A Time-Multiplexed FPGA

Steve Trimberger, Dean Carberry, Anders Johnson, Jennifer Wong

Year of publication: *1997*
Area: *Architecture and Technology*

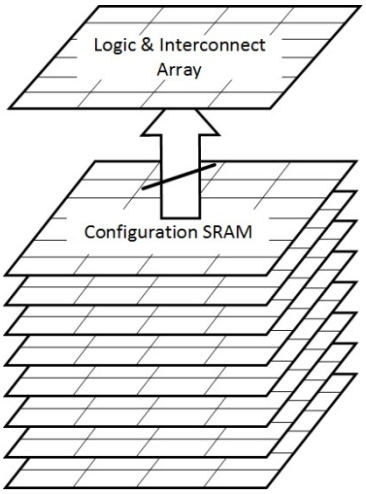

In the early 1990's, it was challenging to fit designs onto the small FPGAs of the day, and several research groups [(Bhat 1993), (DeHon 1995), (Tau 1995)] were exploring rapid reconfiguration capability. The reconfigurable mesh model, which was proposed in the 1980's offered a theoretical view, but no architectural mechanisms to support it. While the community was debating dynamic reconfiguration and the potential value of such a mechanism for application developers, this paper offered a feasible approach to realize such a model in hardware, illustrating potential benefits and some applications of such a feature. The design offered a small number of configurations to be stored on the device and provided architectural support for the device to be rapidly reconfigured by switching between these configurations. The configuration memory itself could be used as on chip storage, predating the current Block RAM (BRAM) idea. This paper pushed the boundary (at that time) by showing programmable devices that could do more than one task rapidly.

The key concept is to enable dynamic re-use of the hardware. The paper proposed to store eight configurations in distributed SRAM on the device. These configurations can be used to configure the logic as well as the interconnects. The authors proposed three modes of operation. In logic engine mode, a single large design can be emulated using multiple smaller configurations. In the time-share mode, designs with multiple communicating FPGAs can be emulated. Finally, they provide a static mode in which the logic is resident on the device all the time and is not reconfigured. The architecture also provided the capability to mix these modes to improve application performance. The configuration memory could also be written from on-chip logic, a feature that later showed up in Xilinx Virtex FPGAs with the ICAP interface. The paper provided detailed designs to enable these capabilities.

The paper stimulated thinking both from the architecture community to realize dynamic reconfiguration in hardware and from the applications community to contemplate novel uses of FPGAs if rapid reconfiguration was feasible. The exposition is crisp, and it makes a good reading in FPGA architectures as well as application acceleration courses. However, the design did not lead to time-multiplexed FPGAs from Xilinx---Moore's Law scaling remained their preferred path to increasing device capacities.

Viktor K Prasanna

DOI: http://dx.doi.org/10.1109/FPGA.1997.624601