Universität Stuttgart

IARIA

IARIA's NexComm 2025

Storage for Science

The bw-SFS Project Journey

- or -

Is the performance only what matters?

Dr.-Ing. Dr. Alexey Cheptsov

What a wonderful day is today

Shortly about myself

Alexey Cheptsov

cheptsov@tik.
uni-stuttgart.de

Head of Data Center and Cloud Technology @ TIK -
IT-service provider of the University of Stuttgart

- **Founded 1829**
  - **Engineering**
  - **Nat. Science**
- **10 faculties**
- **150 institutes**
- **28.000 students**
- **270 professors**
- **3.500 scientific & research staff**

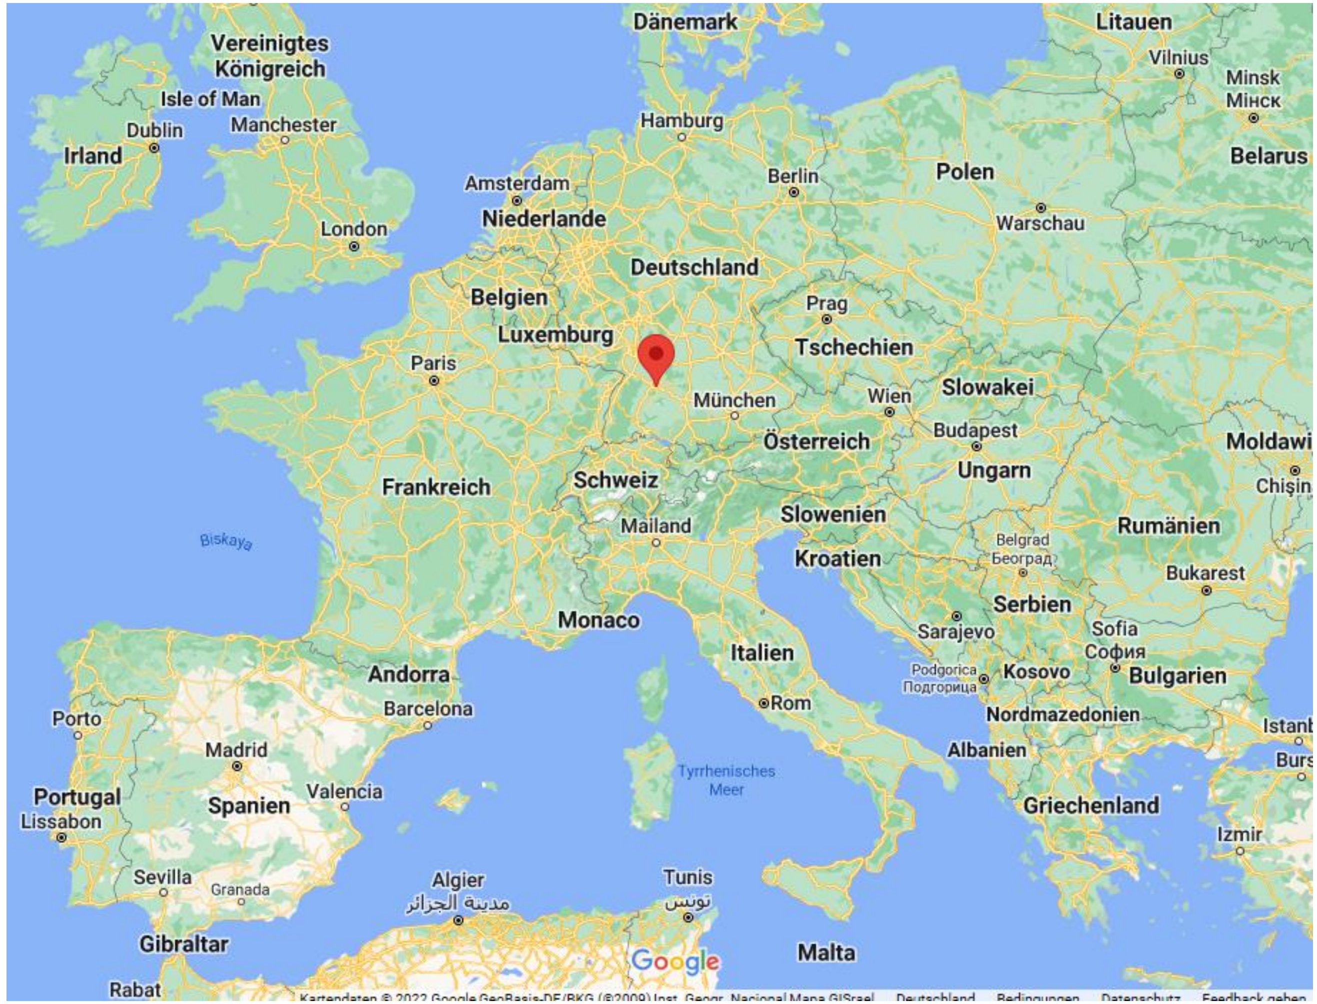

Industrial Background

KÄRCHER

Turnover: ~1-2 bil. €
Staff: ~8000

BOSCH

Turnover: ~50 bil. €
Staff: ~300.000

DAIMLER

Turnover: ~100 bil. €
Staff: 260.100

MAHLE

Driven by performance

Turnover: ~5 bil. €
Staff: ~50.000

STIHL

Turnover: ~2,5 bil. €
Staff: ~11.00

PORSCHE

Turnover: ~14 bil. €
Staff: ~19.000

OUTLINE

1. Introduction to High Performance Computing
   – What is HPC
   – How high is the high performance
   – Application examples

2. Data Science
   – Applications
   – Challenges
   – Data Center solutions

3. Conclusion

OUTLINE

1. Introduction to High Performance Computing
   - What is HPC
   - How high is the high performance
   - Application examples
2. Data Science
   - Applications
   - Challenges
   - Data Center solutions
3. Conclusion

Performance: Definition

- **Number of operations for performing an algorithm**
  - Basic microprocessor operations such as ***Addition***, ***Substraction***, ***Multiplication***, and ***Division***.

$$\text{Performance} = \frac{\text{Nr. Operations}}{\text{Time [s.]}}$$

Performance: Measurement

- **Major metric – FLOP/s**
  - Floating point (double precision – 64 bits) operations per second

196736893.993695017392

x

785.24891512582605623

= ??? [Flop/s]

What‘s about the other systems‘ performance

- **HAWK – a new-generation HPC system of HLRS**
  - HPL (High-Performance Linpack) – solver of a (random) dense system of linear equations in double precision

[Flop/s] 26 $\underbrace{000\ 000\ 000\ 000\ 000}_{\text{Peta}}$

Performance: Way to increase

- **Increase the chip capacities**

  - IBM‘s Microprocessor as of 2018

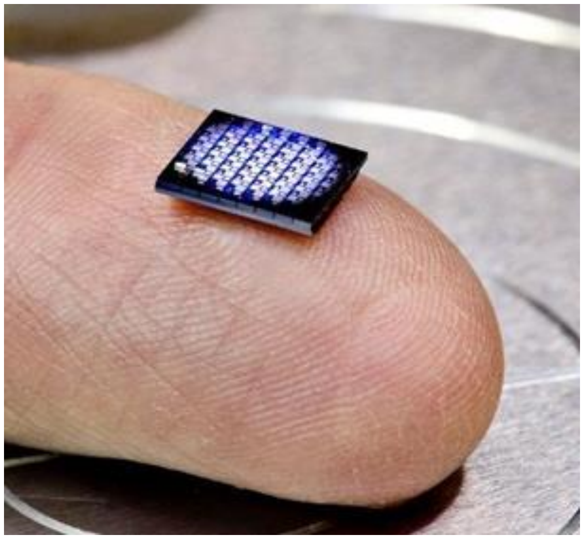

15 MFLOPs

  - Intel‘s Microprocessor Haswell

480.000 MFLOPs

Performance: Way to increase

- **Increase the chip capacities: The Moore's law**

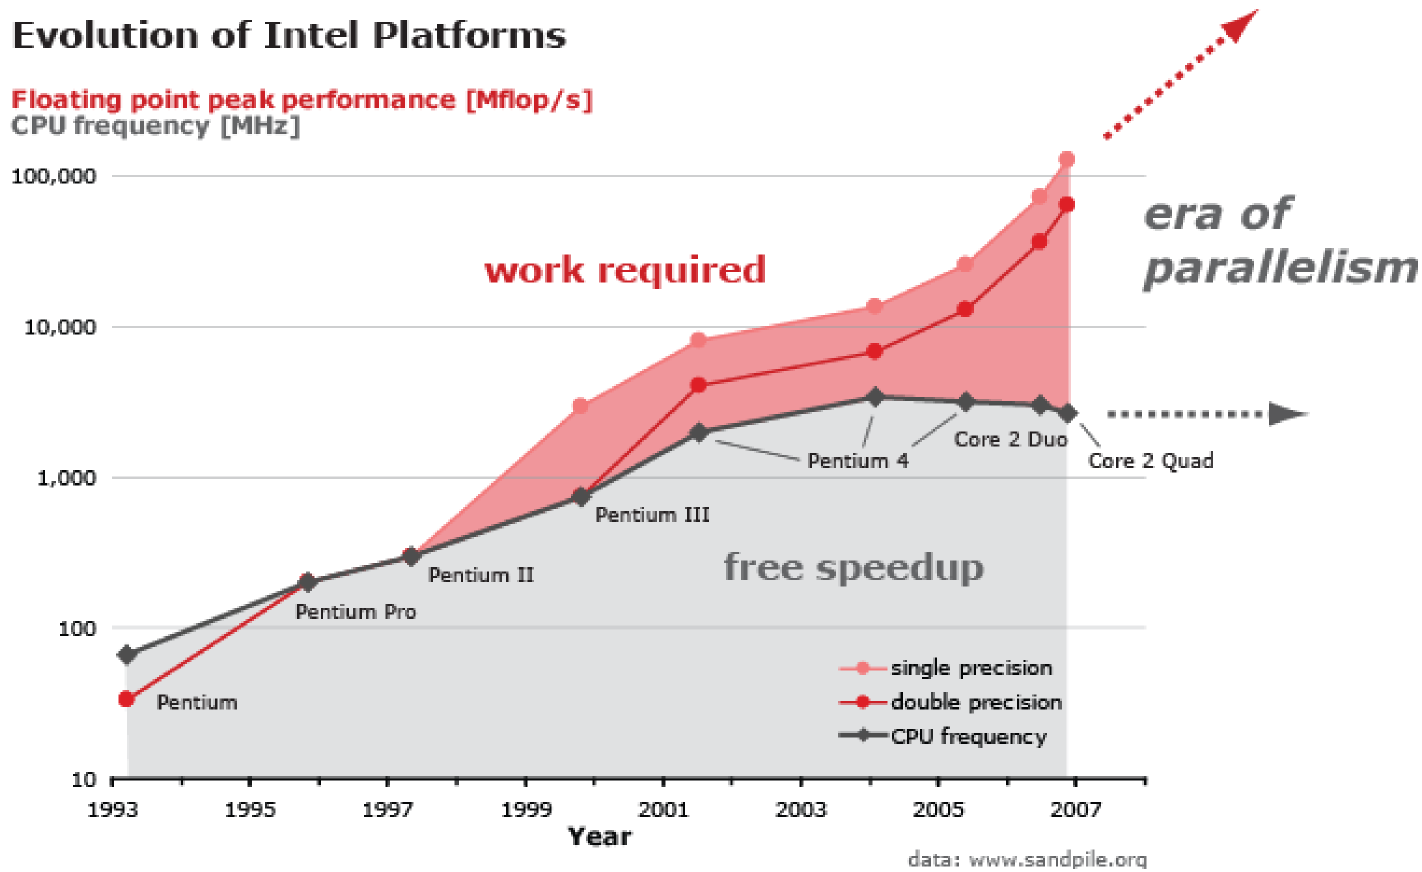

Performance: Way to increase

- **Parallelisation**

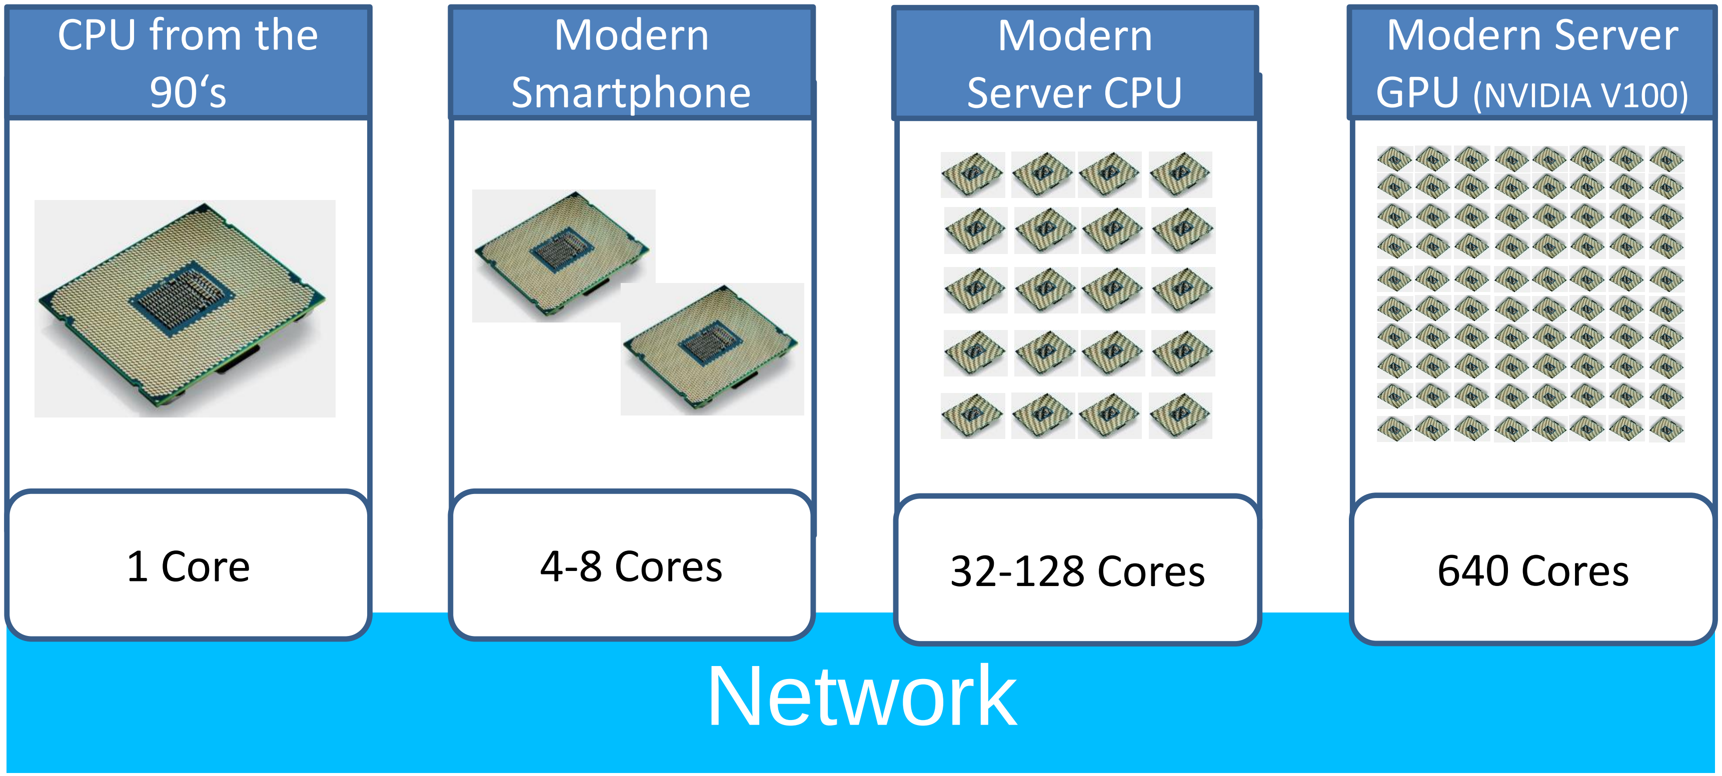

Performance: Historical Development

- **Cray-2 (Stuttgart, 1986)**
  - 4 CPUs
  - 2 GFlop/s
  - 32GB RAM
  - 2,5 Tones
    (1/3 – cooling liquid)

Performance: Historical Development

- **HLRS‘s Hawk (2021)**

  – 5.632 Nodes / 2 CPUs/ 64 Cores
  – 720.896 Cores totally
  – 7,4 PFlop/s
  – 5 MW

Another fancy HPC System

- BSC - Marenostrum

Performance: Leaders

- TOP500.org

Rank	System	Cores	Rmax (PFlop/s)	Rpeak (PFlop/s)	Power (kW)
1	**El Capitan** - HPE Cray EX255a, AMD 4th Gen EPYC 24C 1.8GHz, AMD Instinct MI300A, Slingshot-11, TOSS, HPE DOE/NNSA/LLNL United States	11,039,616	1,742.00	2,746.38	29,581
2	**Frontier** - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE Cray OS, HPE DOE/SC/Oak Ridge National Laboratory United States	9,066,176	1,353.00	2,055.72	24,607
3	**Aurora** - HPE Cray EX - Intel Exascale Compute Blade, Xeon CPU Max 9470 52C 2.4GHz, Intel Data Center GPU Max, Slingshot-11, Intel DOE/SC/Argonne National Laboratory United States	9,264,128	1,012.00	1,980.01	38,698
4	**Eagle** - Microsoft NDv5, Xeon Platinum 8480C 48C 2GHz, NVIDIA H100, NVIDIA Infiniband NDR, Microsoft Azure Microsoft Azure United States	2,073,600	561.20	846.84	
5	**HPC6** - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, RHEL 8.9, HPE Eni S.p.A. Italy	3,143,520	477.90	606.97	8,461

Evolution of Computational Science

$$\text{continuity:}\quad \frac{\partial u}{\partial x}+\frac{\partial v}{\partial y}=0$$

$$\text{x-momentum:}\quad u\frac{\partial u}{\partial x}+v\frac{\partial u}{\partial y}=-\frac{1}{\rho}\frac{\partial p}{\partial x}+\nu\left[\frac{\partial^2 u}{\partial x^2}+\frac{\partial^2 u}{\partial y^2}\right]$$

$$\text{y-momentum:}\quad u\frac{\partial v}{\partial x}+v\frac{\partial v}{\partial y}=-\frac{1}{\rho}\frac{\partial p}{\partial y}+\nu\left[\frac{\partial^2 v}{\partial x^2}+\frac{\partial^2 v}{\partial y^2}\right]$$

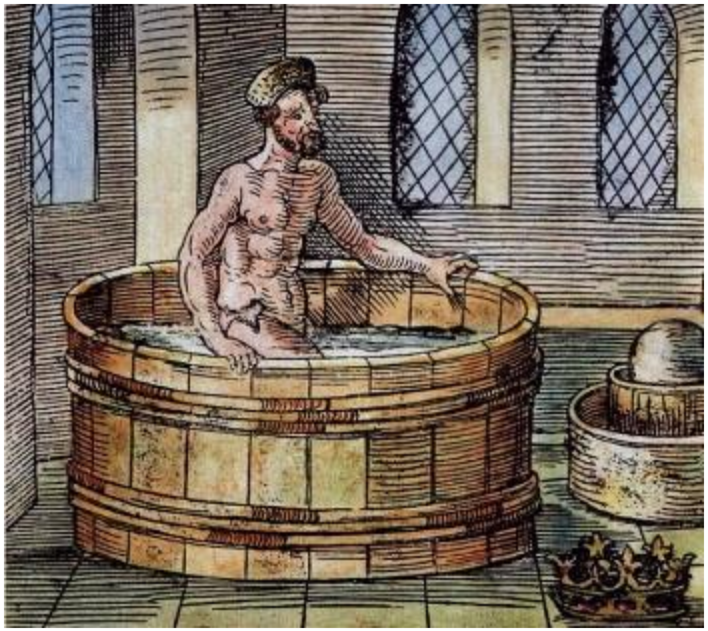

```
f = figure;
set(f,'Color', [0 0 0]);
sphere(200);
axis vis3d off;
h = findobj('Type','surface');
shading interp;
set(h,'FaceColor',[0.5 0.5 0.5]);
light('Position',[-3 -1 3]);
set(h,'DiffuseStrength',1.0);
set(h,'SpecularStrength',1);
set(h,'SpecularExponent',1);
set(h,'AmbientStrength',0.25);
set(h,'BackFaceLighting','lit')
```

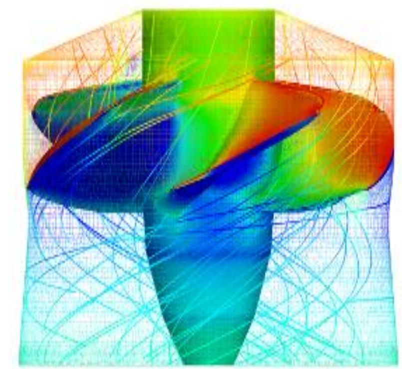

experimental	theory	computation	simulation
-1000	**-100**	**-10**	**2010** Year, A.D.

Simulation: The Way

- Optimization
- Simulation Software
- Mathematical Model
- Abstract Model
- Simulation object

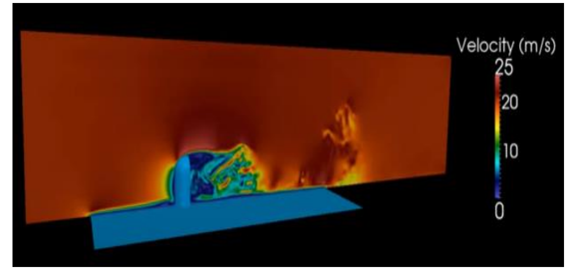

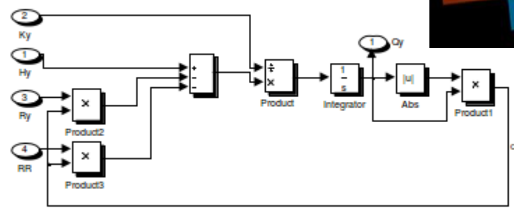

$$\rho\dot{\vec{v}} = \rho\left(\frac{\partial\vec{v}}{\partial t} + (\vec{v}\cdot\nabla)\vec{v}\right) = -\nabla p + \mu\Delta\vec{v} + \left(\zeta + \frac{\mu}{3}\right)\nabla(\nabla\cdot\vec{v}) + \vec{f}.$$

```
defineDynamics(lattice, lattice.getBoundingBox(),
               new CylinderShapeDomain3D<T>(cx,cy,radiu
               new BounceBack<T,DESCRIPTOR>);


createPoiseuilleBoundaries(lattice, parameters, *bounda
lattice.initialize();

// Main loop over time iterations.
for
    // Lattice Boltzmann iteration step.
    lattice.collideAndStream();
}

delete boundaryCondition;
```

Application Domains: Mathematical Background

- CFD – Computational Fluid Dynamics
- MD – Molecular Dynamics

$$\frac{\partial \overline{\upsilon}}{\partial t} + \overline{\upsilon} \cdot (\nabla \overline{\upsilon}) = \overline{F} - \frac{1}{\rho} \nabla p + \nu \Delta \overline{\upsilon}$$

$$\nabla \cdot \overline{\upsilon} = 0$$

$$\rho S \cdot (1 + \xi) \cdot \left(\upsilon^2 \big|_{S_2} - \upsilon^2 \big|_{S_1} \right) = S \cdot \left(p \big|_{S_1} - p \big|_{S_{21}} \right) +$$

$$p \Phi dy - pS \frac{d\upsilon}{dt} dy - \frac{\rho \lambda S}{r} \upsilon^2 dy$$

Applications: The XXL Portfolio

- Climate simulation on a high (a few kms) resolution
  - University of Hohenheim
  - 84.000 cores, 84 hours, 450 TB data

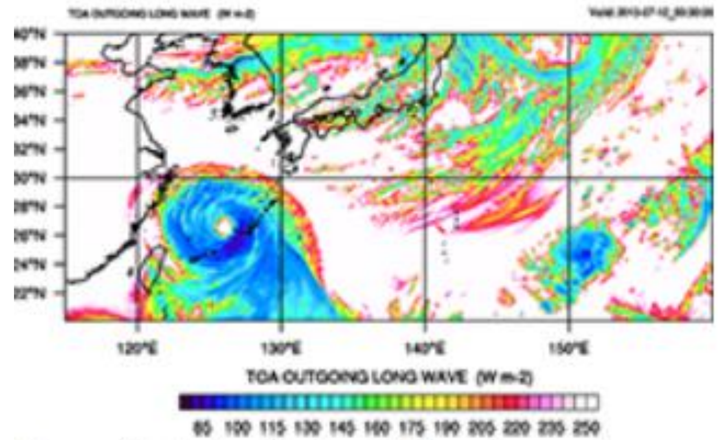

Copyright: *Institute of Physics and Meteorology, Universität Hohenheim*

- Turbulent flow simulation for air- and gas dynamics analysis
  - RWTH Aachen University
  - 92.000 cores, 110 hours, 80 TB data

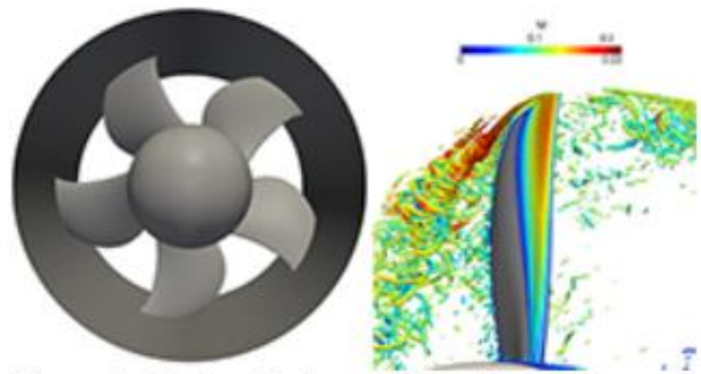

Copyright: *Institute of Aerodynamics, RWTH Aachen University*

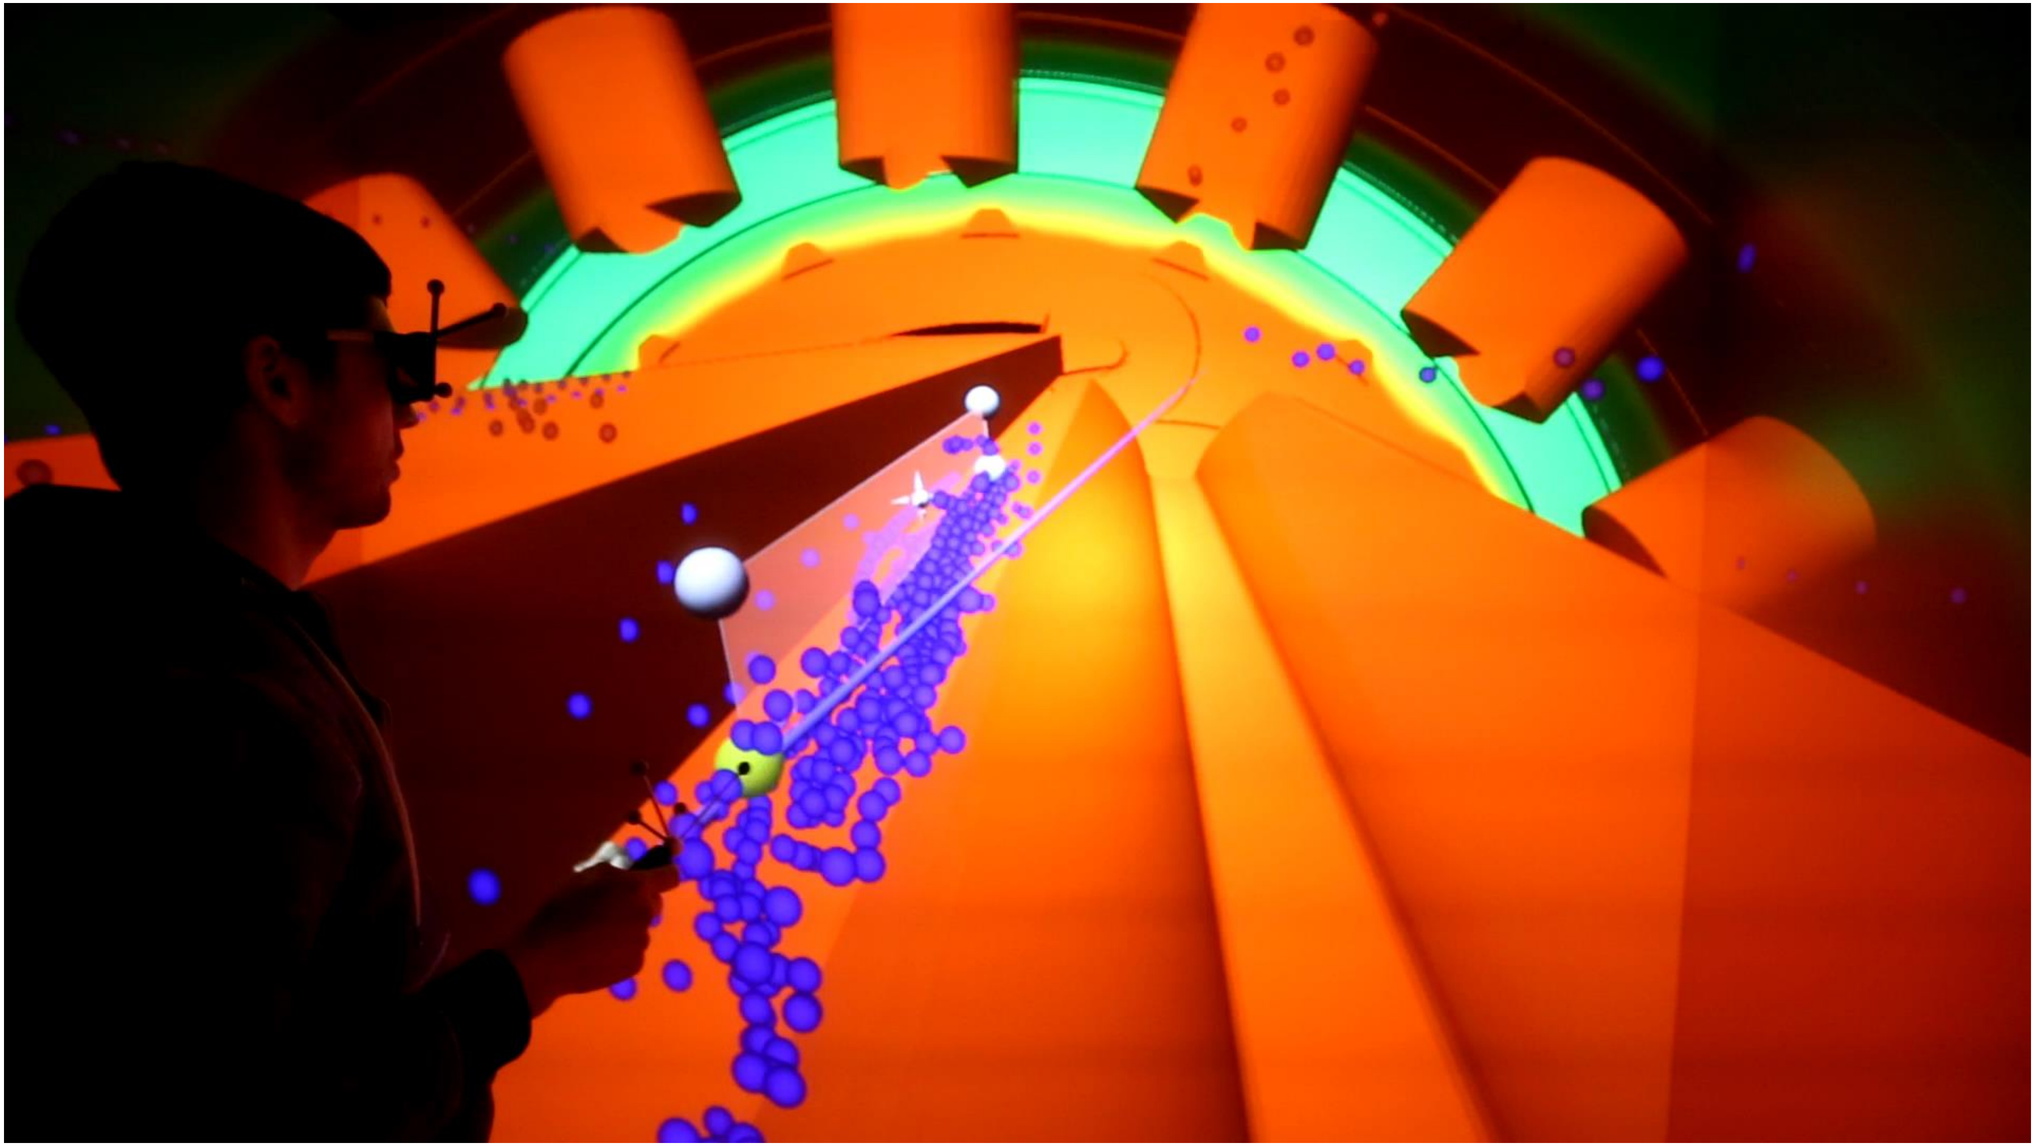

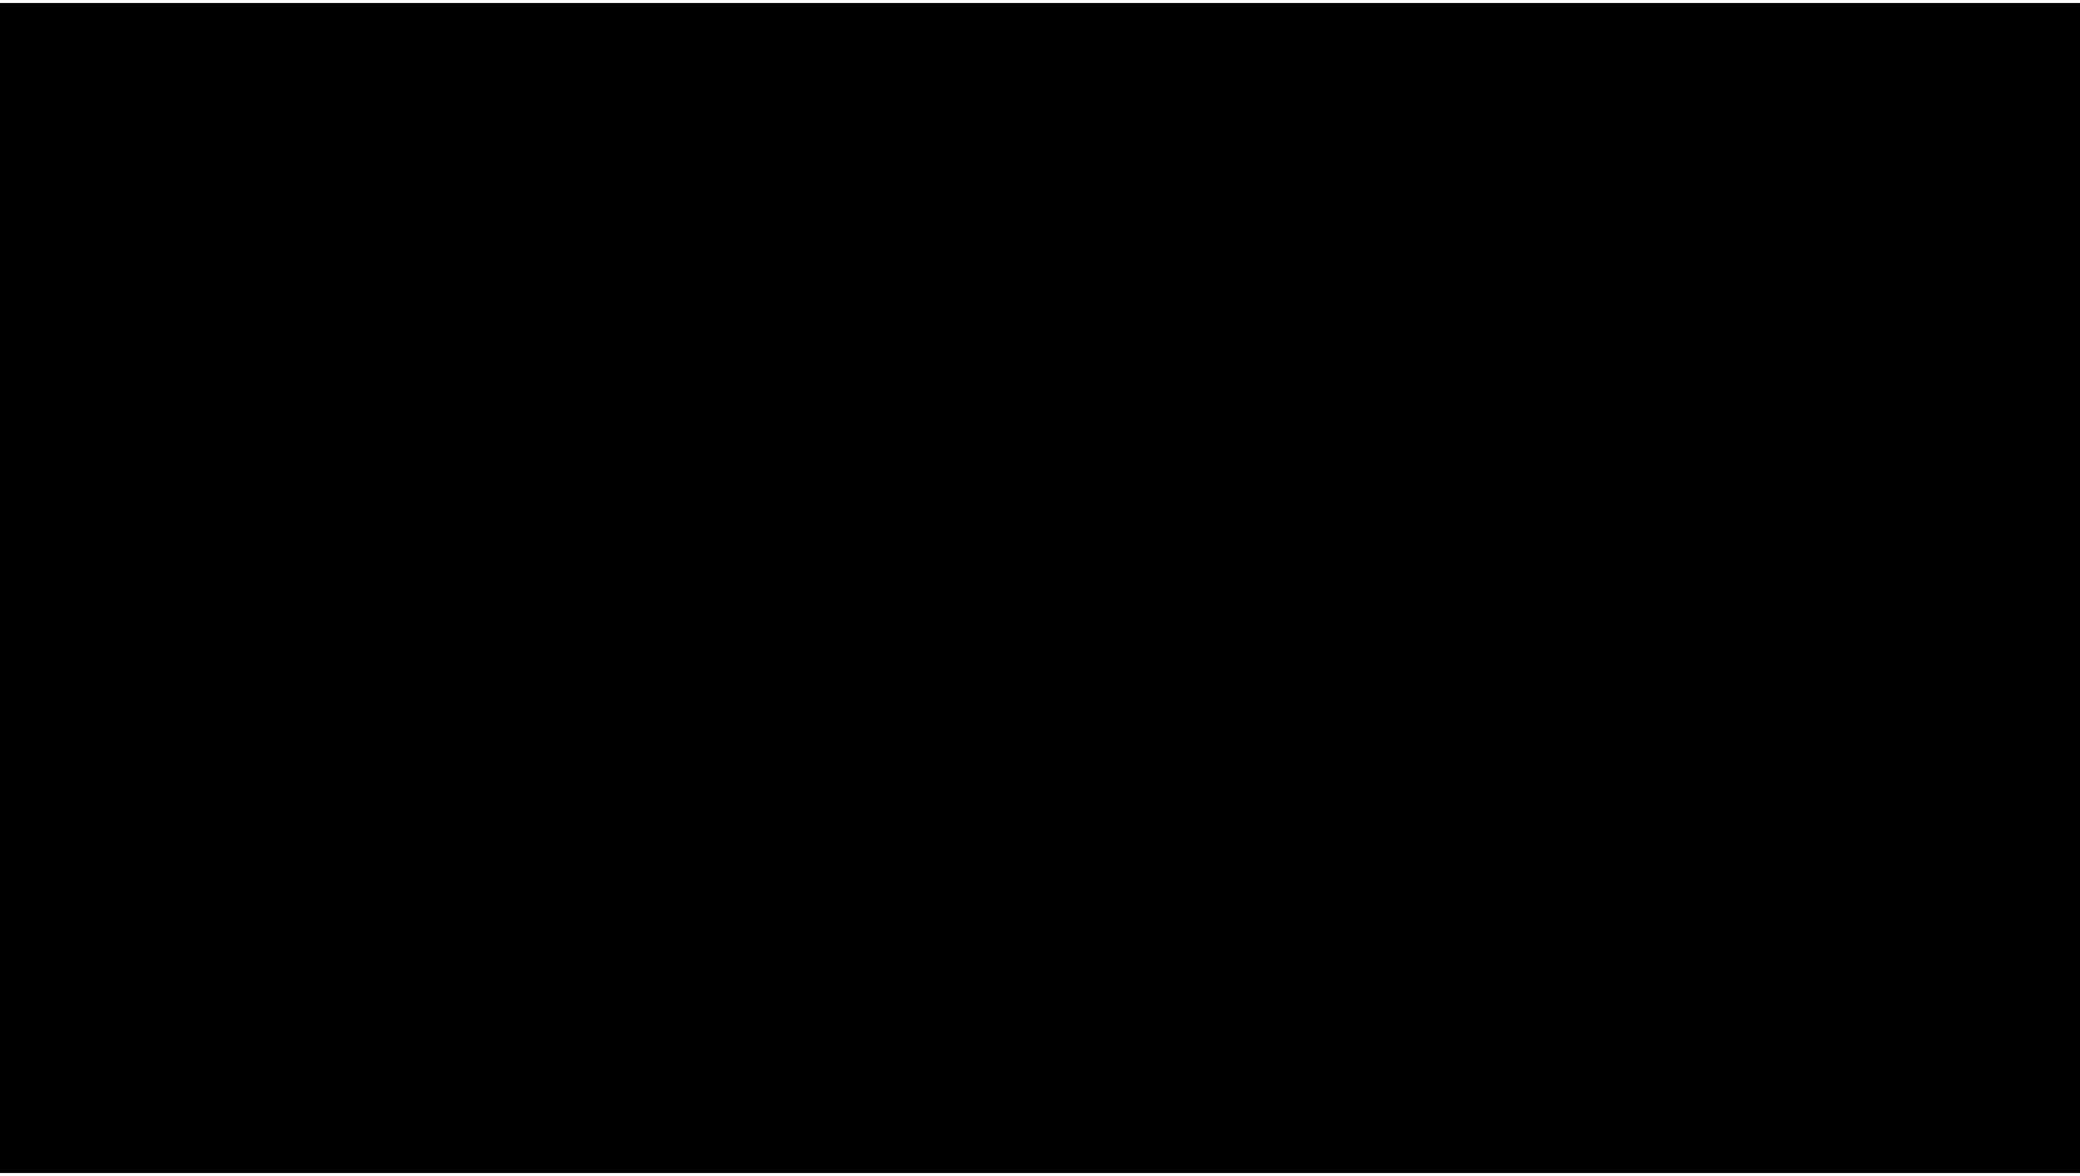

Further Use Cases

- Vehicle aerodynamics

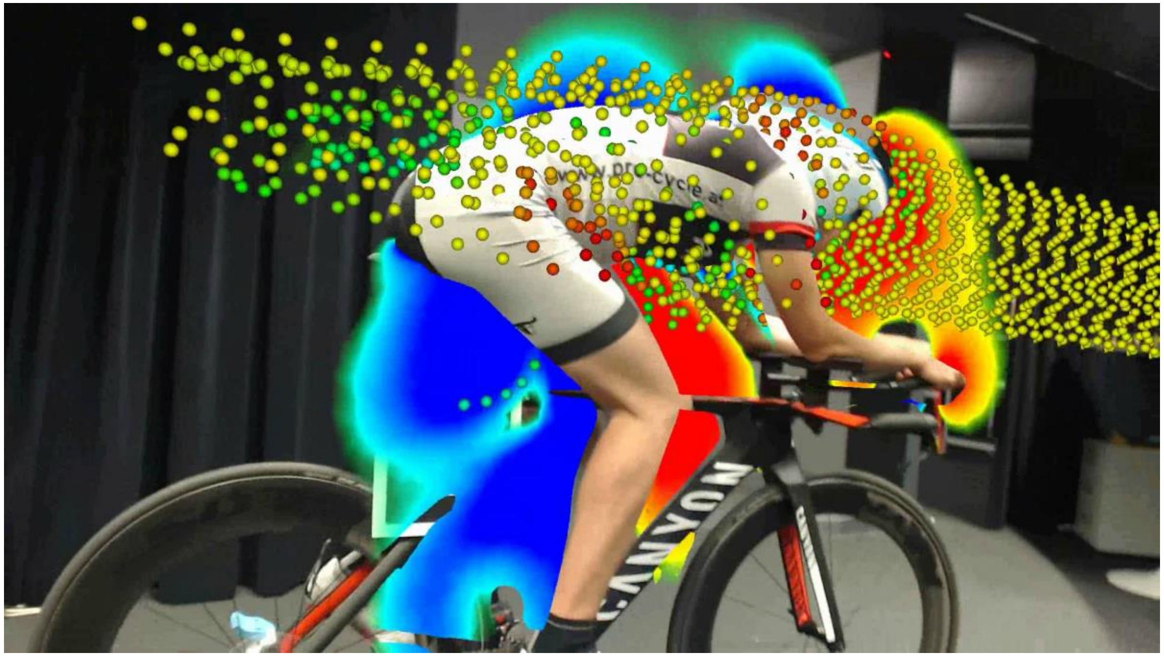

Further Use Cases

- **Vehicle aerodynamics**

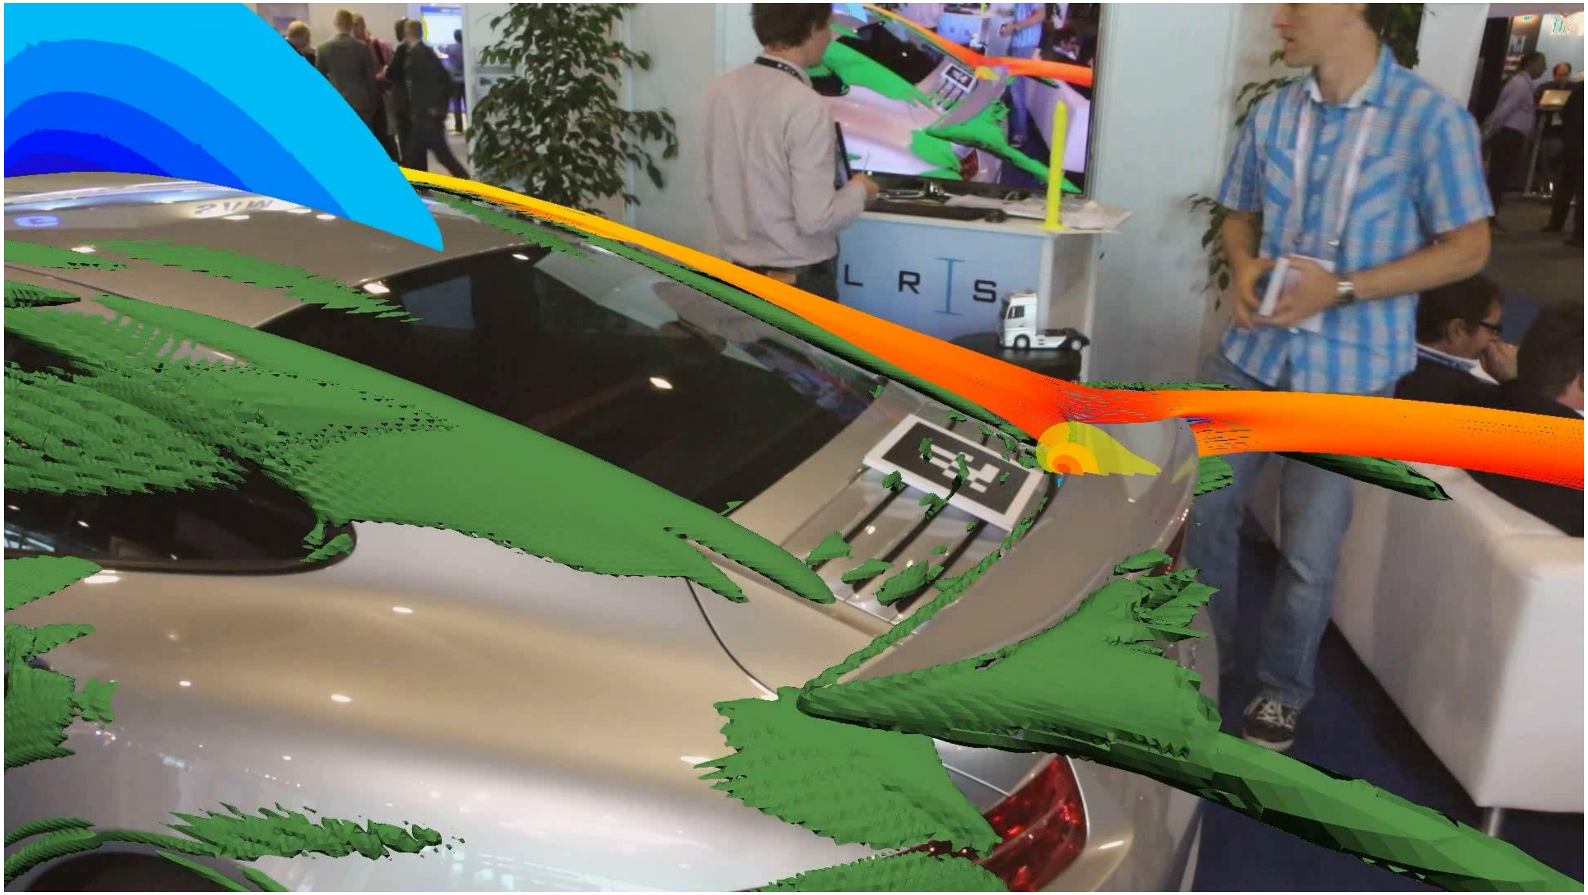

Further Use Cases

- Crash tests

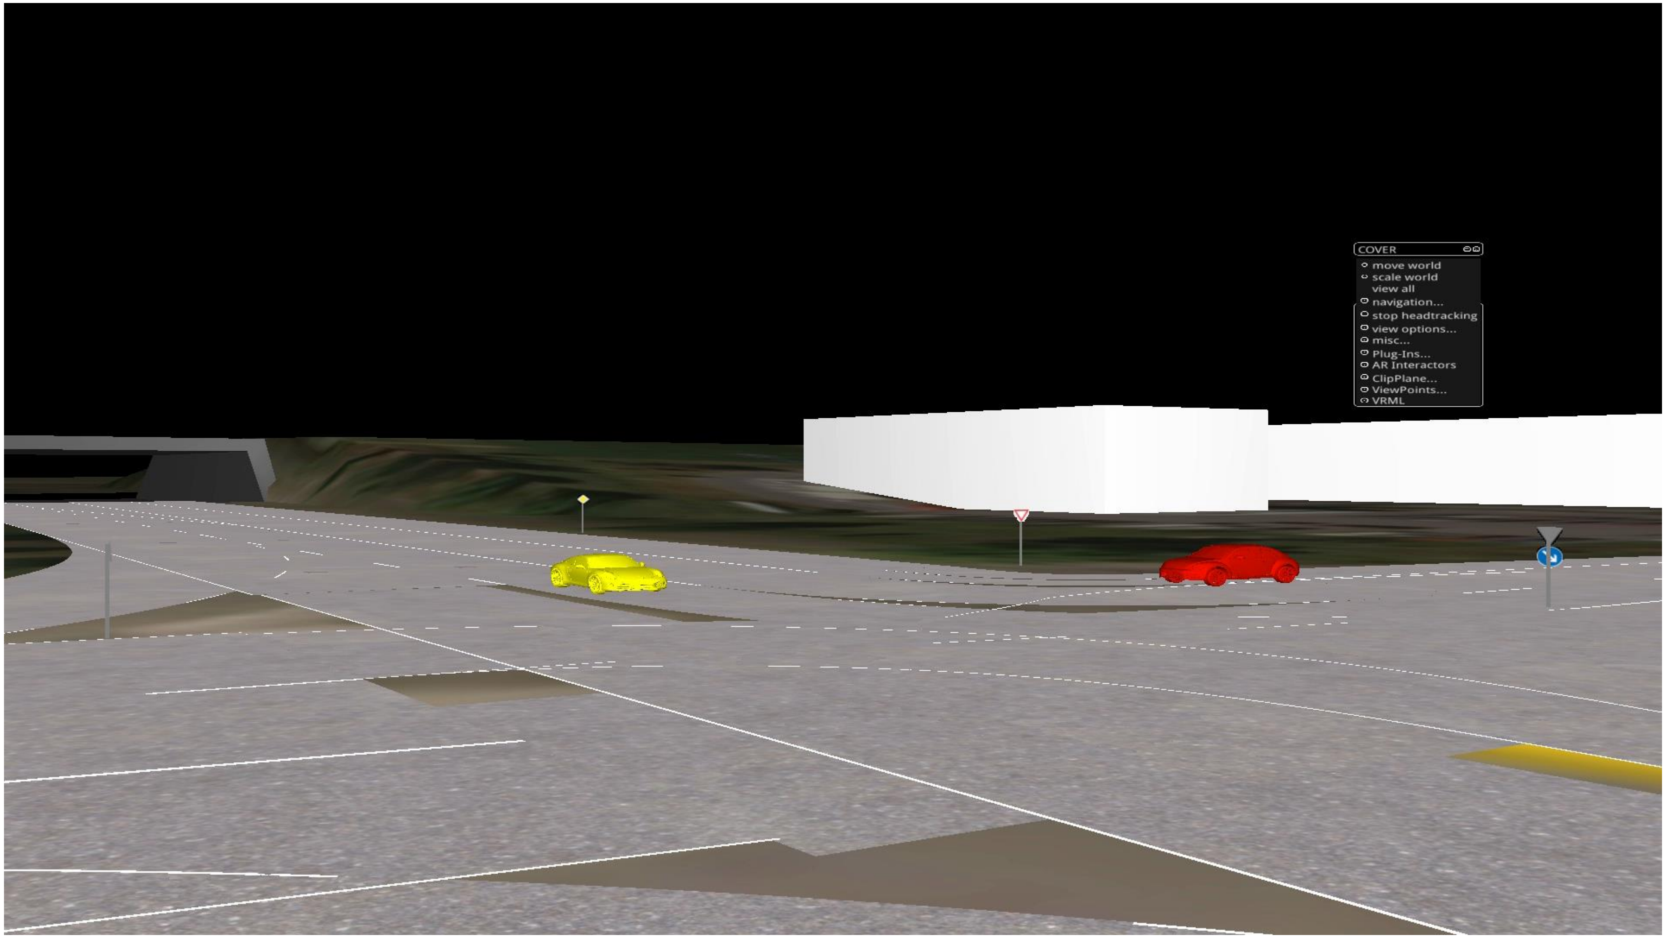

Further Use Cases

- Crash tests

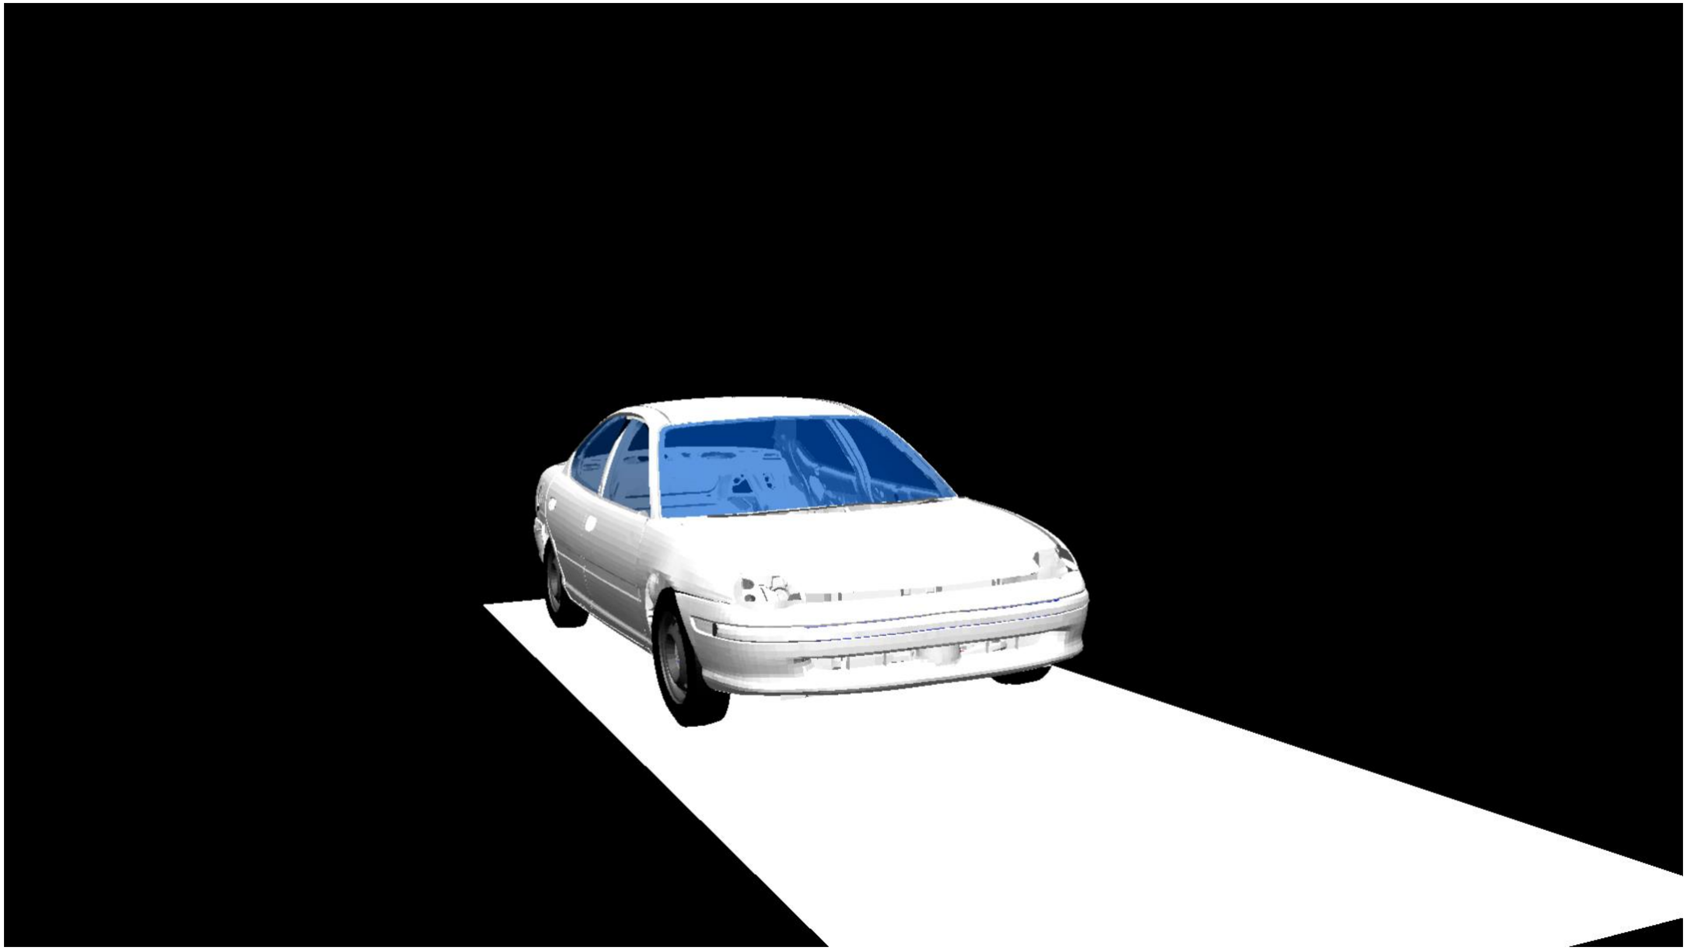

Further Use Cases

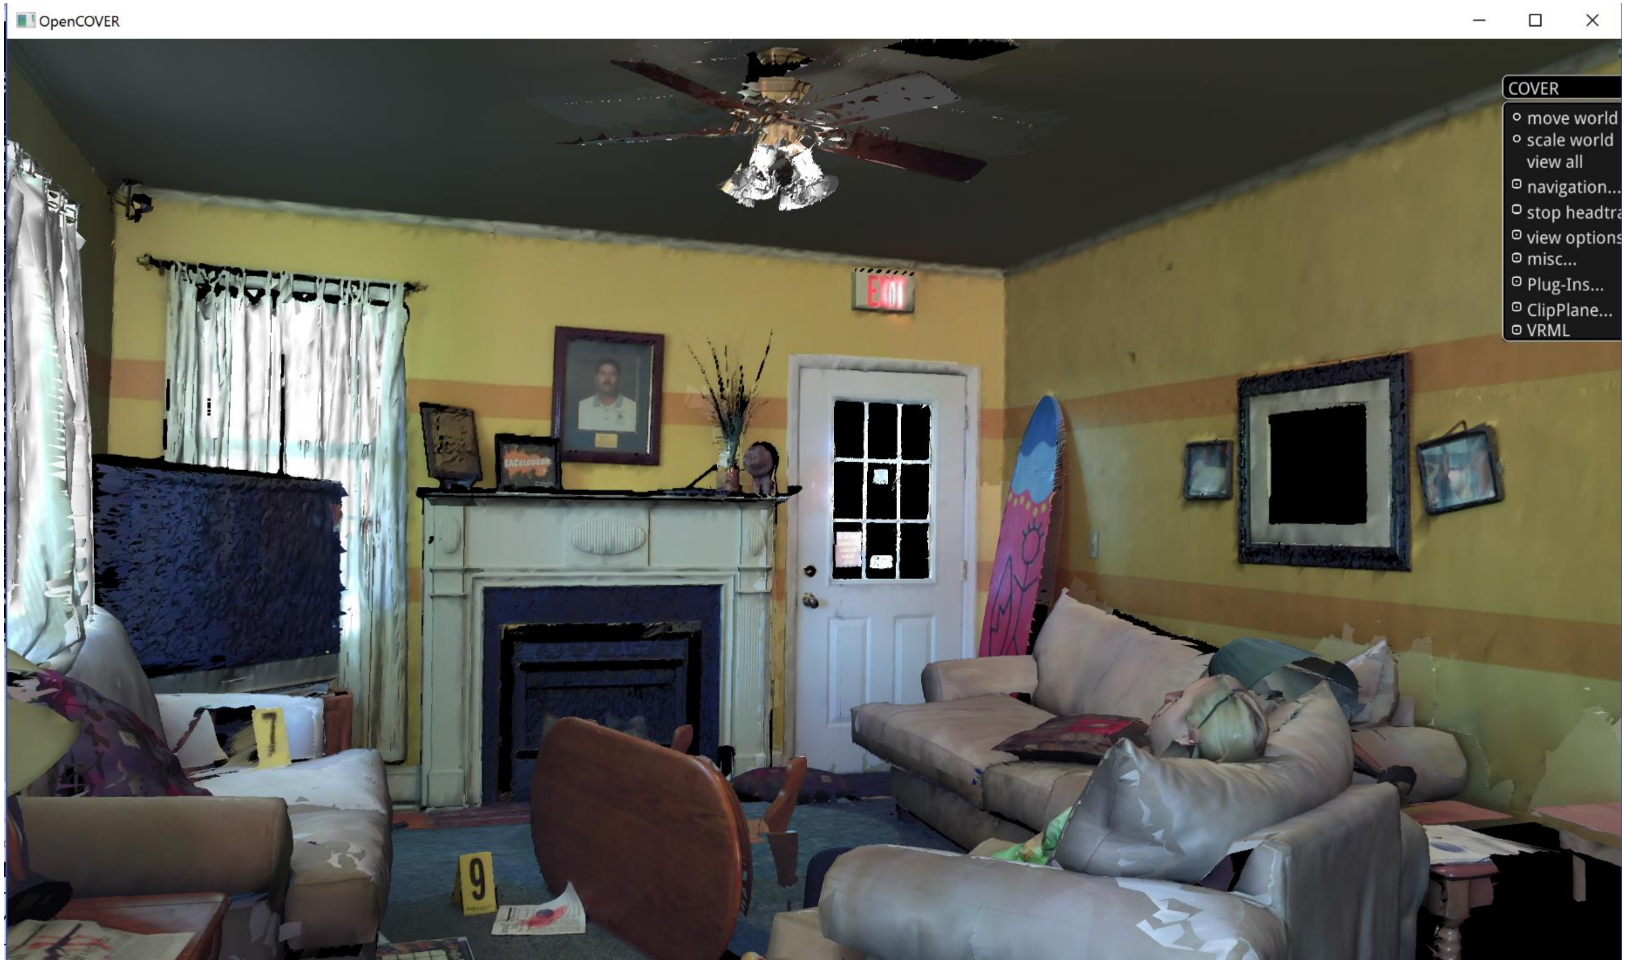

Further Use Cases

- **Crime Scene Reproduction**

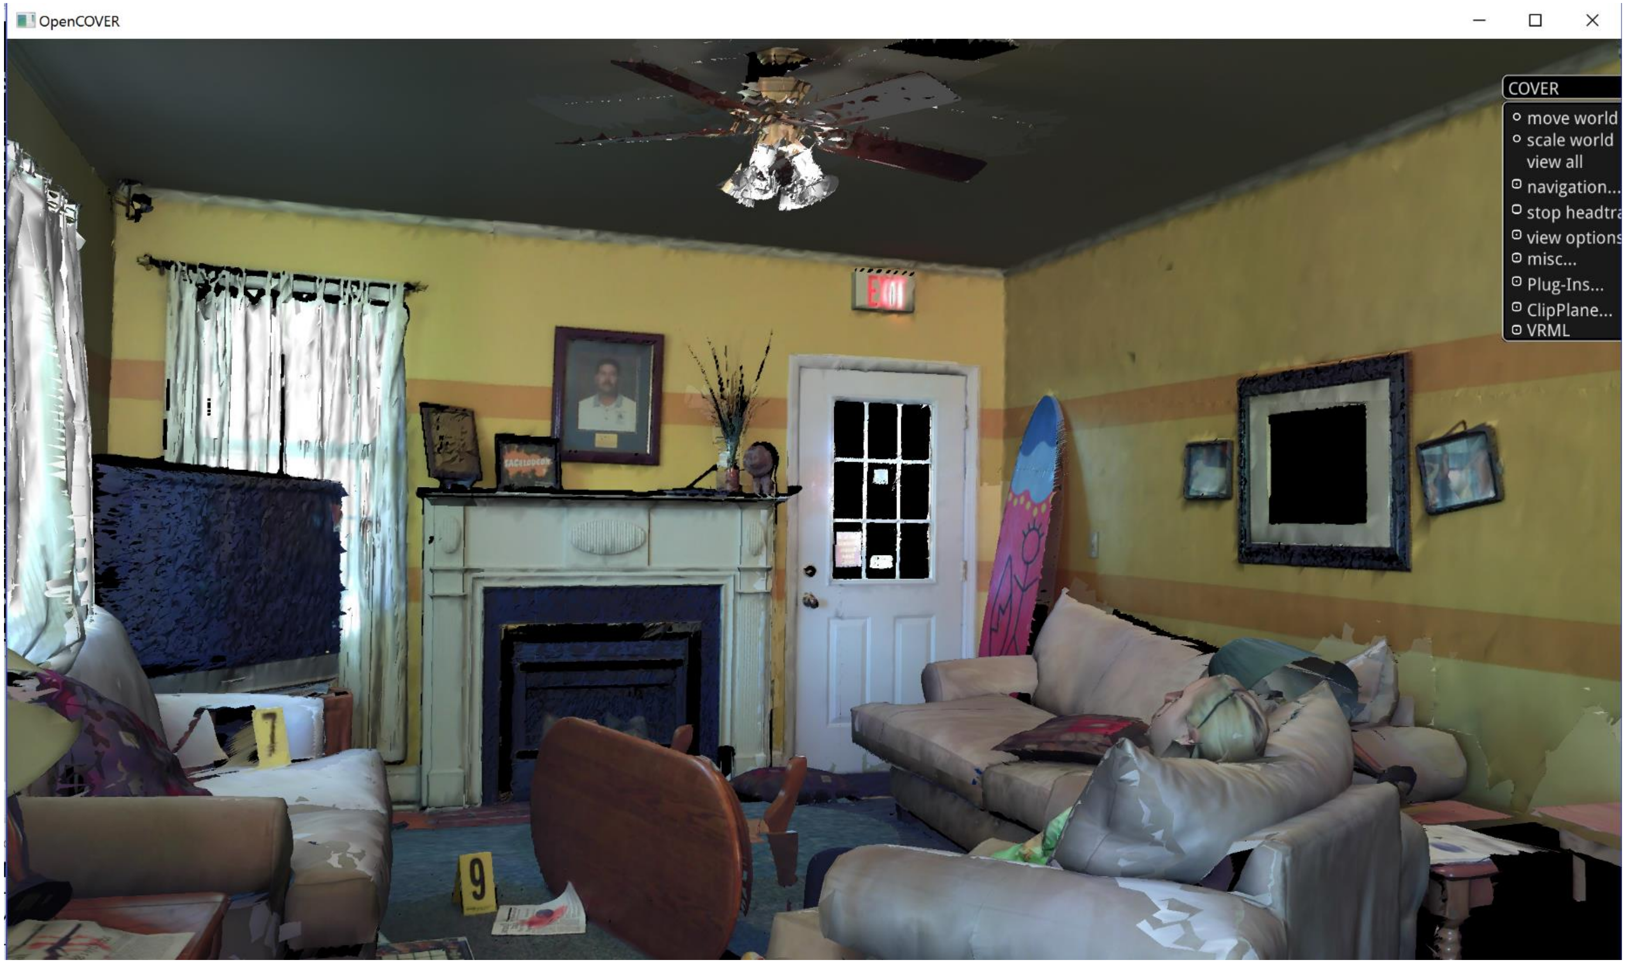

Further Use Cases

- Crime Scene Reproduction

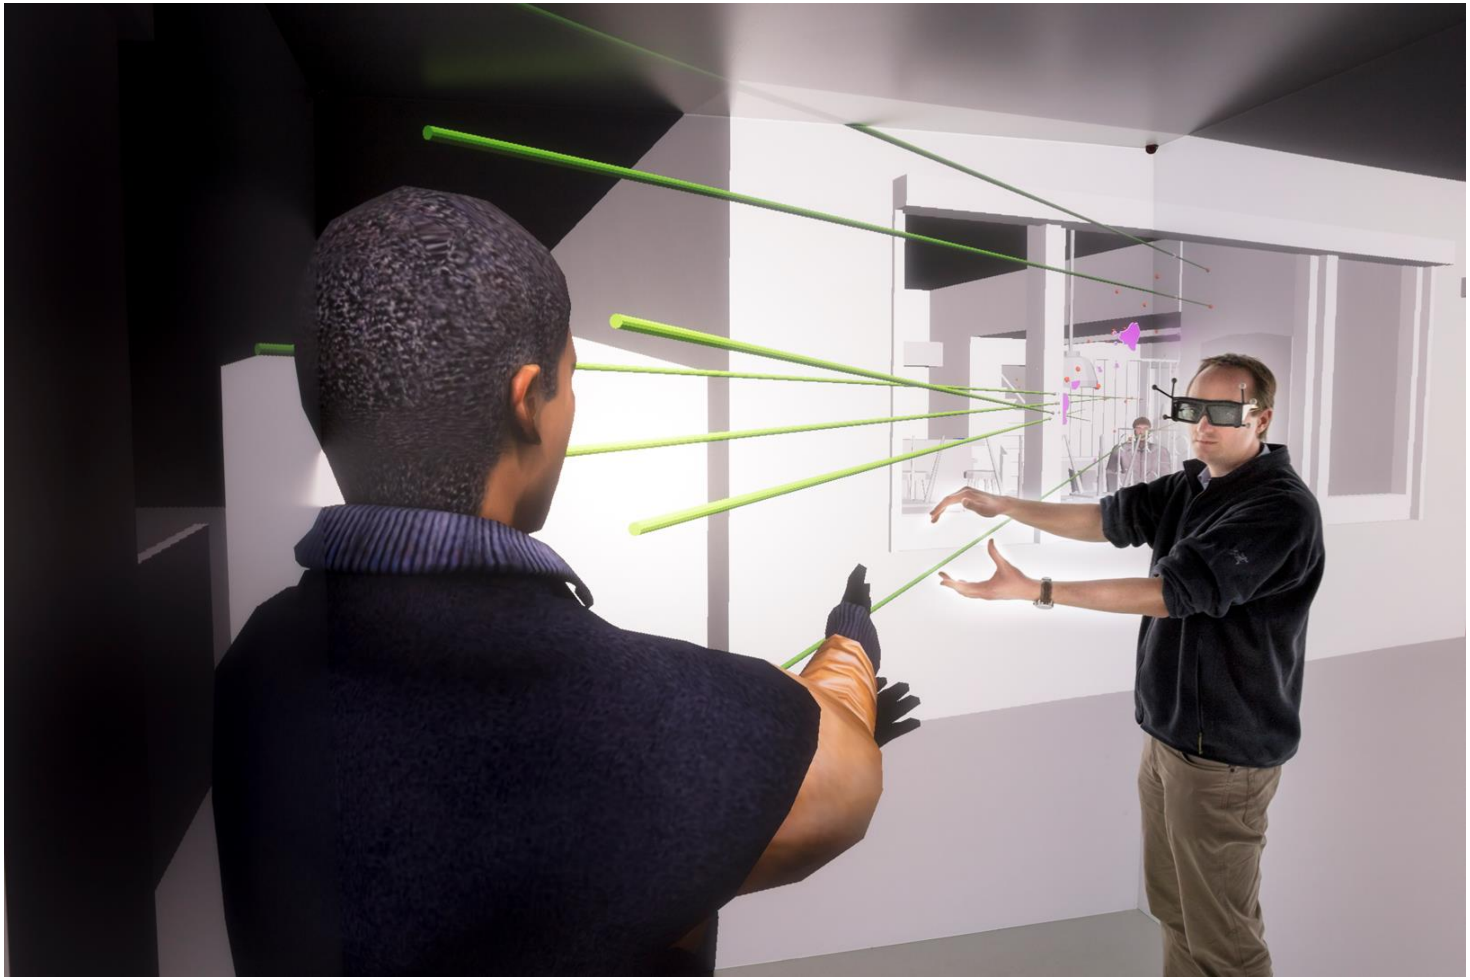

OUTLINE

1. Introduction to High Performance Computing
   - What is HPC
   - How high is the high performance
   - Application examples
2. Data Science
   - Applications
   - Challenges
   - Cloud and Data Center technologies
3. Conclusion

Further Evolution of Computational Science

$$\text{continuity: } \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} = 0$$

$$\text{x-momentum: } u\frac{\partial u}{\partial x} + v\frac{\partial u}{\partial y} = -\frac{1}{\rho}\frac{\partial p}{\partial x} + \nu\left[\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2}\right]$$

$$\text{y-momentum: } u\frac{\partial v}{\partial x} + v\frac{\partial v}{\partial y} = -\frac{1}{\rho}\frac{\partial p}{\partial y} + \nu\left[\frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial y^2}\right]$$

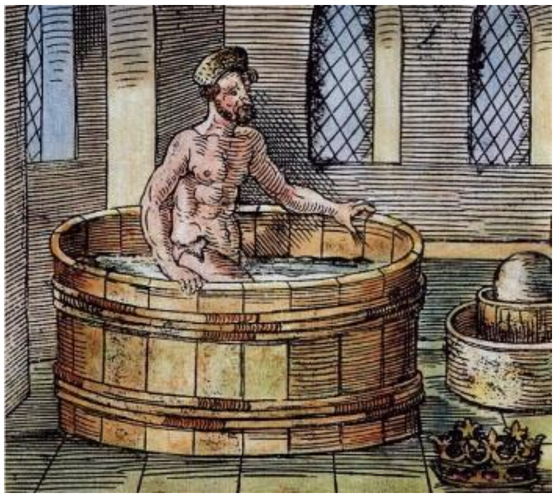

```
f = figure;
set(f,'Color', [0 0 0]);
sphere(200);
axis vis3d off;
h = findobj('Type','surface');
shading interp;
set(h,'FaceColor',[0.5 0.5 0.5]);
light('Position',[-3 -1 3]);
set(h,'DiffuseStrength',1.0);
set(h,'SpecularStrength',1);
set(h,'SpecularExponent',1);
set(h,'AmbientStrength',0.25);
set(h,'BackFaceLighting','lit')
```

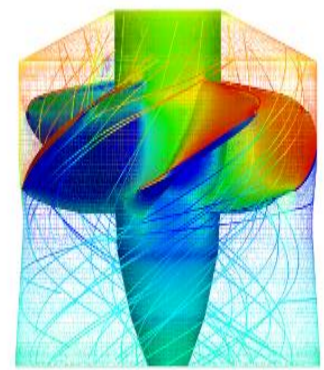

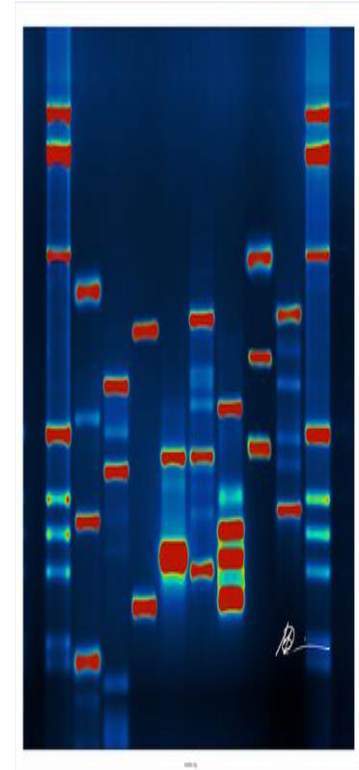

experimental | theory | computation | simulation | data-driven

-1000 | -100 | -10 | 2020 | Year, A.D.

Digital Assets

- Data Volume in Zettabytes (Gartner)

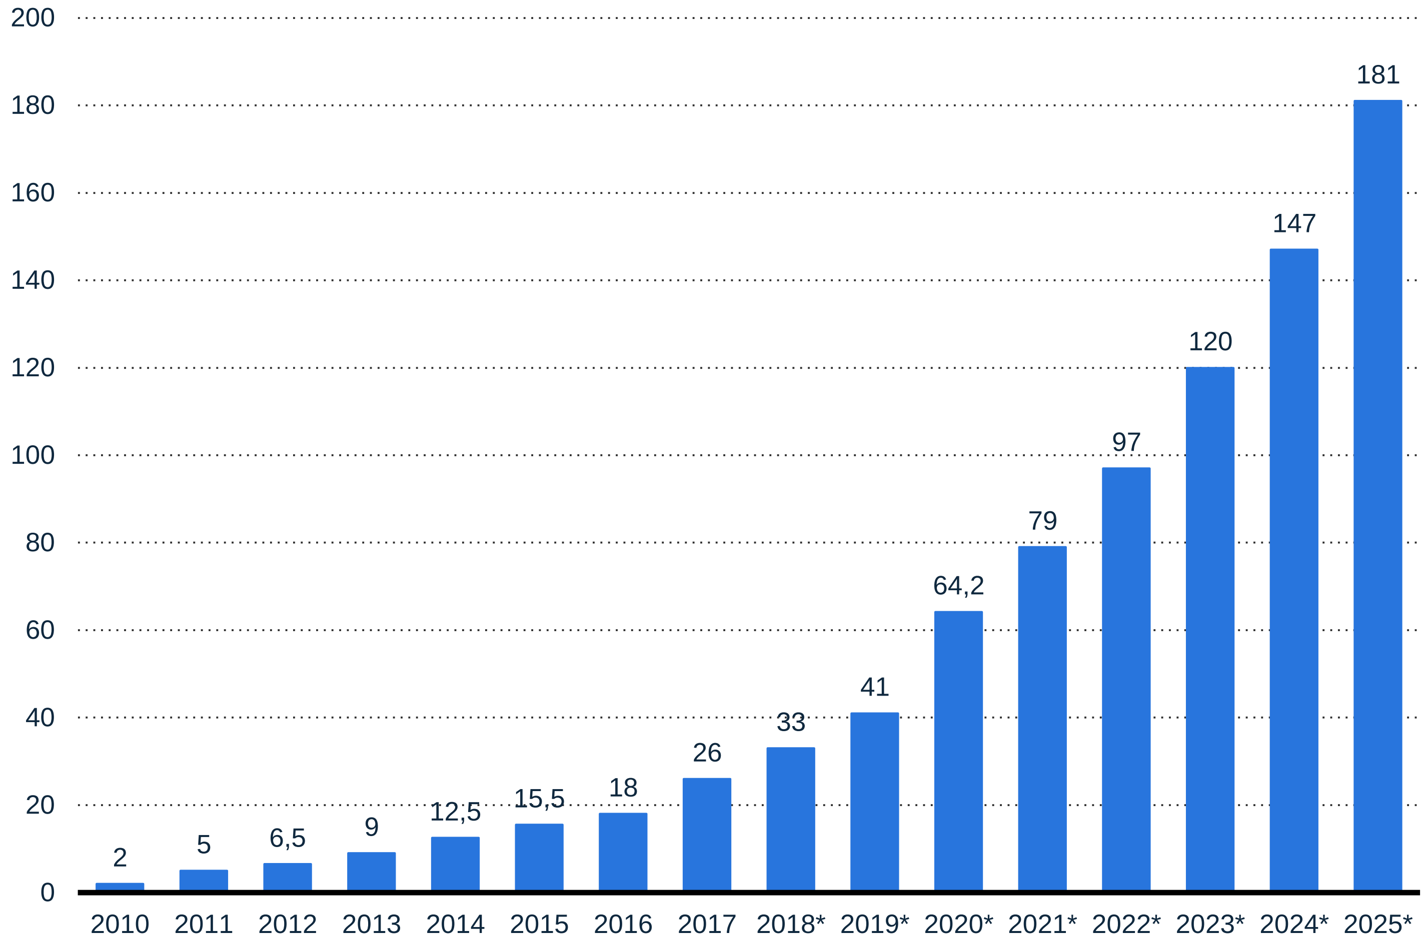

Digital Assets

- Data Volume in Yottabytes (IBM)

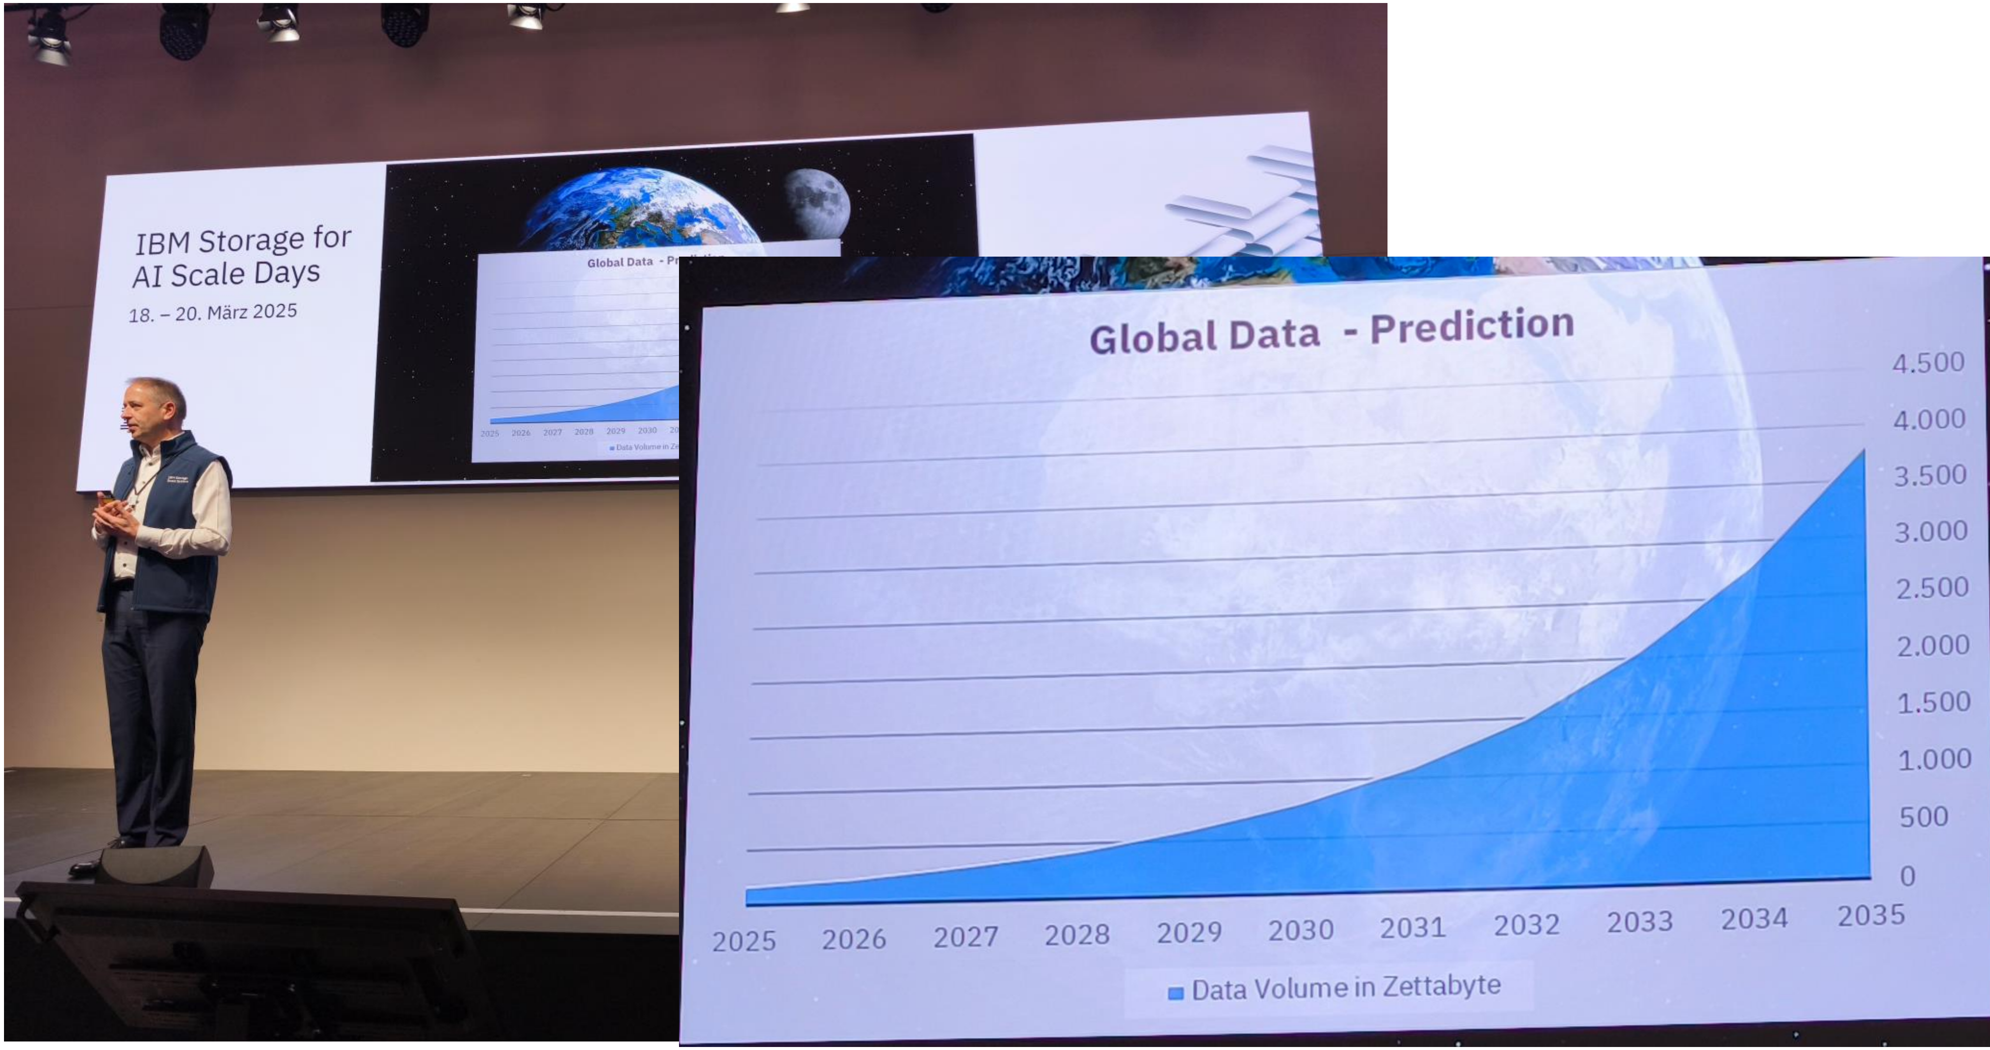

Digital Assets

- **Data Evolution**

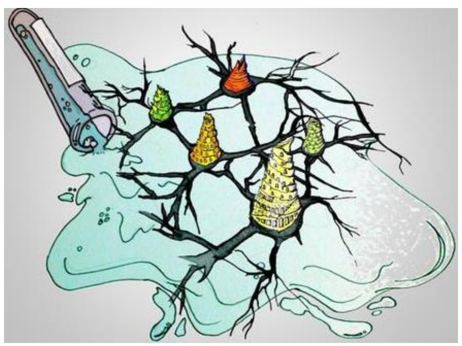

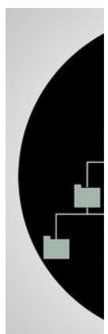

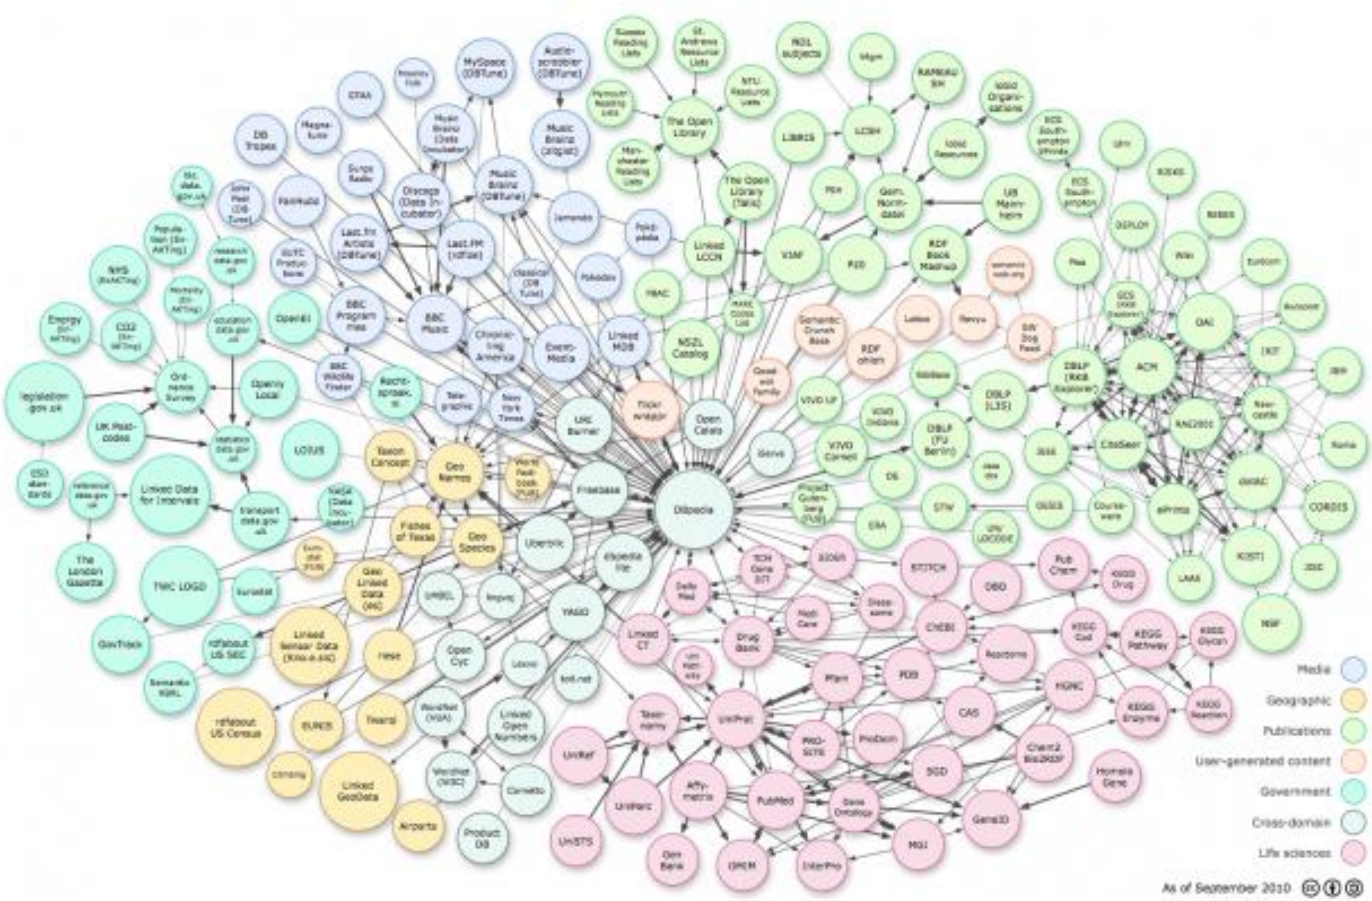

unstructured data | ontologies | linked data | Semantic Web

2000 | **2005** | **2010** | Year, A.D.

Open Data

- **Linked Open Data Cloud**

Legend

- Cross Domain
- Geography
- Government
- Life Sciences
- Linguistics
- Media
- Publications
- Social Networking
- User Generated

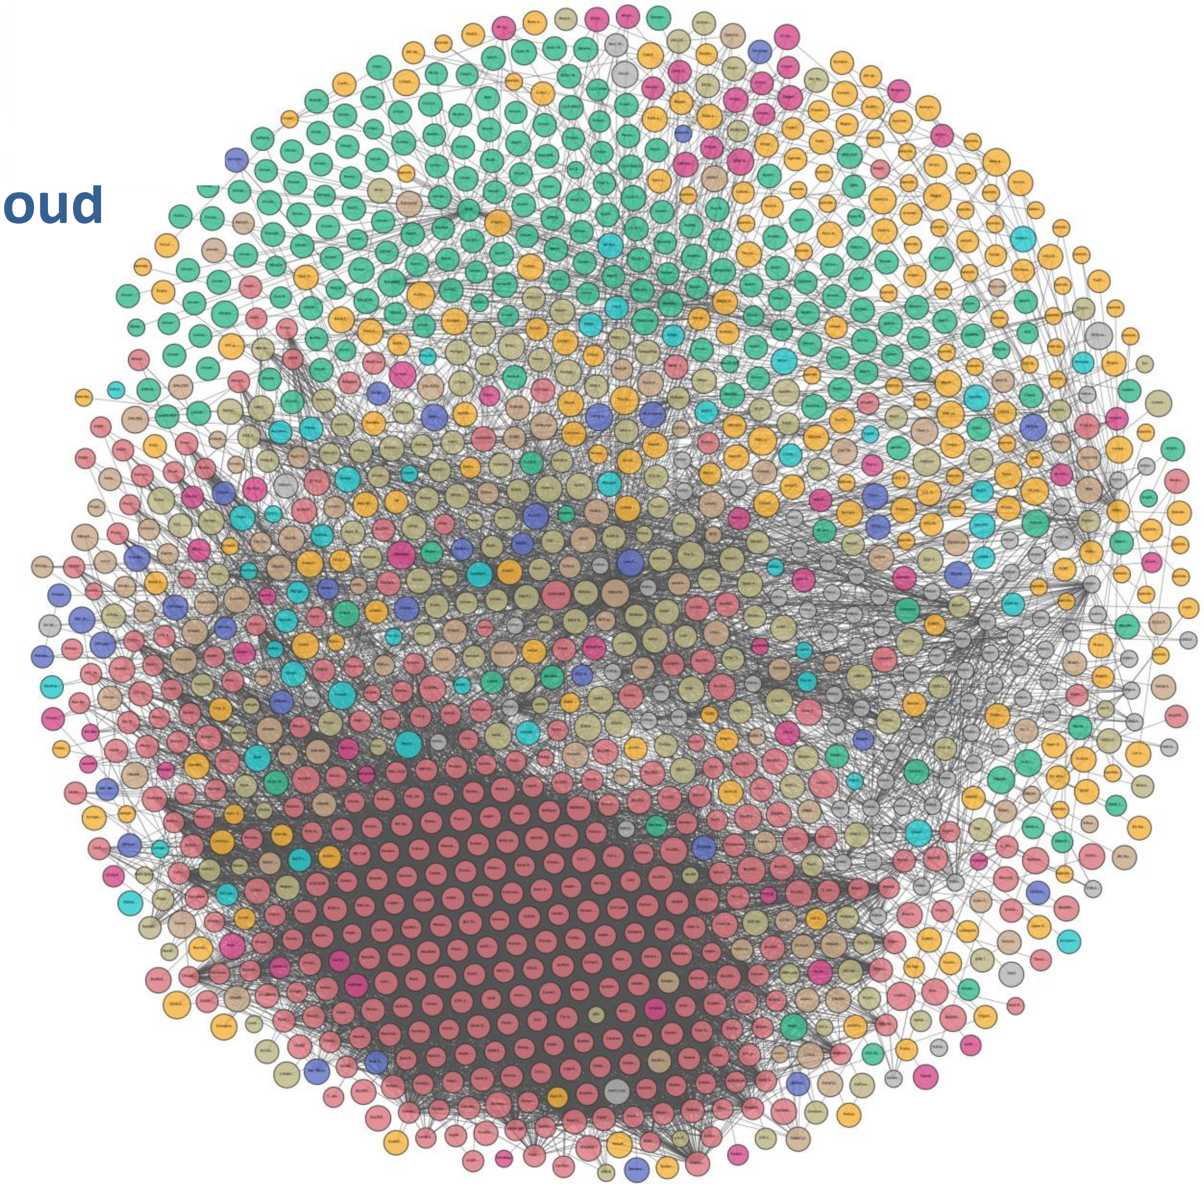

The Linked Open Data Cloud from lod-cloud.net

Data Analytics Applications

- **Semantic Web**
  - A common schema (RDF) for annotation and query of unstructured and unannotated data
    - Google Knowledge Graph
      - **570 million** entities (nodes)
      - **500 billion** facts (connections)
      - **several petabytes** of files
    - Facebook's Social Graph
      - **60 PB** of structured data
      - **1 Trillion** edges
    - Twitter/X's Interest Graph
      - **60 PB** of structured data

Data Analytics Applications

- **Artificial Intelligence**
  – Automated pattern recognition technology based on NN

- **Time to prepare**
- **Number of slides**
- **Duration of speech**

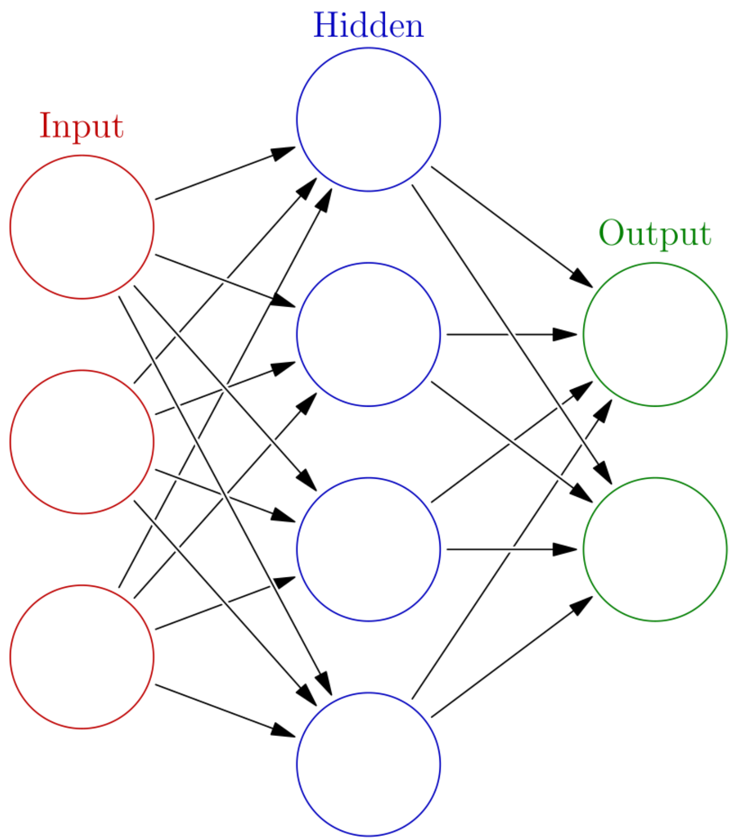

- **A perfect IARIA talk**

Data Analytics Applications

- **Large Language Models**

LLM	Nr. of params	Corpus size (nr. of tokens)	Activated params per token	Training costs, EFLOPS/day	Inference costs
Meta's LLaMA-4 (2025)	• 109 B • 400 B • 2000 B	• 40 T • 22 T • ?	• 17 B • 17 B • 288B	• 71 • 34 • ?	
OpenAI's GPT-4 (2025)	• 1800 B	• 13 T			
DeepSeek-V3	• 671 B				

Big Data is not just about the size, but also about the scale!

Big Data is not just about the size, but also about the scale!

Challenges and Bottlenecks

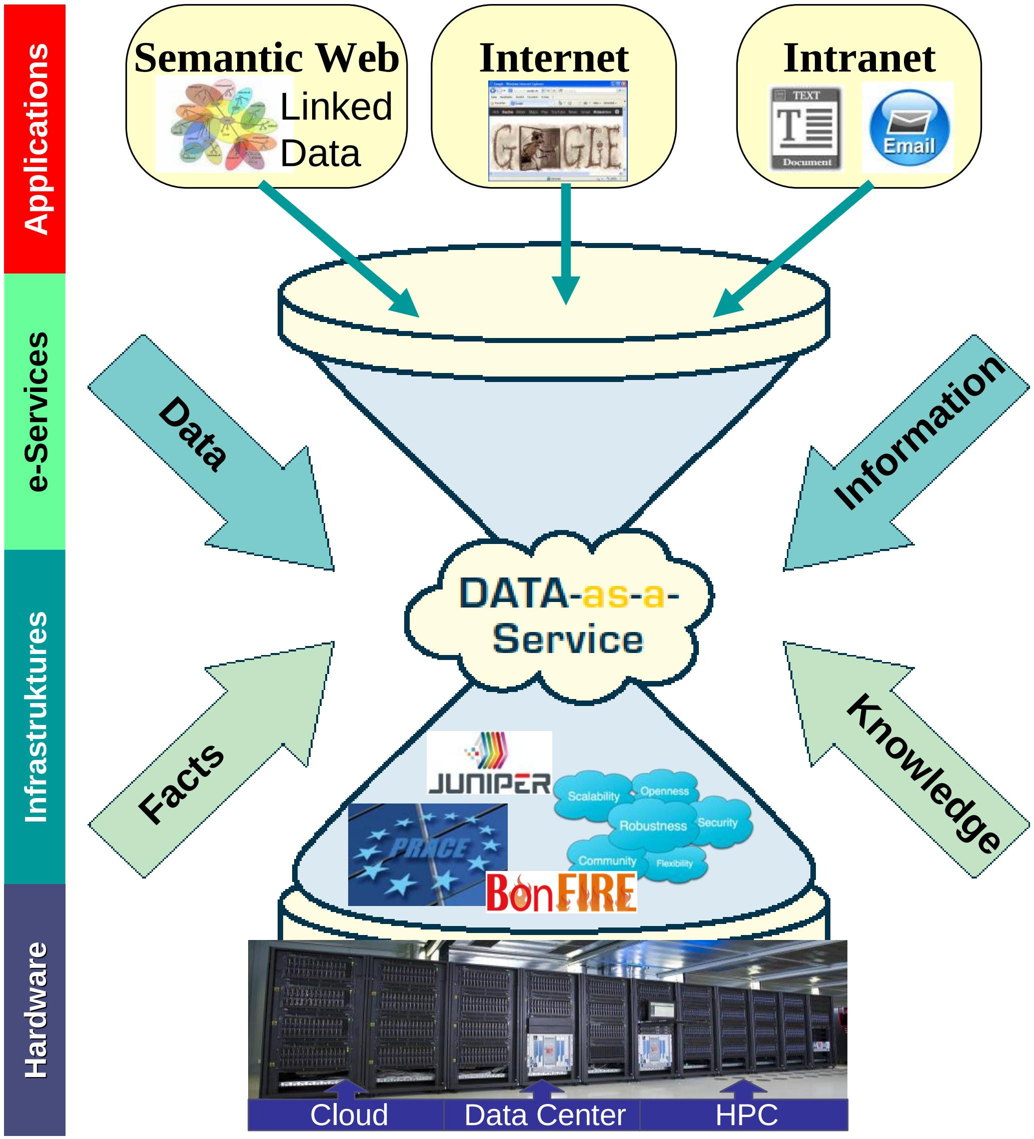

Data Science Infrastructure

- Data Management Issue

BigData-Scientist

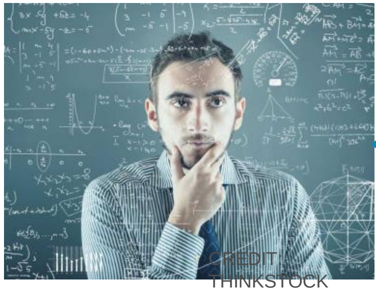

Storage-for-Science Framework

Store

Manage

Publish

Datacenter

HPC

Data Science Infrastructure

- **Hierarchical Storage Management**

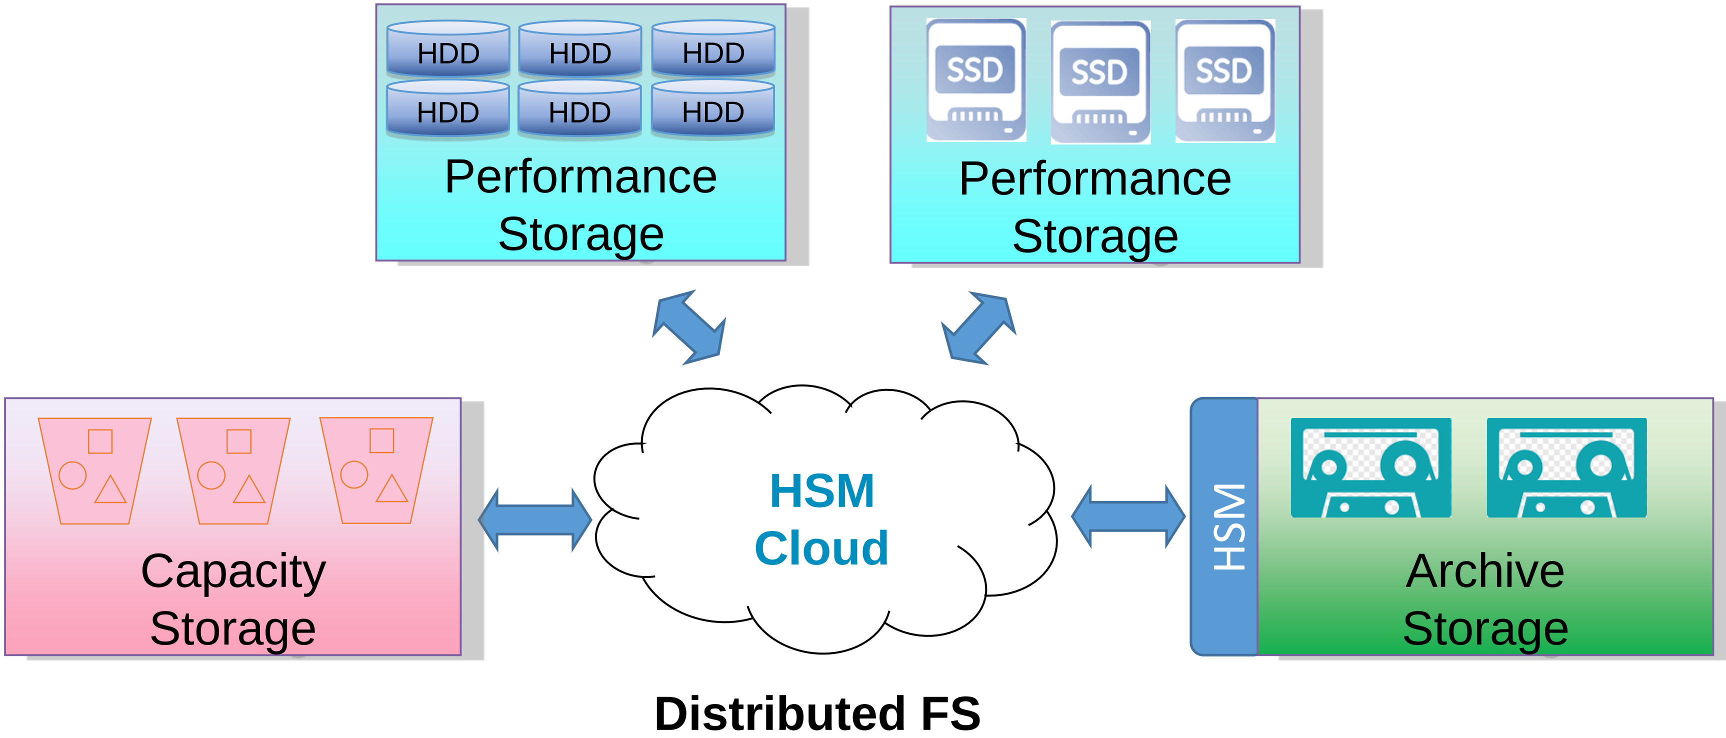

Scientific Data Management

- ## Typical Workflow

Data Infrastructure	Data Management Framework	Application

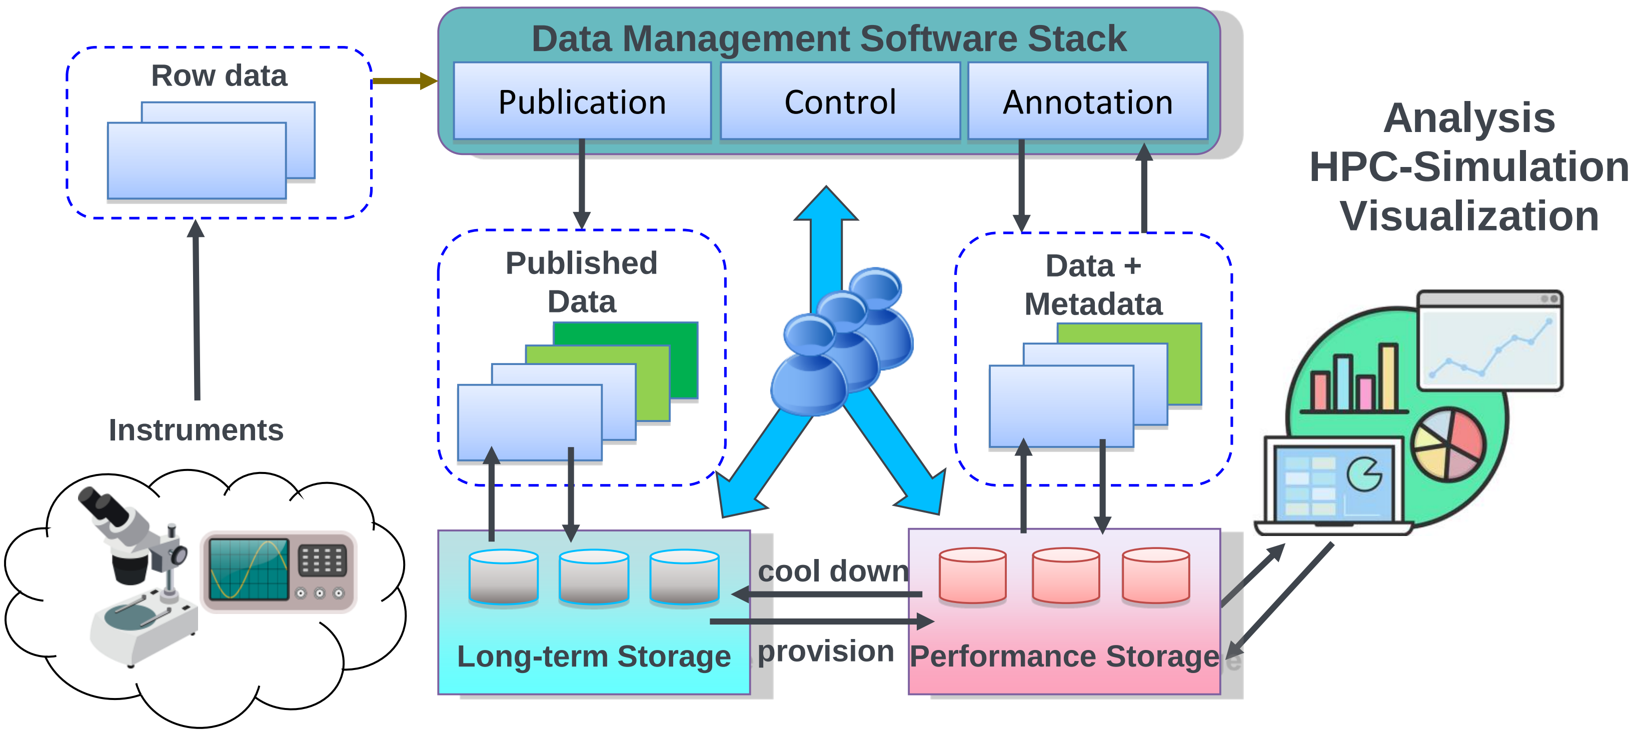

Bw-Storage-for-Science

- **Architecture**
  - Front-end
    - Angular
  - Back-end
    - Golang (OpenAPI)
  - Data Layer
    - Arrango-DB
  - System Software
    - Ansible

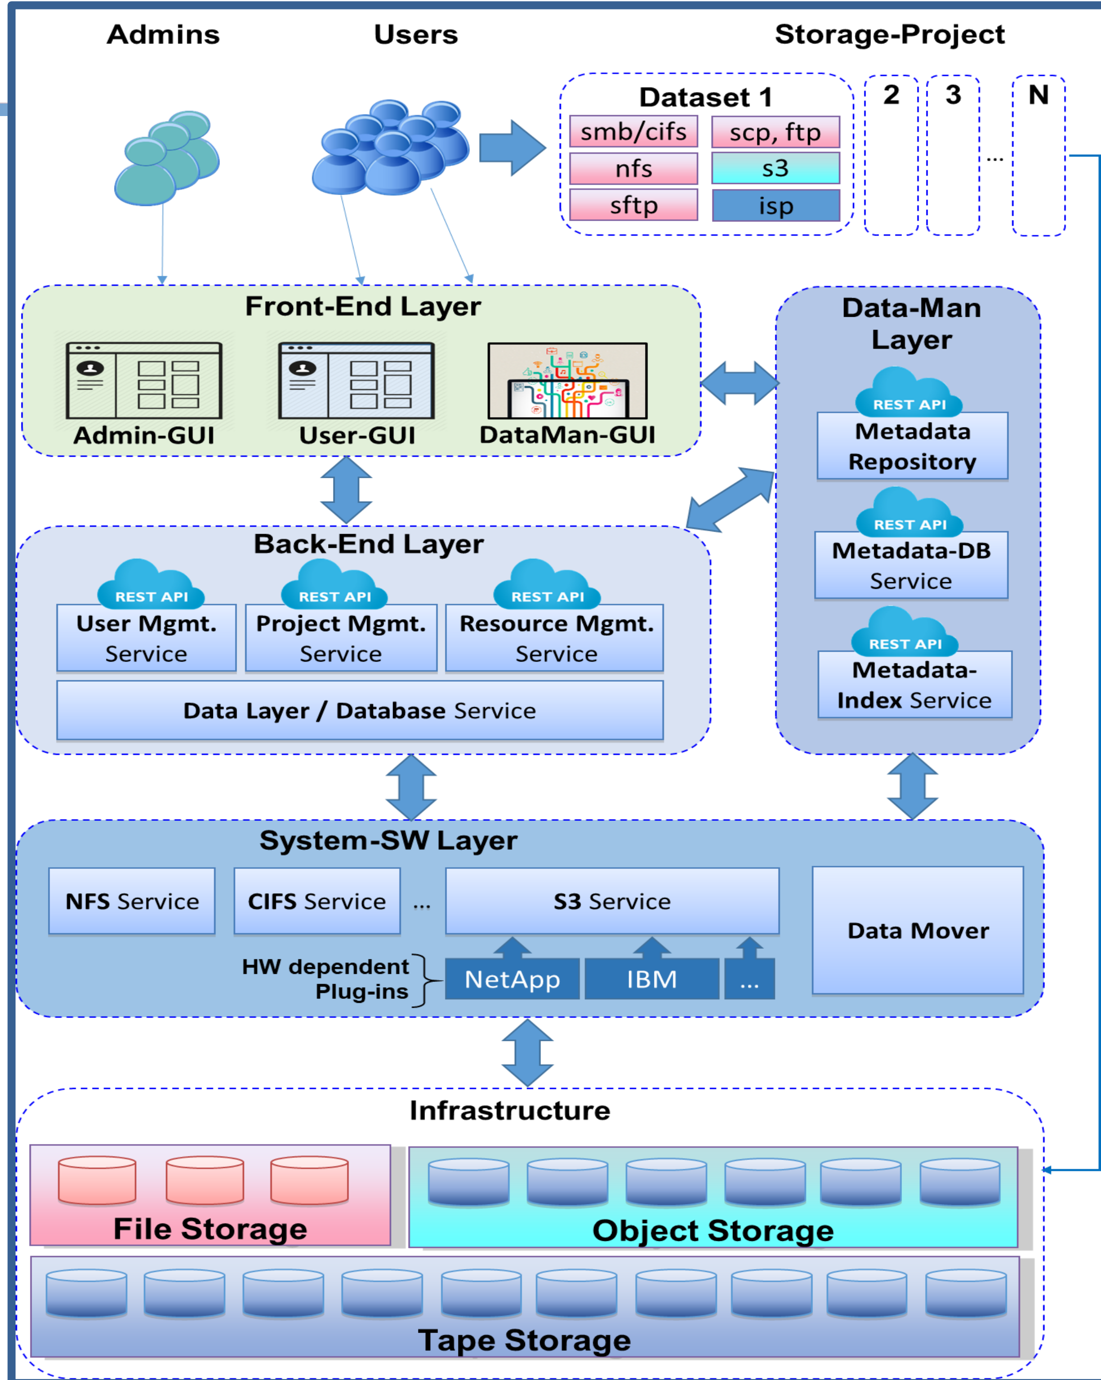

OUTLINE

1. Introduction to High Performance Computing
   - What is HPC
   - How high is the high performance
   - Application examples
2. Data Science
   - Applications
   - Challenges
   - Cloud and Data Center technologies
3. Conclusion

Conclusion

- **Data Centers and Data Science have to find a common language**
  - A lot of integration is required
  - Virtualization infrastructure (*-a-a-S) are beneficial
  - Lack of a common software technology

- **Qualification of researchers**
  - Data Management as a curriculum discipline

Conclusion

- **Duration this talk (1 hour)…**
  - **30.000 hours** of **videos uploaded on Youtube**
  - **1.500 M** of **emails sent**
  - **250 M** of **search queries done on Google**
  - **6.000** of **text messages sent on Whatsup**
  - **20M** of **social network posts published on X**

Thank You!