

From Classical to Runtime Aware Architectures

Abstract: In the last years the traditional ways to keep the increase of hardware performance to the rate predicted by the Moore's Law vanished.

When uni-cores were the norm, hardware design was decoupled from the software stack thanks to a well-defined Instruction Set Architecture (ISA). This simple interface allowed developing applications without worrying too much about the underlying hardware, while computer architects proposed techniques to aggressively exploit Instruction-Level Parallelism (ILP) in superscalar processors. Current multi-cores are designed as simple symmetric multiprocessors on a chip. While these designs are able to compensate the clock frequency stagnation, they face multiple problems in terms of power consumption, programmability, resilience or memory. The solution is to give more responsibility to the runtime system and to let it tightly collaborate with the hardware. The runtime has to drive the design of future multi-cores architectures. In this talk, we will introduce an approach towards a Runtime-Aware Architecture (RAA), a massively parallel architecture designed from the runtime's perspective. RAA aims at supporting the parallel runtime system by enabling fine-grain tasking or managing hybrid memory hierarchies.