



UMIACS Seminar Series: Dr. John Feo

Requirements and Performance of Data Intensive, Irregular Applications

When: 29 Sept, 2010, at 2pm.

Where: Room 3258 AVW

Abstract: Many fundamental science, national security, and business applications need to process large volumes of irregular, unstructured data. Data collection and analysis is rapidly changing the way scientific, national security, and economic communities operate. There are worldwide operational deployments of instruments to detect the proliferation of weapons of mass destruction, monitor terrorist cells, and track the movement of illicit goods and services. In the next 15 years 30% of battle-space defense forces will be autonomous with each advanced robotic device carrying dozens of sophisticated sensors collecting, processing, analyzing and transmitting large amounts of data. American economic competitiveness will depend increasingly on the timely analysis of many Petabytes of data collected in diverse computing clouds charting the social and economic behavior of consumers.

Unlike traditional scientific applications based on linear algebra routines, data analytic applications comprise large, integer-based graph computations with irregular data access patterns, low computation to memory access ratios, and high levels of fine grain parallelism that pass data and synchronize frequently. Traditional architectures optimized to run large-scale floating point intensive simulations are inadequate, and more suitable high-end architectures such as the Cray XMT are needed. In this talk I will discuss the implications of the emerging field of data analytics on high-performance computer systems. I will survey the research at PNNL's Center for Adaptive Supercomputer Software in graph algorithms, software tools, and next generation architecture design. In particular, I will present several key graph algorithms we have developed with an emphasis on structure, use of special hardware features, performance, and scalability.

Vita: Dr. John Feo is the director of the Center for Adaptive Supercomputer Software at the Pacific Northwest Laboratory. Dr. Feo received his Ph.D. in Computer Science from The University of Texas at Austin. He began his career at Lawrence Livermore National Laboratory where he managed the Computer Science Group and was the principal investigator of the Sisal Language Project. Dr. Feo then joined Tera Computer Company (now Cray Inc) where he was a principal engineer and product manager for the MTA-1 and MTA-2, the first two generations of the Cray's multithreaded architecture. After a short two year "sabbatical" at Microsoft where he led a software group developing a next-generation virtual reality platform, he joined PNNL

Dr. Feo's research interests are parallel programming, graph algorithms, multithreaded architectures, functional languages, and performance studies. He has published extensively in these fields. He has held academic positions at UC Davis and is an adjunct faculty at Washington State University.