Center for
HIGH PERFORMANCE COMPUTING
THE UNIVERSITY OF UTAH

Spring 2022

An Essential Need for HPC in Nuclear Engineering Applications

Meng-Jen (Vince) Wang, Glenn E. Sjoden, Department of Civil and Environmental Engineering

The University of Utah's Nuclear Engineering Program (UNEP) has greatly benefitted from the tremendous resources afforded for research by the University of Utah's Center for High Performance Computing (CHPC). The modern reality is that nuclear applications cannot be practically fielded without an extraordinary amount of modeling and simulation carried out prior to actual testing, evaluation, or new engineering endeavor. For this reason, we must maintain access to "premier" supercomputing resources, such as those available via CHPC, simply because models that require detailed radiation transport computations are computationally intense. Overall, radiation transport simulations consist of two main approaches: deterministic solvers that require solutions of many billions of simultaneous equations, or Monte Carlo based solvers that require billions of particle histories to simulate, in addition to proper pre-processing of nuclear data to enable complex system computations. UNEP maintains HPC parallel optimized codes that can utilize both deterministic and Monte Carlo solution approaches; in fact, most often, we must apply both approaches to validate our findings to ensure all nuances of profiling "where the radiation goes" is fully explored. This is where having the resources enabled by CHPC is essential; these HPC frameworks are comprehensively pressed to execute extremely accurate models, enabling us to achieve what others may not be able to, allowing us to proceed with novel research objectives with the knowledge that we can successfully achieve our end goals.

For example, using both NOTCHPEAK and REDWOOD at CHPC, we have successfully executed simulations of highly accurate 3-D models of the University of Utah Training Reactor (UUTR)[1]. The UUTR on the Utah main campus is a 100 kW General Atomic Mark I TRIGA pool nuclear reactor with four main irradiation facilities, including the Central Irradiator (CI) with the highest flux at the center of the core, Thermal Irradiator (TI) with a trapezoidal tank filled with D_2O, a Pneumatic Irradiator (PI) sample-transfer system ("rabbit" system), and Fast Neutron Irradiation Facility (FNIF). The UUTR is used for neutron and gamma irradiation, radioisotope production, electronics hardness testing, neutron activation analysis, and other applications. Figure 1 shows a model central slice through our latest reactor model obtained with the use of the deterministic PENTRAN code executed on NOTCHPEAK. Results (a snapshot of those is depicted in Figure 2) from this work were also validated using the Monte Carlo code MCNP6 on REDWOOD. As a result, we have very high confidence in our results, allowing us to apply the resulting reactor parameters to new experiments to optimize use of radiation fields in our research.

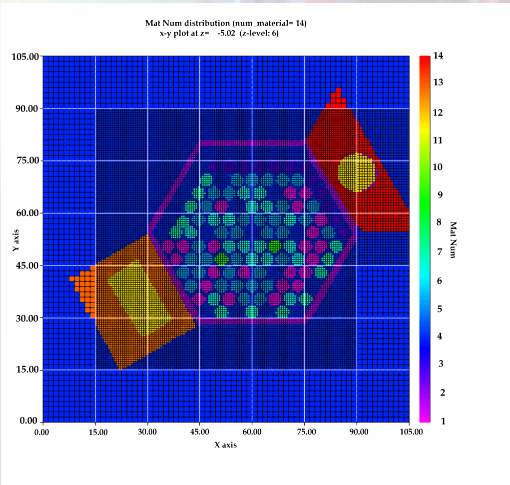

Fig. 1. Cross section of a 3-D deterministic radiation transport model used for a PENTRAN calculation on NOTCHPEAK for determining reactor critical flux.

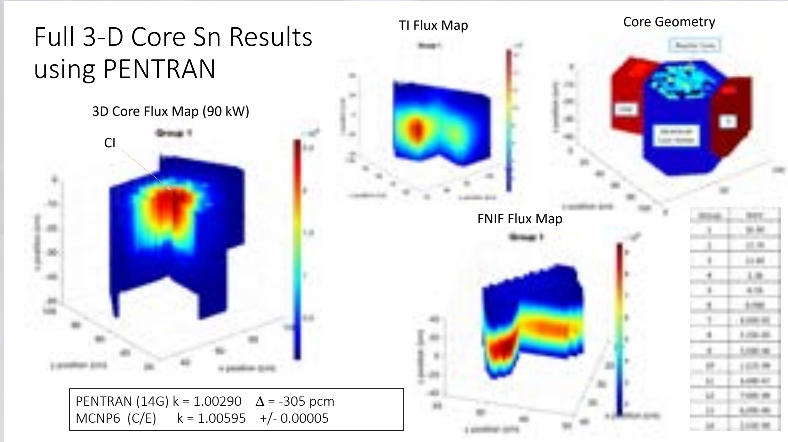

Fig. 2. Results of neutron flux models of the UUTR Reactor computed using CHPC resources supporting nuclear research at the University of Utah.

Based on these results, we were able to complete the design of new neutron beam port and source configurations [2, 3] in our reactor, and complete a new beam stop and imaging chamber design [4], highlighted in Figures 3a and 3b.

These results then enabled us to design, optimize, and perform testing in our UUTR reactor to support a new methodology for Xe-135 gas standards for calibration studies used by the Comprehensive Test Ban Treaty Organization (CTBTO) [5], funded through a grant with Idaho National Laboratory. The new production modality has improved extended utilization of gas standards, shipped worldwide, by more than 30% (Figure 4). Follow-on research is now pending to extend this methodology to other research reactors as a new approach to optimize radio-xenon gas samples supporting nuclear test ban treaty monitoring.

Our UNEP computational group also completed model development to enable full 3-D characterizations of burnup and depletion in nuclear reactors, with dynamic, machine controlled adaptive timesteps, error control, and representation of full neutron energy dependence [6]. This new BSOLVE algorithm, which works with our HPC parallel 3-D deterministic and Monte Carlo neutron transport solvers, was developed and tested using NOTCHPEAK and REDWOOD resources. As this work expands, we expect this capability to enable more accurate models and optimization studies for modern reactor designs. Because nuclear power produces virtually carbon free electricity, there is a new need for these computational tools.

New research work on our immediate horizon include methods to support enhancing nuclear forensics, assessing radiation damage and realizing novel space shielding concepts, producing new isotopes for nuclear medicine, and probing nuclear physics questions, all of which will require HPC based modeling and simulation. CHPC therefore will continue to play a significant role in enabling our research success. As we exploit new research to sustain nuclear energy in pursuit of carbon free energy solutions, computation remains at the heart of our ability to predict experimental outcomes, support reactor operations, and enable our technical edge. We are extremely grateful to have world class HPC support to complete our mission.

Fig. 3. (a, left) angular neutron profile of neutrons emerging from new beamline design for the UUTR. (b, right) CAD model and overlay of Monte Carlo results depicting neutron beamline; CHPC results.

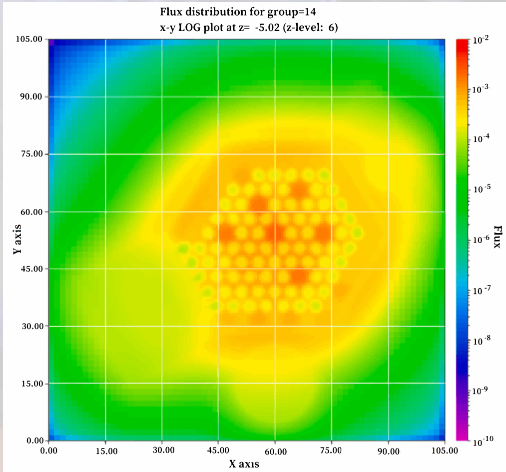

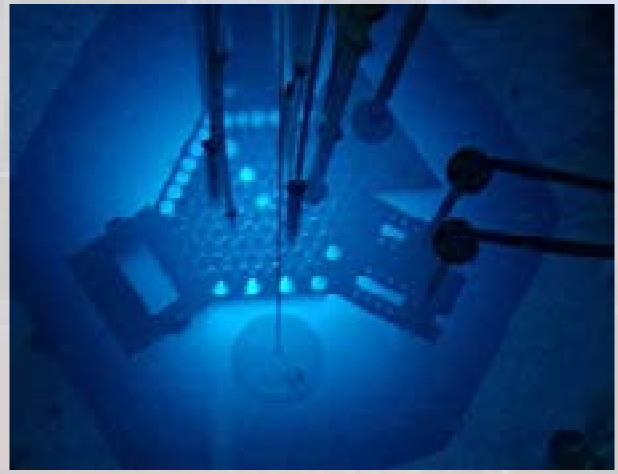

Fig. 4. (Left) Computational model of UUTR reactor with special test chamber for optimized Xe-135 samples rendered using NOTCHPEAK. (Right) Photo of actual test chamber post-design irradiation performed in the UUTR at the University of Utah.

References

1. M. Wang, G. Sjoden, A. Foley, S. Mohanty, "3D SN and Monte Carlo calculations of the Utah TRIGA reactor core using PENTRAN and MCNP6", *Annals of Nuclear Energy*, Vol. 155, 108158, June (2021),

2. M.-J. Wang, G.E. Sjoden, "Experimental and Computational Dose Rate Evaluation using SN and Monte Carlo Method for a Packaged ^{241}AmBe Neutron Source", *Nuclear Science and Engineering*, Vol. 195, Issue 11, pp. 1154-1175, June (2021).

3. M. Hartos, M. Wang, and G. Sjoden, "Computational Design and Optimization of a Neutron Imaging Beamline Using Monte Carlo and Deterministic SN Radiation Transport for the Utah TRIGA Reactor", *Nuclear Engineering and Design*, Vol. 382, 111374, October (2021).

4. M. Hartos, M. Wang, and G. Sjoden, "Design of an Ultra-Compact Imaging Chamber and Radiation Beamstop for a Neutron Radiography System Employing Particle Transport", *Nuclear Engineering and Design*, Vol. 386, 111570, January (2022).

5. Tanner W. Hall, Meng-Jen Wang, G. E. Sjoden, and Matthew Watrous, "Computational and Experimental Optimization of ^{135}Xe Production in Calibration Sources", *Journal of Environmental Radioactivity*, pp. 244–245, 106814, January (2022).

6. G. Sjoden, M. Wang, and N. Kurtyka, "BSOLVE: Energy Dependent Depletion with Algorithm-adapted Error Control for 3-D Transport", Transactions of the ANS, Vol 125, 2021 ANS Winter Meeting, December (2021). Doi.org/10.13182/T125-36621

Virtual DDC Tour Now Available!

Crystal Young, CHPC Student Staff

Ever had a question about where CHPC operated systems are housed? Or questions about how the capabilities and history of the Downtown DataCenter (DDC)?

Pre-COVID we periodically offered in-person DDC tours to showcase the University's data center where the CHPC systems are housed, allowing us to answer questions about the DDC capabilities along with CHPC supported systems and our offerings. While we hope to be able to bring back the in person DDC tours in the near future, we now have a virtual tour to offer. The virtual tour, located at *https://youtu.be/tWqjlL3Zr6g*, is a mix of 3D virtual models and actual metrics from the systems Datacenter personnel and CHPC monitors.

The 3D models were generated in Unreal Engine 4 and sequenced together with in-person images using Adobe Premier, with Adobe Premier and Open Broadcaster Software used to overlay the audio tracks. The main thing you're missing from an in-person tour is trying to hear someone over all the machines.

New Home Directory and Scratch Space in the General Environment

Sam Liston, CHPC System Administrator

The current home directory solution in the general environment will be out of warranty this spring. In addition, the Lustre scratch space, /scratch/general/lustre, is now out of warranty. Therefore, CHPC has been exploring available options to replace both of these file systems.

The solution chosen is from VAST Data, *vastdata.com*, and is an "all flash" storage solution. The solution, including the network topology, is shown in the figure on the next page.

The VAST hardware will consist of three C-Boxes (shown in the middle of the figure), each containing four C-Nodes, along with two D-Boxes, each containing two D-Nodes (shown at bottom of the figure). These twelve C-Nodes and four D-Nodes will be interconnected by a back-end network at 100gbps. The C-Nodes will be outfitted with 3D Xpoint (or non-volatile) memory. This

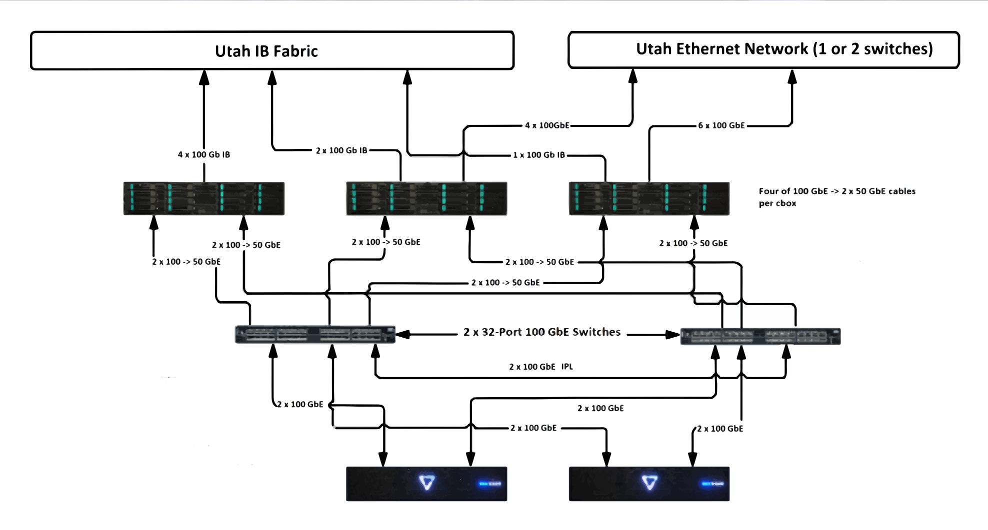

will be the initial landing point for all writes to the file systems. This written data is then massaged, ordered and de-duped before it is written to the D-Nodes, which contain large QLC NVMe drives. By massaging the data and sending it to the D-Nodes in an ordered, streamed manner, VAST is able to extend the life of the QLC flash, a form of flash storage that was created to be a low-cost flash option, but with a side effect of lower durability.

The C-Nodes will be divided between the IB and ethernet networks, with seven of the C-Nodes connected to the IB fabric and five connected to Ethernet at 100gbps each. Connectivity will be offered to both the IB and Ethernet side via NFS. SMB/CIFS access will also continue to be offered via mount paths through `samba.chpc.utah.edu`.

With the all-flash nature of the VAST system, I/O per second (IOPS) have the potential to be monumentally higher than both the current Compellent home directory system and the Lustre scratch system, as both of them rely on spinning media for data storage. With our particular hardware configuration VAST has quoted CHPC the following expected performance characteristics: 63GB/s reads, 10GB/s writes, 605K read IOPS, 227K write IOPS.

Because of these excellent performance characteristics, CHPC will be using this system to provide both home directory and scratch space. Combining these two file systems, which usually are architected differently purposes, seems odd. However, within the VAST system the two personalities will be logically separated and, for the most part have dedicated hardware to support each. The seven IB connected C-Nodes will be used for the writes to the scratch file system and the five Ethernet connected C-Nodes for writes to the home directory space. There will be a few outlier usasge cases that will connect to scratch via ethernet or home via IB.

The flash storage in the D-Nodes will be configured using 80+4 erasure coding. With this parity of 4, the mean time before data loss is estimated at 60 million years. This resiliency coupled with the performance of flash hardware as well as features like composable QoS (Quality of Service) to manage noisy/contentious neighbors and fully online upgrades and expansion, gave us confidence to pursue a combined home and scratch system.

The physical usable capacity of the system will be 1.2 PB. After running a data-profiling suite of software provided by VAST, which provides an estimation of data reduction, the total capacity of the system after dedupe should be around 2.0 PB.

CHPC will continue to run the DDN Lustre system /scratch/general/lustre in an out-of-maintenance mode, for the foreseeable future. It should be used in a strictly 'scratch' mode, as software failure is a possibility and without support recovery may not be possible.

Whereas the Compellent cost was $1250/TB for the 5 year warranty lifetime, this new solution has a cost of $900/TB for the 5 year warranty lifetime. This includes the cost of the space on the VAST storage system and backup of this space. The backup will be to the CHPC object storage, pando, and will be a weekly full backup with nightly incremental backups, with a two week retention window.

CHPC will be taking delivery of the new system in March. After installation and testing we will be reaching out to all groups who have purchased home directory space on our current Compellent home directory, to discuss moving to the new solution. We anticipate that this will be sometime in May.

Cluster OS Upgrade

Brian Haymore, CHPC System Administrator

During March, CHPC will be upgrading the OS of the clusters from CentOS7 to RockyLinux8.

In December 2020, as CHPC was in the process of planning an OS update from CentOS7 to CentOS8, RedHat announced a shift in the CentOS Project.

CentOS (Community Enterprise Operation System) is a

Rocky Linux™

free, community drive rebuild of the commercial RedHat Enterprise Linux (RHEL). While Centos started as an independent organization, it joined RedHat in 2014. The announcement was that the existing CentOS, which is a downstream distribution of the RHEL offering, would be replaced with a new CentOS Stream offering. This new offering would be an upstream distribution of RHEL, or a development version of upcoming RHEL releases. From RedHat's perspective, this provides the value of having additional testing from sites adopting the early release. From the adopters' perspective, the new releases will not be as well tested as the current downstream versions, and could be more prone to having regression or compatibility issues that will need to be addressed. Additional consequences of this change is that there would no longer be a distribution with a long lifecycle (10+ years) as was available with the original CentOS offering, and the changes between versions could result in issues in support requirements for the drivers needed for certain hardware or file systems (such as Infiniband, lustre, GPUs) as well as for some software packages.

The announcement included that CentOS8, which was released in September 2019, would be "End of Life" or EOL as of December 31, 2021, and would not have the expected life cycle commitment of earlier versions. This announcement did not impact the CentOS7 life cycle, which has an EOL June 30, 2024.

With that announcement, CHPC, as well as other users of CentOS which includes many campus HPC centers, started to explore alternatives to CentOS. There are a number of choices, including some new offerings. The choices include the new CentOS Stream, as well as existing Ubuntu/Debian, Oracle Linux, RedHat Enterprise Linux, and OpenSuse. In addition, after the CentOS announcement, two new efforts of interest were announced; AlmaLinux and RockyLinux. The AlmaLinux project was founded by CloudLinux with the intent to create a community enterprise Linux distribution that is 1:1 binary compatible with RHEL. The RockyLinux project was founded by Greg Kurtzer, one of the founders of the CentOS project, and named after Rocky McGaugh who was also a CentOS project co-founder. RockyLinux is also 1:1 binary compatible with RHEL.

The decision was made to stay with a solution that was binary compatible with RHEL as well as to stay with a community driven, no cost solution. The reasoning for the first choice is that CHPC is already familiar with the driver, application, and software support of RHEL binary compatible OS offerings, as well as we already have a large number of end user applications, more than 700 at this time, that work in this OS. In addition, we wanted an option with a lifecycle that would provide a long-term stability. This narrowed the choice to either AlmaLinux or RockyLinux. Initial explorations at CHPC indicated that either would be a viable solution. Looking at the Academic Research Computing Community, RockyLinux, *https://rockylinux.org/*, is the offering that is being favored. This lead CHPC to decide to explore this offering in more detail.

As with past OS upgrades, CHPC first completed the OS change on a small test cluster on which CHPC staff could perform testing. This includes CHPC User Support testing some of the applications we have installed. Our initial tests showed that most of the currently installed installations should have no issues running on the RockyLinux OS installation. However, as it is impossible for us to test each of the many end user applications, so we need to rely on users to do additional testing. Therefore, we also set up a time table which includes upgrading the clusters in phases.

The first upgrade to RockyLinux8 will include the compute and interactive nodes of lonepeak along with the frisco nodes. The upgrade will happen in the next CHPC downtime, scheduled for March 22, 2022. This will allow users to test their workloads on lonepeak. However, as lonepeak does not have an infiniband (IB) network, workloads that are dependent on this cannot be tested. For these workloads we will provide access to this test cluster, which has IB, and we strongly recommend that users test their applications before we upgrade the other production clusters. The schedule for the OS upgrade on the other clusters will be determined based upon the feedback from the userbase on issues that arise from the testing on the lonepeak cluster and the frisco nodes.

Recent Changes in the PE

Anita Orendt, CHPC Scientific Consultant

There were three recently announced changes to the CHPC Protected Environment (PE).

Addition of a "redwood-shared-short" partition

The `redwood-shared-short` partition consists of two nodes, each having 28 physical cores and 128 GB memory. These nodes are available for use by all users, regardless of a user's access to a general allocation; the use of these nodes will not count against any allocation. To use set both the slurm partition and the account to `redwood-shared-short`. The partition has also been added as an option when launching a job using the pe-ondemand instance.

This partition is set up to maximize the throughput of short jobs. As node sharing is being used, users MUST specify the number of cores. In addition, you can specify the amount of memory, unless you are fine with using the default of 2 Gb/core. See *https://www.chpc.utah.edu/documentation/software/node-sharing.php* for additional details on node sharing.

Use of these nodes is limited to the following constraints:

- Maximum wall time is 8 hours
- Maximum running jobs per user is 2
- Maximum cores per user is 8
- Maximum memory per user is 32 GB
- Maximum cores per job is 8
- Maximum memory per job is 32 GB

Addition of two interactive nodes

The new `bristlecone1.chpc.utah.edu` and `bristlecone2.chpc.utah.edu` nodes are the equivalent of the frisco nodes in the general environment. Each of these nodes has 28 physical cpu cores and 64 GB memory.

While arbiter is still running on these nodes, the settings and limits are higher than they are on the `redwood1.chpc.utah.edu` and `redwood2.chpc.utah.edu` interactive nodes. On the bristlecone nodes the cgroups settings are 12 cores and 24 GB memory (versus 4 cores and 8 GB memory on the two redwood interactive nodes available to all users). In addition, the threshold or trigger levels are at 4 cores and 16 GB memory versus 1 core and 4 GB memory, and the time before reaching penalty state is 60 minutes instead of 15 minutes. As with the frisco nodes, these nodes are accessible via the PE's ondemand instance. For more details on the acceptable usage on the generally available interactive or login nodes as well as the arbiter limits see *https://www.chpc.utah.edu/documentation/policies/2.1GeneralHPCClusterPolicies.php#Pol2.1.1*.

Change in the backup policy of the project space directories

As most of you know, CHPC did a refresh of the Protected Environment (PE) in 2017 with funds from a NIH Shared Instrumentation Grant. This award was used for, among other items, the purchase of the home directory, project space, and archive storage in the PE. As the grant funded the hardware, we instituted a policy that all home and project spaces would be backed up free of charge.

Due to the tremendous growth in the number and scope of the projects making use of the PE, since the initial hardware purchases in 2017 CHPC has, with CHPC funds, expanded the resources available. The expansion including growing the original 160 TB usable home/project space to 982 TB usable space (with 830 TB committed) and the archive space from 1 PB raw space to 3 PB raw space.

Now that the grant period (April 2017-2022) is coming to an end CHPC has revisited the backup policies and will be making changes to be in line with the policies for the group space in the general environment.

Effective July 1, 2022, the PE project spaces will no longer be backed up free of charge. Groups that want to continue to have CHPC back up their project space will need to purchase the amount of archive space to allow for this. Note that arrangements can be made for a subset of the data in the project space to be backed up and there are also user driven backup strategies that groups can employ. If you want to explore options, please contact CHPC via *helpdesk@chpc.utah.edu*.

CHPC will continue to back up PE home directories, set at 50 GB per user, on the current weekly full backup with nightly incremental backup schedule. In the PE we do not allow for larger home directory spaces.

For more information on the PE and its policies and resources, see the PE page on the CHPC website *https://www.chpc.utah.edu/resources/ProtectedEnvironment.php*.

Introducing the CHPC Deep Learning Module

Brett Milash, CHPC Scientific Consultant

CHPC encourages users to install and maintain their own copy of python and associated libraries (see *https://bit.ly/3t0KzML*). However, we get many requests for the installation of up-to-date deep learning packages in python for various machine learning projects. Therefore, CHPC has committed to maintaining a module (named deeplearning) with recent versions of python, tensorflow, keras, and pytorch. Our intent is to update this module as needed every six months to keep it relatively current.

Implementation

The deeplearning module is implemented as a set of shell functions that are added to your environment when you load the module (`module load deeplearning`). These shell functions in turn execute commands within a singularity container. The shell functions are:

- `python` – starts a python interpreter with access to the various deep learning libraries
- `jupyter` – starts the jupyter process for executing a jupyter notebook

- `jupyter-lab` – starts the jupyter-lab process for working in the jupyter lab environment

The singularity container that contains the deep learning software is built upon the docker image "jupyter/tensorflow-notebook", and as such is compatible with our Open OnDemand web portal environment. To use the latest deeplearning module for OnDemand jupyter sessions, select "CHPC Deep Learning" from the "Jupyter Python version" menu in OnDemand.

Contents

The deeplearning module contains python version 3.9.7 along with the following python libraries: tensorflow (2.6.2), keras (2.6.0), pytorch (1.10.2+cu102), torchvision, torchaudio, ipykernel, numpy, scipy, sklearn, skimage, seaborn, pandas, and PIL. The module also provides GPU support with the CUDA 11.2 libraries, making this module compatible with all of CHPC's GPUs with the exception of the NVIDIA Tesla K80 GPUs found on the kingspeak cluster.

Documentation

The deeplearning module is documented at: *https://bit.ly/3Kx1cWn*. If you have any questions or comments, please contact us at *helpdesk@chpc.utah.edu*.

Initial Insights from the CHPC User Survey

Brett Milash and Anita Orendt, CHPC Scientific Consultants

CHPC ran its most recent user survey during the Fall of 2021 and Winter of 2022 using the University's Qualtrics license. The survey itself was based on a similar effort (RMACC Cyberteam Cyberinfrastructure Survey Results, *https://osf.io/g3hjs/*) conducted by the RMACC Cyberteam, an NSF-funded multi-institutional support team, but was adapted to the specific resources available at the Center for High Performance Computing. This article highlights a few insights gained from a first look at the results of the survey.

Background

Of the 8342 users invited to take the survey, 458 (5.5%) responded. We believe this low response rate is due in part to the number of inactive users in our mail list, and that the response rate among active users is much higher. Of those responding, 391 (85.4%) are from the University of Utah, 36 (7.8%) are from the Utah State University, and the remaining 33 (7.2%) are from other institutions. The general area of research among the respondents shows a continuing trend in the increasing number of researchers in the life sciences; biology, biomedical research, genomics, and ecology comprised 35.3% of the respondents to the survey.

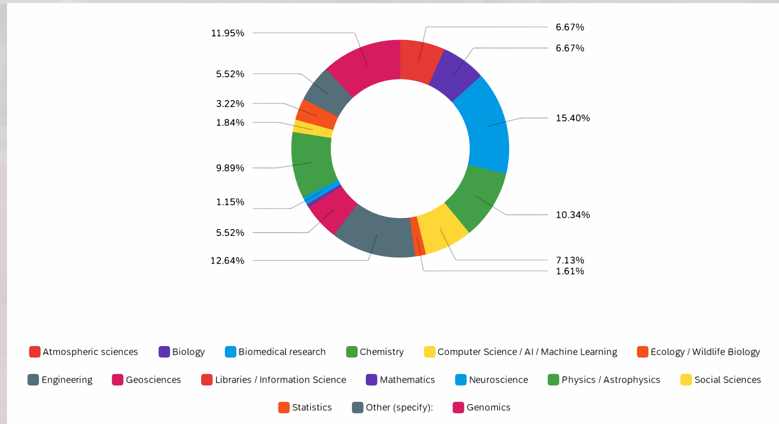

General level of satisfaction

Our survey asked about the general level of satisfaction or dissatisfaction with a variety of resources offered at CHPC.

The survey shows that users are generally satisfied with the resources we offer. A surprising number of users are unaware of some of CHPC's resources, for example the virtual machines, Open OnDemand server, and Windows servers.

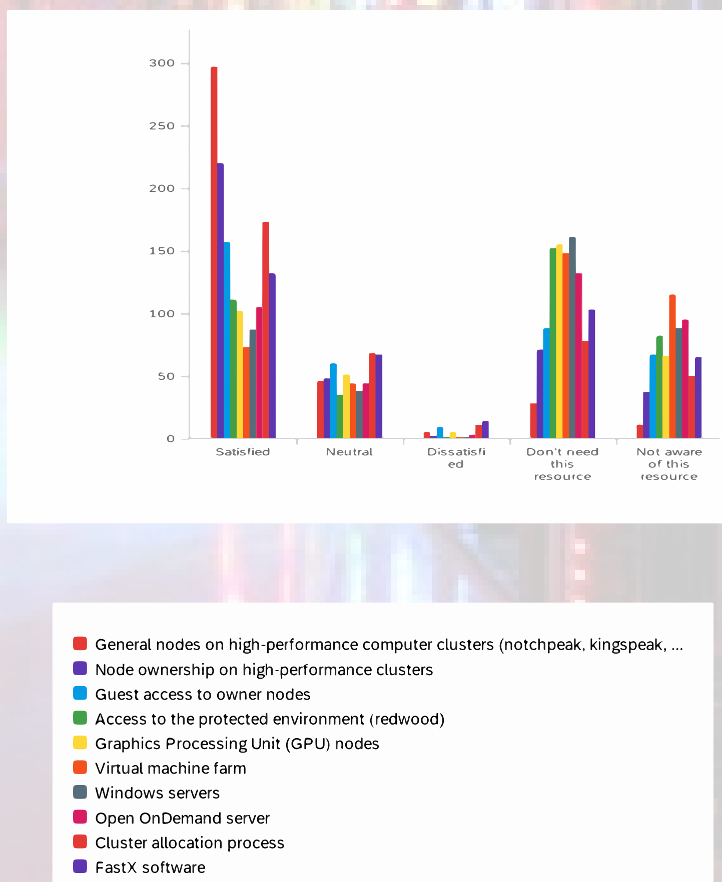

Where are users dissatisfied?

The three areas showing greatest user dissatisfaction are the use of Fastx software, the cluster allocation process, and guest access to owner nodes. The main source of dissatisfaction with Fastx seems to be its perceived lack of reliability. We hope this will be addressed by the recent upgrade from Fastx version 2 to version 3. Dissatisfaction with the cluster allocation process stems from allocations being too small, or the inability to have multiple allocations (for separate projects) at the same time. We have recently increased our maximum allocation size from 250K core-hours to 300K core-hours per quarter, and we hope this will address some users' allocation size concerns. For users with projects requiring more compute time, we frequently suggest they explore the large-scale compute capabilities at XSEDE (*https://xsede.org*) or the Open Science Grid (*https://opensciencegrid.org*). Some users' concerns about owner node access include job run times being too short and the challenges presented by pre-emption by owner's jobs. Since the jobs of owners of the hardware have priority over guest jobs, guests can adapt to job preemption through techniques such as checkpointing and automatic restarting of preempted jobs (*https://bit.ly/3Jkx545*).

Barriers

Our survey asked if "any of the following activities are barriers to successfully executing your research data workflows?" The choices included accessing computer clusters, using the batch system, programming, using Linux, data transfer, using domain-specific application software, data visualization, working with data formats, using digital repositories, and publishing data. The survey shows that visualizing data, programming, and data transfer are the activities most likely to present a barrier sometimes or always.

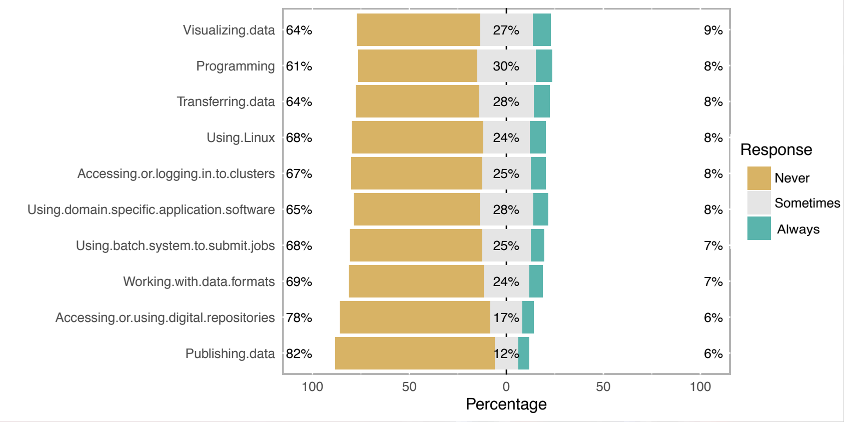

Conclusions

This initial look at results of the survey has revealed some interesting trends and opportunities for improvement at CHPC, for example through alterations of policies or improved outreach to users about the services we provide. We will dig deeper into the survey , and highlight changes at CHPC due to the survey in future newsletters.

Please acknowledge the use of CHPC Resources

If you use CHPC computer time or staff resources, we request that you acknowledge this in technical reports, publications, and dissertations. An example of what we ask you to include in your acknowledgement is:

"A grant of computer time from the Center for High Performance Computing is gratefully acknowledged."

If you make use of the CHPC Protected Environment, please also acknowledge the NIH shared instrumentation grant:

"The computational resources used were partially funded by the NIH Shared Instrumentation Grant 1S10OD021644-01A1."

Please submit copies or citations of dissertations, reports, pre-prints, and reprints in which CHPC is acknowledged by sending to *helpdesk@chpc.utah.edu*.

The University of Utah
University Information Technology
Center for High Performance Computing
155 South 1452 East, Room 405
SALT LAKE CITY, UT 84112–0190

Center for HIGH PERFORMANCE COMPUTING THE UNIVERSITY OF UTAH