

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.  
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) 08-03-2022	2. REPORT TYPE Final Report	3. DATES COVERED (From - To) 1-Mar-2021 - 28-Feb-2022
---	--------------------------------	--

4. TITLE AND SUBTITLE Final Report: Redefining the Concept of Large-Scale Computational Dynamics: Billion-Degree of Freedom Terramechanics Simulation on Affordable Hardware	5a. CONTRACT NUMBER W911NF-21-1-0029
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER 611103

6. AUTHORS	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES University of Wisconsin - Madison Suite 6401 21 N Park Street Madison, WI 53715 -1218	8. PERFORMING ORGANIZATION REPORT NUMBER
--	--

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211	10. SPONSOR/MONITOR'S ACRONYM(S) ARO
	11. SPONSOR/MONITOR'S REPORT NUMBER(S) 77321-MA-RIP.1

12. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.
--

13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.
---

14. ABSTRACT
--------------

15. SUBJECT TERMS
-------------------

16. SECURITY CLASSIFICATION OF:	17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Dan Negrut
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU	19b. TELEPHONE NUMBER 608-890-0914

**RPPR Final Report**  
as of 08-Mar-2022

Agency Code: 21XD

Proposal Number: 77321MARIP

**Agreement Number: W911NF-21-1-0029**

**INVESTIGATOR(S):**

**Name:** Dan Negrut  
**Email:** NEGRUT@ENGR.WISC.EDU  
**Phone Number:** 6088900914  
**Principal:** Y

**Name:** Radu Serban  
**Email:** serban@wisc.edu  
**Phone Number:** 6082656124  
**Principal:** N

Organization: **University of Wisconsin - Madison**

Address: Suite 6401, Madison, WI 537151218

Country: USA

DUNS Number: 161202122

EIN: 396006492

**Report Date:** 31-May-2022

Date Received: 08-Mar-2022

**Final Report** for Period Beginning 01-Mar-2021 and Ending 28-Feb-2022

**Title:** Redefining the Concept of Large-Scale Computational Dynamics: Billion-Degree of Freedom Terramechanics Simulation on Affordable Hardware

**Begin Performance Period:** 01-Mar-2021

**End Performance Period:** 28-Feb-2022

**Report Term:** 0-Other

Submitted By: Dan Negrut

Email: NEGRUT@ENGR.WISC.EDU

Phone: (608) 890-0914

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

**STEM Degrees:**

**STEM Participants:**

**Major Goals:** Goal of the project was to update an existing CPU/GPU cluster that is used for HPC by Negrut's lab.

**Accomplishments:** The cluster has been updated with the hardware listed in the attached document.

**Training Opportunities:** Nothing to Report

**Results Dissemination:** Nothing to Report

**Honors and Awards:** Because of the HPC work enabled by the updated supercomputer, our lab has been one of the four finalists in the Best Poster competition organized in St. Louis, MO, in conjunction with the Supercomputing 2021 conference.

**Protocol Activity Status:**

**Technology Transfer:** Nothing to Report

**PARTICIPANTS:**

**Participant Type:** PD/PI

**Participant:** Dan Negrut

**Person Months Worked:** 1.00

**Funding Support:**

Project Contribution:

National Academy Member: N

**RPPR Final Report**  
as of 08-Mar-2022

**Partners**

,

I certify that the information in the report is complete and accurate:  
Signature: Dan Negrut  
Signature Date: 3/8/22 3:42PM

## **Hardware Upgrade – DURIP Program W911NF1810476**

PI: Dan Negrut – University of Wisconsin-Madison

### **Assets purchased:**

- \$81,000 – 2× Supermicro GPU Server
  - 2× AMD EPYC 7272 CPU
  - 4× NVIDIA A100 40GB SXM4 GPU
  - 4TB NVMe SSD
  - 100 Gb/s Ethernet Adapter
  
- \$11,500 – 1× Virtualization Server
  - 1× AMD EPYC 7282 CPU
  - 2× NVIDIA RTX A5000 PCIe GPU
  - 1TB PCIe 3.1 SSD
  - 100Gb/s Ethernet Adapter
  
- \$107,237.69 – Network Upgrade (to 100 Gb/s)
  - 2× HPE SN3700M 32-Port Switches
  - 20× 100 Gb/s Active Optical Cables
  - 26× Mellanox QSFP Splitter Cables (200 Gb/s → 2×100 Gb/s)
  - 72× 100 Gb/s Ethernet Adapters

### **Euler Upgrade Process 2021**

The entire planning, purchasing, and deploying of both the computational and network components for this upgrade were done by a part-time System Administrator for the Euler Cluster. Upgrading Euler in the midst of the COVID-19 pandemic was challenging owing to parts shortages.

#### **Pre-planning**

The Sysadmin began assimilating information to be used in the cluster upgrade in February 2020 in response to perceived deficiencies, e.g., in the performance of distributed (MPI) tasks, in access to large datasets over the network, and in maintaining a software stack able to support the latest computational hardware. This involved a deep-dive into the capabilities of various types of high-speed networks including InfiniBand, OmniPath, and RoCE.

#### **Planning**

Once instrumentation was awarded in December 2020, the Sysadmin used the prepared information to propose a two-pronged approach to infrastructure upgrades:

1. Approximately half of the funding would go toward acquiring the latest GPU technology from NVIDIA. Due to supply chain issues stemming from the ongoing COVID-19 pandemic and a sharp rise in the popularity of cryptocurrency mining, consumer GPUs such as the Geforce RTX 3080 were selling for more than double their recommended value and were almost impossible to obtain even at those prices. Instead, it was determined that the best option would be to acquire something that was out of reach for cryptocurrency, such as the Tesla A100 datacenter GPU.

2. The other half would be used to fund a complete overhaul of Euler's network infrastructure. This would require new network switches, 100Gb network adapters for each server, and the cables to connect all of the machines to the switches. This would also enable a more sophisticated and higher-performing filesystem which would be deployed during the maintenance window for the network upgrade.

Pandemic-related supply chain issues meant that despite beginning the acquisition process in late January 2021, vendors struggled to predict product availability and lead times. After vendor negotiations concluded in late February, it was estimated that the hardware would arrive at some time during April. As a result, a 12-day maintenance window was planned for the end of May 2021, in which all of the changes would need to be implemented. When this was decided, the University of Wisconsin – Madison campus was still in a partial lockdown; while many buildings could be entered, much of the academic staff had not yet resumed in-person operations because vaccines were still restricted to specific types of essential workers.

Further complicating the approaches detailed above was the catastrophic announcement that Euler's operating system, CentOS 8, would be discontinued nearly seven years before its originally advertised End-of-Life. Purchasing licenses for Red Hat Enterprise Linux would be prohibitively expensive, but maintaining compatibility with the latest hardware would be similarly difficult without a widely-accepted target platform.

### **Preparations**

In the months and weeks leading up to the maintenance window, in addition to his normal maintenance and classroom assistance duties, the Sysadmin conducted exhaustive testing on potential replacements for CentOS, assisted users with the migration and cleanup of roughly 100TB of stored data, drafted new configuration scripts, and collected system backups in order to ensure that the maintenance window could be entirely devoted to implementing the planned system changes.

Several problems arose as the deadline approached, not the least of which was a number of manufacturing issues, shipping delays, and last-minute availability changes that caused some of the network components to arrive as late as just a few days before the deadline.

### **Installation**

Four student volunteers were able to provide time in shifts adding up to roughly 20 person-hours of physical labor. This became invaluable as installing switches in racks and moving servers within range of cables could not be safely completed by a single person. The layout of the server room necessitates that cables are run in overhead channels, requiring the installer to climb atop a ladder while manipulating the cable. An estimated 600 meters of cable would need to be moved in this way in order to perform the upgrade, and another 300 meters would need to be run inside of the server racks. This process took the first six days of the maintenance window, totaling 80 person-hours of labor, 60 of which were performed by the Sysadmin. The second half of the maintenance window was spent performing a complete overhaul of the cluster's software infrastructure, including dual stack networking, a new Identity Management (IdM) deployment, a new Ceph filesystem deployed from scratch onto a cluster of seven servers, a new deployment of Foreman (the upstream of RedHat's Satellite), and the reprovisioning of the system's GPU compute nodes. This period took a little over 60 hours of labor and was performed by the Sysadmin.

### **Post Install and Performance Observations**

After the maintenance window completed, several weeks of time were dedicated to assisting users with the migration of their existing work onto the new platform. The remaining (low traffic, CPU-only) compute nodes were eventually reprovisioned as user demand grew over the summer term. The resulting platform has showed greatly improved network performance, showing latencies as low as 890 nanoseconds during RDMA benchmarks; nearly 1000x faster than the former configuration. Filesystem performance has improved as well with a parallel write throughput of roughly 2.5 GB/s, over 4x previous speeds. Read

throughput has improved even more impressively, delivering data at around 11GB/s, over 20x faster than the original implementation.

### Use of cluster


The equipment that was purchased through this award was added to Euler, a cluster that is more than a decade old but has been constantly upgraded throughout its service. Euler is a co-op that several faculty from the College of Engineering contributed towards. By a large margin, Negrut’s lab has the largest share of hardware (about 80%). Each co-op contributor has priority over his/her share. However, when his/her share is not used, other users can tap into it.

Negrut’s research would not be possible without this cluster, which anchors his lab, which is called “Simulation-Based Engineering Lab (SBEL)”. As the name suggests, the entire lab works off Euler since computer modeling and simulation is our focus. We use the cluster in multi-phase fluid dynamics, granular flows, terradynamics (for ongoing Army & NASA projects), autonomous vehicle simulation, human-in-the-loop simulation. All the papers (journal & conference) that come out of our lab draw on this cluster. As a measure of recognition of our work, the poster we submitted to the 2021 Supercomputing Conference in St. Louis, held in November, was one of the four finalists for the Best Poster Award in the entire Supercomputer Conference. The remarkable thing is that everybody who contributed to that poster was a Mechanical Engineer, being recognized at a Supercomputing Conference where the focus is on the TOP500 supercomputers in the world, exascale computing, etc.

Finally, the cluster is used in instructional activities, where students get accounts for one semester to learn about HPC. For instance, I’ll be teaching in Fall the class *ME759: High Performance Computing for Applications in Engineering*. This class is a PhD level class that draws between 70 and 100 graduate students each time is taught. Because of the DURIP funding, which over the years allowed me to keep this cluster up to date, this highest-level class in the College of Engineering at UW-Madison became the largest graduate course within the entire college (made up of 12 departments).



**Figure A.** Schematic of Euler. The CPU-GPU cluster has more than 80 GPU cards using the Maxwell, Pascal, Volta, and Ampere architectures.



**Figure B.** Partial view of the Euler heterogeneous cluster. Shown are several CPU/GPU nodes.