Computing for FAIR

Thorsten Kollegger

Facility for Antiproton and Ion Research
& GSI Helmholtzzentrum

FAIR Project

New International Accelerator Facility in Darmstadt, Germany

- **Upgraded accelerators from GSI as injectors**

International Participation in FAIR

- FAIR governed by international convention
  - 9 shareholders + 1 assoc. partner (orange)
  - Germany + Russia largest shareholders
- Scientists from all over the world are engaged
  - More than 200 institutions from 53 countries are involved with their scientists (orange + blue)

FAIR Project

- Call for tenders of civil construction of „North Area“ (SIS100) published in Q3+Q4 2016
- Installation and commisioning of Experiments in 2021-2024
- Full completion of FAIR MSV in 2025

FAIR Scientific Pillars

APP

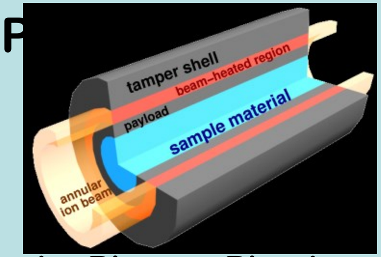

Atomic, Plasma Physics and Applications

CBM

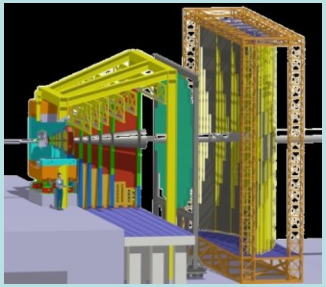

Compressed Baryonic Matter

PAND

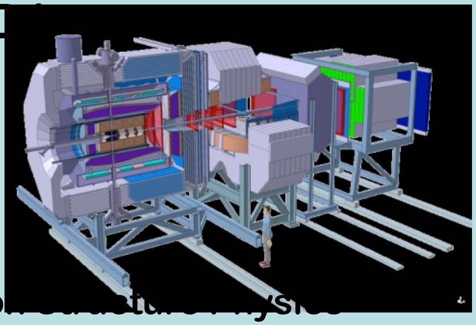

Hadro

NU

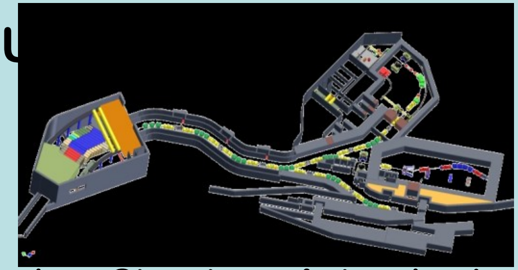

Nuclear Structure, Astrophysics and Reactions

Computing at FAIR: Big Data

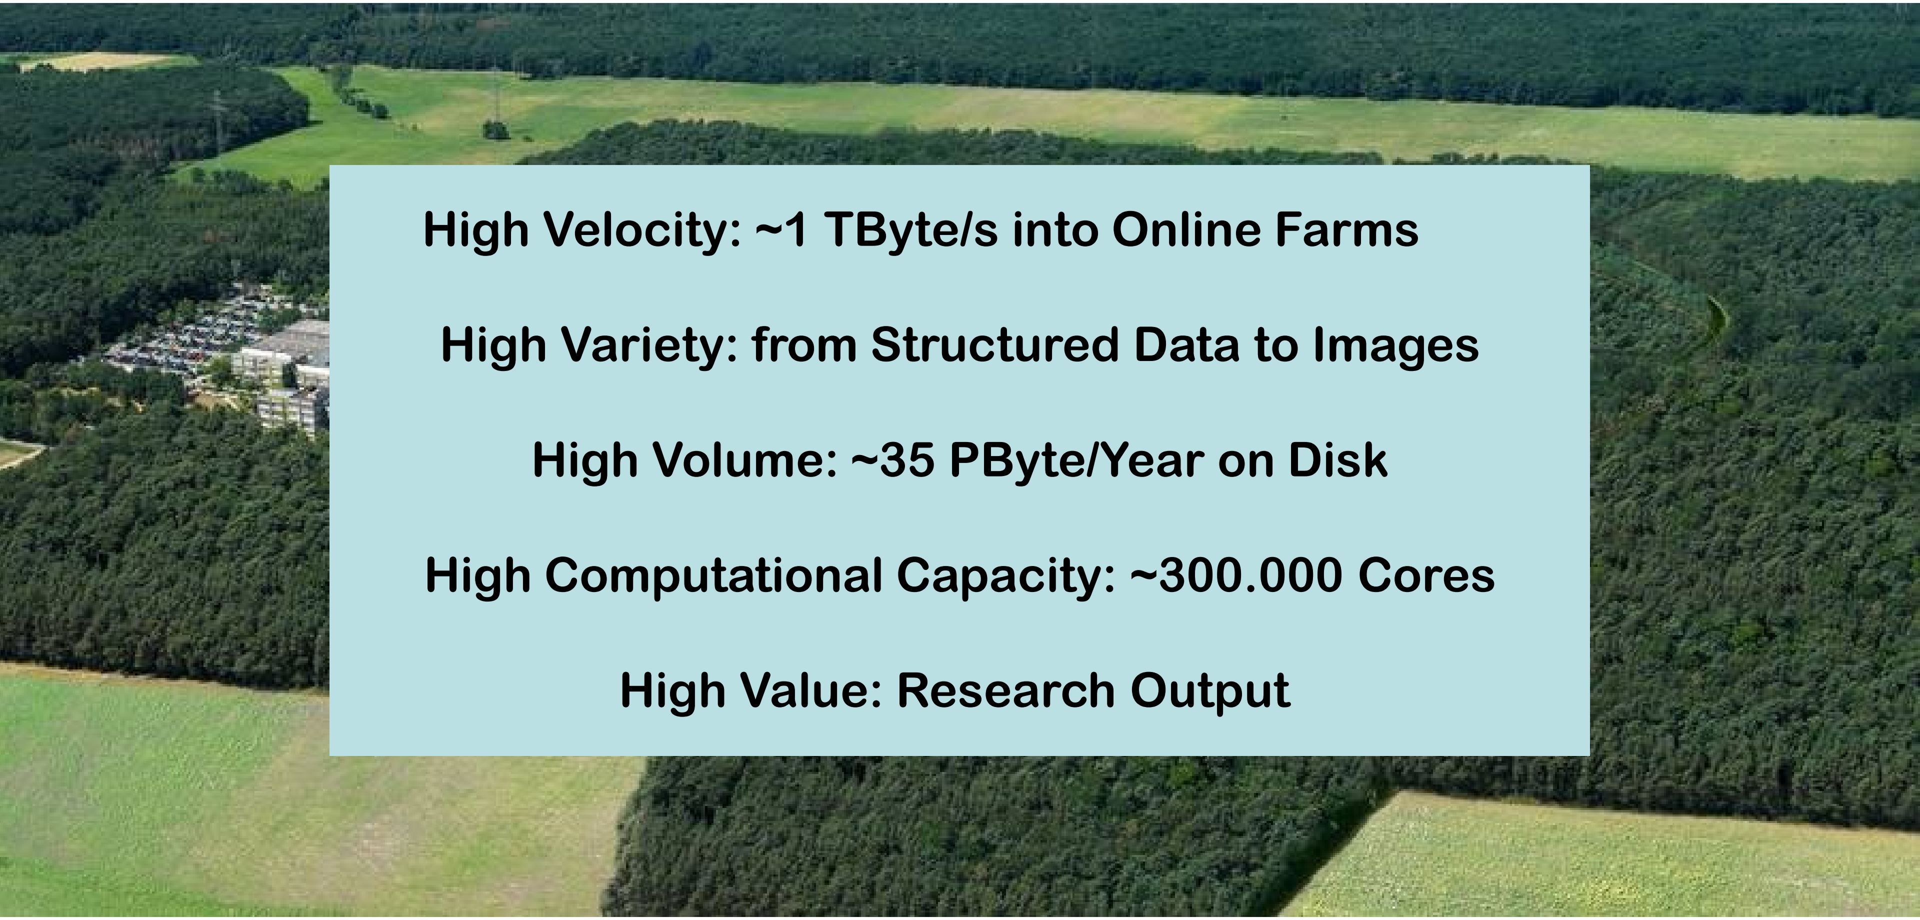

High Velocity: ~1 TByte/s into Online Farms

High Variety: from Structured Data to Images

High Volume: ~35 PByte/Year on Disk

High Computational Capacity: ~300.000 Cores

High Value: Research Output

Computing at FAIR

„Data is inherently dumb.
Algorithms are where the real values lies“
Peter Sondergaard, Gartner, 2015

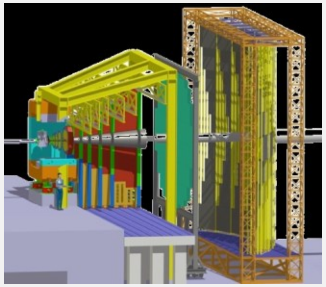

More complex algorithms necessary to exploit the full scientific potential of the instrumentation (i.e. accelerators, detectors)

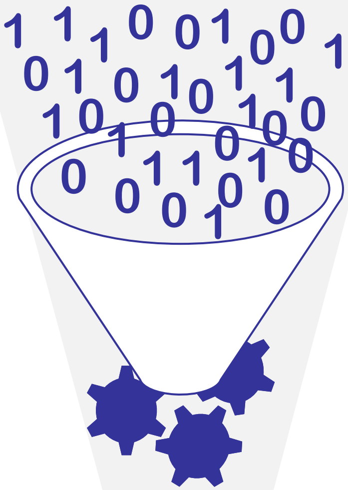

Fundamental change in detector design
From custom electronic to commodity computing systems (“triggerless”)

- **Software defined -> Agility**
- **Huge data rates into (quasi-)real time analysis**
- **FAIR, LHC upgrades, Nuclear Physics …**

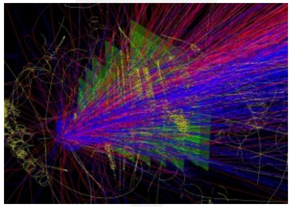

Computing at FAIR

Data-Processing Libraries

Workload Management

International Partners

Software Frameworks

Clouds

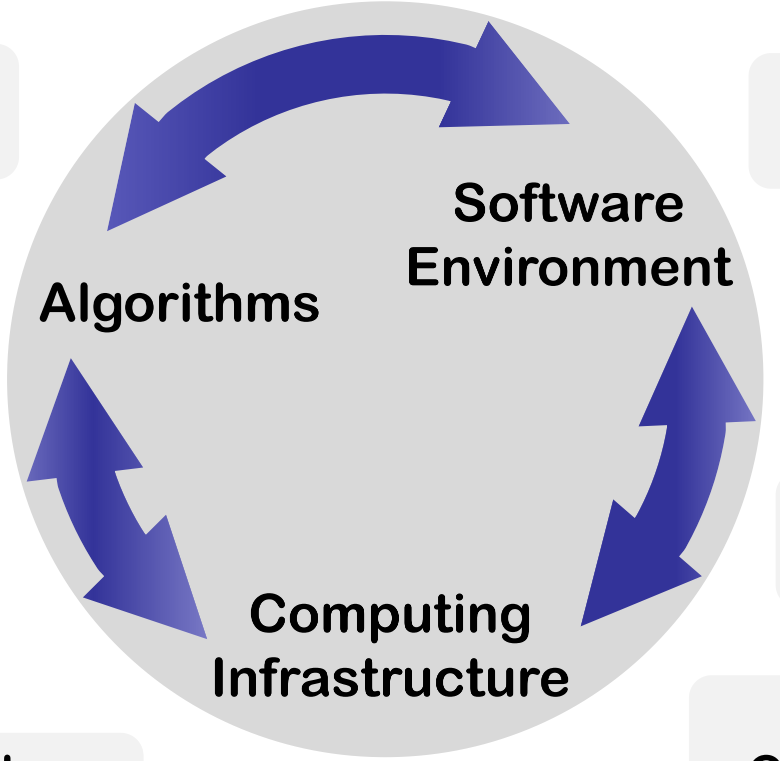

Algorithm Libraries

Networks

Data Storage Systems

Parallel Algorithms

Compute Clusters/HPC

Accelerator Hardware

Compute Centers

GSI/FAIR Green Cube

12 MW common data center for FAIR and GSI
- In operation since Feb ‘16

PUE < 1.07
- Very good Power Usage Efficiency (PUE)
- Reduced CO_2 Emission
- Low Operating Costs

More detailed information:
Jan Trautmann, HEPIX Spring 2016
https://indico.cern.ch/event/466991/contributions/1143585/

Green Cube Cooling Concept

Design: V. Lindenstruth, FIAS+GSI

Evaporative cooling towers

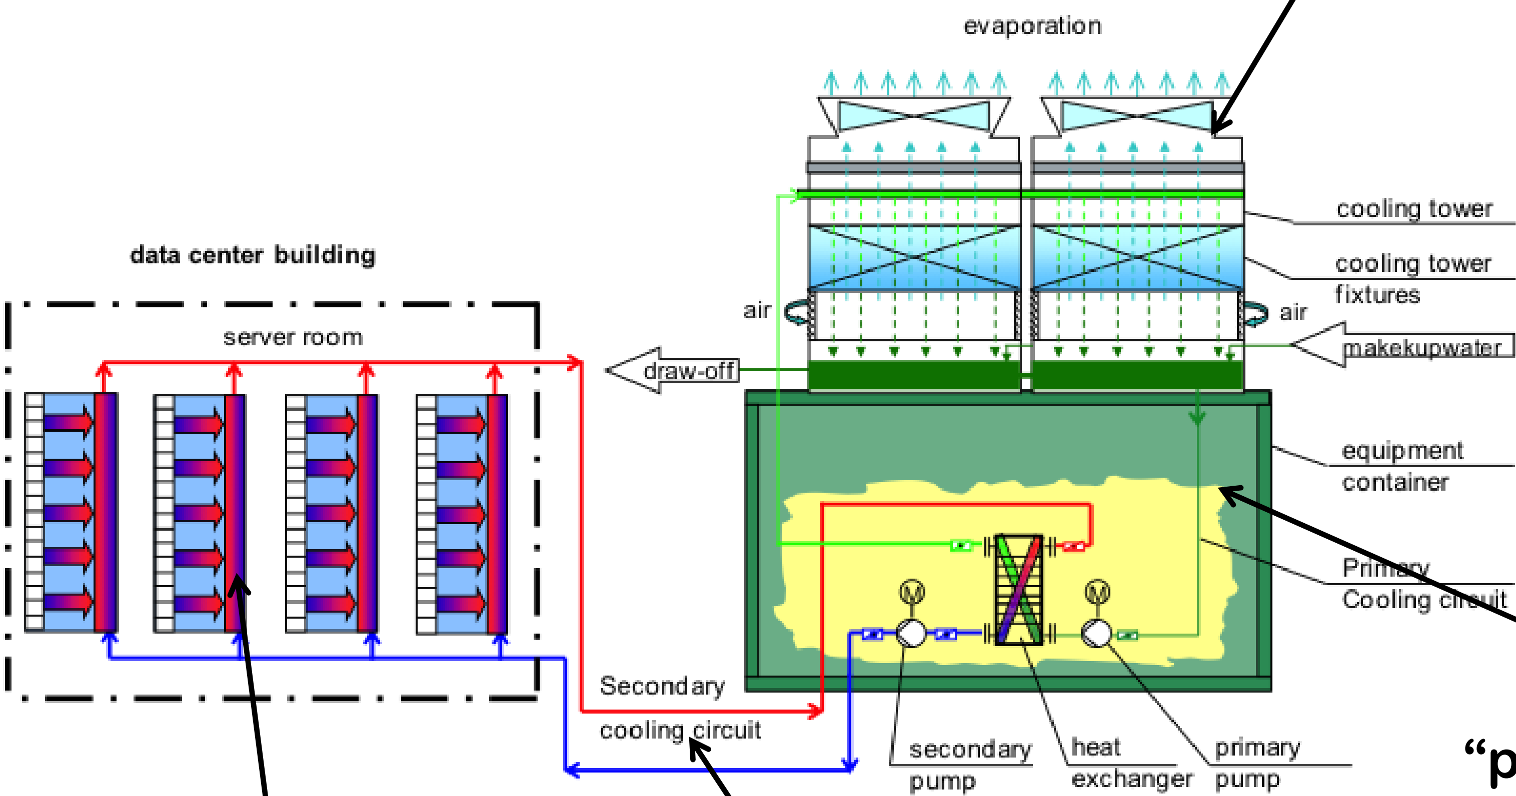

Open "primary" circuit

Passive rear-door heat-exchanger

Closed "secondary" circuit

GSI/FAIR Green Cube

Constructed:	Dec ‘14 – Nov ’15
Building:	6 Floors, 4.645 sqm 768 19” racks (256 racks in 1st stage)
Cooling & Power:	12 MW (4 MW in 1st stage) PUE < 1.07 (Design) PUE ~ 1.04 (Commissioning) Water cooled - Passive rear-door heat-exchanger - Evaporative cooling towers N+1 Redundancy
Cost:	16 M€ (1st stage: 11.5 M€)

Green Cube Buildings

Cooling Towers

GT1 (Cooling)

GC 1 (Computing)

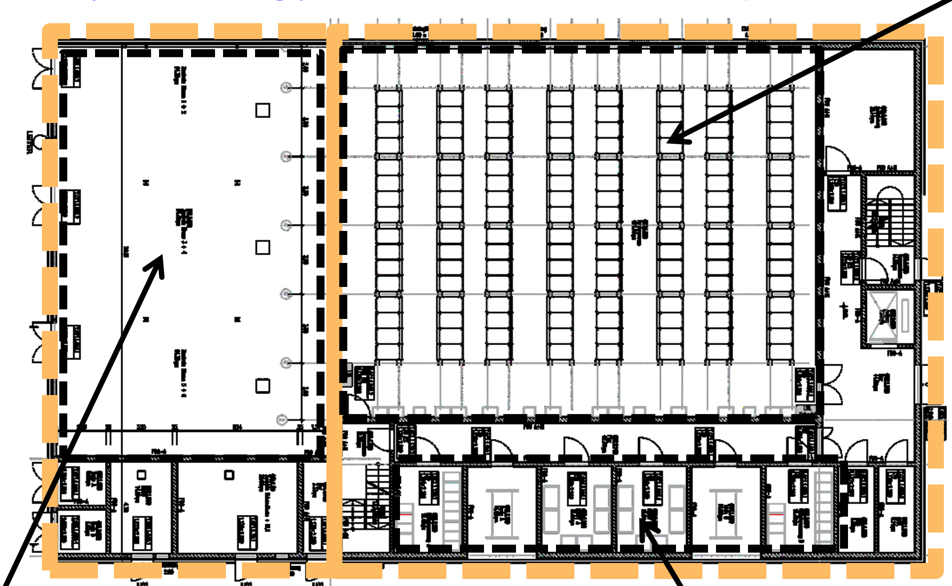

Computing Room

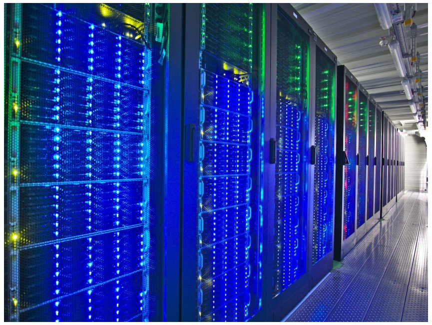

Cooling Circuits Pumps & Heat Exchangers

Electrical Power

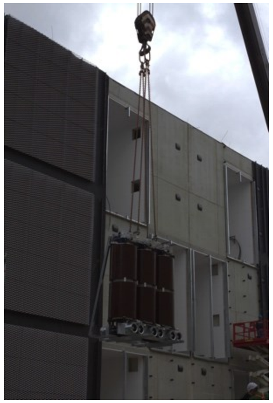

Green Cube CG1

Green Cube CG1

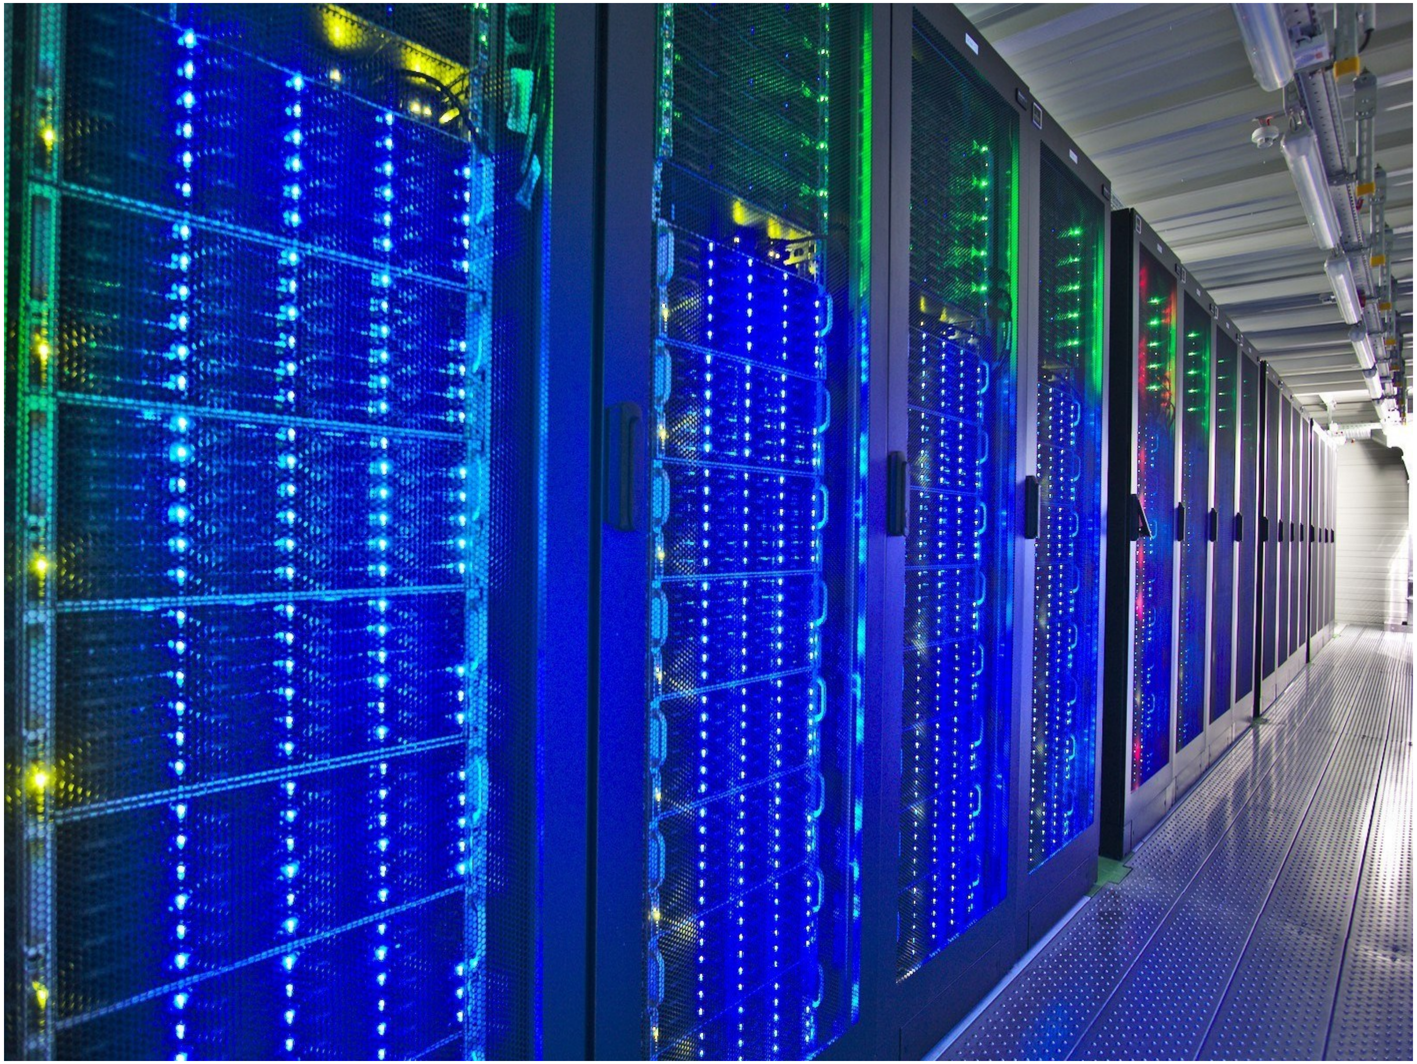

Helmholtz-Rosatom Project:
FAIR Russia Research Center (FRRC) Supercomputer at ITEP

- **GSI/FAIR: B.Sharkov, H.Gutbrod, V.Lindenstruth**
- **ITEP**: A.**Golubev,** M.**Polikarpov,** M.**Sokolov, I.Korolko, V.Kolosov**

FRRC Supercomputer at ITEP

February 2012

Launch of first supercomputer servers

1280 AMD cores

2013-2014

Installation and commissioning of hardware and software:

- AMD cores: 10240
- GPGPU cards TESLA: 40
- MSS: 2 PB
- Ethernet and Infiniband
- EMERSON container
- JAEGGI cooling system

Performance:

CPU - 59.2 TFlops (linpack)

GPU - 36.0 TFlops (linpack)

CBM Standalone First Level Event Selection (FLES) Package

(1) Collision

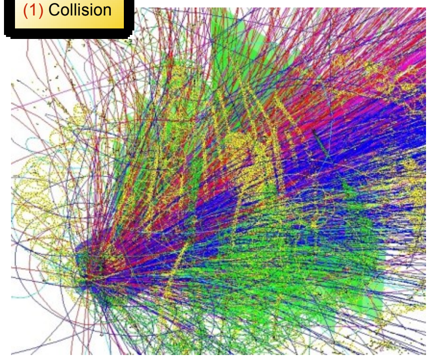

(2) Detection

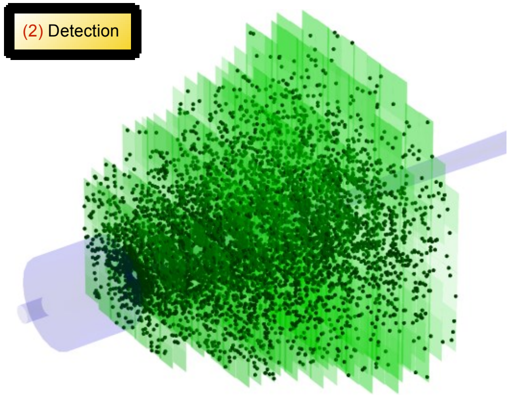

(3) Reconstruction

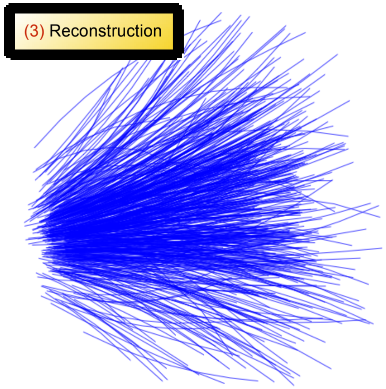

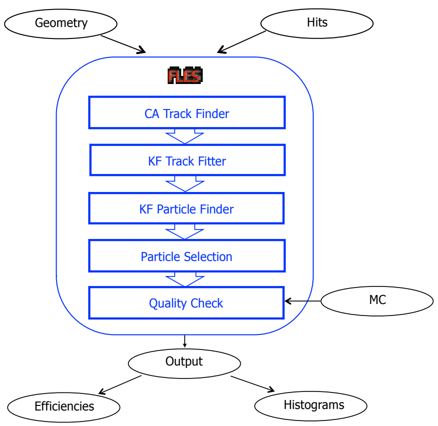

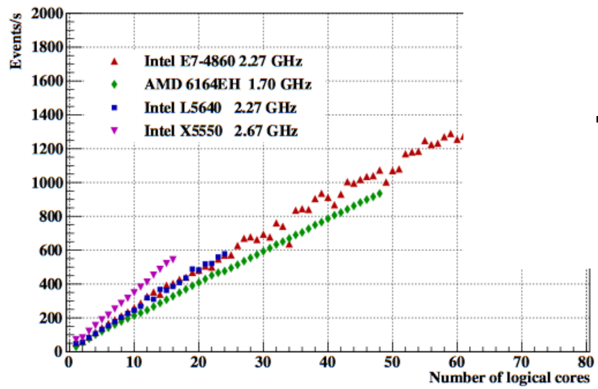

2.2×10^5 events/s

[redacted]

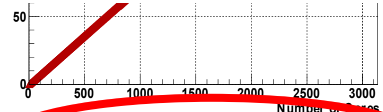

ITEP, Moscow
Cooperation with Kurchatov Institute is planned

The first version of the FLES package is vectorized, parallelized, portable and scalable up to 3 200 CPU cores

Ivan Kisel, Uni-Frankfurt, FIAS, GSI

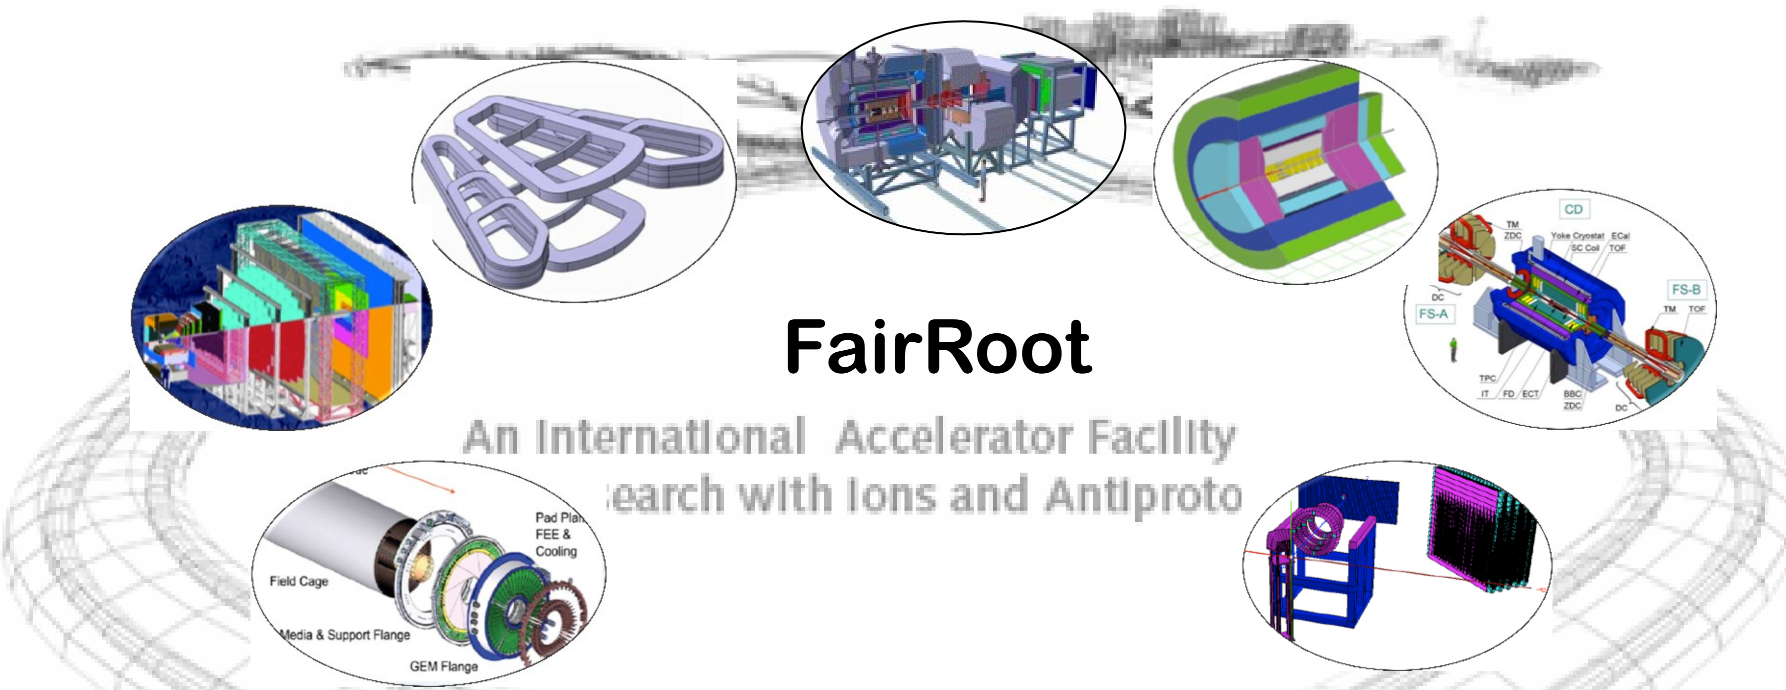

Common simulation, reconstruction and analysis software framework for the FAIR experiments (and beyond)

https://github.com/FairRootGroup/FairRoot

Mul Purpos Detector and MpdRo

The software MpdRoot is developed for the MPD event simulation, reconstruction of experimental or simulated data and following physical analysis of heavy ion collisions registered by the MultiPurpose Detector at the NICA collider.

(based on ROOT and FairRoot)

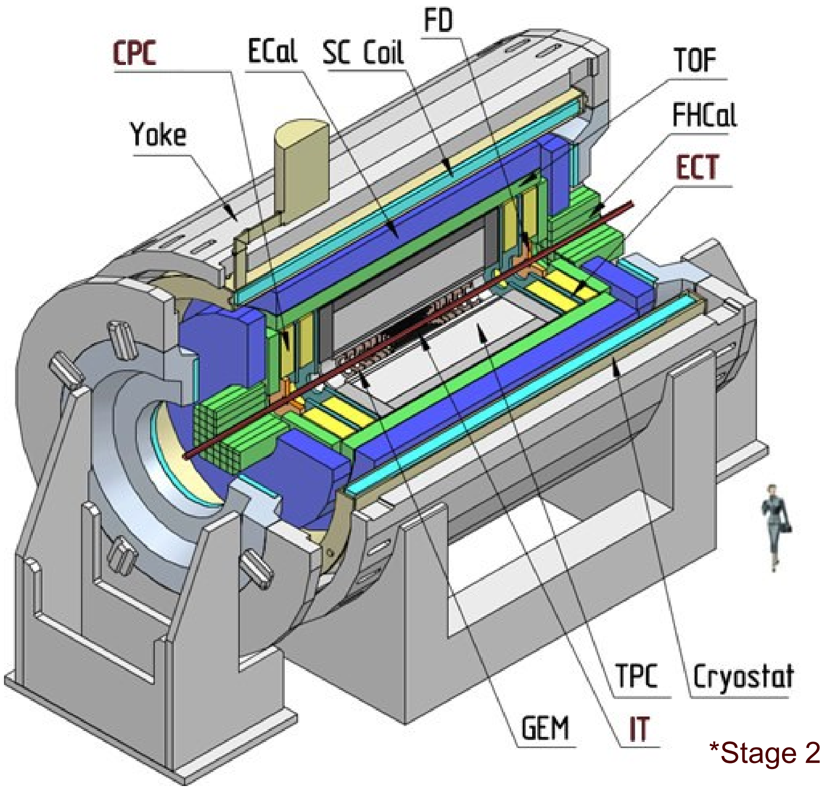

The MpdRoot software is available in the GitLab https://git.jinr.ru/nica/mpdroot

16 November 2016 K. Gertsenberger 3

FairRoot @ NICA, JINR

BM@N and BmnRoot softwar

The software BmnRoot is developed for the BM@N event simulation, reconstruction of experimental or simulated data and following physical analysis of collisions of elementary particles and ions with a fixed target at the NICA collider.

(based on ROOT and FairRoot)

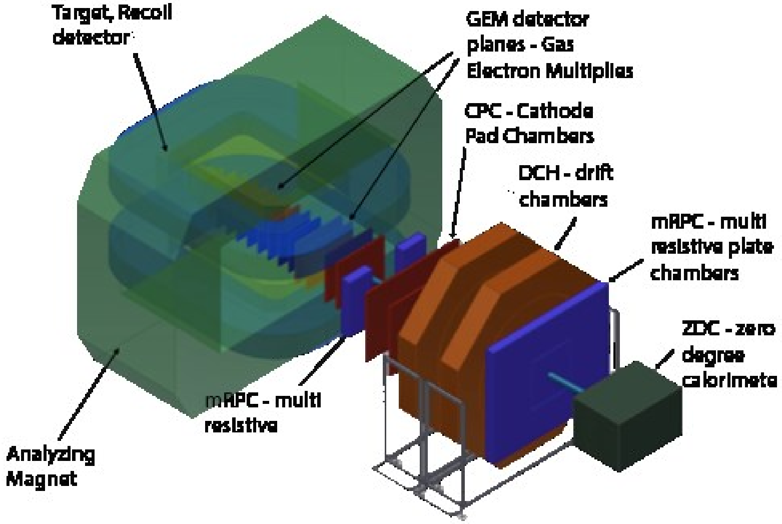

- deutron and C^{11} beams with T = 3 - 4 AGeV, targets: carbon, copper or none
- Trace beams, measure beam profile and time structure
- Test detector response ToF-400, ToF-700, T0+Trigger, DCH-1, DCH-2 and ZDC, ECAL modules, beam monitors (4 run: + GEM & Si)
- Test integrated DAQ and trigger system

The BmnRoot software is available in the GitLab https://git.jinr.ru/nica/bmnroot

16 November 2016 K. Gertsenberger 13

FairRoot @ NICA, JINR

MpdRoot (&BmnRoot) framewor

- Both use FairSoft external packages:

ROOT, XRootD, Pythia, PLUTO, HepMC, MillePede, Geant3/4, VGM, gtest, GSL, boost…

- There is a common class part inherited from FairRoot (GSI)
- Experiment-specific parts are developed for each detector independently
- Advanced detector response functions, realistic tracking and PID were included
- Extended set of event generators for collisions:

PLUT
+ Ur
HAD
simpl

MPD and BM@N homepage: http://mpd.jinr.ru

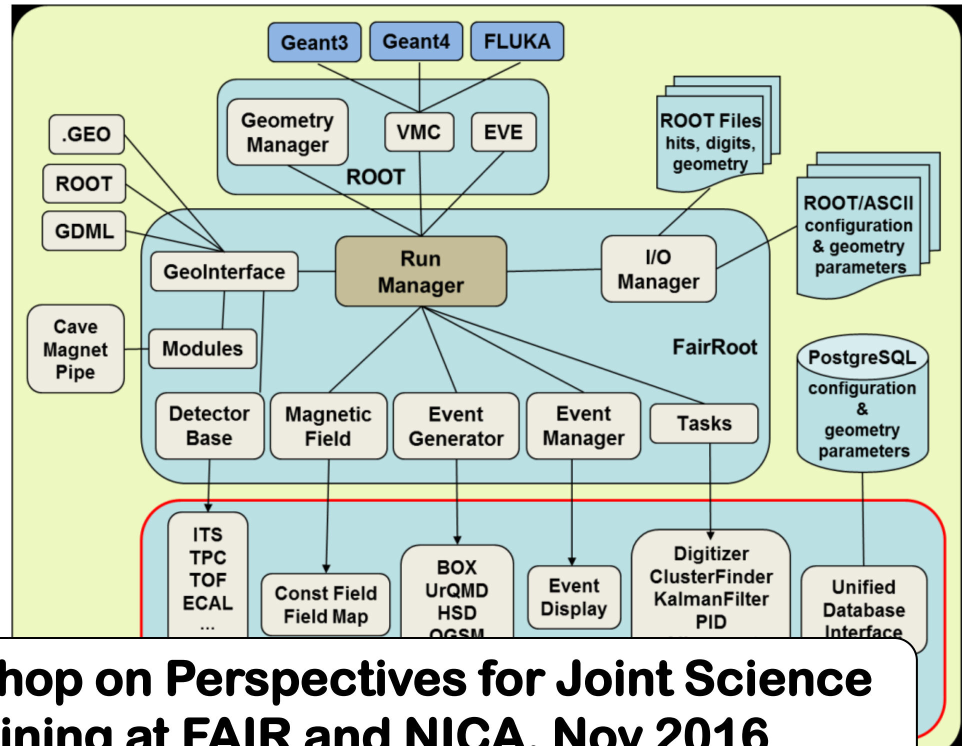

Slides from Workshop on Perspectives for Joint Science and Academic Training at FAIR and NICA, Nov 2016

16 November 2016 K. Gertsenberger 6

FairRoot

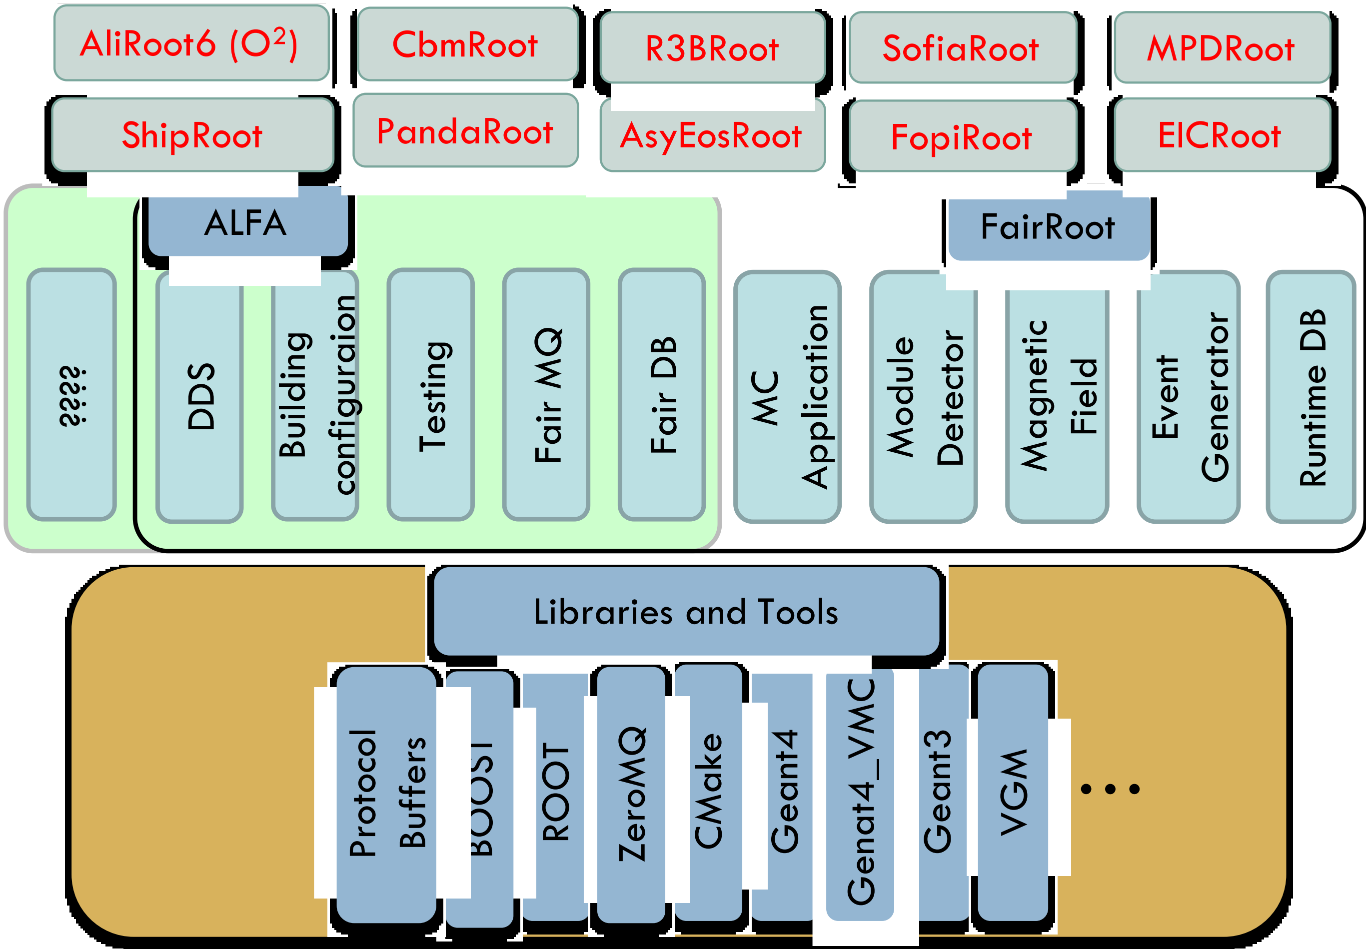

Modularized, LGPLv3 licensed common software stack for the FAIR experiments and others, development steered by GSI

FairMQ

The Data Processing Component of FairRoot

- Multi-process concept (specialized devices)
- Data-flow model: Message queues for data exchange, technology agnostic

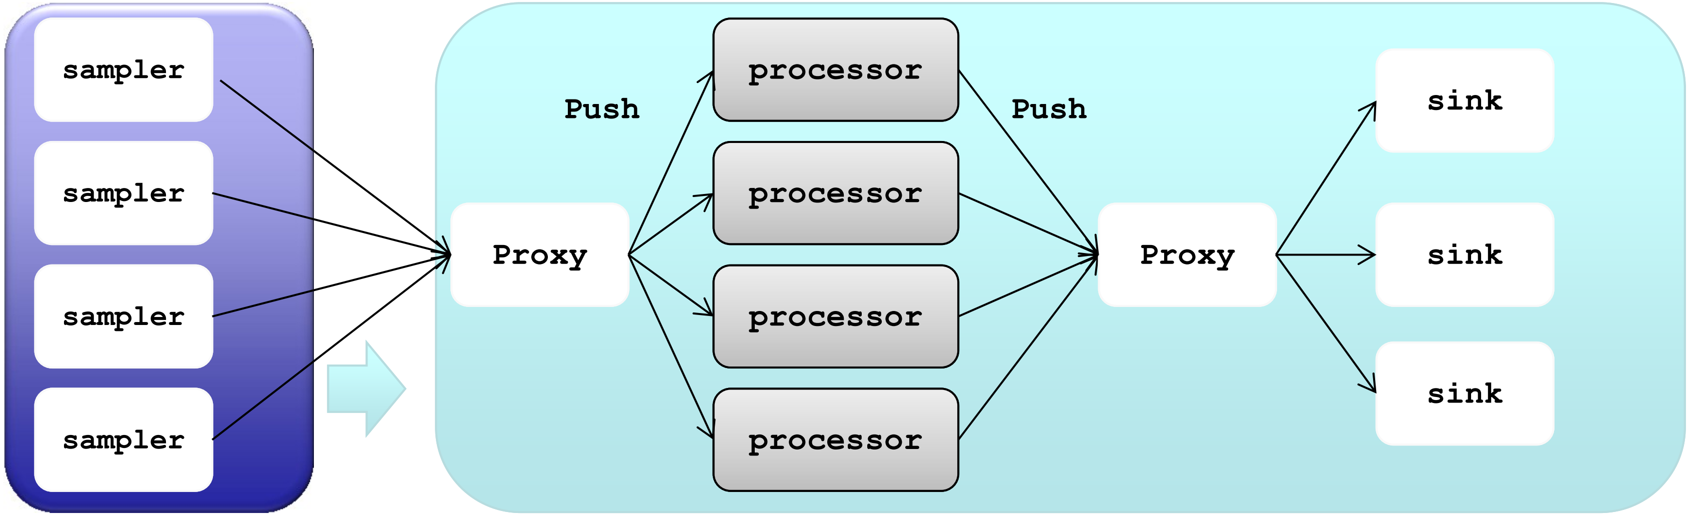

Design Goals

- Scalability, Maintainability, Reliability
  - efficient use of multi-core architectures
- Reusable with common data processing components
  - Reduce cost of new developments, agile development

FairMQ

Looking at the IT landscape: shift towards

- **Microservices**
  - **Unbundled, decentralized modules**
  - **Organized around specific capability**
- **Containers**
- **Algorithm Economy**

These are at the heart of the „cloud/app“ business model/economy

- **driven by scalability and reliability demands**
- **based on multi-process and message exchange**
- **development cost advantage**

FairMQ uses many of these technologies under the hood; replacing custom code (e.g. ALICE HLT framework)

Example Application: PANDA MVD Pixel Detector Prototype

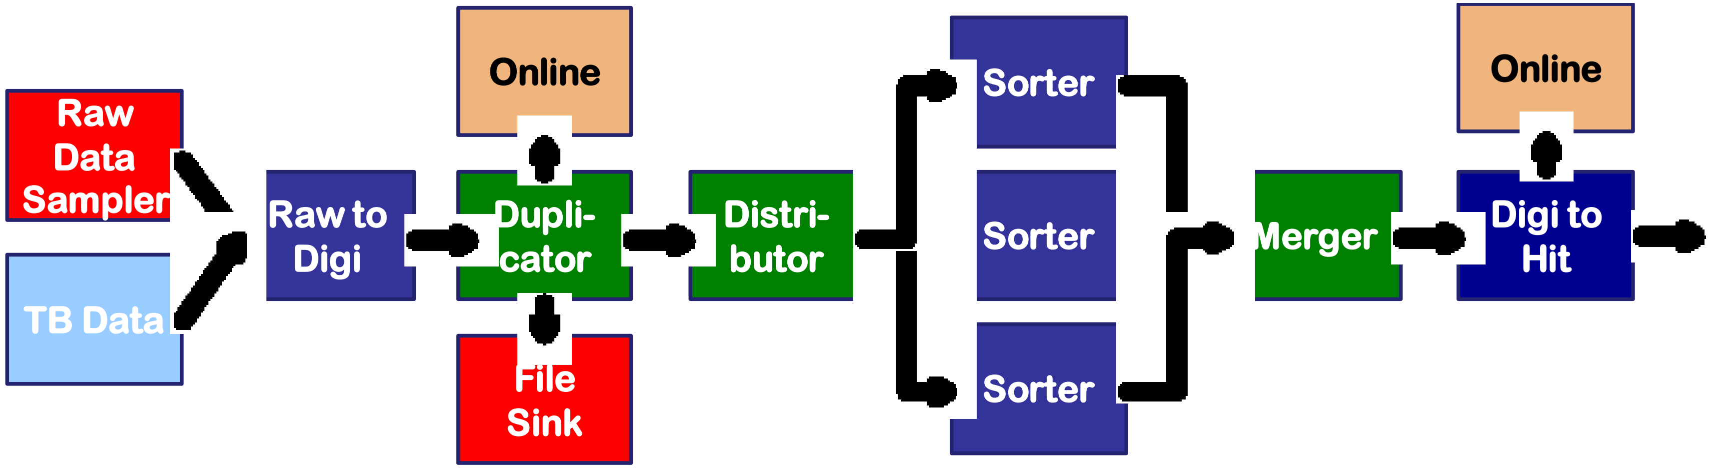

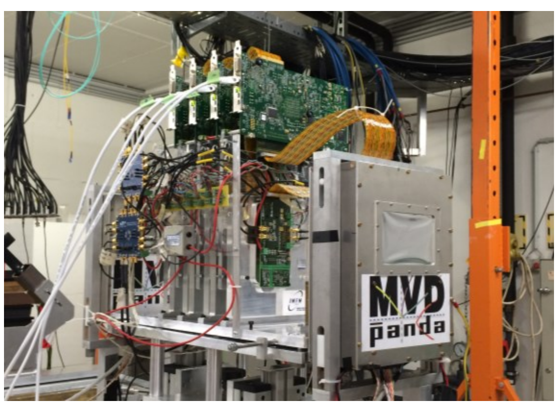

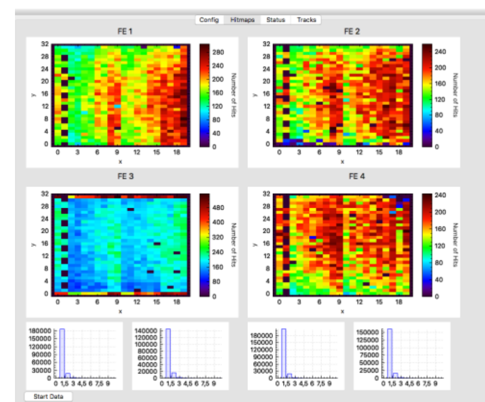

Tobias Stockmanns (FZJ, PANDA)

Driven by needs of FAIR experiments
for online reconstruction

- ~1 TByte/s into online farms
- ~300.000 cores (majority on-site in common compute center)
- 35 PByte/"year" disk
- 30 PByte/year tape

Algorithms and software development

- common frameworks & libraries, e.g. FairRoot