Microsoft

Scaling Performance for Barracuda Virtual Reactor on Azure

Ahmed Taha
Senior Specialists, HPC + AI GBB Americas

Barracuda Virtual Reactor Users' Conference
Chicago, IL
June 30, 2022

Agenda

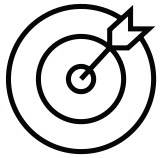

Objectives of the Study

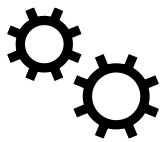

Study Models

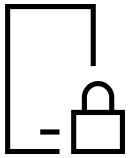

Azure GPUs Used

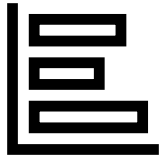

Results

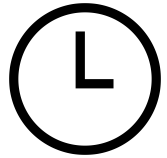

Analysis / Conclusions

Credits

Objectives of the Study

Testing Barracuda Virtual Reactor "BVR" on Azure's state-of-the-art GPU-Compute virtual machines, VMs (V100 & A100).

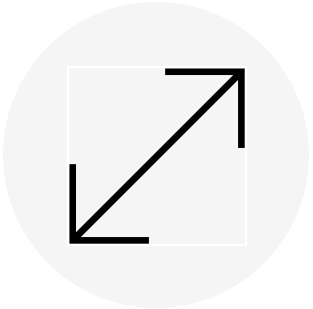

Running the "biggest industrially-relevant model" on Azure's iconic GPUs.

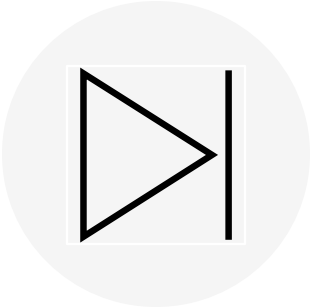

Showcase performance scalability and speedup for BVR on Azure across a variety of models.

Showcase of economic feasibility for "performance-to-price" ratio running on Azure's A100 GPUs.

Study Models

TWO CASE STUDIES/TWO MODELS

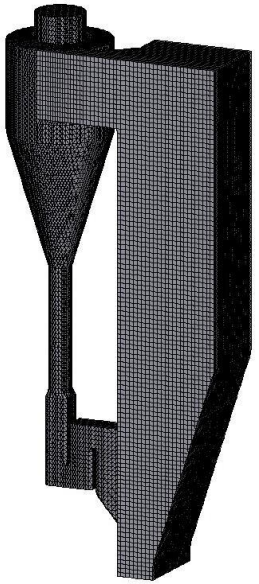

1. BENCHMARK MODEL (CIRCULATING FLUIDIZED BED "CFB" REACTOR).

2. "ENCINA" PRODUCTION MODEL (NON-CIRCULATING FLUIDIZED BED "CFB" REACTOR).

Model-1

1. BENCHMARK MODEL

(CIRCULATING FLUIDIZED BED "CFB" REACTOR).

- **FOUR models: 3 proprietary & 1 generic** (presented here ➔)
- **EIGHT tests with biggest model of 55M particles** (below table ⬇)

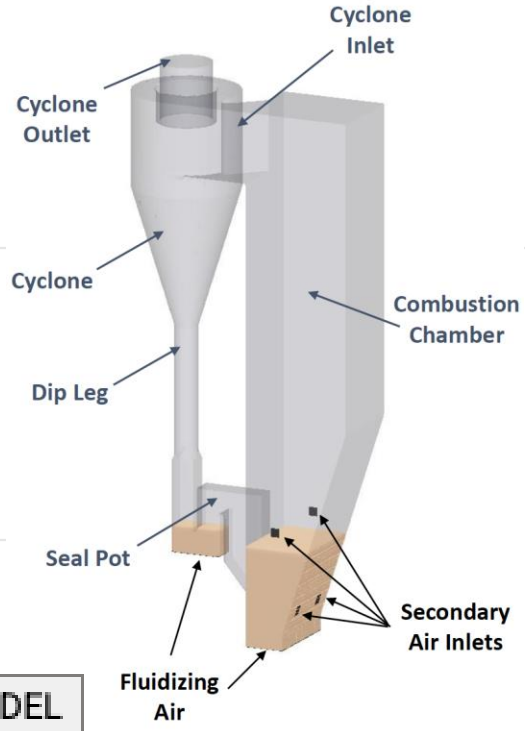

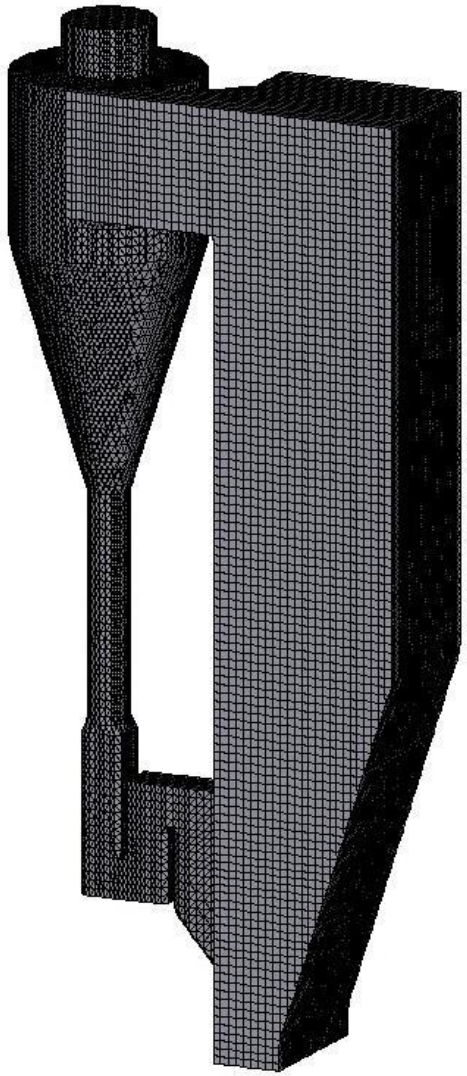

TEST PROBLEM NO	CELL COUNT	NO. OF PARTICLES	CHEMISTRY	THERMAL	P1 MODEL
479	243267	29920300	enabled	enabled	disabled
499	243267	3581140	enabled	enabled	disabled
480	105157	40581300	disabled	disabled	disabled
500	105157	9528670	disabled	disabled	disabled
481	389320	50170500	enabled	enabled	disabled
501	389320	12994400	enabled	enabled	disabled
482	821781	55070200	enabled	enabled	disabled
502	821781	22625700	enabled	enabled	disabled

Published (Jan. 2022): https://techcommunity.microsoft.com/t5/azure-compute-blog/barracuda-virtual-reactor-on-azure/ba-p/3054715

Model-1: Azure GPU Used

1. BENCHMARK MODEL

(CIRCULATING FLUIDIZED BED "CFB" REACTOR).

- **Three Azure GPU SKUs:** NDv4 A100, NCv3-series & NCas_T4-v3-series (Specs detailed below⬇)

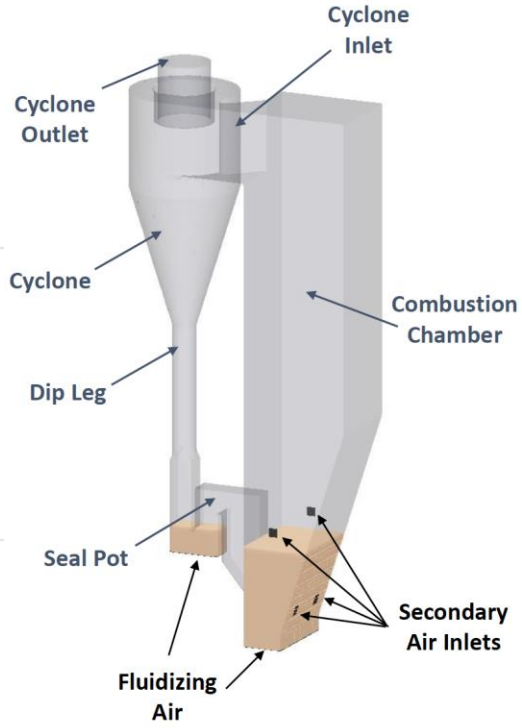

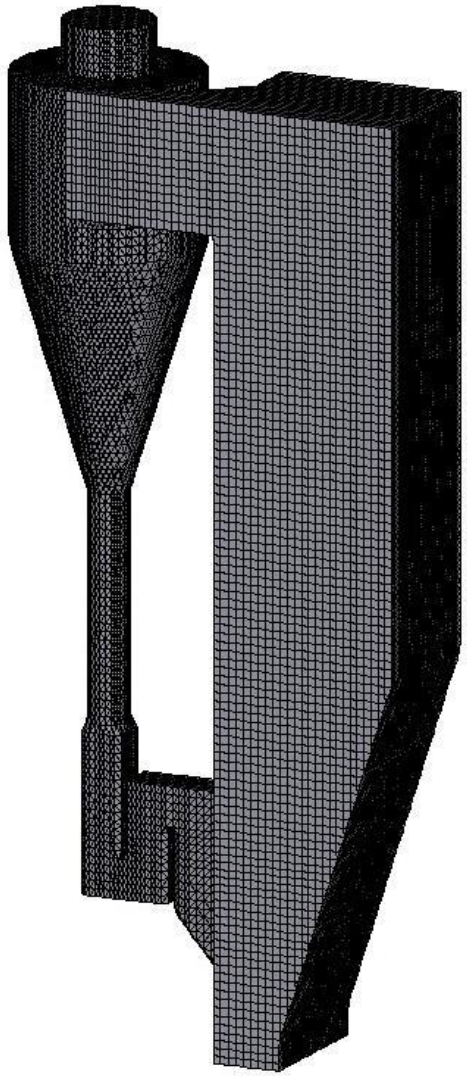

VM NAME	VCPU	MEMORY (GIB)	SSD (GIB)	GPU	GPU MEMORY (GIB)	MAX DATA DISK
Standard_ND96asr_v4	96	921	1024	8 A100	40	32
Standard NC24s_v3	24	448	2048	4 V100	16	32
Standard NC64as_T4_v3	64	448	1024	4 T4	16	32

Published (Jan. 2022): https://techcommunity.microsoft.com/t5/azure-compute-blog/barracuda-virtual-reactor-on-azure/ba-p/3054715

Model-1: Results

1. BENCHMARK MODEL

(CIRCULATING FLUIDIZED BED "CFB" REACTOR)

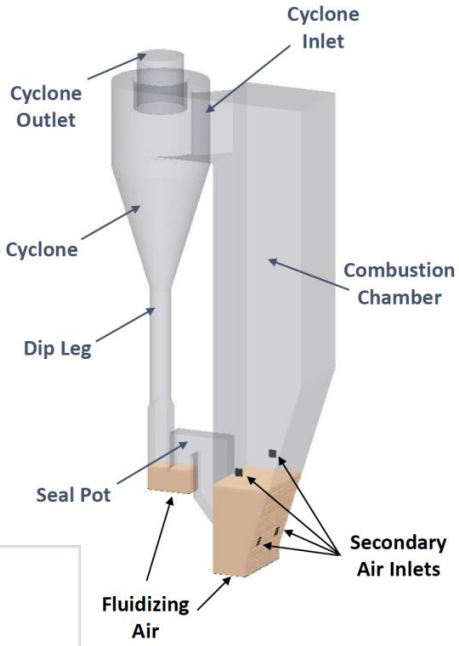

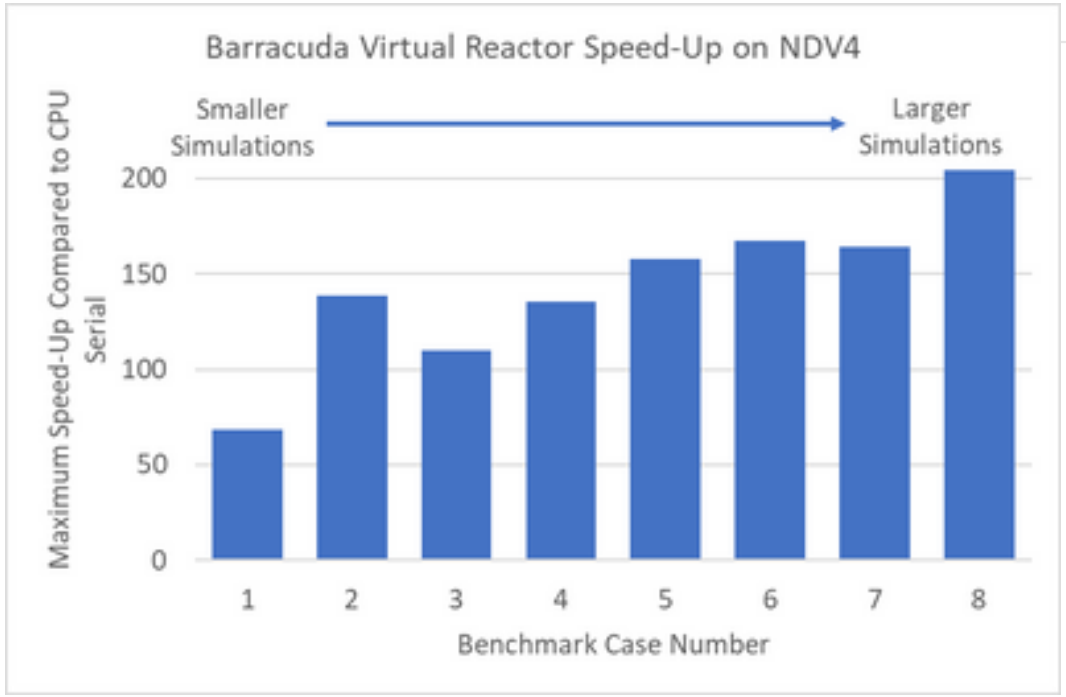

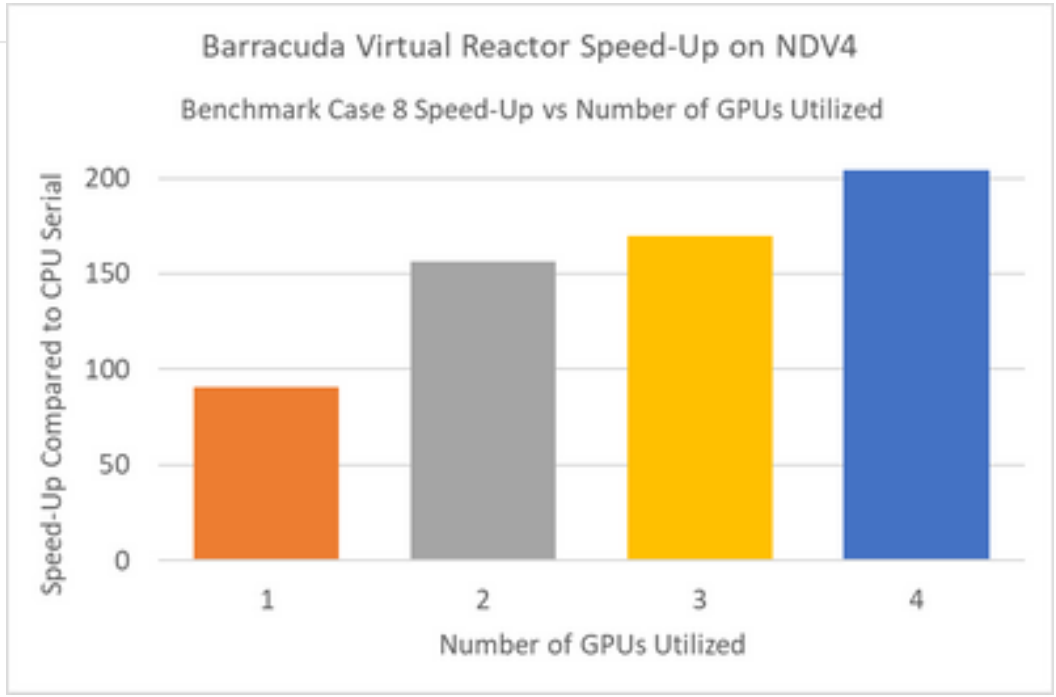

200x speed-up*

(*Compared to CPU Serial)

Published (Jan. 2022): https://techcommunity.microsoft.com/t5/azure-compute-blog/barracuda-virtual-reactor-on-azure/ba-p/3054715

Model-2

2. “ENCINA” PRODUCTION MODEL

(NON-CIRCULATING FLUIDIZED BED REACTOR)

Model Description

- Largest real-world Virtual Reactor simulation ever run (96M particles/clouds).
- Much larger than the simulations most Barracuda users typically run (hardware constraints).
- Taking advantage of Azure’s advanced GPU system (capacity & speed).

Case size details	
Number of real cells	2.9 million
Number of clouds	96 million
Thermal calculations	Enabled
Chemical reactions	Enabled
Memory requirement	72 GB

Model-2: Azure GPU Features

2. “ENCINA” PRODUCTION MODEL

(NON-CIRCULATING FLUIDIZED BED REACTOR)

Azure GPU Platform Features:

i. V100: NDv2 V100 [SKU name: "Standard_ND40rs_v2"]

1. Specs:

- Powered by 8 NVIDIA Tesla V100 Tensor Core NVLINK-connected GPUs, each with 32 GB of GPU memory.
- Each NDv2 VM also has 40 non-HyperThreaded Intel Xeon Platinum 8168 (Skylake) CPU cores and 672 GiB of system memory.

2. Uses:

- Best suited for AI Training, AI Inference, HPC.

Model-2: Azure GPU Features

2. "ENCINA" PRODUCTION MODEL

(NON-CIRCULATING FLUIDIZED BED REACTOR)

Azure Platform Details:

ii. A100: NDv4 A100 [SKU name: "Standard_ND96ars_v4"]

1. Specs:

- Powered by 8 NVIDIA Ampere A100 Tensor Core GPUs with 40 GB of memory each and 96 physical 2nd-generation AMD Epyc™ CPU cores (2.44GHz) with a total memory of 921GB.
- Each GPU features NVLINK 3.0 connectivity for communication within the VM.
- ND A100 v4-based deployments can scale up to thousands of GPUs with 1.6 Tb/s of interconnect bandwidth per VM. Each GPU within the VM is provided with its own dedicated, topology-agnostic NVIDIA Quantum 200Gb/s InfiniBand networking. These connections are automatically configured between VMs occupying the same virtual machine scale set, and support GPUDirect RDMA.
- The scale-out InfiniBand interconnect is supported by a large set of existing AI and HPC tools built on NVIDIA's NCCL2 communication libraries for seamless clustering of GPUs.

2. Uses:

- Best suited for DL Training, DL Inference, HPC, High Performance Data Analytics.

Model-2: Azure GPU Features

2. "ENCINA" PRODUCTION MODEL

(NON-CIRCULATING FLUIDIZED BED REACTOR)

Azure Platform Details: (GPU SKUs specs & features)

ii. A100: NDv4 A100 [SKU name: "Standard_ND96ars_v4"]

3. Special Feature: *(Multi Instance GPU "MIG")*

- A100 with MIG maximizes the utilization of GPU-accelerated infrastructure.
- With MIG, an A100 GPU can be partitioned into as many as ***seven*** independent instances, giving multiple users access to GPU acceleration.
- With A100 40GB, each MIG instance can be allocated up to 5GB, and with A100 80GB's increased memory capacity, that size is doubled to 10GB.
- MIG lets infrastructure managers offer a right-sized GPU with guaranteed quality of service (QoS) for every job, extending the reach of accelerated computing resources to every user.

Model-2: Results

2. "ENCINA" PRODUCTION MODEL

(NON-CIRCULATING FLUIDIZED BED REACTOR)

Performance Comparison: 4A100 vs. 4V100

Azure VM	GPUs Used	Simulation Speed (s/day)	Simulation Speed (s/hour)
ND40rs v2 (8 x 32 GB V100 GPUs)	4 x V100	6.61	0.28
ND A100 v4 (8 x 40 GB A100 GPUs)	4 x A100	20.22	0.84
	Speed-up factor	*3.06*	

Model-2: Results

2. "ENCINA" PRODUCTION MODEL

(NON-CIRCULATING FLUIDIZED BED REACTOR)

Simulation Speed (wall clock seconds/day):

- A100 GPU achieved a simulation speed of 20.22 wall clock seconds/day outperforming the V100 GPU which achieved only 6.61 wall clock seconds/day.

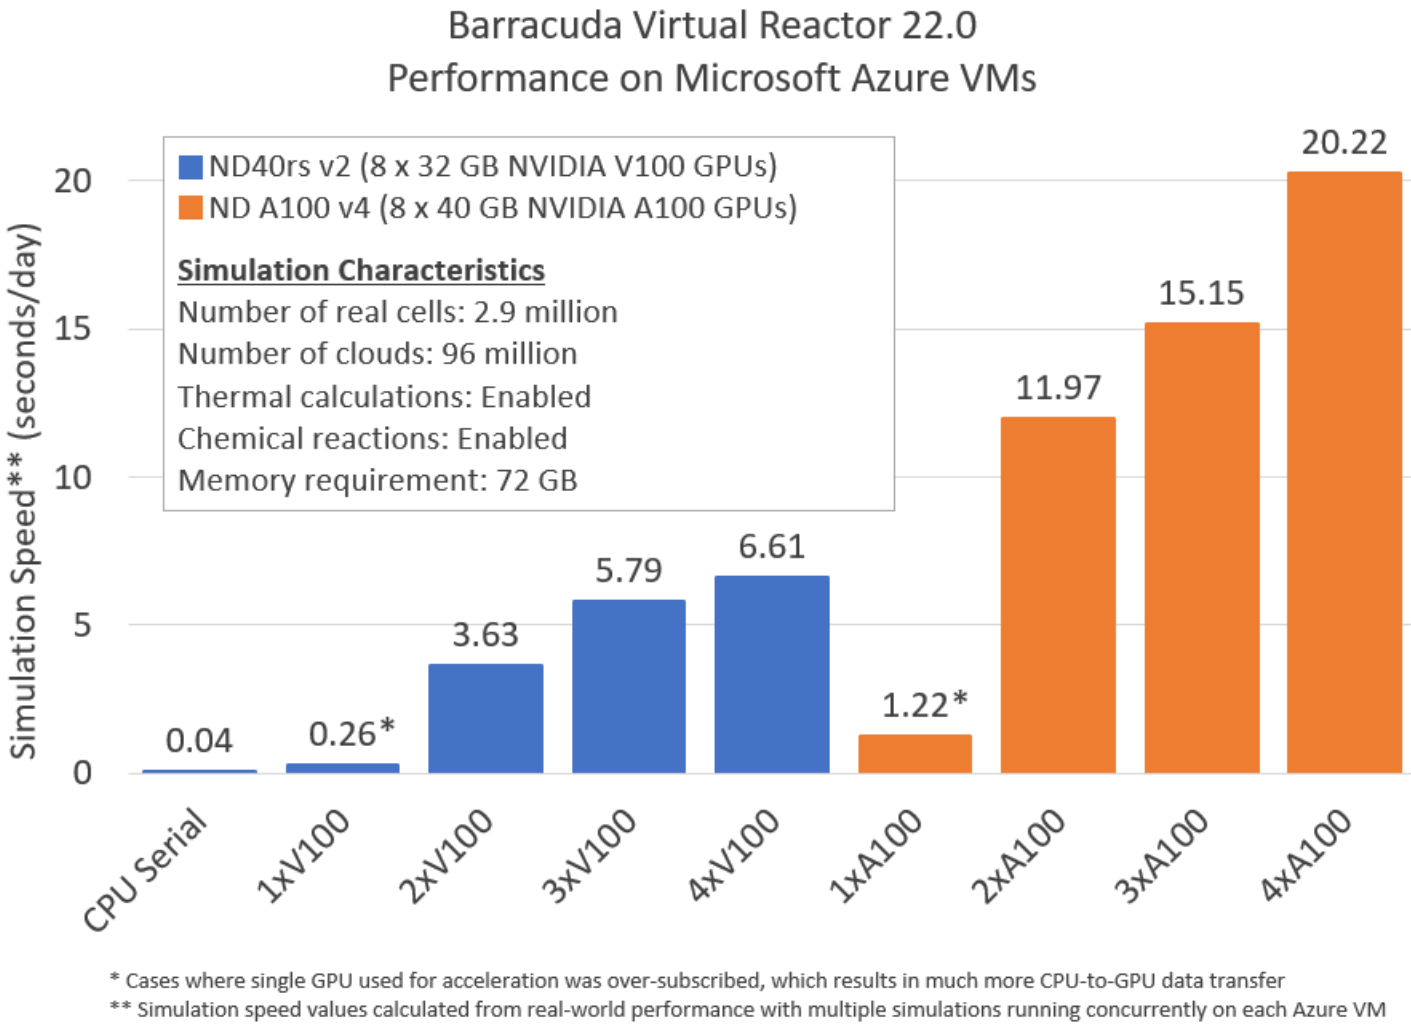

Model-2: Results

2. "ENCINA" PRODUCTION MODEL

(NON-CIRCULATING FLUIDIZED BED REACTOR)

Simulation speedup in comparison to the CPU Serial case: = GPU Simulation Speed/CPU Serial Simulation Speed:

- A100 achieved speed up of **505.6*** = (20.22/0.04) Compared to 165* (6.61/0.04) for V100 as shown in Figure-5.

*Compared to CPU Serial

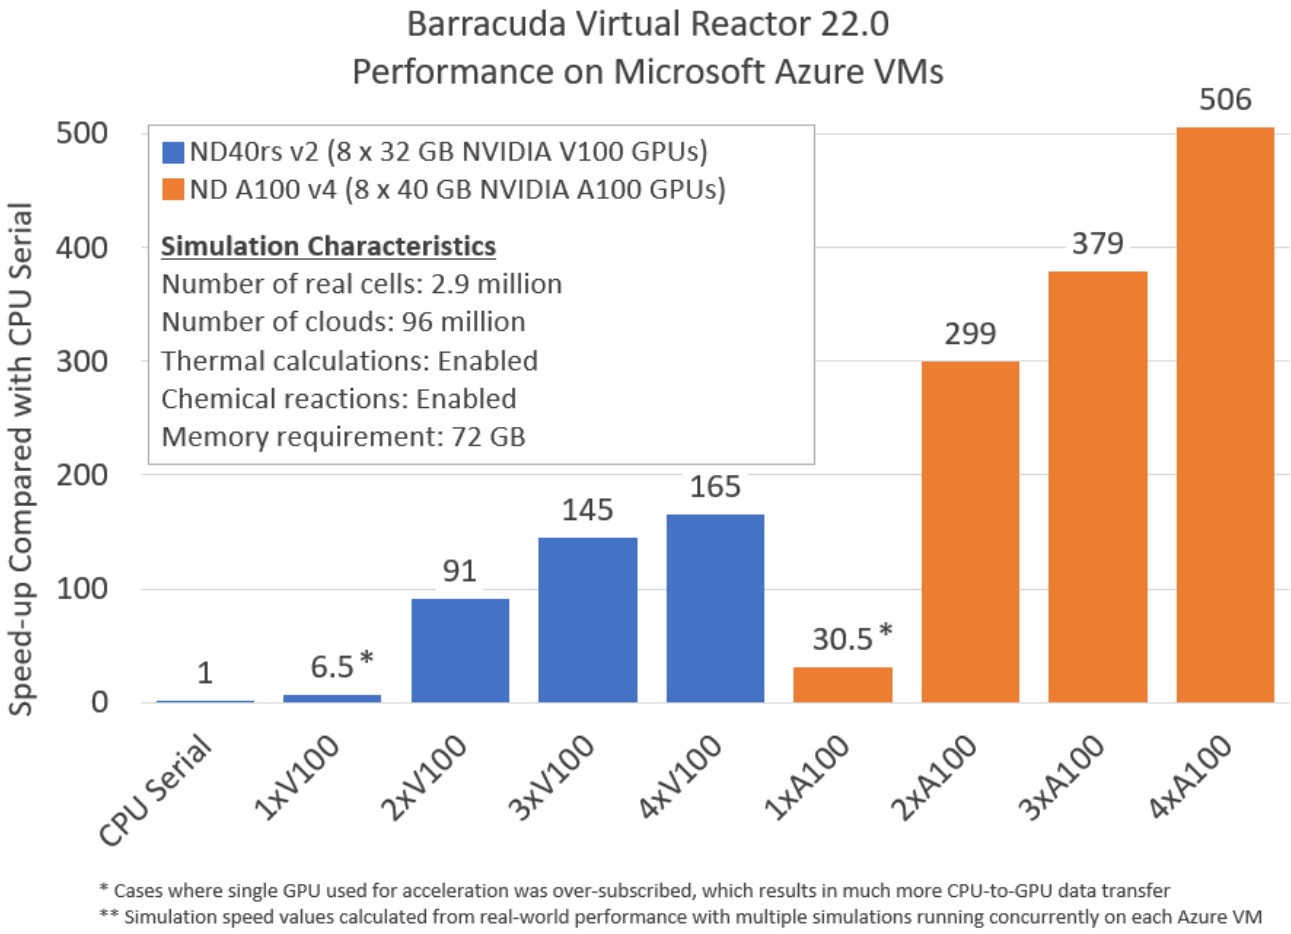

Model-2: Results

2. "ENCINA" PRODUCTION MODEL

(NON-CIRCULATING FLUIDIZED BED REACTOR)

Show case for the compute capabilities of A100: Better speedup for bigger models

Case	Size (Particles)	GPU SKU	# of GPUs	Speedup (Compared to CPU Serial)
Model-1	55M	A100	4	200
Model-2	96M	A100	4	506

Model-1

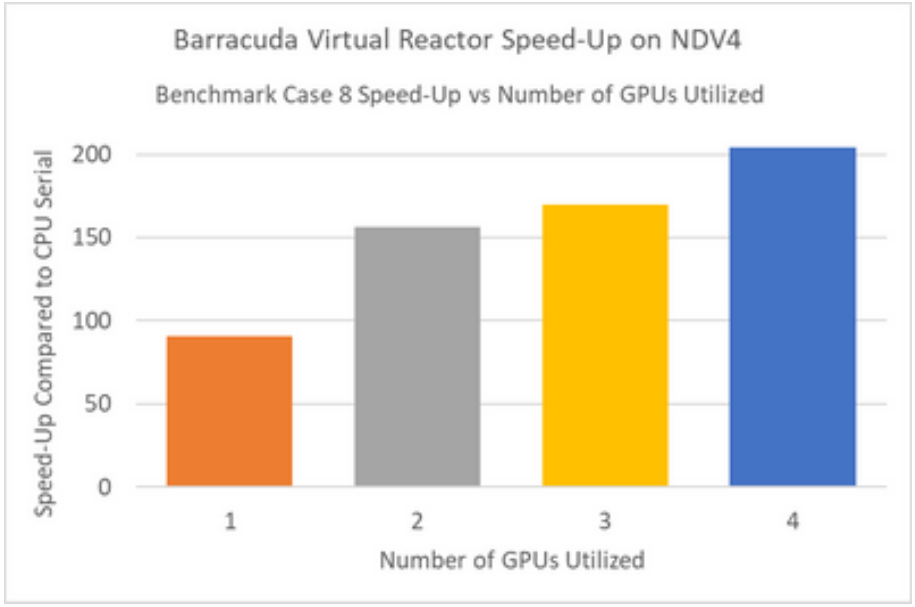

Model-2

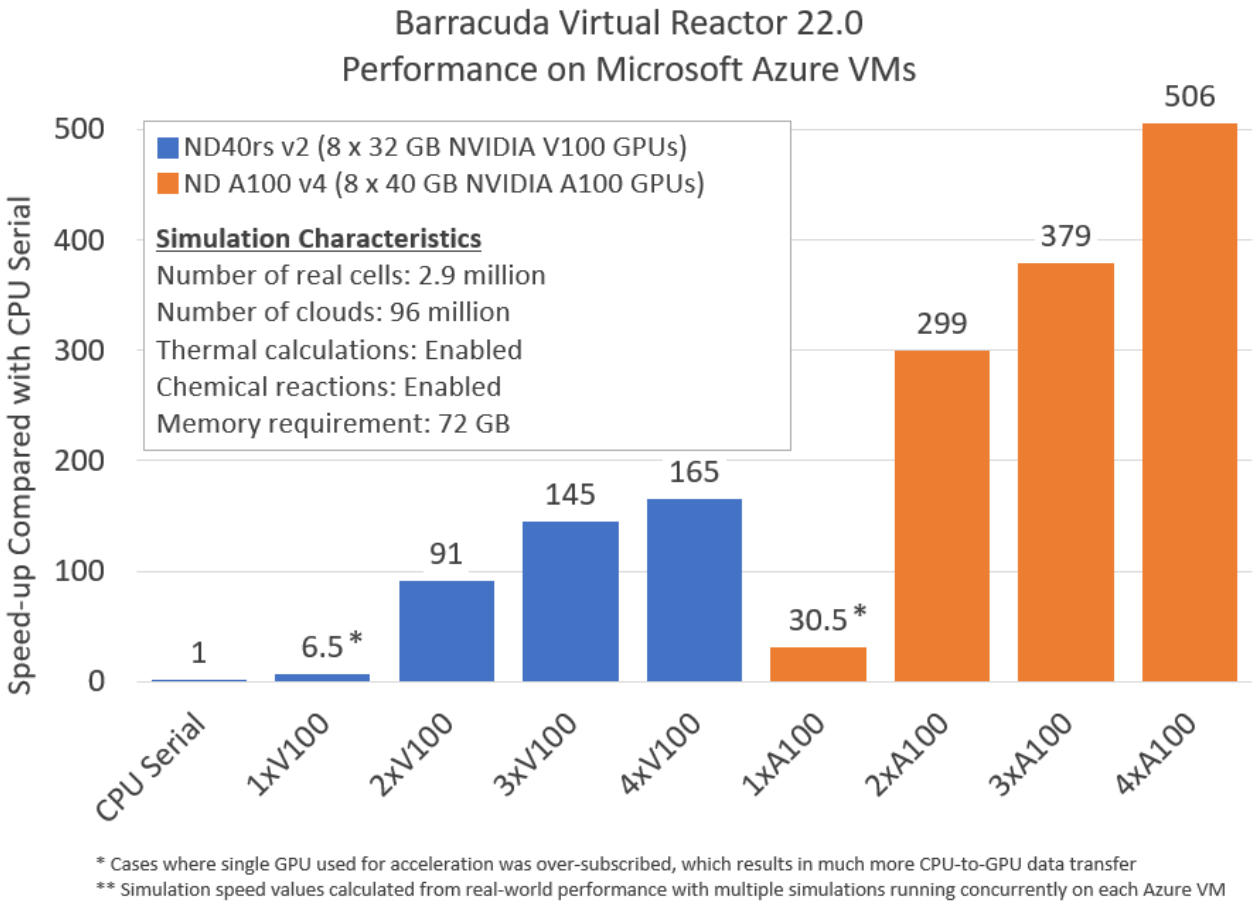

Model-2: Results

2. "ENCINA" PRODUCTION MODEL

(NON-CIRCULATING FLUIDIZED BED REACTOR)

Performance & Cost Comparison: 4xA100 vs. 4xV100

Azure VM	GPUs Used	Simulation Speed (s/day)	Simulation Speed (s/hour)
ND40rs v2 (8 x 32 GB V100 GPUs)	4 x V100	6.61	0.28
ND A100 v4 (8 x 40 GB A100 GPUs)	4 x A100	20.22	0.84
	Speed-up factor	*3.06*	
VM Pricing as of June 8, 2022	**Pay as you go ($/hour)**	**1 year reserved ($/hour)**	**3 year reserved ($/hour)**
ND40rs v2	$26.44	$12.95	NA
ND A100 v4	$32.63	$22.59	$13.05
Assume 2 concurrent simulations	**Pay as you go**	**1 year reserved**	**3 year reserved**
Price per simulation second	**($/simulation second)**	**($/simulation second)**	**($/simulation second)**
ND40rs v2	$48.00	$23.52	NA
ND A100 v4	$19.36	$13.41	$7.75

Model-2: Analysis

2. "ENCINA" PRODUCTION MODEL

(NON-CIRCULATING FLUIDIZED BED REACTOR)

Capacity Optimization Analysis.

- The A100 & V100 systems have huge capacity.
- Since each system tested had 8 GPU cards each, 2 concurrent simulations ran on each system.
- This effectively can cut the price of each "Sim. Sec." into half as long as users have enough simulation work to keep all 8 GPUs fully utilized.

Model-2: Analysis

2. "ENCINA" PRODUCTION MODEL

(NON-CIRCULATING FLUIDIZED BED REACTOR)

Financial Analysis: 4xV100.vs.4xA100 cost analysis and comparison.

Assume 2 concurrent simulations	Pay as you go	1 year reserved	3 year reserved
Price per simulation second	($/simulation second)	($/simulation second)	($/simulation second)
ND40rs v2	$48.00	$23.52	NA
ND A100 v4	$19.36	$13.41	$7.75

- The fastest hardware (A100) is also the best value ($$$/Sim Sec).
- When considering price per simulation second, the A100 system is a better value.
- Given its higher performance to price ratio for Barracuda simulations (compared to the V100-based system), customers can run at maximum speed and save money by using the A100-based system.

Summary

- The current study highlighted the great compute capacity of Azure's latest and greatest GPU SKU's; V100, NDv2 V100 and A100, NDv4 A100.
- The study ran Barracuda Virtual Reactor "BVR" for simulating the largest real-world, industry grade Virtual Reactor simulation ever run making it much larger than any simulation most Barracuda users typically run.
- The A100 GPU system exhibited an outstanding capability managing such massive model scaling it to a speedup of 506X compared to the CPU serial case.
- NVDIA's technology of Multi Instance GPU "MIG", expands the performance and value of the NVIDIA A100 allowing users to run as much as 7 concurrent simulations on the same GPU VM.
- This is not only extending the reach of accelerated computing resources to every user but decreases the cost/simulation making the A100 system of a better value.

Credits

- **Encina:**
  - Song Wang
- **CPFD:**
  - Sam Clark
  - Peter Blaser
- **Microsoft Azure:**
  - Ahmed Taha
  - Gauhar Junnarkar

Thank You !