Programmable Active Memories: Reconfigurable Systems Come of Age

Jean E. Vuillemin, Patrice Bertin, Didier Roncin, Mark Shand, Hervé Touati, and Philippe Boucard

Published: *IEEE Transactions on Very Large Scale Integration Systems, vol. 4, no. 1, pp. 56--69, March 1996.*

The DEC Programmable Active Memories project was one of the first efforts to explore the use of FPGAs for computing. The PAM work started in 1987, just a year or two after the first commercial FPGAs were available and six years before the first IEEE FCCM Workshop. This pioneering work set the stage for the whole field of Reconfigurable Computing that was to come.

This article comes relatively late in the development of PAM (4 years after the main work), providing a comprehensive look back on the work. As such this is the single, must-read paper that describes the entire project and components. A significant piece of the paper describes quantified speedups on a broad set of 11 distinct applications, demonstrating the power of this new computing platform. This was particularly significant at the beginning when it was unclear if these reconfigurable computing platforms really had broad reach, or if they were just lucky on a single problem or two or a small application domain. Their applications covered cryptography, pattern matching, compression, data coding, audio and video processing, physical simulations, and neural networks---a set of application domains where FPGA-based computing has thrived for the past 20 years. To make programming viable at a time when even Verilog/VHDL synthesis was immature, they developed a programming model and generator language to provide a higher-level view of the platform than a collection of gates.

Using a collection of 23 FPGAs, comparable in cost to a high-end workstation of the day, they were able to show performance exceeding contemporary supercomputers. For real-time tasks, they were able to achieve the same performance as full custom designs. For many tasks, the PAM implementation was the fastest of its day, even considering custom VLSI designs. Their entire 23 FPGA PAM design had only 15,000 4-LUTs and 4MB of SRAM; these independent SRAM components were an early hint that FPGAs would eventually need embedded RAM blocks. The article also provides foundational underpinnings for reconfigurable computers by laying out a model for understanding the computational capacity of these FPGA-based-machines as computational substrates, including how they will scale with Moore's Law.

Endorsement by: André DeHon, Professor, University of Pennsylvania

DOI: http://doi.org/10.1109/92.486081