

CosmoFlow: Using Deep Learning to Learn the Universe at Scale

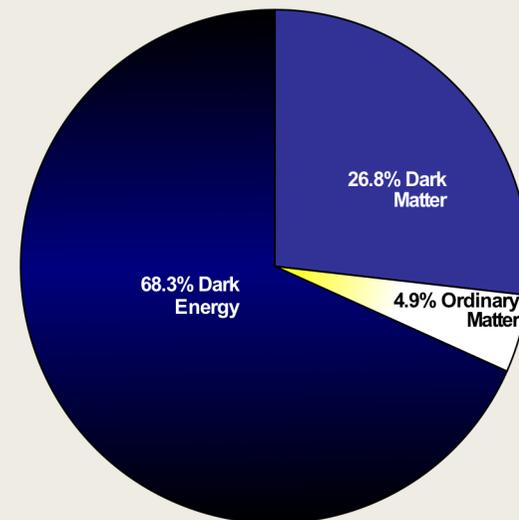
Amrita Mathuriya*, Deborah Bard†, Peter Mendygra‡, Lawrence Meadows*, James Arnemann§, Lei Shao¶, Siyu He**†, Tuomas Kärnä*, Diana Moise‡, Simon J. Pennycook¶, Kristyn Maschhoff‡, Jason Sewall¶, Nalini Kumar¶, Shirley Ho**†, Michael F. Ringenburt‡, Prabhat† and Victor Lee¶

*Intel OR, †LBNL, ‡Cray, §U.C. Berkeley, ¶Intel CA, †Carnegie Mellon University, **Flatiron Institute

presented by Mateusz Bysiek, School of Computing, Tokyo Institute of Technology
in HPC2018 class, 20 December 2018

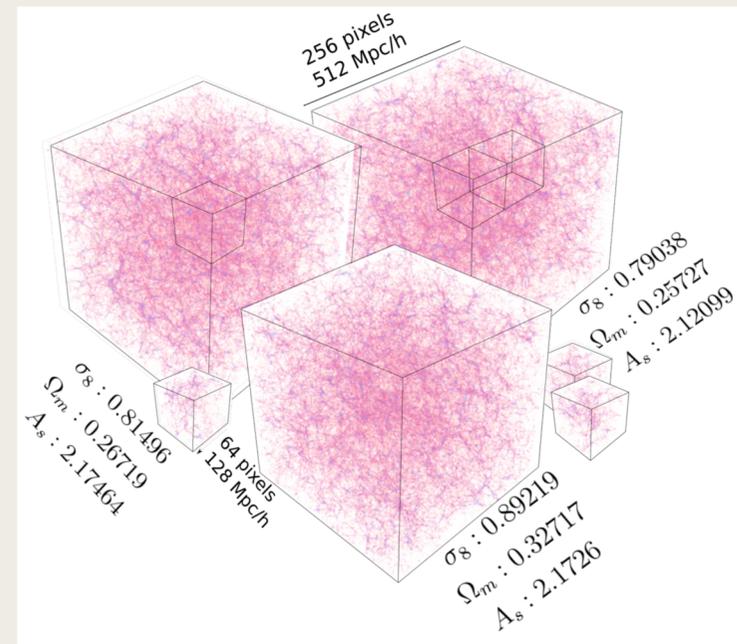
What is the nature of dark energy?

- Unknown force that drives accelerated expansion of the universe
- Impossible to measure directly – only through effects on observable universe
 - *Distribution of matter is an effect of interplay of gravity and dark energy*



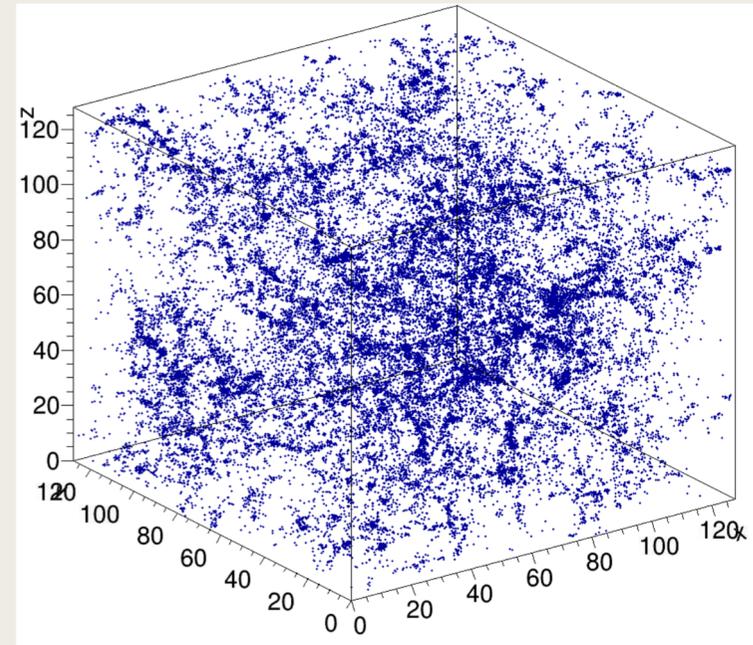
Why use deep learning for cosmology?

- Matter distribution is typically characterised using *few* selected features
- Deep learning networks could capture *all* features
- Simulation-generated matter distributions provide training data



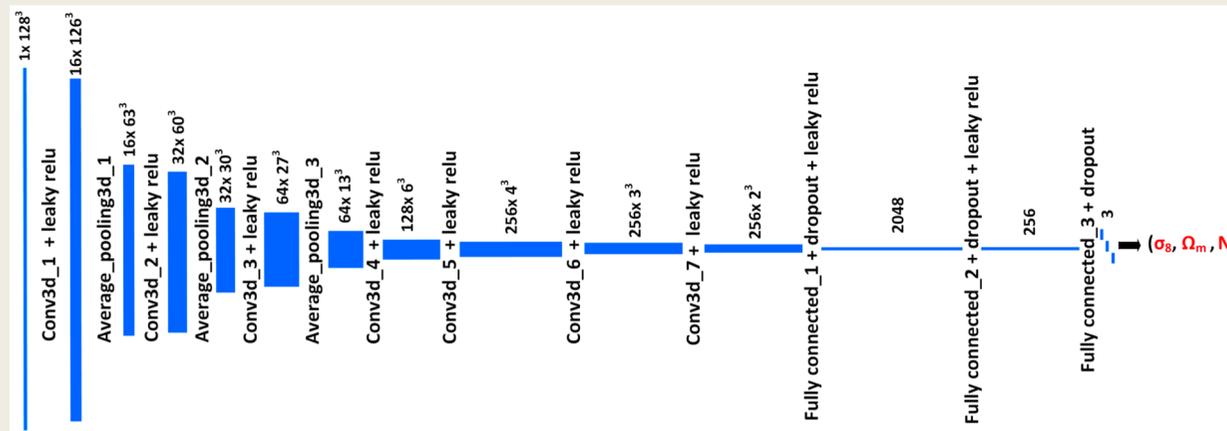
Challenges for deep learning in science

- Scientific data often complex (3+ dimensions and many channels)
- Measured in tera- or petabytes
- Efficient processing at scale essential for relevance of deep learning in science



Modifications in CosmoFlow

- Adapt existing deep learning network to scalable architecture
- Extra convolution layer and average pooling layer (near input) to cope with input size increase
- State of the art work predicted 2 parameters at problem size 64^3 – it did not scale
- Predict 3 cosmological parameters at problem size of 128^3



The predicted cosmological parameters

- Ω_M – Matter density parameter, i.e. proportion of matter in the universe
- σ_8 – The present root-mean-square matter fluctuation averages, i.e. amplitude of mass fluctuations in the universe, over a sphere of radius $8h^{-1}$ Mpc
- n_s – scalar spectral index of the spatial curvature of a comoving slicing of space-time, i.e. how density fluctuations vary with scale

Modifications in CosmoFlow, contd.

- Use batch size of 1 per MPI rank
- Adam Optimizer
- Polynomial learning rate decay schedule

Optimizations: node-level

- Change number of output channels to aid vectorization
- Arrays are blocked by channels in forward propagation
- 3 innermost loops are completely unrolled and vectorized with AVX512
- Similar strategy for backward propagation

```
Require:  $SRC \in \mathbb{R}^{ICB \times ID \times IH \times IW \times 16}$   
Require:  $DST \in \mathbb{R}^{OCB \times OD \times OH \times OW \times 16}$   
Require:  $W \in \mathbb{R}^{OCB \times ICB \times KD \times KH \times KW \times 16 \times 16}$   
1: for  $ocb = 1 \dots OCB$  do // Output channel block  
2: for  $icb = 1 \dots ICB$  do // Input channel block  
3: for  $od = 1 \dots OD$  do // Output depth  
4: for  $oh = 1 \dots OH$  do // Output height  
5: for  $owb = 1 \dots OWB$  do // Output width block  
6:  $d \leftarrow DST[ocb, od, oh, 28owb, 0]$   
7: for  $kd = 1 \dots KD$  do // Kernel depth  
8: for  $kh = 1 \dots KH$  do // Kernel height  
9: for  $kw = 1 \dots KW$  do // Kernel width  
10:  $s \leftarrow SRC[icb, od + kd, oh + kh, 28owb + kw, 0]$   
11:  $w \leftarrow W[ocb, icb, kd, kh, kw, 0, 0]$   
12: for  $ow = 1 \dots 28$  do // Output width  
13: for  $oc = 1 \dots 16$  do // Output channel  
14: for  $ic = 1 \dots 16$  do // Input channel  
15:  $d[16ow + oc] \leftarrow w[16ic + oc]s[16ow + ic]$ 
```

Optimizations: scaling

- Burst buffer: CPE ML plugin
- Reduce straggler effect of Synchronous Stochastic Gradient Descent – hide timing imbalances with non-blocking MPI
- No central/unique master node
- batch size of 1 per MPI rank
Remove batch-norm layers to aid scaling

Require: N = total number of epochs

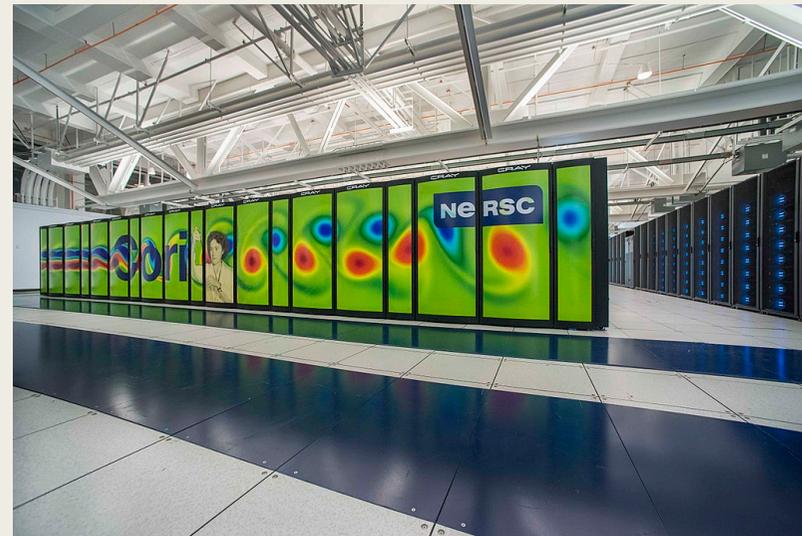
Require: n = total number of training samples

Require: k = number of MPI ranks

```
1: for epoch = 1  $\dots$   $N$  do  
2:   for step = 1  $\dots$   $n/k$  do  
3:      $g_{step} \leftarrow \text{compute\_gradients}(\text{local\_batch}_{step})$   
4:      $G_{step} \leftarrow \text{mc.gradients}(g_{step})$   
5:      $loss_{step} \leftarrow \text{apply\_gradients}(G_{step})$ 
```

Cori: Cray XC40 at NERSC

- 2004 nodes with Intel Xeon Phi E5-2698 v3
9688 nodes with Intel Xeon Phi 7250
 - 68 cores, 16GB memory
 - 32/32KB instr./data L1 cache
 - 2D mesh network
- 96GB of DDR4-2400 DRAM per node
- 288 nodes of Cray DataWarp (i.e. burst buffer)
 - 2x 3.2TB SSD (~1.8PB total)
 - 1.7TB/sec
- Sonexion 2000 Lustre: 248 OSTs, 10168 disks, 30 PB



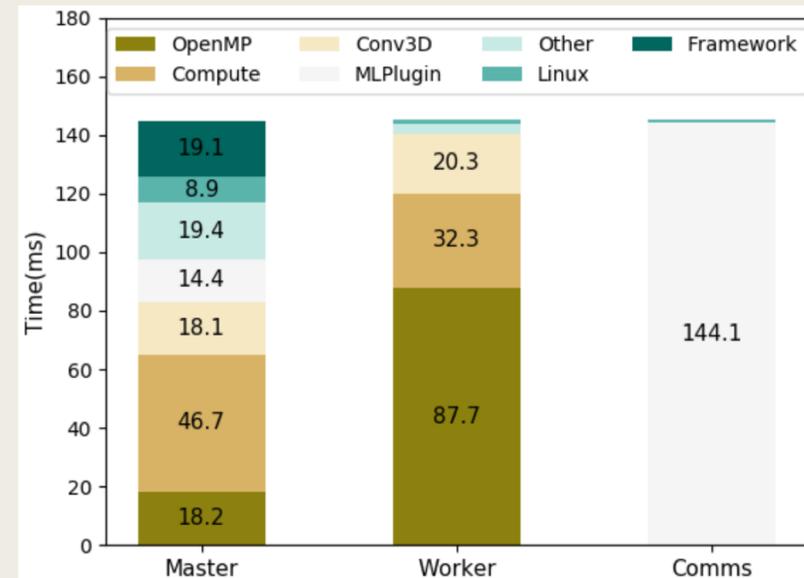
Piz Daint: Cray XC50 at CSCS

- 1431 nodes of 2x Intel Xeon E5-2695 v4
- 5320 nodes of Intel Xeon E5-2690 + NVIDIA P100 PCIe GPU
- Sonexion 3000 Lustre: 40 OSTs, 6.2 PB



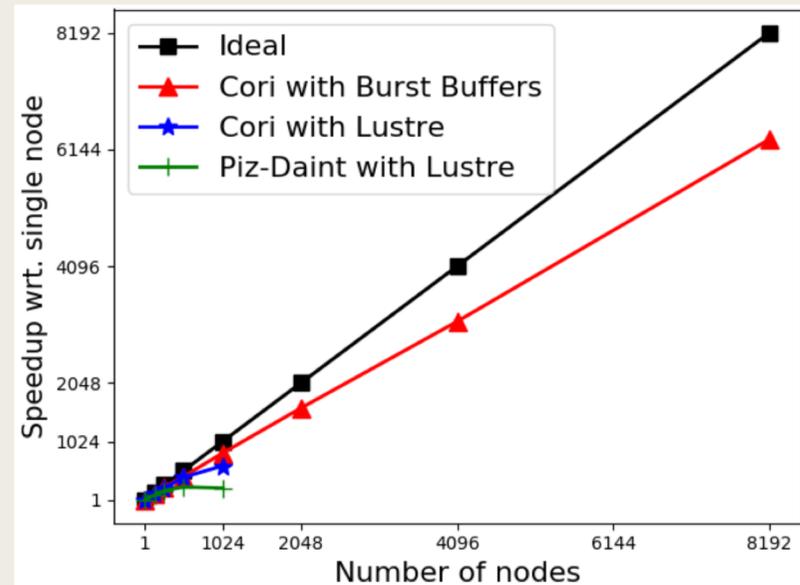
Single node performance results

- Data-parallel training is ideal due to volume of the data
- Times:
OpenMP spin time and overhead,
non-convolutional compute time,
3D convolutions,
CPE ML Plugin,
other time,
Linux kernel time,
TensorFlow framework time

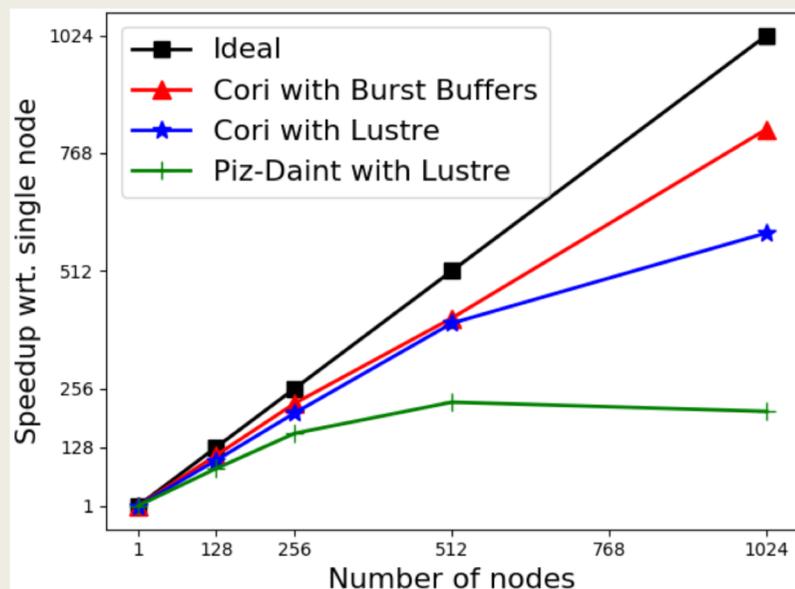


Scalability

- fully-synchronous training on 8192 nodes of Cori
- 77% parallel performance
- sustained 3.5Pflop/s (single precision)

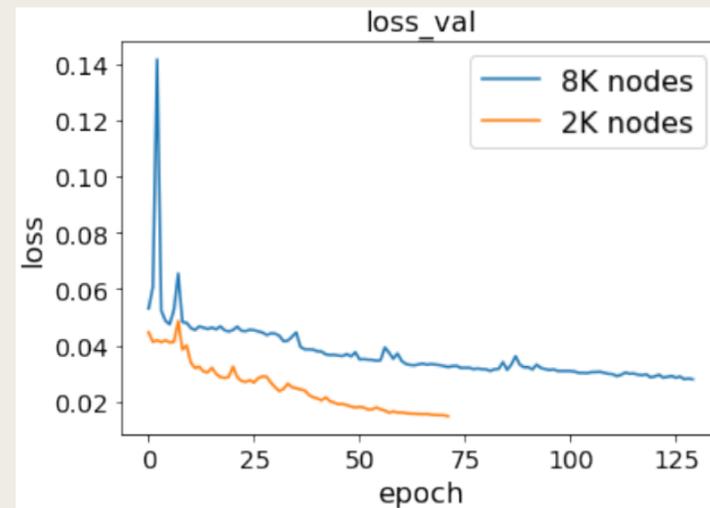
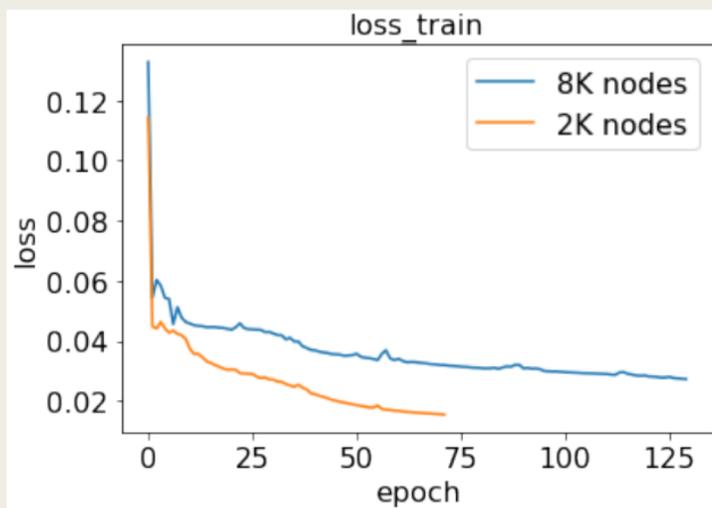


Scalability, contd.



- Lustre-based approaches fail to scale beyond 512~1024 nodes
- Hard to scale beyond anyway without careful optimizer tuning
- IO variability
- CPE ML plugin is MPI-based, as opposed to default inefficient implementation in TensorFlow

Loss function



Physical results

