Delivering High-performance Computing Solutions to Business and Financial Engineering

Yingwen Song

Information Technology Research Institute (ITRI),
National Institute of Advanced Industrial Science and Technology (AIST)
AIST Central 2, 1-1-1 Umezono, Tsukuba, Ibaraki 305-8568 Japan
E-mail: yw-song@aist.go.jp

Abstract

==========

In order to understand the relationship between risks and profitability, researchers and practitioners in business and financial engineering are forced to use complicated models to solve computationally intensive problems with HPC (High Performance Computing) engines. The HPC engines enable them to analyze and run simulation models on complex portfolios and price options, detect fraud, and predict currency shifts. However, it is hard to say that they can use HPC technology efficiently because both fields are traditionally different. The most challenging issue in HPC is the development of applications utilizing the characteristics of platforms due to the wide varieties of computer architectures. Therefore, we will focus this talk on how to solve this most pressing problem.

We start with writing applications to take advantage of the hierarchical structure of memory for conventional scalar CPUs, present some key points in developing coarse-grained OpenMP applications for Multi-core processors and SMPs (Symmetric Multi-Processing machines), and highlight the strategies to achieve load balance and to minimize the communication latency for MPI applications on cluster-based machines. These days, people are pursuing computational power from novel parallel architectures like CBE (Cell Broadband Engine) and CUDA (Compute Unified Device Architecture). We then will present how to achieve performance on CBE and CUDA platforms, attempt to compare their merits and difficulties. Since grid computing is a well-established approach for ultra-scale computation, we will discuss the strategies to achieve fault tolerance for large-scale and long-time computation in a grid environment in the next part. Cloud computing is now a new opportunity to implement low-cost, low-power and high-efficiency systems to deliver scalable infrastructure, we will then discuss the concept, usability, and performance issues of cloud computing in the last part. We finally close this talk with an outlook to future marriage of HPC and business including financial engineering.

Briefly, this talk will discuss the characteristics of platforms ranging from conventional systems to cutting-edge architectures, provide suggestions to develop performance-driving applications in the fields of business and financial engineering.

Short Biography

==========

Dr. Yingwen Song is a research scientist at AIST, Japan. He received his Ph.D. from Saitama University, Japan in 1999. Before he joined AIST, he had been working as a HPC researcher at INRIA, France and The University of Tokyo, Japan. He also has over 3 years industrial experience in developing, parallelizing, and tuning applications. His research interests cover a wide range, particularly include porting and tuning of HPC applications, the use of novel and advanced-computer architectures, grid computing, as well as middleware development. He has contributed over 50 technical papers to various conferences and refereed journals.