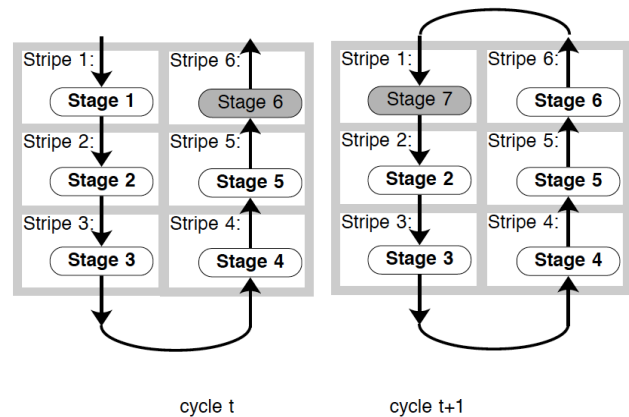


Managing Pipeline-Reconfigurable FPGAs

Srihari Cadambi, Jeffrey Weener, Seth Copen Goldstein, Herman Schmit, Donald E. Thomas

Year of publication: 1998

Area: *Reconfigurable Computing*



How can an FPGA design scale with Moore's Law? This work introduced both a computational model abstracted from the particular size of a specific FPGA and relatively complete details on how you would go about implementing an FPGA to support the model.

In 1998, the largest FPGAs had only 50,000 LUTs; it was not uncommon for users to struggle with designs that would not fit on a single FPGA. Furthermore, even if your design did fit at a particular FPGA generation, when the next generation of FPGAs was introduced, it was necessary to redesign, or at least recompile, the design to exploit the greater capacity available in the next generation parts. While today's million LUT FPGAs make the first problem less of a concern, the lack of compatibility and scalability across generations remains.

PipeRench provided the developer with a logical rather than physical model for run-time reconfiguration. It views the computation as a large, static, feed-forward graph and uses run-time reconfiguration behind-the-scenes to swap different *pieces* of a circuit in and out when that circuit would otherwise be too large to fit in the available hardware. Uniquely, PipeRench showed how it was possible to intimately overlap (pipeline) reconfiguration with computation, allowing rapid configuration switching with modest bandwidth to configuration memory external to the array. The PipeRench model allowed a bitstream compiled for one instance of PipeRench to be loaded into a larger (or smaller) instance of PipeRench, with a corresponding increase (or decrease) in performance.

The PipeRench effort was a pioneering large-scale system project covering a wide range of issues in reconfigurable computing. This paper explains the management of configuration data and presents some of the lower-level design details for the PipeRench platform. It discusses the finer points of the configuration process and mechanisms to ensure data arrives at the appropriate physical location where a virtual pipeline stage is mapped. Anyone beginning a reconfigurable architecture project should read this and other PipeRench papers to get a feel for the wide variety of design problems they may encounter and some innovative ways to solve them.

Katherine Compton and André DeHon

DOI: <http://doi.acm.org/10.1145/275107.275120>