Computing challenges and opportunities for FAIR experiments

university of groningen

kvi - center for advanced radiation technology

Johan Messchendorp, FAIR2014, Worms, Germany, 13 October 2014

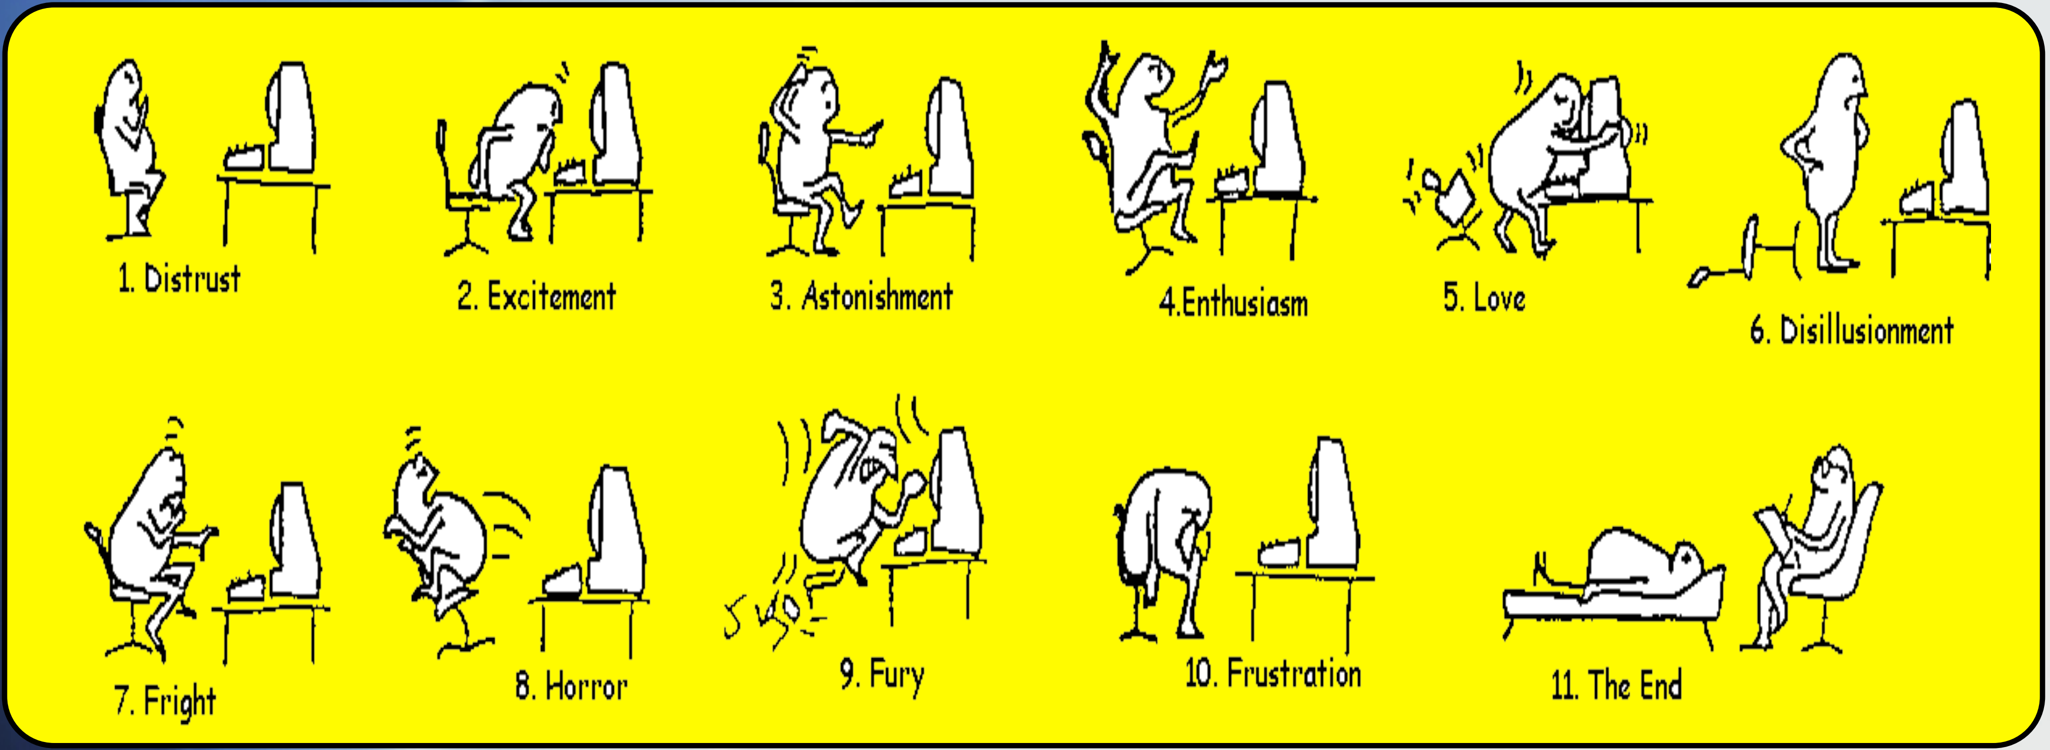
1. Distrust
2. Excitement
3. Astonishment
4.Enthusiasm
5. Love
6. Disillusionment
7. Fright
8. Horror
9. Fury
10. Frustration
11. The End

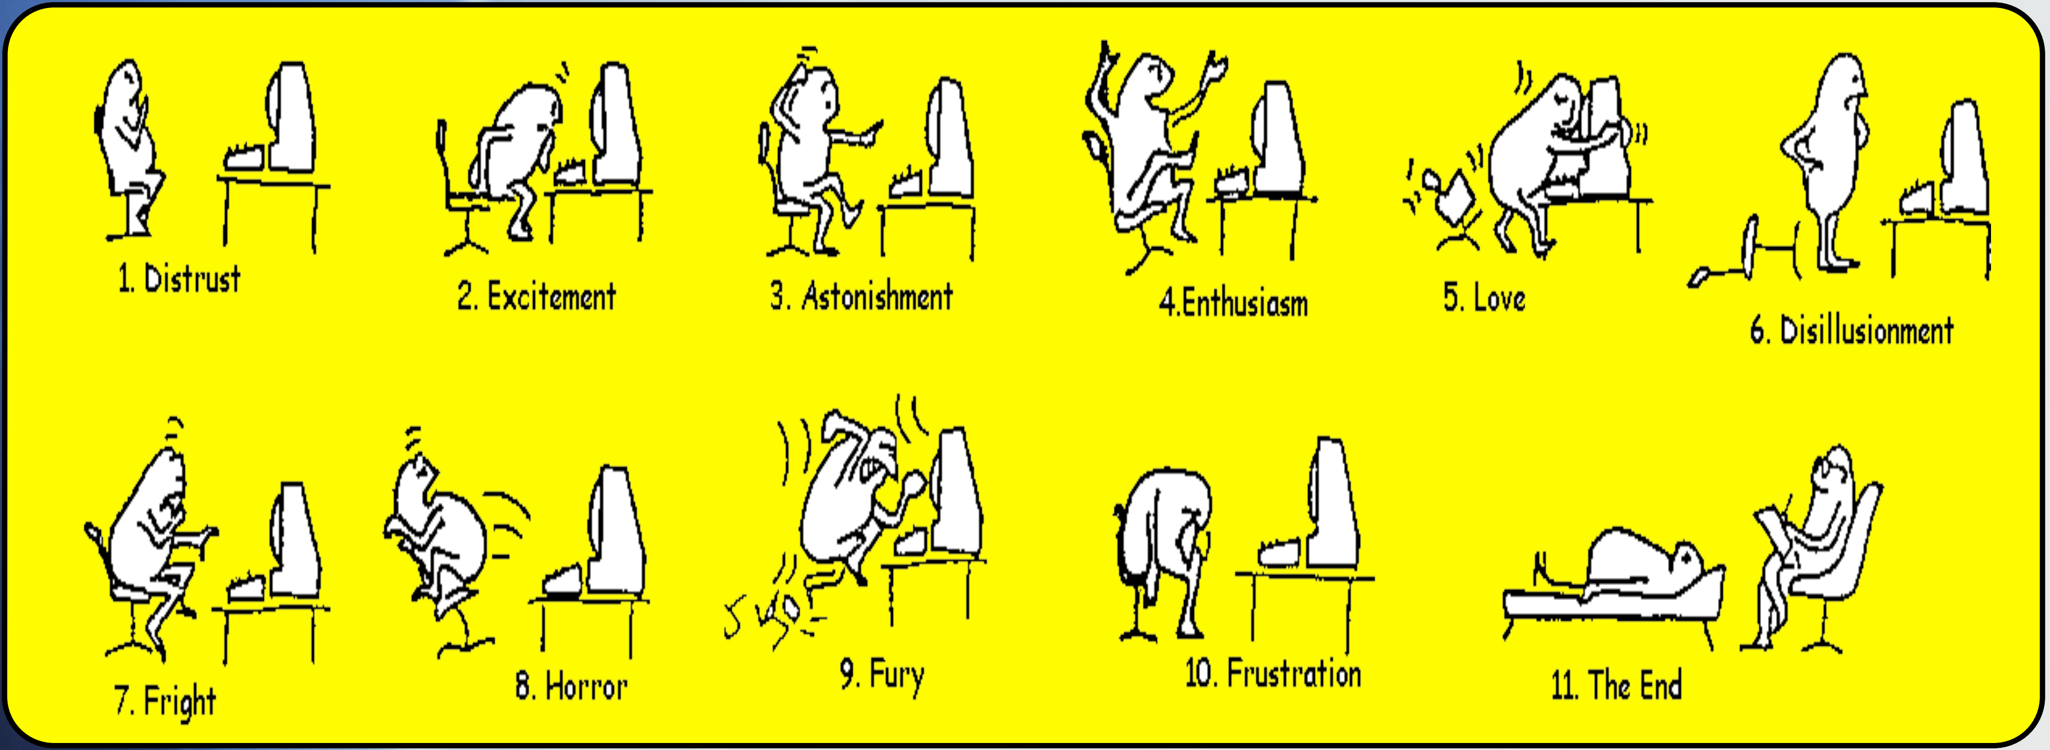

Disclaimer:

The information selected for this talk solely reflects my personal excitement, astonishment, enthusiasm, and love for (FAIR) computing.

Who are you? And why should I listen?

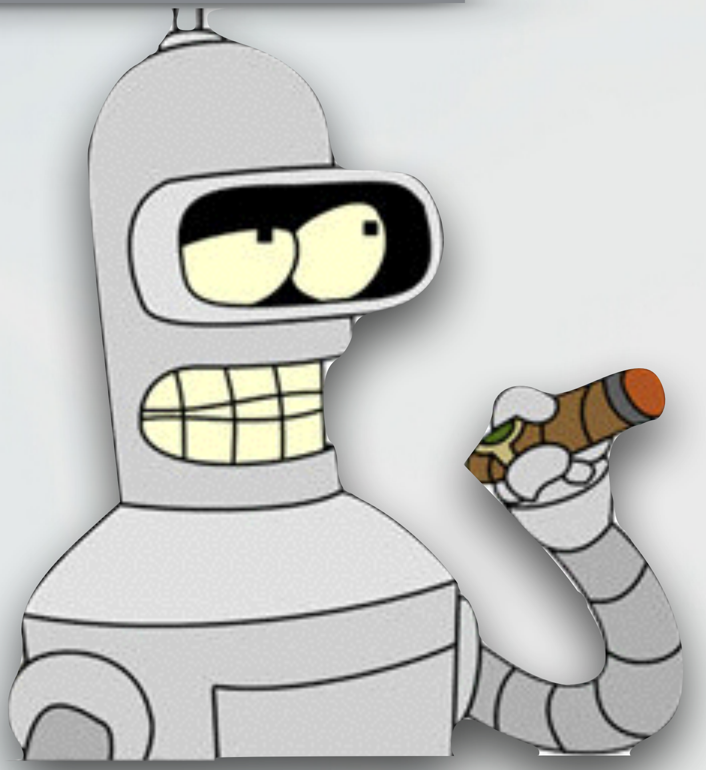

Who are you? And why should I listen?

1. Everyone (dis)likes computing, good for the upcoming dinner discussion.

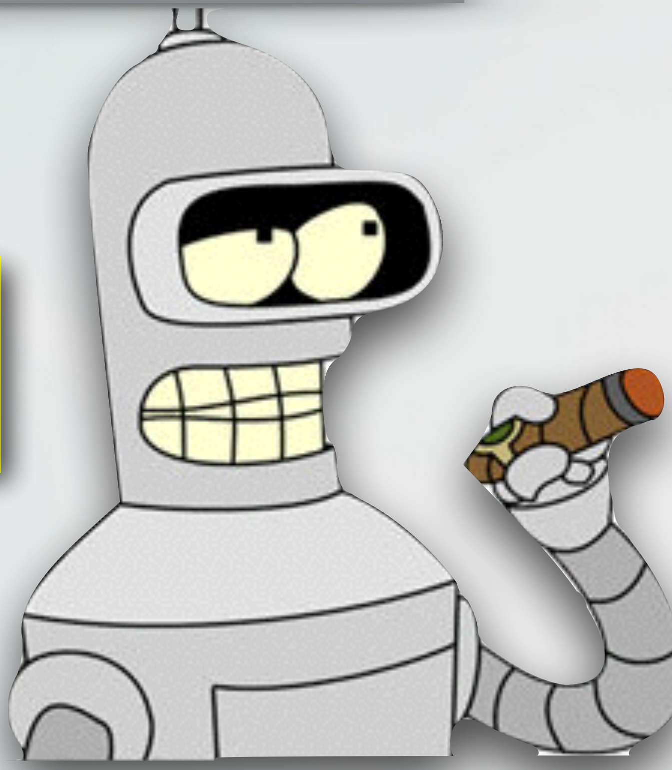

Who are you? And why should I listen?

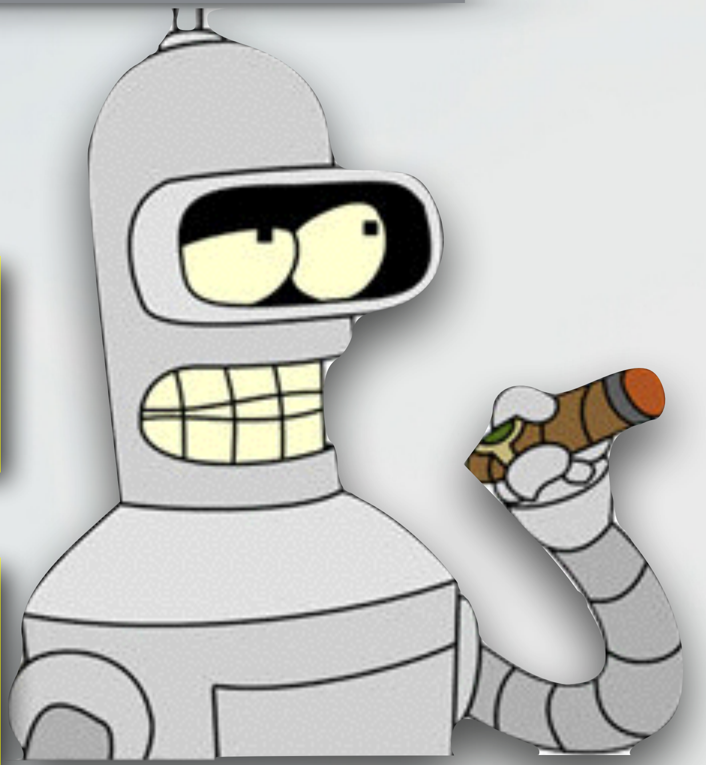

1. Everyone (dis)likes computing, good for the upcoming dinner discussion.

2. Very large chance you have to deal with “FAIR computing” in some way.

Who are you? And why should I listen?

1. Everyone (dis)likes computing, good for the upcoming dinner discussion.
2. Very large chance you have to deal with “FAIR computing” in some way.
3. Computing remains rapidly changing: new developments & opportunities!

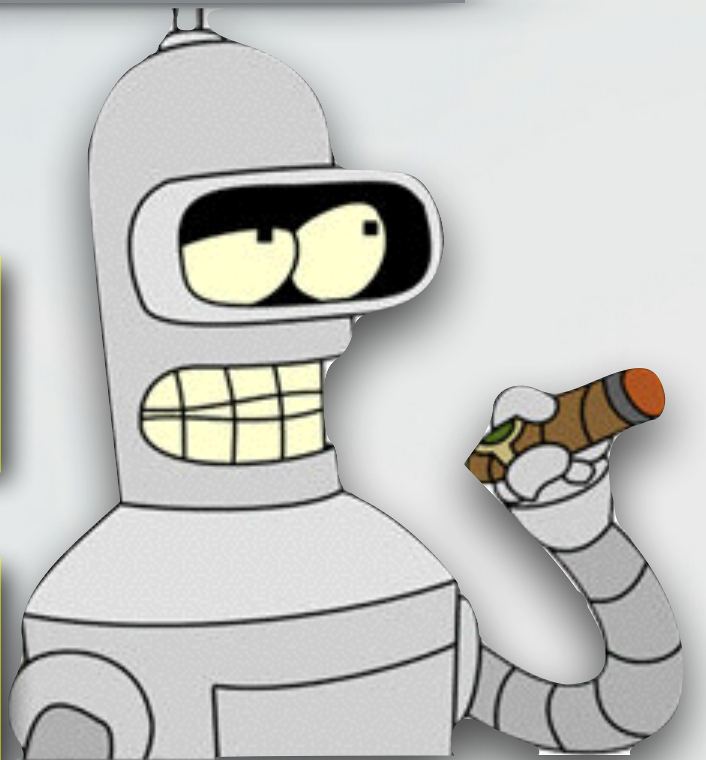


BIG
DATA
USE
ONE
SINCE
SOFTWARE
USING
TYPES
WORKING
GARTNER
MASSIVELY
REQUIRING
UBIQUITOUS
RADIO-FREQUENCY
COMPLEXITY
SAN
SOLID
WIRELESS
BUSINESS
MOVING
ANALYTICS
PERFORMANCE
RELATED
ALSO
BIOLOGICAL
HUNDREDS
CREATED
TOLERABLE
PARALLEL
SIZE
NEEDED
QUALITIES
BURIED
PROCESSING
LOGS
PETABYTES
INTERNET
TECHNOLOGIE
ORGANIZATIONS
DESKTOP
ZETTABYTES
DISK
GENOMICS
COMPUTING
TOOLS
SET
SHARED
TIME
RELATIONAL
CURRENTLY
TENS
PRACTITIONERS
WITHIN
COMBAT
SIGNIFICANT
PROCESS
INCLUDE
SYSTEMS
CAPACITY
FC
DEFINITION
BIOGEOCHEMICAL
NETWORKS
FORMS
USED
DISTRIBUTED
CAPTURE
SEARCH
INFORMATION
OPPORTUNITIES
MAY
MANAGEMENT
MANAGE
GROW
PRESENTATIONS
NOW
CONNECTOMICS
RECORDS
WORLD'S
SOCIAL
LARGE
DEFINING
STORE
COST
CONTINUES
TERABYTES
MPP
RESEARCH
STORAGE
EVERY
CASE
CITATION
COMPLEX
DATABASES
LARGER
INDEXING
DIFFICULTY
TARGET
ABILITY
SETS
EXAMPLES
ELAPSED
CURRENT
THOUGHT
APPLIED
SENSOR
ARCHIVES
AMOUNT
DESCRIBING

The "Big Data" challenges

Doug Lany, META Group Inc.

The "Big Data" challenges

Doug Lany, META Group Inc.

Volume: data volume grows faster than computing power.

Velocity: the rate at which data is processed for fast decision making.

Variety: heterogeneity of data types, representation, and interpretation.

Veracity: degree of certainty about data.

Complexity: data from complex network of many different sources.

FAIR and the “Big Data” challenge

Thorsten Kolleger, Stephane Pietri, Stefano Spataro, Volker Friese

FAIR and the “Big Data” challenge

Thorsten Kolleger, Stephane Pietri, Stefano Spataro, Volker Friese

PANDA/CBM

- monolithic systems
- LHC-like rates & requirements
- up to 1 TB/s data rates
- online data reduction to ~1 GB/s
- total request 3×10^5 cores at Tier0
- distributed computing

FAIR and the "Big Data" challenge

Thorsten Kolleger, Stephane Pietri, Stefano Spataro, Volker Friese

PANDA/CBM

- monolithic systems
- LHC-like rates & requirements
- up to 1 TB/s data rates
- online data reduction to ~1 GB/s
- total request 3×10^5 cores at Tier0
- distributed computing

NUSTAR/APPA

- many smaller experiments
- up to 1 GB/s data rates
- computing power ~4×10^3 cores
- large variety in setups
- high flexibility in data processing
- complexity in database management

FAIR and the "Big Data" challenge

Thorsten Kolleger, Stephane Pietri, Stefano Spataro, Volker Friese

PANDA/CBM

- monolithic systems
- LHC-like rates & requirements
- up to 1 TB/s data rates
- online data reduction to ~1 GB/s
- total request 3×10^5 cores at Tier0
- distributed computing

NUSTAR/APPA

- many smaller experiments
- up to 1 GB/s data rates
- computing power ~4×10^3 cores
- large variety in setups
- high flexibility in data processing
- complexity in database management

Volume
Velocity
Veracity
Variety
Complexity

The data-processing challenge

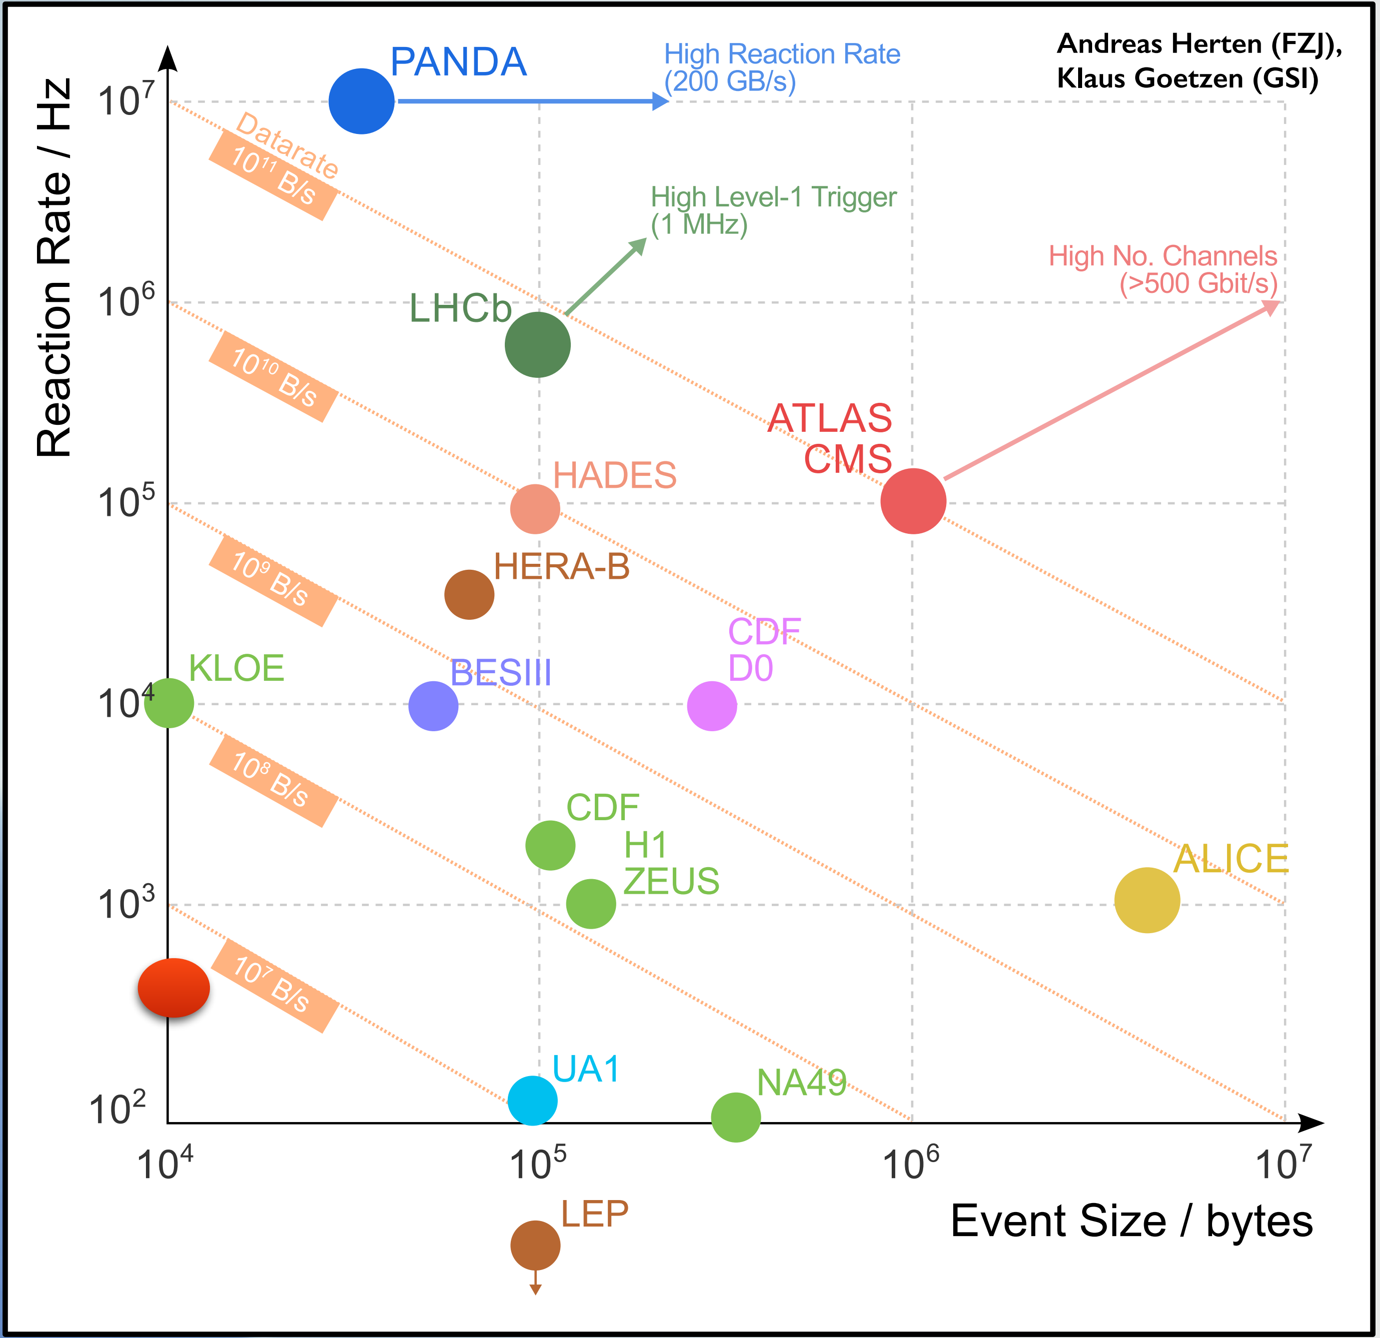

The data-processing challenge

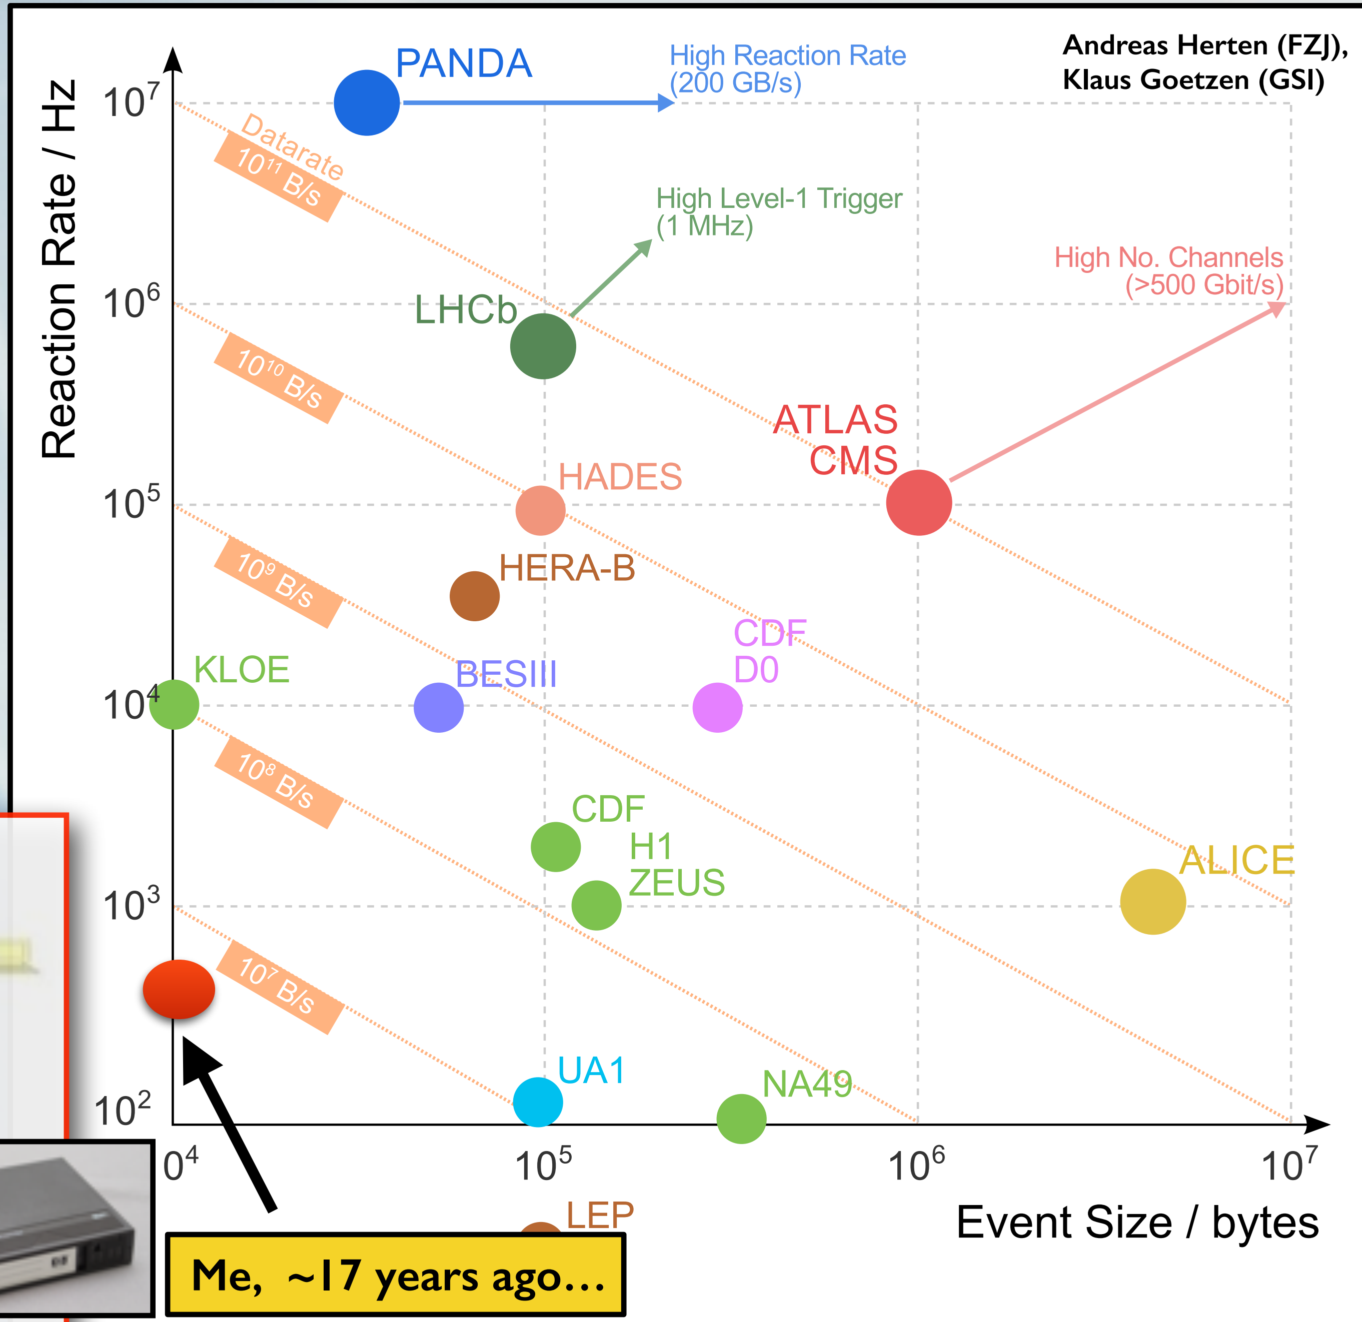

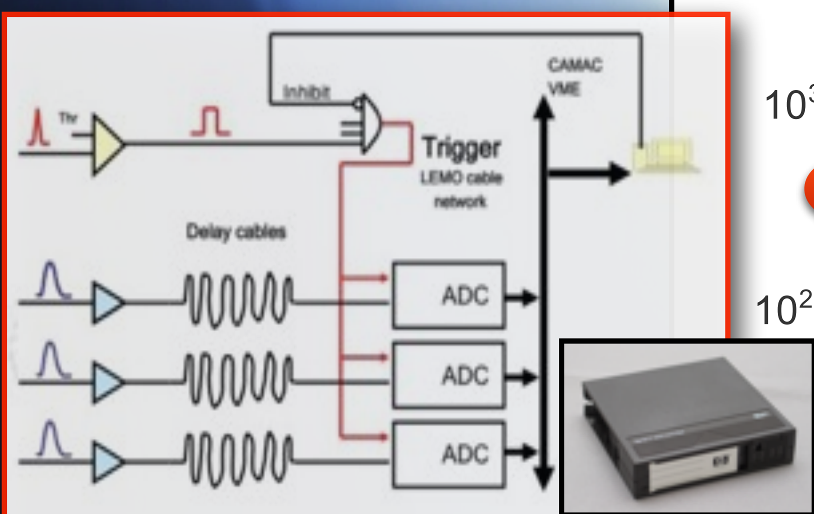

Me, ~17 years ago…

The data-processing challenge

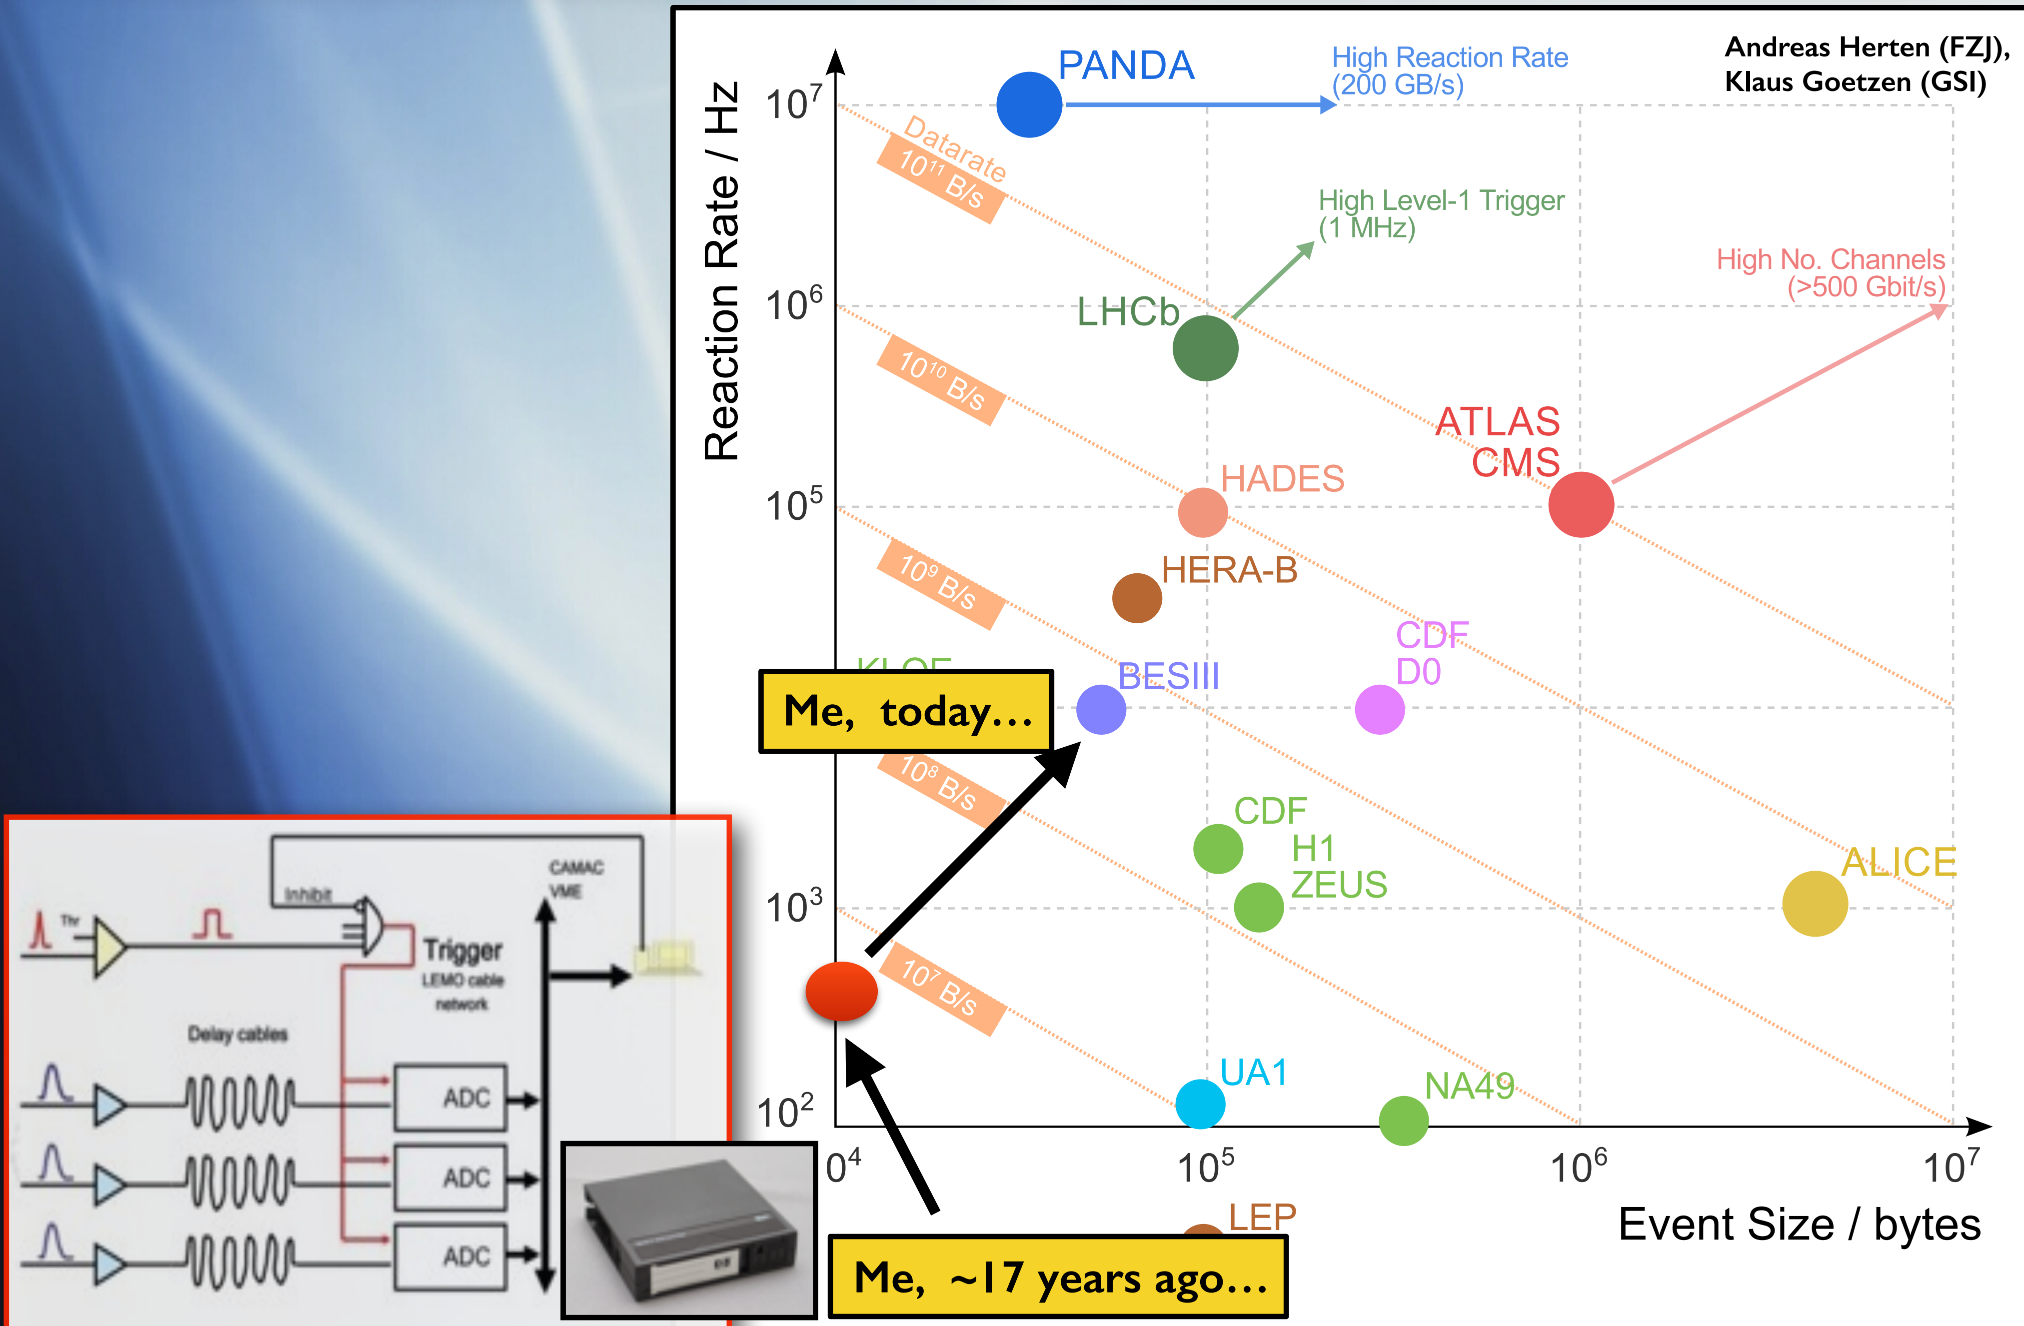

The data-processing challenge

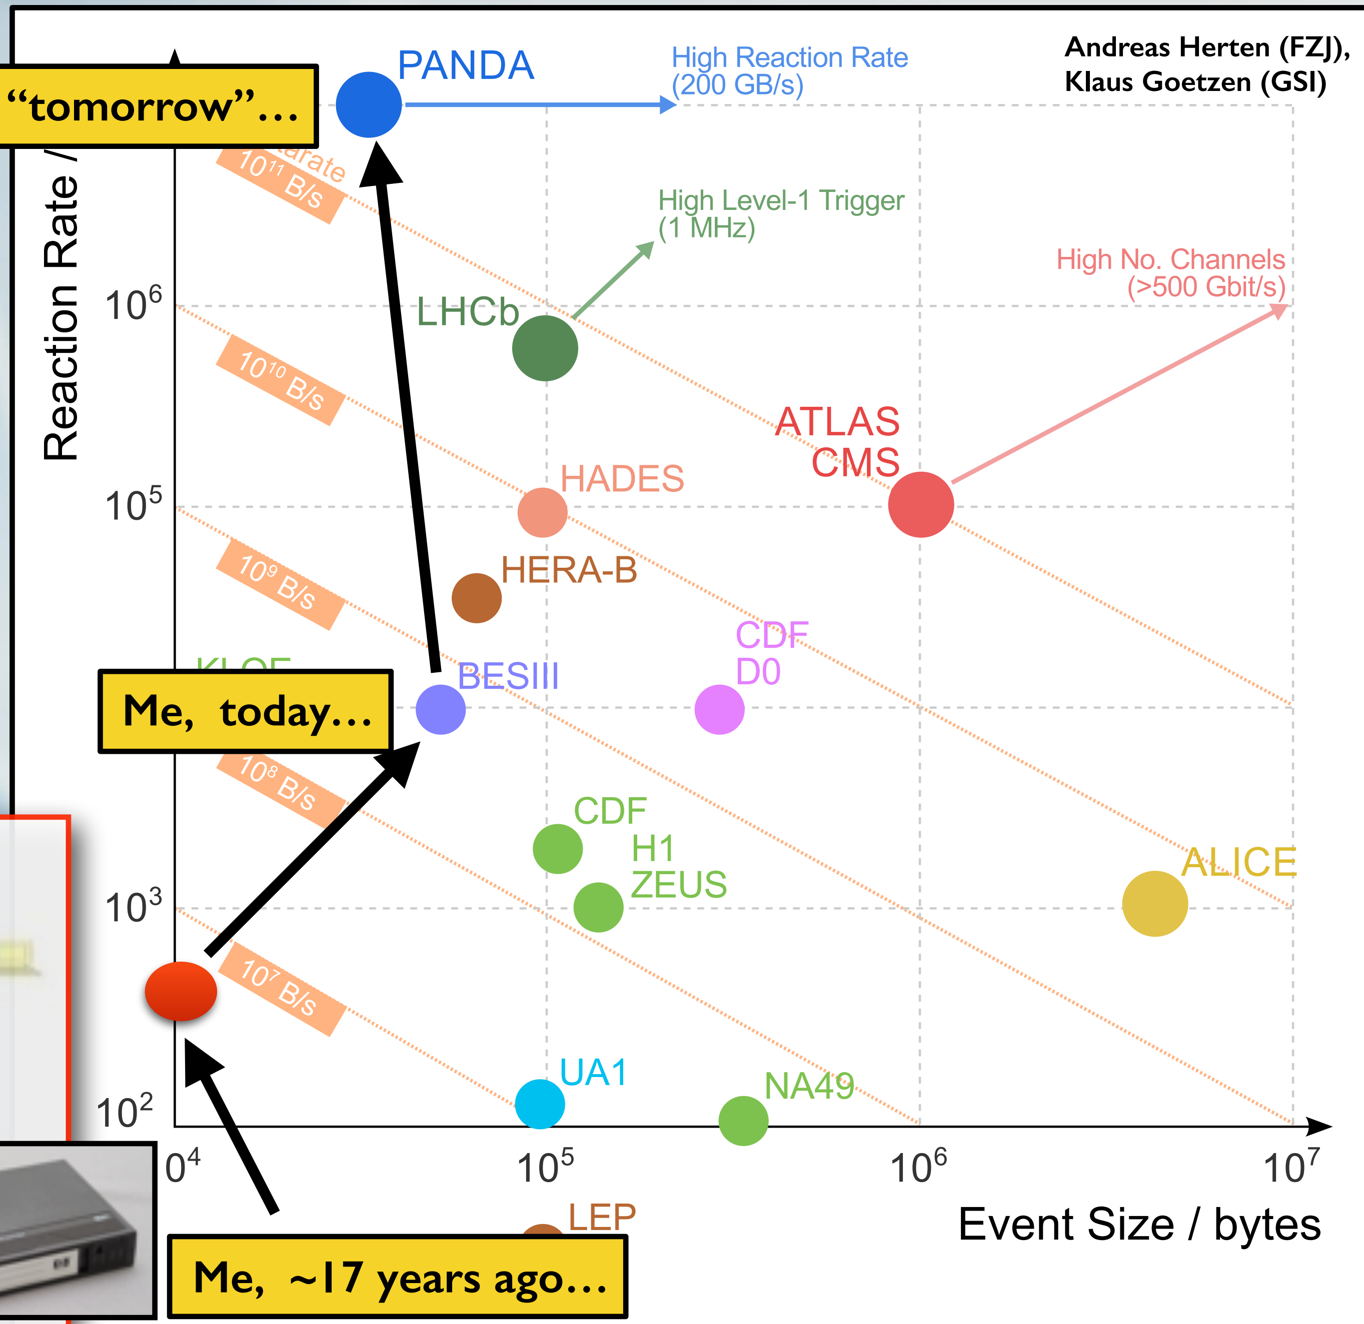

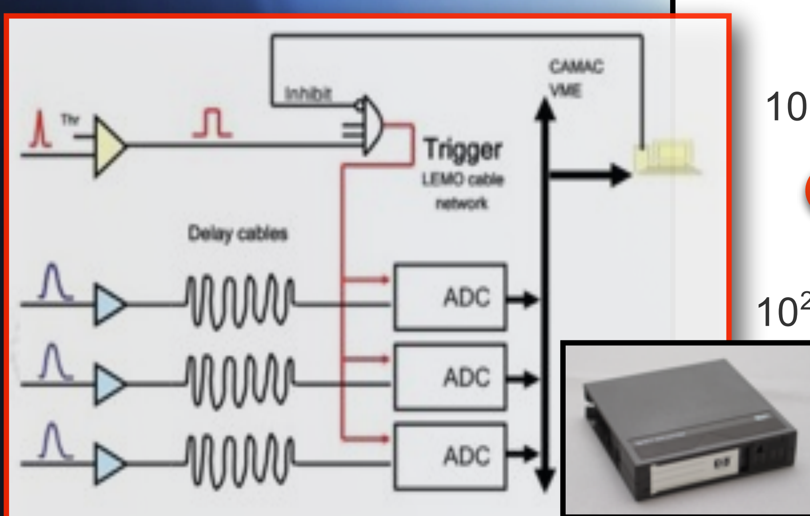

“Tomorrows” experiment

“Tomorrows” experiment

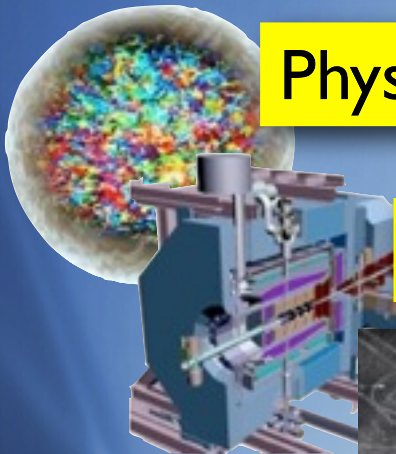

Physics

Detector

Feature extraction

Online computing

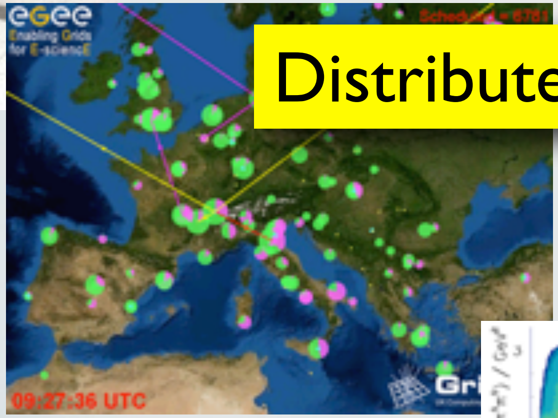

Distributed computing

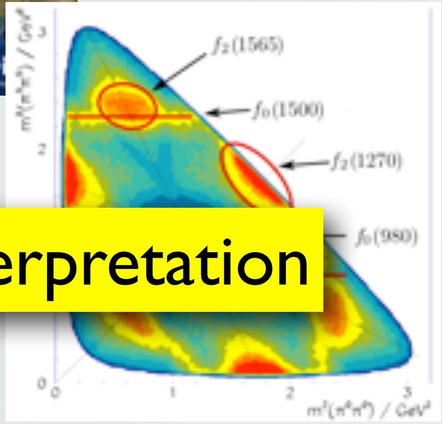

Data interpretation

“Tomorrows” experiment

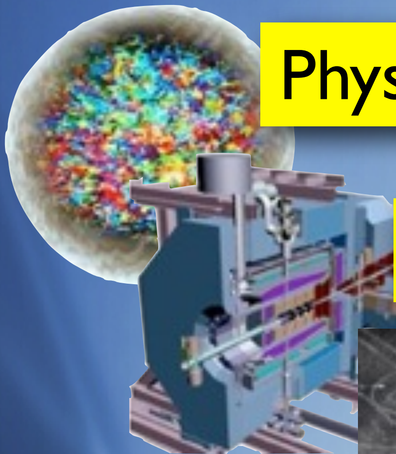

Physics

Detector

Feature extraction

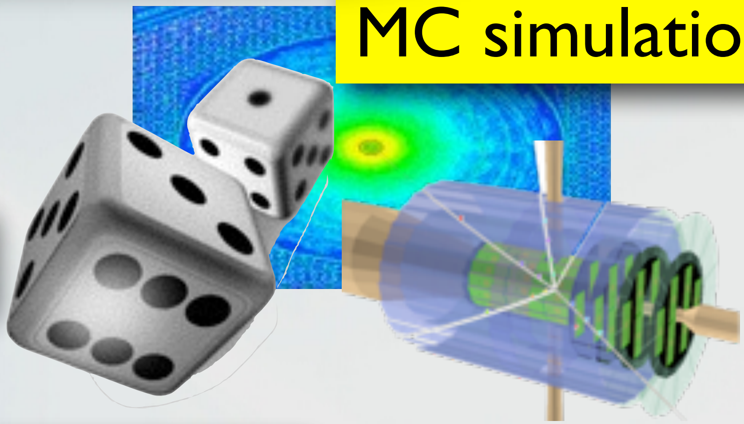

MC simulations

Online computing

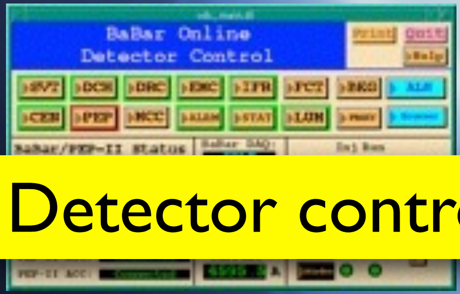

Detector control

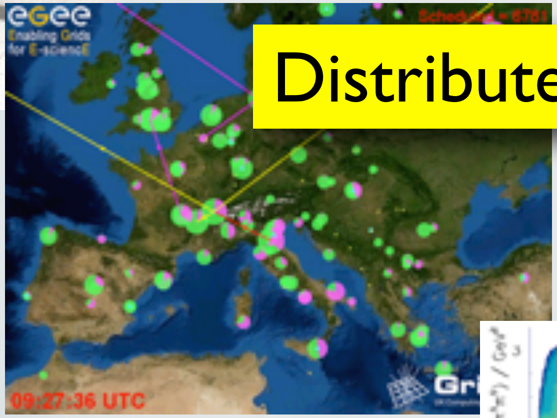

Distributed computing

Database management

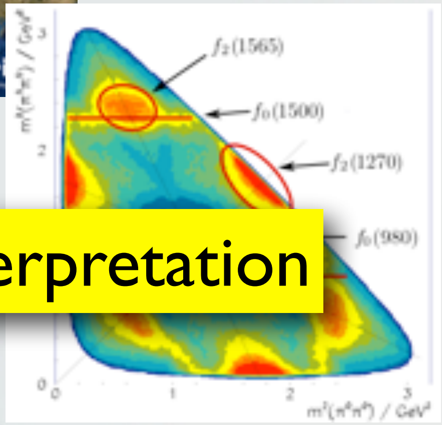

Data interpretation

The "physics" challenge! (PANDA as an example)

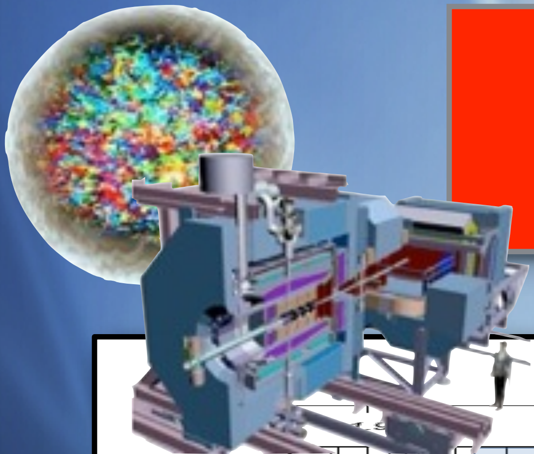

Klaus Goetzen (GSI)

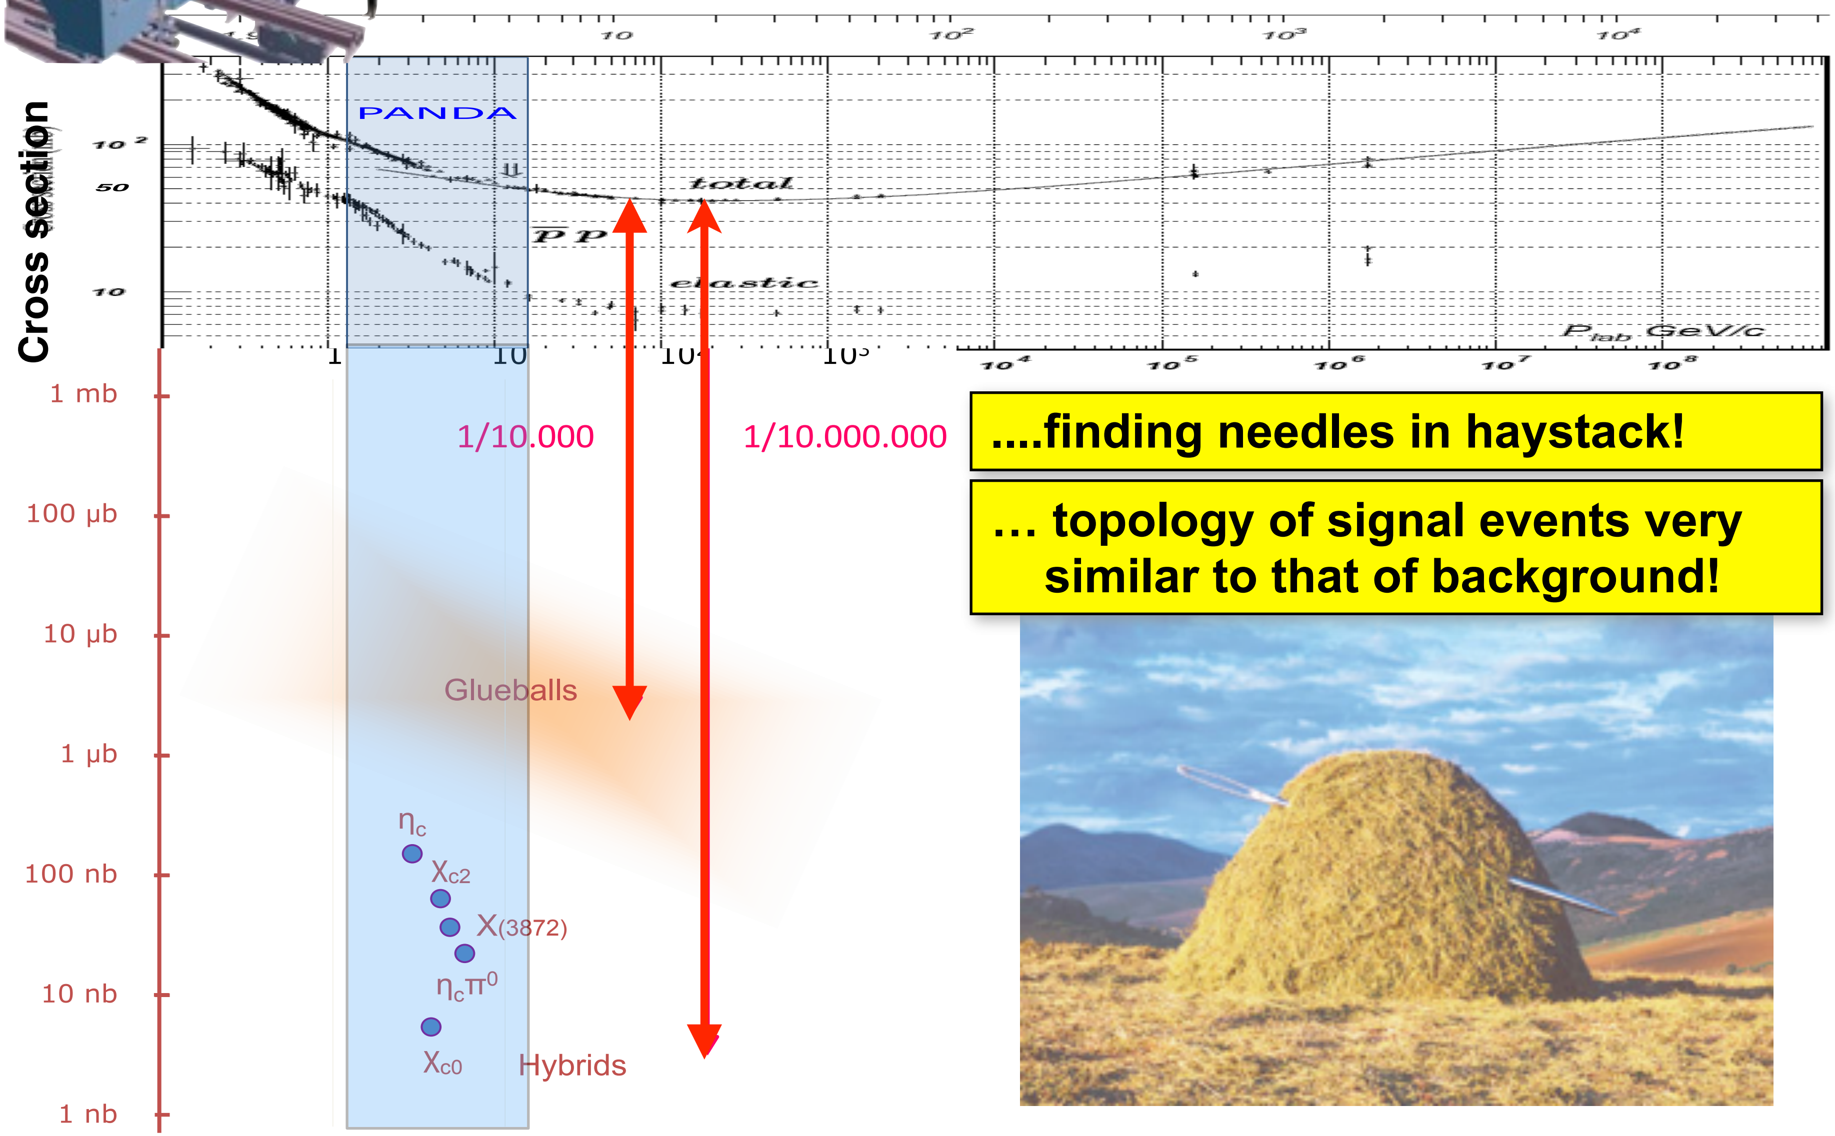

....finding needles in haystack!

… topology of signal events very similar to that of background!

The "physics" challenge! (PANDA as an example)

Klaus Goetzen (GSI)

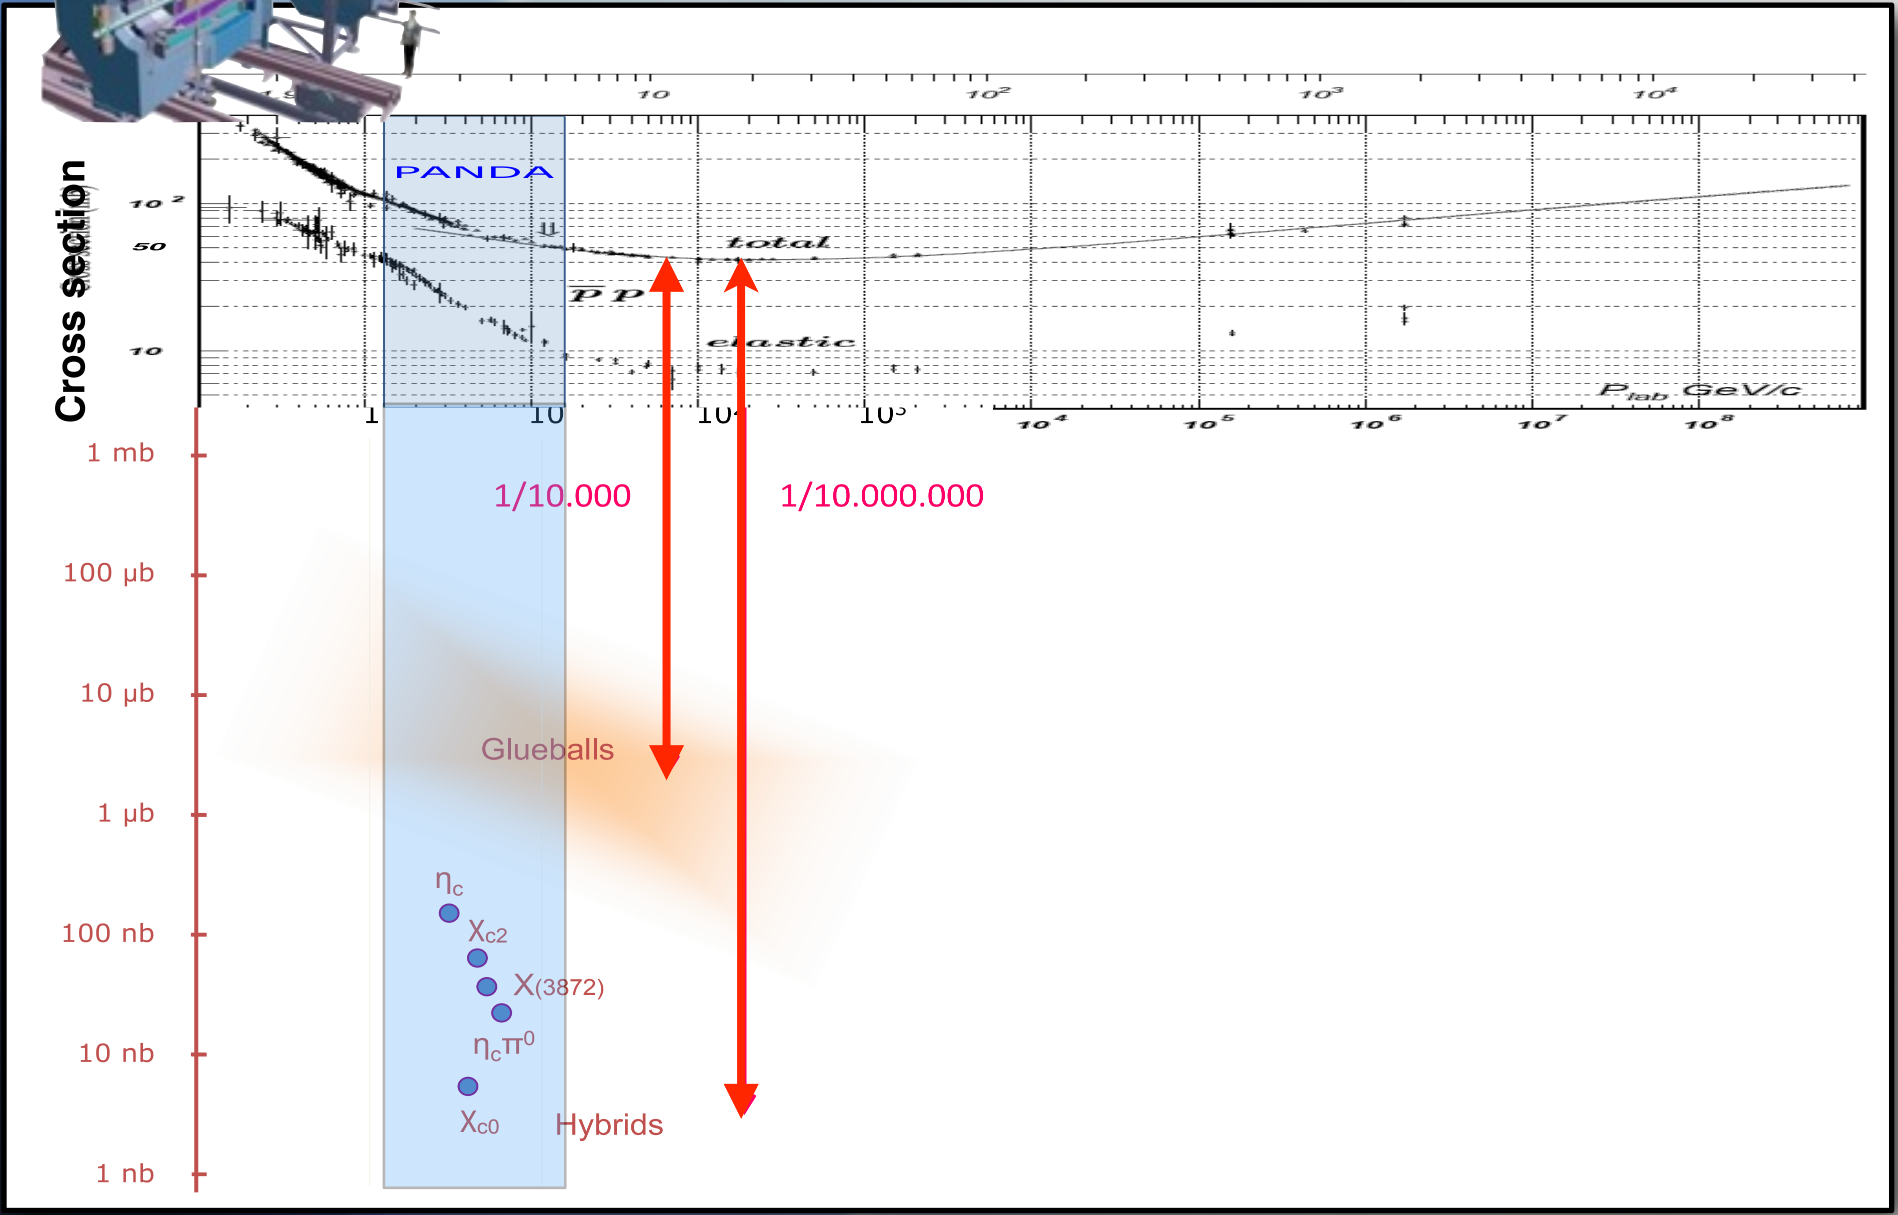

The "physics" challenge! (PANDA as an example)

Klaus Goetzen (GSI)

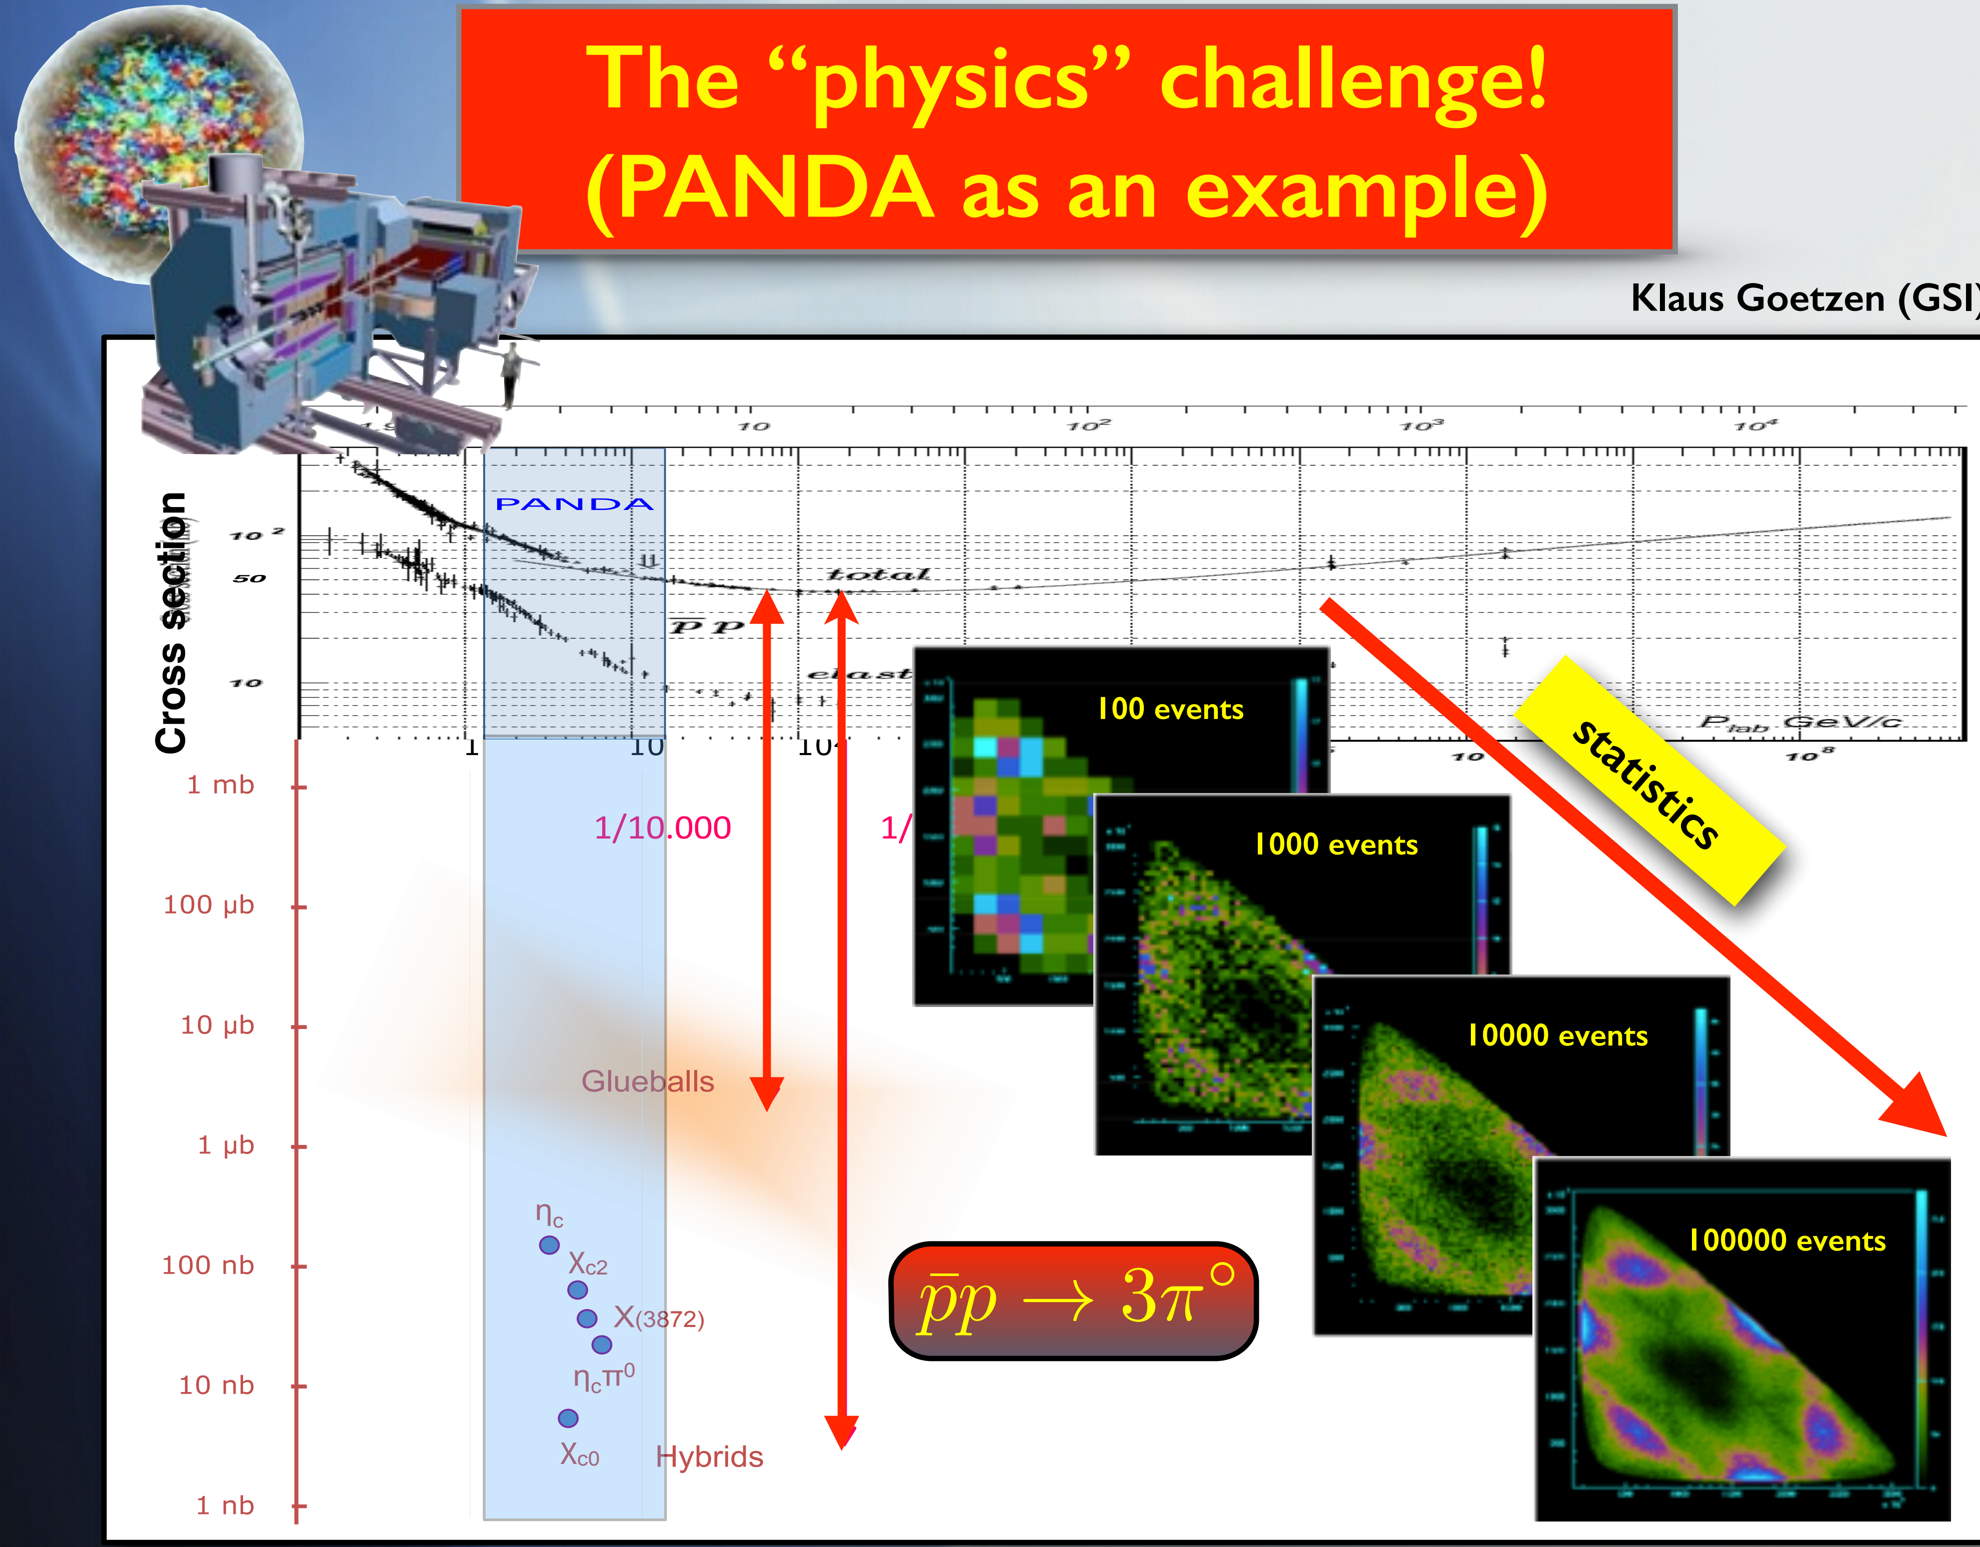

Versatile experiments

Era of large-scale facilities and huge collaborative efforts

Experiments aim to cover a broad physics program

Not one "golden" channel

Experiments need to become capable and flexible to accommodate versatility

Trigger-less Data Acquisition

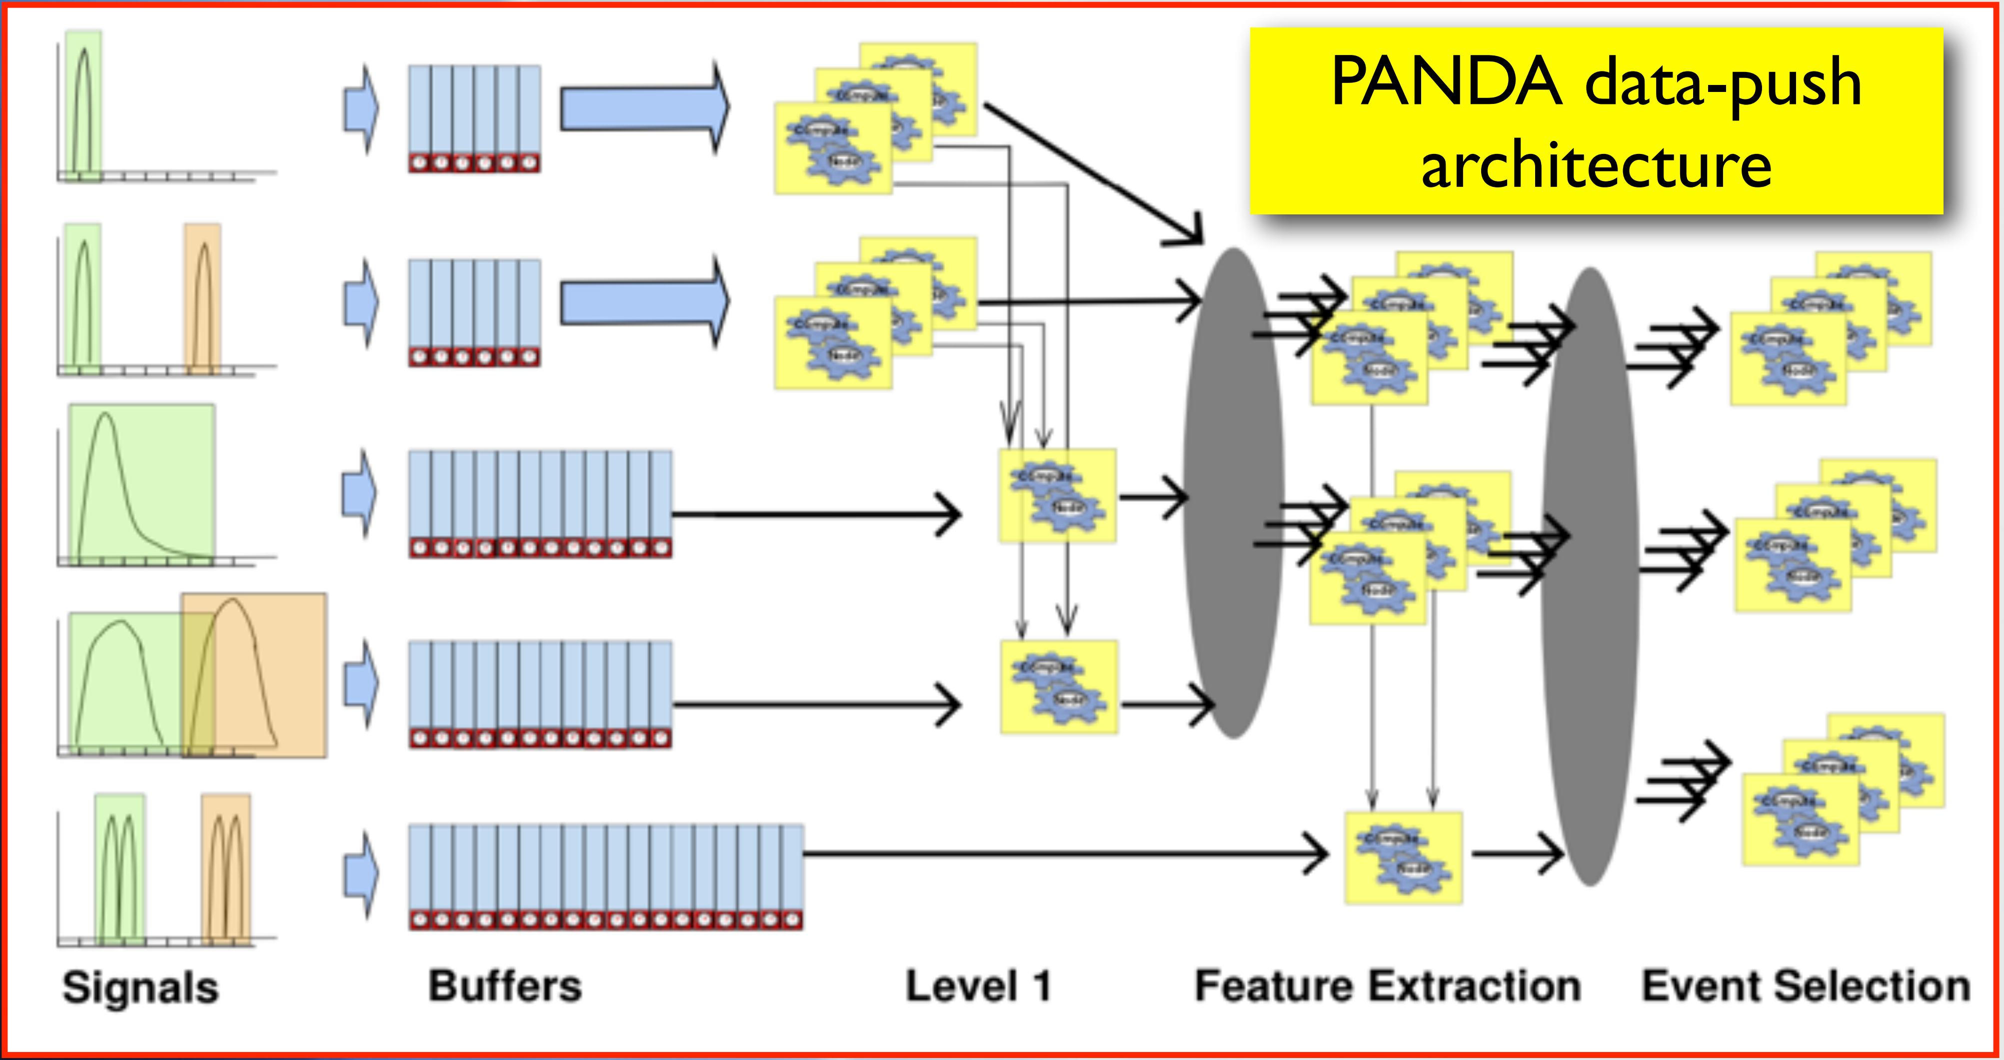

continuous data sampling with self-triggered detector front-end ---> flexible event selection

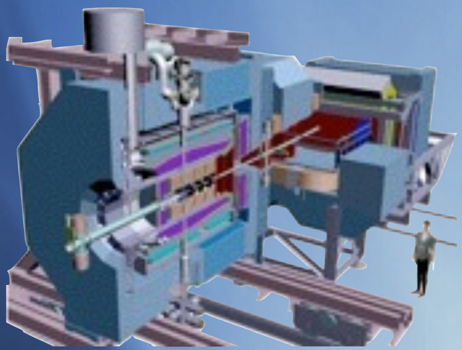

10^7 /sec.

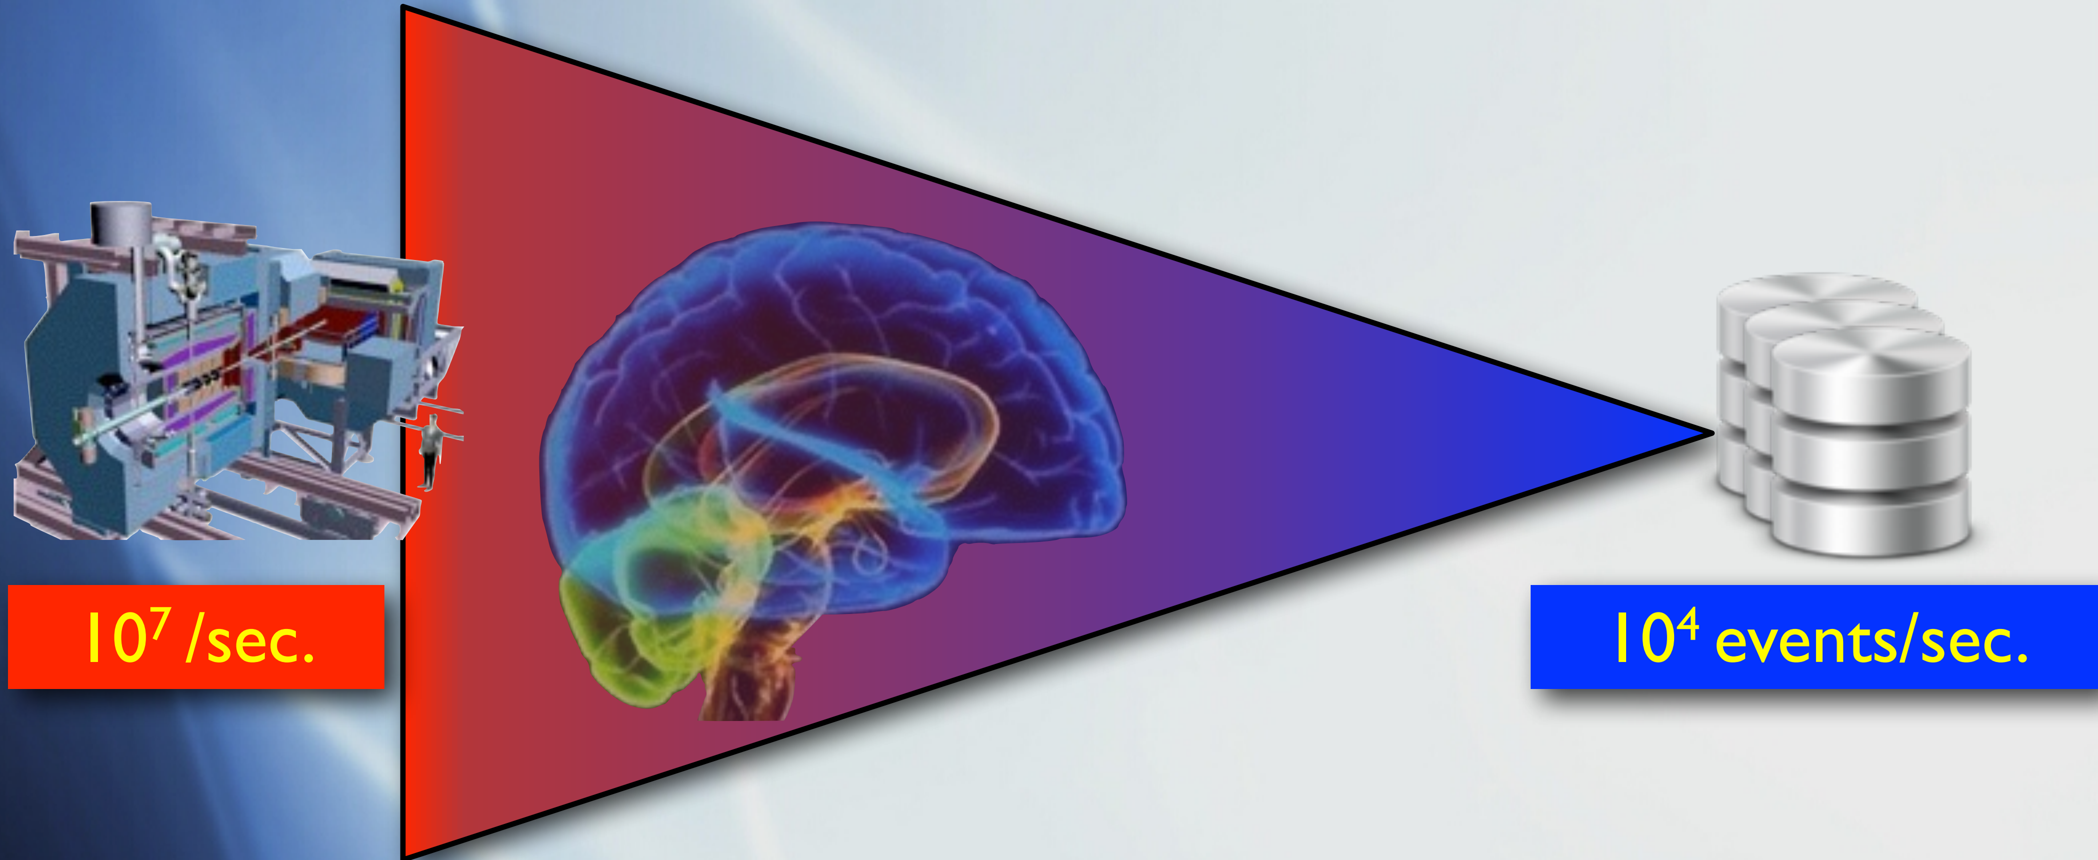
10^7 /sec.
10^4 events/sec.

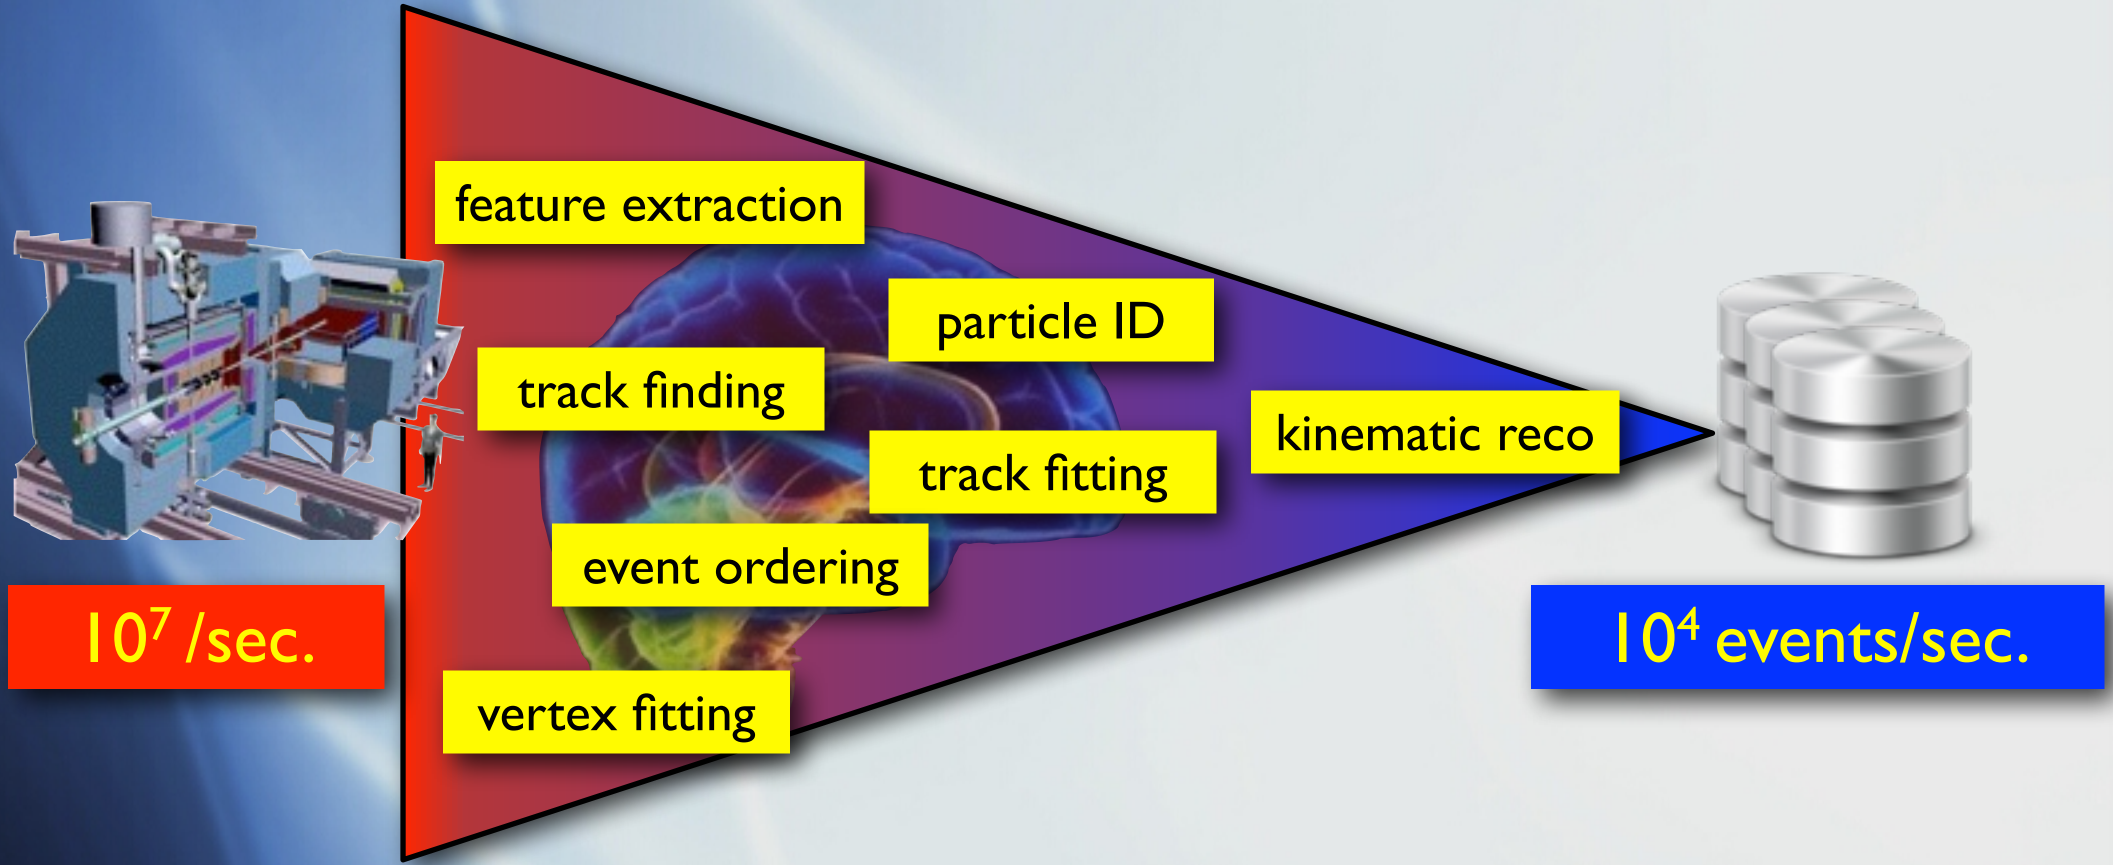
feature extraction
particle ID
track finding
kinematic reco
track fitting
event ordering
vertex fitting
10^7 /sec.
10^4 events/sec.

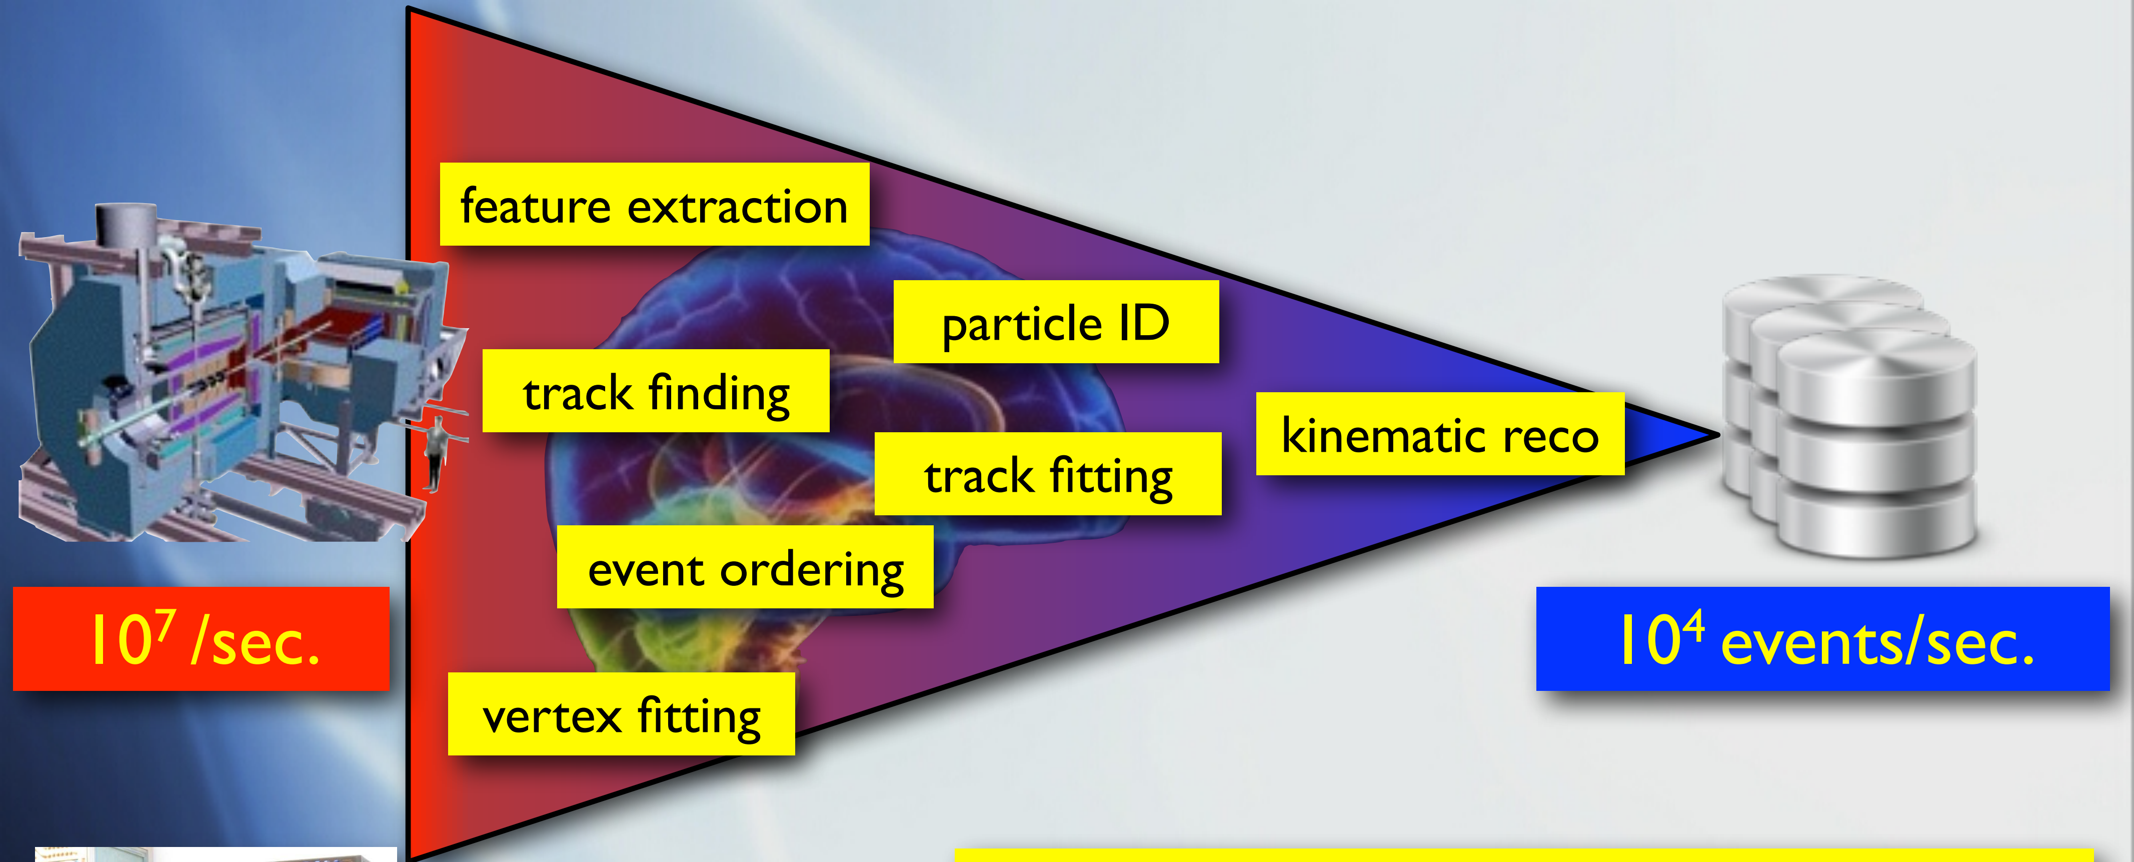

In-situ event reconstruction?

Yes, no other choice! But …

How to implement on hardware?

Which hardware to use?

Streaming data, pileup?

How to adapt existing code?

How to simulate online processing?

How to link components together?

“The free lunch is over”

Lets admit! We were spoiled by the frequency-scaling era

Cores don’t get faster, but more parallel! Requires new approaches!

Even then, not all jobs are easy parallelizable (Amdahl’s law)

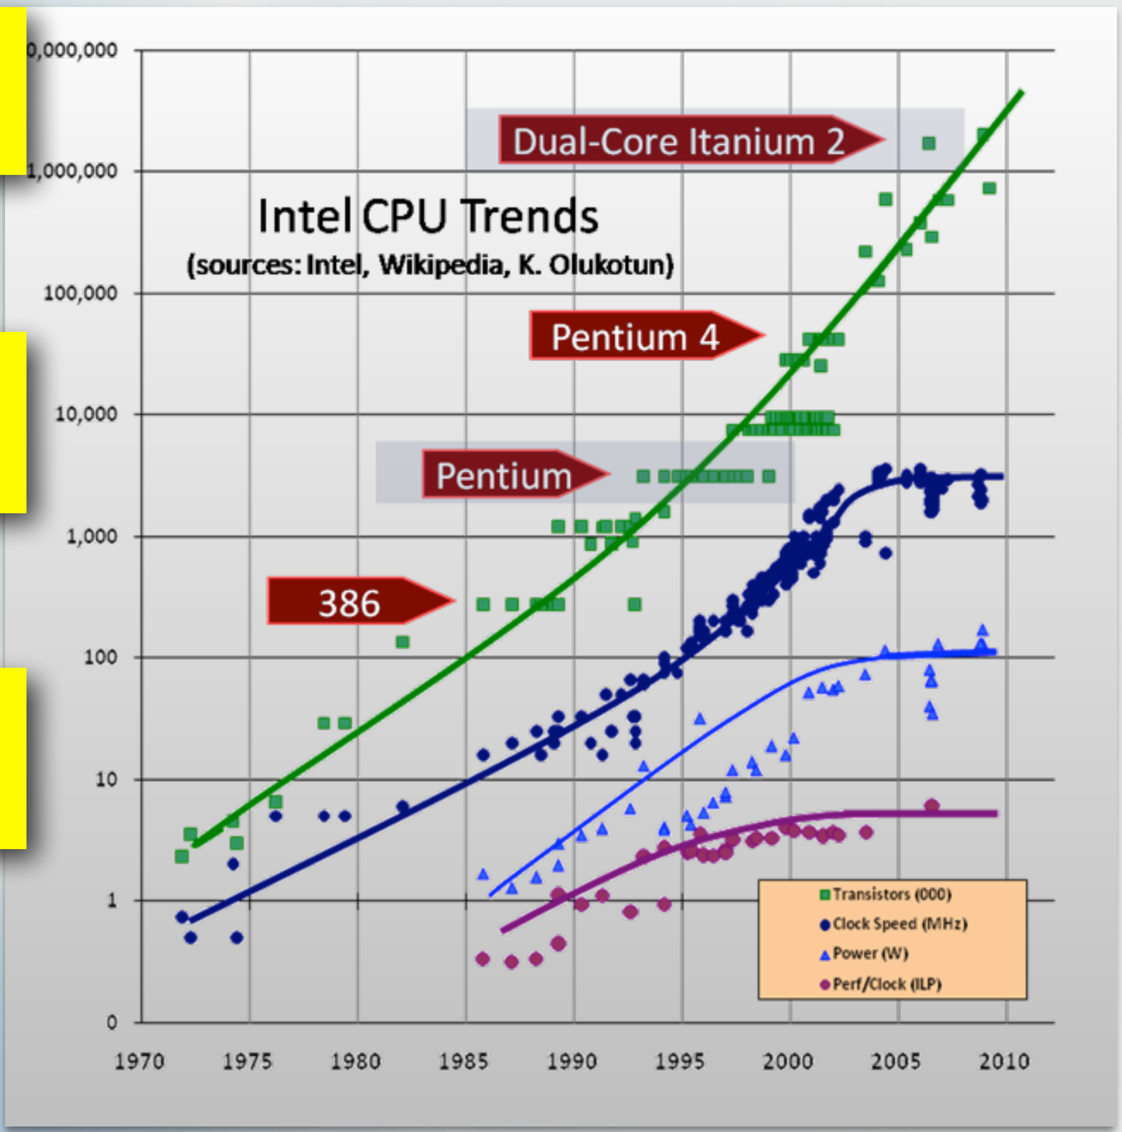

“The free lunch is over”

Lets admit! We were spoiled by the frequency-scaling era

Cores don’t get faster, but more parallel! Requires new approaches!

Even then, not all jobs are easy parallelizable (Amdahl’s law)

Innovative ideas needed!

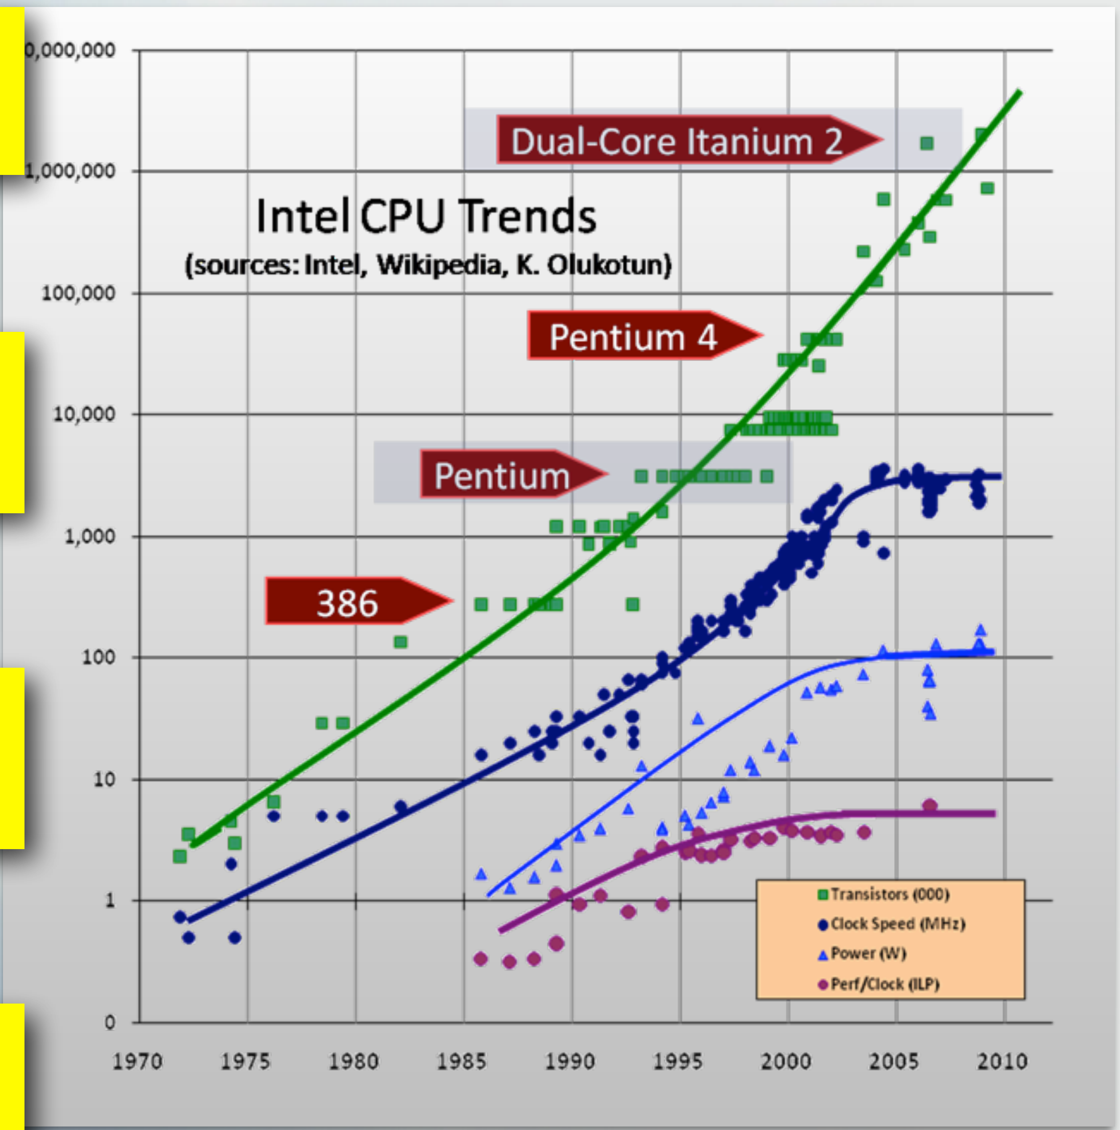

“Intelligence” in practise: GPUPWA

Niklaus Berger

Likelihood, given n data points at Ω_i

$$\mathcal{L} \propto \prod_{i=1}^{n} \frac{I(\Omega_i)}{\int \eta(\Omega) I(\Omega) d\Omega}$$

Product over data events

Normalisation integral over phase space

Detection efficiency

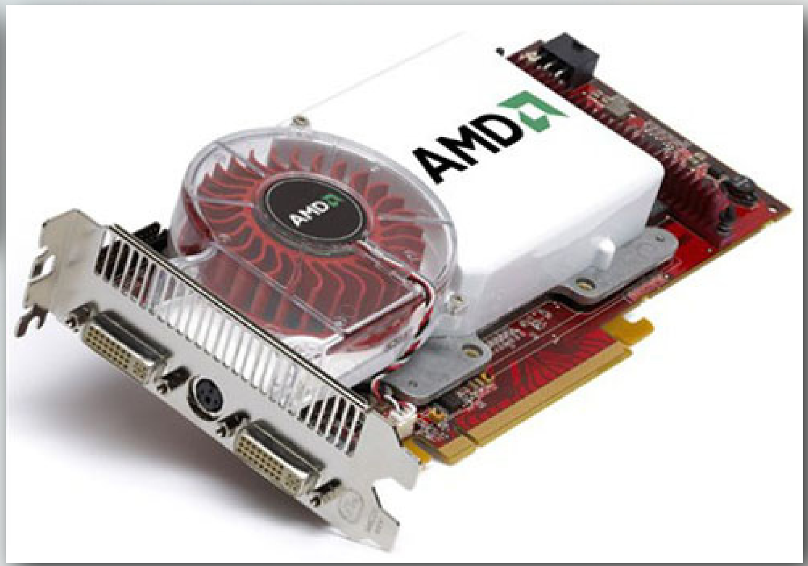

Resonance searches with BESIII data

Partial-wave analyses for a **huge** data samples are extremely time consuming

Innovative idea! Use SIMT capabilities on cheap graphic cards

Massive parallel floating point operations and lookup tables ideally suited for the fitting problem

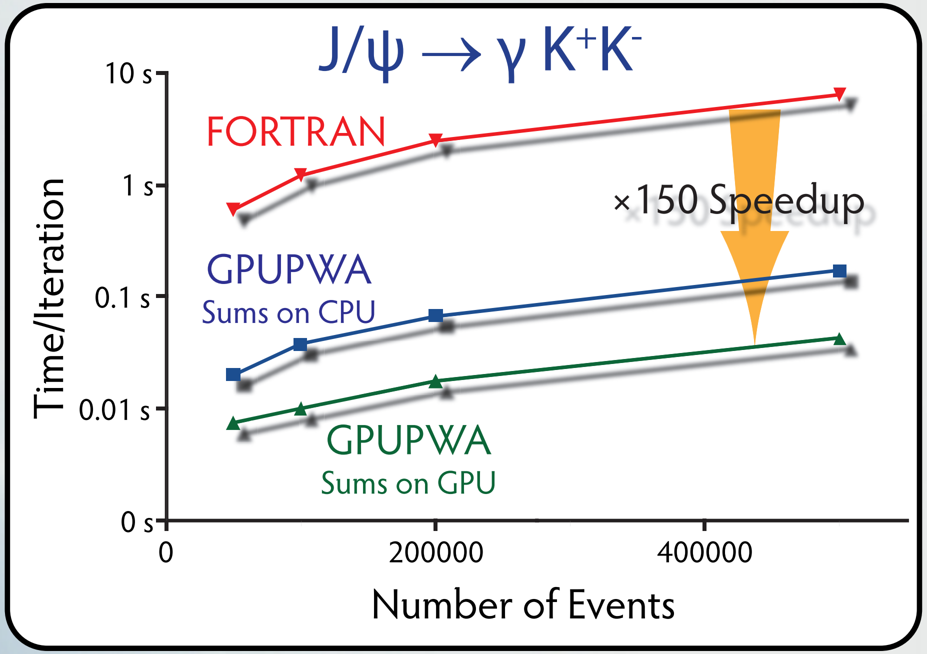

Towards real-time applications

Track finding & fitting using **GPUs**

Inner tracker of PANDA:

- Hough transformations (2D)
- Riemann track finder (3D)
- Triplet finders
- ….

Work in progress…

Andreas Herten, Tobias Stockmanns (FZJ)

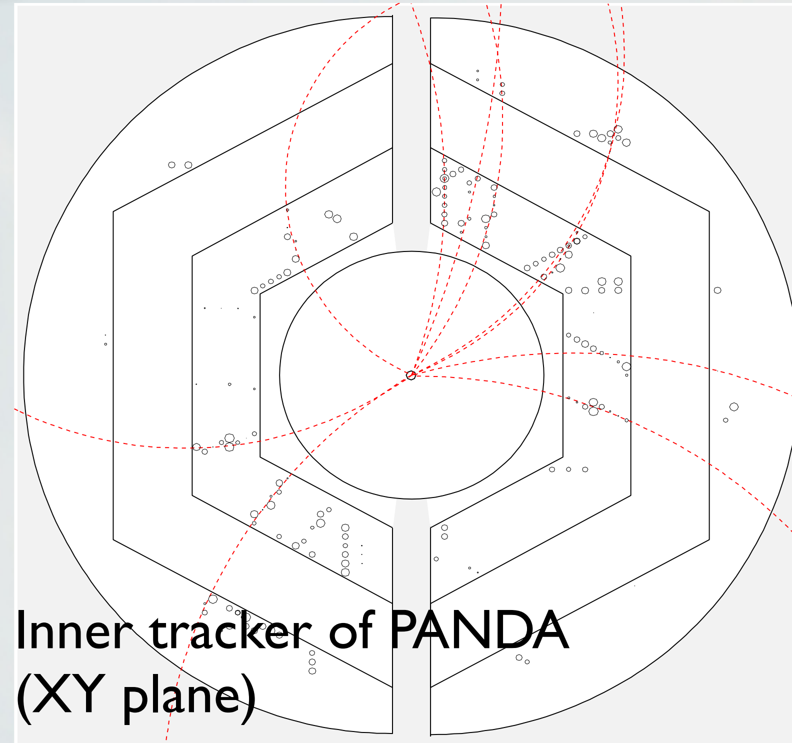

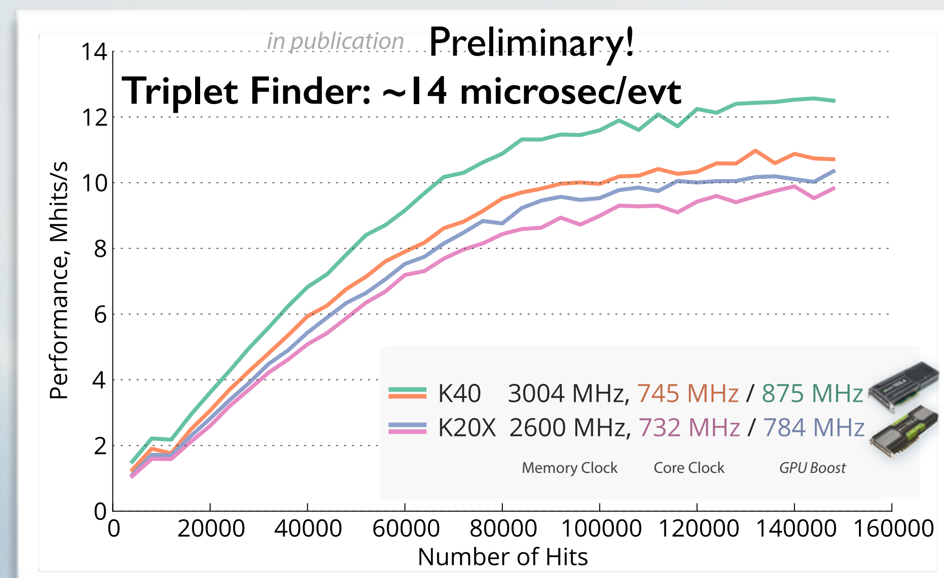

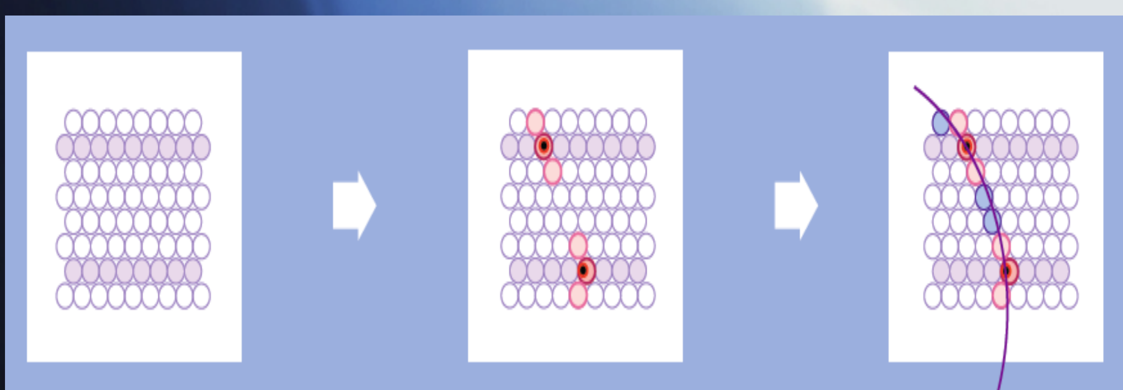

Towards real-time applications

Track finding & fitting using **FPGAs**

VHDL: Very-high-speed integrated circuits Hardware Description Language (dataflow language)

Helix fitter for tracker of PANDA,
1st iteration only :
260 microseconds/event

Yutie Liang (Uni Giessen)

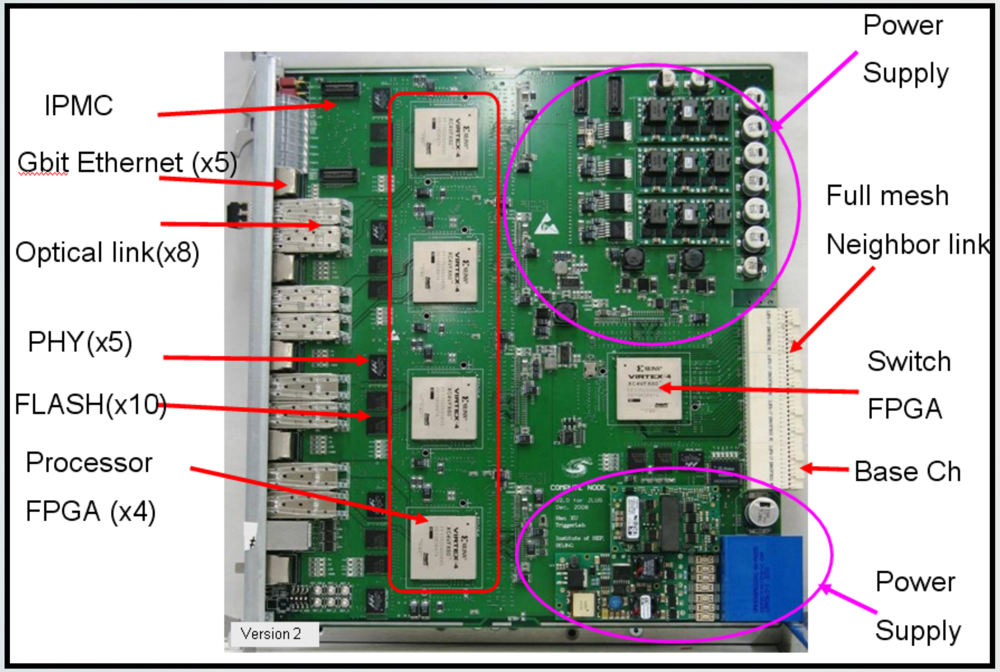

High Performance and scalable **Compute Node**

Device Utilization Summary				
Logic Utilization	**Used**	**Available**	**Utilization**	**Note(s)**
Number of Slice Flip Flops	17,187	50,560	33%	
DCM autocalibration logic	14	17,187	1%	
Number of 4 input LUTs	27,438	50,560	54%	
DCM autocalibration logic	8	27,438	1%	
Number of occupied Slices	18,159	25,280	71%	

5 FPGAs:Virtex-4 FX60
10Gb DDR2 RAM
~32 Gbps bandwidth
2 PowerPCs
ATCA compliant

see talks by Yutie Liang (PANDA) and Jan Michel (CBM) on Thursday afternoon in Tier 6

Consolidating beyond our communities!

Mohammad Babai, Michael Wilkinson, JM (University of Groningen)

Image data are everywhere!

Collaboration with Institute for Mathematics and Computing Science at University Groningen

Vision-based algorithm for track finding for fast track filtering

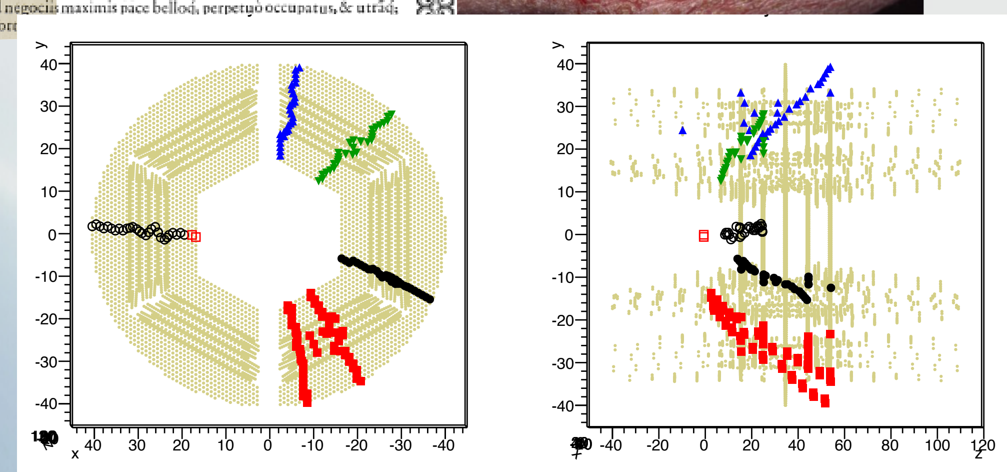

Modern computing is all about FRAMEWORKS

Providing access to technologies

Details hidden for the "user"

Modular, virtual, and versatile

Modern computing is all about FRAMEWORKS

Providing access to technologies

Details hidden for the “user”

Modular, virtual, and versatile

Collaborative developments

Human resource friendly

Geant 4

Example of a success story: FAIRROOT

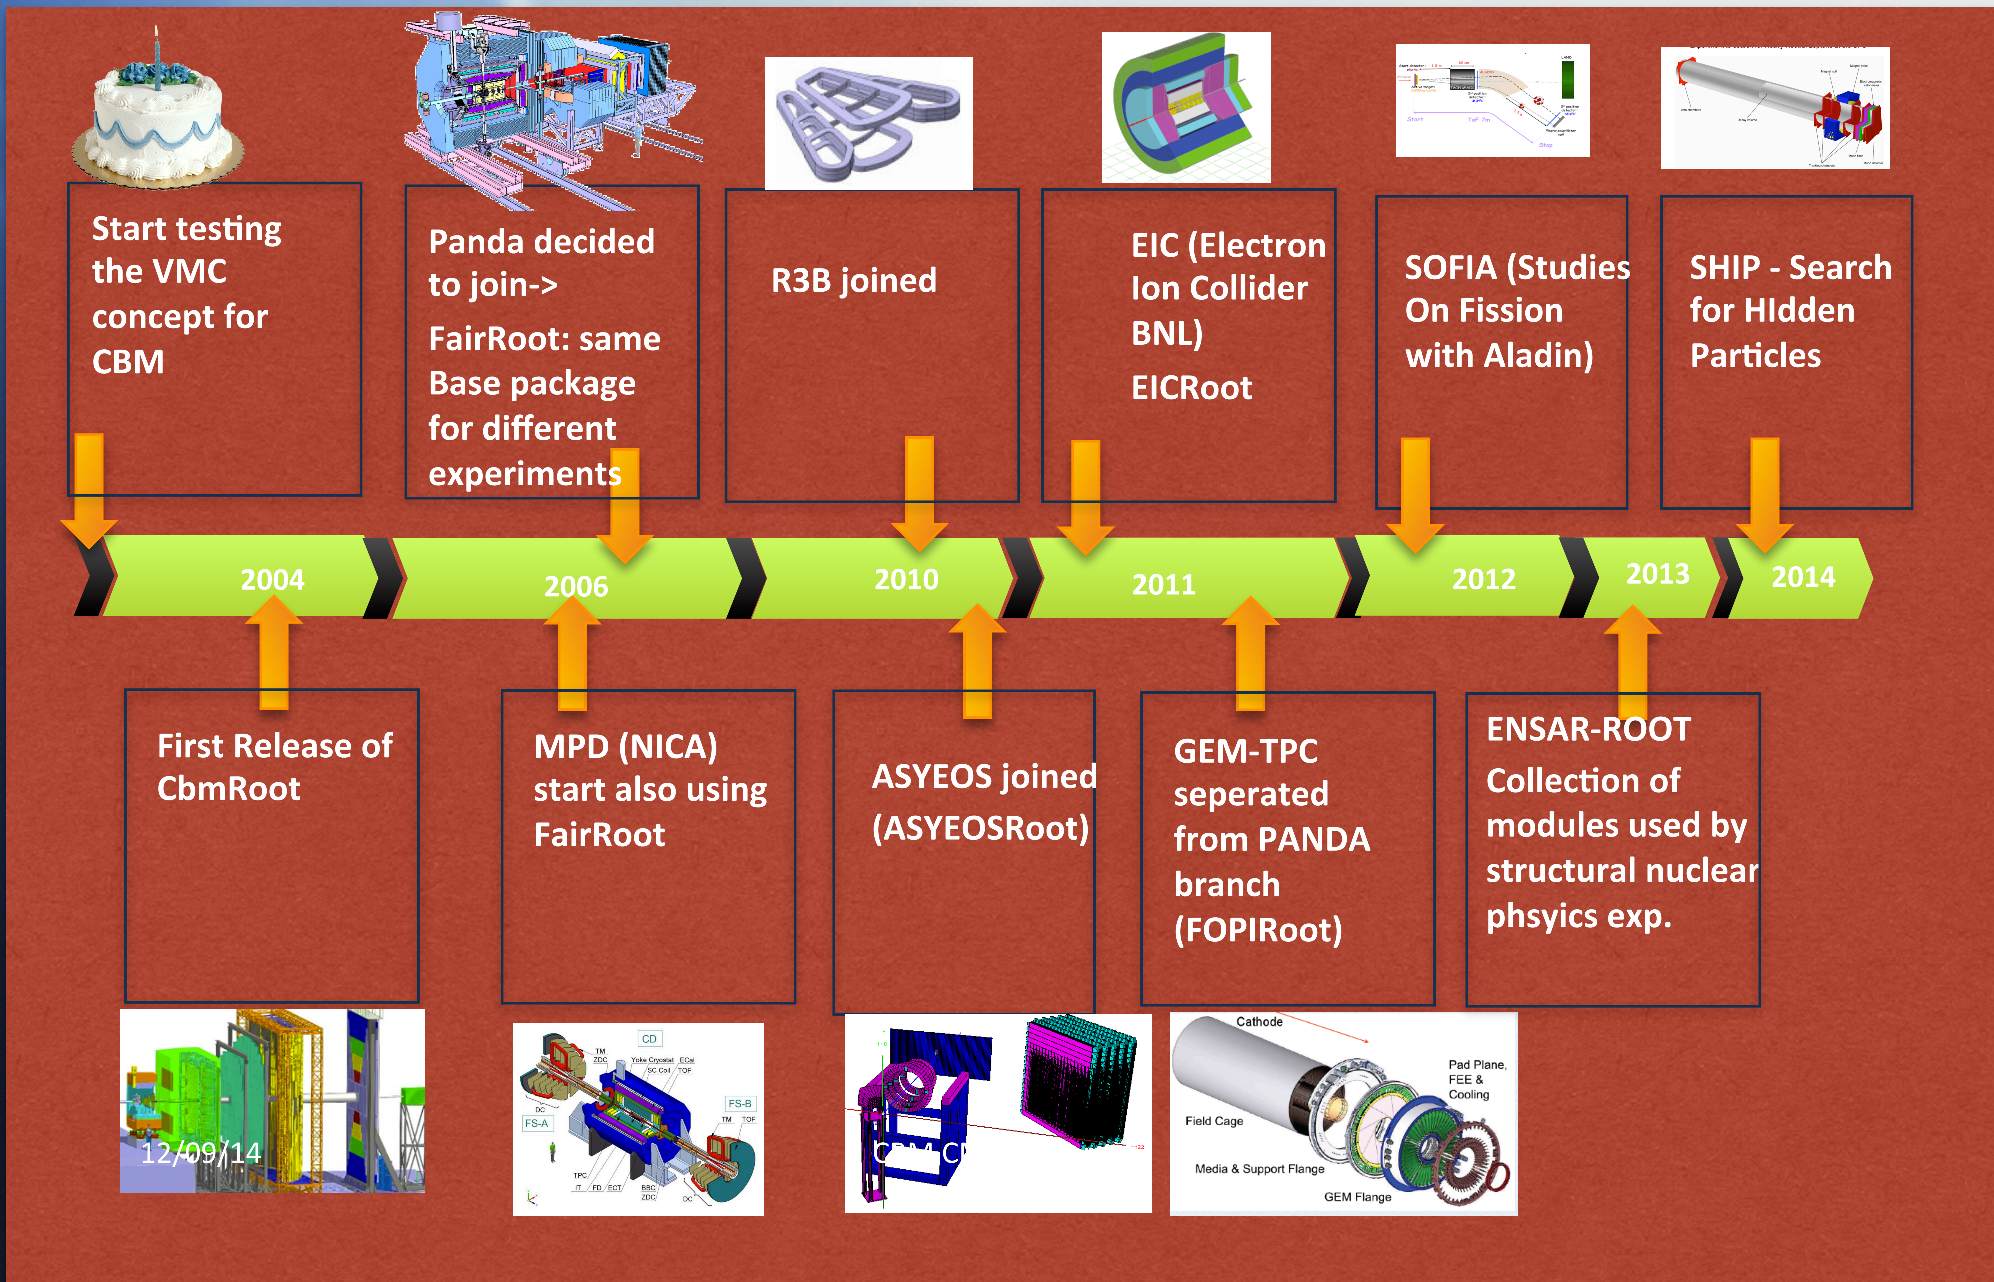

Mohammad Al-Turany, et al.

Example of a success story: FAIRROOT

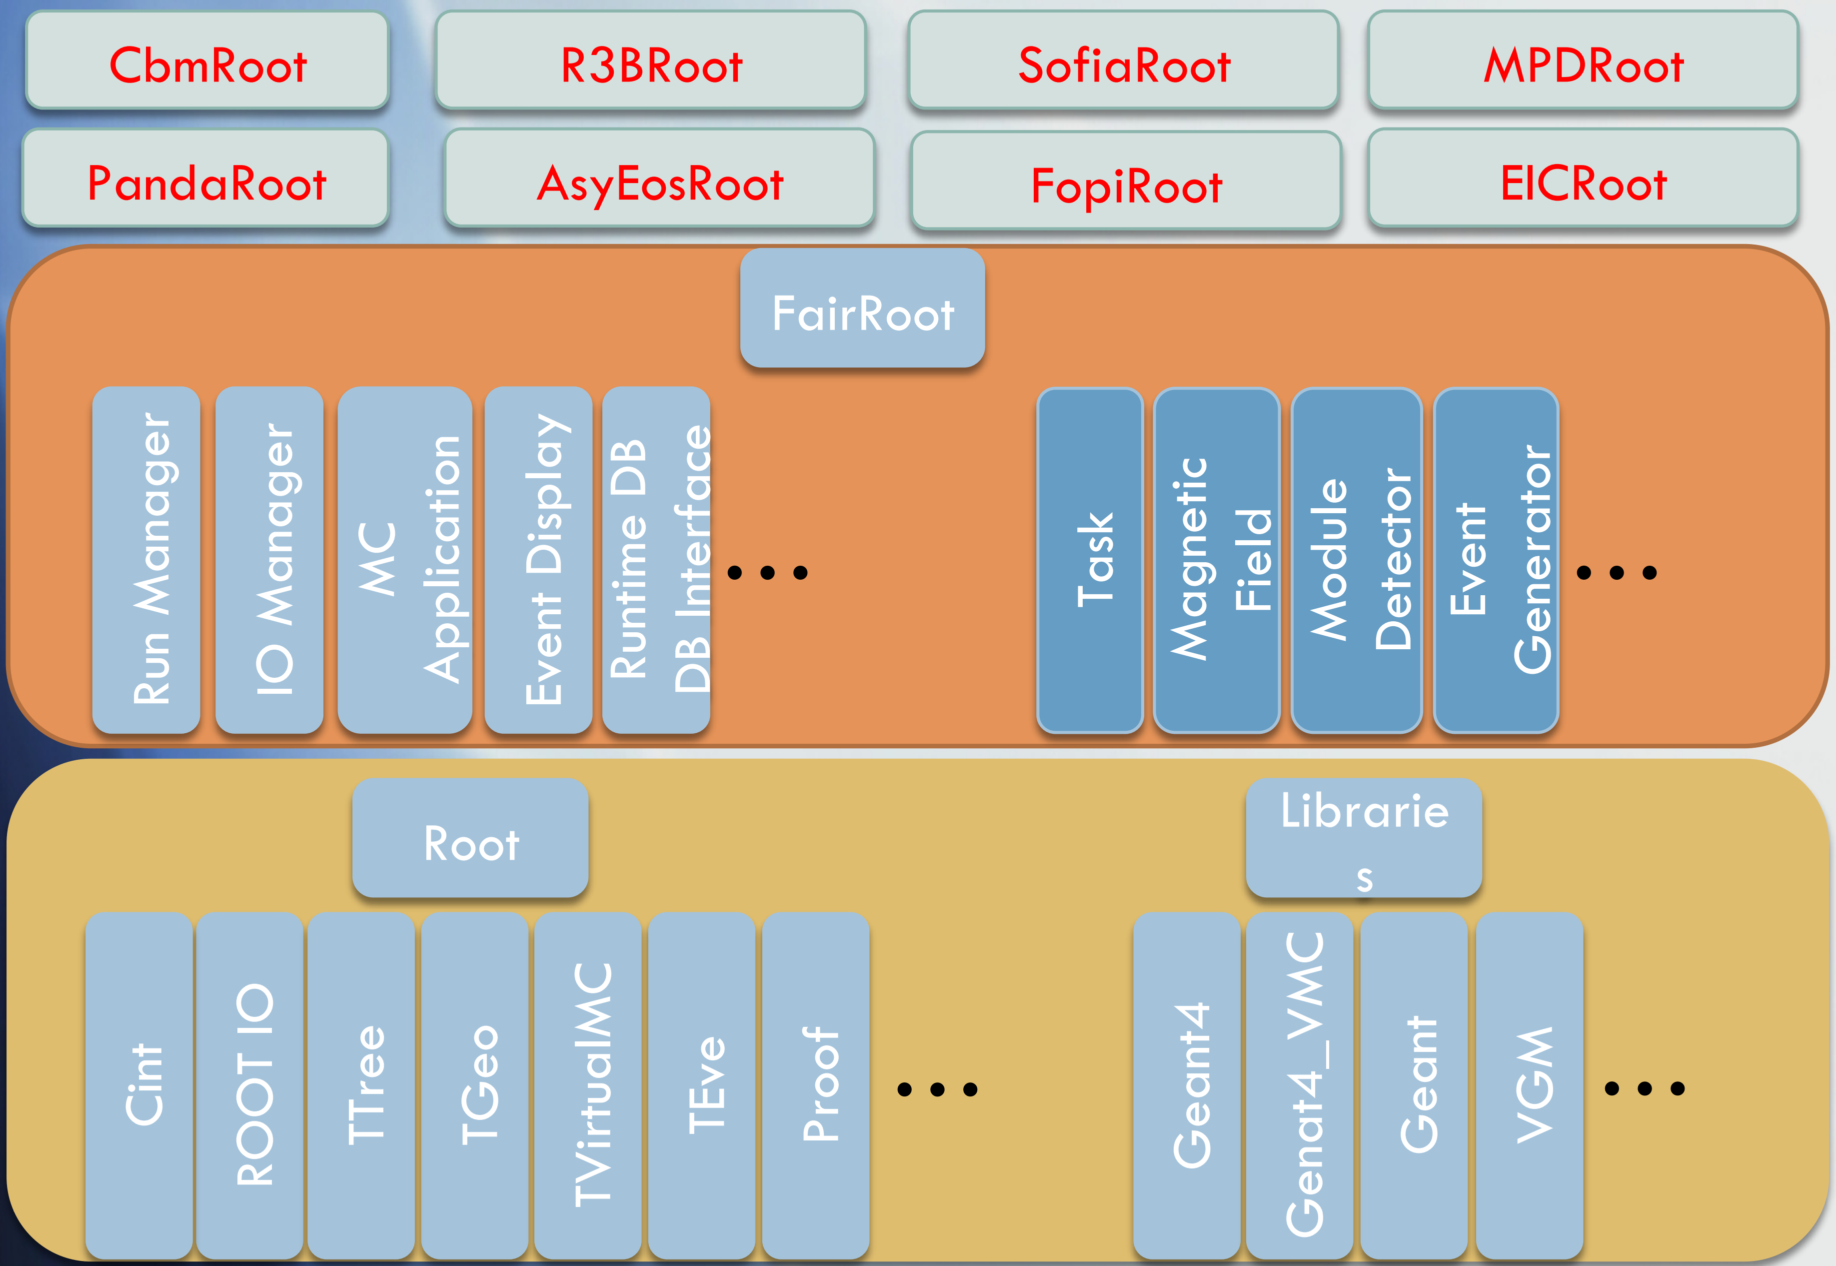

Mohammad Al-Turany, et al.

2013: FAIRROOT meets ALICE O^2

2013: FAIRROOT meets ALICE O^2

ALICE

ALICE O^2:

- DAQ, online & offline with one framework

Mohammad Al-Turany, et al.

2013: FAIRROOT meets ALICE O^2

ALICE O^2:

- DAQ, online & offline with one framework

FAIRROOT:

- Concurrency, merging online and offline

Mohammad Al-Turany, et al.

2013: FAIRROOT meets ALICE O^2

ALICE

FAIR

ALFA

ALICE O^2:

- DAQ, online & offline with one framework

+

FAIRROOT:

- Concurrency, merging online and offline

=

ALFA:

- Join forces in a combined framework!

Mohammad Al-Turany, et al.

2013: FAIRROOT meets ALICE O^2

Why join efforts?

FAIR: testing with real data and existing detectors

ALICE: interested in all the features of FAIRROOT

ALFA

Mohammad Al-Turany, et al.

2013: FAIRROOT meets ALICE O^2

Why join efforts?

FAIR: testing with real data and existing detectors

ALICE: interested in all the features of FAIRROOT

Mission of the framework

Ready for massive data reduction

Tighter coupling between online and offline reconstruction software

ALFA

Mohammad Al-Turany, et al.

Dataflow model based on message queues

- BSD sockets API
- Bindings for 30+ languages
- Lockless and Fast
- Automatic re-connection
- Multiplexed I/O

iMatix Corporation
huge open source community

ALFA

FAIRMQ:

- o Asynchronous messaging toolkit
- o Broad scala of messaging pattern
- o Easy and scalable networking
- o Communication layer: 0MQ, nanoMSG

Mohammad Al-Turany, et al.

Data centers and distributed computing

A common data center
for FAIR (**GreenITCube**)

19 M€ funding from
Helmholtz AI program

12 M€ building cost
7 M€ initial installation

Completion in Q3/2015

Volker Lindenstruth, Jan Trautmann

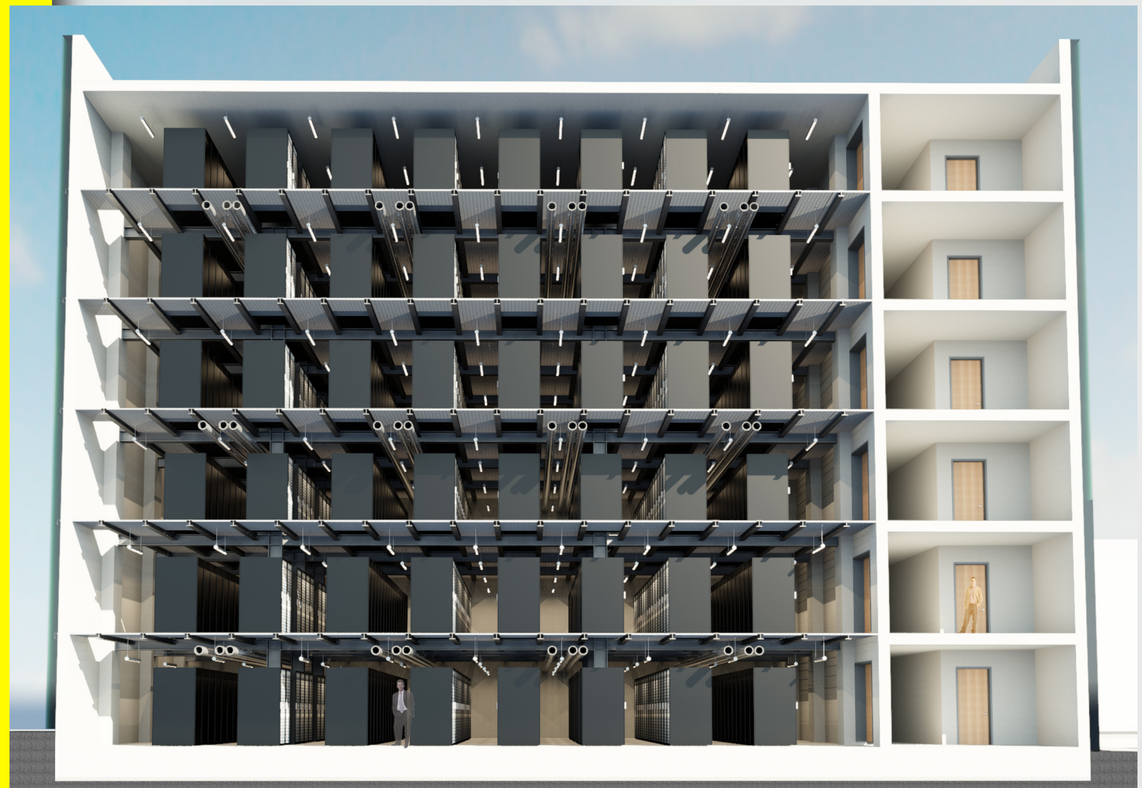

768 19” racks
4 MW cooling (baseline)
Max. cooling power 12 MW
>2300 fibers to exps

Data centers and distributed computing

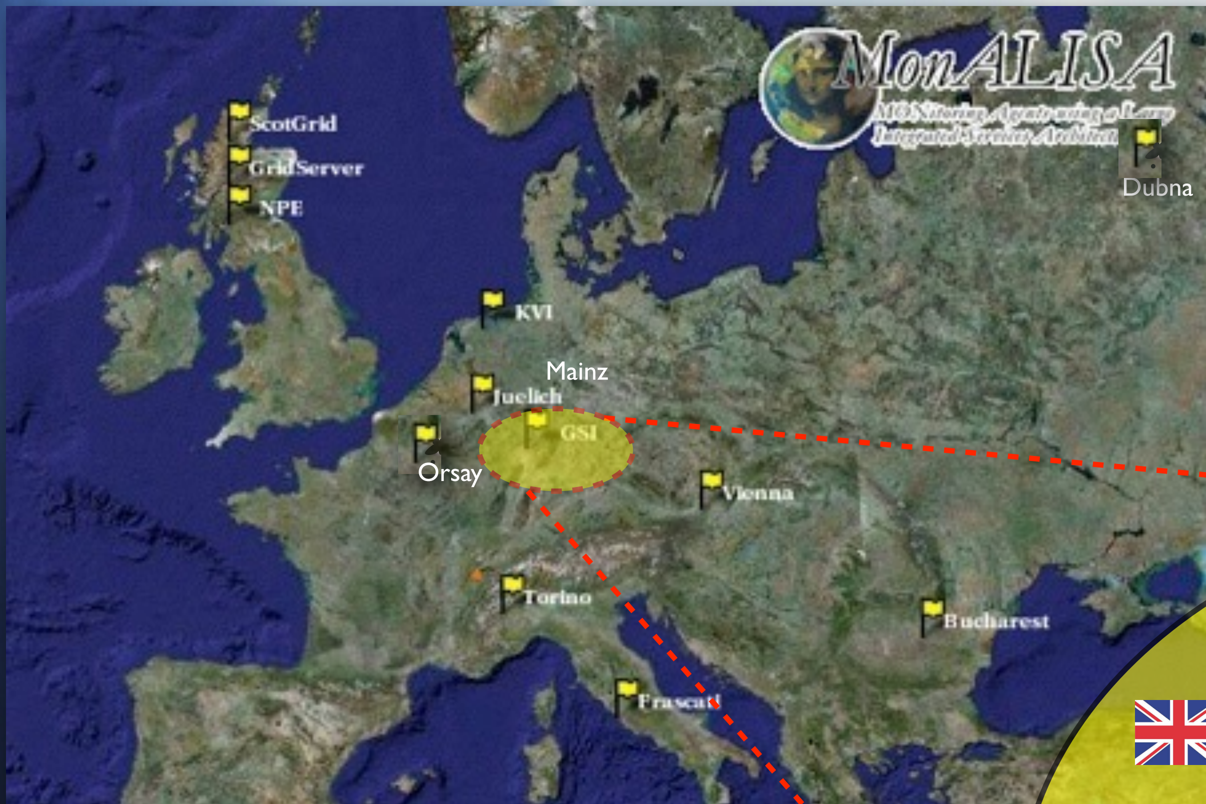

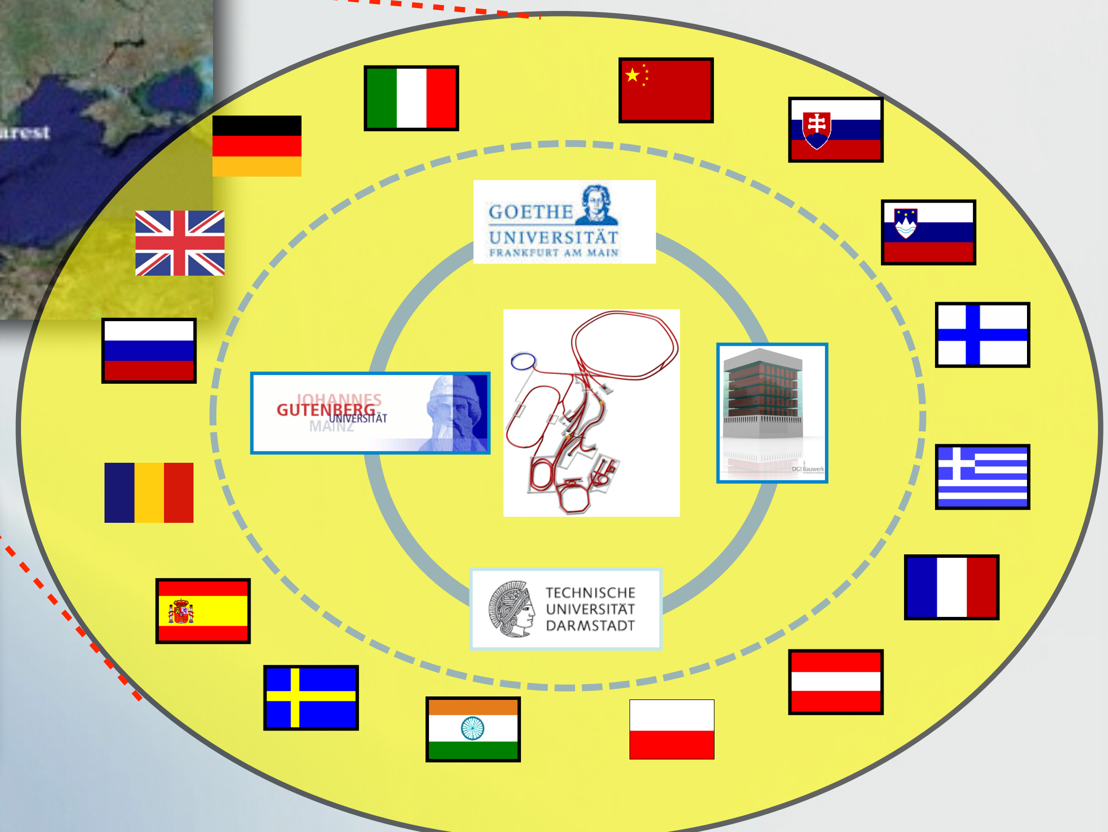

Close integration of HPC computing centers into FAIR T0/T1: Teralink network

Data centers and distributed computing

Distributed computing

- o K.I.S.S. principle
- o small number of large centers
- o possibly heterogeneous
- o usage of standard protocols

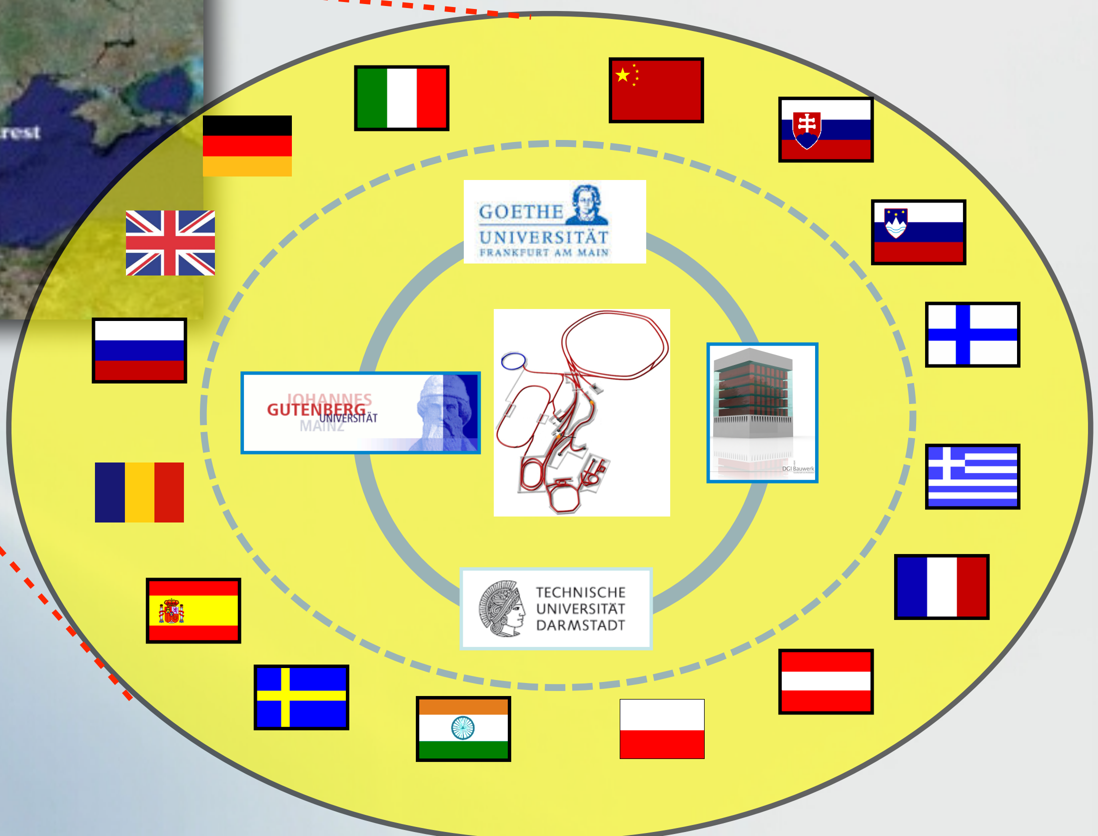

Close integration of HPC computing centers into FAIR
T0/T1: Teralink network

Data, data, and even more data

Data, data, and even more data

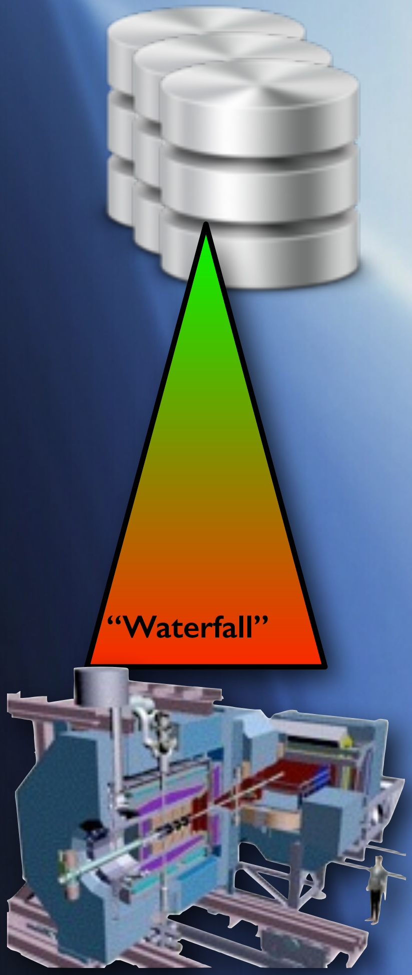

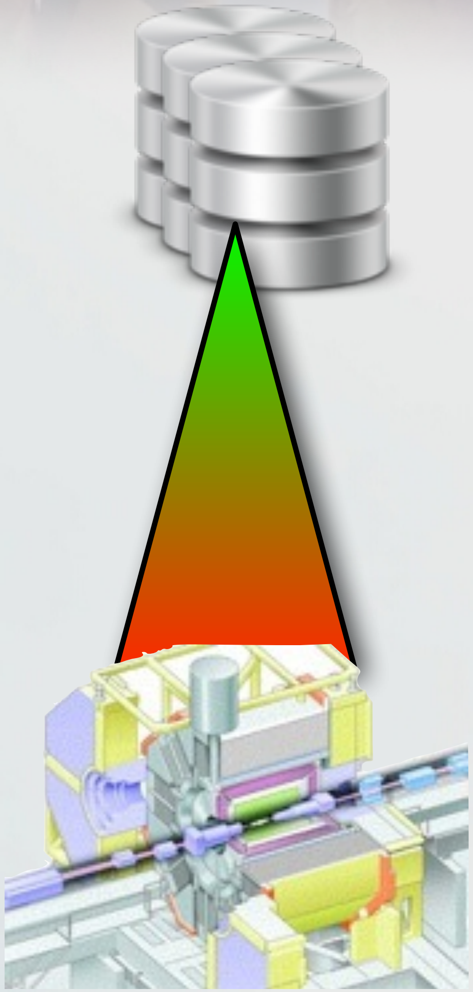

Data, data, and even more data

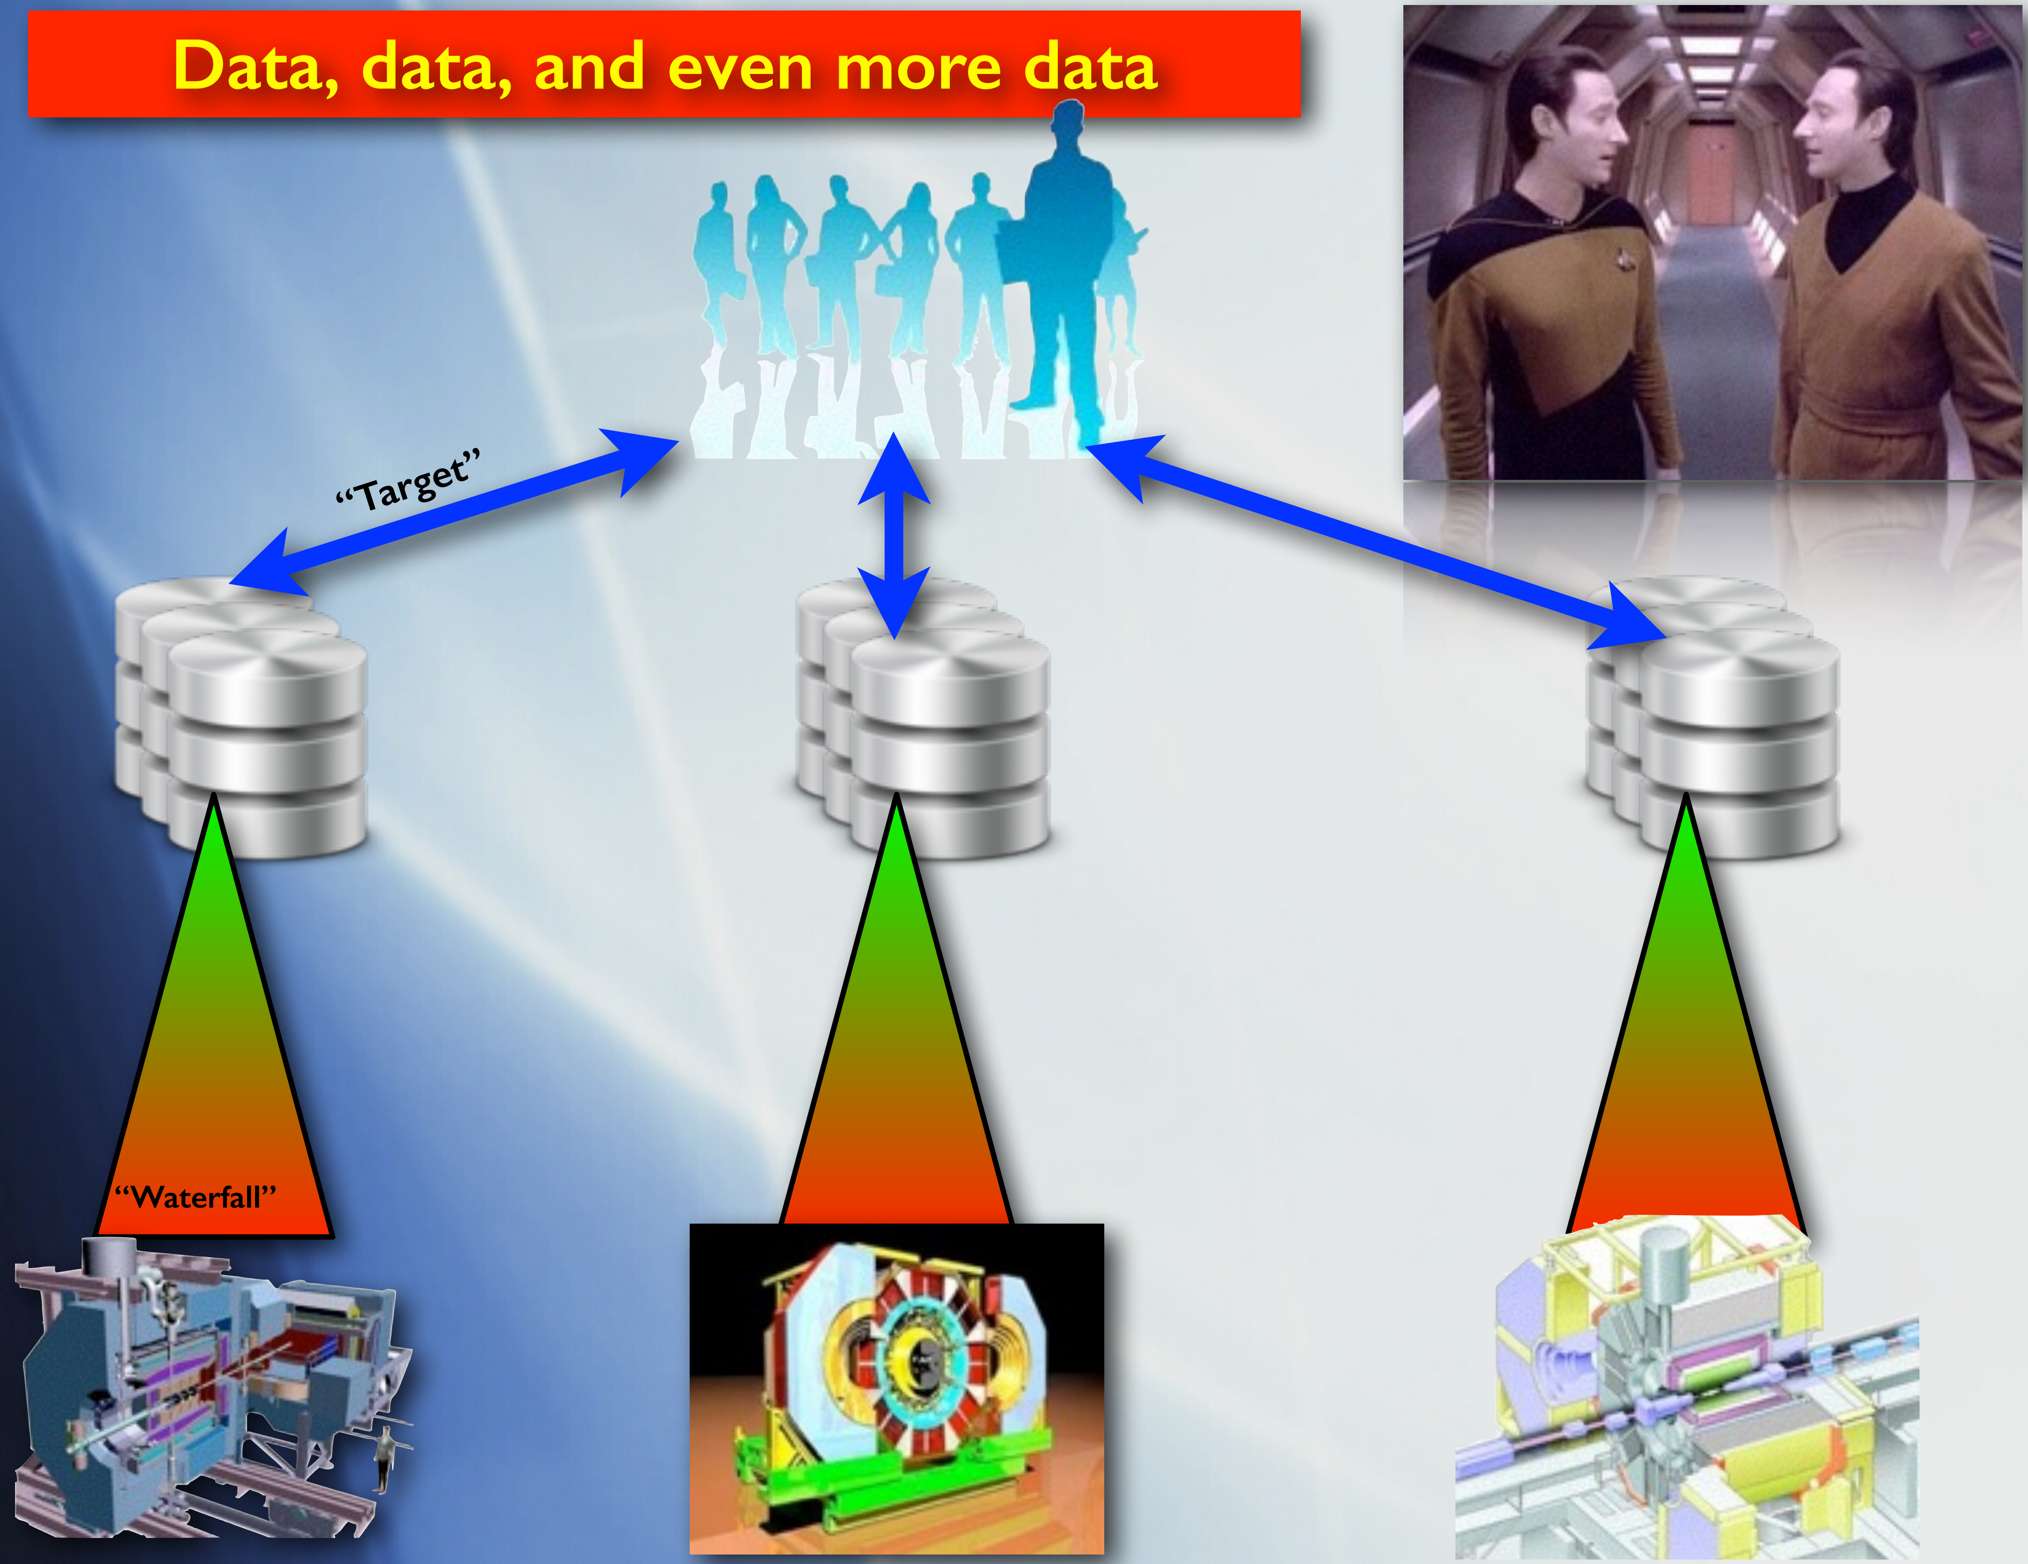

Data, data, and even more data

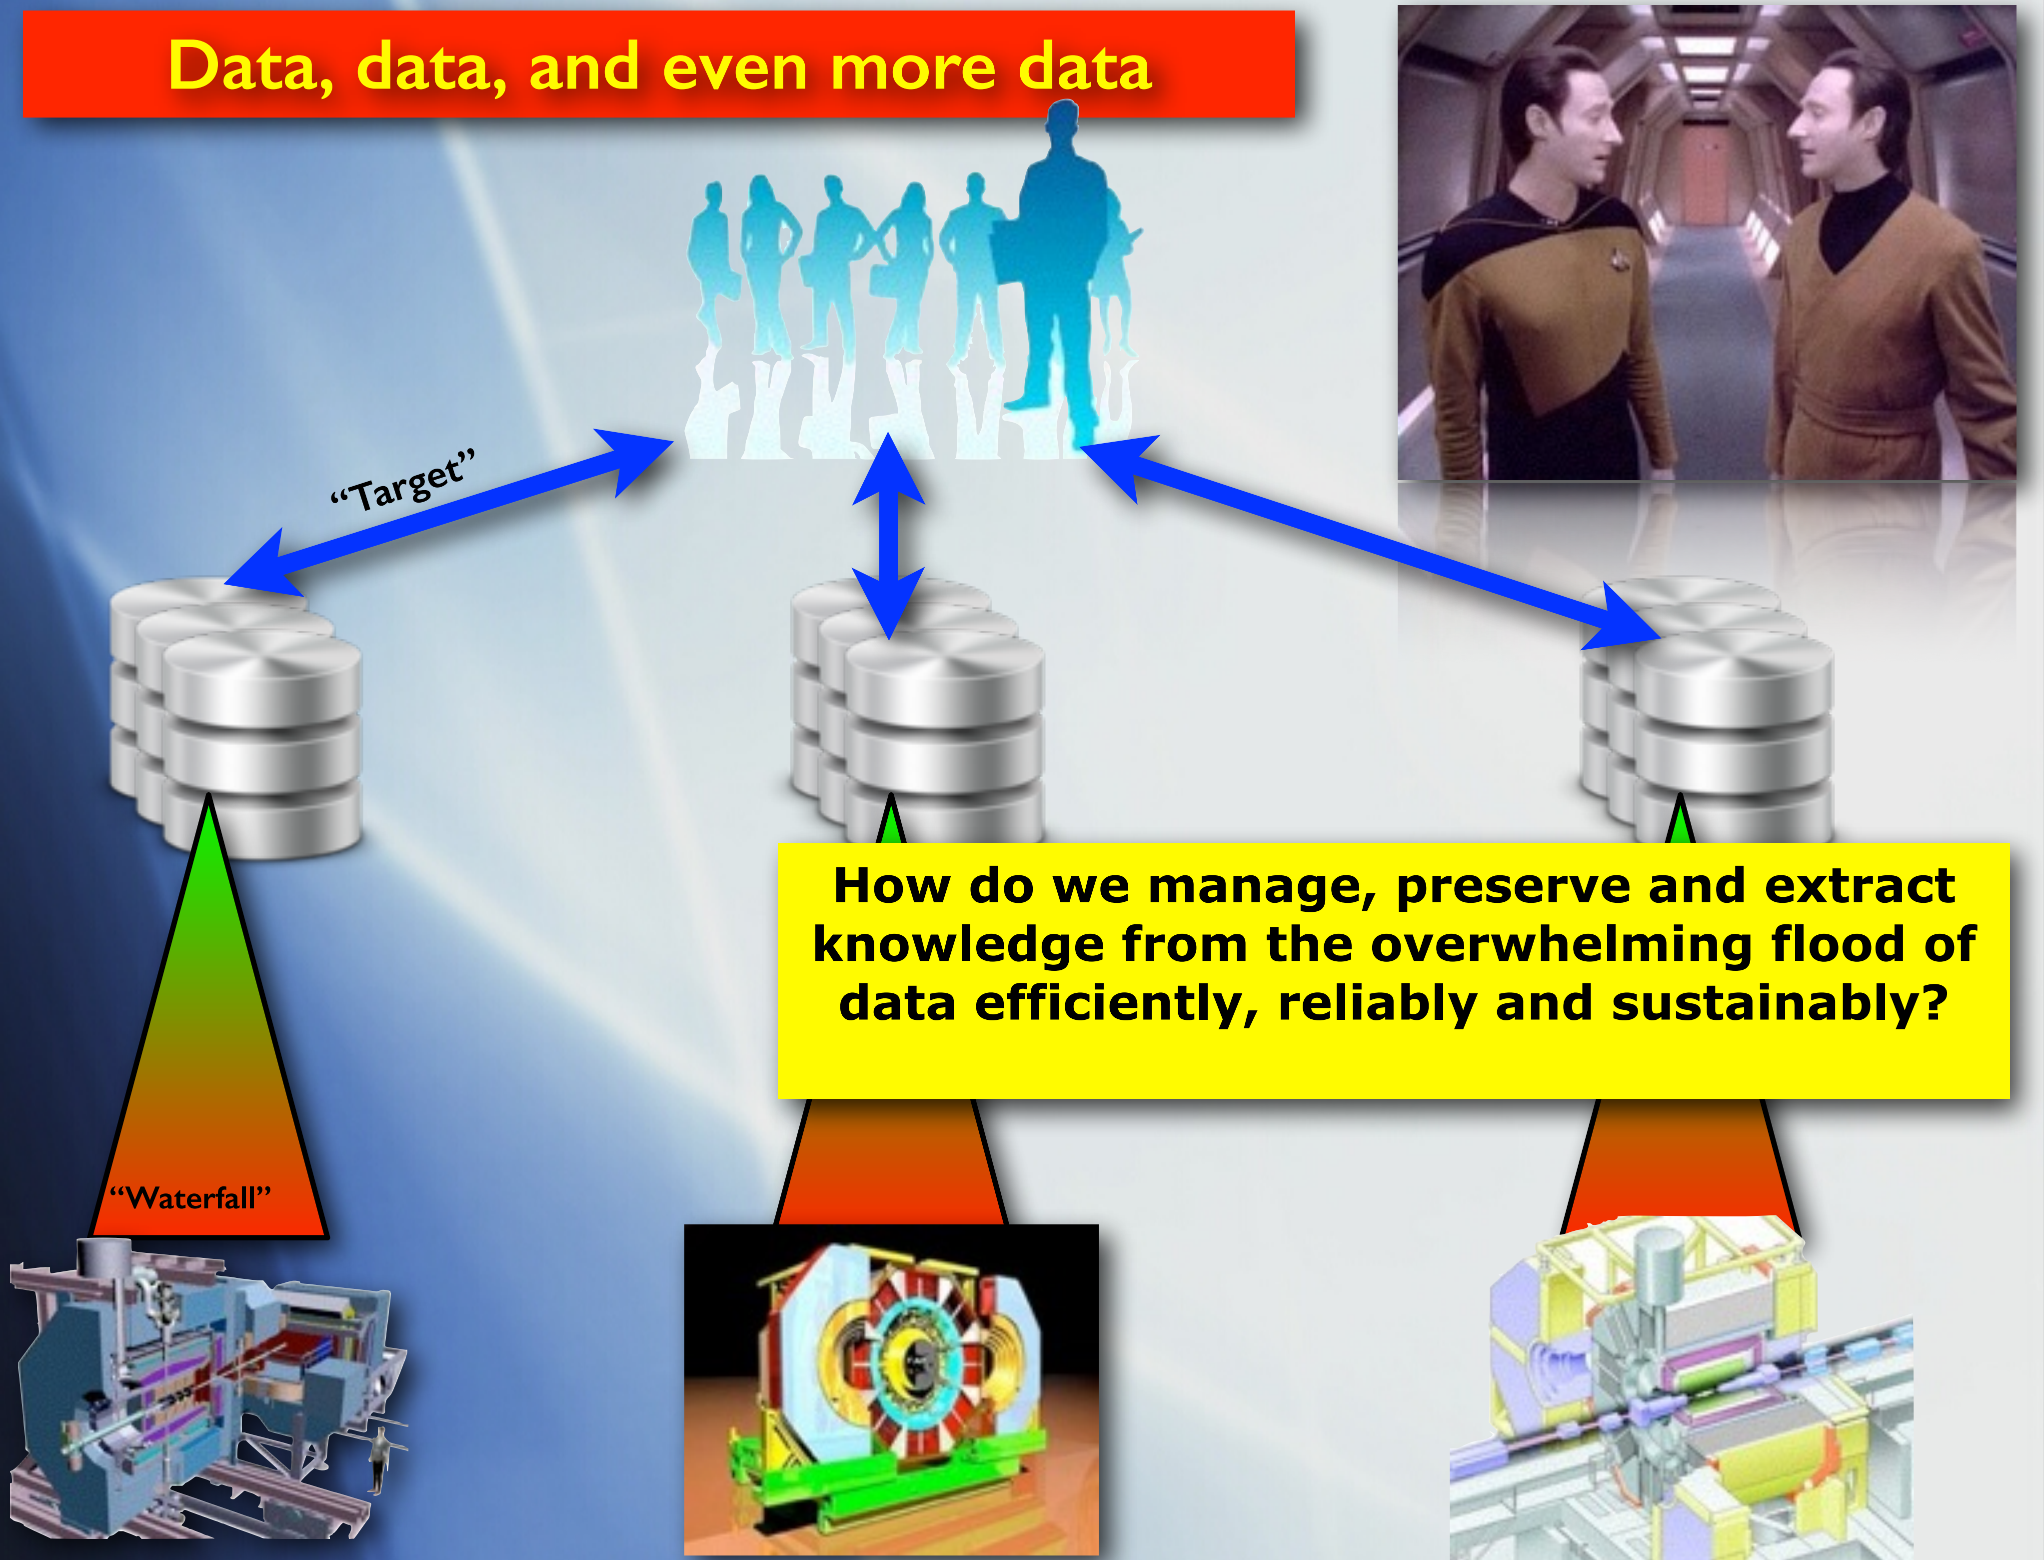

How do we manage, preserve and extract knowledge from the overwhelming flood of data efficiently, reliably and sustainably?

Data Processing, Publishing, and Mining

Edwin Valentijn et al.
(Astrowise, Target)

Data Processing, Publishing, and Mining

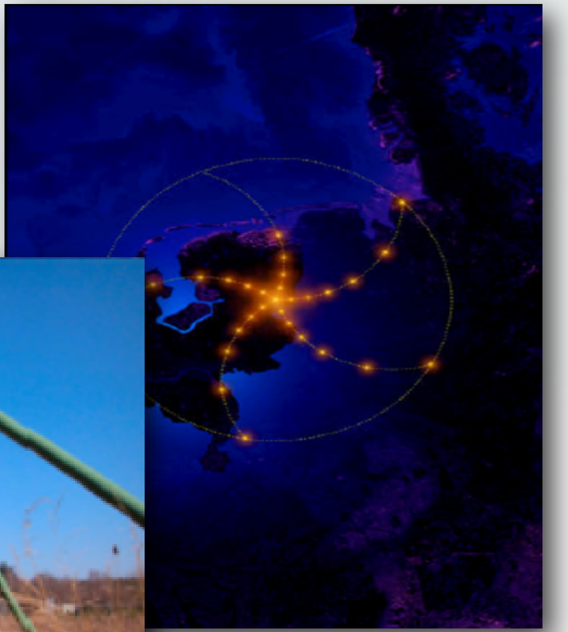

Astronomy

Edwin Valentijn et al.
(Astrowise, Target)

Data Processing, Publishing, and Mining

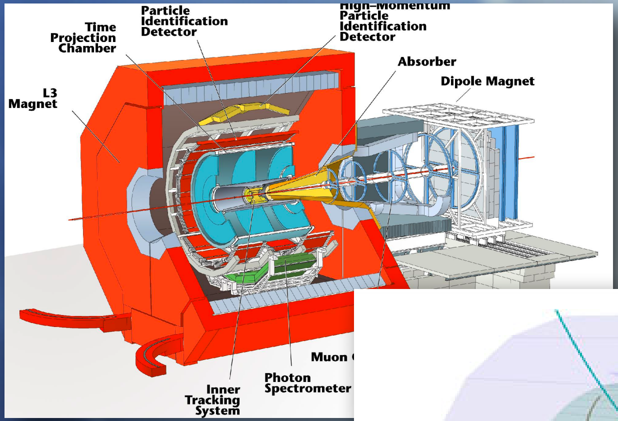

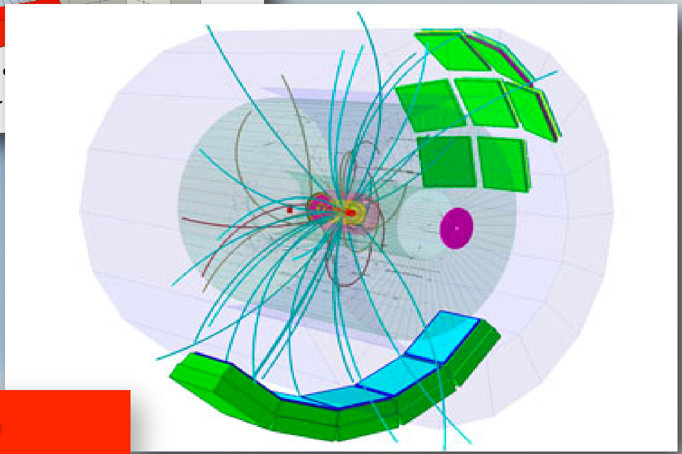

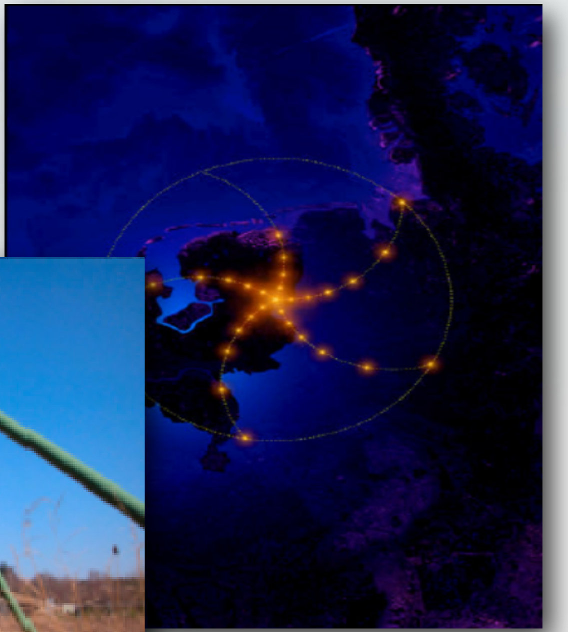

HEP (old style)

Astronomy

“Waterfall”

forward chaining
“tier” architecture
driven by raw data
process in pipeline
operators push data
results in releases
static archive
raw data is obsolete

Edwin Valentijn et al.
(Astrowise, Target)

Data Processing, Publishing, and Mining

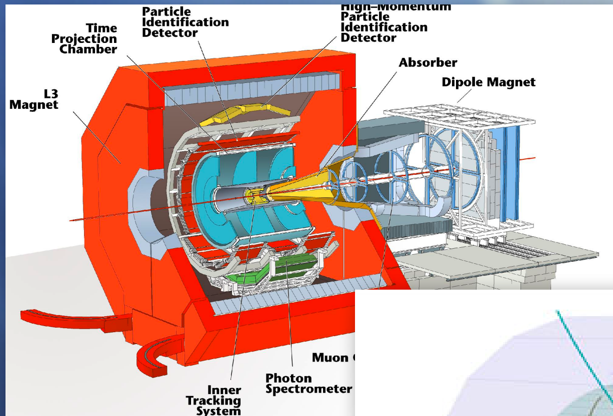

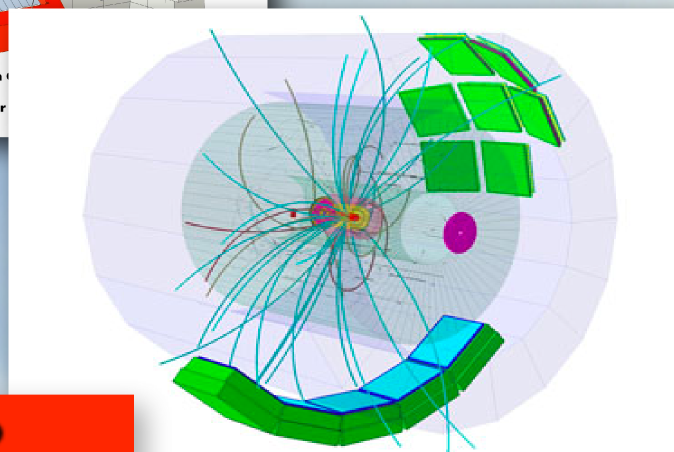

HEP (old style)

Astronomy

"Waterfall"	"Target"
forward chaining	backward chaining
"tier" architecture	"target" architecture
driven by raw data	driven by user query
process in pipeline	process on-the-fly
operators push data	users pull data
results in releases	information system
static archive	dynamic archive
raw data is obsolete	raw data is sacred

Edwin Valentijn et al.
(Astrowise, Target)

The 21st International Conference on Computing in High Energy and Nuclear Physics

CHEP2015

@Okinawa Institute of Science and Technology Graduate University (OIST), Okinawa, Japan

April 13-17, 2015

Keywords for abstract submission

APPLICATION AREAS

•Computing facilities and infrastructures •Computing model
•DAQ •Data analysis •Data preservation •Data stores
•Experiment frameworks for WAN distributed computing
•Middleware and services for production infrastructures
•Monitoring •Multi-discipline •Outreach •Reconstruction
•Simulation •Trigger

TECHNOLOGIES

•Cloud computing •Collaborative tools •Continuous integration systems
•Control systems •CPU architecture, GPU, FPGA •Data algorithms
•Data handling/access •Databases •Event processing frameworks
•High performance computing •Networking •Parallel programming
•Performance and validation tools •Software design
•Software development process •Storage systems •Virtualization

http://chep2015.kek.jp/

The 21st International Conference on Computing in High Energy and Nuclear Physics

CHEP2015

@Okinawa Institute of Science and Technology Graduate University (OIST), Okinawa, Japan

April 13-17, 2015

Keywords for abstract submission

APPLICATION AREAS

•Computing facilities and infrastructures •Computing model
•DAQ •Data analysis •Data preservation •Data stores
•Experiment frameworks for WAN distributed computing
•Middleware and services for production infrastructures
•Monitoring •Multi-discipline •Outreach •Reconstruction
•Simulation •Trigger

TECHNOLOGIES

•Cloud computing •Collaborative tools •Continuous integration systems
•Control systems •CPU architecture, GPU, FPGA •Data algorithms
•Data handling/access •Databases •Event processing frameworks
•High performance computing •Networking •Parallel programming
•Performance and validation tools •Software design
•Software development process •Storage systems •Virtualization

News

2014/10/9: Abstract submission deadline has been extended until October 25.

2014/8/19: 1st bulletin is released.

http://chep2015.kek.jp/

Computing **challenges** and **opportunities** for FAIR experiments

Computing challenges and opportunities for FAIR experiments

FAIR computing is all about the "Big Data" challenge

Computing challenges and opportunities for FAIR experiments

FAIR computing is all about the "Big Data" challenge

Innovative approaches are pursued in this era of many cores & data mining

Computing challenges and opportunities for FAIR experiments

FAIR computing is all about the "Big Data" challenge

Innovative approaches are pursued in this era of many cores & data mining

Collaborative efforts and usage of frameworks pay off royally

Computing **challenges** and **opportunities** for FAIR experiments

FAIR computing is all about the "Big Data" challenge

Innovative approaches are pursued in this era of many cores & data mining

Collaborative efforts and usage of frameworks pay off royally

The ultimate aim: place physics at the center and physicist in the driver seat