MYSTIC: Programmable Systems Research Testbed to Explore a Stack-WIde Adaptive System fabriC

Alexandru Iulian Orhean[1] **Alexander Ballmer**[1] **Travis Koehring**[1] **Kyle Hale**[1] **Xian-He Sun**[1] **Ophir Trigalo**[2] **Nikos Hardavellas**[3] **Sanjiv Kapoor**[1] **Ioan Raicu**[1]

[1]Computer Science Department, Illinois Institute of Technology [2]Office of Technology Services, Illinois Institute of Technology [3]Computer Science Department, Northwestern University

Abstract

The pervasiveness and the continuous advancement of computer systems have determined a sudden upsurge of digital data. The Big Data phenomena has resulted in a widespread adoption of parallel and distributed systems for efficiently storing and accessing data. Likewise computing applications are changing from compute-centric to data-centric, where communication of data comprises a large part of the program and prominently dictates performance. Static configurations of network, memory, and storage under the form of computer clusters show signs of fatigue in the face of increasing scale of modern big data applications. This NSF-funded infrastructure project proposes to build a dynamically configurable cluster called MYSTIC to study system re-configurability across the entire computing stack, from the processor to memory, storage, and the network. It will allow low-level experimentation and reconfiguration at the level of networks-on-chip (NoC), universal memory, and the network interconnect with multidimensional network topologies. Dynamically reconfiguring interconnects, memory and storage will speed-up applications running on a heterogeneous computing environment. The goal of this research is to remove network, memory, and storage performance bottlenecks by dynamically reconfiguring them to the needs of an application. In this testbed the organization of memory, storage, and their connectivity will change dynamically to adapt to the workload to improve overall performance of the executing programs. An overview of both the hardware and software stack will be presented, as well as how the community can get access to this exciting new research testbed.

Hardware Specifications

Located in the Ocient data center in the Illinois Institute of Technology campus.

- Racks: 3
- Nodes: 54
- General purpose cores: 1396
- Accelerators: 40
- Peek power consumption: 60kW
- Memory: 4.3TB DDR4
- SAS HDD storage: 388TB
- SAS/NVMe SSD storage: 120TB
- External network: 10 GbE
- Internal network: 100 GbE Fat-Tree

Computers and Accelerators

Nodes	CPU Model	Sockets	Compute Power	RAM	Disk
1	Intel SP 8160	8	192c (384HT) @ 2.1 GHz	768 GB	22.1 TB
1	Intel SP 8160	2	48c (96HT) @ 2.1 GHz	394 GB	8 TB
1	Intel SP 6130	2	32c (64HT) @ 2.1 GHz	64 GB	9.4 TB
16	Intel SP 4108	2	16c (32HT) @ 1.8 GHz	64 GB	1.4 TB
8	Intel SP 4110	1	8c (16HT) @ 2.1 GHz	32 GB	1.5 TB
1	Intel BW 2620 v4	2	16c (32HT) @ 2.1 GHz	64 GB	2 TB
2	Intel HW 2620 v3	2	12c (24HT) @ 2.4 GHz	32 GB	2 TB
4	Intel Xeon Phi 7210	1	64c (256HT) @ 1.5 GHz	64 GB	0.4 TB
1	AMD Naples 7501	2	64c (128HT) @ 2 GHz	128 GB	2.5 TB
1	ThunderX 88XX	2	96c (96HT) @ 2 GHz	64 GB	0.4 TB
1	ThunderX2 99XX	2	56c (224HT) @ 2 GHz	128 GB	0.4 TB
1	POWER9	2	40c (160HT) @ 3.8 GHz	128 GB	1.9 TB
4	Intel SP 3106	2	16c (16HT) @ 1.7 GHz	48 GB	4.4 TB
4	Intel SP 4110	2	8c (16HT) @ 2.1 GHz	64 GB	97 TB
8	BlueField SoC	1	16c (16HT) @ 2.5 GHz	64 GB	2 TB

Model	Cards
Nvidia Tesla V100 (32GB) - 5120c @ 1.245GHz	8
Nvidia Tesla V100 (16GB) - 5120c @ 1.245GHz	8
Intel 510T Altera Arria 10 FPGA (32GB)	8
Nvidia Tesla K40 (12GB) - 2880c @ 0.745GHz	8
Xeon Phi 3120A (8GB) - 57c @ 1.1GHz	8

Software Stack

Overcloud: OpenStack (PerconaDB, Keystone, Horizon, Cinder, Glance, Nova, Neutron, RabbitMQ, Swift, Ironic, Blazar, LXD)

Undercloud: MaaS, Juju, LDAP, Ceph Cluster, User Login, Monitoring Service (Grafana + InfluxDB)

Protocloud: CentOS 7.6, Gateway, Firewall, IP Failover, DHCP, DNS, HW/SW RAID, KVM Hypervisor, SaltStack

Full Stack Re-Configurability

Allows low-level experimentation and reconfiguration at multiple levels:

- Processor: network-on-chip (NoC), hybrid architectures
- Memory: Deep memory hierarchy
- Storage: NVMe, NVDIMM
- Network: software defined networking, programmable NICs

Leadership

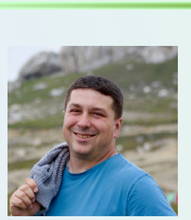
Ioan Raicu

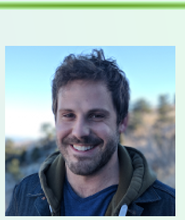
Kyle Hale


Xian-He Sun


Ophir Trigalo


Nikos Hardavellas


Sanjiv Kapoor


Alexandru Iulian Orhean


Alexander Ballmer


Travis Koehring

Acknowledgments

This work is supported by the NSF CNS-1730689 CRI award (Mystic Testbed) and the generous support from NVIDIA. This infrastructure is physically located in the Ocient Computational Center at Illinois Institute of Technology.