

PRIONN: Predicting Runtime and IO using Neural Networks

Michael R. Wyatt, Stephen Herbein, Todd Gamblin,
Adam Moody, Dong H. Ahn, and Michela Taufer



Acknowledgements



Stephen Herbein



Todd Gamblin



Adam Moody



Dong H. Ahn




Michela Taufer

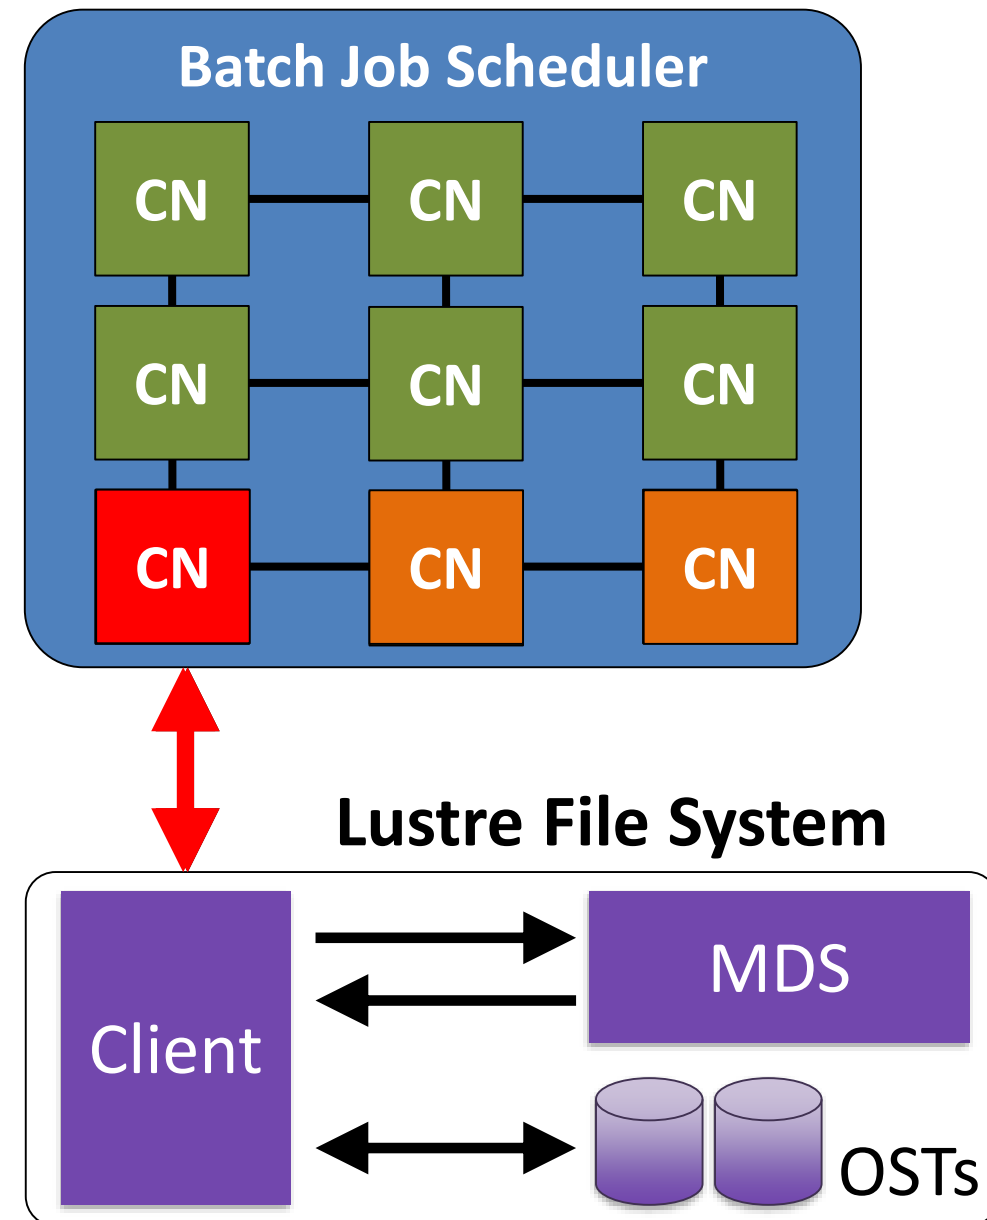
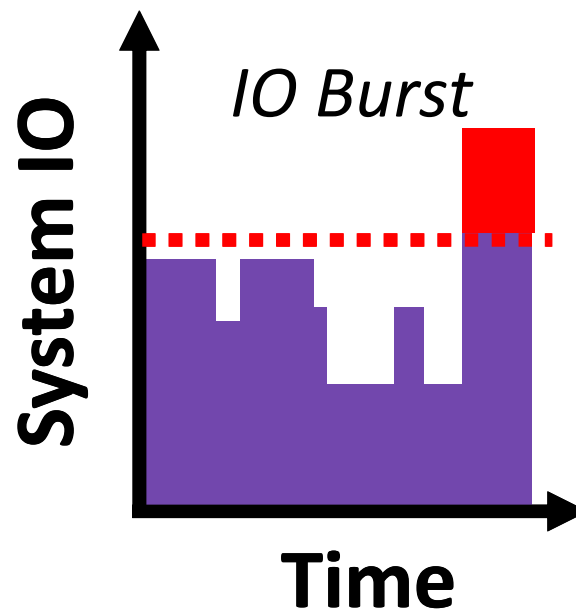
This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.



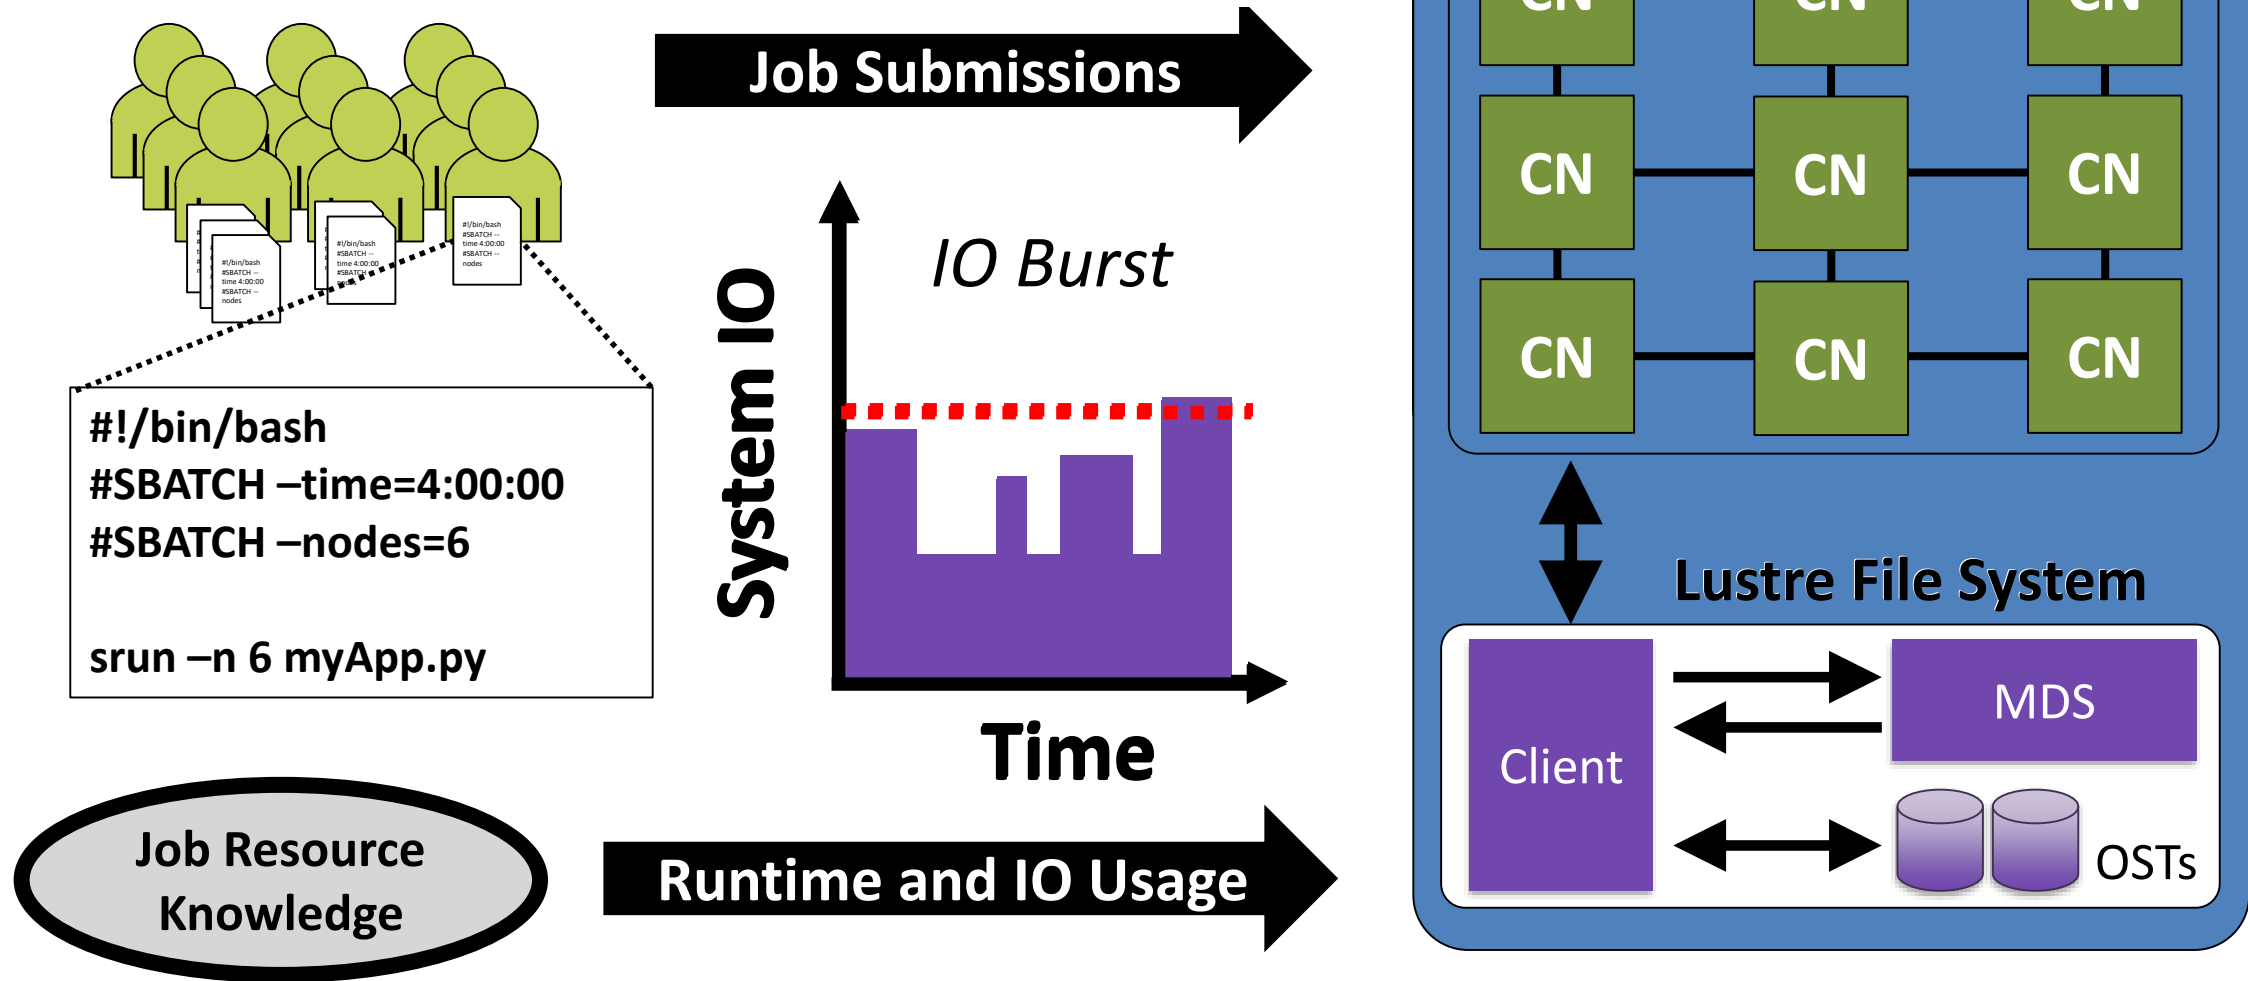
HPC Batch Job Scheduling



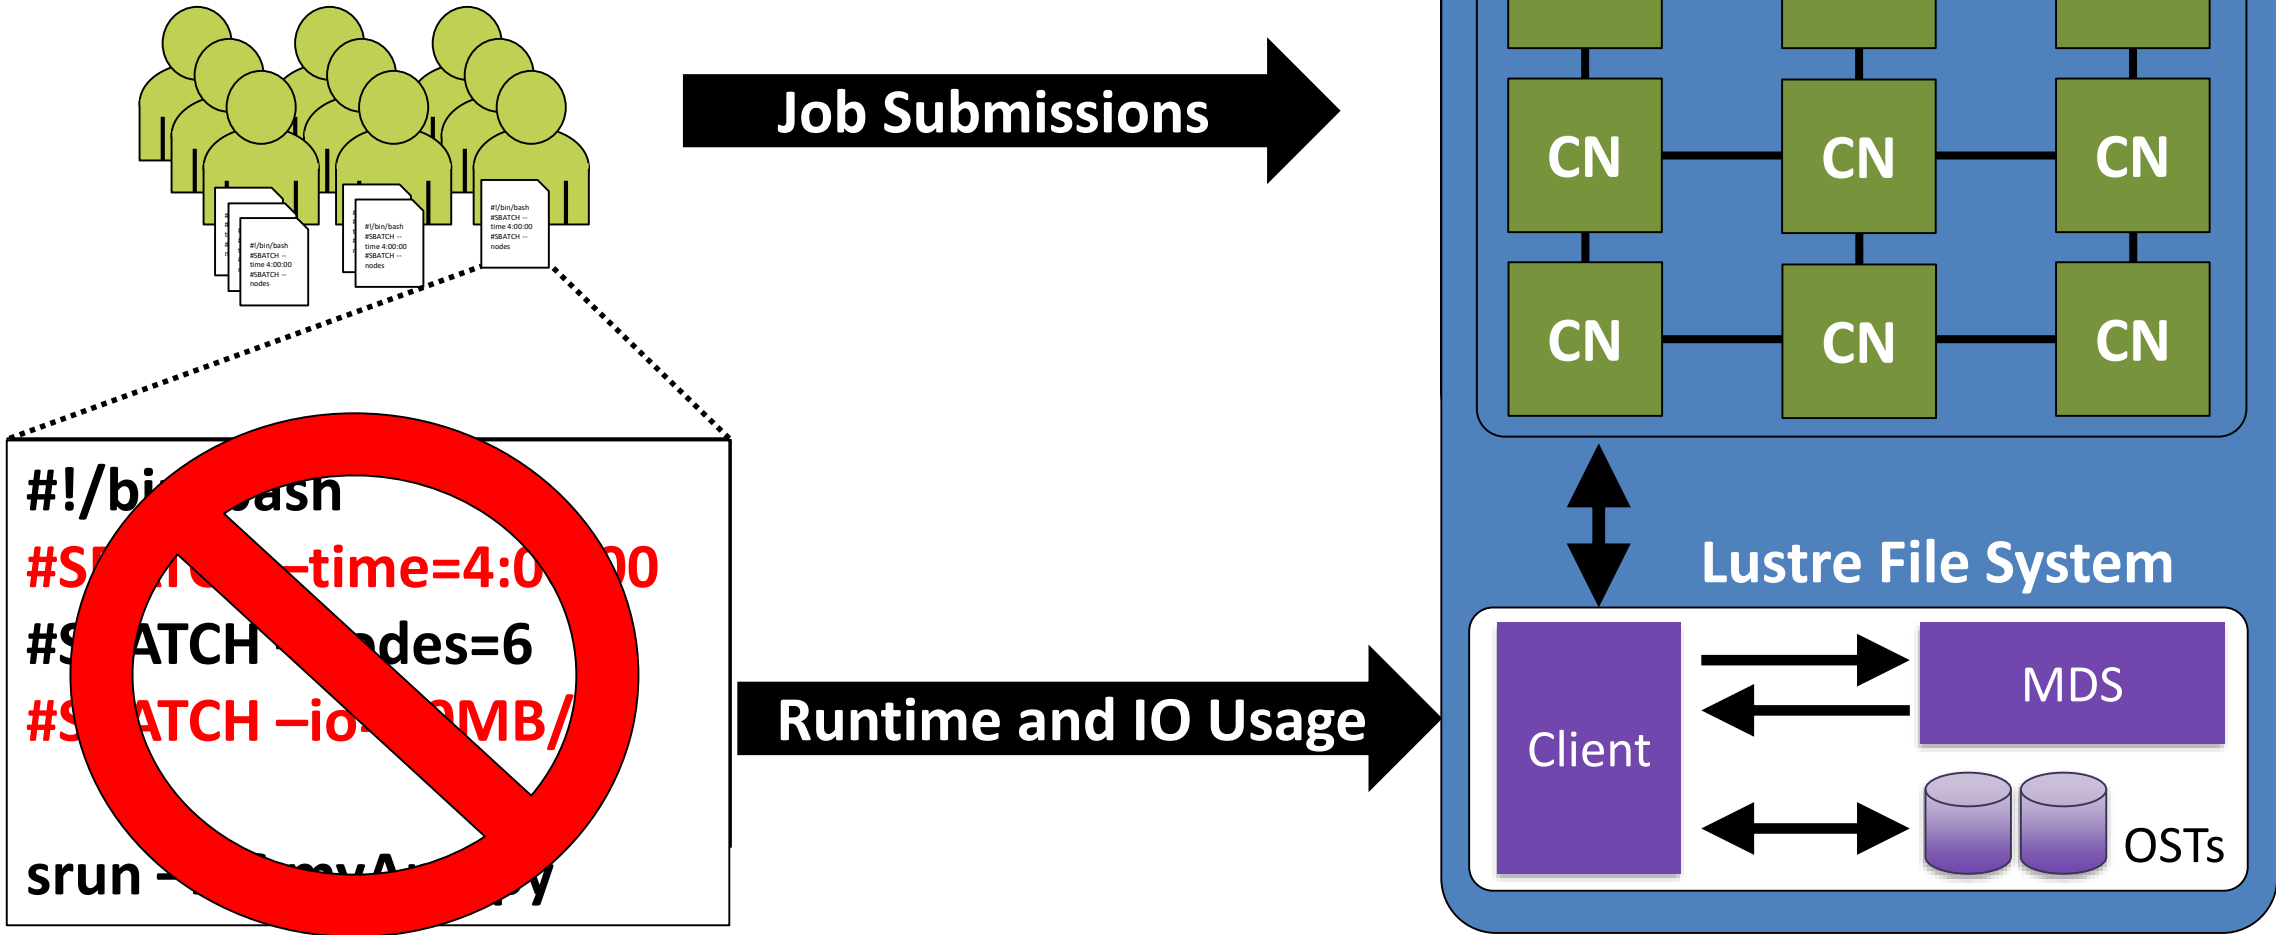
Job Submissions



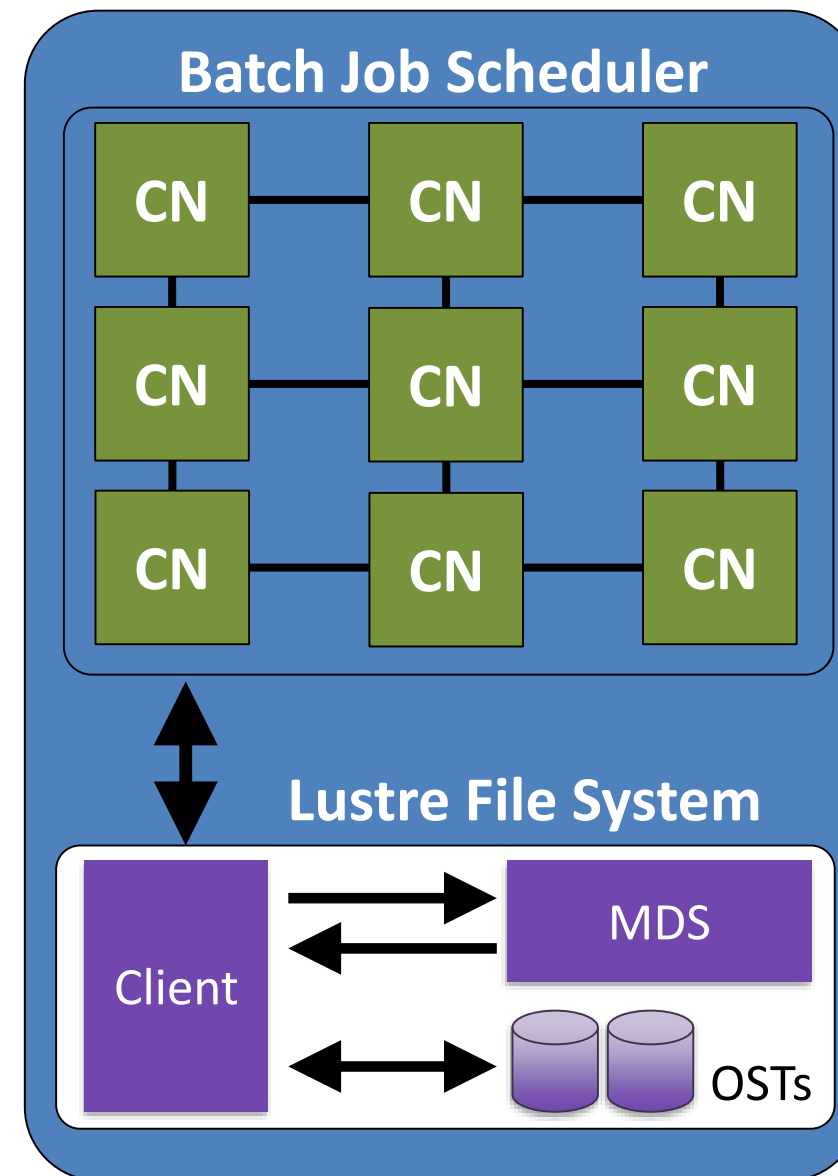
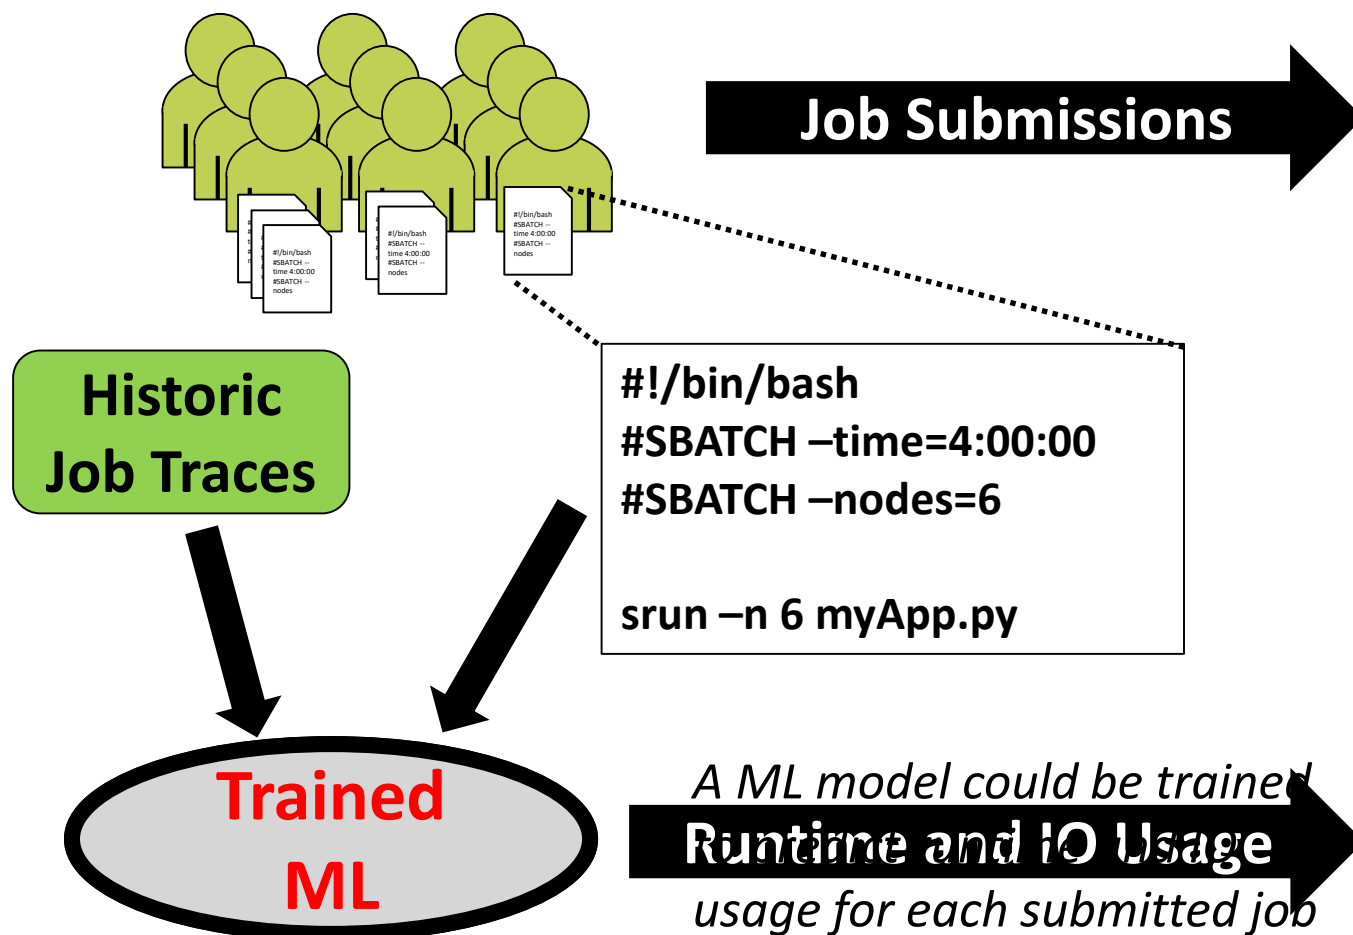
HPC Batch Job Scheduling (Future)



Persuading Job Resource Knowledge




Resource Usage Prediction



PRIONN: From Job Scripts to Predictions

- We want an **automatic**, **general**, and **accurate** tool for resource prediction
 - **Automatic**: avoid development and maintenance of parsers
 - **General**: works with any type of job script
 - **Accurate**: better prediction accuracy than other ML methods



PRIONN: Learning From Entire Job Scripts

- We convert job scripts to image-like representations
- Text is mapped to vectors

Job Script

```
#SBATCH --time=4:00:00  
#SBATCH --nodes=6
```

(Structured)

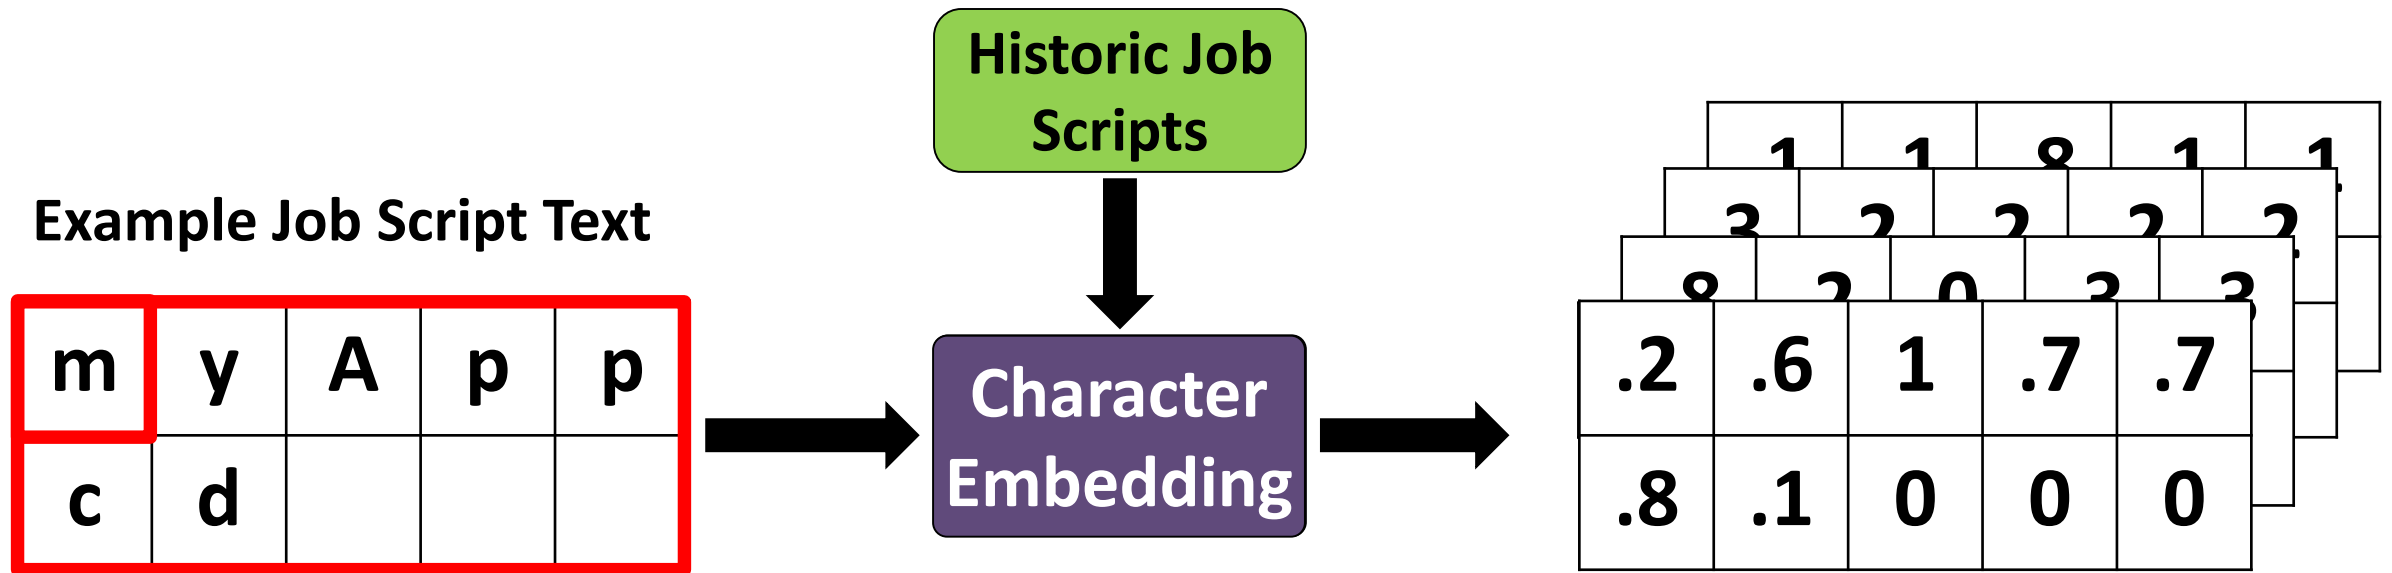
```
cd $HOME/project_A  
mkdir results  
echo "Running simulation"  
for i in [0..8]; do  
    INPUT=data/$i.csv  
    srun -n 48 ./myApp $INPUT  
done
```

(Unstructured)

Image-like Representation




PRIONN: Job Script to Image



- We test several character embedding methods
- **Word2Vec** gave best predictive accuracy

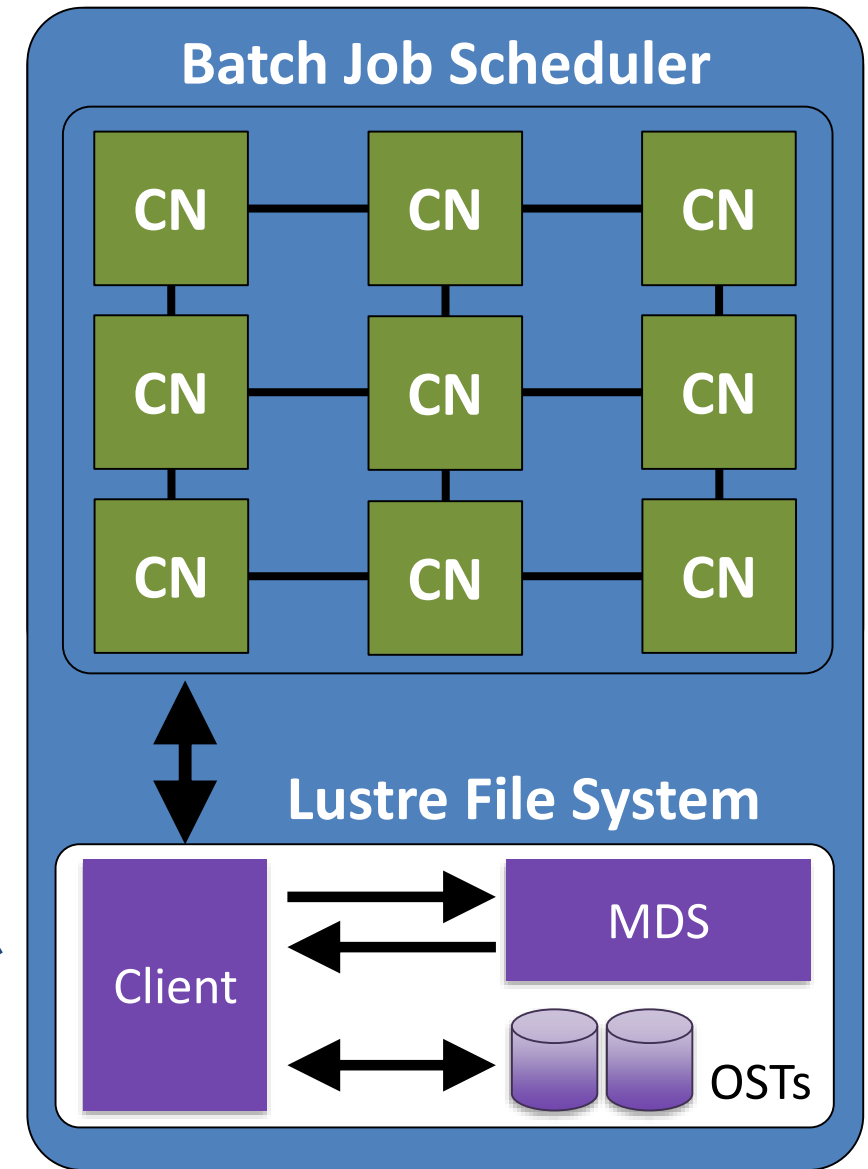

PRIONN: From Job Scripts to Predictions

- We leverage the power of CNNs to learn and predict from our image-like job script representation

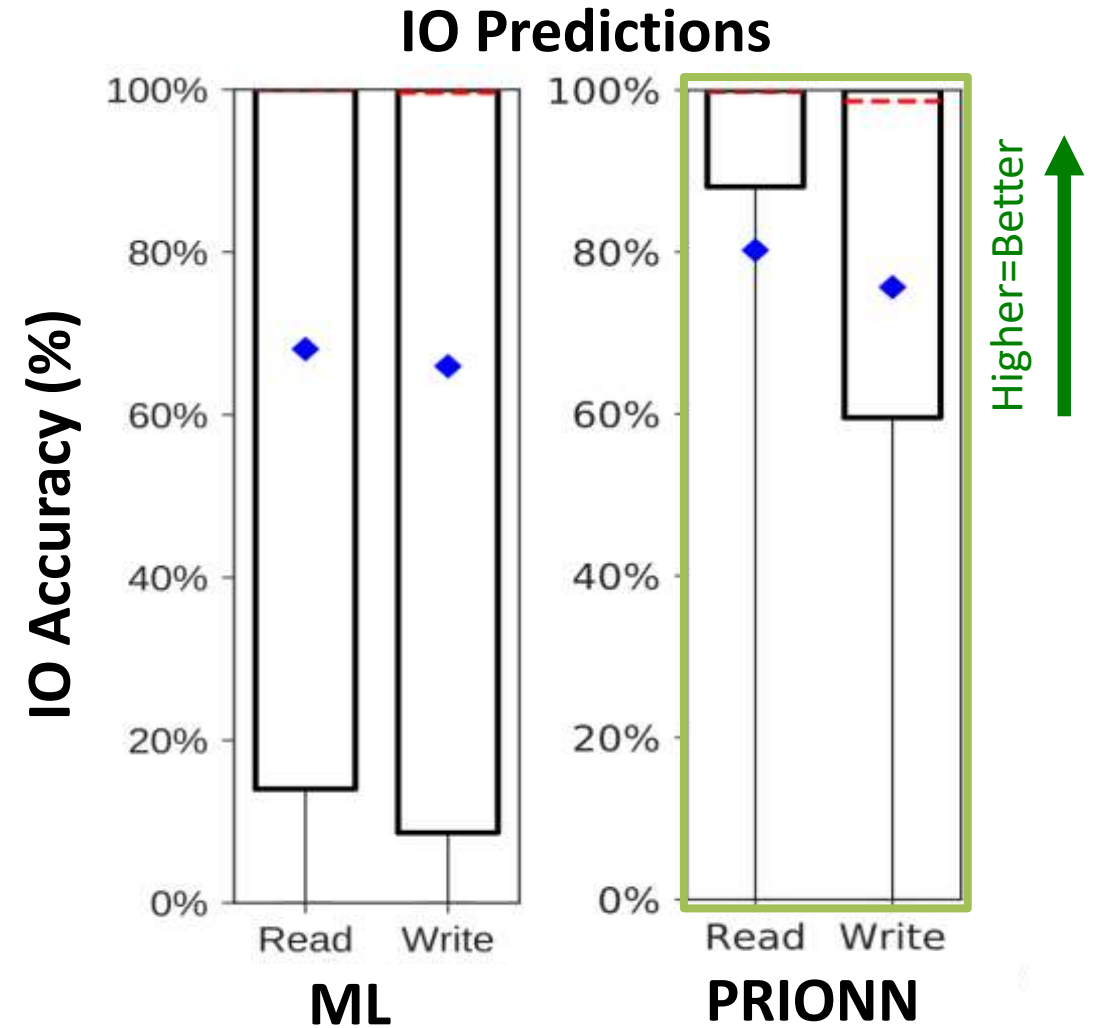
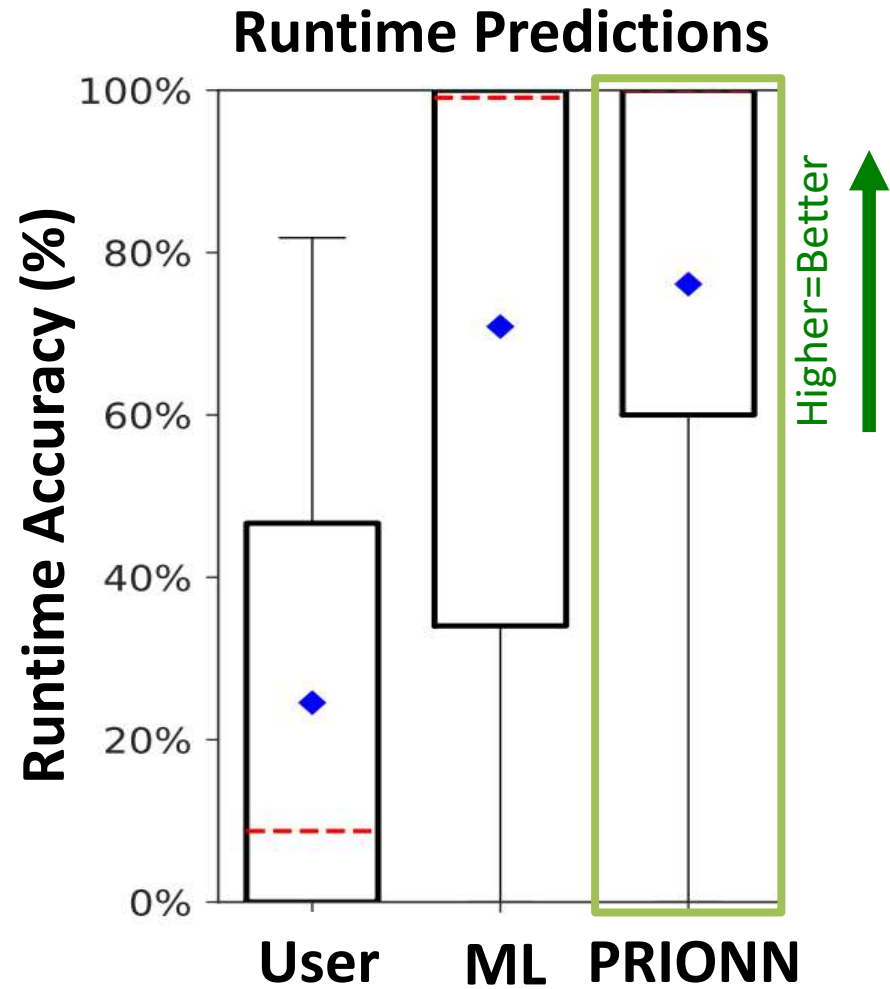


PRIONN Training and Prediction

- We continuously retrain PRIONN with the most recent system data




Per-Job Runtime and I/O Predictions

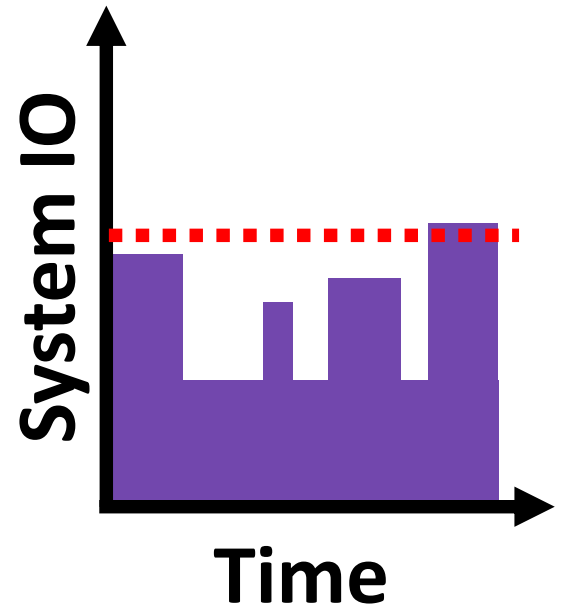


Enabling IO-Aware Scheduling

1. Advance per-job runtime knowledge ✓
2. Advance per-job IO usage knowledge ✓
3. Feed runtime and IO usage to system simulation to model future system IO ✓




Future System IO Prediction



[1] D. Ahn, et al. Flux: A Next-Generation Resource Management Framework for Large HPC Centers (ICPP'14)

[2] S. Herbein, et al. Scalable I/O-Aware Job Scheduling for Burst Buffer Enabled HPC Clusters (HPDC'16)

Evaluating IO Burst Predictions



- We compare our predicted and real system IO
- A correct IO burst prediction falls within the *anticipation window* for the real IO burst
- We tally TP, FP, TN, and FN IO burst predictions

Evaluating IO Burst Predictions

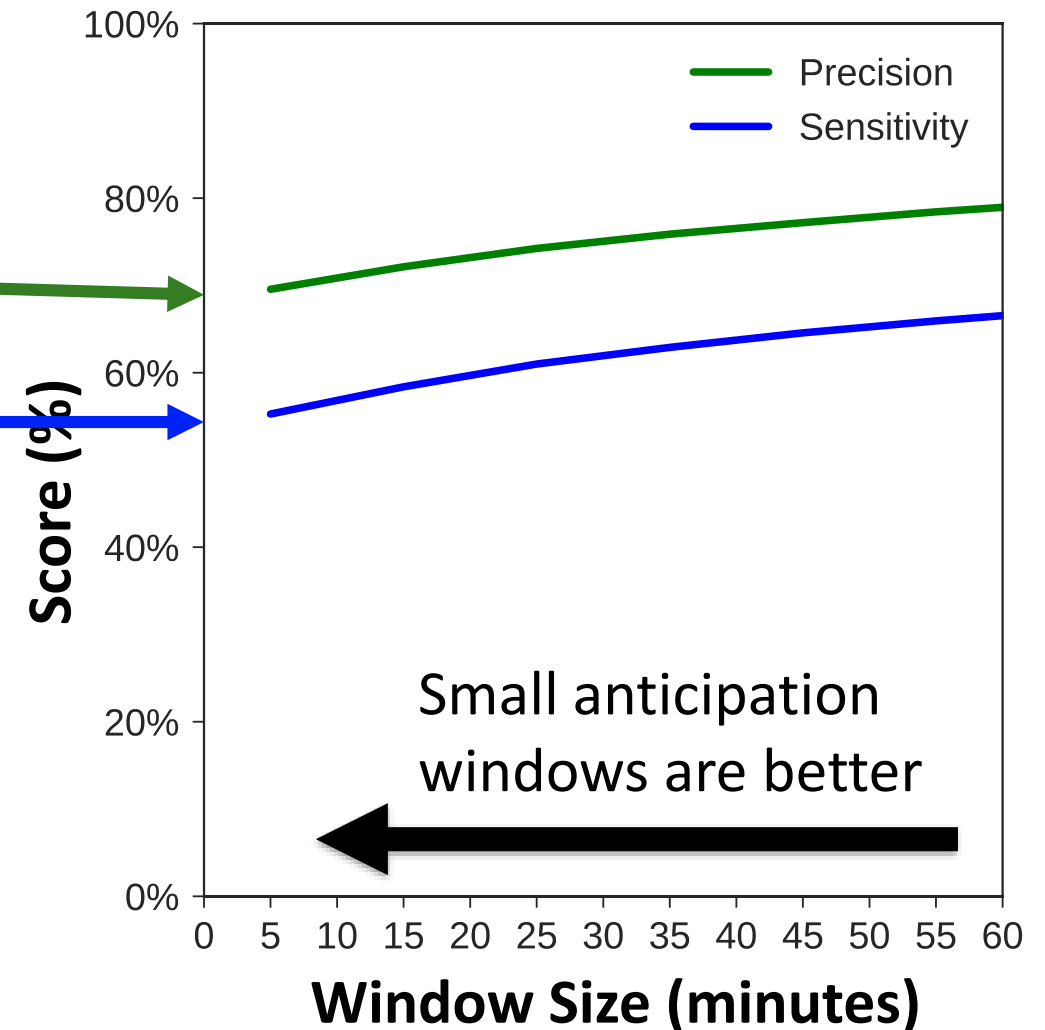
Precision and sensitivity of anticipating IO burst given a window size in minutes

70% of anticipated bursts occur

55% of bursts are anticipated

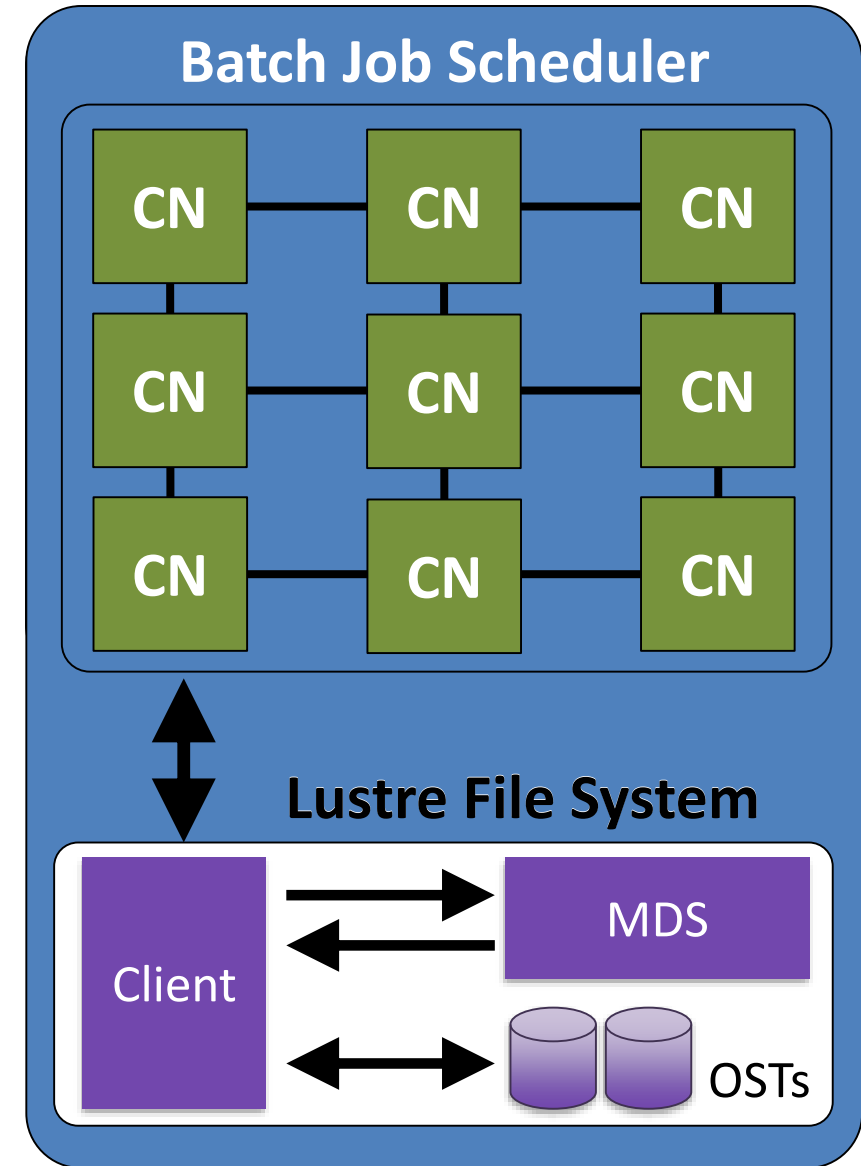
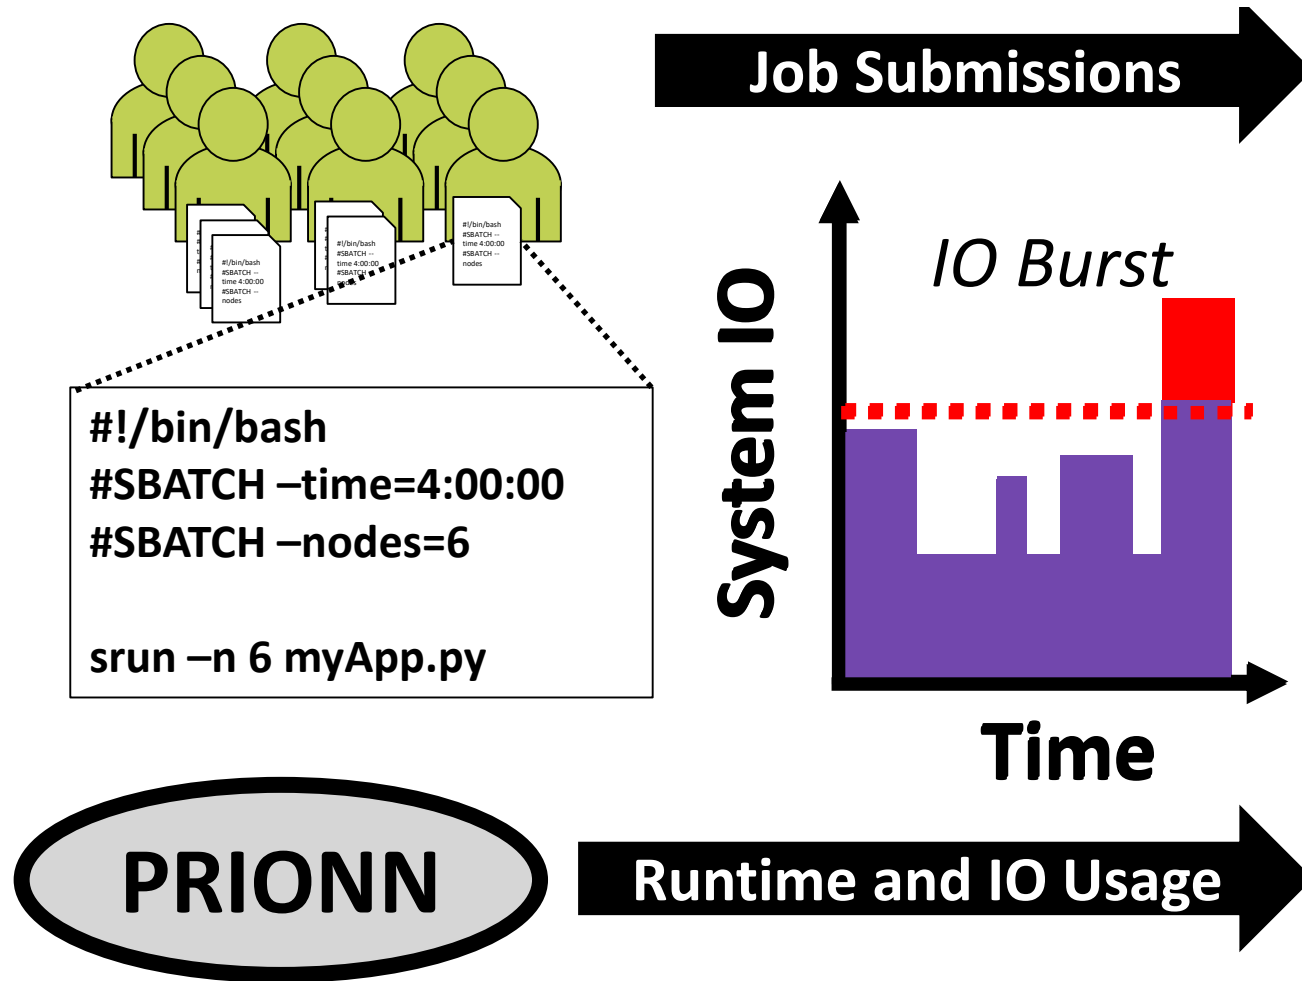
Precision = $\frac{TP}{TP + FP}$ *(How many predicted IO bursts occur?)*

Sensitivity = $\frac{TP}{TP + FN}$ *(How many IO bursts are predicted?)*



Based on: <https://commons.wikimedia.org/wiki/File:Precisionrecall.svg>

HPC Batch Job Scheduling (Future)



Lessons Learned

- Next-generation HPC systems will require IO-aware batch schedulers
- PRIONN is a resource prediction tool for IO-aware schedulers
 - Per-job runtime and IO usage predictions with >75% mean, >95% median accuracy
 - We forecast >55% of system IO bursts before they occur
- PRIONN uses deep learning and does not require job script parsing
 - Can be quickly deployed on many systems
 - Can be used with any job scripts
 - Accurate for resource usage prediction