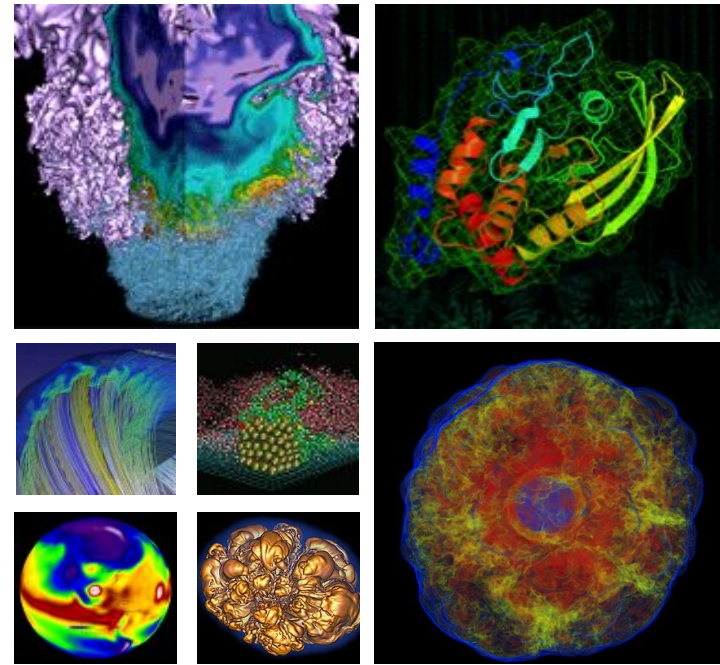


Jupyter at NERSC

Redefining the Interface to HPC



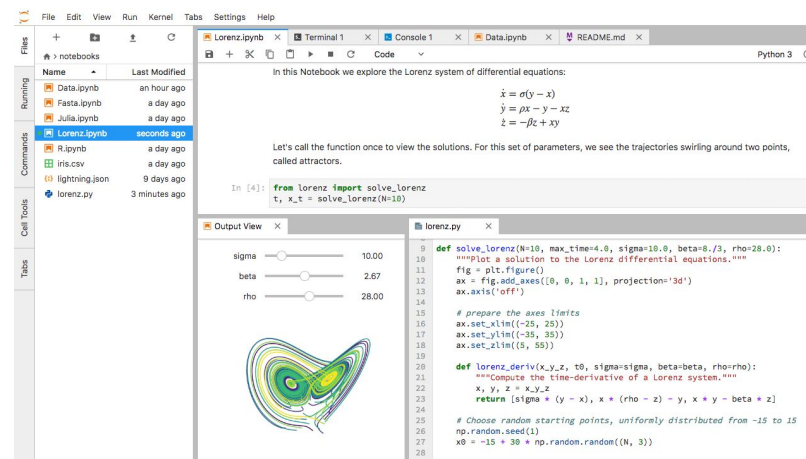
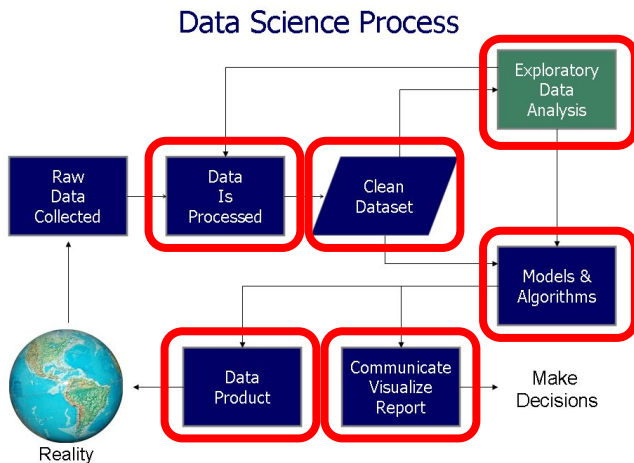
Rollin Thomas
Data and Analytics Services

NERSC User Group Meeting
Rockville MD • 2019-07-19

What is Jupyter?



Tool for reproducible, shareable narratives, literate computing:
Notebook: Document containing code, comments, outputs.
Rich text, interactive plots, equations, widgets, etc.



Goal: Enable exploratory data analytics, deep learning, workflows, and more through Jupyter on NERSC systems.

Why Jupyter, Why Now at NERSC?



Integral part of Big (Data) Science & Superfacility:

**LSST-DESC, DESI, ALS, LCLS, Materials Project
NCEM, LUX/LZ, KBase...**

Generational shift in analytics for science + more:

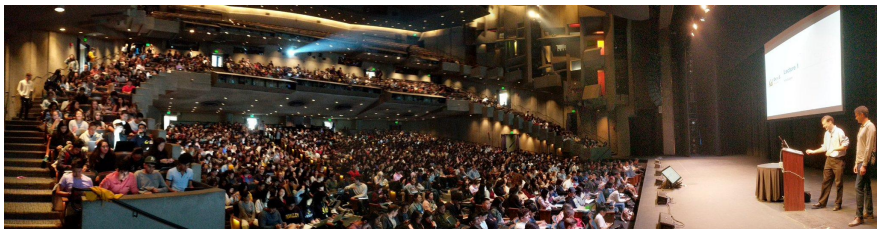
**UCB's Data Science 8 course, entirely in Jupyter
"I'll send you a copy of my notebook"**

Training events adopting notebooks (DL)

Supporting reproducibility and science outreach:

Open source code and open source science

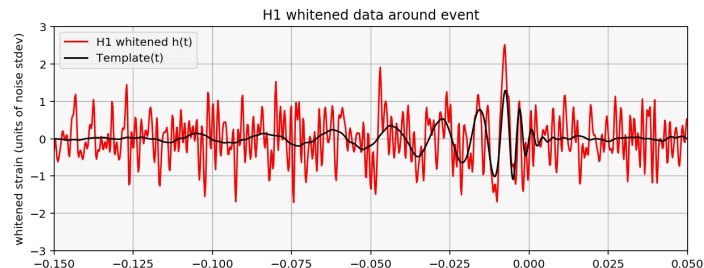
Jupyter notebooks alongside publications (LIGO)



Data 8: Foundations of Data Science, Fall 2018, Zellerbach Hall

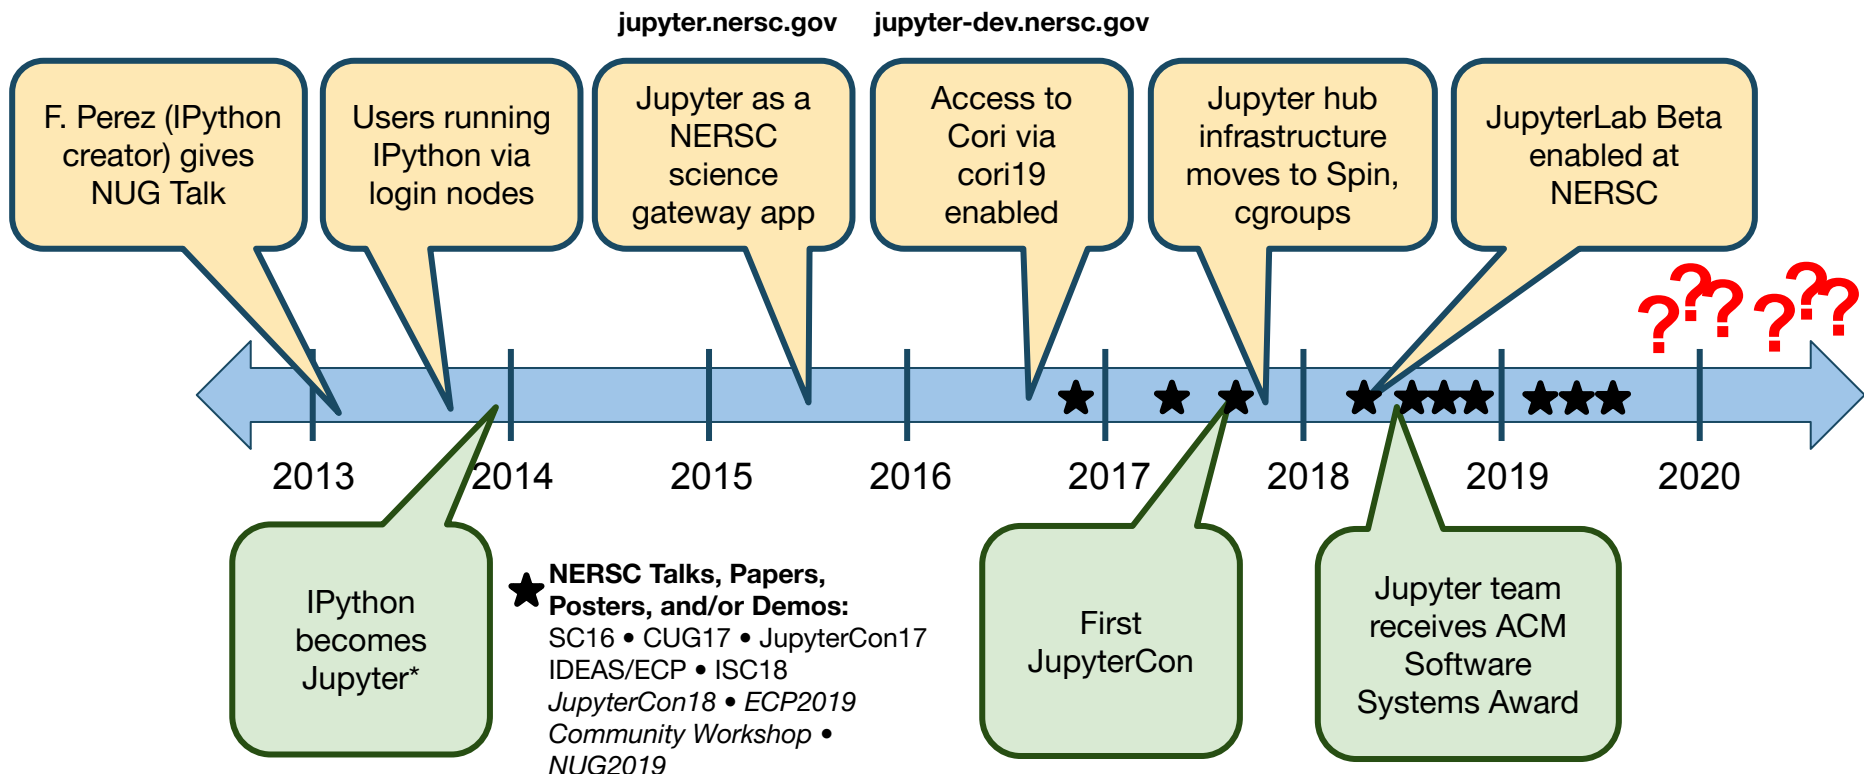


2017 ACM Software System Award: "... *a de facto standard for data analysis in research, education, journalism and industry.* Jupyter has broad impact across domains and use cases. Today more than **2,000,000 Jupyter notebooks are on GitHub**, each a distinct instance of a Jupyter application—covering a range of uses from technical documentation to course materials, books and academic publications."



LIGO Binary BH-BH Merger GW Signature
Figure from LIGO EPO/Publication Jupyter Notebook

Jupyter at NERSC Timeline

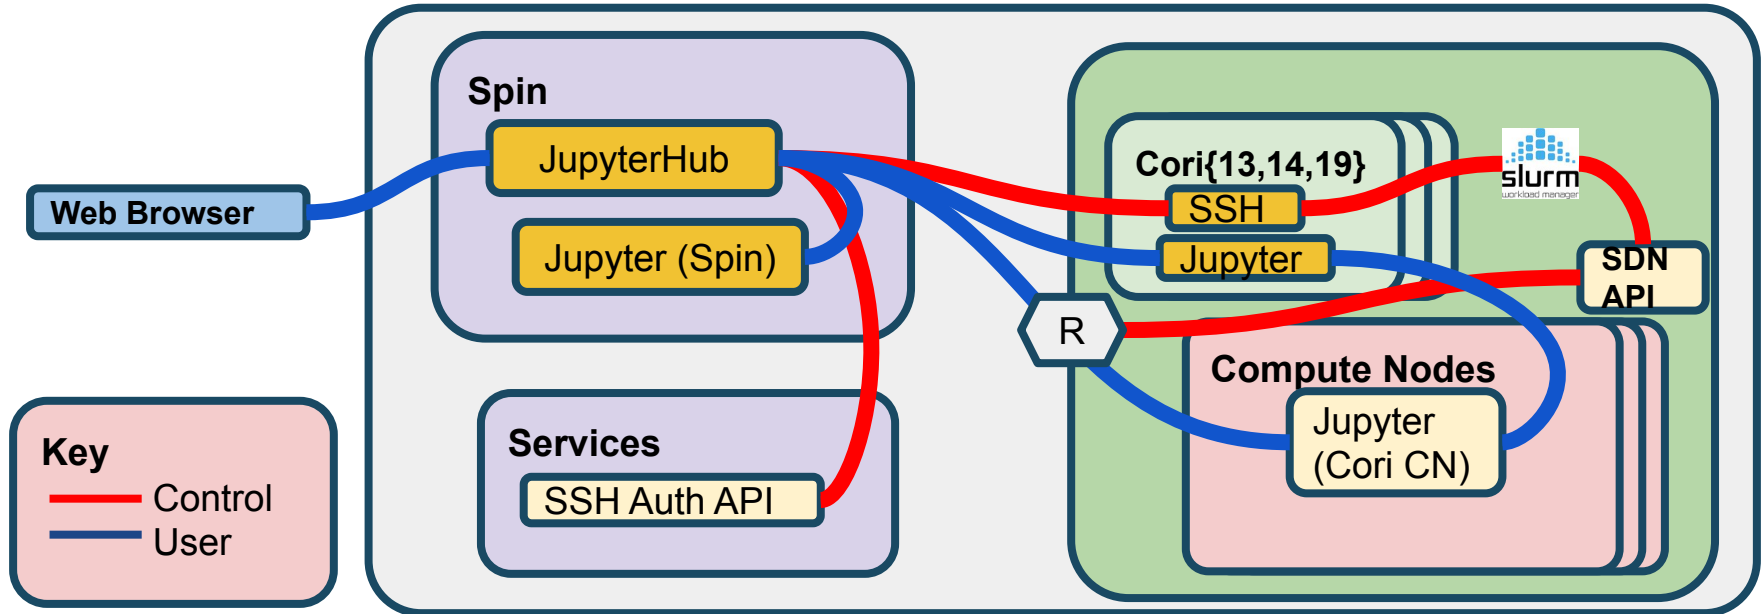


Use Cases & Access Modes @ NERSC



Use Case	Where	Why
Light-weight data analysis and visualization	Spin Container <i>(In production now.)</i>	Usable when other systems are down. Simple, interactive access
Workflow execution and medium-scale data analysis	Cori "Login" Nodes <i>(In production now.)</i>	Access to batch and scratch Larger memory shared node
Heavy weight computation including task frameworks	Cori Compute Nodes <i>(In testing now.)</i>	Dedicated resources (e.g. memory and cores). Ability to launch parallel workloads in the notebook.

Jupyter @ NERSC Architecture



Jupyter Matters to NERSC Users



Users appreciate Jupyter @ NERSC...

“I really like the jupyter interface.”

“New jupyter notebooks are awesome!”

“Great interactive workflow (e.g. for postprocessing) via JupyterHub”

“As mentioned, the ability to access data from the scratch directories through the Jupyter hub is very important to my workflow. The Jupyter hub has been running more and more consistently, but it still seems to lag or stall sometimes. I guess **my only thought on how to improve (currently)** would be to improve the stability of the Jupyter hub.”

[Venkitesh: “... jupyter notebooks are very important for me: **The 3 most important things in life: food, shelter and jupyter... everything else is optional.**”]

“I absolutely love the fact that I can use the Jupyter hub to access the Cori scratch directory. This allows me to analyze data through the browser ... or to quickly check that simulation runs are going as expected without having to transfer data to a different location. **I actually also have access to other supercomputer clusters, but this is one of the biggest reasons I mainly use Cori and Edison for debugging and production runs.**”

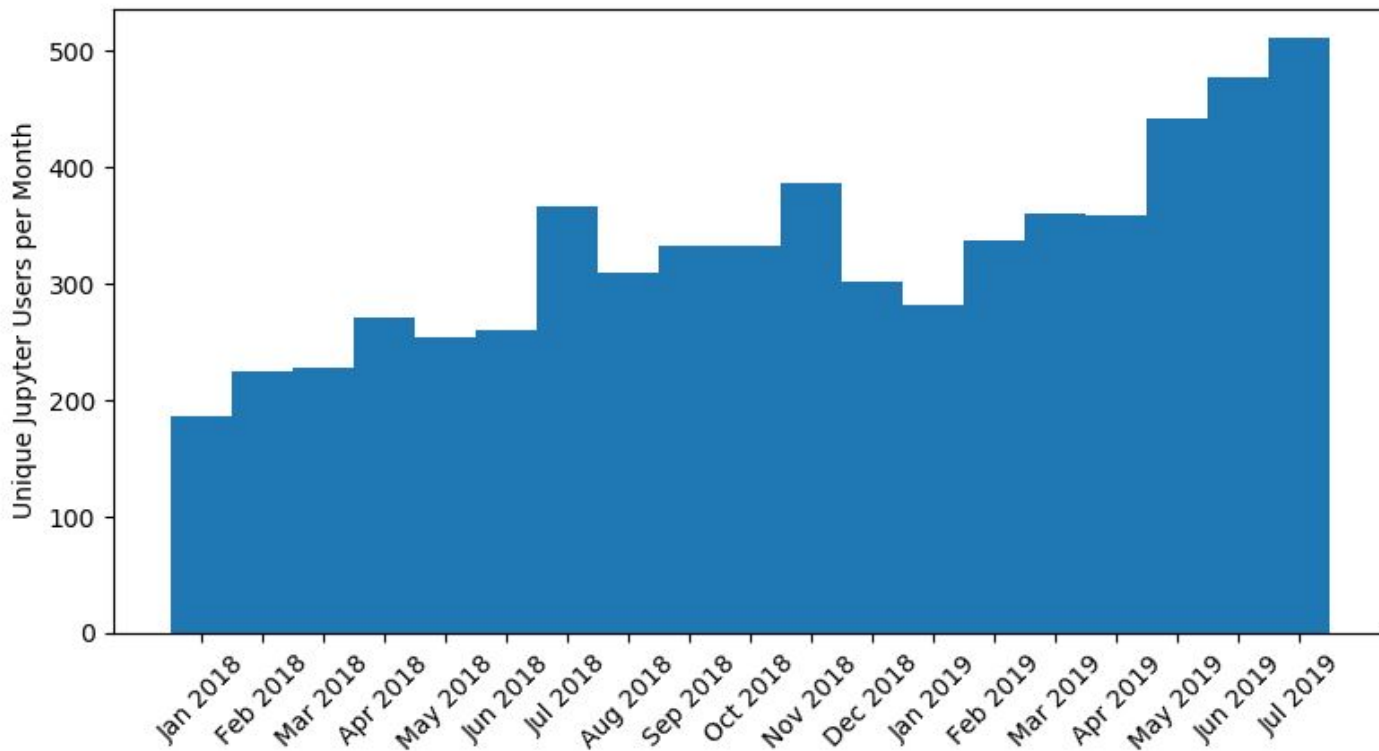
...but need increased stability and to scale up.

“I would really appreciate it if jupyter.nersc.gov wouldn't go down as much as it does.”

(2017 User Survey)

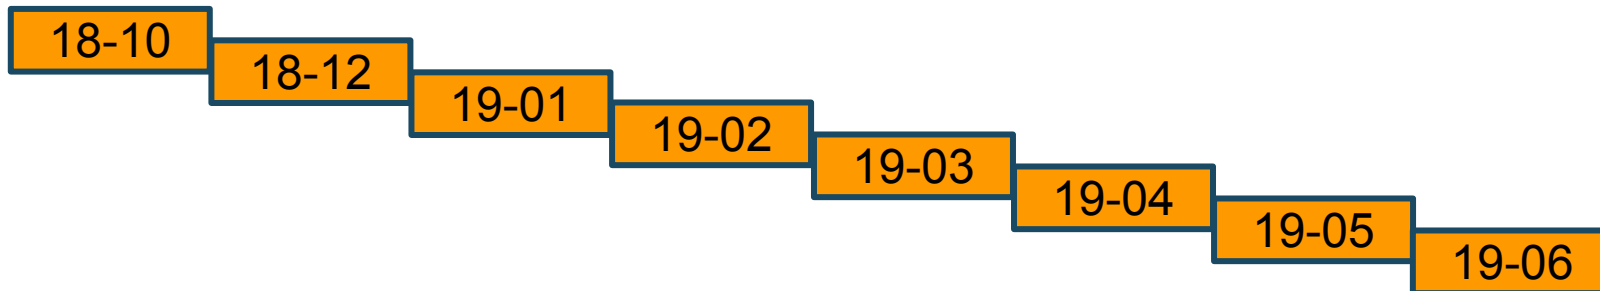
“MPI cannot be used in jupyter notebook as well, where the jupyter hubs run on login nodes (unless when using the compute nodes through SLURM.)”

Jupyter on Cori Usage Numbers



Supporting these users, their current and future needs is a lot of work!

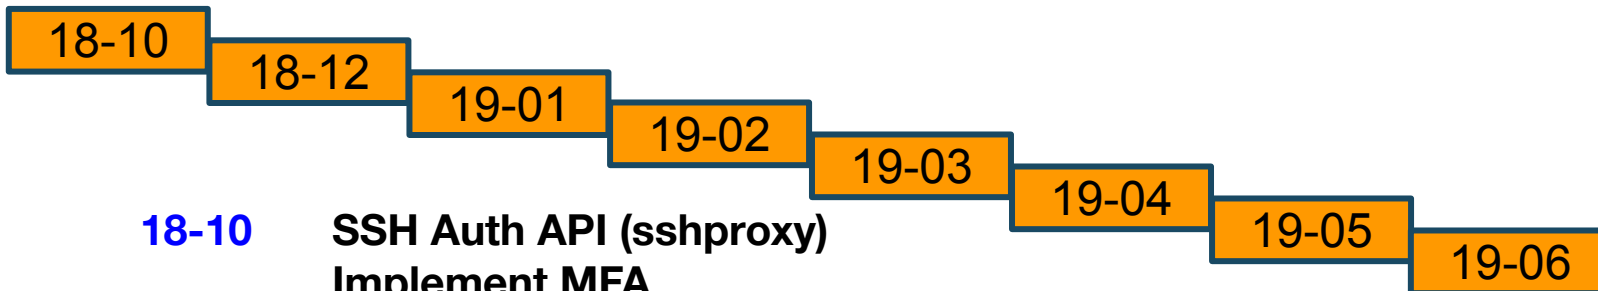
Deployment Schedule



Jupyter at NERSC now has a ~monthly update schedule
Usually follows Cori maintenances by 2-5 days
We post an announcements when it's coming up
Jupyter upgrades themselves are very short (usually...)
Rarely change notebook env, mostly hub changes:



Deployment Schedule



18-10 SSH Auth API (sshproxy)
Implement MFA

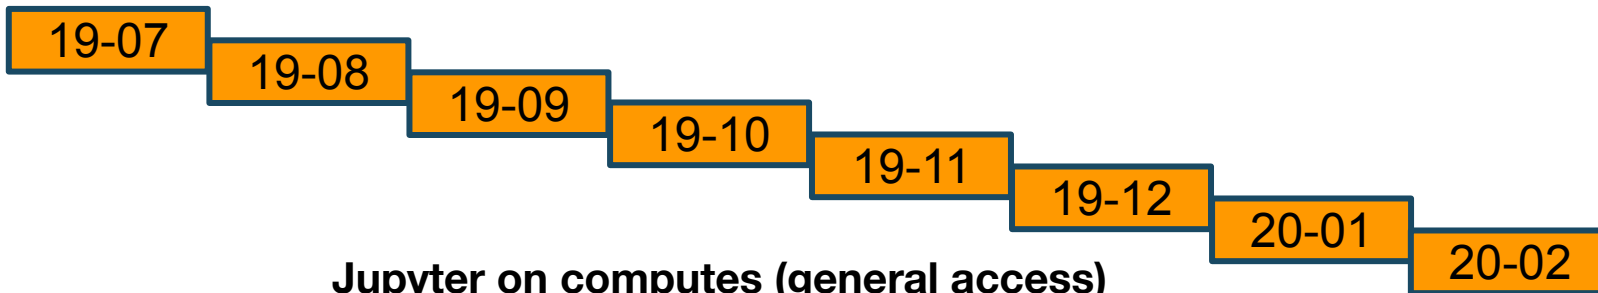
18-12 Add “offline” communication page
“Announcements” service for communication

19-02 Ability to leverage more Cori nodes

19-03 Converged hubs: *jupyter-dev* and *jupyter*

19-05 Computes and GPU node access (special access)
Leverage Iris for role-based access control
JupyterHub 1.0

Deployment Schedule



Jupyter on computes (general access)

Configurable jobs:

Reservation (for training/tutorials/workshops)

Full shifter image support

Integration with NERSC SSO

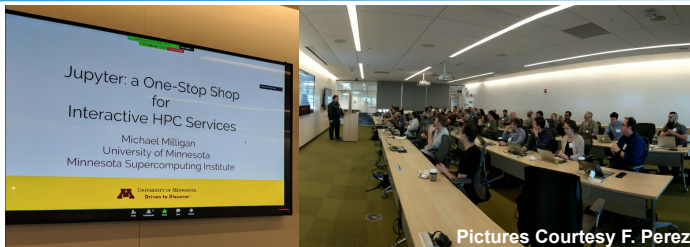
Jupyter Slurm integration

Parameterized notebooks

Managing curated notebooks

...

Workshop: Jupyter for Science User Facilities and High Performance Computing



Pictures Courtesy F. Perez

Joint Workshop w/BIDS: June 11-12 at NERSC, June 13 at BIDS

Committee: Rollin Thomas • Shane Canon • Shreyas Cholia • Kelly Rowland
Debbie Bard • Dan Allan (BNL) • Chris Holdgraf (BIDS)

Part of “Jupyter Community Workshop” Series

Competitive application process • Granted up to \$20K for travel support
Funds from Bloomberg, managed by NumFOCUS and Project Jupyter

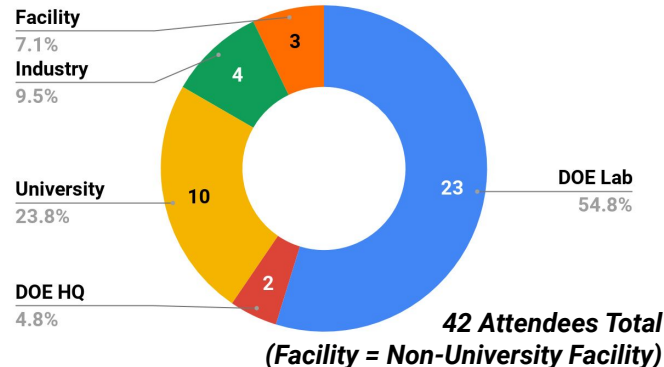
User Facilities, HPC & Data Centers Represented

NSLS-II • LSST • APS • SLAC • JGI • ARM • European XFEL
NERSC • ALCF • TACC • MSI (@UMN) • Compute Canada • ESA

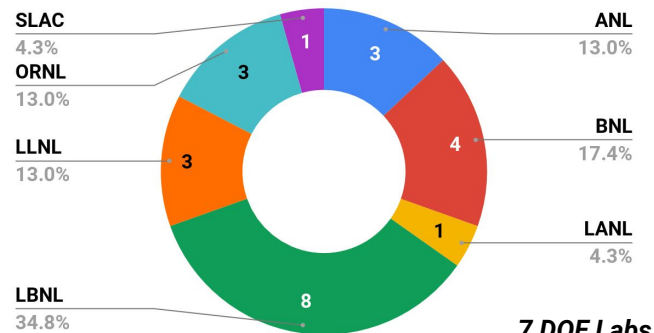
Content

Talks: Deployment • Infrastructure • Extending Jupyter for HPC • Use Cases
Breakouts: Organizing Collaboration • Securing Jupyter • Sharing Notebooks
Reproducibility • Best Practices • Future Plans • Tutorials
Roundtable Meeting with Core Jupyter Developers

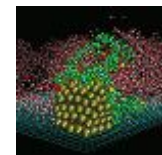
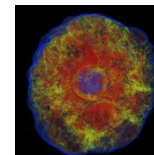
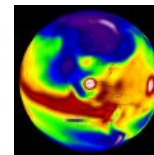
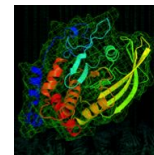
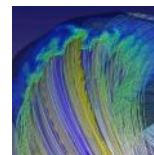
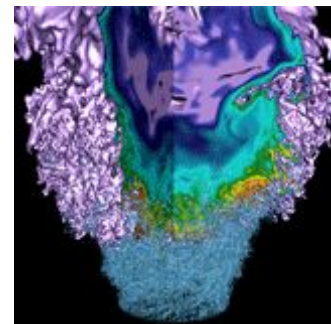
Attendance by Category



DOE Lab Representation



Demo Intro



NCEM Superfacility Highlight Slide

Slide Courtesy Matt Henderson and Shreyas Cholia (LBL CRD)



Scientific Achievement

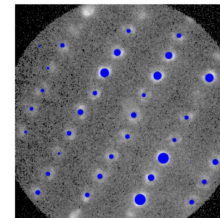
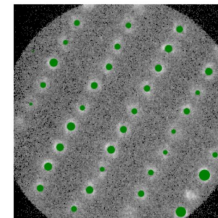
Enable interactive human-in the loop computing for Superfacility Workflows at NERSC

Significance and Impact

Superfacility Projects like ALS and NCEM need a flexible interactive user interface to be able to analyze data from experiments in real time. We use Jupyter Notebooks to enable these projects to combine live code execution, reproducibility, interactive visualizations on HPC resources.

Research Details

- **Command and control center** for integrating distributed workflows under a common platform
- Enable HPC Tools to **allow Jupyter notebooks to execute code in parallel on distributed backend** resources using Dask/IPyParallel/MPI
- Realtime rendering of results **rendered inline in Jupyter to enable live interaction, visualization and job steering**
- Parameterized notebooks can to execute a workflow over multiple datasets and parameters.
- Enable reproducible workflows through project curated notebooks that can be cloned / executed



Find Bragg disks in all diffraction patterns

```
In [5]: ##### Parameters #####
corrPower = 0.8
sigma = 2
edgeBoundary = 20
maxNumPeaks = 70
minPeakSpacing = 50
minRelativeIntensity = 0.001
subpixel = True
verbose = True

#####

bragspeaks = find_Bragg_disks(dc, probe kernel,
                             corrPower=corrPower,
                             sigma=sigma,
                             edgeBoundary=edgeBoundary,
                             minRelativeIntensity=minRelativeIntensity,
                             minPeakSpacing=minPeakSpacing,
                             maxNumPeaks=maxNumPeaks,
                             subpixel=subpixel,
                             verbose=verbose)

bragspeaks.name = 'bragspeaks'
Analyzing: |*****| 100.0% Complete
Analyzed 3600 diffraction patterns in 0h 3m 35s

In [6]: # Show results for a few DPS
xs=(10,31,18)
ys=(45,31,10)
power=0.3
```

Achieved 90x speedup on NCEM Py4DSTEM Notebook running on HPC resources at NERSC:
<https://github.com/py4dstem/py4DSTEM/blob>

Team



Rollin Thomas



Shane Canon



Kelly Rowland



Trevor Slaton



Shreyas Cholia



Matt Henderson

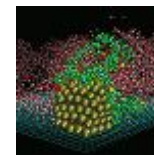
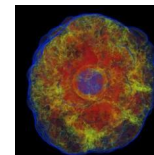
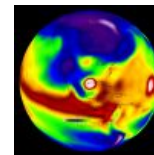
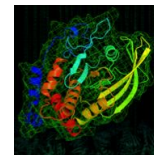
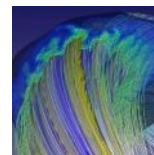
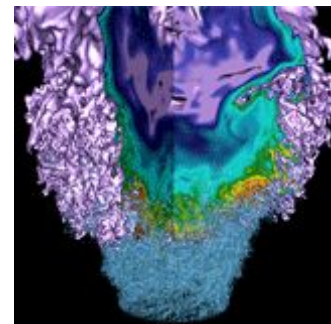


William Krinsman



Jon Hays

Demo



- **Streamline notebook ↔ running job interaction**
- **Exclusive node access (notebook on compute) in real-time-like queue**
- **Experiment with configurable-http-proxy for dynamic routing (not SDN)**
- **Support >300 users/day, maybe add 1-2 more shared nodes, manage via Slurm**
- **Deploy Jupyter solutions at NERSC:**
 - **Sharing notebooks, parameterized notebooks, curated notebooks**
 - **Voila and dashboards**
 - **More HPC-centric extensions and fixes to JupyterLab/Hub.**
- **Engagement with Jupyter and Jupyter-in-HPC communities**
- **User engagements through Superfacility and beyond:
NCEM, ALS, DESI, LSST, ...**



Thank You



U.S. DEPARTMENT OF
ENERGY

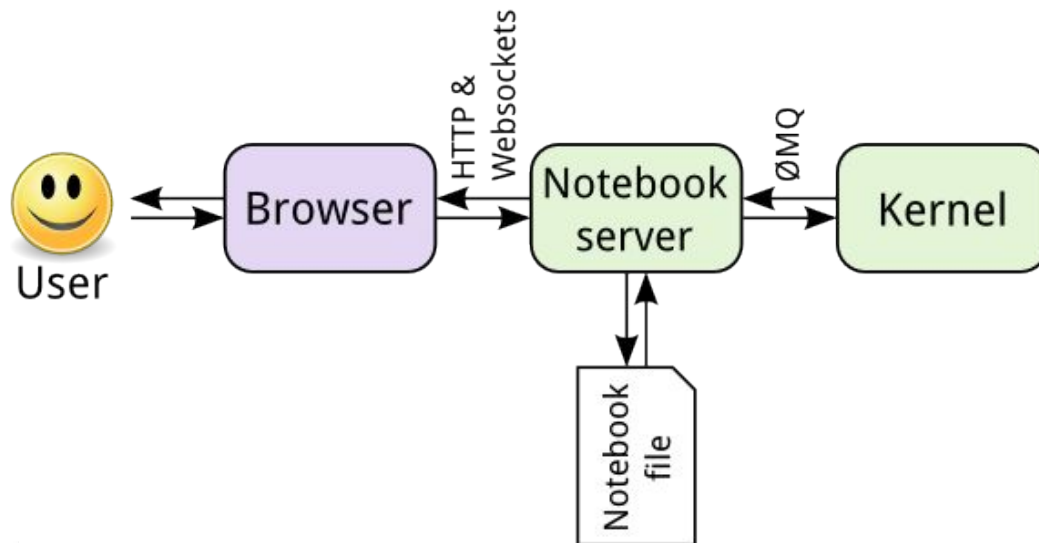
Office of
Science



How It Works



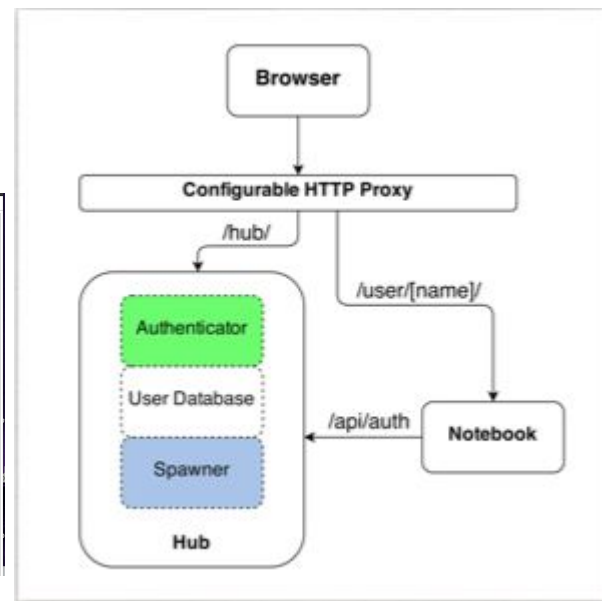
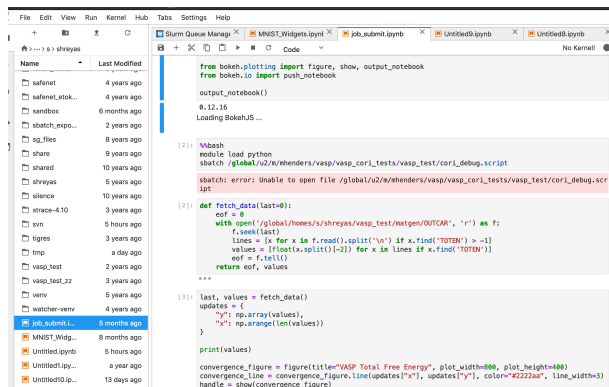
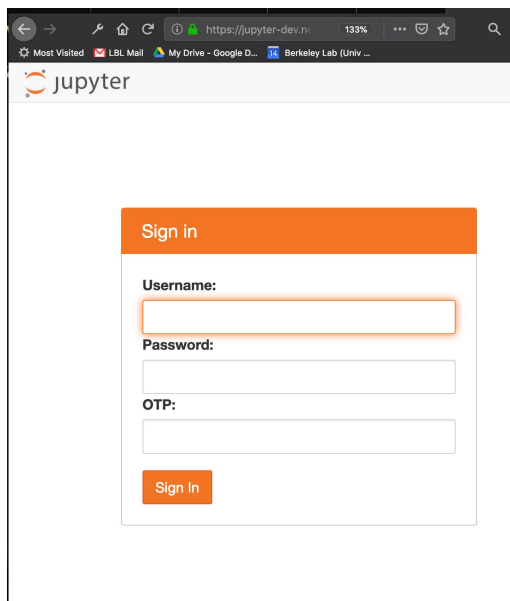
The **Notebook Server** sends code (via **ZeroMQ**) to a language “**kernel**” that executes this code. In addition to running your code, it stores code and output, together with markdown notes, in an editable document called a **notebook**, saved as a **JSON** file with a **.ipynb extension**.



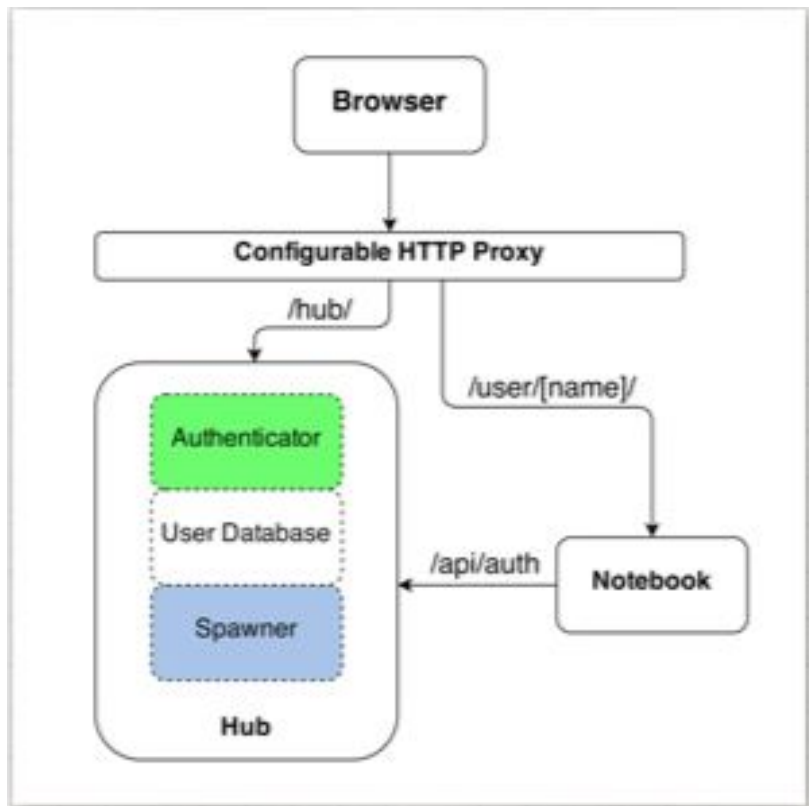
JupyterHub: Jupyter as a Service



- Service to deploy notebooks in a multi-user environment
- Manages user authentication, notebook deployment and web proxies



JupyterHub Architecture



Components are abstracted:

- Authenticator
- Spawner
- Proxy

Pieces we've created:

- GSIAuthenticator (IT IS NO MORE)
- SSHAPIAuthenticator

- SSHSpawner (Had gsissh support)
- NERSCSpawner
- NERSCSlurmSpawner

Pieces we re-use and love:

- WrapSpawner (NERSCSpawner)
- BatchSpawner (NERSCSlurmSpawner)

JupyterLab: Notebooks++

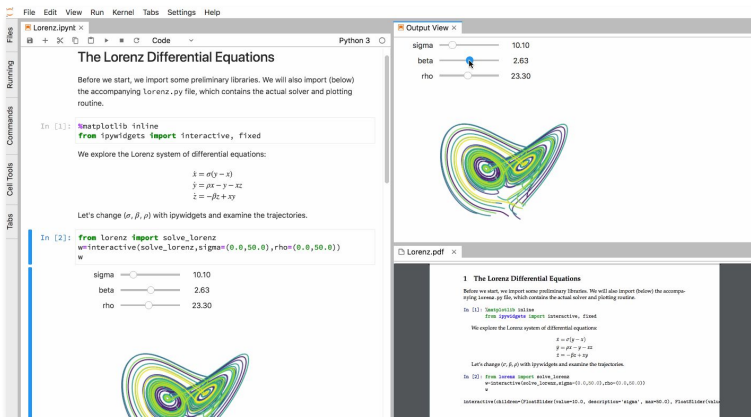
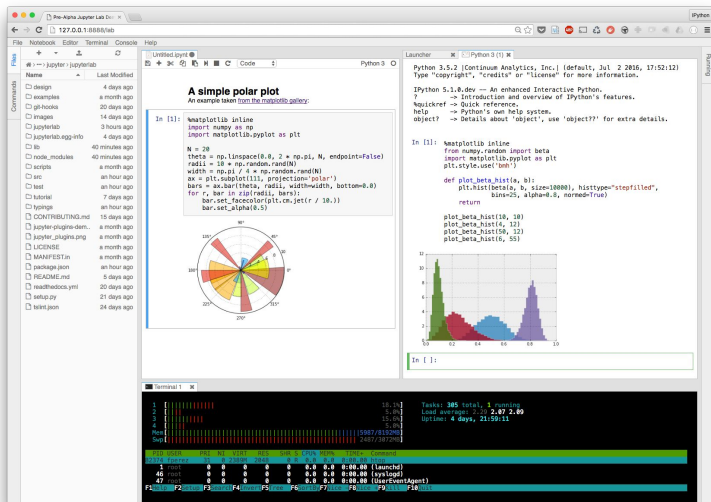


Modern Frontend for Jupyter

- Integrated GUI for Jupyter ecosystem (filebrowser, tabbed panes, notebooks, terminal etc.)
- Common framework to integrate multiple applications - “extensions”
- eg. connect multiple viewers to common underlying kernel

*Extensions for many users?
Right now we manage them all.*

... we've made a Slurm one I'll show later.



Handling Spawner Options



Account: dasrepo

Nodes: 3

Time Limit: 240

QOS: regular

Constraint: haswell

Reservation: DisCoHackathon2

Shifter Image: registry.services.nersc.gov/rthomas/shifter-jupyterlab:latest

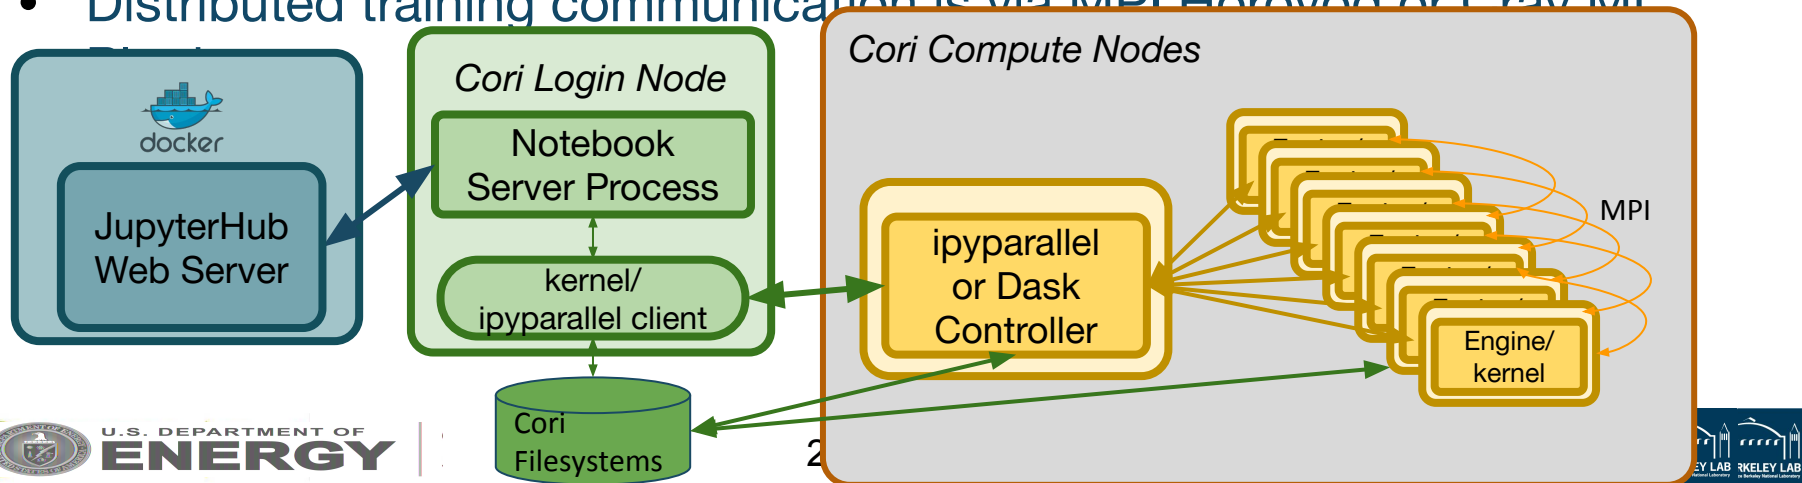
Spawn

Populated by queries against internal REST API's at NERSC (or example ones)

Jupyter architecture



- Allocate nodes on Cori interactive queue and start ipyparallel or Dask cluster
 - Developed %ipcluster magic to setup within notebook
- Compute nodes traditionally do not have external address
 - Required network configuration / policy decisions
- Distributed training communication is via MPI Horovod or Cray ML



Setting up ipyparallel cluster



Via Magic (entire workflow in notebook) or a console script

```
In [1]: import ipcluster_magics
```

```
In [2]: job_name = "isc_ihpc_mnist"
nodes = 1
engines = 1
module = "python/3.6-anaconda-4.4"
conda_env = "/global/cscratch1/sd/sfarrell/conda/isc-ihpc"
```

```
In [3]: %ipcluster -m $module -e $conda_env -N $nodes -J $job_name -t 01:00:00
```

```
salloc: Pending job allocation 13289619
salloc: job 13289619 queued and waiting for resources
salloc: job 13289619 has been allocated resources
salloc: Granted job allocation 13289619
2018-06-21 15:55:55.813 [scheduler] Scheduler started [leastload]
```

```
salloc --qos=interactive -N 1 -C haswell
wbhimji@nid00032:~> ./startCluster.sh

# Use a unique cluster ID for this job
clusterID=cori_${SLURM_JOB_ID}
echo "Launching controller"
ipcontroller --ip="$headIP" \
             --cluster-id=$clusterID &
sleep 20
echo "Launching engines"
srun ipengine --cluster-id=$clusterID
```

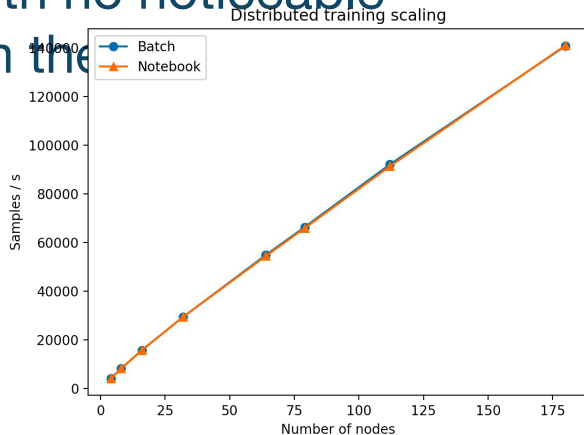
Connect to cluster in notebook

```
In [7]: # Cluster ID taken from job ID above
job_id = 13272466
cluster_id = 'cori_{}'.format(job_id)

# Use default profile
c = ipp.Client(timeout=60, cluster_id=cluster_id)
```

Distributed Training

- Distributed training in notebooks with IPyParallel and Horovod-MPI
- Notebook cells specified for parallel execution using cell magic
 - MPI code in a notebook
- Scales well with no noticeable overhead from the infrastructure



Build and train the model

In [8]: `%%px` → Parallel notebook cell

```
# Model config
h1, h2, h3, h4, h5 = 64, 128, 256, 256, 512
optimizer = 'Adam'
lr = 0.001 * hvd.size()

# Training config
batch_size = 128
n_epochs = 4

# Build the model
model = build_model(train_input.shape[1:],
                    h1=h1, h2=h2, h3=h3, h4=h4, h5=h5,
                    optimizer=optimizer, lr=lr,
                    use_horovod=True)

if hvd.rank() == 0:
    model.summary()
```

Construct model on every worker

[stdout:1]

Layer (type)	Output Shape	Param #
input_2 (InputLayer)	(None, 64, 64, 1)	0
conv2d_5 (Conv2D)	(None, 64, 64, 64)	640
conv2d_6 (Conv2D)	(None, 32, 32, 128)	73856
conv2d_7 (Conv2D)	(None, 32, 32, 256)	295168
conv2d_8 (Conv2D)	(None, 16, 16, 256)	590080
flatten_2 (Flatten)	(None, 65536)	0
dense_3 (Dense)	(None, 512)	33554944
dense_4 (Dense)	(None, 1)	513
Total params: 34,515,201		
Trainable params: 34,515,201		
Non-trainable params: 0		

`%%px`

Train with Horovod on all workers

```
# Train the model
history = train_model(model, train_input=train_input, train_labels=train_labels,
                      valid_input=valid_input, valid_labels=valid_labels,
                      batch_size=batch_size, n_epochs=n_epochs,
                      use_horovod=True)
```

[stdout:0]

Train on 64000 samples, validate on 32000 samples

Distributed HPO - Setup



Easy but powerful setup for random search HPO

- Define HP sets to evaluate
- Define model training function
- Run the HPO tasks with load-balanced scheduler

Launch user-defined training function and arguments

AsyncResult objects can be queried for status, outputs

```
# Define the hyper-parameter search points
n_hpo_trials = 336
h1 = np.random.choice([4, 8, 16, 32, 64], size=n_hpo_trials)
h2 = np.random.choice([4, 8, 16, 32, 64], size=n_hpo_trials)
h3 = np.random.choice([8, 16, 32, 64, 128], size=n_hpo_trials)
conv_sizes = np.stack([h1, h2, h3], axis=1)
fc_sizes = np.random.choice([32, 64, 128, 256], size=(n_hpo_trials, 1))
lr = np.random.choice([0.0001, 0.001, 0.01], size=n_hpo_trials)
dropout = np.random.rand(n_hpo_trials)
optimizer = np.random.choice(['Adadelata', 'Adam', 'Nadam'], size=n_hpo_trials)
```

```
# Load-balanced view
lv = c.load_balanced_view()

# Loop over hyper-parameter sets
results = []
for ihp in range(n_hpo_trials):
    print('Hyperparameter trial %i conv %s fc %s dropout %.4f opt %s, lr %.4f' %
          (ihp, conv_sizes[ihp], fc_sizes[ihp], dropout[ihp], optimizer[ihp], lr[ihp]))
    checkpoint_file = os.path.join(checkpoint_dir, 'model %i.h5' % ihp)
    result = lv.apply(build_and_train,
                     input_dir, n_train, n_valid,
                     conv_sizes=conv_sizes[ihp], fc_sizes=fc_sizes[ihp],
                     dropout=dropout[ihp], optimizer=optimizer[ihp], lr=lr[ihp],
                     batch_size=batch_size, n_epochs=n_epochs,
                     checkpoint_file=checkpoint_file)
    results.append(result)
```

Load-balanced scheduling

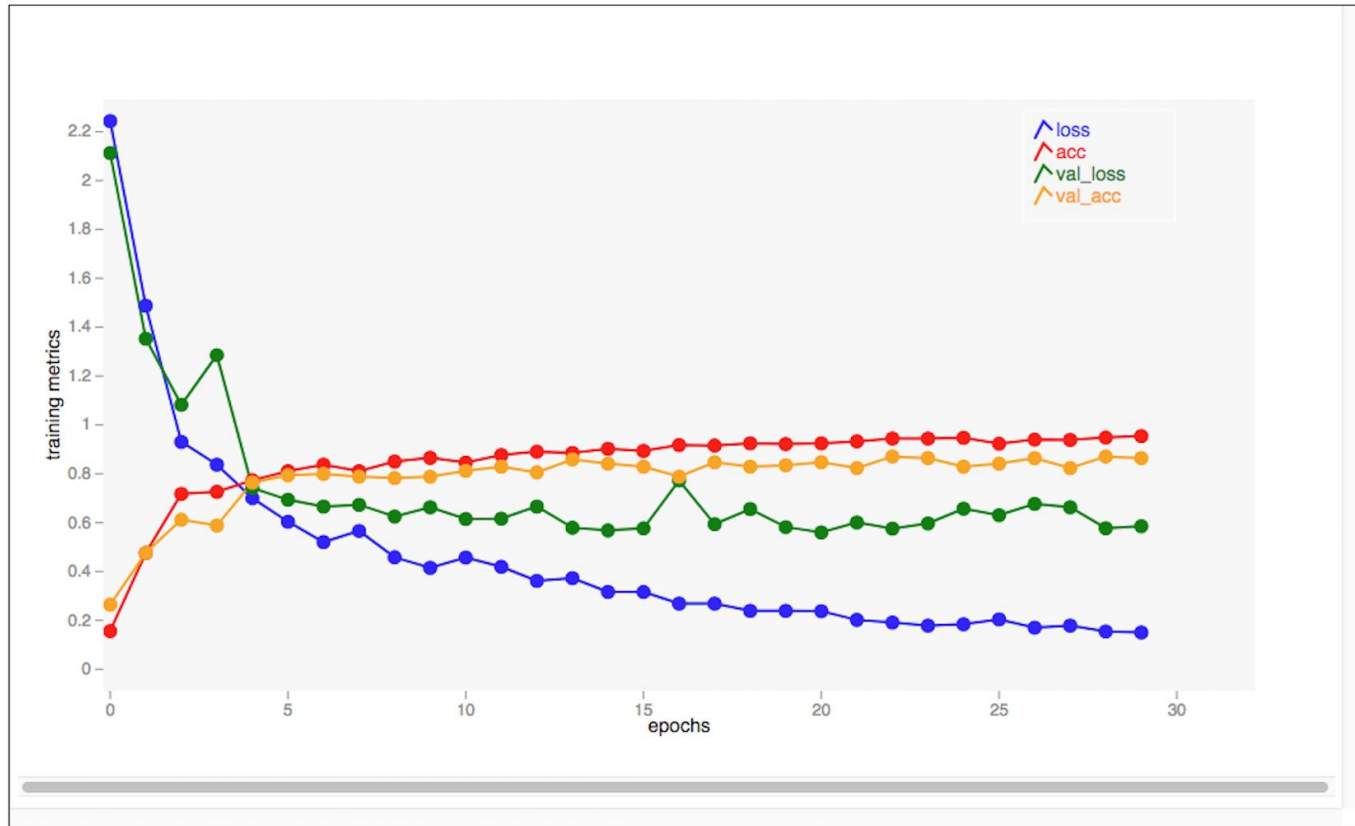
```
Hyperparameter trial 0 conv [ 64 16 128] fc [128] dropout 0.3234 opt Nadam, lr 0.0100
Hyperparameter trial 1 conv [ 4 8 64] fc [64] dropout 0.6747 opt Adadelata, lr 0.0010
```

Plots update live
Table shows
different
configurations:

- Status
- Current loss and accuracy
- Sort

Can add further
quantities to plot
and interaction
buttons

<https://github.com/sparticlesteve/cori-intml-examples/>



Stop selected Restart selected

index	status	epoch	h1	h2	h3	dropout	optimizer	loss	val_loss	acc	val_acc
3	Ended Tra...	31	16	64	16	0.88467	Adam	2.003565...	1.689878...	0.213253...	0.682352...
2	Ended Tra...	31	16	8	8	0.19765	Adam	0.852827...	0.829521...	0.763855...	0.800000...
0	Ended Tra...	31	64	8	8	0.04836	Adadelta	0.157987...	0.579903...	0.944578...	0.870588...
1	Ended Tra...	31	4	8	16	0.03825	Adadelta	0.151153...	0.585079...	0.954216...	0.864705...

```

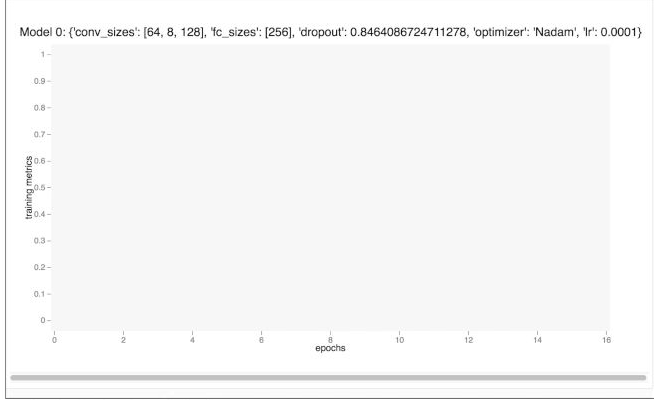
conv_sizes=conv_sizes,
fc_sizes=fc_sizes,
dropout=dropout,
optimizer=optimizer,
lr=lr,
)

paw = ParamSpanWidget(
    compute_func=train_func,
    vis_func=plot_func,
    params=hpo_params,
    jpp_cluster_id=cluster_id
)

paw.submit_computations()

paw

```



Index	status	epoch	conv_size#	fc_sizes	dropout	optimizer	lr	loss	val_loss	acc	val_acc
0	Not Started	-1	[64, 8, 128]	[256]	0.84641	Nadam	0.0001	nan	nan	nan	nan
1	Not Started	-1	[4, 8, 16]	[32]	0.69948	Nadam	0.01	nan	nan	nan	nan
2	Not Started	-1	[32, 8, 32]	[128]	0.29744	Adam	0.0001	nan	nan	nan	nan
3	Not Started	-1	[32, 4, 128]	[128]	0.8138	Adam	0.0001	nan	nan	nan	nan
4	Not Started	-1	[32, 16, 64]	[256]	0.39651	Adadelta	0.01	nan	nan	nan	nan
5	Not Started	-1	[8, 64, 128]	[128]	0.8811	Adadelta	0.0001	nan	nan	nan	nan
6	Not Started	-1	[32, 32, 1...	[256]	0.58127	Adadelta	0.01	nan	nan	nan	nan
7	Not Started	-1	[16, 32, 1...	[256]	0.88174	Nadam	0.01	nan	nan	nan	nan
8	Not Started	-1	[64, 16, 64]	[256]	0.69253	Nadam	0.01	nan	nan	nan	nan
9	Not Started	-1	[4, 64, 128]	[128]	0.72525	Adam	0.0001	nan	nan	nan	nan
10	Not Started	-1	[4, 16, 128]	[64]	0.50132	Adam	0.0001	nan	nan	nan	nan
11	Not Started	-1	[64, 4, 128]	[128]	0.95608	Adadelta	0.0001	nan	nan	nan	nan
12	Not Started	-1	[16, 4, 8]	[128]	0.64399	Adadelta	0.001	nan	nan	nan	nan
13	Not Started	-1	[8, 64, 128]	[256]	0.42386	Nadam	0.01	nan	nan	nan	nan
14	Not Started	-1	[4, 4, 64]	[128]	0.60639	Adam	0.0001	nan	nan	nan	nan
15	Not Started	-1	[8, 64, 32]	[256]	0.01919	Adam	0.001	nan	nan	nan	nan

Curated Notebook Environments



- Browse curated examples managed by project using tools like nbviewer
- Clone and launch notebook into users workspace with appropriate conda environment
- Reproducible Notebooks – Similar to Binder



py4DSTEM / py4DSTEM / notebooks / basic_demos

Name



..

dataobjects.ipynb

filehandling_nativefiles.ipynb

filehandling_non-nativefiles.ipynb

metadata.ipynb

Browse

View

```
datacube.aet_scan_shape(10,10)

In [9]: # And now the data shape looks more reasonable!
print(datacube.dataID.shape)
print(datacube.R_Ry, datacube.R_Rx, datacube.Q_Rx, datacube.Q_Ry)
(10, 10, 512, 512)
10 10 512 512

In [10]: # Basic processing and display

# Get an average image at each real-space scan position, add up all the counts in the x
# pattern. (i.e. # bright field image with a large integrating detector)
avg_in = np.sum(datacube.dataID,axis=2,3)

# Get a single diffraction pattern from a particular beam position
DP = datacube.dataID[5,1,1]

# Display
fig,(ax1,ax2)=plt.subplots(1,2,figsize=(12,6))
ax1.imshow(avg_in,cmap='gray')
ax2.imshow(DP*0.2,cmap='gray')
plt.show()
```

```
Terminal 1
dataobjects.ipynb
Name Last Modified
dataobjects.ipynb seconds ago
filehandling_nativefile... seconds ago
filehandling_non-nativ... seconds ago
metadata.ipynb seconds ago

DataObjects
This notebook discusses how py4DSTEM structures data. It includes 3 secti
1. The DataObjects
2. Searching DataObjects
3. Saving DataObjects

[1]: import numpy as np
import matplotlib.pyplot as plt
import py4DSTEM

Part 1: The DataObjects
In py4DSTEM there are several classes which contain distinct types of data,
basic classes discussed here are:
1. DataCube
2. DiffractionSlice
3. RealSlice
4. PointList
5. PointListArray

[2]: # Import the DataObject classes
from py4DSTEM.file.datastructure import DataCube
from py4DSTEM.file.datastructure import DiffractionSlice, RealS
from py4DSTEM.file.datastructure import PointList, PointListArr
```

Launch

Parameterized Notebooks



- Run the same notebook against
 - different sets of parameters
 - Different datasets
 - Think “data parallel”
- Running as notebook gives you a live document of each task
- Save successful runs

```
In [1]: parameters x ... Add tag
1 # This cell is tagged `parameters`
2 alpha = 0.1
3 ratio = 0.1
```

Executing a Notebook

The two ways to execute the notebook with parameters are: (1) through the Python API and (2) through the comr interface.

Execute via the Python API

```
import papermill as pm

pm.execute_notebook(
    'path/to/input.ipynb',
    'path/to/output.ipynb',
    parameters = dict(alpha=0.6, ratio=0.1)
)
```


Jupyterlab Extensions



- SLURM Extension
- Resource Usage Monitoring Extension

The screenshot shows the JupyterLab interface with the Slurm Queue Manager extension. The sidebar on the left contains a search bar, a console section with various actions like 'Change Kernel...', 'Clear Console Cells', and 'New Console', and a file operations section with 'Autosave Documents' and 'Close All'. The main window displays the 'Slurm Queue Manager' tab, which includes a search bar, action buttons (Reload, Kill Selected Job(s), Hold Selected Job(s), Release Selected Job(s), Deselect all), and a table of job queue information.

JOBID	PARTITION	NAME	USER	ST	TIME	NODES	NOELIST(REASON)
10923303	regular	7zz_chev	fzzhao	PD	0:00	8	(JobHeldUser)
10923311	regular	shortuni	fzzhao	PD	0:00	8	(JobHeldUser)
10923436	regular	7zzaynth	fzzhao	PD	0:00	8	(JobHeldUser)
11062030	regular	shortuni	fzzhao	PD	0:00	10	(JobHeldUser)
1117376	regular	REG_D	mwu	PD	0:00	80	(JobHeldUser)
11476878	regular	REG_C_TH	mwu	PD	0:00	64	(JobHeldUser)
11495724	regular	REG_L_TH	mwu	PD	0:00	64	(JobHeldUser)
13386219	regular	REG_R_TH	mwu	PD	0:00	64	(JobHeldUser)
13386711	regular	Ndpd_k5	fzzhao	PD	0:00	16	(JobHeldUser)
13386772	regular	Ndpd_k6	fzzhao	PD	0:00	16	(JobHeldUser)
13386715	regular	Ndpd_k7	fzzhao	PD	0:00	16	(JobHeldUser)

Below the table, there is a 'Show' dropdown set to 'entries', a status 'Showing 1 to 15 of 18,631 entries 1 row selected', and a pagination control with buttons for 'Previous', '1', '2', '3', '4', '5', '1243', and 'Next'. A checkbox labeled 'Show my jobs only' is also present.