



**EXCELLERAT**



## TARGET AUDIENCE

- Industrial design engineers
- Experts from aeronautics industrial sectors



# Full Airplane Simulations using Alya on Heterogeneous Architectures

## CHALLENGE

Many of the future Exascale systems will be heterogeneous and include accelerators such as GPUs. With the explosion of parallelism, the expected performance of the various computing devices is more variable and the performance of the system components will be less certain. Therefore, leading-edge engineering simulation codes, such as Alya, need to be malleable enough to adapt to the new environment.

## SOLUTION

- Alya rises dynamic load balancing **to increase the parallel efficiency**.
- The load balancing method is **an efficient in-house space filling curve partitioner**.
- Focused on maximising the parallel performance of the mesh partition process **to minimise the load balancing overhead**.

## BENEFITS

- **Well-balanced co-execution** using both the CPUs and GPUs simultaneously.
- **23% faster** than using only the GPUs.
- Increased resilience of the software to system performance variability.

## UNIQUE VALUE

- Parallel mesh adaptation strategy for re-meshing.
- Design for massively parallel supercomputers.
- **Open source** package.

## PRODUCTS / SERVICES

- R&D Services.
- R&D joint collaborations.
- **Alya** (in-house HPC based multi-physics software) as a benchmarking tool.
- Training on HPC optimization.

Learn more about this [Success Story](#).

**Register to the EXCELLERAT Service Portal for free** and request support from EXCELLERAT staff on this application software.

**Visit EXCELLERAT Service Portal.**

## TARGET AUDIENCE

- Industrial R&D engineers
- R&D experts from aeronautics industrial sectors



# Enabling HPC through a Data Exchange & Workflow Portal

## CHALLENGE

Today, organisations and smaller industry partners face various problems in dealing with HPC calculations, HPC in general, or even access to HPC resources. In many cases, the calculations are complex and the potential users do not have the necessary expertise to fully exploit HPC technologies without support.

## SOLUTION

**SSC** and **RWTH** is developing a secure data exchange and transfer platform and tool as part of the EXCELLERAT project to facilitate the use of HPC and to make data transfer more efficient.

## BENEFITS

- **Reduction of HPC complexity** due to web frontend.
- **Time and cost savings** due to a high degree of automation.
- **Efficient, user-friendly and secure post-processing/data analytics.**

## UNIQUE VALUE

- Provision and operation/support of the platform in order to significantly **facilitate the use of HPC.**
- **"Automatic" accounting and pre-calculation of resources.**
- Providing a user-friendly and scalable data analytic tool to **efficiently perform modal decomposition of large-scale simulation data.**

## PRODUCTS / SERVICES

- **Alya, BSC:** Multiphysics simulation code available under license, available as Open Source / SaaS, and Consultancy about parallel mesh adaptation strategy based on the use of Gmsh library for re-meshing.
- **Modal decomposition toolkit, RWTH:** Scalable data analytic tool: modal decomposition of large-scale simulation data.
- **SSC platform, SSC:** Data exchange and transfer platform.
- **HPC resources** to run the simulation including Data Exchange & Workflow.
- **SSC & RWTH:** Consultancy, best practice guide, workshop and training about modal decomposition of simulation data.

Learn more about this [Success Story](#).

[Register to the EXCELLERAT Service Portal for free](#) and request support from EXCELLERAT staff on this application software.

[Visit EXCELLERAT Service Portal.](#)



## TARGET AUDIENCE

- Code developers
- R&D engineers in aeronautics sectors



# Bringing combustion design industrial end-users to Exascale computing

## CHALLENGE

Some physical processes such as soot formation are so CPU-intensive and non-deterministic that it is not yet possible to apply predictive modelling techniques to them, which limits our insights to ad hoc correlations and preliminary assumptions. Moving these runs to Exascale level computing systems will allow **simulations that are longer by orders of magnitudes**, and will **achieve the compulsory statistical convergence that is required for a design tool**.

## SOLUTION

**CERFACS** developed a **predictive model to port, benchmark and optimise the AVBP code** for Arm and AMD architectures. This will ensure early access to these new architectures and prepare our codes for the widespread availability of systems equipped with these processors.

## BENEFITS FOR INDUSTRIAL END USERS: ENGINEERS, ACADEMIC RESEARCHERS, AND CODE DEVELOPERS

- Code ready for the **widespread access of the Arm Rome architecture**.
- Strong and weak scaling measurements **up to 128,000 cores**.
- Initial optimisations for Arm architectures.

## UNIQUE VALUE

**AVBP** offers the possibility of using unstructured grids than can be **automatically adapted to the flow at runtime, allowing for automatic tracking of the zones of interest**. Additionally, mesh quality criteria in adaptive mesh refinement can compensate for any defects on the original grid that can lead to numerical issues.

## PRODUCTS / SERVICES

- Code Simulation AVBP open source.
- Training on MPI/OpenMP high performance computing.
- Consultancy on code optimization and porting.
- Best practice guide on bringing industrial end-users to Exascale computing.
- An industrial-level combustion design tool on 128,000 cores.

Learn more about this [Success Story](#).

**Register to the EXCELLERAT Service Portal for free** and request support from EXCELLERAT staff on this application software.

**Visit EXCELLERAT Service Portal.**

## TARGET AUDIENCE

- Code developers
- R&D engineers from aeronautics sectors



# Enabling parallel mesh adaptation with Treadapt

## CHALLENGE

Advanced meshing software is limited to sequential or shared memory architectures. To bypass bottlenecks, codes have turned to mesh adaptation as a solution, but massively parallel mesh adaptation workflows remain scarce and require efficient load balancing, interpolation and remeshing techniques.

## SOLUTION

**Treadapt** is a library for massively parallel mesh adaptation developed by CERFACS. It generates a partitioned domain where the mesh generation tool MMG can be employed while freezing the parallel interfaces. **Rebalancing and adaptation then occur iteratively until converged within a user defined tolerance.**

## BENEFITS

- Time reduction from three days to one hour to generate a one billion element mesh.
- Reduce result quality dependency on user experience.
- Improve high fidelity results.

## UNIQUE VALUE

**AVBP** with the library Treadapt offers the possibility to use unstructured grids that can be automatically adapted to the flow at runtime allowing for automatic tracking of the zones of interest. Additionally **mesh quality criteria in AMR can compensate for any defects on the original grid that might lead to numerical issues.**

## PRODUCTS / SERVICES

- Treadapt is part of an **open source** package.
- **Code Simulation AVBP** open source.
- Training on **Enabling parallel mesh adaptation with Treadapt.**
- Consultancy on **mesh adaptation methods applied to aeronautics and combustion.**
- Best practice guide on **Enabling parallel mesh adaptation with Treadapt.**

Learn more about this [Success Story](#).

[Register to the EXCELLERAT Service Portal for free](#) and request support from EXCELLERAT staff on this application software.

[Visit EXCELLERAT Service Portal.](#)

## TARGET AUDIENCE

- Application developers



# Accelerating engineering codes using reconfigurable architectures

## CHALLENGE

Many engineering and scientific codes are not entirely compute-bound, but do suffer from bottlenecks including how quickly data can be loaded to the CPU from the main memory. **To fully exploit future Exascale systems**, both techniques and hardware are required to address properties which **constrain application performance and scaling**.

## SOLUTION

We are exploring the benefits of reconfigurable architectures (FPGAs) for our codes. **FPGAs enable a chip to be configured** to electronically represent an application. Different parts of the FPGA can compute on different elements of data, whilst other parts are fetching and reordering data ready for later stages.

## BENEFITS

- **Optimisation of code performance.**
- **Reduced power draw.**
- **Developed algorithm level techniques** that suit dataflow architectures, typically speeding up resultant code by a 1000-times or greater.

## UNIQUE VALUE

- FPGAs typically have an order of **magnitude more bandwidth** than other technologies.
- FPGAs enable **a tailor-made effect and support the effectiveness of the chip** in the execution of the simulation code.

## PRODUCTS / SERVICES

- **Consultancy** on "How to port simulation code to the HPC architecture".

Learn more about this [Success Story](#).

**Register to the EXCELLERAT Service Portal for free** and request support from EXCELLERAT staff on this application software.

**Visit EXCELLERAT Service Portal.**

## TARGET AUDIENCE

- Code developers
- Hardware providers



# Enabling Nek5000 on GPU systems

## CHALLENGE

A collaboration between EXCELLERAT and the Swedish e-Science Research Centre (SeRC) focused on **improving the performance of the EXCELLERAT core code Nek5000** and getting it ready to run on future hardware. Porting and optimising Nek5000 on GPU systems using hardware provided by NVIDIA and AMD was the key challenge of this collaboration.

## SOLUTION

The CPU version of Nek5000 to run on NVIDIA GPUs using OpenACC directives was extended. Additionally, some **kernels were merged** by rewriting them in CUDA. Several routines were found to have significant memory bank conflicts when certain input was used. These conflicts are now removed.

## BENEFITS

- **Significant reduction of kernel runtime**, e.g., from 800  $\mu$ s to 200  $\mu$ s.
- **Duration of kernels** affected by memory bank conflicts **was improved by around 20%**.

## UNIQUE VALUE

Nek5000, a code written in Fortran and C, provides a good example of **how to migrate large legacy codes towards Exascale**. Nek5000 is a long-established code with **very good, strong scaling**, and the experience gained in porting it to GPUs is highly valuable in porting other legacy codes.

## PRODUCTS / SERVICES

- **Training and consultancy** on how to migrate large legacy codes towards Exascale computing on GPUs.

Learn more about this [Success Story](#).

**Register to the EXCELLERAT Service Portal for free** and request support from EXCELLERAT staff on this application software.

**Visit EXCELLERAT Service Portal.**

## TARGET AUDIENCE

- Code developers
- CFD users



# Enabling sustainable GPU acceleration on a Fortran legacy code

## CHALLENGE

AVBP is **cutting-edge software**; however, many of the most powerful HPC systems heavily rely on GPU accelerators, and this trend will likely continue with the upcoming Exascale systems. **Extending AVBP to support GPU architectures** in addition to the current CPU implementation seems mandatory to efficiently use computing resources in the Exascale era.

## SOLUTION

We ported AVBP to GPUs without impacting its accessibility and maintainability by using OpenACC. It **enables offload computations to GPUs** by adding simple directives thereby **reducing the impact in the source code**. Only parts of AVBP require a full rewrite to run on GPU and keep CPU compatibility.

## BENEFITS

- Release of a **GPU-ready CFD code**.
- Leveraged **performance of GPU nodes** on HPC systems superior to pure CPU nodes (by a factor of 5 at least).
- Maintained **accessibility of the code** to non-HPC developers.

## UNIQUE VALUE

- Reduction of time to solution using **GPUs**.
- Compatibility with both **CPU and GPU systems**.
- Single maintainable **source code**.
- Reduce **energy consumption** for a simulation when using GPUs compared to equivalent runtime on CPUs.

## PRODUCTS / SERVICES

- **CFD code for compressible reactive flows** running on GPU systems.
- **Best practices** for GPU porting for legacy code.
- **Best practices** for CFD using GPU hardware.

Learn more about this [Success Story](#).

**Register to the EXCELLERAT Service Portal for free** and request support from EXCELLERAT staff on this application software.

**Visit EXCELLERAT Service Portal.**



## TARGET AUDIENCE

- R&D engineers in aerospace sector



# A novel framework for online estimation of the uncertainties in turbulent flow statistics

## CHALLENGE

The computation of the statistics of turbulent flows contains a level of uncertainty due to the finite number of time samples taken from the CFD simulations. To accurately quantify the uncertainty, it must be taken into account that the time samples for a turbulent flow quantity at any spatial point are autocorrelated up to a generally unknown number of lags. In addition, **storing the time series data of turbulent flows for the uncertainty quantification (UQ) analyses is not computationally efficient.**

## SOLUTION

A framework for in-situ estimation of the time-averaging uncertainties in turbulent flow simulations was developed. The framework consists of **low-storage in-situ UQ algorithms, adapters for sampling from a CFD solver, and interfaces between the CFD and UQ suites.** It is computationally efficient and provides accurate uncertainty estimations that **can be applied to any CFD solver.**

## BENEFITS

- In-situ monitoring of intermediate UQ results.
- **Less than 5% computational overhead** due to the UQ techniques.
- Reduction of data I/O **to <1% during/after** the simulation.

## UNIQUE VALUE

The software tool **provides accurate in-situ estimators for the uncertainty in turbulent flow statistics.** It can be attached to any flow solver providing a ParaView Catalyst interface, **resulting in a strong data I/O reduction** and an increase of the spatial resolution of the output without reducing the accuracy compared to offline UQ estimators.

## PRODUCTS / SERVICES

- **Software tools** for in-situ UQ of turbulent flow simulations.
- **Consulting activities** on in-situ UQ for autocorrelated and cross-correlated time-series data.
- The machinery and framework can be used for **developing an in-situ version of various data-driven techniques.**

Learn more about this [Success Story](#).

**Register to the EXCELLERAT Service Portal for free** and request support from EXCELLERAT staff on this application software.

**Visit EXCELLERAT Service Portal.**

## TARGET AUDIENCE

- Industry users in aerospace
- Research organisations in aerospace



# Transparent Integration of HPC Technologies into the CFD Software CODA

## CHALLENGE

Providing high-fidelity aerodynamics simulations requires vast computational resources. Leading-edge HPC systems are increasingly complex: consisting of millions of heterogeneous processing elements including GPUs. To harness their full potential, CFD software such as **CODA** require not only the efficient utilization of the CPU but must also support emerging technologies such as GPUs.

## SOLUTION

Within EXCELLERAT, the capabilities of Spliss (the Sparse Linear Systems Solver) **have been extended and optimized for GPUs**. Spliss allows CODA the efficient and transparent usage of GPUs for linear systems solving while at the same time hiding the inherent complexity.

## BENEFITS

- Simulation of aircraft aerodynamics with **shorter time to solution** and **higher-quality results**.
- **Readiness for current and future HPC systems** including accelerators.

Learn more about this [Success Story](#).

[Register to the EXCELLERAT Service Portal for free](#) and request support from EXCELLERAT staff on this application software.

[Visit EXCELLERAT Service Portal](#).

# ABOUT EXCELLERAT

The EXCELLERAT project is a single point of access for expertise on how data management, data analytics, visualisation, simulation-driven design and Co-design with high-performance computing (HPC) can benefit engineering, especially in the aeronautics, automotive, energy and manufacturing sectors.

Subscribe to the EXCELLERAT Service Portal



[www.excellerat.eu](http://www.excellerat.eu)

The EXCELLERAT project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823691