Large Scale Orbit Correction for the LHC

Ralph J. Steinhagen,

Accelerators & Beams Department, CERN and

3rd Inst. of Physics, RWTH Aachen,

Grindelwald, 2004-07-12

LHC schematics

- LHC is the first proton collider where a continuous control of the beam position is required during all operational phases (injection, acceleration, final focus, coast for physics)
  - Layout
    - two rings
    - L=26.7 km, $\Delta\mu_{FODO}\approx 90°$, avg. $\beta_{max}/\beta_{min}\approx 6$
    - 4 crossing insertions (coupling!) with detectors
    - ~ BPM spacing $\Delta\mu_{BPM}\approx 45°$
  - 1056 beam position monitors
    - measure in both planes
    - (2112 readings)
    - Precision (nominal beam): $\Delta x \approx 1\ \mu m$
  - 530 correction dipole magnets/plane
    - Time constants: $\tau \approx 203\ s$
    - bandwidth: $f_{bw} \approx 1$ Hz

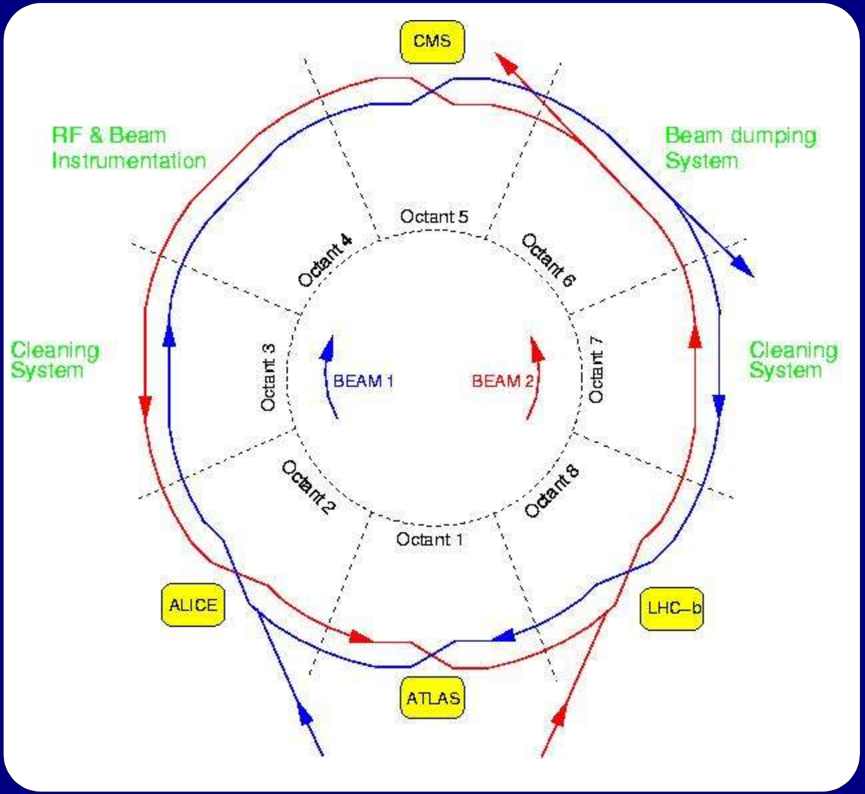

  - Required orbit stability:
    - ~70 µm in cleaning insertions, global constraint ≤ 500 µm
    - (up to 5 µm for TOTEM experiment)

- More on requirements/constraints:
  - J. Wenninger: “Orbit Stabilization at the Large Hadron Collider” (this workshop)

Central Global Orbit Feedback I/II

Central global orbit feeback controller:

- SVD based algorithm for spacial solution: $x(t\rightarrow\infty) \rightarrow x_{ref}$
  - Controller receives and sends data of all BPMs and CODs
- PID controller *D(s)* running @25Hz
  (@50Hz ultimate, imposed by the power converter controller design):
  - steers the (undelayed) power converter *G(s)*
  - Smith Predictor (SP) branch:
    - extension to compensates for constant propagation delay λ due to:
      - computation time, task switching, propagation delay in the network, ...

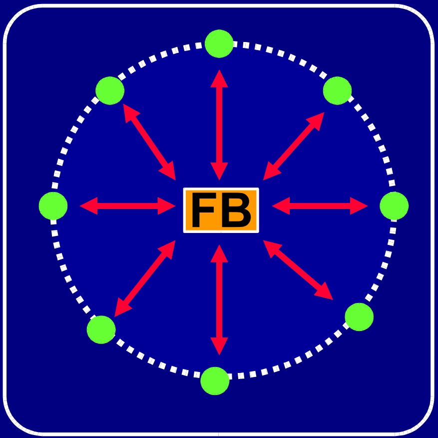

- More: *“LHC Orbit Feedback Control – Prototyping at the SPS”*, EPAC'04, Lucerne, 2004

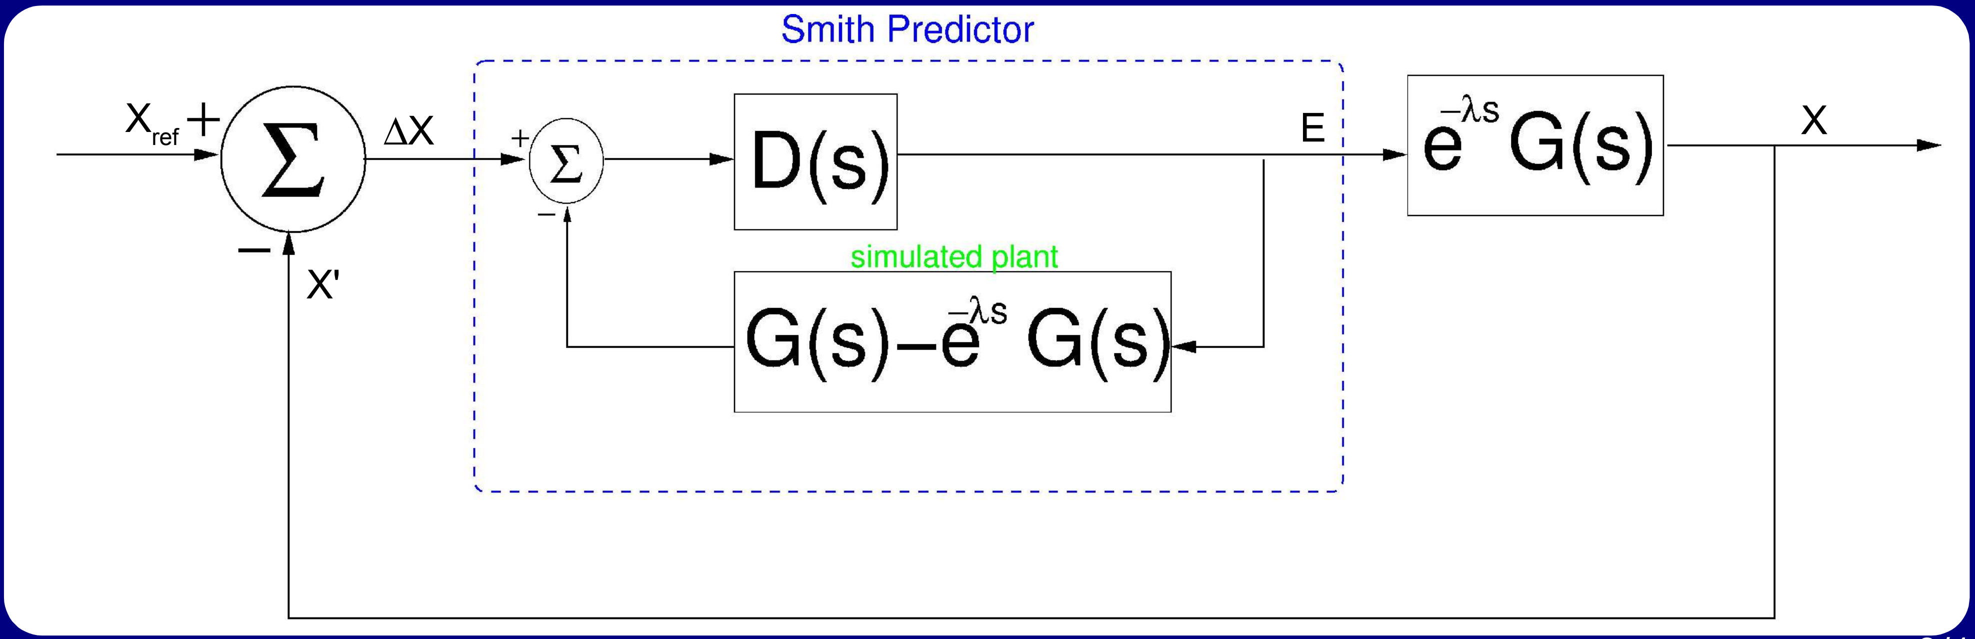

Feedback Control Layout

- BPMs (organised in ~ 70 front-ends):
  - acquisition electronics of up to 18 BPMs (DABs):
  - PowerPC (PPC) running LynxOS 4 for pre-processing of the BPM data
- power converters (~ 50 gateways):
  - WorldFIP gateway (f_{bw}=50Hz) to the up to 16 power converter controller (f_{bw}=1kHz)

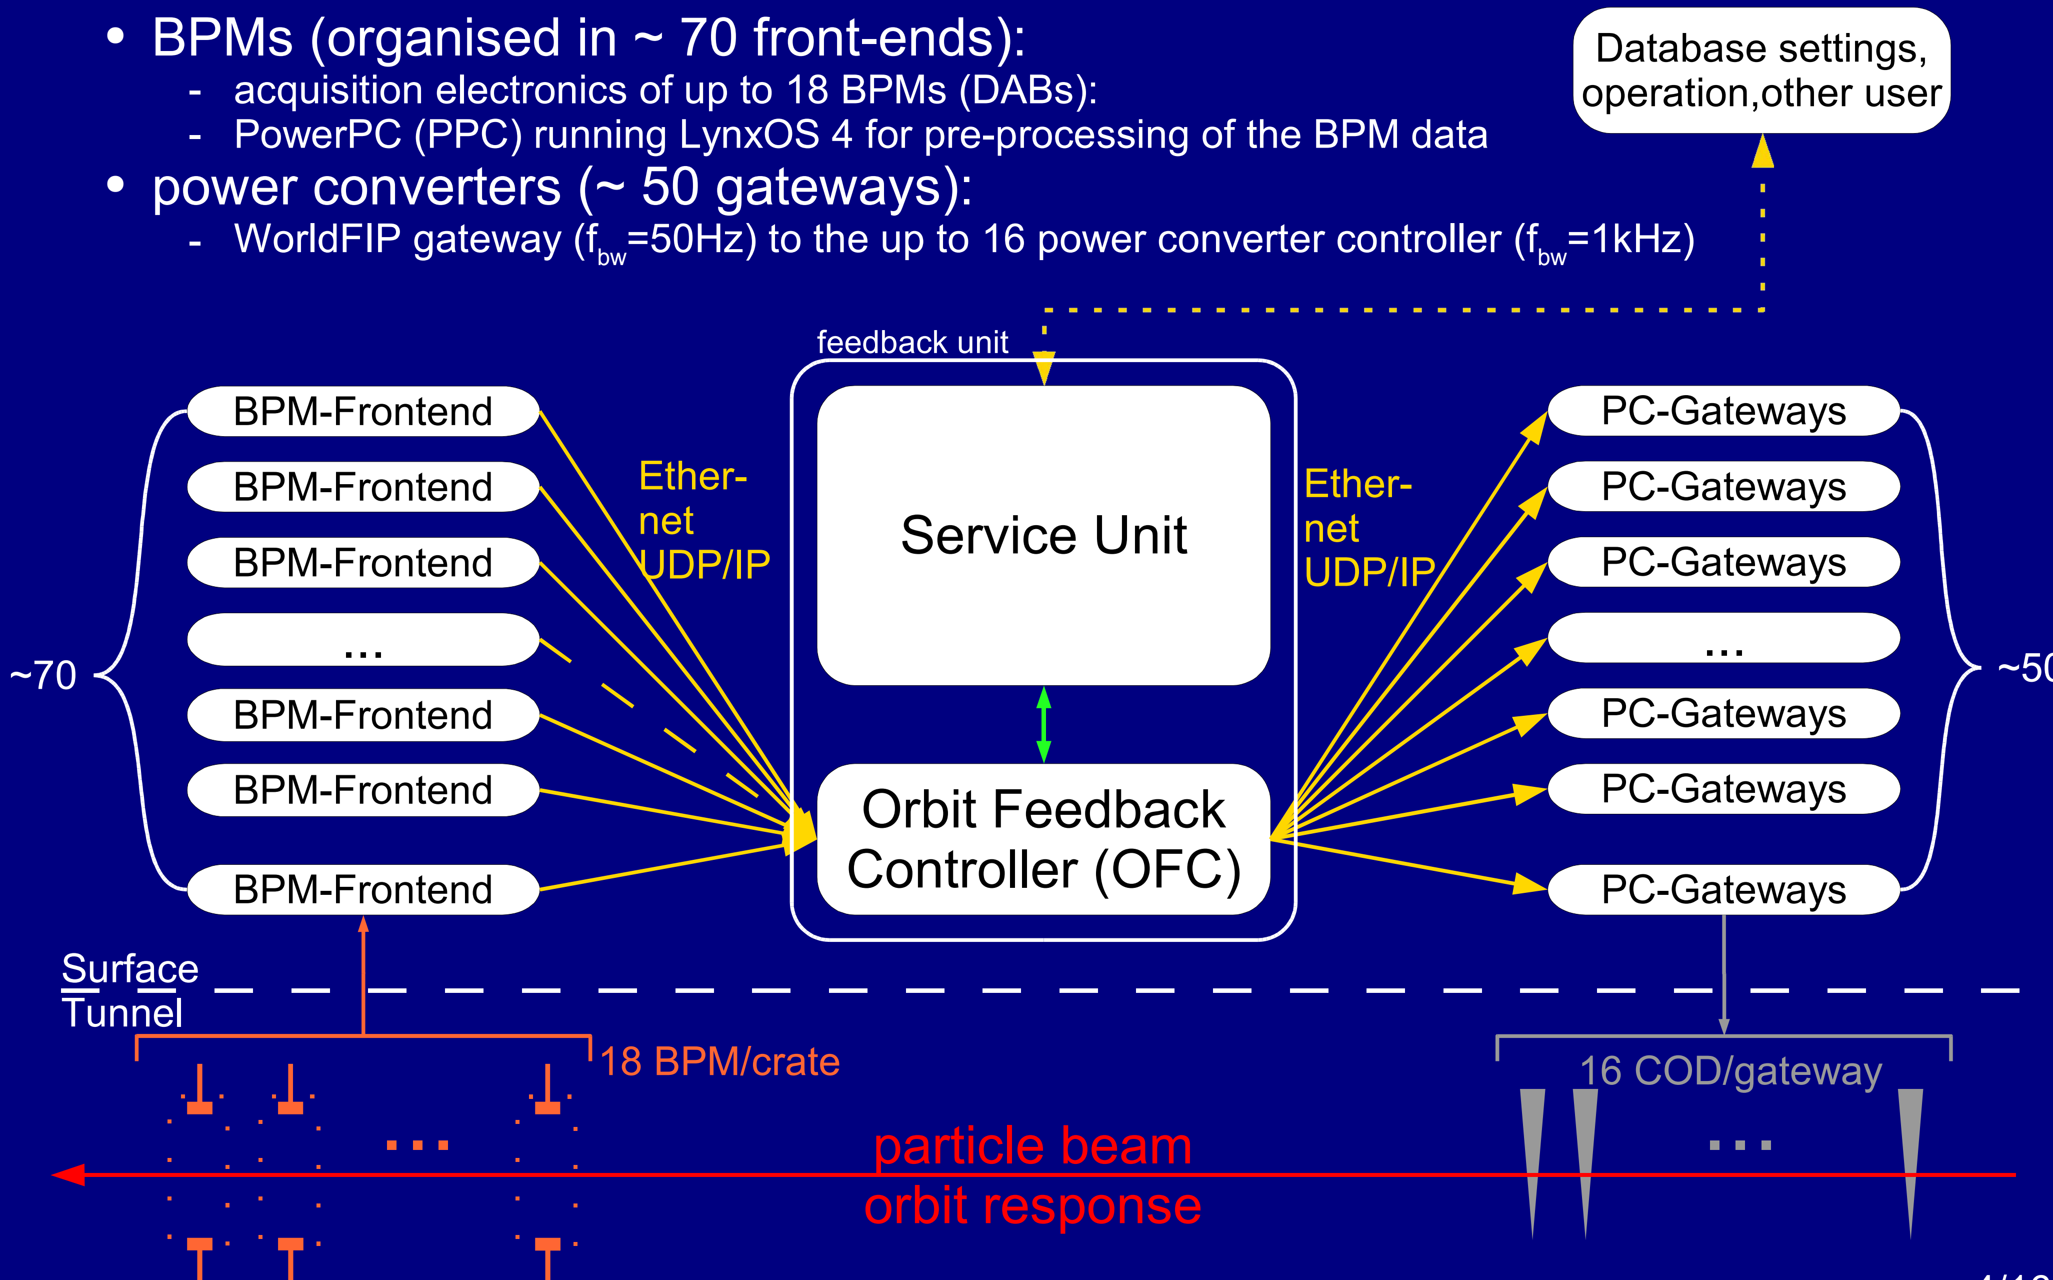

Split of feedback controller into two functional parts

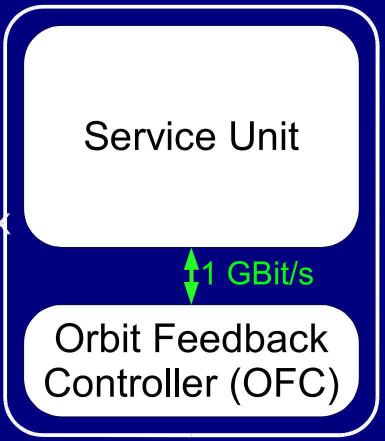

- Orbit Feedback Controller (OFC) - performs the actual fb:
  I. reception of the BPM data (UDP) through the Ethernet
  II. comparison of the measured orbit with a reference
  III. Correction in space: new steady state deflection using SVD based matrix
  IV. Correction in time: selected control algorithm (PID, Smith-Predictor)
  V. Conversion of the deflections angles into currents
  VI. Send the new settings to PC gateways

- Service Unit (SU) to the feedback controller:
  - Monitoring of machine states (energy, optics, mode [injection, ramp, squeeze,..)
  - Interface to LHC controls and other users
  - Interface to machine operation and experts.
  - Data monitoring (logging).
  - Sanity checks (BPM and COD faults)
  - Update of the orbit response matrices and quantities derived from it (SVD decomposition…) whenever the relevant machine or equipment conditions are modified

Both parts (presently) exchange data through a private direct gigabit data link

Feedback Controller: Use case

- Likelihood is high that one of the >2000 involved elements (BPM, COD) fails or is temporarily unavailable for control.
- Split of functions makes feedback system more robust and flexible for those operational cases.
- What happens when a BPM (or COD) fails? on-the-fly recalculation of the SVD based correction matrix:

Orbit Feedback Controller

- Short term (few seconds):
  zero-order holder of BPM position
  - Monitor/COD is marked
  - The position changes slightly during most operational phases:
    - OFC can reuse the last valid value
  - wait till new matrix update
  - Swap active matrix in feedback controller

Gigagbit Link

Service Unit

- Medium term:
  recalculation of the SVD matrix
  - Service unit marks the device
  - Recalculation of the inverse response matrix
  - Transfer the new matrix to the feedback controller

  - The feedback loop does not need to stop while recalculating the new matrix
    - Feedback requirements preserved in most parts of the rings
    - Only a small area and likely only one ring is affected

Orbit Feedback Controller

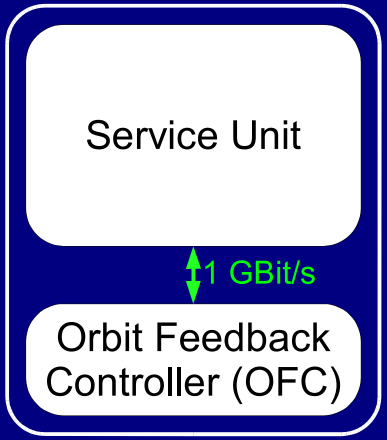

Controller and Service Unit differ by their criticality:

- Service Unit:
  - Dynamic load due to user interaction
  - complex tasks:
    - data monitoring
    - sanity checks of data
    - recalculation of SVD based parameter (several s on high-end CPU)
    - ...
  - soft real-time constraints
  - unavailability will not necessarily stop the feedback.
- Orbit Feedback Controller:
  - less dependent on control software environment
  - (logically) simple streaming task: receiving -> processing -> sending
  - very deterministic constant load
  - hard real-time constraints
  - Feedback loop stops immediately if not operational

- segmentation gives the possibility to change the implementation of the controller while avoiding the necessity of changing the Service Unit as well, in case the performance, reliability etc. must be improved.

Technical network

- CERN's Technical Network is the backbone for the LHC orbit feedback system.
  - Switched network
    - no data collisions
    - no data loss
  - Covers the 27 km circumference
  - Layout corresponds to geographic layout
  - double (tripple) redundancy

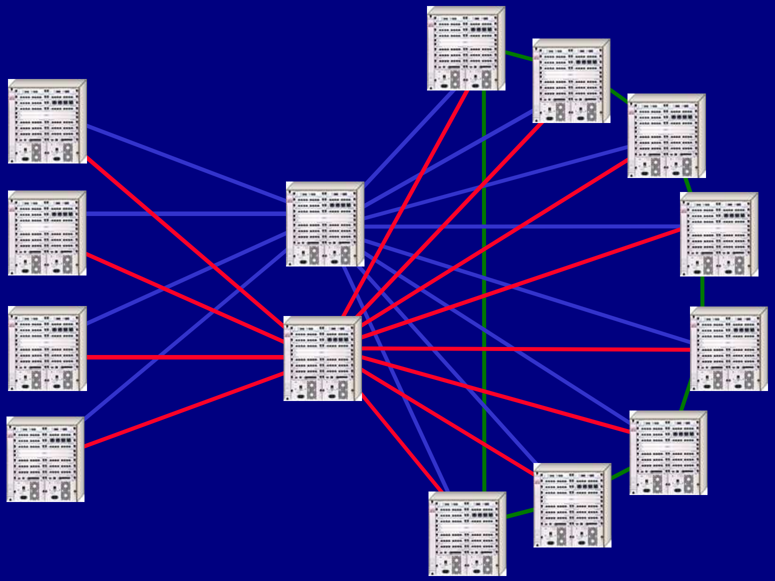

- Core: Enterasys X-Pedition 8600 Routers
  - 32 Gbits/s non-blocking
  - $3 \cdot 10^7$ packets/s troughput
  - QoS on hardware level
  - Routing delay ~ 3 µs
  - longest transmission delay (exp. verified) ~ 320 µs
    (500 bytes, IP5 -> Control room ~5 km)
    - 20% due to infrastructure (router/switches)
    - 80% due to traveling speed of light inside the optic fibre

Bottlenecks: Technical network?

Network seen as by:

- *Isn't Ethernet non-deterministic especially under heavy load?*
- *Can Ethernet technology be used for reliable real-time control loops at all?* Yes!
  - Depends on maximum feedback frequency and worst case delay jitter the control design can cope with.
- CERN's technical network:
  - Quality of Service (QoS):
    - packets are sorted and served based on their priority
    - 4 queues = 4 profiles
    - One queue dedicated to real-time feedback
      - private network for the orbit feedback
      - Orbit feedback 'real-time' data packets are served first
        - “Nearly” deterministic response
        - worst case router delay 300 µs
          (5 MB router cache fills up under 32 Gbits/s load and QoS)
        - Uses ~ 3% of network bandwidth
      - Theoretic total worst case transmission delay <1.5 ms
        (all TN router buffers filled up and need to be processed before serving the targeted packet)
  - Network's latency and jitter are negligible on the time scale of the feedback frequency (25/50 Hz < 1kHz).

Bottleneck I: Network in the high-level front-ends!

- The front-end network interfaces are presently the bottleneck. e.g. feedback controller @ 50 Hz:
- lots of in-/outbound connections:
  - Two types of loads:
    - Real-Time: BPM and COD control data
      - Avg. bandwidth: ~13 Mbit/s
      - short bursts: full I/O load within few ms
        (100 MBit/s resp. 1GBit/s, burst duration desired to be short in order to minimise the total loop delay)
    - Non-Real-Time:
      - transfer of new settings to OFC (correction matrix ~30 MB)
      - PID configuration etc.
      - relaying of BPM and feedback state data (monitoring/logging)
      - ...
      -
  - (Peak) load similar to high-end network servers
    - Nearly constant full load during certain operational phases
- network interface should be scheduled on the device level to provide a Quality of Service (QoS) for real-time data
  - One reserved FIFO queue for feedback data
  - General purpose queue for other data

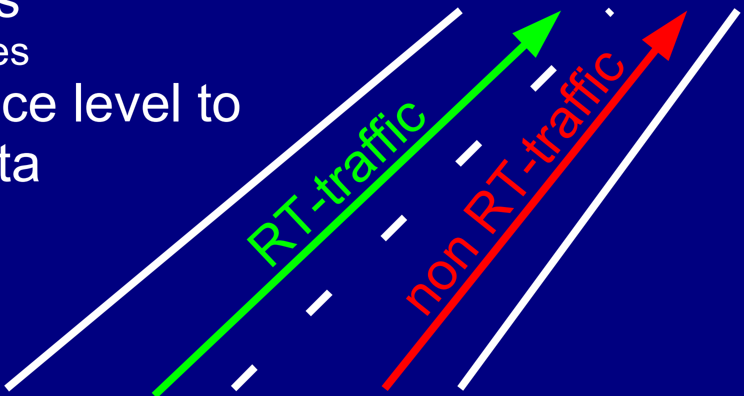

Bottleneck II: Orbit Feedback Controller!

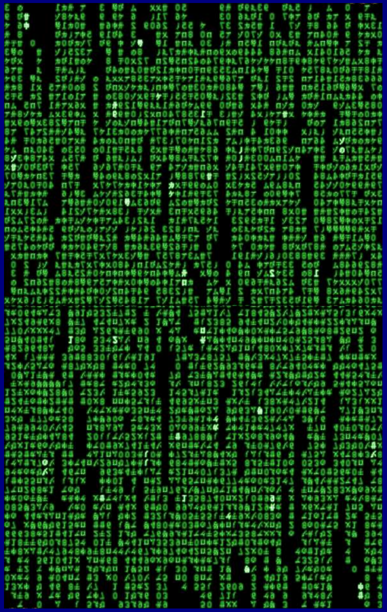

- Controller: must handle large matrices (~30 MB)
  - core of orbit correction:
    - multiplication of inverse orbit response matrix with input position vector: $\sim 4 \cdot 10^6$ *double* multiplications per sample @50Hz: ~ 400 MFLOPS
    - 1.5 GByte/s local memory data transfer
    - several ms processing time on a high-end SMP system
  - Similar requirements as for web, file or database servers:
    - high performance & high reliability, but:
    - hard real-time constraints: total execution time has to be deterministic and less than 20/40 ms to fit the 25/50 Hz feedback frequency requirement
- present test solution:
  - x86 based server: HP Proliant 380 DL, 2.8GHz Xeon SMP, 3 GByte RAM
  - 2 GigaBit Ethernet connection (one dedicated card to service unit)
  - hardware redundancy (2 power supplies, 2 disks, hw monitor, local watchdog, remote hw reboot ...)
  - running CERN's 'Scientific Linux' with recent 2.6 kernel
    - First real-time capable Linux kernel, found at www.kernel.org
    - under heavy (CPU & I/O) load maximum task response latency (TRLT): $\tau < 135\ \mu s$
  - Processing duration per feedback cycle: ~12 ms

Orbit Feedback Controller - Delays

- Poles of COD-System are sufficiently compensated by the PID controller
- Total delay determines the system's actual bandwidth and performance
  (Small variation of execution and propagation time (jitter) is inevitable due to use of buffered multi-threaded OS, Ethernet and technology)
- Delays are inevitable and part of digital control systems. Some sub-systems that contribute to the open loop delay:

– Beam position Monitors:	acquisition (255 turns@f_{rev}~11kHz)	~ 23 ms
	processing and sending	~ 5 ms
	technical network	< 1 ms
	...	
– Feedback Controller:	network inbound (100 MBit/s)	~ 3 ms
	data processing (essentially matrix multiplication)	~ 15 ms
	network outbound	~ 3 ms
	technical network	< 1 ms
	...	
– Correction System:	network inbound	~ 3 ms
	WorldFIP (50 Hz) clock	~ 20 ms
	...	

- Some delays are concurrent: open loop is longer than the feedback delay
  - Data processing has to take less than the feedback delay (20/40 ms)
  - Smith Predictor can be used to compensate worsening effect of the constant open loop delay on the feedback response

Real-Time control requires an upper limit for total loop delay

(reminder: classic Smith Predictor compensates only constant delays)

Two main strategies:

- measurement of actual delay and its dynamic compensation in SP-branch:
  - high numerical complexity, since the branch transfer function has continuously to be modified
  - only feasible for small systems
- Jitter compensation using a periodic external signal:
  - CERN wide synchronisation of events on sub µs scale that triggers:
    - Acquisition of BPM system
    - Reading of receive buffers
    - Processing and sending of data
    - time to apply in the power converter front-ends
  - The total jitter, the sum of all worst case delays, must stay within “budget”.
  - Measured and anticipated delays and their jitter are well below 20 ms.
  - feedback loop frequency of 50 Hz feasible for LHC, if required...

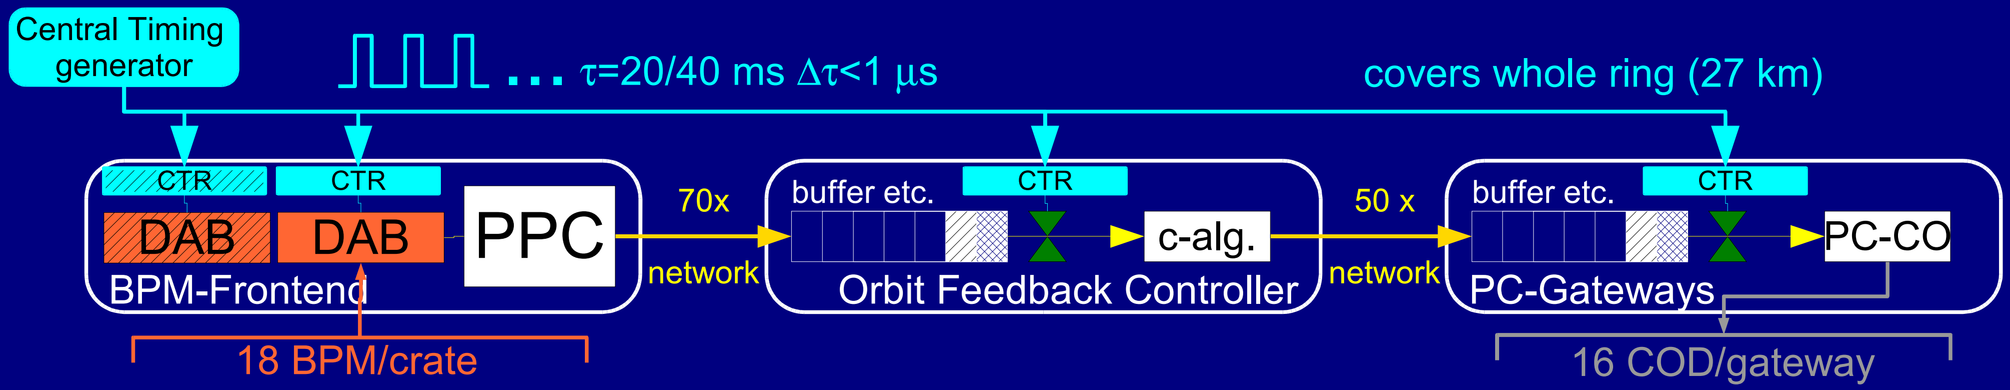

Orbit Feedback Controller – Test Bed

- Test bed complementary to the Orbit Feedback Controller (OFC):
  - Simulates the open loop and orbit response of COD->BEAM->BPM
  - Same data delivery mechanism and timing as in the real front-end
    - transparent for the OFC
    - same code for real and simulated machine:
      - possible and meaningful "offline" debugging for the fb controller
  - Controller strategies can be tested without beam:
    - Decay/Snapback, Squeeze
    - Ground motion simulations
    - other environmental influences
  - Tests of controller implementations (scheduling of CPU, network, timing....)

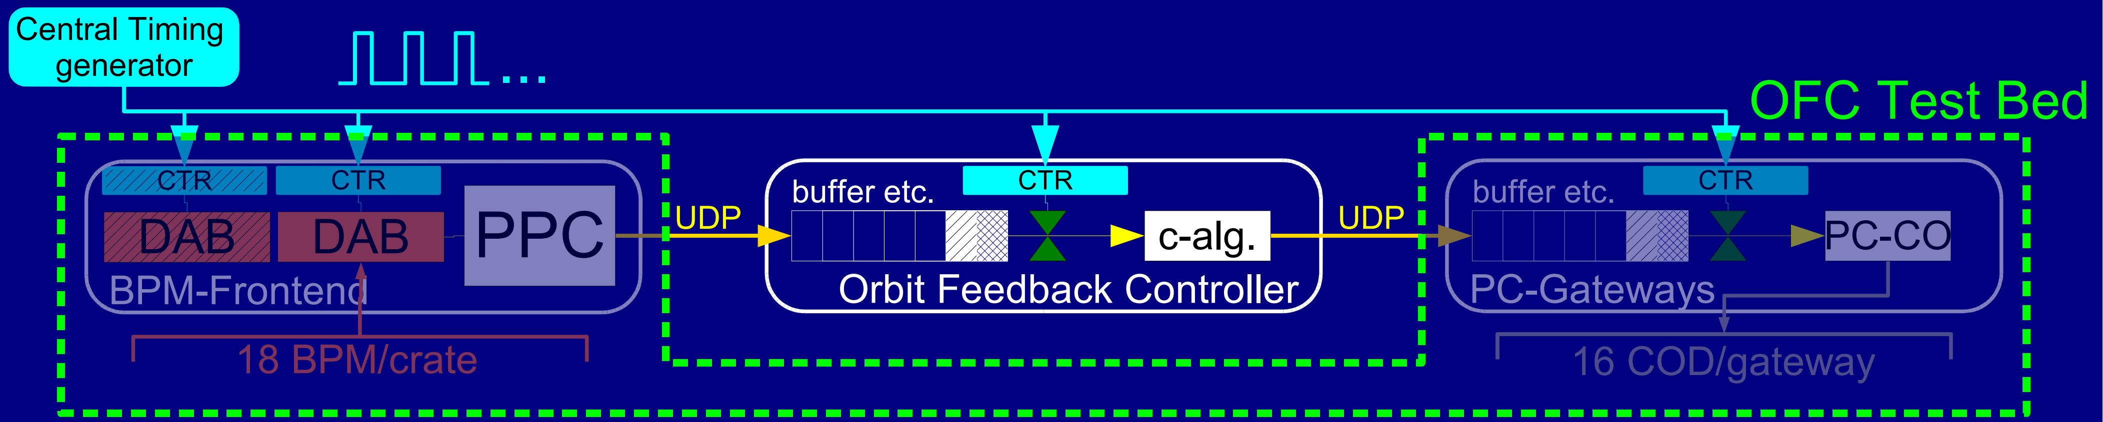

OFC Test Bed -Performance

- The test bed has to run faster and than orbit feedback controller and keep the same real-time constraints

- runs up to 128 Hz in hard real-time for a full LHC orbit simulation
  (including response of quadrupoles and main dipole magnets)
  - 40 times faster than correction dipole magnet bandwidth
  - precise simulation of the plants and closed orbit movements
  - detect and measure OFC limitations
    - actual feedback bandwidth
    - delay response
    - max. orbit stability at given parameters:
      - feedback frequency
      - white noise limit due to BPM
      - ...
- Simulation of failure scenarios:
  - packets loss
  - BPM/COD failures
  - [..]

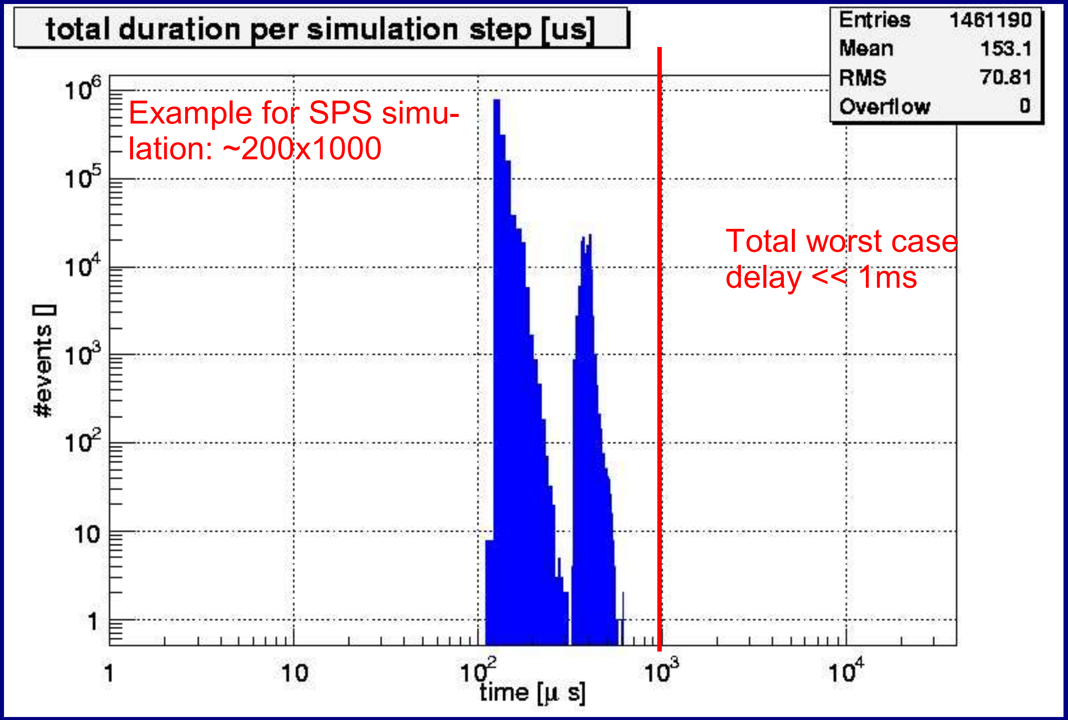

Test Bed total execution delay response including user interaction and simulation of COD, quad and main dipole response

Conclusion

- LHC requires a reliable feedback system during all operational phases.
  - Two Beams
  - Total: 1056 BPM & 1060 CODs
  - ~120 PowerPC front-ends interconnected through switched Ethernet
- Constant delay budget through site wide machine timing system
  - jitter of Ethernet does not pose a problem for given configuration
  - Identified bottlenecks:
    - network interfaces in the front-ends
    - numerical complexity in the x86 based orbit feedback controller
    - Are addressed and limited through:
      - scheduling of the data streams on the network interfaces (QoS)
      - optimised use of multi-processors and network connections
      - real-time task scheduling of operating system
- Development of a 'Test Bed' environment that helps testing and developing of
  - feedback controller implementations (scheduling, response, performance)
  - control strategies
- A machine wide feedback loop is feasible of up to 1 kHz which gives enough operational safety margin for the targeted 25/50 Hz of the orbit feedback loop

- For the SPS prototype studies:
  - 6 dedicated position monitors (*BPMBs*) with full LHC acquisitionin LSS5
  - Power converters of CODs have been enabled to receive real-time reference current changes.
  - The pre-processed BPMB data is sent from the surface building BA5 over Ethernet connection to the Prevessin Control Room (PCR) to a PC that houses the controller performing the correction and sends the steering data back to the COD power converter controller.

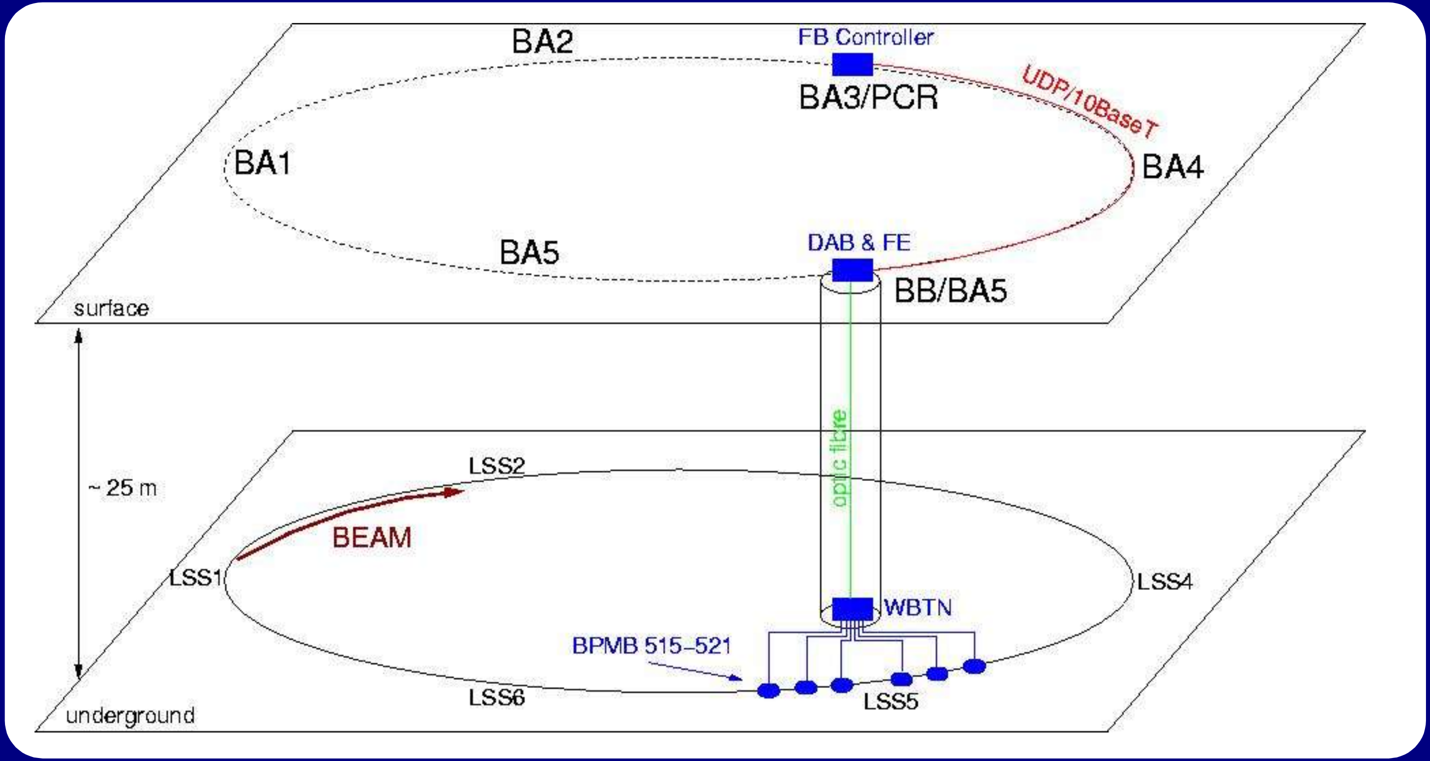

LHC Orbit Feedback Control
Results I/II 2003

Steering example with additional external noise:

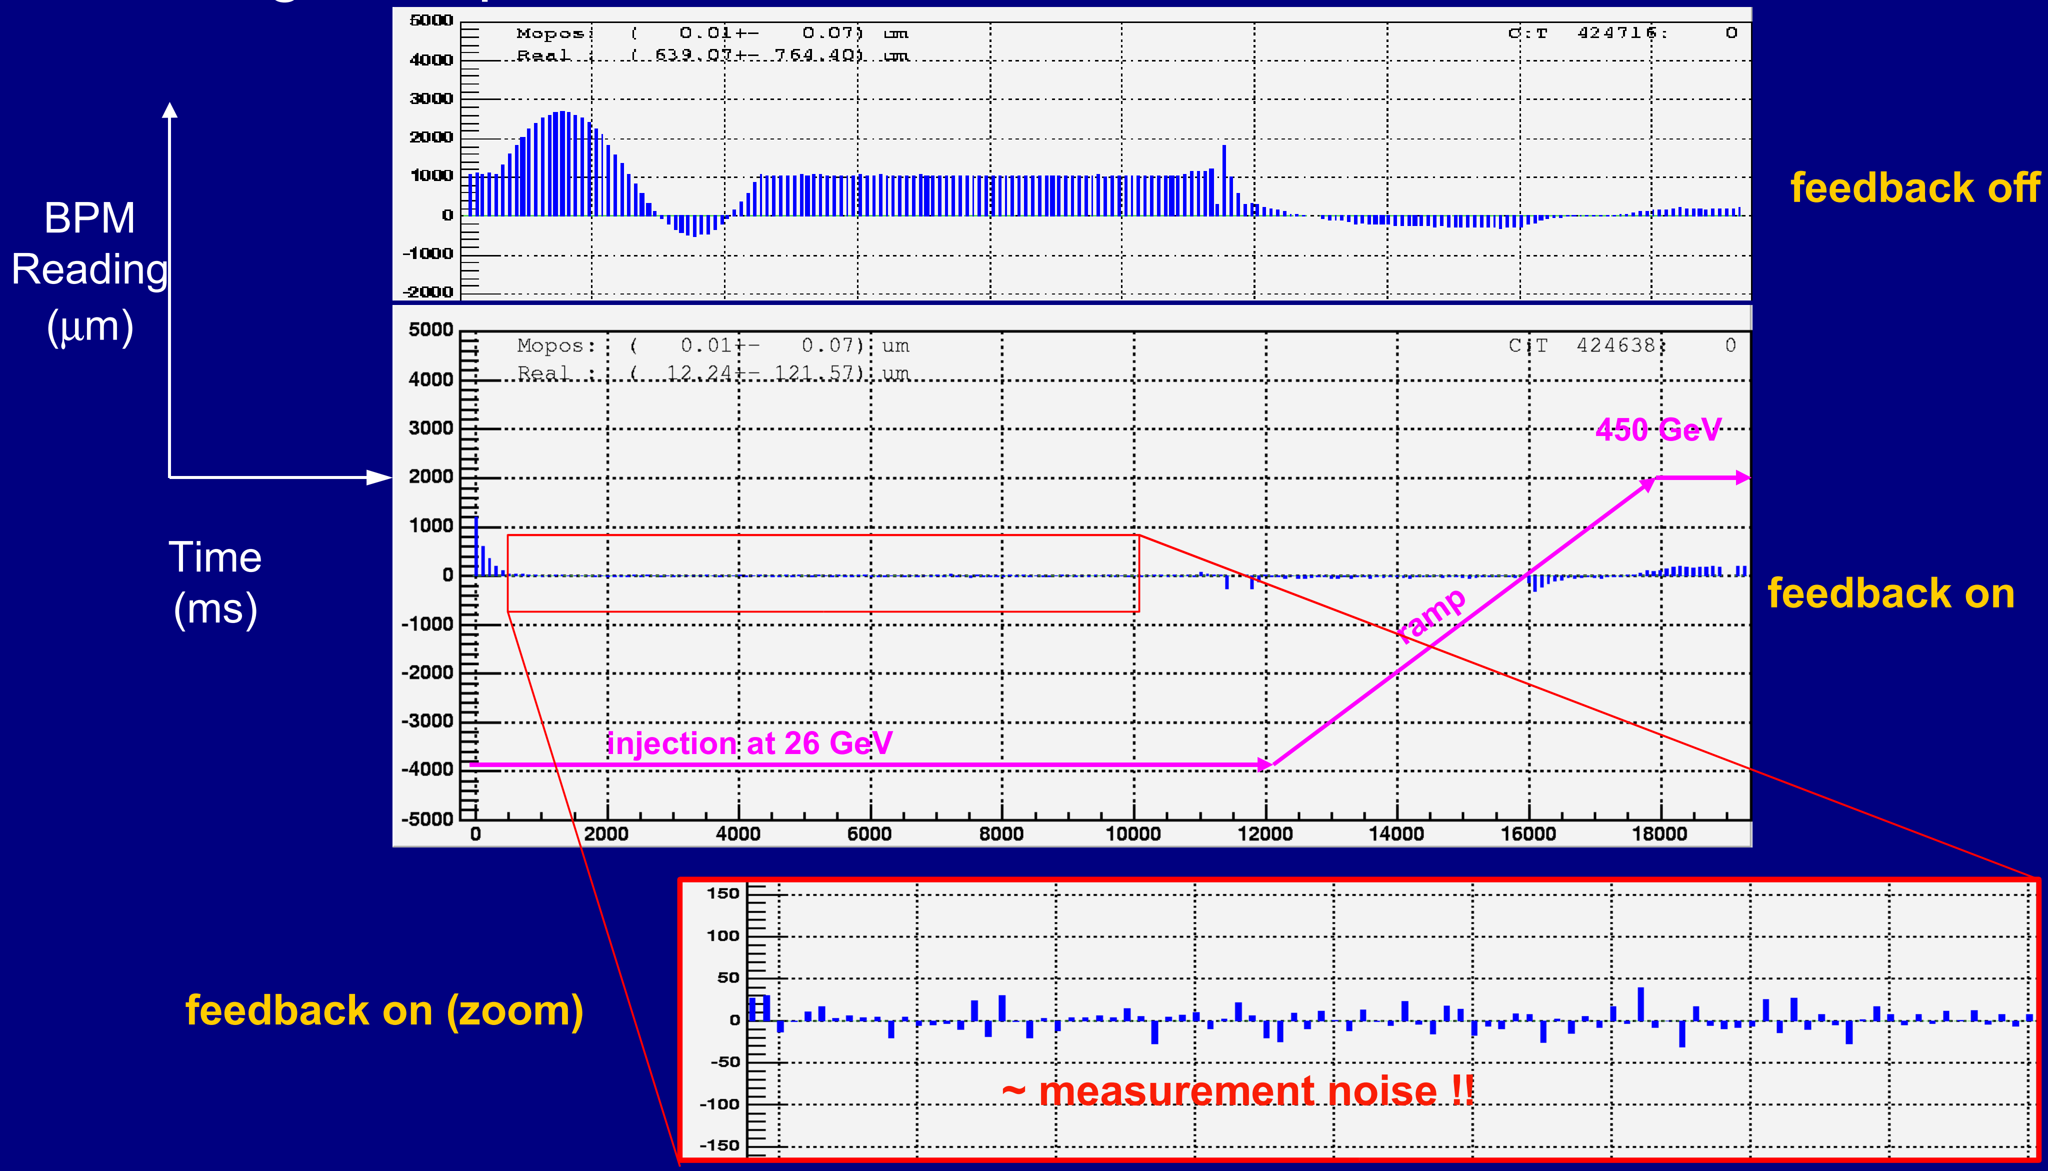

LHC Orbit Feedback Control
Results II/II 2003

- Feedback loop showed an average good performance
  - SPS system in 2003:
    - Stabilised the beam at 4 BPMB
    - Max. feedback sampling frequency 100 Hz
    - position within 8.5 μm

Feedback attenuation

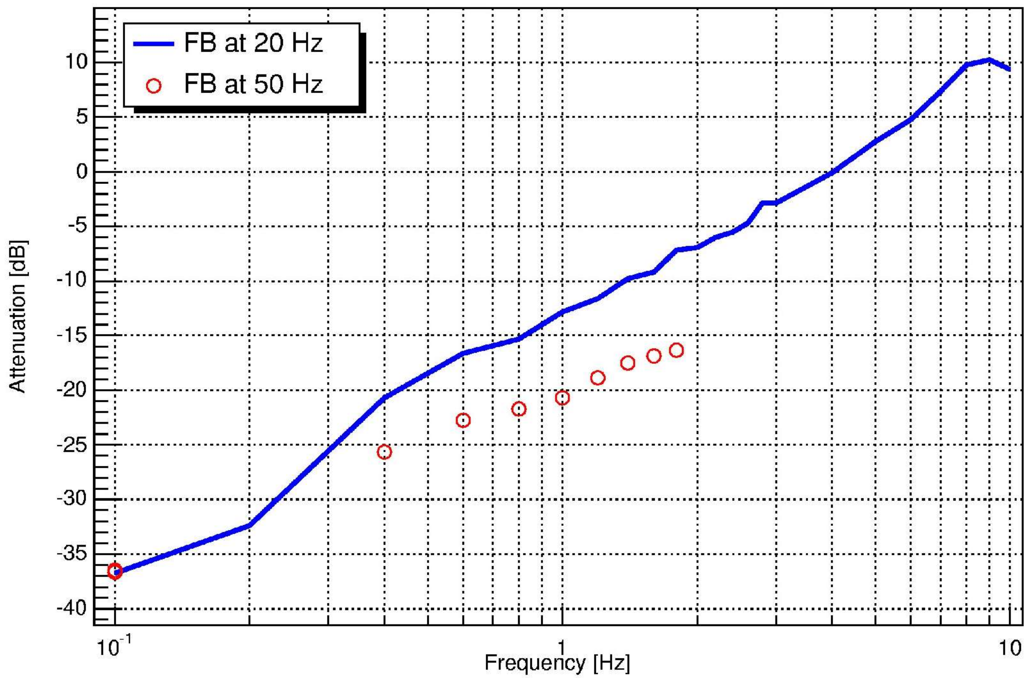

Position distribution @ 100 Hz
σ = 8.5 μm

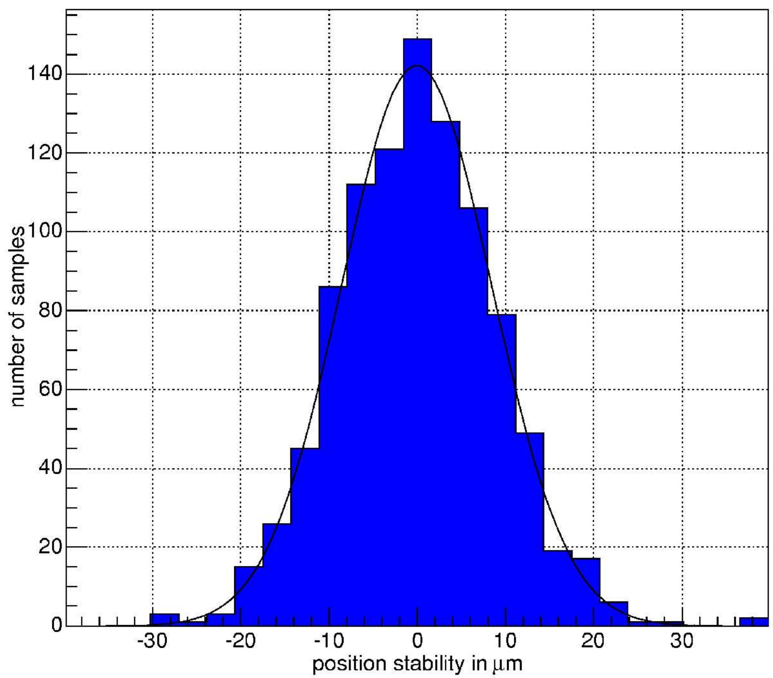

(2004: position distribution @ 25 Hz -> $\sigma < 2$ μm)!

LHC Orbit Feedback Control
Results 2004 – coasting beam at 270 GeV

- Stabilised vertical orbit in the SPS
  - 270 GeV coasting (proton) beam
  - 72 bunches, ~1e11 protons each
  - $\beta_v \sim 100$ m

  –> $\sigma < 2\ \mu$m
- magnitudes better than required
- Target: maintain same longterm stability

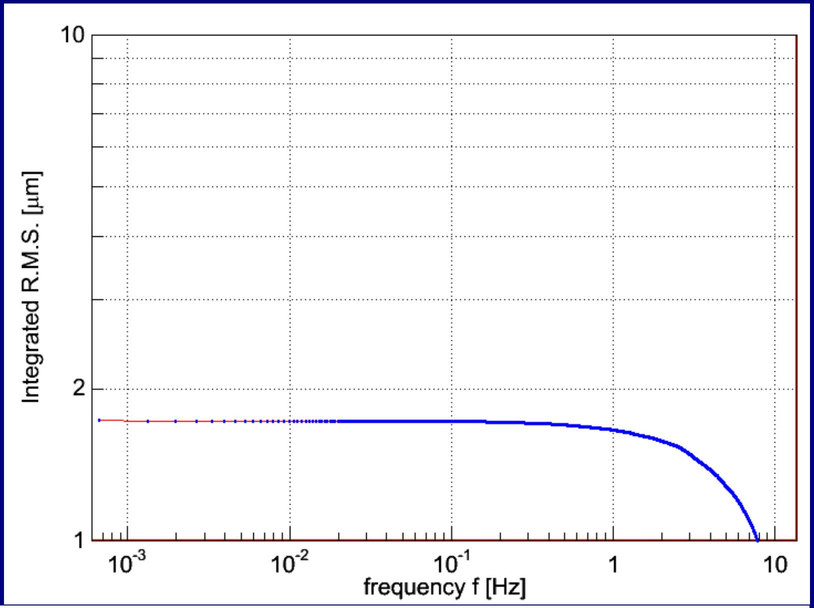

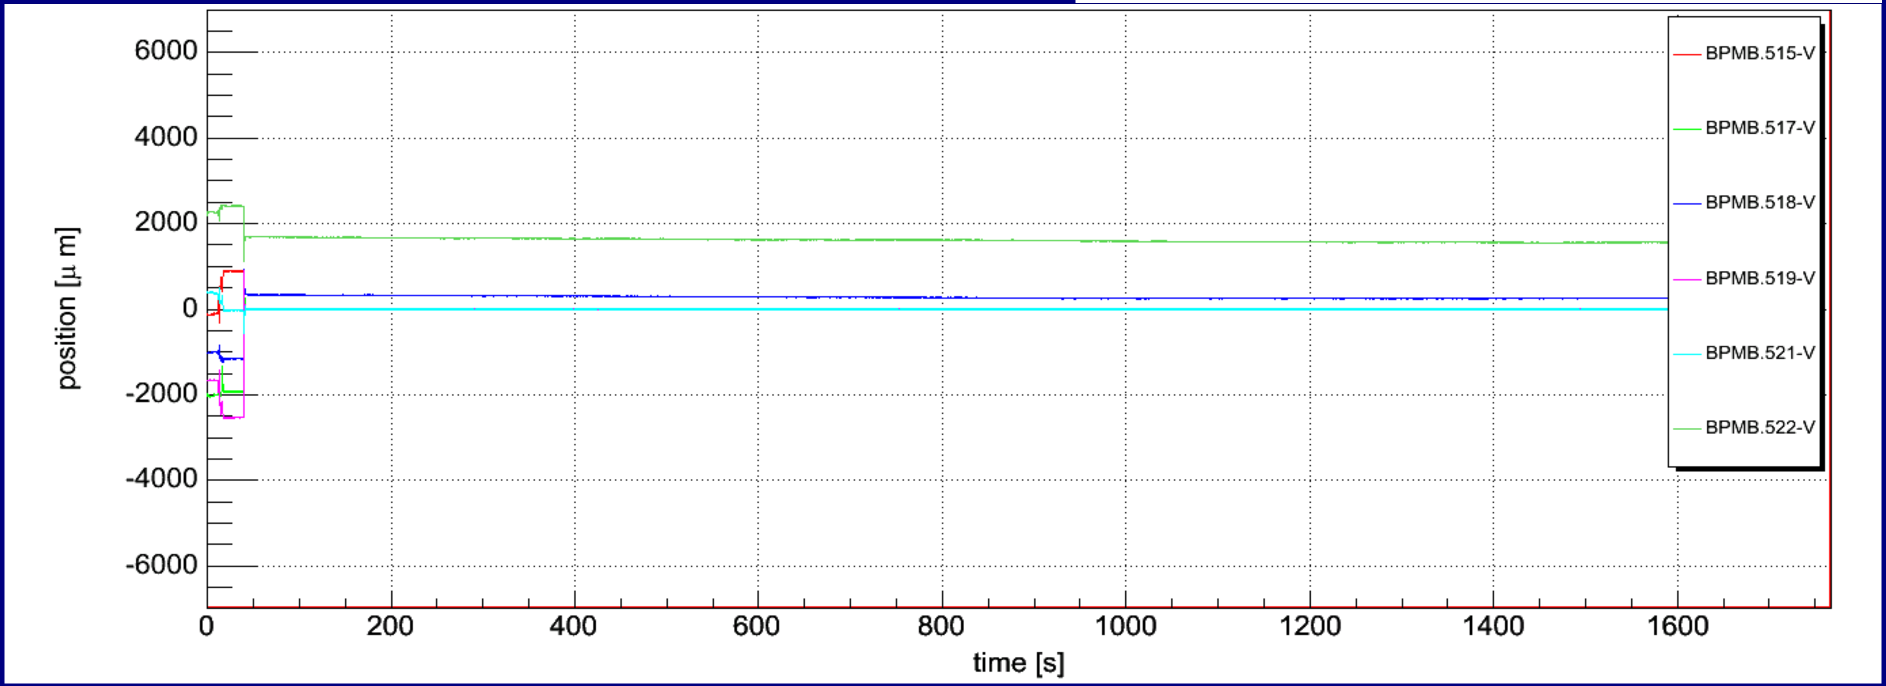

Beam Position Monitor System

- LHC is equipped with > 1200 BPMs
- each BPM:
  - measures horizontal & vertical plane
    - ~ 8 samples per betatron oscillation (redundancy!)
  - measures on a per bunch basis:
    - Wide Band Time Normaliser (WBTN) for h/v button signals
      - position -> analog Δt signal, two laser pulses (analog)
    - Mezzanine card (Integrator and ADC): Δt -> ~ *position*
    - Digital Acquisition Board (DAB): preprocessing
      - conversion of ADC values and compensation for intrinsic the BPM's non-linearities
      - filling the data into 100k turn acquisition, real-time orbit data and "*post mortem*" buffer
    - One PowerPC readout and controls up to 18 BPMs
    - Single shot (1 bunch, 1 turn) precision:
      - pilot bunch (~10^9 protons): <200 μm/shot
      - Nominal bunch (~$1.3\cdot10^9$ protons): ~100 μm/shot
- closed Orbit: 255 turn average of all bunches (in order to get rid of potential 50 Hz aliasing signals)
  - Closed orbit measurement precision:
    - pilot: ~ 13 μm
    - nominal beam (~2800 bunches): ~ 0.1 μm
- BPM data send via Ethernet to fb Controller
  - ~ 22kByte/sample , max 50 samples/s

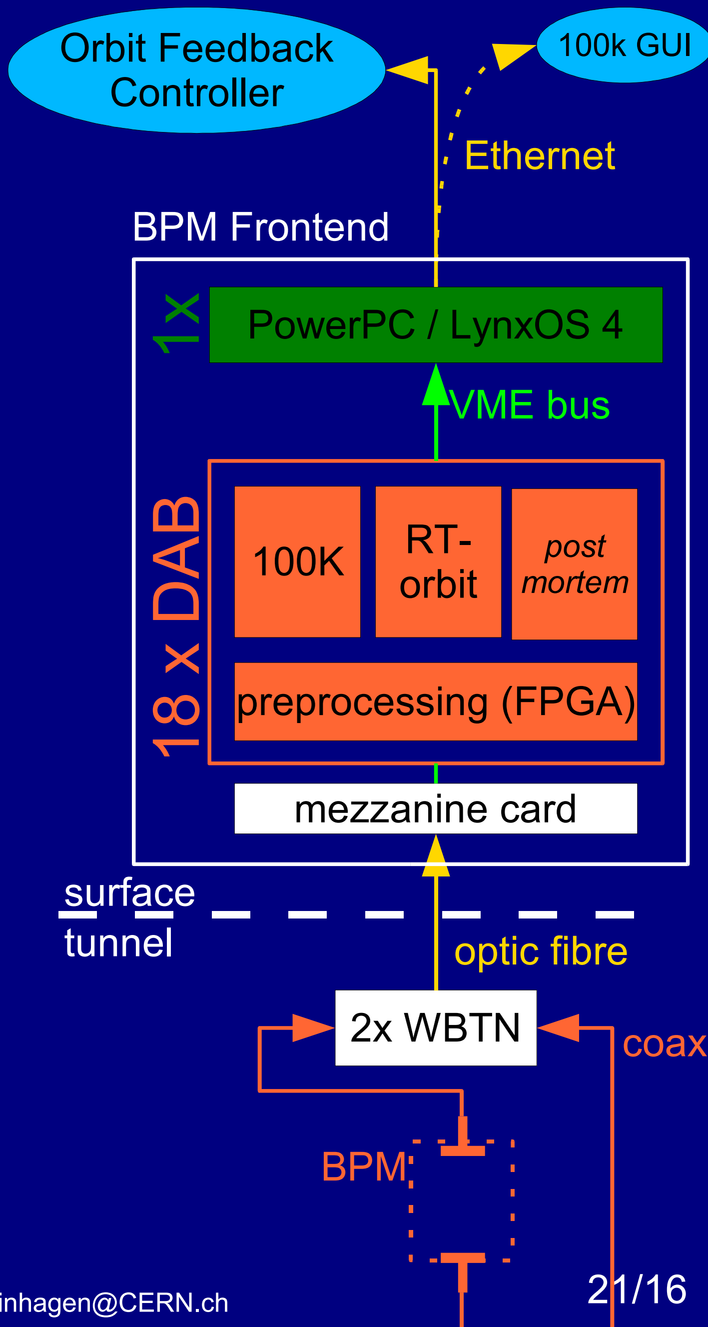

Possible Architectures

Local

- reduced # of network connections.
- numerical processing simpler.
- …
- less flexibility.
- not ideal for global corrections.
- coupling between loops is an issue.
- problem with boundary areas to ensure closure.
- ..

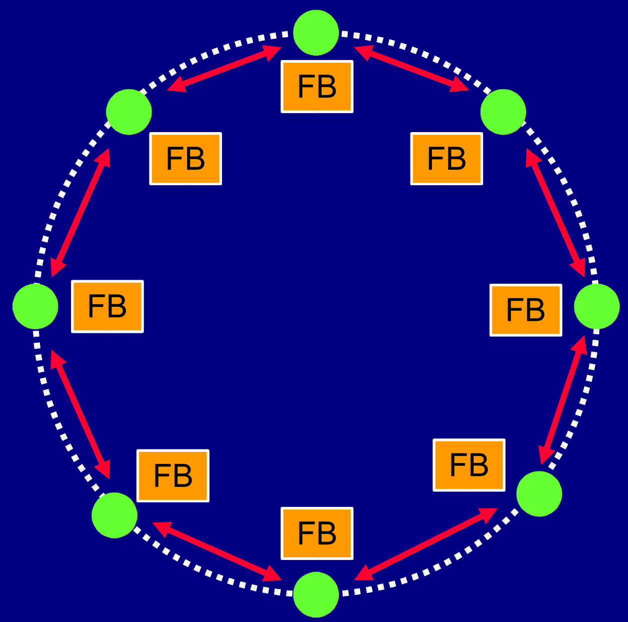

Central

- entire information available.
- all options possible.
- can be easily configured and adapted.
- …
- network more critical – DELAYS !
- large amount of network connections.
- …
-

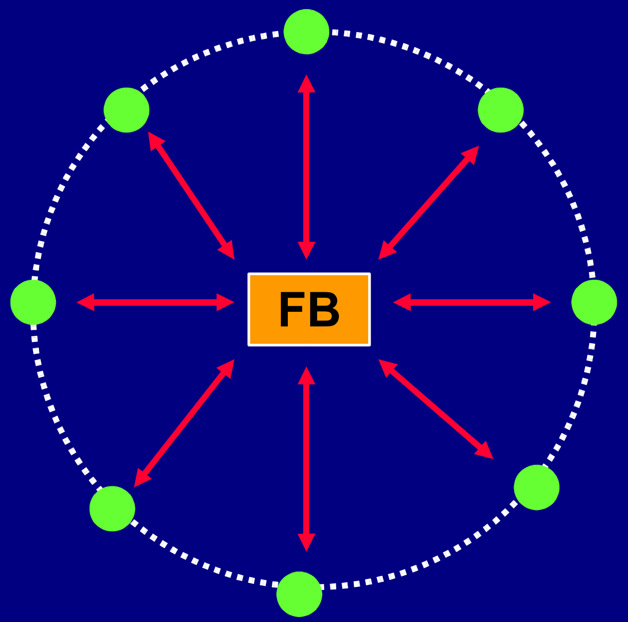

Preferred!

Beam Position Monitors and Data Rates

Hardware:

- both rings covered by 1056 BPMs
- Measure both planes (2112 readings)
- Organised in front-end crates (PowerPC/VME) in surface buildings
  - 18 BPMs (hor & vert) ⇔ 36 positions / VME crate
  - 68 crates in total, 6-8 crates /IR

Data streams:

- Average data rates per IR:
  - 18 BPMs x 20 bytes ~ 400 bytes / sample / crate
  - 1056 BPMs x 20 byte ~ 21 kbytes / sample

  @ 50 Hz: ~ 8.5 Mbit/s

  \+ protocol (UDP/IP, identifier, check bits...) ~ 13 Mbit/s
- Peak data rates (bursts): 100Mbit/s resp. 1Gbit/s (depending on Ethernet interface)

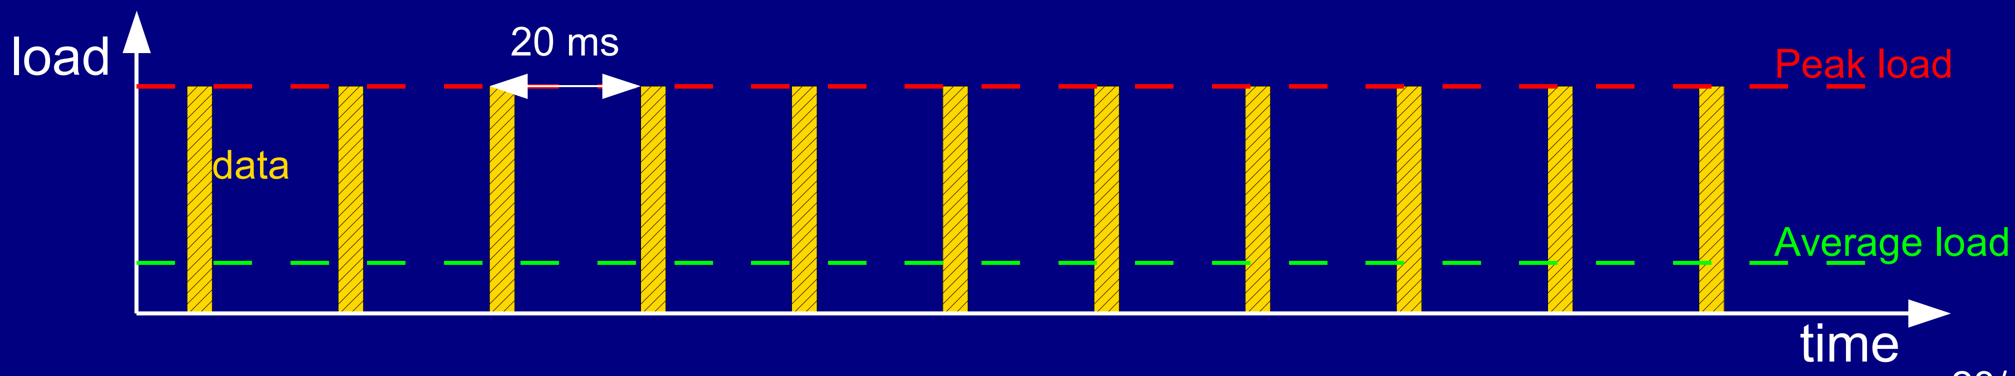

What a Smith-Predictor compensates:

- induces an inhibitor signal to to delay the actuator signal by λ