

Carlos Tavares, Oskar Mencer and Wayne Luk - Imperial College London

Accelerating assembly-level programs

Various techniques have been developed for partitioning an application into two parts: software running on an embedded processor, and hardware in a co-processor. Co-processors have been shown to be effective in accelerating applications for embedded systems. The embedded processor in such systems are often based on Reduced Instruction Set architectures. Recently system-on-chips based on Complex Instruction Set architectures such as x86 have been proposed. This talk addresses the challenges in accelerating assembly-level programs with complex instructions, making use of both static and dynamic analysis. We introduce an infrastructure to support assembly-level program manipulation, and describe its facilities such as those for instruction-level instrumentation