ExaNeSt

European Exascale System Interconnect and Storage

FORTH, ARM,, EnginSoft, eXact Labs, Fraunhofer, Iceotope, INFN, INAF, MonetDB Solutions, University of Mnchester, University of Valencia, Virtual Open Systems

Objectives - Approach

Objectives

Interconnects

- Low-latency, unified compute & storage traffic

Storage

- Fast, distributed, in-node non-volatile memory

Applications

- Real applications: scientific computing data analytics

System Packaging Technology

- Compact, fully-immersed liquid cooling technology

Co-Design Approach

Prototype to be built:

- 1000+ ARMv8 cores
- 100+ M.2 SSD's
- 4TB+ of DDR
- UNIMEM Address Space
- Shared I/O

In collaboration with ExaNoDe & ECOSCALE: FPGA Accelorators.

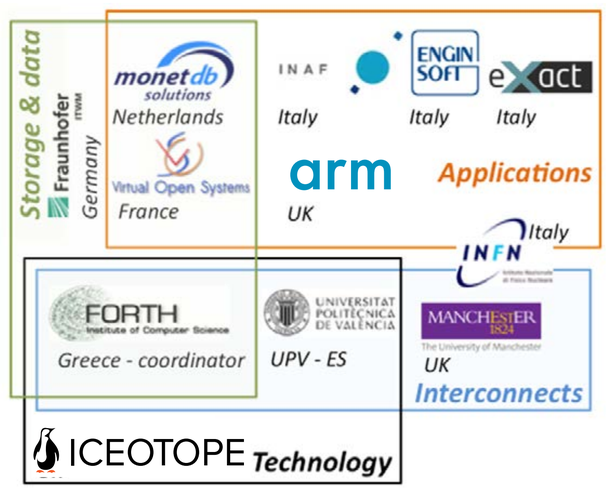

ExaNeST Rack-Scale Prototype

Dense Packaging 1: Quad-FPGA Daugther Board (QFDB)

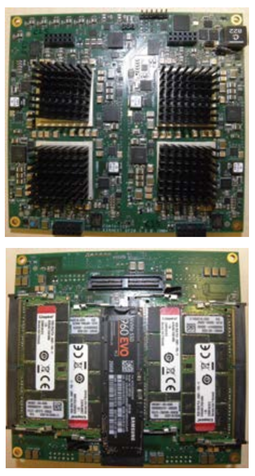

- 4x Zynq FPGA = 16x ARM 64-bit A53 proc. cores + 10k DSP slices + 2.4 million logic elements
- 64 GBy DRAM 2133MHz, ECC
- 250 GBy SSD
- 10 off-board links x 10 Gbps
- 120x130 mm2 board, 12 PCB layers, 1700 components, 46 power supplies, 16 power sensors
- 24 high-speed serial x16 Gbps + 144 LVDS-pair on-board links

Dense Packaging 2: Liquid Cooling

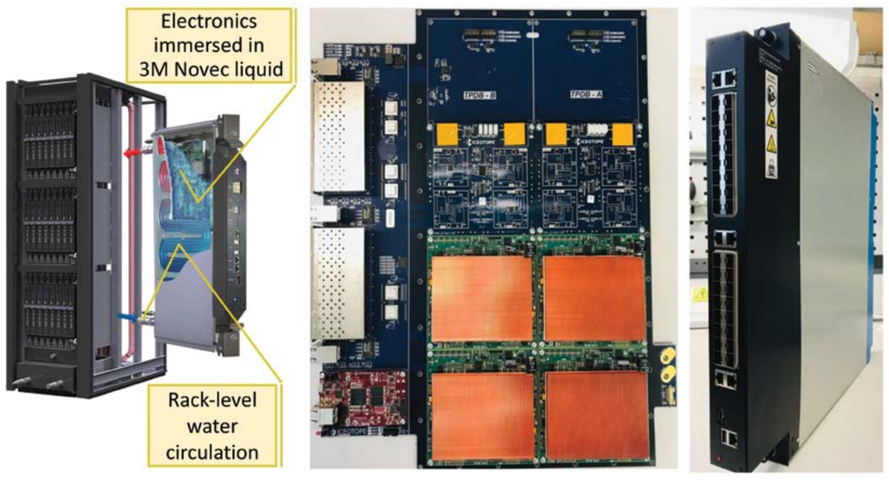

- Currently: vertical blades, fully immersed
- Next Generation: horizontal, sprinkled

The HPC Testbed

- Currently: 6 Blades = 24 QFDBs = 96 FPGAs = 384 cores (64-bit A53) + 1.5 TBy DRAM + 6 TBy SSD
- Runs full systems software stack & HPC jobs mng'mnt
- Runs full, real Applications
- Soon: twice the size

Interconnect

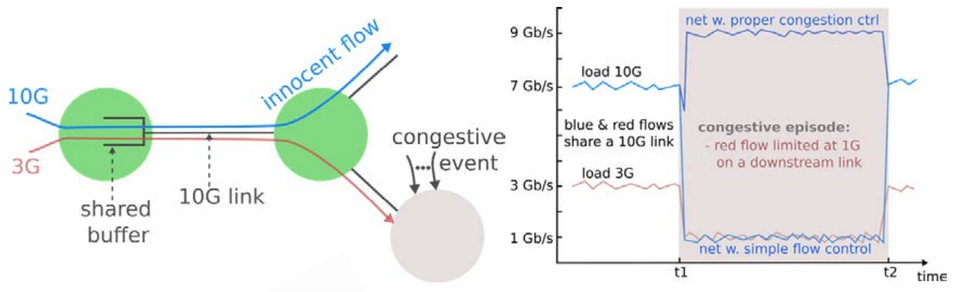

Congestion Management

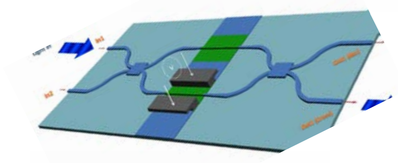

Silicon Photonics

- Multi-tiered Topologies
- Geographical Routing
- Low Latency Communication (APEnet)

Storage

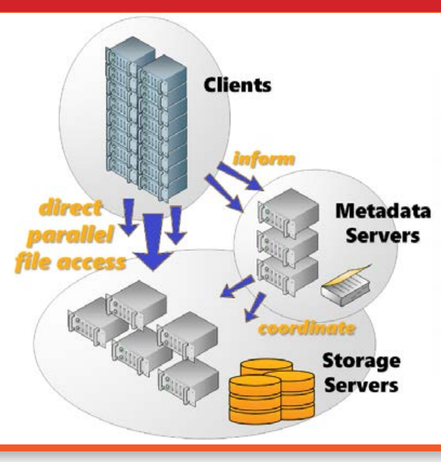

- UNIMEM - PGAS
- BeeGFS Filesystem
- HPC Virtualization
- Profiling Tools
- Checkpointing
- Status Monitoring
- FPGA Acceleration

Applications

System

6x-10x
LESS ENERGY
on HPCG and HPL

Number of FPGA's = Number of Intel Sockets, 4 cores each	HPCG (High Performance Conjugate Gradient		HPL (High Performance Linpack)	
	ExaNeSt [kJ]	Intel Cluster [kJ]	ExaNeSt [kJ]	Intel Cluster [kJ]
4	46	449	44	456
8	83	686	83	713
16	219	1264	168	1852
32	440	2864	300	2617

Gadget

6x
LESS ENERGY
to solution

ExaHiNBody

1.3x
LESS ENERGY
to solution

Oil Reservoir Simulation (Rachford-Rice equation)

8.5 GFLOPS/Watt, 200 GFLOPS (300 MHz)

SGEMM (single precision FP matrix multiply)

17 GFLOPS/Watt, 1100 GFLOPS (at 300 MHz, 82% DSP utilization)

Smart City (real-time video proc., Lucas-Kanade alg.)

One FPGA: 36 ms/frame, 8 Watt Vs NVidia GTX 960 (16 SM): 43 ms/ frame, 75 Watt

Space-CNN (Convolutional Neural Networkweight compression for space data classification)

QFDB: 265 GFLOPS, (at 250 MHz) Vs NVidia Quadro K2200: 123 GFLOPS

THE FRAMEWORK PROGRAMME FOR RESEARCH AND INNOVATION
HORIZON 2020

ExaNeSt (H2020-ICT-671553), Horizon 2020 Program

www.exanest.eu
twitter: @exanest_h2020

Contact: Prof. Manolis Katevenis (email: kateveni@ics.forth.gr)
Institute of Computer Science, FORTH, Greece