HIGH PERFORMANCE COMPUTING AT ROLLS-ROYCE

Todd Simons, April 8 & 9, 2025, HPC User Forum in Santa Fe, NM

OUTLINE

- Introduction
- Company Overview – who we are and who we are not
- US Footprint of Rolls-Royce
- The Physical Cloud, Cloud Computing, and Quantum
- HPC at Rolls-Royce
- How Adjoint CFD methods resemble Deep Neural Networks

WE DON'T MAKE CARS

Used with permission from bmwblog.com

Rolls-Royce is a power and propulsion company

The automotive division was spun off in 1971

- There are no discounts, no company cars, no loaners, no chauffeurs
  - Until recently I drove a Toyota Camry

Rolls-Royce is a **power and propulsion** company – Our mission is to innovate efficient and sustainable power solutions to meet customer's operational requirements and to protect our planet, secure our world and explore the universe.

ROLLS-ROYCE BUSINESS GROUPS

Civil

73% Large Engines
20% Business
3% Regional
4% other

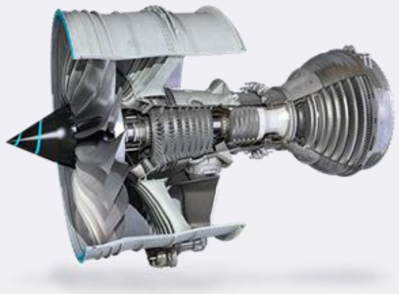

Defense

31% Transport
34% Combat
22% Submarines
8% Naval
5% Helicopters

Power Systems

39% Power Gen
25% Governmental
12% Marine
24% Industrial

New Markets

Emerging markets include Small Nuclear Reactors (SMR) driven by energy security and decarbonization.

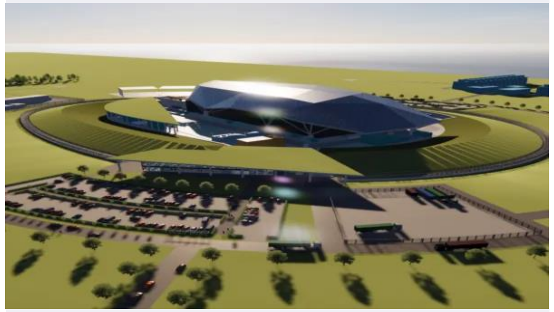

US FOOTPRINT

Rolls-Royce directly employs more than **5,600** diverse high-skill US teammates and supports tens of thousands of jobs across our downstream network with **$2.1B Annual US Supplier Spend**.

Since 2011, we have invested more than **$1.5B in our US facilities**.

We also operate multiple University Technology Centers (UTCs) and support several apprentice and workforce development programs at our flagship sites.

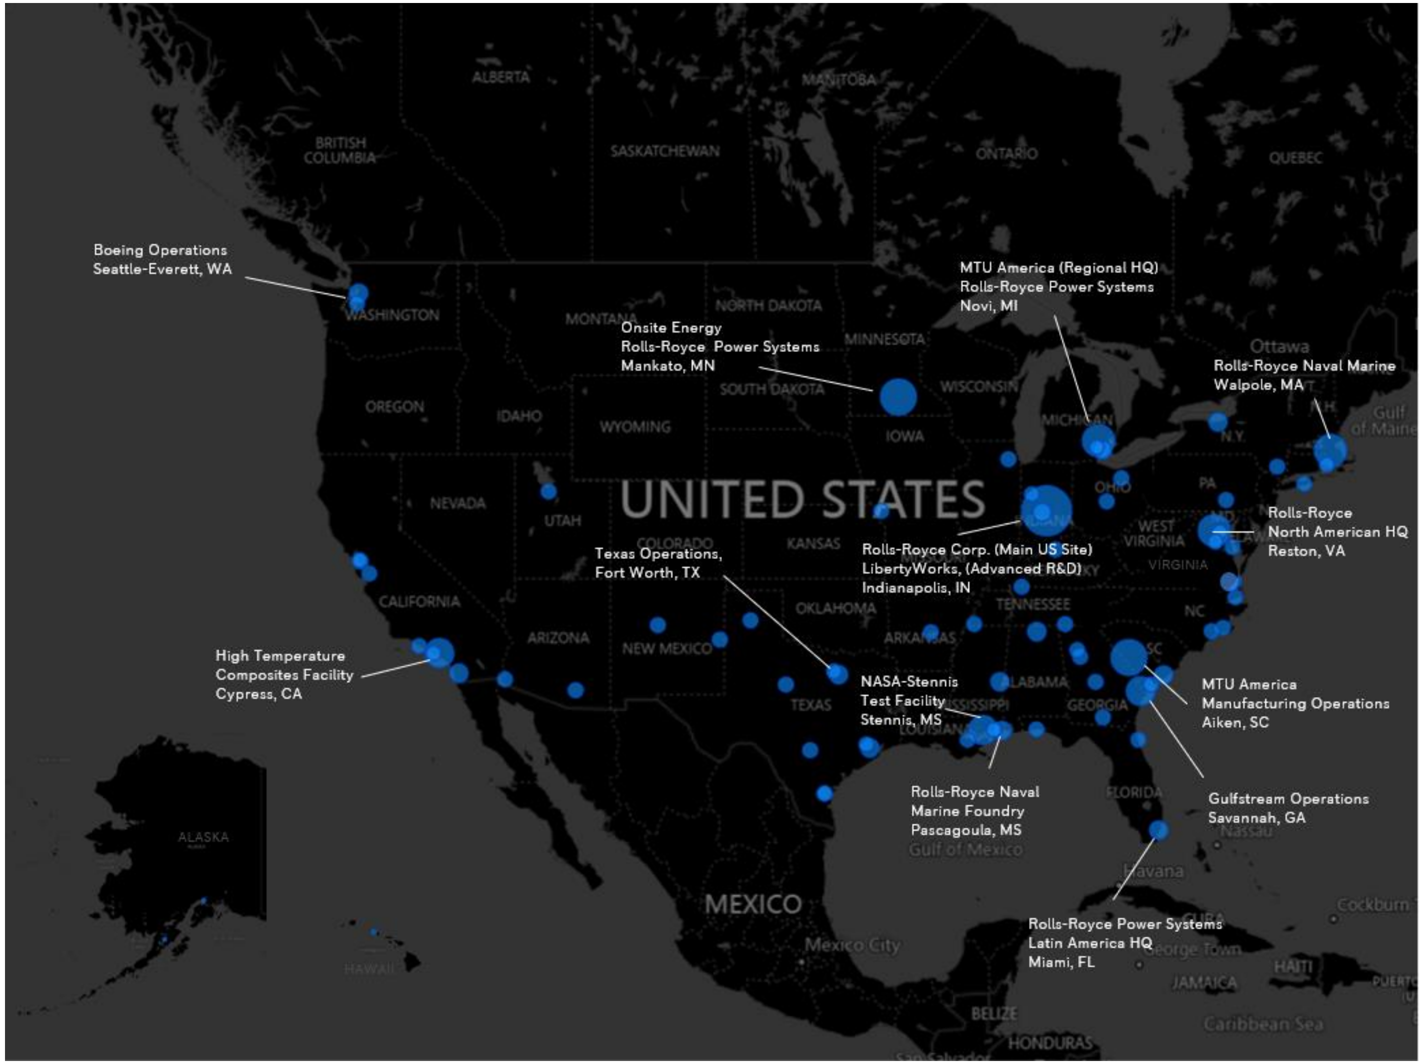

US DEFENSE

In the US Rolls-Royce is a US Defense contractor located in Indianapolis, IN

F35-B

Earlier in my career as a CFD analyst I had the privilege to support a re-design the LiftFan propulsion system VaneBox that gives the F35-B the capability to take off and land vertically.

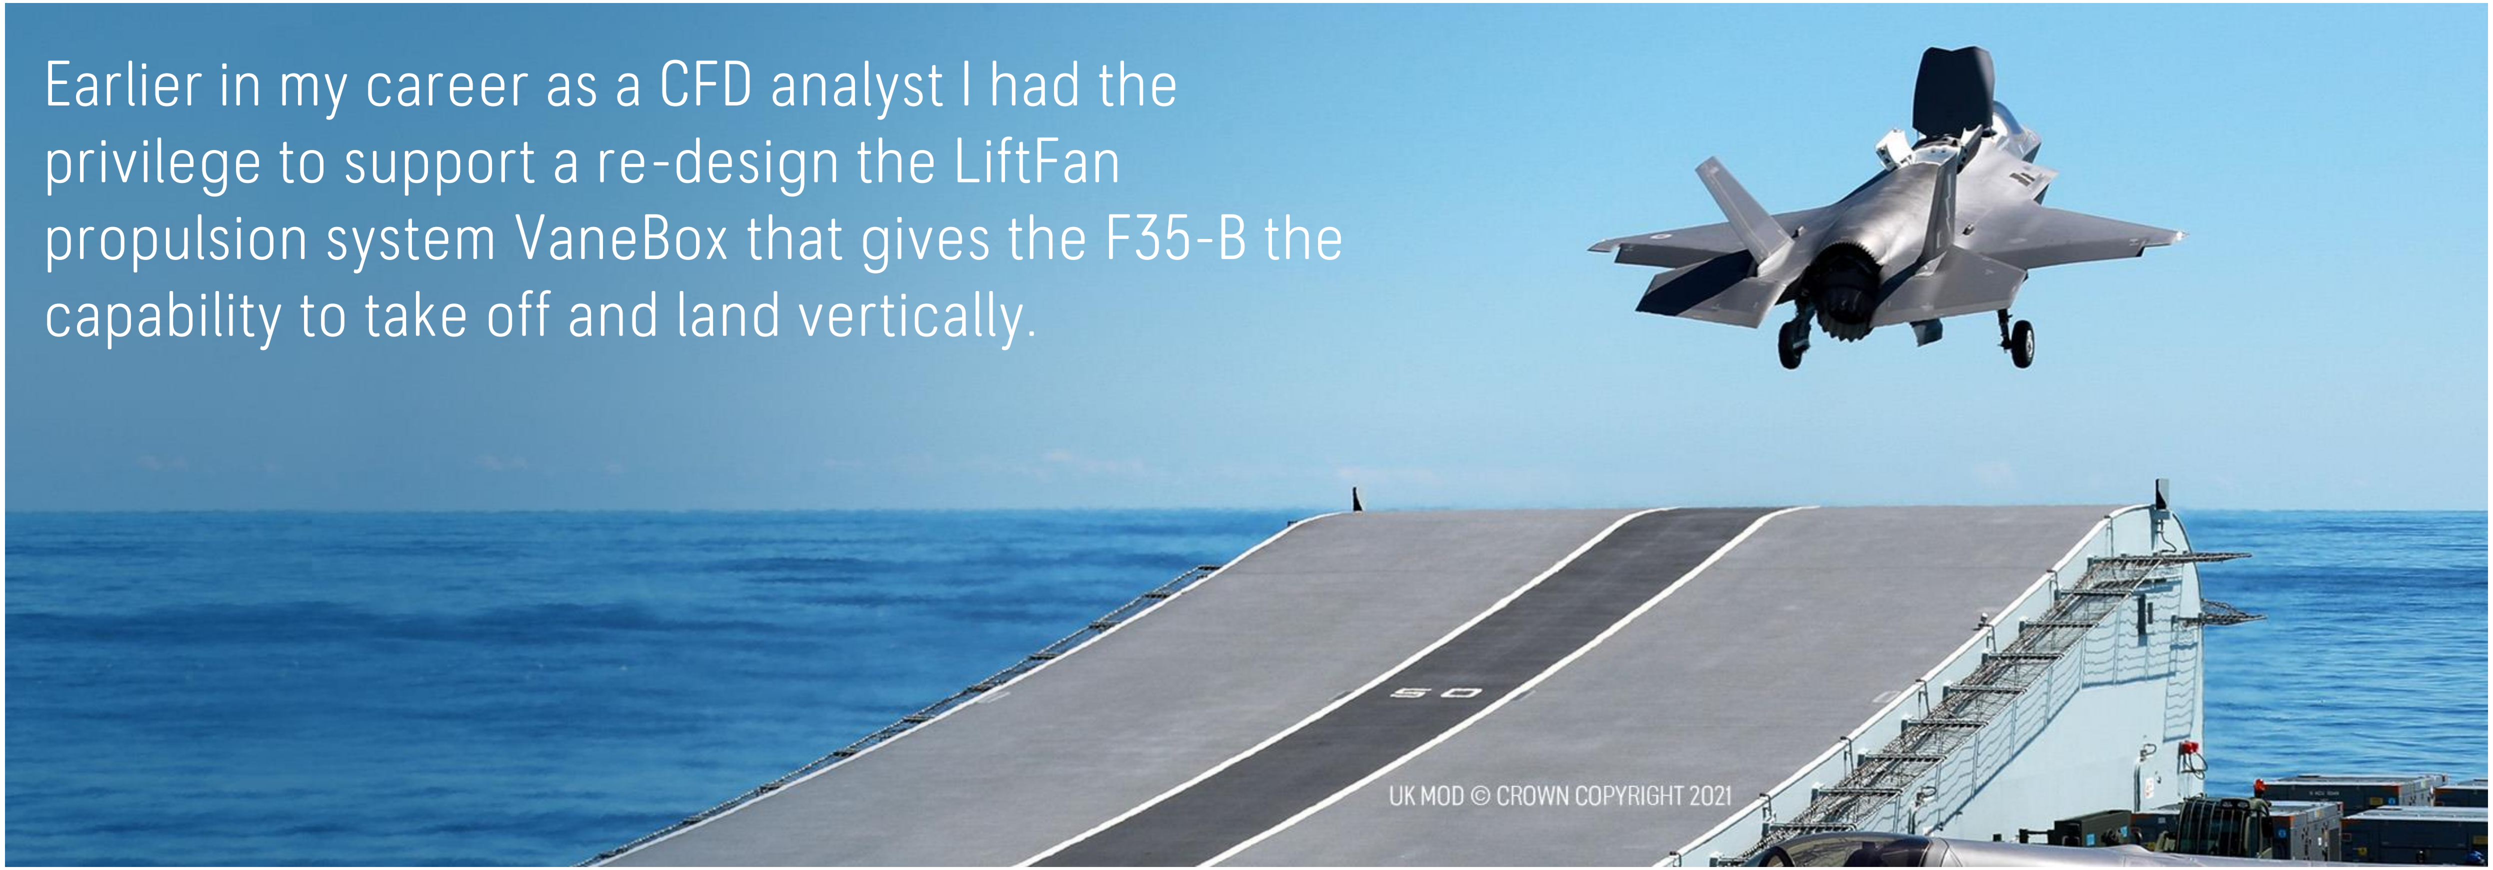

ENDURANCE

In the 70 years since its first flight, the T56 engine has amassed more than 220 million service hours, a rare occurrence in aviation history. First flying in 1954, the T56 powers C-130H Hercules, which is still flying as part of the longest continuous military aircraft production run in history.

Since entering service in 1994, the AE 2100 engine has reached 10 million hours, powering the tactical airlift workhorse of the Air Force: The C-130J Super Hercules.

The AE 2100 is part of the AE engine family – a versatile, proven engine core with more than 88 million flight hours known for a record of dependable and efficient service across military and commercial fleets around the world.

US DEFENSE

In February the Rolls-Royce U.S. Defense team delivered an instrumented AE 3007N engine to Boeing for the MQ-25 Stingray program. The MQ-25 Stingray is the U.S. Navy’s first aircraft carrier-based unmanned air vehicle. The aircraft will provide unmanned aerial refueling and intelligence, as well as surveillance and reconnaissance capabilities.

Manufactured in Indianapolis, this engine is the first of eight to be delivered to support flight testing later this year. The comprehensive testing program will help ensure readiness in an aircraft carrier environment, specifically measuring how the aircraft and engine perform with the systems that launch and recover aircraft from aircraft carriers (catapults and arresting gear).

ROLLS-ROYCE POWERS THE WEB

One in every three online clicks is powered by Rolls-Royce mtu diesel gensets

Today, there are more than 85,000 units installed around the globe, many of them providing vital backup power for data centers. Alongside keeping vital services online, they offer a future ready solution due to their ability to run on sustainable fuels like Hydrotreated Vegetable Oil (HVO).

ROLLS-ROYCE AND THE CLOUD

Our engines are now in the cloud.

The latest EHM release now includes the Trent 7000 engines powering the Airbus A330 Neo fleets, following on from the Trent 1000, Trent 500, and Pearl 700 engine fleets previously migrated.

We are now able to manage new and more complex engineering diagnostic capabilities to prevent in-service issues from occurring and strengthen our cyber security.

Our teams are also embracing new technologies to analyze more data and bring more analytics into service much more rapidly.

NUCLEAR ENERGY – SMR AND MICRO

Rolls-Royce provides the nuclear power plants for the UK submarines fleet.

Rolls-Royce is part of a DOD micro-reactor program, Pele, that has broken ground in Idaho.

Rolls-Royce secured funding from the UK Space Agency to research how a micro-reactor could support a moon base for astronauts.

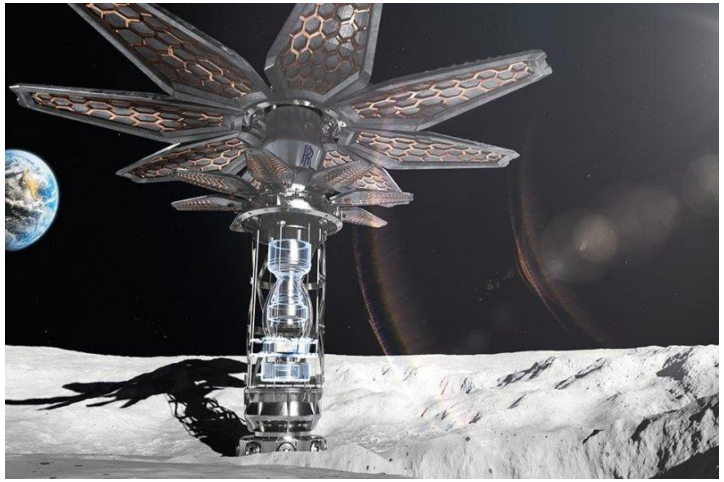

Moon based MW micro-reactor

Micro-reactor

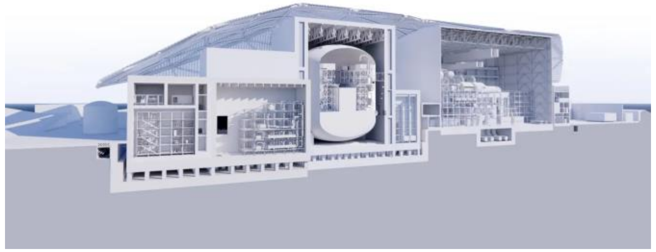

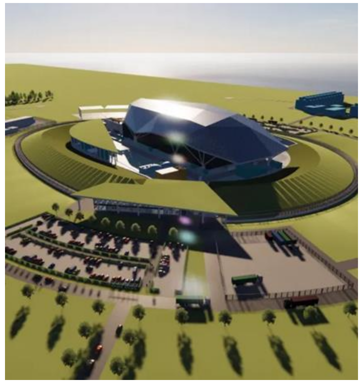

Giga-watt SMR technology (470 MWe)

Our SMR team as partnered with ČEZ Group (CEZ) to deploy 3-GW of power in the Czech Republic in 2024

QUANTUM COMPUTING

Our quantum efforts began in 2018.

We are working with a number of companies and quantum approaches

We have an Industry-led collaboration, Catalyst, working with

- Universal Quantum (using ion trap technologies)
- Riverlane, a UK-based quantum error correction company
- Xanadu, a Canadian quantum company focusing on photonic quantum computing
- Pasqal, a Canadian quantum company focusing o neutral atom technology
- ATOS QLM machines with up to 38 qubits
- IBM

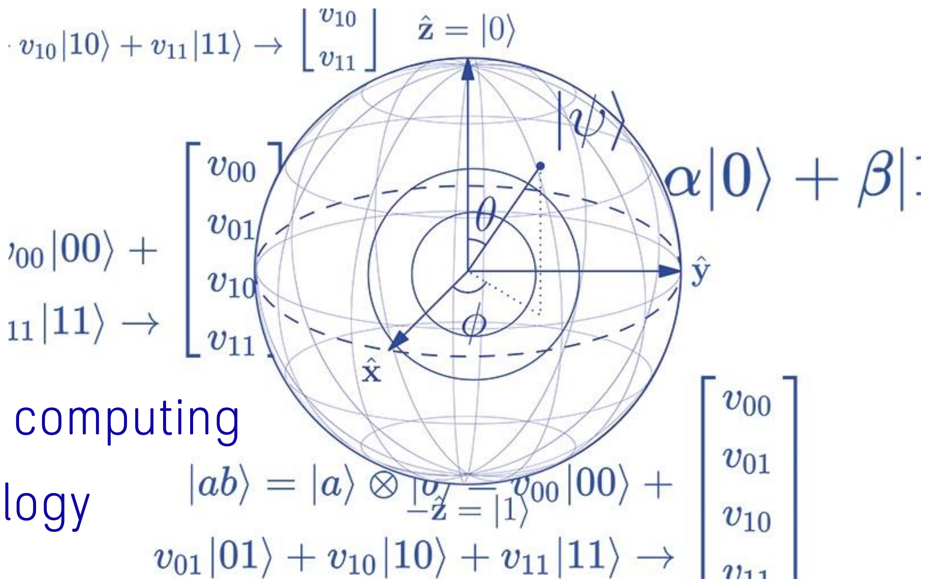

Our internal focus is on hybrid algorithms that bridge classical computing and quantum computing. We are developing quantum algorithms for linear equations in CFD - ultimately aimed at solving the Navier-Stokes equations.

HPC IN ROLLS-ROYCE DESIGN

- HPC is used in detailed design for aerodynamics, structural mechanics, combustion and materials and process modeling.
- We have large traditional HPC clusters with CPUs and some GPUs for physics-based modelling and simulation.
- We use both in-house/bespoke software as well as commercial software.
- We continue to use 1-D, 2-D, and coarse 3-D CFD analysis in the design process. We then use high fidelity models in later in detailed component design.


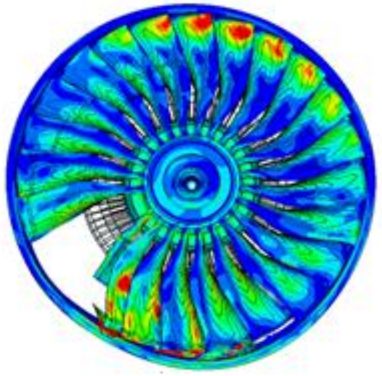
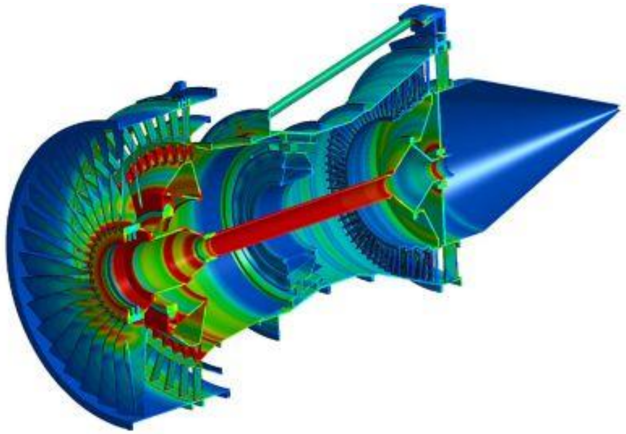
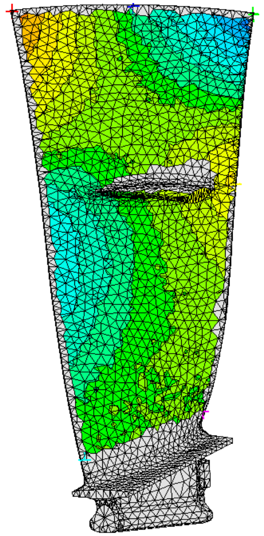

HYDRA CFD CODE

We celebrated the 25th anniversary of the Rolls-Royce Hydra CFD code this year.

Hydra has become our main CFD corporate turbomachinery code. During the 1970s and 80s, the aerothermal design process for new engines was largely experimental, guided by empiricism and re-scaling components from previous successful designs.

3-D CFD codes made their appearance during the 1980s and grew in impact through the 1990s.

Hydra consists of four main codes:

1. Nonlinear steady & unsteady Navier-Stokes solver
2. Linearized harmonic flow NS equations
3. Steady **adjoint** form of NS for design optimization
4. Adjoint harmonic flow equations for aeroelasticity

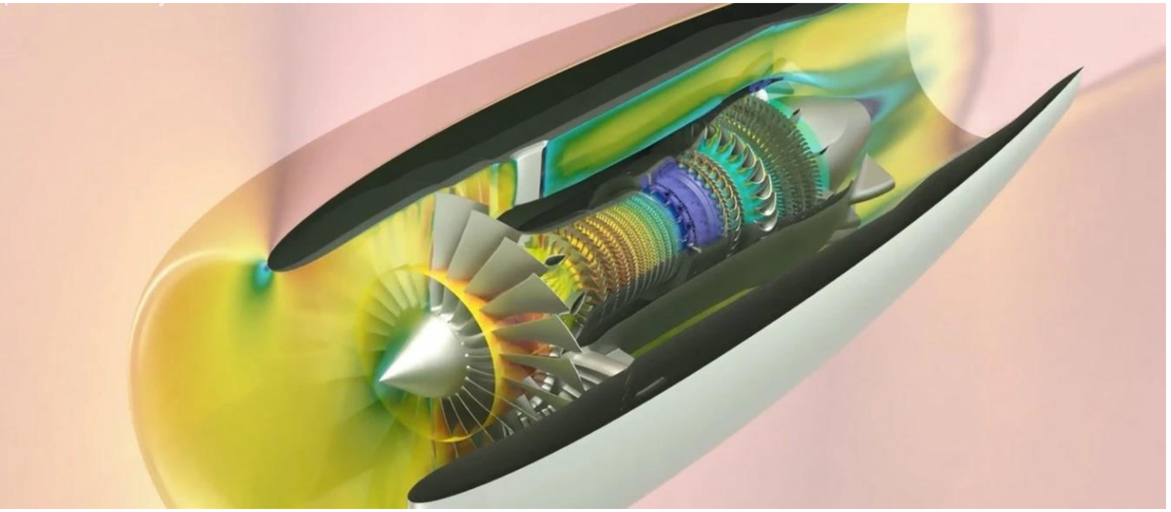

HYDRA IMPROVEMENTS OVER TIME

In the early 2,000's we used the explicit solvers in Hydra with multi-grid, typically five-stage Runge-Kutta solvers.

Design was performed with multiple runs to map the design space using Design of Experiments (DOE) methods.

We now use the implicit solver with Adjoint to explore design space.

This is an order of magnitude faster.

We've also developed an open-source parallel framework that allows Hydra to run on GPUs – both NVidia and AMD GPUs at this point.

This can be extended to FPGAs.

Near term plans include scaling up the AMD GPU version to demonstrate parallel scaling of the software.

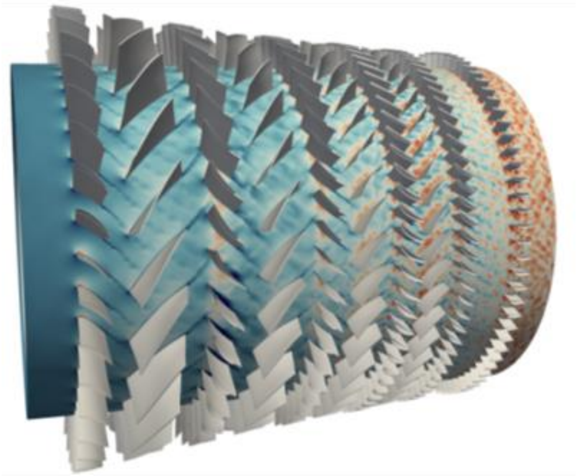

This plot shows the axial velocity contours for a 4.5 billion grid point unsteady simulation of the German Aerospace Center (DLR) 4.5 stage research compressor, Rig 250.

This won the SC22 Best Visualization Award for Full Aero-Engine compressor visualization.

COMPARISON OF NEURAL NETWORKS AND CFD

There are interesting similarities in the mathematics of Deep Neural Network and Adjoint CFD Simulations.

Caveats:

This is not a talk on solving the CFD analysis with Neural Networks or modeling turbulence with Neural Networks. We'll save that for another presentation.

In the next few charts I will discuss some parallels of DNN classification and Adjoint CFD – which is an interesting story in the context of this conference.

PARALLELS OF DEEP NEURAL NETWORK & ADJOINT CFD

Deep Learning / Neural Networks

The backpropagation algorithm is the workhorse of learning in neural networks and is the expresses for the partial derivative of the cost function.

The partial derivatives of the target function are computed with the chain rule similar to the adjoint formulation in CFD.

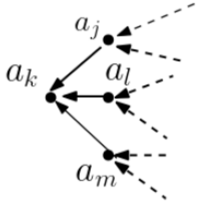

$$\frac{\partial f}{\partial a_k} = \frac{\partial f}{\partial a_j}\frac{\partial a_j}{\partial a_k} + \frac{\partial f}{\partial a_l}\frac{\partial a_l}{\partial a_k} + \frac{\partial f}{\partial a_m}\frac{\partial a_m}{\partial a_k}$$

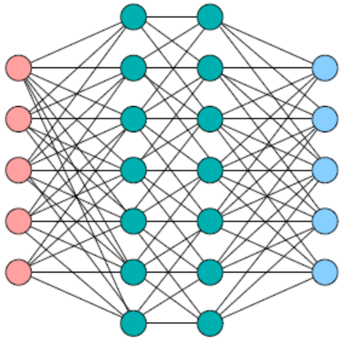

Adjoint Navier-Stokes calculation

CFD equations (primal field)

$$\begin{aligned}(\mathbf{v}\cdot\nabla)\,\mathbf{v} &= -\nabla p + \nabla\cdot(\nu\,\nabla\mathbf{v}) - \alpha\mathbf{v} \\ \nabla\cdot\mathbf{v} &= 0\end{aligned}$$

Adjoint CFD (dual field)

$$\begin{aligned}-(\nabla\mathbf{u})\,\mathbf{v} - (\mathbf{v}\cdot\nabla)\,\mathbf{u} &= -\nabla q + \nabla\cdot(\nu\,\nabla\mathbf{u}) - \alpha\mathbf{u} \\ \nabla\cdot\mathbf{u} &= 0\end{aligned}$$

Given a set of design variables, α, which control the geometry of the airfoil, wing or aircraft being designed, and a set of flow variables at discrete grid points, U, the aim is to minimize a scalar objective function J(U,α).

For a design variable, we can linearize about a base solution with the equation:

$$\frac{\mathrm{d}J}{\mathrm{d}\alpha} = \frac{\partial J}{\partial U}\frac{\mathrm{d}U}{\mathrm{d}\alpha} + \frac{\partial J}{\partial \alpha}$$

DNN & CFD – SIMILARITIES IN TRAINING AND INFERENCE

Training a DNN

For classification, for example the MNIST data set, we stream multiple data sets to train the network.

Training is the process of computing the matrix of assembled differential terms for inferencing.

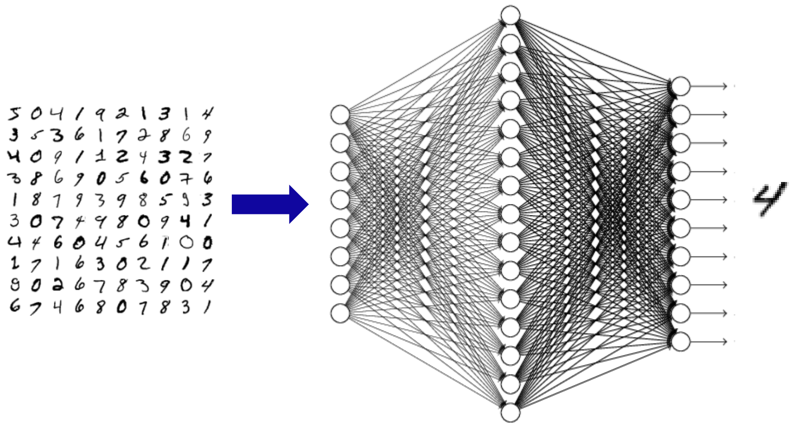

The Inferencing uses this matrix to quickly classify new input.

Design optimization with Adjoint CFD

Primal solution is the solution to the nonlinear steady Navier Stokes equations. The adjoint or dual solution solves the sensitivities of aerodynamic design quantities at each geometric point.

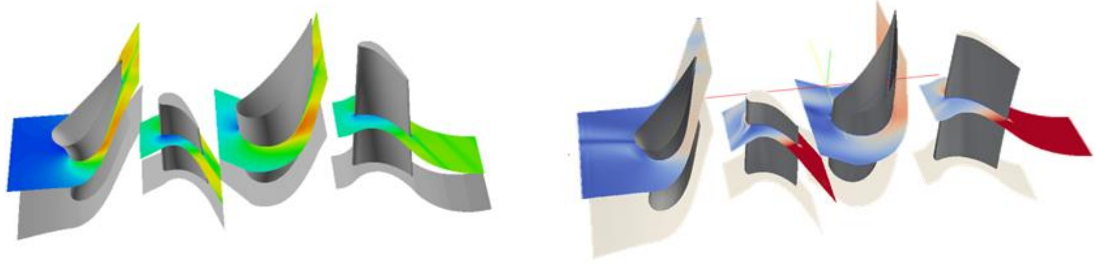

Once calculated, the design of experiments can be done very quickly by using traditional blade transformations, lean, tilt, and stacking, or more advanced optimization to optimize the aerodynamics of the airfoils.

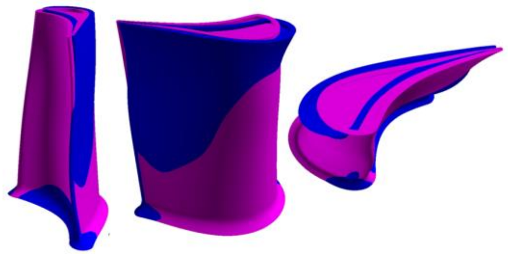

OTHER PARALLELS BETWEEN AI AND CFD

AI – LLMs and DNN

Precision: Training the matrix does not require high precision floating point mathematics. NVidia is investing significant chip real estate to low precision circuits.

Hardware: Large DNNs and Large Language Models (LLMs) and training takes place on specialized hardware or GPUs.

Frameworks: These tools use Python, PyTorch and other frameworks.

Large neural networks have emergent behavior that allow interaction with LLMs such as ChatGPT. Ask a question and you can get an answer – but it may hallucinate!

Adjoint CFD

Precision: CFD analysis uses double precision 64-bit floating point.

Hardware: Our Hydra CFD code runs on both CPUs and GPUs.

Frameworks: The Hydra CFD code is mostly Fortran along with C/C++ libraries. We use an abstraction library for MPI calls (Oplus) as well as an open-source (OP2) framework that allows us to run on accelerated hardware, such as NVidia or AMD GPUs

Aerospace engineers that intuitively think in geometry perturbations and pressure fields can quickly interact with aerodynamic solutions - **ChatCFD** anyone?