

3rd ATTRACT TWD Symposium in Detection and Imaging

Tripolis, 31 May-1 June 2017

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

Evangelos Angelakos, Spiros Poulis, Grigoris Dimitroulakos, Konstantinos Masselos

Nanotronix Inc, USA

Computer Systems Design Group, University of Peloponnese

Reminder

Towards Exascale Computing:
The Energy Bottleneck

European Technology Platform for High Performance Computing target: 250 Pflops using up to 15 MW of power in 2019

A published goal is Exascale computing using 20 MW of power existing circuits consume an order of magnitude too much power to meet this goal

The overall power consumption of just one "Exascale" supercomputer in 2018 (if possible), will be in the order of 10 GW

Trends

Why are FPGAs successful as accelerators?

Compared to software FPGAs provide power/energy results which are about up to 4 orders of magnitude better

The computational model is more efficient since it is data stream-based and not instruction-stream-based

Key investments, suggesting FPGA use in next-gen computing


- \$16.7bn Intel's purchase of Altera (June 2015)

- Microsoft's Catapult, accelerating Bing searches with FPGAs

- Amazon Plugs Xilinx FPGA into its Cloud (2 Dec. 2016)


Accelerator Interfacing Options

Closely coupled
to CPU
(INTEL-Altera)



- Focused on CPU data crunching (narrow application field)
- For XEONs CPUs in 2017 (targeting high-end market)
- 1-2 CPUs / system (?) (not scalable)

Expansion buses
(as a PCIe card)



- An expansion card in the mobile, tablet, USB-C era ?!
- Expensive
- Driving software required
- 3-4 cards per system (?) (not scalable)

Network approach




- Crunching CPU & Sensor data (wide application field)
- Lower cost
- Distributed/sharable resource
- Ultra-scalable

Trends & Wishes

Network based High Performance
Computing Architecture

Node-X board



FPGA/ASIC-based stand-alone acceleration board

Conventional/high-speed Ethernet to interconnect

Limiting per-node processing elements (FPGAs or ASICs) to one, allows building highly granular, Ethernet based, computing fabrics

Node-X Computing Architecture

Node-X boards are designed to operate in Single Input Single Output (SISO) mode

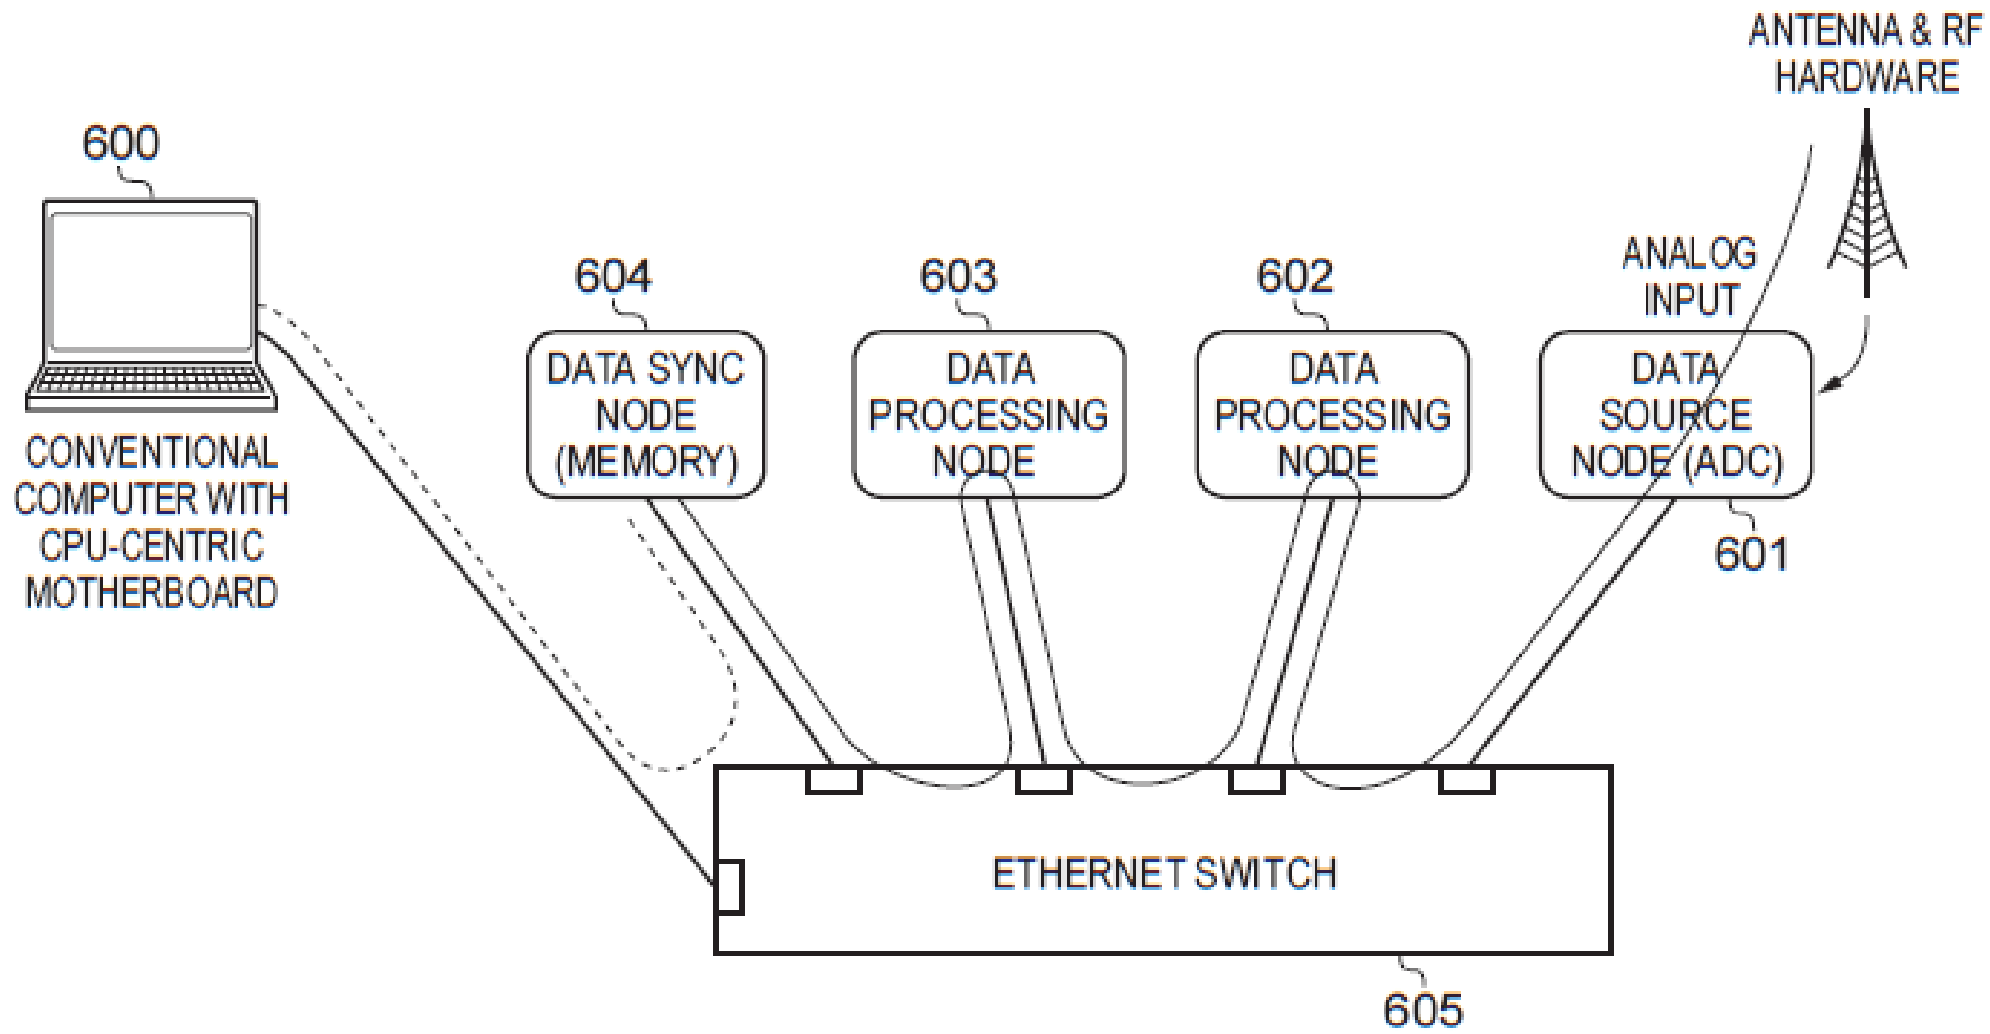
Two possible hardware implementations:

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

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

Dreams

Node-X Computing



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

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

Thank you for your interest