Contribution ID: **82** Type: **not specified**

Node-X: A networked architecture for energy efficient high performance computing and data acquisition

Introduction

Energy dissipation currently represents the major bottleneck towards exascale computing. High end instruction set processors including many/multi –core processors and graphics processing units (GPUs) adopt an instruction stream based/temporal model of computation which is energy inefficient. Field Programmable Gate Arrays (FPGAs) and Application Specific Integrated Circuits (ASICs) allow implementation of a data driven/spatial model of computation which is more energy efficient. Despite interconnect overheads FPGAs are about 4 times more energy efficient than state-of-the-art many/multi-core CPUs and GPUs.

Different options exist for interfacing accelerator components (such as FPGAs and ASICs) to host CPUs. These include tightly coupled approaches such as integration of CPU and accelerator in a single integrated circuit and CPU motherboards with FPGA sockets. Such approaches are cost inefficient and present serious scalability constraints as far as the number of accelerator components over host CPUs employed is concerned. Expansion buses such as PCIe bus form an attractive choice for loosely coupled accelerators due to their high throughput capacity. However such approaches are expensive and require significant code development efforts both on the accelerator (FPGA) and host CPU side while scalability remains an issue (the number of PCIe slots in a CPU motherboard is limited).

Node-X computing architecture aims at overcoming limitations of state of the art computing architectures in terms of energy efficiency, scalability, cost and programmability.

The Node-X architecture

Node-x computing architecture deploys FPGA/ASIC-based, stand-alone acceleration boards and uses conventional/high-speed Ethernet to interconnect them. Limiting per-node processing elements (FPGAs or ASICs) to one, 'Node-X nodes'build highly granular, Ethernet based, computing fabrics.

Early work on analyzing data-flow patterns in modern data-centers showed that, in certain (common) processing scenarios, moving data through Ethernet (between processing nodes) is not introducing energy consumption or latency penalties given the distributed nature of client-server/API based programming paradigms becoming dominant. Additionally, 4x10G Ethernet interconnected nodes were measured to process and round-trip data in more predictable and efficient ways (compared to locally processed data), without evidently stressing per-unit (node) Bill of Materials (BoM), given the popularity and availability of mature 10G Ethernet silicon.

Node-X nodes are architectured to operate in Single Input Single Output (SISO) mode, foreseeing two possible hardware implementations:

i. Processing nodes: Collect input data from Ethernet –process using FPGAs or ASICs –and release output again on Ethernet and

ii. Data provision nodes: Source/Sink data between Ethernet and Memory or I/O

Abstracting data provisioning from processing nodes allows data residing in volatile/non-volatile memories to be handled in equal footings with data coming from high-speed I/O (ADC/DAC), making node-X an ideal candidate for DAQ applications while alleviating programming from handling data concurrency challenges.

Furthermore, by having node-X nodes built around a single FPGA or ASIC (per node), further gains in lifting programming complexities and in increasing performance predictability are observed, together with critical benefits in assisting fault tolerance scenarios to be deployed, enhancing hardware level granularity, as well as sizing-down per-node manufacturing costs.

Signal processing, data acquisition

System integration and engineering

Computing

Software and imaging

Primary authors: Mr ANGELAKOS, Evangelos (Nanotronix Inc - University of Peloponnese); Mr POULIS, Spiros (University of Peloponnese); Dr DIMITROULAKOS, Grigoris (University of Peloponnese); Dr MASSELOS, Konstantinos (University of Peloponnese)

Presenter: Mr ANGELAKOS, Evangelos (Nanotronix Inc - University of Peloponnese)

Track Classification: Front/Back-End electronics