On Scaling Deep Learning

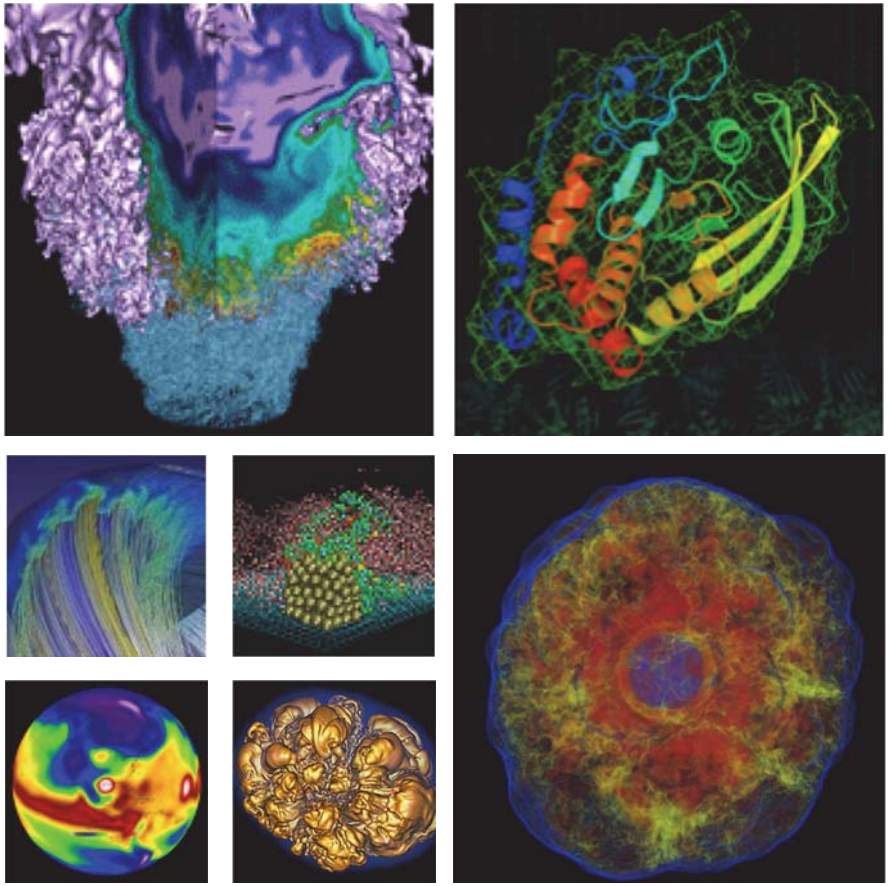

NERSC

Prabhat

11/6/2018

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Outline

- **Introduction to NERSC**
  – Why Scale Deep Learning?
- **Case Studies**
  – Deep Learning @15 PF
  – CosmoFlow @3.5 PF
  – Deep Learning @1 EF
- **Open Challenges**
- **Conclusions**

NERSC: the Mission HPC Facility for DOE Office of Science Research

NERSC

U.S. DEPARTMENT OF ENERGY | Office of Science

Largest funder of physical science research in the U.S.

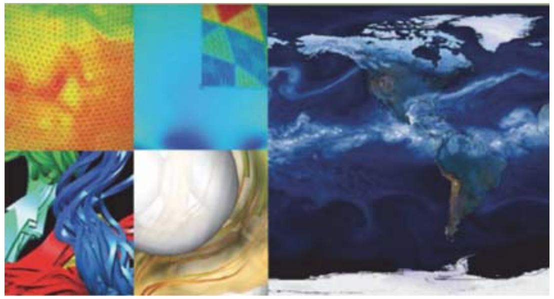

Bio Energy, Environment

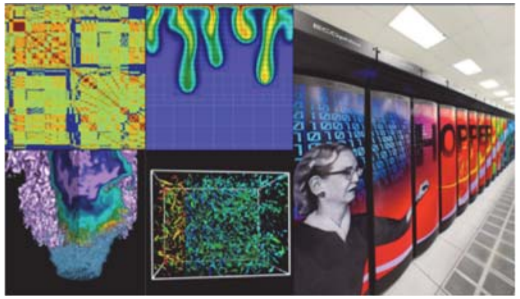

Computing

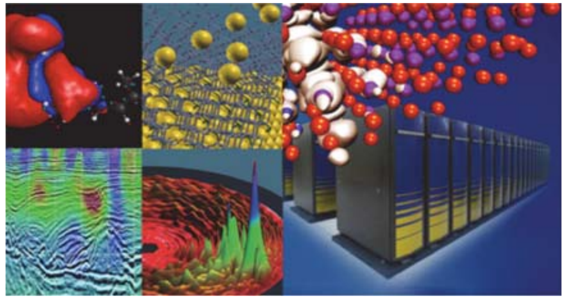

Materials, Chemistry, Geophysics

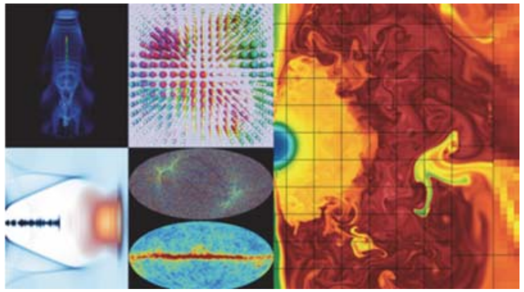

Particle Physics, Astrophysics

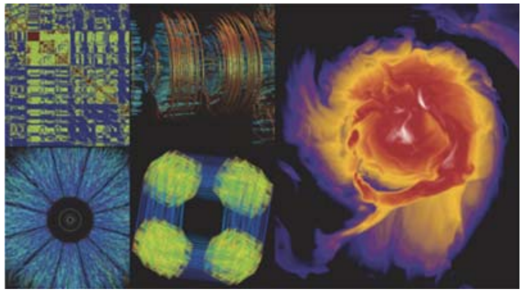

Nuclear Physics

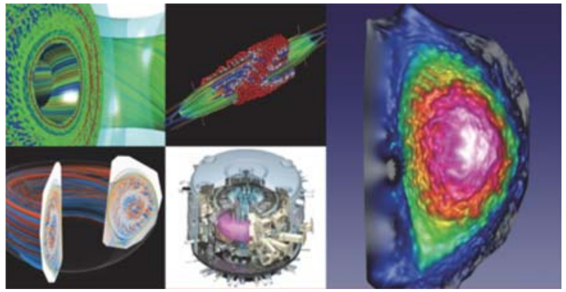

Fusion Energy, Plasma Physics

7,000 users, 800 projects, 700 codes, 48 states, 40 countries, universities & national labs

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

NERSC Hardware

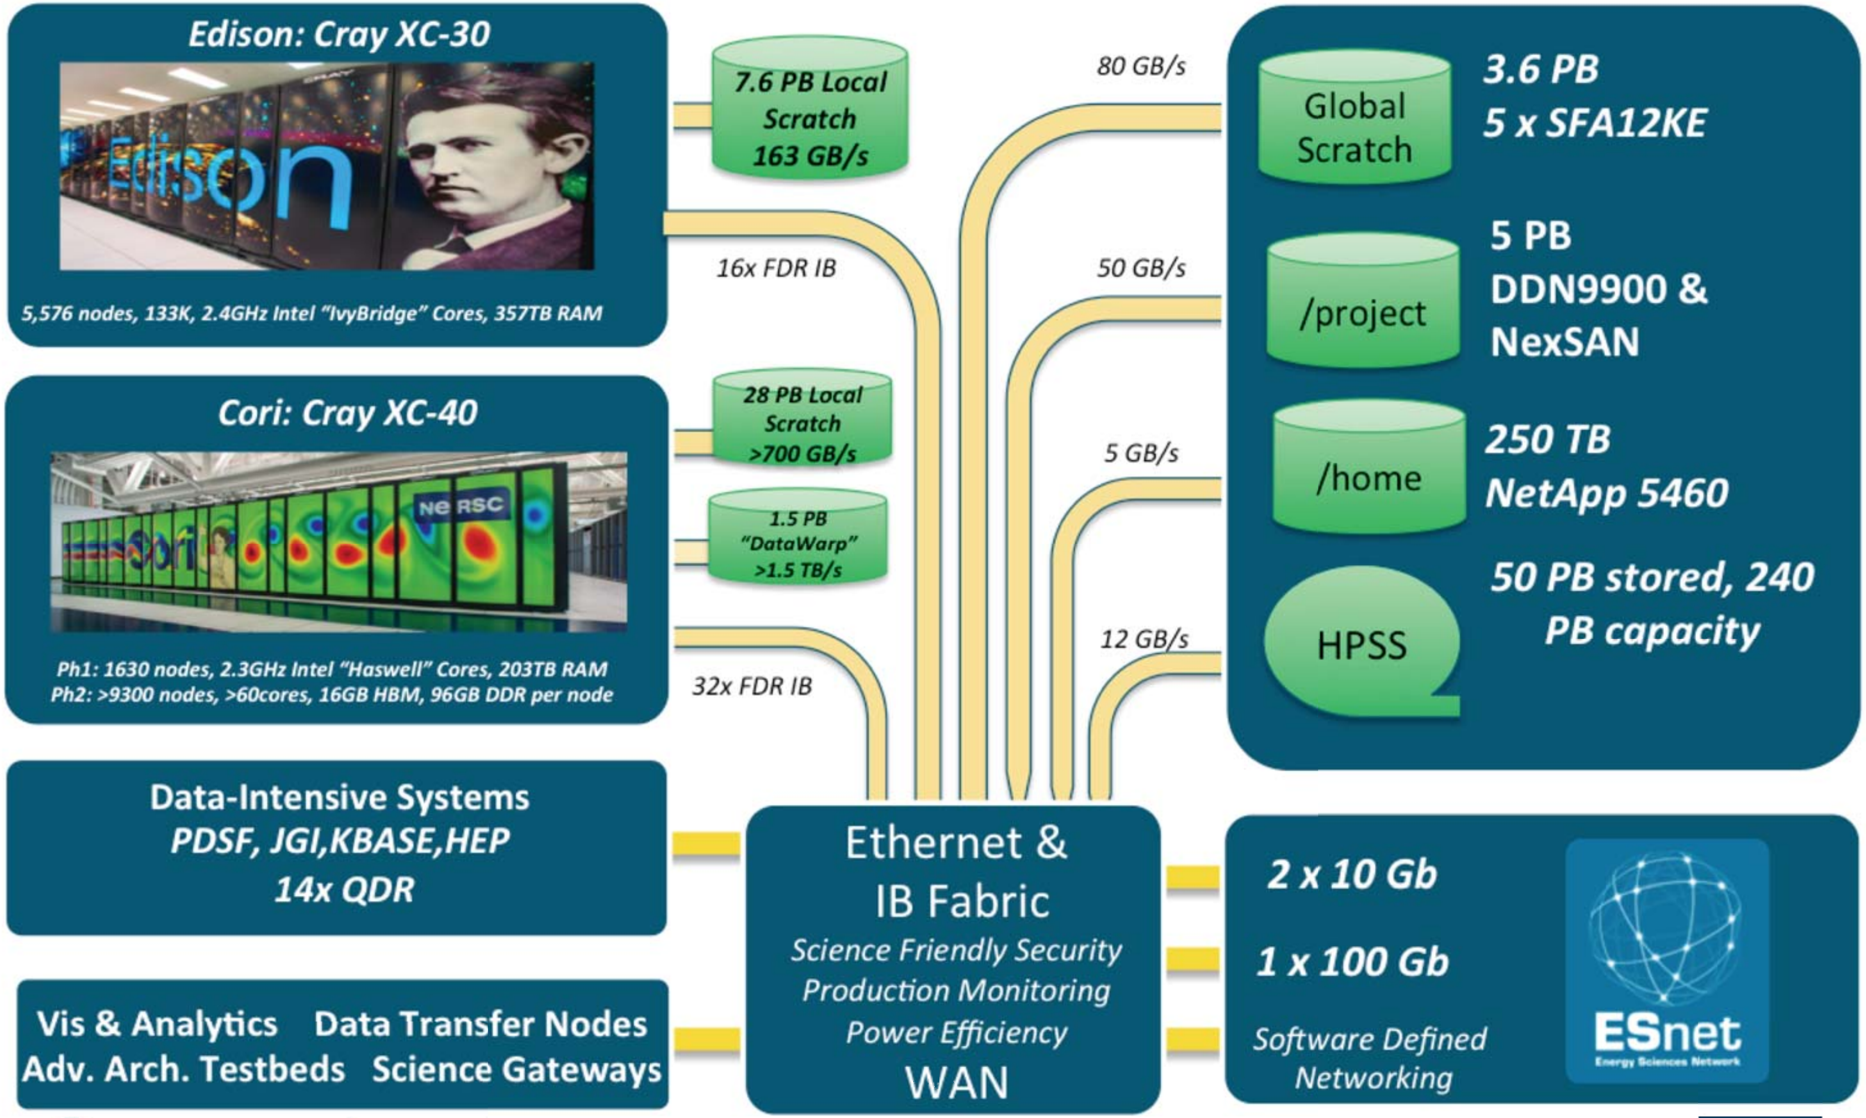

	HEP			BER		BES		NP	FES
	Astronomy	Cosmology	Particle Physics	Climate	Genomics	Light Sources	Materials	Heavy Ion Colliders	Plasma Physics
Classification	X		X	X	X	X	X	X	X
Regression	X	X	X	X	X	X	X	X	X
Clustering	X	X	X	X	X	X	X	X	X
Dimensionality Reduction				X				X	
Surrogate Models	X	X	X	X			X	X	X
Design of Experiments		X		X		X	X		X
Feature Learning	X	X	X	X	X	X	X	X	X
Anomaly Detection	X		X	X		X		X	

Data Analytics Methods

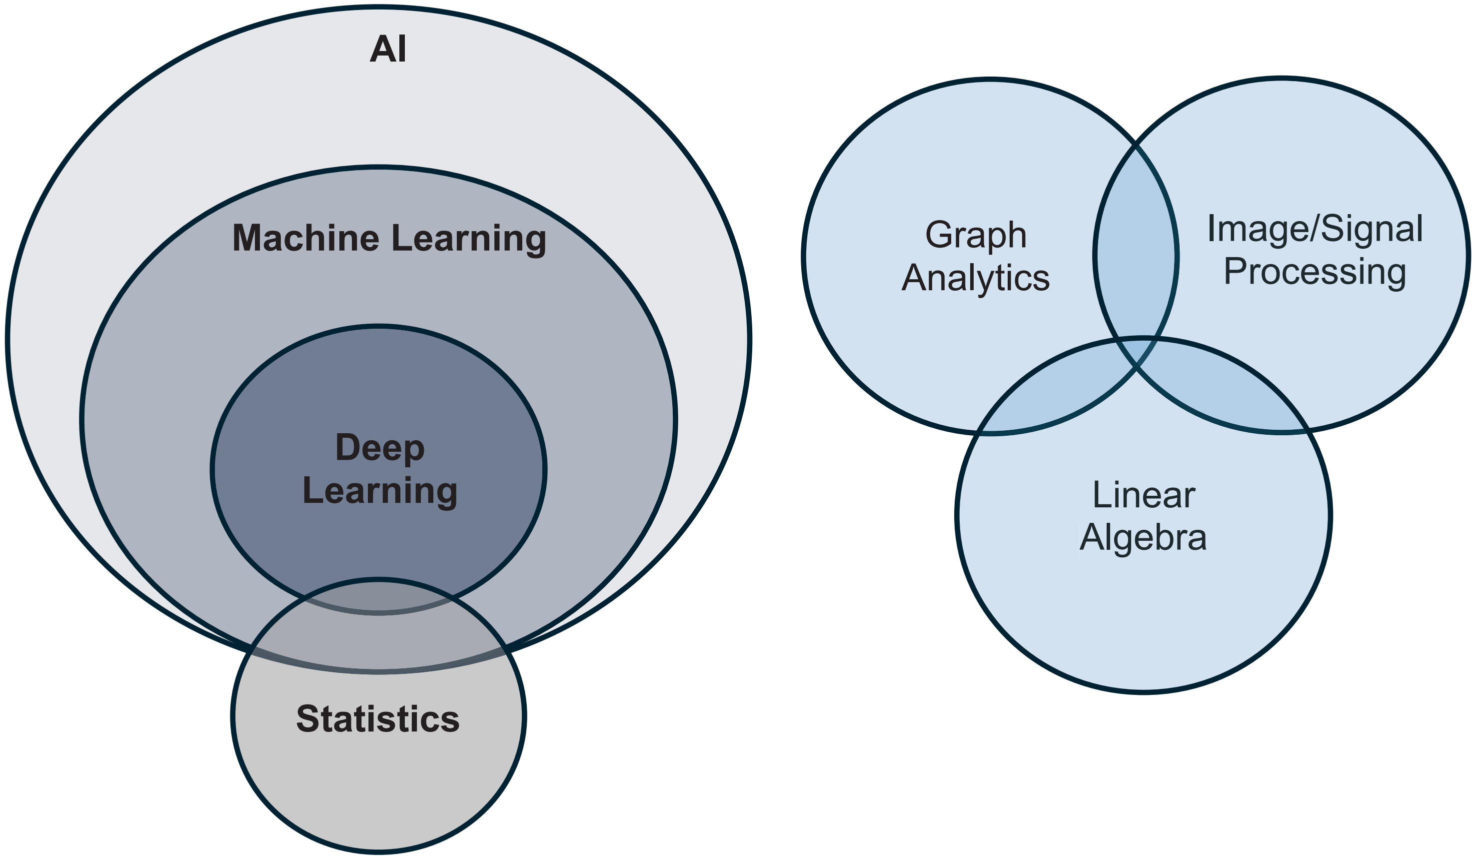

NERSC Big Data Stack

Capabilities	Technologies
Data Transfer + Access	globus online, GridFTP, jupyter, python django, newt
Workflows	FireWorks, taskfarmer
Data Management	HDF, netCDF, ROOT, mongoDB, MySQL, PostgreSQL
Data Analytics	python, Spark, TensorFlow, PYTORCH, R, julia, Caffe, K
Data Visualization	visit, ParaView

Deep Learning Stack

	Technologies
Deep Learning Frameworks	K, TensorFlow, Caffe, PYTORCH, CNTK, MXNet, Neon, …
Multi Node libraries	Cray Plugin, Horovod, GRPC, MLSL, MPI
Single Node libraries	MKL-DNN, CuDNN
Hardware	CPUs (KNL), GPUs, FPGAs, Accelerators

Deep Learning for Science

Modeling galaxy shapes

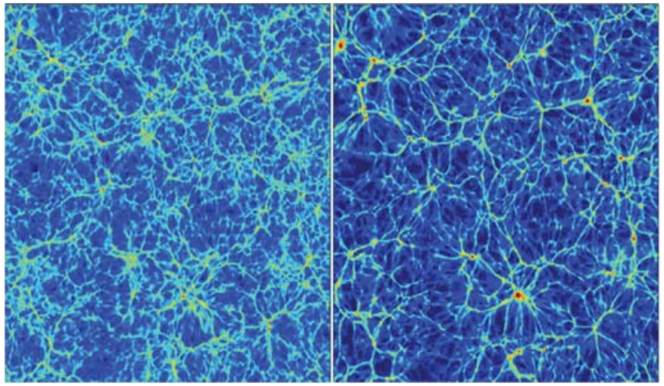

Generating cosmology mass maps

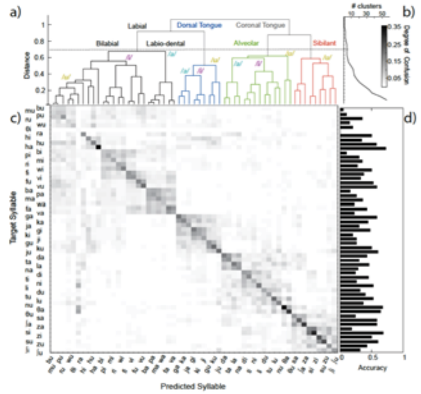

Decoding speech from ECoG

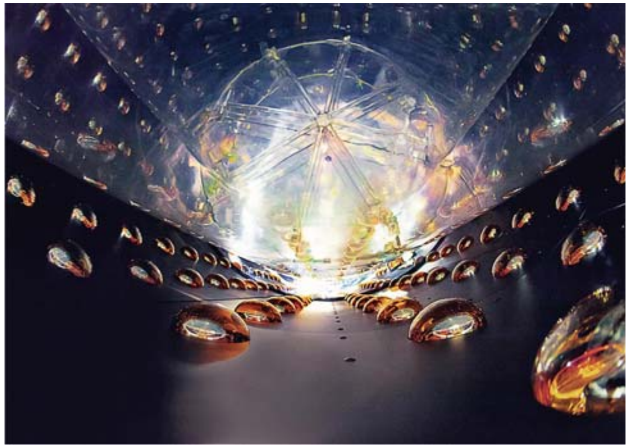

Clustering Daya Bay events

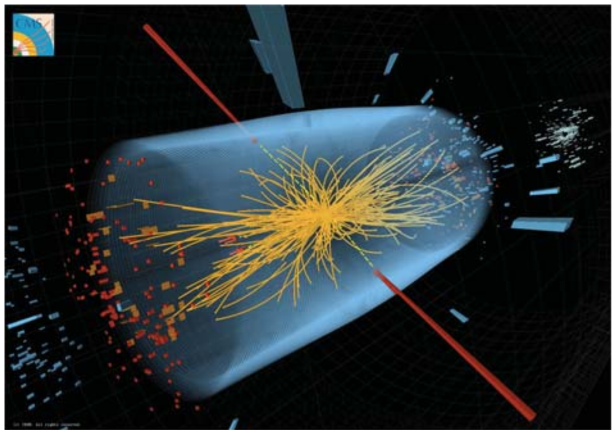

LHC Signal/Background classification

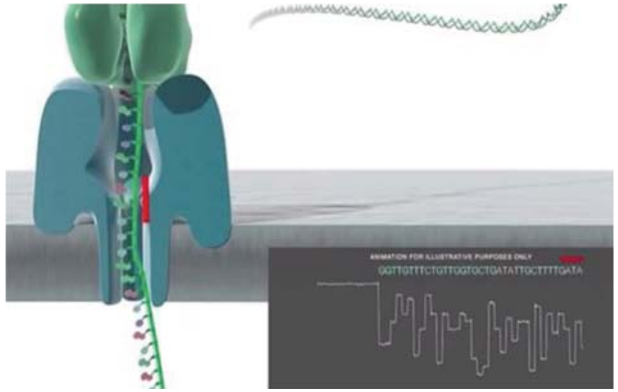

Oxford Nanopore sequencing

U.S. DEPARTMENT OF ENERGY | Office of Science

https://www.oreilly.com/ideas/a-look-at-deep-learning-for-science

BERKELEY LAB

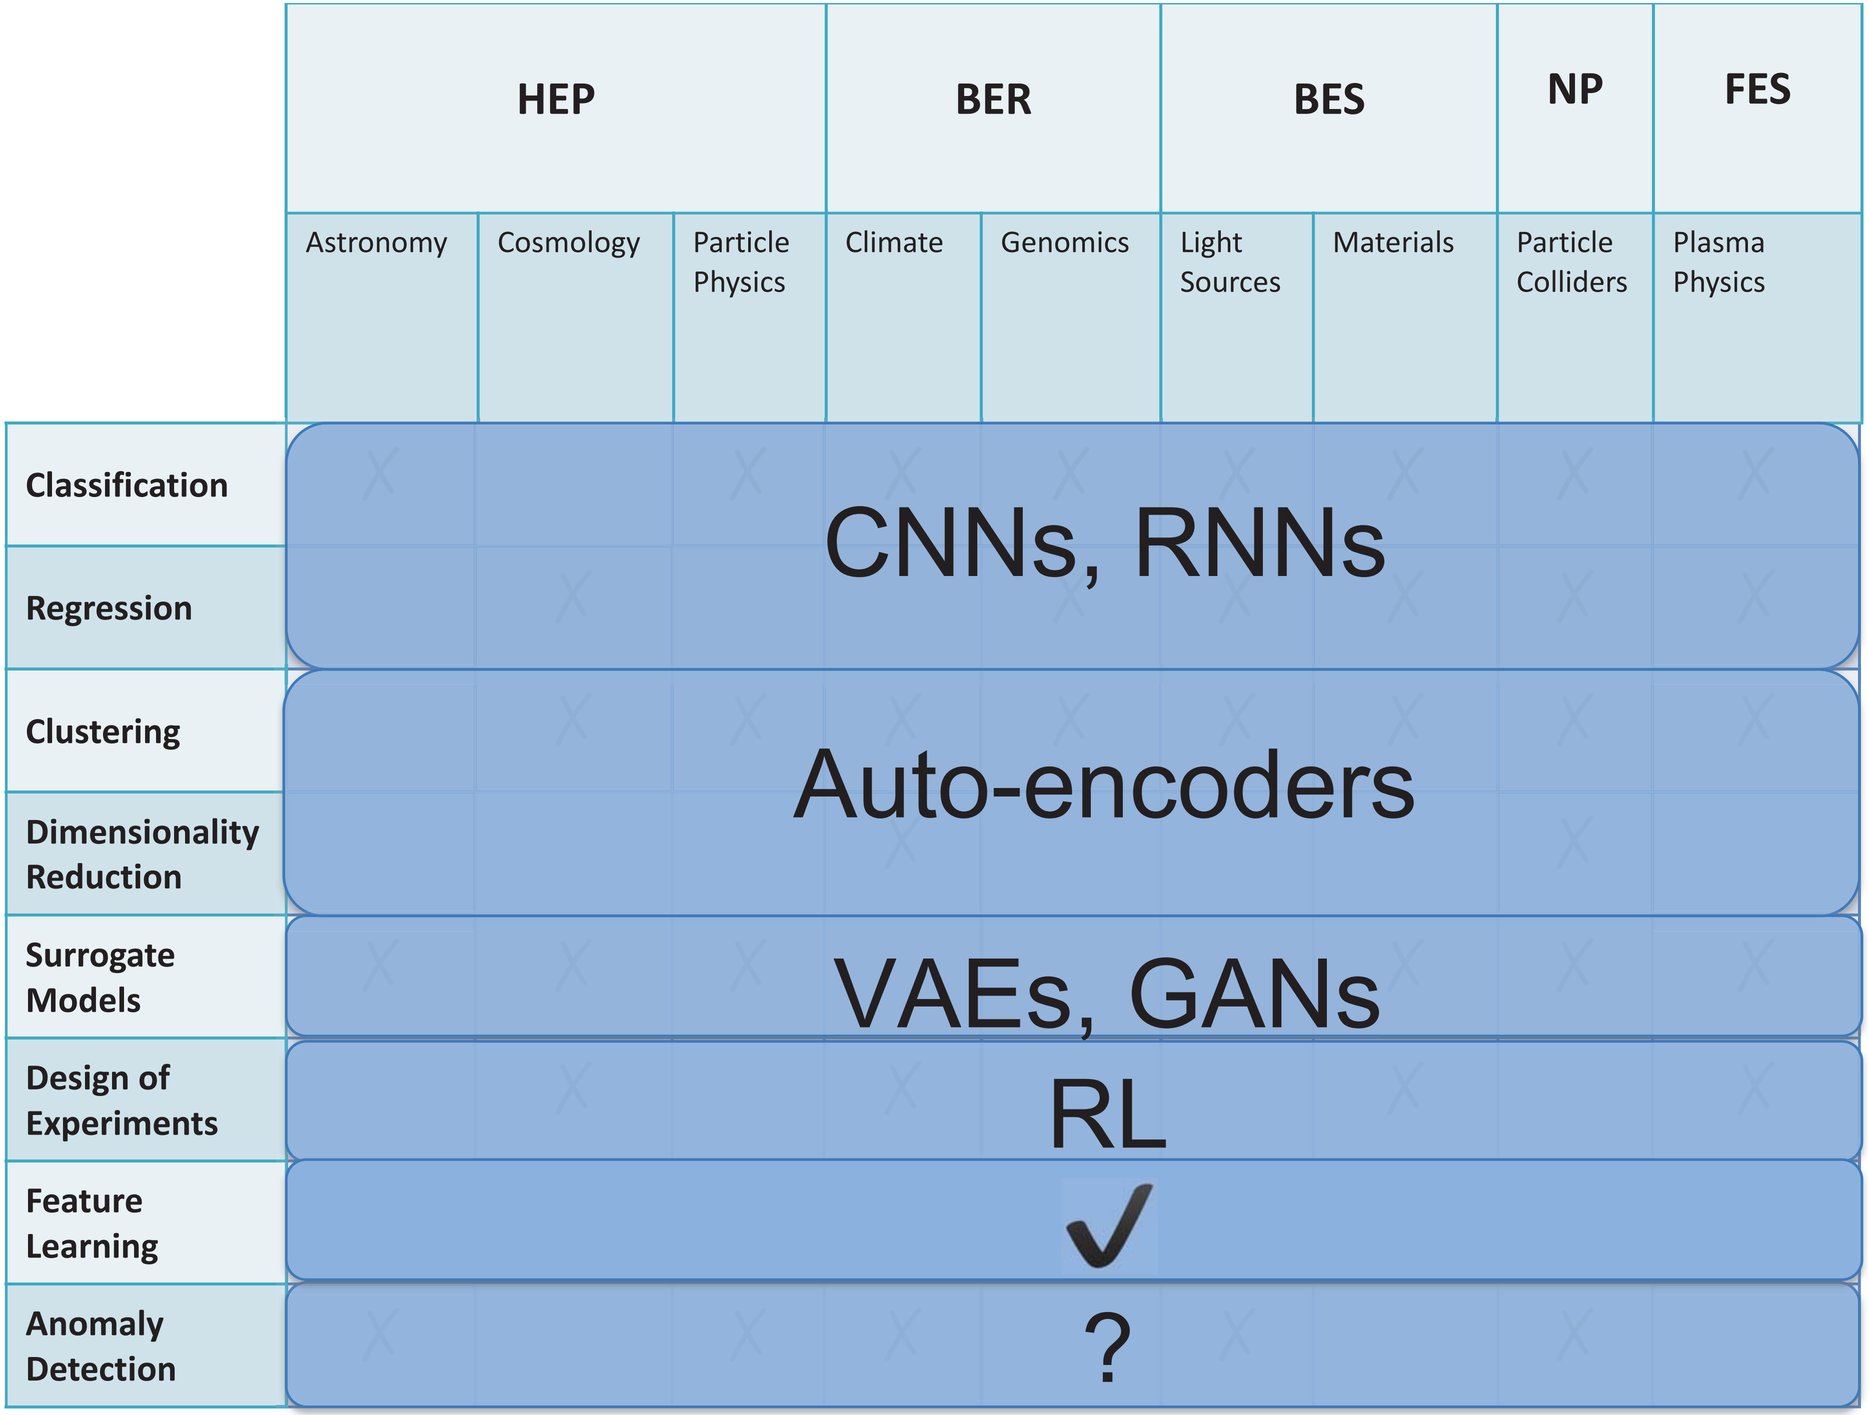

HEP
BER
BES
NP
FES
Astronomy
Cosmology
Particle Physics
Climate
Genomics
Light Sources
Materials
Particle Colliders
Plasma Physics
Classification
Regression
Clustering
Dimensionality Reduction
Surrogate Models
Design of Experiments
Feature Learning
Anomaly Detection
CNNs, RNNs
Auto-encoders
VAEs, GANs
RL
✓
?

DL adoption by NERSC user community

- 150-200 users exercising the DL stack
- 160 respondents to recent ‘ML@NERSC’ survey

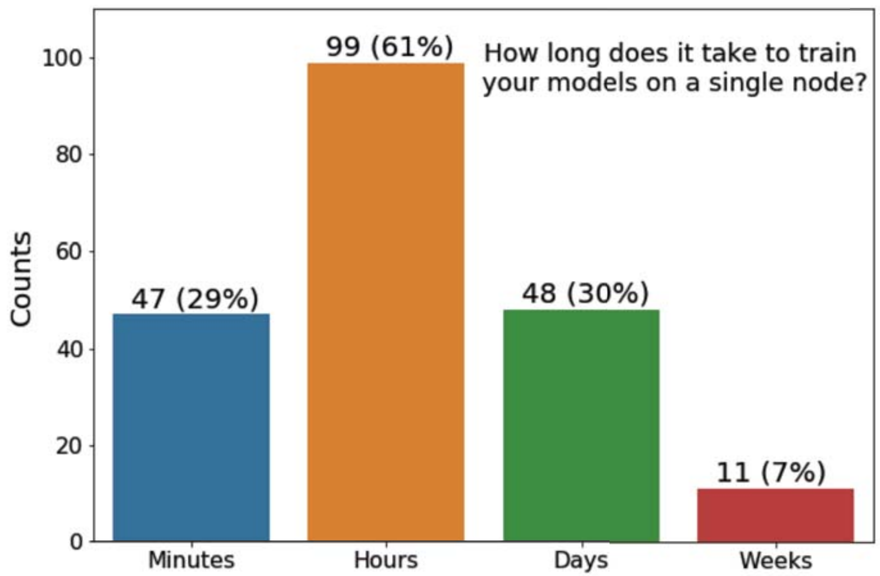

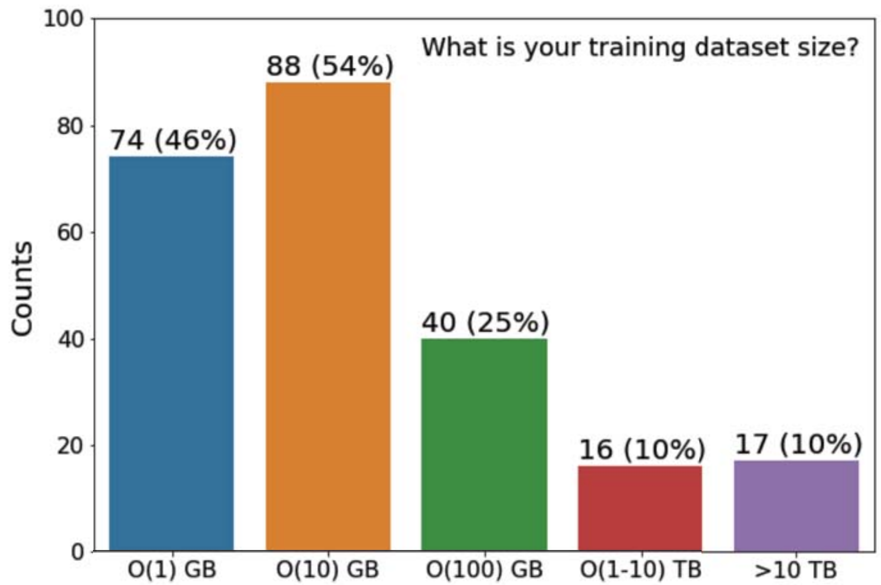

Why Scale Deep Learning?

- **Day/Week-long runtimes for O(100) GB - O(1) TB sized datasets**
  - 'Classical' convolutional architectures
  - More advanced architectures (Hybrid CNN + LSTM, spacetime convolutions)
- **Hyper-Parameter optimization is important**
- **Large computational demands**

- **Problem is well suited for HPC systems**

Outline

- Introduction to NERSC
  - Why Scale Deep Learning?
- **Case Studies**
  - Deep Learning @15 PF
  - CosmoFlow @3.5 PF
  - Deep Learning @1 EF
- Open Challenges
- Conclusions

Deep Learning @15 PF

- ACM/IEEE Supercomputing 2017 Paper
- Thorsten Kurth, Jian Zhang, Nadathur Satish, Evan Racah, Ioannis Mitliagkas, Md. Mostofa Ali Patwary, Tareq Malas, Narayanan Sundaram, Wahid Bhimji, Mikhail Smorkalov, Jack Deslippe, Mikhail Shiryaev, Srinivas Sridharan, Prabhat, Pradeep Dubey

https://arxiv.org/abs/1708.05256

Characterizing Extreme Weather in a Changing Climate

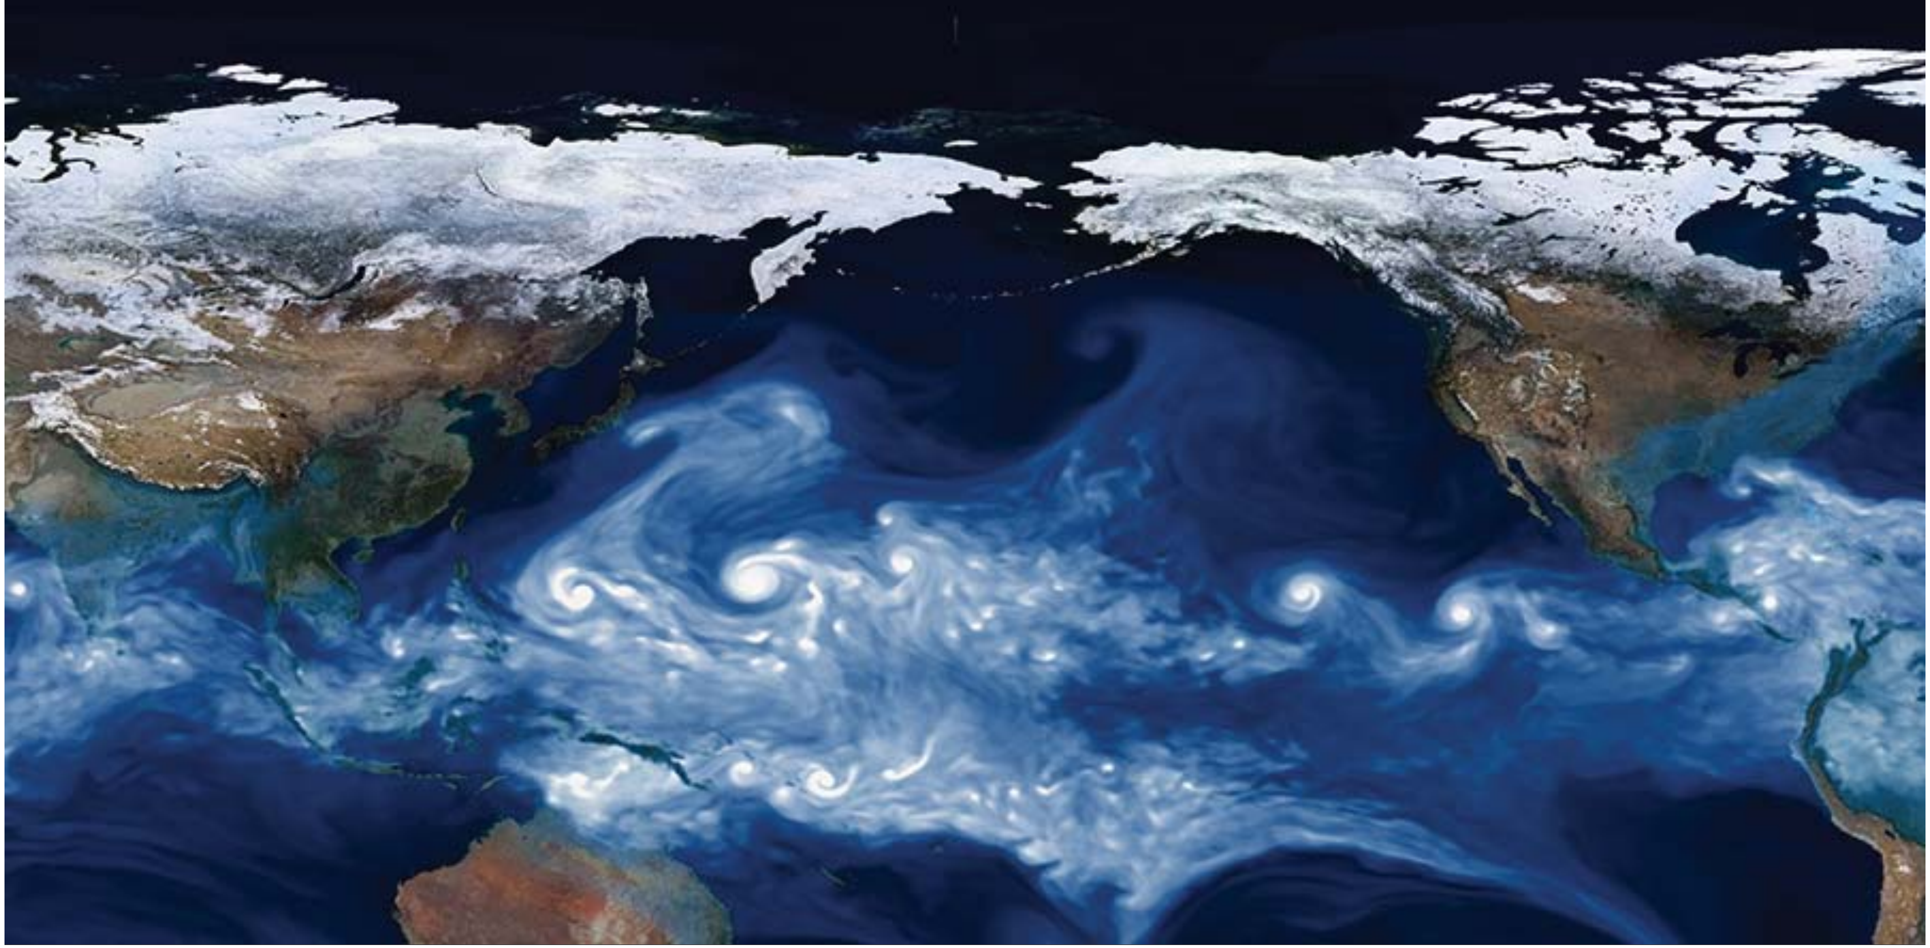

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Extreme Weather

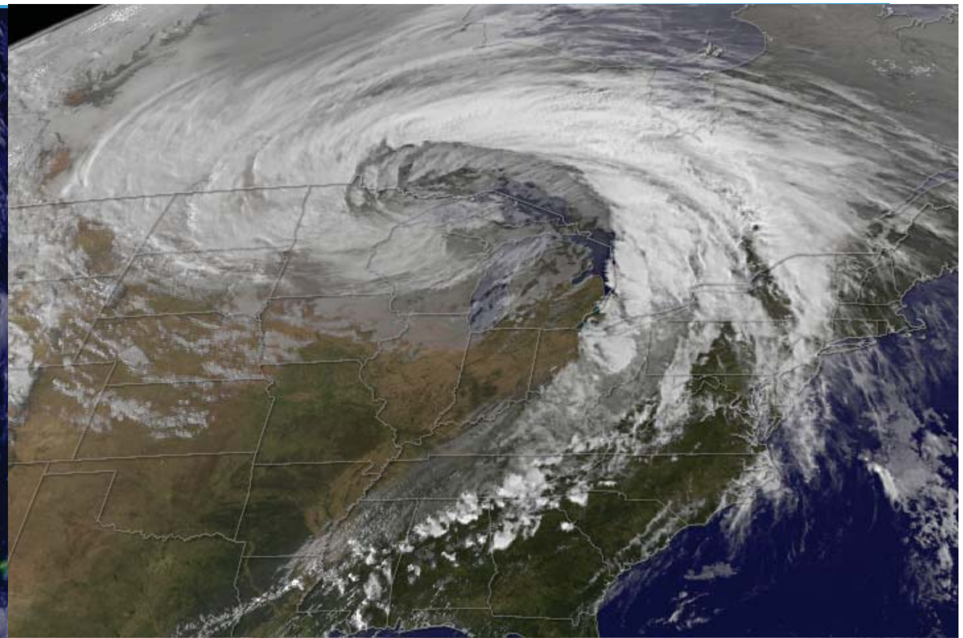

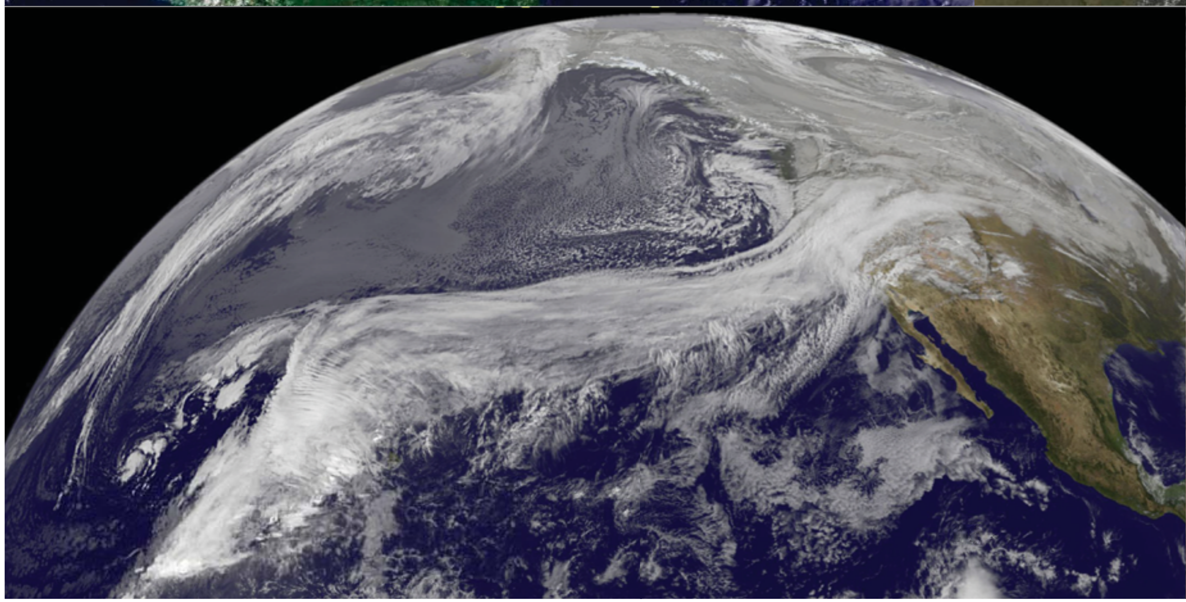

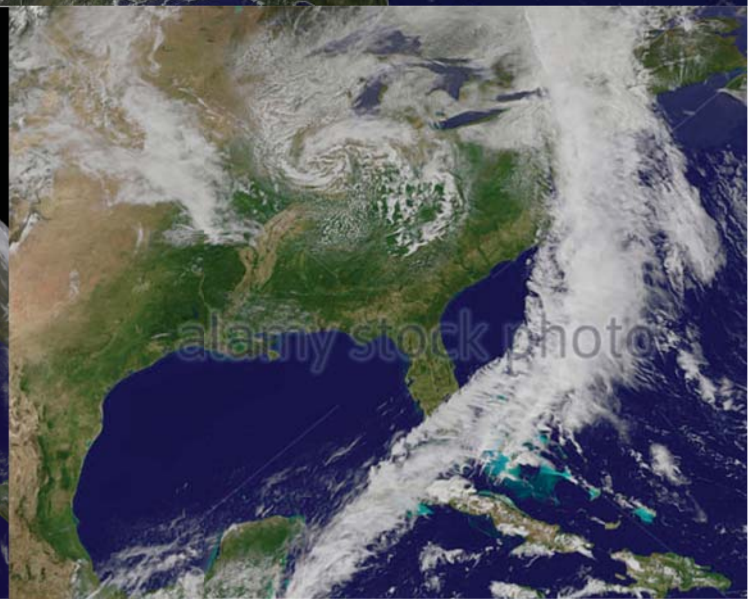

Climate Science Tasks

Classification	Classification + Localization	Object Detection	Instance Segmentation

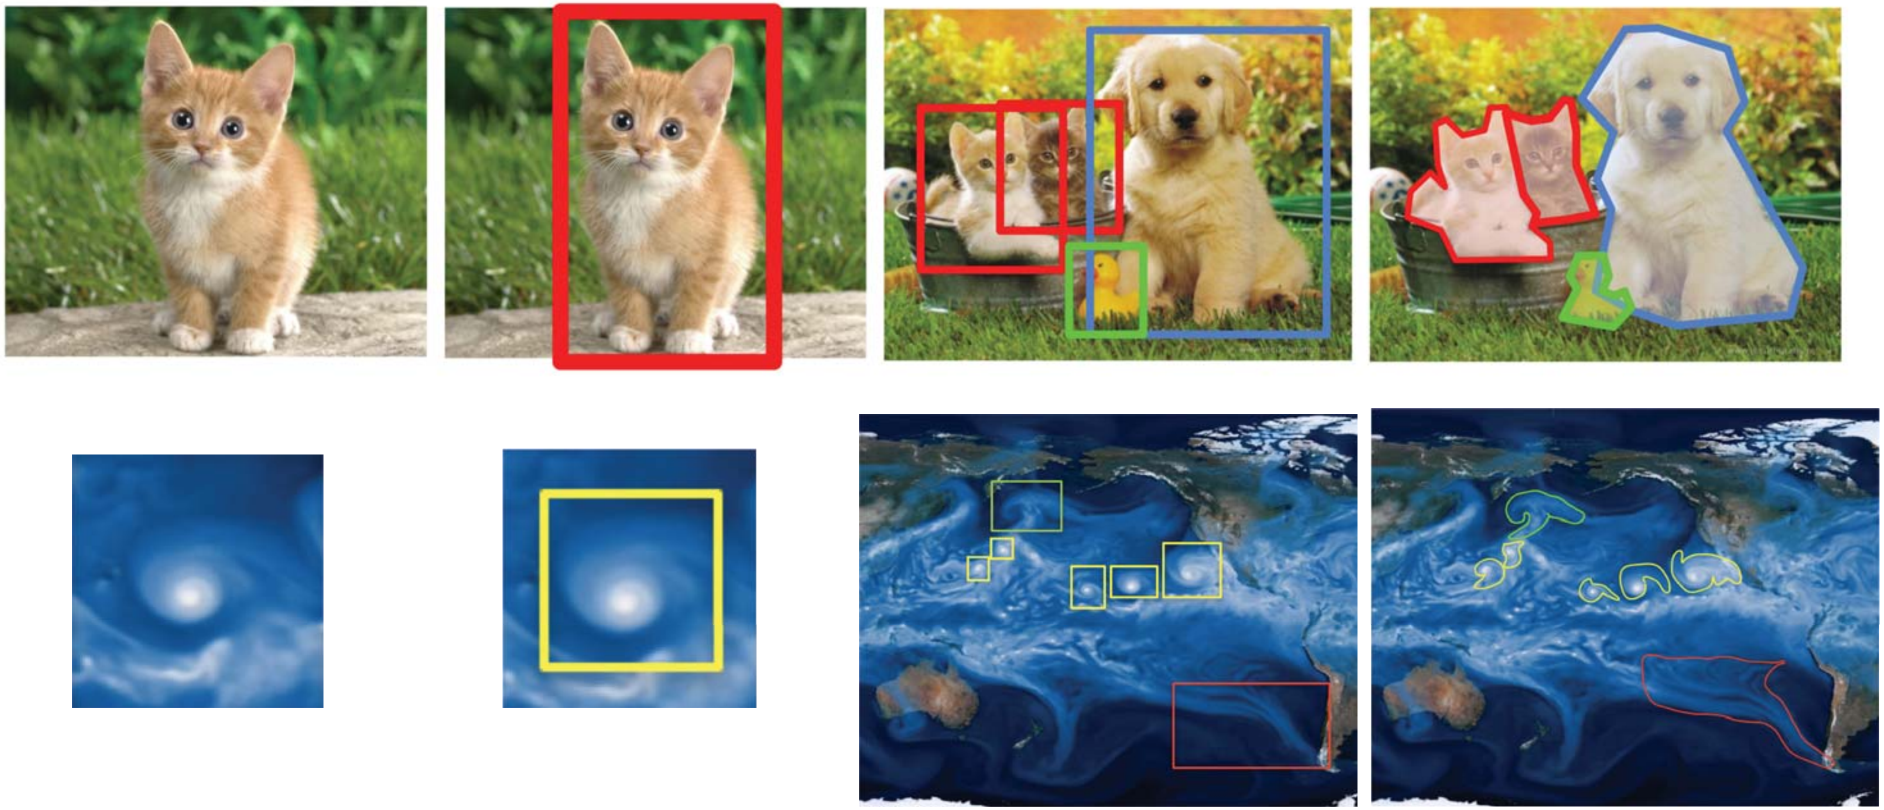

Semi-Supervised Convolutional Architecture

Encoder **Decoder**

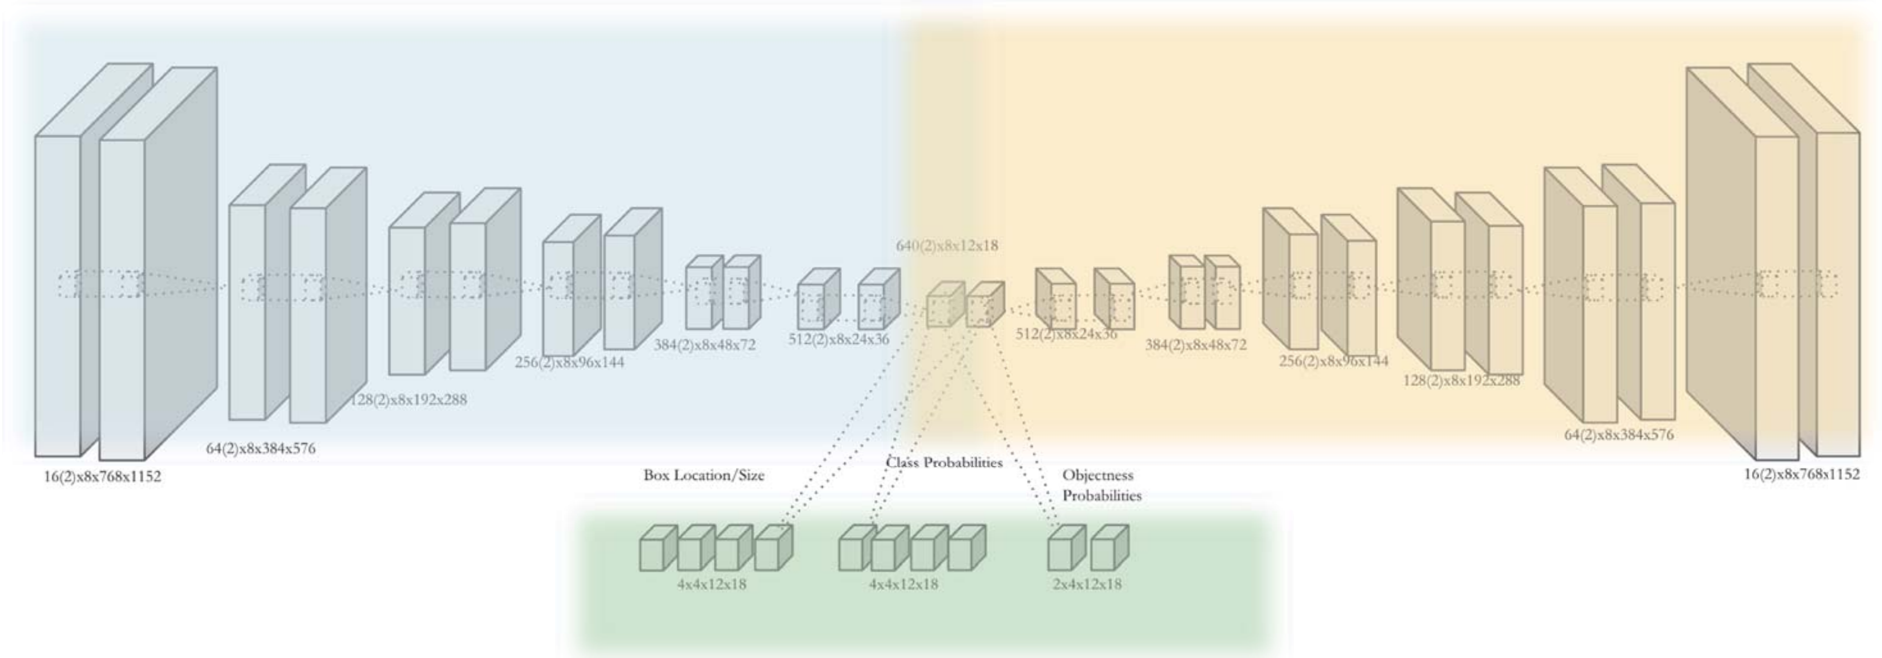

Classification + Bounding Box Regression

Contributors: Evan Racah, Chris Pal, Chris Beckham, Samira Kehou, Tegan Maharaj (MILA)

DL Software

	Technologies
Deep Learning Frameworks	K, TensorFlow, Caffe, PYTORCH, CNTK, MXNet, Neon, …
Multi Node libraries	Cray Plugin, Horovod, GRPC, MLSL, MPI
Single Node libraries	MKL-DNN, CuDNN
Hardware	CPUs (KNL), GPUs, FPGAs, Accelerators

Single Node Performance

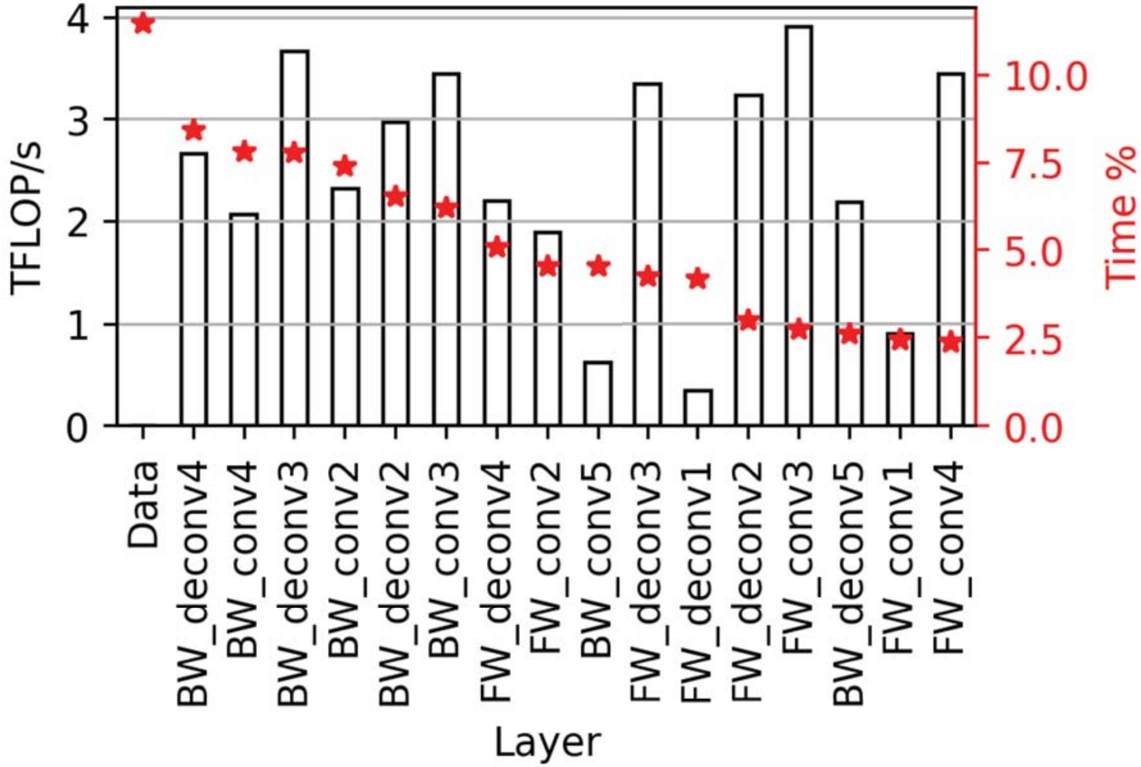

- **Optimizations for conv and deconv layers in MKL-DNN**
- **Obtained 2.09 TF; Theoretical max ~6 TF**

Multi-Node Strategy

- **Data/Batch Parallelism**

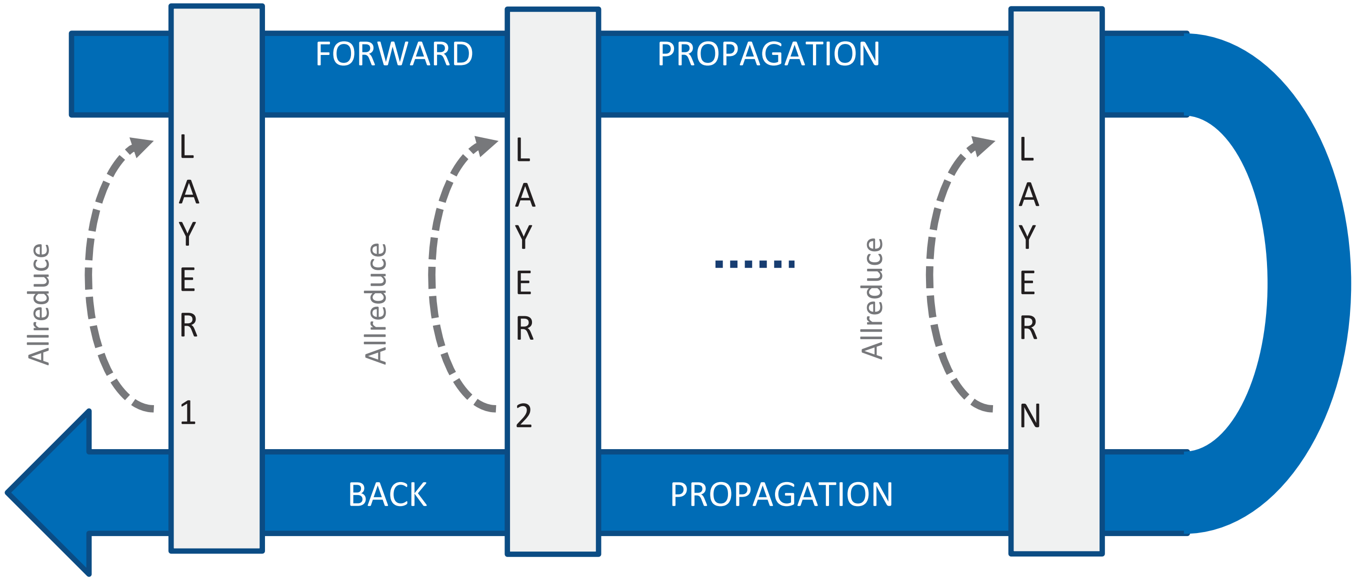

From Pradeep Dubey, "Scaling to Meet the Growing Needs of Artificial Intelligence (AI), IDF 2016
https://software.intel.com/en-us/articles/scaling-to-meet-the-growing-needs-of-ai

Multi-Node Strategy

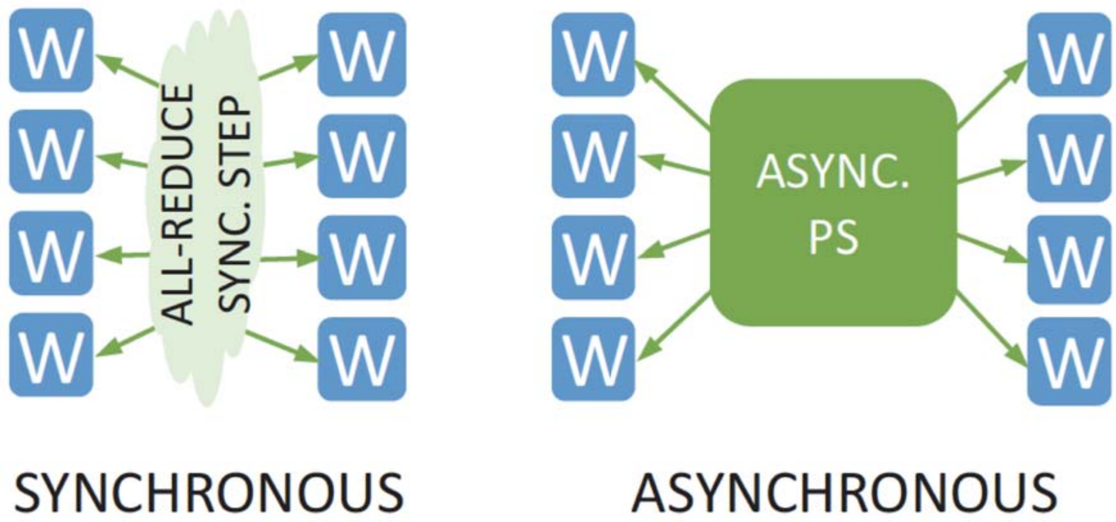

	SYNCHRONOUS	ASYNCHRONOUS
Pros	• **Stable convergence** • **Same # iterations to converge as serial implementation**	• **Faster Iterations** • **Robustness to node failures** • **Better control of batch size**
Cons	• Straggler effect • Susceptible to node failure • Effective Batch size grows with # nodes	• Parameter Server can be bottleneck • Stale gradients can negatively impact convergence rate

Hybrid Synchronization

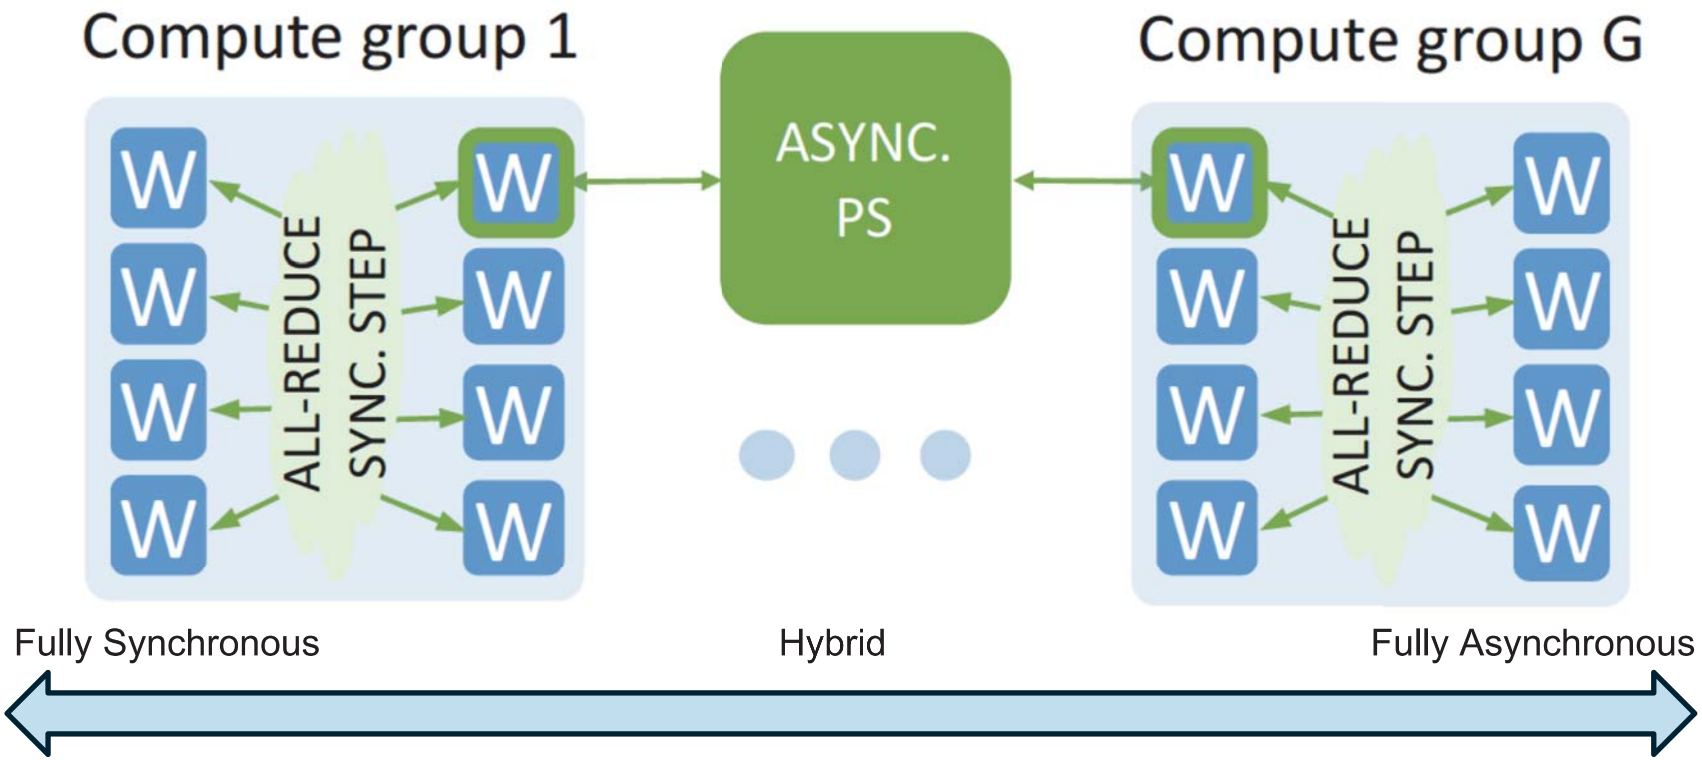

Pros	• **Reduced Impact of Stragglers** • **Finer control of batch size**
Cons	• Group size needs to be tuned

Scaling Results

Strong Scaling

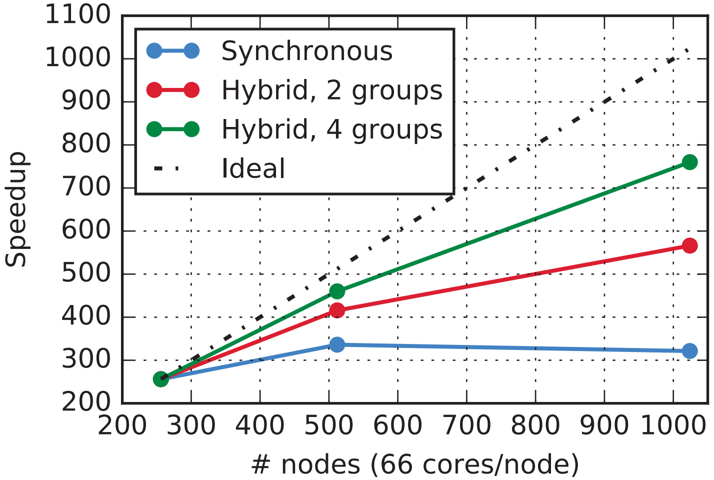

- **Batch size = 2048 per group**

Weak Scaling

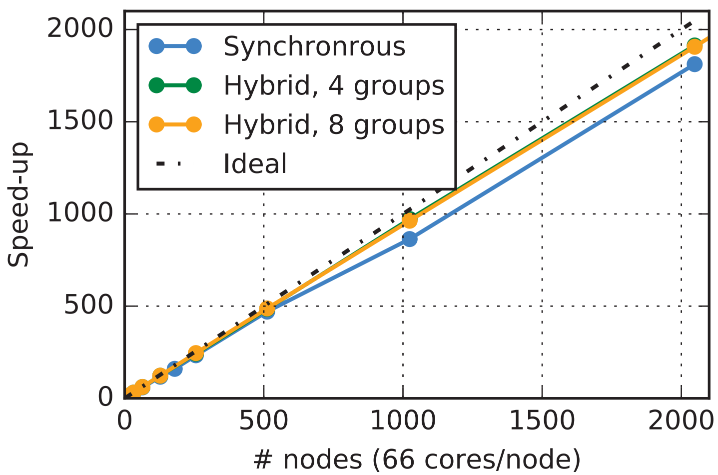

- **Batch size = 8 per node**

Final Results

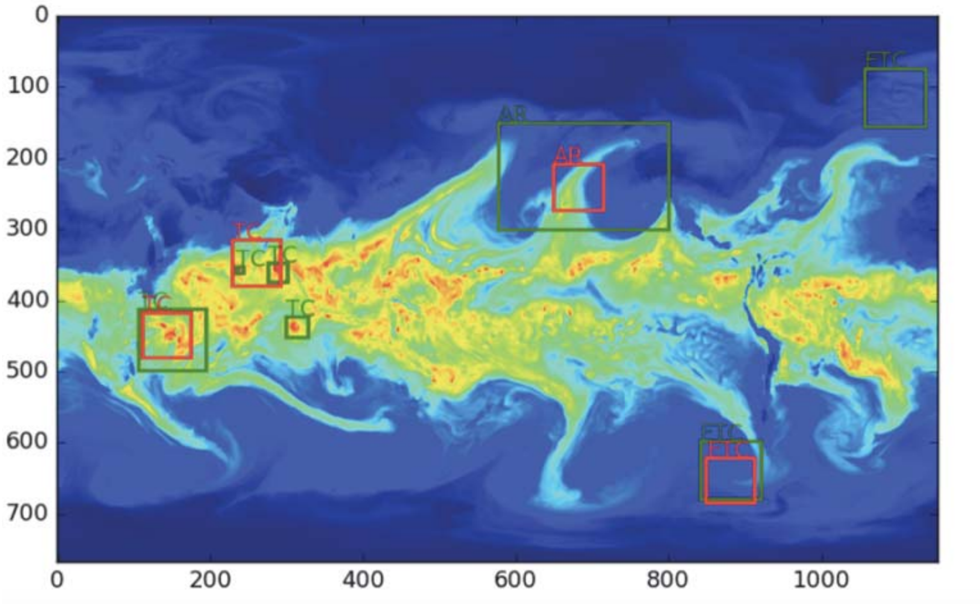

Ground Truth
Prediction

- Reasonable detection results for Climate App
- *Largest Caffe run on CPU-based HPC system*
  – 13.3 PF sustained, 15.1 PF peak on 9600 nodes
- Performance enhancements to MLSL, Intel Caffe and MKL-DNN released to the broader community

CosmoFlow @3.5 PF

- **ACM/IEEE Supercomputing 2018 Paper**

https://arxiv.org/abs/1808.04728

CosmoFlow Team

Amrita Mathuriya
Intel

Debbie Bard
NERSC

Pete Mendygral
Cray

Lawrence Meadows
Intel

James Arnemann
UC Berkeley

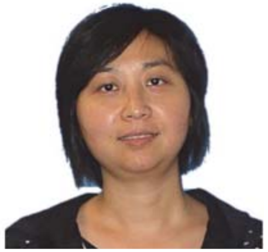

Lei Shao
Intel

Siyu He
LBNL/CMU

Tuomas Karna
Intel

Diana Moise
Cray

Simon Pennycook
Intel

Kristyn Maschhoff
Cray

Nalini Kumar
Intel

Shirley Ho
LBNL/CMU

Mike Ringenburg
Cray

Prabhat
NERSC

Victor Lee
Intel

Determining the Fundamental Constants of Cosmology

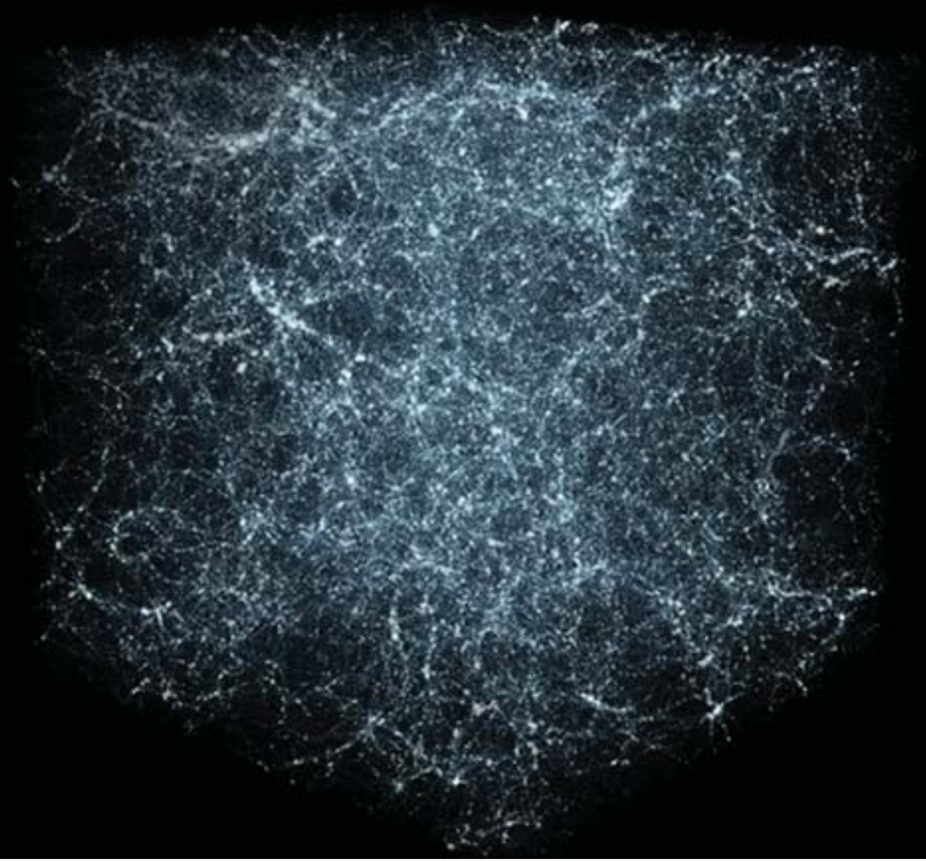

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

3D Convolutional Network

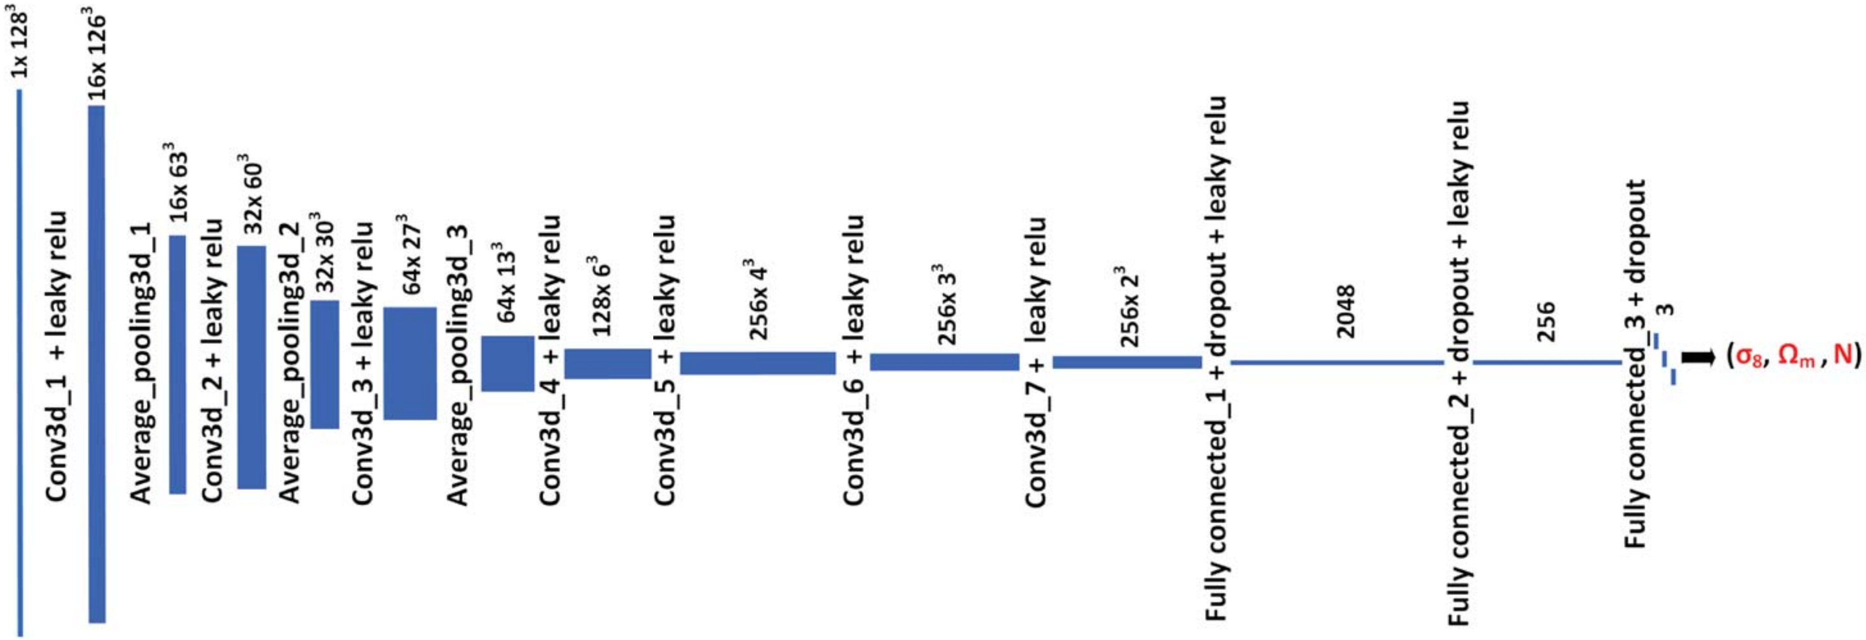

- **Based on design in Ravanbaksh et al. (PMLR'16)**
- **7 convolution layers followed by pooling layers**
- **3 fully-connected layers**
- **All layers use leaky Relu as activation function**

DL Software

	Technologies
Deep Learning Frameworks	TensorFlow, Caffe, PyTorch, K; CNTK, MXNet, Neon, …
Multi Node libraries	Cray Plugin, Horovod, GRPC, MLSL, MPI
Single Node libraries	MKL-DNN, CuDNN
Hardware	CPUs (KNL), GPUs, FPGAs, Accelerators

Single-node Performance

- **Optimized 3D convolutions and pooling in MKL-DNN**
  - Larger convolutions achieve >1 TF
- **Overall 535 GF performance on a single KNL node**
  - Includes I/O and the Cray Plugin overhead

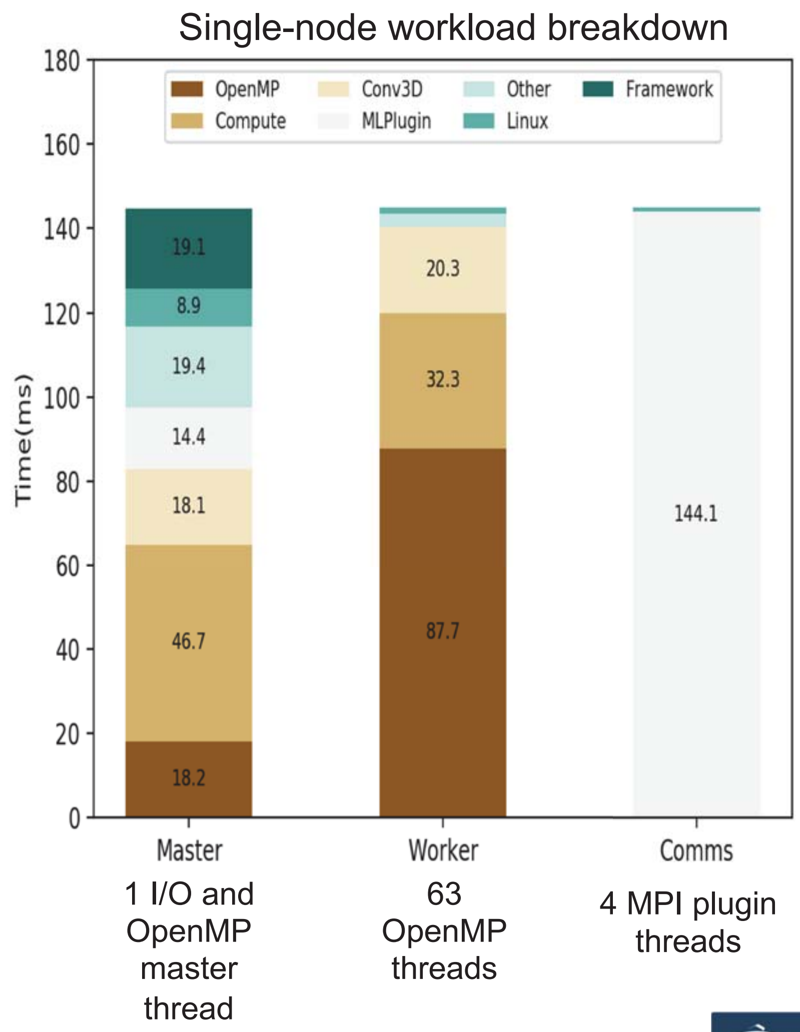

Multi-node Optimizations

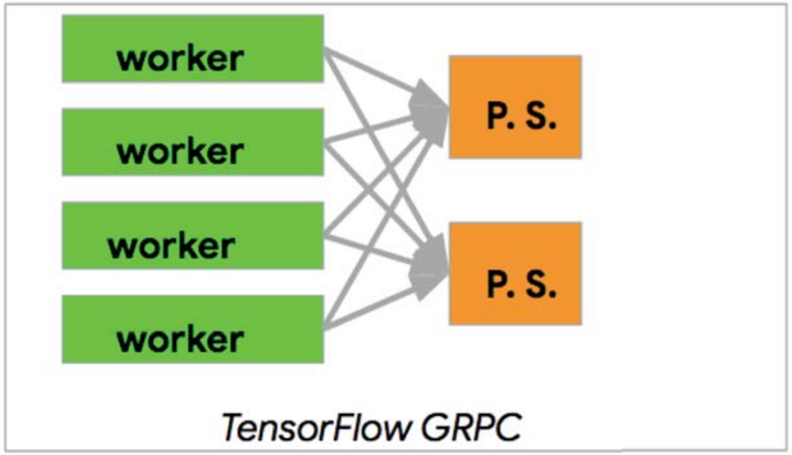

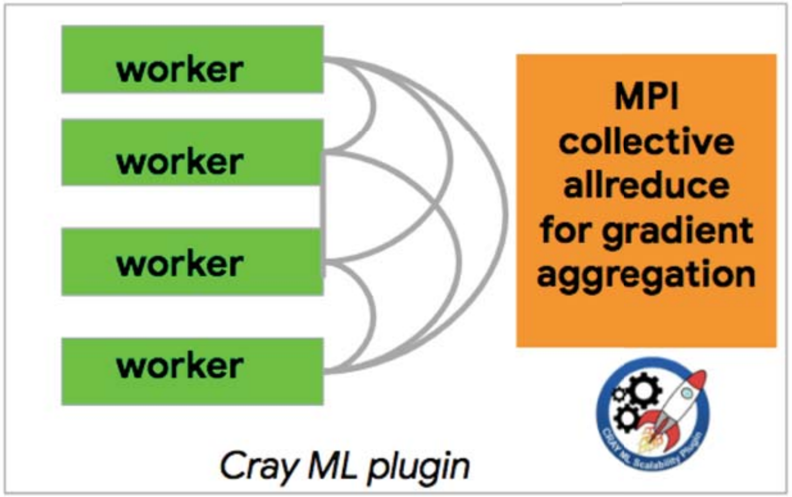

- **Data Parallel, synchronous updates**
- **Skip parameter servers, update worker nodes directly**
- **Cray CPE ML Plugin**
  - MPI-based, framework-independent plugin
  - No modification to Tensorflow, python loopback calls
  - Dedicated thread pool and custom reduction algorithm

Scaling Performance

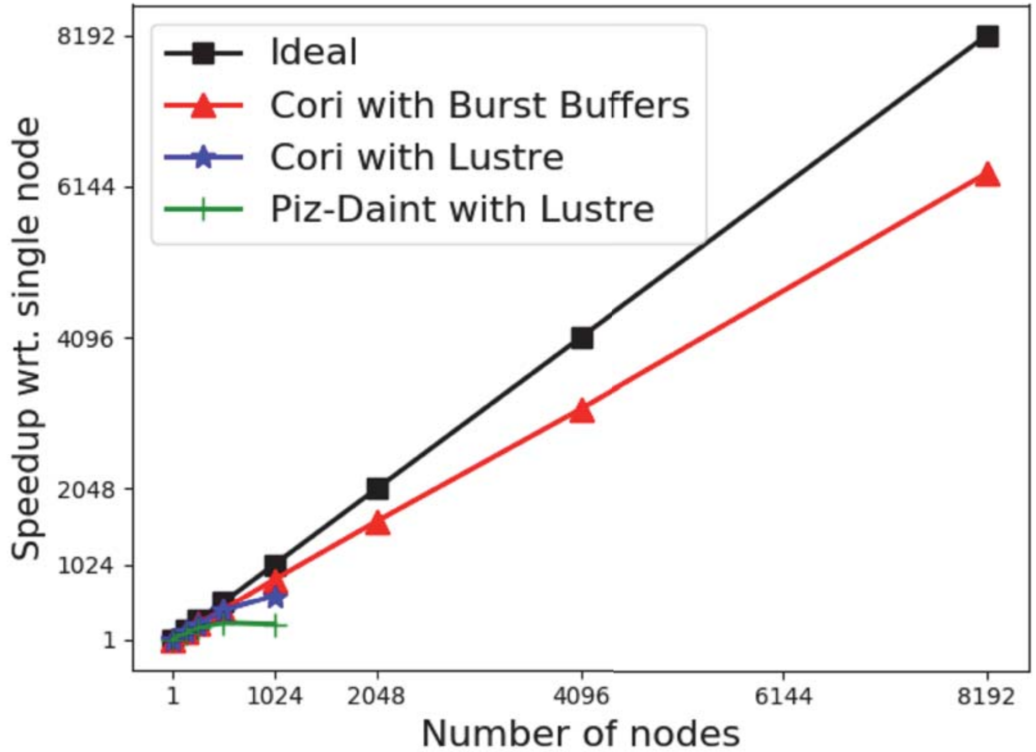

- **Measure overall walltime per epoch (throughput)**
- **Achieve 77% scaling efficiency at 8192 nodes**
  - Global batch size = 8192, weak scaling

Final Results

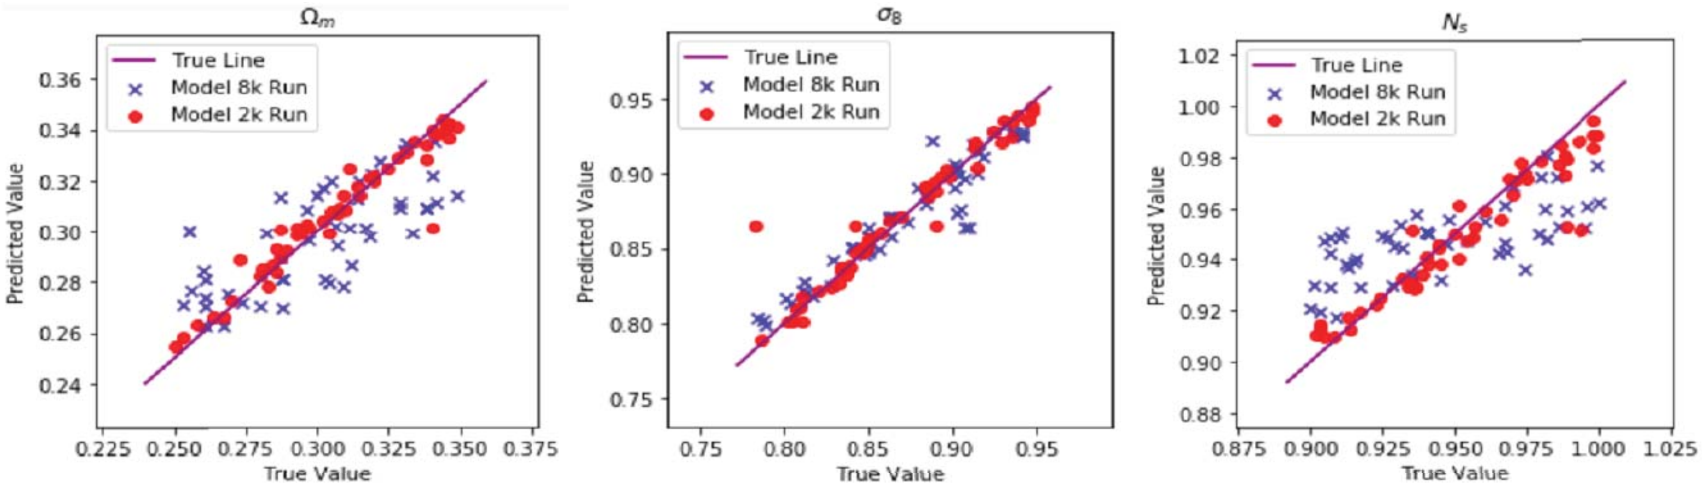

- Parameter estimation comparable to best experimental uncertainty for Ω_m and σ_8, almost 5x smaller for N_s
- CosmoFlow scaled out to 8192 nodes; 77% scaling efficiency, 3.5PF sustained performance
- *Largest application of TensorFlow on CPU-based system* with fully-synchronous updates
- Cray plugin and MKL-DNN enhancements deployed in production

Deep Learning @ 1EF

- ACM/IEEE Supercomputing 2018 Gordon Bell Finalist

DL@1EF Team

Thorsten Kurth

Sean Treichler

Joshua Romero

Mayur Mudigonda

Nathan Luehr

Everett Phillips

Ankur Mahesh

Michael Matheson

Jack Deslippe

Massimiliano Fatica

Prabhat

Michael Houston

Climate Science Tasks

Classification	Classification + Localization	Object Detection	Instance Segmentation

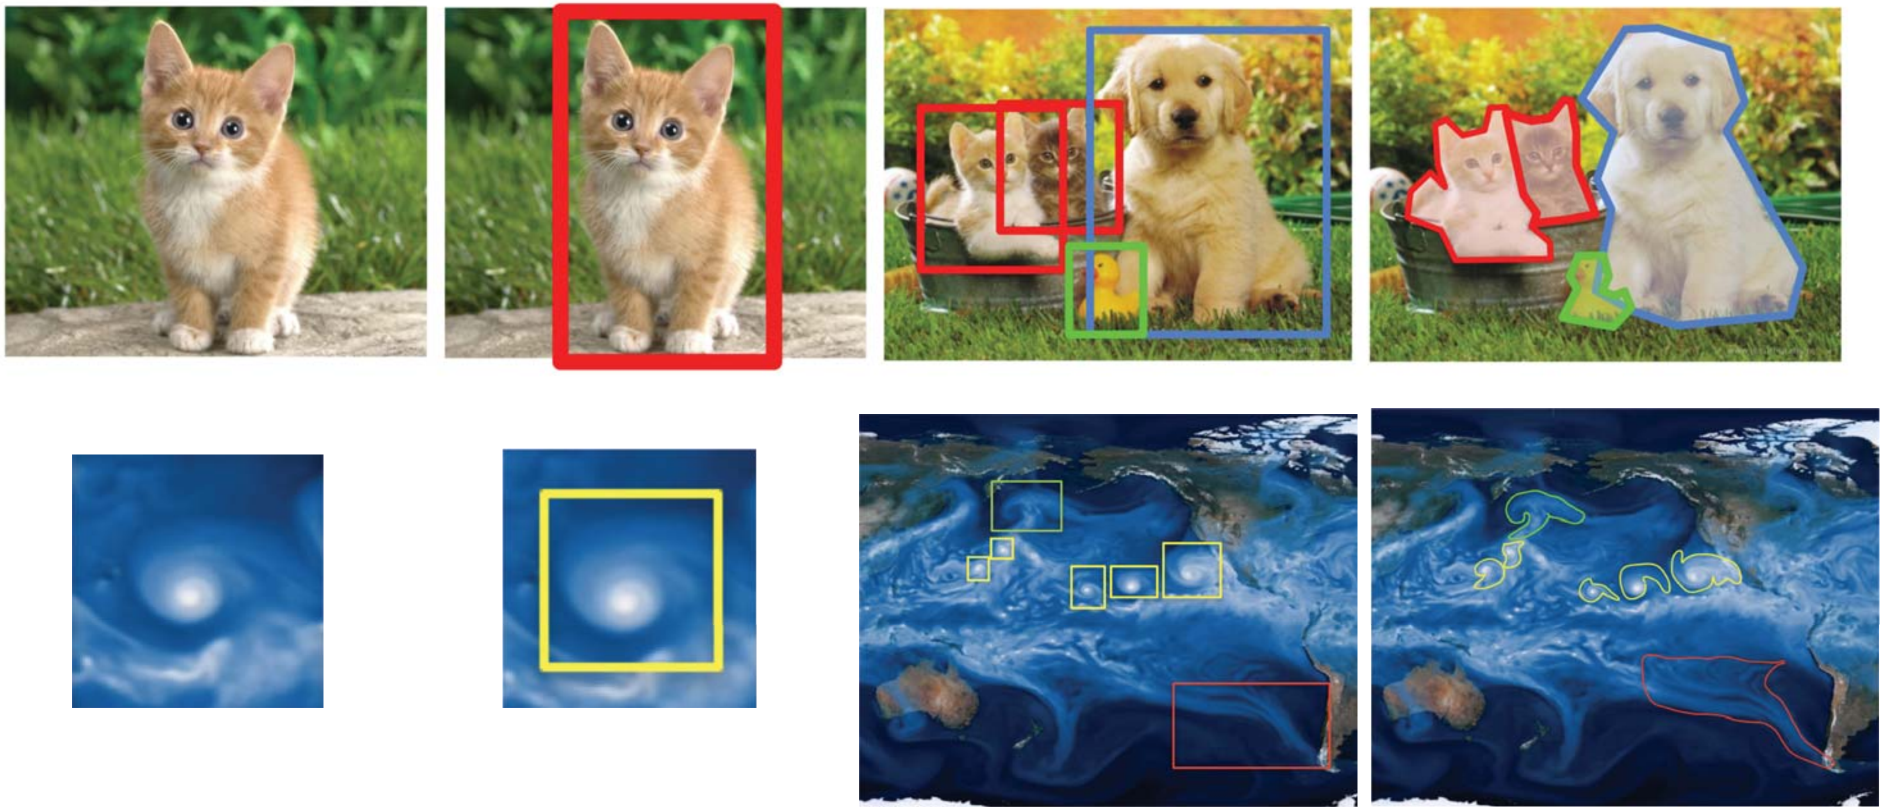

DeepLabv3+ Segmentation Architecture

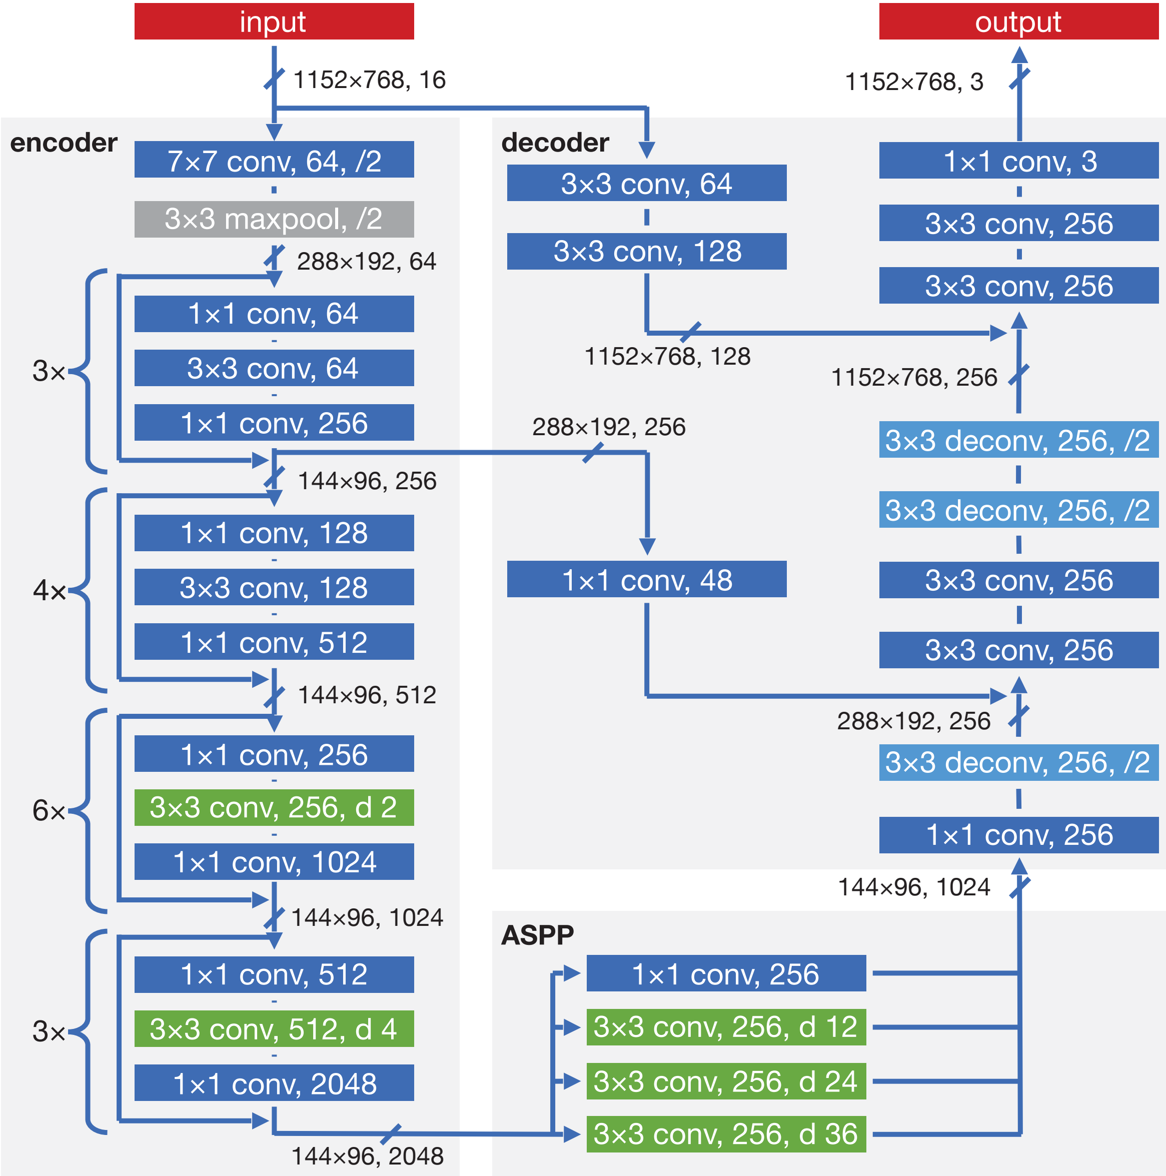

DL Software

	Technologies
Deep Learning Frameworks	K, TensorFlow, Caffe, PYTORCH, CNTK, MXNet, Neon, …
Multi Node libraries	Cray Plugin, Horovod, GRPC, MLSL, MPI, NCCL
Single Node libraries	MKL-DNN, CuDNN
Hardware	CPUs (KNL), GPUs, FPGAs, Accelerators

Single Node Performance

Category		FP32 Training							FP16 Training						
		# Kern	Time (ms)	Math (TF)	Mem (GB)	% Time	% Math	% Mem	# Kern	Time (ms)	Math (TF)	Mem (GB)	% Time	% Math	% Mem
Forward	Convolutions	71	172.4	1.40	100.0	31.4	51.7	64.4	95	105.5	2.79	96.1	25.3	21.2	101.2
	Point-wise	563	43.6	< 0.1	32.2	7.9		82.1	564	51.1	< 0.1	35.3	12.2		76.8
Backward	Convolutions	95	270.5	2.79	153.2	49.2	65.7	62.9	113	159.7	5.58	95.8	38.3	28.0	66.7
	Point-wise	113	4.1	< 0.1	2.2	0.7		59.6	123	11.6	< 0.1	5.0	2.8		47.9
Optimizer		1056	3.0	< 0.1	0.7	0.5		25.9	1056	3.0	< 0.1	0.9	0.7		33.3
Copies / Transposes		388	30.5	-	19.8	5.5		78.0	530	51.5	-	28.2	12.3		60.8
Allreduce (NCCL)		25	28.2	< 0.1	0.4	5.1		1.6	30	22.4	< 0.1	0.7	5.4		3.5
Type Conversions									143	0.5	-	0.1	0.1		22.2
GPU Idle										12.0			2.9		
Total		2311	549.9	4.19	308.5		48.5	62.3	2654	417.3	8.38	262.1		16.1	69.8

- 16-bit and 32-bit implementations
- Obtained 39 TF in 16-bit; Theoretical Max: 125 TF

Multi-Node Optimization

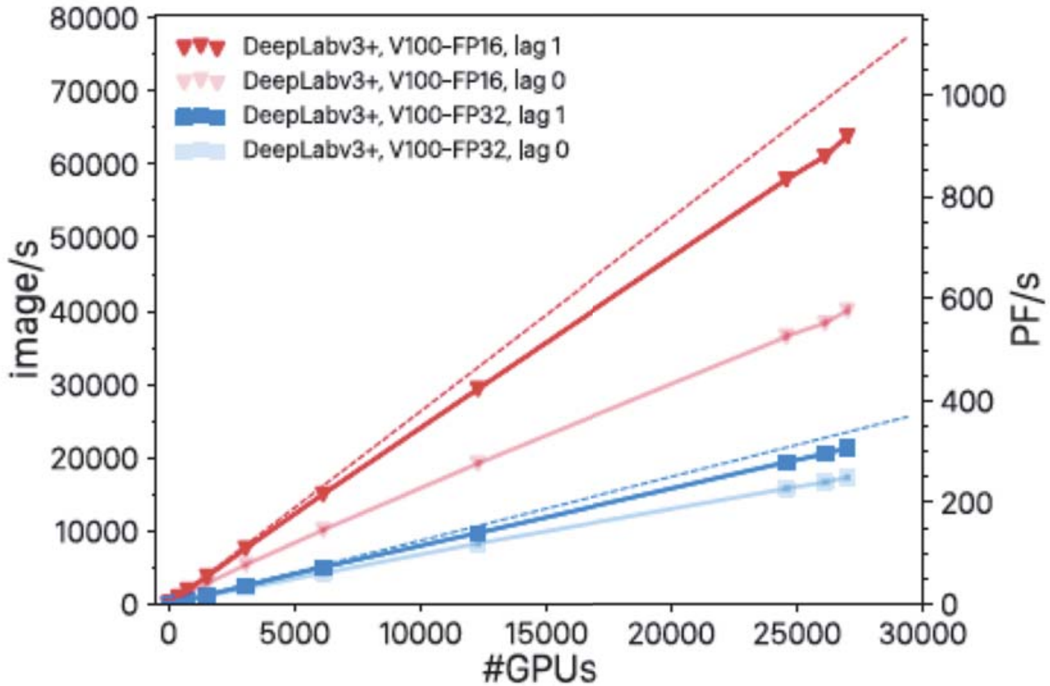

- **Data Parallel, synchronous and lagged updates**
- **Hierarchical scheme for local updates (over NCCL/NVLink) and global updates (over MPI/Infiniband)**

Final Results

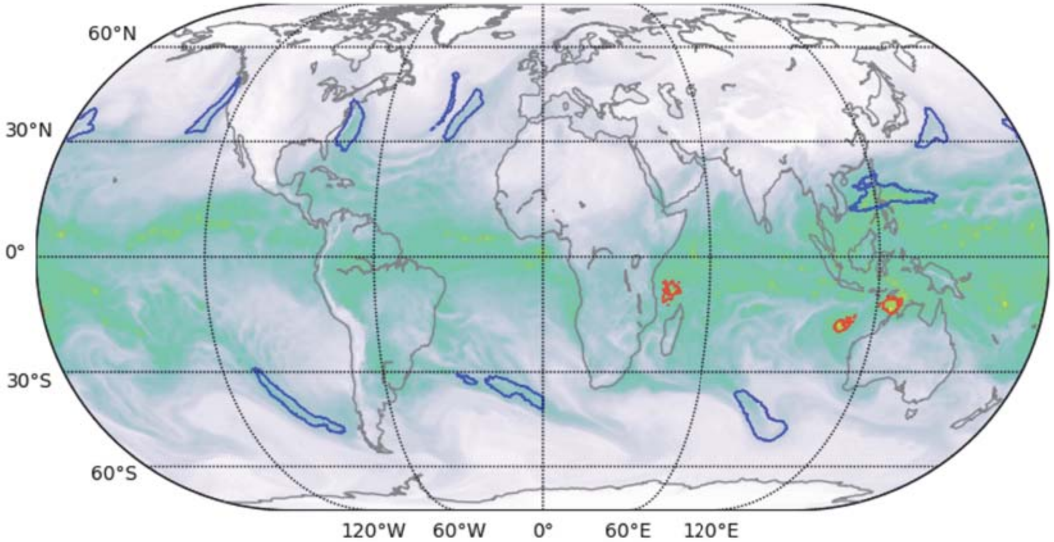

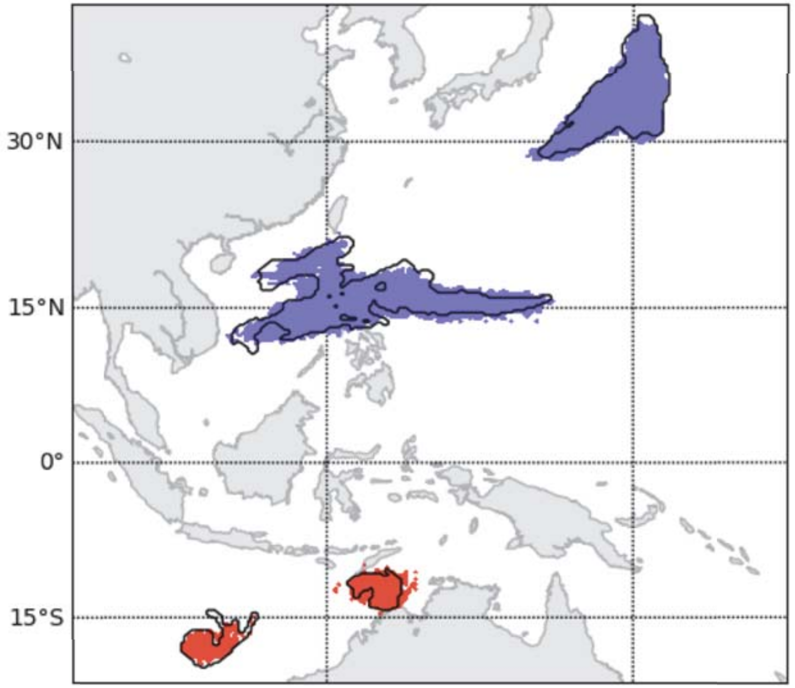

- High quality segmentation results obtained for climate data
- Network scaled out to 4560 Summit nodes (27,360 Volta GPUs)
- 1.13 EF peak, 0.999 EF sustained performance in 16-bit precision
- *Largest application of TensorFlow on GPU-based system, first exascale Deep Learning application*
- TensorFlow and Horovod enhancements deployed in production

Convergence…

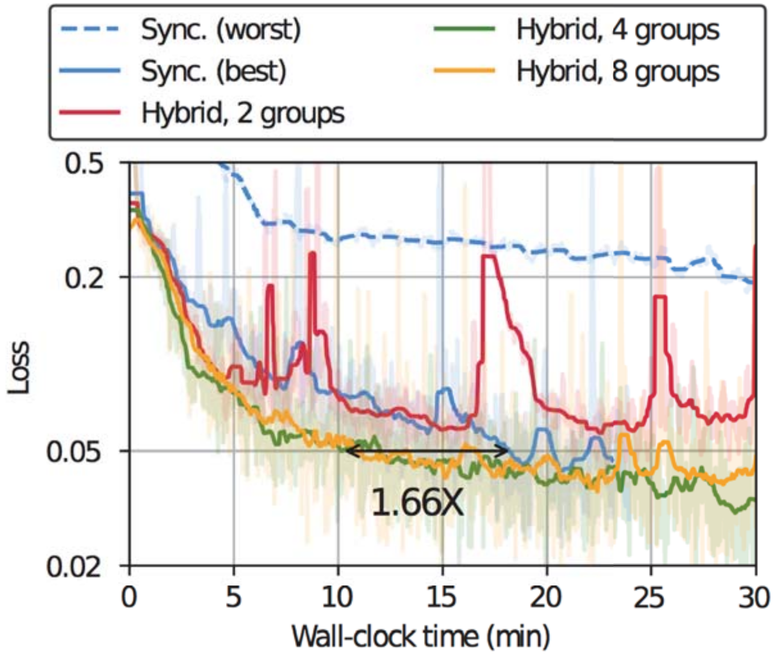

DL @ 15PF

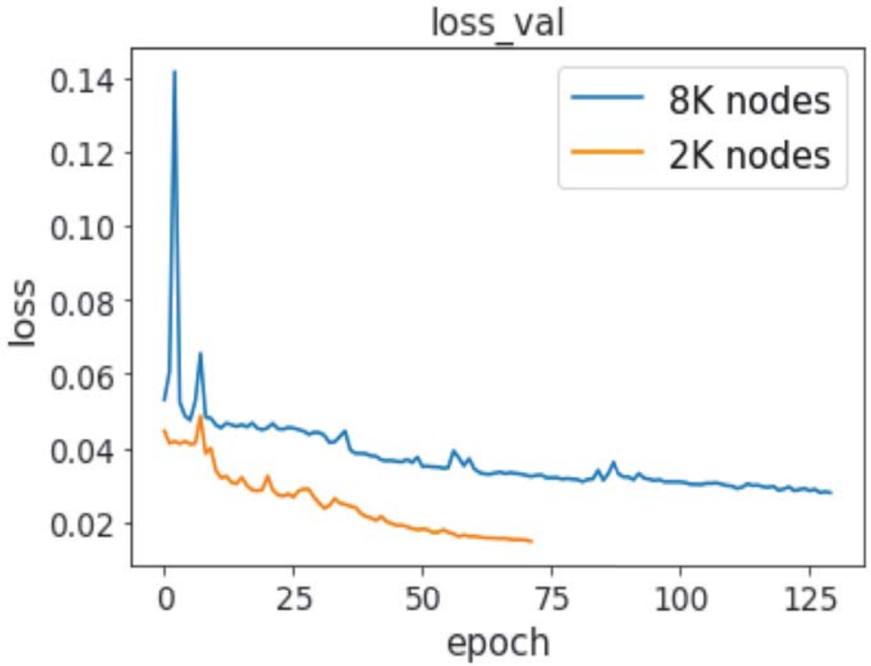

CosmoFlow

Convergence…

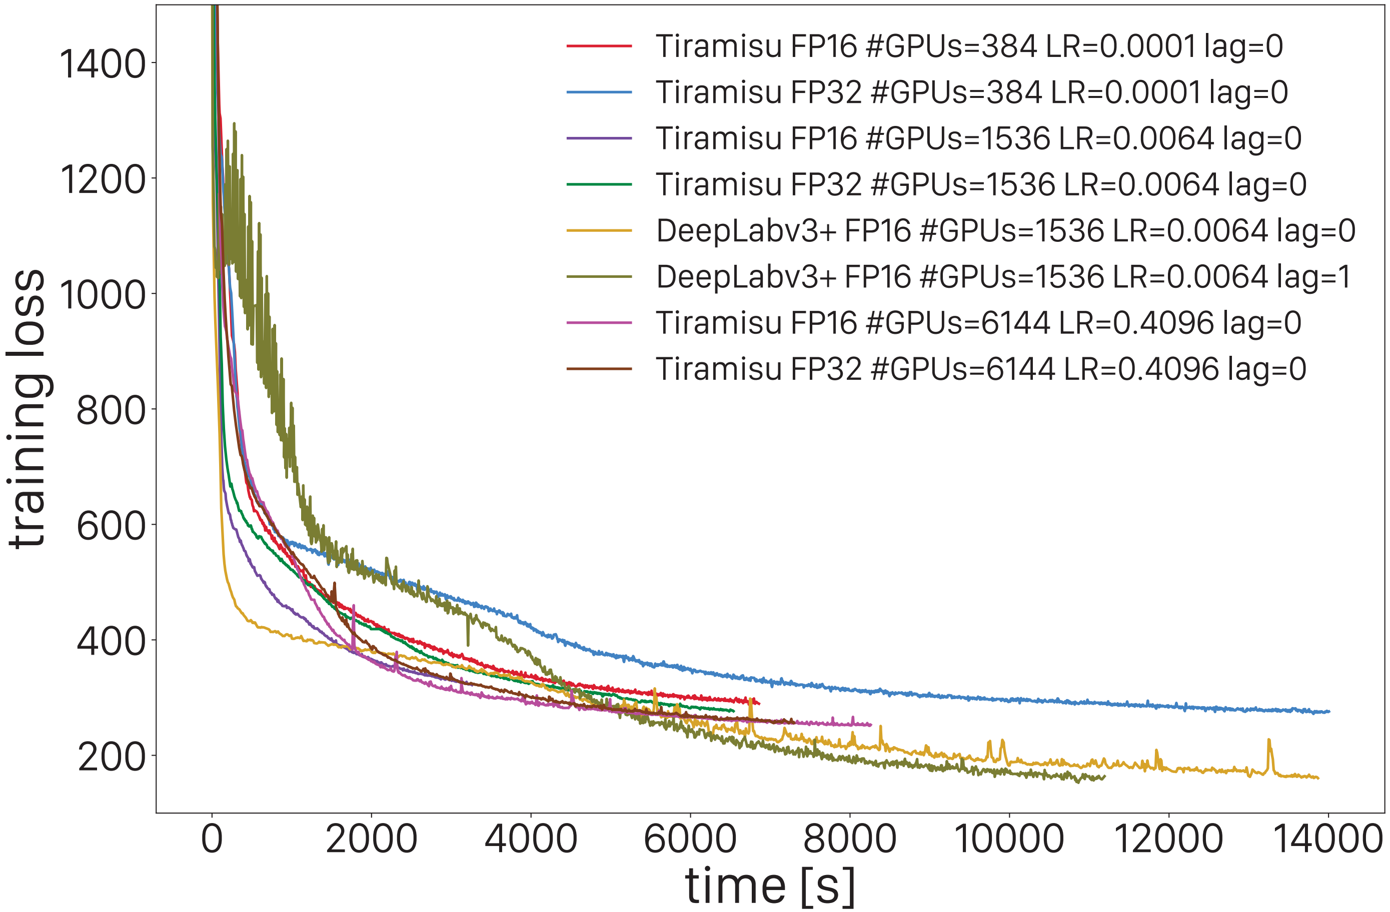

DL @ 1EF

Outline

- Introduction to NERSC
  - Why Scale Deep Learning?
- Case Studies
- **Open Challenges**
  - Data Management
  - Hyper-Parameter Tuning
  - Convergence
- Conclusions

Data Management at Scale

- **DL workloads are extremely demanding (Data and Metadata)**
  - Read-only, Random shuffles vs. Contiguous reads/writes
  - O(10) of TBs spread across O(100,000) files
- **Lustre and GPFS typically can't keep up**
  - Burst Buffer and node-local NVMe storage was critical
- **Ingest pipelines for loading scientific data (HDF5, NetCDF, ROOT) into DL frameworks are not optimized**
  - multi-threaded support
- **I/O middleware for modern DL workloads might need to be redesigned**

https://www.nextplatform.com/2018/10/09/hpc-file-systems-fail-for-deep-learning-at-scale/

Hyper-Parameter Tuning at Scale

- **Tuning learning rate + curriculum at scale is an empirical exercise**
  - smaller scale experiments need to inform large scale runs, need theory
  - linear / square-root scaling for learning rate followed by decay
- **Choice of optimizer (SGD + Momentum, ADAM, LARS/LARC) is important**

Convergence at Scale

- **Computational Efficiency vs. Statistical Convergence**
  - Time to Solution will depend on both
- **Parameter update schemes: Sync vs. Async, Lagged**
- **Effect of Batch Size on convergence**
  - Start with low batch size and ramp up
- **Higher-order optimization schemes**

Conclusions

- **Deep Learning is applicable for scientific problems**
- **Large datasets and complex architectures require performance and scaling**
  - HPC systems are a good match
- **Success in scaling DL to largest CPU and GPU-based HPC systems (15PF, 3.5PF, 1EF) with productive frameworks**
- **Open Challenges**
  - Data Management, Hyper-Parameter Tuning, Statistical Convergence
- **Exciting area, open to collaboration!**

Questions?

prabhat@lbl.gov