


これからの計算工学に NVIDIA GPUがもたらすものとは

プロメテックグループ
(プロメテック・ソフトウェア/GDEPソリューションズ)

NVIDIA






これからの計算工学に NVIDIA GPU がもたらすものとは - NVIDIA HPC SDK と NVIDIA Modulus の紹介 -

Shinnosuke Furuya, Ph.D., HPC Developer Relations

TOP500 List

Top 500 Supercomputer with Accelerator



<https://top500.org/>

Applications Accelerated on NVIDIA Platforms

Applications Accelerated on NVIDIA Platforms

The Accelerated Apps Catalog features DPU- and GPU-accelerated solutions. Find applications, developer tools, plugins, and more for AI, data science, design, and beyond and discover how they benefit from the latest NVIDIA technologies.

Workload Industry Type Acceleration AI Accelerated ☐ Sort A-Z Share

Search apps

Display 15 per page

1 - 15 of 1299 items

1

Shenzhen Rayvision Technology Co Ltd

3D CAT.live

3D CAT.live is a real-time rendering cloud service for 3D applications that processes heavy image rendering calculations and streams output to the terminal device synchronously.

- > Cloud XR SDK

Sunvega

3D Cloud Design

Online 3D home furnishing interior design service. Multimedia content restores the true texture and provides thousands of customers with the look and feel they have always dreamed of.

- > Minimalist operation, easy to use
- > Intelligent identification of apartment types
- > Smart placement
- > RTX real-time ray tracing rendering technology produces 2K images in 1 minute, which subversively accelerates rendering

Slicer

3D Slicer

3D Slicer is an open-source software platform for medical image informatics, image processing, and three-dimensional visualization. Slicer brings free, powerful cross-platform processing tools to physicians, researchers, and the general public.

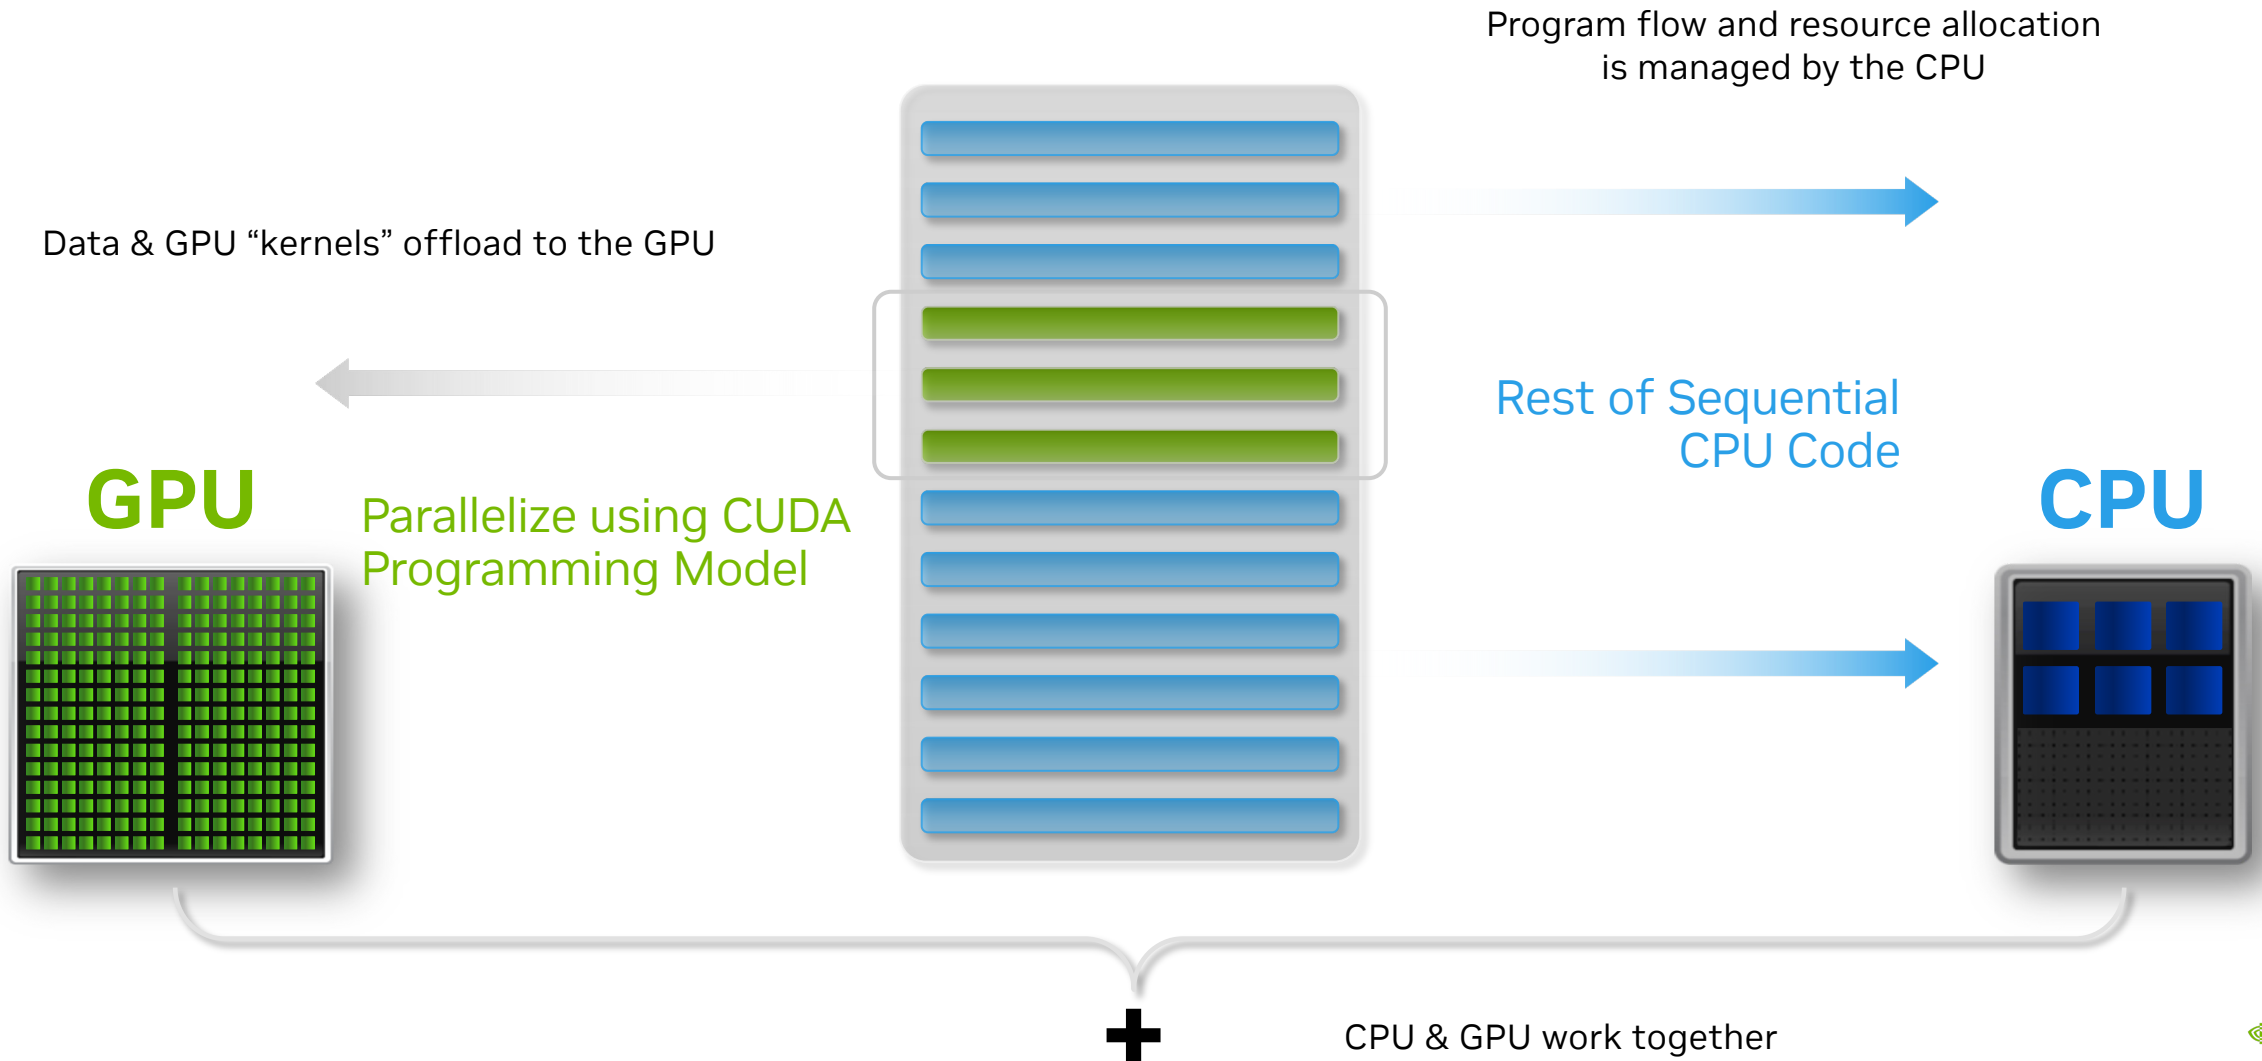
- > NVIDIA Clara AI-assisted Annotation
- > Supports multi organs, from head to toe
- > Multi-modality imaging (MRI, CT, US, nuclear medicine, and microscopy)
- > Bidirectional interface for devices

The background features a complex pattern of glowing green lines and shapes against a solid black field. On the left, numerous thin, parallel lines radiate outwards. On the right, there are larger, more intricate structures that resemble stylized, glowing green leaves or perhaps a microscopic view of a plant, with multiple layers and sharp edges. The overall effect is one of high-tech, futuristic energy.

NVIDIA HPC SDK

GPU Computing in a Nutshell

All GPU Programming Models Follow This Pattern



Programming the NVIDIA Platform

CPU, GPU, and Network

ACCELERATED STANDARD LANGUAGES

ISO C++, ISO Fortran

```
std::transform(par, x, x+n, y, y,  
    [=](float x, float y){ return y + a*x; }  
);
```

```
do concurrent (i = 1:n)  
    y(i) = y(i) + a*x(i)  
enddo
```

```
import cunumeric as np  
...  
def saxpy(a, x, y):  
    y[:] += a*x
```

INCREMENTAL PORTABLE OPTIMIZATION

OpenACC, OpenMP

```
#pragma acc data copy(x,y) {  
    ...  
    std::transform(par, x, x+n, y, y,  
        [=](float x, float y){  
            return y + a*x;  
        });  
    ...  
}  
  
#pragma omp target data map(x,y) {  
    ...  
    std::transform(par, x, x+n, y, y,  
        [=](float x, float y){  
            return y + a*x;  
        });  
    ...  
}
```

PLATFORM SPECIALIZATION

CUDA

```
__global__  
void saxpy(int n, float a,  
    float *x, float *y) {  
    int i = blockIdx.x*blockDim.x +  
        threadIdx.x;  
    if (i < n) y[i] += a*x[i];  
}  
  
int main(void) {  
    ...  
    cudaMemcpy(d_x, x, ...);  
    cudaMemcpy(d_y, y, ...);  
  
    saxpy<<<(N+255)/256,256>>>(...);  
  
    cudaMemcpy(y, d_y, ...);  
}
```

ACCELERATION LIBRARIES

Core

Math

Communication

Data Analytics

AI


Quantum

NVIDIA Math Libraries


Linear Algebra, FFT, RNG and Basic Math




cuBLAS




cuSPARSE




cuTENSOR




cuSOLVER




CUTLASS



AMGX




cuRAND



cuFFT



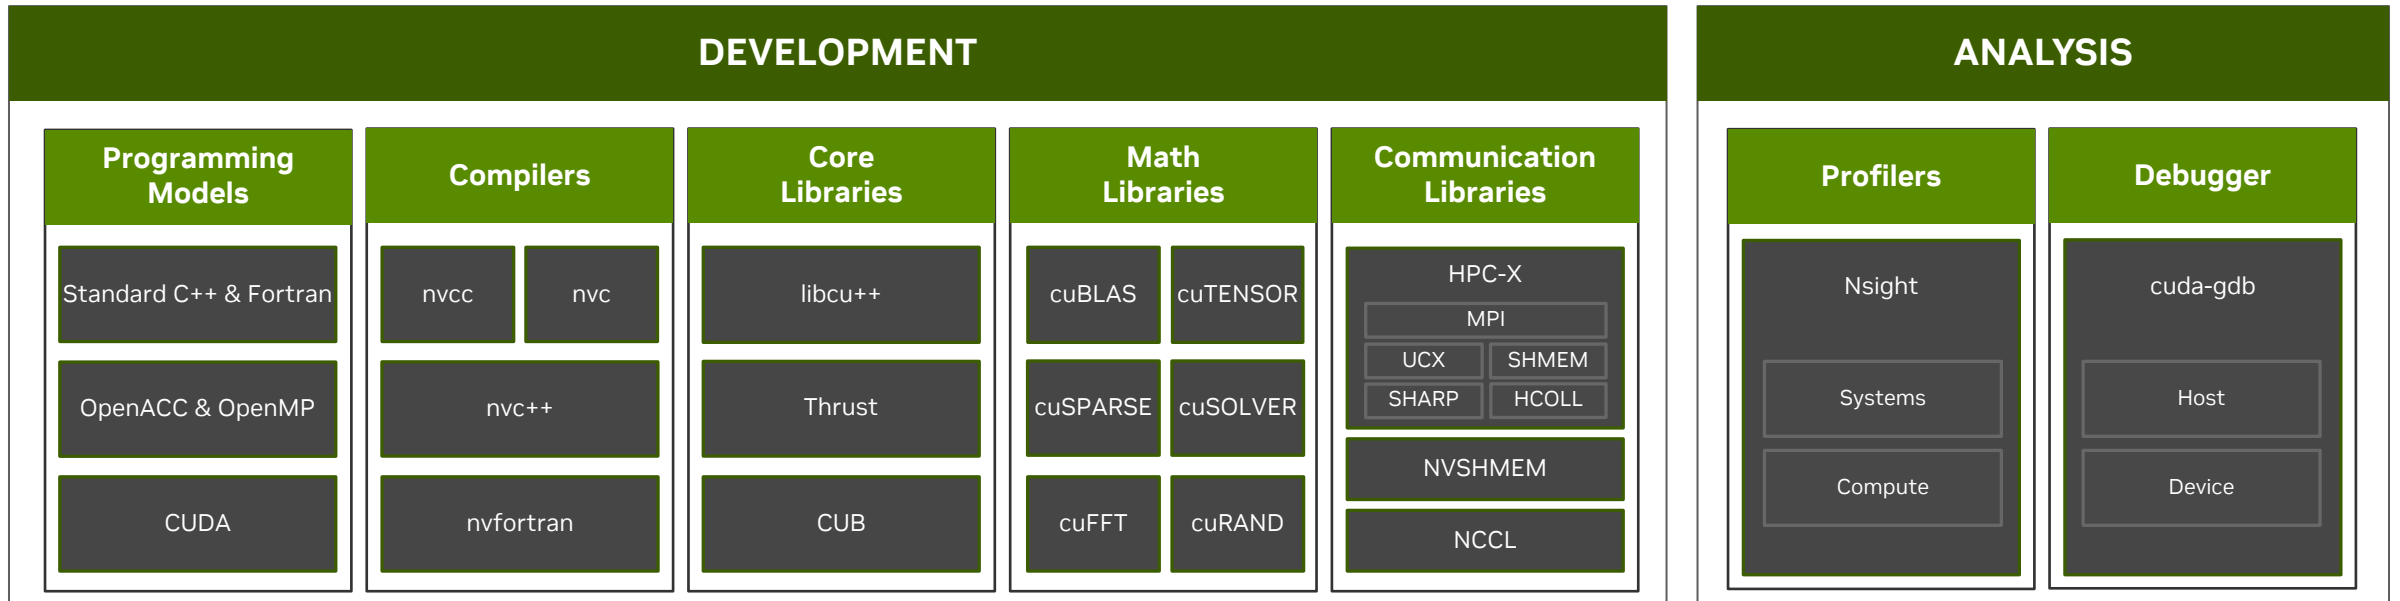
CUDA Math API



Legate

NVIDIA HPC SDK

Available at developer.nvidia.com/hpc-sdk, on NGC, via Spack, and in the Cloud



Develop for the NVIDIA Platform: GPU, CPU and Interconnect
Libraries | Accelerated C++ and Fortran | Directives | CUDA
x86_64 | AArch64 | OpenPOWER
7-8 Releases Per Year | Freely Available

Choose a Programming Model

They can be ~~only~~ **more than** one

Libraries	Standard Languages	Compiler Directives	CUDA Languages
<ul style="list-style-type: none">• Accelerate common operations with little/no code changes• Expert-tuned performance• Forward support guarantees	<ul style="list-style-type: none">• Strong cross-platform support• Single source code for multiple platforms• Reduced learning curve	<ul style="list-style-type: none">• High cross-platform support• Single source code for multiple platforms• Reduced learning curve• Additional programmer control	<ul style="list-style-type: none">• Exposes full GPU capabilities• Trades portability for performance• Distinct GPU/CPU code paths• Full programmer control
Programmer Productivity			Programmer Control



By design these approaches are interoperable,
so you can choose the right balance for your needs

Magnetohydrodynamics Simulation

Eliminating Compute Bottleneck with cuFFT


MHD3D with cuFFT

- Incompressible Hall MHD sim bottleneck is 3D FFTs
- MultiGPU offers speedups of 13x and 21x
- 1D FFT of FFTW replaced with cuFFT
- Accelerate transpose operations from Alltoallv to NCCL
- Bottleneck moved to visualization




*Wisteria-Aquarius – Xeon Platinum 8360Y | A100 40GB

Getting Started with NVIDIA HPC SDK



Home > High Performance Computing > High Performance Computing HPC SDK



NVIDIA HPC SDK

A Comprehensive Suite of Compilers, Libraries and Tools for HPC

The NVIDIA HPC Software Development Kit (SDK) includes the proven compilers, libraries and software tools essential to maximizing developer productivity and the performance and portability of HPC applications.


DEVELOPMENT					ANALYSIS	
Programming Models	Compilers		Core Libraries	Math Libraries	Communication Libraries	
Standard C++ & Fortran	nvcc	nvfortran	libxc++	cuBLAS	cuTENSOR	Open MPI
OpenACC & OpenMP	nvcc	nvfortran	libxc++	cuSPARSE	cuSOLVER	NCCL
CUDA	nvcc	nvfortran	libxc++	cuFFT	cuRAND	NCCL

The NVIDIA HPC SDK C, C++, and Fortran compilers support GPU acceleration of HPC modeling and simulation applications with standard C++ and Fortran, OpenACC® directives, and CUDA®. GPU-accelerated math libraries maximize performance on common HPC algorithms, and optimized communications libraries enable standards-based multi-GPU and scalable systems programming. Performance profiling and debugging tools simplify porting and optimization of HPC applications, and containerization tools enable easy deployment on-premises or in the cloud. With support for NVIDIA GPUs and Arm, OpenPOWER, or x86-64 CPUs running Linux, the HPC SDK provides the tools you need to build NVIDIA GPU-accelerated HPC applications.

[Download Now](#)

[Get Container](#)

Go to <https://developer.nvidia.com/hpc-sdk> and click “Download Now”



Home > High Performance Computing > NVIDIA HPC SDK > NVIDIA HPC SDK Current Release Downloads

NVIDIA HPC SDK Current Release Downloads

Select Target Platform

Click on the green buttons that describe your target platform. Only supported platforms will be shown. By downloading and using the software, you agree to fully comply with the terms and conditions of the [HPC SDK Software License Agreement](#).

☐ I accept the license agreement

[Documentation](#)

[HPC SDK Releases](#)

Read the “HPC SDK Software License Agreement”, click “I accept the license agreement”, and proceed to select your platform, etc.

Getting Started with NVIDIA HPC SDK


Resources

- Latest Version: 23.5
- Product Page
 - <https://developer.nvidia.com/hpc-sdk>
- NGC Container Image
 - <https://catalog.ngc.nvidia.com/orgs/nvidia/containers/nvhpc>
- Documentation
 - <https://docs.nvidia.com/hpc-sdk/index.html>

NVIDIA Modulus




AI Powered Computational Domains



Computational Eng.

Solid & Fluid Mechanics,
Electromagnetics,
Thermal, Acoustics,
Optics, Electrical,
Multi-body Dynamics,
Design Materials,
Systems




Earth Sciences

Climate Modeling,
Weather Modeling,
Ocean Modeling,
Seismic Interpretation




Life Sciences

Genomics,
Proteomics



Computational Physics

Particle Science,
Astrophysics




Computational Chemistry

Quantum Chemistry,
Molecular Dynamics

Process/Product Design,
Manufacturing, Testing,
In-Service

Developing Digital Twins with Physics-ML


Industrial Digital Twins



Siemens Energy HRSG
PINN | Coupled Flows, Physics



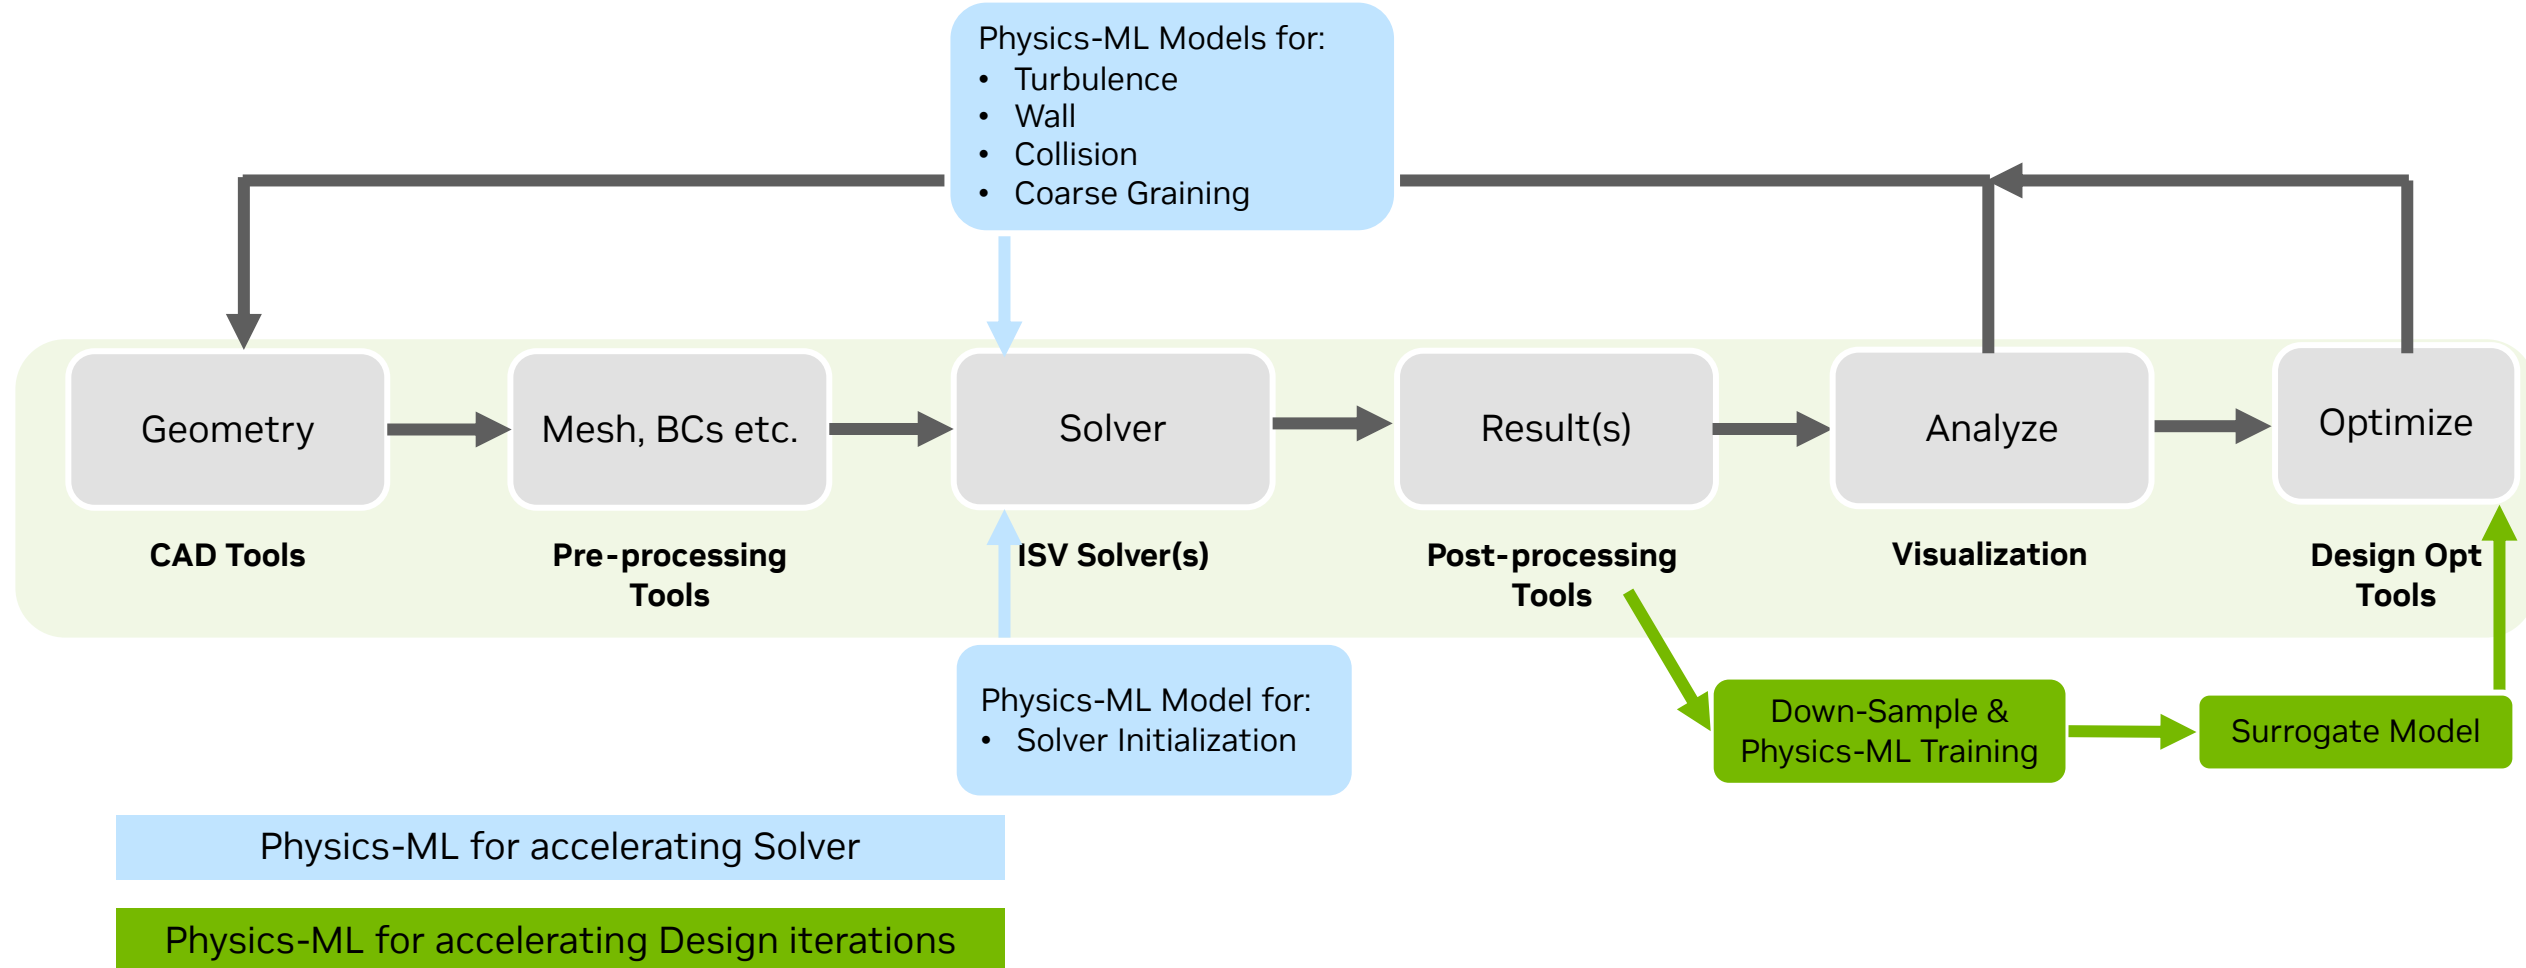
Siemens Gamesa Windfarm
PINN/GAN | Super-Resolution



NETL Power Plant Boiler
PINN | Multi-Physics, Custom Training

Simulation Acceleration with Modulus

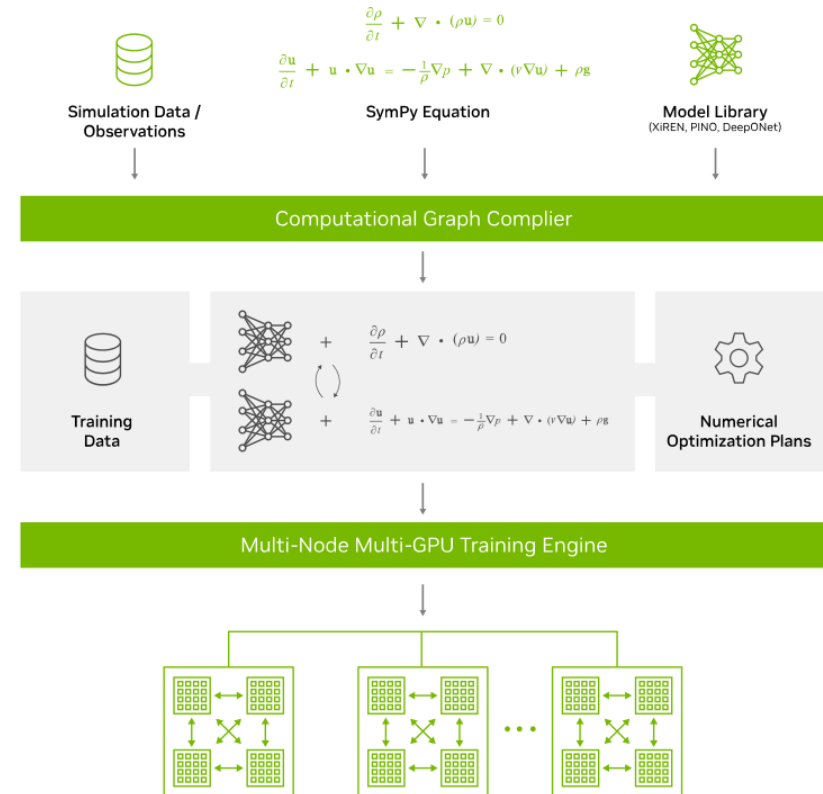
Physics-ML Augmented Simulation Workflows



Open-Source Toolkit for Physics-ML

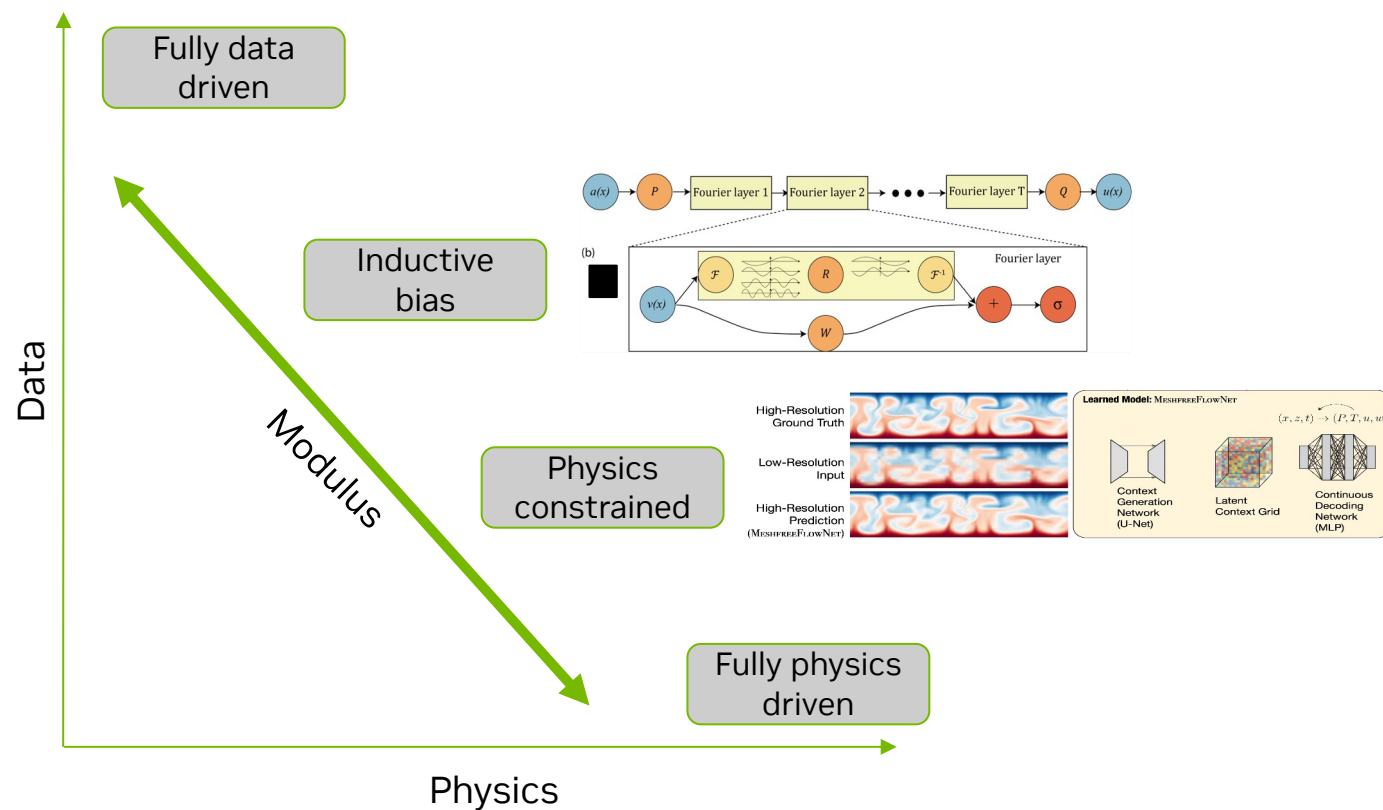
NVIDIA Modulus

- A customizable platform - training and inference pipeline - using Physics (governing equations) and Data (simulation/observations)
- Python based APIs for ease of use
- Facilitates open collaboration within the Physics-ML scientific community
- Well documented features and functionality for ease of use
- Open-source code – easier to understand and customize
- Import PyTorch models from research for your custom application
- Source code:
 - <https://github.com/NVIDIA/Modulus>
 - <https://github.com/NVIDIA/modulus-launch>
 - <https://github.com/NVIDIA/modulus-toolchain>
 - <https://github.com/NVIDIA/modulus-sym>



Open-Source Toolkit for Physics-ML

Novel NN architectures



- Diverse Physics-ML approaches - Model Zoo:
 - PDE driven Physics-ML recipes
 - Data driven Physics-ML recipes
 - Hybrid (Data + PDE) Physics-ML recipes
- PDE Driven - PINNs:
 - Fourier Feature Network
 - Spatial-temporal Fourier Feature Networks
 - Super Resolution Net ...
- Data Driven - Neural Operators:
 - Fourier Neural Operator family (FNO, AFNO, Nested)
 - DeepONet
- GNNs:
 - MeshGraphNet
 - GraphCast
- Hybrid: PINO, ..

HRSG FLUID ACCELERATED CORROSION SIMULATION — SIEMENS ENERGY

Use Case

- Detecting and predicting point of corrosion in heat recovery steam generators (HRSGs)

Challenges

- Using standard simulation to detect corrosion, it took SE at least couple of weeks, and the overall process took 14-16 weeks for every HRSG unit.

Solution

- Using NVIDIA Modulus Physics-Informed Neural Network, SE simulates the corrosive effects of heat, water and other conditions on metal over time to fine-tune maintenance needs.
- SE can replicate and deploy HRSG plant digital twins worldwide with NVIDIA Omniverse.


NVIDIA Solution Stack

- Hardware: NVIDIA V100 & A100 Tensor Core GPUs
- Software: NVIDIA Modulus, NVIDIA Omniverse

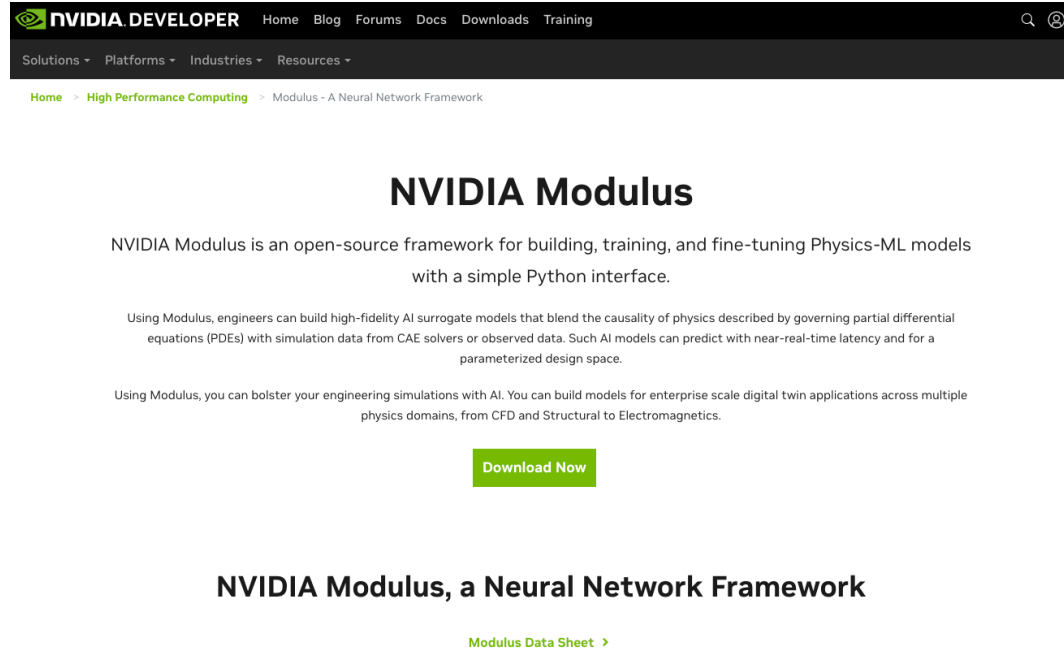
Outcome

- 10,000X speed-up and inference in seconds can reduce downtime by 70%, saving the industry \$1.7 billion annually

[Link to Demo](#)



Getting Started with NVIDIA Modulus



The screenshot shows the NVIDIA Developer website. The header includes the NVIDIA Developer logo and navigation links: Home, Blog, Forums, Docs, Downloads, Training. Below the header is a breadcrumb trail: Home > High Performance Computing > Modulus - A Neural Network Framework. The main heading is "NVIDIA Modulus". Below it, a paragraph states: "NVIDIA Modulus is an open-source framework for building, training, and fine-tuning Physics-ML models with a simple Python interface." Another paragraph follows: "Using Modulus, engineers can build high-fidelity AI surrogate models that blend the causality of physics described by governing partial differential equations (PDEs) with simulation data from CAE solvers or observed data. Such AI models can predict with near-real-time latency and for a parameterized design space." A third paragraph says: "Using Modulus, you can bolster your engineering simulations with AI. You can build models for enterprise scale digital twin applications across multiple physics domains, from CFD and Structural to Electromagnetics." A green "Download Now" button is visible. At the bottom, the text "NVIDIA Modulus, a Neural Network Framework" is displayed, followed by a link "Modulus Data Sheet >".

NVIDIA Modulus

NVIDIA Modulus is an open-source framework for building, training, and fine-tuning Physics-ML models with a simple Python interface.

Using Modulus, engineers can build high-fidelity AI surrogate models that blend the causality of physics described by governing partial differential equations (PDEs) with simulation data from CAE solvers or observed data. Such AI models can predict with near-real-time latency and for a parameterized design space.

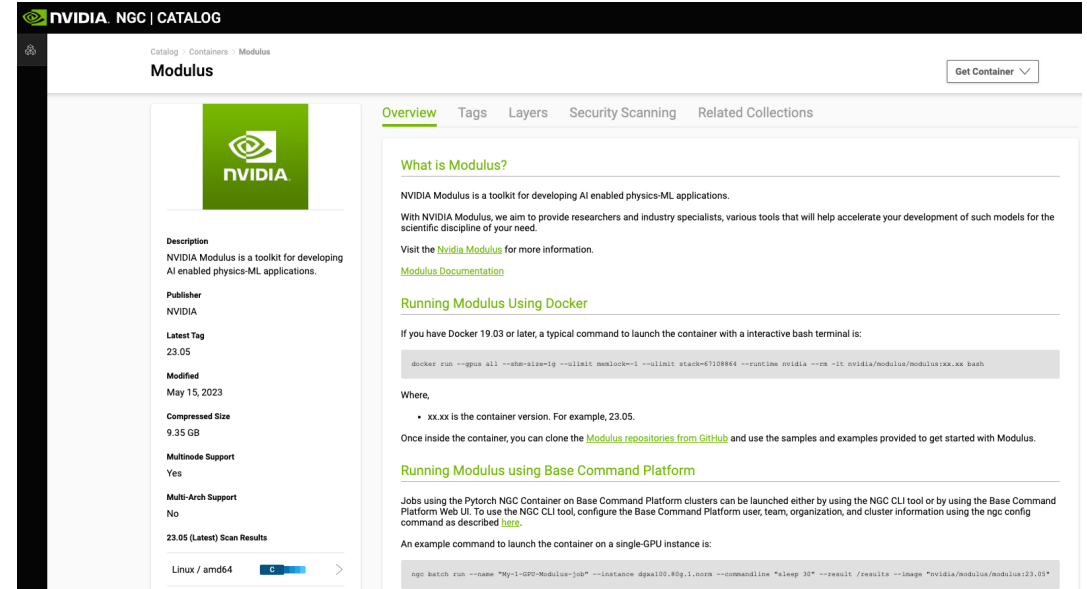
Using Modulus, you can bolster your engineering simulations with AI. You can build models for enterprise scale digital twin applications across multiple physics domains, from CFD and Structural to Electromagnetics.

[Download Now](#)

NVIDIA Modulus, a Neural Network Framework

[Modulus Data Sheet >](#)

Go to <https://developer.nvidia.com/modulus> and click “Download Now”



The screenshot shows the NVIDIA NGC Catalog page for the Modulus container. The header includes the NVIDIA NGC | CATALOG logo and navigation links: Catalog, Containers, Modulus. Below the header is a breadcrumb trail: Catalog > Containers > Modulus. The main heading is "Modulus". To the right of the heading is a "Get Container" button. Below the heading is a sidebar with the NVIDIA logo and a description: "NVIDIA Modulus is a toolkit for developing AI enabled physics-ML applications." The sidebar also lists the publisher as NVIDIA, the latest tag as 23.05, the modified date as May 15, 2023, the compressed size as 9.35 GB, and multi-arch support as Yes. The main content area has tabs for Overview, Tags, Layers, Security Scanning, and Related Collections. The Overview tab is selected. It contains the following sections: "What is Modulus?" (NVIDIA Modulus is a toolkit for developing AI enabled physics-ML applications. With NVIDIA Modulus, we aim to provide researchers and industry specialists, various tools that will help accelerate your development of such models for the scientific discipline of your need. Visit the [Nvidia Modulus](#) for more information. [Modulus Documentation](#)), "Running Modulus Using Docker" (If you have Docker 19.03 or later, a typical command to launch the container with an interactive bash terminal is: `docker run --gpus all --shm-size=1g --ulimit memlock=-1 --ulimit stack=67108864 --runtime nvidia --rm -it nvidia/modulus/modulus:xx.xx bash`), "Where," (xx.xx is the container version. For example, 23.05. Once inside the container, you can clone the [Modulus repositories from GitHub](#) and use the samples and examples provided to get started with Modulus.), "Running Modulus using Base Command Platform" (Jobs using the Pytorch NGC Container on Base Command Platform clusters can be launched either by using the NGC CLI tool or by using the Base Command Platform Web UI. To use the NGC CLI tool, configure the Base Command Platform user, team, organization, and cluster information using the ngc config command as described [here](#). An example command to launch the container on a single-GPU instance is: `ngc batch run --name "My-1-GPU-Modulus-job" --instance dgsai00.80g.1.ncr --commandline "sleep 30" --result /results --image "nvidia/modulus/modulus:23.05"`).

Modulus

[Get Container](#)

[Overview](#) [Tags](#) [Layers](#) [Security Scanning](#) [Related Collections](#)

What is Modulus?

NVIDIA Modulus is a toolkit for developing AI enabled physics-ML applications.

With NVIDIA Modulus, we aim to provide researchers and industry specialists, various tools that will help accelerate your development of such models for the scientific discipline of your need.

Visit the [Nvidia Modulus](#) for more information.

[Modulus Documentation](#)

Running Modulus Using Docker

If you have Docker 19.03 or later, a typical command to launch the container with an interactive bash terminal is:

```
docker run --gpus all --shm-size=1g --ulimit memlock=-1 --ulimit stack=67108864 --runtime nvidia --rm -it nvidia/modulus/modulus:xx.xx bash
```

Where,

- xx.xx is the container version. For example, 23.05.

Once inside the container, you can clone the [Modulus repositories from GitHub](#) and use the samples and examples provided to get started with Modulus.

Running Modulus using Base Command Platform

Jobs using the Pytorch NGC Container on Base Command Platform clusters can be launched either by using the NGC CLI tool or by using the Base Command Platform Web UI. To use the NGC CLI tool, configure the Base Command Platform user, team, organization, and cluster information using the ngc config command as described [here](#).

An example command to launch the container on a single-GPU instance is:

```
ngc batch run --name "My-1-GPU-Modulus-job" --instance dgsai00.80g.1.ncr --commandline "sleep 30" --result /results --image "nvidia/modulus/modulus:23.05"
```

Follow the instructions described here:

```
$ docker run --gpus all ...
```

Getting Started with NVIDIA Modulus


Resources

- Latest Version : 23.05
- Product Page
 - <https://developer.nvidia.com/modulus>
- NGC Container Image
 - <https://catalog.ngc.nvidia.com/orgs/nvidia/teams/modulus/containers/modulus>
- Documentation
 - <https://docs.nvidia.com/modulus/index.html>
- GitHub
 - <https://github.com/NVIDIA/modulus>
- Japanese Page **<- NEW!**
 - <https://developer.nvidia.com/ja-jp/modulus>
- Resource Center
 - <https://resources.nvidia.com/l/en-us/modulus-pathfactory-explore-page>

「AI サロゲートモデルでシミュレーションを高速化する方法とは？」

ソフトウェア ウェビナー シリーズ Vol. 3

ソフトウェアウェビナーシリーズ 3つのシミュレーション関連 注目技術をわかりやすく解説



現在、機械や建設での設計をはじめ幅広いものづくりにおいて、様々な領域で CAE が用いられています。しかしながら、技術の進歩により、シミュレーションを行うパラメータ数は増え続け、数値解析の結果の出力がただちに得られないケースも増えてきました。

そこでシミュレーションの一部を AI に置き換えるサロゲートモデルの活用が提唱されています。NVIDIA では Physics-ML を開発するためのフレームワーク NVIDIA Modulus をご提供しており、すでに一部の企業がこれを用いて風力発電機やプラントのシミュレーションを行っています。

本ウェビナーでは、Modulus によって何が可能になるのか？得意としている領域とは？ご利用になるための日本語の情報をご紹介します。

【日 程】2023 年 7 月 27 日 (木) 14:00 - 15:00 (60 分)


【対 象】大学や企業で CAE 活用を研究の方、CAE に Physics-ML 導入をご検討の方

【主 催】エヌビディア 合同会社

【参加費】無料 / 事前登録制

【配信方法】ON24 Simulive (Q&A はテキストにてライブでご対応いたします)

【お問い合わせ】NVIDIA セミナー事務局 (NVIDIAJapanEvent@nvidia.com)



丹愛彦

エヌビディア合同会社 ソリューションアーキテクチャ & エンジニアリング

シニアソリューションアーキテクト

製造業の研究所にて数値流体解析の研究開発に従事したのち、エヌビディア合同会社入社。現在は HPC 分野を中心に技術支援を担当。



柴田良一 教授

独立行政法人 国立高等専門学校機構

岐阜高等専門学校 建築学科

建設系機械系を含めた幅広いものづくりを対象に、オープン CAE による構造解析や流体解析、これらの連成解析を研究分野としている。さらに最近は、数値解析技術と人工知能技術との融合に興味を持ち、PINNsの可能性の検証を進めている。

▼詳細はこちら▼

<https://event.on24.com/wcc/r/4217685/BBE6855EE53BEA8B0AA3A67FB8FEC3B6/4725196>



The background is a black field filled with numerous thin, bright green lines that curve and flow across the frame. On the right side, there are larger, more complex green structures that resemble stylized, glowing leaves or perhaps the intricate wiring of a microchip. These structures have a layered, translucent appearance with internal details visible.

NVIDIA GPUs

NVIDIA GPUs at a Glance

	Fermi (2010)	Kepler (2012)	Maxwell (2014)	Pascal (2016)	Volta (2017)	Turing (2018)	Ampere (2020)	Hopper (2022)	Ada Lovelace (2022)
Data Center GPU	Compute (FP64/FP32)	K80		P100	V100		A100 A30	H100	L40 L4
	M2090		M40			T4	A40 A2		
	Compute (FP32)								
	VDI (FP32)	K1	M10				A16		
RTX / Quadro				GP100	GV100				
	ProVis (FP32)	K6000	M6000	P5000		RTX 8000	RTX A6000		RTX 6000 Ada Generation
GeForce				TITAN Xp	TITAN V	TITAN RTX			
	Gaming (FP32)	GTX 780	GTX 980	GTX 1080		RTX 2080 Ti	RTX 3090 Ti		RTX 4090


NVIDIA H100 Tensor Core GPU

- HPC / DL Training / DL Inference / HPDA
 - Exascale HPC / LLM Inference
- Two form factors
 - SXM for HGX / PCIe
- FP64 / FP32
- 4th Generation Tensor Core
 - FP64 / TF32 / BF16 / FP16 / FP8 / INT8
- 4th Generation NVLink
 - 900GB/s (SXM) / 600GB/s up to 2 GPUs via NVLink Bridge (PCIe)
- High-Bandwidth Memory
 - 80GB HBM3 (SXM) / 80GB HBM2e (PCIe) / 188GB HBM3 (NVL; total)
- Transformer Engine
- 2nd Generation Multi-Instance GPU (MIG)



NVIDIA A100 Tensor Core GPU

- HPC / DL Training / DL Inference / HPDA
- Two form factors
 - SXM for HGX / PCIe
- FP64 / FP32
- 3rd Generation Tensor Core
 - FP64 / TF32 / BF16 / FP16 / INT8
- 3rd Generation NVLink
 - 600GB/s (SXM) / 600GB/s up to 2 GPUs via NVLink Bridge (PCIe)
- High-Bandwidth Memory
 - 80GB HBM2e
- Structural Sparsity
- Multi-Instance GPU (MIG)



NVIDIA H100 Gen-to-Gen Comparison

	NVIDIA H100			NVIDIA A100	
GPU Architecture	Hopper			Ampere	
Form Factor	SXM (SXM5)	PCIe (PCIe Gen5)	NVL (2x PCIe Gen5)	SXM (SXM4)	PCIe (PCIe Gen4)
FP64 FP32 TFLOPS	34 67	26 51	2x34 2x67	9.7 19.5	
TF32 TC BF16 TC TFLOPS	494* 989*	378* 756*	2x494* 2x989*	156 312	
FP16 TC FP8 TC TFLOPS	989* 1979*	756* 1513*	2x989* 2x1979*	312 NA	
Memory	80GB HBM3	80GB HBM2e	2x94GB HBM3	80GB HBM2e	80GB HBM2e
Memory Bandwidth	3.35TB/s	2TB/s	2x3.9TB/s	2039GB/s	1935GB/s
Max TDP	Up to 700W (configurable)	300-350W (configurable)	2x 350-400W (configurable)	400W	300W
MIG	Up to 7 @10GB	Up to 7 @10GB	Up to 14 @12GB	Up to 7 @10GB	Up to 7 @10GB
Interconnect	NVLink: 900GB/s PCIe: 128GB/s	NVLink: 600GB/s PCIe: 128GB/s	NVLink: 600GB/s PCIe: 128GB/s	NVLink: 600GB/s PCIe: 64GB/s	NVLink: 600GB/s PCIe: 64GB/s

* Double when using sparsity

NVIDIA RTX 6000 Ada Generation

- Professional Visualization
- 3rd Generation RT Core
- 4th Generation Tensor Core
- 48GB GDDR6 ECC memory
- PCIe Gen4 x16
- 4x DisplayPort 1.4
 - 4x 4096 x 2160 @ 120Hz
 - 4x 5120 x 2880 @ 60Hz
 - 2x 7680 x 4320 @ 60Hz
- Virtualization-Ready



NVIDIA RTX 6000 Ada Generation Gen-to-Gen Comparison

	NVIDIA RTX 6000 Ada Generation	NVIDIA RTX A6000
GPU Architecture	Ada Lovelace	Ampere
CUDA Cores	18176	10752
Tensor Cores	568	336
RT Cores	142	84
Memory Size	48GB GDDR6 ECC	48GB GDDR6 ECC
Memory Bandwidth	960GB/s	768GB/s
NVLink	Not supported	2-way
Virtual Workstation	Yes	Yes
Media Acceleration	3 NVENC (+1 AV1 encode) 3 NVDEC (+1 AV1 decode)	1 NVENC 2 NVDEC (+1 AV1 decode)
Display Connections	4x DP 1.4	4x DP 1.4
Max TDP	300W	300W
Graphics Bus	PCIe Gen4 x16	PCIe Gen4 x16

NVIDIA Japan Social Media Directory

Find Us Online!




- Twitter
 - <https://twitter.com/NVIDIAJapan>
- Facebook
 - <https://www.facebook.com/NVIDIA.JP>
- YouTube
 - <https://www.youtube.com/user/NVIDIAJapan>




- Twitter
 - <https://twitter.com/NVIDIAAIJP>
- Facebook
 - <https://www.facebook.com/NVIDIAAI.JP>



- Facebook
 - <https://www.facebook.com/NVIDIANetworkingJapan>



- Twitter
 - <https://twitter.com/NVIDIAGeForceJP>
- Facebook
 - <https://www.facebook.com/NVIDIAGeForceJP>
- Instagram
 - <https://instagram.com/nvidiageforcejp>
- YouTube
 - <https://www.youtube.com/@nvidiageforcejapan44>
- Twitch
 - <https://www.twitch.tv/nvidiajapan>



- Twitter
 - <https://twitter.com/NVIDIAStudioJP>
- Facebook
 - <https://www.facebook.com/NVIDIAStudioJP>
- Instagram
 - <https://instagram.com/nvidiastudiojp>
- YouTube
 - <https://www.youtube.com/@nvidiastudiojapan1621>

