A Glimpse on Future Server Architectures and Their Implications on System Software

Bernd Winkelsträter, Fujitsu Technology Solutions GmbH

Abstract

In the next five to ten years, commercial computer architectures will undergo fundamental changes. In CPUs, we will see lots of parallelism with many, highly power managed cores, and virtualization becoming the default execution mode in the enterprise. Non-x86 architectures will gain new traction due to the special needs of the hyper-scale datacenter market. Today's HPC coprocessors will become the main execution engine – the current approach using PCI-Express cards is no real fit for the Exaflop era. The memory paradigm will change dramatically - not only with exploding capacities and growing importance of ccNUMA optimization in software, but also resembling a new tier of persistency, completely blurring the lines between storage and main memory. Since current SMP coherency mechanisms are based on CPU caches and not the memory tier, fundamental software changes are needed to exploit the full potential of non-volatile main memory. Many kinds of in-memory applications will benefit from this paradigm shift, but it often means to rework their code. Driven by hyper-scale data centers and HPC, parallel optical interconnects will become affordable mainstream, enabling new levels of system modularity, configurability, and scalability. In all that, operating systems play a central role and need to adapt to these rapid and fundamental changes. Some of their algorithms like the I/O stack already feel the pain; they had to be optimized to handle current SSDs with 10,000s IOPs and latencies around 100 microseconds. But what if persistency latencies will go down further, below 10 or even 1 microsecond with SCM? What does it mean for the OS storage paradigm, if SCM is located in a fabric, accessed in new ways and shared by multiple hosts? Emerging new memories and chip packaging technologies will enable memory capacities approaching 100 Terabytes and 10s of Gigabytes of high speed on-package memory. In combination with future chip-level photonics integration, this will enable radical new computer architectures, where data and CPUs co-located in the same package, cutting latencies and power dramatically. However, all such technology advances are just one side of the story. The real driving factor in this technology race is "economy of scale": whatever is mass-produced, will become cheap and gain mainstream adoption. No matter how bright an invention might be, someone has to pay for its development to make a commercial product. As long as there is no market with a killer application and it is not backed by big industry players, it will stay in research or a business niche. Hence, without having economics in mind, it is therefore impossible to predict the future of server technology.