What about Big Data and all that fuzz around the buzz?

Big Data Meets Accelerator Controls

Tim Wilksen
Beschleuniger-Betriebsseminar Travemünde, 22.02.2019

Topics

Big Data Meets Accelerator (Controls)

01 Brief excursion
- Ever growing data volume
- (Futile) attempt to define Big Data

02 Zooming in on DESY and its accelerators
- Classic approach of acquiring data in controls
- A different way - FLASH Data Acquisition
- Archives and databases

03 Beyond FLASH
- EuXFEL implementation of the FLASH-Style DAQ
- Some Stats and Experiences
- Shortcomings and such
- A look into the future (of the EuXFEL DAQ)

04 Reaching out
- Projects and Initiatives
- Helmholtz POF IV – MT and DMA

05 Summary
- Conclusions

Brief Excursion


BIG
DATA
Reinforcement Learning
Data Science
Unsupervised Learning
Storage
Data Management
Fault
Analysis
Timeseries Databases
Supervised Learning
Parallel Programming
Accelerator Controls
Meta Databases
Availability
Feedback
System Modeling
Quality Management
Optimization
Data Acquisition
Network
Predictive Maintenance
Machine Learning
Algorithms
Exabytes
Anomaly Detection
Reliability
Streaming
Neural Networks
Near-Realtime Analysis
Fault Tolerance

Ever Growing Data Volume

Business et al.

- **Global Datasphere =**
  cloud and data centers + enterprise infrastructure + IoT, PCs, Phones will reach 175 ZB by 2025 up from 33 ZB in 2018
- Real-time data up to 30% of global datasphere by 2025
- Shipped storage device capacity in the last 20 years was about 7 ZB – this will expand to 22 ZB in the next seven years
- Google: 10 – 15 EB managed data (estimated)
- Amazon AWS/S3: 2 x 10^{12} Objects (2013)
- Facebook: 500 TB/d input, 300 PB stored (2016)

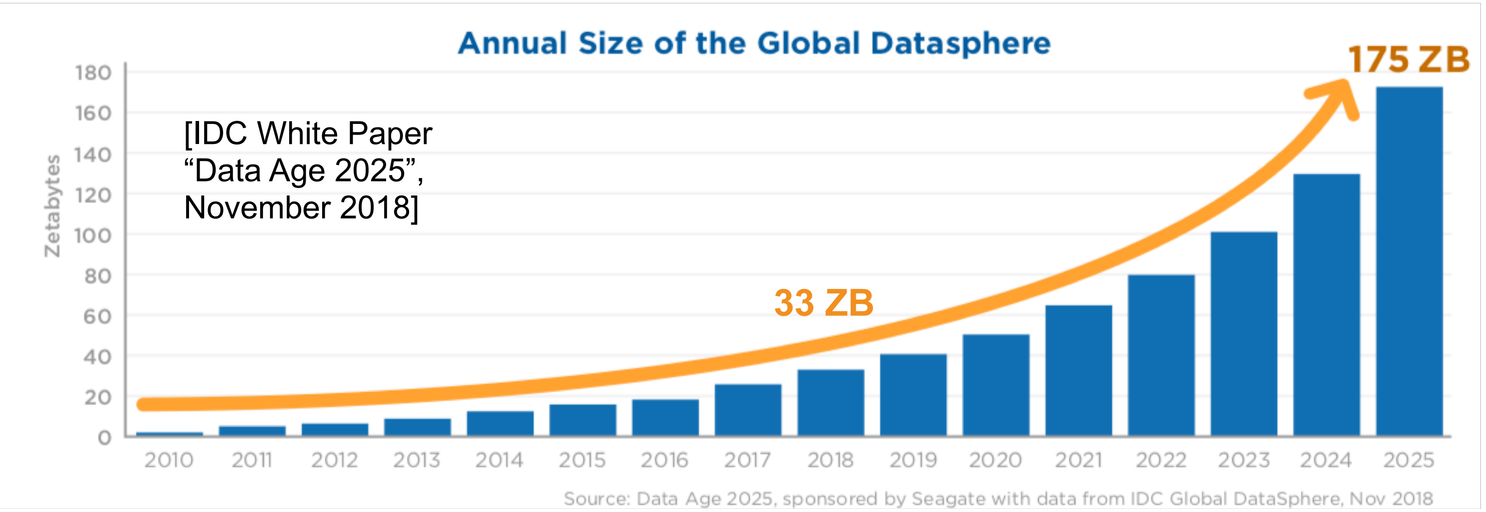

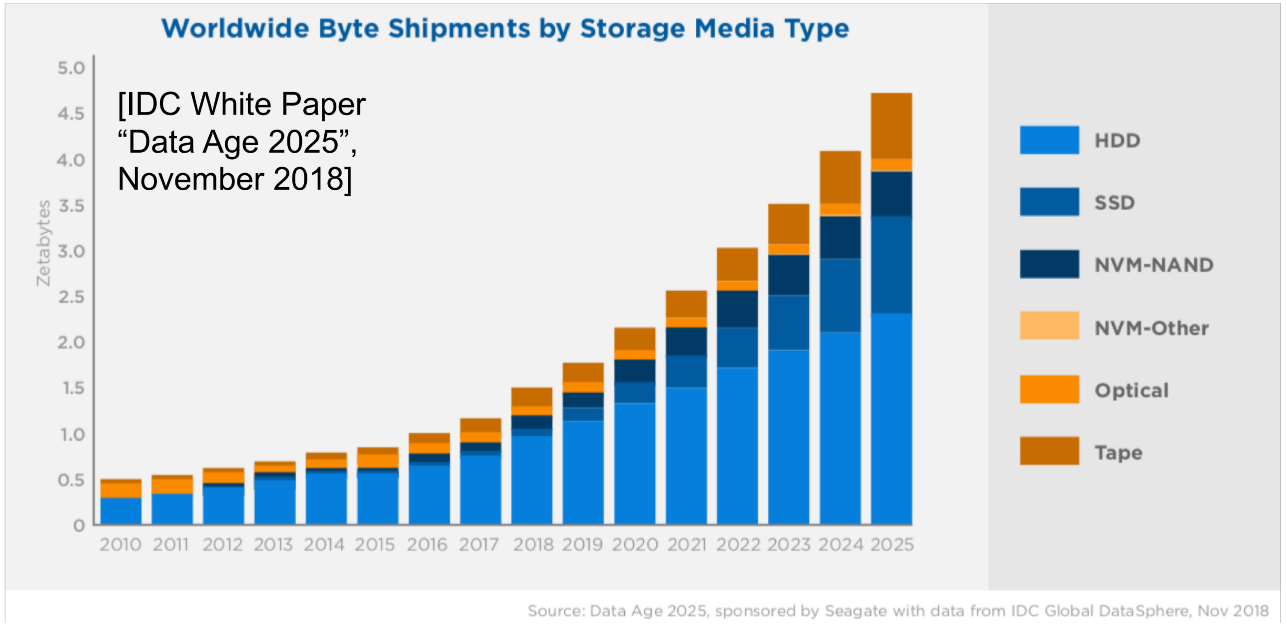

Ever Growing Data Volume (cont'd)

HEP et al.

- LHC HL Upgrade

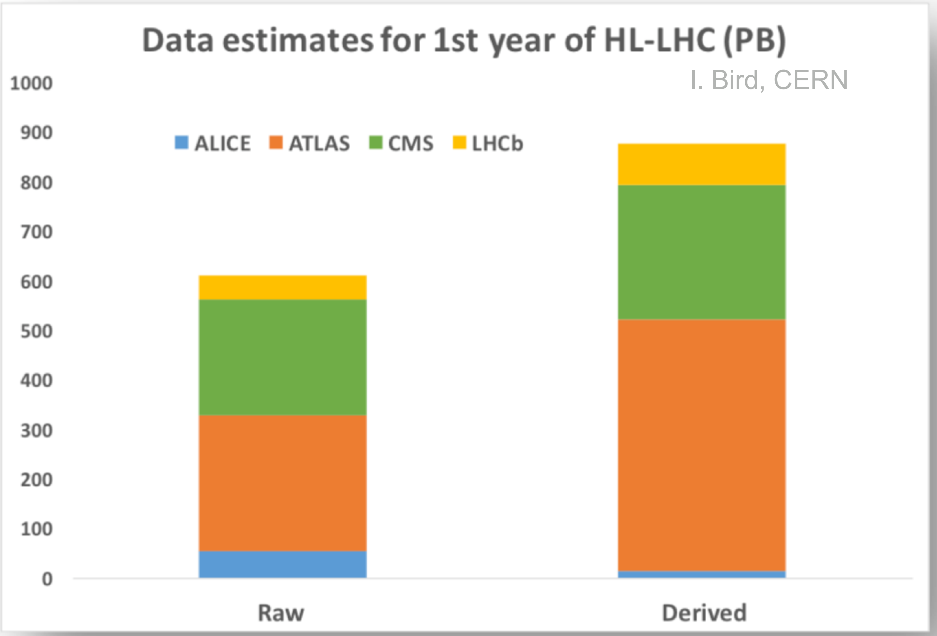

Data:
- Raw 2016: 50 PB → 2027: 600 PB
- Derived (1 copy): 2016: 80 PB → 2027: 900 PB

- SKA – Square Kilometer Array Telescope

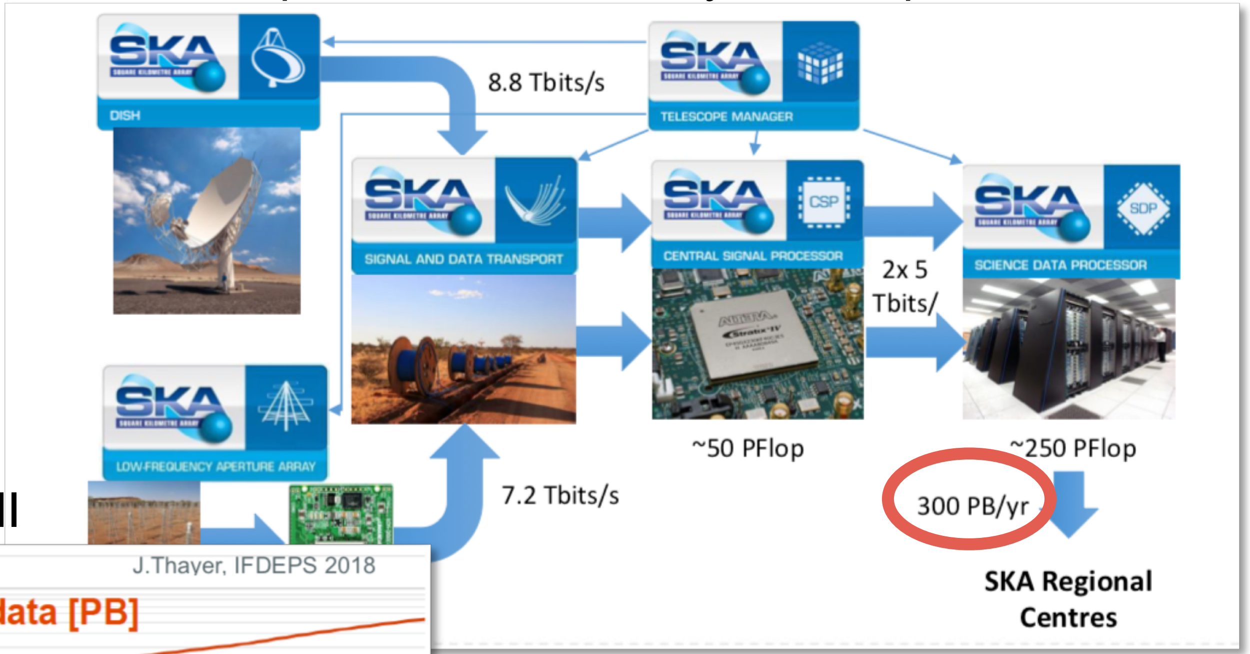

- LCLS II

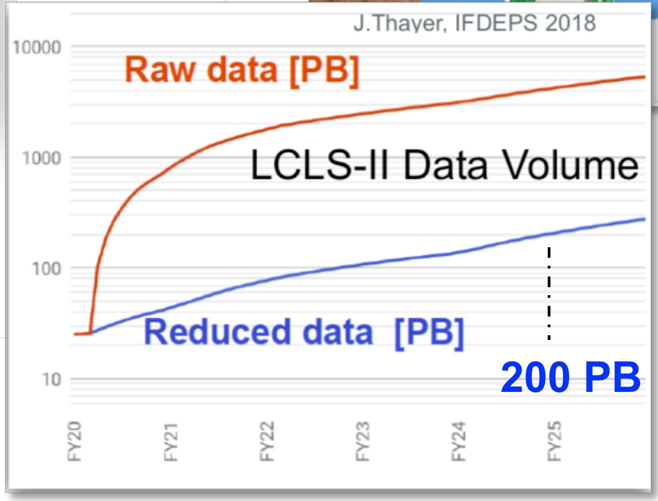

Big Data

A (futile) attempt for a definition

The term "big data" isn't really new, but was popularized in the 90's by SGI chief scientist, J Mashey. [„The Origins of 'Big Data': An Etymological Detective Story", S. Lohr, NYTimes, 02/01/2013]

First stop, Wikipedia:
"Big data refers to **data sets** that are **too large or complex for traditional data-processing application software** to adequately deal with. [...] Big data challenges include capturing data, data storage, data analysis, search, sharing, transfer, visualization, querying, updating, information privacy and data source."

Gartner, 2011: "Big data is **high-volume, high-velocity and/or high-variety** information assets that demand cost-effective, innovative forms of information processing that enable enhanced insight, decision making, and process automation."

The four* V's:

- Quantity and size of acquired and stored data
- Speed at which the data is being generated and processed

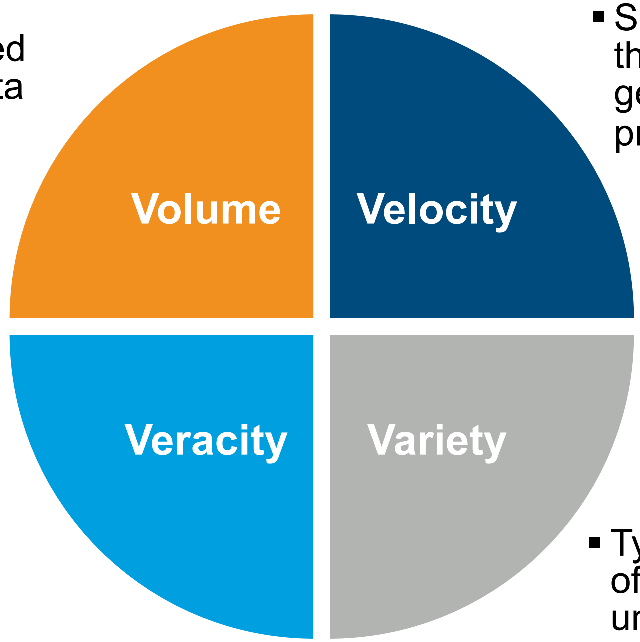

- Data quality, uncertainty and integrity
- Type and nature of data; unstructured and structured data

* IBM added 2011 **Veracity** to Gartners definition. Sometimes a fifth V, **Value** is used in the context of business intelligence

Big Data life cycle

- Ingest data
- Persist data in storage
- Compute and analyze
- Visualize

Big Data (cont'd)

A (futile) attempt for a definition

Gartner's Hype Cycle 2018:

So there is Deep Learning, Artificial General Intelligence, AI PaaS, Knowledge Graphs, Digital Twin(!), IoT Platform, asf.

But…

"Big Data" term is gone on last year's hype cycle:
It is rather a set of technologies and methods now than an emerging technology per se.

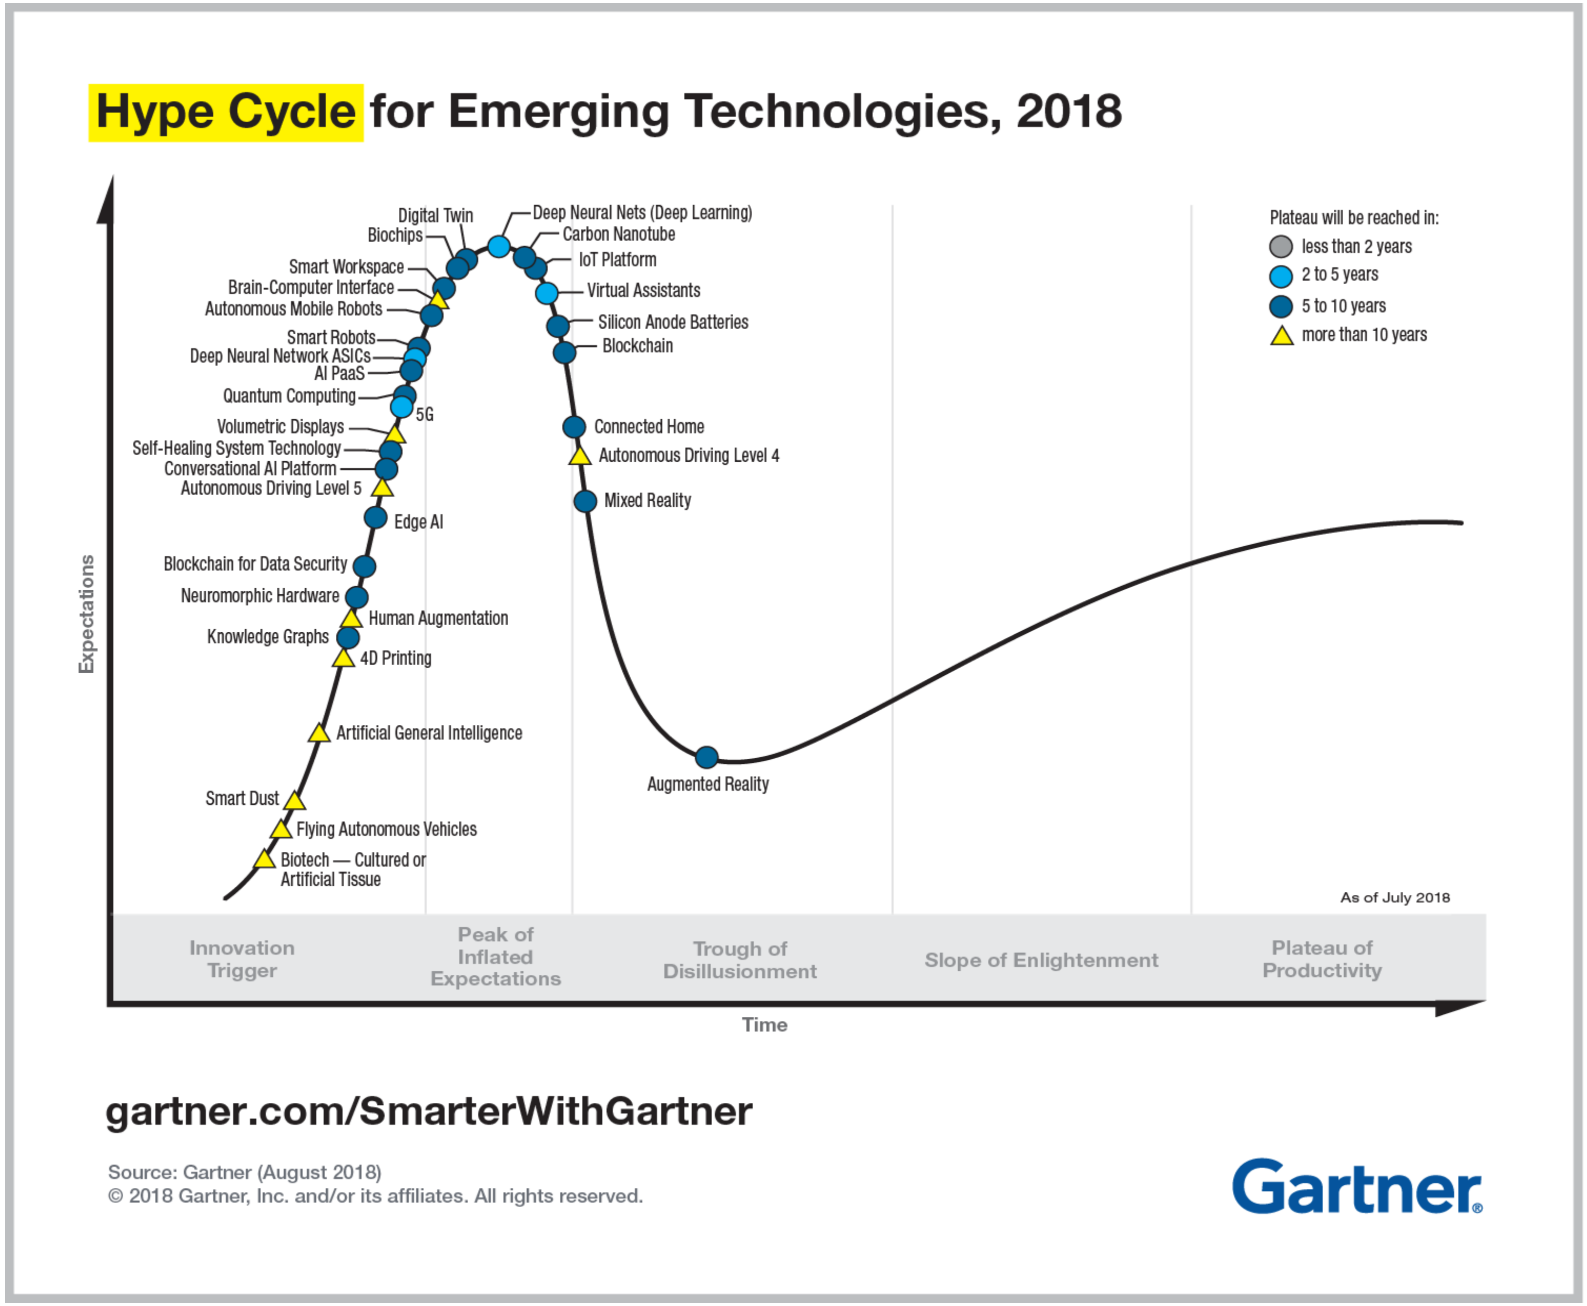

Big Data (cont’d)

A (futile) attempt for a definition

Now, what about particle accelerators and “Big (Science) Data”?

Zooming in on DESY and its accelerators

Accelerator Controls Data*

Let's do a step back first …

- Monitoring and recording data from accelerator and experiments is essential for commissioning, operations, performance analysis and any kind of R&D work.
- Implementations often seen, can be categorized as follows:

a) The "classic approach"

- Clients collect data from any device server or any other kind of data source.
- Client can be rich client application, another server, middle layer server or a graphical interface (operator)
- Flat channel-based data model – not necessarily accelerator-based
- Archiving, logging facility or database.

b) The "subsystem approach"

- Use standalone acquisition for individual subsystems (e.g. turn-by-turn BPM DAQ, OTR image acquisition), specifically for fast systems and bulk data.
- Often ommercial packages and tools for standalone acquisition
- Potentially complex data model (structures)
- File-based storage

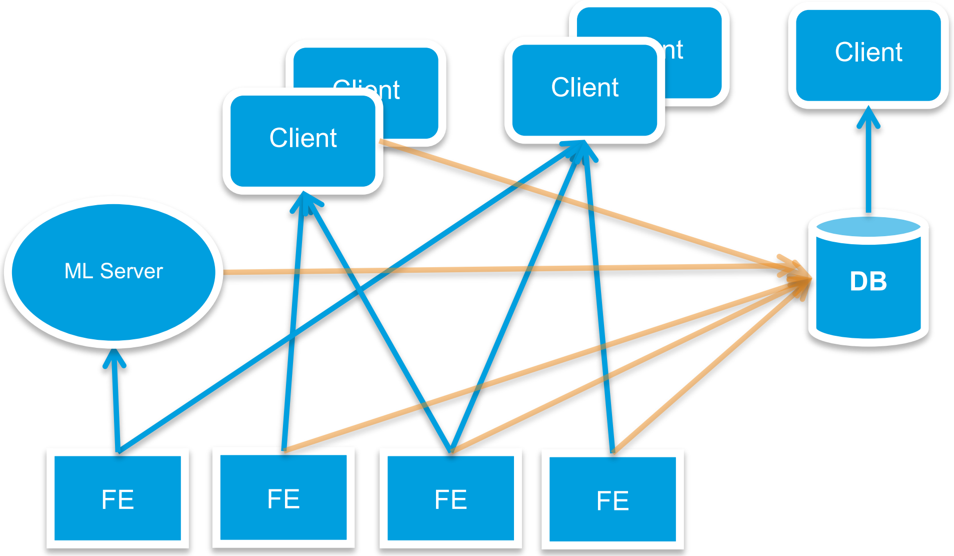

Works, but synchronization and efficiency can be an issue (a), or combination of too many different data formats and sources (b).
⟹ Revisit control system model and data acquisition approach …

* not considering accelerator simulations and its data here

FLASH Data Acquisition

Integrate DAQ into control system middle layer

- Derived conceptually from a high-energy physics experiment DAQ
- Data has event-like structure **recorded at shot repetition rate** (10 Hz)
- Data tagged with shot/train ID allows for easy **synchronization**
- Data from fast, triggered sources sync'd together with slow information
- **Event record** holds all **train related and bunch-resolved data**
- Central **shared memory** keeps brief history of shot/pulse data
- Middle layer applications perform near **real-time/online analysis**
- Results can be inserted into event records in addition to other data
- RAID cluster for intermediate storage w/ proprietary file format
- C++-, Java-, MATLAB-API for data extraction and batch processing; Java-based data browser and extraction tool
- Run-based configuration and bookkeeping of data

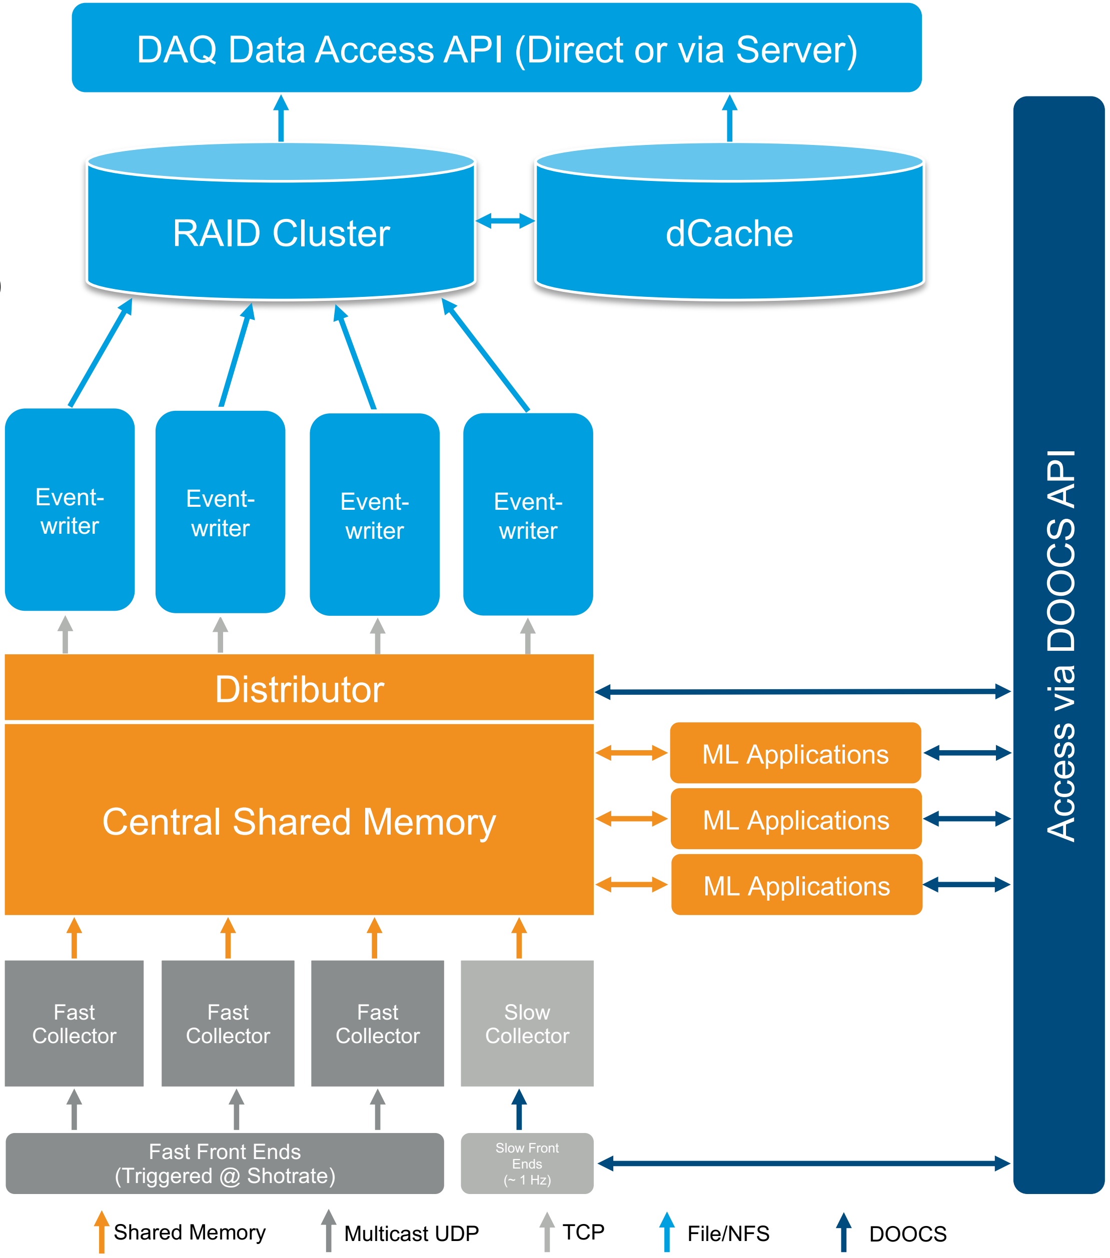

Archives and Databases

Local, central, DB or not to DB

- DESY accelerator control systems utilize full range of archiving types:
  - **Local archives on disk** with server applications – data is accessible through server API itself
  - **Central archives** (file-based, RDB) with data being collected from various data sources, retrievable via dedicated tool
  - **Configuration databases** (files or RDB) used for configuration of machine or controls devices - available through dedicated API
  - Fast and bulk data from FLASH-type **data acquisition** – proprietary file format, available from disk or dCache (tape) through Java-based tools (no batch mode) or C++/MATLAB API (batch mode)
  - Other data acquisition like tools, e.g. MATLAB, Python, SDDS usually stored as flat files on disk

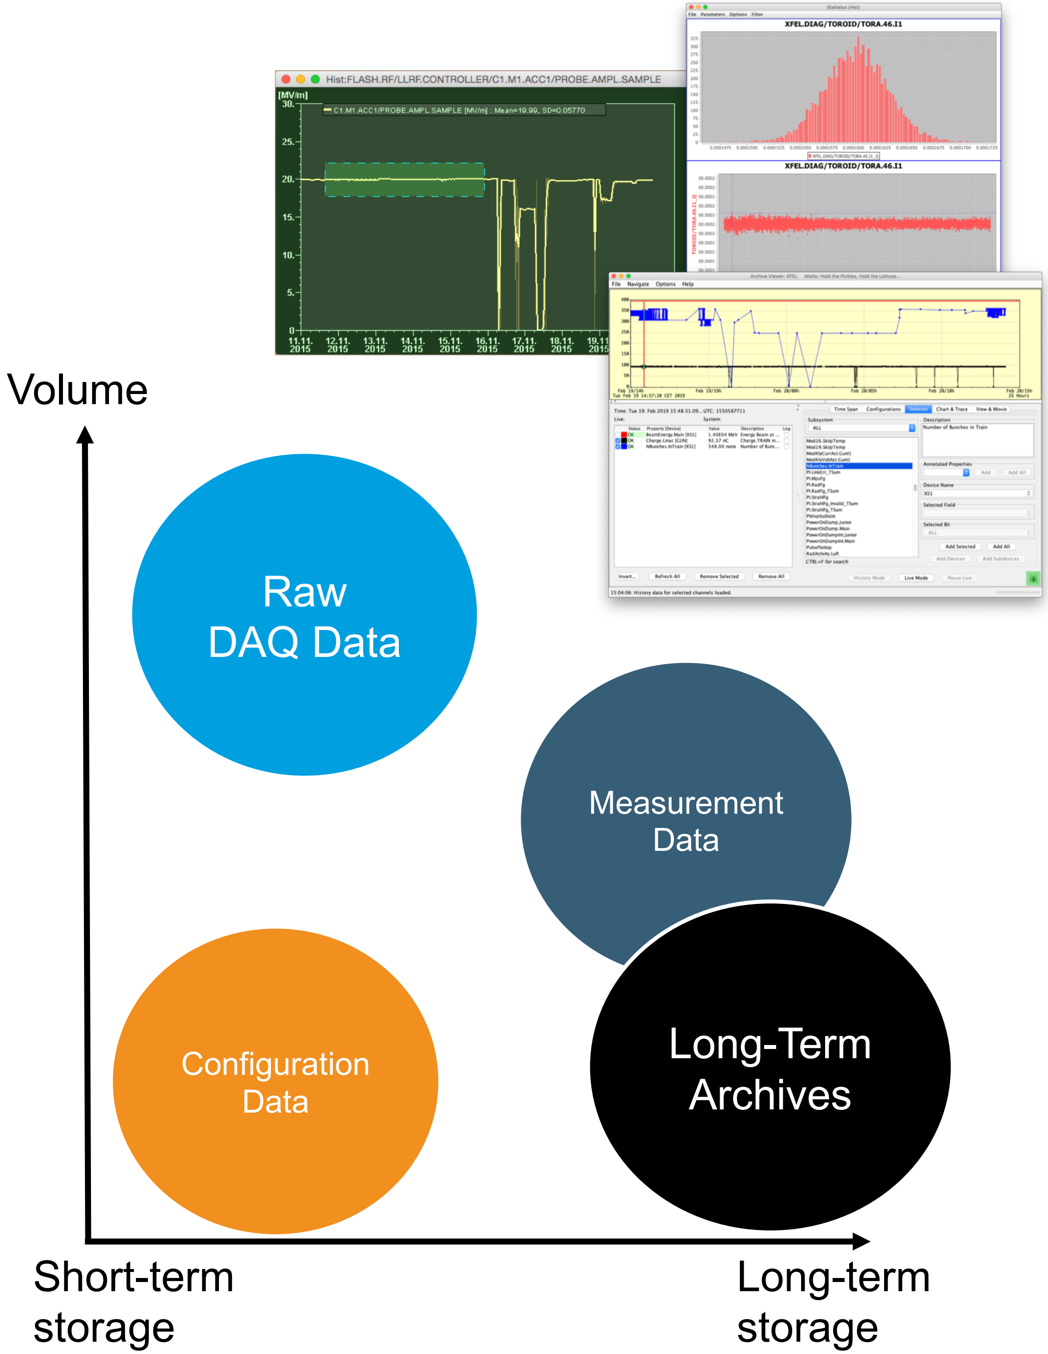

FLASH Data Acquisition …

… and Archives

What about large-scale accelerators, like the EuXFEL?

Beyond FLASH

EuXFEL DAQ

Implementation of FLASH DAQ at EuXFEL

- Same concept, just larger …
- Five instances plus test one
  - Diagnostics DAQ
  - LLRF DAQ x 2
  - Virtual XFEL instance
  - Laser System DAQ
- Multiple input streams
- Mission-critical ML services
- Clustered disk pool for data
- dCache for dedicated storage

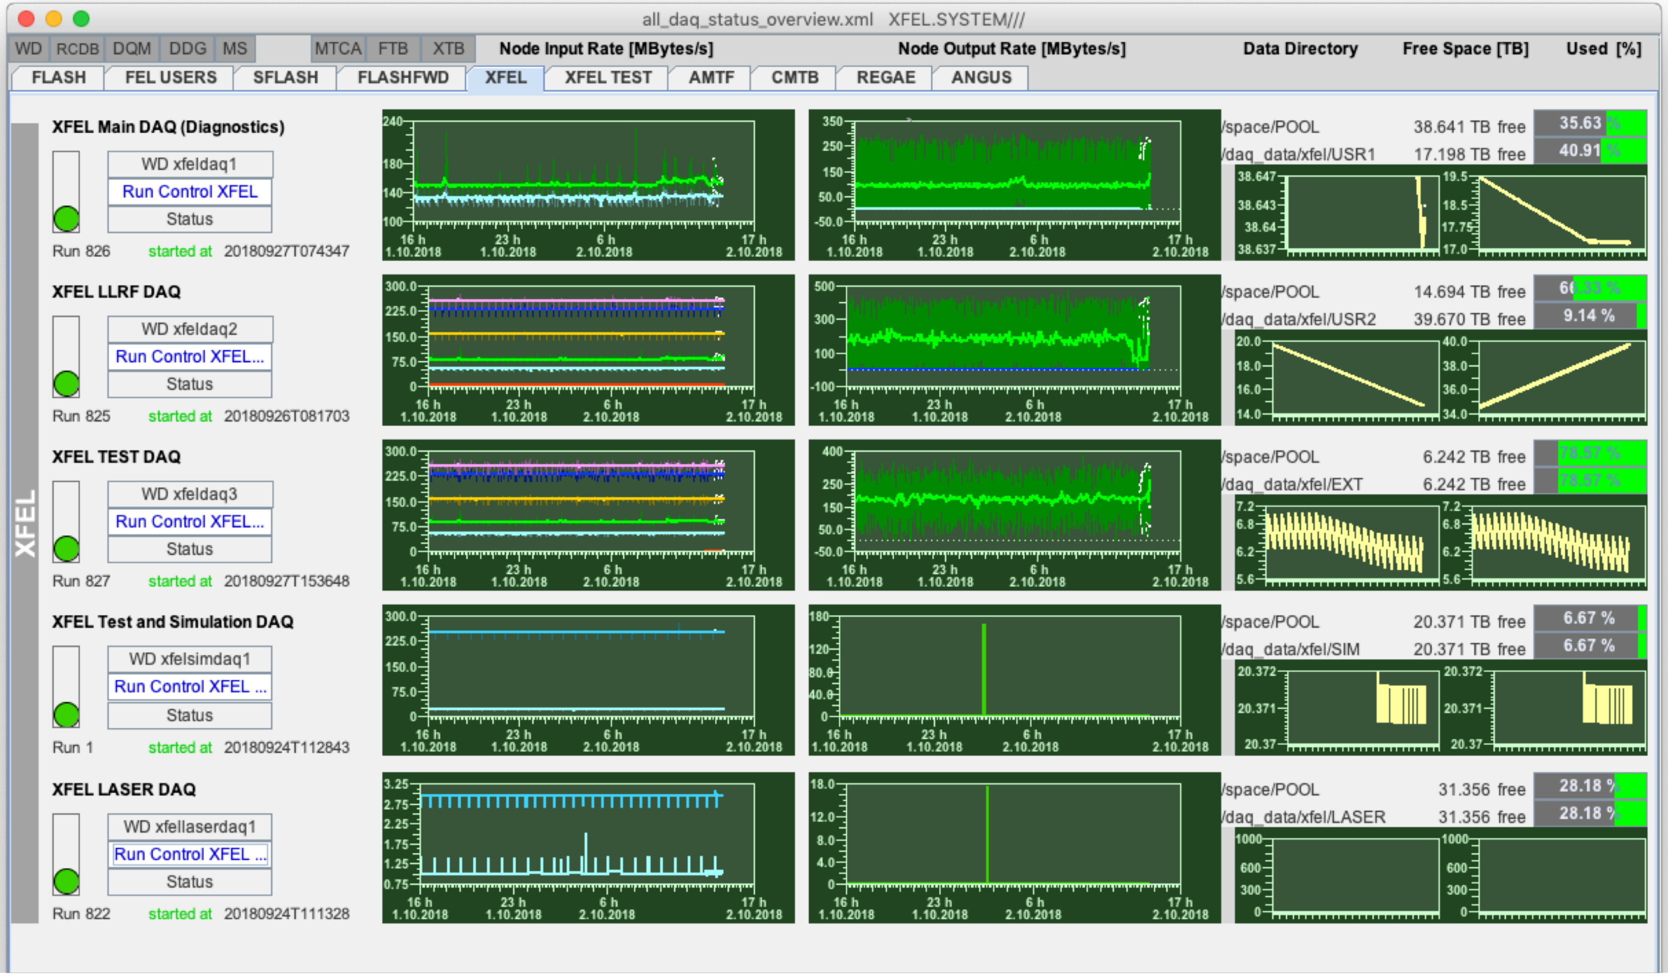

EuXFEL DAQ (cont'd)

Some Stats and Experiences

- More than **10 million** DOOCS control system in accelerator namespace
- About **700.000 local** DOOCS **archives** plus large TINE **central archive**
- Timing system provides synchronization for data acquisition
- About 250 MicroTCA systems online as DAQ FE sender with 20000 channels
- About **1.5 GB/s input rate** from FE producing 130 TB/d uncompressed data or up to **30 TB/d compressed data**

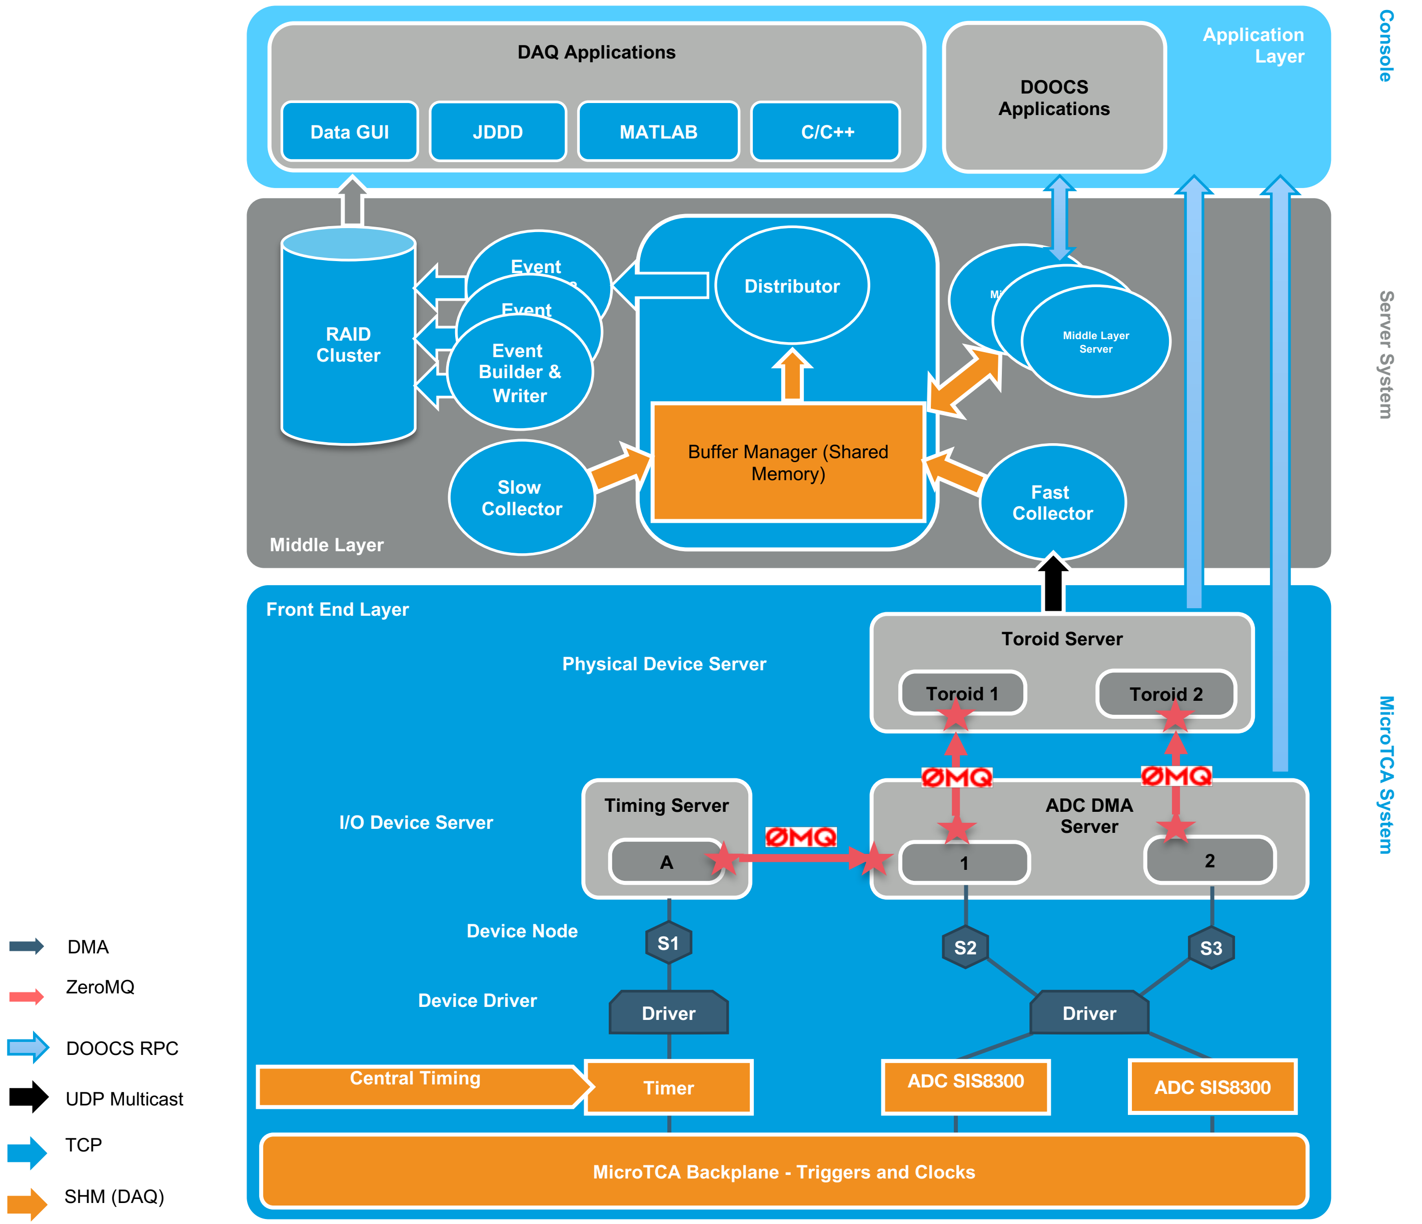

EuXFEL DAQ (cont'd)

Shortcomings and such

- A rate of 30 TB/d would result in roughly 9 PB/y …
- … we can't keep the data! Diagnostics data for 2 weeks online, LLRF data a couple of days …
  - Requires dedicates storage system O(PB) preferably on IT premises with easy access on dedicated computing nodes
- No dedicated Ethernet infrastructure for DAQ
  - Interference with controls and user services
- Clumsy interface for accessing and extracting offline data
  - Historically the more convenient interfaces were lacking performance, re-attempt to utilize HDF5 format
- Quite static configuration capabilities. Not possible to add data sources on-the-fly, requires restart.
- Run configuration and file management is only half the story, missing meta database with search capability
- Unused potential: snapshot, post-mortem analysis, predictive analytics, preventive maintenance, trend and fault analysis, anomaly detection …

→ Revisit and re-evaluate concepts

PETRA IV

CW-EuXFEL

Revised Archiving Concepts

Future archiving and data acquisition

- Huge amount of data to come from **novel diagnostics** (i.e. current CRISP and KALYPSO detectors)
- Increasingly complex and intelligent front-end electronics enables low-level processing capabilities
- **Performance**, **reliability and availability** are key parameters for machine operation

  ➡ significantly enlarge information collection for monitoring and analytical purposes:
  - Configuration management of hardware (electronics, utilities, …) and firm-/software
  - Machine configuration management
  - Performance tracking
  - Fault detection and post-mortem analysis
  - Predictive analytics and preventive maintenance
  - Automation for reproducibility
- Don’t forget to categorize and classify the data, **data reduction** is crucial with this high volumes!
- Scalability of data collection, processing and archiving
- Easiness of accessing and using online and archived data

➡ Leverage “Big Data” technologies and methods

Revised DAQ Concepts (I)

MicroTCA-based DAQ

- High rate (PETRA IV) and continuous wave (CW-XFEL) operation:
  - → Requires capable CPU AMC (memory, processing power, throughput) -> MCH bandwidth
  - → Requires faster Ethernet - CPU/MCH connection via >100GB Ethernet
  - → More CPU power per system - distribute AMC readout among several CPU modules
  - → More powerful FPGA - shifts applications from CPU to FPGA for data handling and processing

➡ MicroTCA-based DAQ

- Use additional layer for acquiring data set up with high-performance MicroTCA-based systems and CPUs
- Requires upgraded MicroTCA standard covering PCIe 5th gen and beyond
- More flexibility in directing, configuring and distributing data flow to server farm
- Allows for intermediate, online data access of data subsets

Revised DAQ Concepts (II)

Streaming Architecture i.e. …

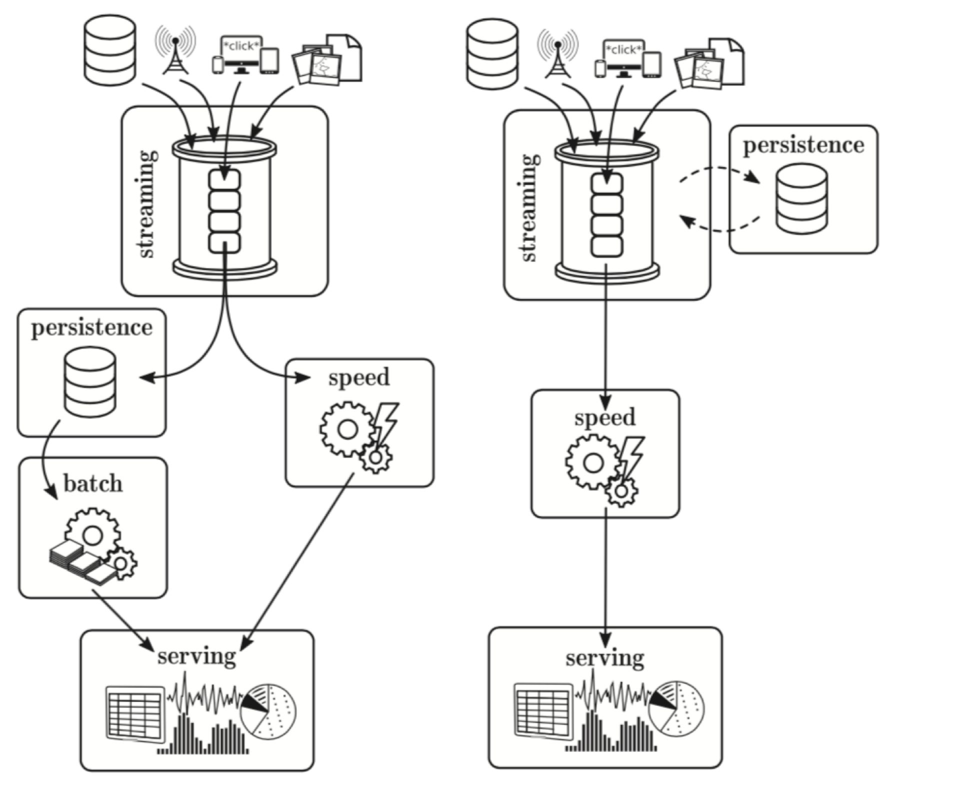

(a) The Lambda Architecture. (b) The Kappa Architecture.

high throughput ⟷ low latency

Wolfram Wingerath*, Felix Gessert, Steffen Friedrich, and Norbert Ritter
Real-time stream processing for Big Data

Utilize "Big Data" technologies:

- **Classification of data** crucial to get out valuable information
- **Streaming architecture** optimized for low latency or high throughput
- **Consolidate data archiving** concepts - time series databases and in-memory databases are prominent choices here; suitability has to be evaluated
- **Near-online analysis** capabilities
- Control-system independent frameworks enable usage across accelerator and experiment facilities
- Several approaches already made, in the context of high-energy physics experiments, accelerator controls (CERN), photon science experiments

⇨ collaborate and share

Revised DAQ Concepts

“Big Data” and such

We might not be alone with this …

Reaching out

Reaching Out

More Projects on Big Data, Machine Learning, DAQ, …

- Machine Learning Workshop
  - SLAC 2018
  - PSI 2019

- **Amalea** – Helmholtz Innovation Pool Project (2019/2020) - DESY, HZB, HZDR, KIT
  - ML application to projects in HEP, photon science and accelerators (I. Agapov)
- Helmholtz OCPC (Office of China Postdoc Council)
  - FS-FLASH proposal on Deep Reinforced Learning using machine and experiment data at FLASH (K. Tiedtke et al.)
- **DASHH** - “Data Science in Hamburg – Helmholtz Graduate School for the Structure of Matter“ – Uni HH, TUHH, DESY + 5
- **CDCS** - “Center for Data and Computing Science“
- **Helmholtz-Inkubator** Information & Data Science
  - Think-tank, interdisciplinary and innovative pilot projects e.g. Helmholtz Analytics Framework

Data Management and Analysis - DMA

Matter and Technology Topic in POF IV

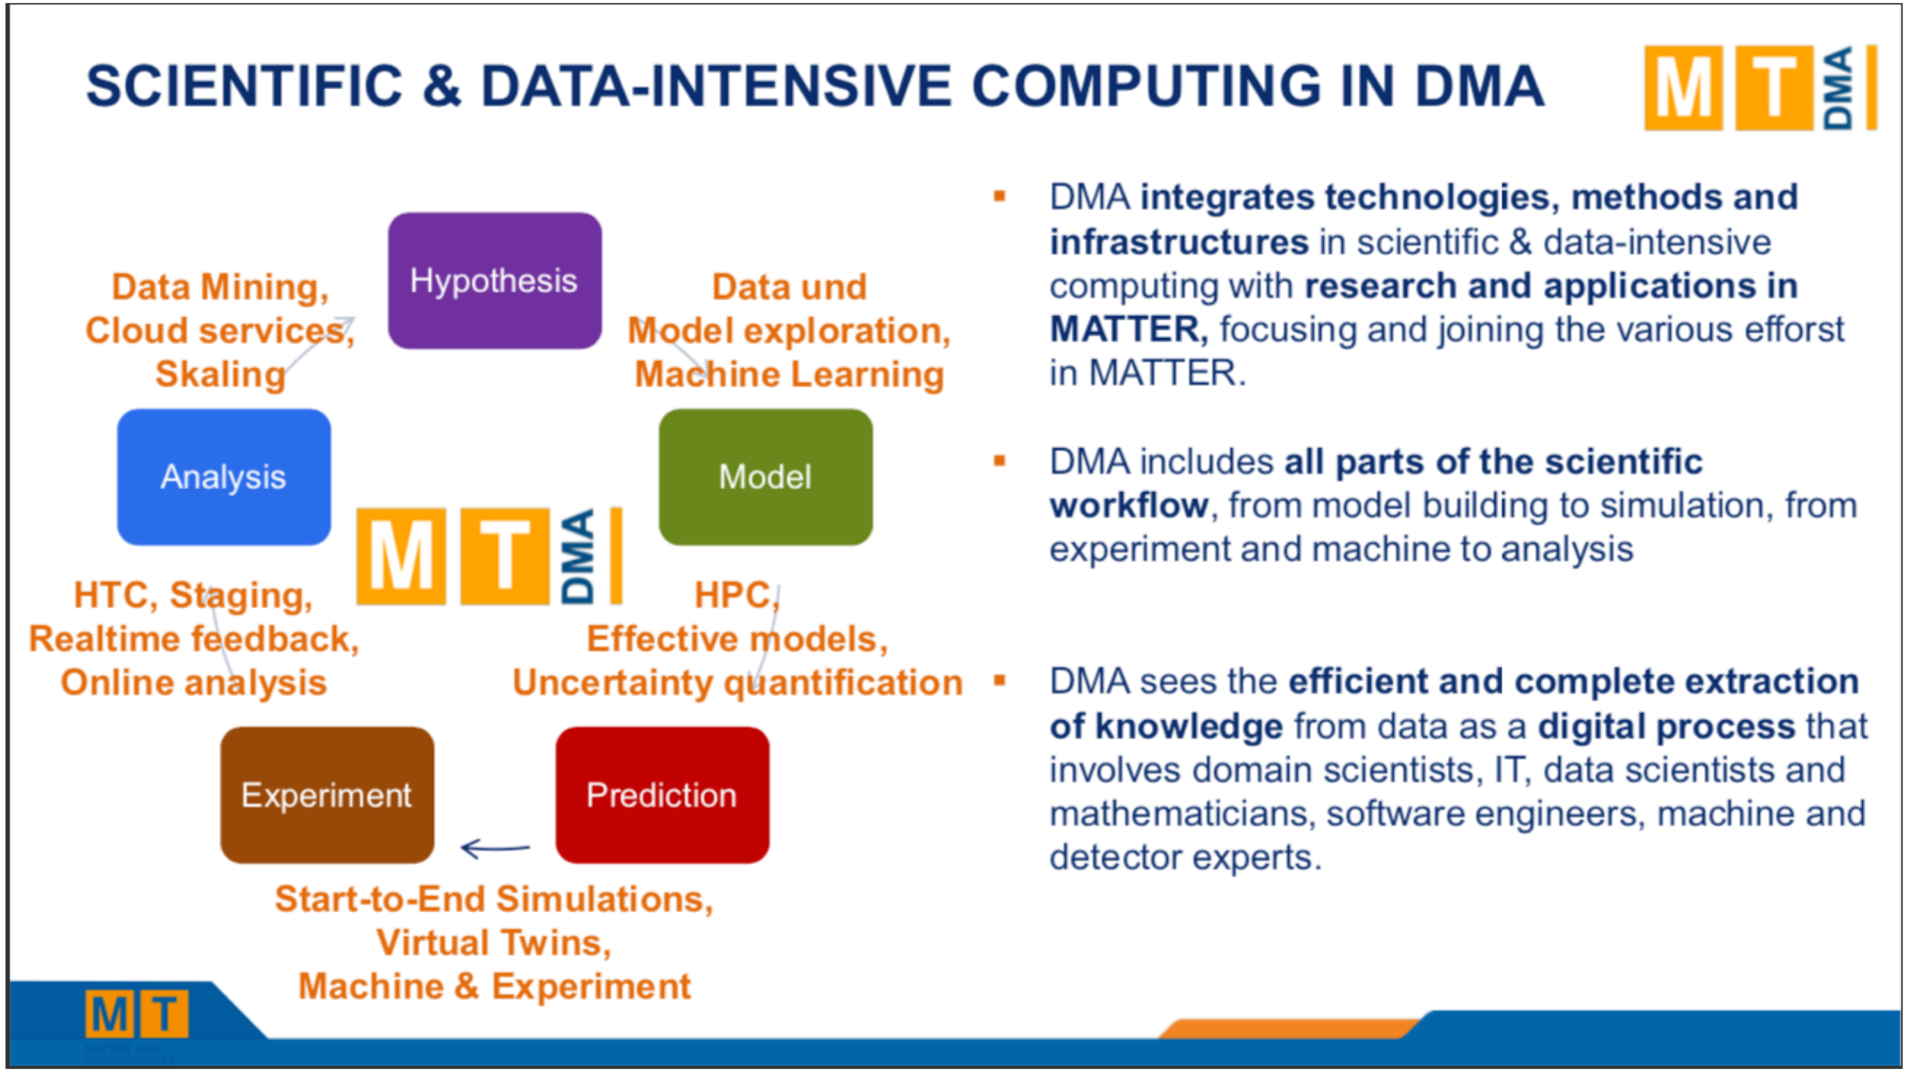

M. Bussmann, V. Gülzow

Data Management and Analysis - DMA

DMA STRUCTURE

ST1: The Matter Information Fabric

- IT infrastructure (Hard- & Software) for facilities
- Automization of Data Lifecycle Management (LK II)
- Solutions für Communities

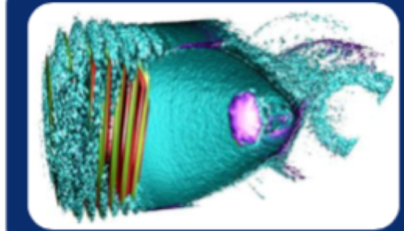

ST2: The Digital Scientific Method

- Matter-specific research in Data Analysis & Simulation methods
- e.g. Machine Learning, Simulation, Visual Analytics, Scientific Workflow
- Developing methods für heterogeneous HPC, HTC, I/O for Matter applications

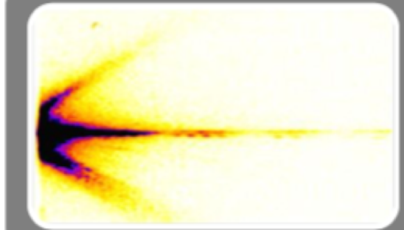

ST3: The Digital Experiment and Machine

- Start-to-End Simulations (Machine/Interaction/Detectors)
- Fast feedback & machine control („Human in the Loop“)
- Quantifying data quality, meta data acquisition & analysis

M. Bussmann, V. Gülzow

Conclusions

Conclusions

… and Outlook

- "Big data" has matured, yet accelerator (controls) community has to catch up with industries and business
- Necessity to **re-think accelerator archiving concepts** in the view of PETRA IV and EuXFEL (CW)
- **Huge potential in accelerator controls data** within the context of performance, reliability and availability
  - Fault prediction and detection, anomaly detection, health monitoring, performance analysis, …
- Major functionality of a modern, large-scale accelerator control and data acquisition system has to be implemented in **software utilizing massive FPGA, GPU and server processing** power – huge effort with respect to required resources
- Upcoming projects and collaborations with other Helmholtz centers and international labs
- Future large-scale accelerator control systems present a challenging yet exciting area for research and development