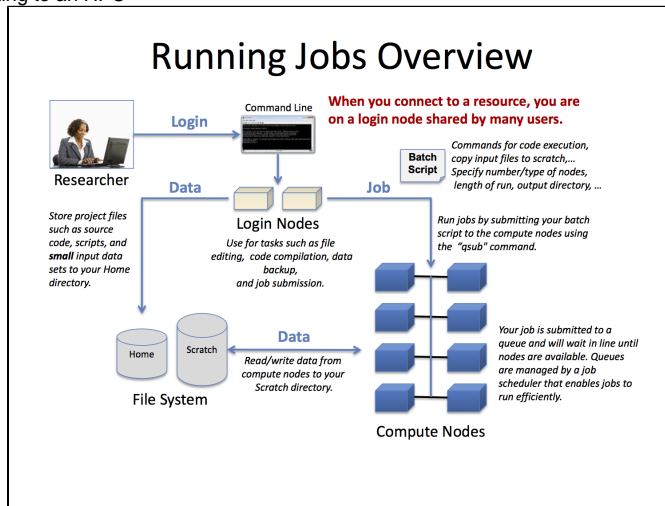
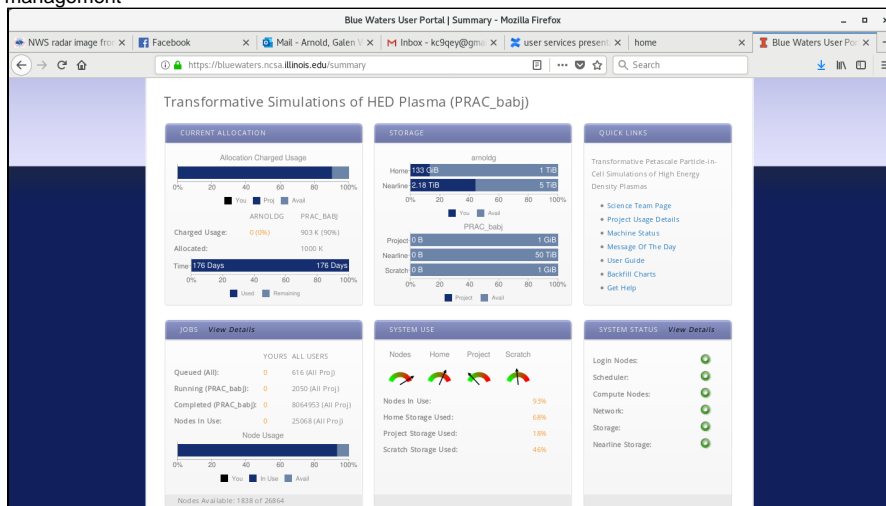


User Services Intro.

- Connecting to an HPC



- Account management



environment (modules)

5. Managing Your Environment (Modules)

The `module` command is a user interface to the Modules package. The Modules package provides for the dynamic modification of the user's environment via `modulefiles` (a `modulefile` contains the information needed to configure the shell for an application). Modules are independent of the user's shell, so both `tsh` and `bash` users can use the same commands to change the environment.

Useful Module commands:

Command	Description
<code>module avail</code>	lists all available modules
<code>module list</code>	lists currently loaded modules
<code>module help modulefile</code>	help on module <code>modulefile</code>
<code>module display modulefile</code>	Display information about <code>modulefile</code>
<code>module load modulefile</code>	load <code>modulefile</code> into current shell environment
<code>module unload modulefile</code>	remove <code>modulefile</code> from current shell environment
<code>module swap modulefile1 modulefile2</code>	unload <code>modulefile1</code> and load <code>modulefile2</code>

To include particular software in the environment for all new shells, edit your shell configuration file (`$HOME/.bashrc` for `bash` users and `$HOME/.cshrc` for `tsh` users) by adding the module commands to load the software that you want to be a part of your environment. After saving your changes, you can source your shell configuration file or log out and then log back in for the changes to take effect.

Note: Order is important. With each module load, the changes are prepended to your current environment paths.

- account usage, project participants, quotas
- Programming environment
 - traditional HPC languages (scale well, high efficiency): c and fortran

```
MPI Numerical Integration Reduction Example - Mozilla Firefox
https://www.dartmouth.edu/~rc/classes/intro_mpi/Numerical_Integration...
int myid, source, dest, tag;
MPI_Status status;
float my_result;

pi = acos(-1.0); /* = 3.14159... */
a = 0; /* lower limit of integration */
b = pi*1/2; /* upper limit of integration */
n = 100000; /* number of increment within each process */

dest = 0; /* define the process that computes the final result */
tag = 123; /* set the tag to identify this particular job */

/* Starts MPI processes ... */

MPI_Init(&argc, &argv); /* starts MPI */
MPI_Comm_rank(MPI_COMM_WORLD, &myid); /* get current process id */
MPI_Comm_size(MPI_COMM_WORLD, &np); /* get number of processes */

h = (b-a)/n; /* length of increment */
num = n/p; /* number of intervals calculated by each process */
my_range = (b-a)/p;
my_a = a + myid*my_range;
my_result = integral(my_a, num, h);

printf("Process %d has the partial result of %f\n", myid, my_result);

/* Use an MPI sum reduction to collect the results */
MPI_Reduce(&my_result, &result, 1, MPI_REAL, MPI_SUM, 0, MPI_COMM_WORLD);
MPI_Finalize(); /* let MPI finish up ... */

float integral(float a, int n, float h)
{
    int j;
    float h2, aij, integ;

    integ = 0.0; /* initialize integral */
    h2 = h/2;
    for (j=a; j<n; j++) { /* sum over all 'j' integrals */
        aij = a + j*h; /* lower limit of 'j' integral */
        integ += fct(aij+h2)*h;
    }
    return (integ);
}
```

```
Mozilla Firefox
https://computing.litl.gov/tutorials/mpisamples/fortran/mpl_latency.f
double precision i, i2, sum, deltaT
character msg
call MPI_INIT(ierr)
call MPI_COMM_RANK(MPI_COMM_WORLD, rank, ierr)
call MPI_COMM_SIZE(MPI_COMM_WORLD, ntasks, ierr)
if (rank.eq. 0) and (ntasks.ne. 2) then
    print *, 'Number of tasks = ', ntasks
    print *, 'Only need 2 tasks - extra will be ignored...'
endif
tag=1

if (rank.eq. 0) then
    print *, 'task', rank, 'has started...'
    print *, 'Beginning latency timing test. Number of reps= ', REPS
    print *, '-----'
    print *, 'Rep# T1 T2 deltaT'
    sumT = 0.0
    dest = 1
    source = 1
    do 10 n = 1, REPS
        Get start time
        T1 = MPI_WTIME()
        call MPI_SEND(msg, 1, MPI_CHARACTER, dest, tag,
            MPI_COMM_WORLD, ierr)
        call MPI_RECV(msg, 1, MPI_CHARACTER, source, tag,
            MPI_COMM_WORLD, status, ierr)
        Get ending time
        T2 = MPI_WTIME()
        deltaT = T2 - T1
        write(*,9) n, T1, T2, deltaT
        format(14, F22.8, F22.8, F12.8)
        sumT = sumT + deltaT
    10 continue
    avgT = (sumT * 1000000) / REPS
    Print final average from all round trips
    print *, '-----'
    print *, 'Avg round trip time=', avgT, 'microseconds'
    print *, 'Avg one way latency=', avgT/2, 'microseconds'
endif

if (rank.eq. 1) then
    print *, 'task', rank, 'has started ...'
```

python, machine-learning

What is PyTorch? — PyTorch Tutorials 1.1.0.dev20190501 documentation - Mozilla Firefox
https://pytorch.org/tutorials/beginner/blitz/tensor-tutorial.html

Reinforcement Learning

Reinforcement Learning (DQN) Tutorial

Extending PyTorch

Creating Extensions Using Numpy and Scipy

Custom C++ and CUDA Extensions

Extending TorchScript with Custom C++ Operators

Production Usage

Model Parallel Best Practices

Getting Started with Distributed Data Parallel

Writing Distributed Applications with PyTorch

PyTorch 1.0 Distributed Trainer with Amazon AWS

ONNX Live Tutorial

Loading a PyTorch Model in C++

PyTorch in Other Languages

Using the PyTorch C++ Frontend

Tutorials > Deep Learning with PyTorch: A 60 Minute Blitz > What is PyTorch?

All the tensors on the CPU except a `char` tensor support converting to Numpy and back.

CUDA Tensors

Tensors can be moved onto any device using the `.to` method.

```
# let us run this cell only if CUDA is available
# We will use 'torch.device' objects to move tensors in and out of GPU
if torch.cuda.is_available():
    device = torch.device("cuda") # a CUDA device object
    y = torch.ones_like(x, device=device) # directly create a tensor on GPU
    x = x.to(device) # or just use strings 'to('cuda')'
    z = x + y
    print(z)
    print(z.to("cpu", torch.double)) # 'to' can also change dtype together!
```

Out:

```
tensor([-0.4743], device='cuda:0')
tensor([-0.4743], dtype=torch.float64)
```

Total running time of the script: (0 minutes 5.866 seconds)

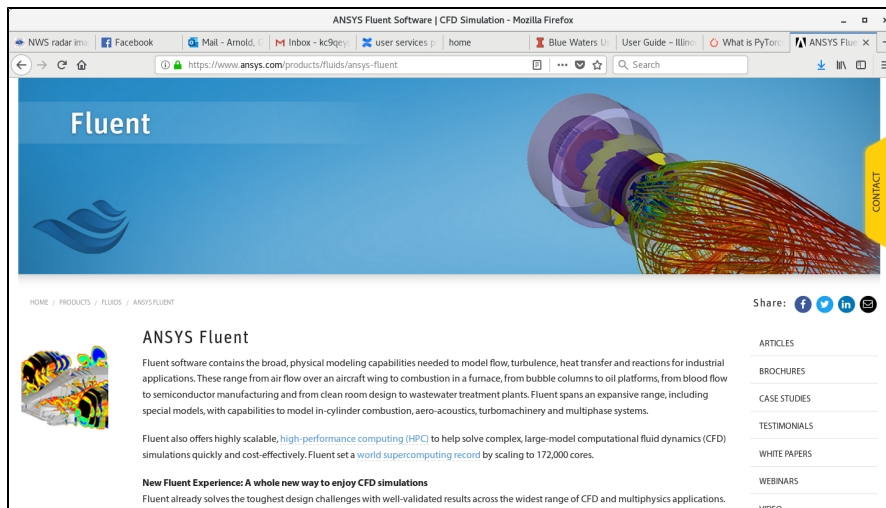
What is PyTorch?

+ Getting Started

+ Numpy Bridge

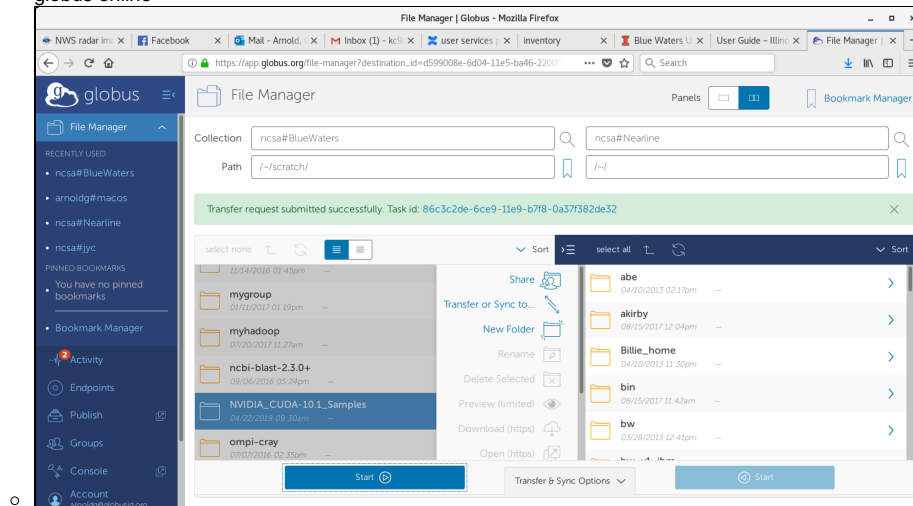
CUDA Tensors

commercial software

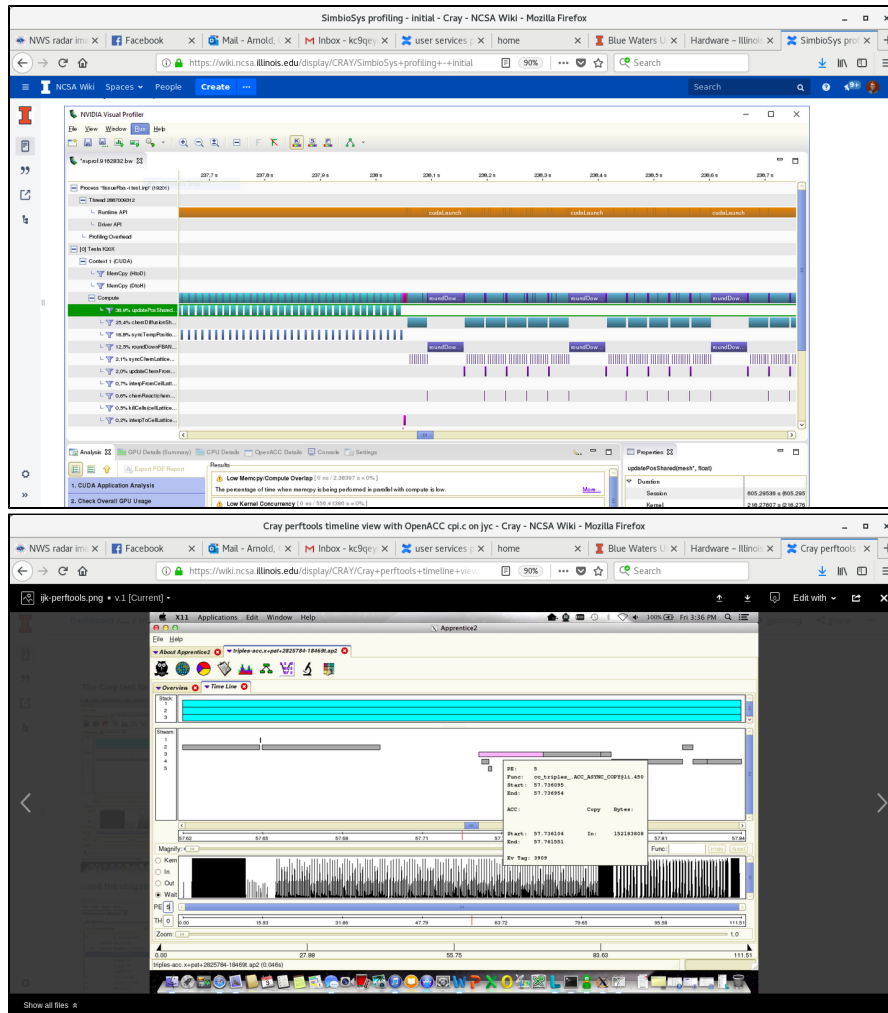


- Data Transfer with an HPC


- globus online



- ssh/scp/rdist
- Thinking about your workflow end to end
 - optimize performance
 - on the HPC
 - performance profiling



■ interacting with the HPC



- data transfer bottlenecks (last mile)
 - WiFi

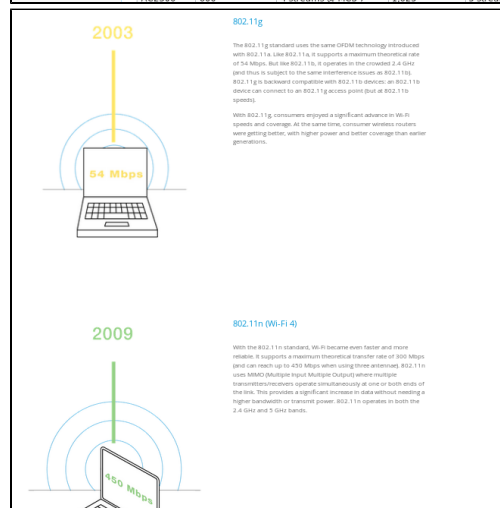
IEEE 802.11ac - Wikipedia - Mozilla Firefox

160 MHz channels, and thus the throughput might be unusable in some countries/regions due to regulatory issues that allocated some frequencies for other purposes.

Advertised Speeds [\[edit\]](#)

802.11ac-class device wireless speeds are often advertised as AC followed by a number, that number being the highest link rates in Mbits of all the simultaneously-usable radios in the device added up. For example, an AC1900 access point might have 600 Mbps capability on its 2.4 GHz radio and 1300 Mbps capability on its 5 GHz radio. No single client device could connect and achieve 1900 Mbps of throughput, but separate devices each connecting to the 2.4 GHz and 5 GHz radios could achieve combined throughput approaching 1900 Mbps. Different possible stream configurations can add up to the same AC number.

Type	2.4 GHz band ^[c] Mbit/s	2.4 GHz band config [all 40 MHz]	5 GHz band Mbit/s	5 GHz band config [all 80 MHz]
AC450 ^[16]	-	-	433	1 stream @ MCS 9
AC600	150	1 stream @ MCS 7	433	1 stream @ MCS 9
AC750	300	2 streams @ MCS 7	433	1 stream @ MCS 9
AC1000	300	2 streams @ MCS 7	650	2 streams @ MCS 7
AC1200	300	2 streams @ MCS 7	867	2 streams @ MCS 9
AC1300	400	2 streams @ 256-QAM	867	2 streams @ MCS 9
AC1300 ^[17]	-	-	1,300	3 streams @ MCS 9
AC1350 ^[18]	450	3 streams @ MCS 7	867	2 streams @ MCS 9
AC1450	450	3 streams @ MCS 7	975	3 streams @ MCS 7
AC1600	300	2 streams @ MCS 7	1,300	3 streams @ MCS 9
AC1700	800	4 streams @ 256-QAM	867	2 streams @ MCS 9
AC1750	450	3 streams @ MCS 7	1,300	3 streams @ MCS 9
AC1900	600 ^[d]	3 streams @ 256-QAM	1,300	3 streams @ MCS 9
AC2100	800	4 streams @ 256-QAM	1,300	3 streams @ MCS 9
AC2200	450	3 streams @ MCS 7	1,733	4 streams @ MCS 9
AC2300	600	4 streams @ MCS 7	1,625	5 streams @ MCS 7



USB-N with N < 3 ?

Understanding USB: Speed, Connectors and Functionality | Tripp Lite - Mozilla Firefox

1. USB Speed Standards

Your USB connections are about to become faster! Products compatible with the new USB 3.1 Gen 2 standard will start to hit the market soon. This latest generation promises to take connectivity to the next level, with data transfer rates twice as fast as USB 3.0 (also known as USB 3.1 Gen 1):

USB Standard	Data Transfer Speed	Also Known As
USB 1.1	12 Mbps	Full Speed
USB 2.0	480 Mbps	Hi Speed
USB 3.0	5 Gbps	USB 3.0 or SuperSpeed
USB 3.1 Gen 1	5 Gbps	USB 3.1 Gen 1 or SuperSpeed
USB 3.1 Gen 2	10 Gbps	SuperSpeed+ or SuperSpeed 10

2. USB Connector Standards

USB Type-A

The standard, universal connector found on virtually every desktop PC and laptop in use to day, as well as TVs, game consoles and media players. Although USB 3.0 Type-A connectors have more internal pins, the form factor is the same, so it can operate in any Type-A port, even USB 1.1. Data transfer becomes much faster at the speed of the data connection.

- on a campus lan
 - gigabit ethernet (show bandwidth and latency from a test)
- on a "foreign" lan (home, coffee shop, another continent...)
 - vpn ?

PC REVIEWS BEST PICKS HOW-TO NEWS SMART HOME BUSINESS SHOP

The chart below breaks down the results from the ten fastest VPN services we tested. This information was based on testing performed over the course of the month of December, 2018. Note that the chart at the top of this piece is also sorted by order of speed.

	Download Speed Percent Change	Upload Speed Percent Change	Latency Percent Change
	(Lower Is Better)	(Lower Is Better)	(Lower Is Better)
HideIPVPN	52.6%	53.0%	300.0%
TunnelBear VPN	74.7%	61.0%	100.0%
Hide.me VPN	75.8%	68.4%	0.0%
Trust.zone	77.7%	58.5%	325.0%
Private Internet Access	80.7%	76.3%	0.0%
IPVanish VPN	81.0%	78.8%	0.0%
TorGuard VPN	81.8%	75.5%	66.7%
NordVPN	82.6%	77.7%	0.0%
Buffered VPN	83.3%	78.6%	175.0%
IVPN	85.8%	74.4%	66.7%

- performance vs security
- debugging codes and issues that arise
 - your code ?
 - system problem ?
 - [Debugging on Blue Waters](#)

References:

<https://bluewaters.ncsa.illinois.edu/user-guide>

<https://campuscluster.illinois.edu/resources/docs/user-guide/>