possible energy savings. Despite the inherent attractiveness, there are still barriers to overcome, and one of these, according to the authors is the need for a standard or framework to manage both operations and IT services.

They write that "this framework needs to provide the structure necessary to ensure any cloud implementation meets the business needs of industry and academia and include recommendations of best practices which can be adapted for different domains and platforms."

Their work examines service portability for a private cloud deployment. Storage, backup and data migration and data recovery are all addressed. The paper presents a detailed case study about cloud storage and bioinformatics services developed as part of the Cloud Computing Adoption Framework (CCAF). In order to illustrate the benefits of CCAF the authors provide several bioinformatics examples, including tumor modeling, brain imaging, insulin molecules and simulations for medical training. They believe that their proposed solution offers cost reduction, time-savings and user friendliness.