NSF Presidential Faculty Fellow (PFF) award from President Clinton in 1995 and the Governor's Medal for Science and Technology from Governor Michael Leavitt in 1999. In 2003 he received the Distinguished Professor Award from the University of Utah. In 2004 he was elected a Fellow of the American Institute for Medical and Biological Engineering (AIMBE) and in 2005 he was elected a Fellow of the American Association for the Advancement of Science (AAAS).

Session 1: Peer-to-Peer Algorithms

Session 2: Science, Finance and Combinatorial Applications

Session 3: Cluster and Server Architectures

Session 4: Software Support for Large Scale Scientific Computing

Session 5: Scheduling Algorithms

Session 6: Search, Text and Web Applications

Session 7: Processor Architecture

Session 8: Performance Analysis and Optimization

Session 9: Complexity of Algorithms

Session 10: Power and Energy Aware Computing

Session 11: Performance Modeling and Evaluation

Session 12: Middleware and Tools

Symposium Evening Tutorial: High-performance Computing Methods for Computational Genomics
Presenters: Srinivas Aluru, David A. Bader, and Ananth Kalyanaraman

Abstract: As biomolecular sequence data continue to be amassed at unprecedented rates, the design of effective computational methods and capabilities that can derive biologically significant information from them has become both increasingly challenging and imperative. In this tutorial, the audience will be first introduced to the different types of biomolecular sequence data and the wealth of information they encode. Following this technical grounding, high-performance computing approaches developed to address some of the most computationally challenging problems in genomics will be described. The contents will be presented in three parts: (i) In the first part, we will describe methods that were designed to query a sequence against a large sequence database. Two popular parallel approaches, mpiBLAST and ScalaBLAST, implementing the NCBI BLAST suite of programs will be described. (ii) Next, we will describe PaCE, which is a parallel DNA sequence clustering algorithm. As direct applications, we will discuss the clustering of large-scale Expressed Sequence Tag data and the assembly of complex genomes. (iii) Finally, we describe GRAPPA, which is a high-performance software suite developed for phylogenetic reconstruction of a collection of genomes or genes.

Throughout the tutorial, emphasis will be on both scalability and effectiveness in exploiting large-scale state-of-the-art supercomputing technologies. The intended audience are academic and industry researchers, educators, and/or commercial application developers, with a computational background. No background in biology is assumed.

Wednesday, March 28, 2007

Keynote Speech: Avoiding the Memory Bottleneck through Structured Arrays
Michael J. Flynn, Stanford University

Abstract: Basic to parallel program speedup is dealing with memory bandwidth requirements. One solution is an architectural arrangement to stream data across multiple processing elements before storing the result in memory. This MISD type of configuration provides multiple operations per data item fetched from memory. One realization of this streamed approach uses FPGAs. We'll discuss both the general memory problem and some results based on work at Maxeler using FPGAs for acceleration.

Bio: Michael Flynn is Senior Advisor to the Maxeler Corporation, an acceleration solutions company based in London. He received his Ph.D. from Purdue University and joined IBM working there for ten years in the areas of computer organization and design. He was design manager System 360 Model 91 Central Processing Unit. Between 1966 and 1974 Prof.