


Several aspects of fusion energy physics are still not well understood.



Experimental methods are essential for gathering new operational modes. But **simulations** are used to validate basic theory, plan experiments, interpret results on present devices, and ultimately to design future devices.

**CGYRO** is a premier tool for multi-scale plasma turbulence simulation and has been in use by the fusion community for several years. It is an Eulerian gyrokinetic solver and relies heavily on global FFT computations, which require **low-latency** and high-throughput networking.

<https://gafusion.github.io/doc/cgyro.html>

Textbook HPC application

Current cutting-edge simulations require several TeraBytes of memory and  $O(100)$  TFLOPS worth of compute to complete in a reasonable timeframe. This used to be the domain of leadership-class HPC centers but can now be also found in the Microsoft operated Azure commercial Cloud.

**NDv2 Series** <https://azure.microsoft.com/en-us/pricing/details/virtual-machines/linux/>  
 NDv2-series virtual machine is the latest, fastest and most powerful addition to the GPU family specifically designed for the cutting edge demands of distributed HPC, AI and machine learning workloads. The VMs feature 8 NVIDIA V100 Tensor Core GPUs with NVIDIA NVLink interconnected GPUs with 32 GB of memory each, 40 non-hyperthreaded Intel Xeon Platinum 8168 processor cores, and 672 GiB of system memory. The NDv2-series VMs also feature 100 Gb/sec EDR InfiniBand with support for standard Mellanox OFED drivers and all MPI types and versions. With total of 256 GB of GPU memory and 100 Gb/sec InfiniBand interconnect NDv2-series VMs allowing for easy scale-up to supercomputer-sized clusters that can run the most demanding ML models and distributed AI training workloads utilizing Cuda, TensorFlow, Pytorch, Caffe, and other Frameworks.

Add to estimate	Instance	Core	RAM	Temporary storage	GPU	Pay as you go	1 year reserved (% Savings)	3 year reserved (% Savings)	Spot (% Savings)
+	ND40rs v2	40	672 GiB	2,900 GiB	8X V100 (NVlink)	\$26.438/hour	\$12.9548/hour (~51%)	N/A	\$4.6373/hour (~82%)

**HBv2-series**  
 HBv2-series VMs are optimized for applications driven by memory bandwidth, such as fluid dynamics, explicit finite element analysis, and weather modeling. HB VMs feature 120 AMD EPYC™ 7002-series CPU cores, 4 GB of RAM per CPU core, and no hyperthreading. HBv2-series VMs provide up to 350 GB/sec of memory bandwidth, which is 45-50 percent faster than x86 alternatives and 3x faster than what most HPC customers have in their datacenters today. HBv2-series VMs also feature 200 Gb/sec HDR InfiniBand from our technology partners Mellanox.

Add to estimate	Instance	Core	RAM	Temporary storage	Pay as you go	1 year reserved (% Savings)	3 year reserved (% Savings)	Spot (% Savings)
+	HB120rs v2	120	480 GiB	900 GiB	\$3.96/hour	\$2.9700/hour (~25%)	\$1.9800/hour (~50%)	\$0.6317/hour (~84%)



Unlike HPC centers, Microsoft Azure does not provide an HPC batch system or a shared file system out of the box.

Microsoft does offer **CycleCloud** as a free application, which we used. Note that CycleCloud is not tightly integrated into Azure and needs additional installation and configuration steps.

CycleCloud provides several batch system options; we chose **SLURM**, mostly due to our familiarity with that system.

CycleServer allows for push-button deployment of a fully configured SLURM submit node, with ssh access and auto-scaling capabilities. More advanced options, like requesting spot instances for worker nodes and using a dedicated, high-performance disk partition for shared storage, requires the use of their API.

We compiled and benchmarked CGYRO on both CPU-only and GPU-equipped instance types, using gcc 9 on CPU-only instances and PGI compilers on GPU-providing instances. OpenMPI v3 was used on all instances. We did not encounter any significant problems during setup or compilation.

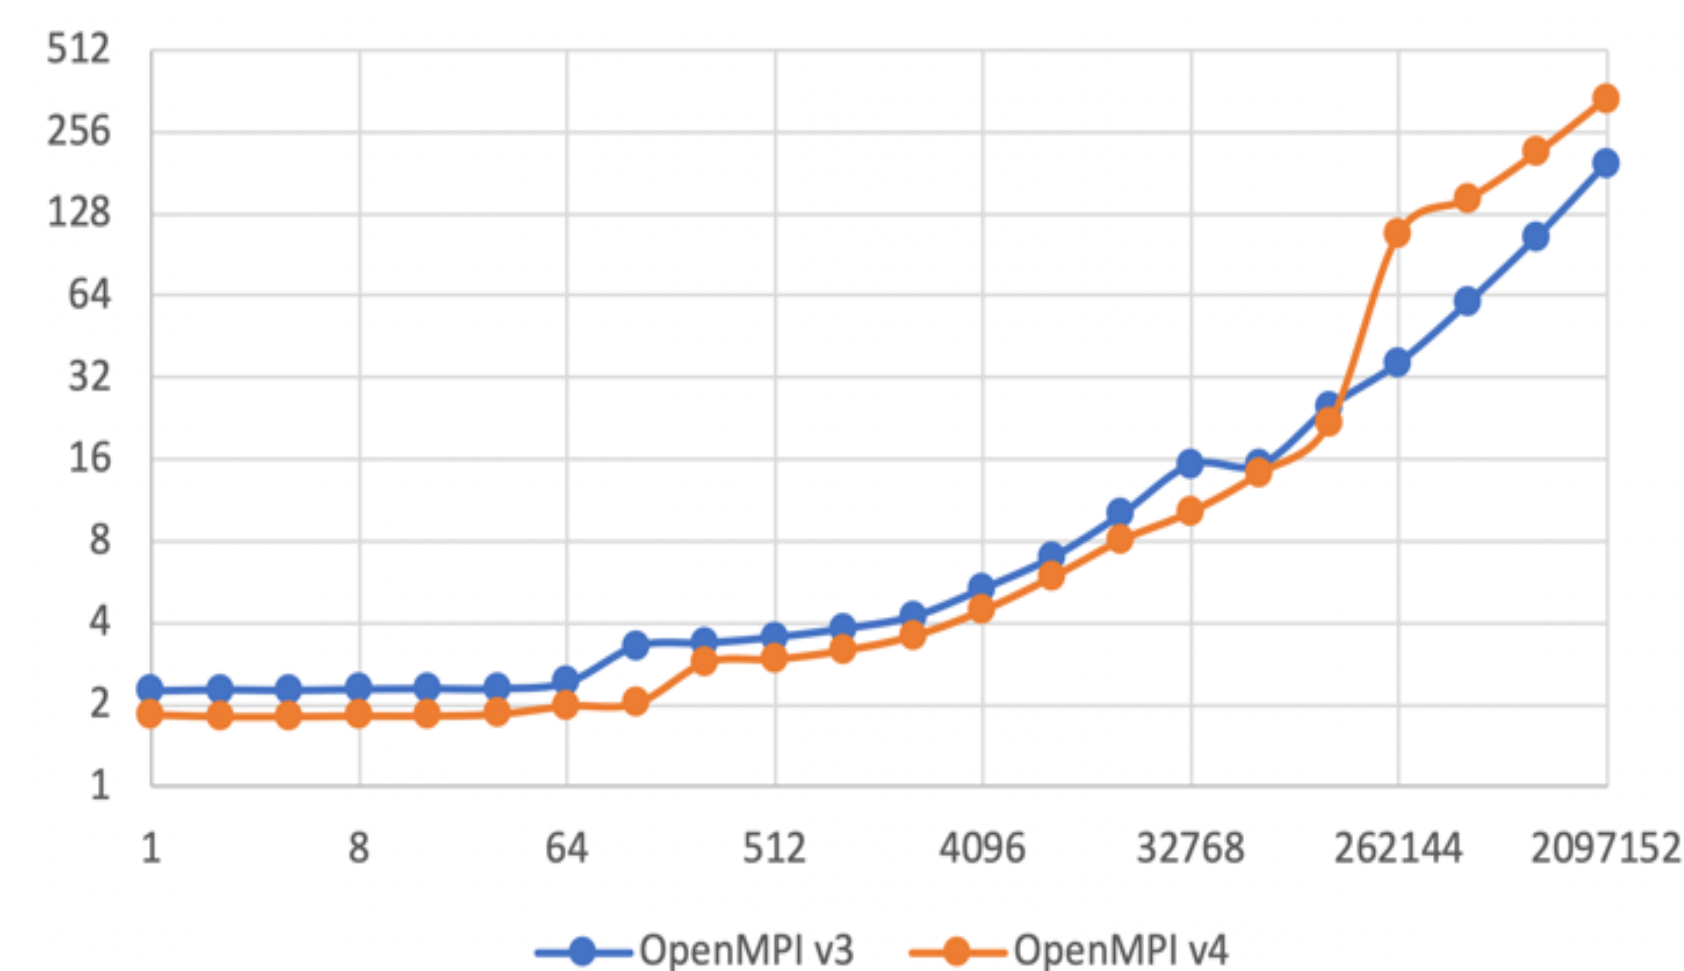
Note that we did also try other MPI implementations, including OpenMPI v4 and MVAPICH2; the performance difference was typically within 10%, with OpenMPI v3 generally providing the best results.



Azure offers three instance types that provide infiniband networking;

- NDv2 instances provide eight NVIDIA Tesla V100 32GB GPUs,
- HC instances provide dual Intel Xeon Platinum 8168 CPUs and
- HBv2 instances provide dual AMD EPYC 7002 CPUs.

The first two come with 100 Gbps EDR and the last comes with 200 Gbps HDR. As a result, the network latencies between nodes are very low.



## Cutting-edge CGYRO simulation

[https://github.com/scidac/atom-open-doc/blob/master/2020.11-SC20/multiscale\\_input/input.cgyro](https://github.com/scidac/atom-open-doc/blob/master/2020.11-SC20/multiscale_input/input.cgyro)

All numbers represent one typical step during the simulation. Cost is computed using spot instance pricing.

Instance	Nodes	Total time	Comm. time	Total cost
NDv2	24	161s	139s	\$5.24
NDv2	8	369s	316s	\$4.00
HBv2	36	272s	104s	\$1.81
HBv2	18	441s	190s	\$1.47
HC	35	416s	113s	\$2.42
HC	18	763s	151s	\$2.28

The GPU-providing NDv2 instances are drastically faster on a per-node basis, but AMD CPU-based NBv2 instances provide a much better value per dollar.

The main reason for GPU-based instances to trail in cost-effectiveness is due to the amount of time spent in network communication; a single 100 Gbps EDR infiniband link is a poor match for the compute power of 8 V100 GPUs.

Note the Intel CPU-based NC instances trail the other two in the benchmarks and are thus not a good fit for users using CGYRO.

## Typical size CGYRO simulation (nl03 benchmark case)

All numbers represent one typical step during the simulation. Cost is computed using spot instance pricing.

Instance	Nodes	Total time	Comm. time	Total cost
NDv2	16	121s	92s	\$2.64
NDv2	4	397s	293s	\$2.15
HBv2	36	87s	45s	\$0.58
HBv2	9	289s	64s	\$0.48
HC	24	223s	60s	\$0.89
HC	12	431s	96s	\$0.86

Using **spot instances** implies

- lower availability and
- potential preemption during runtime.

CGYRO is instrumented to deal with preemption by implementing checkpointing at a tunable interval; in our tests, we checkpointed every couple of hours, resulting in negligible overhead.

Availability of a large number of nodes can be a problem. We were never able to provision more than 24 NDv2 nodes or 36 HBv2 nodes. The experienced preemption rate was however very low for smaller node counts. We ran a simulation using 8 NDv2 nodes over a 5-day period with only 3 preemption events and had several 8 hour runs on 18 HBv2 nodes without a single preemption. Given that spot instances cost less than 20% compared to the more reliable on-demand instances, the trade-off is well worth it for the small node counts.

In order to compare the performance of CGYRO on Cloud resources, we ran the same cutting-edge simulation also on

- OLCF Summit, which is a GPU system, and
- NERSC Cori, which is a CPU-only system.

System	Nodes	Total time	Comm. time
Summit	32	86s	67s
Cori	128	165s	62s
Cori	48	339s	160s

Azure resources easily match the performance of the NERSC HPC system.

The slowest result on Summit is however significantly faster than anything we measured on Azure, which is due to Summit nodes both having dual-rail 100 Gbps infiniband and only six NVIDIA Tesla V100 16 GB GPUs; we could not perform the CGYRO simulation with less than 32 nodes, due to GPU memory requirements.

System	Nodes	Total time	Comm. time
Summit	16	82s	46s
Cori	64	112s	46s
Cori	16	372s	120s

Azure resources again match performance of resources provided by HPC centers.

## Summary

We explored the feasibility of running CGYRO on Azure HPC resources in spot mode. We observed both that CPU-only resources were very efficient, and that running in spot mode was doable, with minimal side effects. The GPU-enabled resources were less cost effective but allowed for higher scaling.