Reconfigurable Logic for Systems on a Chip

W. Shields Neely
National Semiconductor
w.shields.neely@nsc.com

Abstract

The electronic systems of the future will be implemented in terms of multi-million gate "systems on a chip". These systems will require an enormous investment in design and manufacturing; yet the pace of technological change (e.g., new algorithm development, new processor and memory designs) and ever changing requirements puts them in danger of obsolescence soon after they are created – applications always want to take advantage of new technical advances and must meet changed requirements. What is needed are single chip systems that are *designed to be adaptable to a family of applications*. The emerging technology of configurable logic offers the promise of large-scale silicon systems that are adaptive after manufacture, with little or no sacrifice in execution efficiency compared to hard-wired systems.

But the application of configurable logic to single chip systems requires a radical rethinking of the approach, with new architectures and design tools to address the issues of complexity and scale. No longer will configurable logic systems consist entirely or even significantly (in terms of silicon area) of FPGA's. Instead, adaptive single chip systems must be designed with just the right amount of configurable logic in just the right places.

In order to deal with the complexity of designing and manufacturing multi-million gate silicon systems, we must design in terms of pre-designed, well-understood, highly optimized blocks of fixed logic functions. Configurable logic will never have the area efficiency of optimally placed fixed logic used in VLSI chips due to the overhead of routing paths and the size of programmable structures. On the other hand, configurability provides a critical degree of flexibility while retaining far better efficiency than is possible through software work-arounds. By mixing blocks of fixed logic functions of the correct granularity with programmable components, we can construct silicon systems that approach the area efficiency of standard VLSI while retaining many of the advantages of configurable logic systems.

Creating the best balance of fixed and configurable logic for a given application is a challenging systems design problem that can only be addressed through radically new architectures and the design environments to create and support them. These environments must support exploration and trade-off of design dimensions such as hardware-software partitioning, bus architecture and partitioning, and circuit size/power versus speed. Such a methodology will enable commonality of chips for different applications within a given domain, allow post-production bug fixes, and support dynamic reconfiguration for different modes of operation.

The system design problem for recofigurable systems is a function of the configuration requirements, which vary by application. Both static configuration, which occurs at power up time, and is used to perform optioning, such as bus interface types, and modal optioning, in which a custom co-processor is configured for each major mode of operation, such as DSP function selection, can provide significant post-layout benefits with even small amounts of configurable logic placed strategically with an architecture. For example, a video graphics processor with a configurable bus interface could be used to replace an existing family of parts. Dynamically configured systems, in which custom instructions are loaded for time critical tasks, can provide lower power solutions for signal processing applications. For example, the National Napa1000 Reconfigurable Processor, which provides a 32 bit RISC core and a adaptive logic processor on a single chip, is a step in the direction of mixed configurable and fixed logic for signal processing.