

Designing the Computing for the Future Experiments

Stefano Spataro



Tuesday, 15th October, 2013

A new experiment, many old questions

How to design a Reconstruction Framework?

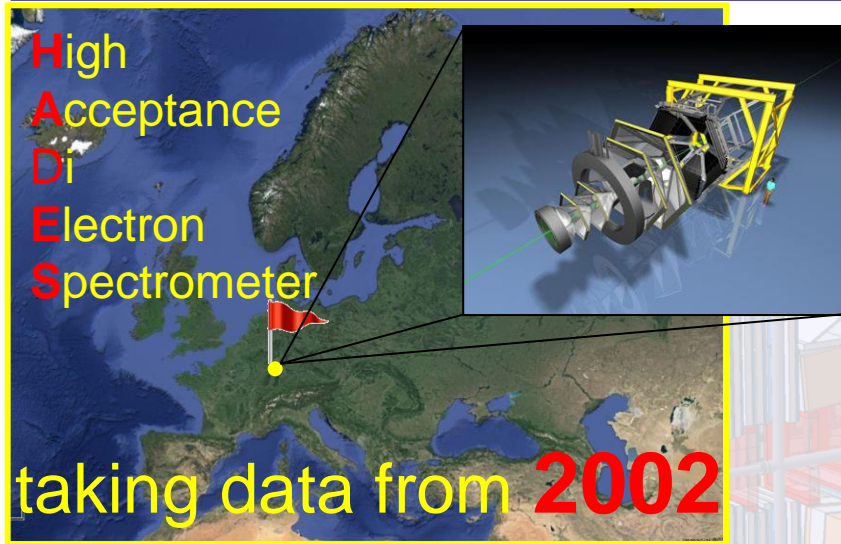
starting from scratch, “high efficiency”, limited manpower

How to design the best Data Acquisition?


highest data rate, smartest trigger selections, wider physics program

How to use the new computing technologies?

Parallel computing, FPGAs, GPUs, ARM?



Hades **sY**stem for **D**ata
Reduction and **A**nalysis



A good name to scare students...

HADES

HSpectrometer

HDetector

HDataSource: Data IO

HRuntimeDb: Parameter handling


HParlo: Parameter IO

HTaskSet

HReconstructor

- Written in C++ and OO (from scratch)
- Based on Root objects
- Geant3 (patched)
- Data input: hdl, geant root, reco root
- Parameters: ASCII, ROOT, Oracle

Hades sYstem for Data Reduction and Analysis



the beginner problems
(me as Diploma student @ 2000)

- Installation of external code
 - CERNLIB
 - GEANT3
 - ROOT
 - Oracle client...
- Installation of HYDRA
- Linux compatibility
- Gcc compatibility
- Different results in different systems


(now it is much better!)

Not easy for a newcomer!

The solution	→	work remotely at GSI (vi rules)
Production	→	GSI batch farm

People involved in analysis had to develop parts of reconstruction


Let's speak about future experiments



In the upcoming future (taking data from 2018) ...

$\bar{p}p$, $\bar{p}A$ collisions
1.5 \Rightarrow 15 GeV/c (\bar{p} momentum)

Anti-Proton
ANihilation at
DARMstadt



- Charmonium ($c\bar{c}$) spectroscopy
- Open charm spectroscopy
- Search for gluonic excitations
(hybrids - glueballs)
- Charmed hadrons in nuclei
- Drell-Yan
- Single and double Hypernuclei
- Parton Dist., EM Form Factor...

2017: Detector Assembly
2018: First Data taking

More than 500 physicists from more than 54 institutions in 17 countries

Few considerations about “software” and “manpower”

not easy to have new people working on software developements
for **not running** experiments

How a boss can “divert” potential developers...

- w/o data no publications, no conferences, no good CV
- In your PhD you should do physics and not programming
- The “detector” developments are more attractive, experimental tests...
- Programming is only for “experts”
- Do something else, wait till somebody will prepare the code for you
(analysis of some running experiments)

How a boss usually attacks the currently working developers...

- **Why the reconstruction is not ready yet???**
(my student has to finish his thesis !!!)

Software manpower is extremely limited, busy also with other activities,
but the things to be implemented are endless...

What do you need for a reconstruction software?

- ✓ Data objects format
- ✓ Geometry handling
- ✓ I/O Manager
- ✓ Database connection (which DB?)
- ✓ Simulation of physics processes (G3, G4, Fluka, ?)
- ✓ Event Display
- ✓ Advanced Analysis Tools

The basic features **do not depend** on detector specifics

Somebody else has already done the job

No need to reinvent the wheel!



The (Panda) Solution




Use a framework already used by other experiments

- Less software developments for your computing group
- More people using the same code → better debug
- Share of the same tools by larger community

Up to now PANDA has changed software framework two times

The third seems the “long-lasting” solution

└→ PandaRoot since 2006



Start testing
the VMC
concept for
CBM

Panda decided
to join->
FairRoot: same
Base package
for different
experiments

R3B joined

EIC (Electron
Ion Collider
BNL)
EICRoot

SOFIA
(Studies On
Fission with
Aladin)

2004

2006

2010

2011

2012

2013





First Release of
CbmRoot


MPD (NICA)
start also using
FairRoot

ASYEOS joined
(ASYEOSRoot)

GEM-TPC
separated
from PANDA
branch
(FOPIRoot)

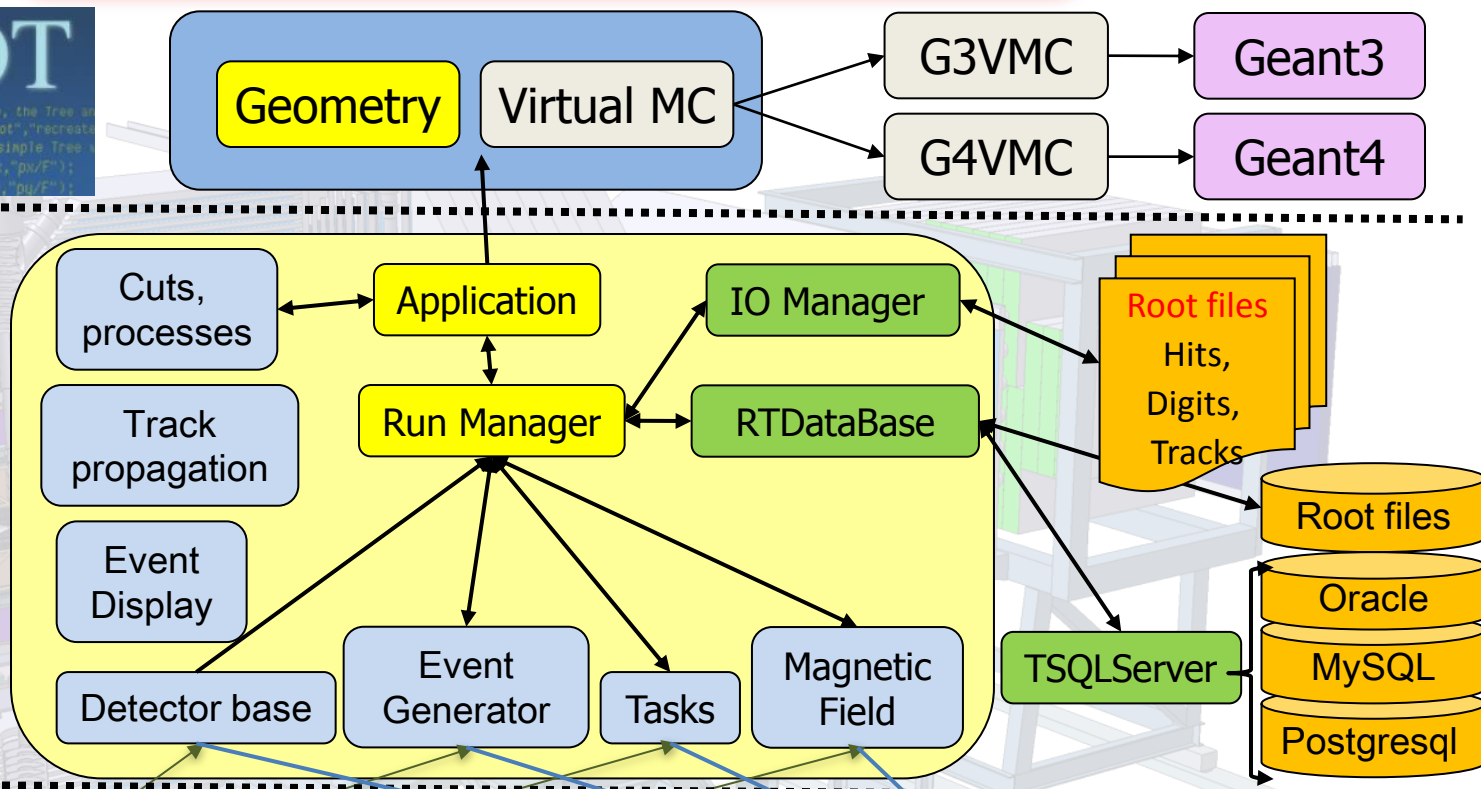
ENSAR-ROOT
Collection of
modules used by
structural nuclear
physics exp.



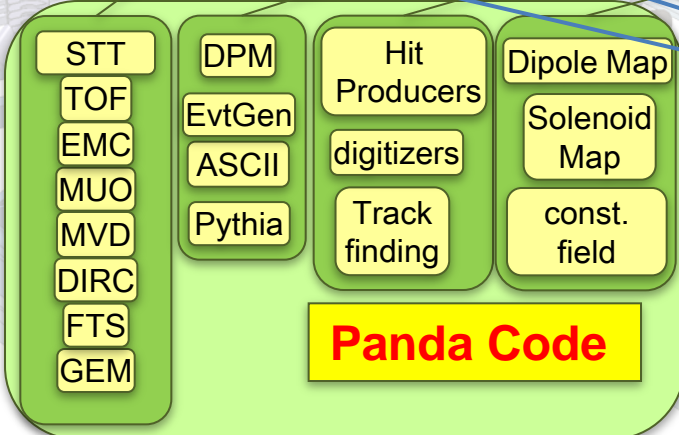


FairRoot

*M.Al-Turany,
D.Bertini,
F.Uhlig,
R.Karabowicz*



PandaRoot



CbmRoot
R3BRoot
MPDRoot (NICA)
ASYEOSRoot
EICRoot

- ✓ No executable

Root macros to define the experimental setup, the tasks for reco/analysis, the configuration

- ✓ No fixed simulation model

Different simulation models with the same user code (VMC)

- ✓ No fixed output structure

Dynamic event structure based on Root TFolder and TTree

For more information:

- ✓ **Extending the FairRoot framework to allow for simulation and reconstruction of free streaming data**

Mohammad AL-TURANY on 14 Oct 2013 from 16:10 to 16:30

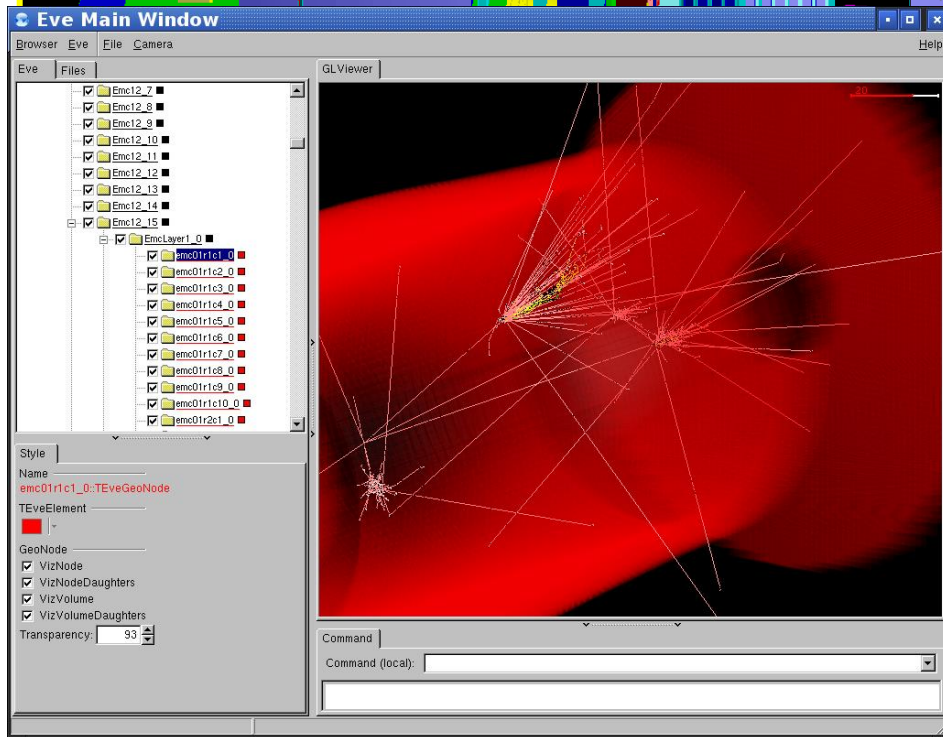
- ✓ **An Event Building scenario in the trigger-less PANDA experiment**

Radoslaw KARABOWICZ on 14 Oct 2013 from 16:45 to 17:05


ROOT Geometry and Event Display

Oh Yes! IT'S
FREE

TGeoManager



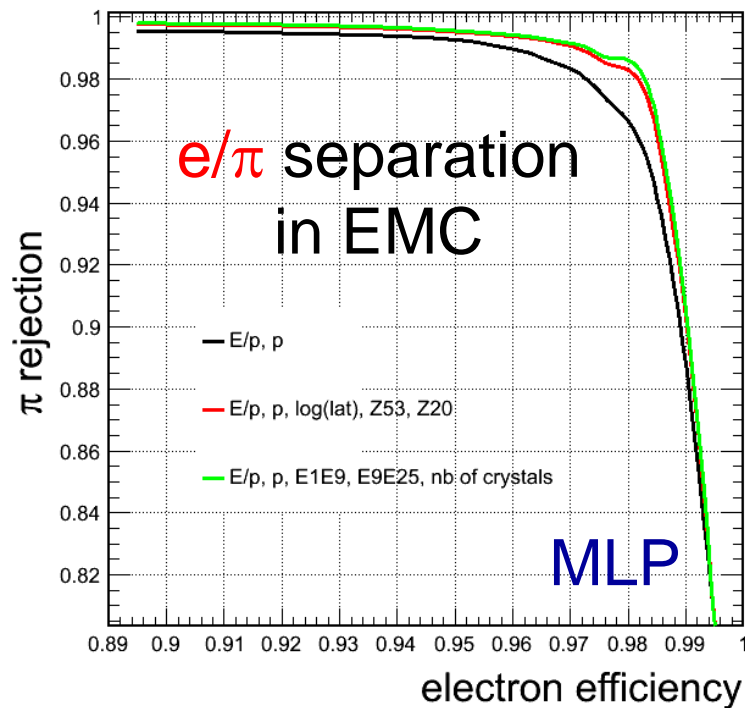
TEve



MultiVariate Particle Identification

Oh Yes! IT'S
FREE

e/π separation
in EMC




implementation of
ROOT **TMVA** methods

- ✓ K- Nearest Neighbors (KNN)
- ✓ Learning Vector Quantization (LVQ)
- ✓ Multi Layer Perceptron (MLP)
- ✓ ...

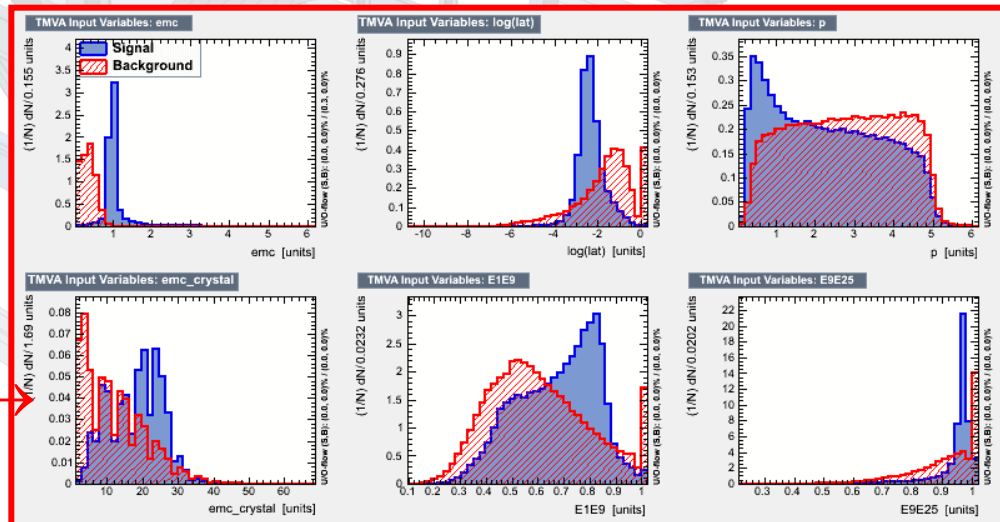
➤ **EMC** shower shape analysis

Correlation Matrix (signal)




Correlation


Variables




SPEED UP FACTOR DEPENDENCE ON NUMBER OF EVENTS



SPEED UP FACTOR DEPENDENCE ON NUMBER OF EVENTS



TIME SPEEDUP VS NUMBER OF WORKERS



a lot of work to modify the code to make it “Proof compatible”

PoD on external CPUs with SSH
(4CPUs)+(8CPUs)+(8CPUs)

Compiled and running on many Linux distributions and on MAC OS X

- The release was tested on:
 - MAC OS X 10.6.x with gcc 4.2.1 (64 bit)
 - MAC OS X 10.7.x with gcc (64 bit)
 - MAC OS X 10.8.x with llvm-gcc 4.2 and gfortran-4.7 from hpc.sourceforge.net (64 bit)
 - MAC OS X 10.8.x with clang 4.0 and gfortran-4.7 from hpc.sourceforge.net (64 bit)
 - Suse 12.1 with gcc 4.6.2 (32 and 64 bit)
 - Suse 12.2 with gcc 4.7.1 (32 and 64 bit)
 - Fedora 17 with gcc 4.7.0 (32 and 64 bit)
 - Fedora 17 with gcc 4.7.2 (64 bit)
 - Ubuntu 12.04 with gcc 4.6.3 (32 and 64 bit)
 - Ubuntu 12.10 with gcc 4.7.2 (64 bit)
 - Debian Squeeze with gcc 4.4.5 (64 bit)
 - Debian wheezy with gcc 4.7.2 (64 bit)

Everybody in his
desktop, laptop, local
farm can run the code
w/o problems (hopefully)

Using a set of self-configuring scripts (CMake)
and regular checks (DashBoard)

Central SVN repository

Dedicated test server



2. SVN triggers test server

3a. Update of local copy

1. Developer commit code



3b. Configure, build and test on local machine

4. Send results automatically to central web page

5. Dashboard prepare and display results



6. In case of problems Dashboard sends an E-mail to Developer and Administrator

7. Developer check results

Log in: All Dashboards Thursday, May 09 2013 10:34:08 CEST

PandaRoot

Dashboard Calendar Previous Current Next Project

Information about stability and validity of the newest PandaRoot subversion revision.


No file changed as of Wednesday, May 08 2013 - 00:00 CEST Show Filters Advanced View Auto-refresh Help

Nightly		Update		Configure		Build		Test		Build Time	
Site	Build Name	Files	Error	Warn	Error	Warn	Not Run	Fail	Pass		
pcha03.physik.uni-giessen.de	Ubuntu12.04-Linux-x86_64-gcc4.7-fairsoft_install	0	1	0	1	0	0	0	0		May 08, 2013 - 02:00 CEST
demac004	MacOSX-darwin-i386-gcc4.2.1-fairsoft_apr13	66	0	0	30	500	0	23	4		May 08, 2013 - 06:01 CEST
fb07-nucular.physik.uni-giessen.de	SLC_5.5-Linux-x86_64-gcc4.1.2-fairsoft_install_2012_08_07	66	0	0	30	500	0	23	4		May 08, 2013 - 02:01 CEST
ln0303.gsi.de	Squeeze-Sid64-Linux-x86_64-gcc4.4.5-fairsoft_apr13	0	0	0	11	3	0	23	4		May 08, 2013 - 16:17 CEST
ln0304.gsi.de	Debian_Squeeze64-Linux-x86_64-gcc4.4.5-fairsoft_fan12	66	0	0	0	500	0	5	22		May 08, 2013 - 09:02 CEST
lnraeos	Ubuntu12.04.2-Linux-x86_64-gcc4.6-fairsoft_apr13	50	0	0	0	500	0	4	23		May 08, 2013 - 02:01 CEST

Continuous		Update		Configure		Build		Test		Build Time	
Site	Build Name	Files	Error	Warn	Error	Warn	Not Run	Fail	Pass		
ln0103.gsi.de	Squeeze64-GNU-Linux-x86_64-gcc4.4.5-fairsoft_sep12	1	0	0	5	3					23 hours ago
ln0103.gsi.de	Squeeze64-GNU-Linux-x86_64-gcc4.4.5-fairsoft_sep12	2	0	0	5	3					May 08, 2013 - 17:57 CEST
ln0103.gsi.de	Squeeze64-GNU-Linux-x86_64-gcc4.4.5-fairsoft_sep12	1	0	0	0	0					May 08, 2013 - 17:56 CEST
ln0103.gsi.de	Squeeze64-GNU-Linux-x86_64-gcc4.4.5-fairsoft_sep12	1	0	0	0	0					May 08, 2013 - 17:07 CEST
ln0103.gsi.de	Squeeze64-GNU-Linux-x86_64-gcc4.4.5-fairsoft_sep12	3	0	0	0	0					May 08, 2013 - 17:05 CEST
ln0103.gsi.de	Squeeze64-GNU-Linux-x86_64-gcc4.4.5-fairsoft_sep12	1	0	0	0	0					May 08, 2013 - 17:03 CEST
ln0103.gsi.de	Squeeze64-GNU-Linux-x86_64-gcc4.4.5-fairsoft_sep12	3	0	0	0	0					May 08, 2013 - 16:52 CEST
ln0103.gsi.de	Squeeze64-GNU-Linux-x86_64-gcc4.4.5-fairsoft_sep12	61	0	0	0	2					May 08, 2013 - 14:20 CEST

Experimental		Update		Configure		Build		Test		Build Time	
Site	Build Name	Files	Error	Warn	Error	Warn	Not Run	Fail	Pass		
lnraeos	Ubuntu12.04.2-Linux-x86_64-gcc4.6-fairsoft_apr13	0	0	0	0	0	0	2	25		May 08, 2013 - 17:25 CEST
lnraeos	Ubuntu12.04.2-Linux-x86_64-gcc4.6-fairsoft_apr13	0	0	0	0	31	0	5	22		May 08, 2013 - 17:17 CEST
host075-pod027	MacOSX-darwin-i386_64-gcc4.2.1-fairsoft_apr13	4	0	0	0	360	0	10	17		May 08, 2013 - 17:09 CEST

Where to run my simulation?




GSI Batch Farm

Not very “flexible”



Physics Book
(older framework)



Several batch farms
(Bochum, GSI, Lyon, Orsay)

analyses
running locally

Distributed computing

In 2003 the future was GRID!

The usual computing problems

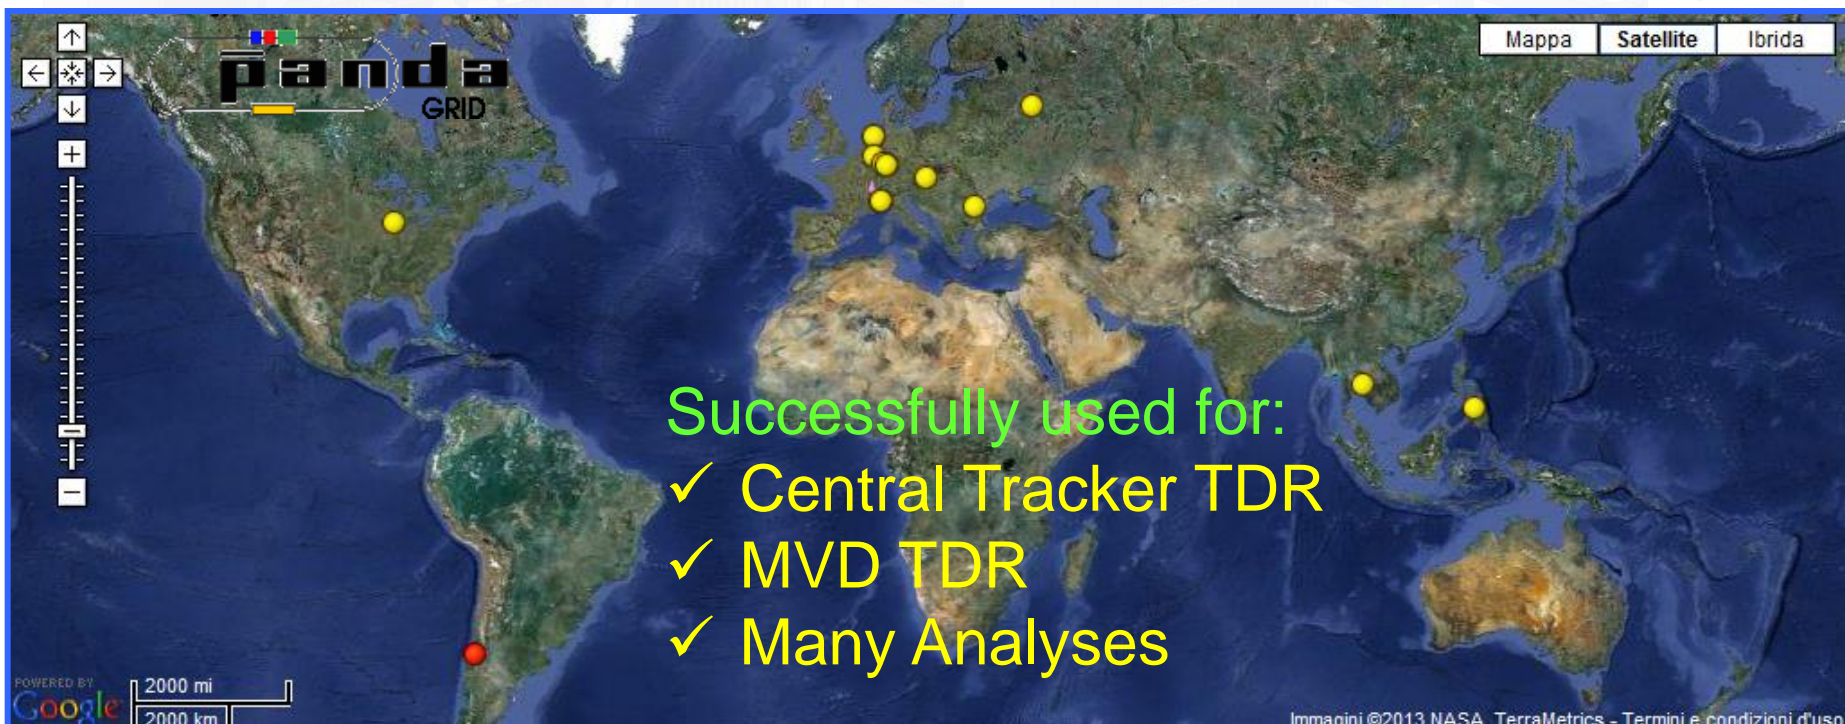
- Not a lot of manpower to develop middleware
- Too early to buy machines for the production
- Need to run simulation years before data taking

Why Alien?

- It can run on all platforms (source distribution)
- Several Panda institutions were hosting Alien sites


- ✓ “Reuse” of currently existing manpower
- ✓ Use of parts of already existing resources
- ✓ Strong collaboration with Alien developers

Beta-tester
(now 2.20)



What is now the future of our distributed computing?

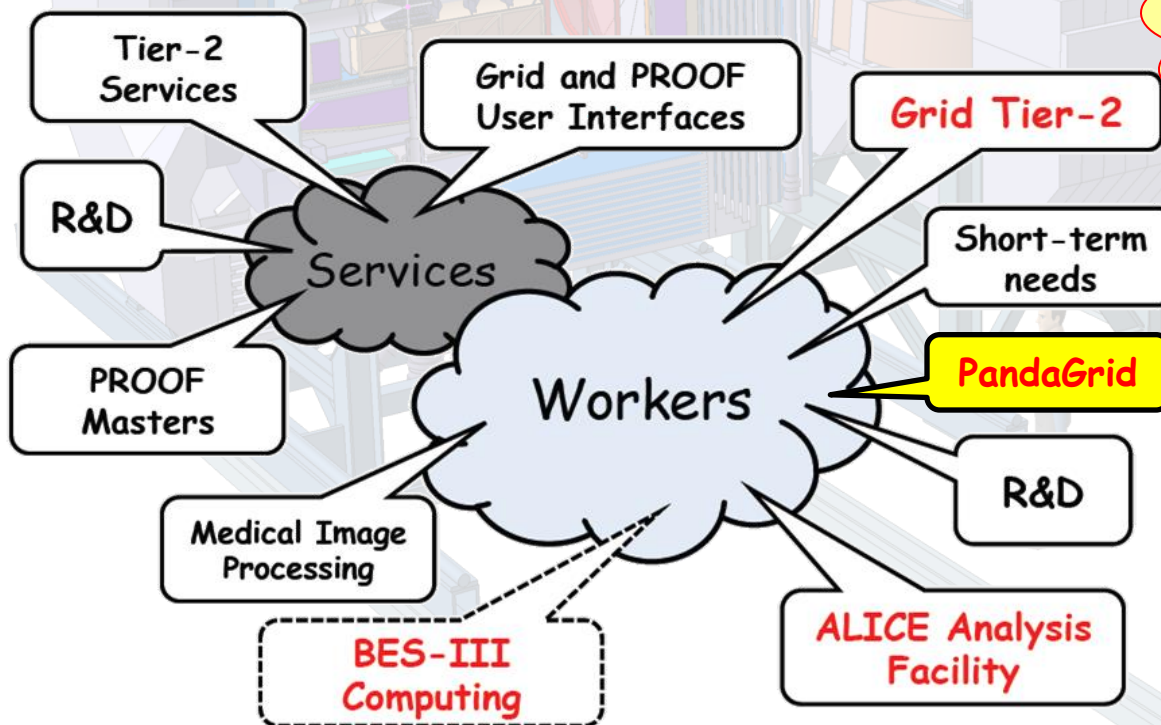
No further developments for Alien2



- Alien3?
- PanDa?
- Big batch farm?

Still some time before taking decisions for PandaGrid

Cloud computing
will help us



Torino Private Cloud
(S. Bagnasco et al.)

Distributed T0/T1 centre embedded in Grid/Cloud



APPA
CBM
LQCD
NuStar

in 2018

300000 cores + grid

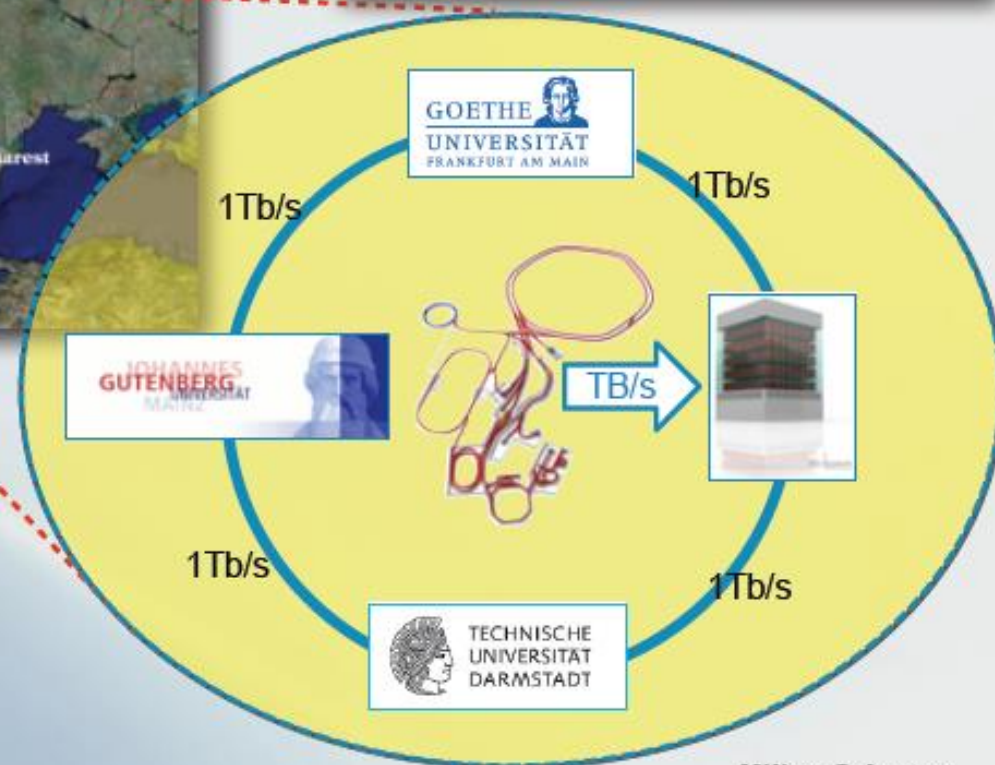
40 PB disk

40PB/y archive

Panda (66k cores, 12PB disks, 12PB/y tape)

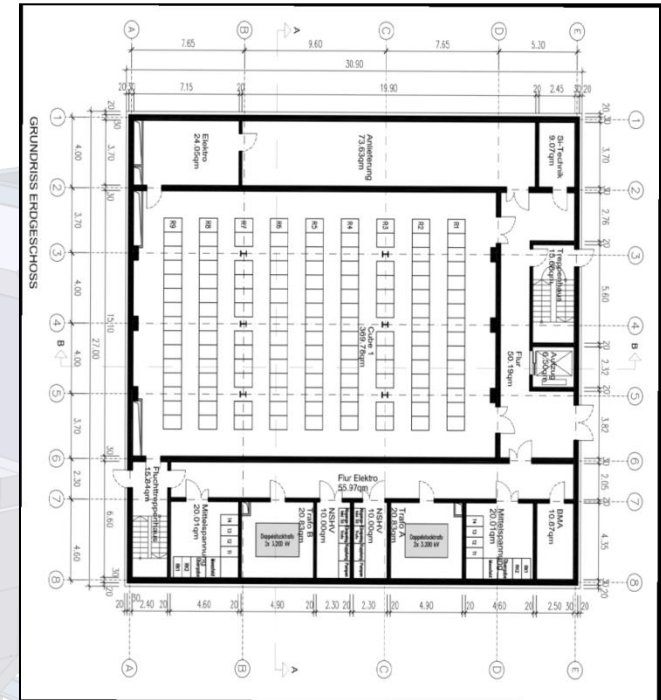
computing center for high-energy, hadron, nuclear, and atomic physics

serve >20 collaborations



Killian Schwarz

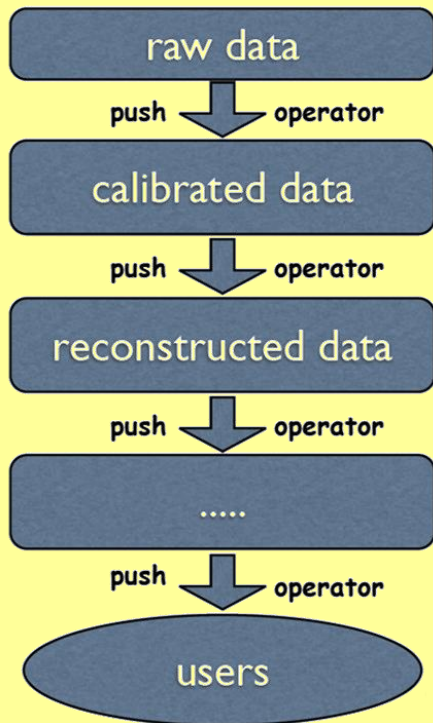
Reducing power consumption,
CO₂ emissions



Alternative ways for data processing and mining?

not only HEP experiments deal
with large amount of data

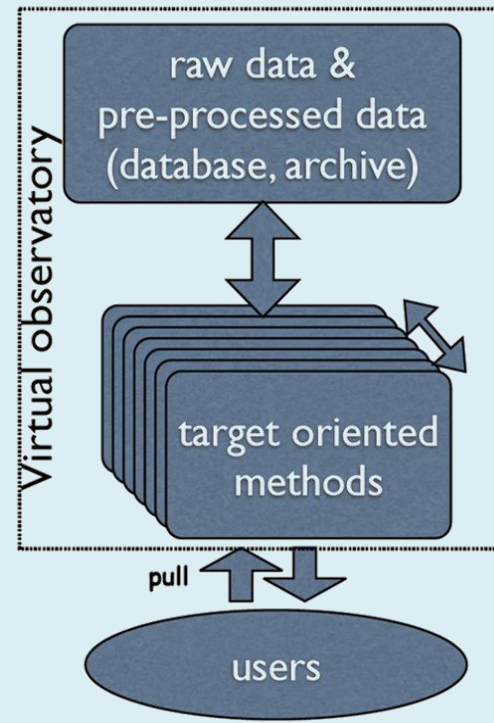
The *Waterfall* model



Forward chaining
“Tier” architecture
Driven by raw data
Process in pipeline
Operators push data
Results in release
Static archive
Raw data is obsolete

Standard in HEP

The *Target* model



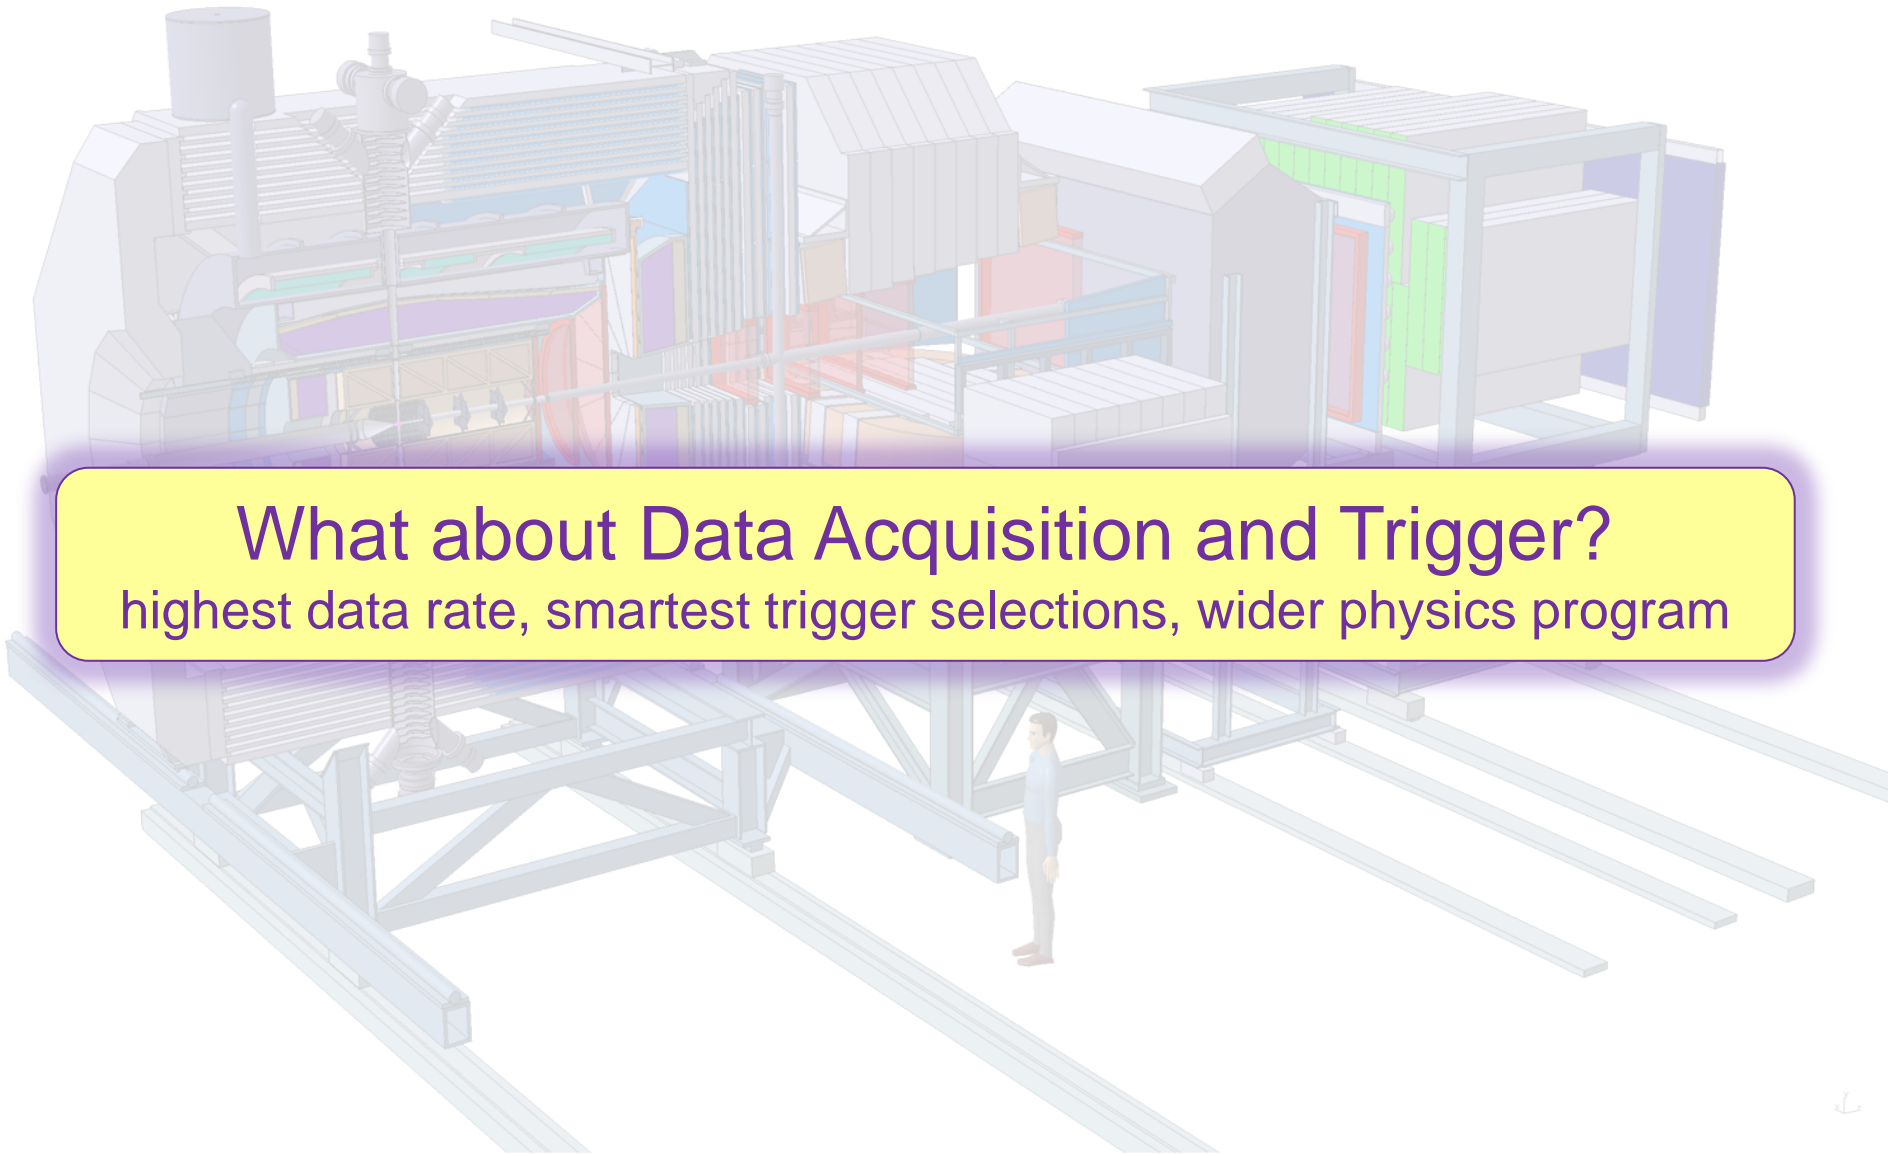
Backward chaining
“Target” architecture
Driven by user query
Process on-the-fly
Users pull data
Information system
Dynamic archive
Raw data is sacred

Used in Astronomy

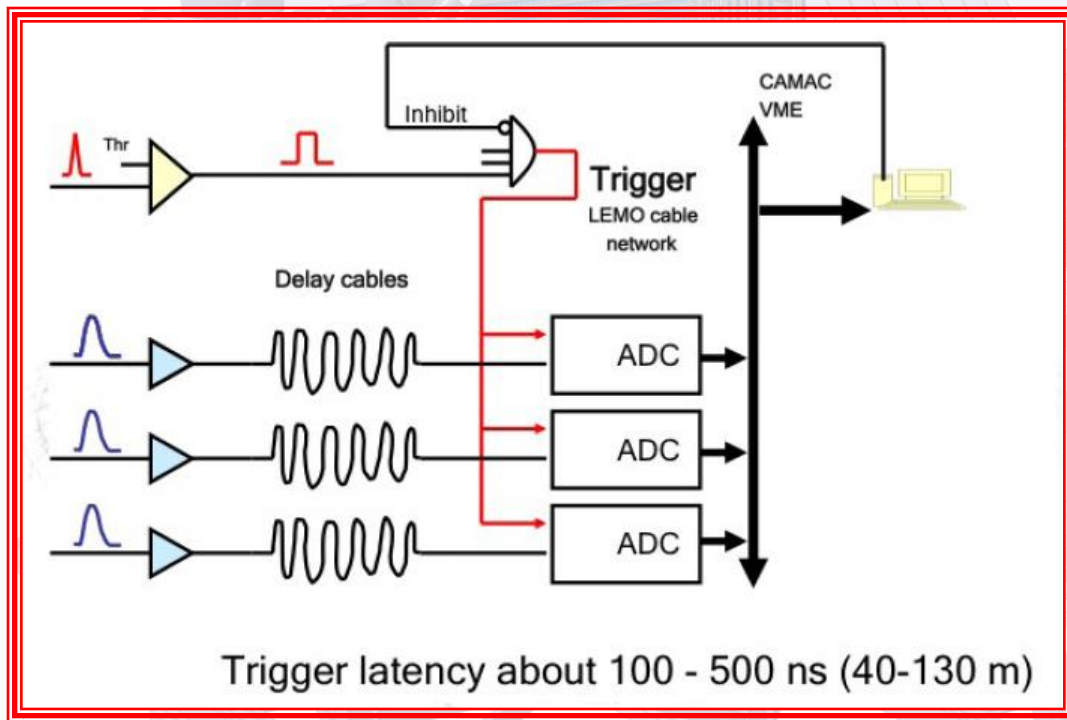
Could it work also for HEP?

What about Data Acquisition and Trigger?

highest data rate, smartest trigger selections, wider physics program



Traditional Data Acquisition



increasing the interaction rate






the number of pile-ups increases



- A first fast selection creates the event
- High level algorithms decide which event to store

- Interaction rates of 20 MHz (50 MHz peak)
- Event size of ~15 kB
- Data rates after front end preprocessing: 80GB/s - 300GB/s





Events/Data acquired by DAQ (temporarily buffered)

*Software Trigger
Algorithms*

„Trickle“ of events
stored on disc

- Required reduction factor: $\sim 1/1000$ (all triggers in total)
- A lot of physics channel triggers \rightarrow even higher reduction factor required

Selection criteria used in the Physics Report Analyses

Channel	TRK	NEUT	Excl.	mult	PID	p	E	ang.	inv M	dist cut	veto	4C	Vtx C	Mass C	Sig Eff[%]
J/psi pi+ pi-	4	0	x		e, pi				x				J/psi pi+ pi-		30
J/psi pi0 pi0	2	4	x		e		g		x		J/psi eta pi0	x	J/psi		17
chi_c1,2 gam	2	2	x		e		g		x			x	J/psi		30
J/psi gam	2	1	x		e				x			x	J/psi		40
J/psi eta	2	2	x		e				x			x	J/psi		40
h_c -> 3gam	0	3	x	3n			g	h_c	x			x			8
h_c -> 2phi gam	4	1	x		K		g		x		pi0	x			8
D+ D-	6	0	x		?	D			x	z(D)		x	D+-		8
D*+ D*-	6	0	x		?	D*			x	z(D*)		x	D0	D0	14
eta_c1 eta	2	7	x		e				x			x		chi, pi0, eta	7
eta_c1 eta	4	8	x		K, pi				x		>1 comb/ev	x	K pi	D0, D0*, eta, pi0	5
J/psi omega	4	2	x		e, pi				x			x	J/psi pi+pi-	J/psi, pi0	15
f2(2230) -> 2phi	4	0	x		K				x			x	phi		20
Ds Ds(2317)	3	0			K, pi			K	x				Ds, phi		20
Xi- Xi+ pi0	6	2	x		p, pi		g		x	d(IP-Xi)	>1 comb/ev	x	Lam, Xi+-	Xi Xi pi0	16
Lam Lam	6	0	x		p, pi				x	d(IP-Xi)			Lam		11 ... 23
Xi- Xi+	6	0	x		p, pi				x				Lam, Lam pi		19

no
multiplicity cuts

PID!

few kinematic cuts
(except mass)



a lot of fitting!

Physics Book criteria:


- J/psi (\rightarrow base for many charmonia)
 - Invariant Mass: **Tracking/Momentum**
 - Electron ID: **Tracking**, cluster energy, track/cluster match
 - Muon ID: **Tracking**, Muon detector information
 - Vertex: **Tracking**
- D/Ds Mesons
 - Pi0s: **EMC** clusters
 - Inv. Mass: **Tracking**
 - Kaon, Pion ID: **dE/dx**, **DIRC** info (w/ track match), ToF (track match)
 - Vertex: **Tracking**
- Baryons
 - Inv. Mass: **Tracking**
 - proton, pion ID: **DIRC** info (w/ track match)
 - Vertex: **Tracking**
- Full events: **4C** fitting

**Tracking & momentum
 \rightarrow key information**

Data Acquisition




Signals



Buffers



Level 1



Feature Extraction

Event Selection


continuous sampling DAQ
flash ADC



each signal get a time stamp (155.52 MHz)
high quality clock distributed

self-triggered
detector front-end

front-end feature extraction
(signal amplitude, shape, ...)

slow transition from **event based** to **time ordered** simulation



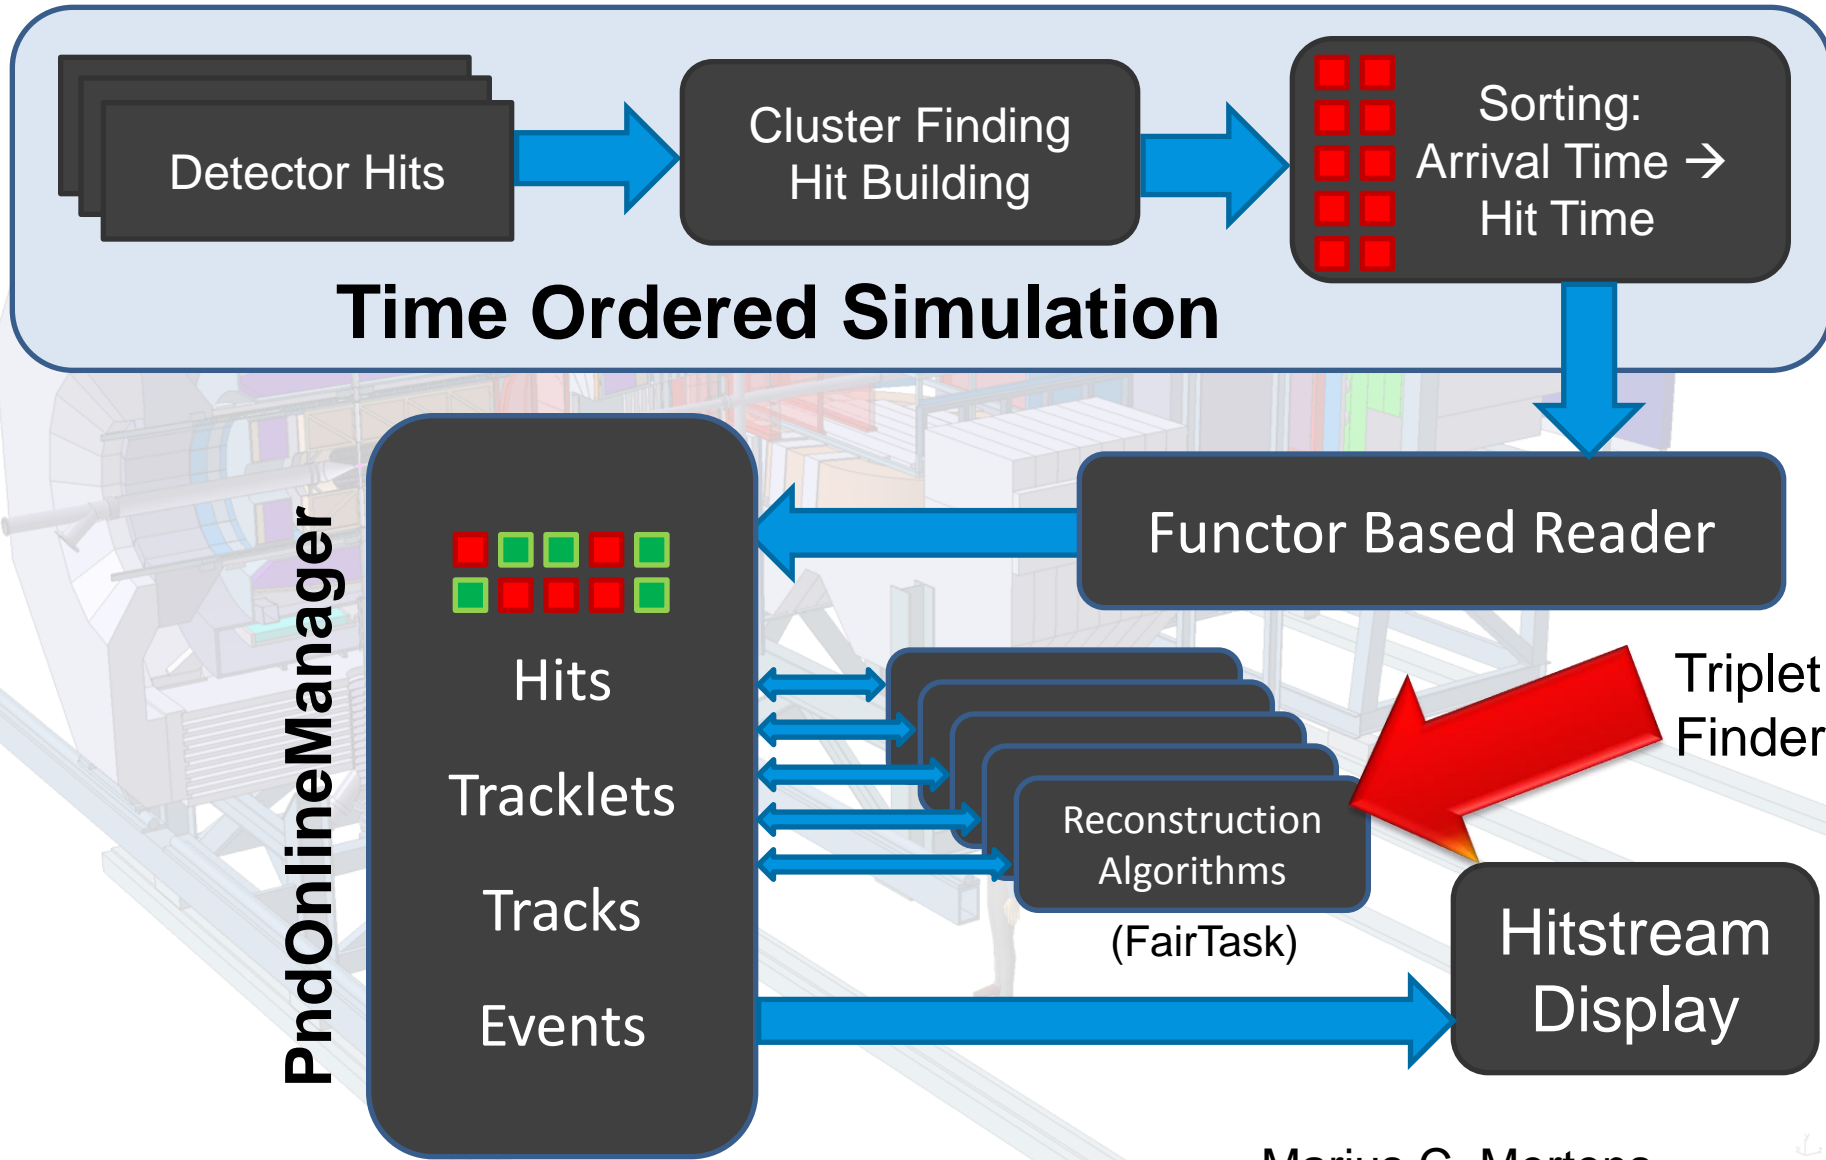


- Randomized Digi Data
- Sorting Digi Data using **Time Stamps + Drift Time, ToT...**
- Event Building – t_0 determination




✓ **An Event Building scenario in the trigger-less PANDA experiment**

Radoslaw KARABOWICZ on 14 Oct 2013 from 16:45 to 17:05



15 GeV/c DPM, 50 ns mean time

XY-View



0 ns

Dual Parton Model (DPM):
Standard $\bar{p}p$ background generator

Black circles: Early isochrone
Blue circles: Early skewed isochrone
Green circles: Close isochrone
Red circles: Late isochrone
Black dots: MVD hits
Green dots: MVD hits $r/z > 0.3$
Black+Red dots: Triplets/Skewlets
Yellow tracks: Vetoed
Blue tracks: Accepted

How to deal with continuous data stream?

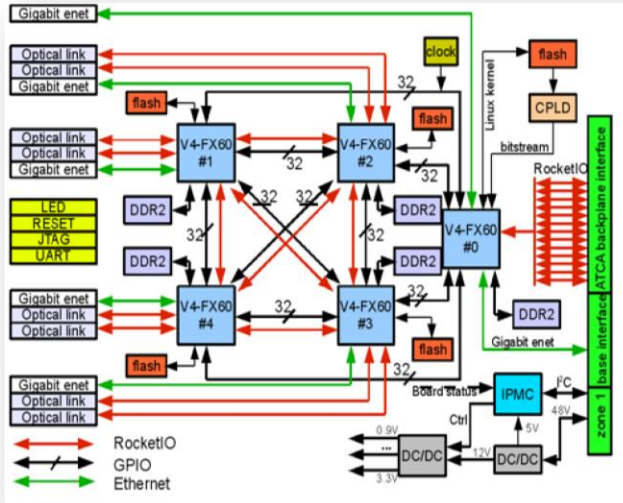
No possibility to pre-filter events (lvl1)

algorithms run continuously


- More power to online computing
- The (almost) whole offline reconstruction should run also online
- Algorithms as fast as possible
- (of course) **Concurrency** is the key!

need for more intelligence in online computing







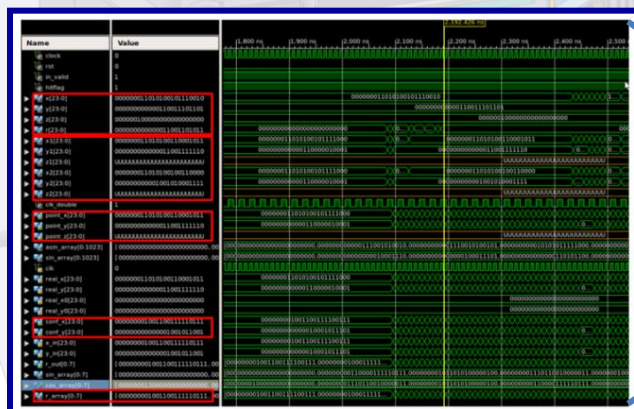
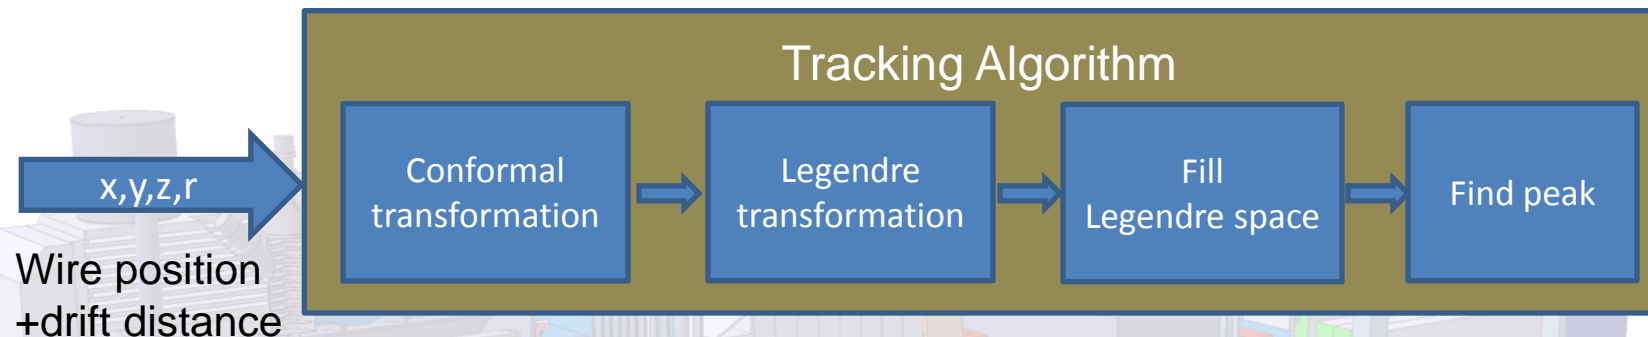
5x Virtex-4 FX60-10/-11 FPGA
 13x 2/3.125Gbps to backplane
 for interconnection
 5x Gigabit Ethernet
 8x 2/6.25Gbps Optical Links for
 data input
 2 GB 400MHz DDR2 SDRAM
 Real time Linux/vxworks




An **universal high performance** platform prepared for multiple applications .

ATCA standard (Full Mesh topology in backplane) and FPGA-based


the actual version for Panda online computing
 what is intended to be used later in Panda still to be decided



And tests with PC as data source and receiver



- GPU threads are extremely **lightweight**
- CPUs can execute **1-2 threads** per core, while GPUs can maintain up to **1024 threads** per multiprocessor (8-core)
- CPUs use **SIMD** (single instruction is performed over multiple data) vector units, and GPUs use **SIMT** (single instruction, multiple threads) for scalar thread processing. SIMT does not require developers to convert data to vectors and allows arbitrary branching in threads.



In the past...

✓ GPUs for event reconstruction


CHEP 2010 – Mohammad Al-Turany

✓ Track Finding in a High-Rate Time Projection Chamber Using GPUs

CHEP 2010 – Felix Böhmer

✓ Track finding and fitting on GPUs, first steps toward a software trigger

CHEP 2012 – Mohammad Al-Turany



Currently

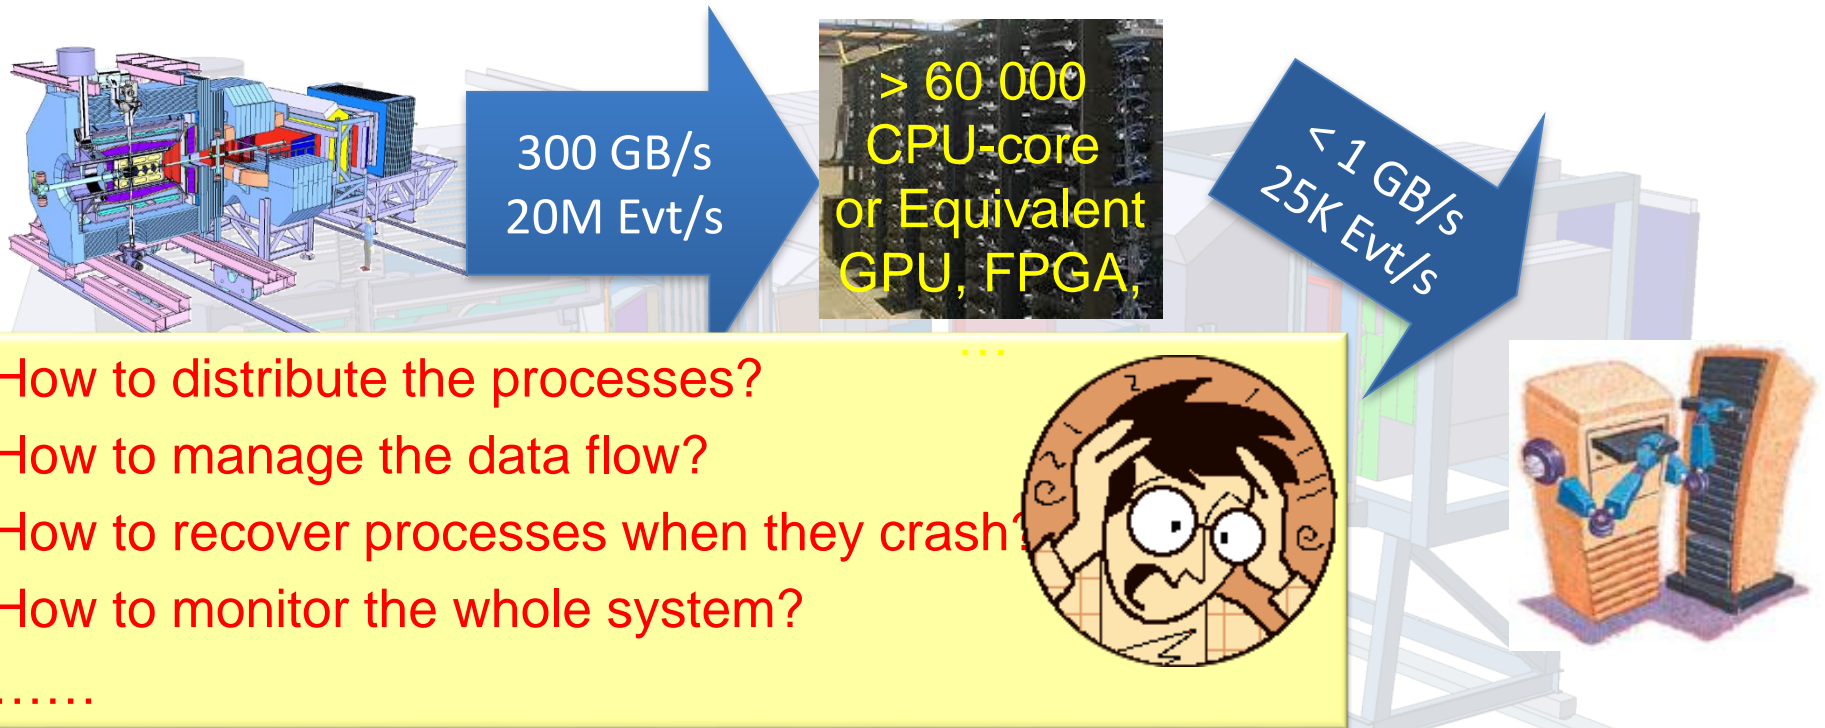
✓ Possibility to run Cuda directly from PandaRoot (FairCUDA)

✓ Direct collaboration with NVIDIA

✓ Algorithm Developments of:

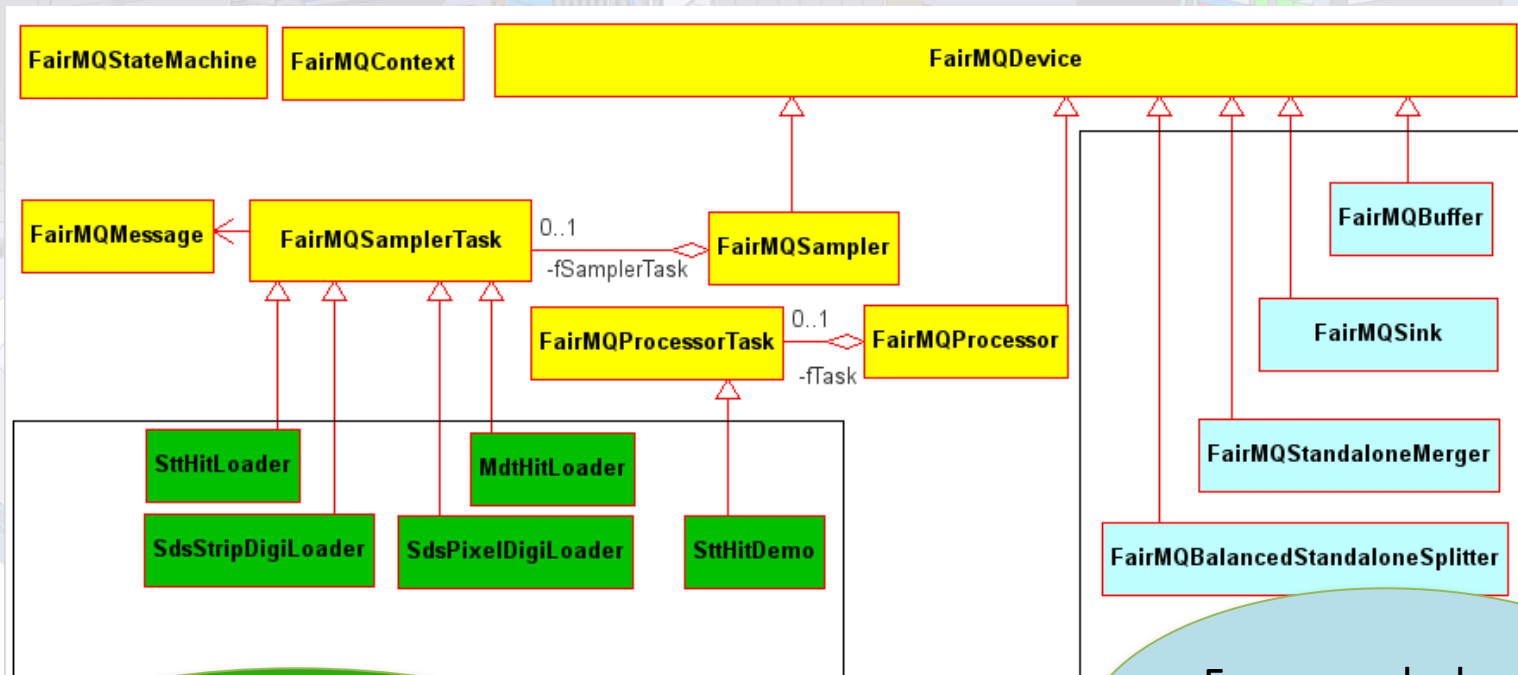
- GPU Hough Transform Tracker
- GPU Riemann Track Finder
- GPU Triplet Finder

The Online Reconstruction and analysis



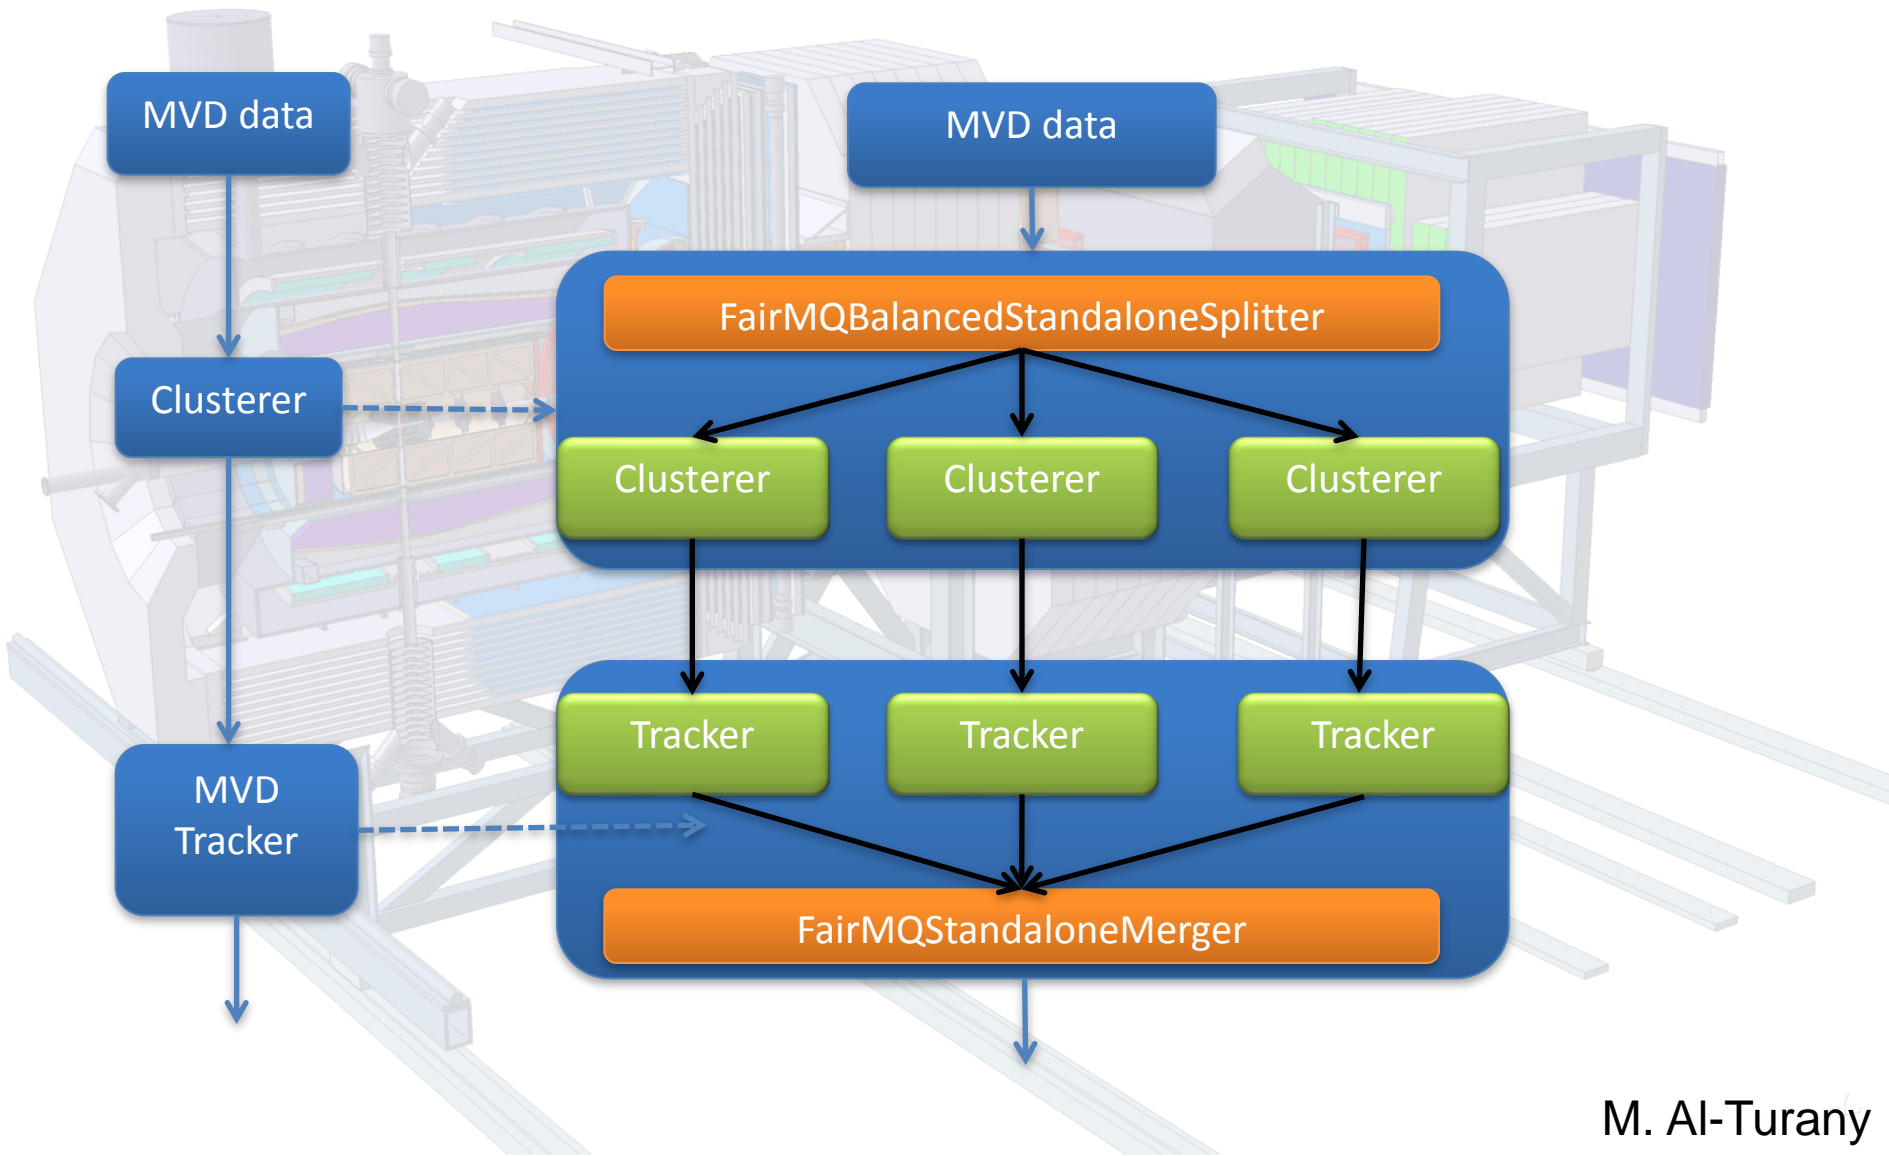
- Highly flexible: different data paths should be modeled.
- Adaptive: Sub-systems are continuously under development and improvement
- Should work for **sim** and **real** data: developing and debugging the algorithms
- It should support all possible hardware (**CPU**, **GPU**, **FPGA**, **ARR?**)
- It has to **scale** to any size! With minimum or ideally no effort.

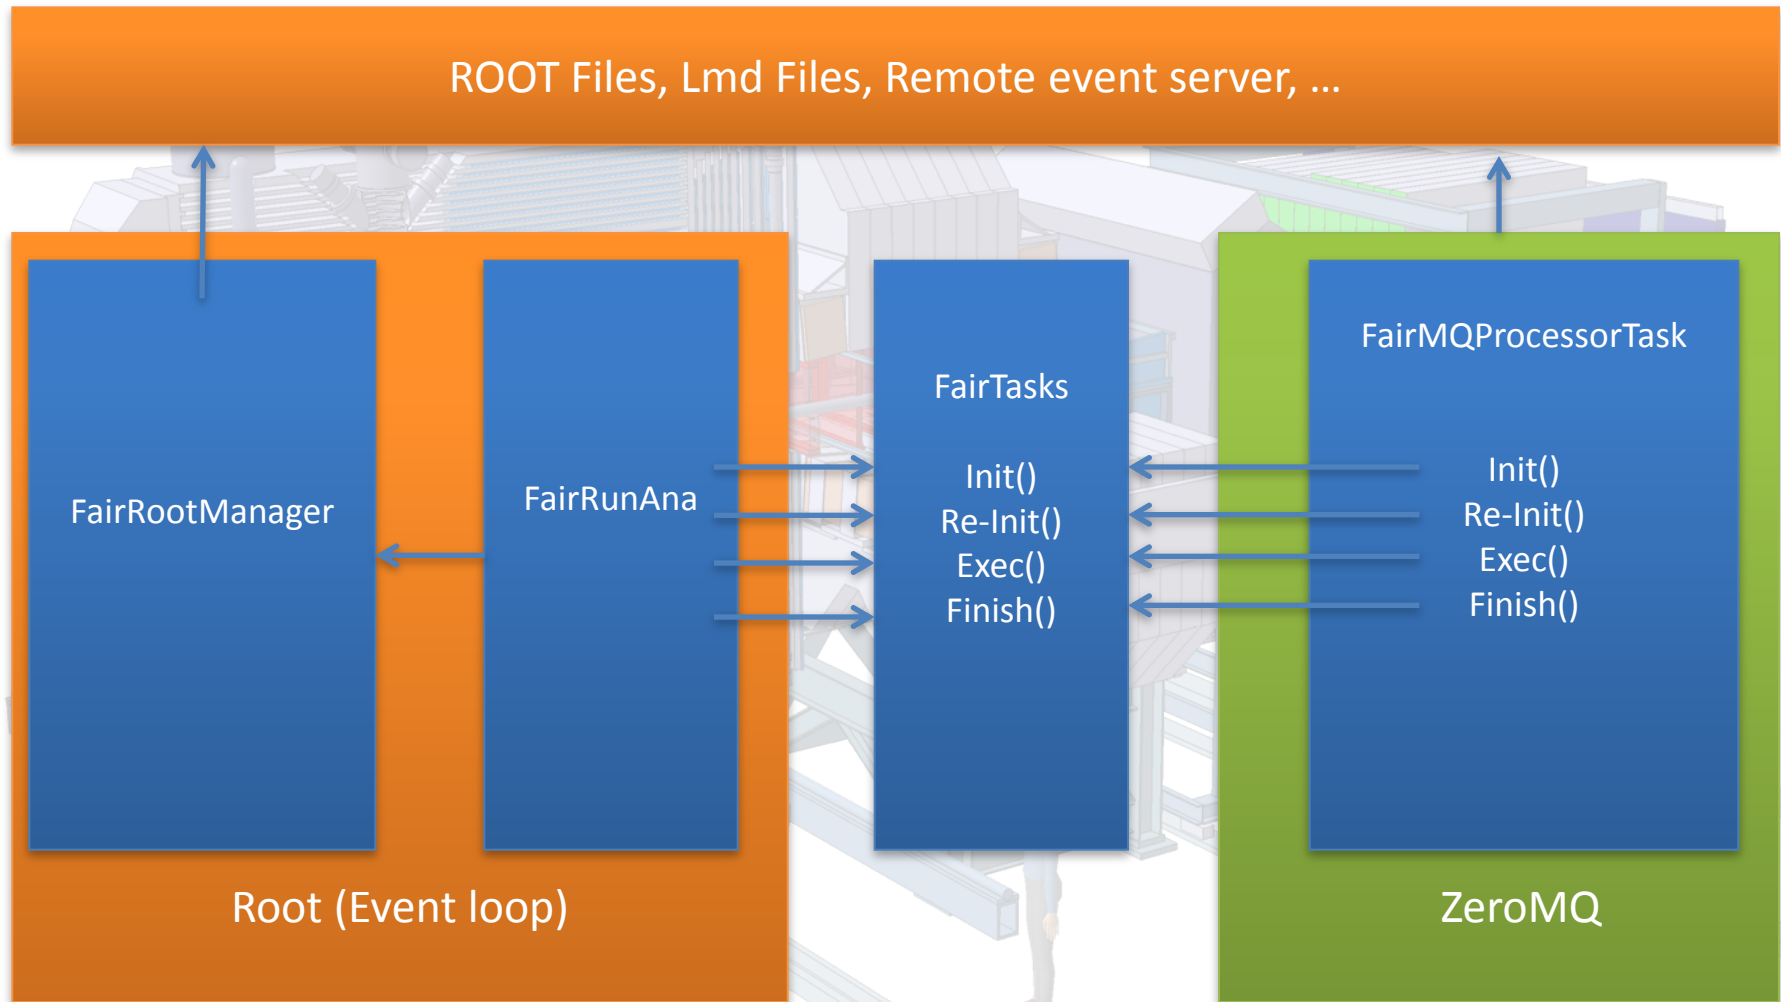
- A socket library that acts as a concurrency framework
- Messaging library, allowing to design a complex communication system without much effort



Experiment/detector
specific code

Framework classes
that can be used
directly





✓ **Extending the FairRoot framework to allow for simulation and reconstruction of free streaming data**

Mohammad AL-TURANY on 14 Oct 2013 from 16:10 to 16:30

- ✓ Panda benefits from the LHC experiences and from the new IT technologies
- ✓ Taking data from 2018, still some time to take final decisions
- ✓ The trigger-less data acquisition is the real challenge

Reconstruction

- PandaRoot is our framework for simulation, reconstruction and analysis
- Dynamic data structure, macro driven, supported on many OS
- Advantages from a large developer community and from 3rd part packages
- Time based simulation under realization (new concept!)
- High importance of Online algorithms

Computing Model

- The MONARC model good starting point but updated by new technologies
- Grid, Cloud, Proof, computing on FPGAs and on GPUs...
- Multi-core CPUs and many-core GPUs → importance of scalable software
- More democratic and flexible models

- ✓ With LHC upgrade higher data rates and more need of software parallelization
- ✓ Many points in common with LHC experiments, mutual benefits?

The 9 kinds of physics seminar

