



Workshop #4

Automatic Tuning of Libraries and Applications

Organizer: Ken Kennedy (Rice University)

For many years, retargeting of applications for new architectures has been a major headache for high performance computation, requiring many person-months (or even years) of effort to retune for each new architecture, and even each new model of an established architecture. Automation of this retuning process has now become a fertile area of computer science research. Most of this work is based on the strategy of using large amounts of computation time to explore a space of different variants of a loop nest, running each variant on the target architecture, and picking the best one. One example of this strategy is the Atlas system, which uses substantive amounts of computation to provide versions of a computational linear algebra kernel that are tuned in advance to different machines. If this approach can be extended more generally to components and whole programs, it would help avoid the enormous human costs involved in retargeting applications to different machines. A major research issue is how to keep tuning time to manageable levels, given that the number of variants in a complete application can be enormous. The purpose of this workshop, which follows a 2005 LACSI Symposium Workshop, is to report on the ongoing research efforts in this area, to solicit feedback from and collaboration with the application development community, and to exchange ideas on the future directions for this work. One specific subgoal will be to initiate an activity to develop a standard set of benchmarks for use in automatic tuning research.