The Totem Project

by Katherine Compton and Scott Hauck

The Totem Project (http://acme.ee.washington.edu/totem) at the University of Washington and Northwestern University focuses on the automatic generation of reconfigurable hardware customized for a particular domain or set of applications. These application domains could include cryptography, DSP (or a subdomain), specific scientific data analysis, or any other compute-intensive area. The target environment for this hardware is system-on-a-chip (SoC) design, where the need for custom fabrication provides the opportunity for tailoring the reconfigurable architecture to the intended uses of the SoC. We expect perhaps 10% - 20% of an SoC's die area to be dedicated to this type of reconfigurable logic.

Reconfigurable hardware is ideal for use in SoCs as it executes applications in hardware instead of software, and yet maintains a level of flexibility not available with more traditional custom circuitry. This flexibility allows for both hardware reuse and post-fabrication modification. Hardware reuse allows a single reconfigurable architecture to implement many potential applications, rather than requiring a separate custom circuit for each. Furthermore, post-fabrication modification can allow for alterations in the target applications, bug-fixes, and reuse of the SoC across multiple similar deployments to amortize design costs.

A widely available form of reconfigurable hardware is the generic field-programmable gate array (FPGA), and the reconfigurable hardware within an SoC could be patterned after these designs. However, FPGAs tend to be somewhat fine-grained, target the general case, and sacrifice performance and area in order to achieve a high degree of flexibility. While this flexibility has its place for situations where the computational requirements are either not known in advance or vary considerably among the needed applications, in many cases this extreme level of flexibility is unnecessary and would result in significant area, delay, and power overheads.

Instead of fully generic FPGAs, if any characteristics of the target application set are known in advance, they can be used to customize the reconfigurable hardware—removing unneeded flexibility in order to reduce area and power requirements and increase performance. Essentially, depending on the needs of the algorithms and the designers, these custom architectures can vary in terms of flexibility from a design close to an ASIC, to something close to a traditional FPGA. Very constrained computations would be primarily fixed ASIC logic, while extremely unconstrained domains would require near-FPGA functionality.

Specialized reconfigurable architectures, while beneficial in theory, would be impractical in practice if they had to be created by hand for each group of applications. Instead, the Totem Project focuses on the automatic creation of customized reconfigurable architectures, including high-level design, VLSI layout, and associated custom place and route tools. Architecture Generation thus far has included methods for the automatic

design of near-ASIC as well as more flexible routing structures, showing significant improvement over existing systems.

The work on the automatic layout engine is concentrating on multiple different methods, including standard-cell (with specialized cells), template reduction (removing unneeded circuitry from a full-custom template), and circuit generators. Meanwhile, the place and route tool provides architecture-adaptive routing, with similar placement techniques in development. Pipelined routing, a key to very high-performance FPGAs, is also being developed.

Efforts are underway to develop domain-specific synthesis systems for 2D island-style FPGAs, product-term PALs and CPLDs, and 1D datapath-oriented reconfigurable systems. We believe these techniques will achieve nearly ASIC-level area, power, and performance, yet with the flexibility of FPGAs for a given domain.