Séminaire quantique EDF TQCI
11 Janvier 2023

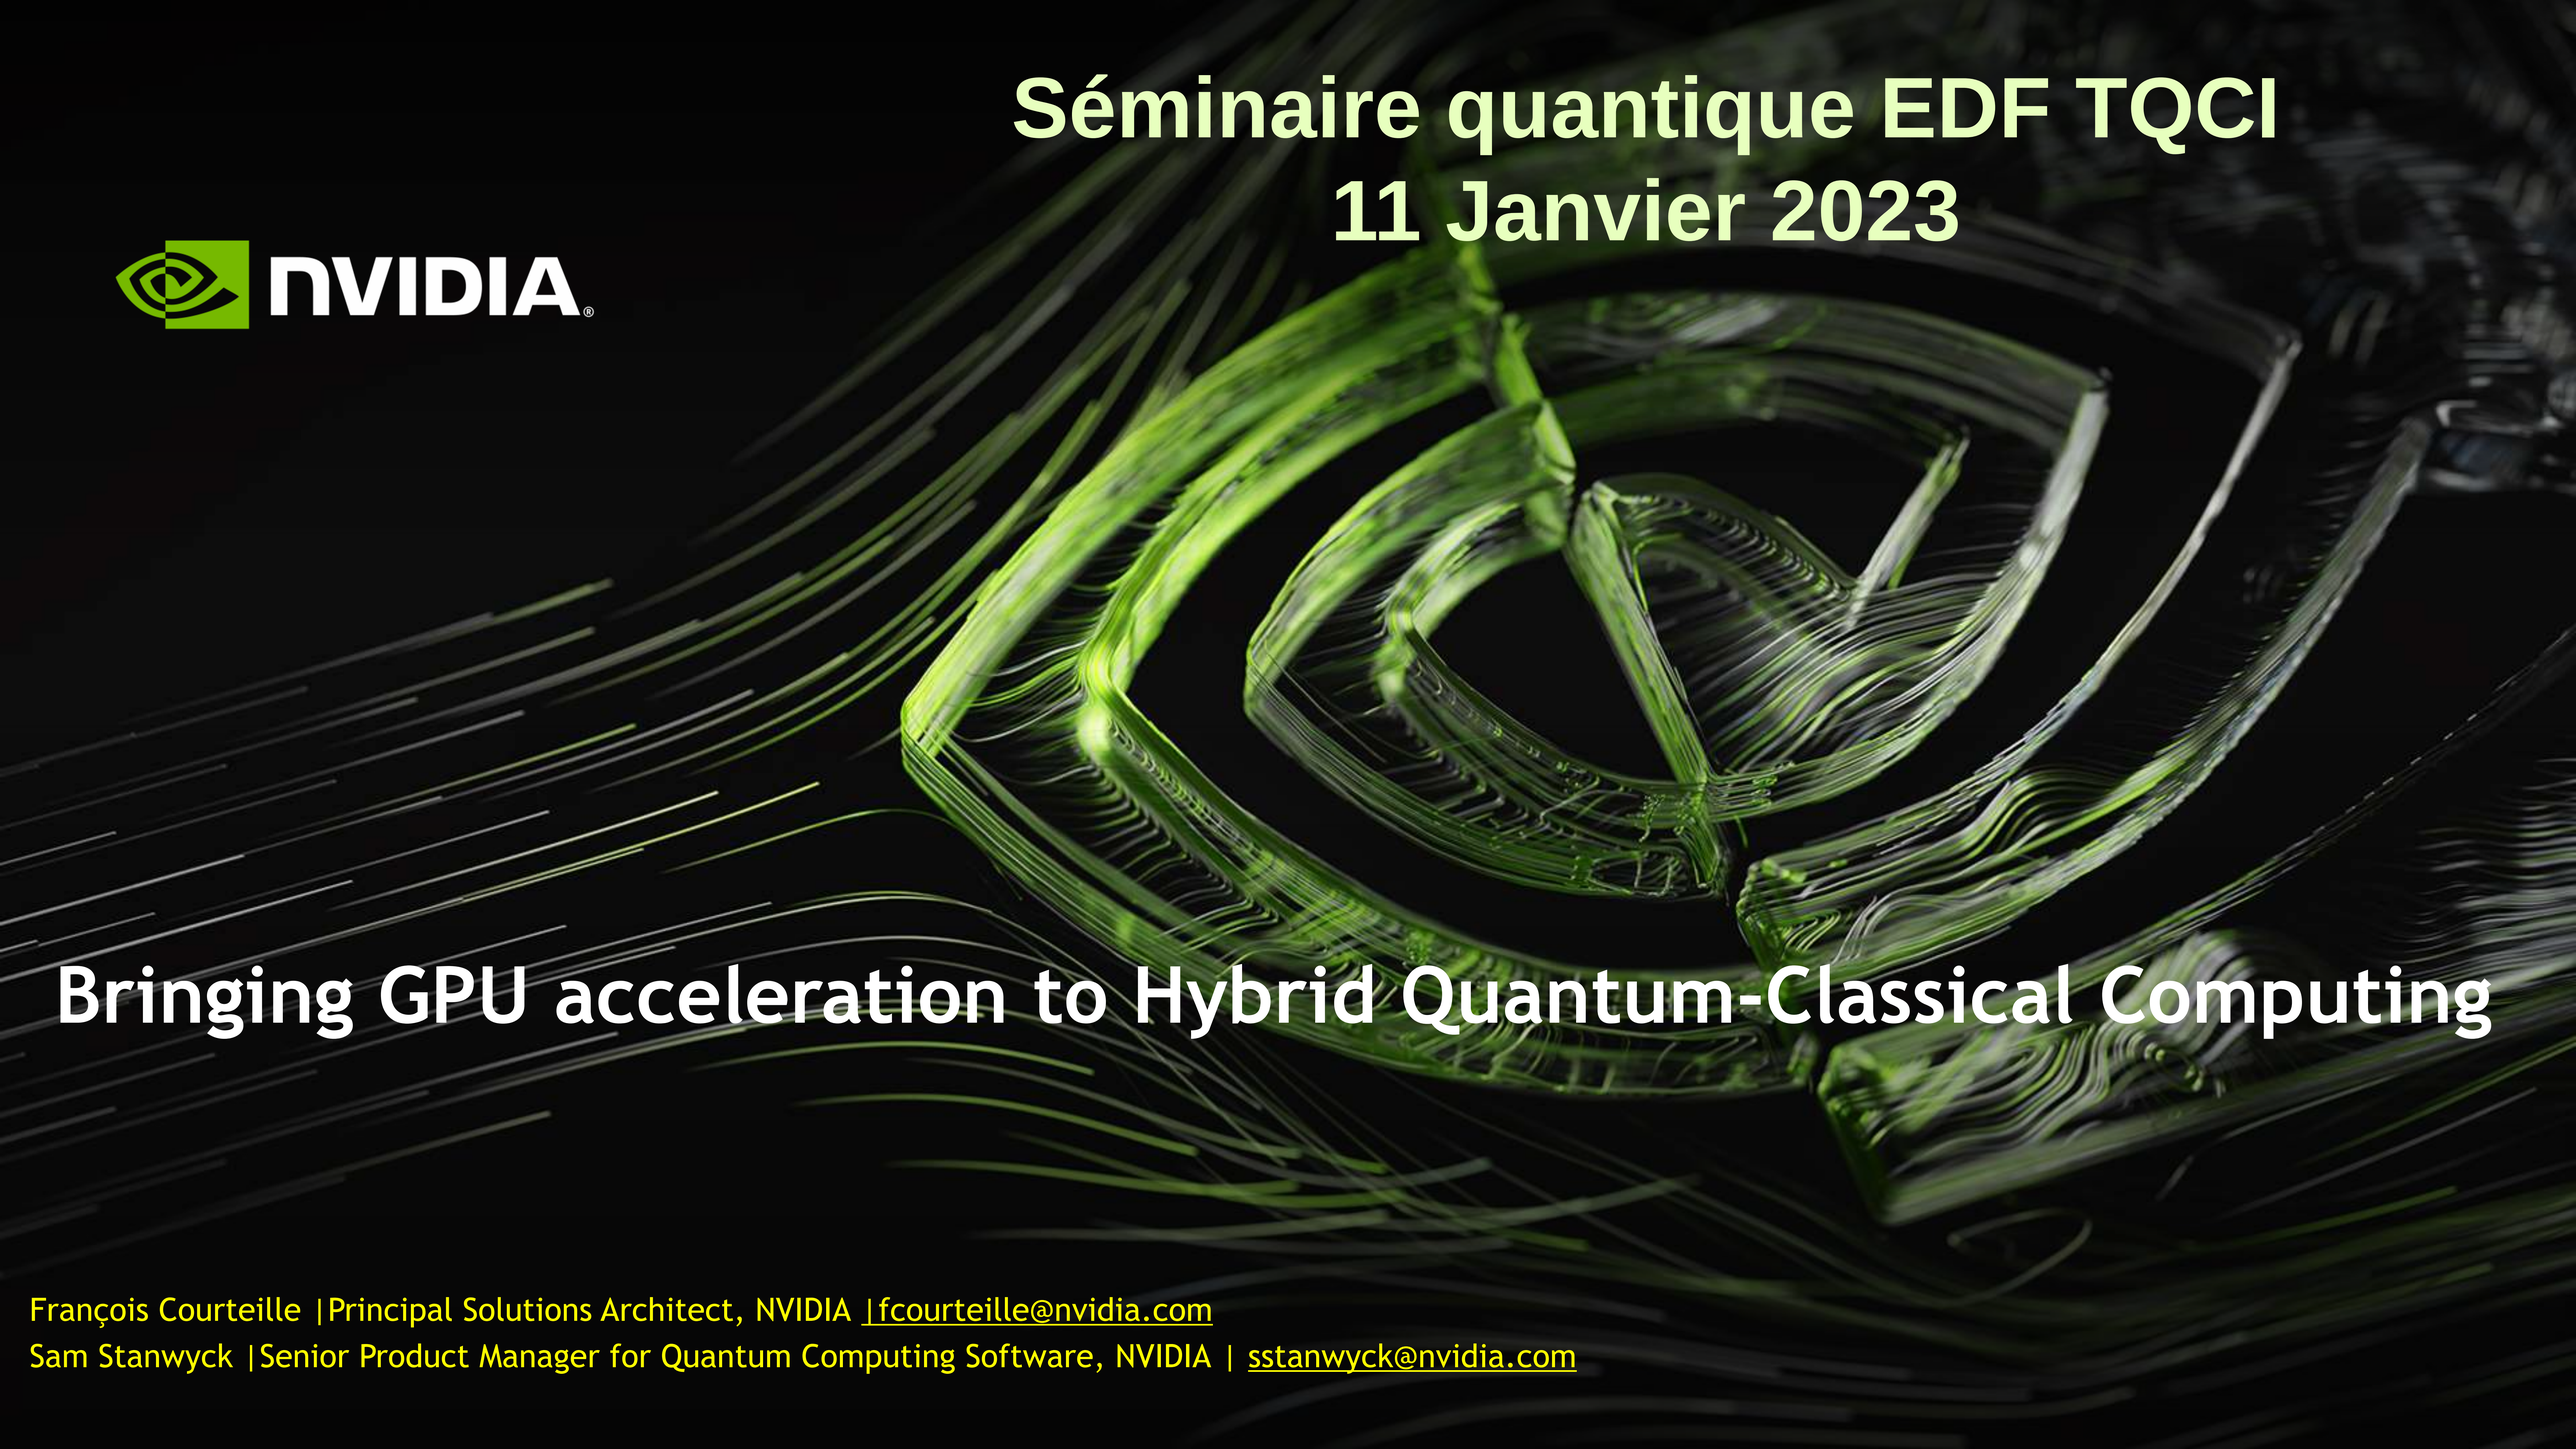

Bringing GPU acceleration to Hybrid Quantum-Classical Computing

François Courteille |Principal Solutions Architect, NVIDIA |fcourteille@nvidia.com
Sam Stanwyck |Senior Product Manager for Quantum Computing Software, NVIDIA | sstanwyck@nvidia.com

POTENTIAL NEAR-TERM QUANTUM USE-CASES

Applications with near-term potential, but quantum advantage is an open question

Quantum Machine Learning

Quantum Generative Models

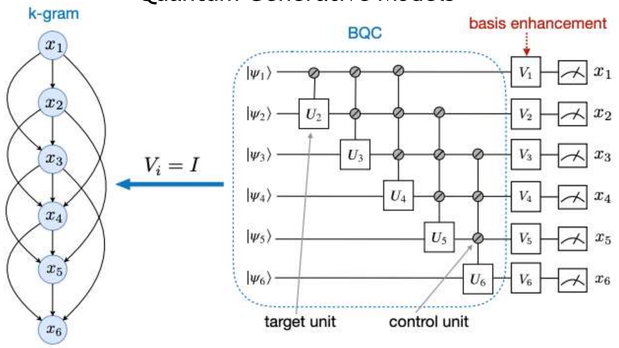

Quantum Support Vector Machine

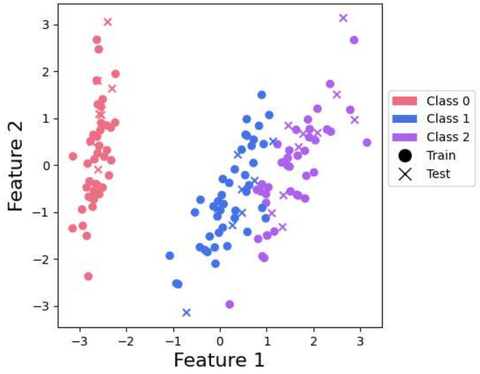

Gao, et al, Phys. Rev. X **12**, 021037
Pennylane.ai

Quantum Chemistry

Variational Quantum Eigensolver for carbon capture

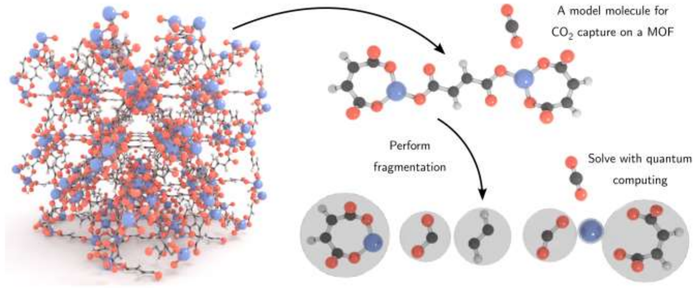

Greene-Diniz, et al, arXiv:2203.15546,

Protein folding

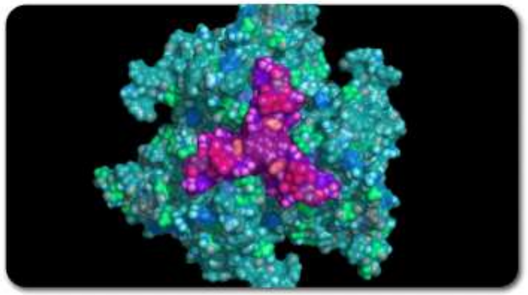

Menten.ai

Combinatorial Optimization

QAOA for resource allocation

Logistics optimization

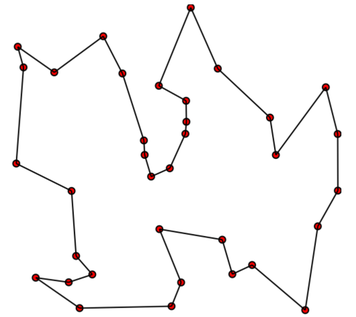

Image from ibm.com
Wikipedia.com

Quantum Computing @ NVIDIA

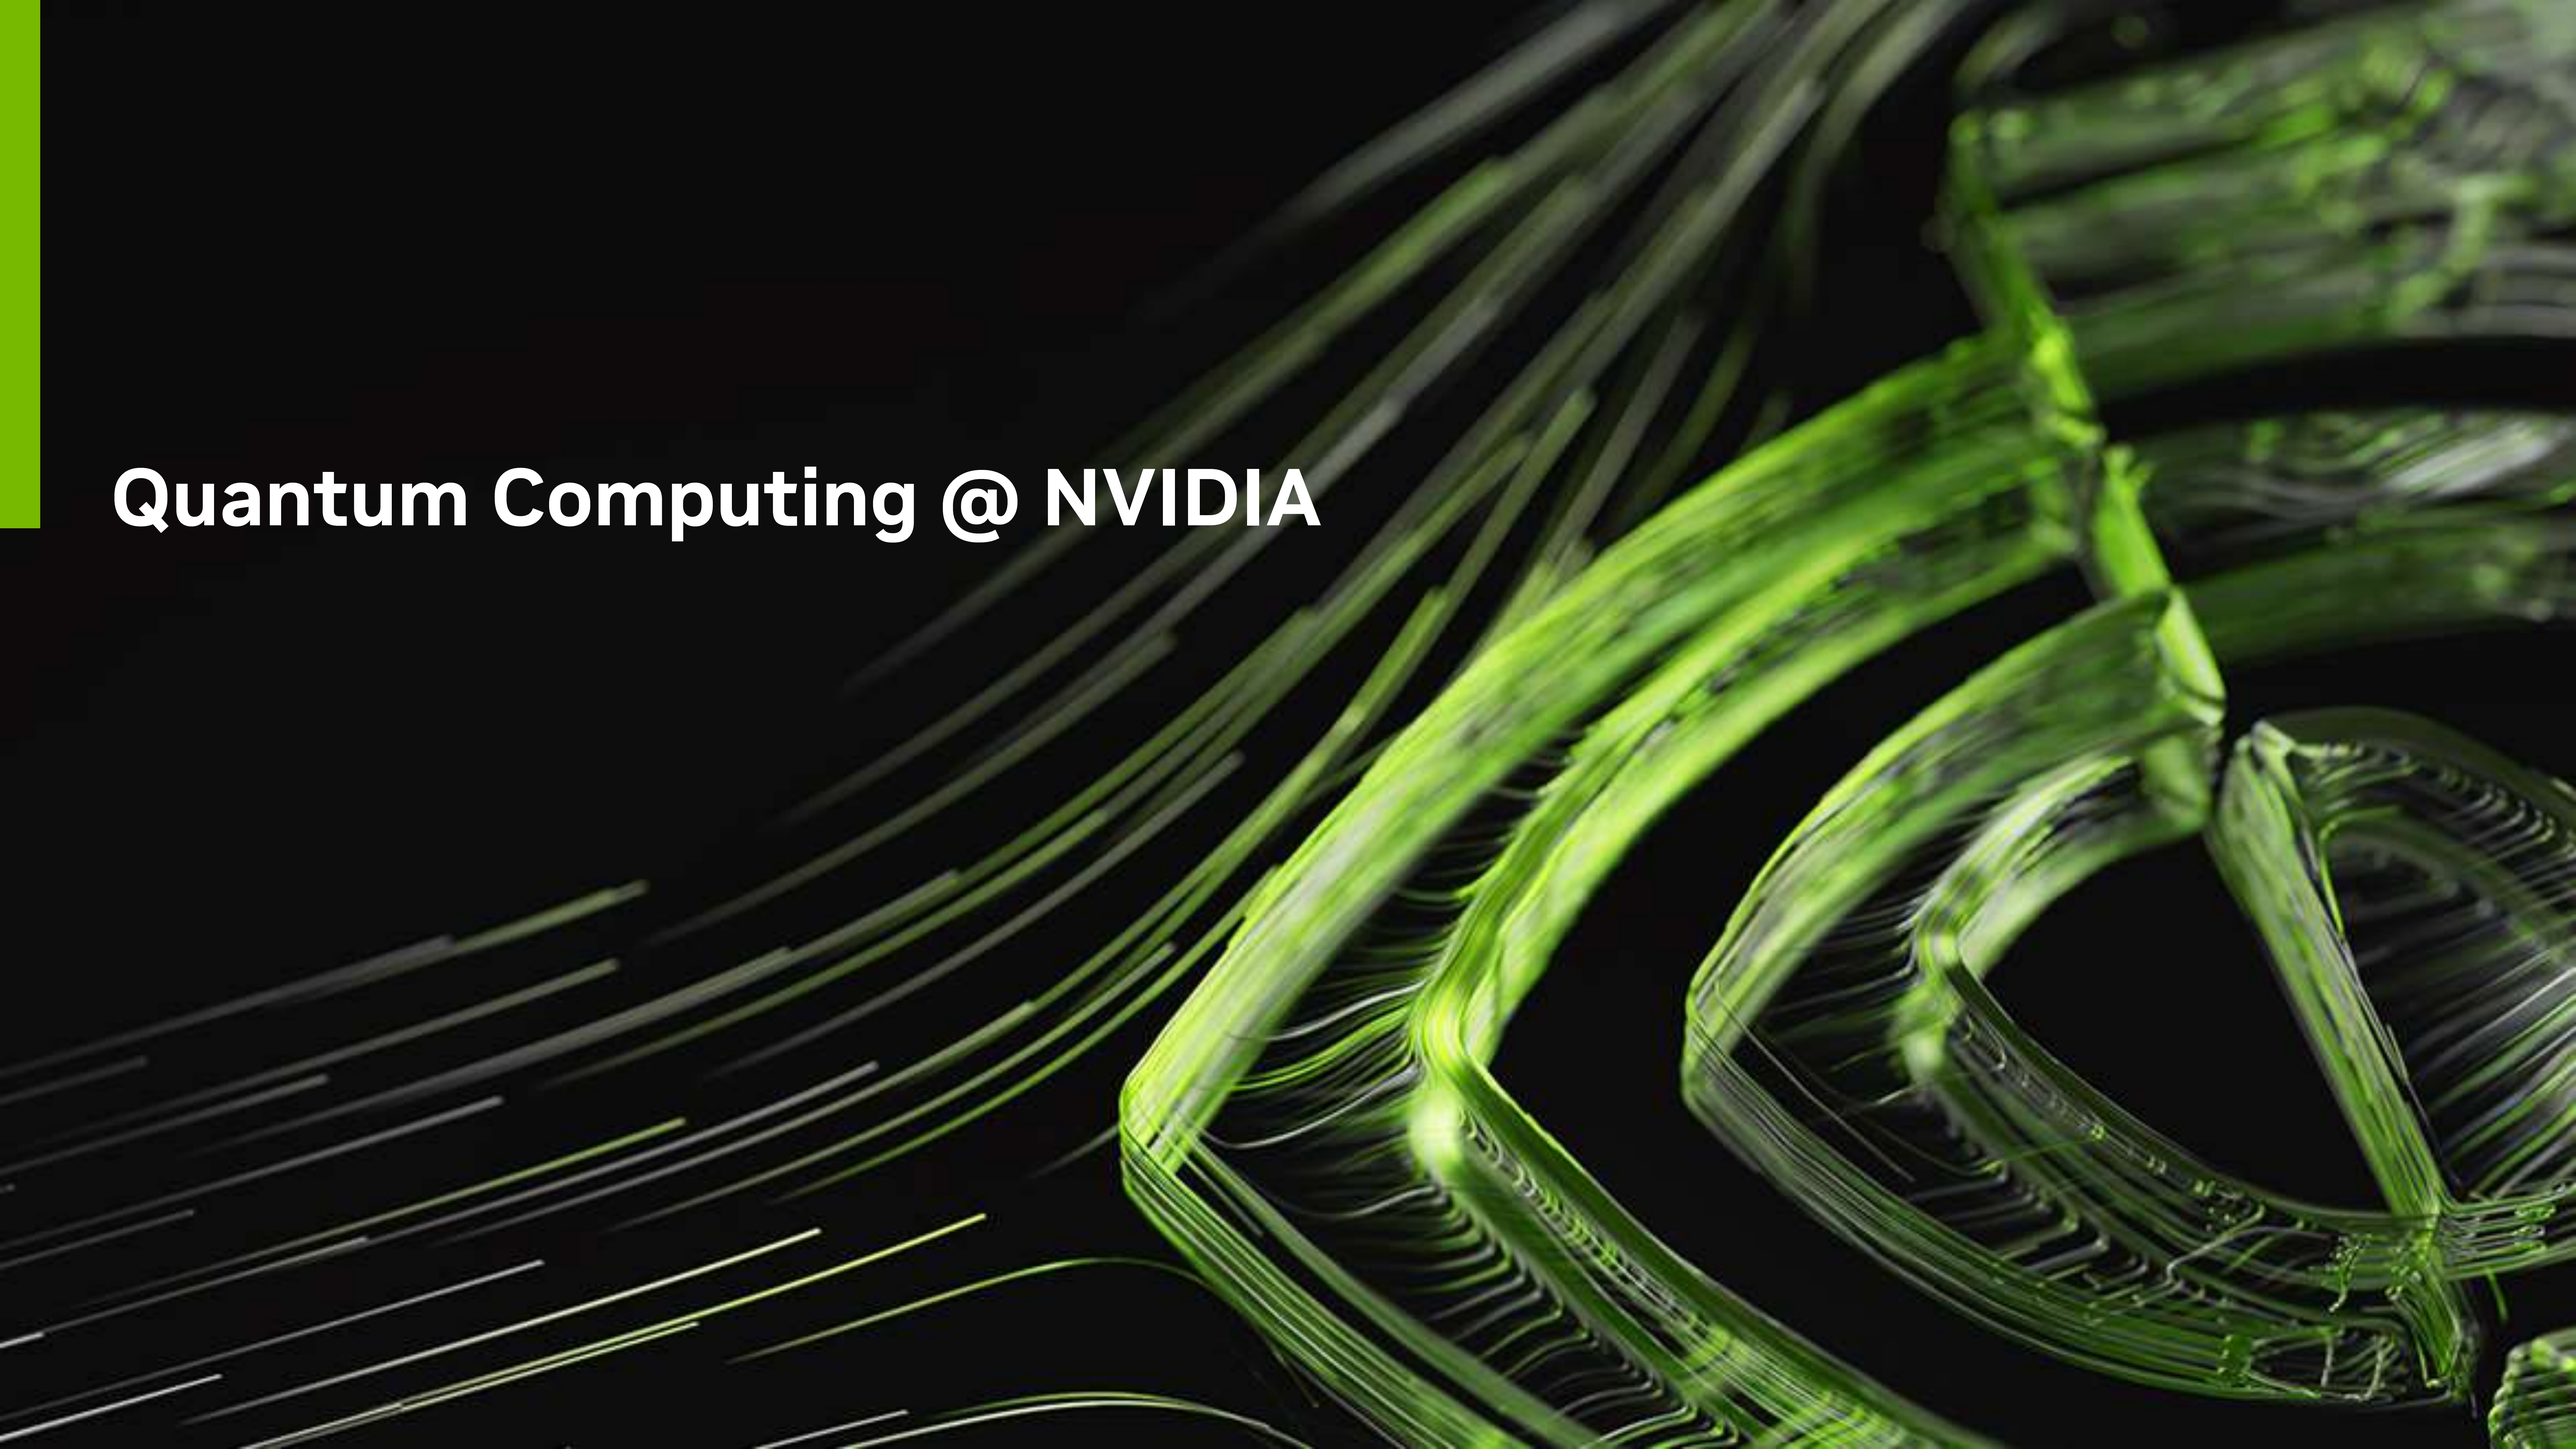

Quantum Computing at NVIDIA

Researching the Quantum Computers of Tomorrow with the Supercomputers of Today

QUANTUM SIMULATION - CUQUANTUM

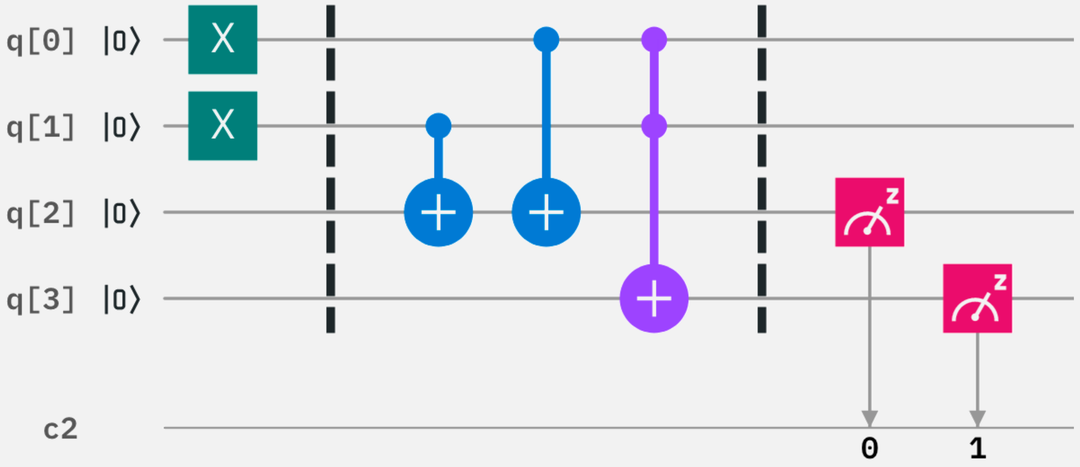

- Develop algorithms at scale of valuable quantum computing
- Discover use cases with quantum advantage
- Design and validate future hardware

HYBRID QUANTUM-CLASSICAL COMPUTING - QODA

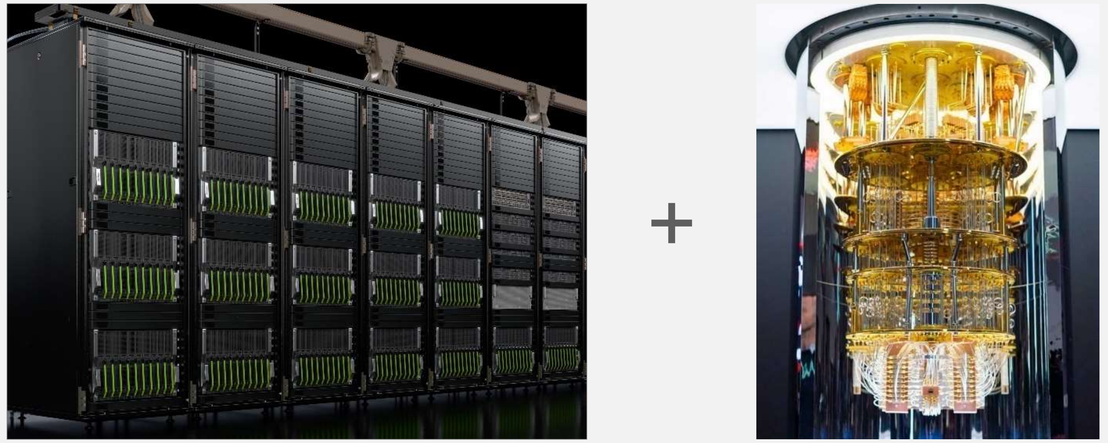

- Develop quantum applications by integrating quantum into leading accelerated applications
- Build a platform that is familiar to domain scientists
- Unparalleled performance and scientific productivity using the best resource for the task

POWERFUL SIMULATIONS WITH cuQuantum

Researching & Developing the Computers of Tomorrow Requires Powerful Simulations Today

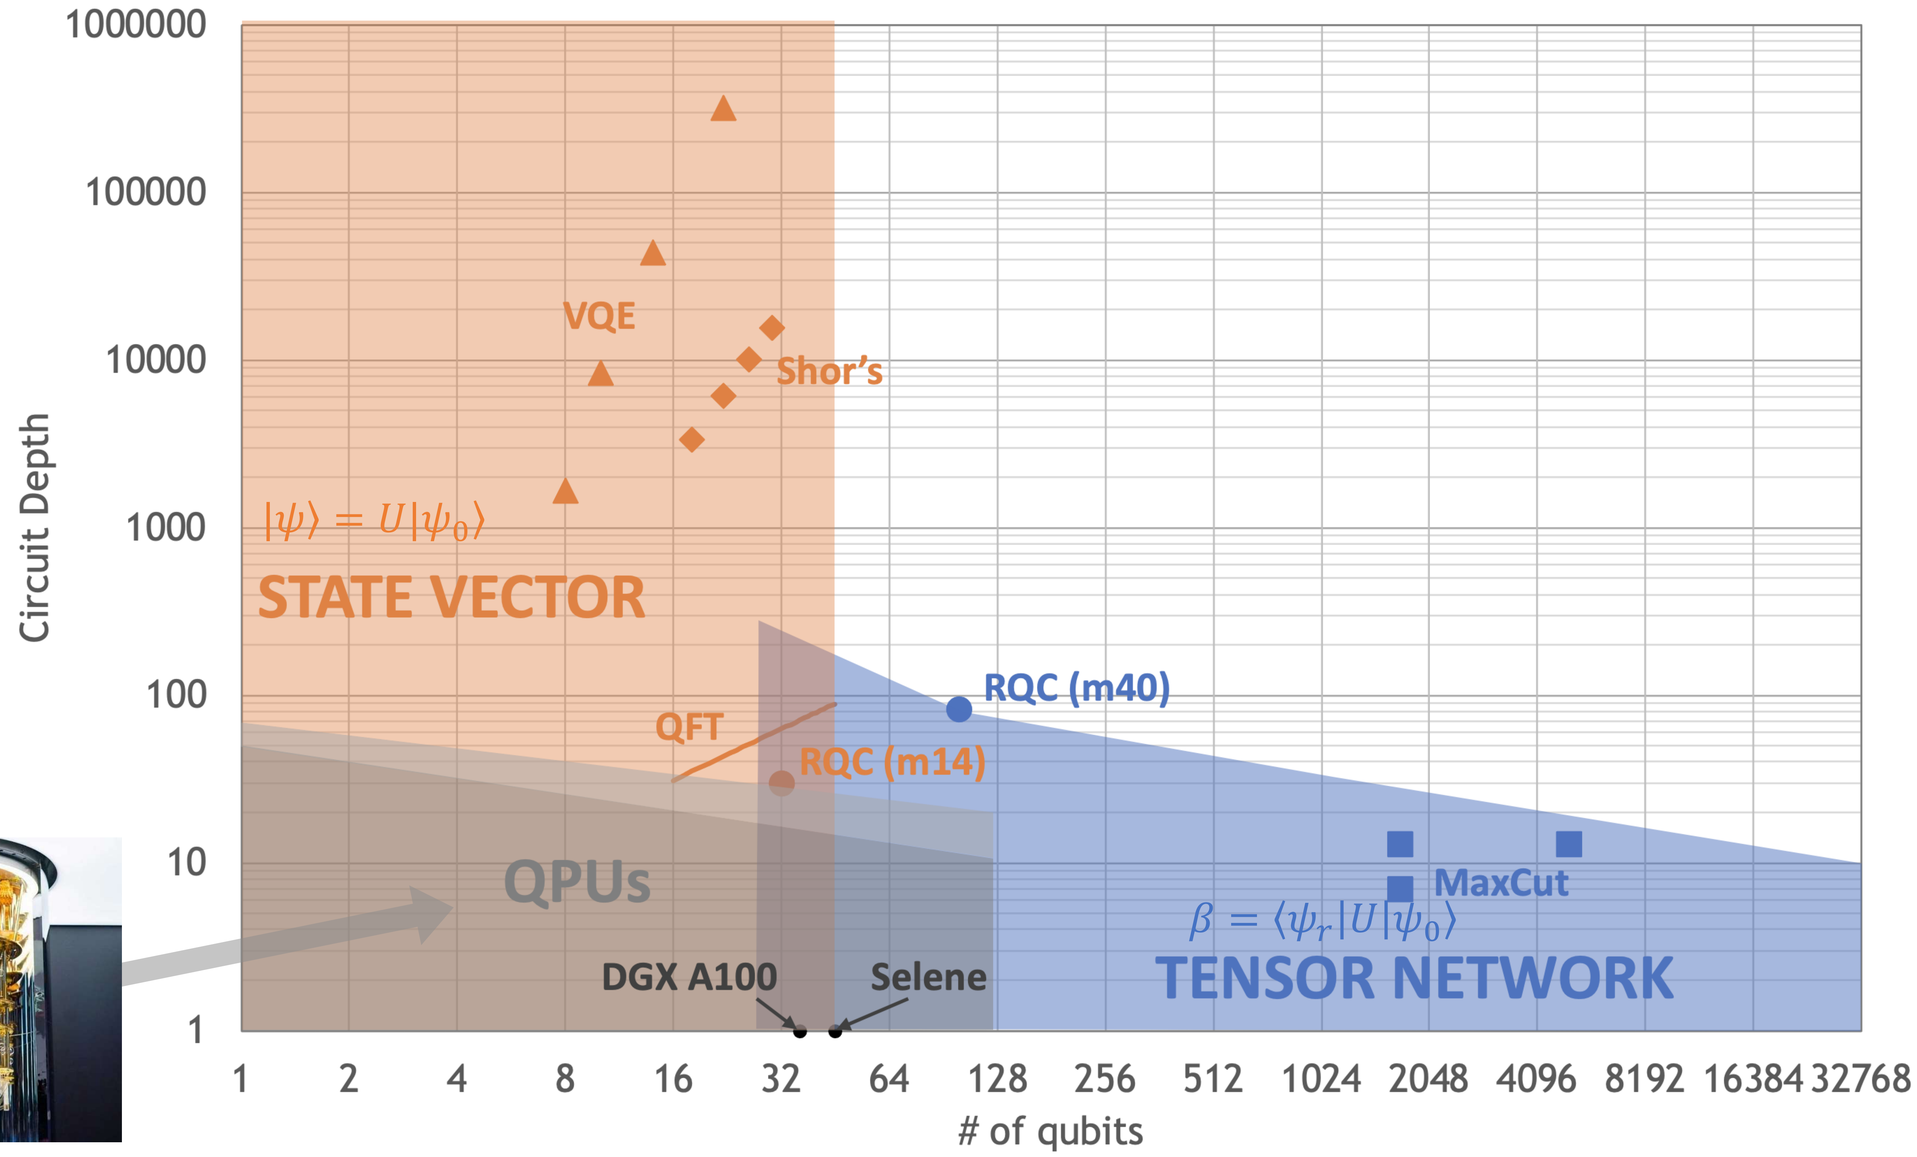

NVIDIA

cuQuantum

- **SDK for Quantum Circuit Simulation**
  - Accelerate Quantum Circuit Simulators on GPUs
  - Simulate ideal or noisy qubits
  - Enable algorithms research with scale and performance not possible on quantum hardware, or on simulators today
- **cuQuantum available now**
  - Integrated into leading quantum computing frameworks Cirq, Qiskit, and Pennylane
  - C and Python APIs
  - Available today at developer.nvidia.com/cuquantum

Cirq

Qiskit

PENNYLANE

aws Amazon Braket

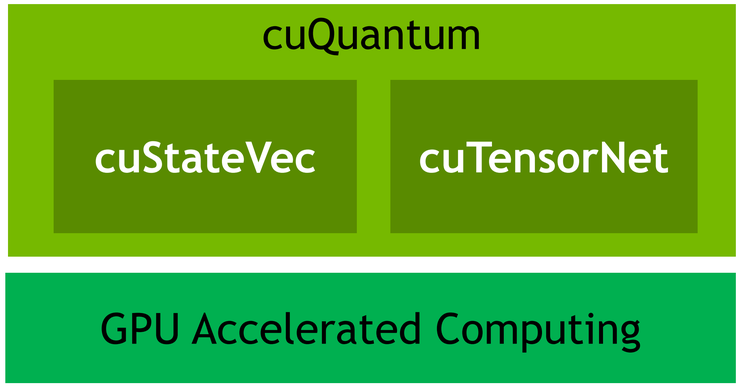

NVIDIA

cuQuantum Appliance

Continuous Performance Improvement

Fully integrated quantum simulation solution
- State-of-the-art ***performance***
- Unmatched simulation ***scale***

Reduce the simulation time by ***orders of magnitude***

Simulate ***thousands*** of perfect or noisy qubits

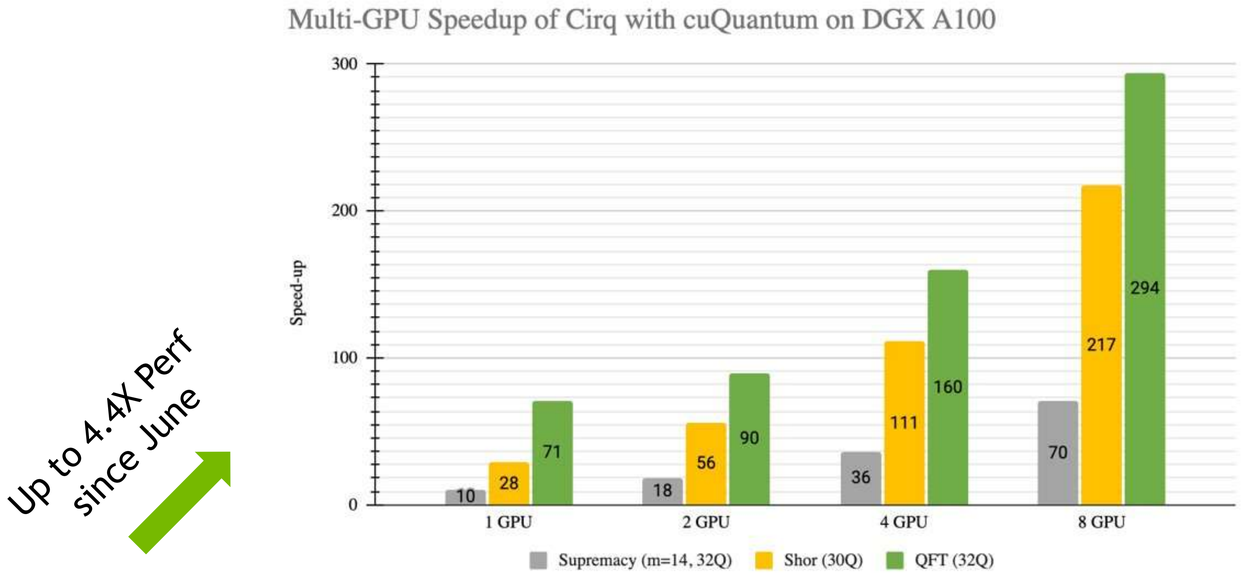

DGX cuQuantum Appliance

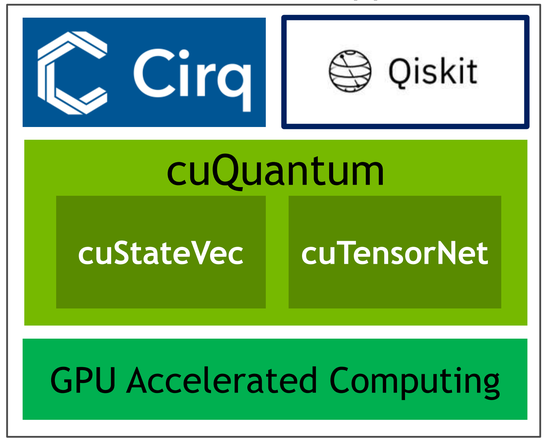

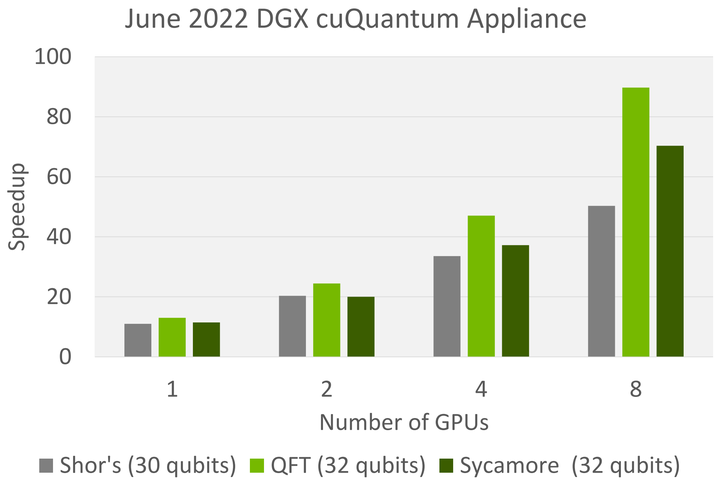

Speedup over AMD EPYC 7742 with 64 cores

Accelerating MultiNode SV Simulation

MPI Dominates Processing Time

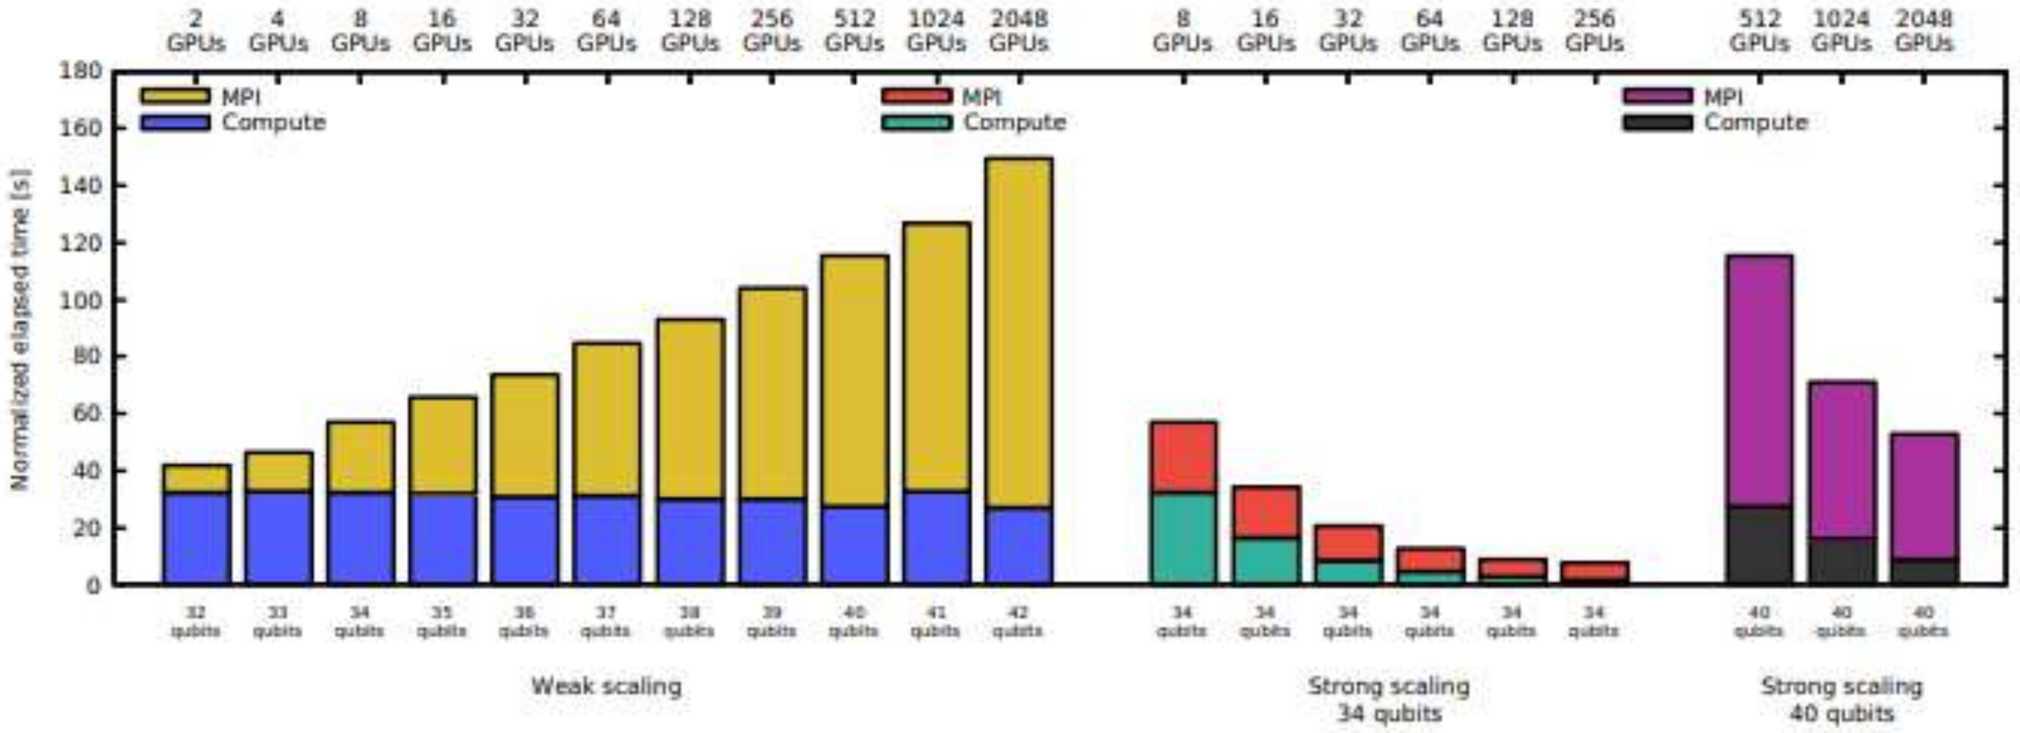

***GPU-accelerated simulations of quantum annealing and the quantum approximate optimization algorithm,** D. Willsch, M. Willsch, F. Jin, K. Michielsen and H. De Raedt, Computer Physics Communications 278, 108411 (2022)

SCALE-UP WITH NVLINK NETWORK

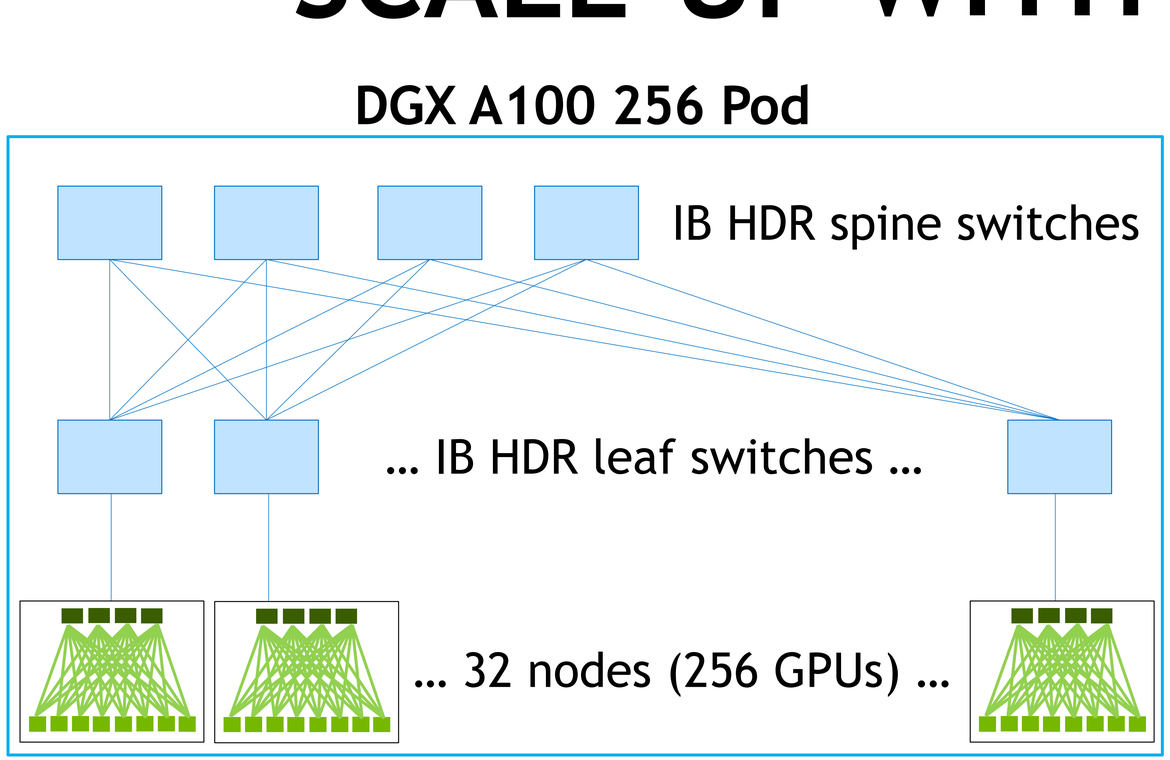

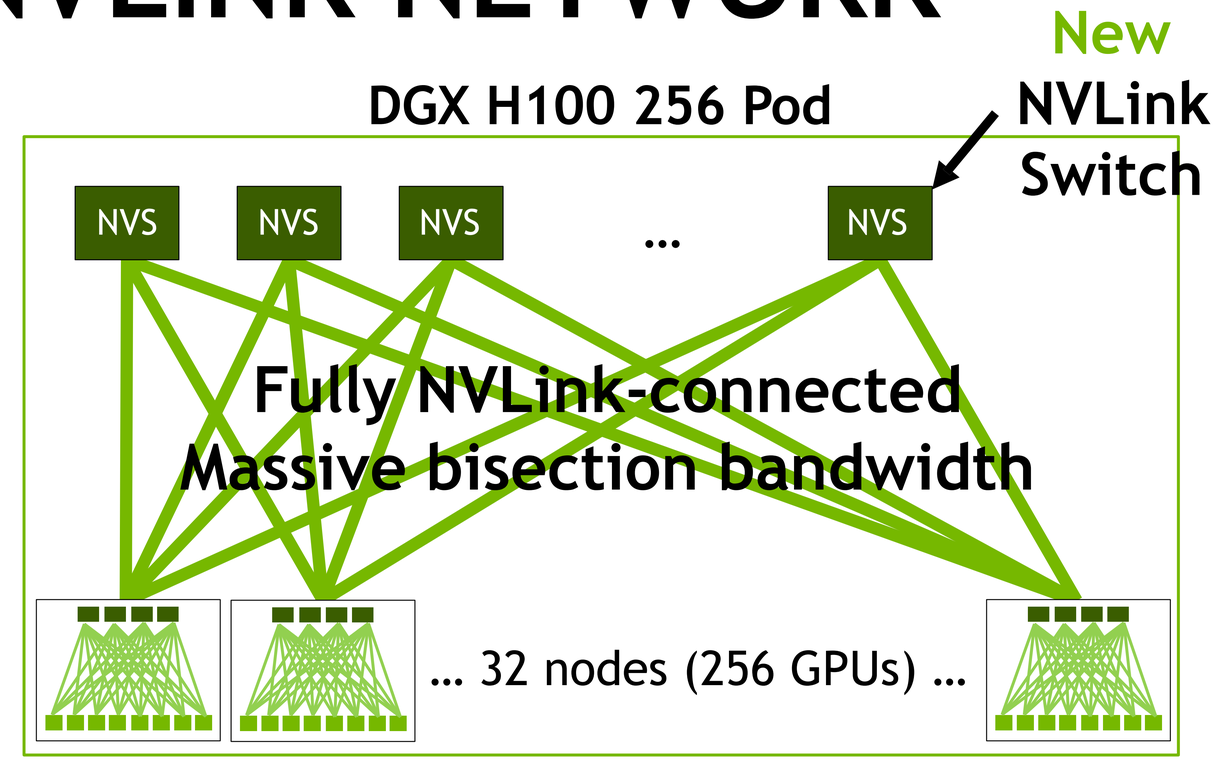

	A100 SuperPod			H100 SuperPod			Speedup	
	Dense PFLOP/s	Bisection [GB/s]	Reduce [GB/s]	Dense PFLOP/s	Bisection [GB/s]	Reduce [GB/s]	Bisection	Reduce
1 DGX / 8 GPUs	2.5	2,400	150	16	3,600	450	**1.5x**	**3x**
32 DGXs / 256 GPUs	80	6,400	100	512	57,600	450	**9x**	**4.5x**

COMING SOON: MULTINODE CUQUANTUM

- Multi-node simulation limited by internode gate operations requiring slower data transfer
- Our APIs will use index bit swapping along with recursive gate grouping to move as many gates as possible to within node, minimizing these transfers
- Can also use smart scheduling, overlapping gates with data transfer on other bits

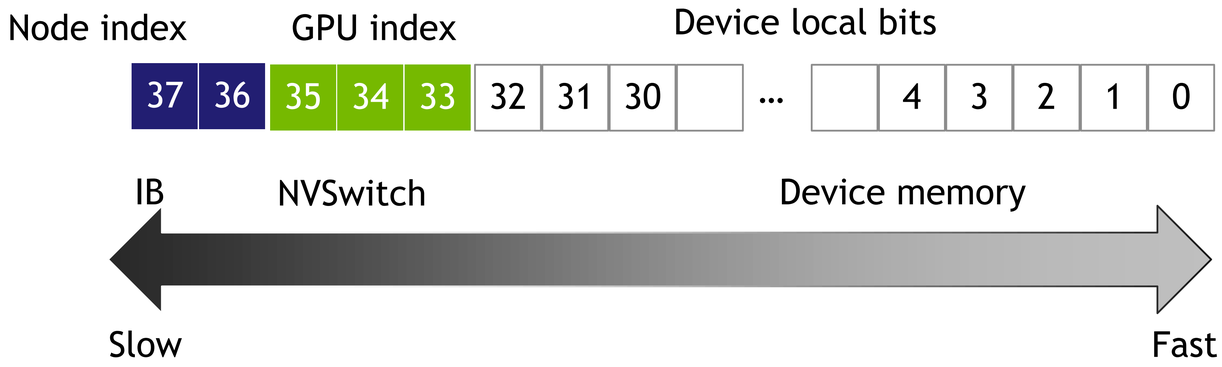

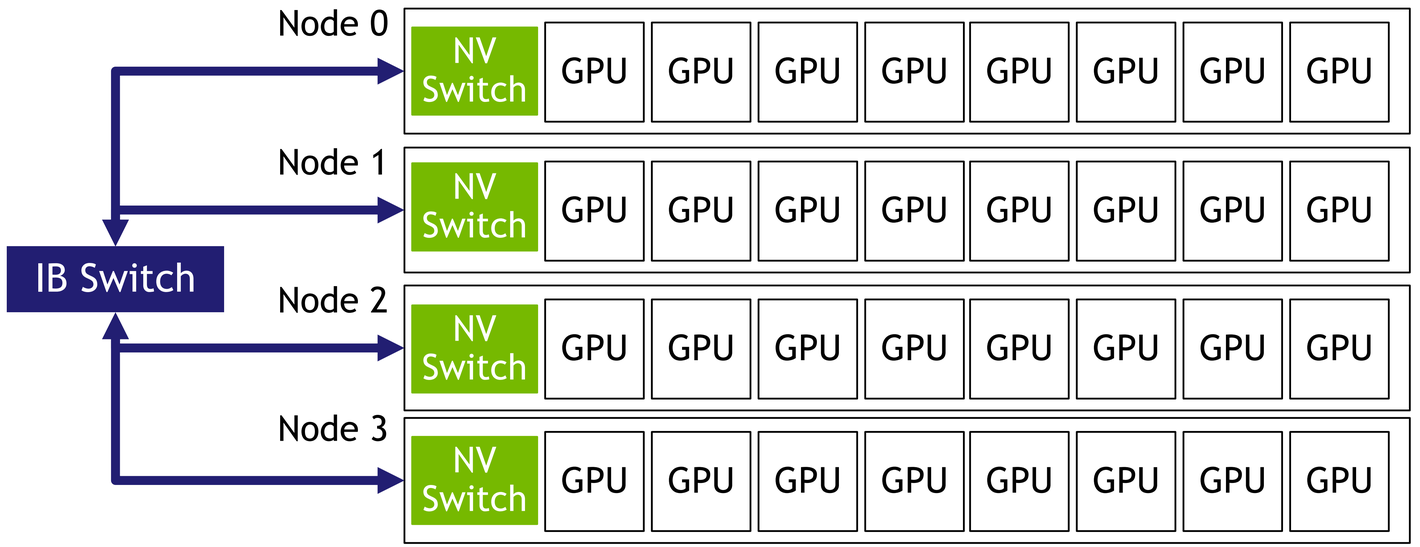

SCALING SIMULATIONS TO A SUPERCOMPUTER

NVIDIA's Selene DGX SuperPOD based supercomputer

- Using NVIDIA's Selene supercomputer
- Solved a 3,375 vertex problem (1,688 qubits/896 A100) with 97% accuracy
- Solved a 10,000 vertex problem (5,000 qubits) with 93% accuracy

Patti et. al.(2021)/NVIDIA Research
Paper: https://arxiv.org/pdf/2106.13304.pdf
Code: https://github.com/tensorly/quantum

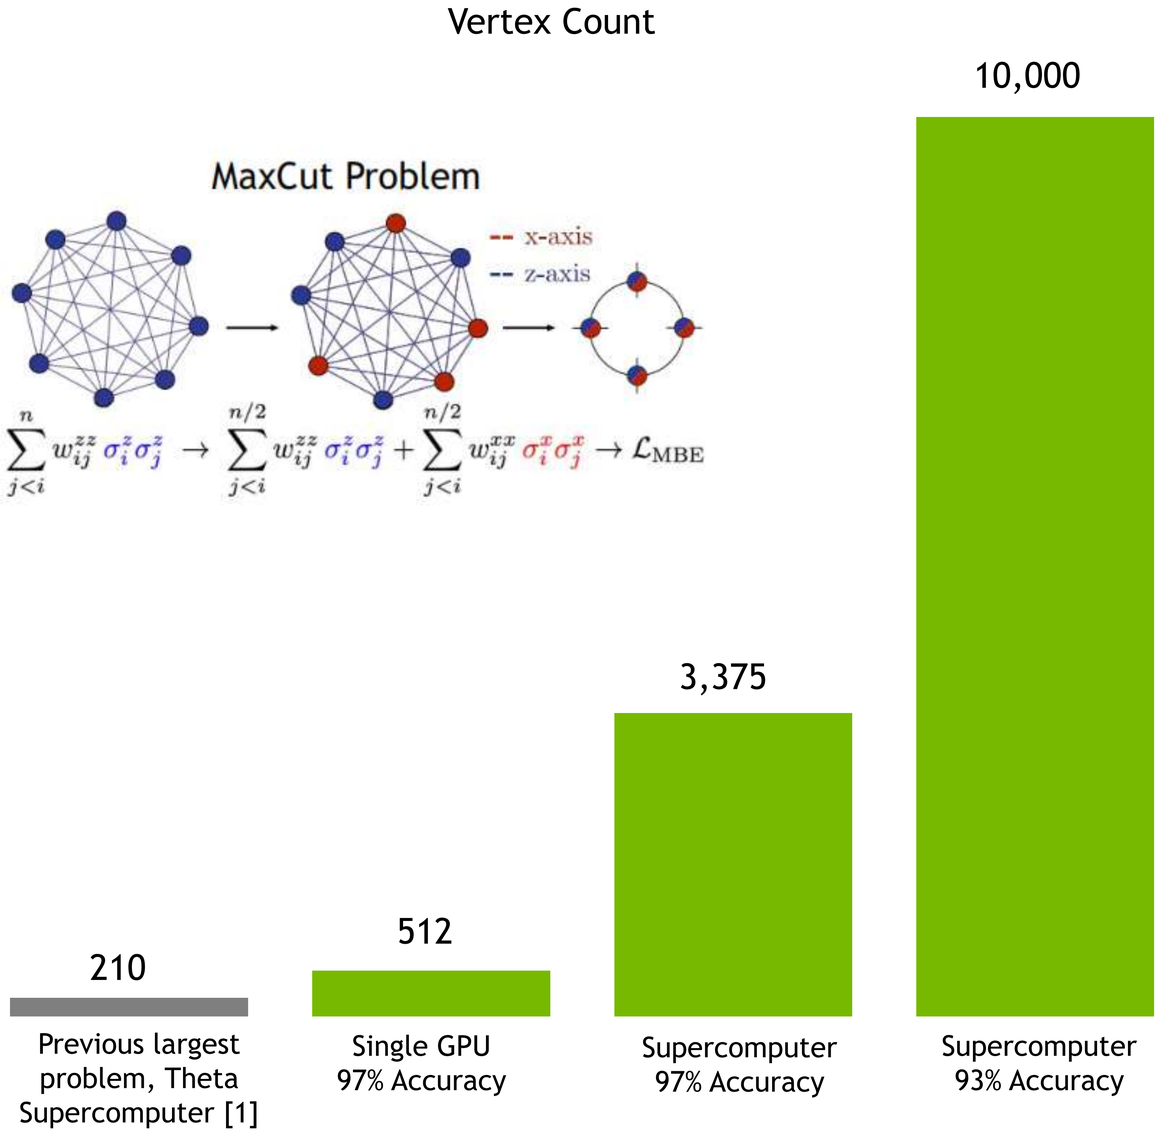

[1] Danylo Lykov et al, Tensor Network Quantum Simulator With Step-Dependent Parallelization, 2020
https://arxiv.org/pdf/2012.02430.pdf

NVIDIA

WHAT ARE USERS DOING?

ABCI
AI Bridging Cloud Infrastructure

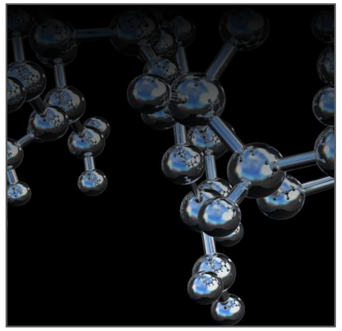

Grand Challenge

41 qubit full statevector simulation
on 512 GPUs

XANADU

129 Qubits

Novel cut technique enables 129
qubit QAOA simulation on
Perlmutter

Johnson&Johnson

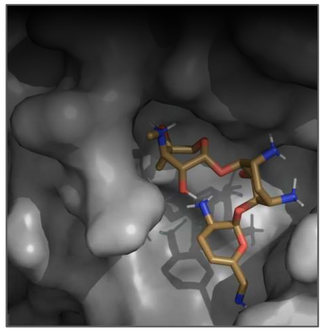

VQE

100X Speedup for VQE applied to
7MER Protein Folding in
collaboration with Strangeworks

NVIDIA cuQuantum Ecosystem

WHAT TO EXPECT NEXT?

- November 2022 - **cuQuantum 22.11** with
  - Support for approximate TN contractions (Matrix Product State formulation)
  - Container for multinode state vector simulation (Accelerating Qiskit-aer)

- March 2023 - **cuQuantum 23.03**
  - Library APIs for multinode state vector simulation
  - Support for batched simulation - enabling parallelized noisy simulations
- June 2023 – cuQuantum 23.06
  - Grace Hopper-optimized quantum circuit simulations

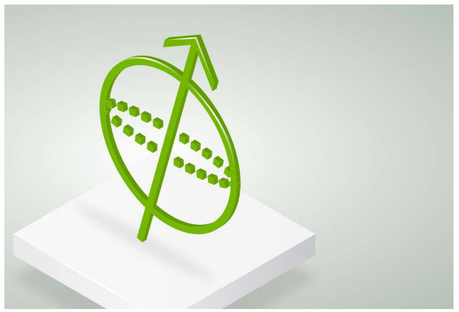

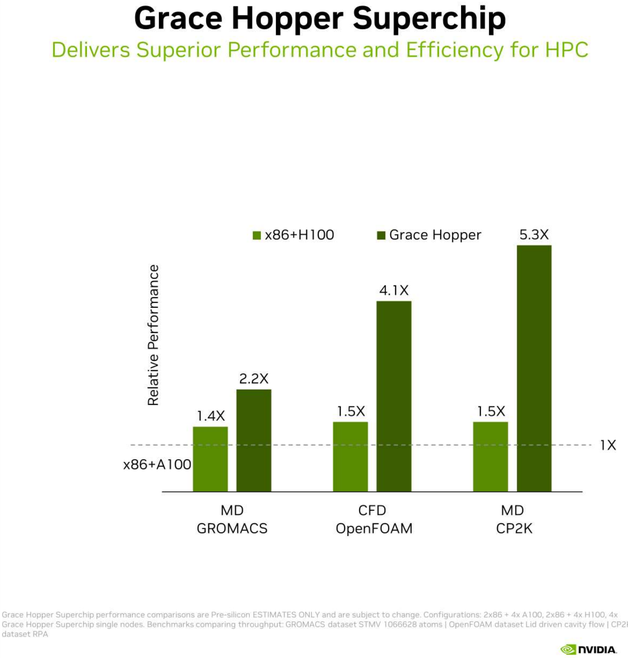

ECOSYSTEM GAPS LIMIT THE PROGRESS OF HYBRID QUANTUM APPLICATIONS

Classical Supercomputer

Quantum Computer

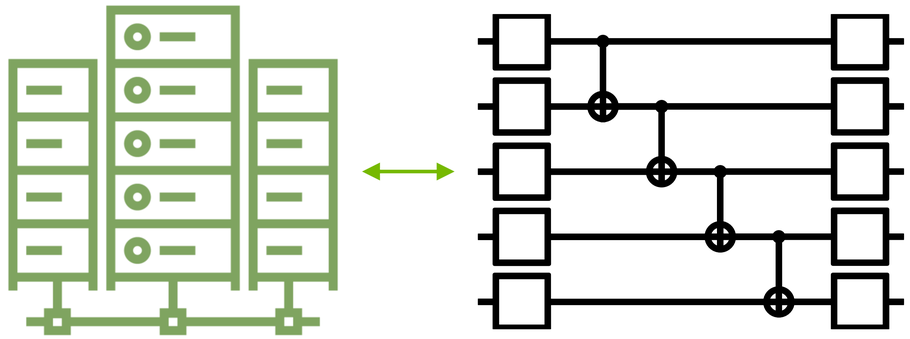

ECOSYSTEM CHALLENGES

- No Performant Software Stack
- Not Accessible To Domain Scientists
- Hybrid System Bottlenecks

ANNOUNCING NVIDIA QODA

THE PLATFORM FOR HYBRID QUANTUM-CLASSICAL COMPUTING

NVIDIA QODA PLATFORM

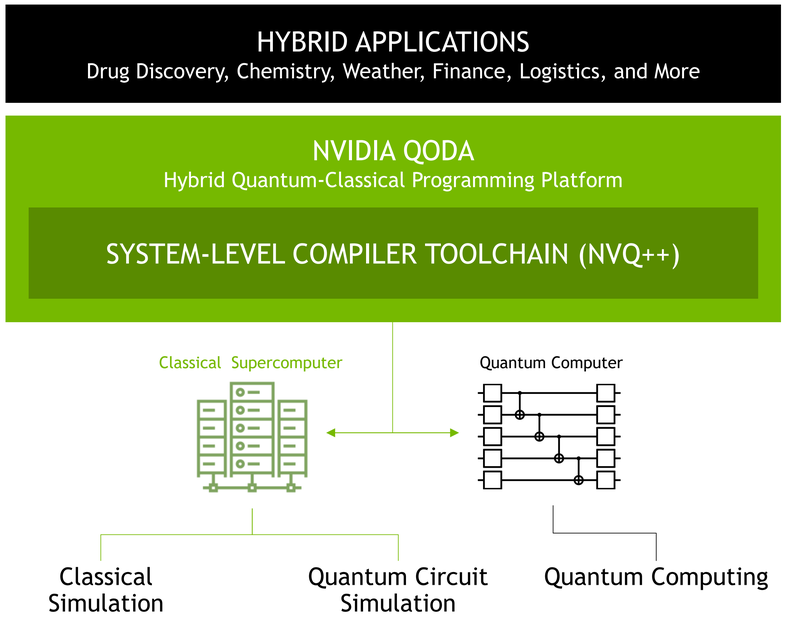

QODA FEATURES

Open and interoperable with today's applications

QPU-agnostic , emulated or physical

Compiler toolchain for hybrid systems

Kernel-based programming model

Single-source C++ and Python

Standard library of quantum algorithmic primitives

Unlocks performance, developer productivity, and cross-platform Interoperability

ANNOUNCING NVIDIA QODA

THE PLATFORM FOR HYBRID QUANTUM-CLASSICAL COMPUTING

NVIDIA QODA PLATFORM

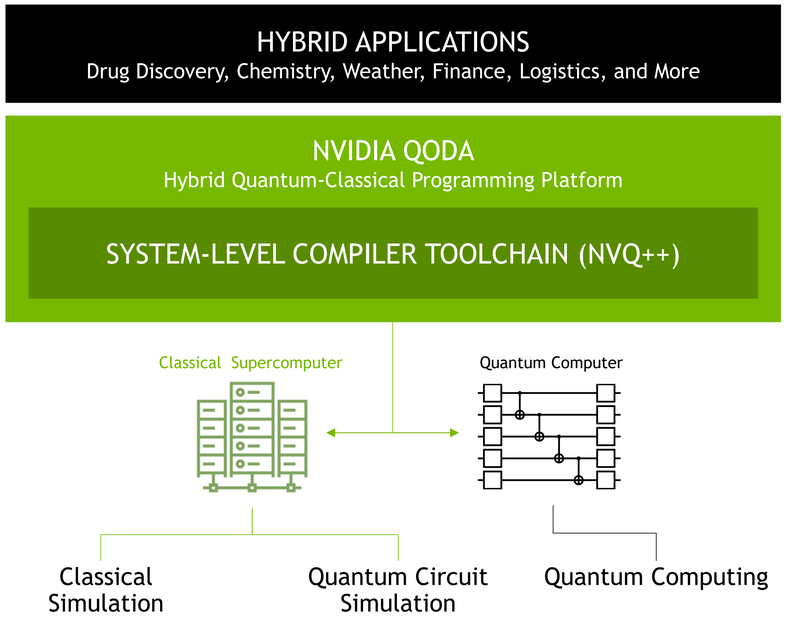

QODA PLATFORM ECOSYSTEM

IQM
PASQAL
QUANTUM BRILLIANCE
QUANTINUUM
rigetti
XANADU
Microsoft
Qubit pharmaceuticals
CLASSIQ
QCWARE
ZAPATA
JÜLICH Forschungszentrum
NERSC
OAK RIDGE National Laboratory

ANNOUNCING NVIDIA QODA

THE PLATFORM FOR HYBRID QUANTUM-CLASSICAL COMPUTING

NVIDIA QODA PLATFORM

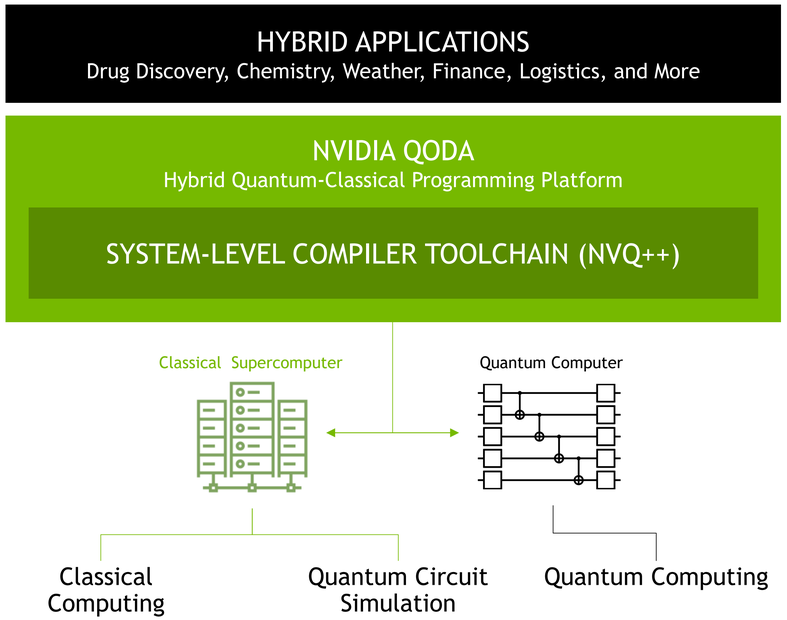

Familiar to Domain Scientists

```
// Define a QODA Quantum Kernel.
auto ansatz = [] (double theta) __qpu__ {
  qoda::qbit q, r;
  x(q);
  ry(theta, r);
  cnot(r, q);
};

// Define the Hamiltonian via Pauli tensor products.
qoda::spin_op H = 5.907 - 2.1433 * x(0) * x(1)
                - 2.1433 * y(0) * y(1)
                + .21829 * z(0) - 6.125 * z(1);

qoda::nlopt::cobyla opt;

// Run Variational Quantum Eigensolver with 1 param.
auto [min_e, opt_p] = qoda::vqe(ansatz, H, opt, 1);
```

ANNOUNCING NVIDIA QODA

THE PLATFORM FOR HYBRID QUANTUM-CLASSICAL COMPUTING

NVIDIA QODA PLATFORM

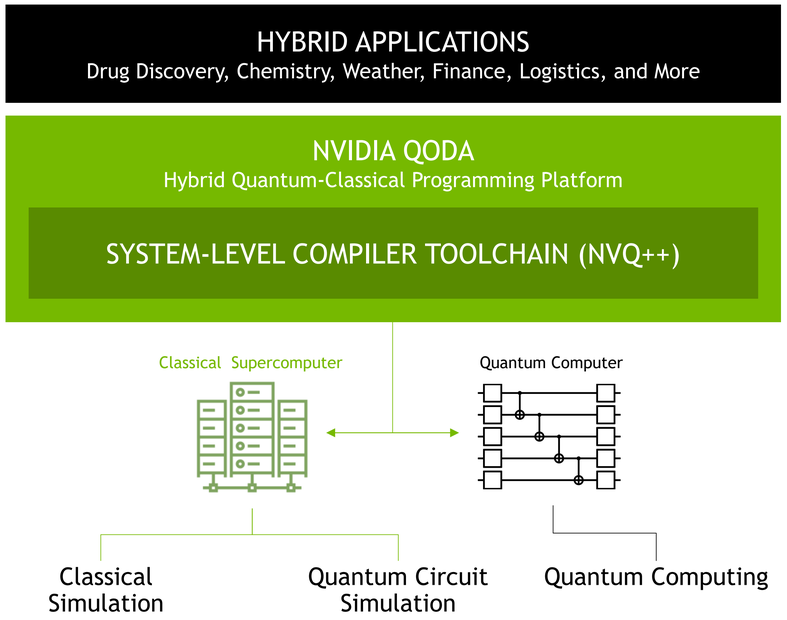

Interoperable with GPU Supercomputing

```
auto cnts = qoda::sample(q, …);
```

Standard Parallelism

```
std::sort(std::execution::par, …);
```

CUDA

```
kernel<<<…>>>(…);
cudaDeviceSynchronize();
```

OpenMP

```
#pragma omp target teams loop
for (…) …
```

OpenACC

```
#pragma acc parallel loop
for (…) …
```

ANNOUNCING NVIDIA QODA

THE PLATFORM FOR HYBRID QUANTUM-CLASSICAL COMPUTING

NVIDIA QODA PLATFORM

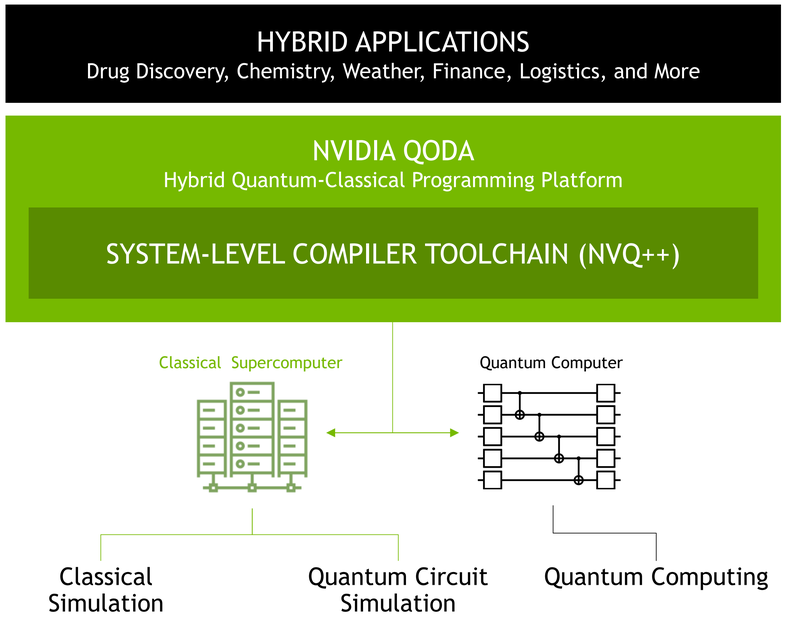

Interoperable with GPU Supercomputing

```
// Compute expectation values with QPU.
qoda::spin_op h = ...;
std::vector<double> sig_exps;
for (auto& pauli_op : generate_pauli_permutations(h.n_qubits()))
  sig_exps.push_back(qoda::observe(qite, pauli_op, h.n_qubits()));
```

```
...
// Compute LU Factorization of S_mat on the GPU.
auto dim = std::pow(2, h.n_qubits());
cusolverDnXgetrf(handle, params, dim, dim, CUDA_C_64F, S_mat,
                 lda, NULL, CUDA_C_64F, buffer_on_device,
                 bytes_on_device, buffer_on_host,
                 bytes_on_host, info);
```

ANNOUNCING NVIDIA QODA

THE PLATFORM FOR HYBRID QUANTUM-CLASSICAL COMPUTING

NVIDIA QODA PLATFORM

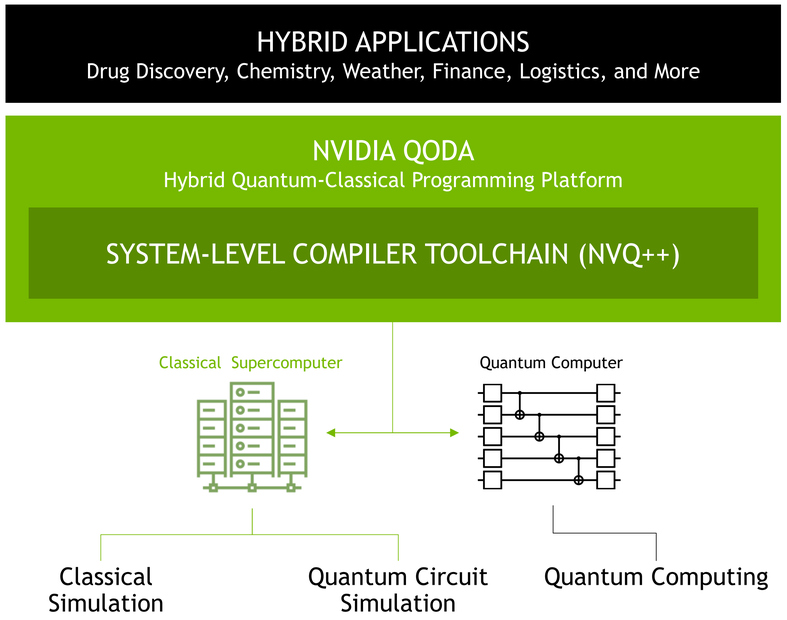

ENGINEERED FOR PERFORMANCE AND SCALE

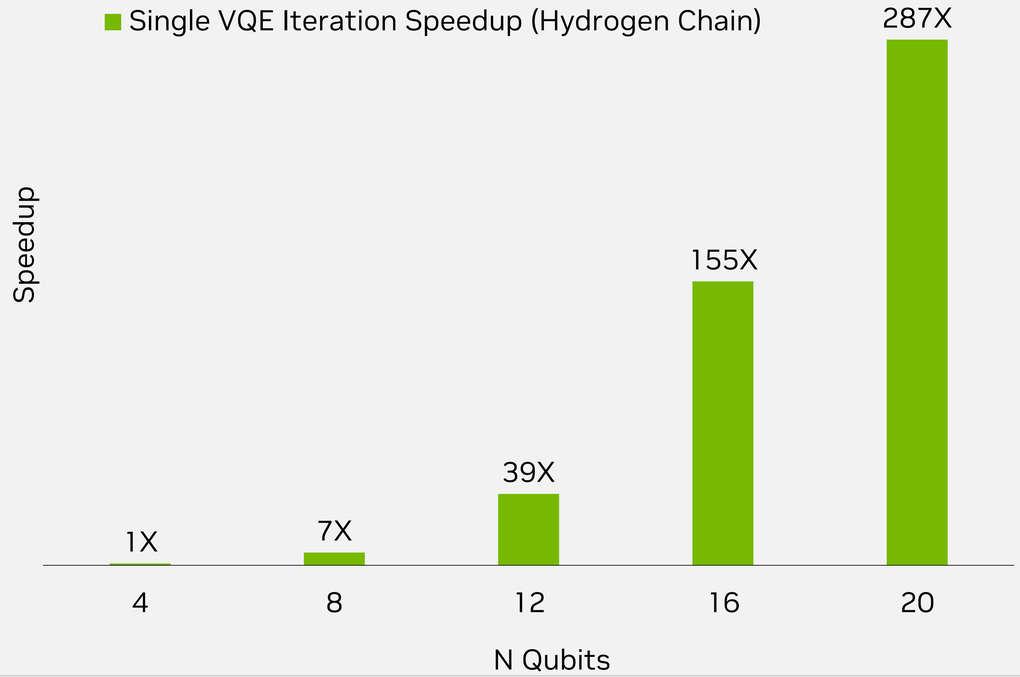

VQE: NVQ++/cuStateVec on A100 GPU vs Leading Pythonic Framework and SImulator with Thrust on A100 GPU

QODA ON QUANTINUUM

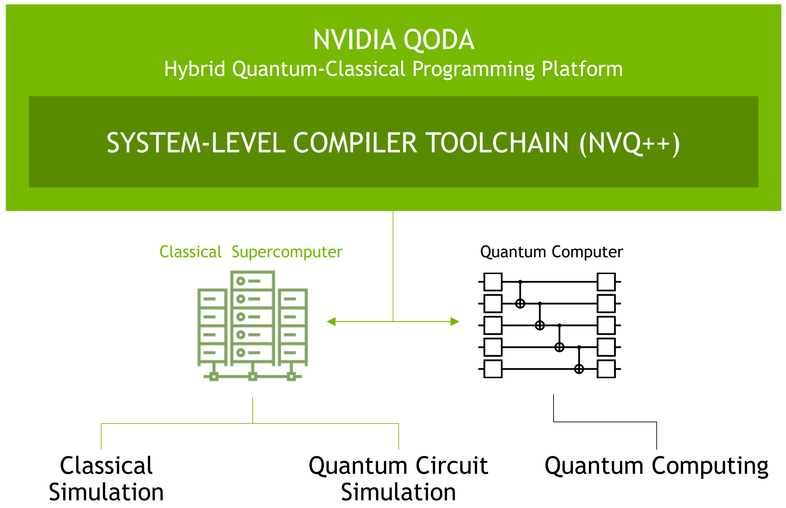

nvq++ -qpu=cuquantum vqe.cpp

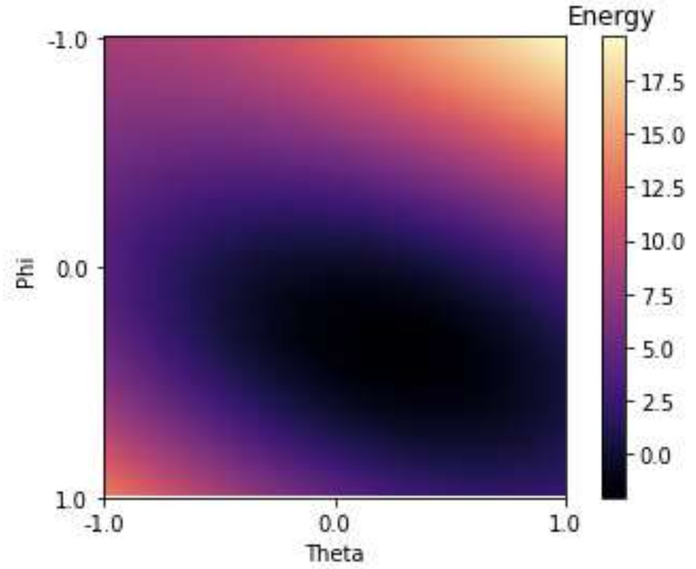

nvq++ -qpu=quantinuum:h1 vqe.cpp

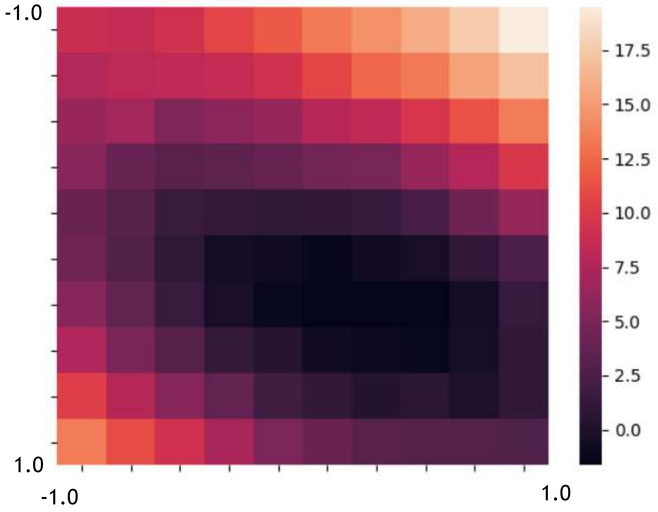

Introducing NVIDIA QODA

Enabling Innovative Quantum Systems Research

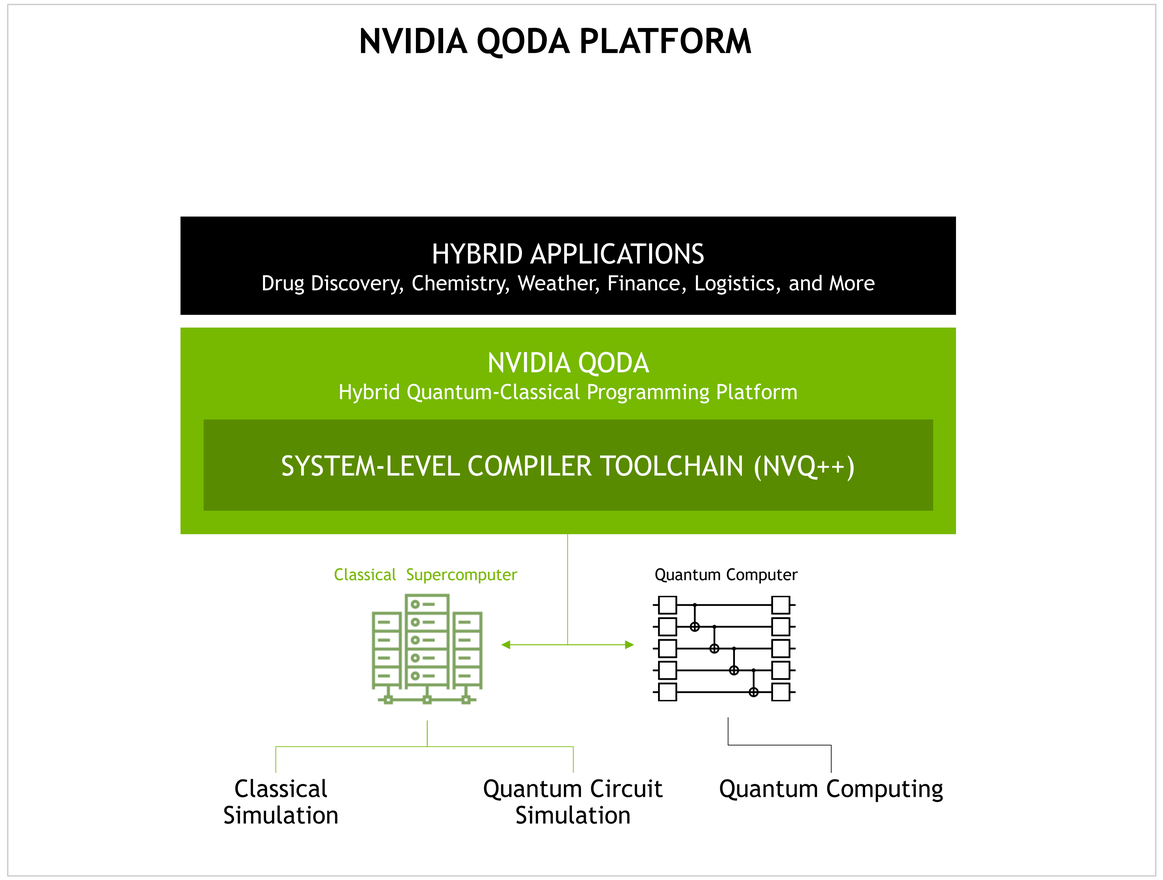


```
std::vector<qoda::observe_sender<double>> subs;
for (auto qpu : qoda::all_qpus()) {
 auto sub_H = H.subspan(qpu.idx() * terms_per_qpu, (qpu.idx() + 1) *
terms_per_qpu);
  subs.emplace_back(
    qoda::observe_async(qpu, sub_H, ansatz, ...));
}
auto sum = stdr::reduce(std::execution::par, qoda::when_all(subs), 0.0);
```

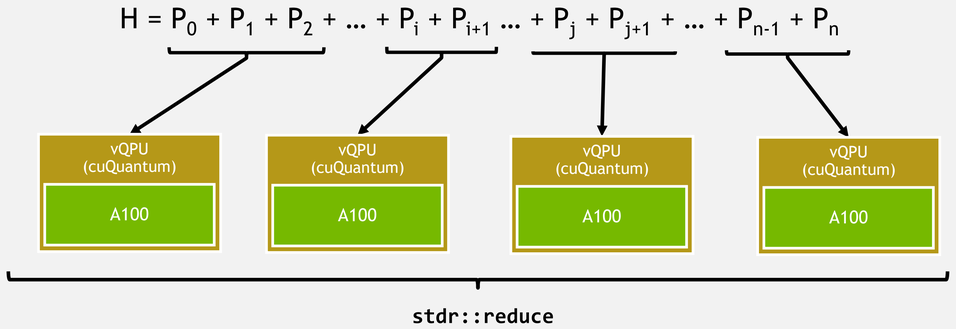

VQE: NVQ++/cuStateVec on A100 GPU vs Leading Pythonic Framework and SImulator with Thrust on A100 GPU

Introducing NVIDIA QODA

Enabling Innovative Quantum Systems Research

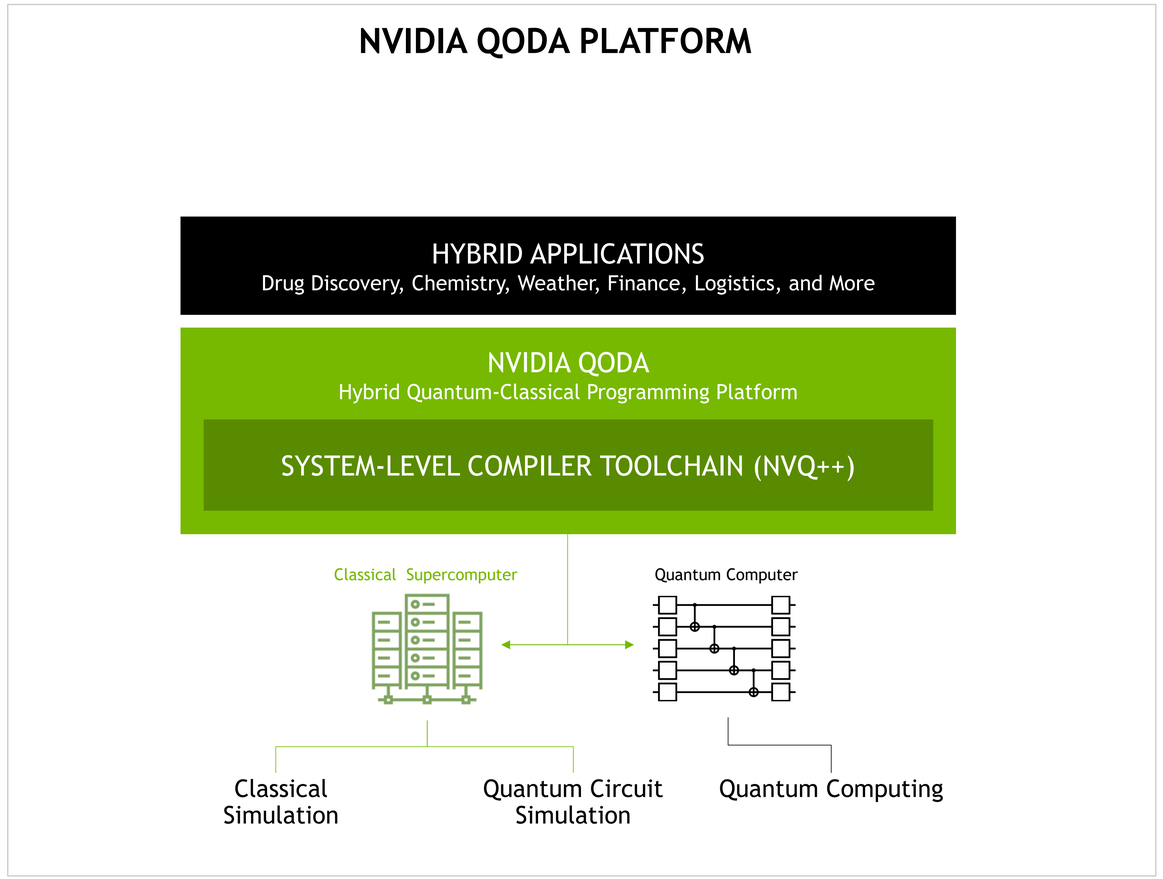

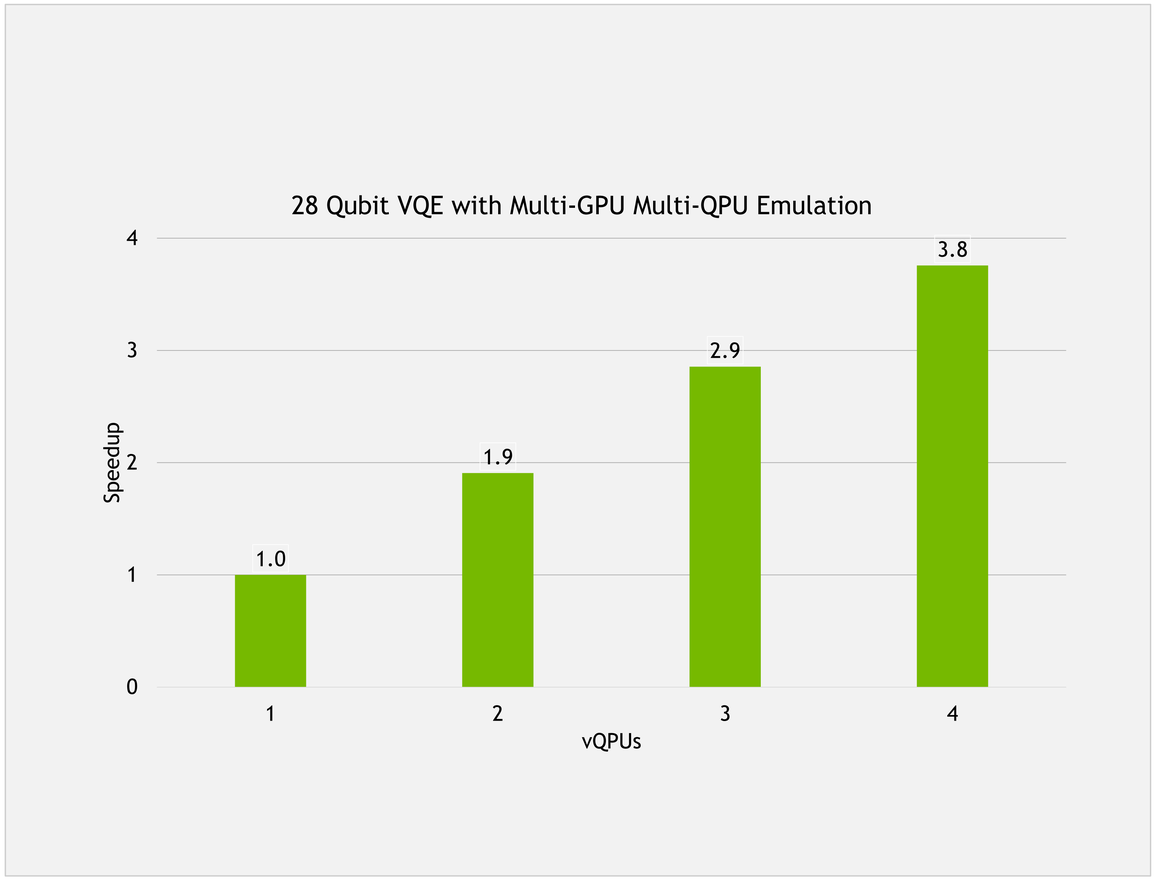

VQE: NVQ++/cuStateVec on A100 GPU vs Leading Pythonic Framework and SImulator with Thrust on A100 GPU

NVIDIA DGX QUANTUM APPLIANCE AND QODA DEVELOPMENT PLATFORM

A Scalable Platform for Hybrid Quantum Computing Research and Deployment

CLOUD PROVIDERS

Alibaba Cloud · aws · BAIDU AI CLOUD · Google Cloud · Microsoft Azure · ORACLE Cloud Infrastructure · Tencent Cloud

SYSTEM PROVIDERS

Atos · DELL Technologies · FUJITSU · GIGABYTE · H3C · Hewlett Packard Enterprise · inspur · Lenovo · Nettrix宁畅 · SUPERMICRO

ANNOUNCING NVIDIA QODA

THE PLATFORM FOR HYBRID QUANTUM-CLASSICAL COMPUTING

Classical Supercomputer

Quantum Computer

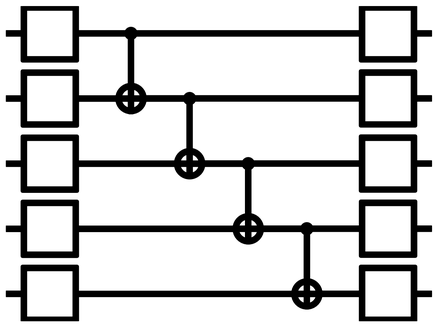

SUMMARY

- Simulating hybrid QC-classical systems requires large scale infrastructure with fast interconnect

- All Valuable Quantum Applications Will Be Hybrid
- A Platform For Real, Valuable Quantum Computing
- For Entire Quantum Computing Ecosystem

To go further:

- Qoda Blog
- Quoda-early-access
- Download CuQuantum

NVIDIA


NVIDIA