

Hardware/Software Co-Design for Efficient Microkernel Execution

Martin Decky, Ph.D., Huawei Technologies, German Research Center

While the performance overhead of IPC in microkernel multiserver operating systems is no longer considered a blocker for their practical deployment (thanks to many optimization ideas that have been proposed and implemented over the years), it is undeniable that the overhead still does exist and the more fine-grained the architecture of the operating system is (which is desirable from the reliability, dependability, safety and security point of view), the more severe performance penalties due to the IPC overhead it suffers. This talk discusses some specific hardware/software co-design ideas to improve the performance of microkernel multiserver operating systems.

One reason for the IPC overhead is the fact that current hardware and CPUs were never designed with microkernel multiserver operating systems in mind, but they were rather fitted for the traditional monolithic operating systems. This calls for an out-of-the-box thinking while designing instruction set architecture (ISA) extensions and other hardware features that would support (a) efficient communication between isolated virtual address spaces using synchronous and asynchronous IPC primitives, (b) efficient handling of hardware interrupts in user space, and (c) treating object references (e.g. capability references) as first-class entities on the hardware level. A good testbed for evaluating such approaches (with the potential to be eventually adopted as industry standard) is the still unspecified RV128 ISA (128-bit variant of RISC-V).
