

Semi-Automatic Microarchitecture Configuration of Soft-Core Systems
Shobana Padmanabhan, John Lockwood, Ron Cytron, Roger Chamberlain, and
Jason Fritts
Washington University
cytron@acm.org

The major goal of our project is to measure and improve application performance, by providing an easily and efficiently reconfigurable architecture along with software support to expedite its use.

Our FPGA-based system is implemented as an extensible hardware module on the Field-programmable Port Extender (FPX) platform at Washington University. The FPX platform provides an environment where a circuit implemented in FPGA hardware can be developed, uploaded, and tested with live data using web and network interfaces.

Our research begins with LEON, a soft-core implementation of a SPARC instruction set. We can easily add instructions and microarchitecture enhancements to LEON and evaluate their performance using a cycle-accurate profiling circuit that we have designed to operate with LEON.

Systems such as LEON contain many parameters that affect the performance of the resulting soft-core system. While the parameters serve to tune the synthesized system's performance, the parameters present a great number of choices that must be made: (data and instruction) cache sizes, line length, and associativity; multiple implementations of features such as floating point; pipeline depth, store-buffer size, and branch-history mechanisms—all of these can dramatically affect a given application's performance. Moreover, the resulting area of the synthesized system is constrained by the FPGA upon which it will be deployed.

The configuration space is relatively large and time-consuming to explore. Standard optimization approaches are unsuitable, as they assume that evaluating a given configuration is relatively inexpensive; for LEON, synthesis of one configuration can take up to 30 minutes on a modern processor.

In this talk, we present an overview of our approach for automating the configuration of LEON for a given application. We present results from hand-tuning cache size for an application and compare that with our automated results. This work is supported by the National Science Foundation under grant ITR-0313203.