Liquid Computing

Stephan Wong

Abstract:

After the Dennard scaling ended (2006-2008), the industry has been searching for other ways to efficiently use the increasing amount of transistors provided by the semiconductor industry that is still following the Moore's Law. Without much thought, the industry decided to rely on multiplying many (simple) computing elements to boost performance in the hope that enough task-level parallelism can be found to drive and sustain the newly introduced resources. However, the industry quickly shifted from homogenous to heterogeneous computing elements to better adapt the computing to the targeted applications. However, this trend is not sustainable either, and reconfigurable computing has been touted as promising new direction to use these resources in a world of changing requirements and environments where computing is needed.

Change is the only constant. This is not only true in life (by Heraclitus), but especially true in current-day computing where the change is no longer required from one processor generation to the next, but continuously during run-time and lifetime of the processor. While in the past, progress regarding performance and energy efficiency were made between generations; we now need to build intelligence into our processors to adapt to a changing environment during operation - e.g., ranging from being safe at an airport to being constantly radiated in the air or even in space. Moreover, with the decrease of component reliability already in play, the adaptation is not only towards the "outside" world but also towards performing at best with the available (sometimes, even less reliable) resources at hand.

A new research direction, called Liquid Computing, was started at the Computer Engineering Laboratory at the Faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology, The Netherlands, to address all issues related to run-time adaptivity of computing systems. Our approach is to investigate possibilities for hardware components (processors, memories, and network-on-chips (NoCs)) to adapt themselves at run-time to the applications that need to be executed to provide adequate performance without much hardware resource waste and, in turn, reduce energy consumption. For example, we designed the Delft reconfigurable and parameterized VLIW processor called rho-VEX (rVEX in short). We implemented the processor on an FPGA as an open-source softcore and made it freely available. Using the rVEX, we intend to bridge the gap between general-purpose and application-specific processing through parametrization of many architectural and organizational features of the processor.

Bio:

Stephan Wong is currently working as an associate professor at the Computer Engineering Laboratory at the Delft University of Technology (TU Delft), The Netherlands. He has considerable experience in the design of embedded reconfigurable processors. His research interests include embedded systems, multimedia processing, complex instruction set architectures, reconfigurable and parallel processing, and distributed/grid processing. He has served in many conference, and workshop PCs as well as participated in the organization of many conferences and was editor of several proceedings and journals. He is a senior member of the IEEE. Visit http://www.ce.ewi.tudelft.nl/wong/ for a more up-to-date overview of all his activities.