Packet Switched vs. Time Multiplexed FPGA Overlay Networks

Nachiket Kapre, Nikil Mehta, Michael deLorimier, Raphael Rubin, Henry Barnor, Michael J. Wilson, Michael Wrighton, André DeHon

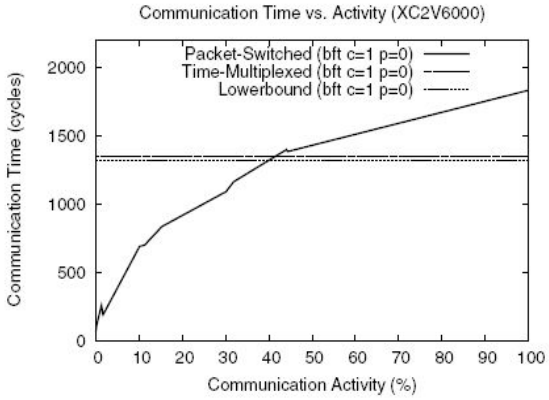

Year of publication: *2006*
Area: *Architecture and Technology*

This paper provided a study of two different high-level approaches to interconnecting complex processing elements on an FPGA, a compile-time-scheduled, time-multiplexed interconnect and a packet-switched interconnect with communication patterns determined at run-time.

One of the distinguishing features of field-programmable gate arrays is the user's ability to customize FPGA resources to exactly the functionality that is required. Since the introduction of the first commercial FPGAs, on-chip interconnect, rather than logic, has been the limiting resource in most devices. However, until the mid-2000s, most reconfigurable computing research projects focused on the best ways to customize computation for available logic resources, while leaving its interconnection as an afterthought.

The advent of the network-on-chip era of inter-processor communication brought greater focus onto the interconnection of FPGA processing elements, such as processor cores or multi-LUT functional blocks. This paper contrasts two distinctly different approaches to providing high-level on-chip interconnection between FPGA processing elements. One approach involves a compile-time analysis of all inter-element communication. A static communication schedule for each inter-element router is then determined. A second approach assumes that a router has sufficient knowledge to dynamically determine a data packet's route as the application executes. Although this second approach often requires more router hardware than the first, in some cases when a large communication schedule must be stored, the hardware requirements for the statically-scheduled approach may be greater than those of the packet-switched. The paper fully evaluates the tradeoffs for a series of graph-based applications.

The paper provides a thought-provoking analysis of high-level FPGA interconnection models. Since its publication numerous follow-on papers have examined both improved FPGA network-on-chip architectures and custom interconnect architectures for statically-scheduled on-chip communication.

Russell Tessier

DOI: http://dx.doi.org/10.1109/FCCM.2006.55