

## **Adapting Convergent Scheduling Using Machine-Learning**

Diego Puppim Mark Stephenson Saman Amarasinghe Martin Martin Una-May  
O'Reilly

To appear at the *16th Workshop on Languages and Compilers for Parallel  
Computing (LCPC03)*, College Station, TX, 2-4 October 2003

### **Abstract**

Convergent scheduling is a general framework for instruction scheduling and cluster assignment for parallel architectures. A convergent scheduler is composed of many independent passes, each of which addresses a particular issue of scheduling on spatial architectures. In addition, each of the passes share a common interface, which allows the passes to be run multiple times, and in any order. Because of this, a convergent scheduler is presented with a vast number of legal phase orderings. In this paper we use machine-learning techniques to automatically search for good phase orderings. Our system found different phase orders for four different VLIW architectures. The architecture-specific phase orderings yield speedups ranging from 12% to 95% (over the order presented in our previous work). Furthermore, cross validation studies that we perform in this work show that our automatically generated orderings perform well beyond the benchmarks on which they were trained.