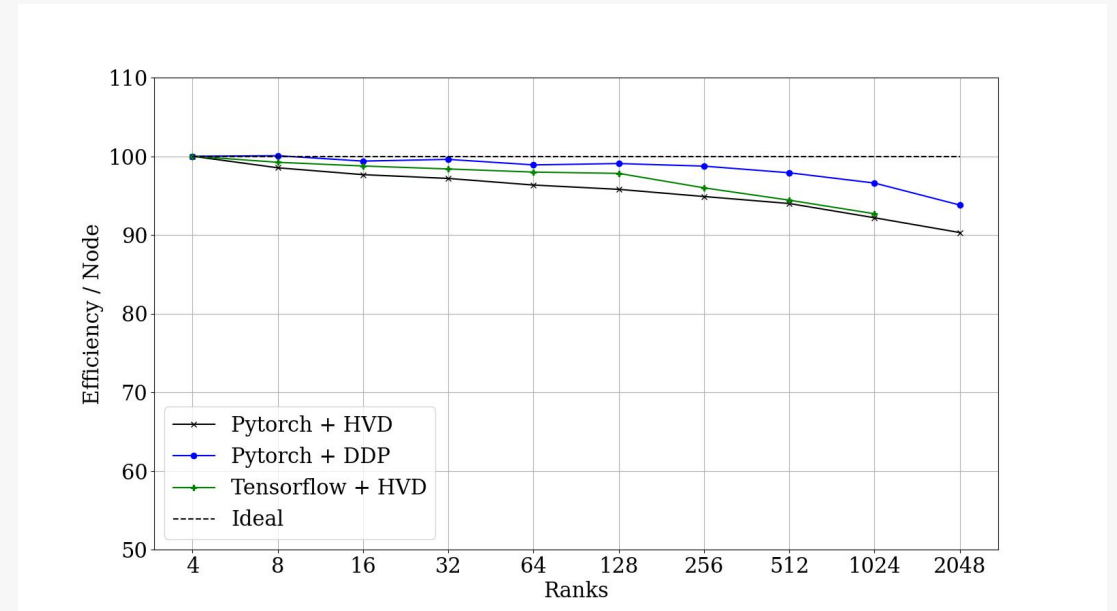


Scalable Frameworks on ALCF Systems

- Many applications target AI frameworks while requiring massive scale of applications.
- Scale-out on Aurora will be at unprecedented levels.
- On Aurora, we will support
 - Learning frameworks with Tensorflow, Torch, and JAX
 - Scale out with Horovod, DDP, DeepSpeed, mpi4py
- ALCF Datascience group has been working with Intel for many years focusing on optimal single-device, scale-up and scale-out performance for Aurora.

Interest in collaborating to scale out learning applications to the largest scales on Aurora.



The largest datasets and largest models will need to use hundreds to thousands of Aurora nodes effectively.

(~187 Aurora nodes = 1 Polaris system)

Fine tuning of scaling and good scale-out design will be critical on Aurora

Focus on: JAX

- One framework that has emerged as uniquely suited for AI for Science Applications: JAX
- Competitive advantages: “numpy + autograd on gpu”, but also has incredibly useful functional transformations (vmap, pmap, jit). Also has scipy operations available.
- Growing ecosystem of tools (flax, orbax, optax) make life easy and reduce boiler plate.
- Have scaled out JAX applications to nearly full Polaris runs (400+ nodes) efficiently.

- Application Highlight: **Differentiable Simulation**

- Full particle-physics detector simulation implemented in JAX.
- Trainable parameters have *physical meaning* (ie, “diffusion”)
- Small, hard-to-simulated pieces replaced by small MLPs
- Technique is applicable to any simulation effort with large datasets to tune against.

