ARCHITECTURAL DESIGNS AND SERVICES FOR BIG DATA

ScaDS – Competence Center for Scalable Data Services and Solutions Dresden/Leipzig

Dr. René Jäkel

ZIH, Technische Universität Dresden

www.scads.de

@ScaDS

MOTIVATION – BIG DATA CHALLENGES

Increasing Data Sizes

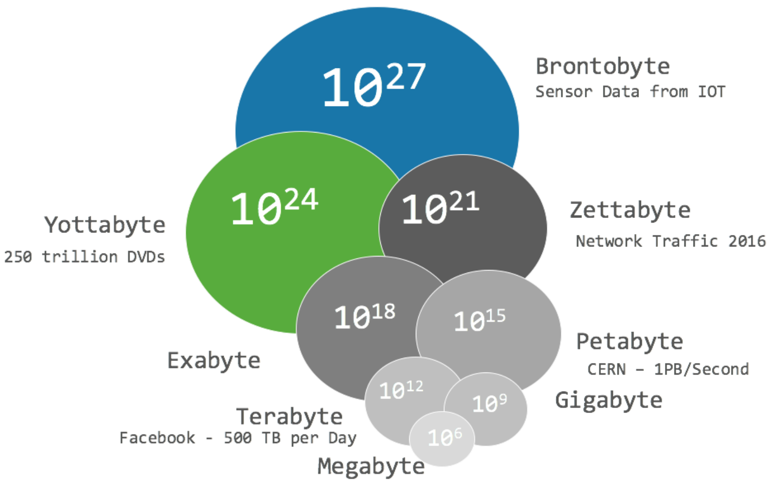

+Technology proliferation

+Veracity, Velocity, Variety

Source: http://api.ning.com

MOTIVATION – BIG DATA CHALLENGES

Increasing Data Sizes

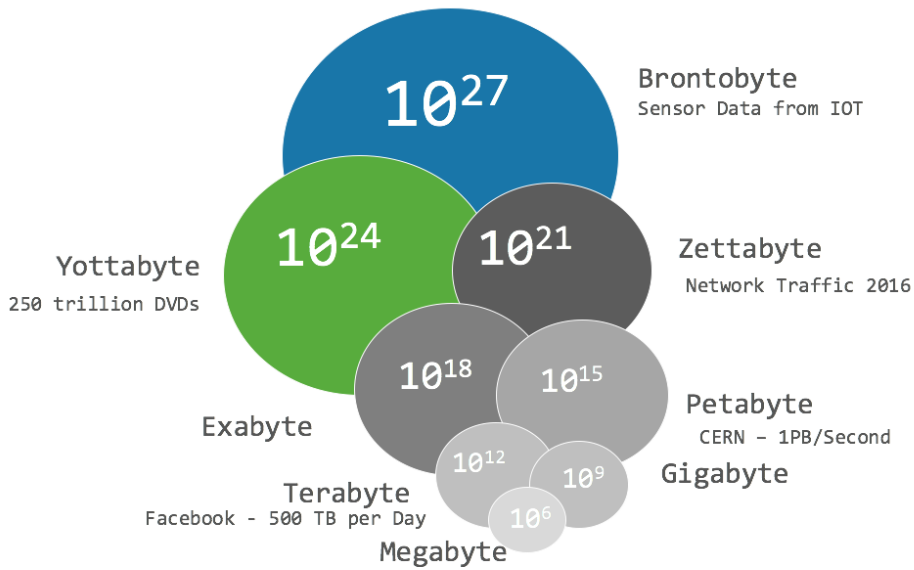

+Veracity, Velocity, Variety

Source: http://api.ning.com

Long tail of “Science“

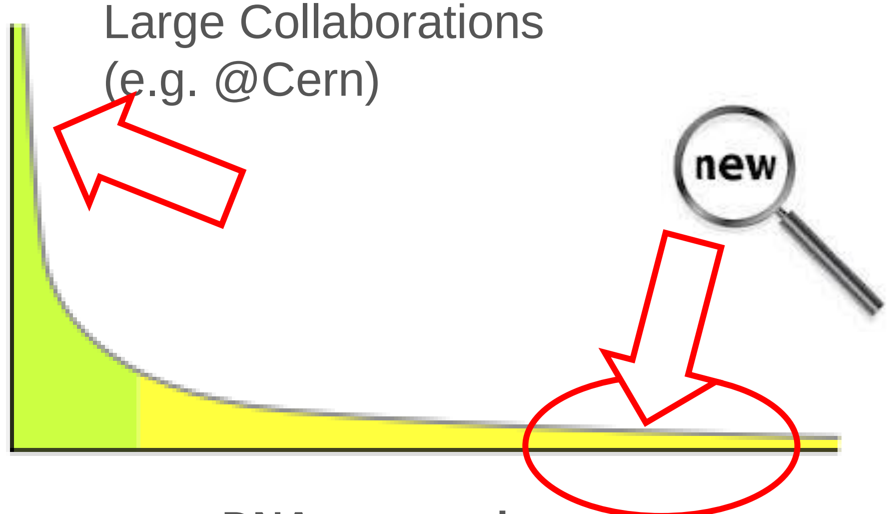

DNA sequencing

Engineering

Transportation

IS THERE A JACK OF ALL TRADES?

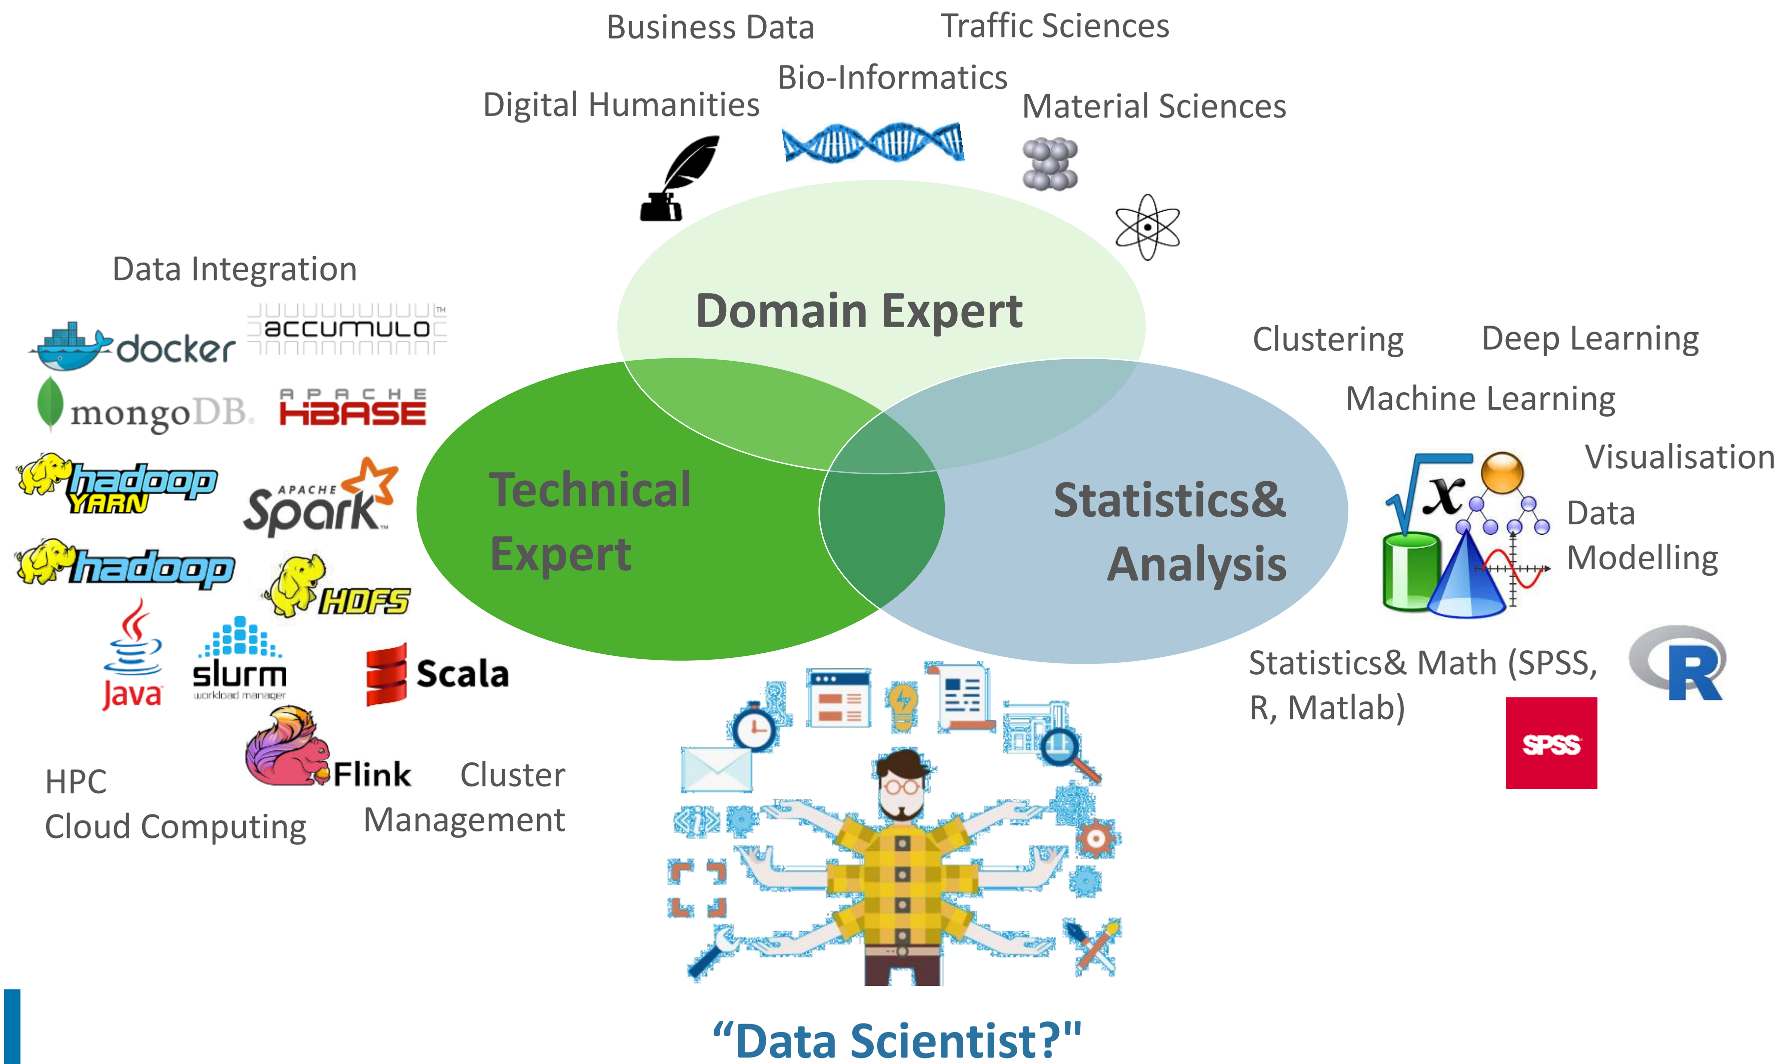

https://www.semrush.com/blog/community-manager-a-jack-of-all-trades/

IS THERE A JACK OF ALL TRADES?

Business Data
Traffic Sciences
Bio-Informatics
Digital Humanities
Material Sciences

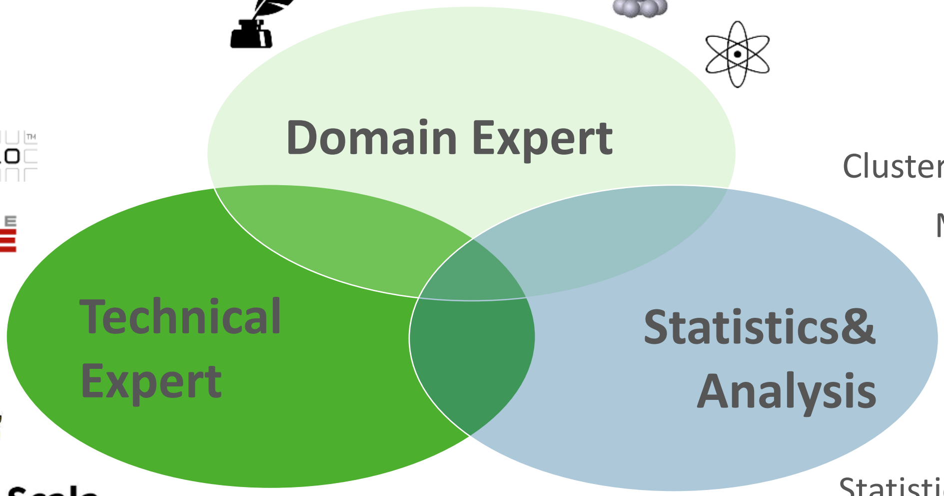

Data Integration

Clustering
Deep Learning
Machine Learning
Visualisation
Data Modelling

Statistics& Math (SPSS,
tlab)

HPC
Cloud Computing

Lets have a team of experts!

https://www.semrush.com/blog/community-manager-a-jack-of-all-trades/

ScaDS DRESDEN LEIPZIG

NATIONAL BIG DATA COMPETENCE CENTER AND ASSOCIATED PARTNERS

Specialists from computer & domain sciences

Focal point for new research activities

Collaborative big data research

Associated Partners

TECHNISCHE UNIVERSITÄT DRESDEN

UNIVERSITÄT LEIPZIG

ScaDS DRESDEN LEIPZIG

CBG Max Planck Institute of Molecular Cell Biology and Genetics

Leibniz-Institut für ökologische Raumentwicklung

HELMHOLTZ ZENTRUM FÜR UMWELTFORSCHUNG UFZ

HZDR HELMHOLTZ ZENTRUM DRESDEN ROSSENDORF

SLUB Wir führen Wissen.

InfAI Institut für Angewandte Informatik

TECHNISCHE UNIVERSITÄT CHEMNITZ

Fraunhofer IVI

Fraunhofer IWS

GISA IT. Mehr als Standard.

SCIONICS computer innovation

ecemp european center for emerging materials and processes dresden

Universitätsklinikum Carl Gustav Carus DIE DRESDNER.

HfTL Hochschule für Telekommunikation Leipzig University of Applied Sciences

LANDESAMT FÜR UMWELT, LANDWIRTSCHAFT UND GEOLOGIE | Freistaat SACHSEN

TIQ.solutions Total Information Quality

[netresearch]

data virtuality

Avantgarde Labs

Netzwerk Logistik

STRUCTURE OF THE COMPETENCE CENTER

Life Sciences

Material and Engineering Sciences

Environmental and Traffic Sciences

Digital Humanities

Business Data

Service Center

Big Data Life Cycle Management und Workflows

Data Quality/
Data Integration

Knowledge Extraction

Visual Analysis

Efficient Big Data Architecture

COMPUTER SCIENCE RESEARCH (METHODS)

- Efficient Big Data Infrastructure
  - Hardening computation infrastructure (Security)
  - Flexible cluster management
  - **Big Data Framework Execution & Monitoring on HPC**
  - Geo-temporal data storage
  - Cloud-based service support for analysis of travel data
- Big Data Lifecycle and Workflows
  - **Execution of large data-driven workflows** (KNIME-workflow integration @HPC)
  - Time series management and forecasting
- Holistic data integration
  - **Privacy-Preserving Data Matching**
  - Deduplication (in Graphs)
- Visual Analysis
  - Improve visualization of large particle data
  - **Multi-scale visualization for engineering data**
- Knowledge Extraction
  - Porting computer vision algorithms on GPUs
  - Knowledge Extraction on biological and environmental data
  - **Deep Learning & structure recognition in spatial planning**

See Talk of Prof. Rahm

SERVICE CENTER AS FOCAL POINT

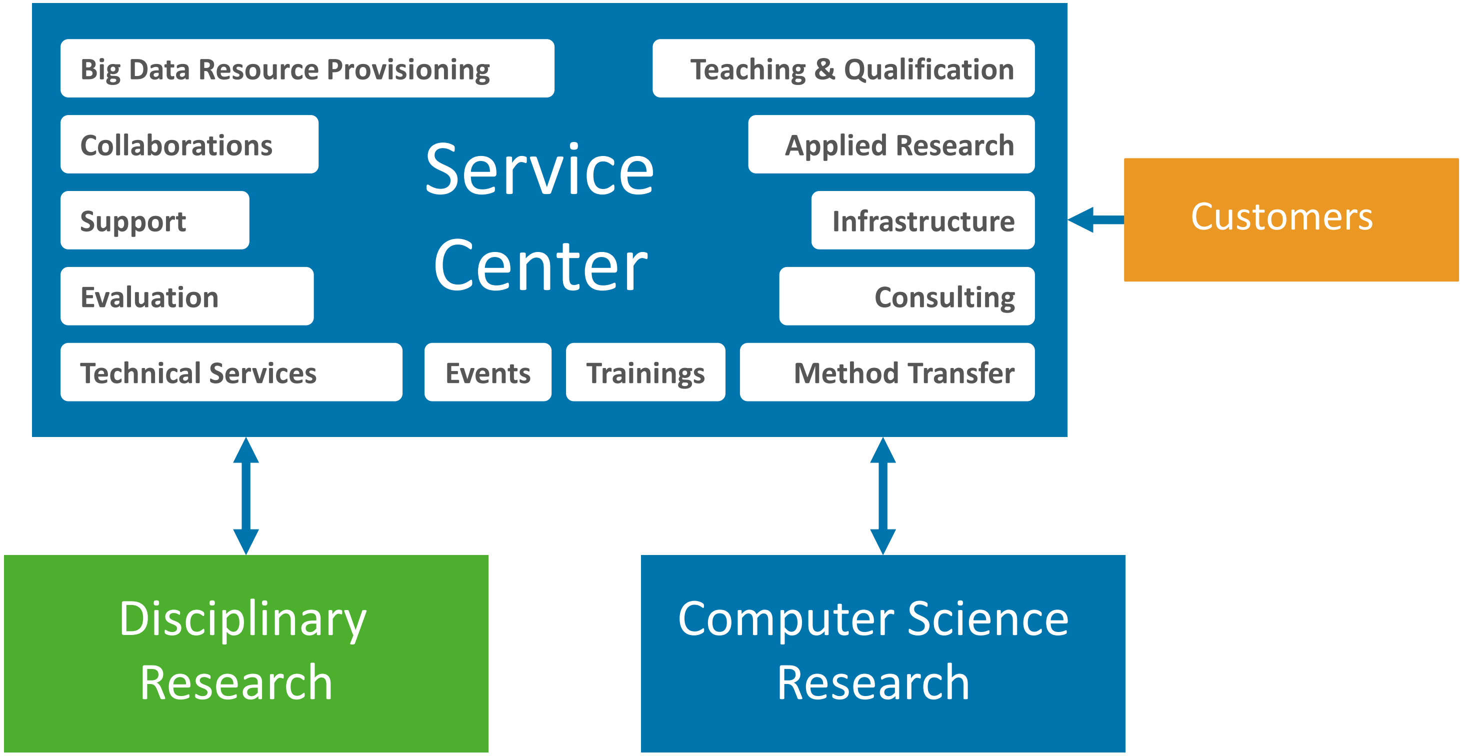

- 10 new projects contributing directly to ScaDS Dresden/Leipzig (new colleagues in ScaDS labs in Dresden and Leipzig)
- 13 further ScaDS-associated project acquisitions: 26 positions

SERVICES AND DEMONSTRATORS

- Goal: open services for scientific communities
- Starting as demonstration services
  - **Visualization for multiscale simulations in engineering**
  - Focus and context methods for point-based data
  - ECAST service
  - Entity-Augmentation
  - **Binary image segmentation**
  - Sierra Platinum: Peak-Calling
  - Imputation service
  - **Analytics service for time series** See Talk of Jan Frenzel
  - Innoplan service
  - Wind anomaly detection
  - Text repository & mining services for Digital Humanities (CTS)
  - Graph Analytics Service

CONTACTS TO INDUSTRY AND WORKSHOPS

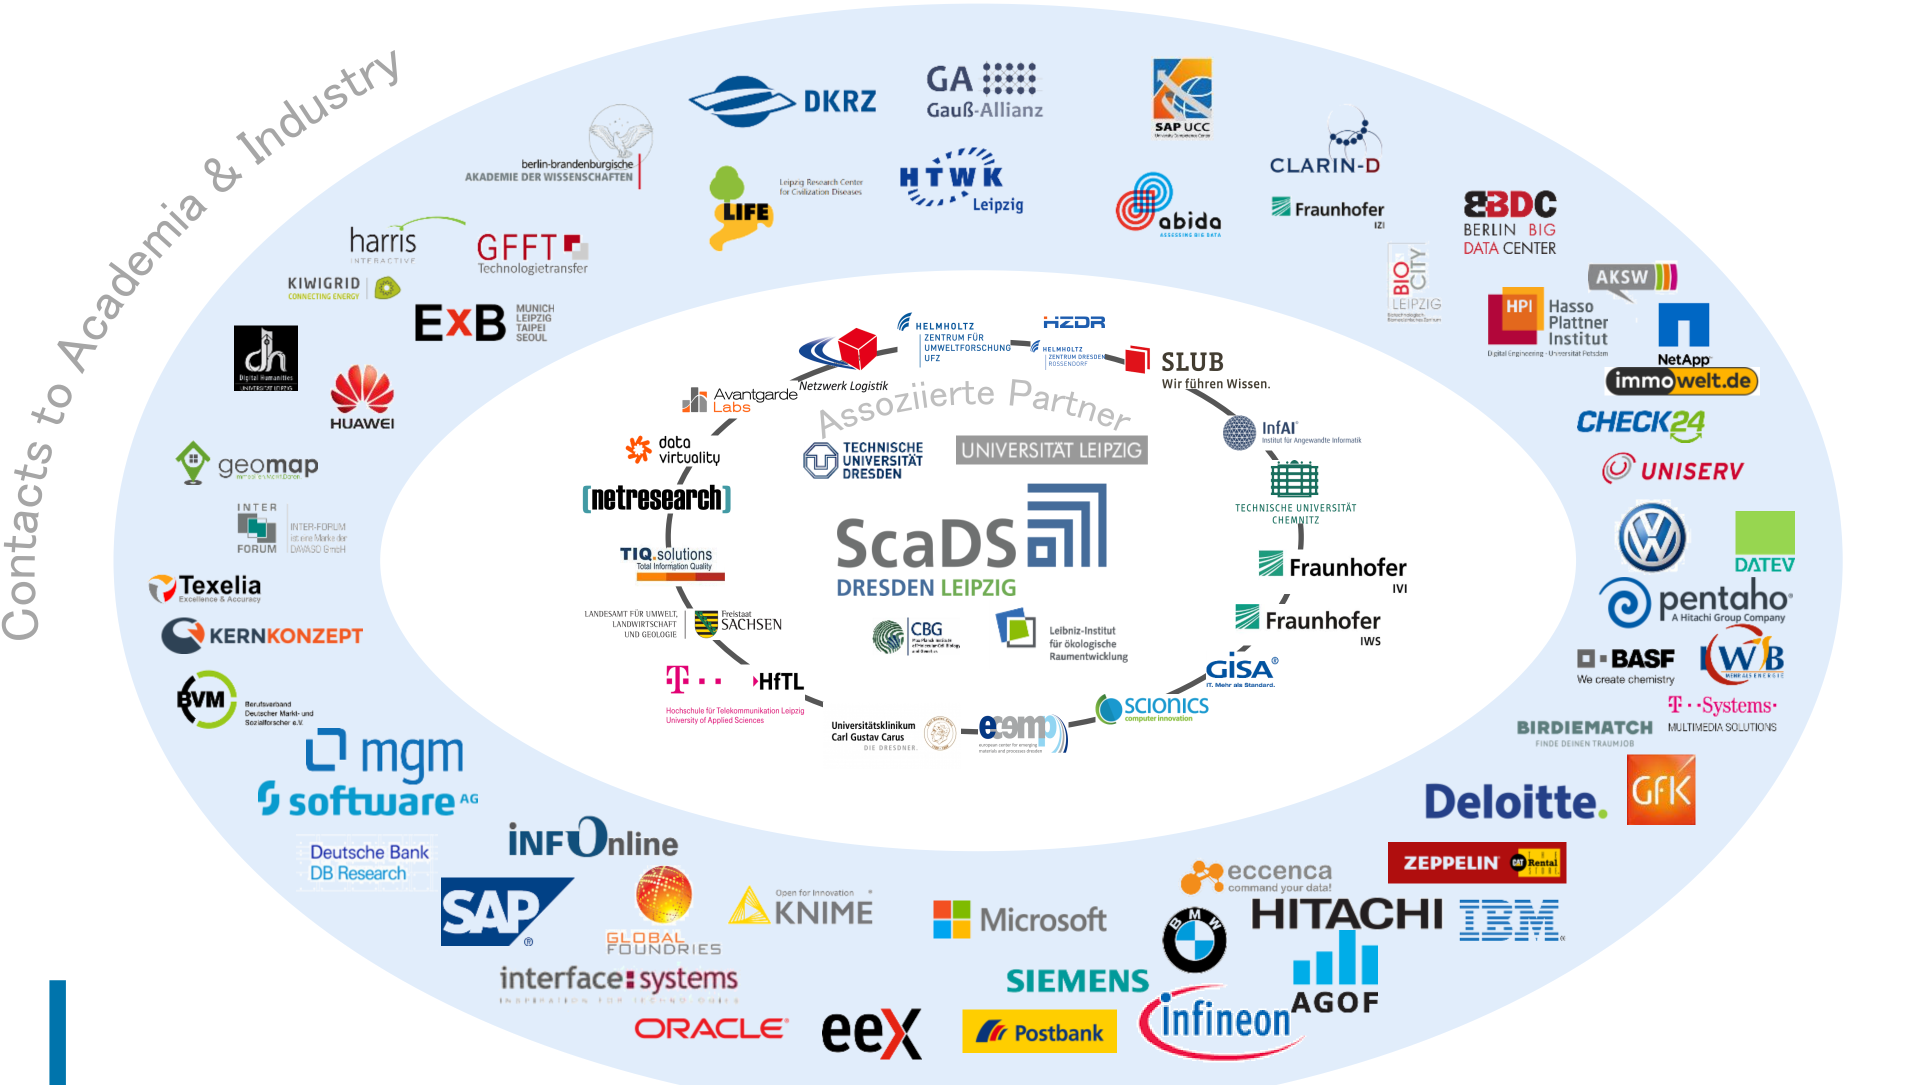

SCADS EVENTS & GUEST PROGRAM

- 3 Big Data in Business (BiDiB) Workshops
- 3 successful international summer schools (Dresden, Leipzig, Munich)
  - more than 250 national/international guests
- Big Data All-Hands-Meeting in Dresden, June 2016
- 30 renowned experts in guest program (21 short-term, 6 mid-term, 3 long-term)

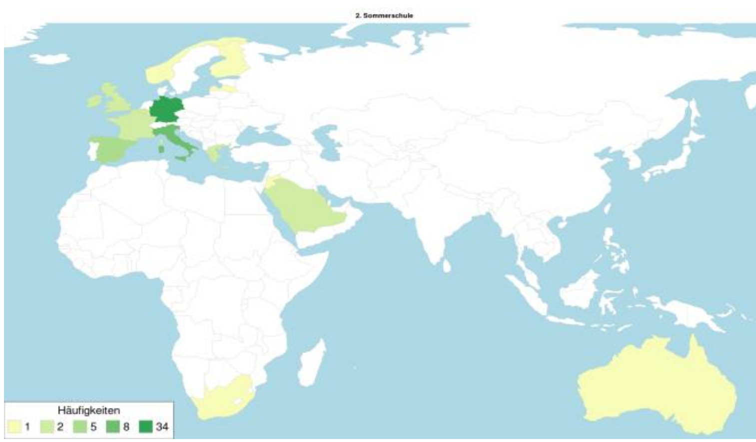

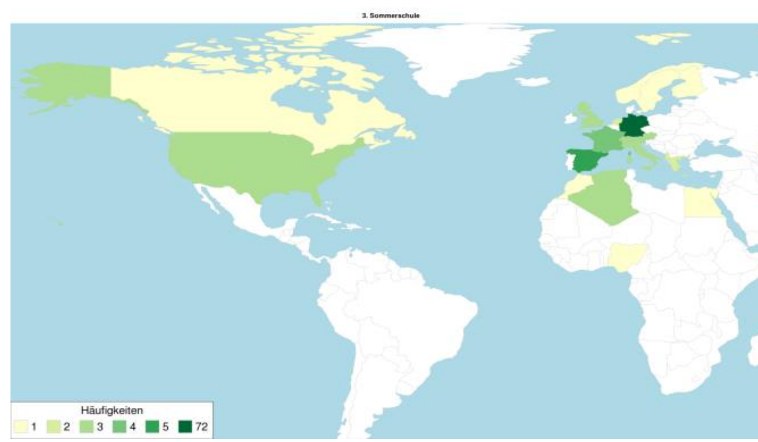

INTERNATIONAL VISIBILITY& OUTREACH

- >120 of publications, > 200 talks worldwide
- Industry talks: i.e. Data2Day, Bitkom Big Data Summit, Fosdem, Flink Forward

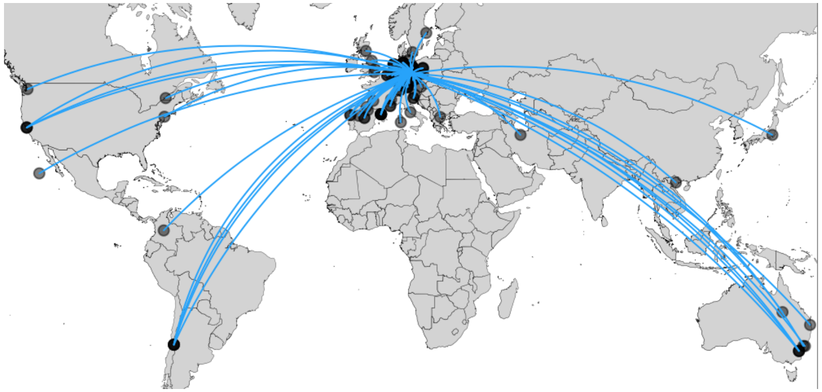

- Awards:
  - Best Science Paper Award der British Machine Vision Conference (BMVC) (Cooperation Prof. G. Myers und Prof. Carsten Rother)
  - Winner of SciVis-Contest IEEE VIS (Group of Prof. Gumhold)
  - Best Demo Award BTW 2017 (Gradoop), 3rd place Data Science Challenge
  - Staatspreis für Innovation, Category „Transfer“ Dr. Stefan Kühne

ORGANIZED WORKSHOPS & SUMMERSCHOOLS

ScaDS Big Data Industry Forum

Overview | Agenda | Blending Tools and Data in KNIME | Linked Big Data Analytics using IBM System G

1. Big Data All-Hands-Meeting in Dresden

Fachgespräch „BigData – Konzepte zur Analyse komplexer Infrastrukturen“

ScaDS DRESDEN LEIPZIG

STRUCTURE OF THE COMPETENCE CENTER

Life-Sciences

Material Sciences

Environmental and Traffic Sciences

Digital Humanities

Business Data

Service-Center

Big Data Life Cycle Management und Workflows

Data Quality/ Data Integration

Knowledge Extraction

Visual Analysis

Efficient Big Data Architecture

ARCHITECTURES FOR DATA ANALYTICS

HPC AND BIG DATA INFRASTRUCTURE

- New machine room and HPC-infrastructure:
  HRSK-2 inauguration May 13th 2015
- Current HPC installation:
  > 1 PetaFlop/s, > 5 PB HDD, > 40 TB SSD
  > 130 TB main memory
- Further systems suited for various purposes:
  - SGI UV 2000 (Venus)
  - “Galaxy” cluster @Leipzig+Dresden;
    90 nodes “Shared-Nothing” architecture
  - Research-Cloud: 13 nodes with OpenStack
    (64 cores, 64GB RAM, 250GB lokal disk)

DEUTSCHER RECHENZENTRUMSPREIS
2014
1. PLATZ KATEGORIE 8: ENERGIE- UND RESSOURCENEFFIZIENTE RECHENZENTREN

BIG DATA ARCHITECTURES - MOTIVATION

Big Data / Data Analytics

- Active research field for almost all scientific domains
  - What insights can we get out of broad data base?
  - Usually only prototypic ideas/solutions
- Limited technological and methodical knowledge in domains present
- High **potential** to characterize value within data by using state-of-the-art methods within frameworks (Hadoop-Ecosystem, deep learning, statistics, …)

HPC

- Highly **specialized hardware**; efficient use requires special knowledge
- Traditional **batch system** based interaction, hardly to be integrated into complex workflows – good for large parallel applications
- „**fixed system**“: no standard methods to shape environments for special needs
- Often mainly **centralized** storage

MOTIVATION: HOW TO SUPPORT USERS WITH INFRASTRUCTURES

- HPC vs. Data Analytics
  - Bring computing to data, or data to computing (data mover)?
  - Systems and infrastructure should support users, not forcing them to follow rigid regiments
  - Let user pick up approach, which is best for individual use case
  - HPC: traditional rather monolithic usage, e.g simulations
  - Big Data analytics: more data centric, but not all and every analysis is embarrassingly parallel, iterative models still induce large data movements
- There is no unique Big Data blueprint!
- From the users perspective – which way to follow?
  more HPC like approach or dynamic possibilities of big data frameworks?

BIG DATA AND HPC – CONVERGENCE PATTERNS

Requirements to support Big Data workloads on HPC

- Support frameworks: more versatile software stacks
- Fast access to data: not just self-production of data (simulation), but also use 3rd-party data (open data, domain repositories)
- Support different data processing paradigms on very same system:
  - Batch vs. Streaming
  - In-Memory and iterations
- Better support of evaluation of (temporary) results, e.g. visualization frontends
- Service orientation (working environments)

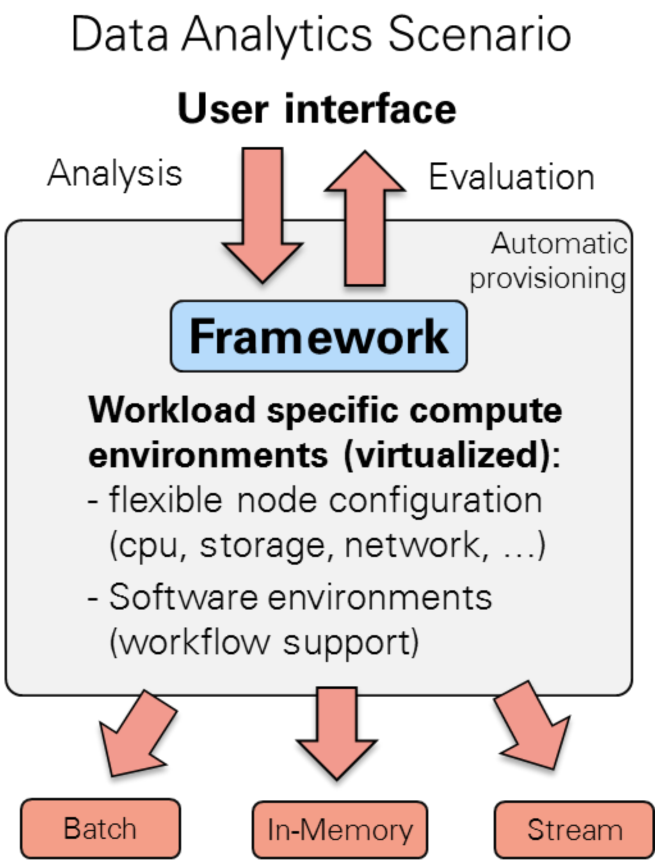

EXTENSION OF HRSK-II FOR HPC DATA ANALYTICS (HPC-DA)

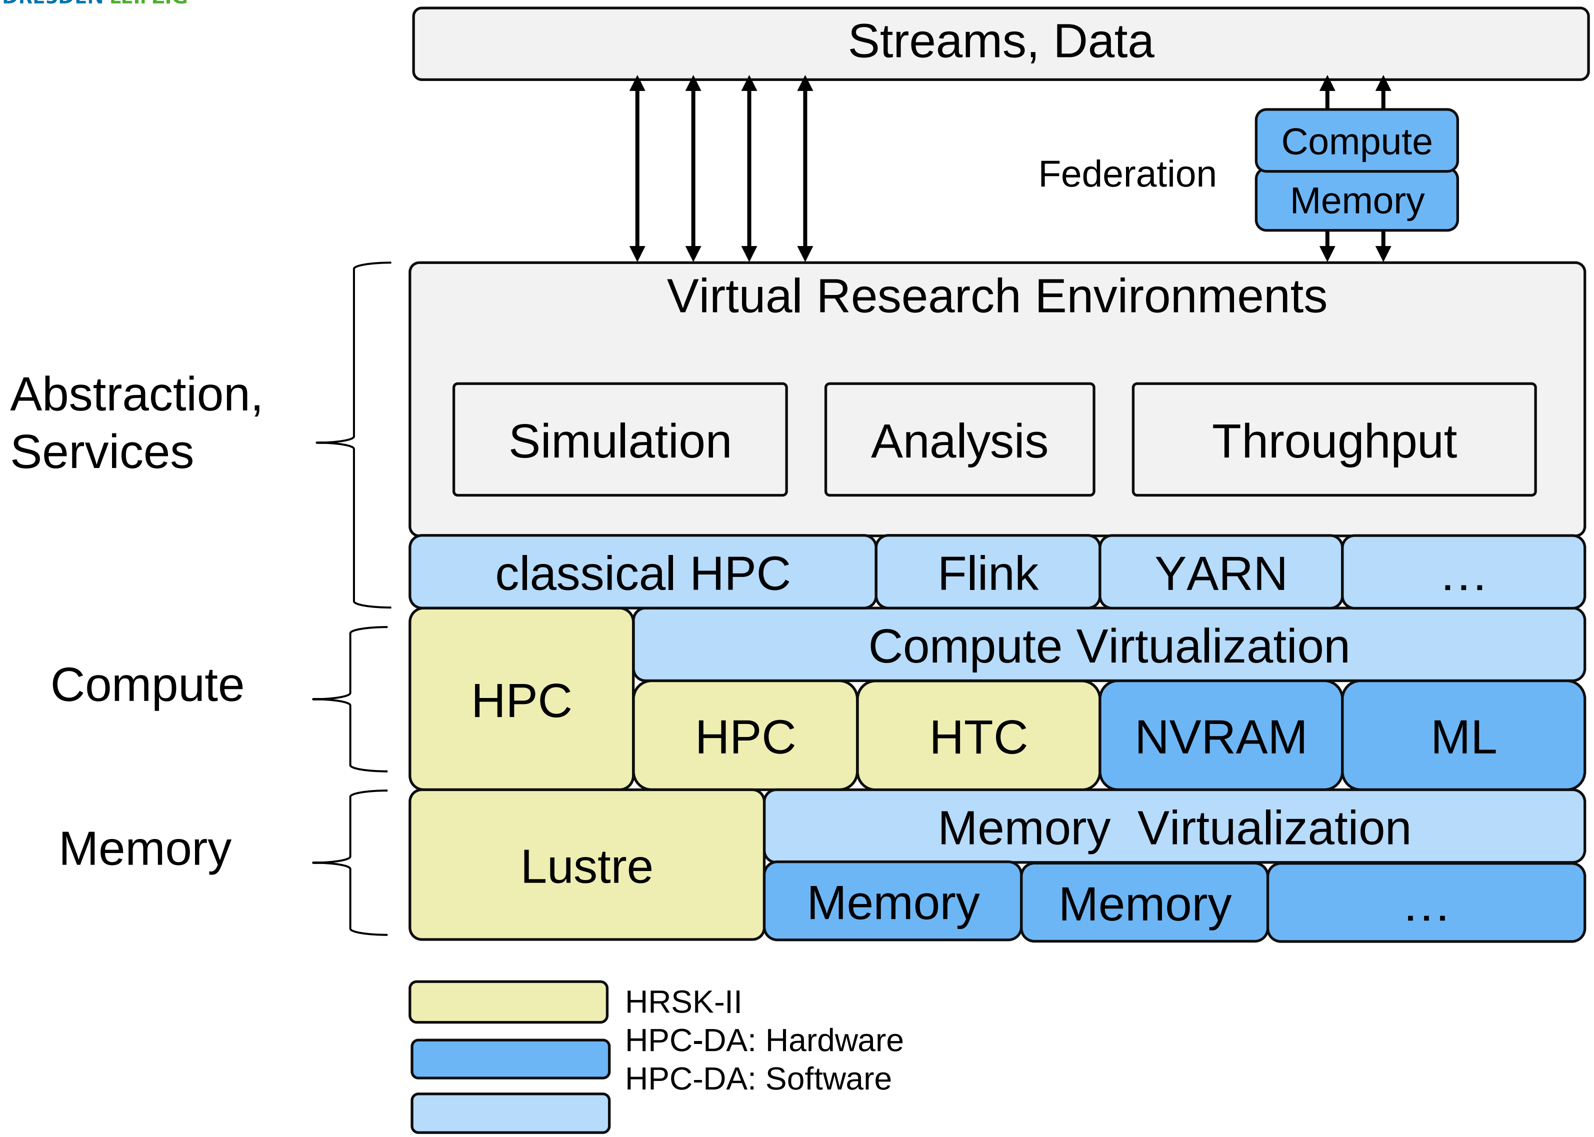

HRSK-II HARDWARE EXTENSIONS (PHASE 1 + 2)

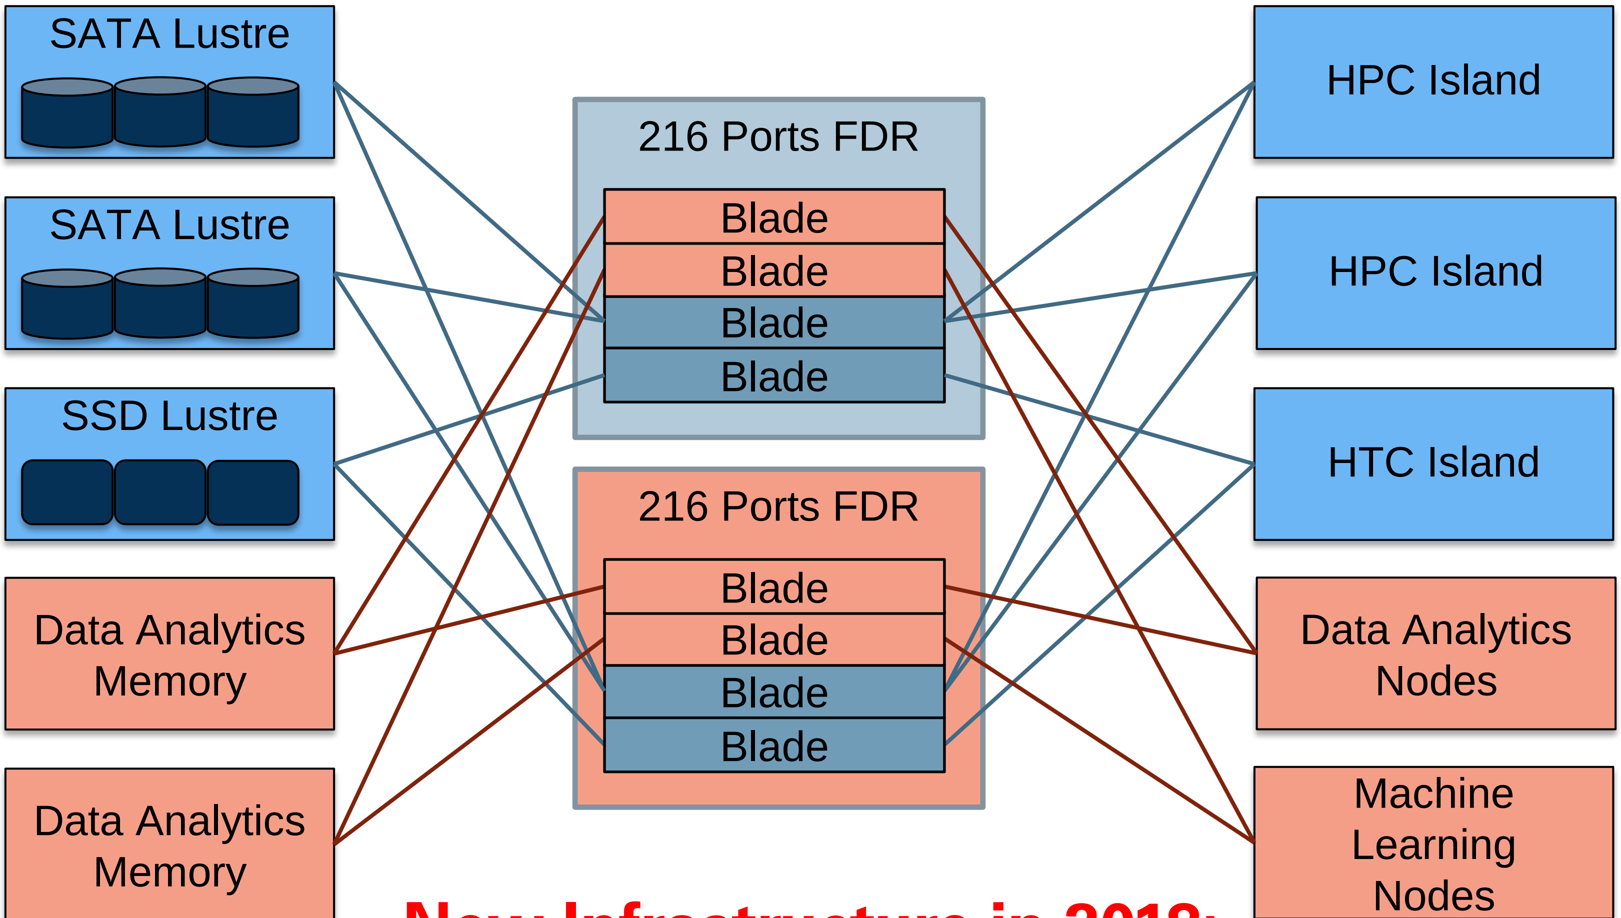

New Infrastructure in 2018: data analytics and ML nodes

PROVISION OF DATA ANALYTICS@HPC

- Provisioning of required environments (Hadoop, Spark, Flink, ML-frameworks …)
- Big Data session created on demand
- Run directly as analytics service at HPC site
- Adoptable to other frameworks/applications

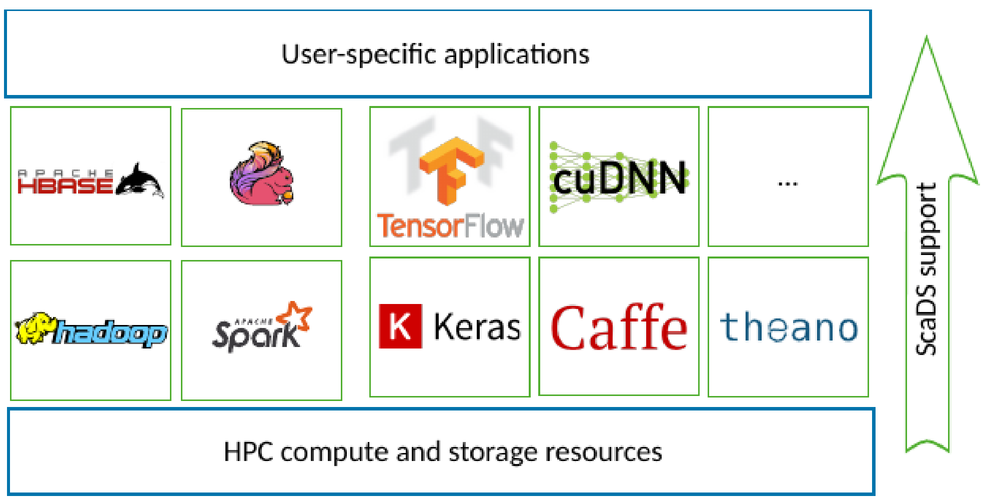

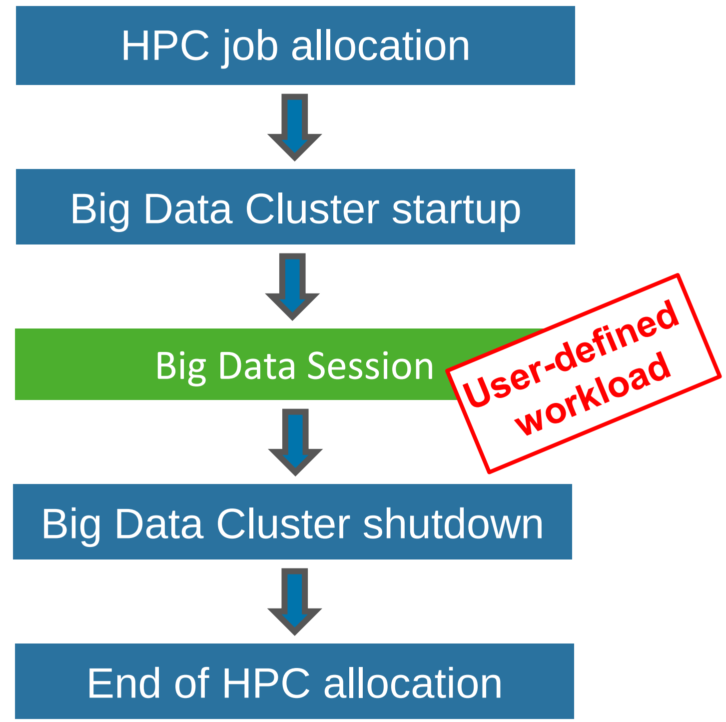

ANALYTICS USE CASES

EXECUTION OF LARGE DATA-DRIVEN WORKFLOWS

- Use Case: processing pipeline for cell tracking (bacteria E.coli) over time

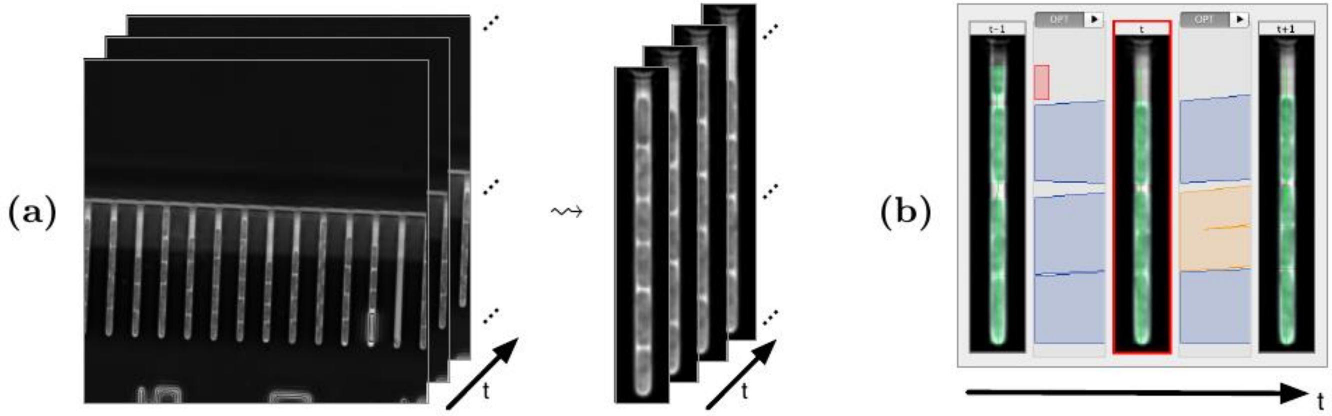

- Challenge: support execution of data-intensive user workflows in HPC environment
  - No prior HPC-knowledge required on user side
  - Formulation of workload directly in workflow environment
- Solution: combination of well-known and widely used tools
  - KNIME for workflow formulation
  - Middleware UNICORE used for HPC interaction

EXECUTION OF LARGE DATA-DRIVEN WORKFLOWS

- First: export of workflow and its input data
- Second: automatic generation of compute jobs and execution on HPC system

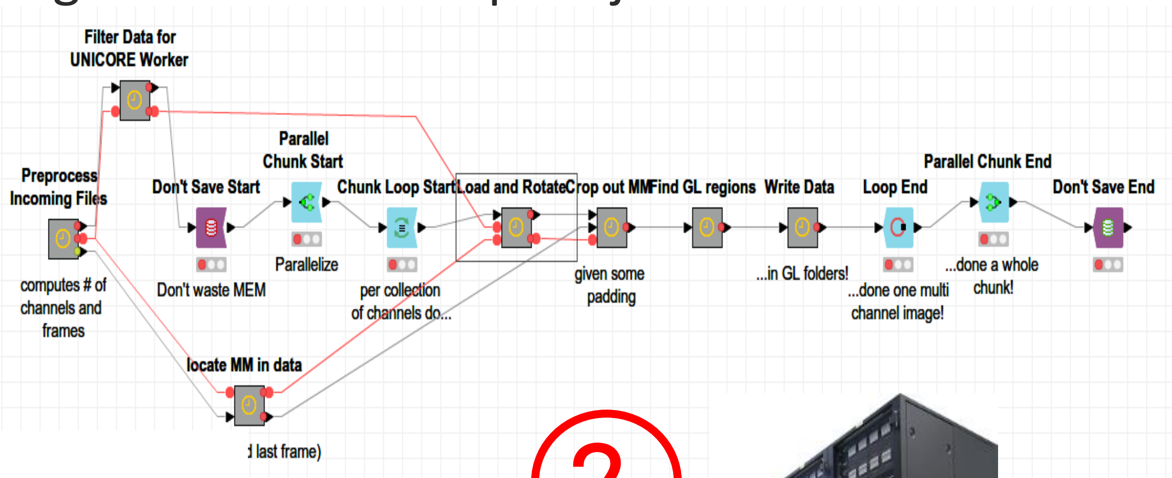

1.

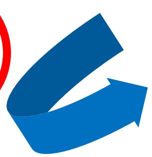

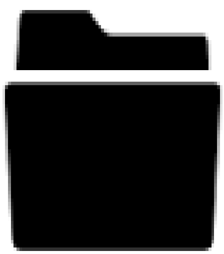

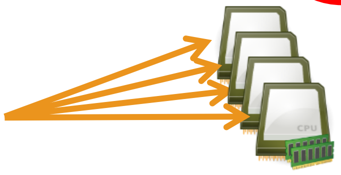

2.

Automatic generation of thousands of computing jobs if required

EXECUTION OF LARGE DATA-DRIVEN WORKFLOWS

- Evaluation data set:
  1,8 TB in ~7,5 M files
- Runtime improvement:
  previously 17d on 4 cores
  now 2h on 800 cores
- **200x** faster parallel execution
- Next steps:
  fully automated pipeline connecting microscope with HPC environment and research data repository

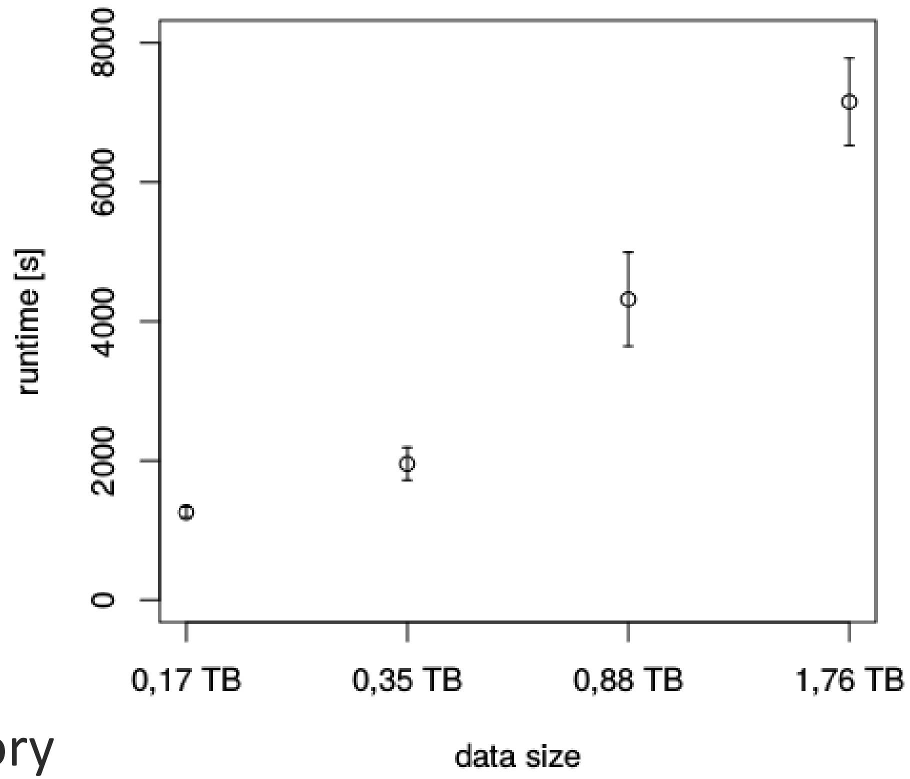

DATA ANALYTICS FOR INTRAOPERATIVE THERMAL IMAGING

- Application area: low delay operation support using thermal imaging processing
- Quasi real-time data processing required in decision support during surgery – University Hospital Dresden (UKD)
  - Neural activity monitoring require long-term intraoperative measurements (~10 minutes)
  - Fast preprocessing required to decrease delay for subsequent analysis workflows and result presentations
    => minimize overall OP delay
  - Iterative process: 3000 frames (5.4 GB) have to be processed every minute (50 Hz sampling rate)

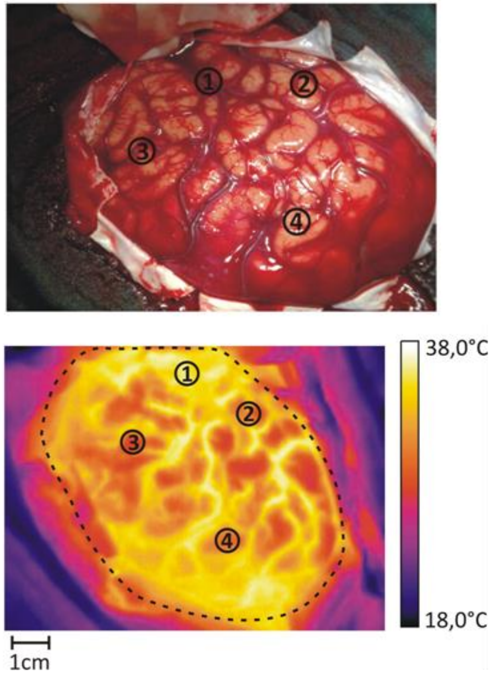

Thermal image of acute subdural hematoma

DATA ANALYTICS FOR INTRAOPERATIVE THERMAL IMAGING

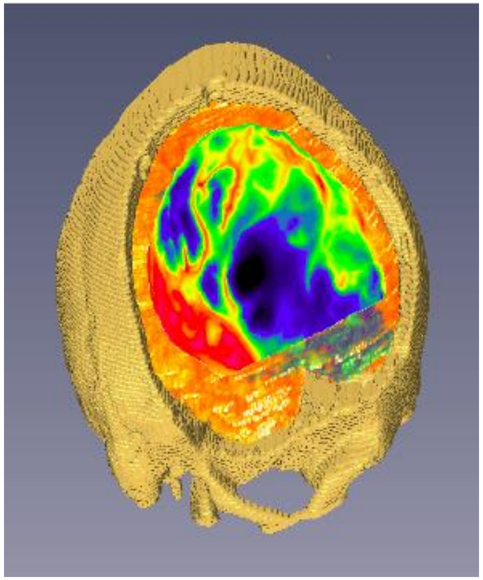

- Application area: low delay operation support using thermal imaging processing
- Quasi real-time data processing required in decision support during surgery – University Hospital Dresden (UKD)
- Solution: Provision of Spark-Cluster @HPC
  - Fast SSD-backend to speed-up IO; fail-safe storage of imaging data
- Runtime improvement:
  - UKD-workstation: ~7000s/30.000 images
  - Spark cluster @Taurus: ~32s/30.000 images

→ ~**220**x faster

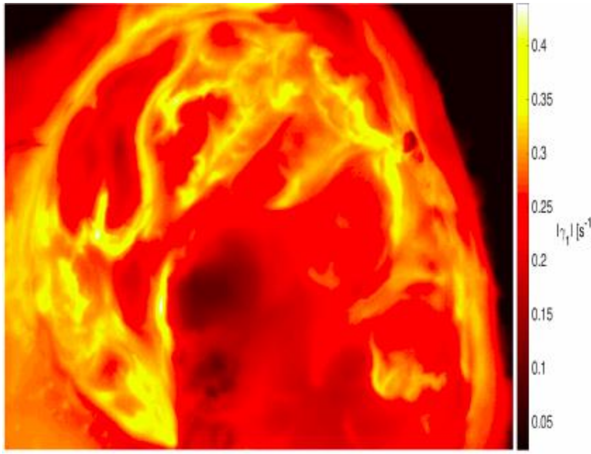

SETTLEMENT STRUCTURES IN TOPOGRAPHIC MAPS

- Application area: environmental sciences and urban modelling
- Challenges:
  - Analysis of maps to trace the development of settlement areas and their internal structure over time

Micro-Level		Meso-Level		Macro-Level
Single building	Building region	surrounding (100m search radius)	urban block	settlement body
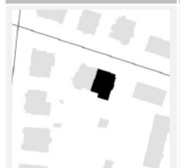	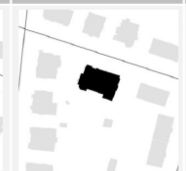		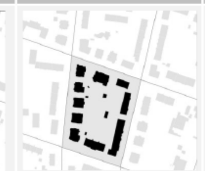	
e.g. Area, Perimeter, Height, Use	e.g. Area, Perimeter, No. of adjacent buildngs	e.g. No. of buildings, mean distance from building to building	e.g. building coverage, ATKIS®-Land Use Information	z. B. area, perimeter, compactness

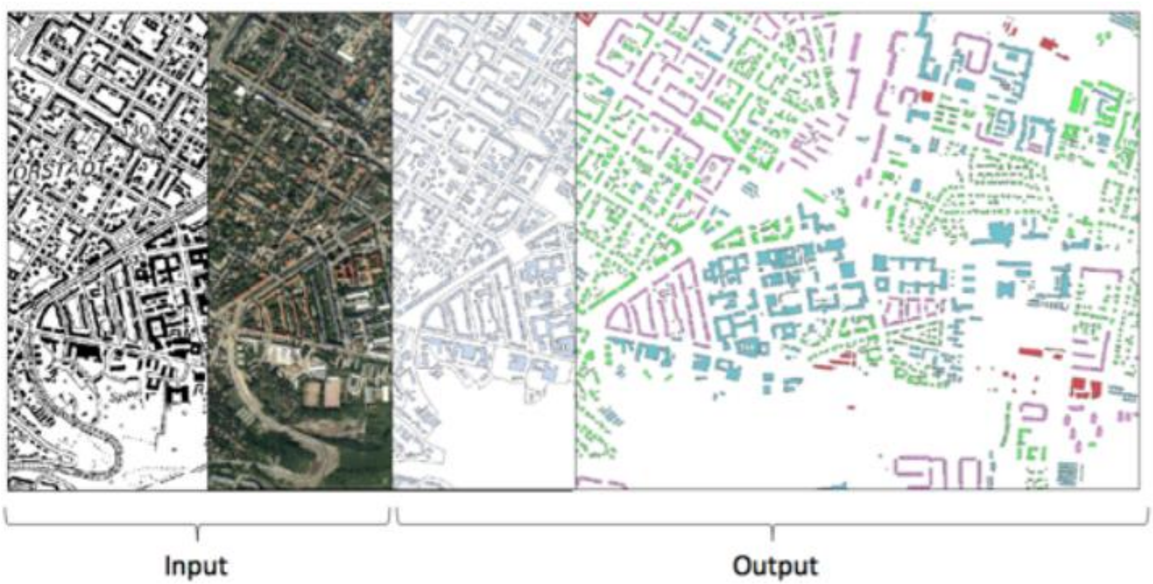

SETTLEMENT STRUCTURES IN TOPOGRAPHIC MAPS

Settlement area detection, Messtischblaetter 1875-1943

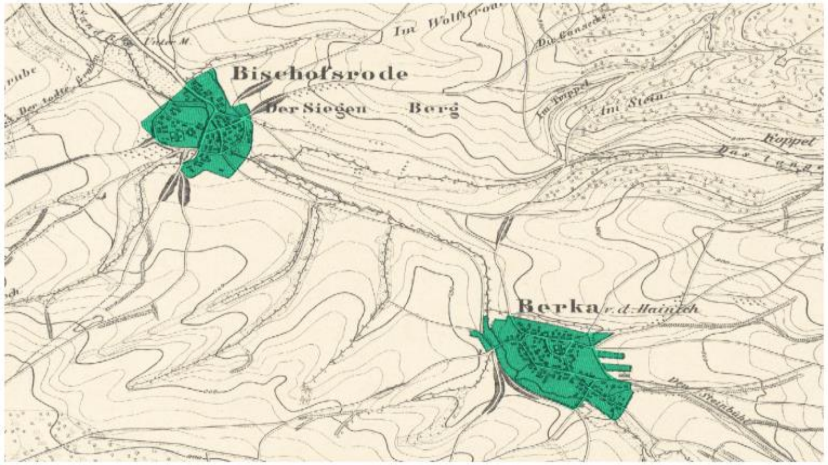

Collaboration between:

Computer Vision Lab Dresden (CVLD)
ZIH, TU-Dresden
IOER

SETTLEMENT STRUCTURES IN TOPOGRAPHIC MAPS

- Application area: environmental sciences and urban modelling
  - Scenario:
    - Analysis of historic maps (“Messtischblätter”): Good coverage of Germany in 1:25000 scale (1875-1945)
    - Thorough evaluation is desired (over time)
    - Accurate training sets required
  - Solution
    - Usage of image segmentation algorithms in data processing
    - Avoid previously required labor intensive manual work

Example settlement areas

SETTLEMENT STRUCTURES IN TOPOGRAPHIC MAPS

- Results:
  - Automatic and new method for settlement detection in historic maps available using Random Forest (RF) + Conditional Random Field (CRF)
  - Scalable data processing of large quantity of input maps possible

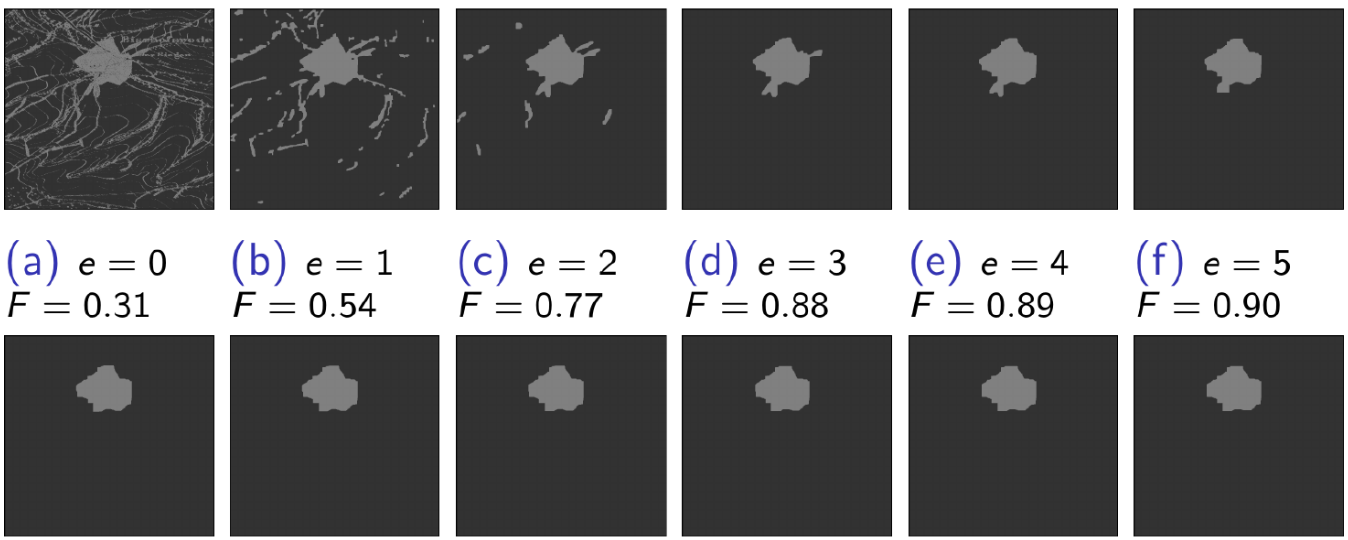

(a) $e = 0$ $F = 0.31$
(b) $e = 1$ $F = 0.54$
(c) $e = 2$ $F = 0.77$
(d) $e = 3$ $F = 0.88$
(e) $e = 4$ $F = 0.89$
(f) $e = 5$ $F = 0.90$
(g) $e = 6$ $F = 0.93$
(h) $e = 7$ $F = 0.93$
(i) $e = 8$ $F = 0.93$
(j) $e = 9$ $F = 0.93$
(k) $e = 10$ $F = 0.93$
(l) $e = 11$ $F = 0.93$

SETTLEMENT STRUCTURES IN TOPOGRAPHIC MAPS

- Results:
  - Automatic and new method for settlement detection in historic maps available
  - Scalable data processing of large quantity of input maps possible
- Runtime improvement:
  - serial processing on ordinary workstation: ~780 minutes (13 hours)
  - Parallel execution: <4min → ~**200**x faster

Input

Correct output labels

MULTI-SCALE VISUALIZATION OF ENGINEERING DATA

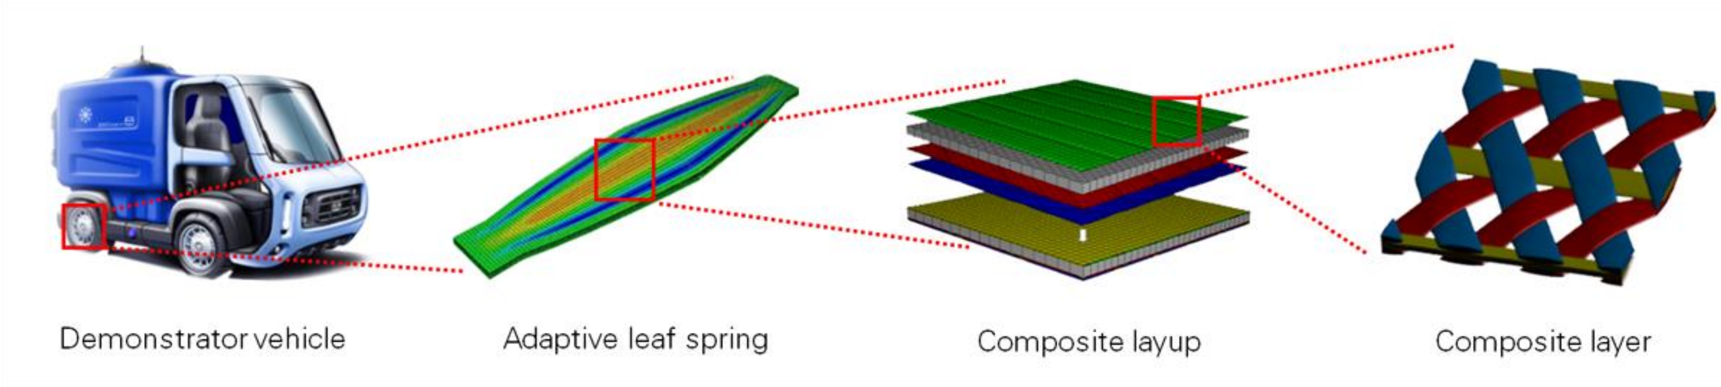

- Follow simulation of processes over time on different scales
- Easy user-interface for direct interaction with data (web-browser based)
- Preparation of simulations on different scales at HPC-site and output presentation via visual analysis

MULTI-SCALE VISUALIZATION OF ENGINEERING DATA

Methods:

- Time-dependent finite element simulation of complete component on 3 scales
- 200 time steps per scale using Abaqus tool; simulations need to be aligned

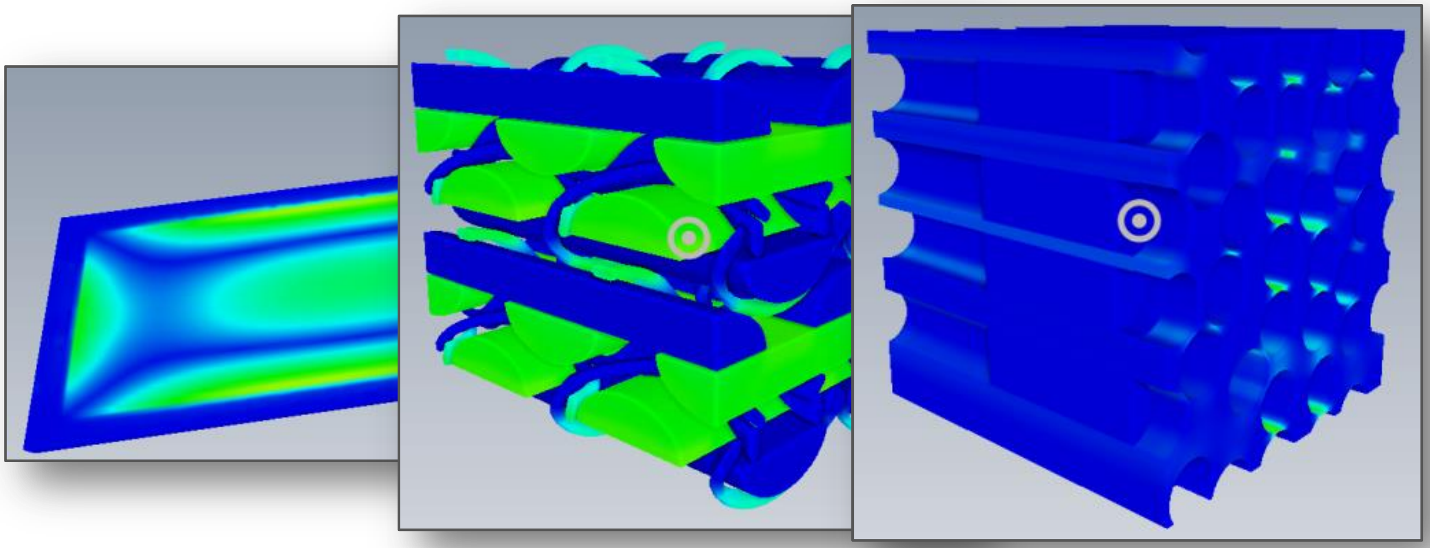

ScaDS
DRESDEN LEIPZIG

Multi-scale visualization – The key for a deeper understanding of materials

TECHNISCHE UNIVERSITÄT DRESDEN

Computer Graphics and Visualization

Institut für Leichtbau und Kunststofftechnik

GEFÖRDERT VOM
Bundesministerium für Bildung und Forschung
DLR Projektträger

SUMMARY AND OUTLOOK

ScaDS DRESDEN LEIPZIG

SUCCESS STORY SCADS DRESDEN/LEIPZIG

Strong scientific output and competence (>120 publications)
i.a. Big Graph Analytics, Sierra Platinum , CTS, data intensive workflows for HPC, settlement recognition in historic maps, Interactive Multi-Scale Visualization...

Service Center for Big Data with high impact
Numerous interdisciplinary big data application projects and industry collaborations & transfer in industry

Many project aquisitions
> 11 Mio Euro (Exploids, BIGGR, TIQ-Graph, KOBRA, MASI, GERDIE, EMUDIG4.0..)

National & international outreach & visibility
200 keynotes/talks worldwide , 3 successful summer schools, 30 proven experts in guest program, 3 successful Big Data in Industry workshops

Successful training and education program
“Big-Data-Schwerpunkt”: lectures/ seminars/ trainings/ PhD seminars
Hundreds of Graduates with Big Data Expertise (Master)
>10 PhDs in Big Data close to finishing

SUMMARY II

- Convergence of HPC and Big Data offers great opportunities in data analysis
- There is no unique big data usage pattern
  - Many different aspects are of interest (not just “volume”)
  - But: transparency for users is very important
- HPC systems will support an extremely large main memory, which will result in huge input/output data (size and/or number of files)
- Other, more distributed approaches still valid, e.g. for Hadoop-like workloads, but more iterative methods needed (machine learning)
- Still depending on use-case requirements – user needs to adopt current workloads

➢ Big Data Analytics at the push of a button … will take a while

SCADS IN REGIONAL/NATIONAL NETWORKS

Opening of the HPC infrastructure for German Big Data Community

Smart Infrastructure Hub

Smart Systems Hub

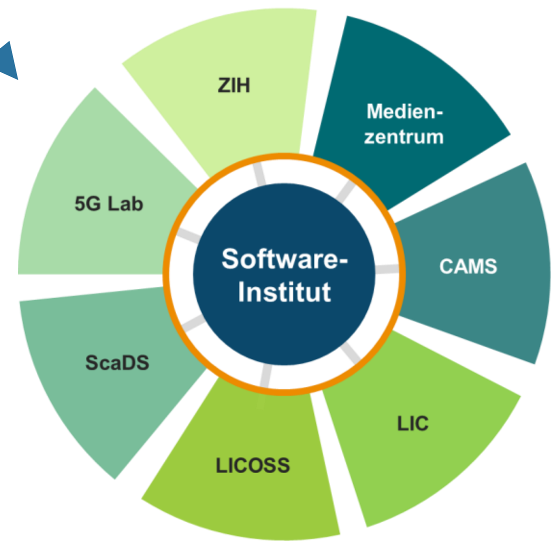

Lehmannzentrum

https://www.staatsregierung.sachsen.de/kompetenz-und-innovationszentrum-smart-systems-hub-4256.html

ScaDS
DRESDEN LEIPZIG

THANK YOU

CONTACT

SCIENTIFIC COORDINATOR

Prof. Dr. Wolfgang E. Nagel

Technische Universität Dresden
Zentrum für Informationsdienste und Hochleistungsrechnen
01062 Dresden

Telefon: +49 351 463-35450
E-Mail: wolfgang.nagel@tu-dresden.de

SCIENTIFIC CO-COORDINATOR

Prof. Dr. Erhard Rahm

Universität Leipzig
Fakultät für Mathematik und Informatik
Augustusplatz 10
04109 Leipzig

Telefon: +49 341 97-32221
E-Mail: rahm@informatik.uni-leipzig.de

MANAGEMENT DIRECTOR

Dr. René Jäkel

Technische Universität Dresden
Zentrum für Informationsdienste und Hochleistungsrechnen
01062 Dresden

Telefon: +49 351 463-42331
E-Mail: rene.jaekel@tu-dresden.de

SERVICE CENTER LEIPZIG

Dr. Eric Peukert
Universität Leipzig

Big Data Zentrum ScaDS Dresden/Leipzig
Ritterstrasse 9-13
04109 Leipzig

Telefon: +49 341 97-39524
E-Mail: peukert@informatik.uni-leipzig.de