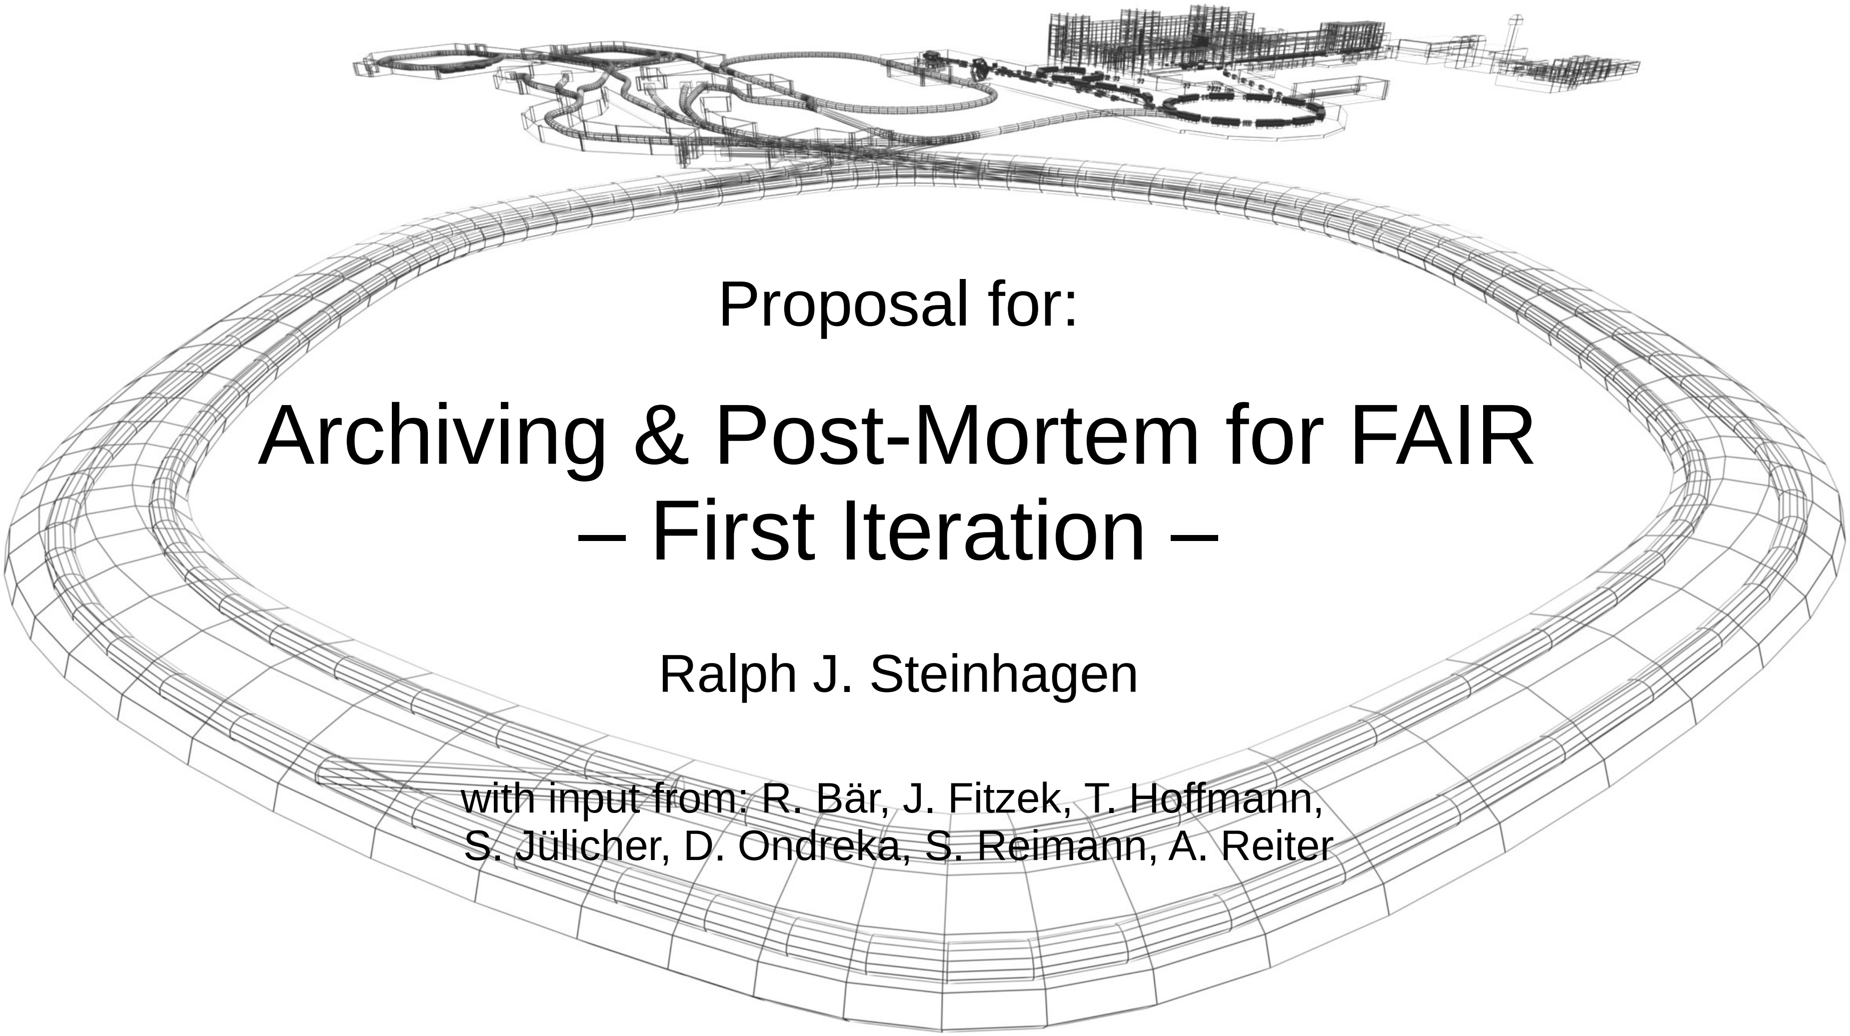

Proposal for:

Archiving & Post-Mortem for FAIR – First Iteration –

Ralph J. Steinhagen

with input from: R. Bär, J. Fitzek, T. Hoffmann, S. Jülicher, D. Ondreka, S. Reimann, A. Reiter

Archiving

Post-Mortem

Aims of Archiving & Post-Mortem

- *Collect and store pertinent accelerator data centrally to permit analysis of the accelerator performance and its proper function*

- *Terminology – distinguish*
  - 'Data Archiving' – continuous (“slow”) data acquisition (aka. 'Logging')
    - periodic, reduced-frequency data (lower bandwidth)
    - purpose: reconstruct machine/beam condition for later refined re-analysis and/or correlation with other experiments
      - example: intensity, emittance, orbit, Q/Q', PC currents, …
  - 'Post-Mortem' – transient data recording
    - asynchronous, high-frequency data (high bandwidth)
    - Purpose: reveal cause of emergency beam abort / possible equipment damage & validate correct funtioning of protection systems
      - example: QPS, PC currents, turn-by-turn BI data after quench

Typical Use-Cases
Classic Examples

op. period #1 | op. period #2

Case #1: Cyc#1 Cyc#2 Cyc#3 Cyc#4

Case #2: $t_0 \rightarrow t_1$ $t_2 \rightarrow t_3$

what's being logged

Q1: Is performance of period

$F_{perf.}(t_0 \rightarrow t_1)$

**better,
worse, or
equal?
'>' | '<' | '='?**

$F_{perf.}(t_3 \rightarrow t_4)$

Q: performance indicator computed in high-/low-level (GUI vs. daemon)?

Q2: Stability of subsequent/selected cycles?

Cyc#1
Cyc#2
Cyc#3
Cyc#4

<Cyc> <Cyc²>

cycle-to-cycle
average & overall variation

→ important for cycle-to-cycle feedbacks
→ important for data reduction (discussed later)

Typical Use-Cases
Cycle-to-Cycle Feedback I/II

- Cycle-to-cycle feedback/ optimisations via archiving system?
  - A priori very similar to use-case discussed before, but with additional more sophisticated cycle selection:
    - Intensity/transmission range
    - correlation with emittances
    - More detailed/complex beam measurement cuts (de-select BPMs, etc.)
  - Initial prototyping: GUI-level implementation
  - Long-term: background daemon process

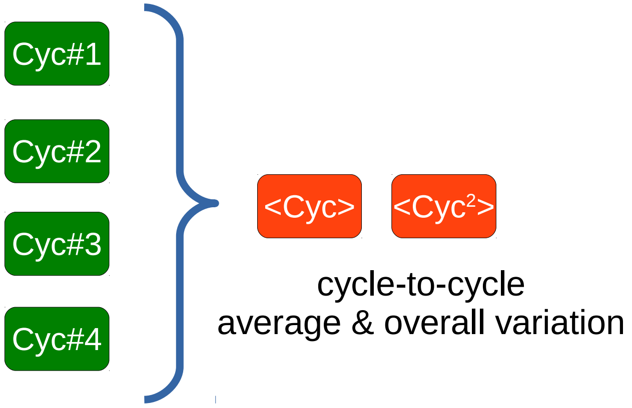

Typical Use-Cases
Cycle-to-Cycle Feedback II/II

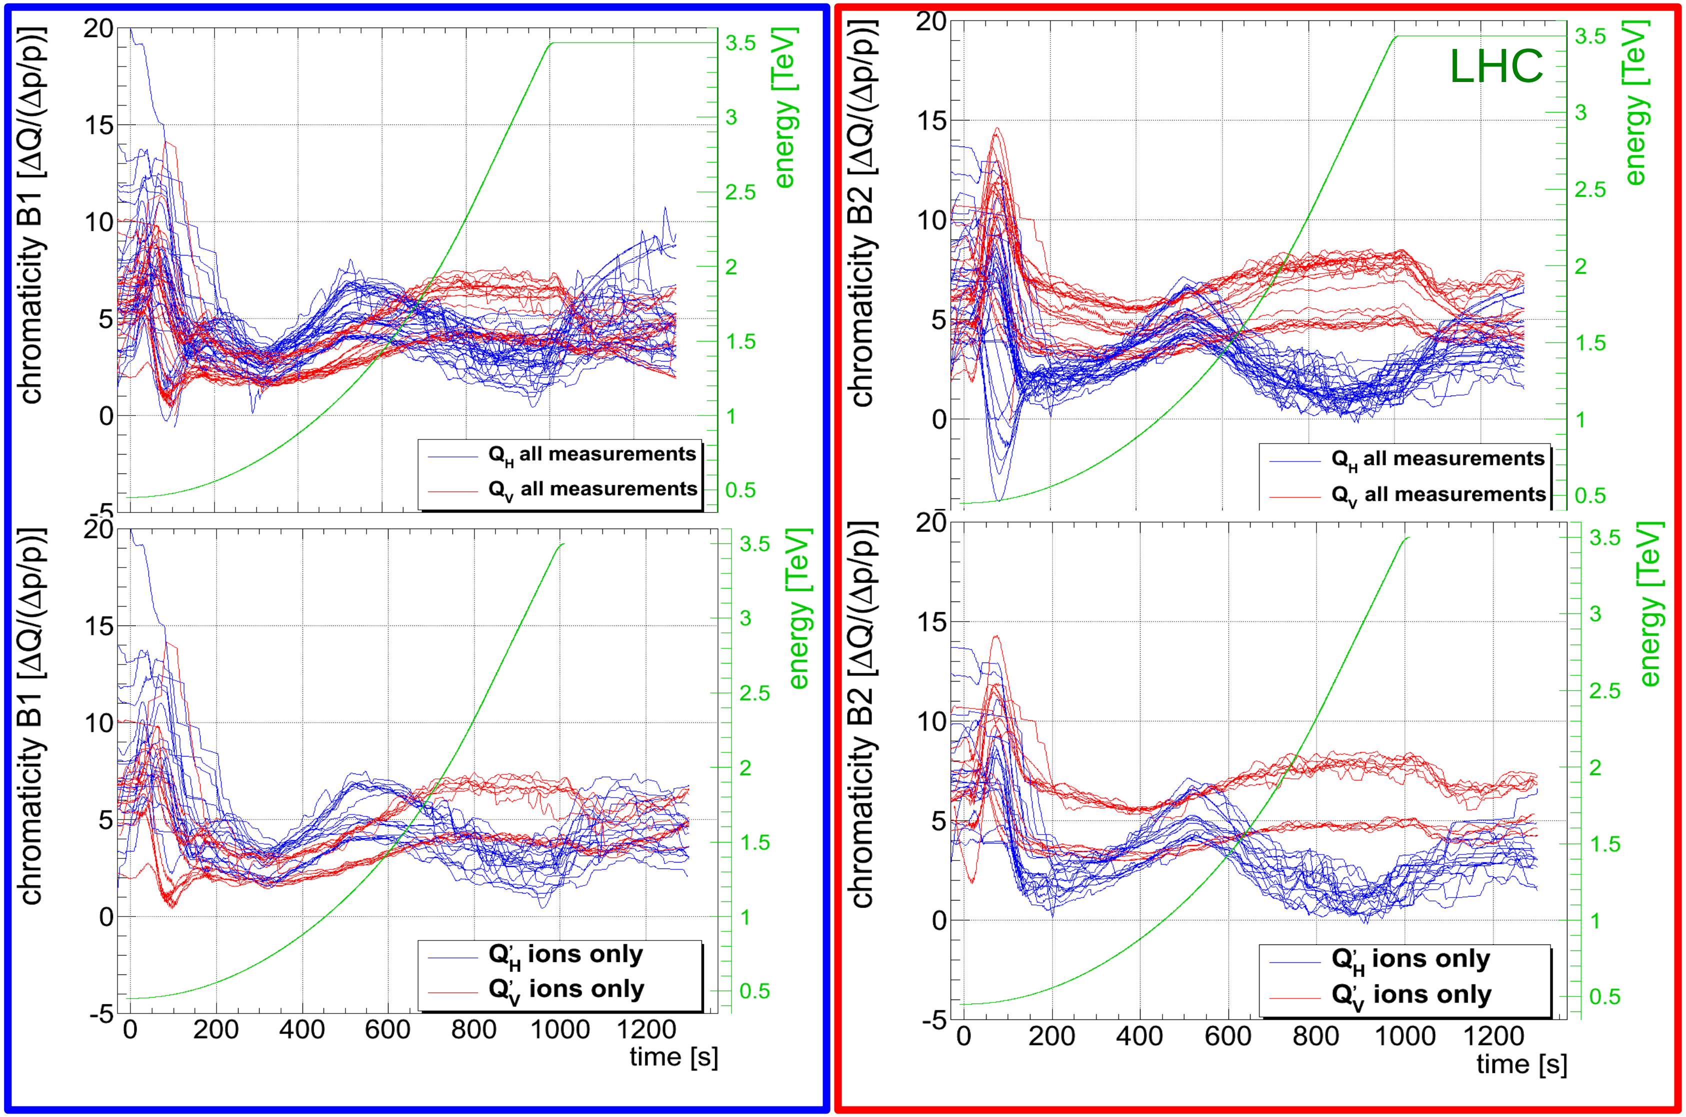

Typical Use-Cases
Analysis of 'Near-misses'/Rare Events I/II

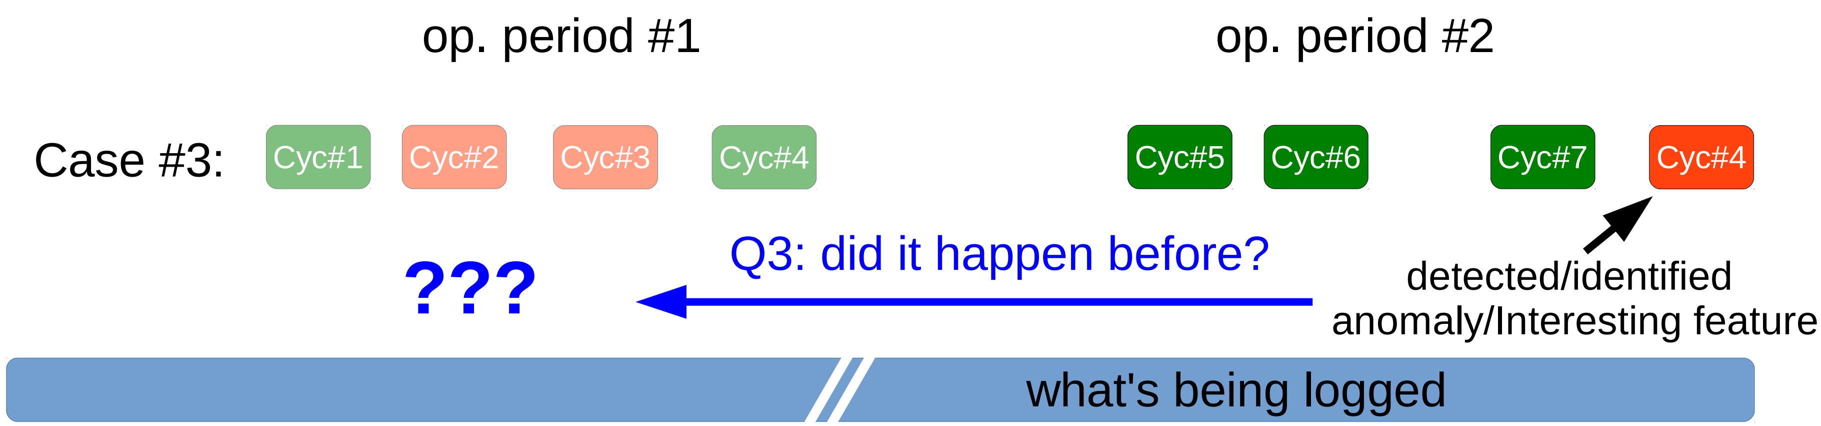

- Permits a quantitative analysis on events which may trigger mitigations or improvements based on the frequency of the individual events, e.g.
  - analysis of near-misses not caught by PM/MP.
  - Comparison with earlier performances & availabilities:
    - “Have we been able to do it better in the past?”
    - “What are the main systems/parameters impacting availability?”

Typical Use-Cases
Analysis of 'Near-misses'/Rare Events II/II

- Studying of such cases helps improving operation by minimising these events and/or proposing new counter measures intercepting these at an earlier (/safer) stage
- Some examples:
  - Poor control of orbit/emittances (e.g. slow drifts over few cycles) → increasing losses in ring/transfer lines → eventual quench
  - … → cryo temperature increase → reduced temperature margin → quench
  - PM occurring & masked SW interlock
  - but also: establishing reliable reference and tolerance thresholds ↔ optimising machine availability
    - e.g. beam-loss (/transmission) thresholds driven by past “good” cycles ↔ reject acceptable tail of cycle-to-cycle distribution

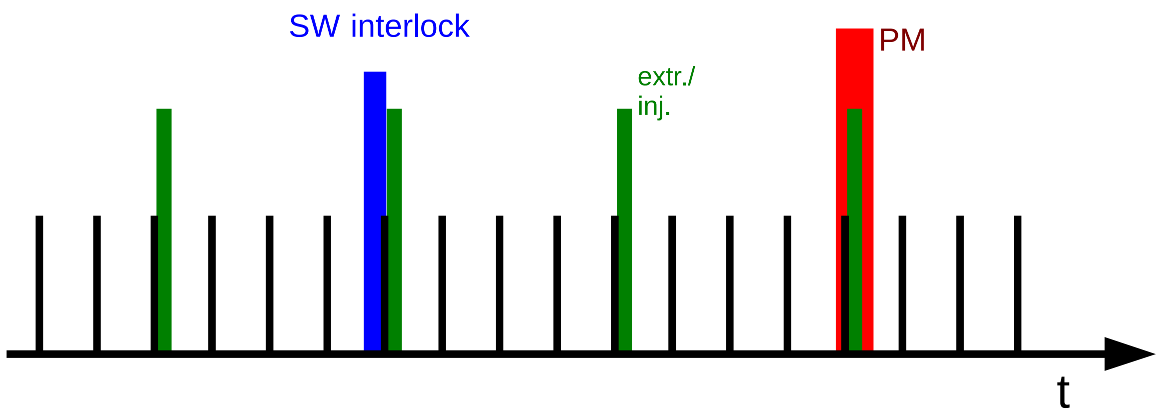

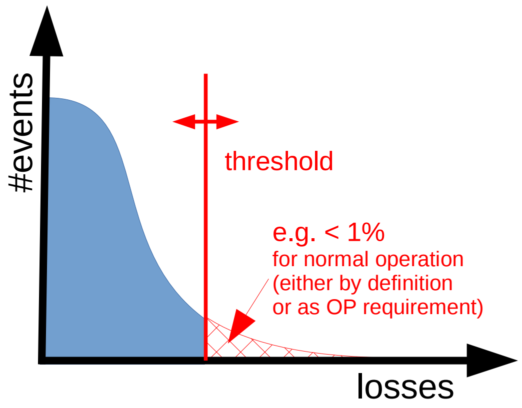

Hierarchy & Coverage
Preparing for the Unknown

“Reports that say that something hasn't happened are always interesting [..], because as we know, there are known knowns; there are things we know we know. We also know there are known unknowns; [..] But there are also unknown unknowns – the ones we don't know we don't know. [..] the latter category that tend to be the difficult ones. The absence of evidence is not evidence of absence, or vice versa.”, Donald Rumsfeld in 2002

- *We do not know in advance which data might be interesting/useful later*
  - 'Somebody’s noise is somebody else’s signal'
- MD studies often reveal unexpected effects (that's why we do them)
  - relevant data might not have been recorded
- Analysis of rare events
  - typically not easily reproducible
  - might have a history of ‘near misses’
  - might not be detectable by post-mortem
- Repetition of study would waste beam time and create unnecessary costs

→ Should aim at a broad coverage of accelerator monitoring to open the chance of providing evidence for identifying difficult operational or machine protection scenarios.

Hierarchy & Coverage
Archiving and Post-Mortem

- Conceptually both fulfil a similar (/the same?) purpose: improve operation and verify proper function of protection systems.
  - Most systems providing PM need to provide also measurement data*
- Historic view: 'Post-Mortem' was introduced for LHC because granularity of logging (archiving) was not adapted for fast transients & vastly different time scales between regular operation and failures*
  - LHC: six orders of magnitudes:
    - normal operation: seconds to minutes (N.B. sc. magnets)
    - failure modes: ms (slow magnet failures) → µs (fast kicker failures, beam instabilities, …)
  - SIS18/100: normal operation: ~ ms (e.g. ramp ~ 200 ms) vs. failure: ~ 1 ms (1000 turns)
- Historically: large & costly gap between “slow” logging and PM
  - constraints have substantially changed for FAIR

*N.B. Aviation industry slowly moves away from transient recorders and to an online streaming approach for redundancy and reliability reasons (lack of data coverage, to-short data coverage time scales, difficulties of recovering recorder, etc.)

Hierarchy & Coverage
Cost of Memory and Storage

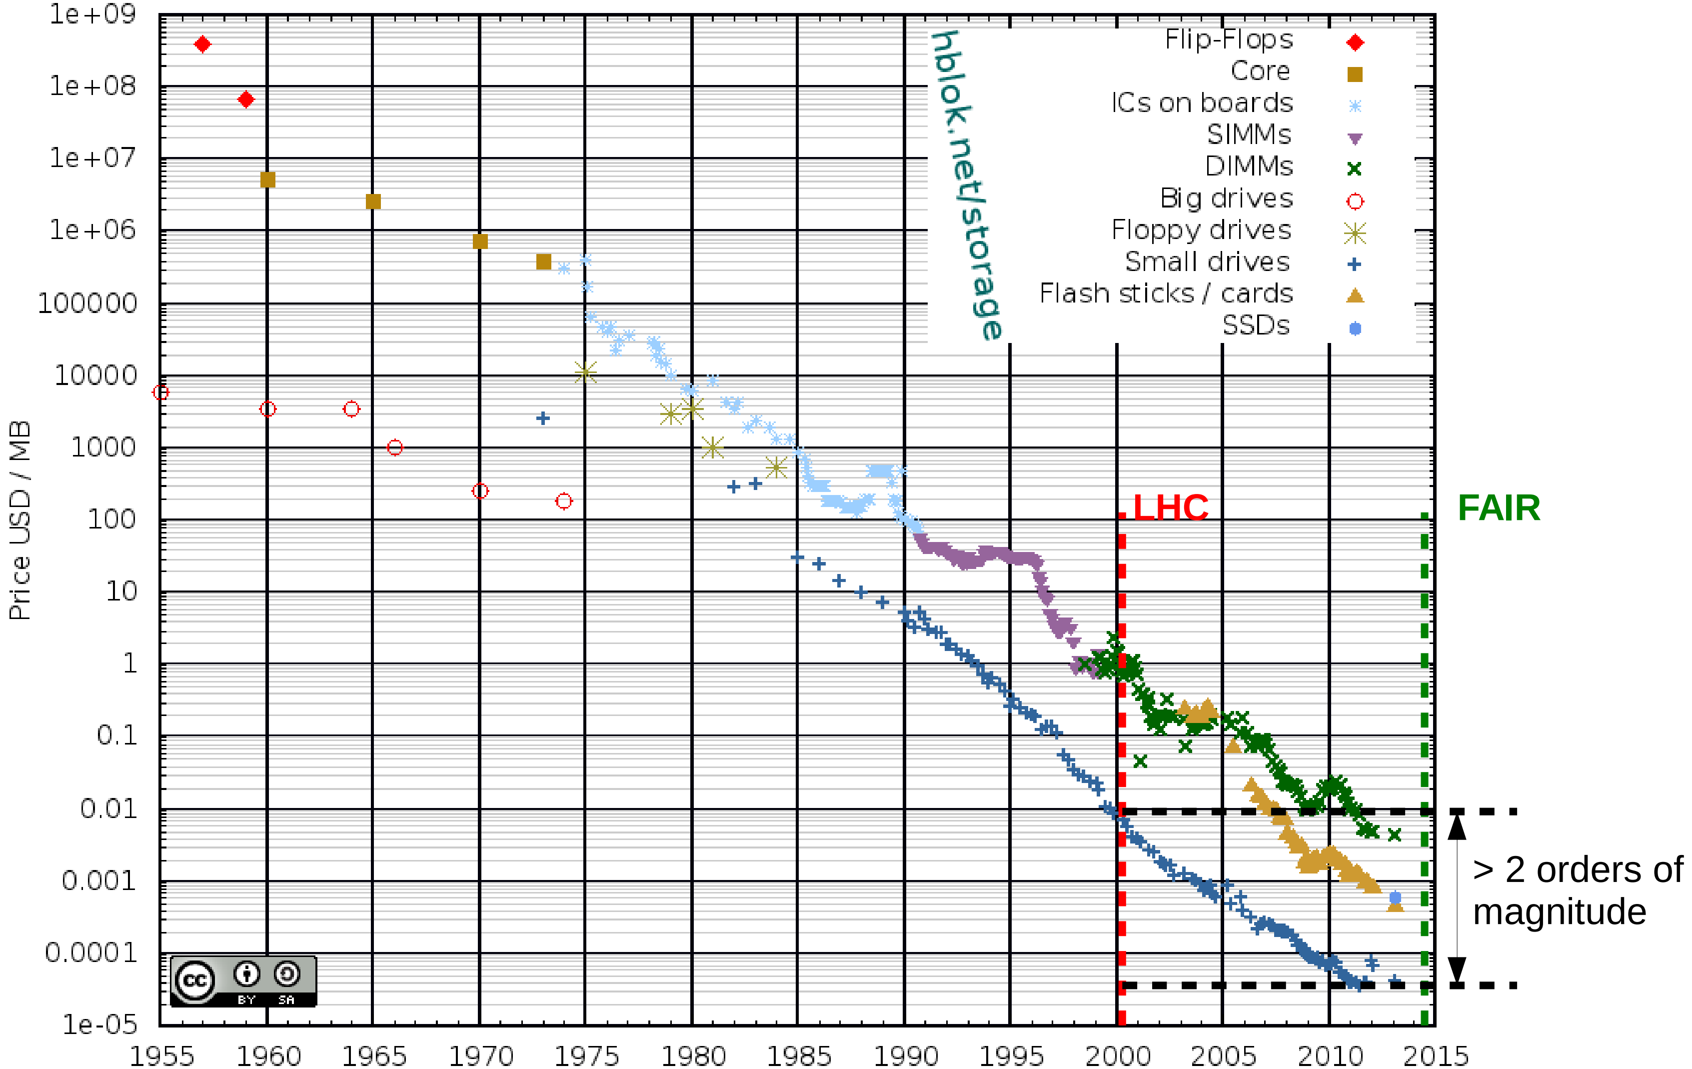

Hierarchy & Coverage
Example: by machine (complementary: by system group)

Root:

- CR
- ESR
- FRS
- HESR
- p-Linac
- SIS18
- SIS100
  - ...
- SFRS
- …
- UNILAC

+ Transfer-Lines (implies synoptic overview to find the one of interest!)

- Beam Instrumentation
  - Beam Intensity
  - Beam Losses
  - Beam Position: Orbit & Trajectory (first/last turn, ...)
  - Beam Profile: Ionisation Profile Monitors (IPMs), Screens, Wall-current-monitor, wide-band current transformer
  - Tune, Chromaticity & Optics
- Beam Feedbacks
- Beam Absorbers & Collimators
- Cryogenics
- Kickers
  - Exciters (AC-dipole, aperture kicker, tune kicker, KO exciter)
  - Injection, extraction & emergency beam dump kicker
- Power Converters & Quench-Protection
  - Corrector magnets
  - Lattice sextupoles
  - Main dipole, main quadrupole
  - Septa magnets
- Radio-Frequency Cavities
  - Main cavities, bunch merging, bunch splitting
- Radiation Monitors (N.B. personnel safety, not BLMs)
- Timing & (HW/SW) Interlocks
- Vacuum
- Access System, general infrastructure (power, ethernet, ...)

Hierarchy & Coverage
More specific Example

- Beam Current Transformer (DCCT, FCT)
  - Avg. Intensity, Status (bit-mask),
    - average losses: e.g. SIS18@1kHz , SIS100@100Hz , CR@10Hz?
    - instantaneous losses:
      - normal: first 1000 turns (injection losses), before/after start ramp (capture losses), extraction
      - Post-mortem: e.g. last 10000 turns
- Beam Position Monitors
  - Position, StDev, Sum-Signal, Status (bit-mask), Mask (enabled/disabled)
    - Orbit rate: e.g. SIS18@1kHz , SIS100@100Hz , CR@10Hz?
    - Trajectory rate:
      - normal: first 1000 turns (injection steering), transition crossing, last 10 turns (extraction steering)
      - Post-mortem: e.g. last 10000 turns
- Beam Loss Monitors (ionisation chambers, cryo-collimators, etc.) ↔ probably similar to BPMs
- All power converter
  - Current, RefCurrent, Voltage, RefVoltage, Status, PCTemperature, CableTemperature
    - Average rate: e.g. SIS18@1kHz , SIS100@100Hz , CR@10Hz?
- Kicker: injection/extraction kicker waveforms ↔ probably similar also long. diagnostics (FCTs)
    - Average rate: typically 2 x 1 turn, e.g. nanosecond resolution ↔ 2 x 4000 data points

Need to establish list to realistically assess coverage and implication on data rates/volumes! → Excel-Sheet

Some Numbers – First Order
(rough upper ball-park figure)

- Total number of devices:

Rough #device estimates:	Total	SIS18	SIS100	CR	HESR	HEBT	Super-FRS
Total	**2032**	**166**	**436**	**129**	**129**	**912**	**260**
Power Converter	**1317**	100	300	70	70	600	177
Beam Instrumentation	**435**	30	100	23	23	200	59
Vacuum	**220**	24	24	24	24	100	24
RF	**60**	12	12	12	12	12	0

- Very rough estimate of number of variables: ~ 50k (factor 2 margin)
  - N.B. power converter or BPMs: typ. 10 variables per device (including time stamp)
    - N.B. does not include profiles measurements (large volume, low rate)
  - → Bandwidth: ~ 0.2 MB/s @ 1 Hz (slow) or 20 MB/s @ 100 Hz (fast)
    - N.B. single hard drive: ~ 50 – 120 MB/s
  - → Volume: ~ 6 TB/year (slow) or 11 TB/week (fast)
- For comparison: LHC 10^5...10^6 variables @ 0.1-1 Hz (< 8 MB/s, 2004)

Some Numbers – Max. Rate I/II

What's needed from a Physics/Control Theory Point of View

- Sampling 'Rule-of-thumb' (control theory)

$$\boxed{f_s > 20 \ldots 40 \cdot f_{bw}}$$

  - N.B. min. of phase est. errors, $f_{bw} \sim 1$ / 'time constant' τ of the physics process
  - Applies for basically all feedback loops

- For comparison:
  - SIS18: → $f_s \geq 257$ Hz e.g. 1 kHz
    - 10 T/s, 1.73 $T|_{max}$ → $\tau_{(10\text{-}100\%)} \approx 0.16$ s
  - SIS100: → $f_s \geq 94$ Hz e.g. 100 Hz
    - 4 T/s, 1.9 $T|_{max}$ → $\tau_{(10\text{-}100\%)} \approx 0.4$ s
  - CR: → $f_s \geq 1.5$ Hz e.g. 10 Hz
    - 0.054 T/s, 1.6 $T|_{max}$ → $\tau_{(10\text{-}100\%)} \approx 27$ s
  - HEBT & Super-FRS:
    - single shot/slow extraction → $f_s \approx 0.1$ or 1 Hz OK? 10/100/1k Hz for slow-extraction?

Some Numbers – Max. Rate II/II

What's needed from a Physics/Control Theory Point of View

- Fold in more realistic rates per machine:

		SIS18	SIS100	CR	HESR	HEBT	Super-FRS	Total
Total variables:		**166**	**436**	**129**	**129**	**912**	**260**	
proposed sampling	[Hz]	1000	100	10	10	1	1	
Data rate	[MB/s]	6.33	1.66	0.05	0.05	0.03	0.01	
Data amount (weekly)	[TB/week]	3.65	0.96	0.03	0.03	0.02	0.01	**4.69**
Data amount (year)	[TB/year]	190.45	50.02	1.48	1.48	1.05	0.30	**244.77**

  - Instantaneous load appears to be less of an issue
    - Total bandwidth: < 10 MB/s (N.B. DVD: 1-2 MB/s, Blu-ray: 5-6 MB/s)
    - Required buffer: ~ 5 TB/week
  - Should aim at a total storage amount of less than 5-10 TB/year
    → 25-50 x data reduction needed

Data Reduction Possibilities I/III
Global/Machine-State dependent Options

- Only short-term data storage
  - Some variables loose importance after a few weeks
  - 'expert-only' variables
    - could change to long-term storage if deemed necessary
- Drop BI data for periods w/o beam
  - maybe trivial but worthwhile mentioning
  - relies on reliable beam current information
- 'Machine-Mode' dependent data rates
  - Shut-down, access, 'machine check-out', beam operation
- 'Beam mode' dependent data rates
  - No-beam/prepare, pilot beam (setup with safe intensities), intensity ramp-up (setup with unsafe beam), production operation

proposal for next FC^{2}WG meeting

- Beam-process dependent data rates:
  - no-beam, Injection, ramp, flat-top/(fast) extraction, slow-extraction, coasting, ramp-down

Data Reduction Possibilities II/III
Short-Term vs. Long-Term storage

- 1 kHz for few weeks → 1 Hz
  - Option: keeping high-frequency data on request for MD purposes
  - Option: beam mode dependence: 'high' (1 kHz) during setup/intensity ramp up, 'low' (10/100 Hz) during production runs?
- On-change data reduction
  - Easy for boolean/integer values (e.g. 'on','off', statuses etc.), non-trivial for floating point values ('*somebody's noise is somebody's signal*')
  - Possible compromise – is this acceptable for most systems:
    a) Store <min/mean/stdev/median/max> of cycles for given periods (with or w/o beam), keep full snap-shot every e.g. 1 hour
    b) Store transients if e.g. <actual value> outside 'n·<stdev> band' for given cycle
- Exception:
  - Full beam profiles at injection & extraction
  - PM data causing beam-abort or quench → keep last minute before PM

Data Reduction Possibilities III/III
On-change Data Reduction – Example

Raw Data Rate:
– what is written to short-term DB

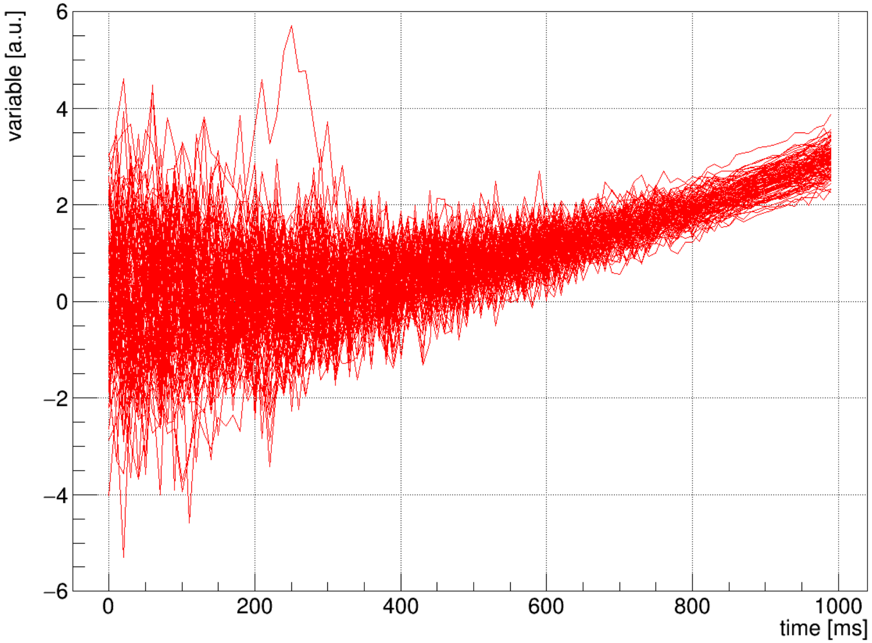

Reduced Data:
– what is written to long-term DB

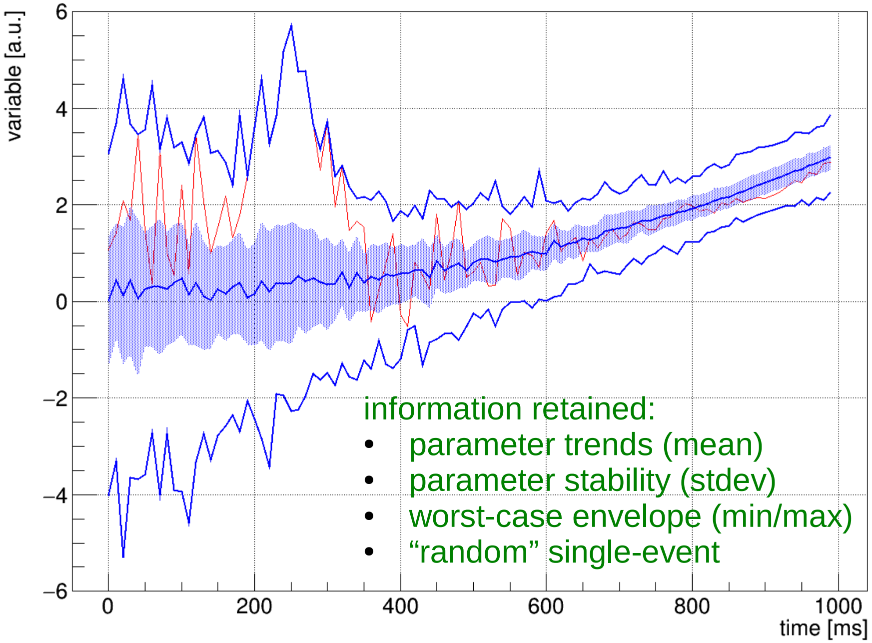

~ factor 30 data reduction

Data Storage Paradigm & Required Read-Bandwidth

- Required 'write' bandwidth probably < 10 MB/s
- May need much more read-bandwidth:
  - Typical use-case is a 'b-tree' type search:
    1) Find given cycle/events within a large time span
    2) read sequentially detailed/full data before/after the event
    3) Compare data sets across large time scale (e.g. long-term temperature drifts, linear time-scale)
    4) Compare cycle-to-cycle evolution (pseudo linear-timescale with large gaps)
  - multiple user and random access (different datasets/interest)
    - Number of parallel user: 10, 20, … ,100?
    - Can expect 100 MB/s from typical HDD!?!
  - option: reduce supported bandwidth/redundancy gradually over past years (↔ users will less likely access older datasets)
- Required CPU performance:
  - Local vs. distributed data analysis?

Data Extraction I/II
Selection Criteria

- Time selection – keep selection of multiple windows/intervals/cycles and collate
  - Time window: $t_{start} \rightarrow t_{stop}$
    - sub-selection: last day, last week (with operational relevant end day/periods definition → for statistics)
  - Time interval: $t_{reference} - \Delta t \rightarrow t_{reference} + \Delta t$
  - Cycle range: t_1(serial id #1) → t_2(serial id #2)
- Sub option:
  - Only select cycles/beam processes with beam
    - ie. ignore beam-out, ramp-down, cycles with missing linac injections
  - Only retrieve data between two events
    - e.g. injection→extraction, start-ramp→end ramp
  - Only select cycles for a given beam type
    - U^{28+}-only, U^{92+}-for-CBM, 'protons for $\bar{p}$-production' → rationalisation needed to minimise the number of options at hand → summarise in meta-information
- Meta-Information (filter conditions)
  - Beam parameter
    - Ion-species, beam target (experiment), actual intensity, actual beam transmission, targeted intensity range, ...
  - Machine parameter
    - Ramp-rate, min./max. rigidity, injection/extraction energy, slow/fast extraction, cycle/store length
  - Cycles with(-out) post-mortem events
  - Operational ranges
    - OP year, between technical stops, OP week, OP day, etc.

Data Extraction II/II
Some Open Questions

- Data extraction should be driven use case examples
- 'Object oriented' vs. variable-by-variable design?
- Data export and low-level interfaces
  - Mainly expert analysis scripts
  - OS support? Java, C++, both? ROOT? Others?
- Who adds new data sources, ensures naming convention, … ?
- How to treat data that is affected by calibration curves? Experiment requirements?
  - e.g. FastBCT, emittance, optics, …: Keep best value? Update recently stored values with more accurate version?
  - Should one aim at storing values always with an error/uncertainty estimates?
- To be kept in mind: UTC vs. local time issue (↔ correlation with logbook)

Additional Features

- Instrumentation of Archiving System
  - Who read/writes what?
  - Amount of data? Throughput & response times?
- Ability to save/load preferred or commonly used list of variables (e.g. beam current & beam energy)
- Java GUI Data Extraction Tool – "Face" of the Archiving System
  - Should aim at implementing most (90%) of the data analysis use-cases described before (also could serve as a documentation/example for other developers on how-to retrieve data, etc.)
- Website with standard variables might be useful
  - intensity/beam energy vs. time, transmission vs. time, beam brilliance vs. time availability for physics vs. time, ...

Next Steps

A Specification for circulation/approval by Q4-2015

B Are there other important use-cases

C Get better data rate/volume estimates:

- detailed list of devices
- variables per device class (1st iteration)
- Post-mortem data

D Meta-Information (filter conditions)

- Need more input/confirmation from potential users

Prepare and circulate spread-sheet to MPL, equipment groups, etc. detailing the number of variables, types, expected logging rates, ...

E Check logging time constants

- particularly for smaller machines (CR, HESR, …)

F Is proposed data reduction paradigm acceptable?

G First conceptual proposal for architecture/required HW?

Appendix

Boundary Conditions
(from a data management point of view)

A) Data Coverage:
- Shall/can we aim at full coverage? What should be logged?

B) How to acquire/write data to file
- Bandwidth, latencies?
- Pushing vs. pulling of data?

C) How to organise the data
- Hierarchy & naming convention?
- Short-term vs. long-term storage & performance?
- Data reduction & pre-processing?
- Data proxying?

D) How to retrieve & analyse data
- Bandwidth, latencies?
- Standard/default use-cases?
- GUIs, low-level interfaces (data format, supported programming languages, etc.), ...

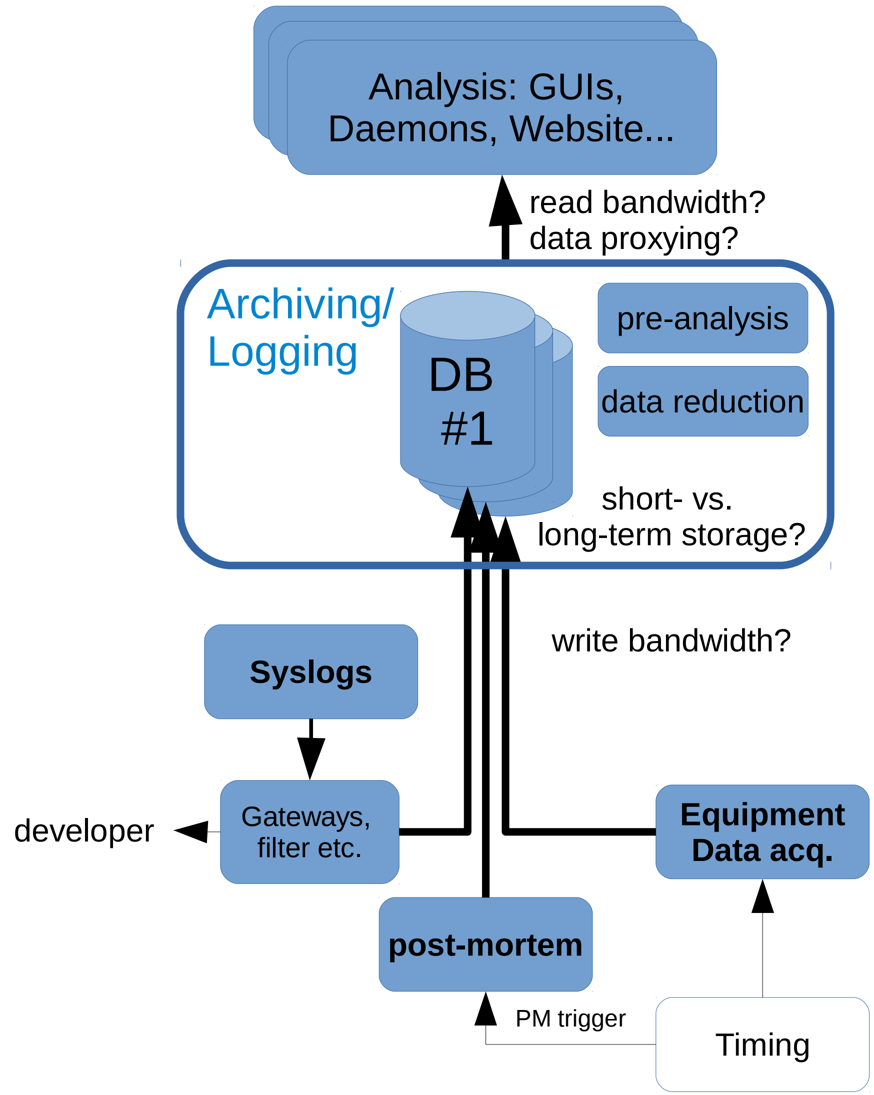

Aim of Logging/Archiving

- Collect and store pertinent accelerator data centrally to permit analysis of the accelerator performance and its proper function

- “Detailed Specification of the FAIR Accelerator Control System Component”:
  - C. Handel, R. Mueller, “*[..] 'Diagnostic Logging System'*”, F-DS-C-10e, 2012-08-10
  - L . Hechler, “*[..] 'Archiving System'*”, F-DS-C-11e, 2011-08-10
  - J. Fitzek, “*[..] 'Post Mortem System'*”, F-DS-C-13e, 2011-08-10

- 'Functional specifications' → more details needed:
  - Q1: Data coverage? Targeted bandwidths? Targeted long-term data retention? Short- to long-term performance requirements?
  - Q2: Storage formats? Data retrieval, interfaces, etc?

Typical Use-Cases
Examples from GSI Operation (D. Ondreka)

- FAIR accelerators are partly aliens
  - We have strong expectations about their behavior
  - Will (hopefully) largely turn out true
  - Be prepared to reveal the hidden 'features':
    Log data as much as you can
- Examples from GSI operation
  - Mysterious reduction of SIS18 current:
    - No transmission change in UNILAC
    - By accident UNILAC profile grids had been printed
    - Vertical beam position changed
    - Traced to change of beam request timing
  - Sudden pressure rise in SIS18 extraction sector
    - All vacuum valves closed
      (logged, but not order nor source of vac. interlock)
    - FRS suspected guilty, but no hard evidence
  - Unexpected activation of H=2 cavity in SIS18
    - Comparison of beam loss patterns might help chasing down the source, but no data available
  - Dynamic vacuum questions
    - Topic often comes up in analysis of MD studies
    - Of course, nobody thought about recording…

Why we should log data as much as we can:

- We don't know in advance which data might be interesting/useful later
- Performance evaluation
  - Why was beam XY better/worse than before?
  - Search for long-term drifts and their causes
- Collect data routinely
  - No risk of forgetting to record data
  - Accumulation of data not only during MDs
  - Large data sets for all kinds of analysis
  - Somebody's noise is somebody else's signal
- New machines/operation modes
  - Backup data in case of unexpected problems
- MD studies often reveal unexpected effects
  - Relevant data might not have been recorded
  - Repetition of study would waste beam time
- Analysis of rare events
  - Typically not easily reproducible
  - Might have a history of 'near misses'
  - Might not be detectable by post-mortem

GSI Helmholtzzentrum für Schwerionenforschung GmbH D. Ondreka, FAIR Operation, 1st FCCWG Meeting 20.05.2015

- Smart and poor choices
  - Oracle number format
    - Flexible, often human readable, precision configurable
    - Fixed-point arithmetic – need to know the range of values in advance
    - Slow, wasteful for large amount of data
  - Floating point
    - Not human readable
    - Some standard issues (endianess not necessarily defined)
    - Precision adapted to scale of value
    - Often no conversion for further numerical processing

“Logging” of Text Information

- … debugging during development phases
  - loop counters, info messages etc.
- … but also examples relevant to operation, e.g.:
  - Software Interlock System
    - Boolean flag ('yes'/'no') indicating a global interlock
    - text information indicating which system/group issued the interlock and also which other active interlocks were masked → useful for machine availability & fault analysis
    - Later: inspect specific device causing the interlock
  - (Semi-)automated daemon/background processes
    - Cycle-by-cycle feedback
    - Real-time feedbacks

- N.B. @GSI: “logging” for 'textual' and “archiving” for 'measurement' data

→ propose to use common definition? e.g.: http://en.wikipedia.org/wiki/Data_logger