2014-05-02 ICE Automation Seminars:
Workshop on Advanced Control Engineering

Design and Performance of the LHC beam-based Feedback Control Systems

– There and back again –

Ralph J. Steinhagen
Accelerator & Beams Department, CERN

27 km Circumference – 1232 LHC dipole magnet

- B field 8.3 T (**11.8 kA**) @ 1.9 K (super-fluid Helium)
- two-in-one magnet design:
  two beam tubes with an opening of 56 mm (210 mm separation)

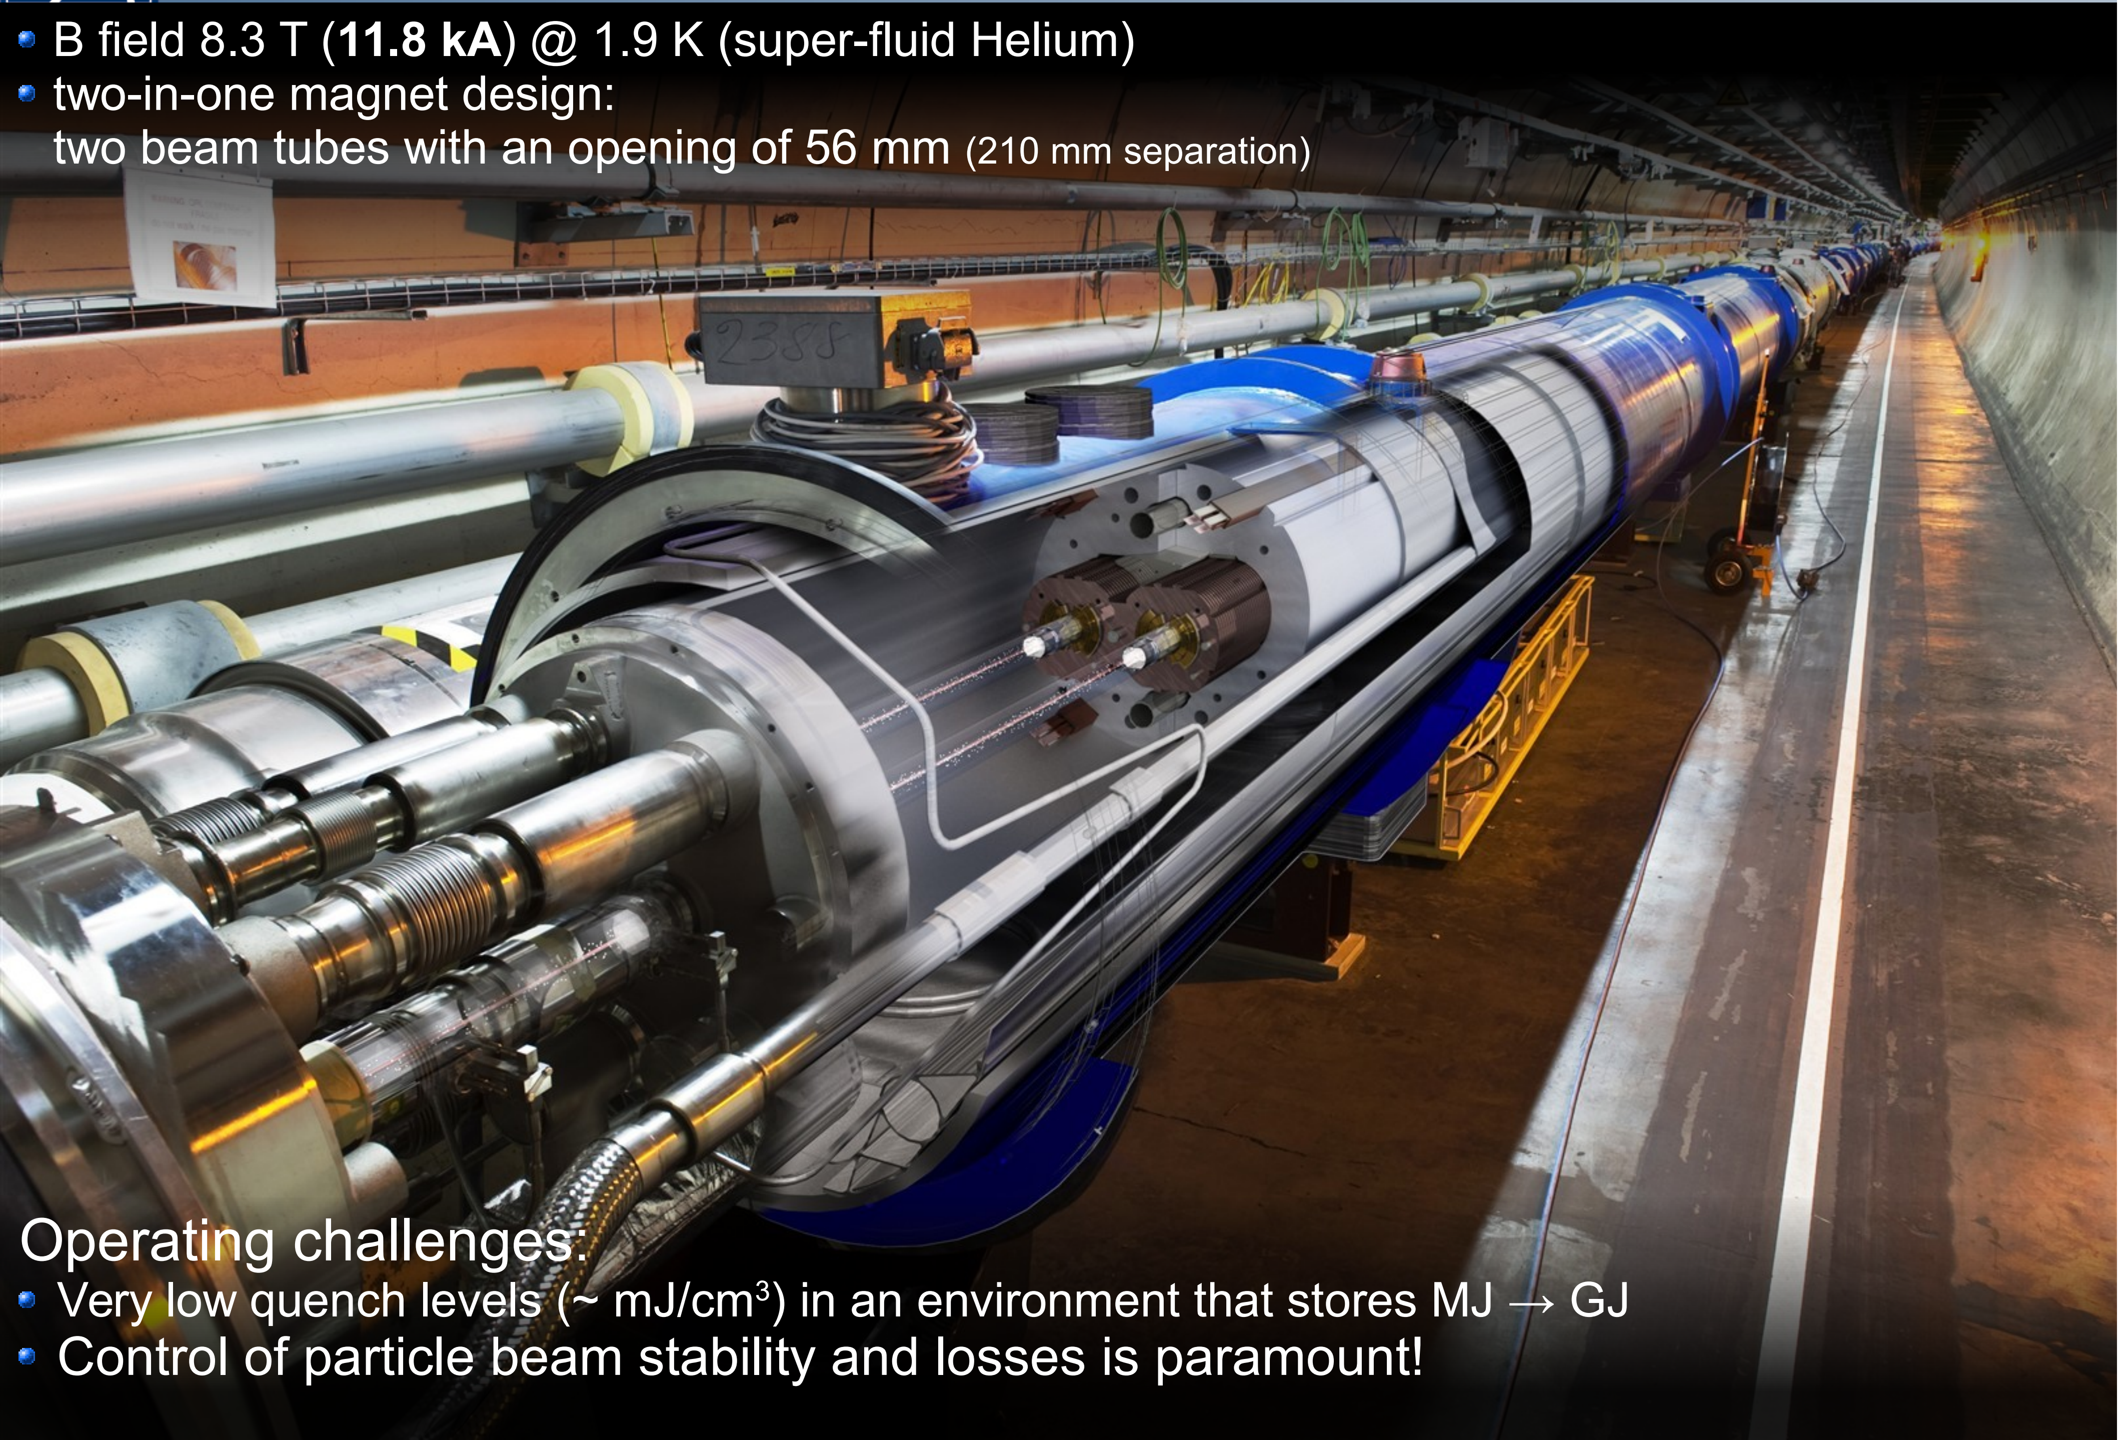

Operating challenges:

- Very low quench levels (~ mJ/cm^3) in an environment that stores MJ → GJ
- Control of particle beam stability and losses is paramount!

Beam Parameter Stability in Hadron Machines

- Traditional requirements on beam stability...

... to keep the beam in the pipe!

- LHC's increased stored intensity and energy much tighter requirements on beam stability:
  1. Capability to control particle losses
     - Machine protection (MP) & Collimation
     - Quench prevention
  2. Commissioning and operational efficiency

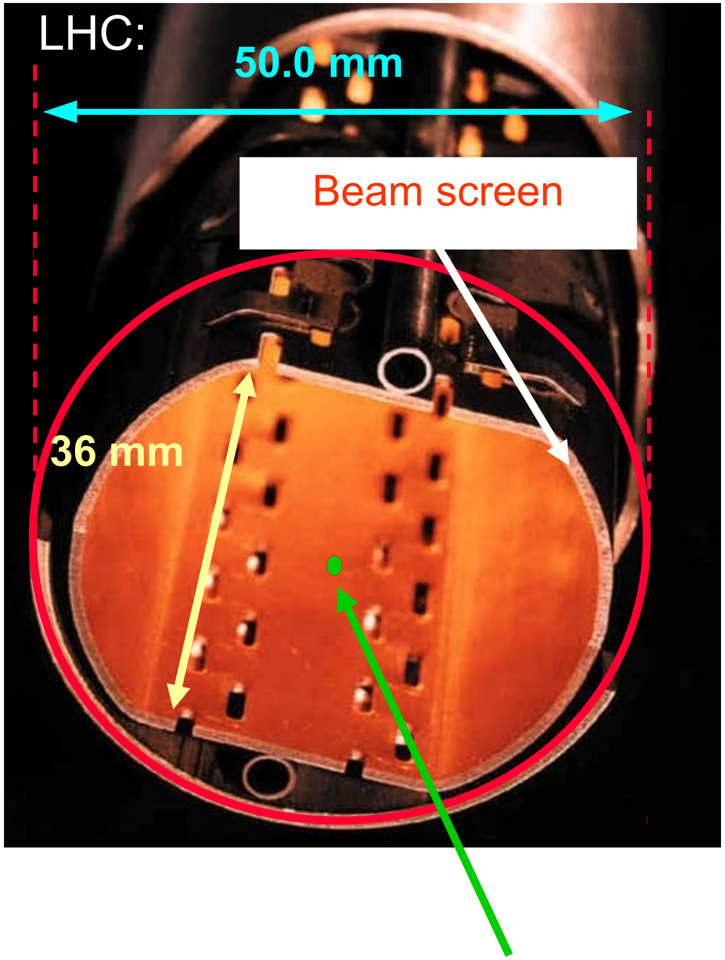

Beam 3 σ envel.
~ 1.8 mm @ 7 TeV

- FBs became a requirement for safe and reliable nominal LHC operation
  – implications on controller reliability, availability and system integration

Requirements on Orbit – machine protection

Combined failure[1]: Local orbit bump and collimation efficiency (/kicker failure) → bumps may potentially compromise collimation function

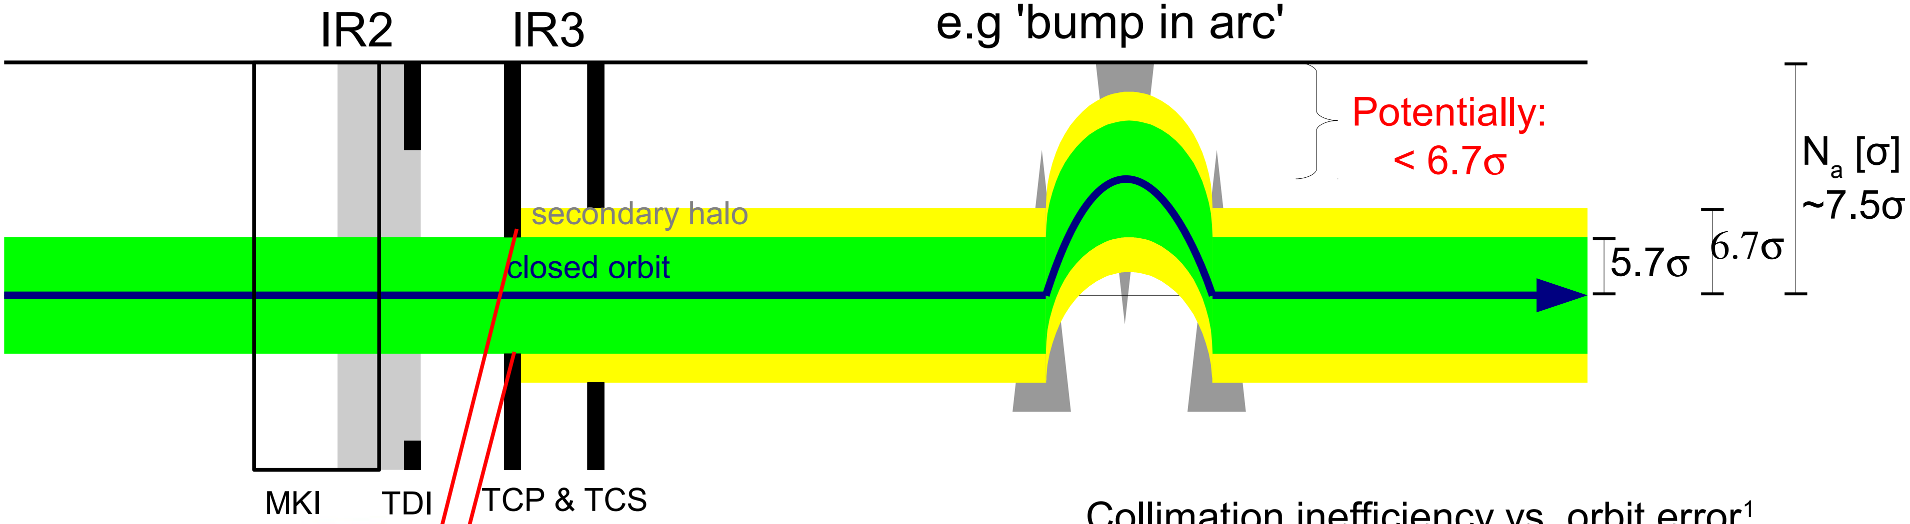

Tight settings (2012):
~2.2 mm gap at primary collimator

Collimation inefficiency vs. orbit error[1]

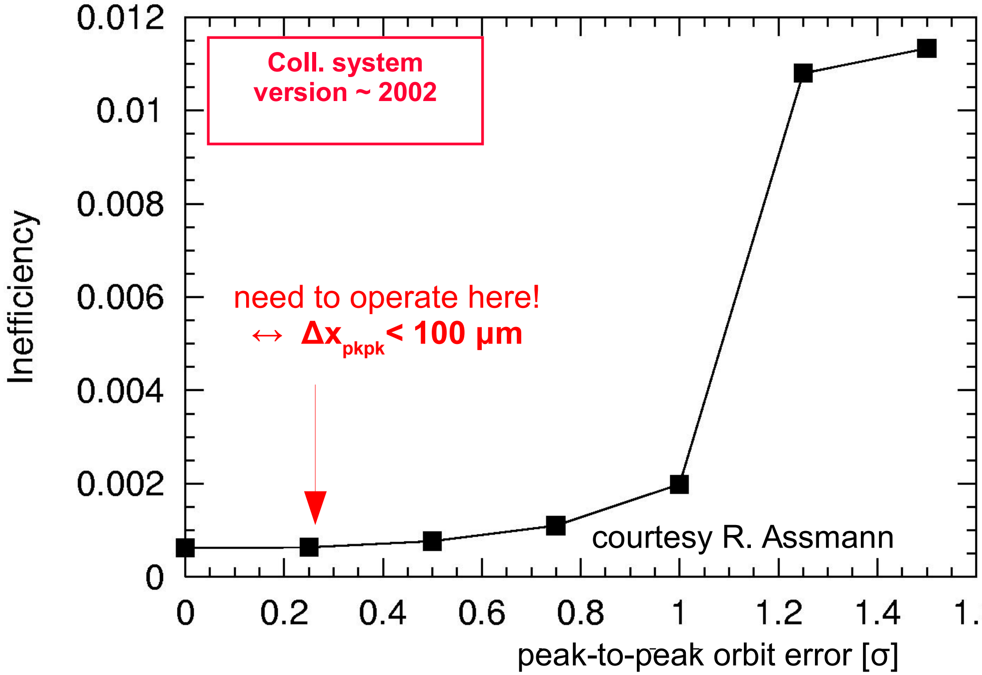

LHC Feedback Operation – Example

- Orbit feedback used routinely and mandatory for nominal beam

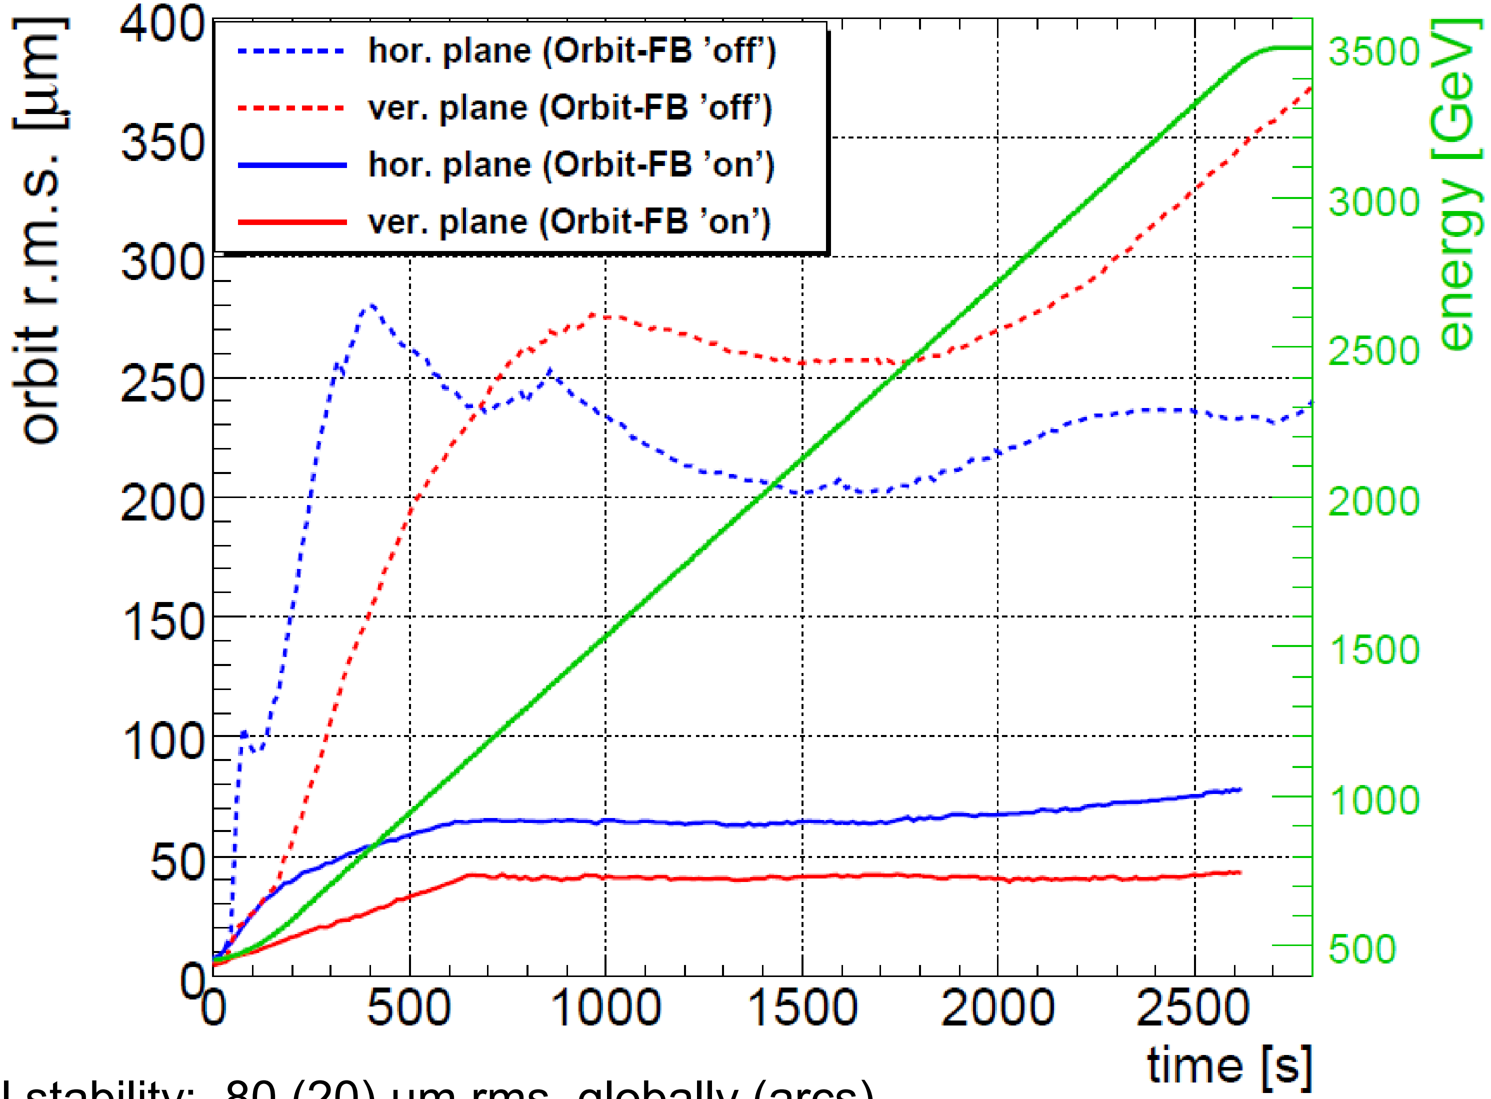

- Typical stability: 80 (20) µm rms. globally (arcs)
- Most perturbations due to Orbit-FB reference changes around experiments

2009 LHC Commissioning

- … somewhat a surprise: 3rd ramp without Tune-Feedback

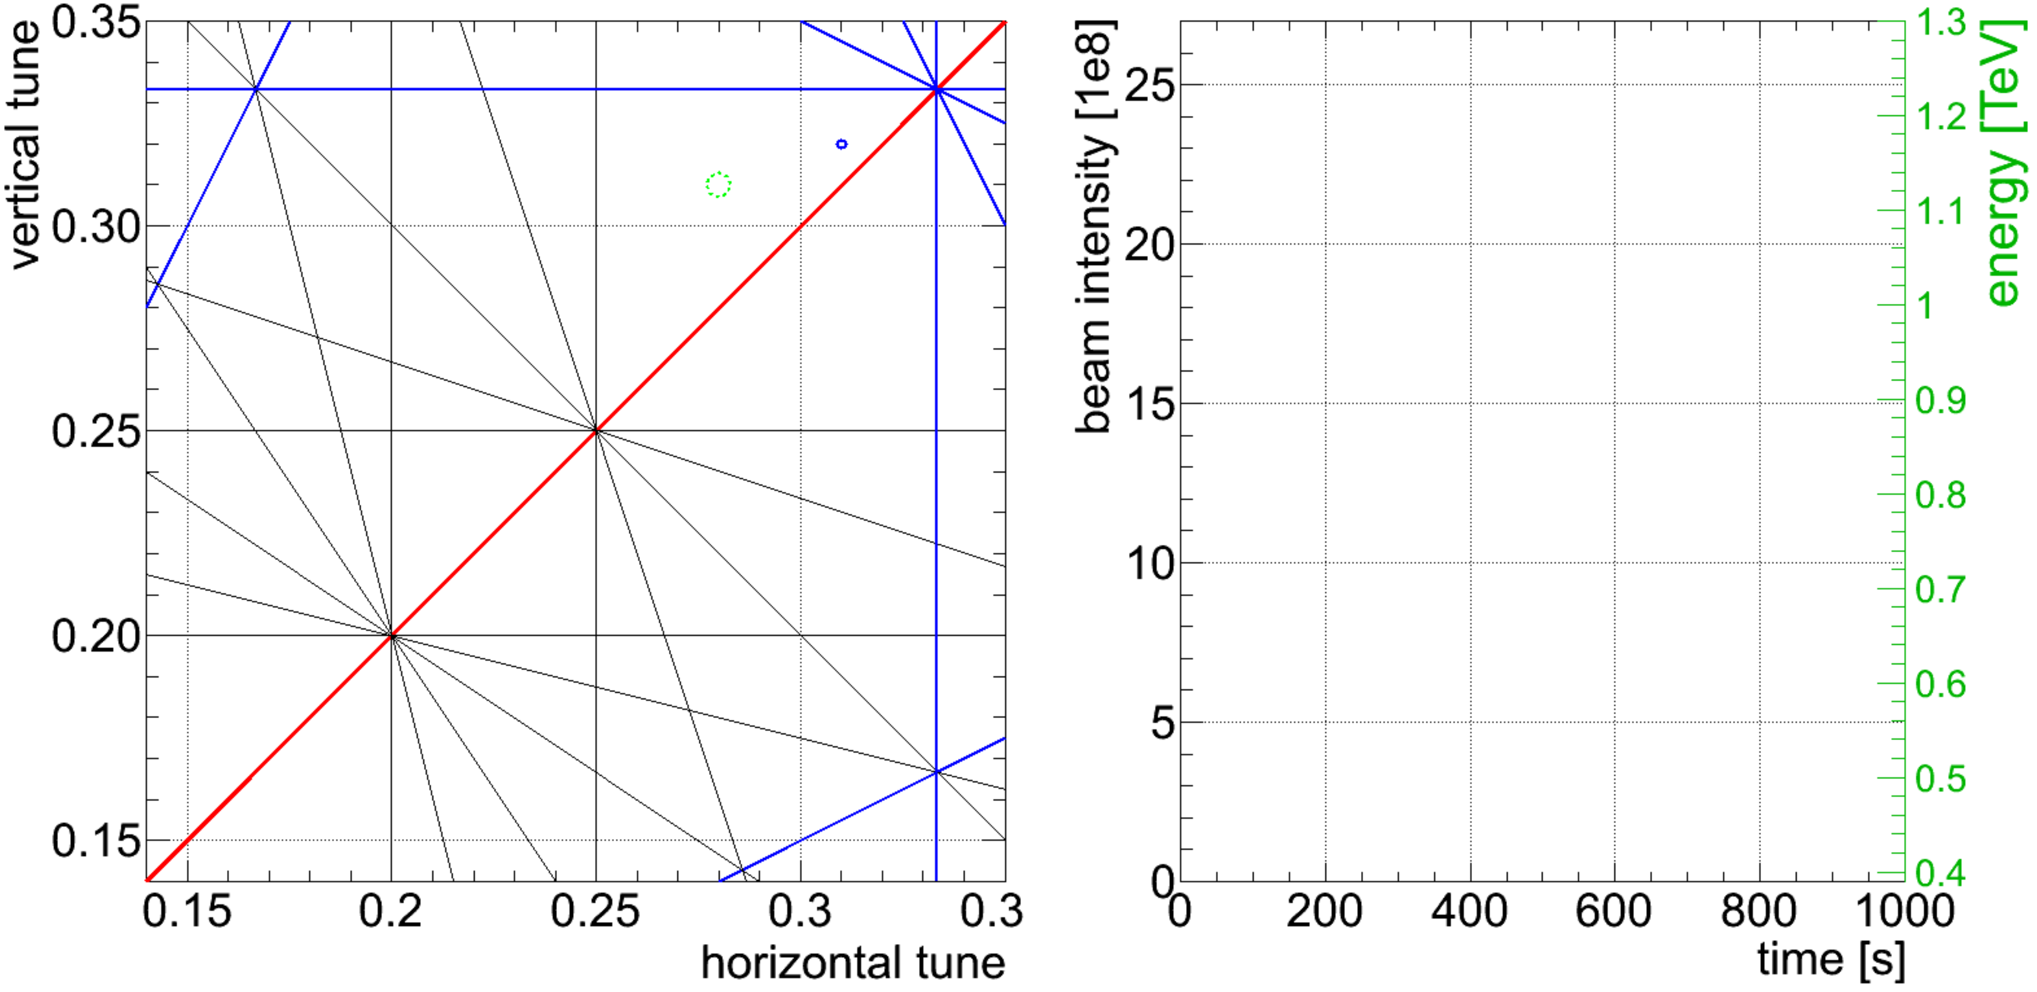

Control Paradigms I/III
Parameter control, either through...

- Feed-Forward: (FF)
  - Steer parameter using precise process model and disturbance prediction

- Feedback: (FB)
  - Steering using rough process model and measurement of parameter
  - Two types: within-cycle (repetition Δt<<10 hours) or cycle-to-cycle (Δt>10 hours)

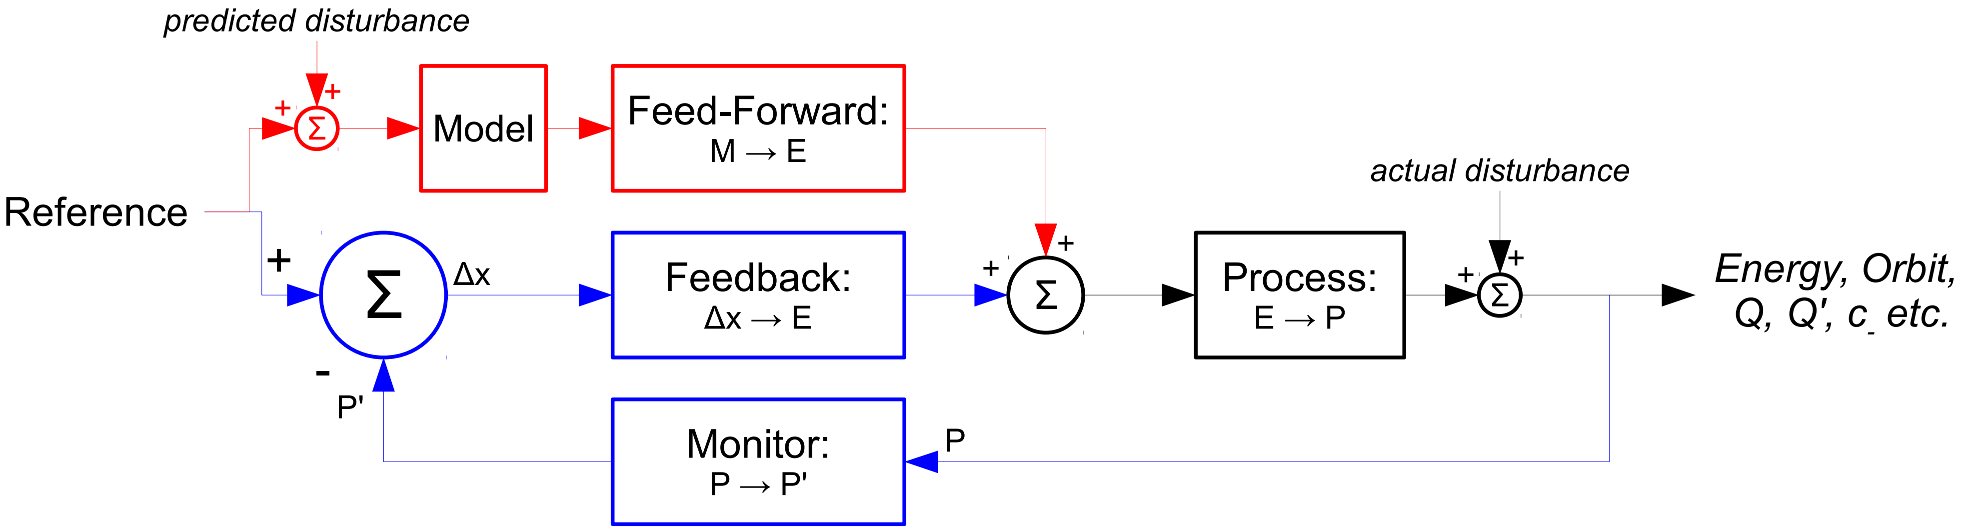

Orbit-Feedback as Prototype for all LHC Beam-Based Feedback Systems

- Orbit-Feedback is one of the largest and most complex LHC feedback:
  - 1088 BPMs → 2176+ readings @ 25 Hz from 68 front-end computers
  - 530 correction dipole magnets/plane, distributed over ~50 front-end computers
  - Total >3500 devices involved

- Specific requirements fairly distributed → opted for central global feedback system

- One central controller (OFC + hot spare):
  - higher numerical load
  - higher network load (↔ ~120 front-ends)
  - dependence of machine operation on single device
  - easier synchronisation between front-ends and FBs
  - flexible correction scheme changes and gain-scheduling
  - most efficient to handle cross-talk and (de-)coupling between FBs

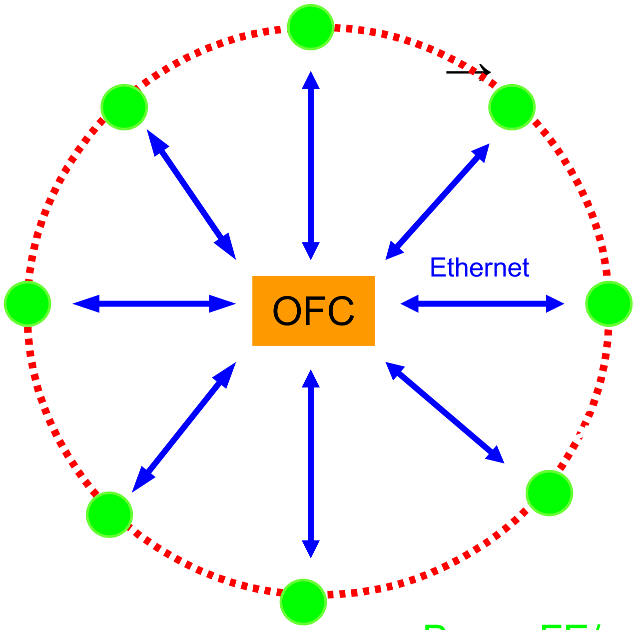

Beam FE/ corr. circuits

Common Feedback/Feed-forward Control Layout

- Feedback Controller (OFC) performing actual feedback controller logic
  - Simple streaming task (10% of total load)
  - Beam data quality checks and real-time filtering (80% of total load)
  - Server running Real-Time Linux OS with <u>periodic constant load</u>
    - multi-core, highly redundant – MTBF > 22 yrs (spec, 120 yrs meas.)
  - Technical Network as robust communication backbone

- Service Unit: Interface to high-level software control and interlock systems
  - Proxies user requests, handles asynchronous non-RT tasks

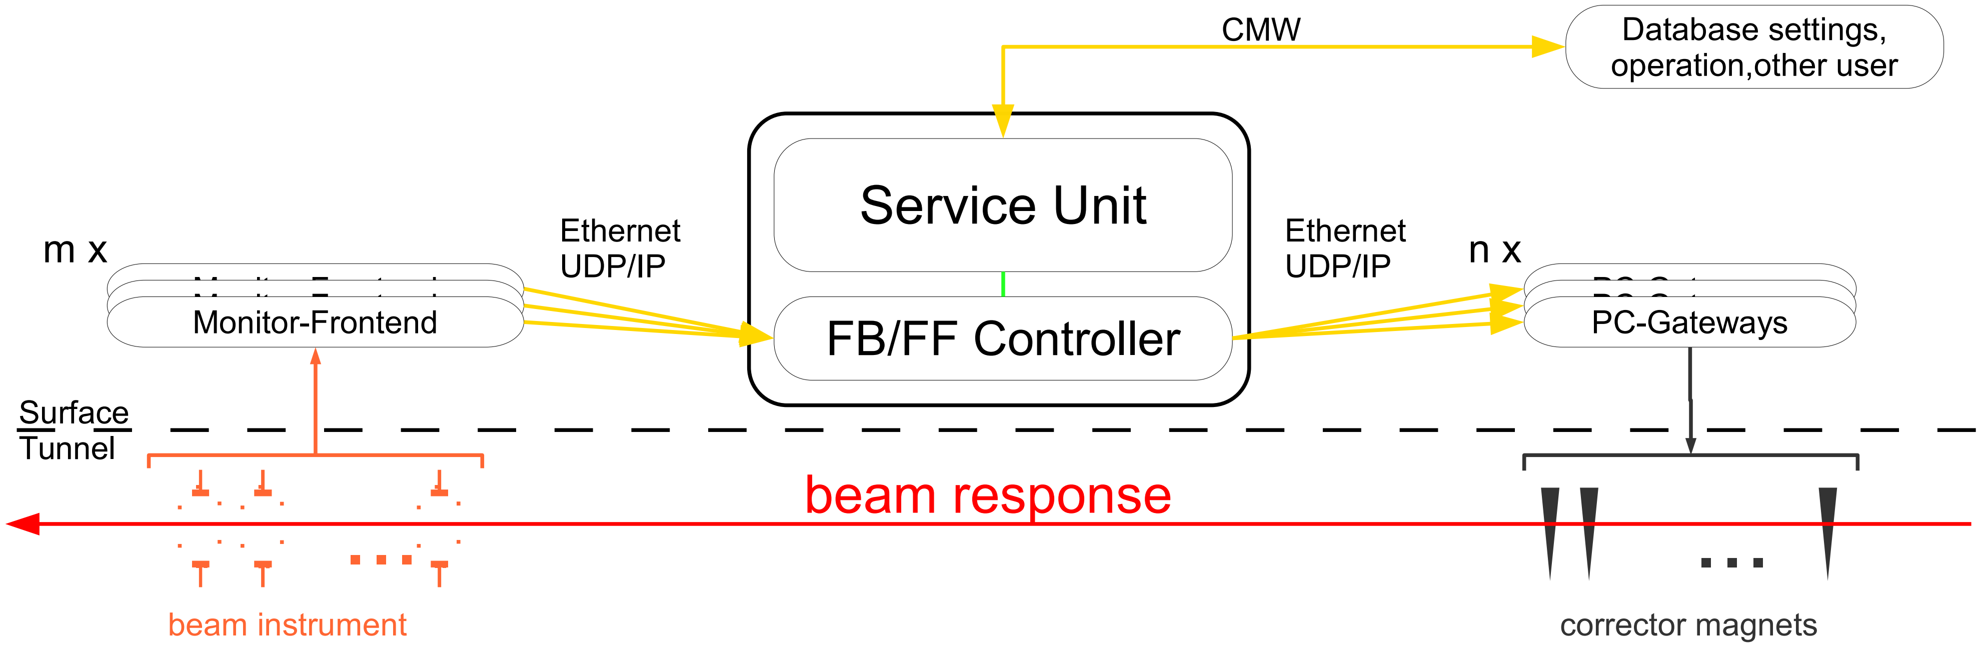

Multiple-Input-Multiple-Output (MIMO) Process Control

- 'Divide and Conquer' feedback controller design approach:
  1. Compute steady-state corrector settings $\vec{\delta}_{ss}=(\delta_{1,}\ldots,\delta_{n})$ based on measured parameter shift $\Delta x=(x_1,\ldots, x_n)$ that will move the beam to its reference position for $t\to\infty$.
  2. Compute a $\vec{\delta}(t)$ that will enhance the transition $\vec{\delta}(t=0)\to\vec{\delta}_{ss}$
  3. Feed-forward: anticipate and add deflections $\vec{\delta}_{ff}$ to compensate changes of well known and properly described sources

space domain

time domain

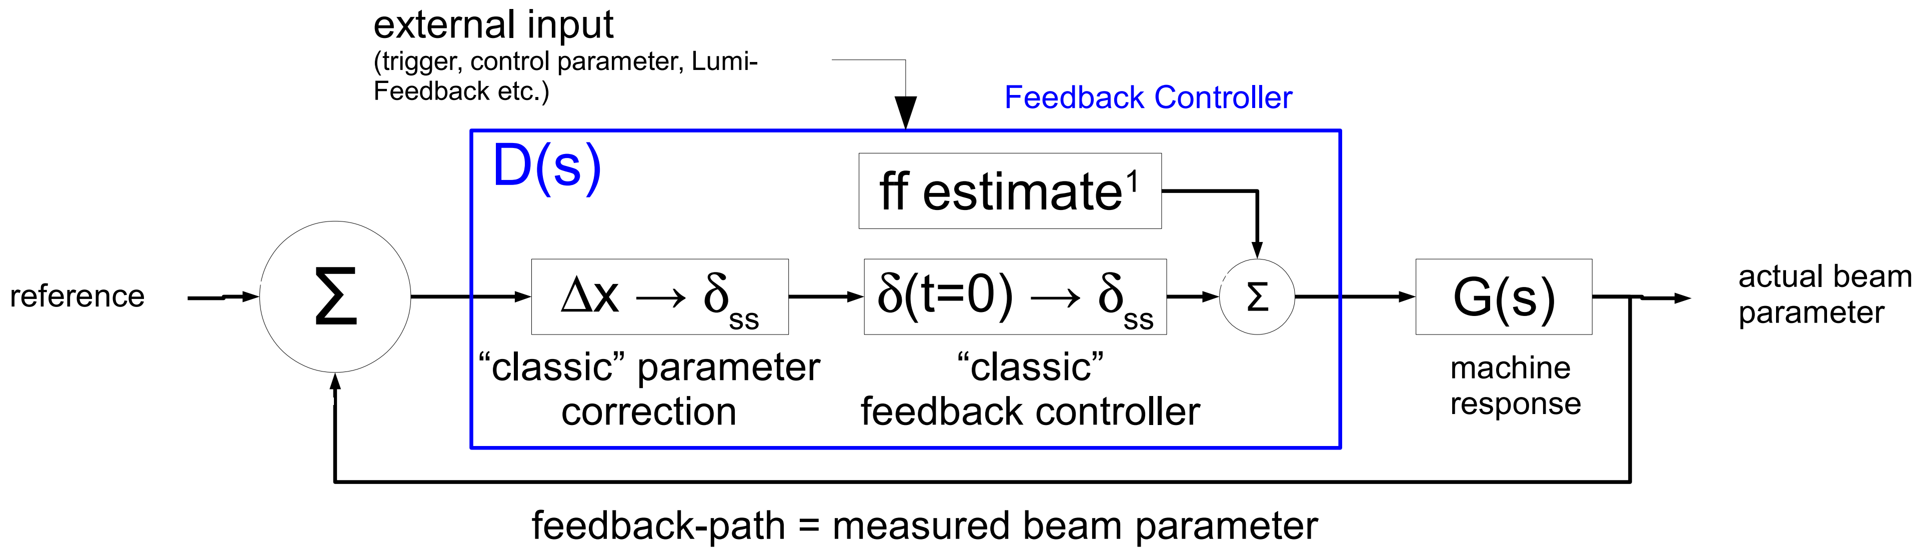

external input

(trigger, control parameter, Lumi-Feedback etc.)

Feedback Controller

D(s)

ff estimate^1

reference

Σ

$\Delta x \to \delta_{ss}$

"classic" parameter correction

$\delta(t=0) \to \delta_{ss}$

"classic" feedback controller

Σ

G(s)

machine response

actual beam parameter

feedback-path = measured beam parameter

- (N.B. here G(s) contains the process and monitor response function)

Space-Domain: Singular Value Decomposition (SVD) on a slide

Linear algebra theorem*:

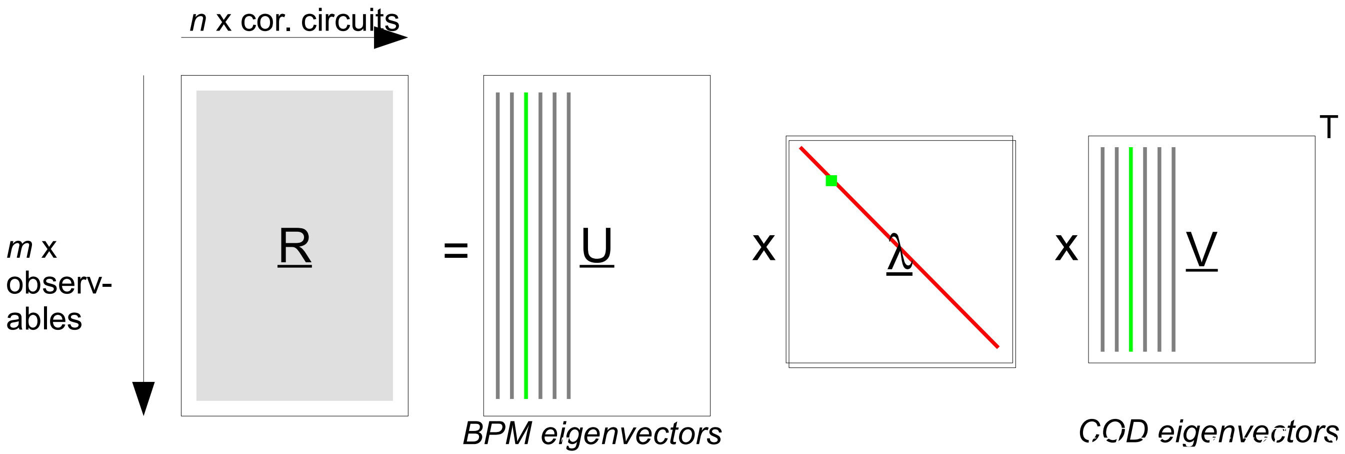

eigen-vector relation:

$$\lambda_i \vec{u}_i = \underline{R}\cdot\vec{v}_i$$

$$\lambda_i \vec{v}_i = \underline{R}^T\cdot\vec{u}_i$$

- though decomposition is numerically more complex final correction is a simple vector-matrix multiplication:

$$\boxed{\vec{\delta}_{ss} = \tilde{R}^{-1}\cdot\Delta\vec{x} \quad with \quad \tilde{R}^{-1} = \underline{V}\cdot\underline{\lambda}^{-1}\cdot\underline{U}^T} \quad \Leftrightarrow \quad \vec{\delta}_{ss} = \sum_{i=0}^{n}\frac{a_i}{\lambda_i}\vec{v}_i \quad with \quad a_i = \vec{u}_i^T\Delta\vec{x}$$

- numerical robust, minimises parameter deviations Δx <u>and</u> circuit strengths δ
- Easy removal of singularities, (nearly) singular eigen-solutions have λ_i~0
  - to remove those solution: if $\lambda_i \approx 0$ → '$1/\lambda_i$:= 0'
  - discarded eigenvalues corresponds to solution pattern unaffected by the FB

*G. Golub and C. Reinsch, "*Handbook for automatic computation II, Linear Algebra*", Springer, NY, 1971

Space Domain:
LHC BPM eigenvector #50 $\lambda_{50} = 6.69 \cdot 10^2$

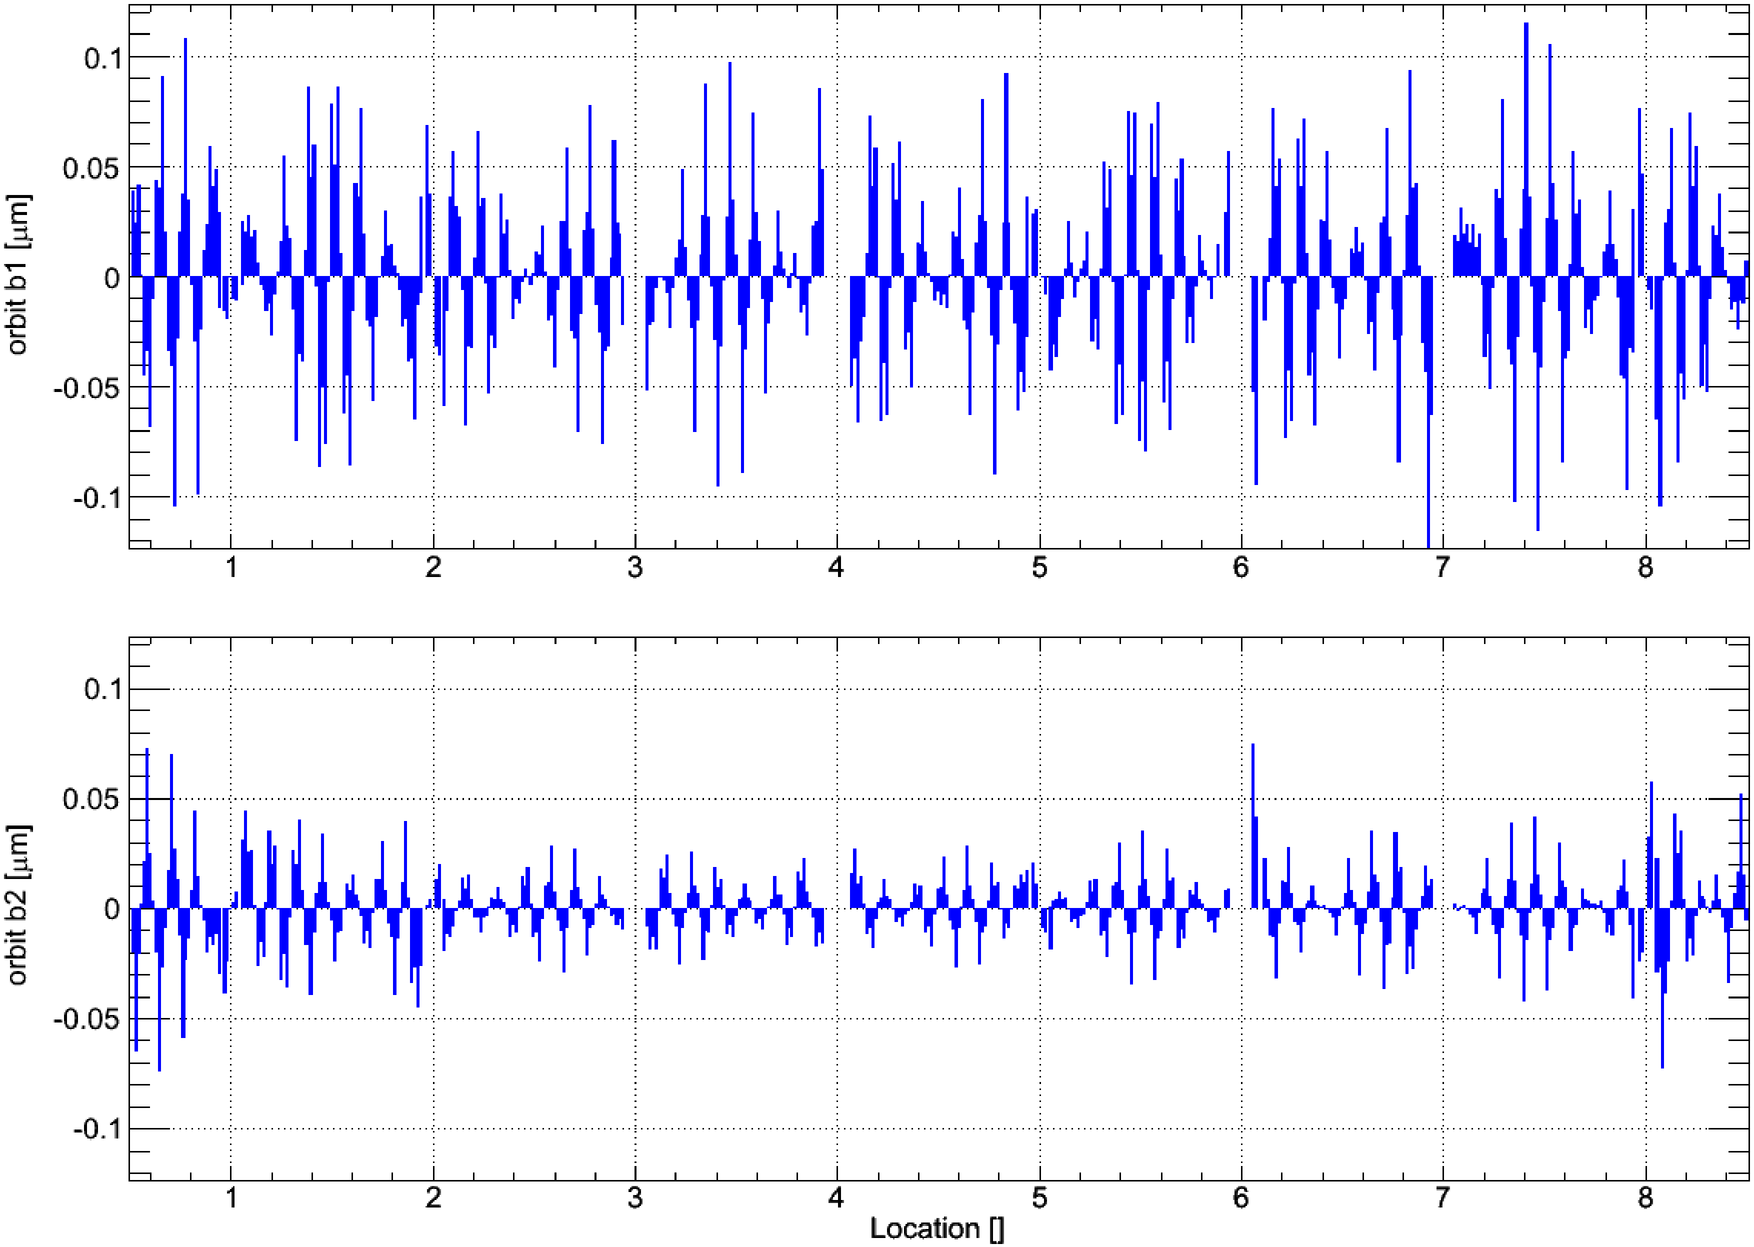

Space Domain:
LHC BPM eigenvector #100 $\lambda_{100} = 3.38 \cdot 10^2$

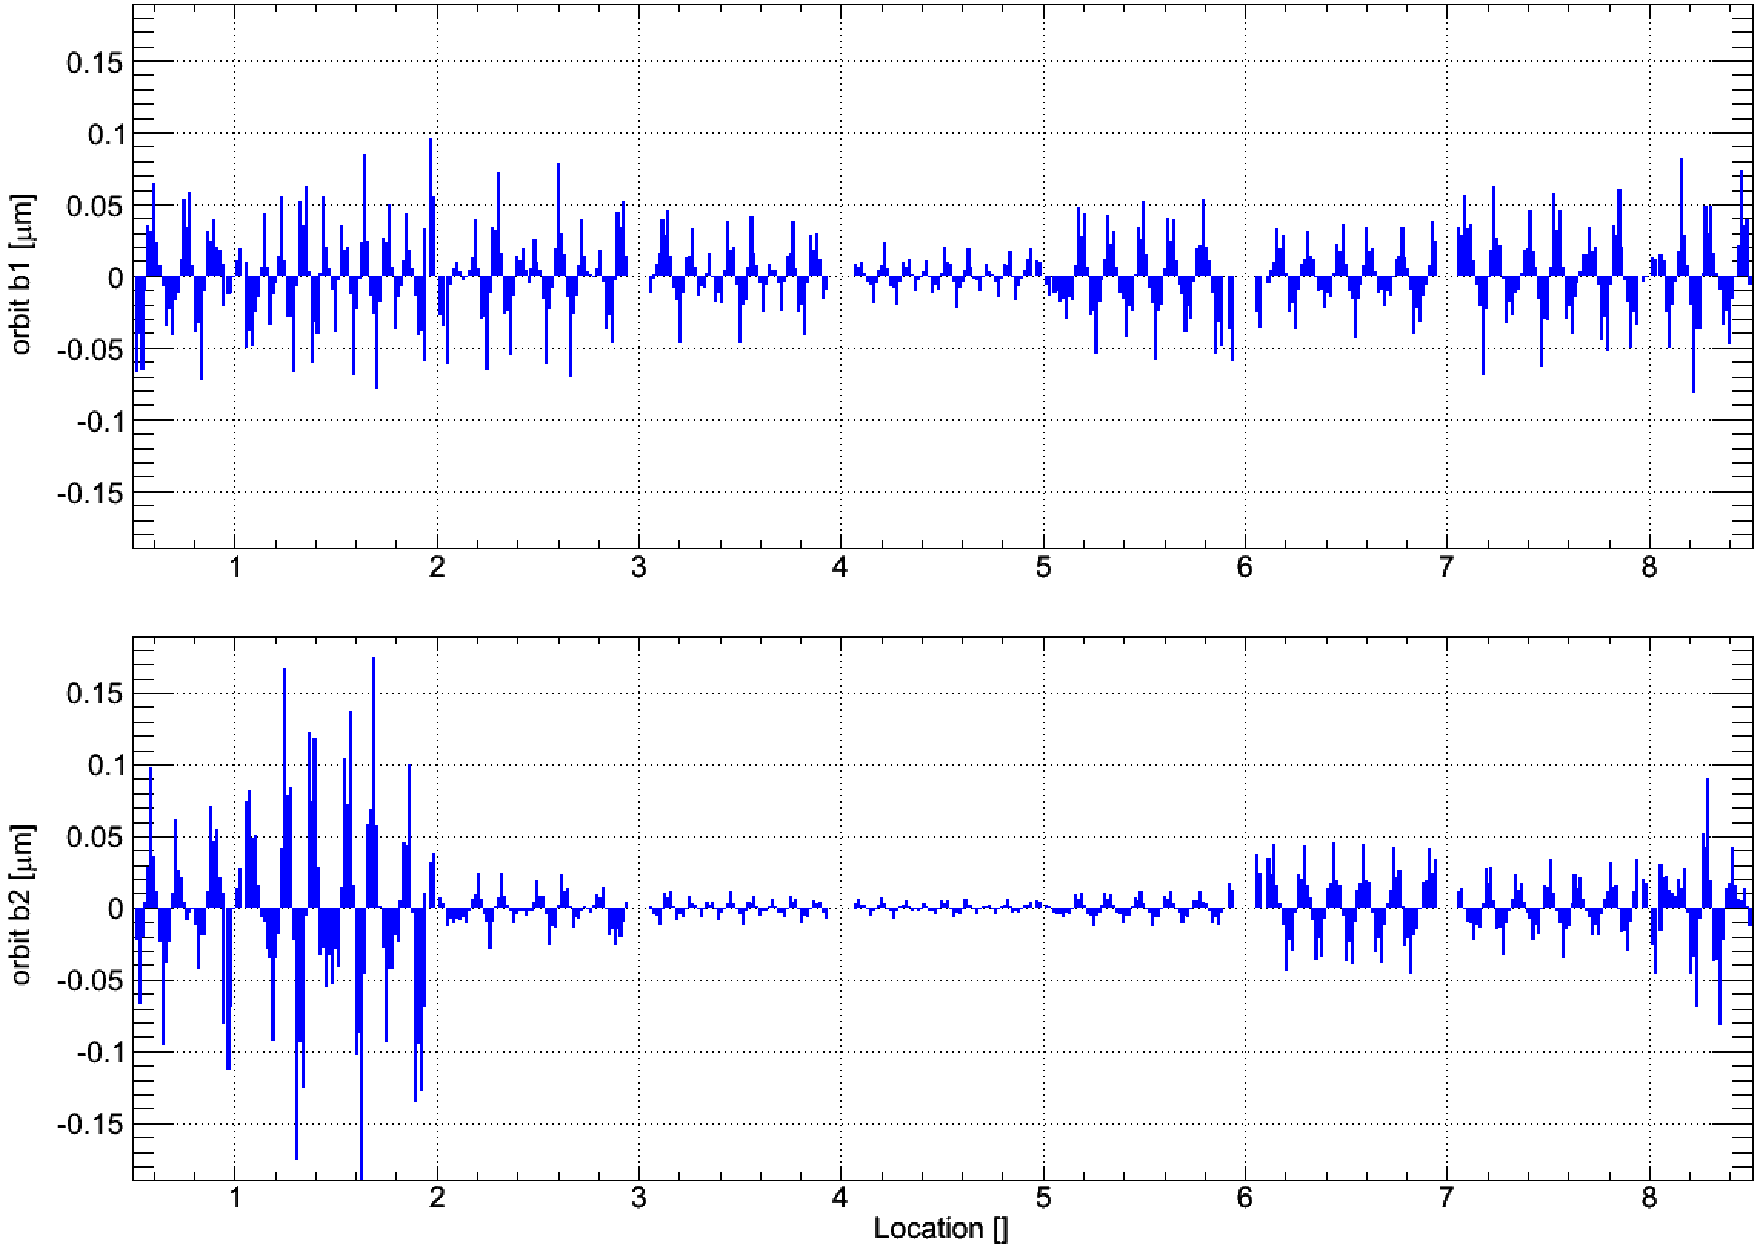

Space Domain:
LHC BPM eigenvector #291 $\lambda_{291} = 2.13 \cdot 10^2$

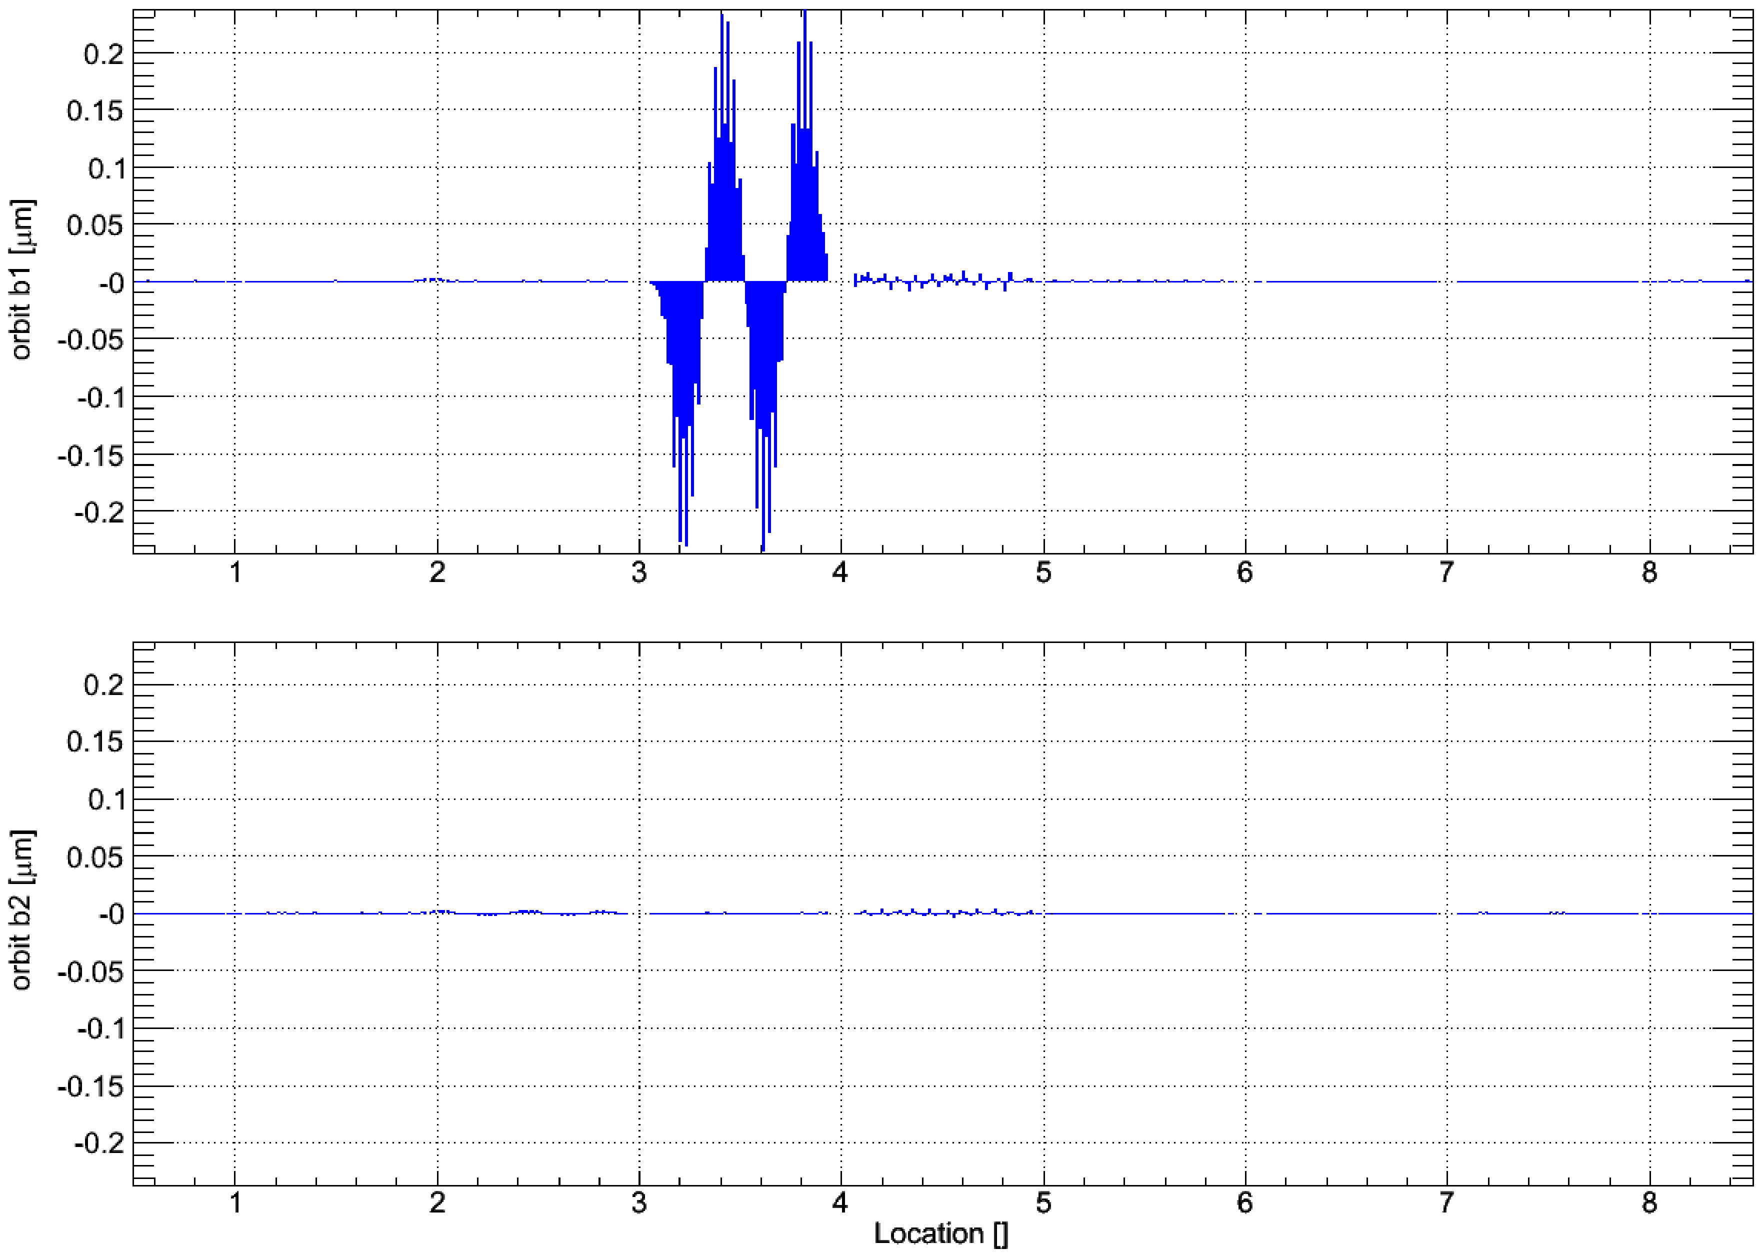

Space Domain:
LHC BPM eigenvector #449 λ_{449}= 8.17•10^1

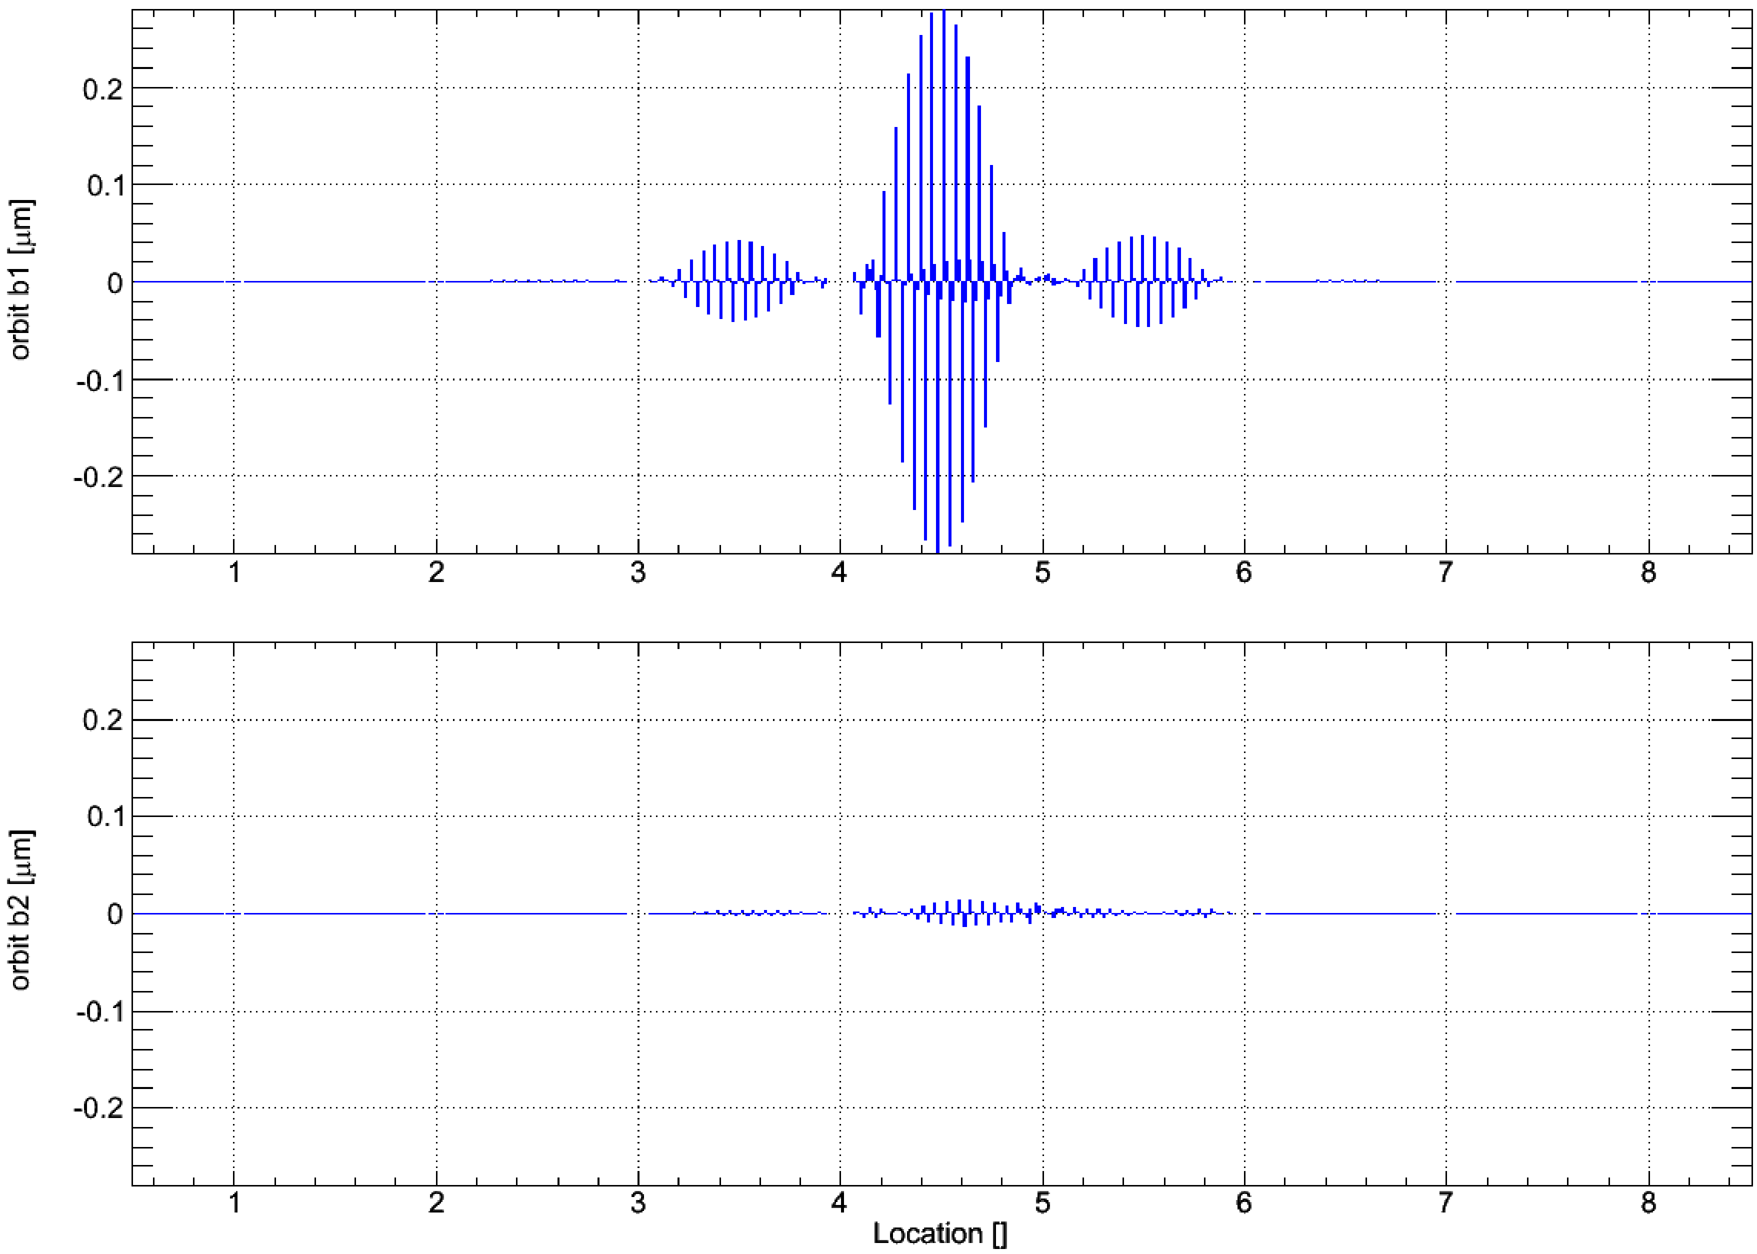

Space Domain:
LHC BPM eigenvector #521 λ_{521}= 1.18•10^0

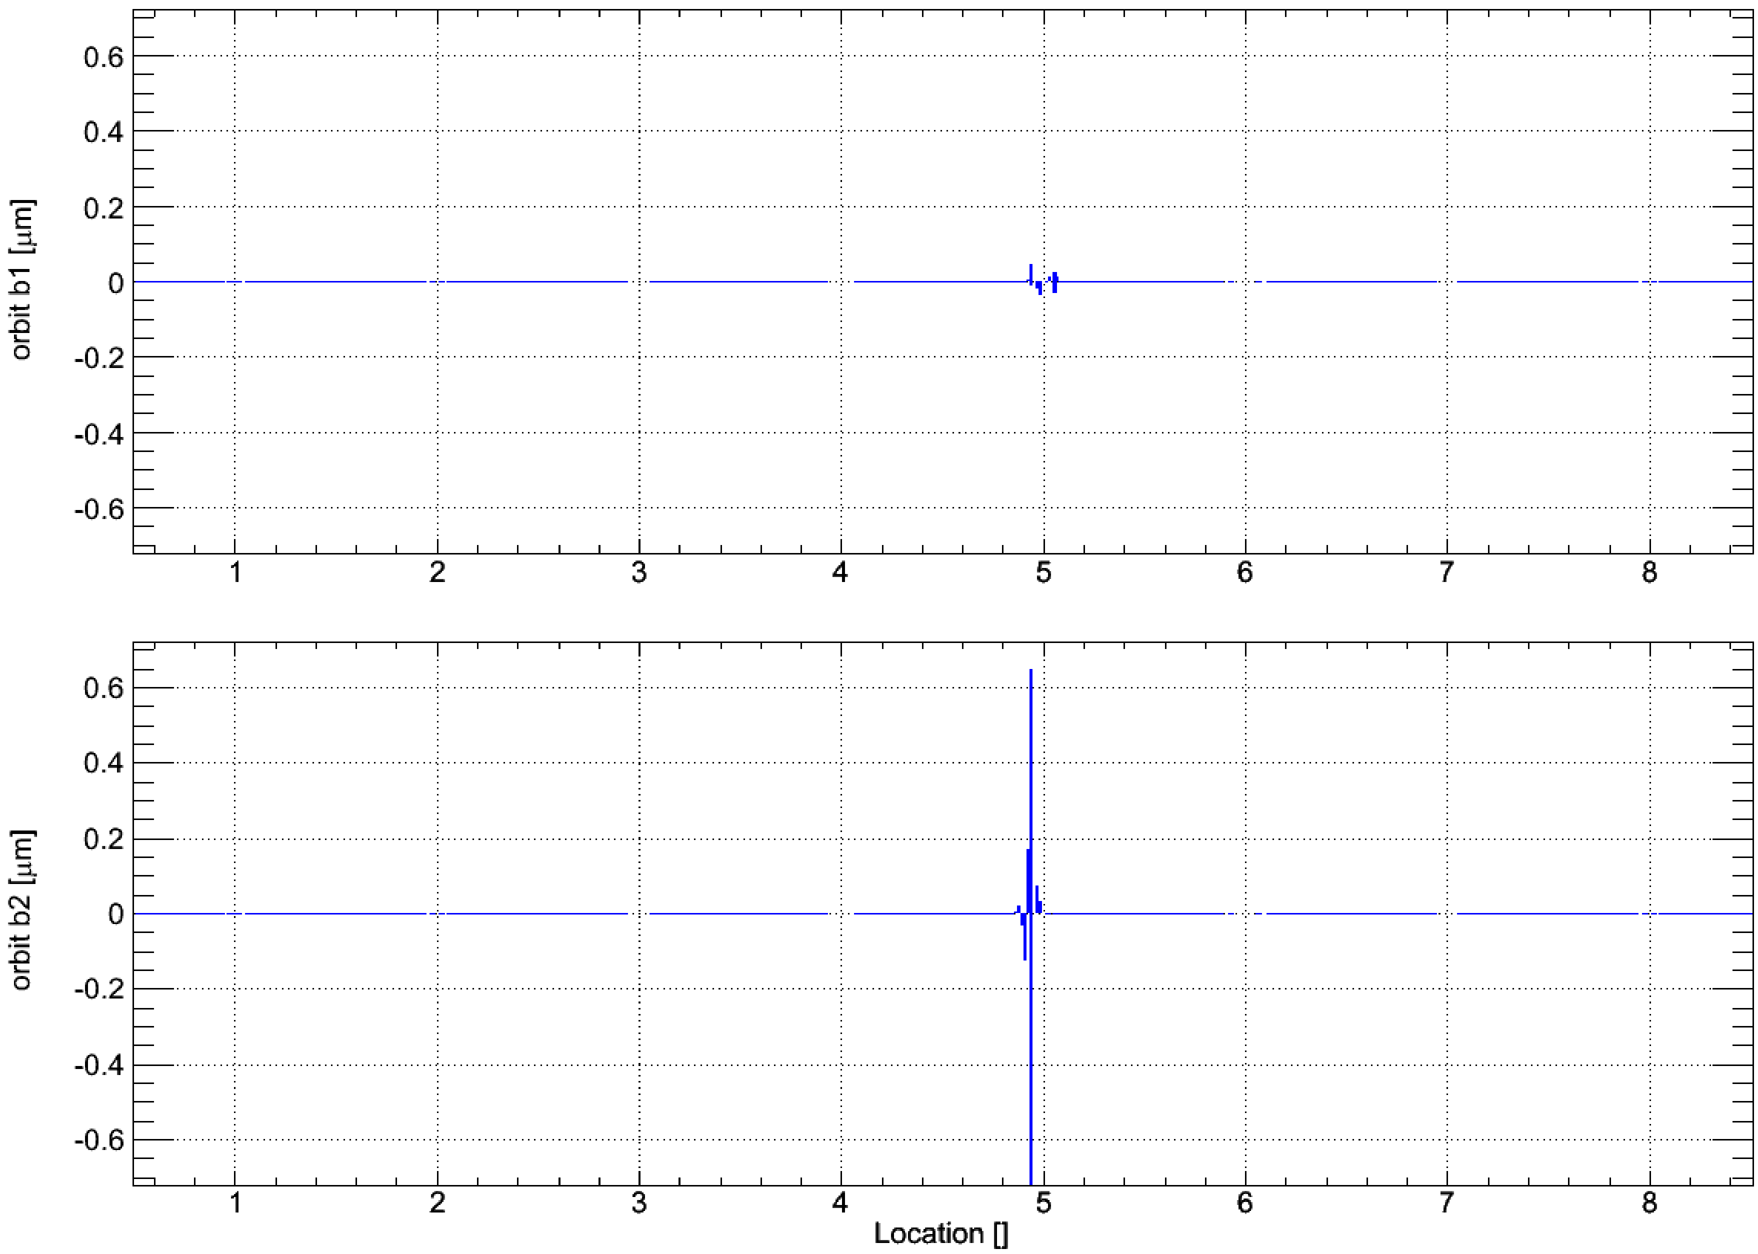

Standard Singular-Value-Decomposition based Orbit Correction

- Initially: Truncated-SVD (set $\lambda_i^{-1} := 0$, for i>N)
  - not without issues: removed λ_i allowed local bumps creeping in (e.g. collimation)

- Regularised-SVD (Tikhonov/opt. Wiener filter with $\lambda_i^{-1} := \lambda_i/(\lambda_i^2 + \mu)$, $\mu>0$)
  - more robust w.r.t. optics errors and mitigation of BPM noise/errors

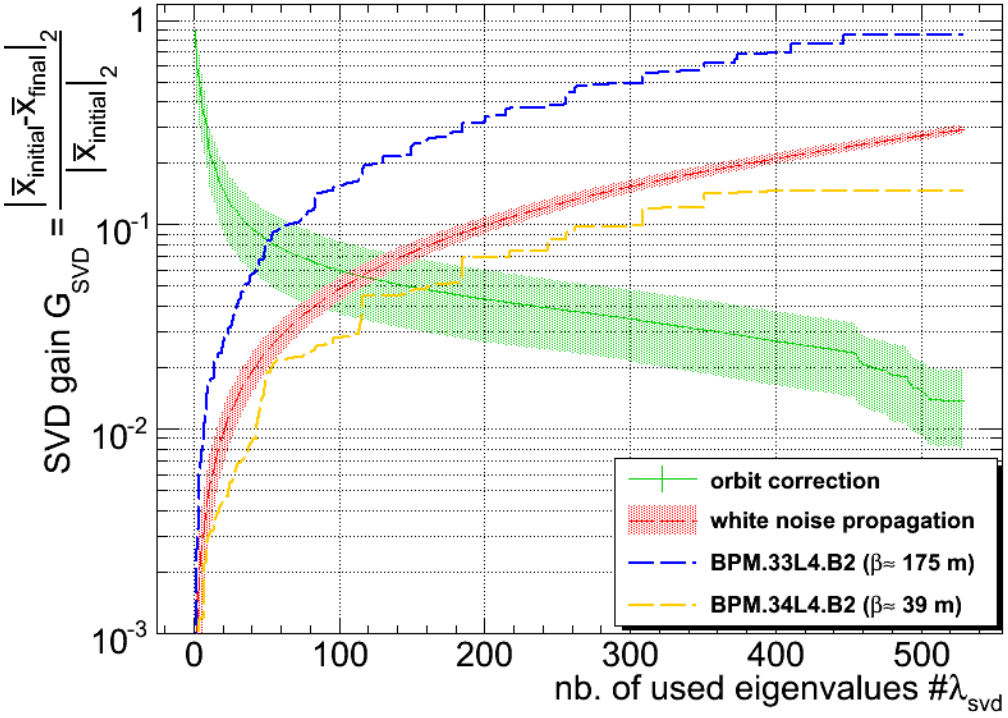

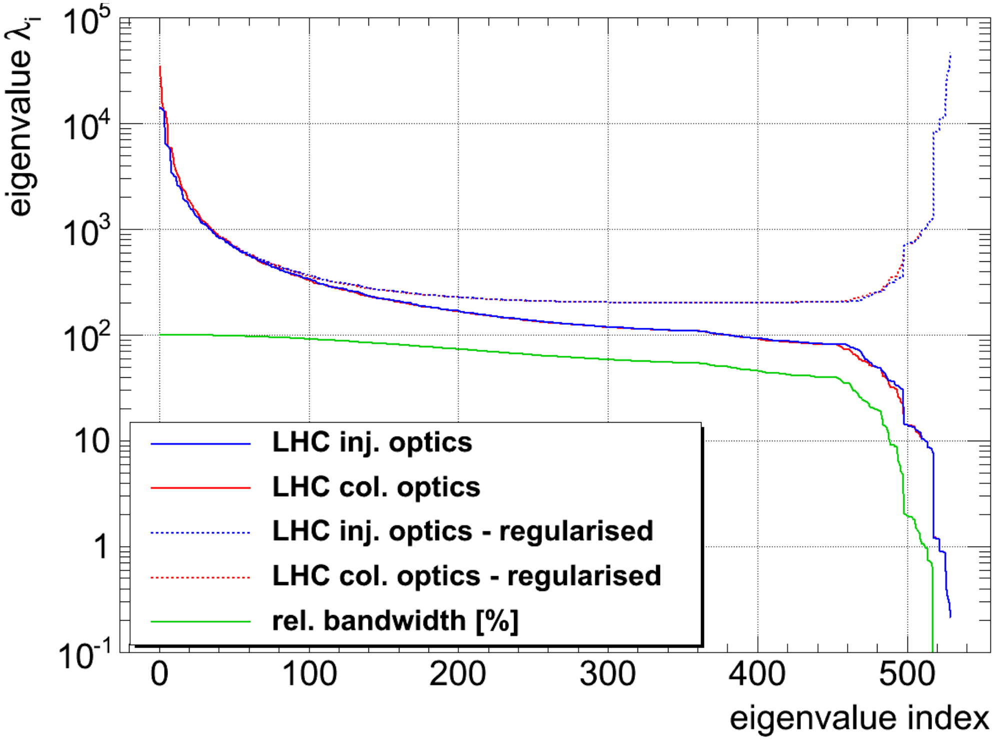

Multiple-Input-Multiple-Output (MIMO) Process Control

- 'Divide and Conquer' feedback controller design approach:
  1. Compute steady-state corrector settings $\vec{\delta}_{ss}=(\delta_1,\ldots,\delta_n)$ based on measured parameter shift $\Delta x=(x_1,\ldots, x_n)$ that will move the beam to its reference position for t→∞.
  2. Compute a $\vec{\delta}(t)$ that will enhance the transition $\vec{\delta}(t=0)\rightarrow\vec{\delta}_{ss}$
  3. Feed-forward: anticipate and add deflections δ_{ff} to compensate changes of well known and properly described sources

space domain

time domain

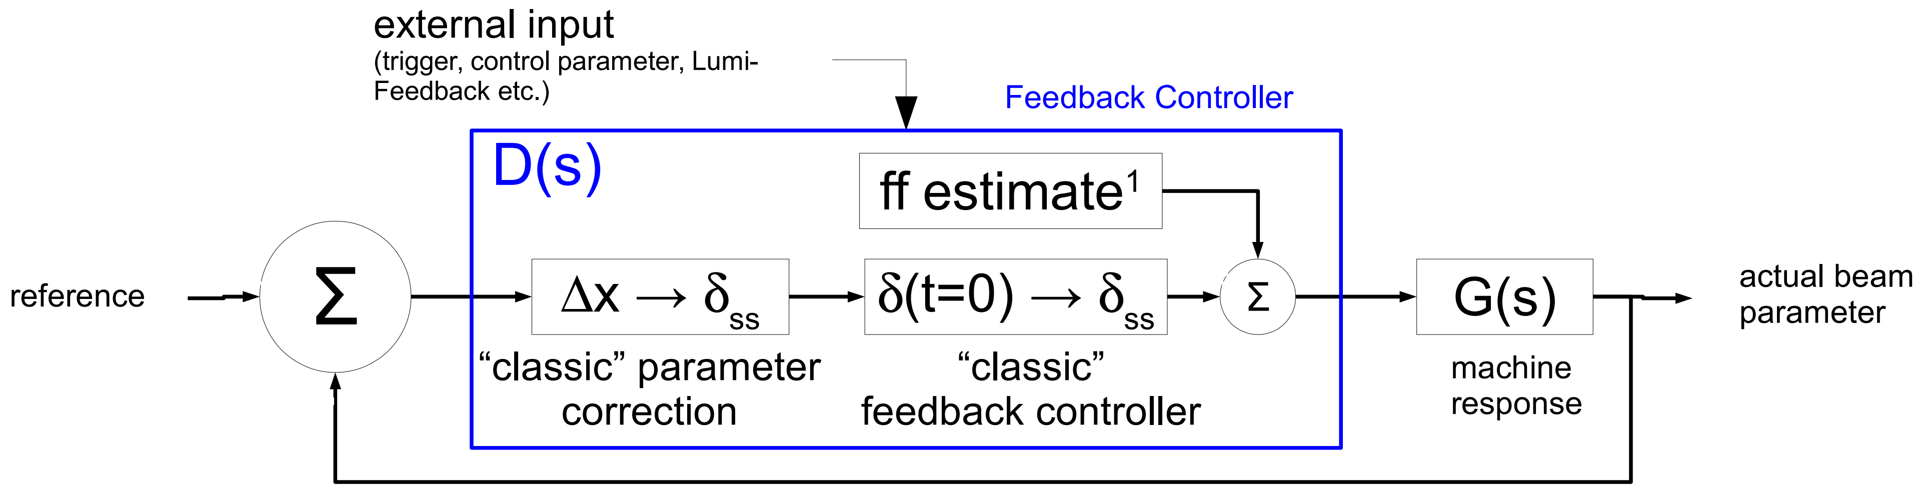

external input
(trigger, control parameter, Lumi-Feedback etc.)

Feedback Controller

D(s)

ff estimate[1]

reference

Σ

$\Delta x \rightarrow \delta_{ss}$

“classic” parameter correction

$\delta(t=0) \rightarrow \delta_{ss}$

“classic” feedback controller

Σ

G(s)

machine response

actual beam parameter

feedback-path = measured beam parameter

- (N.B. here G(s) contains the process and monitor response function)

Time-Domain: Optimal Controller Design II/II

- Optimal control[1] for the 'small-signal response' yields classic PI-controller
  - Single free parameter α > τ…∞
    - facilitates trade-off between speed and robustness
    - adaptive gain-scheduling based on operational scenario

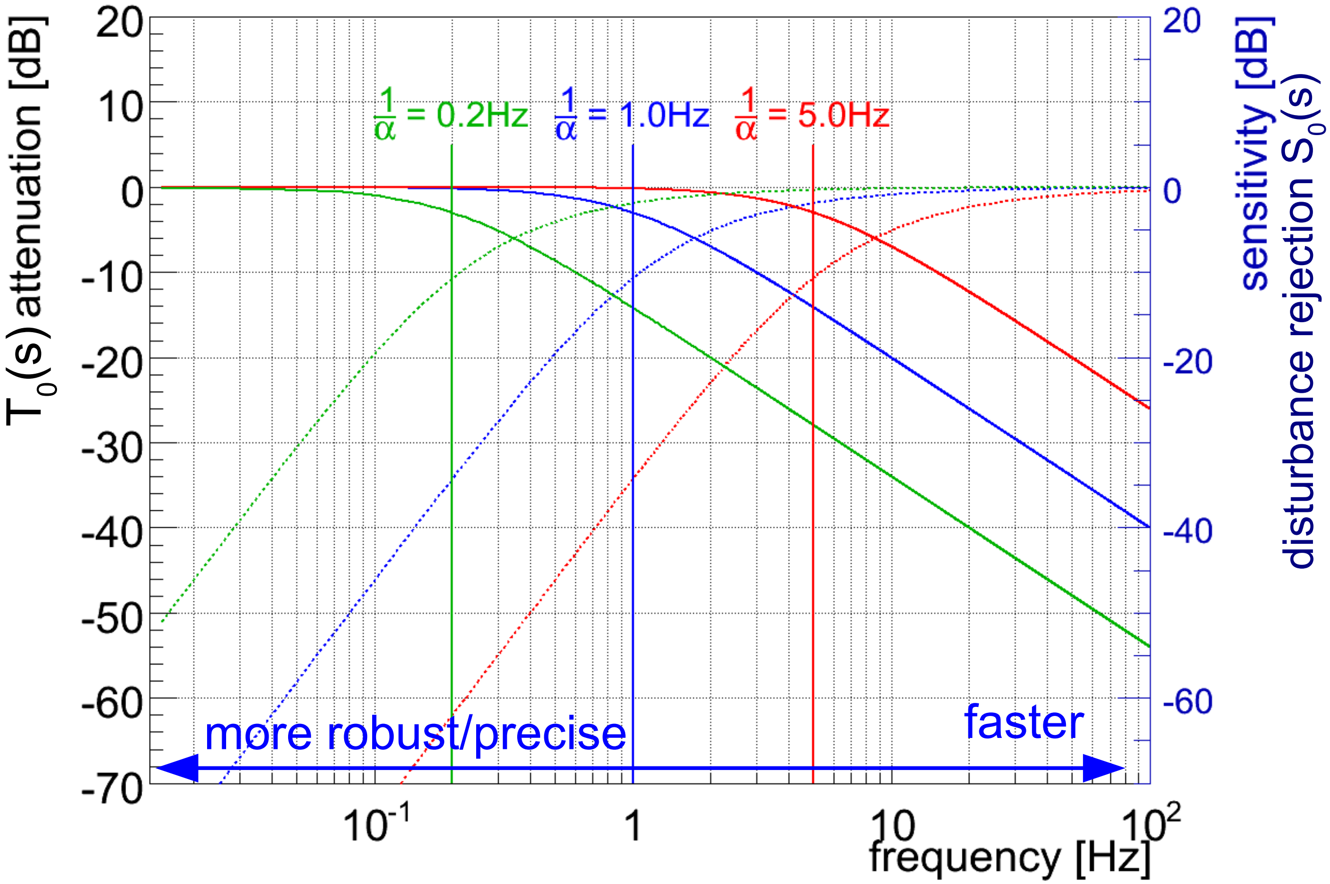

---

[1] D. C. Youla et al., *"Modern Wiener-Hopf Design of Optimal Controllers"*, IEEE Trans. on Automatic Control,1976, vol. 21-1,pp. 3-13 & 319-338

Time-Domain: Non-Linearities I/III

Two non-linear effects that need to be addressed by the controller:

- Delays: computation, data transmission, dead-time, etc.
- Rate-Limiter: limited slew rate of corrector circuits (due to voltage limitations)
  - e.g. LHC: ±60A converter: $\Delta I/\Delta t|_{max} < 0.5$ A/s

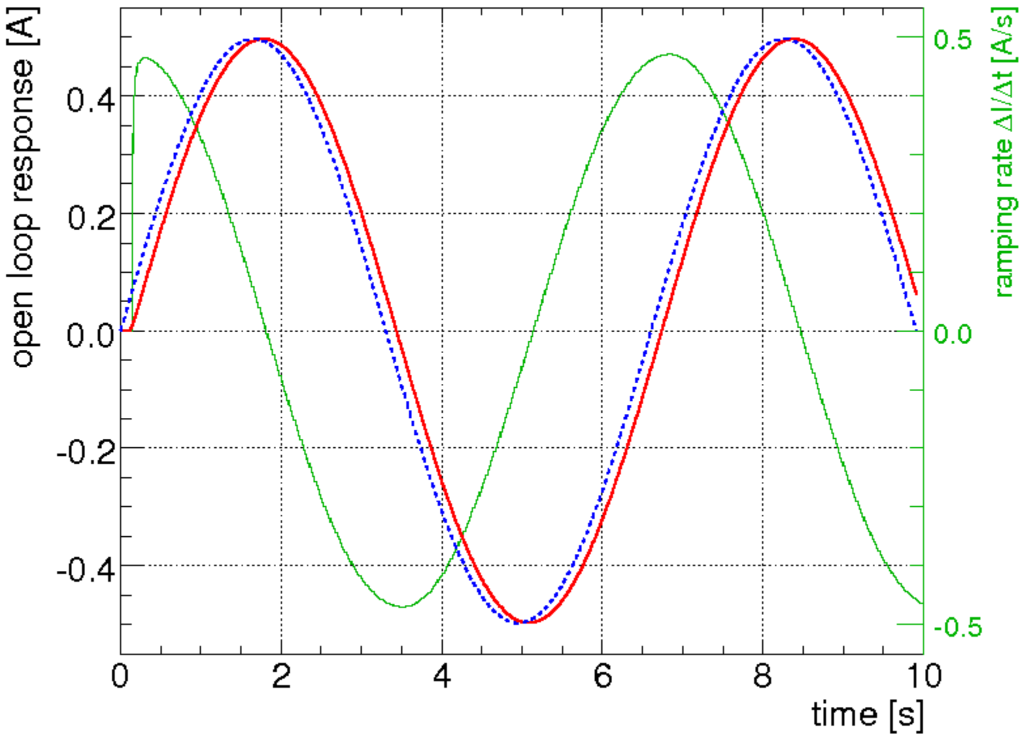
slow perturbation: perfect tracking

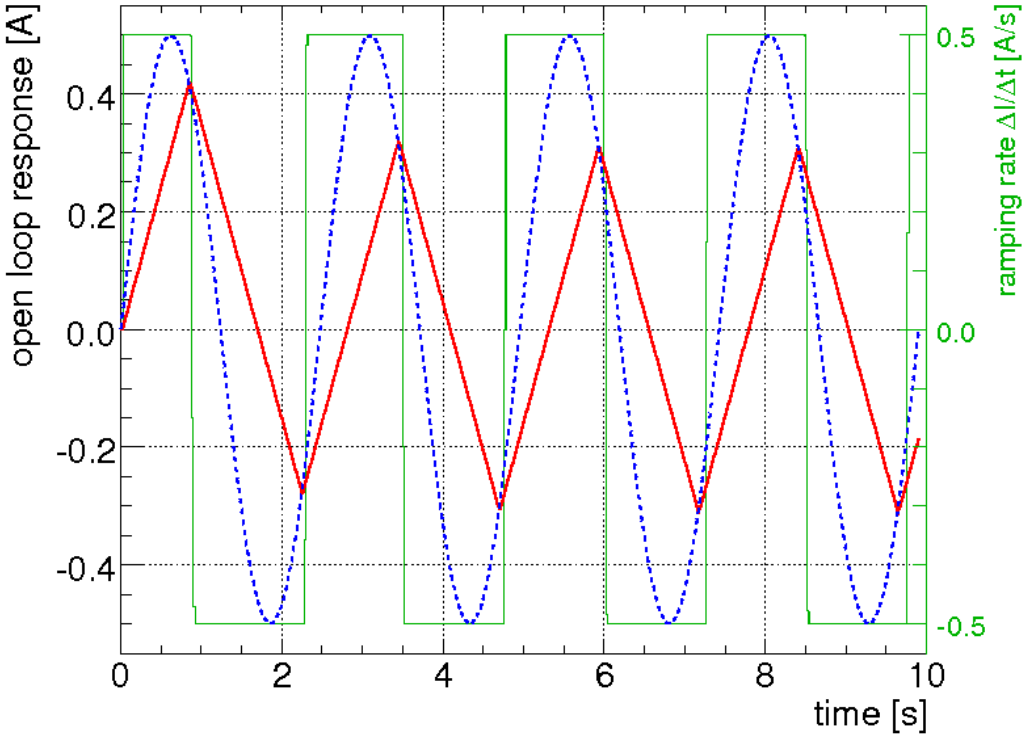
fast perturbation: saw-tooth

Time-Domain: Non-Linearities II/III

- Rate-limiter in a nut-shell:
  - additional time-delay Δτ that depends on the signal amplitude
  - secondary: introduces harmonic distortions

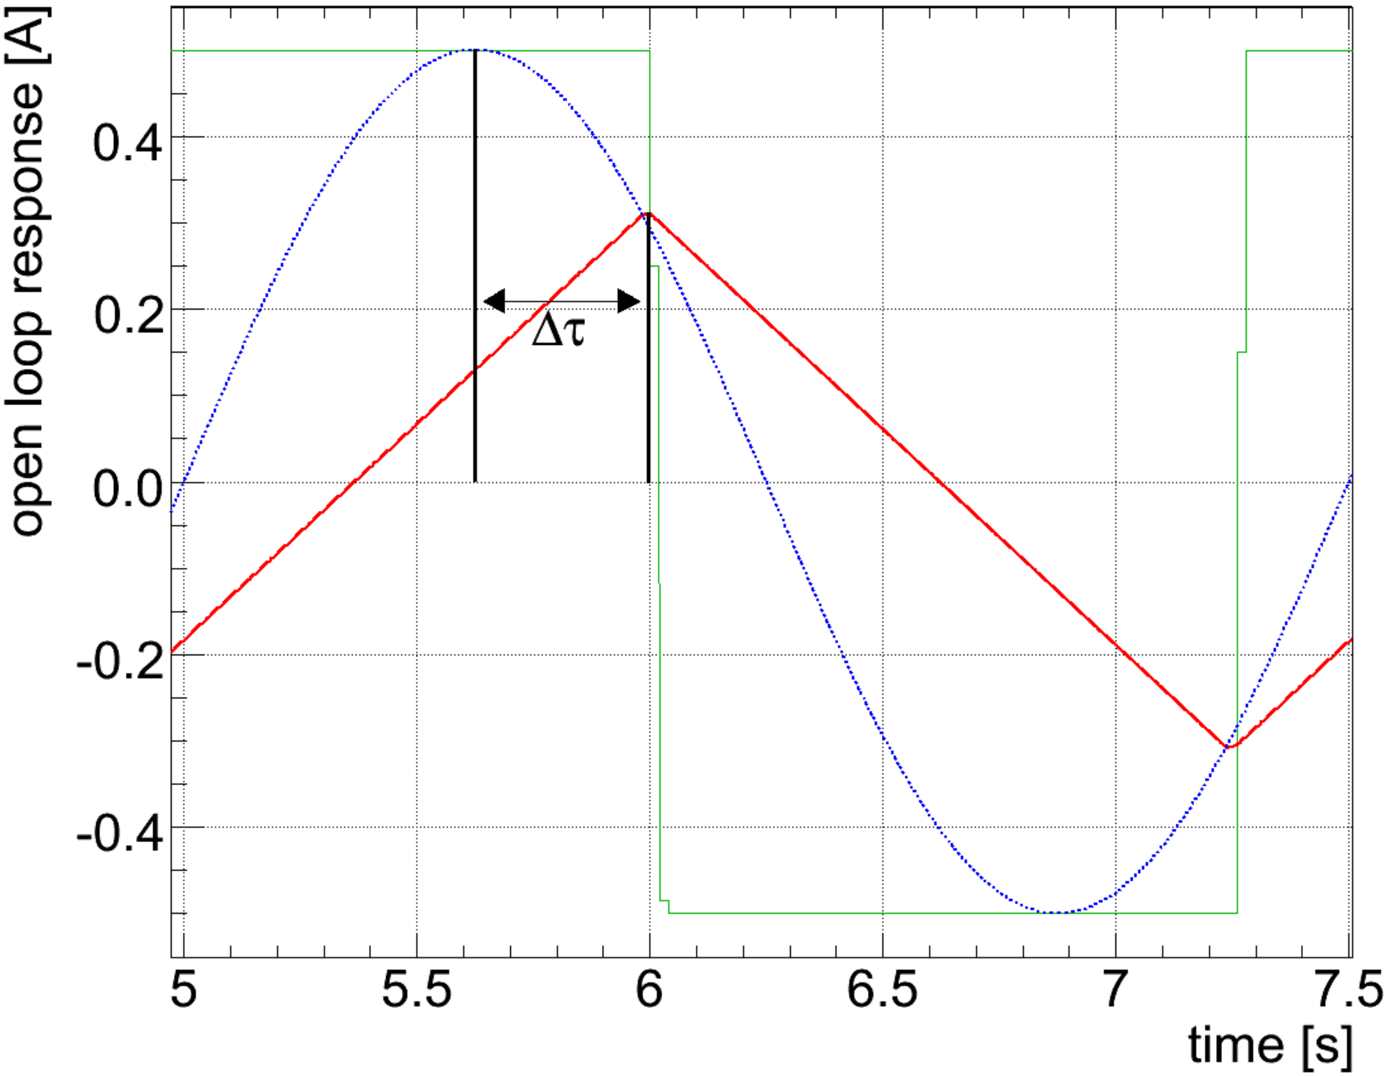

Time-Domain: Non-Linearities III/III
Orbit Feedback Bandwidth vs. Actual Orbit Perturbations

- Closed-loop bandwidth and phase margin depend on the excitation amplitude:
  - + non-linear phase once rate-limiter is in action...

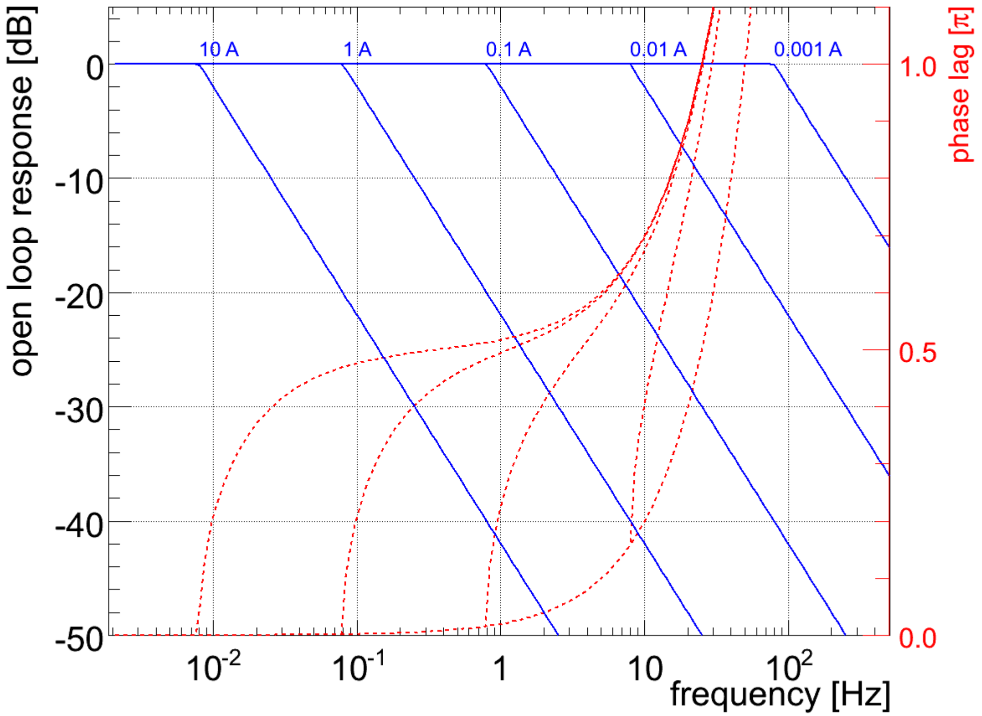

$\Delta I = 0.1A \leftrightarrow \Delta x \approx 32$ μm@β=180m

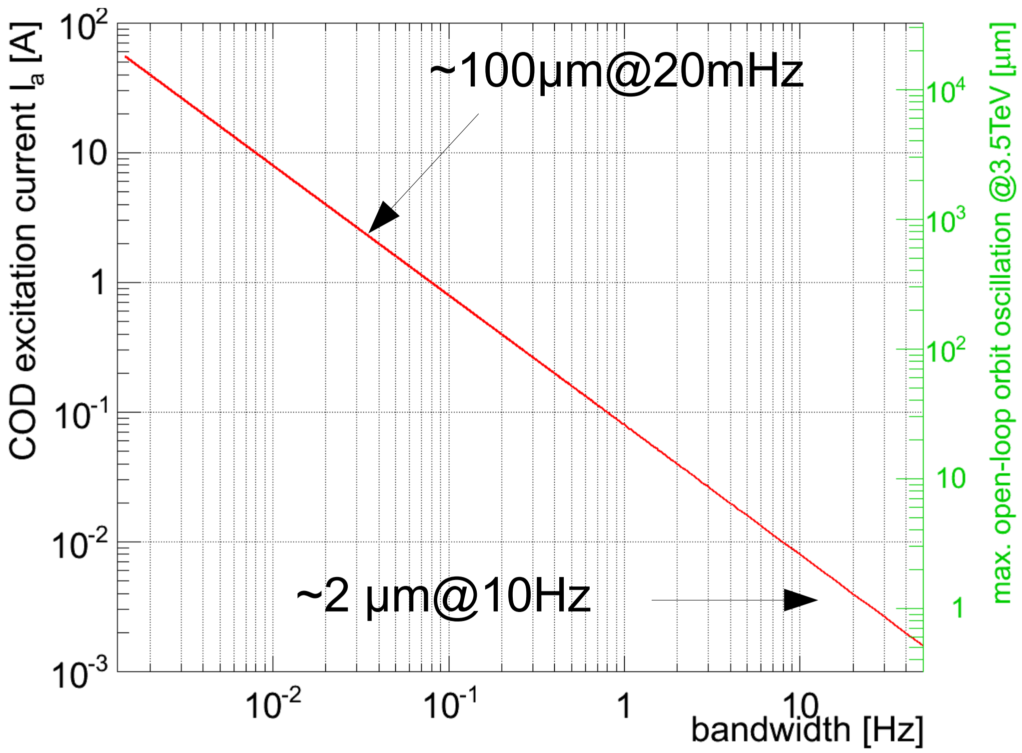
~100μm@20mHz

~2 μm@10Hz

Functional Description of the LHC Feedback Controller
PID Controller & Delays + Rate-Limiter

- In essence, the functional OFC description

$$G(s)=\frac{e^{-\lambda s}}{\tau s+1}G_{NL}(s)$$

  - τ: the power converter time constant and

→ Smith-Predictor and Anti-Windup paths:

$$\boxed{D(s)=\frac{Q(s)}{1-Q(s)G(s)}}$$

$$G^{i}(s)=\frac{\tau s+1}{1}$$

$$T(s)=F_Q(s)\cdot e^{-\lambda s}G_{NL}(s)$$

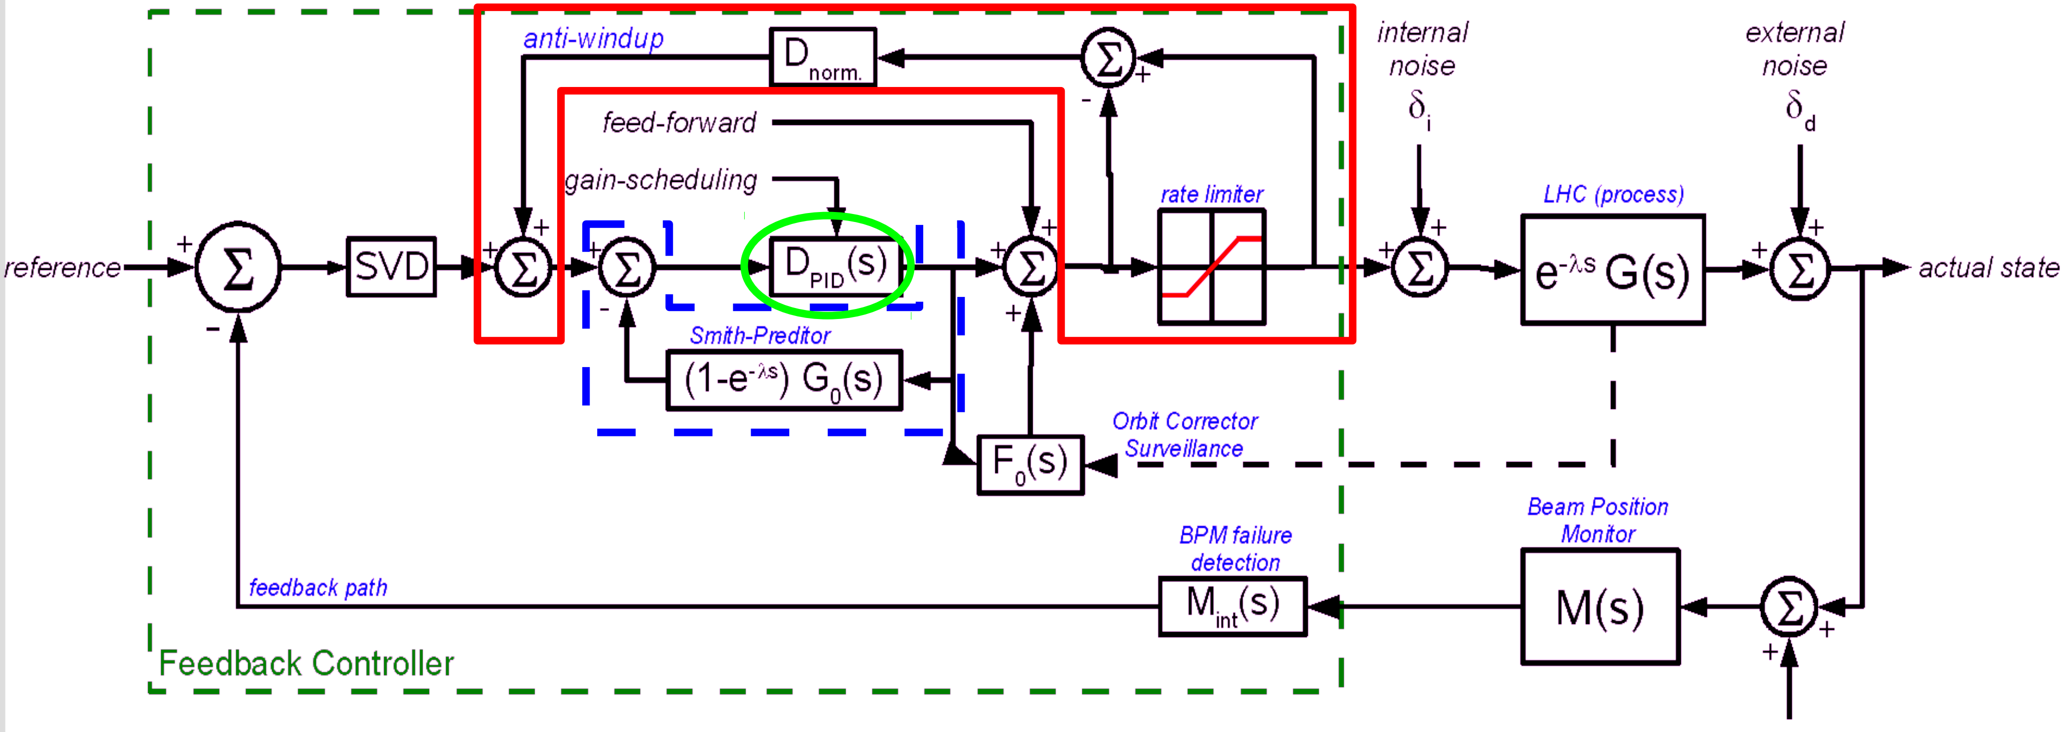

- “simple loop design” but large number of control parameters (→ adds complexity to OFC/OFSU design)

Motivation for Delay and Rate-Limiter Compensation Example: LHC orbit (Q,Q',C⁻, ...) feedback control

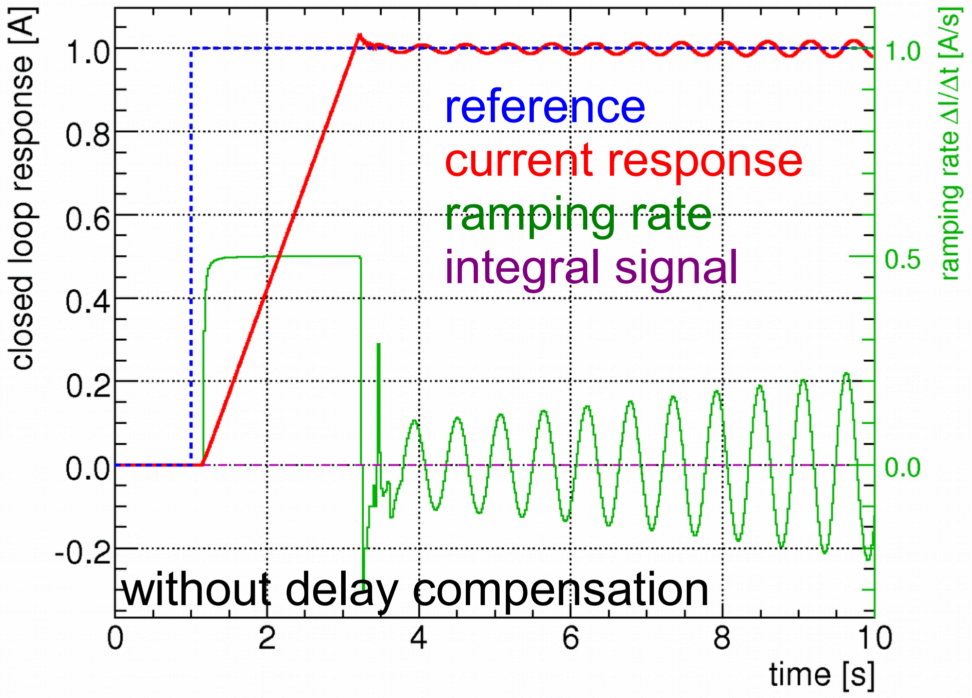

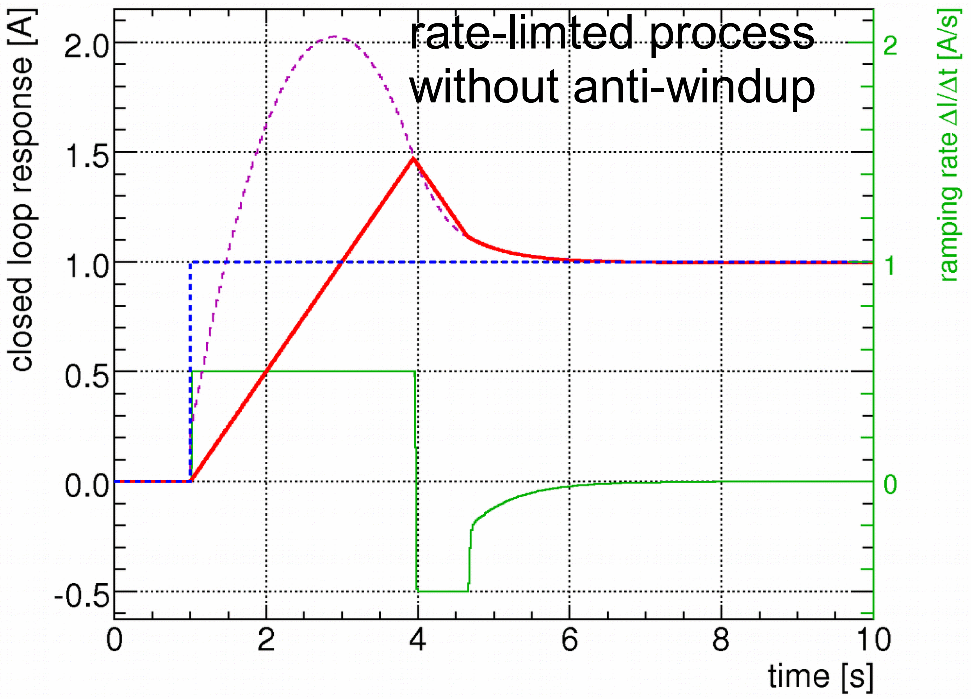

with full delay and windup compensator scheme:

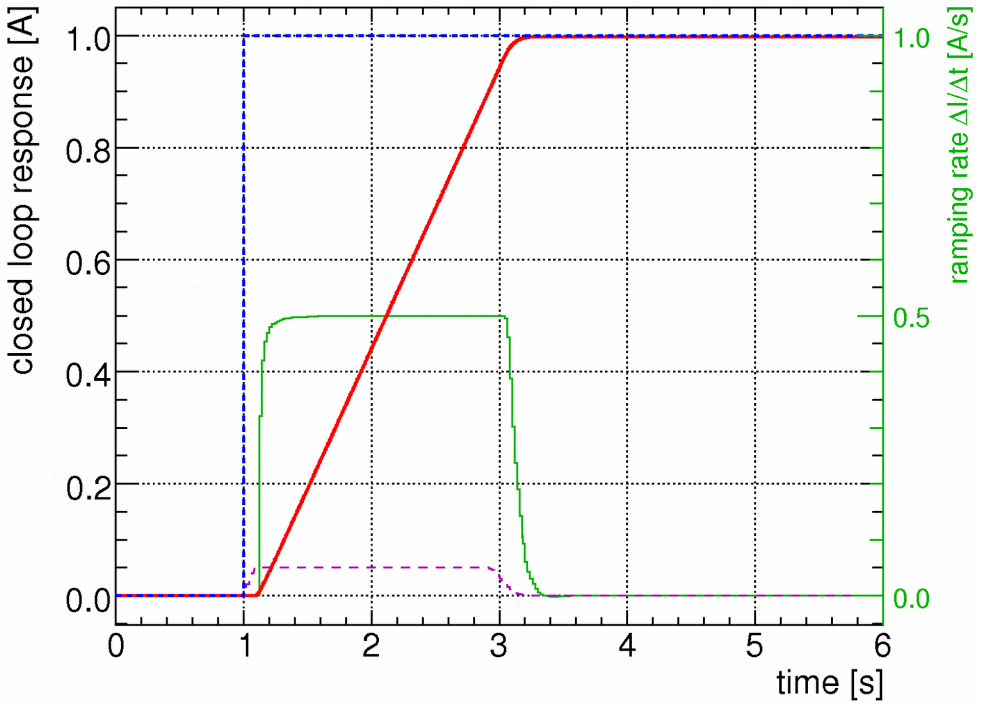

To avoid inherent Cross-Talk between FBs…
… Cascading between individual Feedbacks

- Main strategy: derive measurement from FB control variable
  - Q'-tracker using '$Q_{raw} = Q_{meas} - Q_{trim}$'
  - Sub. Δp/p-mod. from Radial-Loop & Orbit-FB reference

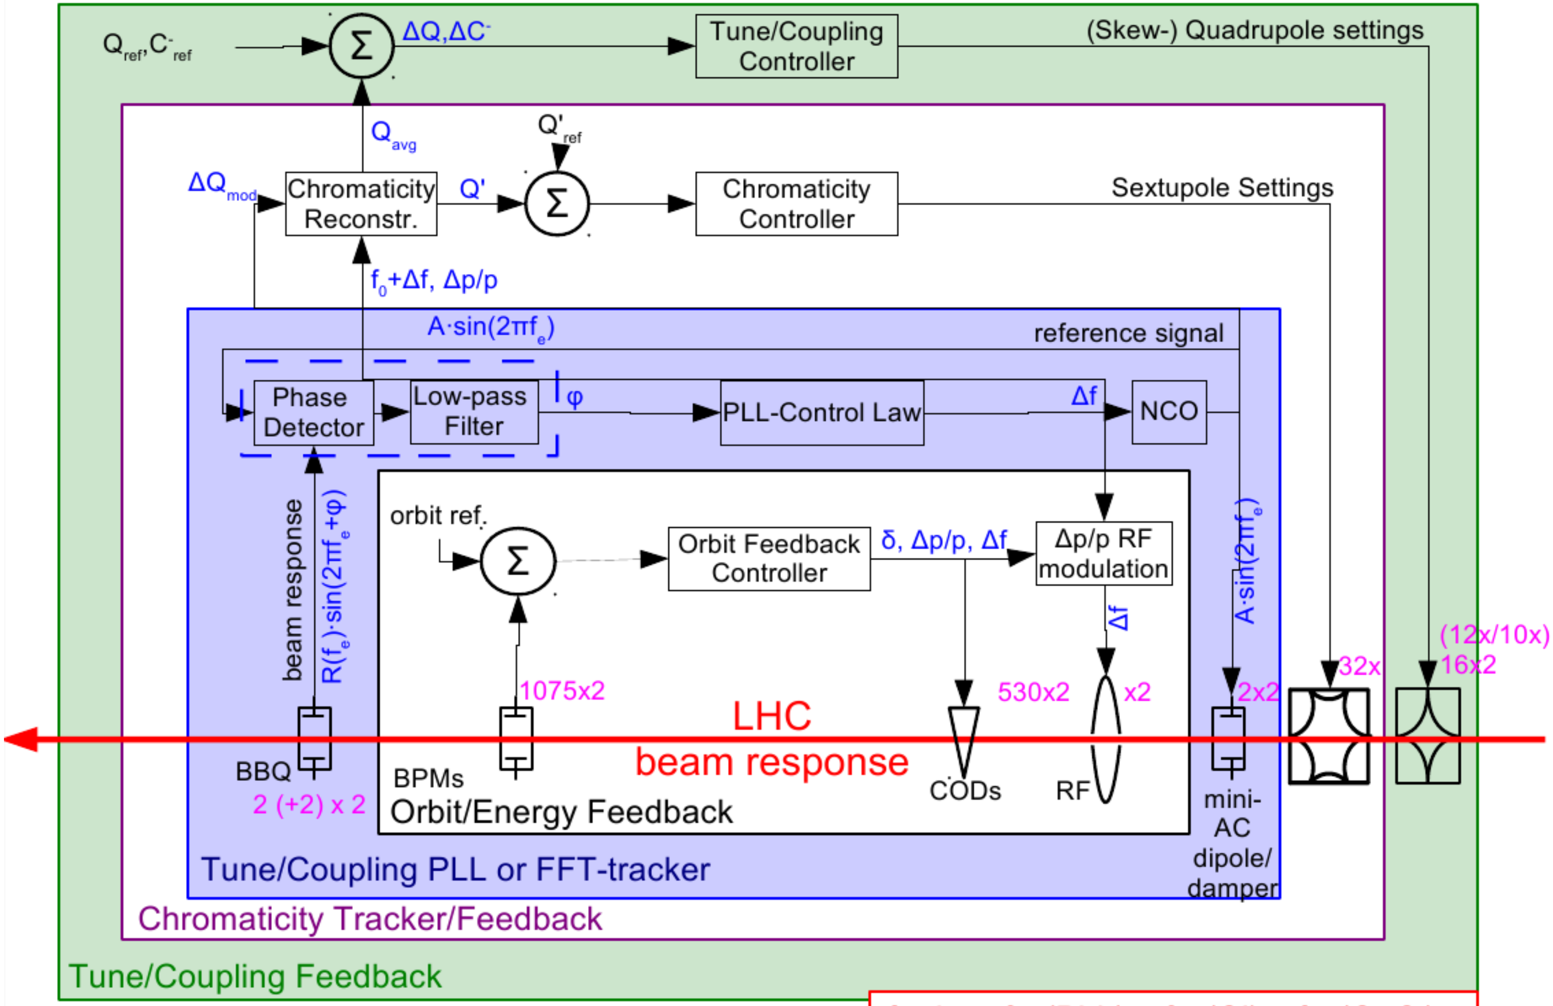

further: $f_{BW}(PLL) \gg f_{BW}(Q') \geq f_{BW}(Q, C^{-})$

Common Feedback/Feed-forward Control Layout
Control implementation split into two sub-systems:

- LHC feedback systems most visible faces are:
  - Feedback Controller (OFC): actual feedback controller logic
  - Service Unit (OFSU): Interface to control system/the world

- However 3500+ devices (~130 FE) and many technical services
  - Overall strength depends on the reliability of the weakest link

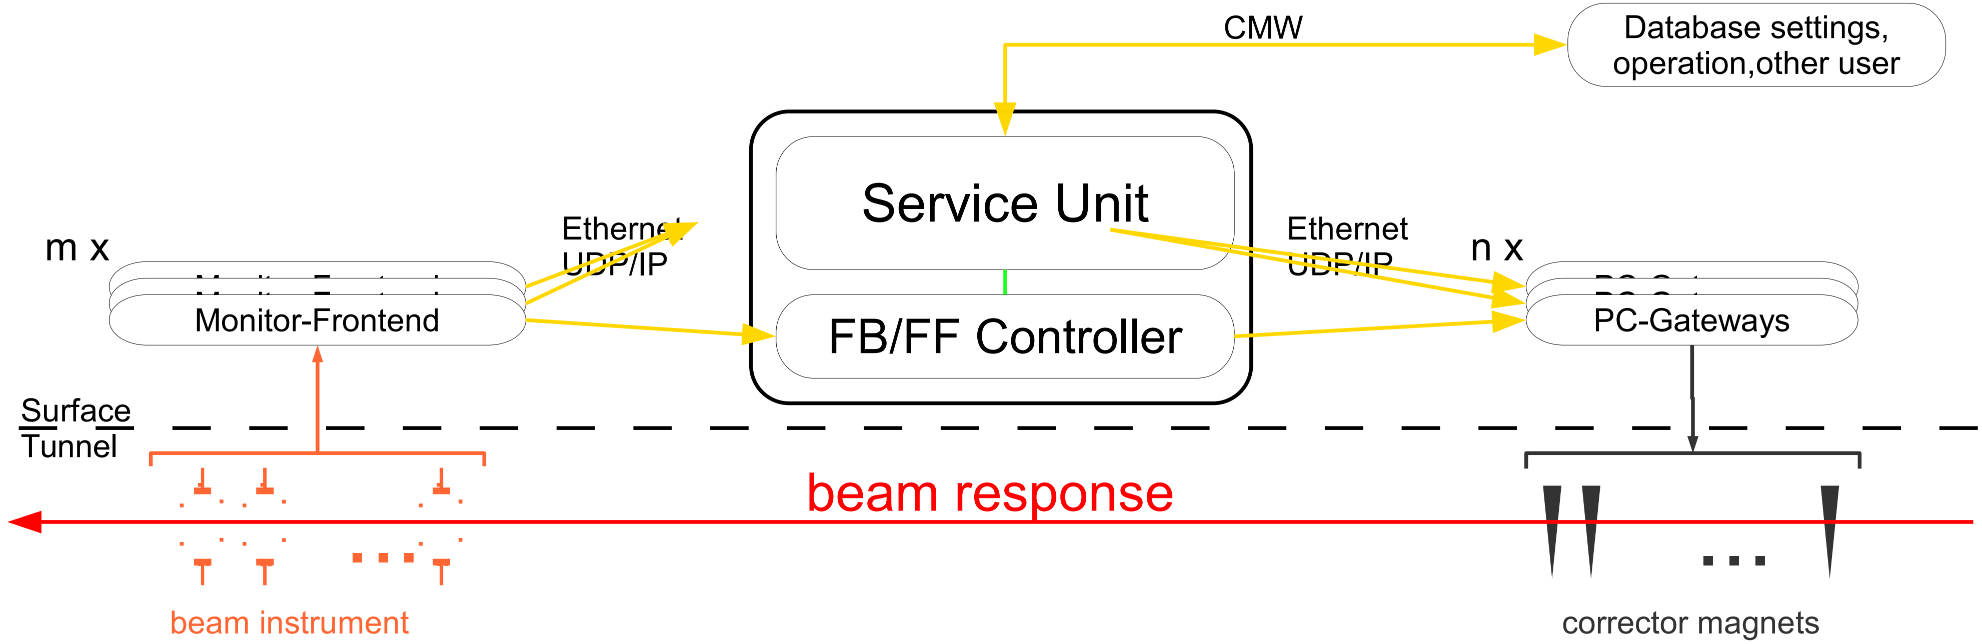

Control Paradigms III/III
Digital Control System & 'Real-Time'

- LHC feedbacks are 'firm real-time systems'
  - some (limited) margin on occasional missing data
  - additional latencies are critical for loop stability

$$\Delta\varphi = 2\pi f_{bw} \cdot \Delta t_{delay}$$

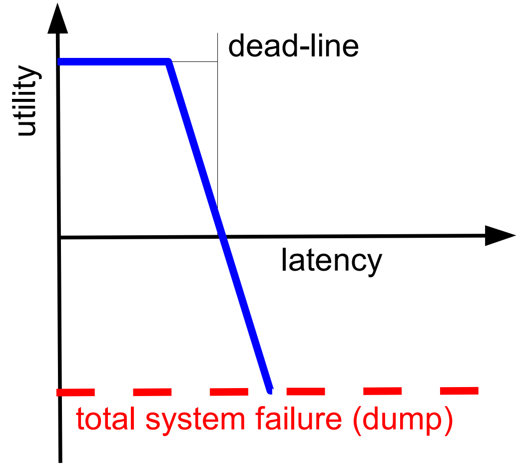

- “How much phase stability is required (i.e. @1 Hz)?”

perturbation, phase error
Correction, res. error

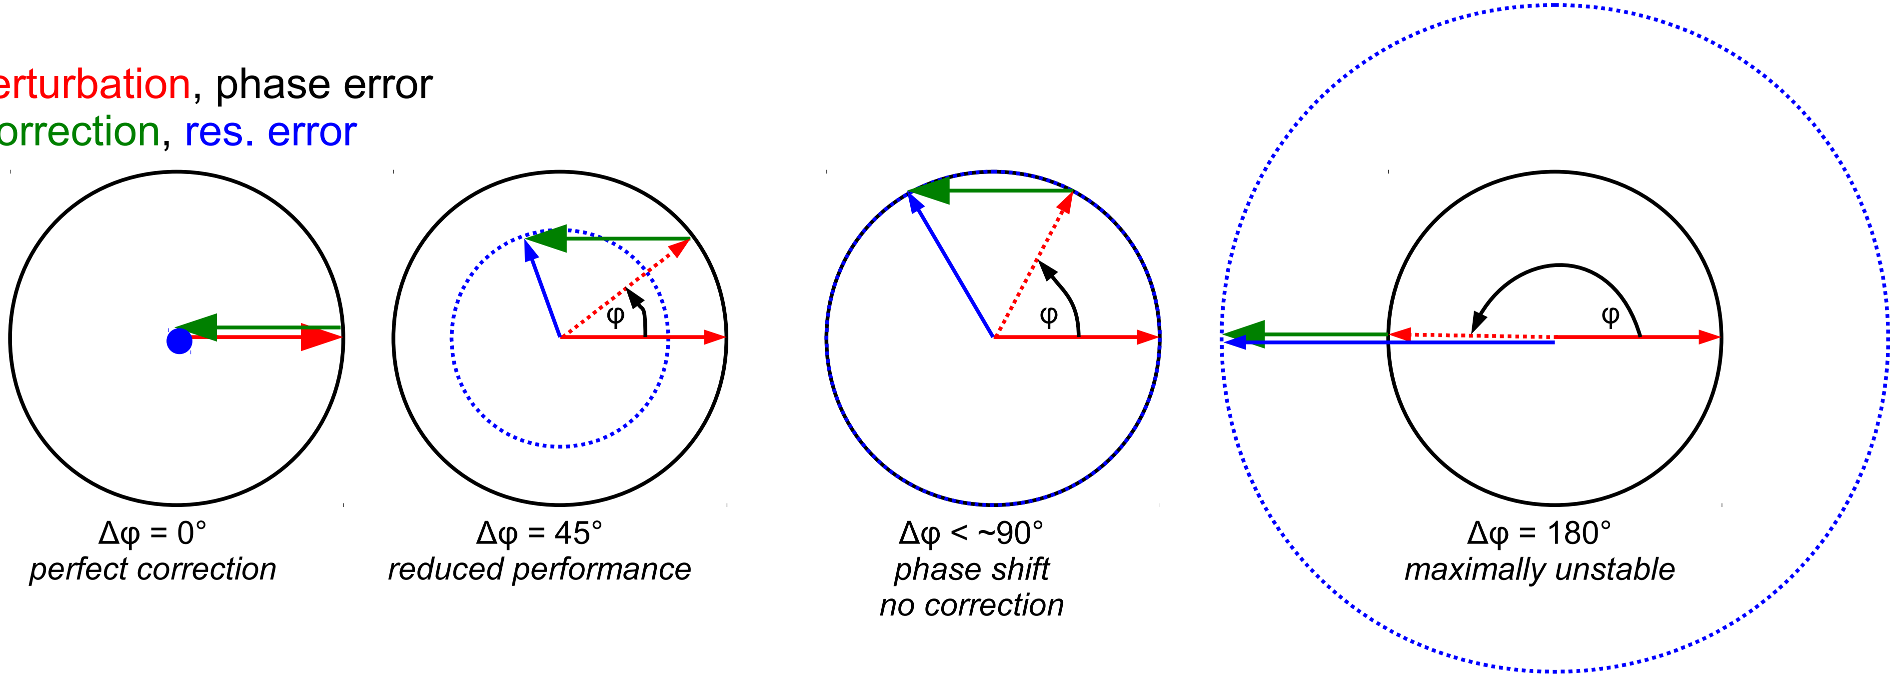

Δφ = 0°
perfect correction

Δφ = 45°
reduced performance

Δφ < ~90°
phase shift
no correction

Δφ = 180°
maximally unstable

Jitter Compensation
Data Concentration Paradigm I/II

Two main strategies:

- actual delay measurement and dynamic compensation in SP-branch:
  - only feasible for small systems

- Jitter compensation using a periodic external signal:
  - CERN wide synchronisation of events on sub ms scale
  - The total jitter, the sum of all worst case delays, must stay within “budget”.
  - feedback loop frequency of 50 Hz feasible for LHC, if required...

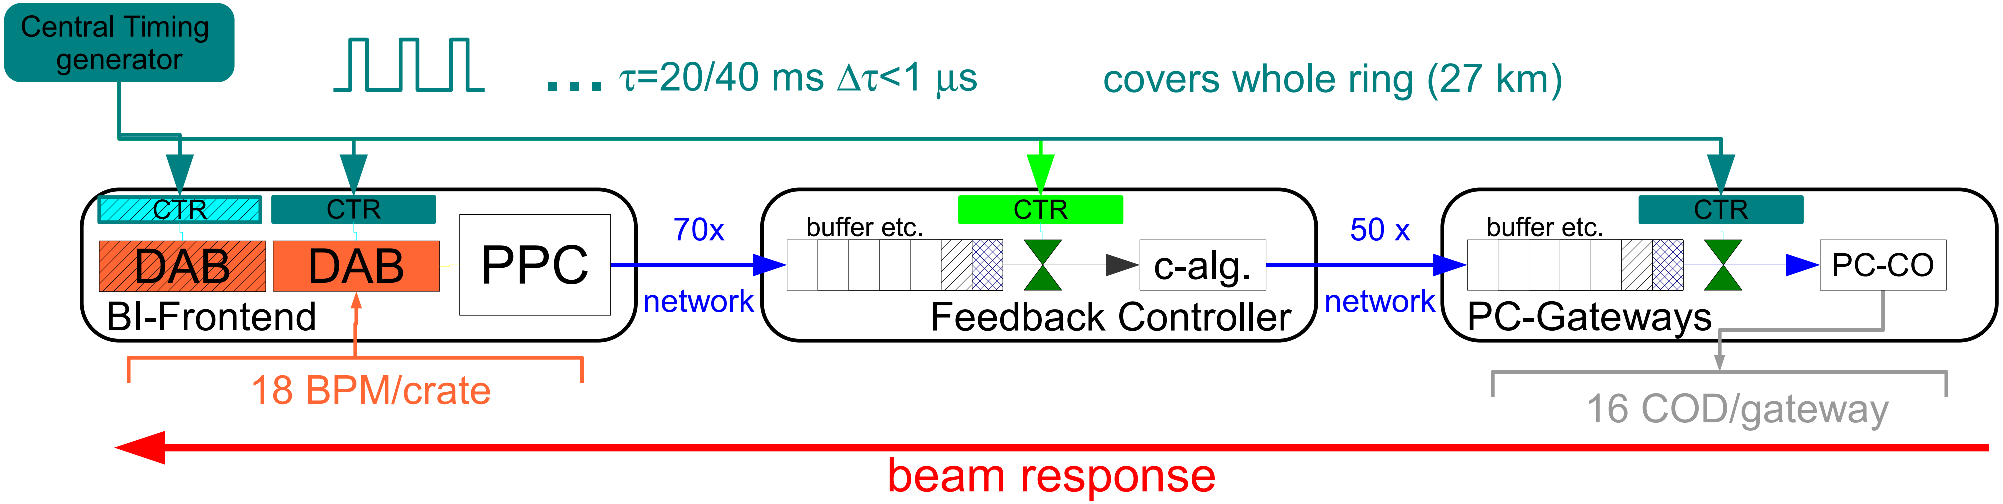

Verification of Real-Time Behaviour
– Measuring how well Real-Time Constraints are being kept

- TechNet round-trip tests: difference between standard and RT-Linux Kernel
  - Important: measure probability & upper-bound (worst-case) latency

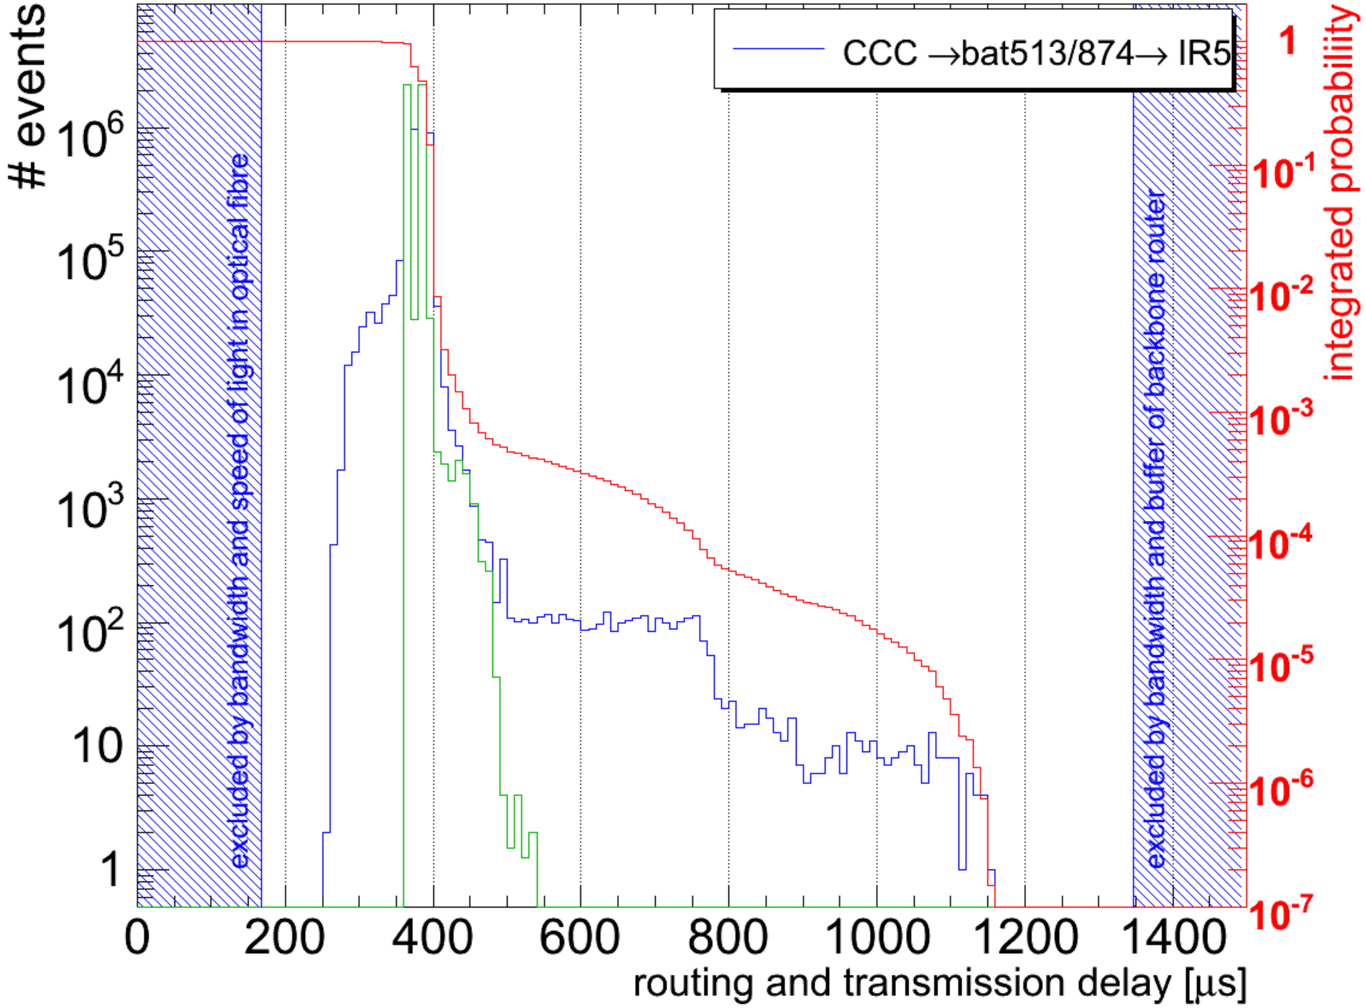

- Latencies dominated by speed-of-light in fibers but nevertheless << 1 ms

Tune Spectra during High-Intensity Operation

- Snapshot of the day with removed mains harmonics

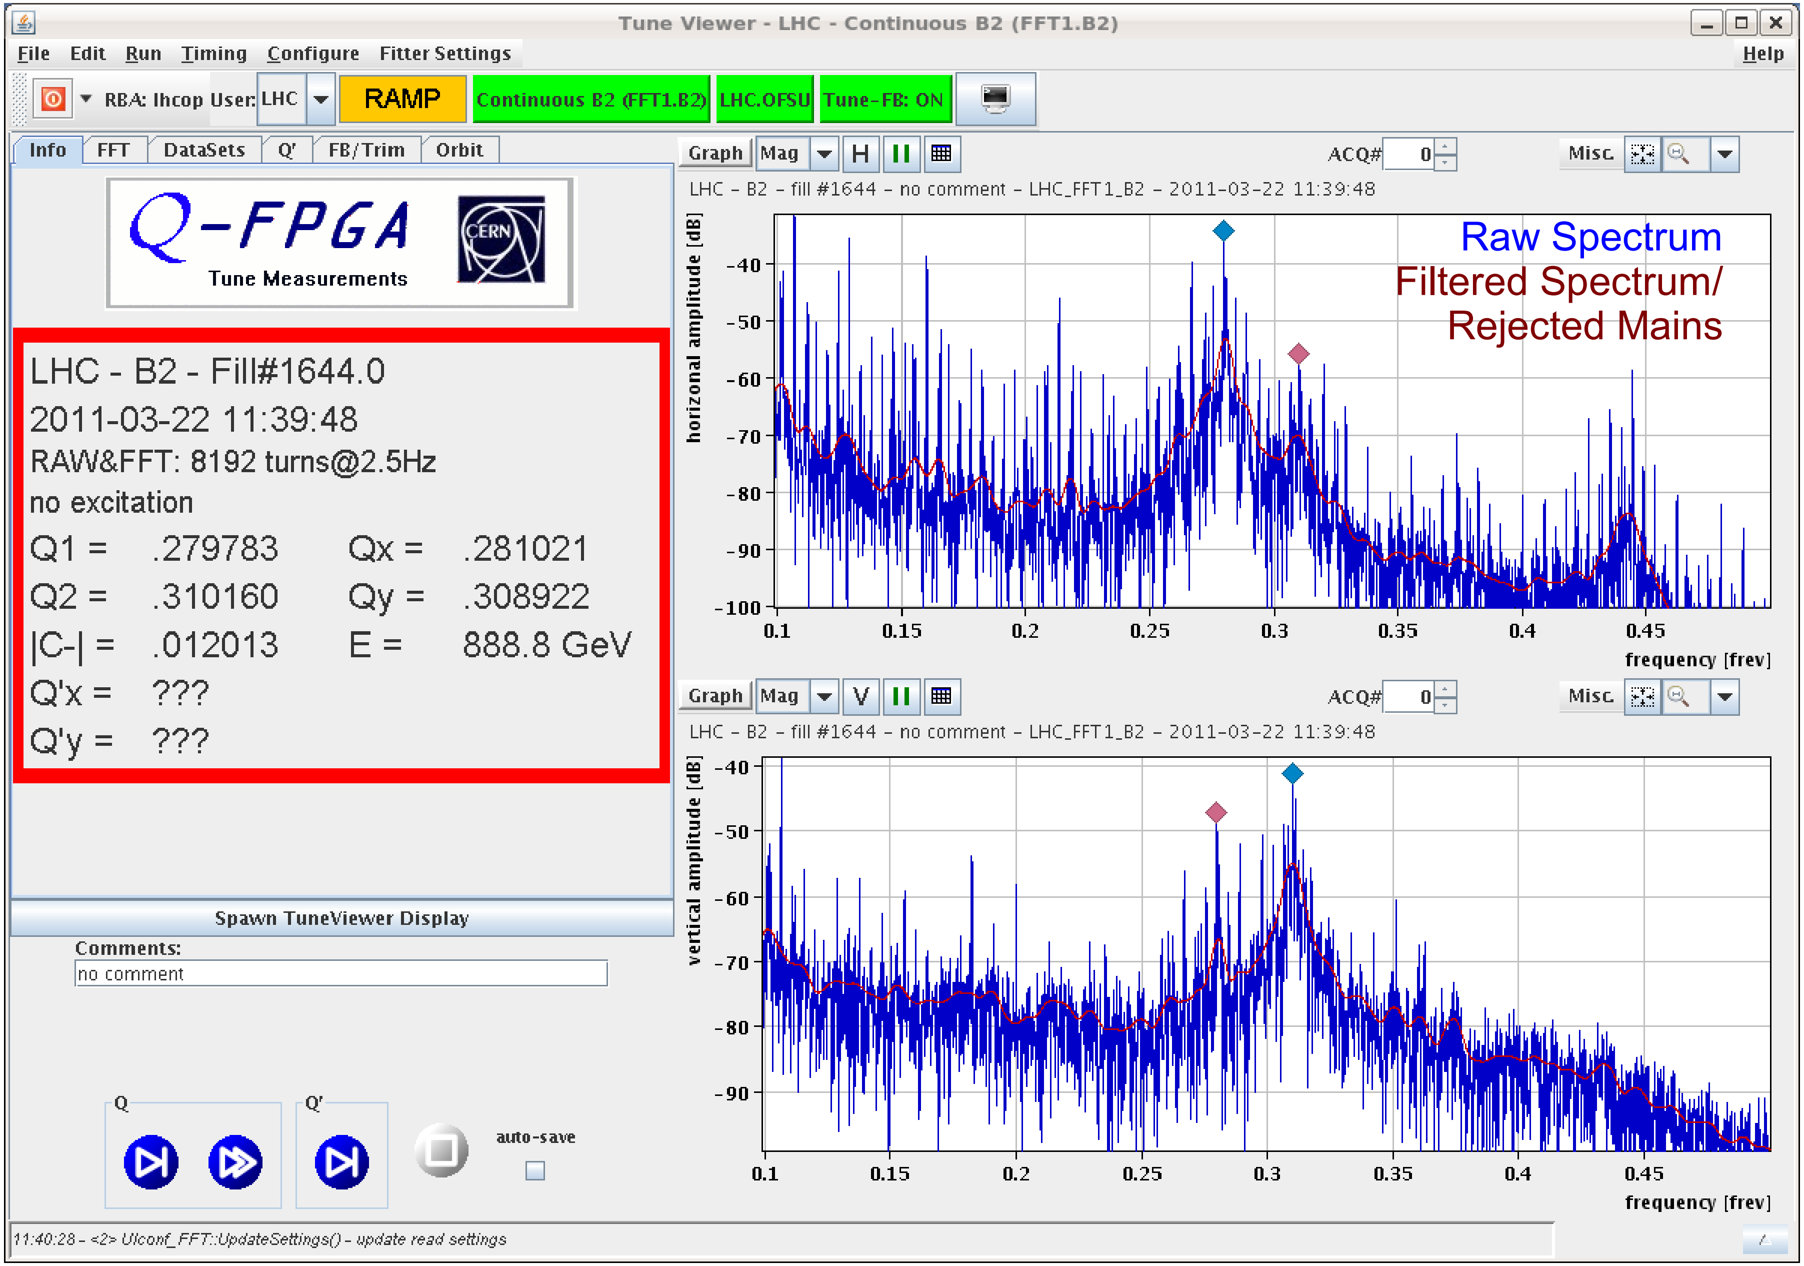

Tune Viewer - LHC - Continuous B2 (FFT1.B2)

LHC - B2 - Fill#1644.0
2011-03-22 11:39:48
RAW&FFT: 8192 turns@2.5Hz
no excitation

Q1 =	.279783	Qx =	.281021
Q2 =	.310160	Qy =	.308922
|C-| =	.012013	E =	888.8 GeV
Q'x =	???		
Q'y =	???		

Comments: no comment

LHC - B2 - fill #1644 - no comment - LHC_FFT1_B2 - 2011-03-22 11:39:48

LHC - B2 - fill #1644 - no comment - LHC_FFT1_B2 - 2011-03-22 11:39:48

11:40:28 - <2> UIconf_FFT::UpdateSettings() - update read settings

BPM Dependencies other than Position

- Measurement drifts ~100 um/h w/o significant temperature changes
  → Orbit-FB may convert these measurement errors into real orbit shift

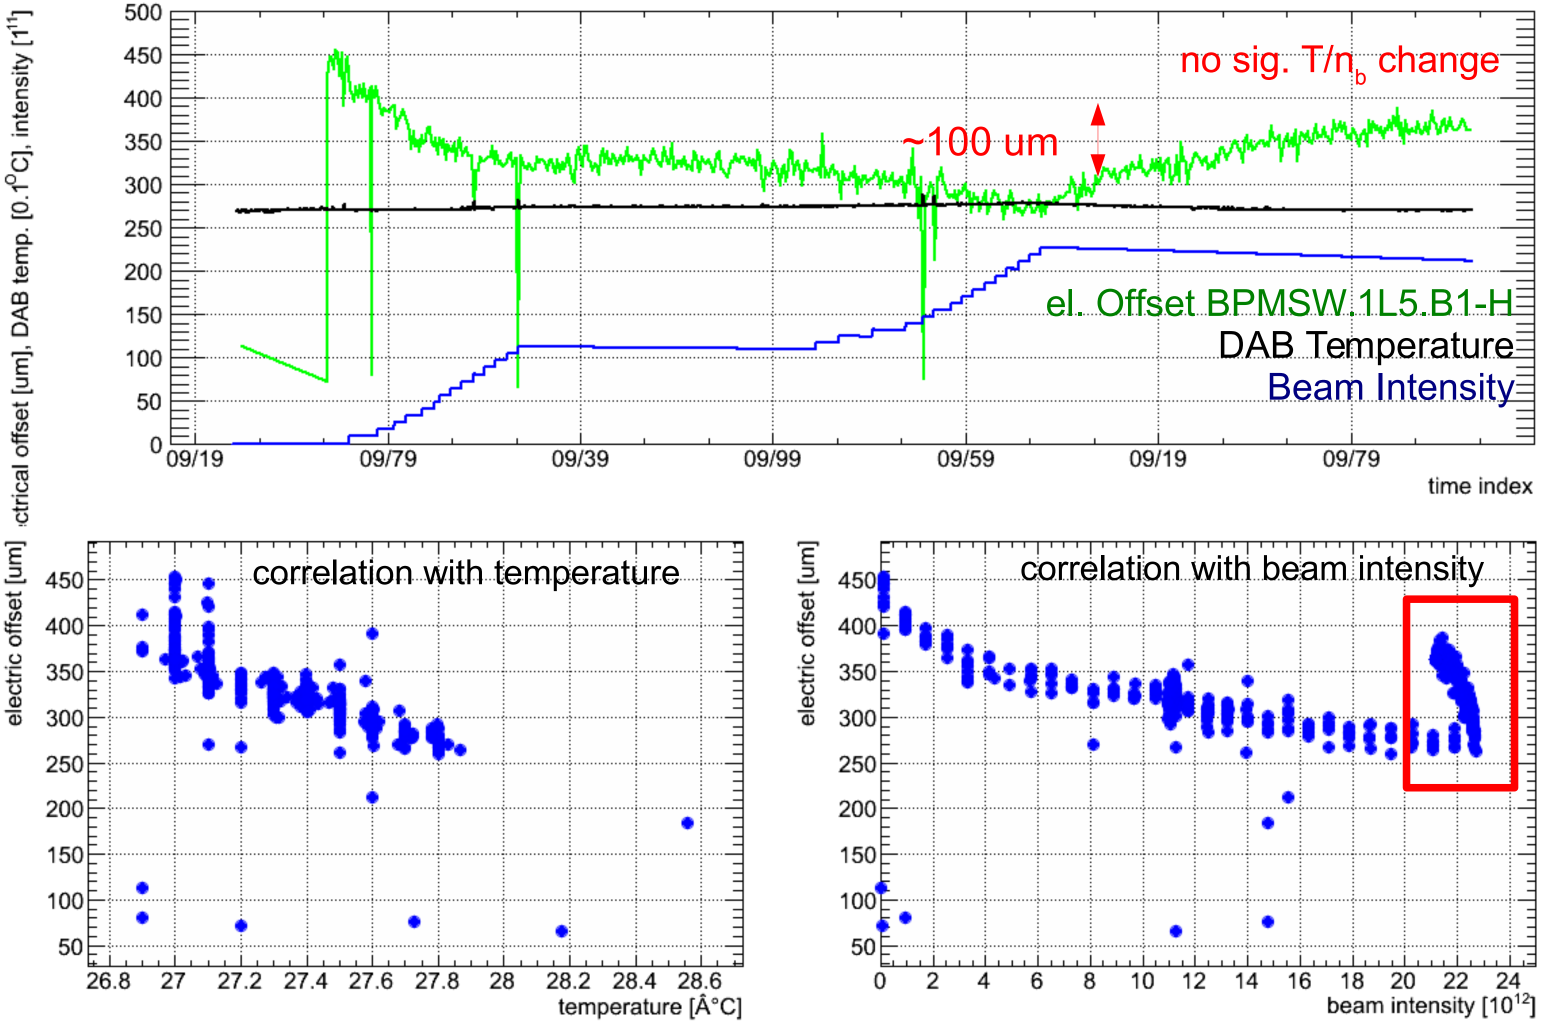

- now mitigated by temperature controlled racks

LHC Feedback Operation

- Trims became de-facto standard to assess the FB and machine performance

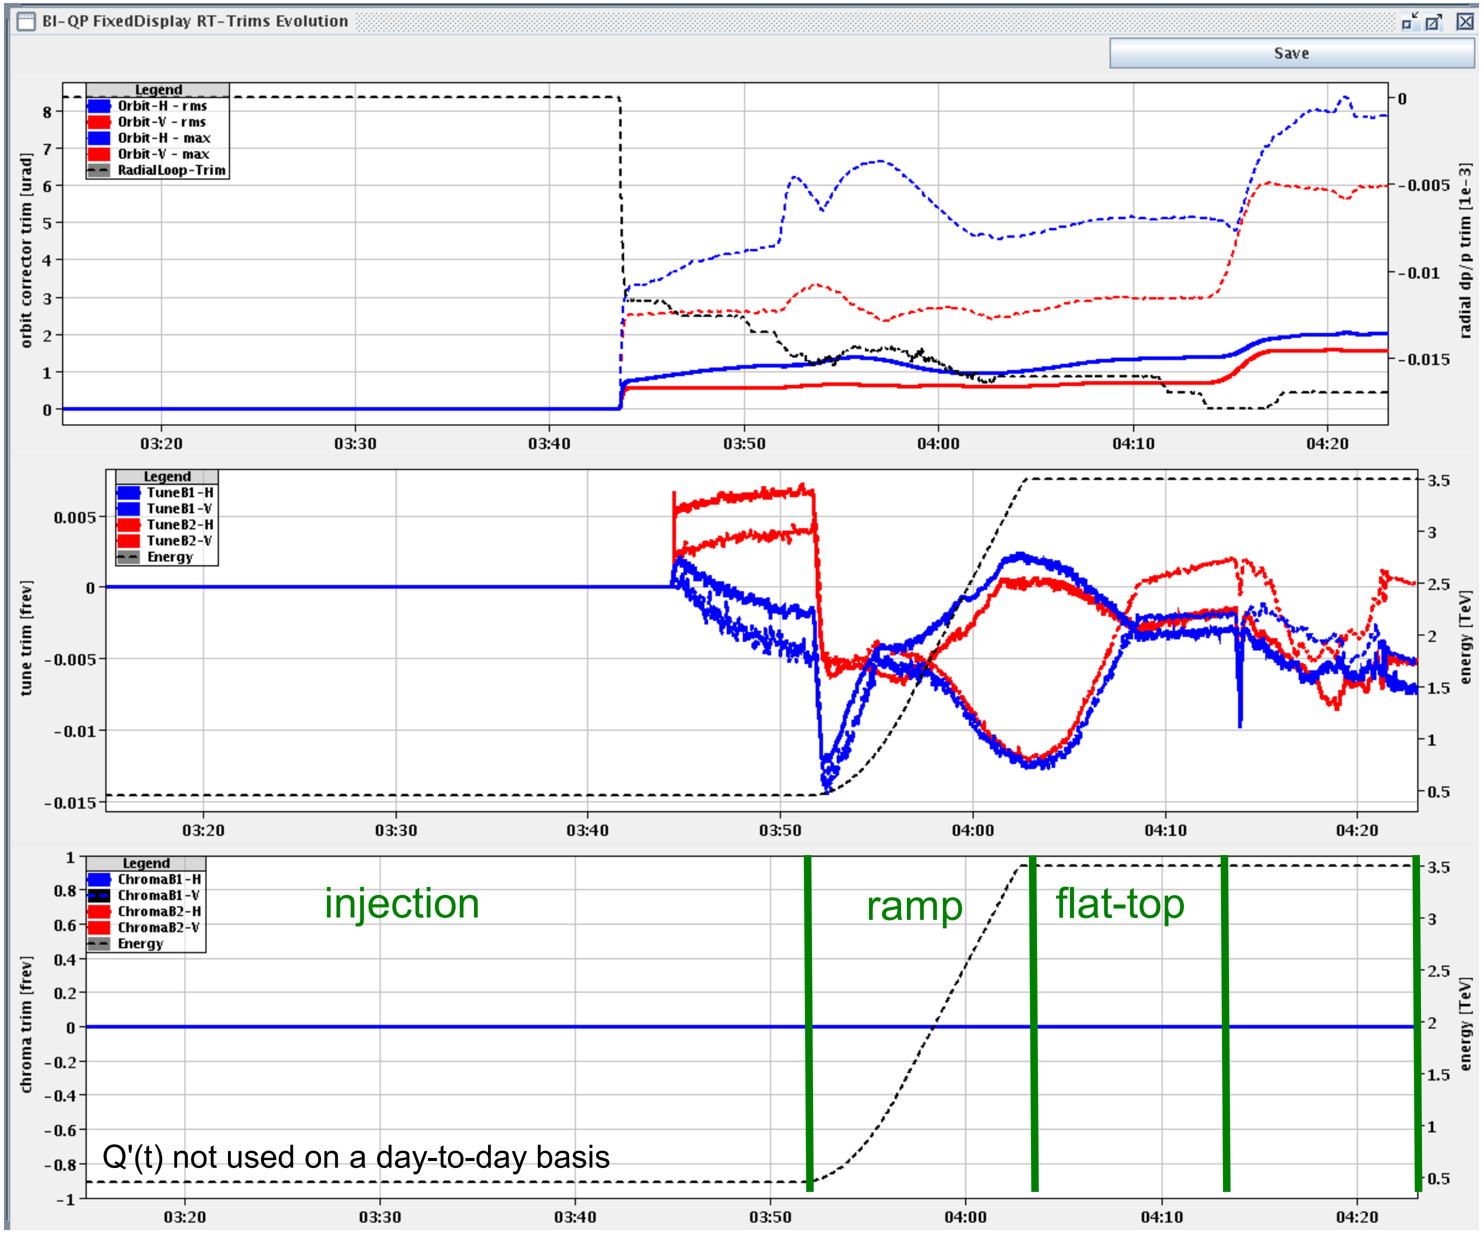

injection

ramp

flat-top

Q'(t) not used on a day-to-day basis

β*-squeeze

Orbit-FB & Radial-Loop Trims (μrad)

Tune-FB trims

Q'(t)-FB trims

Energy (TeV)

LHC Feedback … one of the more visible systems in the CCC – Importance of Controls Integration and SW Usability

LHC - Fill#1574
2011-03-03 19:09:51
Q1 = .309714 Qx = .310523
Q2 = .319568 Qy = .318759
|C-| = .005410 E = 3500.0 GeV
Q'x = +16.2 ± .1
Q'y = +14.0 ± .3

LHC - Fill#1574
2011-03-03 19:09:51
Q1 = .310105 Qx = .310434
Q2 = .320267 Qy = .319938
|C-| = .003598 E = 3500.0 GeV
Q'x = ???
Q'y = +11.9 ± .4

CO - P 1380.120 GeV/c - Fill # 3549 ADJUST - 10/02/13 21-51-44
Mean = -0.004 / RMS = 0.157 / RMS-dp = 0.157 / Dp = -0.0071

CO - P 1380.120 GeV/c - Fill # 3549 ADJUST - 10/02/13 21-51-44
Mean = -0.003 / RMS = 0.104 / Dp = -0.0071

V/c - Fill # 3549 ADJUST - 10/02/13 21-51-50
014 / RMS = 0.139 / Dp = 0.0101

BI-QP Fixed-Display

BI-QP Fixed Display Fill-Number: 1009 Energy: 3566.64 [GeV] RAMPDOWN 2010-03-31 09:40:55

System:	Beam1		Beam2	
Continuous FFT System:	ON	8192 turns@2.0 Hz	ON	8192 turns@2.0 Hz
On-Demand FFT System:	OFF		OFF	
Tune-PLL System:	OFF		OFF	
Tune-FB:	OFF	RT-Trims	OFF	RT-Trims
Chroma-FB:	OFF	RT-Trims	OFF	RT-Trims
Coupling-FB:	OFF	RT-Trims	OFF	RT-Trims

Orbit Feedback - LHC

File Run Configure

RBA: rstein LHC.OFSU TInterlink

Optics - Q/Q' | Energy | RF Mole | Settings | BPM | Optics - Orbit

Settings
OrbitFB Server: LHC.OFSU
Averaging: AverageNone
OrbitFB State: OFF
Radial Loop State: OFF
RadialMod State: OFF
Reset OrbitFB: Engage
TuneFB State: OFF OFF
Chroma-FB State: OFF OFF
Coupling-FB: OFF OFF
Reset Q/Q'-FB: Engage
Sensitivity: Manual
B1-HIGH B2-HIGH
BCT System: NONE
Spawn: Single OrbitViewer
Paired OrbitViewer

Comments: no comment

08:49:31 - <0> OrbitFeedbackSubscription::valueReceived():position vector dimension changed from 100 to 73

Orbit-FB & Radial-Loop Trims (µrad)

Tune-FB trims

Q'(t)-FB trims

Energy (TeV)

β*-squeeze

Earth Tides dominating Orbit Stability during Physics:

- Known effect from LEP → changes the machine circumference/energy

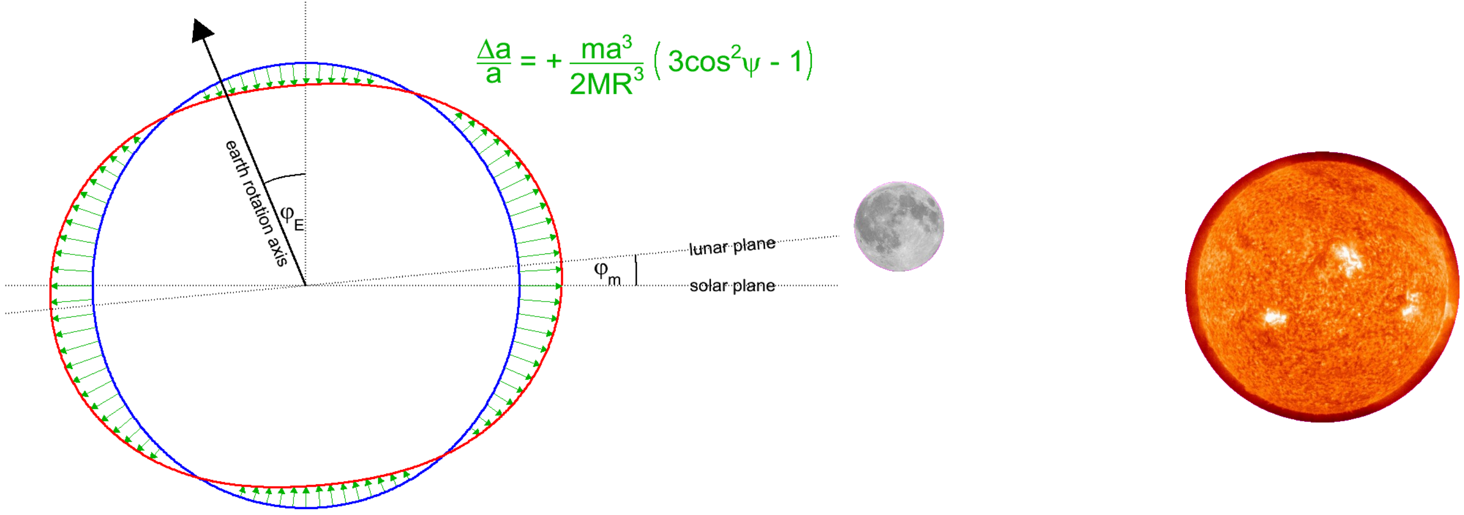

$$\frac{\Delta a}{a} = +\frac{ma^3}{2MR^3}\left(3\cos^2\psi - 1\right)$$

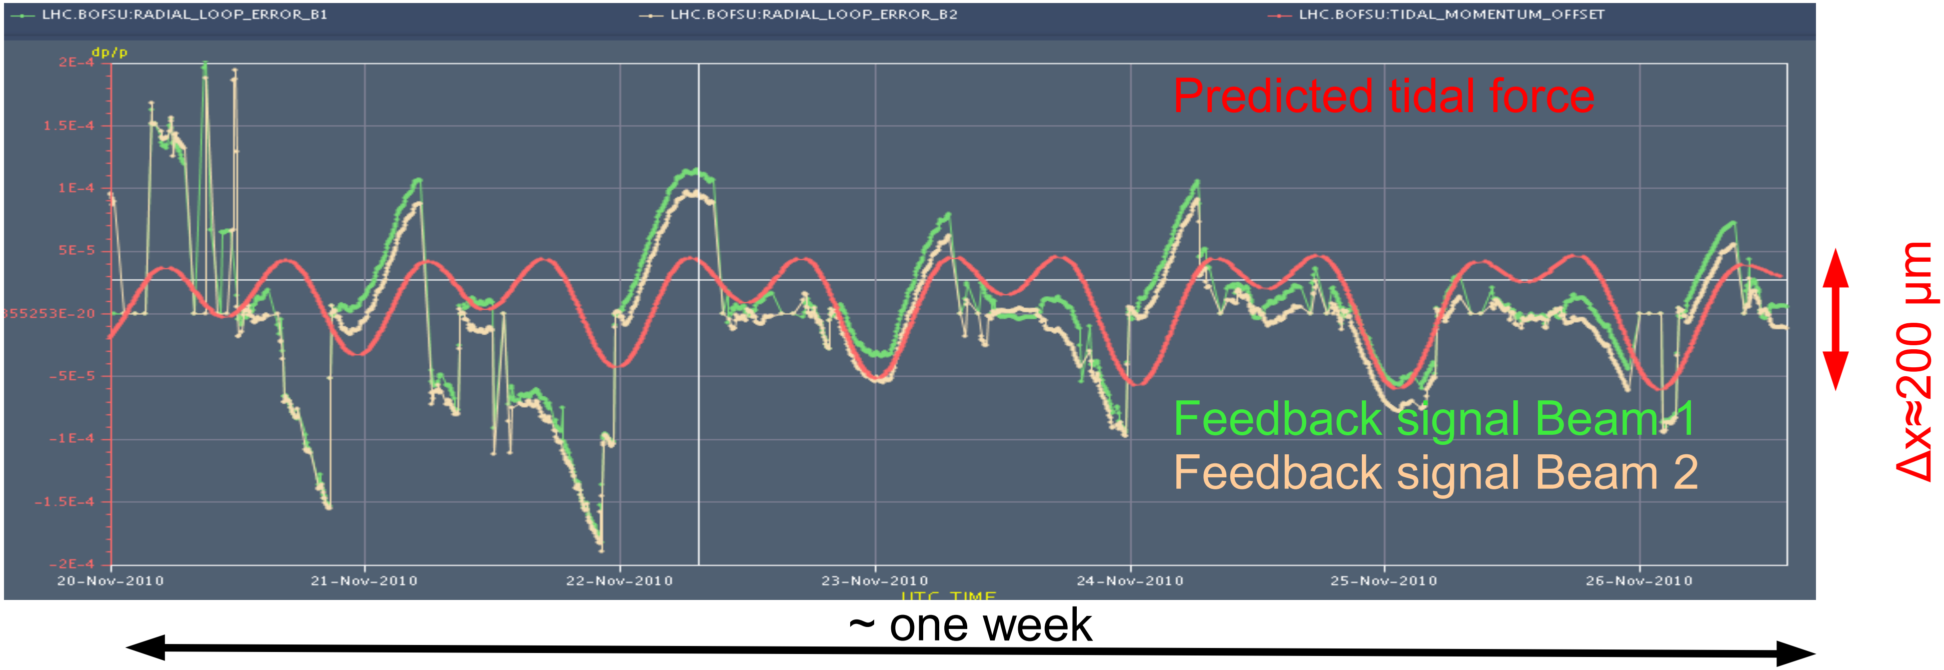

  – Testimony to LHC alignment and beam stability!

Summary

- Generally, feedback performed their designed job.

- Pushing LHC machine parameter envelope also implied increased performance constraints on Feedback operation (notably orbit stability during squeeze)
  → will improve FB sub-systems to keep up with LHC progress post-LS1

- Main paradigms:
  - Split 'space/time' domain design to leverage different user background experiences
  - Nested-loop design to minimise inter-loop cross-talk
  - Central simple input-processing-output feedback controller
  - Managed by service unit (OFSU, settings management, data proxy)
  - LHC Technical-Network as communication backbone
  - 'Firm real-time' constraints using Real-Time capable Linux

Reserve Slides

-

Expected Dynamic Perturbations vs. Requirements – or: Design Assumption vs. Operational Reality

- From Decay/Snap-back expected dynamic perturbations

	Orbit [σ]	Tune [$0.5\cdot f_{rev}$]	Chroma. [units]	Energy [Δp/p]	Coupling [c_]
Exp. Perturbations ('06):	~ 0.5	0.014	~ 70	± 1.5e-4	~0.01
Nom. Requirements:	± 0.15	±0.001	2 ± 1	± 1e-4	« 0.01
Achieved Stability ('13):	~ 0.1	~ 0.001	± 2 (7)	~1e-5	< 0.003

- Initial assumptions and plans (2006-2009):
  - Chromaticity considered as most critical parameter
  - FB Priority list: Chromaticity → Coupling/Tune → Orbit → Energy

- What turned out to be needed operationally
  - 2009 → 2011: **Tune** → Orbit & Energy/Radial-Loop → Q'(t) →…→ C^-
    - impressive Q'(t), C^- and beta-beat stability/reproducibility
  - In 2012: **Orbit** & Tune (snap-back, instabilities)
    - Higher energy & smaller-β^* → much tighter collimator settings → convert smallest orbit deviations into losses/dumps

Why the notion/split between 'space' and 'time' domain?

- Separates specific accelerator physics from specific control theory
  - can test the two domains independently
  - ample expertise/resources in one but not the other domain
    → KISS principle: keep it simple and safe

- Multiple-Input-Multiple-Output (MIMO) in space-domain
  - Can modify correction algorithm without having to worry about whether overall loop remains stable
  - Maintains physical meaning of the individual control variables
  - Basically relying on inversion of response matrices → SVD

- Quasi-Single-Input-Single Output (SISO) in time-domain
  - Similar control problem/laws as e.g. for power converters
  - Time-domain controller identical for orbit, energy, Q/Q' vs. integrated/more complex 'Kalman' or 'Youla-Kucera-Klein'-based method

Time-Domain: Optimal Controller Design
Youla's affine parameterisation II/II

Using Youla's method: “design closed loop in a open loop style”:

Youla showed[1] that all stable closed loop controllers D(s) can be written as:

$$\boxed{D(s)=\frac{Q(s)}{1-Q(s)G(s)}} \qquad (1)$$

Example: first order system

$$G(s)=\frac{K_0}{\tau s+1} \qquad \text{with } \tau \text{ being the circuit time constant} \qquad (2)$$

Using for example the following ansatz:

$$Q(s)=F_Q(s)G^i(s)=\frac{1}{\alpha s+1}\cdot\frac{\tau s+1}{K_0} \qquad (3)$$

Response/optimality can be directly deduced by construction of $F_Q(s)$

$G^i(s)$, pseudo-inverse of the nominal plant G(s)

(1)+(2)+(3) yields the following PI controller:

$$\rightarrow T_0(s)=\frac{1}{\alpha s+1}$$

$$\boxed{D(s)=K_P+K_i\frac{1}{s} \qquad with \qquad K_p=K_0\frac{\tau}{\alpha} \wedge K_i=K_0\frac{1}{\alpha}}$$

---

[1] D. C. Youla et al., *“Modern Wiener-Hopf Design of Optimal Controllers”*, IEEE Trans. on Automatic Control,1976, vol. 21-1,pp. 3-13 & 319-338

LHC Feedback Operation – Example III/IV

- Tune-FB driving and accelerating early commissioning in 2009-2011
  - Tunes kept stable to better than 10^{-3} for most part of the ramp and squeeze

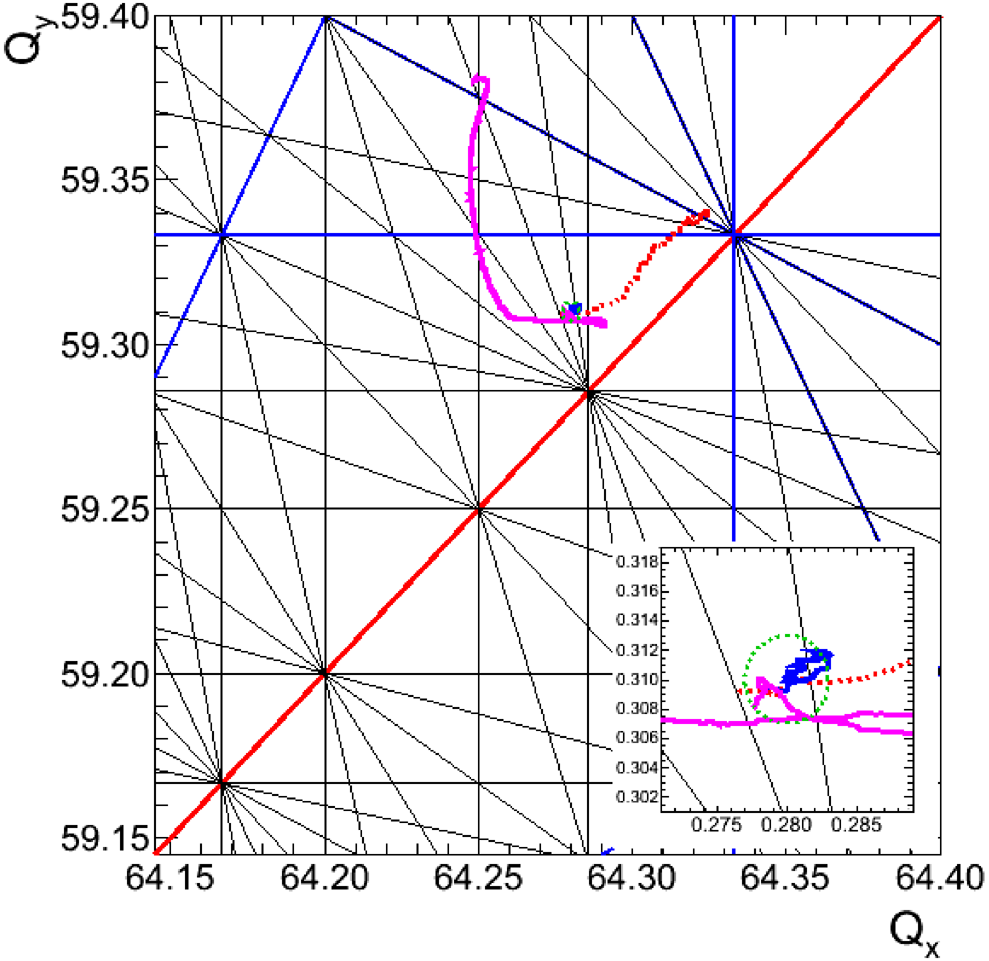

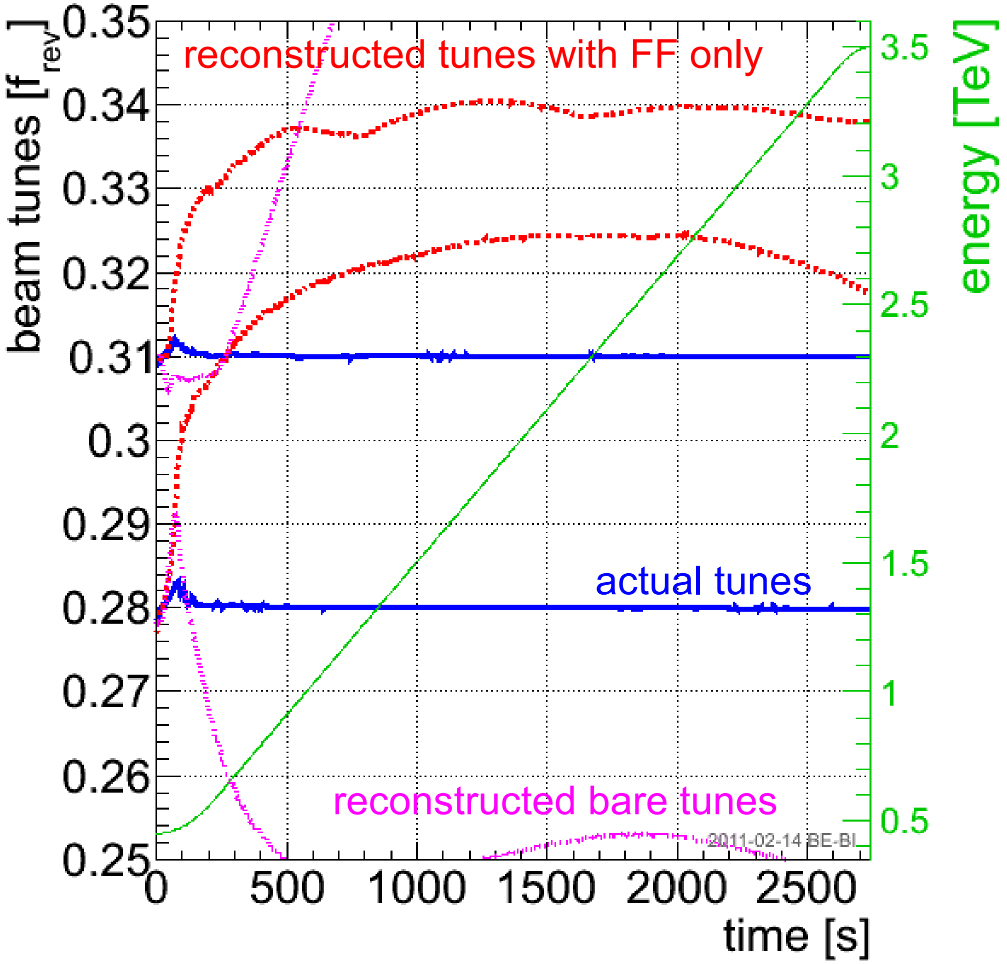

Feed-Forward Back-Bone – Field-Description for LHC (FiDeL) Example: b_3 – Compensation – Static Part

Based on magnet measurements:

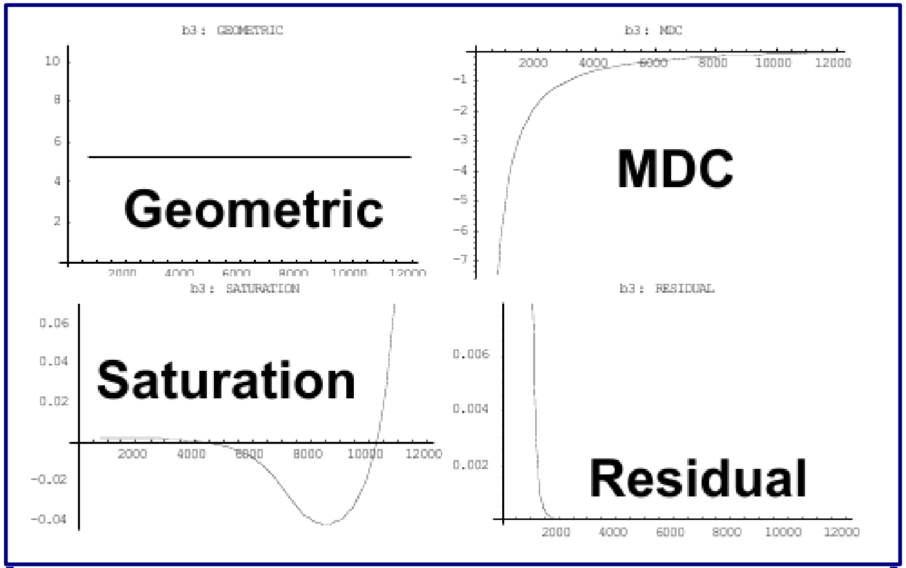

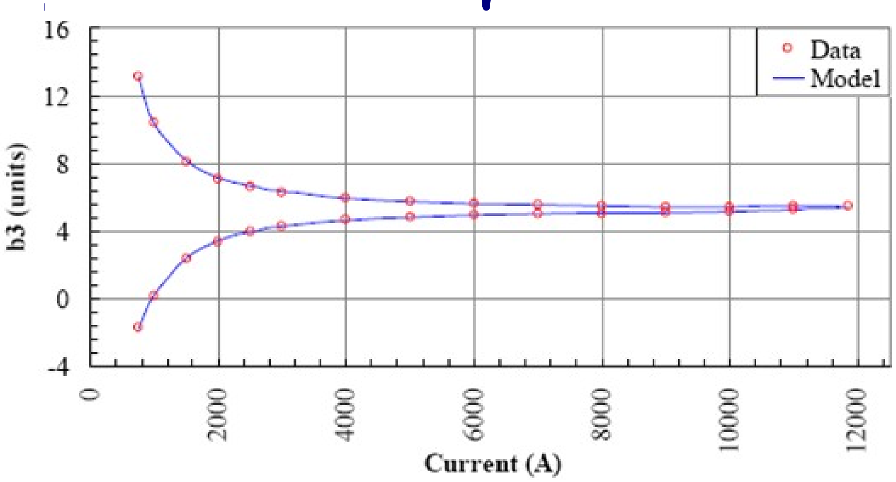

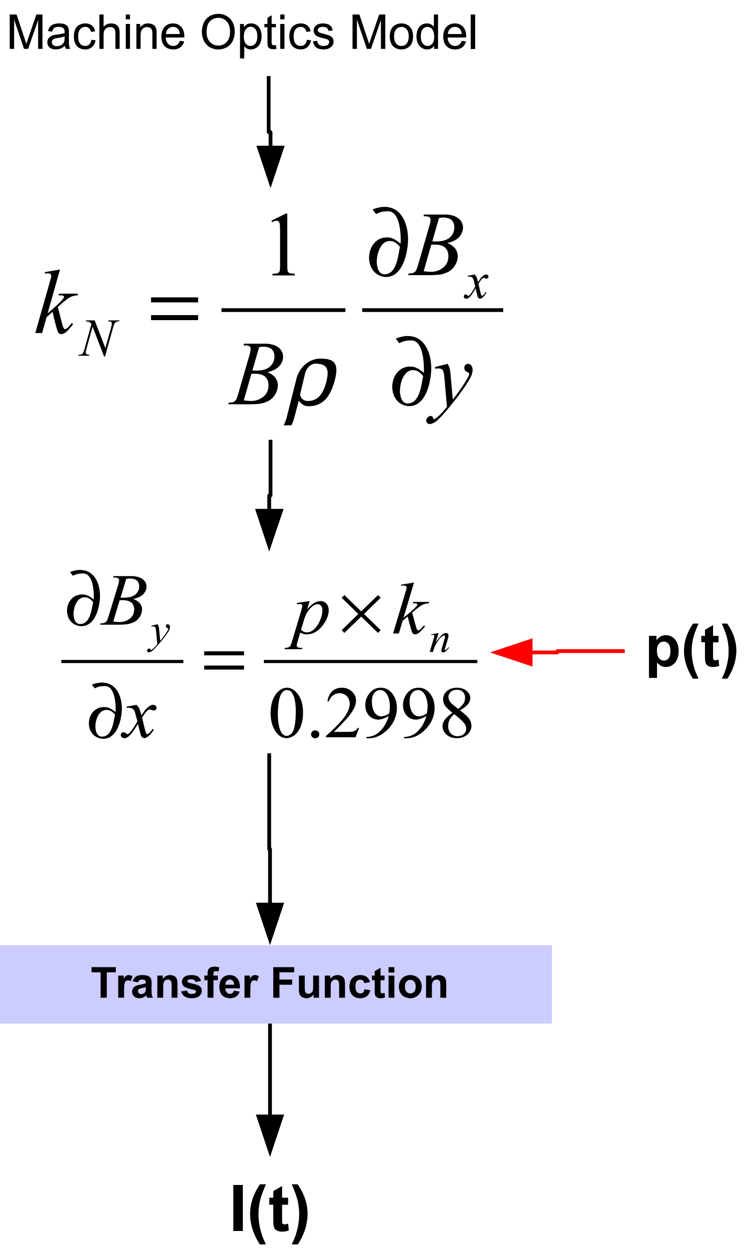

Machine Optics Model

$$k_N = \frac{1}{B\rho}\frac{\partial B_x}{\partial y}$$

$$\frac{\partial B_y}{\partial x} = \frac{p \times k_n}{0.2998}$$

p(t)

Transfer Function

I(t)

Courtesy M. Lamont

Control Paradigms II/III

- Machine imperfections cause steady-state offset ε_{ss} and scale error ε_{scale}:

$$\Delta x(s)=R_i(s)\cdot\delta_i \;\rightarrow\; \Delta x(s)=R_i(s)\cdot(\epsilon_{ss}+(1+\epsilon_{scale})\cdot\delta_i)$$

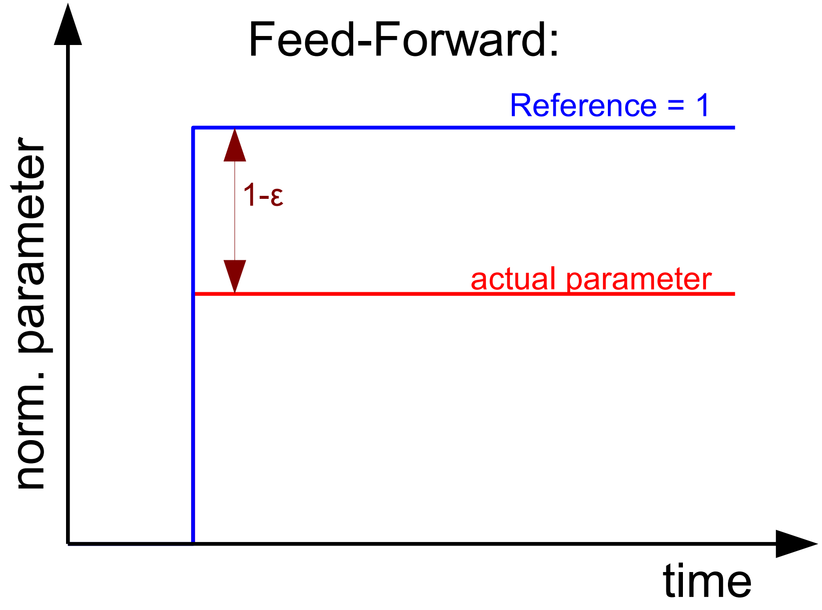

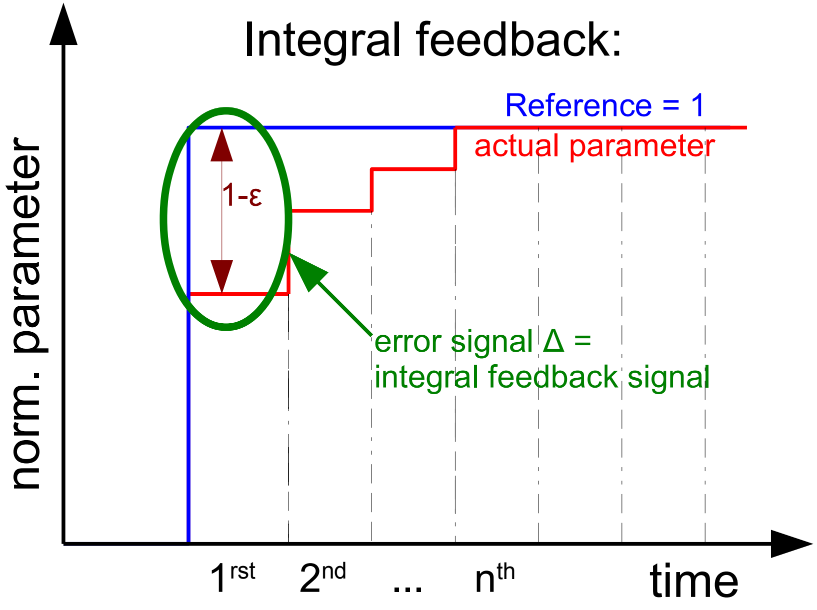

- Uncertainties and scale error of beam response function affects convergence speed (= feedback bandwidth) rather than achievable stability

- Feed-forward of Q'(t)-Feedback signal for next fill turned out to be sufficient!
  - enforced by strict pre-cycling following physics, access or circuits 'off '...

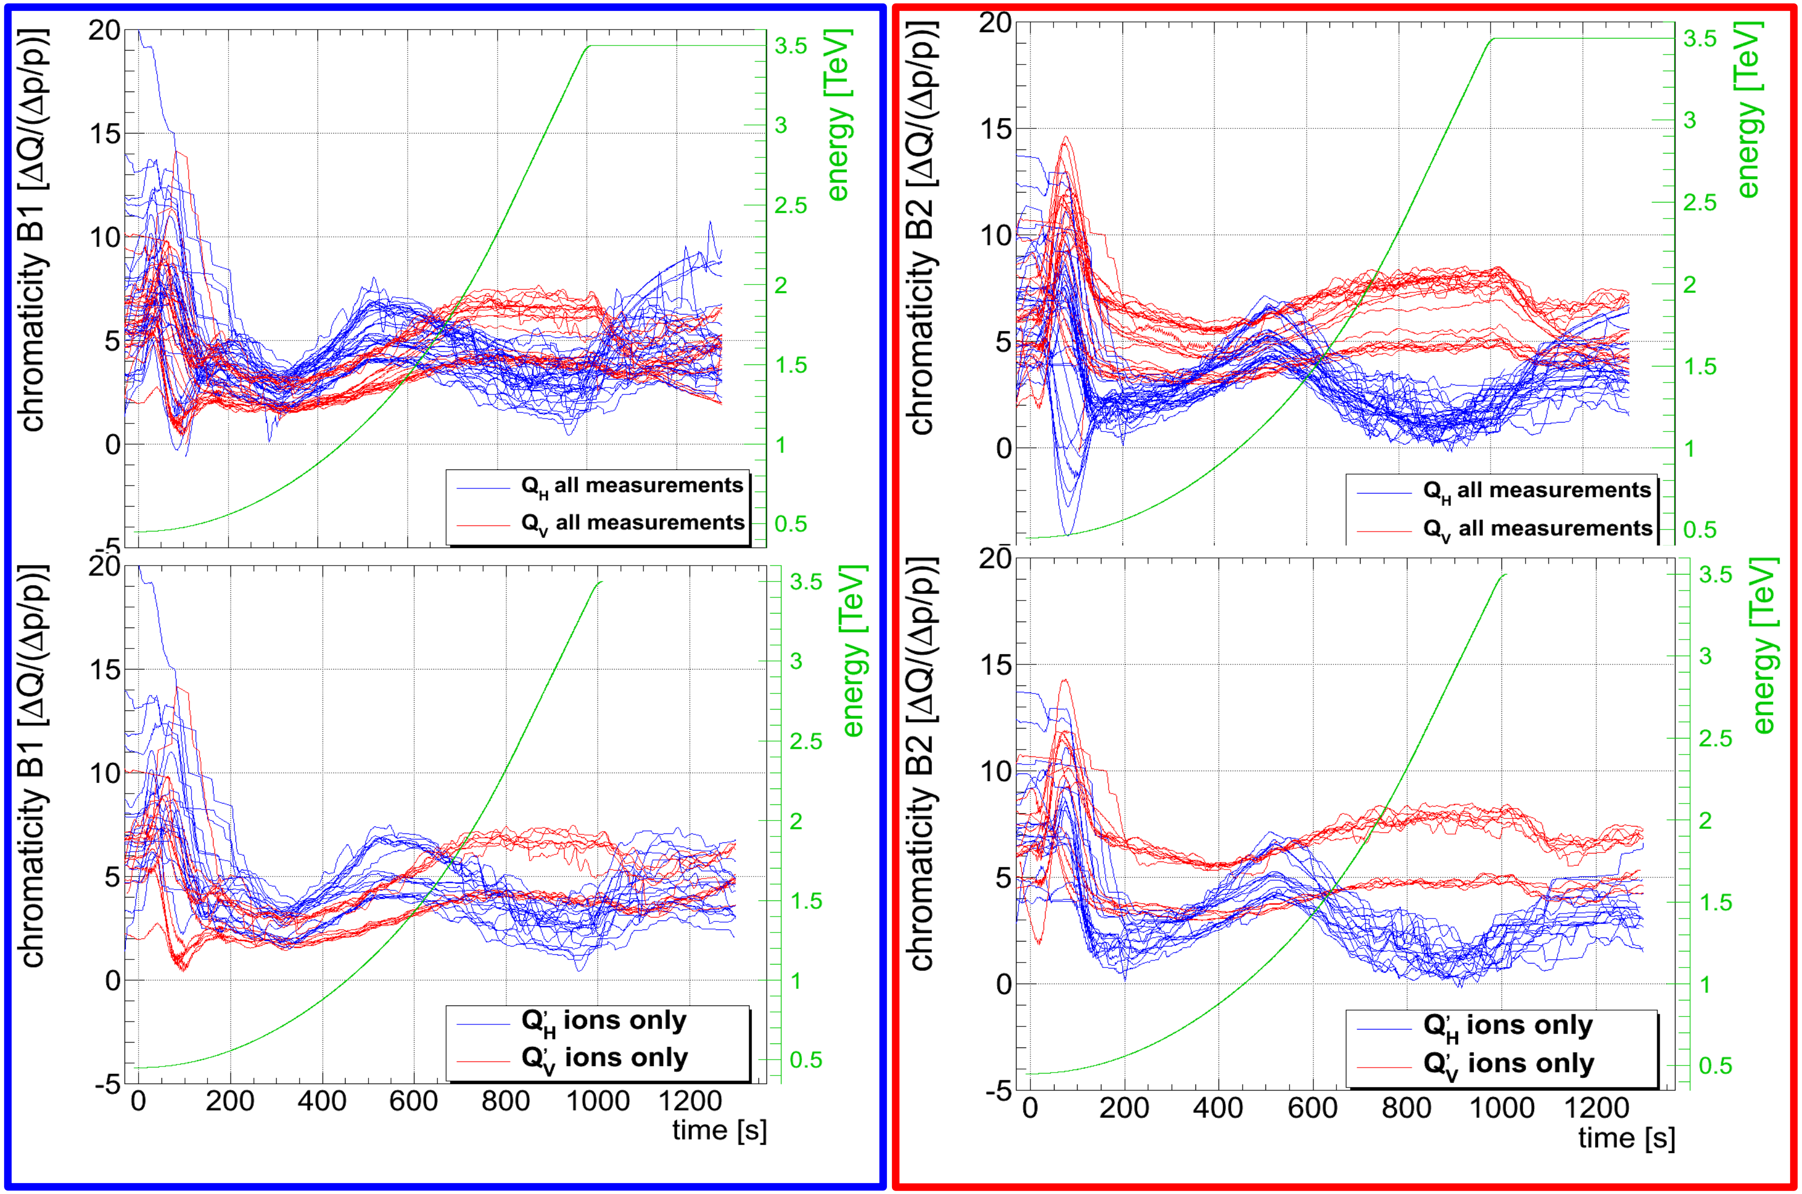

Orbit Feedback Bandwidth vs. Actual Orbit Perturbations

- Closed-loop bandwidth and phase margin depend on excitation amplitude:
  - + non-linear phase once rate-limiter kicks in...

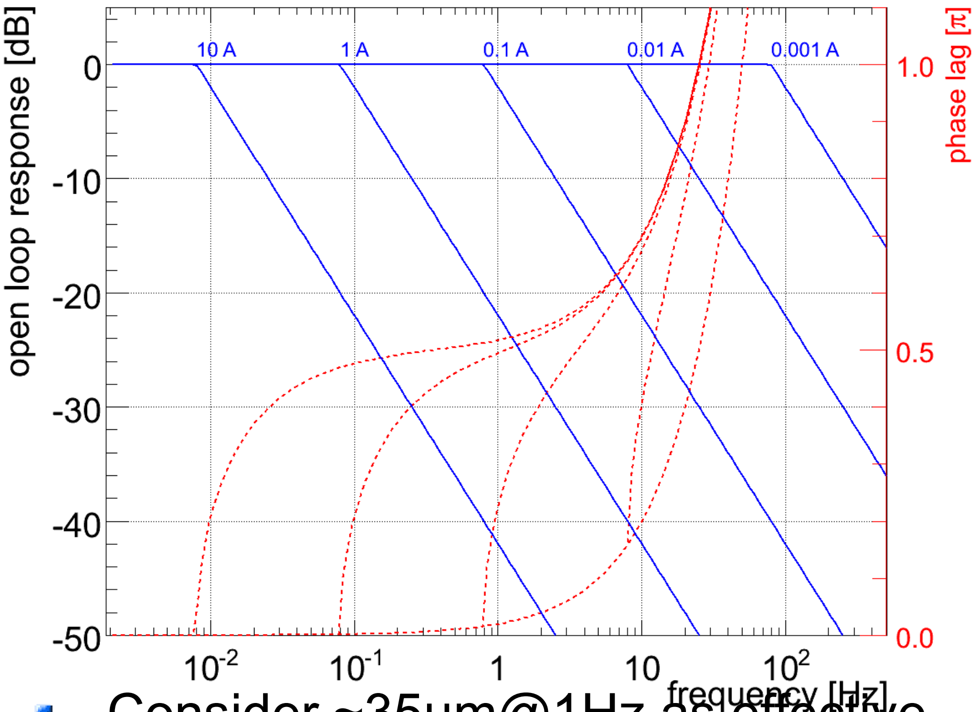

- Consider ~35µm@1Hz as effective bandwidth @4TeV (assuming 3C bump)
- Many latencies become a non-issue 0.1 Hz

ΔI=0.1A ↔ Δx≈32 µm@β=180m

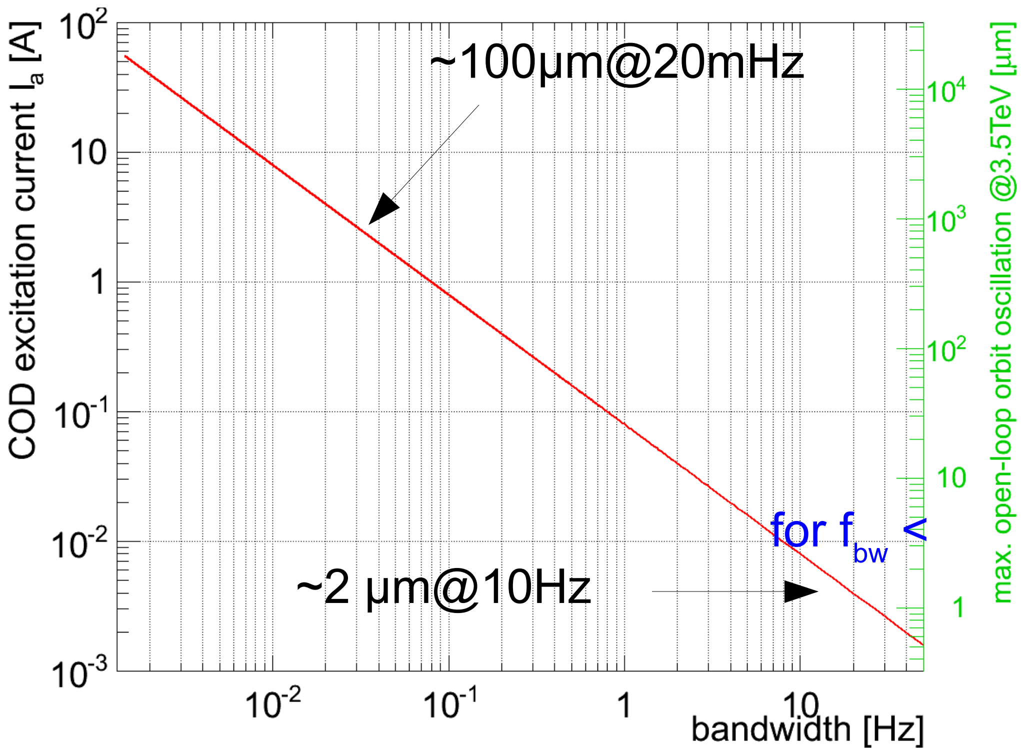
~100µm@20mHz

~2 µm@10Hz

for f_{bw} <

Orbit Stability during β*-Squeeze

- Losses and orbit movement at H-TCP.C6R7.B2 well correlated

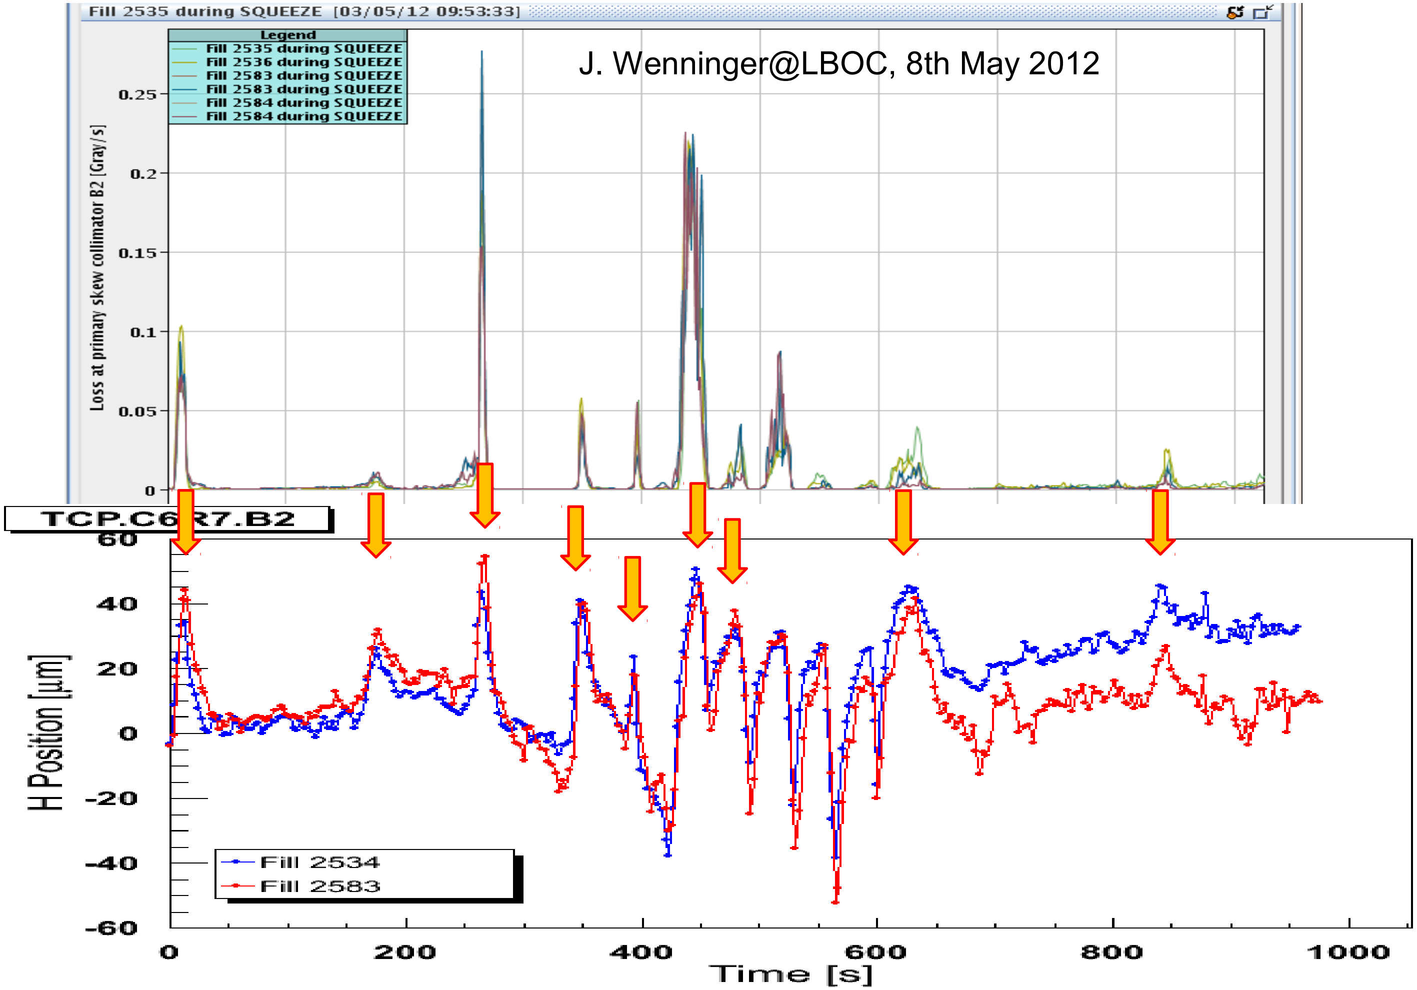

J. Wenninger@LBOC, 8th May 2012

- Maximum drift rates of 40 um/s → (close to) limit of Orbit-FB at 4 TeV
  - Underpinned by FB instability observation for 5x bandwidth increase
- At this speed, OFC needs to operate with correct optics

Correction during Squeeze with imperfect Optics

- Bandwidth modifier w.r.t. eigenvalue index (<1 more stable, >1 diminishes stability margin)

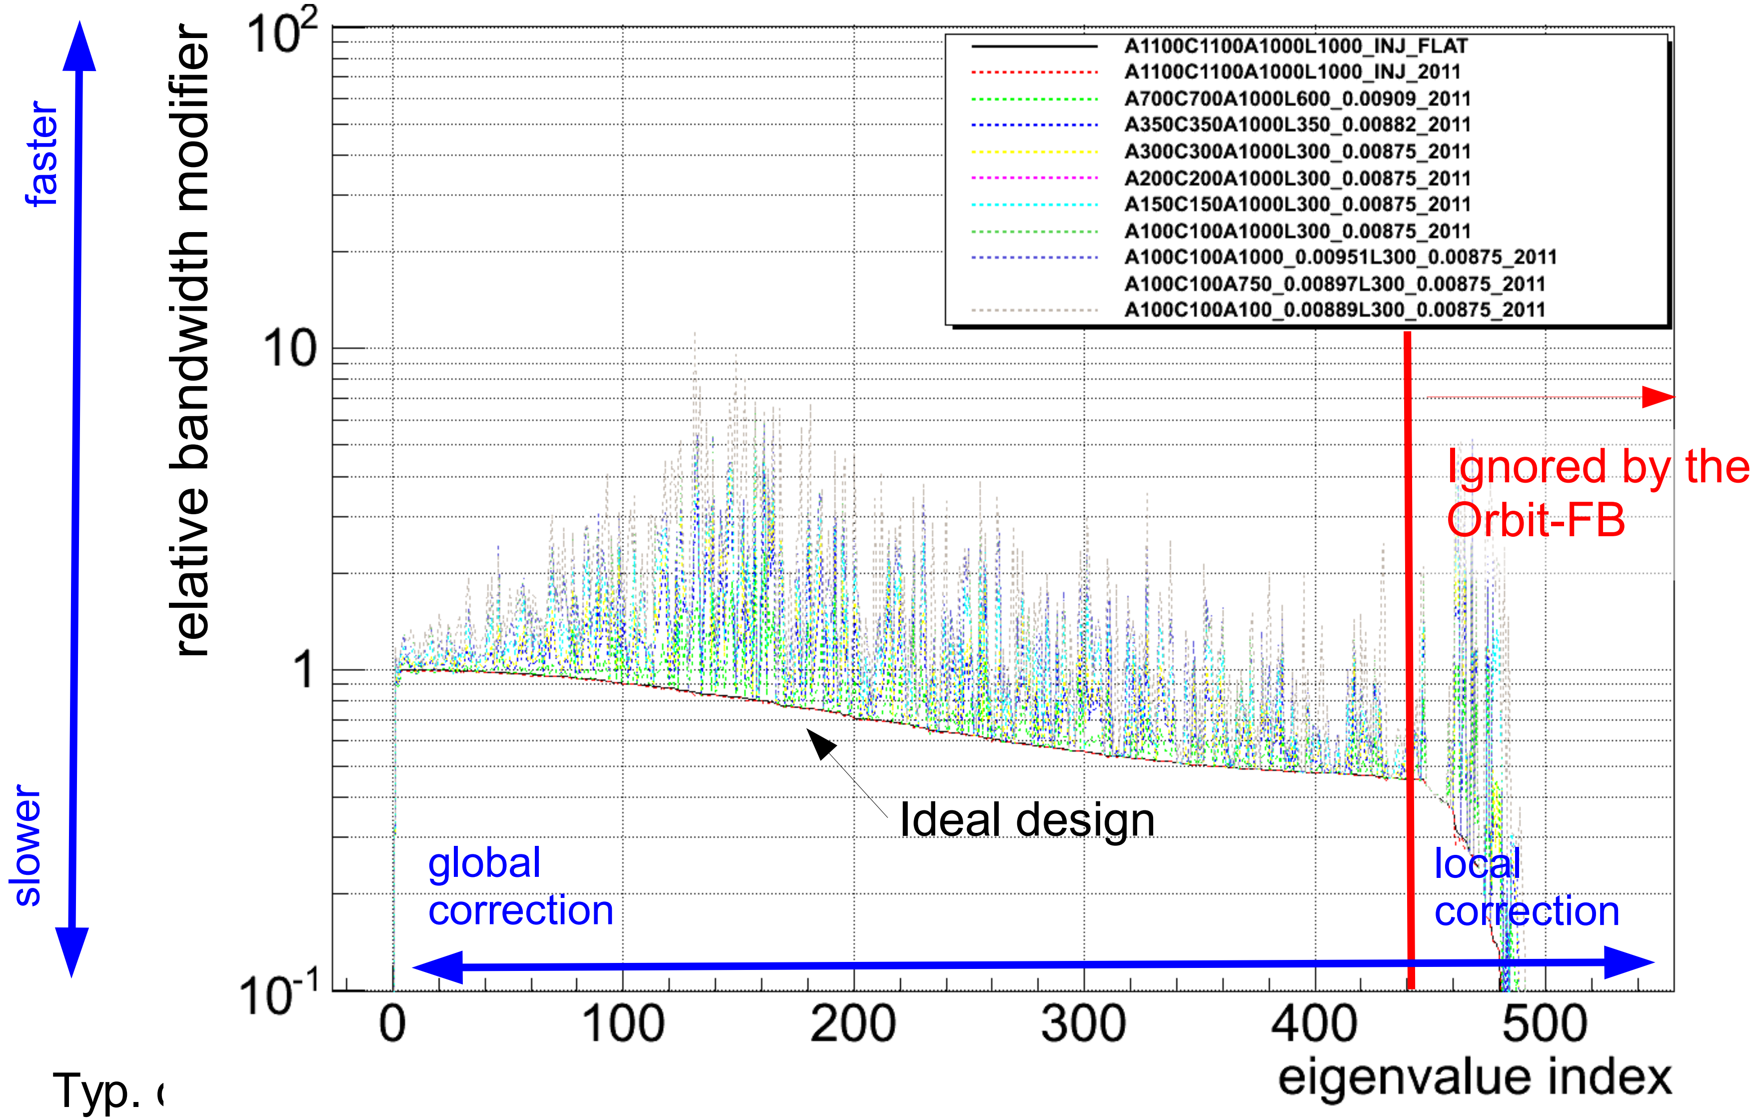

- Typ. (

Tune Evolution during Physics

Quirky side effect:

- Machine circumference changes are propagated via Q' also to the tune

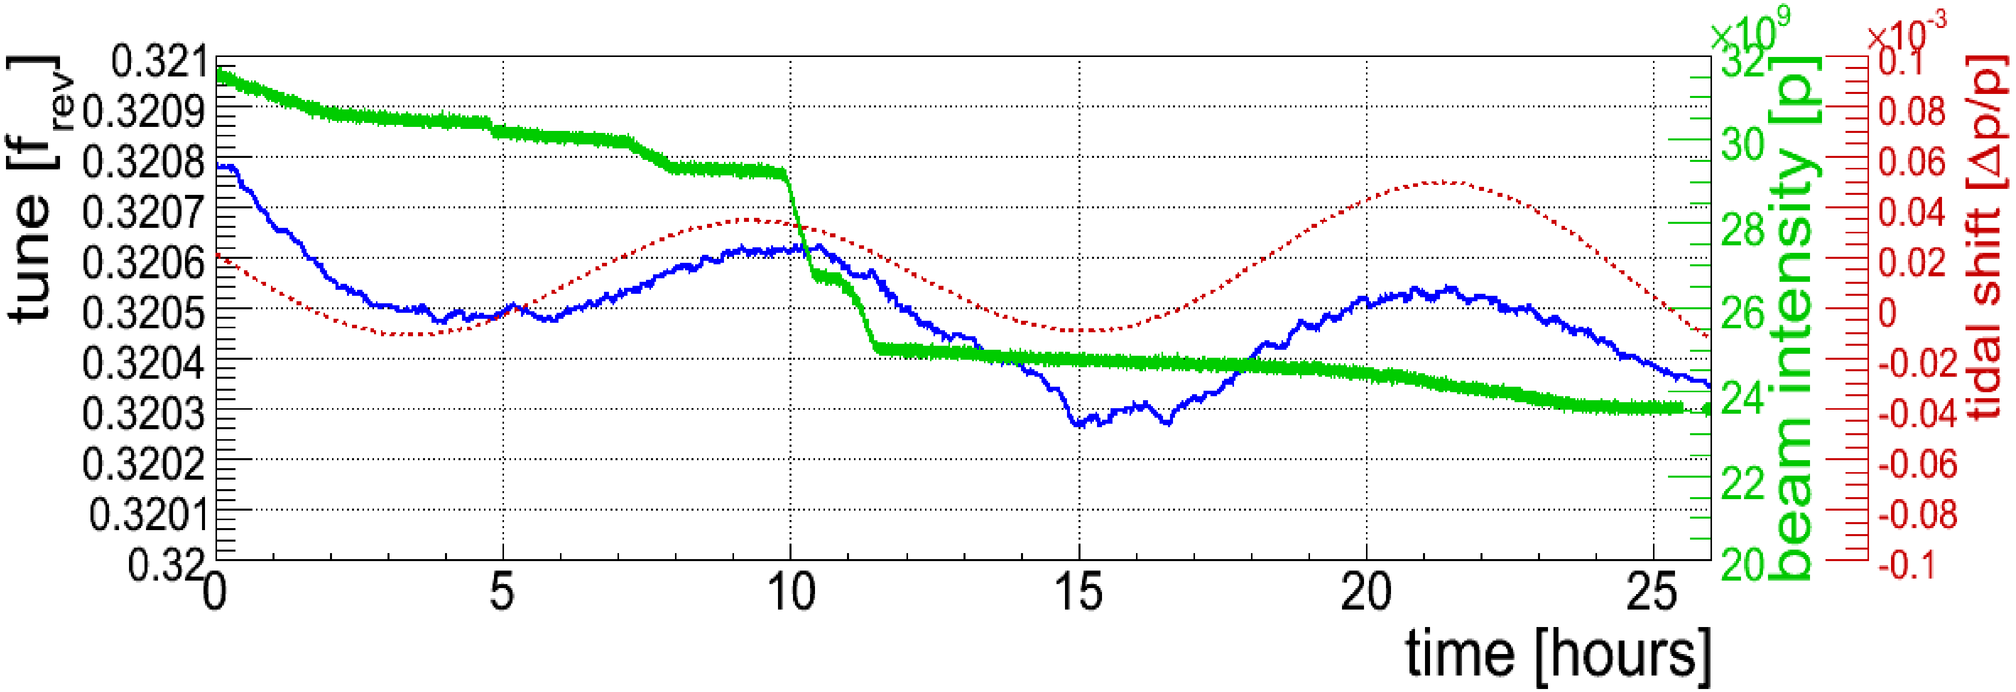

- Probably the slowest high-precision Q' measurement in the World
  - Short-Term Tune-Stability of ~10^{-6}!
- However, stability during nominal physics operation is typically driven by impedance and beam-beam related effects.

Typical Q/Q'(t) Control Room View
2010 Statistics: Out of 191 Ramps...

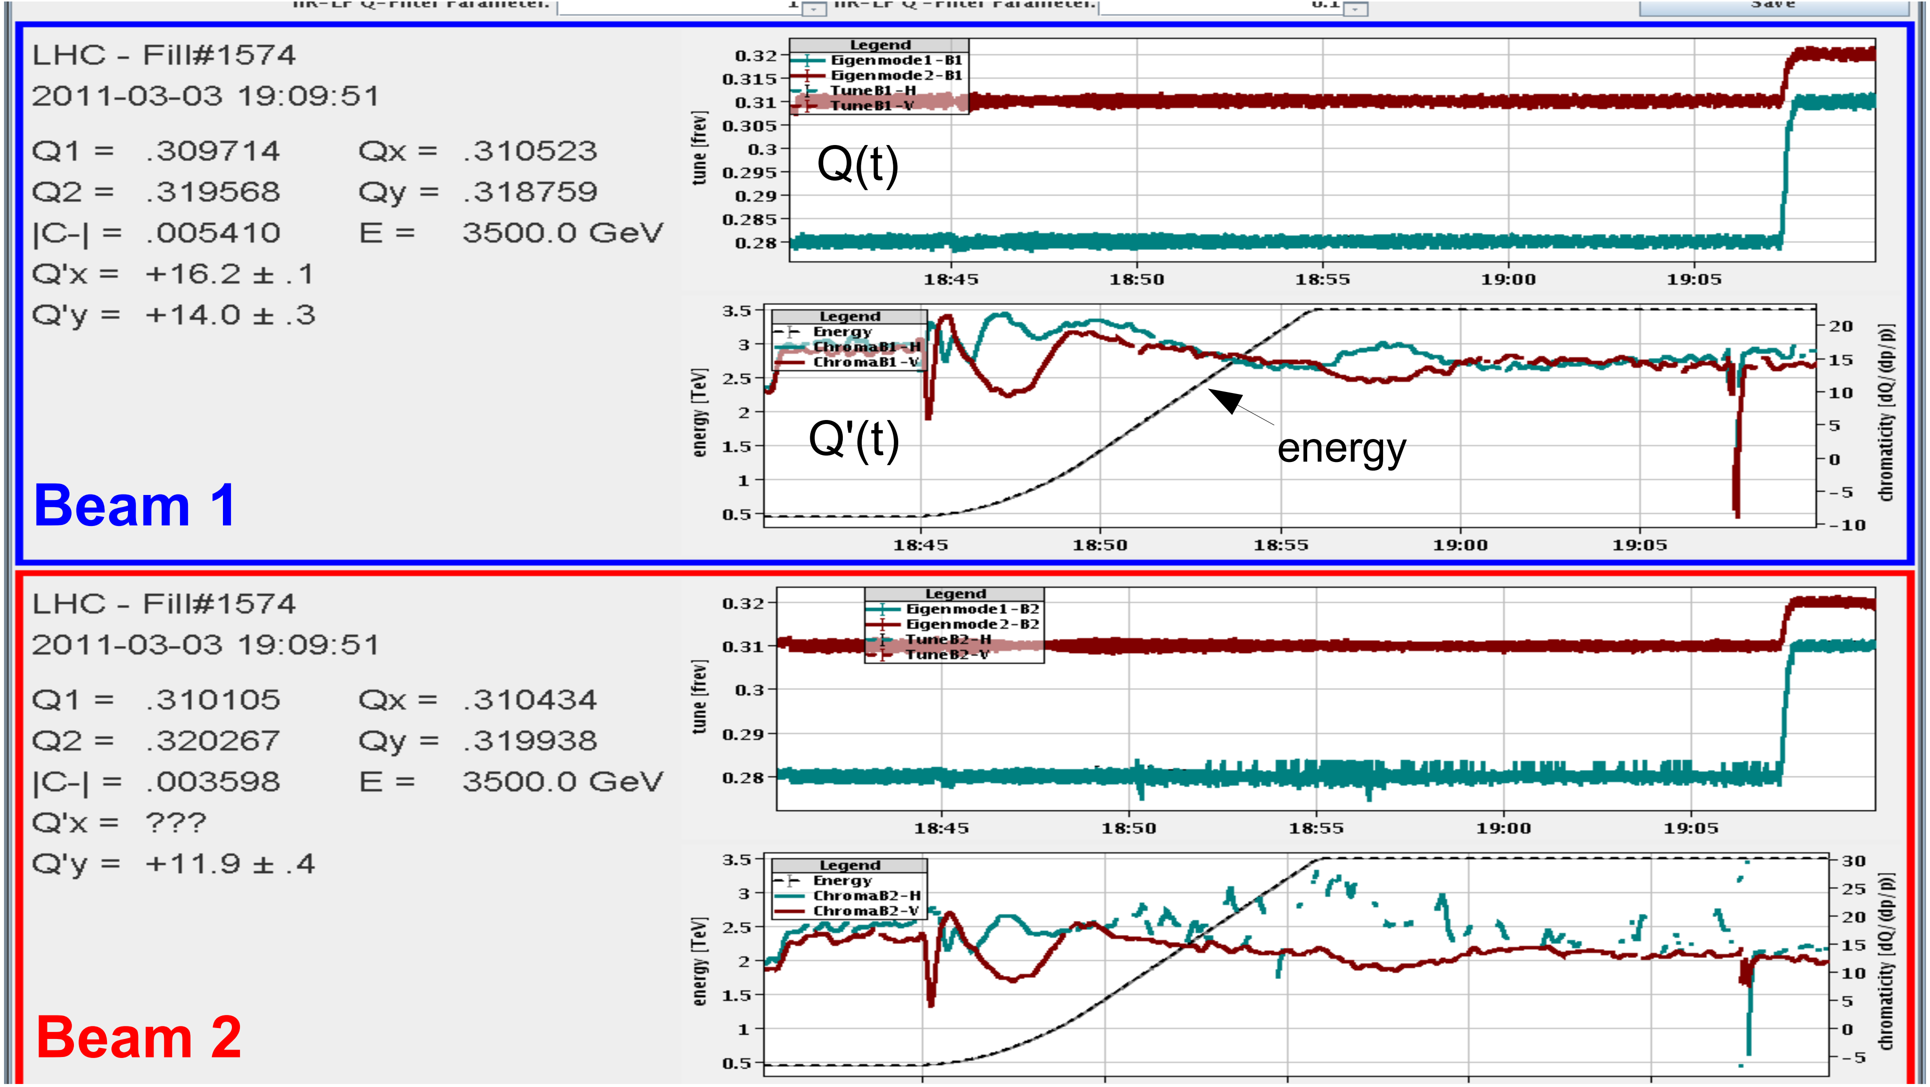

LHC - Fill#1574
2011-03-03 19:09:51

Q1 = .309714 Qx = .310523
Q2 = .319568 Qy = .318759
|C-| = .005410 E = 3500.0 GeV
Q'x = +16.2 ± .1
Q'y = +14.0 ± .3

Beam 1

LHC - Fill#1574
2011-03-03 19:09:51

Q1 = .310105 Qx = .310434
Q2 = .320267 Qy = .319938
|C-| = .003598 E = 3500.0 GeV
Q'x = ???
Q'y = +11.9 ± .4

Beam 2

- … 155 ramps with > 99% transmission, 178 ramps with > 97% transmission
- … only 12 ramps lost with beam (6 with Tune-FB during initial 3.5 TeV comm.)
- … "if without FBs": 83 crossings of 3rd, 4th or C^- resonance, 157 exceeded |ΔQ|>0.01
- Impressive performance for the first year of operation and low-ish intensities:

Control Paradigms III/III
Digital Control System & 'Real-Time'

- … "A system is said to be real-time if the total correctness of an operation depends not only upon its logical correctness, but also upon the time in which it is performed. [..] are classified by the consequence of missing a deadline:
  - Hard – Missing a deadline is a total system failure.
  - Firm – Infrequent deadline misses are tolerable, but may degrade the system's quality of service. The usefulness of a result is zero after its deadline.
  - Soft – The usefulness of a result degrades after its deadline, thereby degrading the system's quality of service."

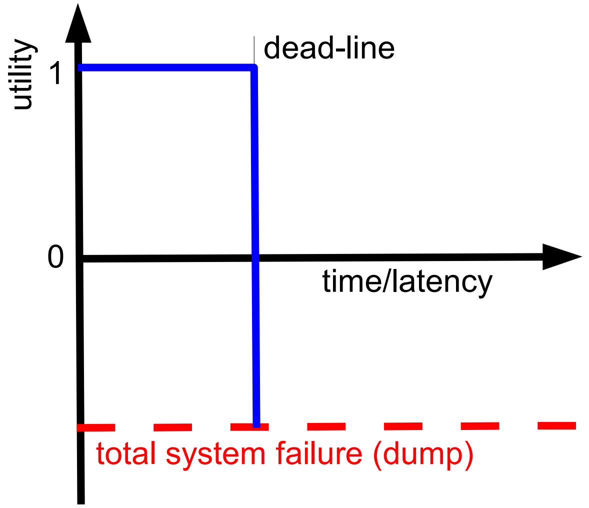

"hard"

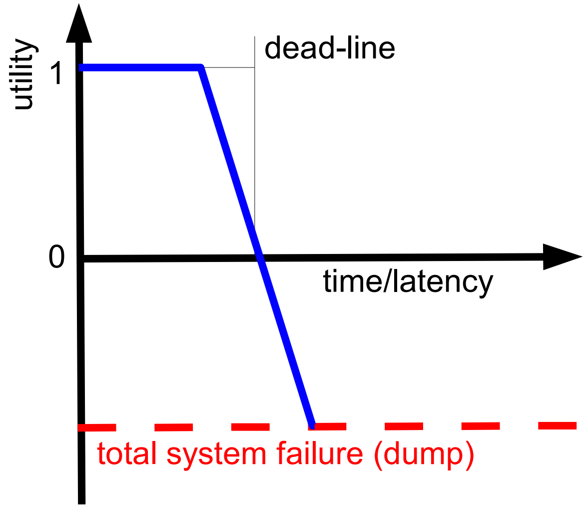

"firm"

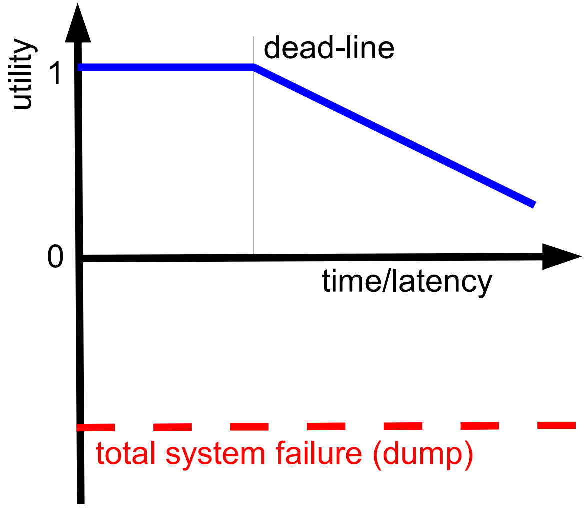

"soft"

Some common Misconceptions/Fallacies about 'Real-Time' (RT)[1]

1. *“There is no science in real-time-system design”*
2. *“Advances in supercomputer hardware will take care of RT requirements.”*
3. *“[..] is equivalent to fast computing.”*
4. *“[..] research is performance engineering.”*
5. *“[..] systems function in a static environment.”*
6. *“[..] is assembly coding, priority IRQ programming, and device driver writing.”*
7. *“[..] all been solved in other areas of computer science or operations research.”*
8. *“It is not meaningful to talk about guaranteeing RT performance, because we cannot guarantee that the hardware will not fail and the software is bug free or that the actual operating conditions will not violate the specific design limits.”*

- Obviously, the above is wrong but seems to be sometimes forgotten when discussing the specific technical implications.

---

[1] John A. Stankovic, “Misconceptions about real-time computing: a serious problem for next-generation systems”, IEEE Computer, Vol. 21 #10, 1988