DARWIN - A Resource for Computational and Data-intensive Research at the University of Delaware and in the Delaware Region

Rudolf Eigenmann, Benjamin E. Bagozzi, Arthi Jayaraman, William Totten, and Cathy H. Wu

2021

University of Delaware

Overview

Cyberinfrastructure (CI) drives progress in science, engineering, and business in a major way, as stated prominently in the 2015 presidential National Strategic Computing Initiative [1]. CI - consisting of computer, data, networking and related resources - enables new ideas and products to be explored “virtually,” without building expensive prototypes and physical experiments. CI also powers almost all business processes today. The DARWIN project, illustrated in Figure 1, implements a major computational and data resource at the University of Delaware (UD), enabling and accelerating progress in all sciences and addressing grand challenges facing our society. The resource is sponsored by the National Science Foundation (NSF) through its Major Research Instrumentation (MRI) program (award # OAC-1919839 [2]) with substantial cost sharing by UD. Through partnerships with regional universities, colleges, health institutions and the private sector, the new resource also boosts research, development, and education in the greater Delaware region.

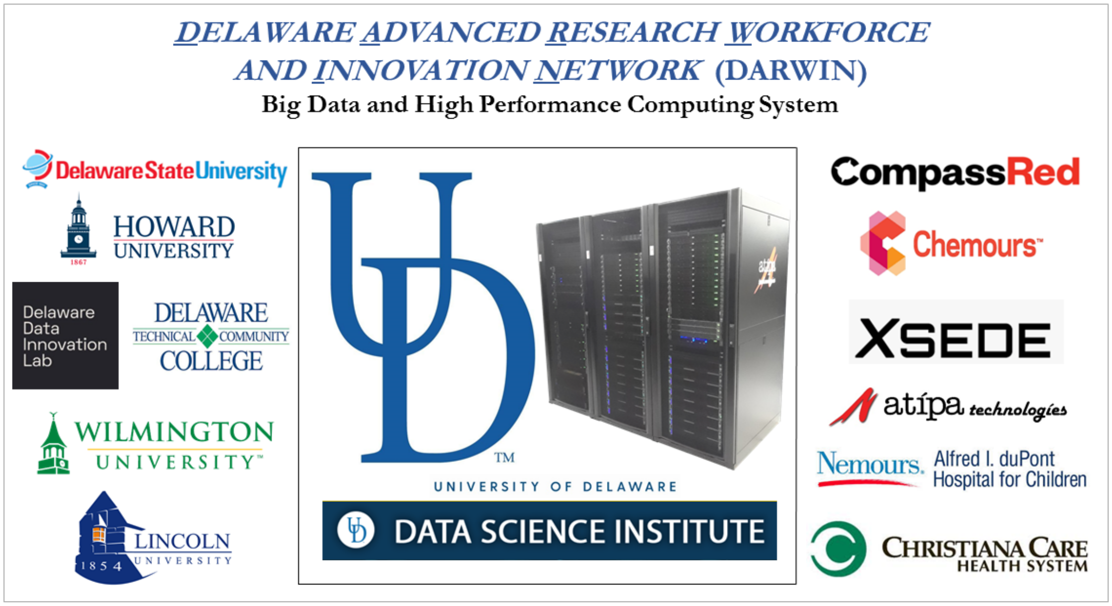

Figure 1: The DARWIN Project

The DARWIN "instrument" is a compute and storage cluster, including diverse compute nodes and storage capacity, supporting the breadth of research of its user community. It serves as a critical and transformative upgrade to Delaware's cyberinfrastructure, enabling research and educational activities for a large number of faculty across all UD colleges as well as for users from academic and industrial partners within the broader Delaware region. The instrumentation is designed to enable research broadly across disciplines with diverse software and hardware needs, such as problems that scale to large numbers of processors and data sets, involve large data transfers, use advanced graphics accelerators, and require new operating modes. It also serves to train students and researchers on computational and data-intensive methods and to enhance these skills in the greater Delaware region. DARWIN, which stands for Delaware Advanced Research Workforce and Innovation Network, will complement ongoing UD initiatives focused on improving and enhancing networking, storage, and compute infrastructure. This addition is timely, providing synergy with the new UD Data Science Institute and the faculty network in high-performance computing, led by the project investigators. Additionally, it serves as a focal point to bring together researchers across the Delaware region.

Participants and Projects

The DARWIN project is a partnership of researchers at UD, UD Information Technologies, and participants in the greater Delaware region.

Research Participants and Thrusts

DARWIN was motivated by the needs of a large community of researchers at the University of Delaware (UD) and in the Delaware region. The project is led by the authors of this paper. While the resource is available to all research groups at UD and DARWIN partner institutions, Table 1 lists key participants and contributors to the NSF proposal that funded DARWIN. The indicated lead investigators are faculty members at UD, unless noted. Column "Enabling Features" points to those machine characteristics that are of particular relevance to the research projects. Section "DARWIN Features and Architecture" describes these characteristics.

Table 1: Initial DARWIN Participants and Projects

Project and Lead Investigator	Discipline	Enabling Features
Improving Government Accountability & Responsiveness with Big Data on Citizen-Government Interactions (Prof. Benjamin Bagozzi)	Political Sciences	1,2,5
Advanced Image Analysis and Processing (Profs. Dawn Elliott, Keith Schneider, Gonzalo Arce, and Jeff Caplan; Prof. Hacene Boukari, DSU)	Image Analysis	1,2
Scalable Platforms for Citizen Science Linking the Cloud to the Crowd (Prof. Austin Brockmeier)	Citizen Science	1,2,4
Enhancing Wesley's Educational Mission in Data Science: (Prof. Malcolm D'Souza, Wesley College	Education	9
Accelerating Simulation of Flow of Neutrons in a Nuclear Reactor (Prof. Sunita Chandrasekara)	Physics / Computer Science	1,3
Mobility-Aware Edge Computing: Prof. Lena Mashayekhy	Computer Science	6
Center for Environmental Monitoring and Analysis – CEMA (Prof. Dan Leathers	Environmental Science	1,8
Water in the Changing Coastal Environment of Delaware – WiCCED (Prof. Kent Messer)	Environmental Science	1,2
Simulations of Complex Fluid Flows in Blood Streams (Prof. Antony Beris)	Material Science	1,3
Topological Analysis of Longitudinal Human Neuroimaging (Prof. Chad Giusti)	Neuroscience	1,2
Clustering of Universal Protein Sequences (Profs. Chuming Chen and Hongzhan Huang)	Bioinformatics	1,2
Leveraging Urban Observations for the Multi-Disciplinary Science of Cities (Prof. Gregory Dobler)	Science of Cities	2,5,6
Finite element discretization and preconditioning of mixed variational formulations (Prof. Bacuta)	Mathematics	1,9
Subsurface modeling of fluids in geological media (Prof. Wallace)	Earth Sci.	1,8
Multi-scale coarse-grained modeling of protein-protein and protein-solute interactions (Prof. Roberts)	Biomolecular Engineering	1,3,8,9
Precision agriculture, Carbon understanding in soil, carbon dynamics (Prof. Vargas)	Plant Sciences	4,5,9
AI for brain analysis to understand food decision making. (Prof. Beheshti)	Health Sci.	1,3,4,8
Macrosystems ecology at a continental scale. AI for Biological information from weather radar (Prof. Buler)	Ecology	2,4,5,8
Exploring individual differences in brain networks supporting person perception and social cognition (Prof. Cloutier)	Cognitive Sci.	1,4,8

Machine Learning, Classification, Information Retrieval (Prof. H. Fang)	Engineering	1,4,7
Large-scale simulations of 3D nonlinear ocean waves (Prof. Guyenne)	Oceanography	1
Reveal turbulent coherent structures in coastal zones (Prof. Hsu)	Civil Enr.	1,3,8
Low-cost, mobile brain imaging in language and literacy (Prof. Kaja)	Linguistics	1,4,7
Modeling diffusion in geological materials: Crystal clocks as a window into Earth's past (Prof. Warren, Prof Lynn)	Earth Sciences	1,7,8
Material properties from Schroedinger equations (Prof. Szalewicz)	Physics	1,2,8
Computational Intergroup Decision-making – computational modeling of social decision-making (Prof. Kubota)	Social Sciences	1,4,7
Thermo-hydra-mechanical material properties (Prof. Kaliakin)	Civil Engr.	1,4
New Data-Characteristic Driven Parallel Signal Processing Paradigm (Prof. Li)	Engineering	1,3,8
Coastal-glacier interactions (Prof. Moffat)	Oceanography	1,7
Understanding the Changing Environment in South Asian Forests (Prof. Mondal)	Earth Science	1,4,7,8
First-principles quantum modeling of nanoelectronic, optoelectronic and nanospintronic devices (Prof. Nikolic)	Physics	1,2,8
Life-cycle of infectious diseases (Prof. Perilla)	Biochemistry	1,7,8
Online social network analysis (Prof. X. Fang)	Social Science	1,2,3
Genetic correlates of visual and auditory statistical learning in ASD, Brain bases of typical and atypical language development (Prof Zhenghan)	Linguistics	1,4,7,8
High-dimensional multi-armed bandit approach with variable selection and optimality guarantees (Prof. Qian)	Economics	1,4
Development of a relativistic atomic code for accurate treatment of complex correlations (Prof. Safranova)	Physics	1,2,8
Parallel computational tools for simulating waves on solids (Prof. Sayas)	Mathematics	1
Large-scale kinetic plasma simulation (Prof. Shay)	Physics	1,7,8
Anxiety in Transition: The Role of Myelination & Testosterone on Pathological Anxiety (Prof. Spielberg)	Health Sci.	1,3,4
Estimating Structural Life-Cycle Models of Consumption, Saving, and Borrowing (Prof. Tobacman)	Economics	1,2,4
Applying Kinetic Gas theory to natural sea spray problems (Prof. Veron)	Geography	1,8
Data science-enabled artificial intelligence for chemical sciences (Prof. Vlachos)	Chemical Engineering	1,7
Distributed Learning: Theory and Algorithms (Prof. Wu)	Data Science	1,8
Atomistic origin of osteoporotic fracture in bone using large scale high performance calculations (Prof. Hossain)	Biomechanical Engineering	1,7,8
Acoustic wave propagation and scattering in oceanography (Prof Badiey)	Oceanography	1,2,4,5,7
Multi-scale modeling of synthetic and biologically relevant polymer systems (Prof. Jayaraman)	Chemical Engineering	1,3,7,8

Operation through UD IT-Research Cyberinfrastructure

Data Center Facility: The DARWIN compute system is located in the UD's main campus data center. This facility services the bulk of UD's processing needs and is home to both institutional (core network, ERP, IaaS, etc.) and educational/research (HPC, Data Storage Services, etc.)

computing services for the university community. The data center is staffed 24/7/365 and secured with multi-factor entry panels. Power is conditioned with UPS and an on-site generator powers the datacenter in the event of a disruption; UPS and generator power are tested monthly.

This facility is one of UD's Internet end-points and is part of a SONET ring shared between UD and the State of Delaware. With the aid of NSF award 1827138: CC* Networking Infrastructure: CyberInfrastructure Technology Advancement for Delaware (CITADel) - 100 Gb/s Connection Upgrade to Internet2, UD's ScienceDMZ is now connected to Internet2 through NYSERNet at 100Gb/s.

Staffing for Operations and Maintenance: The DARWIN machine is operated by the UD IT-Research Cyberinfrastructure team, with an additional staff member added to the team to support DARWIN. The team's experience includes managing three generations of university-wide HPC clusters (following the Purdue Community Cluster model), three departmental HPC clusters, and a number of bespoke research computing systems.

External Partners

DARWIN's external partners are listed in Figure 1. Through these external partnerships, the acquired cyberinfrastructure considerably enhances the collaborative usage of computational and data-intensive technologies within the state of Delaware and beyond. External partner activities both expand and diversify the involved knowledge base across the Mid-Atlantic region, with Delaware serving as the fulcrum for this transformation. Initially 20%, allocated through XSEDE [3], of the resources are allocated to external partnerships. Most external partners are based in the state of Delaware, and include longstanding UD relationships as well as several expanding opportunities of new multidisciplinary collaboration. Together these external partnerships underscore DARWIN's role as a critical resource for all of Delaware, while also ensuring that the system directly contributes to the 'public good' in the state of Delaware and nationally across areas as varied as the environment, public health, education, diversity, and the Delaware region's workforce skilled in computational and data science. External partners interact with each other, and with collaborators at the University of Delaware, not only through the usage of the DARWIN resource itself but also through annual symposia and training opportunities.

DARWIN Features and Architecture

DARWIN was conceived to offer critical new features to enable novel and emerging computational and data-intensive science. These features, previously not accessible to the UD research community through local or national resources, drove the DARWIN architecture design, including:

1. **Scalability:** DARWIN allows researchers to scale their applications and algorithms to computational node counts and data sets beyond what is available today. While the aggregate power of local clusters at UD, such as the Caviness system, is substantial, the

resources are purchased by individual research groups, and scheduling is optimized to support funding stakeholders' needs. DARWIN enables improved scalability, with two important outcomes: (1) Faster research results - many projects are now handling larger problems, enabled by complex computational models and big data sets. The execution time can be excessive, however. DARWIN will be a critical enabler of such research. (2) Staging for national resources - allocation proposals for national resources, such as XSEDE-allocated machines [3], require demonstrations of application scaling to systems on large campus parallel systems. DARWIN will enable such demonstrations.

2. **Large memory and storage:** Data-intensive projects are large in how they utilize datasets, create intermediate files, produce output and/or utilize data memory space. DARWIN is a unique enabler of such work, providing large file storage and compute nodes that include large to very large memory.

3. **Multi-GPU applications:** Several UD projects make use of multiple advanced Graphics Processing Units (GPUs) to accelerate their computations. A number of DARWIN nodes are equipped with GPUs and NVIDIA NVLink technology, allowing these computations to exploit multiple GPUs both within system nodes as well as across nodes.

4. **Data-compute co-location:** Transfer times of large data volumes have become significant time factors in current research projects. By co-locating data storage and compute facilities, connected through fast networks, transfer times can be reduced substantially. DARWIN is physically located in the same building as the main UD storage facility hosting many of the large data volumes used by the involved researchers.

5. **On-demand, interactive, and real-time access:** Batch execution is the common operating mode of large computational and data resources today. Research projects, in particular those processing real-time data and responding to external stimuli, increasingly depend on real-time access to large resources. The new instrument will enable such modes being developed by the campus IT, including fast access to available resources, quasi real-time, automatic scheduling of large resource requests, and preemption for high-priority requests.

6. **Flexibility of resource allocation:** Special projects can be key enablers of advanced research ideas, broadened participation, and partnerships. Even startup allocations on national resources have a response time of one to several weeks. DARWIN provides fast decision processes of allocation requests for emerging science pilot projects, activities that lead to broad participation, requests that may enable new operating modes, and initiatives for new partnerships.

7. **Advanced Visualization:** The ability to inspect and reason about data sets visually is a need that has become critical in computational and in particular in big-data applications. In many image processing and artificial intelligence problems visual analytics has become a primary method of investigation, requiring high-power graphics capabilities of the

underlying hardware. DARWIN is the first UD system designed with the goal of providing advanced visualization accelerators.

8. **Training access:** Training the next generation of computational & data scientists and engineers is a key need expressed by researchers at UD, our regional partners, and nationally. DARWIN enables trainees to experiment with large-scale computational and data-intensive applications. While educational access to national resources is typically associated with special training events, trainees at UD and partners will be able to request flexible educational access to the new instrument, as needs arise.

Table 1 indicates the relevance of these nine features for the UD research projects that drove the DARWIN proposal. The features prompted a system architecture that includes 105 compute nodes with a total of 6672 cores, 22 GPUs, 100TB of memory, and 1.2PB of disk storage.

Compute Nodes: Figure 2 shows the configuration of the 105 compute nodes. 92 of these nodes include of 2 AMD processors (Epyc™ 7502, 2.5GHz, 32 cores/64 threads, 128MB cache), each. They offer four sizes of memory: 48 standard nodes with 512 GiB, 32 large-memory nodes with 1024 GiB, 11 extra-large memory nodes with 2048 GiB, and one extended-memory node with 1024 GiB of regular and almost 3TiB of Intel Optane memory. 13 nodes feature GPU accelerators: 3 nodes with 4 V100 Nvidia GPUs each, 9 nodes with a T4 Nvidia GPU, and one node with an AMD MI50 GPU. The 3 V100 nodes include Intel processors; the other 10 GPU nodes include the same AMD processor configuration as the non-GPU nodes.

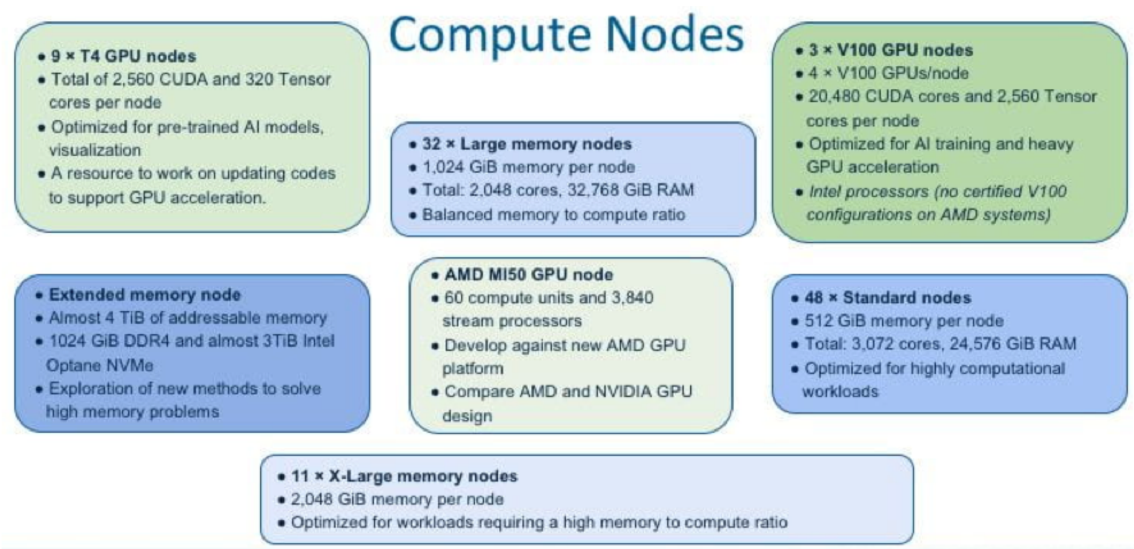

Figure 2: Types and Configuration of DARWIN Compute Nodes

Storage: A shared Lustre scratch is available to all compute nodes. The Lustre scratch file system achieves 36,000 IOPS and provides 978 TiB usable data space (11 PiB RAW). The Lustre OSS nodes are configured in a high availability configuration to ensure data integrity in the event of an OSS failure. This filesystem supports all scientific data on the system. The management nodes share an NFS filesystem to support the standard software library.

Network: The backbone network supporting DARWIN is Mellanox HDR-100. The InfiniBand network ensures high bandwidth and low-latency workloads can take full advantage of each node's 100Gb/s network capability. This network supports both compute and data functions.

Software Environment: DARWIN uses the SLURM scheduler to manage and track workload and allocation usage. All nodes run the latest CentOS 7 operating system. Ganglia is used to monitor and report node status and XDMoD [4] allows for analysis of scheduler performance. The login nodes have SSH configured for system access. Research Software, such as compilers, scientific libraries, and other common software components supporting both computational and data-intensive applications is being installed by IT support resources in support of device operation and research needs.

DARWIN Allocation Procedures

80% of the resources are allocated to users at UD at no cost through a competitive proposal process (dsi.udel.edu/core/computational-resources), with regular proposals being accepted twice a year, and startup as well as educational requests being accepted at any time. 20% of the resources go to external partners via the XSEDE XRAC (xsede.org/allocations) allocation mechanism. Allocations to not-for-profit organizations are at no cost. Resources to for-profit partners are provided on a case-by-case basis. These allocations procedures will be reviewed annually by the DARWIN Advisory Committee.

Conclusions

DARWIN is both a machine resource for computational and data science at the University of Delaware (UD) and a project to advance these sciences in the greater Delaware region. DARWIN is sponsored by the National Science Foundation (NSF) with cost sharing from UD. It leverages the recently established Data Science Institute (dsi.udel.edu) as well as UD's research network in computational science (sites.udel.edu/computational-science). The machine features 105 heterogeneous 64-core compute nodes, offering a variety of memory sizes and GPU accelerators in addition to disk storage of approximately 1 Petabyte. DARWIN supports the breadth of science across all UD colleges and their partners and is expected to significantly boost computational and data-intensive research.

References

[1] **The National Strategic Computing Initiative (NSCI),** https://www.nitrd.gov/nsci/

[2] National Science Foundation award 1919839, **MRI: Acquisition of a Big Data and High Performance Computing System to Catalyze Delaware Research and Education,** https://nsf.gov/awardsearch/showAward?AWD_ID=1919839

[3] John Towns, Timothy Cockerill, Maytal Dahan, Ian Foster, Kelly Gaither, Andrew Grimshaw, Victor Hazlewood, Scott Lathrop, Dave Lifka, Gregory D. Peterson, Ralph Roskies, J. Ray Scott, Nancy Wilkins-Diehr, "**XSEDE: Accelerating Scientific Discovery**", Computing in Science & Engineering, vol.16, no. 5, pp. 62-74, Sept.-Oct. 2014, doi:10.1109/MCSE.2014.80

[4] Jeffrey T. Palmer, Steven M. Gallo, Thomas R. Furlani, Matthew D. Jones, Robert L. DeLeon, Joseph P. White, Nikolay Simakov, Abani K. Patra, Jeanette Sperhac, Thomas Yearke, Ryan Rathsam, Martins Innus, Cynthia D. Cornelius, James C. Browne, William L. Barth, Richard T. Evans, "**Open XDMoD: A Tool for the Comprehensive Management of High-Performance Computing Resources**", *Computing in Science & Engineering*, Vol 17, Issue 4, 2015, pp. 52-62, doi:10.1109/MCSE.2015.68