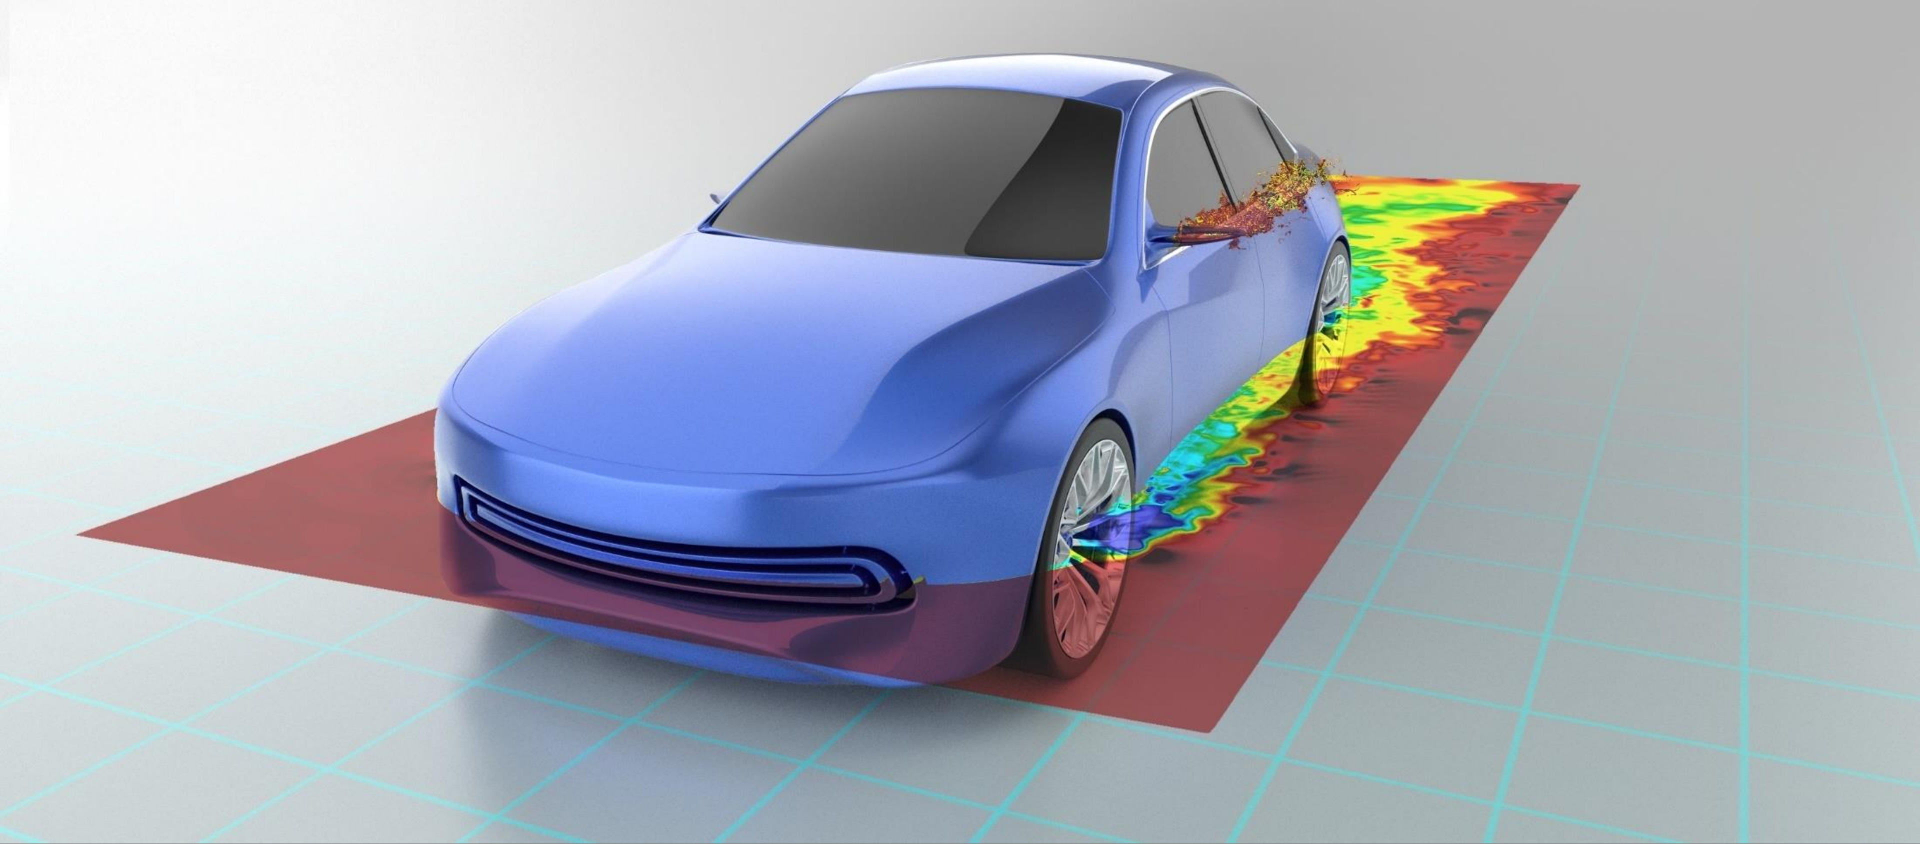

S9719 – ULTRA-FAST HIGH-FIDELITY CFD SIMULATIONS FOR AUTOMOTIVE AERODYNAMICS

Christoph Niedermeier – Development Manager

GTC San Jose, March 20, 2019

Altair

AUTOMOTIVE AERODYNAMICS

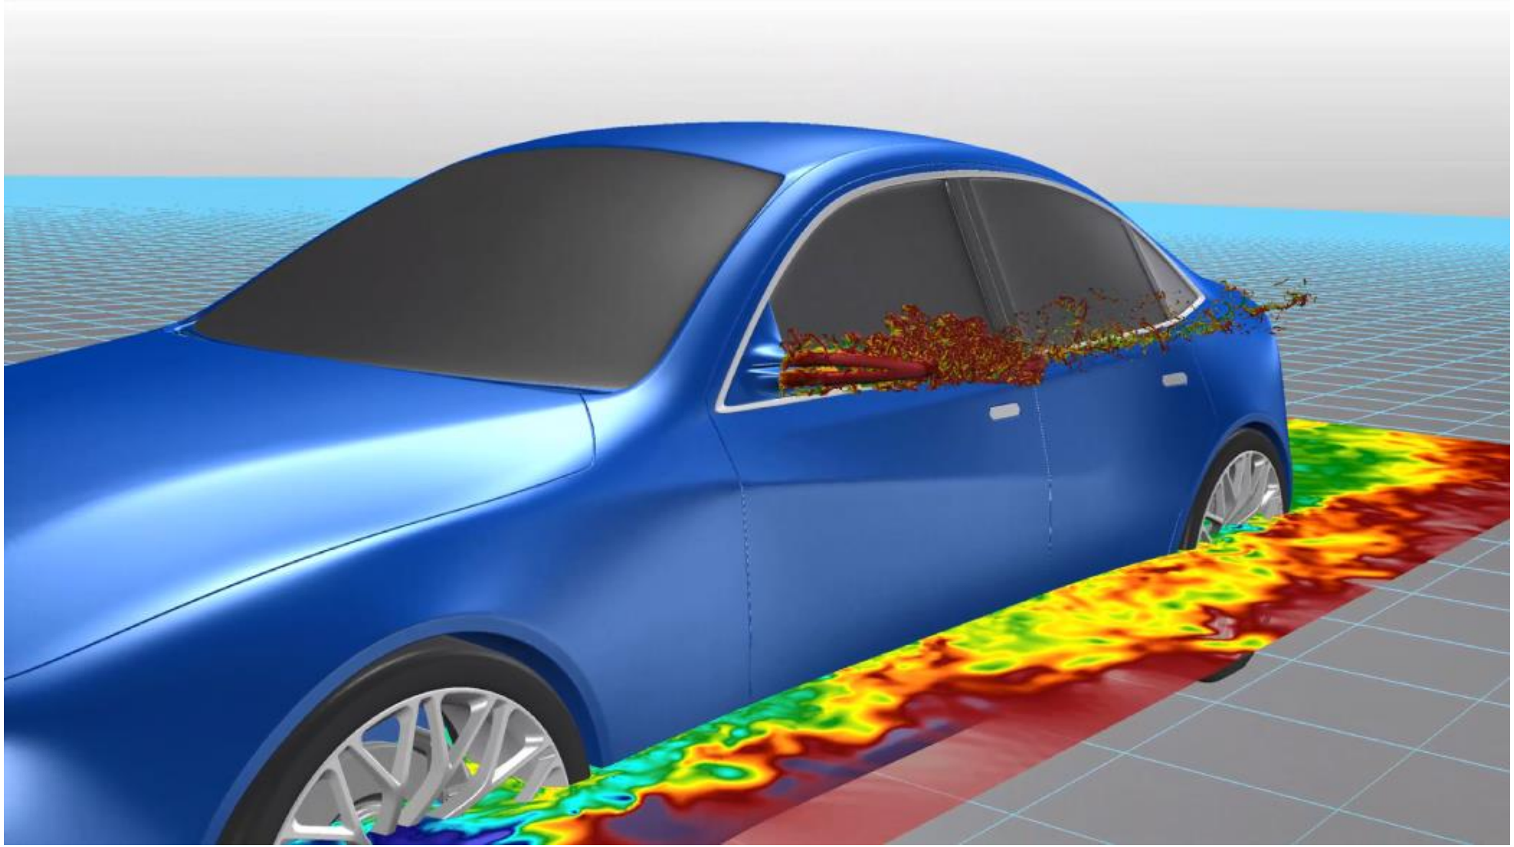

COMPUTATIONAL EFFORT PER SECOND

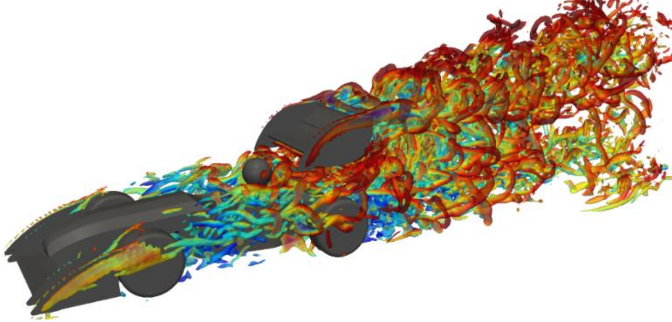

1E8

Voxels

1E5

Time steps

1E13

Node updates

1E18

FLOP

AVAILABLE PERFORMANCE

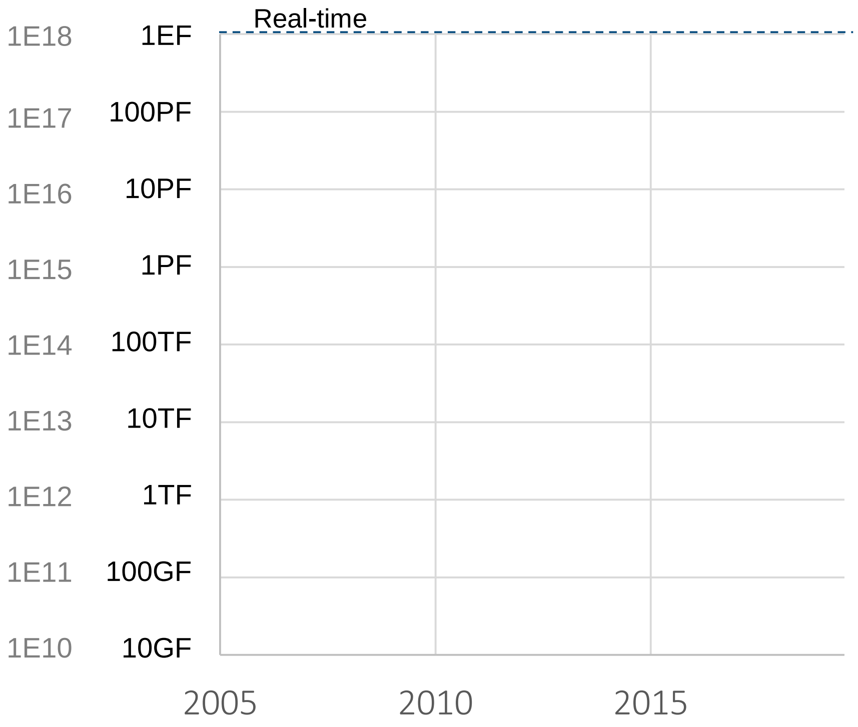

AVAILABLE PERFORMANCE

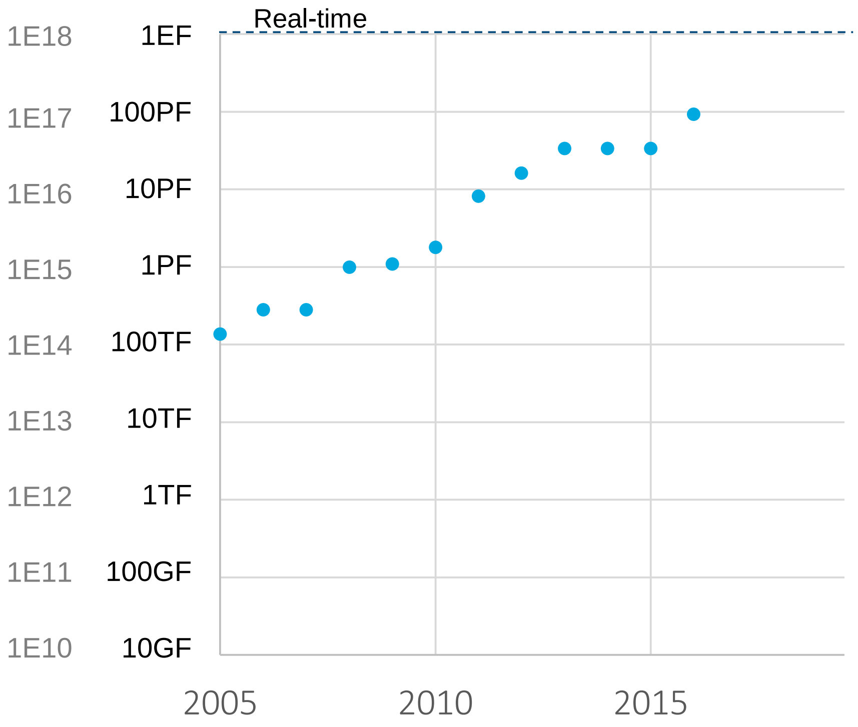

#1 TOP 500

~10s per second

AVAILABLE PERFORMANCE

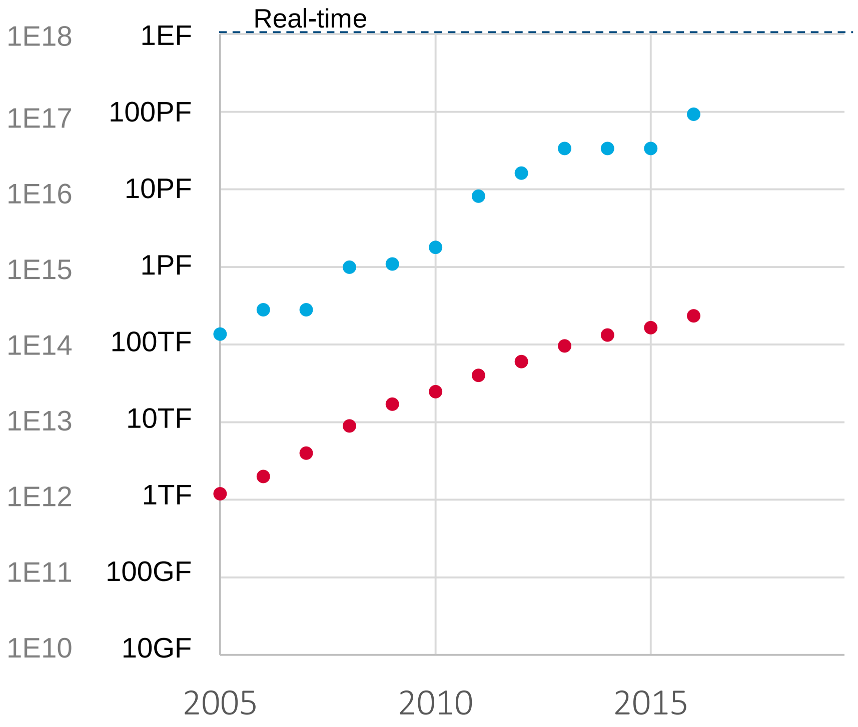

#1 TOP 500 ~10s per second

#500 TOP 500 ~1 hour per second

AVAILABLE PERFORMANCE

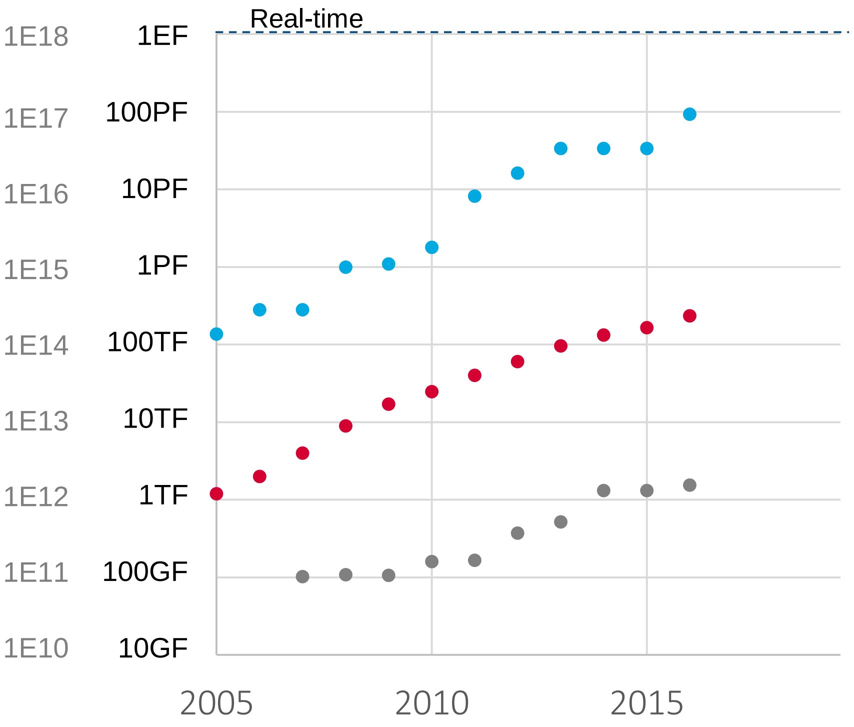

#1 TOP 500 ~10s per second

#500 TOP 500 ~1 hour per second

CPUs (Intel) ~1 week per second

AVAILABLE PERFORMANCE

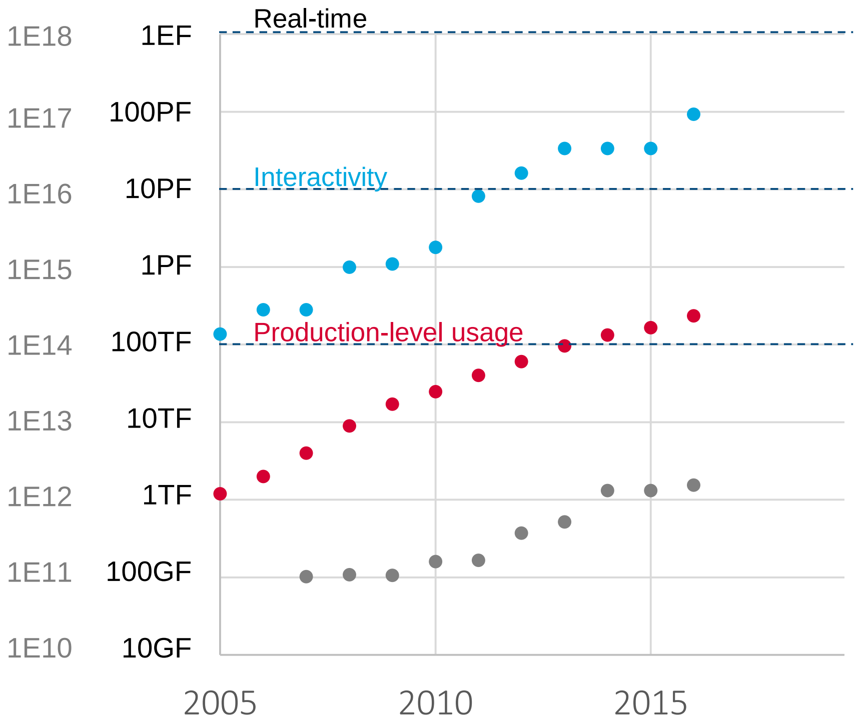

#1 TOP 500 ~10s per second

#500 TOP 500 ~1 hour per second

CPUs (Intel) ~1 week per second

AVAILABLE PERFORMANCE

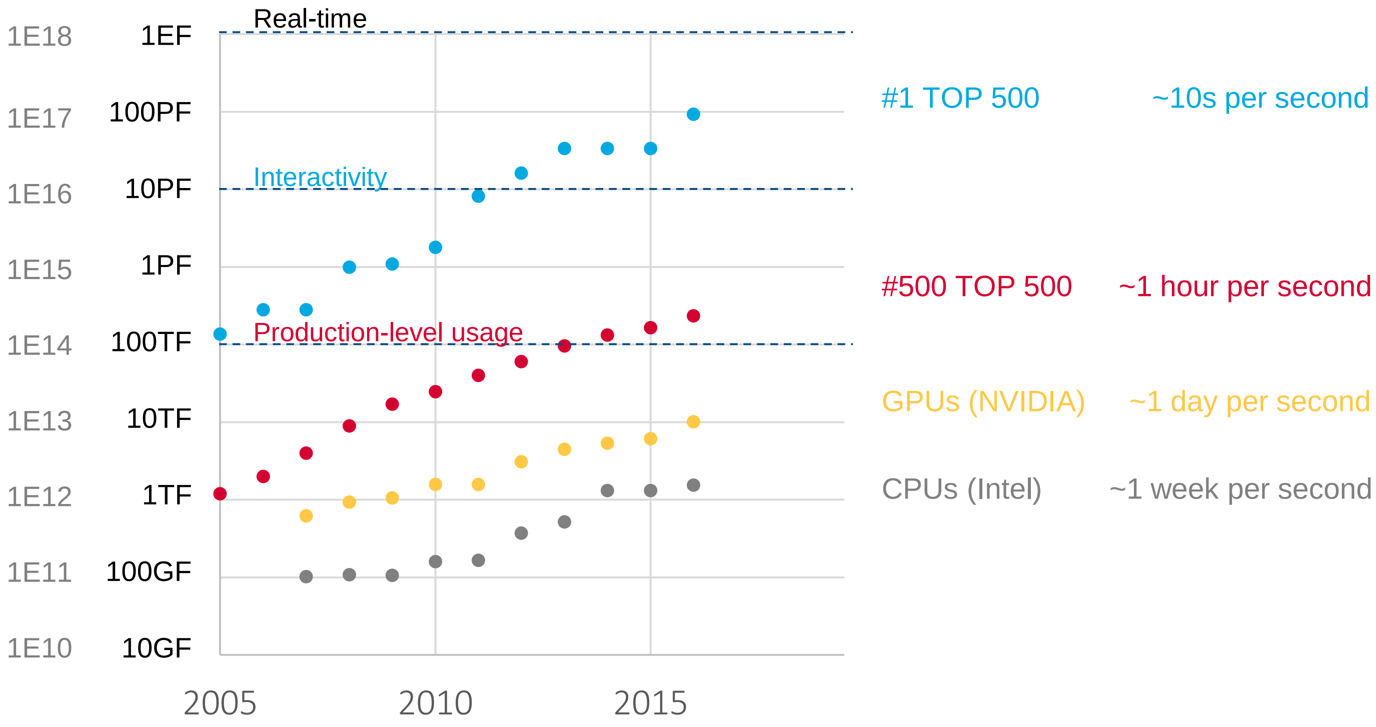

HIGH-FIDELITY CFD NEAR REAL-TIME

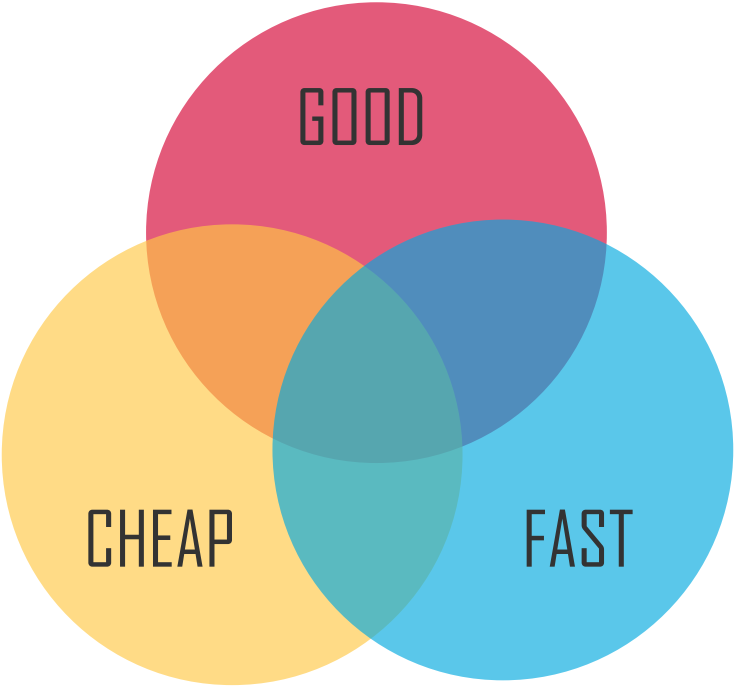

HIGH-FIDELITY CFD NEAR REAL-TIME

Accurate wall-modeled LES

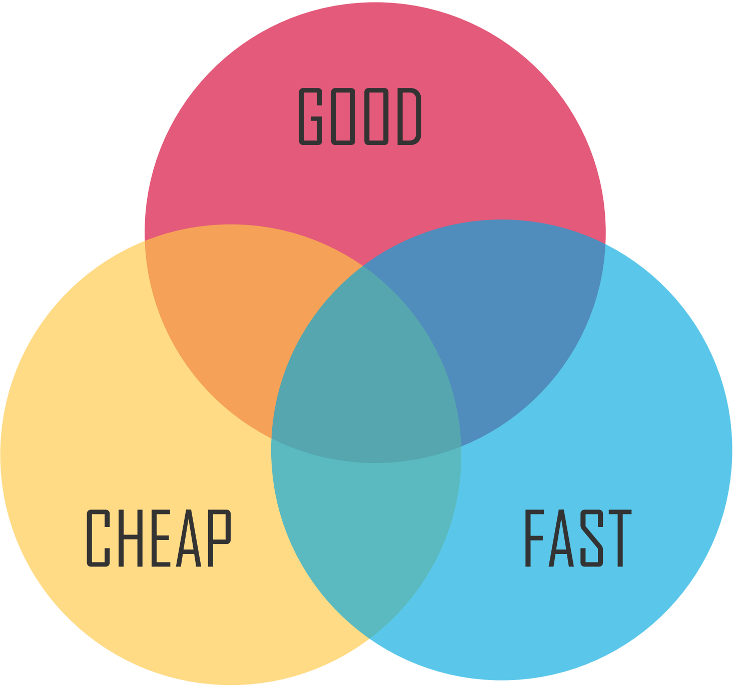

Costs of O(500) USD
per simulation

Interactive simulations
near real-time

HIGH-FIDELITY CFD NEAR REAL-TIME

Accurate wall-modeled LES

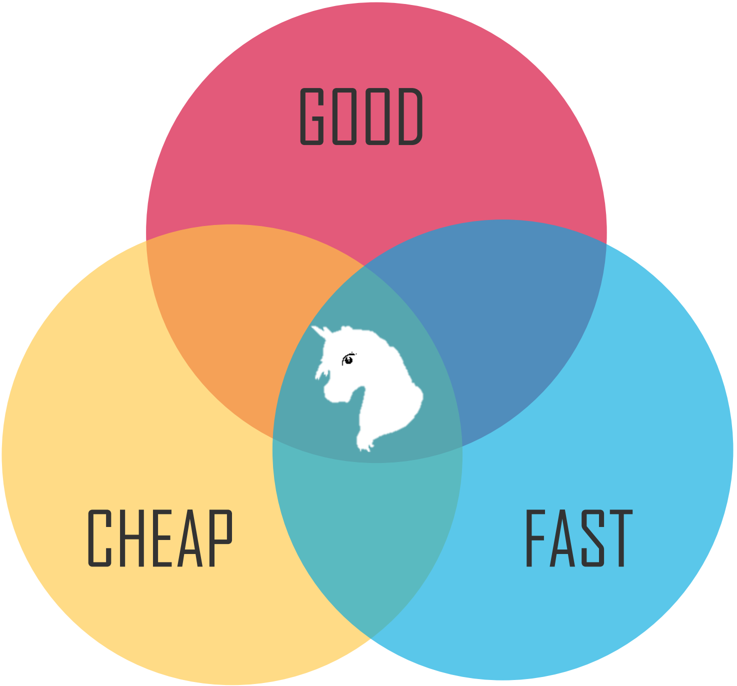

Costs of O(500) USD
per simulation

Interactive simulations
near real-time

METHODOLOGY

HOW TO USE AI SUPERCOMPUTERS FOR HIGH-FIDELITY CFD

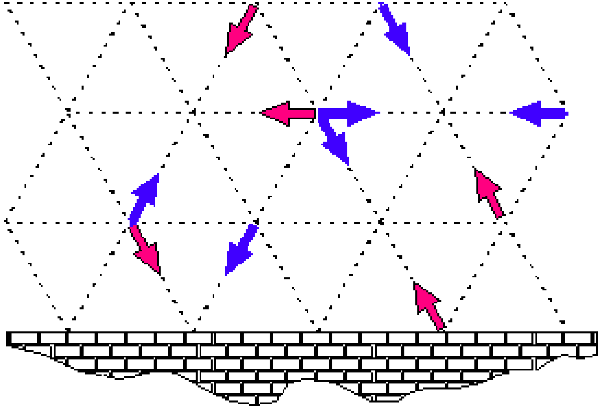

A GLIMPSE OF LBM

- First comparable numerical approach for CFD: Lattice Gas Automata
  - HPP model (1976), FHP model (1986)
  - Bool variable: 0 or 1 particle per discrete direction
  - Collision rules, propagation rules

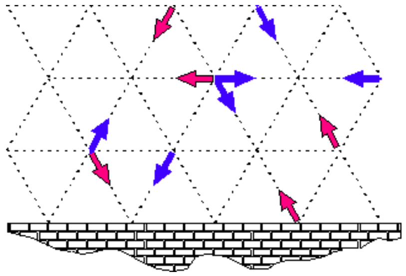

A GLIMPSE OF LBM

- First comparable numerical approach for CFD: Lattice Gas Automata
  - HPP model (1976), FHP model (1986)
  - Bool variable: 0 or 1 particle per discrete direction
  - Collision rules, propagation rules

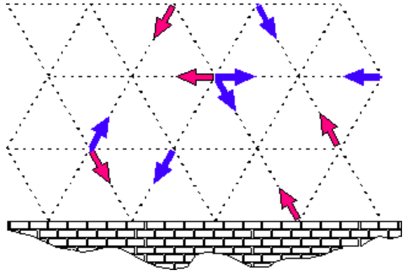

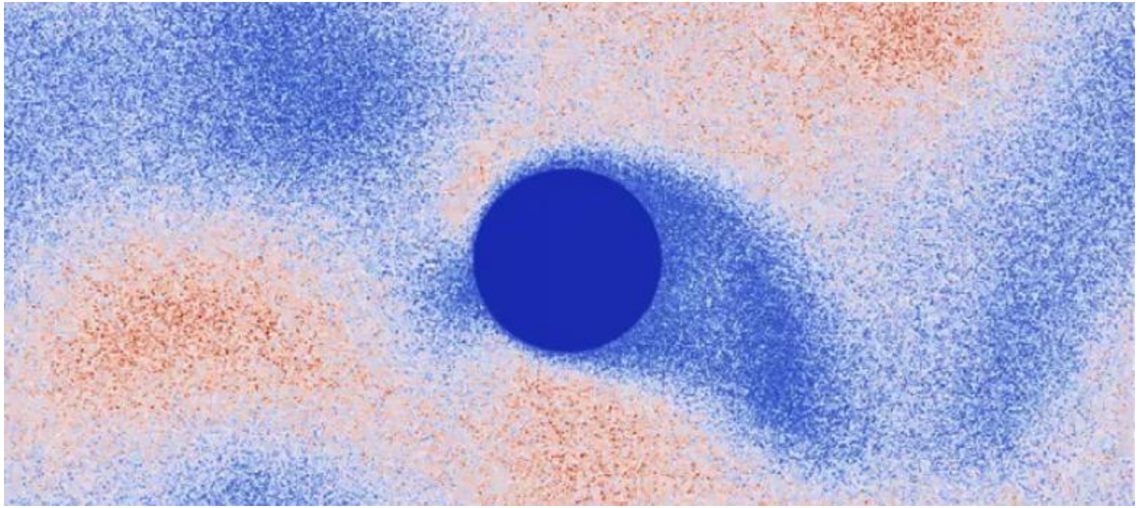

https://keva92.github.io/lgca/

A GLIMPSE OF LBM

- Lattice Boltzmann Method (LBM)
  - Essentially replaces the Boolean variable with a particle distribution function $f(t, \mathbf{x}, \xi)$
  - The algorithmic nature of the method remains unchanged

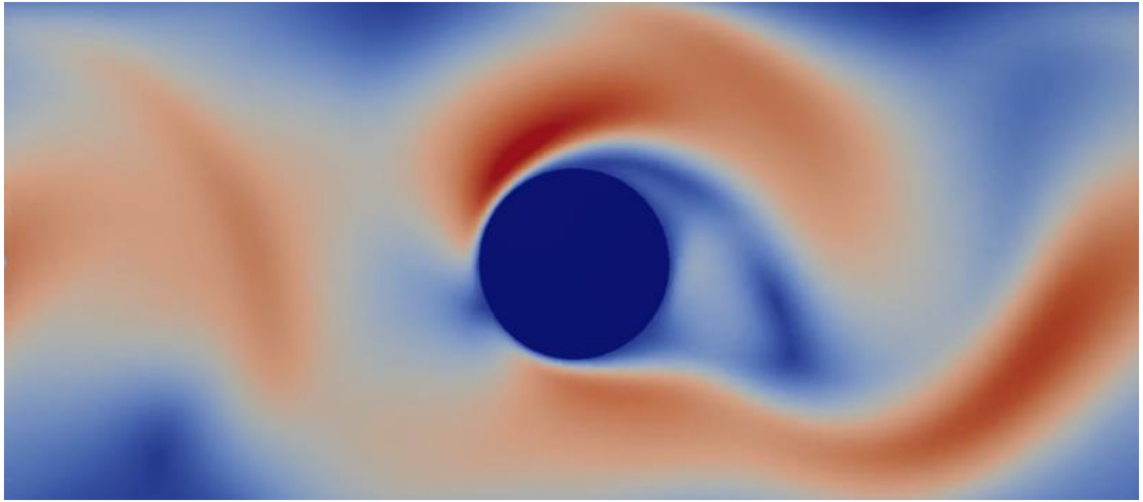

A GLIMPSE OF LBM

- Lattice Boltzmann Method (LBM)
  - Essentially replaces the Boolean variable with a particle distribution function $f(t, \mathbf{x}, \xi)$
  - The algorithmic nature of the method remains unchanged

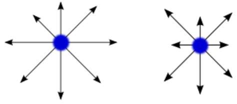

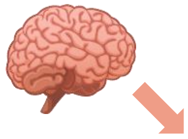

$$\bar{f}_i(t, \mathbf{x}) = f_i(t, \mathbf{x}) + \Omega_i$$

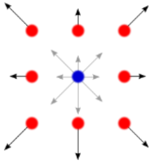

$$f_i(t + \Delta t, \mathbf{x} + \mathbf{e}_i \Delta t) = \bar{f}_i(t, \mathbf{x})$$

Collision

Propagation

A GLIMPSE OF LBM

- Lattice Boltzmann Method (LBM)
  - Essentially replaces the Boolean variable with a particle distribution function $f(t, \mathbf{x}, \xi)$
  - The algorithmic nature of the method remains unchanged

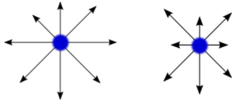
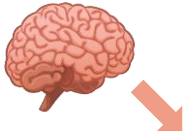

$$\bar{f}_i(t, \mathbf{x}) = f_i(t, \mathbf{x}) + \Omega_i$$

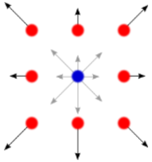

$$f_i(t + \Delta t, \mathbf{x} + \mathbf{e}_i \Delta t) = \bar{f}_i(t, \mathbf{x})$$

Collision

Propagation

- Many very sophisticated LBM collision operators are available, e.g. Cumulant-based operators with very low numerical diffusion, high stability and implicit subgrid-scale modeling

M. Geier, A. Pasquali, M. Schönherr, *Journal of Computational Physics* **348**(1):862-888, 2017
M. Geier, A. Pasquali, M. Schönherr, *Journal of Computational Physics* **348**(1):889-898, 2017

HOW TO USE AI SUPERCOMPUTERS FOR HIGH-FIDELITY CFD

- GPUs deliver tremendous speed for explicit, local algorithms – LBM is a perfect match

- High pace of GPU technology development leads to big performance advantage of GPU-based LBM on recent hardware (~ 1 GPU-h vs. ~ 100 core-h of CPU-based LBM)

- Similar or even lower investment costs and a lower energy consumption compared to CPU-clusters of same performance

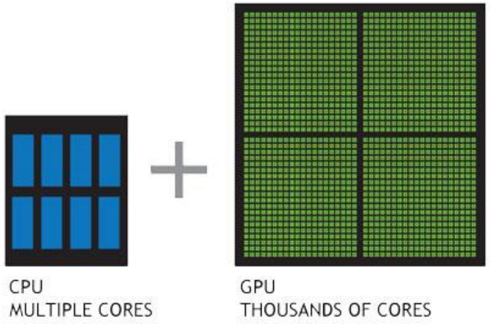

HIGHLY ACCURATE CFD ON MULTIPLE GRAPHICS PROCESSING UNITS

MULTI-GPU SIMULATIONS OF EXTERNAL AERODYNAMICS

Altair ultraFluidX™

Commercial

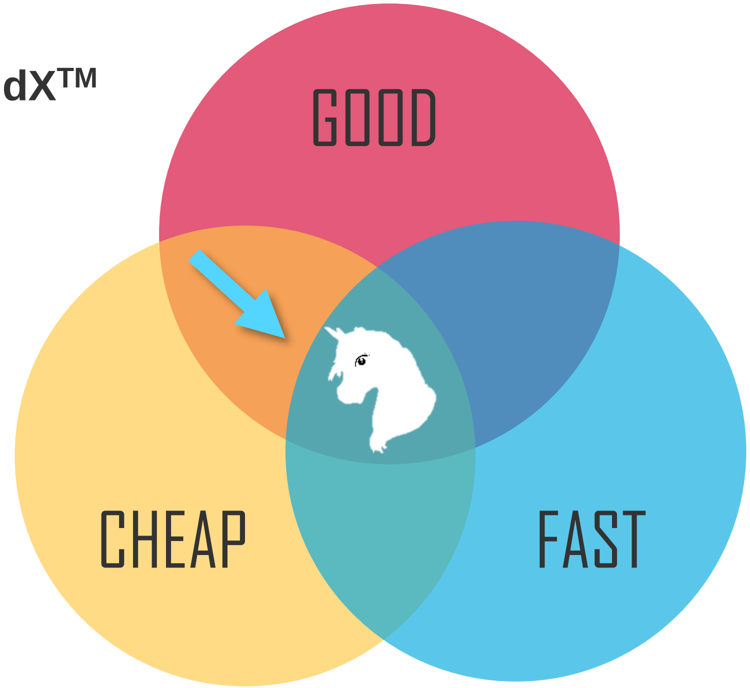

ALTAIR ULTRAFLUIDX™

- High-fidelity flow field computations based on the Lattice Boltzmann method
  - Cumulant collision operator, Smagorinsky LES turbulence model, TBLE-based wall model

- Tailor-made for external aerodynamics
  - Easy and robust handling of complex automotive geometries
  - Extended automotive features, including belt systems, boundary layer suction, porous media support
  - Rotating wheel modeling based on wall-velocity boundary conditions and MRF approaches
  - Advanced post-processing capabilities (window averaging, spatial force contribution, probe outputs, …)

- Aiming at overnight simulations of external aerodynamics
  - CUDA-aware MPI support for multi-GPU and multi-node usage

PREPROCESSING

- Automated volume mesh generation with local grid refinement and support for baffle parts
- Octree-based grids with 4^3 lattice nodes per block

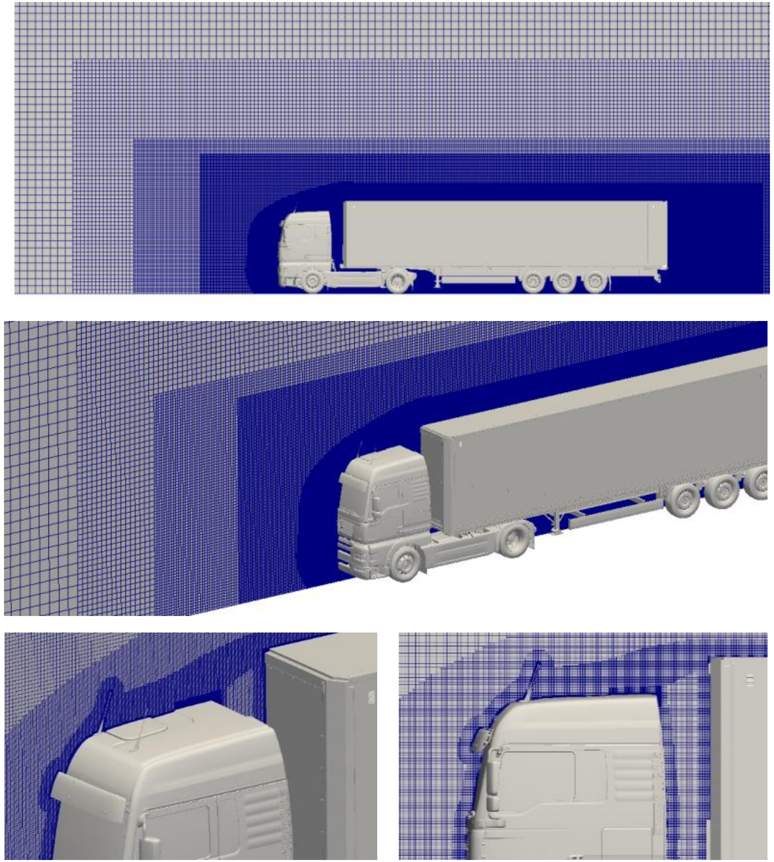

PREPROCESSING

- Automated volume mesh generation with local grid refinement and support for baffle parts
- Octree-based grids with 4^3 lattice nodes per block
- Kinks and thin parts are fully resolved in the volumetric preprocessing

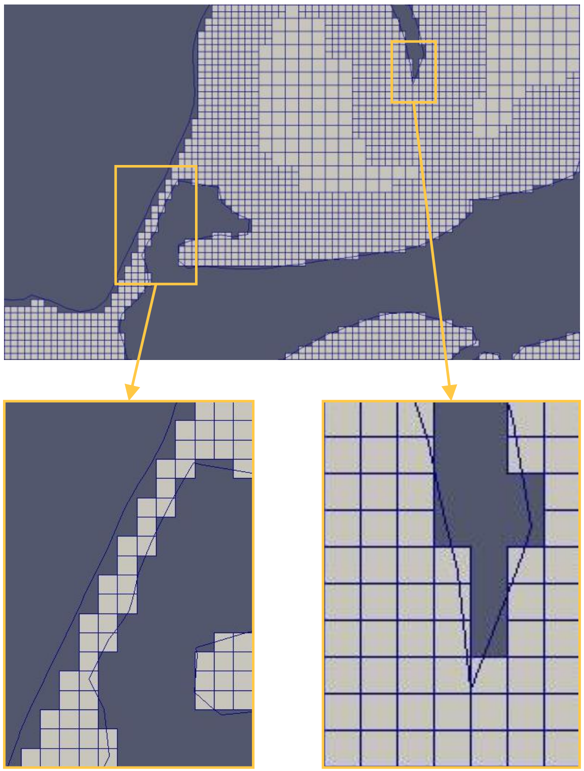

PREPROCESSING

- Automated volume mesh generation with local grid refinement and support for baffle parts
- Octree-based grids with 4^3 lattice nodes per block
- Kinks and thin parts are fully resolved in the volumetric preprocessing
- Fast generation of production-level volume meshes in < 1 hour (CPU-based, MPI-parallel)
- Domain decomposition with space-filling curves, separately on each refinement level

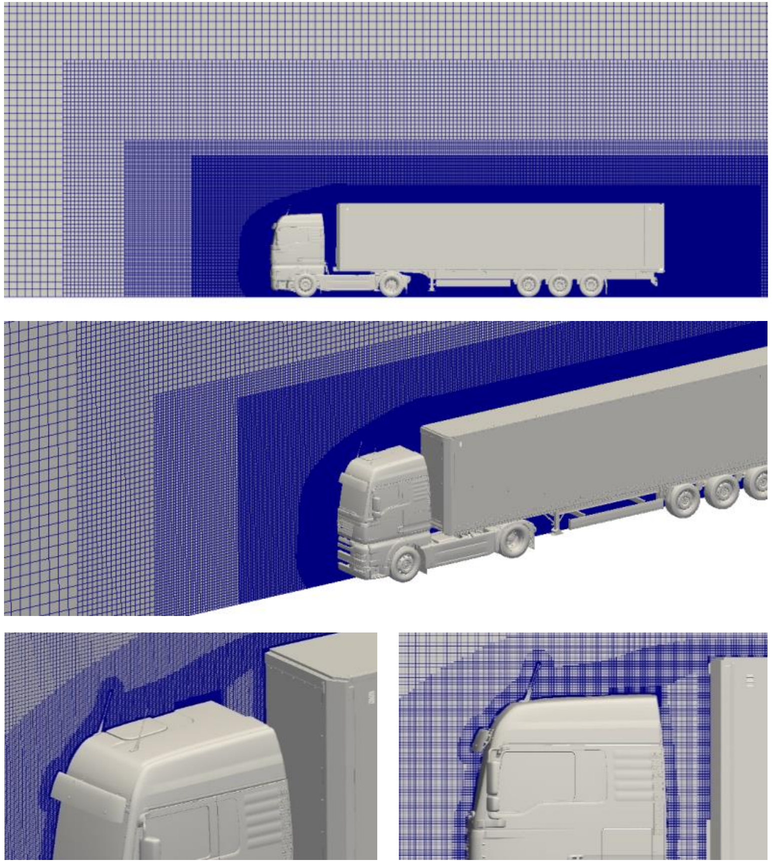

VALIDATION EXAMPLE

- 40% scale DrivAer Notchback with engine bay flow (http://www.drivaer.com)

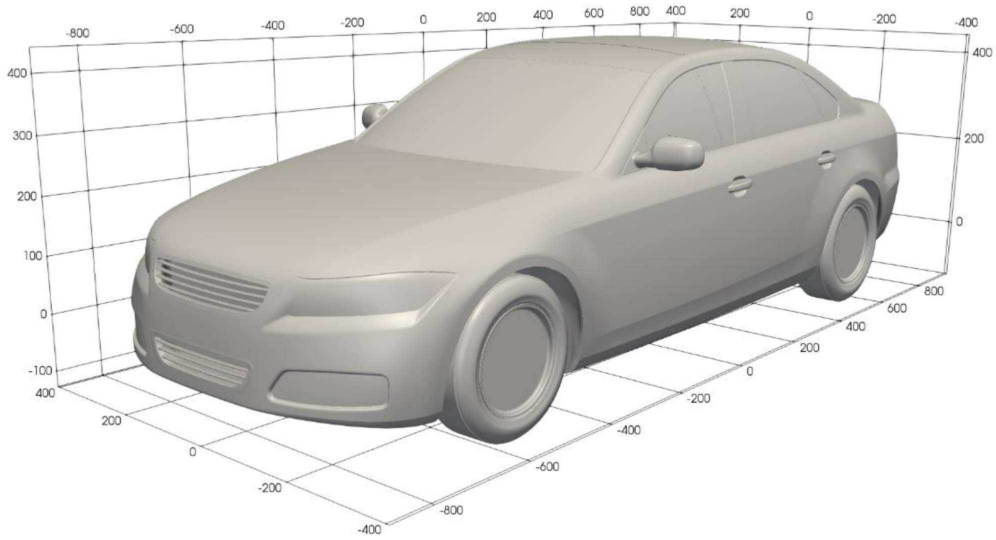

Tested for different radiator opening ratios

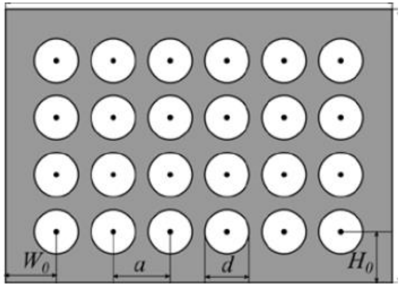

Collin et al. (2017)

VALIDATION EXAMPLE

- 40% scale DrivAer Notchback with engine bay flow (http://www.drivaer.com)

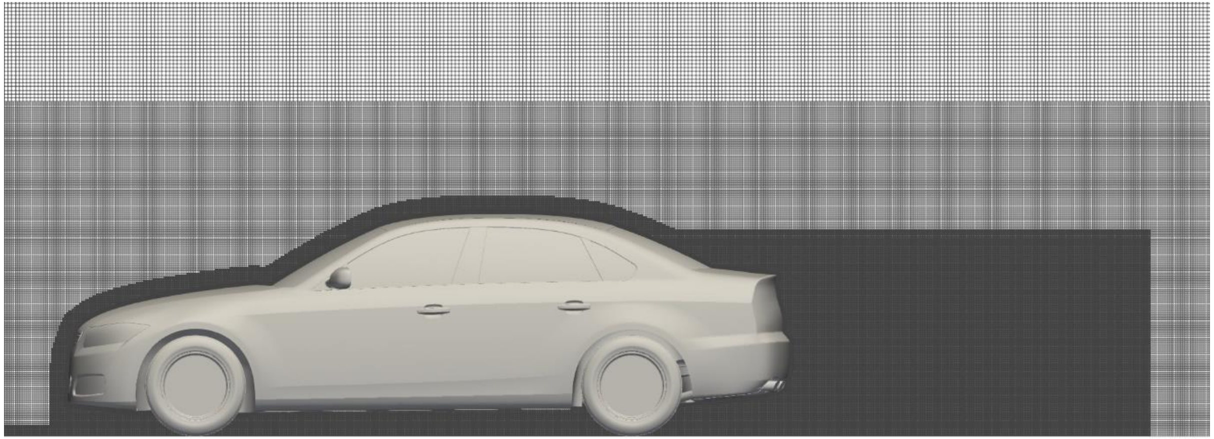

156 M voxels
30 flow passes
(> 3 s full scale)

Ultra-fast, High-fidelity Computational Fluid Dynamics on GPUs for Automotive Aerodynamics
C.A. Niedermeier, C. Del Bene, C.F. Janßen, B. Schnepf, M. Ratzel, T. Indinger, WCCM 2018

VALIDATION EXAMPLE

- 40% scale DrivAer Notchback with engine bay flow (http://www.drivaer.com)

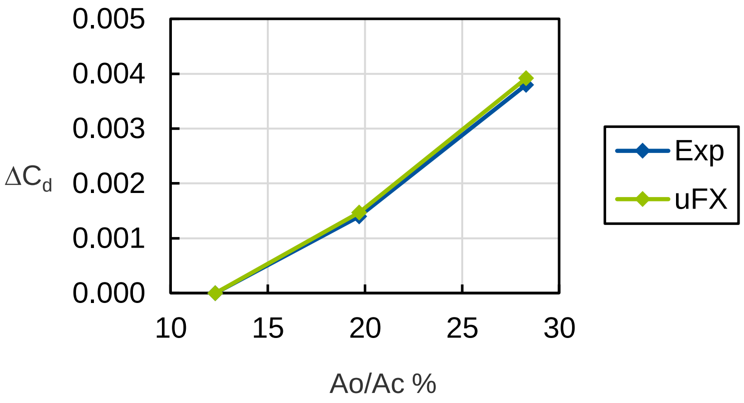

Ultra-fast, High-fidelity Computational Fluid Dynamics on GPUs for Automotive Aerodynamics
C.A. Niedermeier, C. Del Bene, C.F. Janßen, B. Schnepf, M. Ratzel, T. Indinger, WCCM 2018

TOWARDS PRODUCTION-LEVEL USAGE

Can we reach the goal of overnight simulations without compromising on resolution?

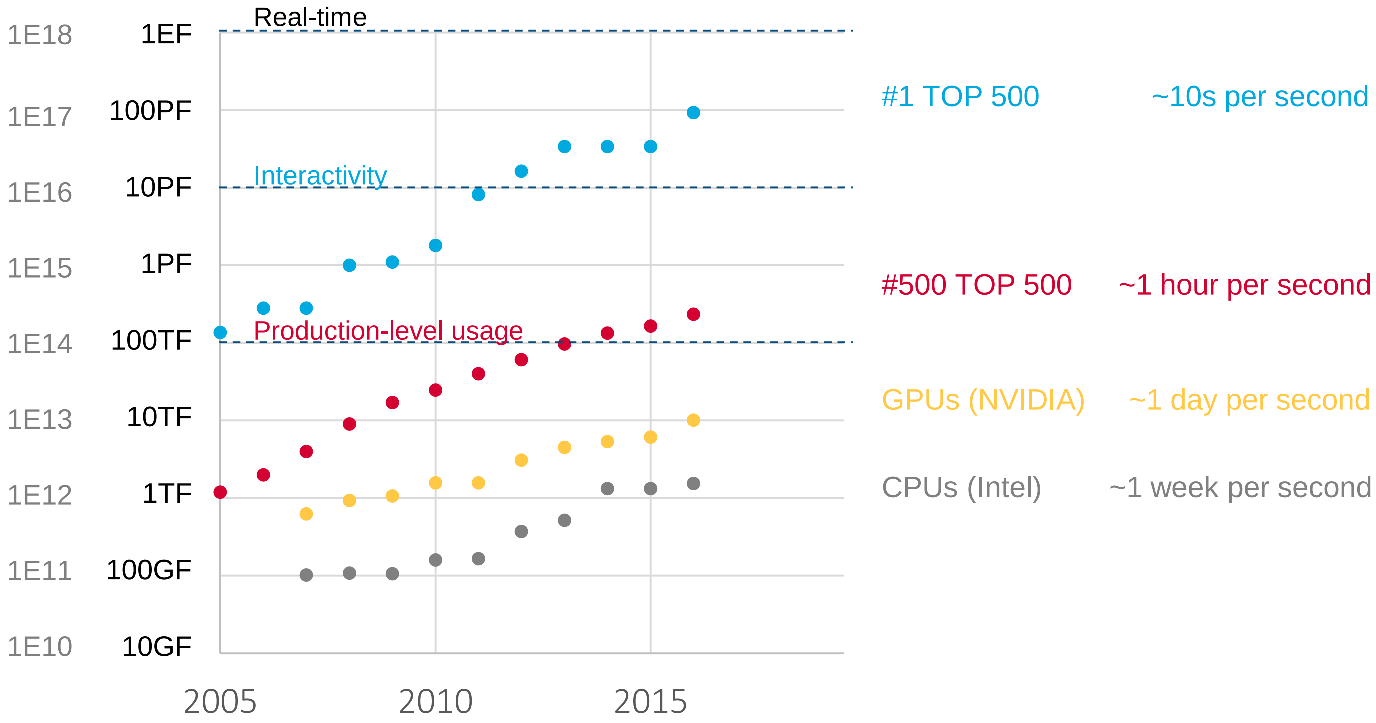

TOWARDS PRODUCTION-LEVEL USAGE

Can we reach the goal of overnight simulations without compromising on resolution?

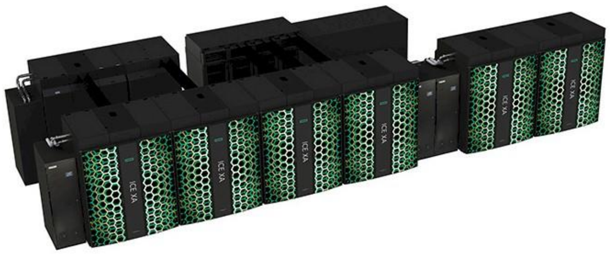

Single-node P100

Local machine with 10x P100 PCIe
(single root complex, 32 GB/s)

Multi-node P100

Up to four nodes of Tsubame 3.0 with
4x P100 SXM2 each, NVLink 1.0 and Omni-Path
(100 Gb/s)

Single-node V100

DGX-1 with 8x V100 SXM2 and
NVLink 2.0 (300 GB/s)

Massively-parallel Multi-GPU simulations for fast and accurate automotive aerodynamics
C. A. Niedermeier, C.F. Janßen and T. Indinger, ECFD 2018

ALTAIR ROADSTER CASE

- Altair roadster with grid refinement (6 refinement levels down to 2 mm and dt = 3.3E-6 s)
- Performance target for overnight simulations: 2000 MNUPS (corresponds to ~ 4 h/s)

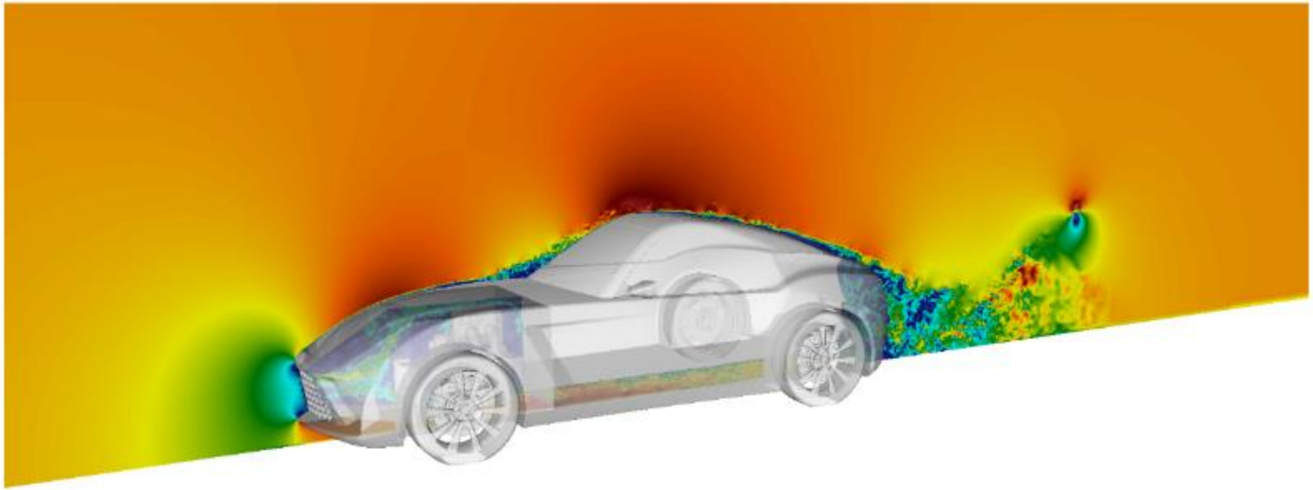

Massively-parallel Multi-GPU simulations for fast and accurate automotive aerodynamics
C. A. Niedermeier, C.F. Janßen and T. Indinger, ECFD 2018

ALTAIR ROADSTER CASE

- Altair roadster with grid refinement (6 refinement levels down to 2 mm and dt = 3.3E-6 s)
- Performance target for overnight simulations: 2000 MNUPS (corresponds to ~ 4 h/s)

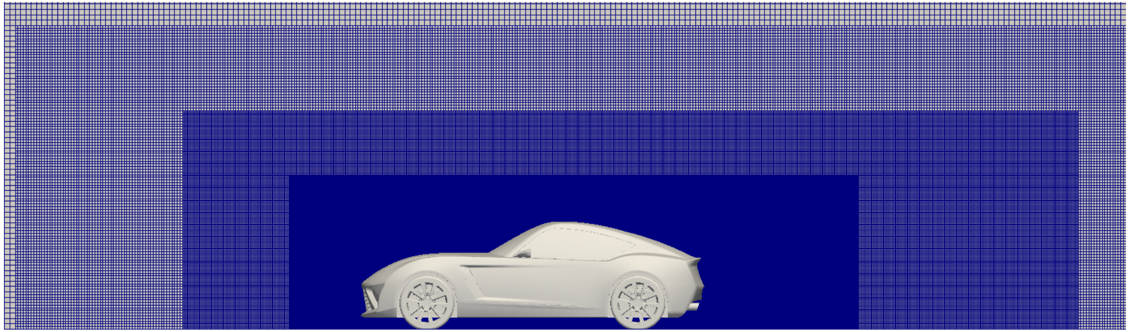

Massively-parallel Multi-GPU simulations for fast and accurate automotive aerodynamics
C. A. Niedermeier, C.F. Janßen and T. Indinger, ECFD 2018

ALTAIR ROADSTER CASE

- Altair roadster with grid refinement (6 refinement levels down to 2 mm and dt = 3.3E-6 s)
- Performance target for overnight simulations: 2000 MNUPS (corresponds to ~ 4 h/s)

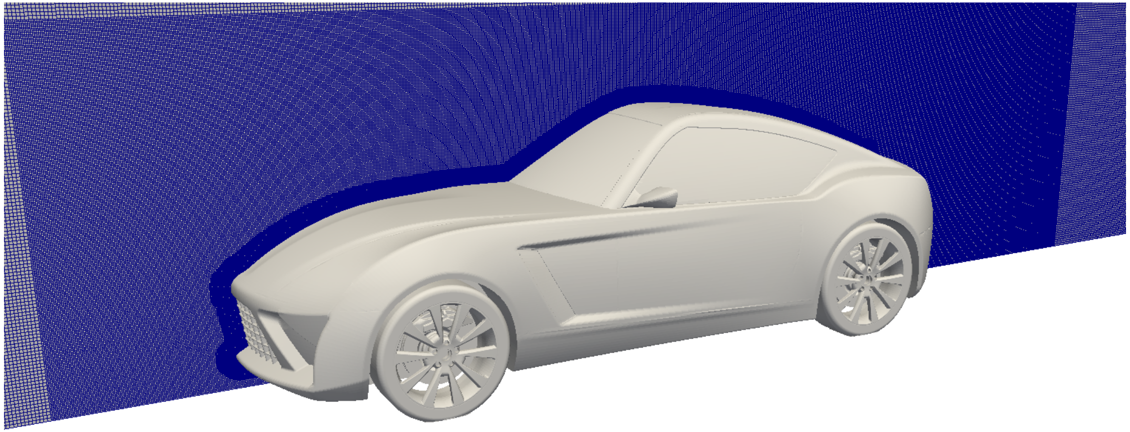

Massively-parallel Multi-GPU simulations for fast and accurate automotive aerodynamics
C. A. Niedermeier, C.F. Janßen and T. Indinger, ECFD 2018

ALTAIR ROADSTER (STRONG SCALING)

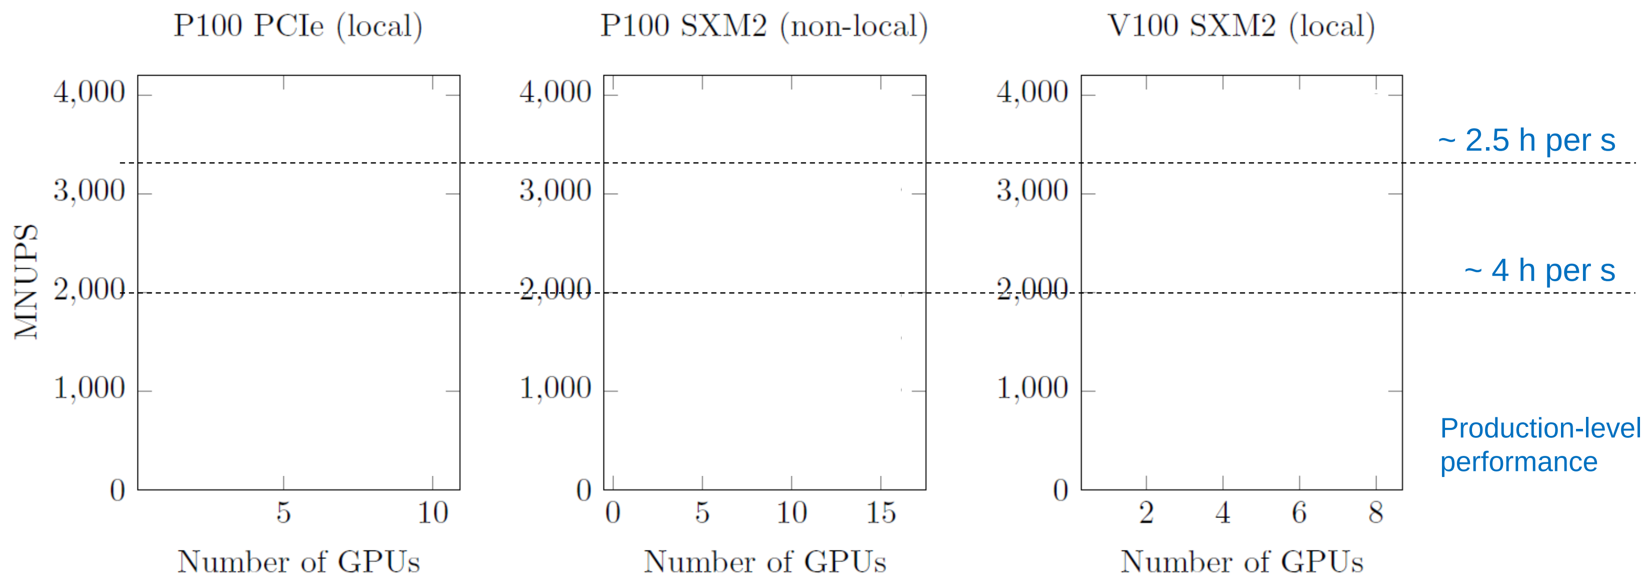

ALTAIR ROADSTER (STRONG SCALING)

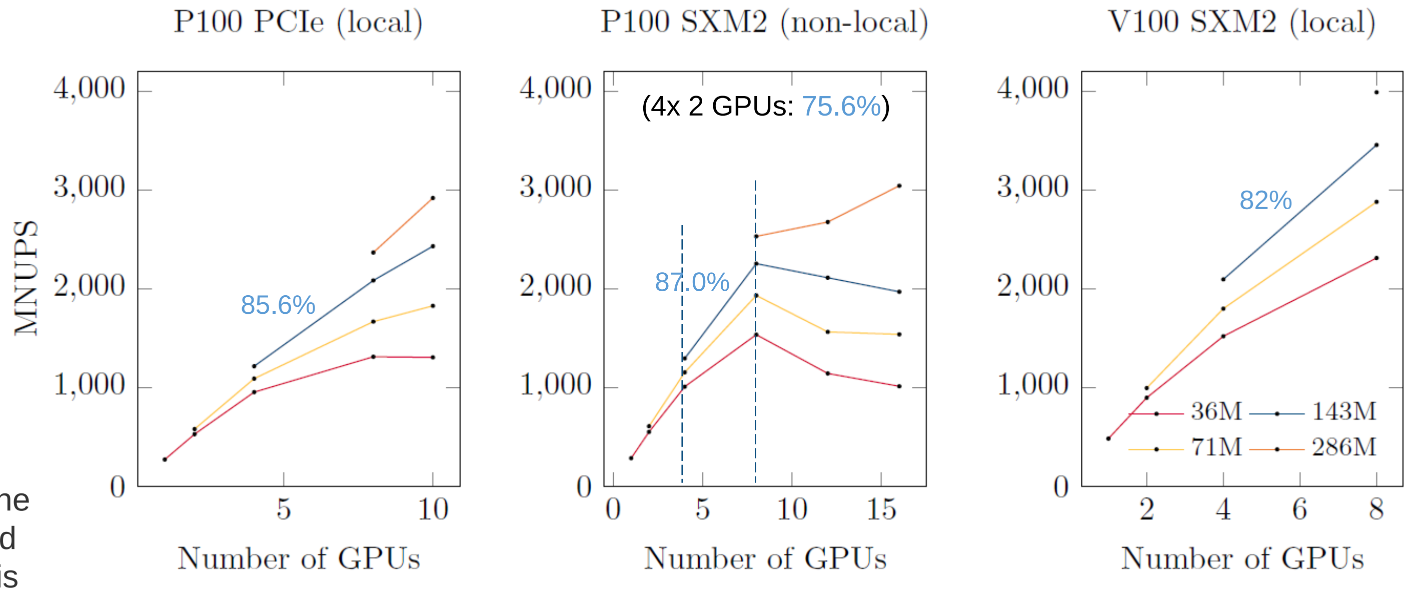

Parallel efficiency for the 143M case and 4 vs. 8 GPUs is between 82.4 and 87.0%

OUTLOOK

One DGX-1 already allows for the simulation of up to 1 s of physical time in 2.5 h – what‘s next?

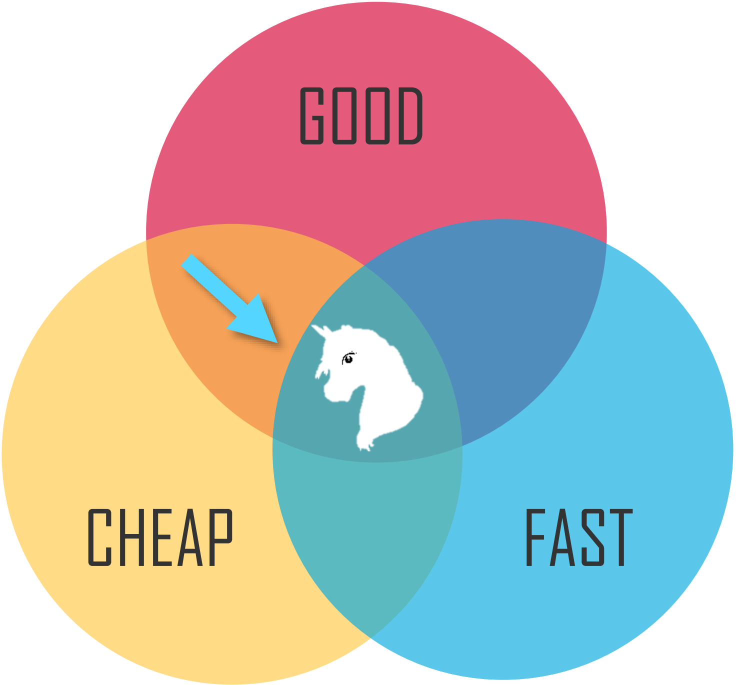

THE MORE YOU USE, …

- Example use case: detailed analysis of 10 design variations before taking decisions for next iteration

- Fully transient simulation:
  3 s physical time, 230 M voxels

THE MORE YOU USE, …

- Example use case: detailed analysis of 10 design variations before taking decisions for next iteration

- Fully transient simulation:
  3 s physical time, 230 M voxels

- Solved on ten 8x V100 machines in the cloud

THE MORE YOU USE, …

- Example use case: detailed analysis of 10 design variations before taking decisions for next iteration

- Fully transient simulation:
  3 s physical time, 230 M voxels

- Solved on ten 8x V100 machines in the cloud

- 11 h wall clock time per individual run

- Overnight results; CPU-based CFD solution would take ~ 4 days on 1,000 cores

DGX-2


NVIDIA® DGX-2™
CFD ~~AI~~ Scale on a Whole New Level
NVIDIA
NVIDIA

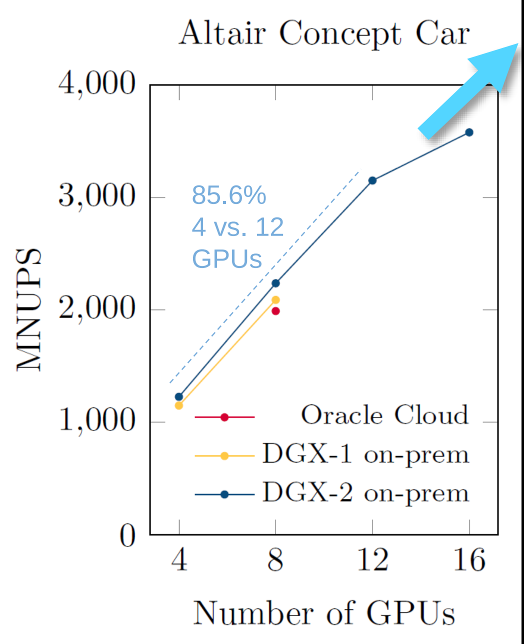
Altair Concept Car
MNUPS
4,000
3,000
2,000
1,000
0
4
8
12
16
Number of GPUs
85.6%
4 vs. 12
GPUs
Oracle Cloud
DGX-1 on-prem
DGX-2 on-prem

CONCLUSIONS AND OUTLOOK

High-fidelity simulations of external aerodynamics in real-time would require access to exa-scale machines.

Multi-GPU architectures can deliver sufficiently enough performance for overnight runs of an automotive production case.

The DGX-1 delivers enough performance to simulate a production-level case overnight – a game changer for simulation-based design.

The DGX-2 will further push the trend towards simulation-based design and at the same time allow for adding additional physics (rotating wheels, water management, aeroacoustics, …).

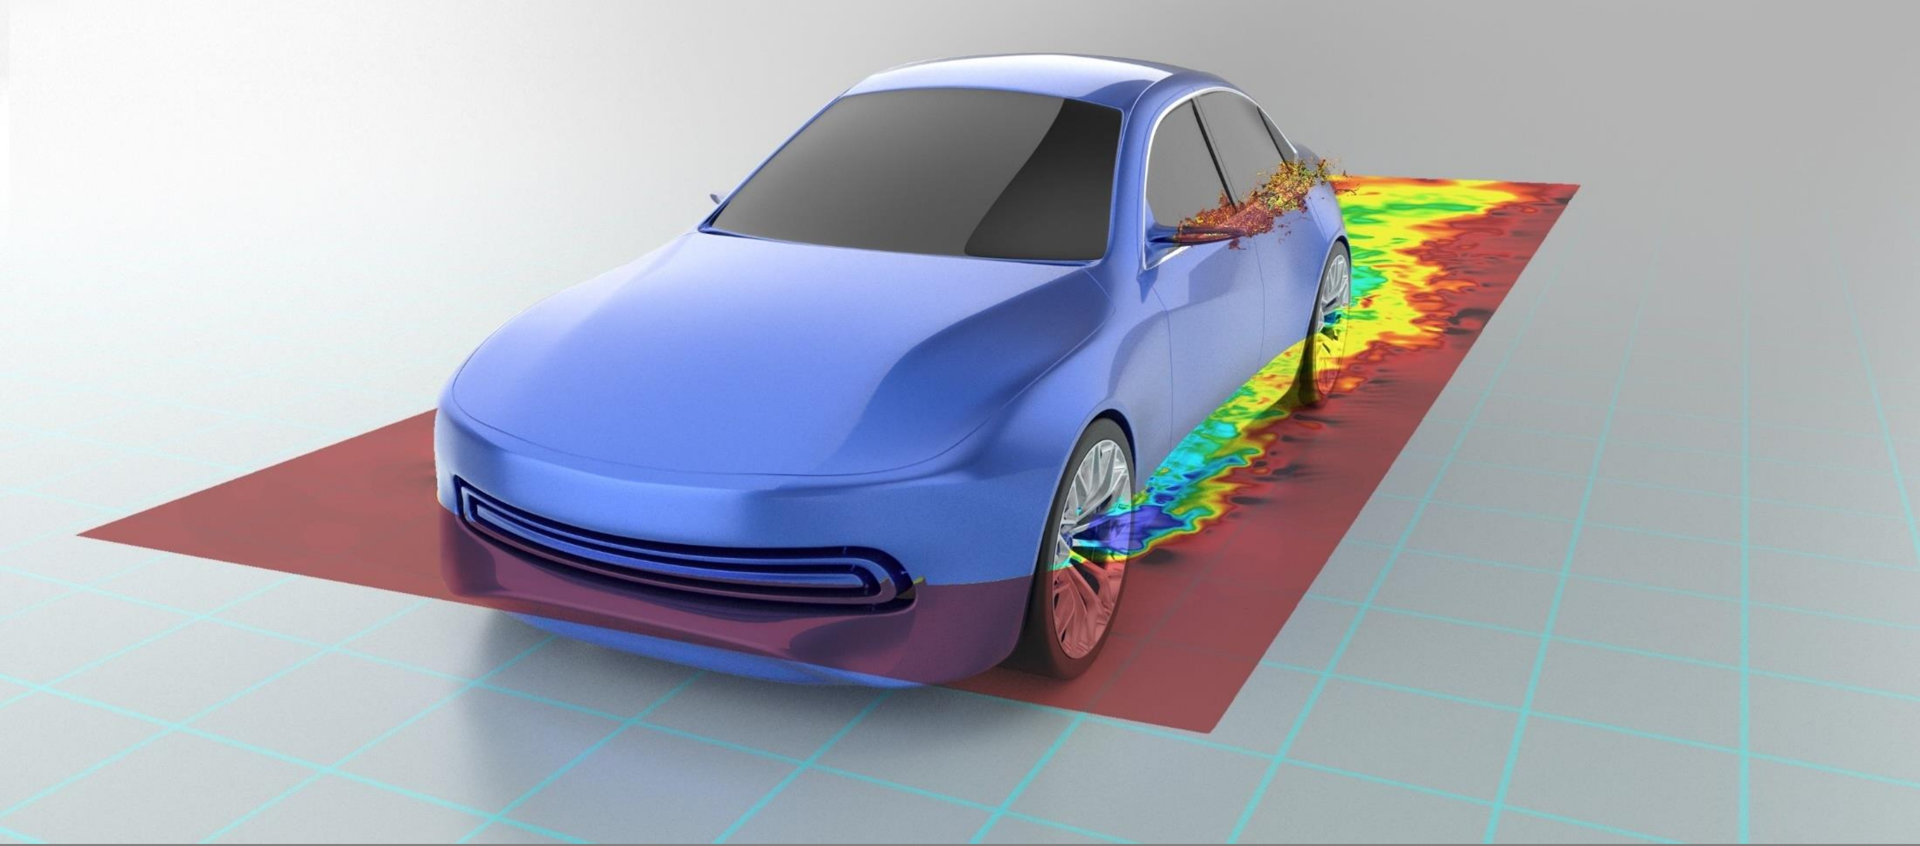

S9719 – ULTRA-FAST HIGH-FIDELITY CFD SIMULATIONS FOR AUTOMOTIVE AERODYNAMICS

Christoph Niedermeier – Development Manager

GTC San Jose, March 20, 2019

Altair