Introduction to the Future of Reconfigurable Computing and Processor Architectures

Luigi Carro¹ and Stephan Wong²

¹ UFRGS, Brazil

² TU Delft, The Netherlands

As technology scales up, the design productivity slows down and new approaches must be sought after to reach the design productivity (or bridge the gap) that is nowadays common for general-purpose processor design. New tools, technologies, languages, operating systems, and design approaches are likely to be needed. Nowadays, extra transistors are utilized for acceleration purposes (next to prototyping circuits), like they have been used in the past for floating-point operation, and it is today for the case of special blocks embedded into computer architectures like the MMX, SSE, and GPU. In this special session, selected papers add to the discussion on how this symbiosis will evolve, showing how general-purpose or multicore computing can benefit from reconfiguration, how can one generalize current accelerators, and how this all will affect the way compilers produce code or deals with virtualization, and the integrating role of the OS. Moreover, one should look at how technology evolution will cope with current possible show stoppers, like the communication problem and the interface to the operating system. Five papers covering these different aspects are presented in this session. The first one discusses the role of the OS on managing reconfigurable resources. The next two papers discuss the role of reconfigurable devices being used as accelerators in new application domains, followed by an analysis of the role of reconfigurable computing in the current search for parallelism exploitation. Finally, the session ends with a survey on the use of reconfiguration in multithread architectures.