Mystic

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

Ioan Raicu
Illinois Institute of Technology
Argonne National Laboratory

Trends in HPDC Workshop 2019
March 14th, 2019

WHO AM I?

- **History**
  - 1997-2002: BS/MS in CS at **Wayne State University**; MS thesis in IPv6 Network Protocols under Sherali Zeadally
  - 2003-2009: PhD in CS at **University of Chicago** in Many-Task Computing under Ian Foster
  - 2009-2010: Postdoc at **Northwestern Univ.** with Alok Choudhary
  - 2016-2017: Sabbatical at **Northwestern Univ.** with Peter Dinda
- **Current Affiliations**
  - Associate Professor in CS at **Illinois Institute of Technology**
    - Director of Data-Intensive Distributed Systems Laboratory
  - Guest Research Faculty in MCS at **Argonne National Laboratory**
  - Advisory Board Member at **Ocient LLC & FusionBlock**

OCIENT™

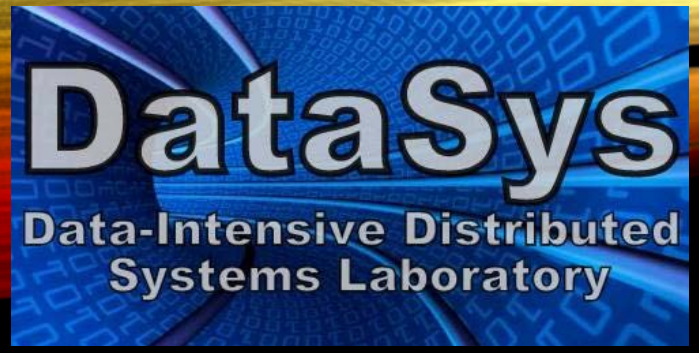

DATASYS:
DATA-INTENSIVE DISTRIBUTED SYSTEMS LABORATORY

Research Focus

Emphasize designing, implementing, and evaluating systems, protocols, and middleware with the goal of supporting data-intensive applications on extreme scale distributed systems, from many-core systems, clusters, grids, clouds, and supercomputers.

WHAT IS MYSTIC?

- $1M NSF-funded institution infrastructure grant at Illinois Institute of Technology
- A testbed for experimenting with system re-configurability across the entire computing stack
- Allow 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, programable NICs
- More information
  - http://mystic.cs.iit.edu
  - https://www.nsf.gov/awardsearch/showAward?AWD_ID=1730689&HistoricalAwards=false

WHAT RESEARCH DOES MYSTIC ENABLE?

- Nautilus Aerokernel Light-weight Operating System
- Xtask: eXTreme fine-grAined concurrent taSK invocation runtime
- Universal Memory through byte-addressable non-volatile memory
- XSearch: Distributed Indexing and Search in Large-Scale Storage Systems
- OS abstractions on programmable NoCs to adapt to application workloads
- Multipath routing protocols in multi-dimensional Torus networks
- Integrated data access system to support PFS/DFS

MYSTIC LEADERSHIP

- Ioan Raicu
  - Principle Investigator: explore practical aspects to realizing universal memory with byte-addressable non-volatile memory where applications compute directly on persistent memory
- Kyle Hale
  - Co-Principle Investigator: build new OS abstractions on top of programmable NoCs to adapt the on-chip network to applications' communication topologies and to enforce performance isolation and QoS between specialized OSes that space-share the chip
- Xian-He Sun
  - Co-Principle Investigator: develop an integrated data access system, Dynamic PortHadoop to support both parallel file systems for data-coherent and MapReduce/Spark systems

MYSTIC COLLABORATORS & ADMINS

- Ophir Trigalo
  - Collaborator: provide resources and expertise needed to bring the MYSTIC testbed online, and to maintain it over the course of the testbed lifetime
- Nikos Hardavellas
  - Collaborator: provide expertise in specialized operating systems (OS) and field programable gate arrays (FPGA)
- Sanjiv Kapoor
  - Collaborator: provide expertise in improving network performance through new dynamic routing algorithms, leveraging rich multi-path topologies of multi-dimensional networks
- Alexandru Iulian Orhean
  - Lead System Administrator: provides day-to-day operational support, maintenance, and technical support
- Alexander Ballmer
  - System Administrator: provides day-to-day operational support, maintenance, and technical support





MYSTIC HARDWARE OVERVIEW

- Physically at IIT in Chicago
- 55-nodes
- 3 racks
- 45kW of power
- OpenStack bare-metal provisioning

MYSTIC HARDWARE OVERVIEW

- Computing capabilities (240TF/s)
  - 1396 general purpose cores (Intel Xeon SP/Phi x86, AMD Epyc x86, Cavium ARM, IBM Power)
  - 100K+ accelerator cores (Intel Arria FPGA, Intel Xeon Phi, NVIDIA Kepler/Volta)
- Memory
  - 4.7TB DDR4
- Storage
  - 375TB HDD SAS
  - 90TB SSD SAS
  - 40TB SSD NVMe (Intel Optane 900P, Samsung EVO 970)
- Network
  - External: 10Gb/s
  - Internal: Mellanox 100GbE Fat-tree

SPECIAL HARDWARE

SPECIAL HARDWARE

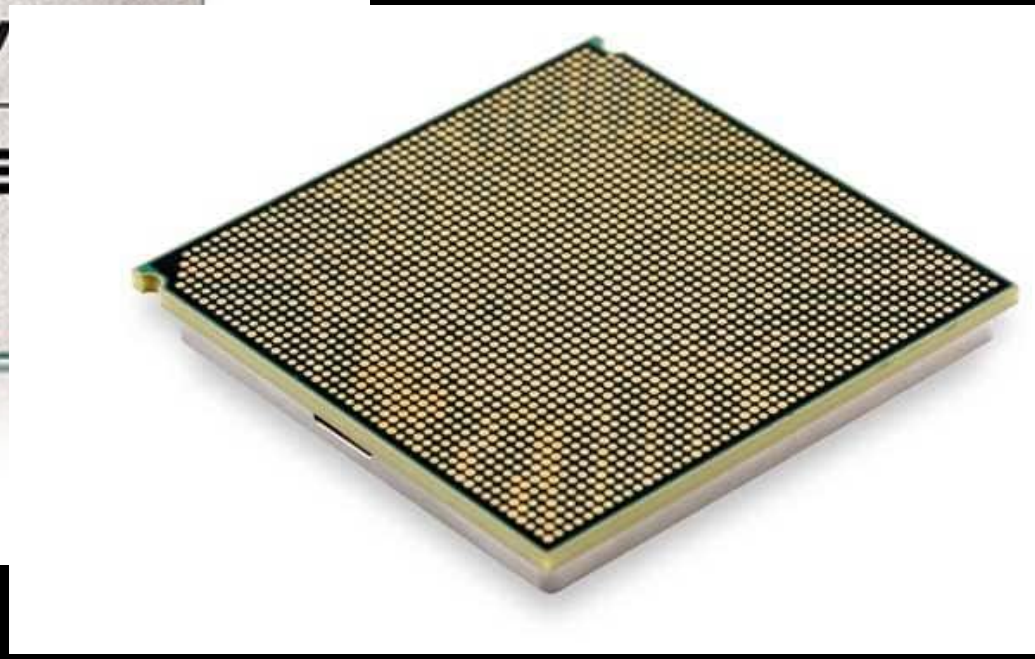

SPECIAL HARDWARE

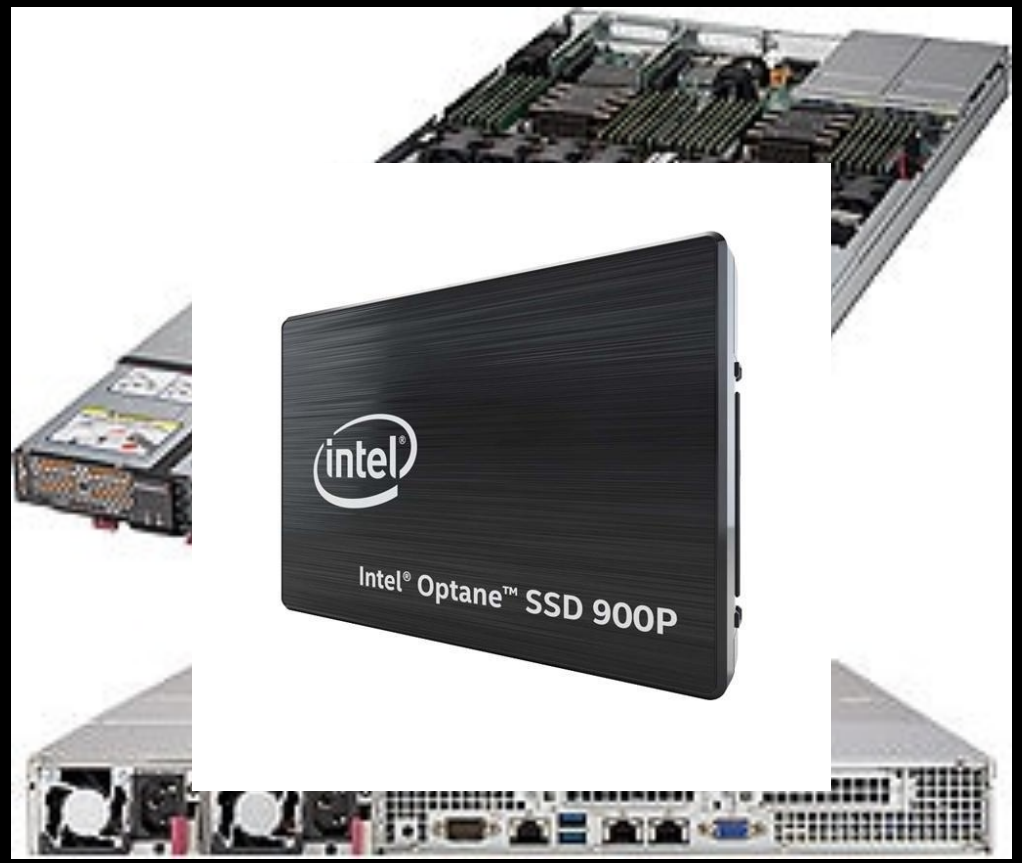

SPECIAL HARDWARE

NETWORK

PLANS FOR FUTURE HARDWARE

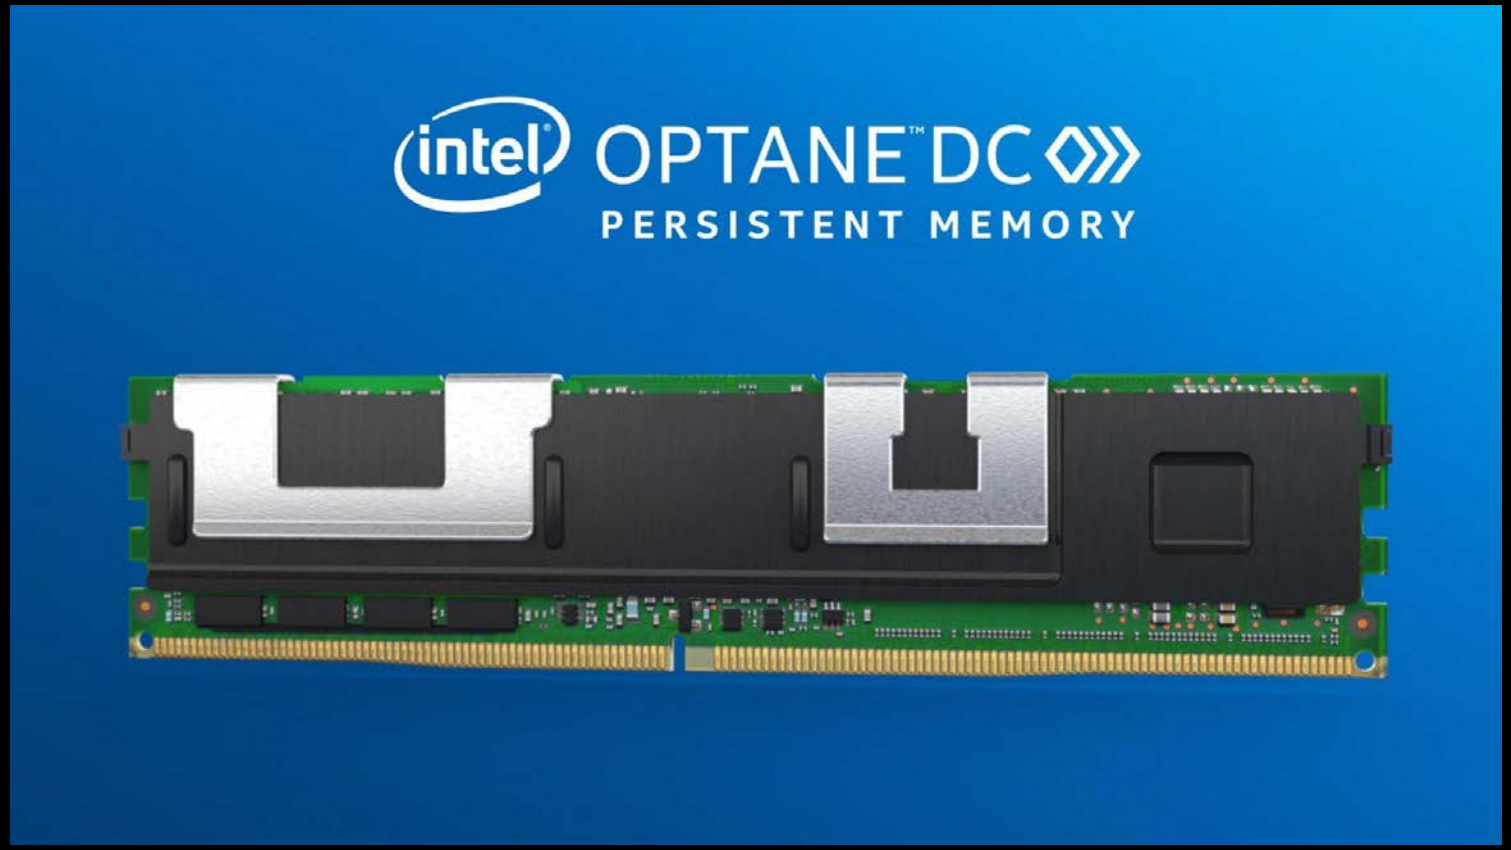

MYSTIC SUMMARY

- All major architectures
  - x86, ARM, Power9, GPUs, FPGAs
- 1-socket to 8-socket systems
- 8-cores to 192-cores shared memory systems
- Storage nodes with 64-lanes of PCIE bandwidth across 32 NVMe drives

INTERESTED IN GETTING ACCESS TO MYSTIC?

- We encourage the community to use Mystic
- Planning to be online in May 2019
  - More information will be announced at: http://mystic.cs.iit.edu
- Planning on running a workshop at IIT in June 2019
- Write me at iraicu@cs.iit.edu if you are interested

RELATED ACTIVITIES

- **NSF REU Site**
  - **BigDataX: From theory to practice in Big Data computing at eXtreme scales**
- **Competitions**
  - **Student Cluster Competition at IEEE/ACM Supercomputing**

BIGDATAX REU SUMMER PROGRAM

STUDENT CLUSTER COMPETITION AT IEEE/ACM SUPERCOMPUTING/SC

FUNDING/COLLABORATIONS ACKNOWLEDGEMENTS

- National Science Foundation, CISE Research Infrastructure (CRI)

- NVIDIA

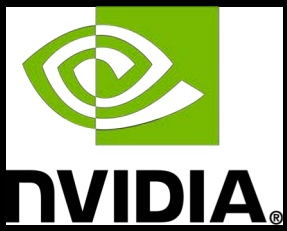

- Chameleon

QUESTIONS

- Contact:
  - iraicu@cs.iit.edu
- More information:
  - http://www.cs.iit.edu/~iraicu/
  - http://datasys.cs.iit.edu
  - http://mystic.cs.iit.edu