DATA SHEET

CPFD Software's Barracuda Virtual Reactor

Making Multi-GPU Simulation Accessible to All with DGX A100

Fluid-particle systems are at the heart of virtually every industry, including refining, petrochemicals, materials, basic chemicals, power generation, manufacturing, and other energy-intensive industries, and are increasingly being used for sustainability, renewables, and decarbonization applications. These systems are used every day by organizations everywhere to build the physical infrastructure and materials that comprise and fuel the foundation of the world we live in.

Simulation of fluid-particle systems, such as fluidized bed reactors, are game changers for organizations looking to maximize their efficiency, reliability, throughput, and impact, no matter if their focus in on accelerated R&D, commercialization, scale up, or operations. Simulation accelerates performance by helping to determine root cause of phenomena, reducing the risk of changes through virtual testing, and identifying additional optimization opportunities.

CPFD Software's award-winning Barracuda Virtual Reactor software platform simulates 3D, transient behavior in fluid-particle systems including the multiphase hydrodynamics, thermal, and chemical reactions in industrial units. It utilizes physics-based modeling with NVIDIA GPU and Multi-GPU parallelization in order to provide an industry-leading simulation experience, no matter the size or type of simulation or industry use case needed. Specifically, the recent release of Multi-GPU parallelization is a groundbreaking capability for those organizations with large or very large simulations, empowering them to complete simulations in minutes or hours, which previously took days or weeks.

KEY FEATURES

- 50x-400x speedups for Single GPU and Multi-GPU Simulations
- Barracuda Virtual Reactor appliance
  - Powered by NVIDIA DGX A100 Station or Server
  - Provided by Mark III Systems
  - Integrated onsite or pre-stage

SPECIFICATION

CPFD Software	**Barracuda Virtual Reactor software platform**
NVIDIA	**DGX A100 Server** - 8 GPU Server **DGX A100 Station** - 4 GPU Workstation
Mark III Systems	**BVR Appliance Integration Services** **Onsite Assistance** **Onboarding Services**

DGX Platform solutions Overview

NVIDIA DGX platform solutions provide Barracuda Virtual Reactor users world class GPU performance to tackle the most demanding simulation workloads. Consisting of the DGX A100 enterprise server and the DGX Station workstation, DGX platform solutions deliver world class performance at data center scale or smaller office environments. The DGX Station is a GPU accelerated workstation built around 4 NVIDIA A100 enterprise GPUs. With 512GB of system memory and 320GB of GPU memory, the DGX Station packs the horsepower for demanding workloads or small working groups. The DGX Station provides clients with the same enterprise GPUs running in the world's most powerful supercomputers, in a unit that runs quietly under your desk. The DGX A100 Enterprise Server is the best choice for Barracuda users needing optimum performance and scale. This NVIDIA engineered enterprise server provides optimal configuration for multi-GPU and multi node workloads. With 8 NVIDIA A100 GPUs, 4TB of system memory, 640GB of GPU memory, and 4 Tb/sec of peak bi-directional bandwidth, the DGX A100 Server will accommodate the most resource intensive simulations. This allows Virtual Reactor users to explore more complex simulations and simulate orders of magnitude faster than CPU or single GPU configurations. Leveraging the NVIDIA designed Reference Architectures for DGX A100 Servers, allows organizations to optimize performance at the scale. Each DGX platform solution provides a common management environment and tool kit that streamlines deployment from dev/test to production. DGX Platform Solutions for CPFD Barracuda Virtual Reactor, offers the right solution options for organizations looking to easily deploy GPUs, support multiple users, decrease simulation times, and increase simulation frequency. Overall driving more value from the Barracuda platform and helping you meet your organizational goals.

Barracuda Virtual Reactor on NVIDIA DGX A100

CPFD Software ran a battery of benchmarks that compared Barracuda Virtual Reactor simulations on NVIDIA A100 GPUs to CPUs. A 50x-400x speedup was observed, where the acceleration was dependent upon the size and type of each simulation and whether single GPU or multi-GPU parallelization was utilized. The larger and more complex the simulation, the greater the impact of multi-GPU parallelization and the greater the overall speedup.

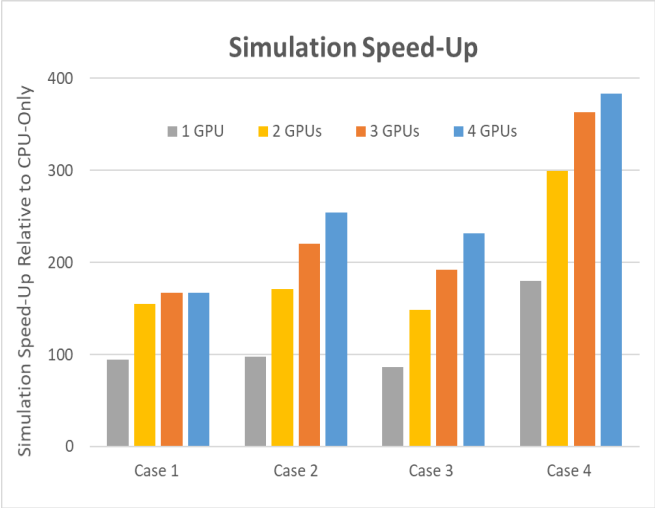

In terms of absolute time and business impact, this often means that simulations that would take days or weeks can now be done in hours or even minutes. Many organizations now run hundreds more simulations within the same time frame as was previously needed for a single case, resulting in broader R&D and accelerated commercialization and scale up. Practically speaking, the speed enables engineers to create higher-fidelity models than ever before, and previously impossible problems can now be solved in industrially relevant time scales.

CPFD Software + NVIDIA + Mark III Systems

Mark III Systems, an NVIDIA NPN Elite Partner, has partnered with NVIDIA and CPFD Software to build a Barracuda Virtual Reactor appliance (powered by DGX A100 Station and Server) to accelerate the time-to-value and simplify the onboarding and ongoing user experience for organizations leveraging Barracuda Virtual Reactor and simulation to transform their work. Regardless of if you need to run Barracuda Virtual Reactor quietly under your desk on a DGX A100 Station or with maximum multi-GPU parallelization in the datacenter with DGX A100 Server, Mark III can customize a bundled package for with Barracuda Virtual Reactor and DGX A100, complete with white glove services for DGX installation onsite at your office or datacenter, or custom pre-staging and pre-integration at Mark III's Global Integration Center in Houston, TX, based on your unique requirements.

For More Information Contact:
Andy Lin
VP of Strategy and Innovation
Email: andy.lin@markiiisys.com

Systems	Performance	Compute	Memory/Storage	Connectivity
DGX A100 Enterprise Server	19.5 TFLOPS of FP64 for HPC 4.8TB/s bi-directional bandwidth	**CPU** Dual AMD Rome 7742, 128 cores total, 2.25 GHz (base), 3.4 GHz (max boost) **GPU** 8x NVIDIA A100 80GB Tensor Core GPUs	**CPU** 2TB System Memory **GPU** 640GB total GPU Memory **Internal Storage** Internal Storage: 30TB (8x 3.84 TB) U.2 NVMe drives	**Node Cluster** Eight single-port NVIDIA ConnectX-6 **Storage** Up to Two dual-port NVIDIA ConnectX-6 Both Supporting HDR/HDR100/EDR InfiniBand default or 200GigE
DGX A100 Station Workstation	2.5 petaFLOPS AI 5 petaOPS INT8 200GB/s bi-directional bandwidth between any GPU pair, almost 3x compared to PCIe Gen4	**CPU** Single AMD® Epyc® CPU 7742, 2.25GHz to 3.4GHz, 64 cores/128 threads, **GPU** 4x NVIDIA A100 80GB Tensor Core GPUs	**Sys Memory** 512GB DDR4 RDIMM, 3200MT/s **GPU** 320GB total GPU Memory **Internal Storage** 1 x 7.68TB U.2 NVME	Dual 10GBASE-T (RJ45) 4x Mini DisplayPort for display out• Remote management1GbE LAN port (RJ45)