Mitglied der Helmholtz-Gemeinschaft

PANDA Computing Numbers

for FAIR Computing CDR

29. Mai 2018 | Tobias Stockmanns PANDA FEE/DAQ Workshop

Timeline for FAIR Computing CDR

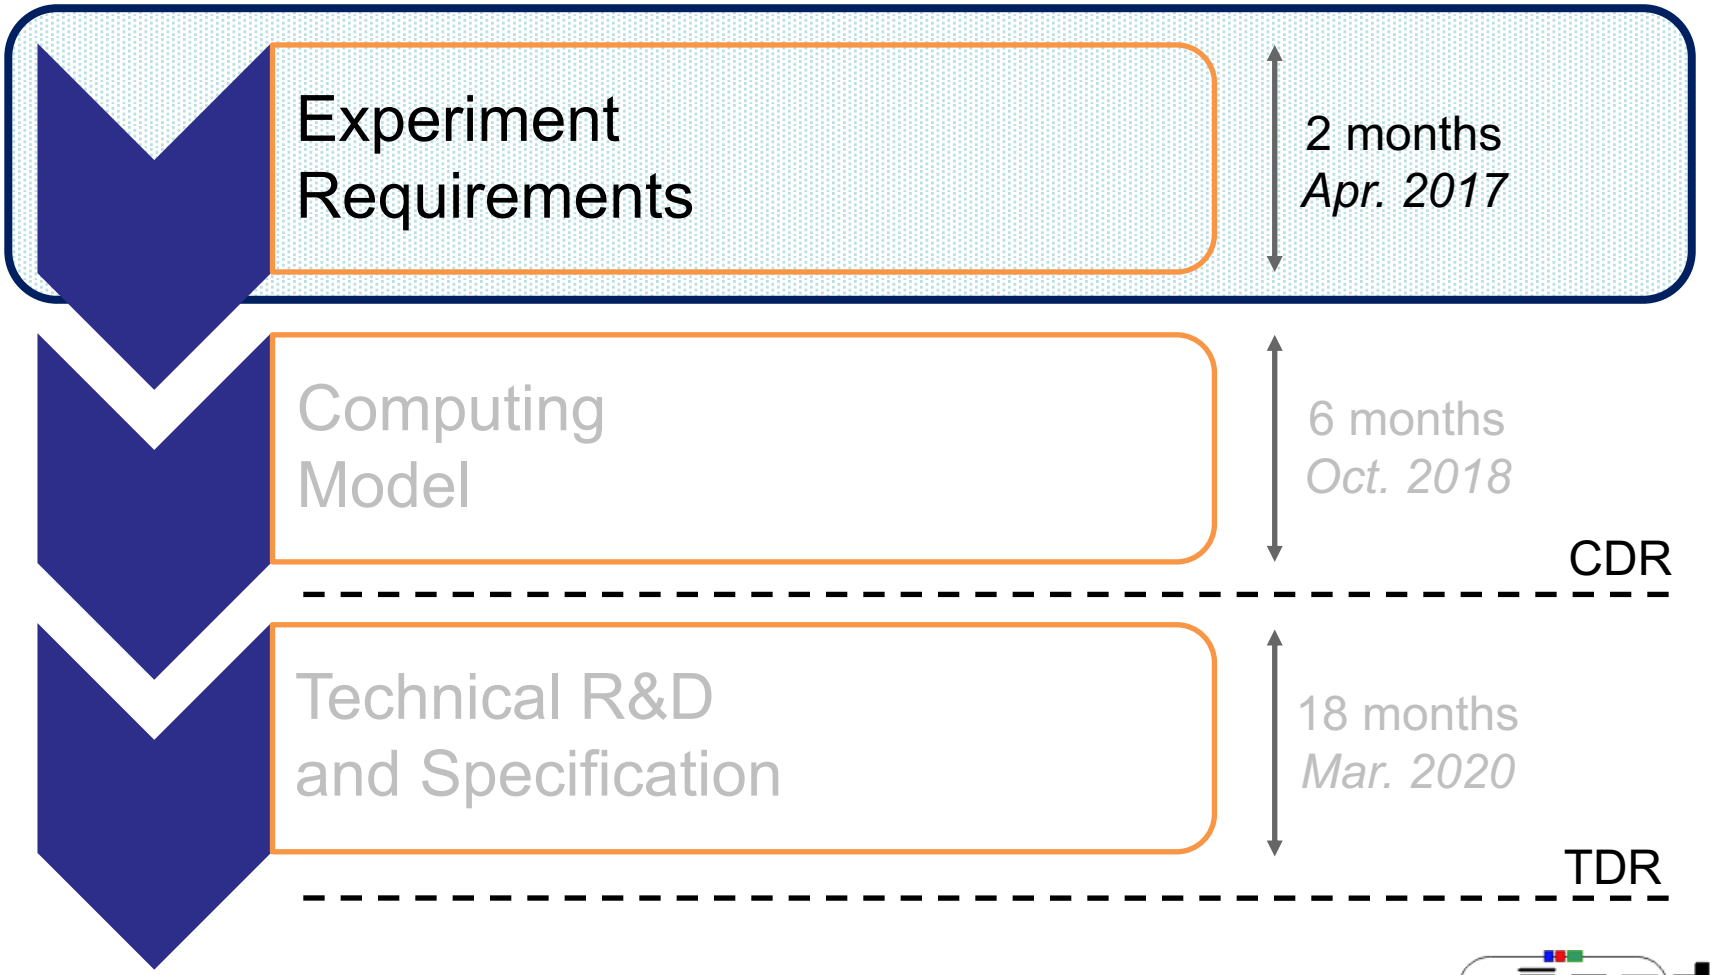

Status: 2009 Estimate

For reference: 2009 numbers

Programm	2010 CPU # cores	2010 On-line PB	2010 Tape PB/a	2013 CPU # cores	2013 On-line PB	2013 Tape PB/a	2016 CPU # cores	2016 On-line PB	2016 Tape PB/a
APPA	100	< 0,1	0,1	200	1,0	2,0	1.000	1,0	2,0
CBM	250	< 0,2	0,2	7.000	15,0	11,0	60.000	15,0	11,0
NUSTAR	100	< 0,1	0,0	300	2,0	2,0	2.000	2,0	2,0
PANDA	250	< 0,2	0,2	8.000	12,0	12,0	66.000	12,0	12,0
ALICE	3.000	1,5	0,0	6.600	3,5	0,0	14.500	8,0	0,0
FOPI	200	0,1	0,2	400	0,2	0,3			
HADES	1.000	1,0	1,5	1.000	2,0	3,0			
Gitter-QCD	2.500	< 0,3	0,3	20.000	0,5	0,5	100.000	1,0	2,0
Theorie	1.000	< 0,2	0,2	5.000	0,5	0,5	20.000	1,0	2,0
rest	500	< 0,3	0,3	2.000	1,5	1,5	10.000	3,0	3,0

Connection to the Green Cube

- 3500 optical fibers for CBM
- 500 optical fibers for PANDA
- 4 empty tubes

(10Gb/s fiber)

PANDA→ 2 X 310 GB/s

Maximum data rate from Panda data concentrators:

300 GB/s

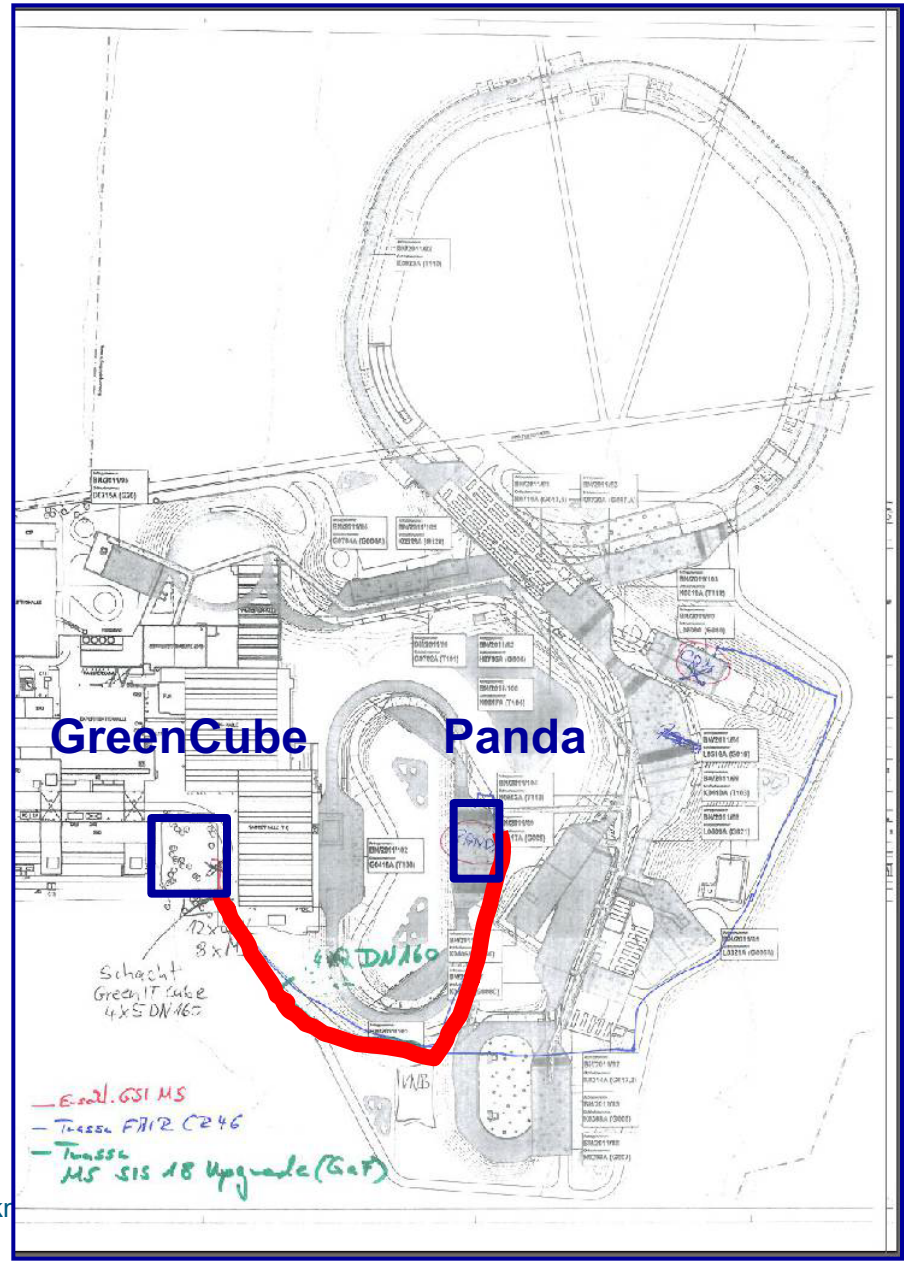

Maximum Average Luminosity

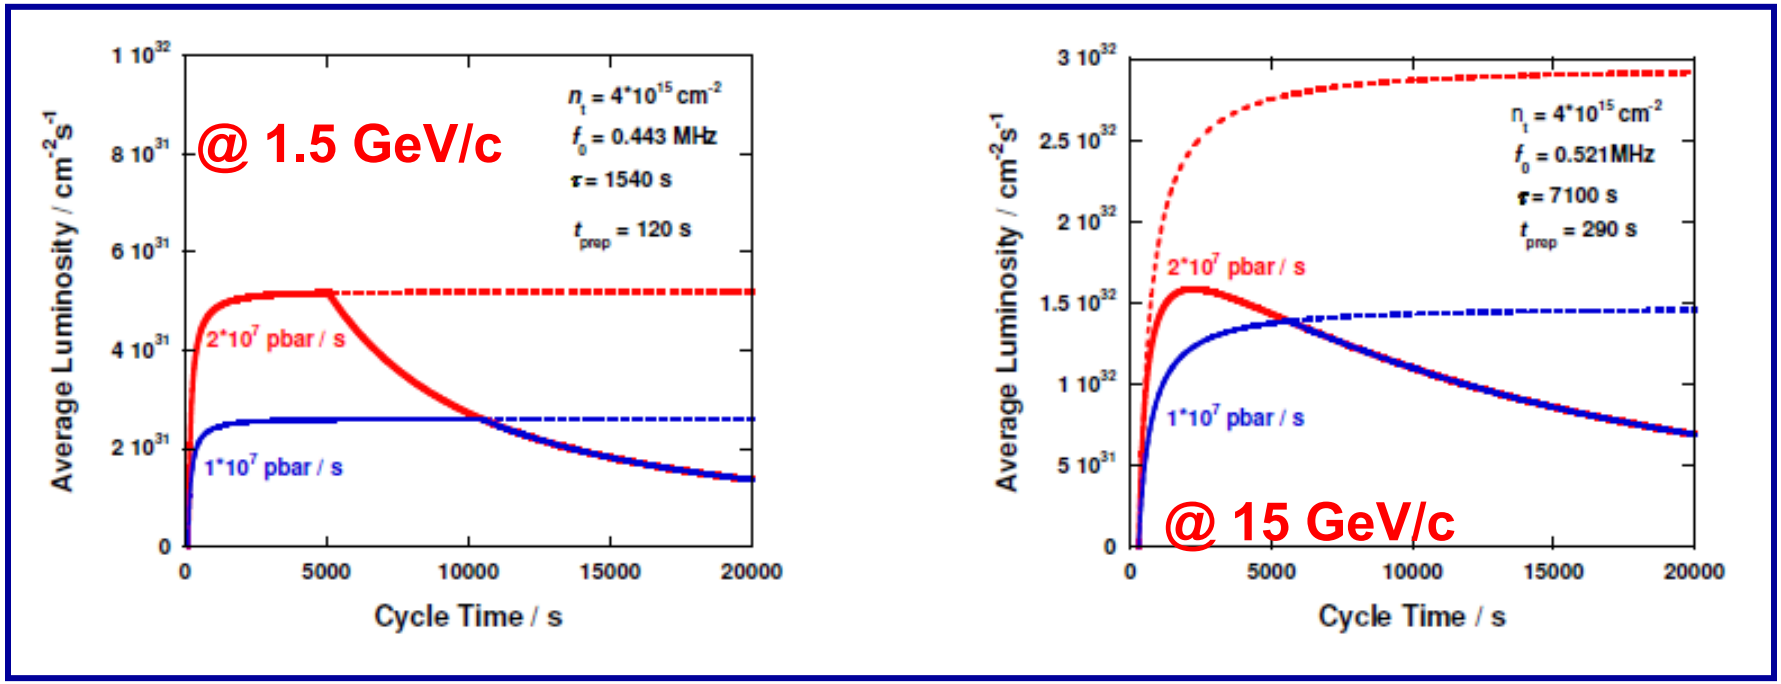

Beam Momentum	Max Average Luminosity	Cross section	Interaction Rate
1.5 GeV/c	0.5 10^{32} cm^{-2}s^{-1}	100 mb	5 MHz
15 GeV/c	1.6 10^{32} cm^{-2}s^{-1}	51 mb	8 MHz

At PANDA start we will have **1.2 10^{31} cm^{-2}s^{-1} ➔ ≈ 800 kHz**

Scenarios

- **Online:** Full rate from detector with suppression factor of 10 from Compute-Nodes. Reduced resolution for higher processing speed.
  - Rate: 800 kHz -> 80 kHz (due to FPGA processing)
  - Reduced resolution (e.g. no Kalman filter)
  - Processing stages: Reco + Tracking + PID + event selection
  - Suppression to achieve: 10

- **Offline:** Re-processing of data. Digis exist from experiment. Highest resolution incl. calibration and alignment to produce physics results.
  - Rate 80 kHz -> 8kHz (due to event filter)
  - Full resolution
  - Processing stages: Reco + Tracking + PID

- **MC Simulation:** Like offline but includes generation of MC Tracks, hits and digis by simulation.
  - Rate: How much more MC data than exp. data do we need?
  - Processing stages: Sim + Digi + Reco + Tracking + Pid

Evaluation of Requested Data

- Two simulations:
  - lxbk[0201-0552].gsi.de
    - *Qualified machines with defined HEPSPEC06 values*
    - *Simulation of 1, 10, 100, 1000 DPM events*
    - *1.5, 4.5, 7, 12, 15 GeV/c beam momentum*
    - *Geant 3 and Geant 4*
    - *Full PandaRoot simulation based on master macros*
    - *Simulation time and memory consumption*
  - MacBook Pro
    - *Modified macros to give “clean” digi, reco file sizes*
    - *15 GeV/c FTF with Geant 3*

Summary of processing times

	Sim	Digi	Reco	Tracking with Kal.	Tracking wo Kal.	PID
1.5 GeV/c	0.42	0.03	0.03	0.34		0.03
7 GeV/c	1.24	0.04	0.04	0.54		0.07
15 GeV/c	2.82	0.06	0.05	0.64	0.05	0.12

	Sim	Digi	Reco	Tracking with Kal.	Tracking wo Kal.	PID	Event Filter	Sum [s / evt]
Online			0.05		0.05	0.12	0.01	**0.23**
Offline			0.05	0.64		0.12		**0.81**
Sim	2.82	0.06	0.05	0.64		0.12		**3.7**

- Processing times dominated by Kalman filter and MC Simulation

Summary of processing times

	Time per event	HS06
Online	0.23 s	5.06
Offline	0.81 s	17.82
Sim	3.7 s	81.4

	# events / year	HS06 sec (1 event)	HS06 sec (all)	CPU Efficiency	HS06	# gen	Cores (22 HS06/core)
Online (80 kHz)	$7.0\ 10^{11}$	5.06	$3.5\ 10^{12}$	85 %	480,000 (100 days)	1	21,600
Offline (8 kHz)	$7.0\ 10^{10}$	17.82	$1.2\ 10^{12}$	85%	48,000 (350 days)	4	8,700
Sim (8 kHz) x 2	$14\ 10^{10}$	81,4	$11\ 10^{12}$	95%	392,000 (350 days)	4	72,000

- ✓ 100 effective days of data taking
- ✓ 800 kHz event rate from detector
- ✓ Suppression factor for online 10 from FPGAs
- ✓ Suppression factor for offline 100 (FPGAs + CPUs)
- ✓ 350 days of continuous CPU activity / 100 days for online
- ✓ CPU efficiency for sim 95%, for real data 85% (lower)

Event size for Simulation

FTF 15 GeV/c

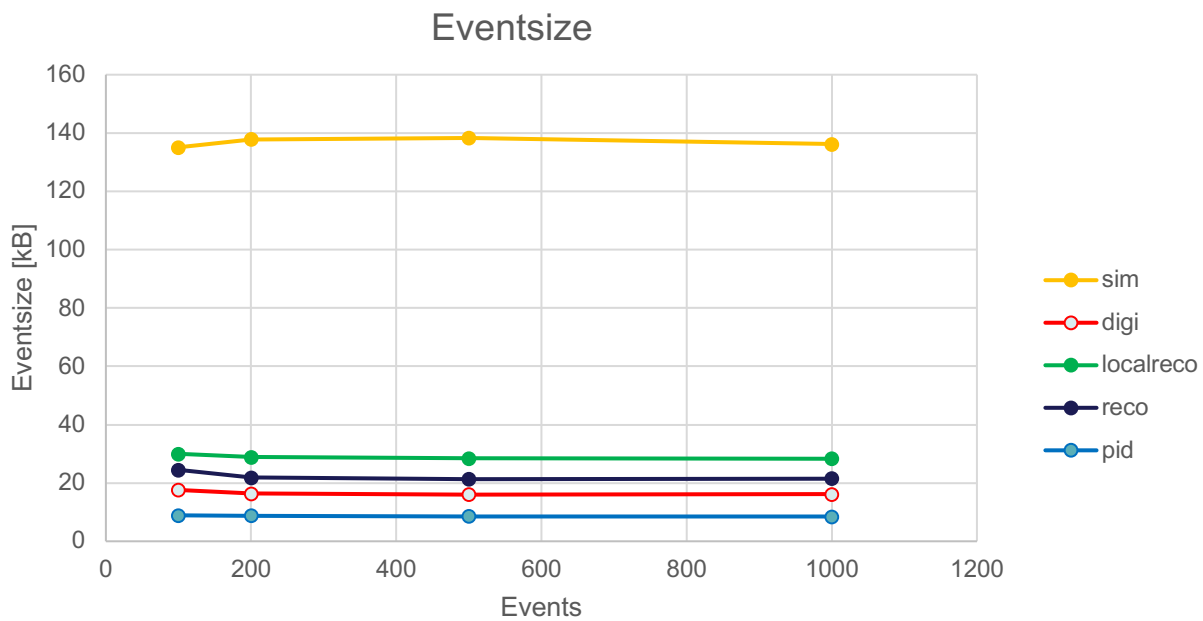

- No difference seen between G3 and G4
- 15 GeV/c produces largest files (1.5 GeV/c about half)
- Still large uncertainties:
  - Data structures not optimized for size
  - MC truth propagation
  - Each track should be fitted 5 times for different pid assumptions
  - No time-based sim but should not be a problem at low lumi

Summary of File Sizes

File Size per	Sim track + point	Sim only track	Digi	Reco ESD	Tracking ESD	PID ESD/AOD	SUM
Event **[kByte]**	137	4.2	16.5	28.9	22.3	8.7	213.4
Second at 8 kHz **[MByte]**	1,070	32	130	226	174	68	1,668
Online / year			1,062				1,062
Offline / year				1,860	1,436	560	3,856
Full Sim	17,638		2,124	3,721	2,871	1,120	27,474
Reduced Sim	17,638					1,120	18,758
Small Sim		529				1,120	1,649
Small Sim + Digi		529	2,124			1,120	3,774

[TByte]

- Online creates Digis which are reprocessed in the offline process 4 times
- Offline keeps data from a single reprocessing run. 4-times more if we keep all.
- Sim is 2-times the data of the experiment
  - Full Sim: contains all processed data
  - Reduced Sim: Only MC tracks, MC points and particle cand.
  - Small Sim: Only MC tracks and particle candidates

Summary of File Size II

- Sum of data:
  - Online + Offline + Full Sim: 32,392 TByte/a
  - Online + Offline + Reduced Sim: 23,674 TByte/a
  - Online + Offline + Small Sim: 6,567 TByte/a
  - Online + Offline + Small Sim + Digi: 8,692 TByte/a

- Small Sim sufficient for data analysis but no reprocessing with different reconstruction algorithms possible
- Small Sim + Digi: Allows reprocessing but no MC point matching
- Reduced Sim allows reprocessing of data but all steps have to be repeated
- Full Sim allows reprocessing of data without repeating all steps

Open questions

- Online: FPGA reduction by 10 reachable?
- Online: Precision w/o Kalman sufficient?
- Online: How many parallel event selection lines?
- Online: Online re-alignment and calibration?
- Simulation: How much more sim than measured?
- Alignment and calibration not in
- How long do we keep data on disk?
- Future processing steps?
- Reduction of data size
- Optimization of algorithms

Final Numbers for FAIR CDR

Compute Requirements (in HEPSpec06)		
Compute Class	**Day 1 configuration**	**Full MSV configuration**
I.a	-	-
I.b (*)	-	-
I.c	-	-
I.d (*)	-	-
II.a	70 k cores	100 k cores
II.b (*)	50 k cores (6 months)	100 k cores (6 months)

I: responsibility of experiment (on site, in GreenCube, exlclusive, shared)
II: responsibility of FAIR (in GreenCube, secured as average, exclusive)

Storage Requirements (I – Data Taking)

Storage Requirements (I – Data Taking)			
		Day 1 configuration	**Full MSV configuration**
Experiment to GreenCube/RZ1	#fibers	500	500
	Bandwidth (MB/s)	300 GB/s	300 GB/s
Bandwidth to permanent storage	Peak (MB/s)	20 000	20 000
	Average (MB/s)	10 000	10 000
Permanent storage/year (TB/year)		18 000	18 000
Additional disk storage (TB)		100	100

Remarks:

- #fibers only if dedicated fibers from Experiment to GreenCube/RZ1 are planned; other n/a
- Additional disk storage: additional temporary disk storage needed during data taking, e.g. caching for delayed event filtering