Architectural Review of the LHC Orbit & Tune Feedback Systems, May 7th, 2013: After three Years of LHC Operation

Specific Orbit- and Q/Q'-Feedback Architecture

Ralph J. Steinhagen,

CERN, Beam Instrumentation Group

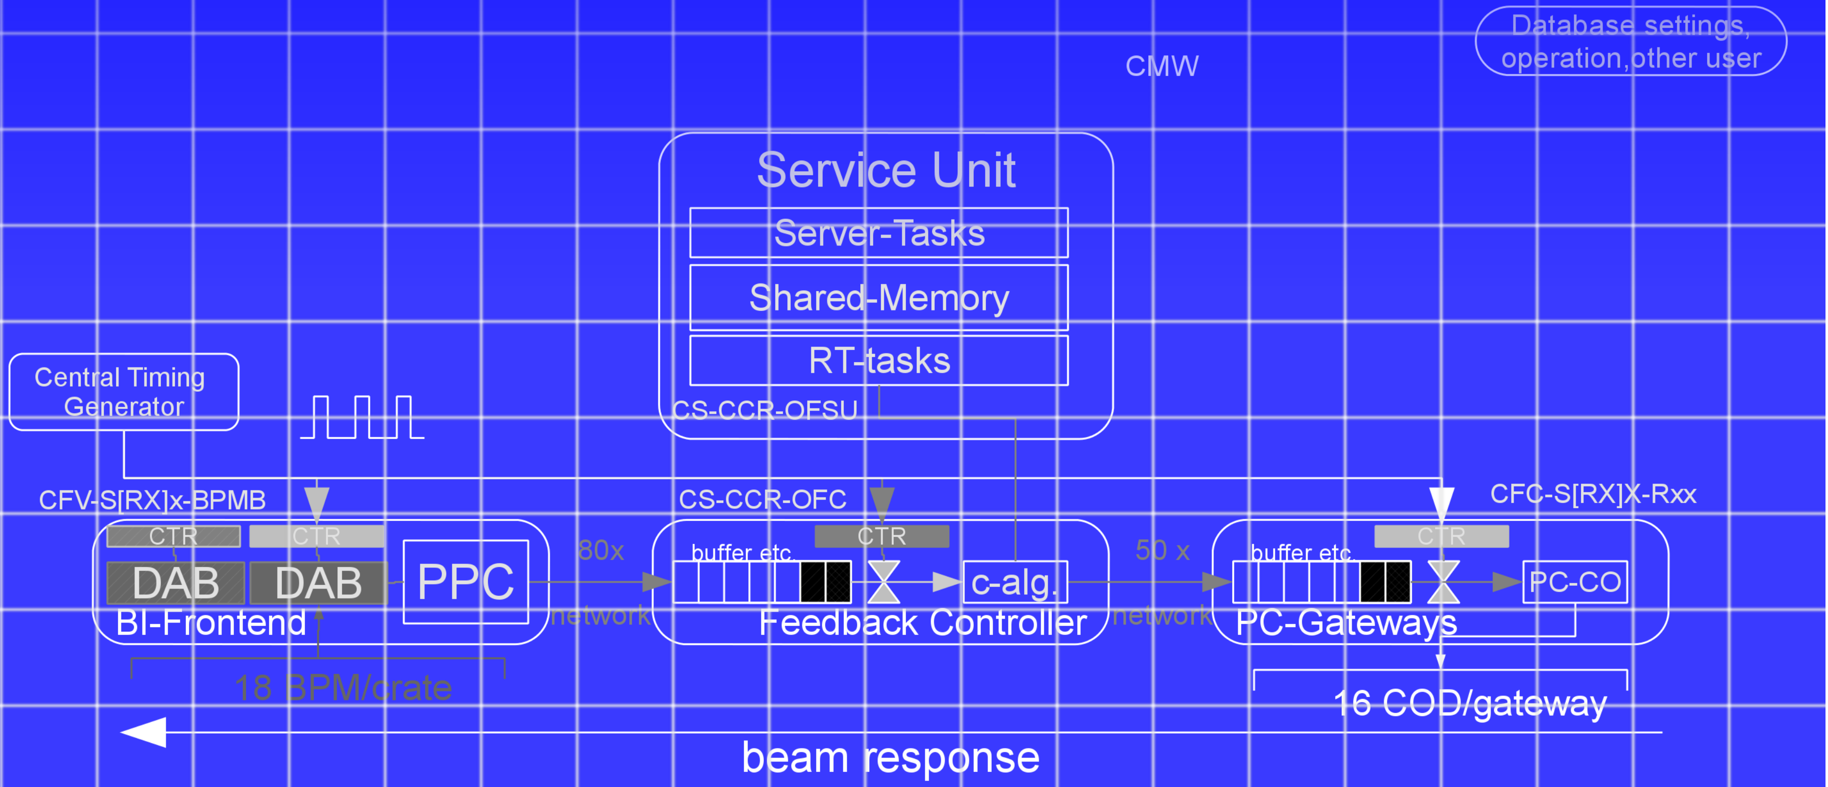

Multiple-Input-Multiple-Output (MIMO) Process Control

- 'Divide and Conquer' feedback controller design approach:
  1. Compute steady-state corrector settings $\vec{\delta}_{ss} = (\delta_1, \dots, \delta_n)$ based on measured parameter shift $\Delta x = (x_1, \dots, x_n)$ 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 $\vec{\delta}_{ff}$ to compensate changes of well known and properly described sources

space domain

time domain

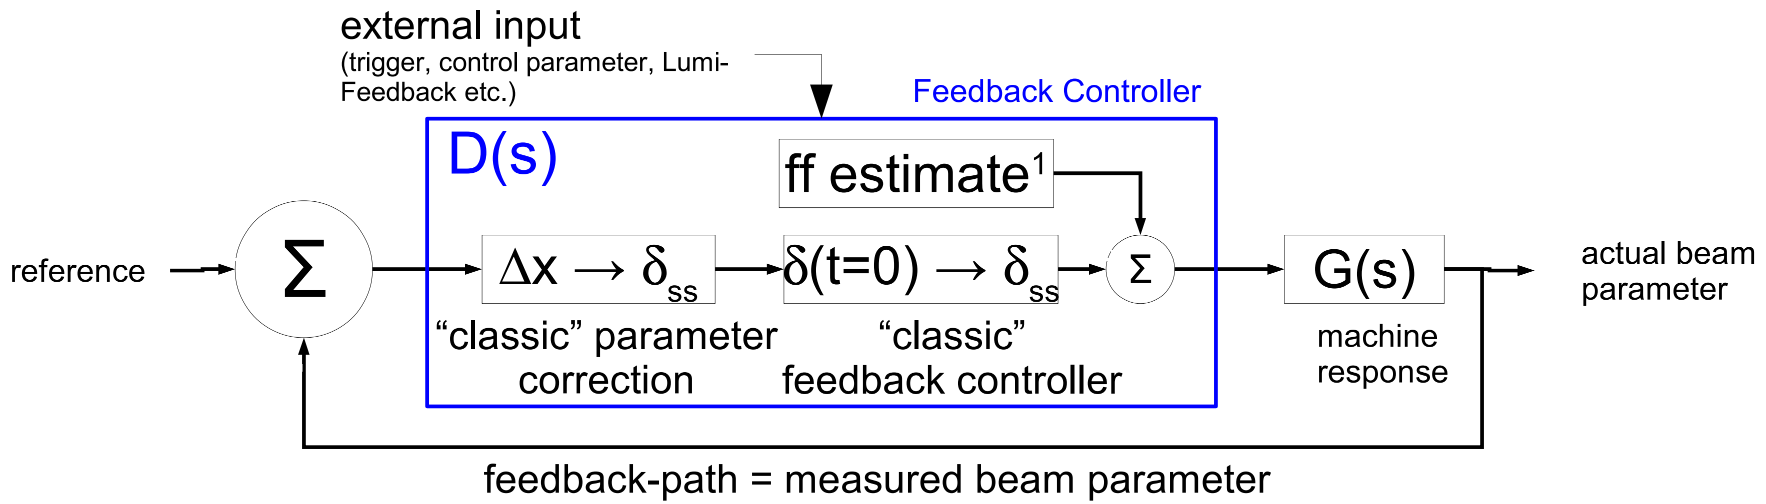

- (N.B. here G(s) contains the process and monitor response function)

Space Domain:

- Effects on orbit, Energy, Tune, Q' and C$^-$ can essentially cast into matrices:

$$\boxed{\Delta \vec{x}(t) = \underline{R}\cdot\vec{\delta}(t)} \quad with \quad R_{ij} = \frac{\sqrt{\beta_i\beta_j}}{2\sin(\pi Q)}\cdot\cos(\Delta\mu_{ij}-\pi Q) + \frac{D_i D_j}{C(\alpha_c - 1/\gamma^2)}$$

matrix multiplication

  - LHC matrices' dimensions:

$$\underline{R}_{orbit}\in\mathbb{R}^{1070\times 530} \qquad \underline{R}_{Q}\in\mathbb{R}^{2\times 16} \qquad \underline{R}_{Q'}\in\mathbb{R}^{2\times 32} \qquad \underline{R}_{C^-}\in\mathbb{R}^{2\times 10/12}$$

  - control consists essentially in inverting these matrices:

$$\boxed{\left\|\vec{x}_{ref}-\vec{x}_{actual}\right\|_2 = \left\|\underline{R}\cdot\vec{\delta}_{ss}\right\|_2 < \epsilon \quad\rightarrow\quad \vec{\delta}_{ss} = \tilde{R}^{-1}\Delta\vec{x}}$$

- Some potential complications:
  - Singularities = over/under-constraint matrices, noise, element failures, spurious BPM offsets, calibrations, ...
  - Time dependence of total control loop
- After this abstraction: Orbit, Energy, Q/Q'-FBs are basically the same

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

Linear algebra theorem*:

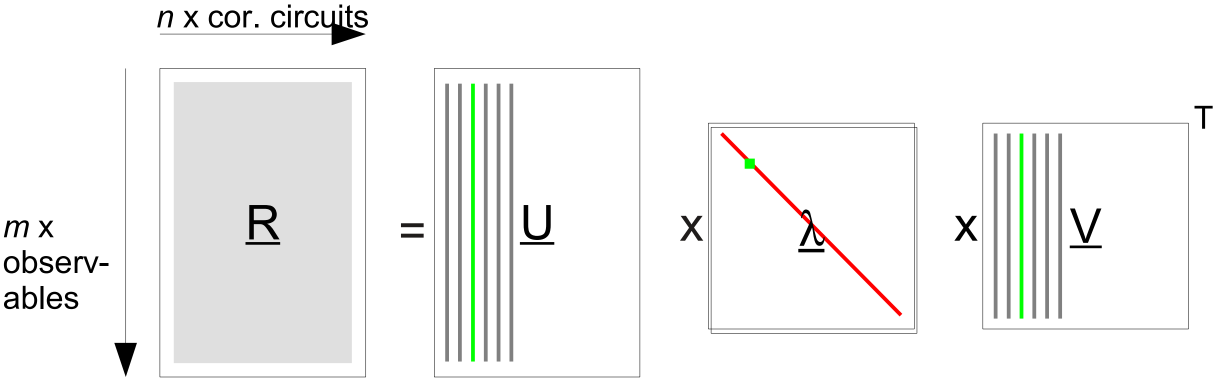

eigen-vector relation:

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

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

- 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 and circuit strengths δ
- Easy removal of singularities, (nearly) singular eigen-solutions have $\lambda_i \sim 0$
  - to remove those solution: if $\lambda_i \approx 0 \rightarrow$ '$1/\lambda_i := 0$'
  - 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

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
    → allowed re-using same ORM for injection, ramp and 10+ squeeze steps

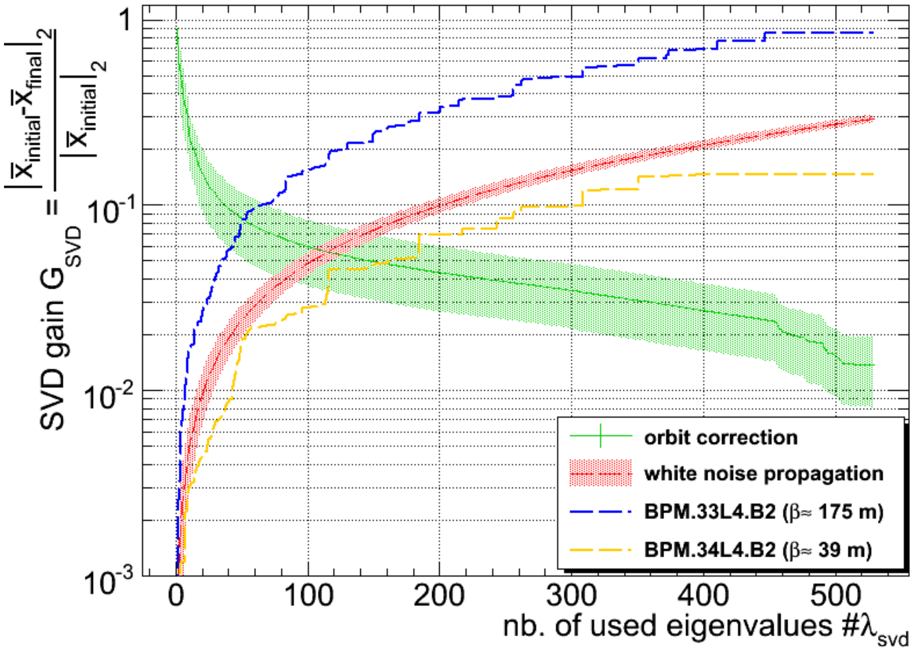

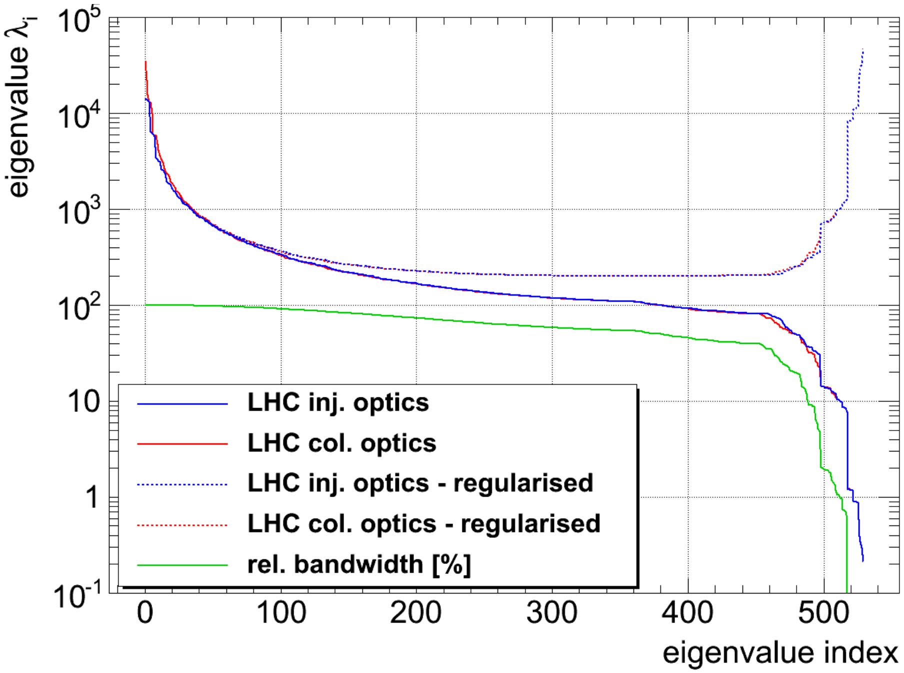

Multiple-Input-Multiple-Output (MIMO) Process Control

- '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

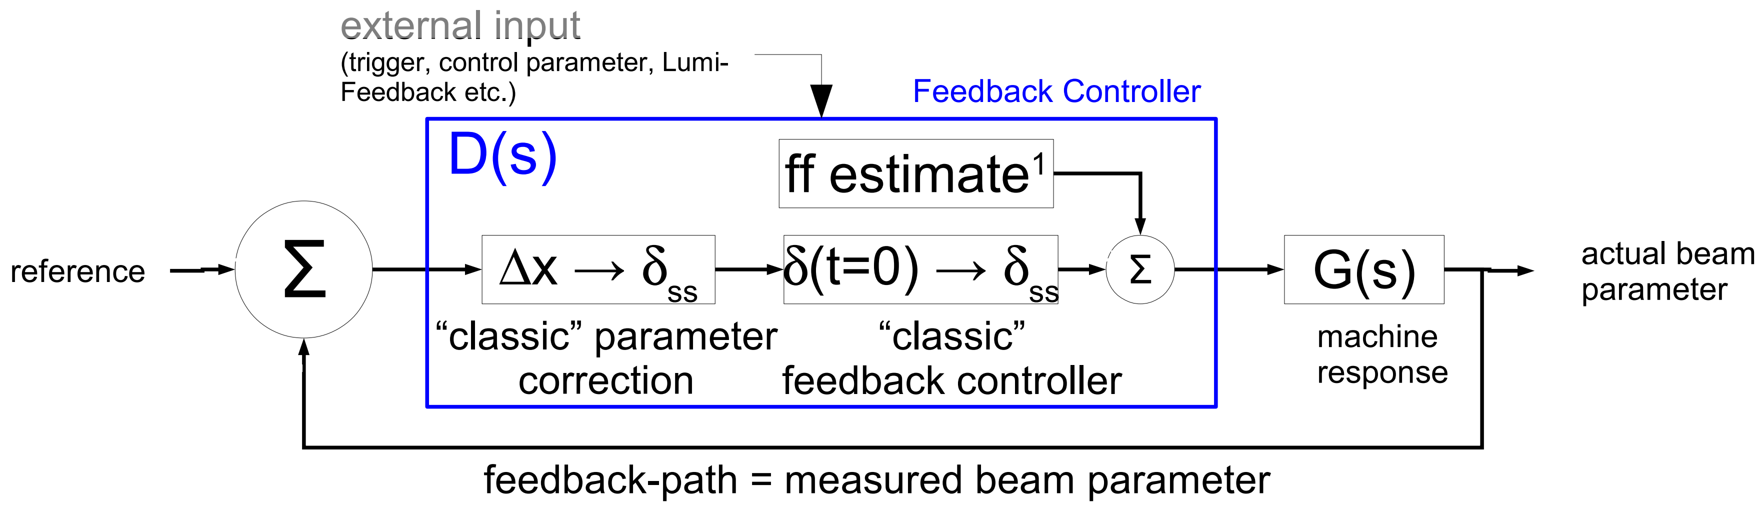

external input
(trigger, control parameter, Lumi-Feedback etc.)

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)

Time-Domain: Optimal Controller Design I/II

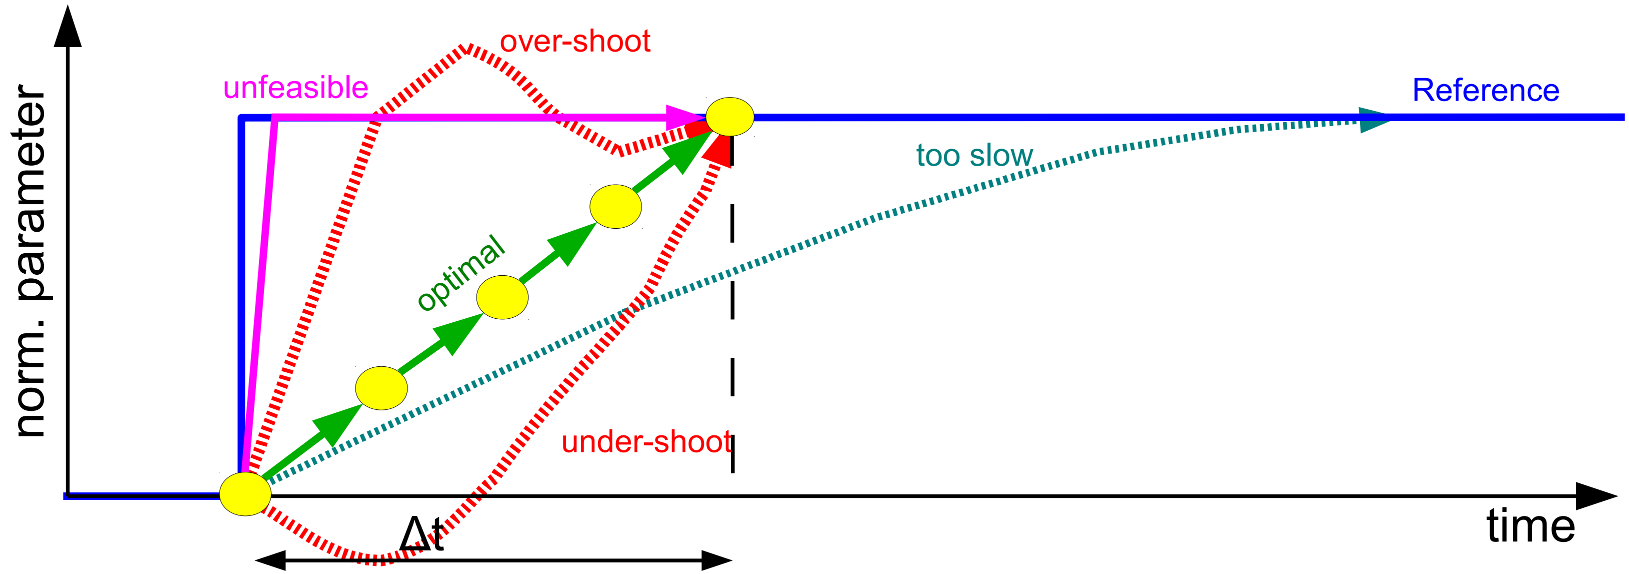

- Optimal control [or design] ...

  "... deals with the problem of finding a control law for a given system such that a given optimality criterion is achieved. A control problem includes a cost functional that is a function of state and control variables."

  – Common criteria: closed loop stability, minimum bandwidth, minimisation of action integral, power dissipation, ...

- classic closed loop: $T_0(s) = \frac{D(s)G(s)}{1+D(s)G(s)} \rightarrow 1$

Time-Domain: Optimal Controller Design II/II

- Optimal control[1] for the 'small-signal response' yields classic PI-controller
  - Single free parameter $\alpha > \tau \ldots \infty$ (one per FB loop)
    - facilitates trade-off between speed and robustness
    - adaptive gain-scheduling based on operational scenario (implemented, but not exploited operationally)

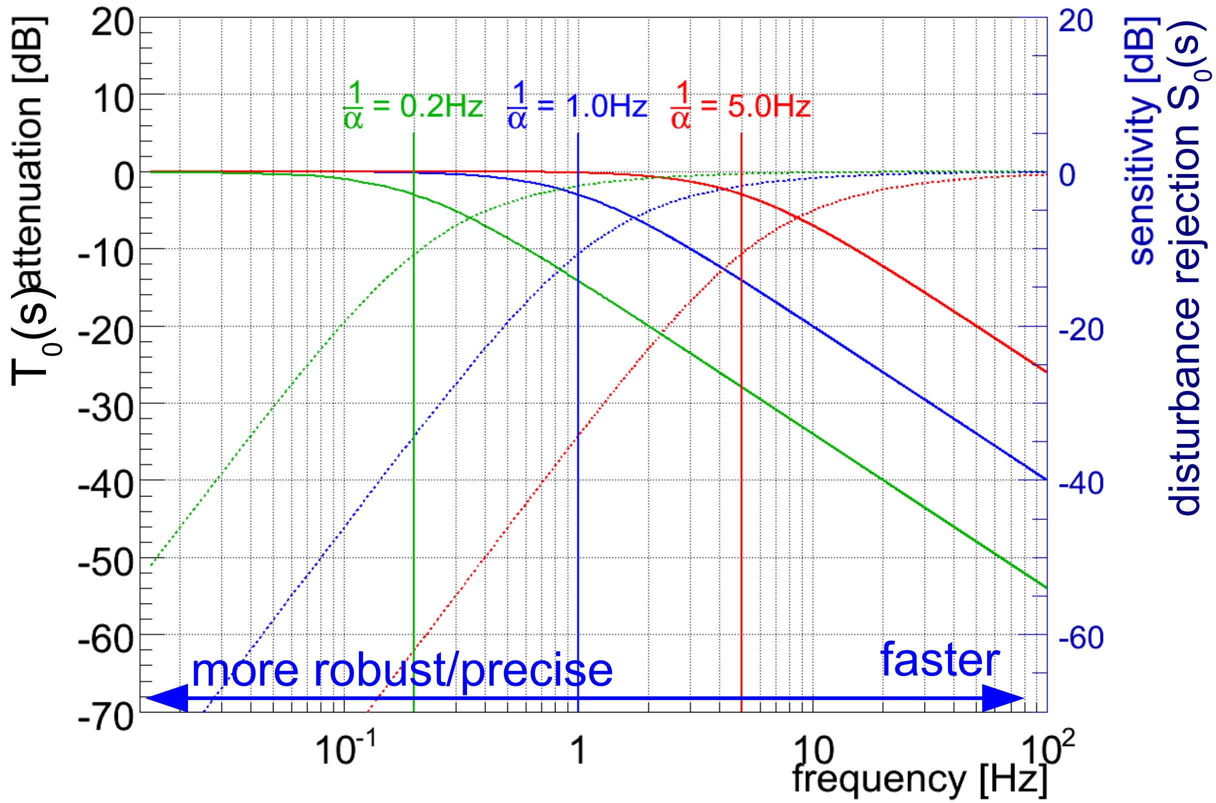

---

[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

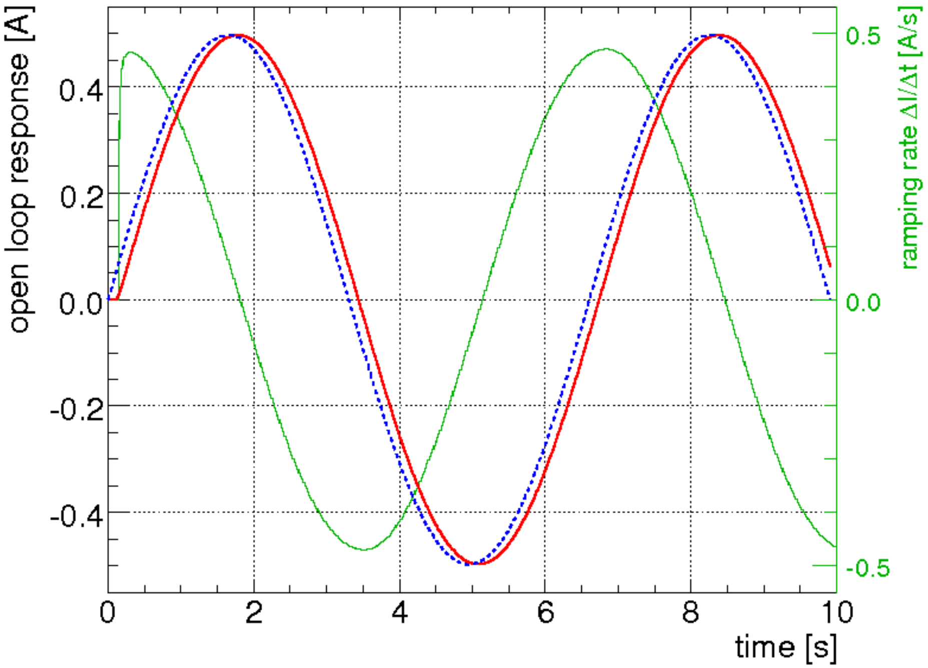
slow perturbation: perfect tracking

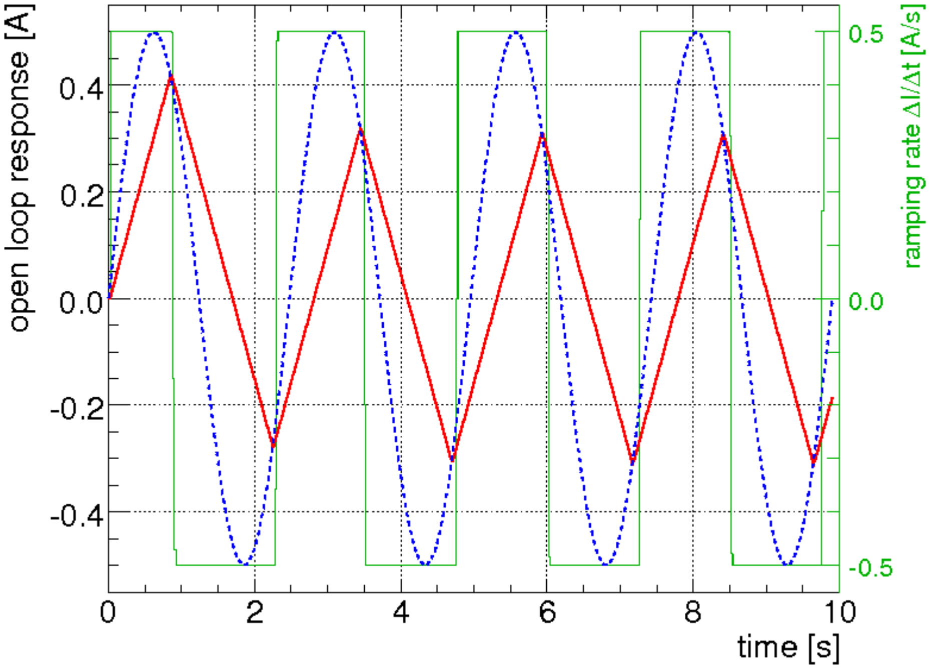
fast perturbation: saw-tooth

Time-Domain: Non-Linearities II/III

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

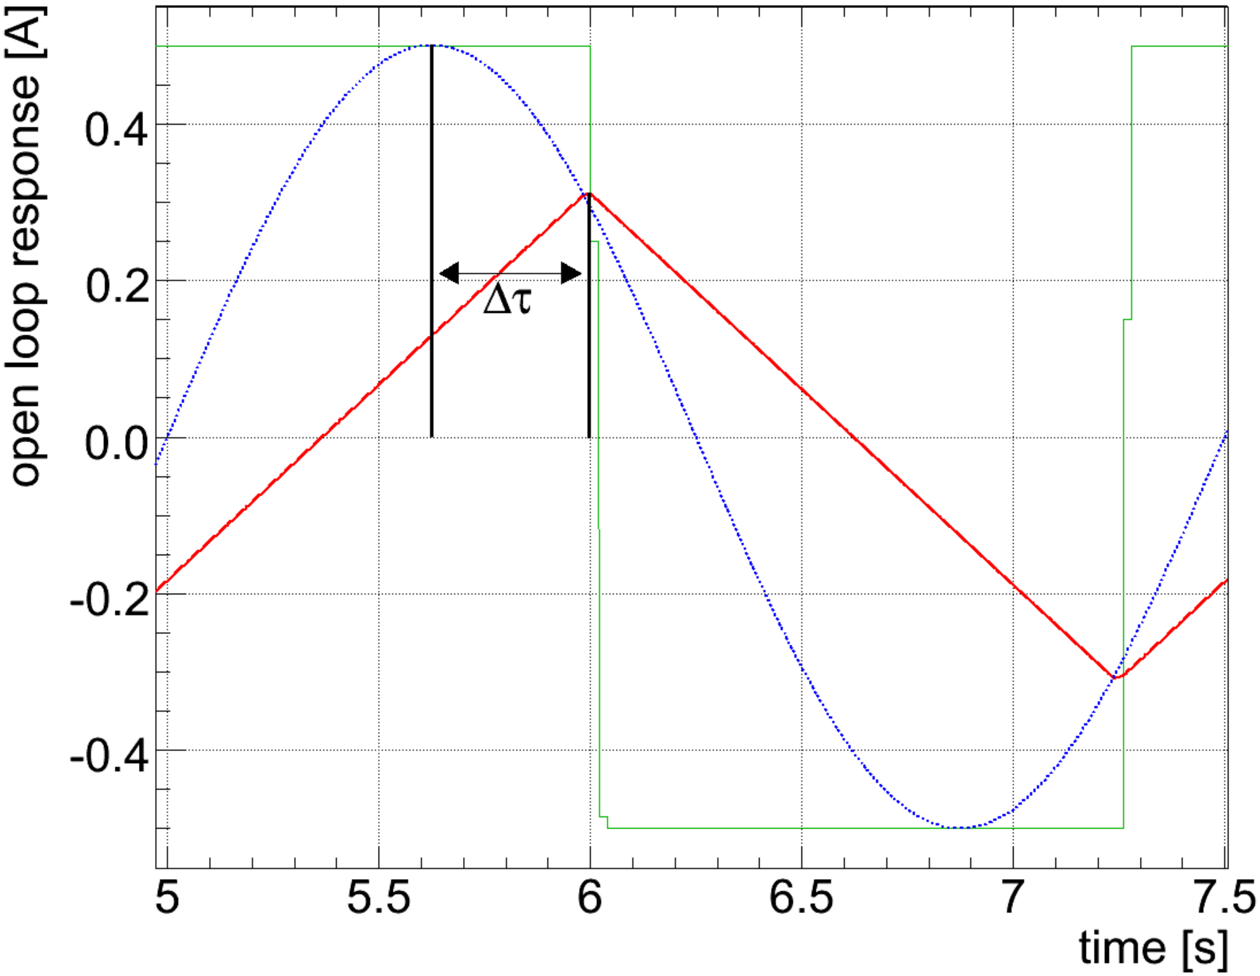

Time-Domain: Non-Linearities III/III
Orbit Feedback Bandwidth vs. Actual Orbit Perturbations

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

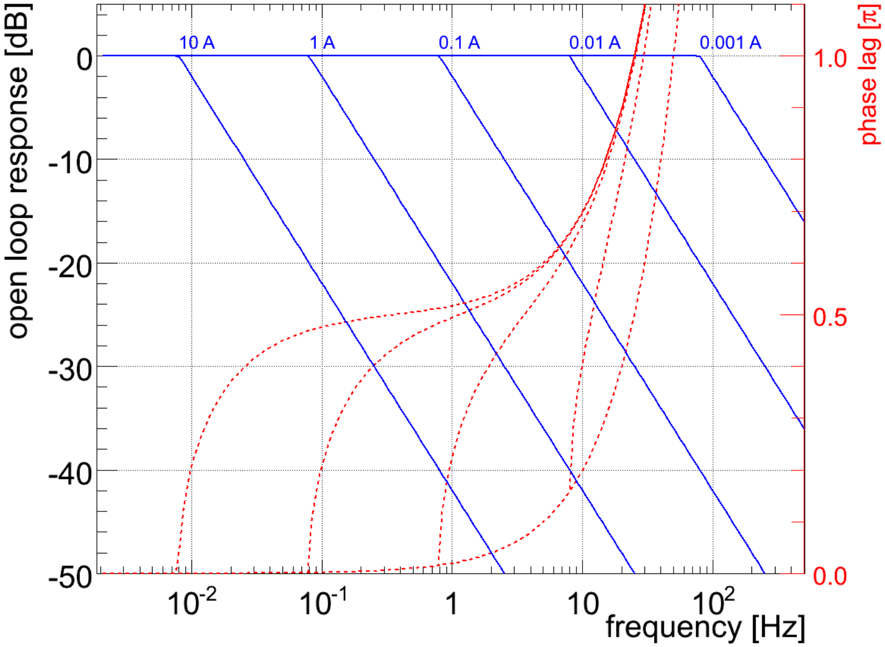

ΔI=0.1A ↔ Δx≈32 µm@β=180m

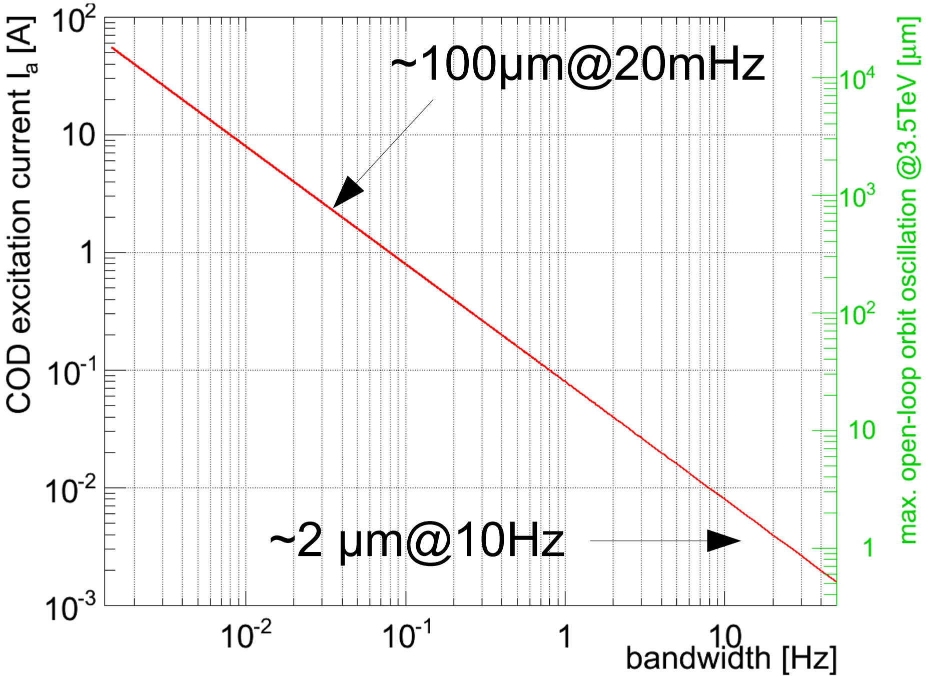

- Consider ~35µm@1Hz as effective/ practical bandwidth @4TeV

  (assuming 3C bump)

- Not much margin to push bandwidth post-LS1@6.5 TeV
  - Orbit-FB operated close to this limit already in 2012 (N.B. Δφ contribution!)

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

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

- 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

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

→ Smith-Predictor and Anti-Windup paths:

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

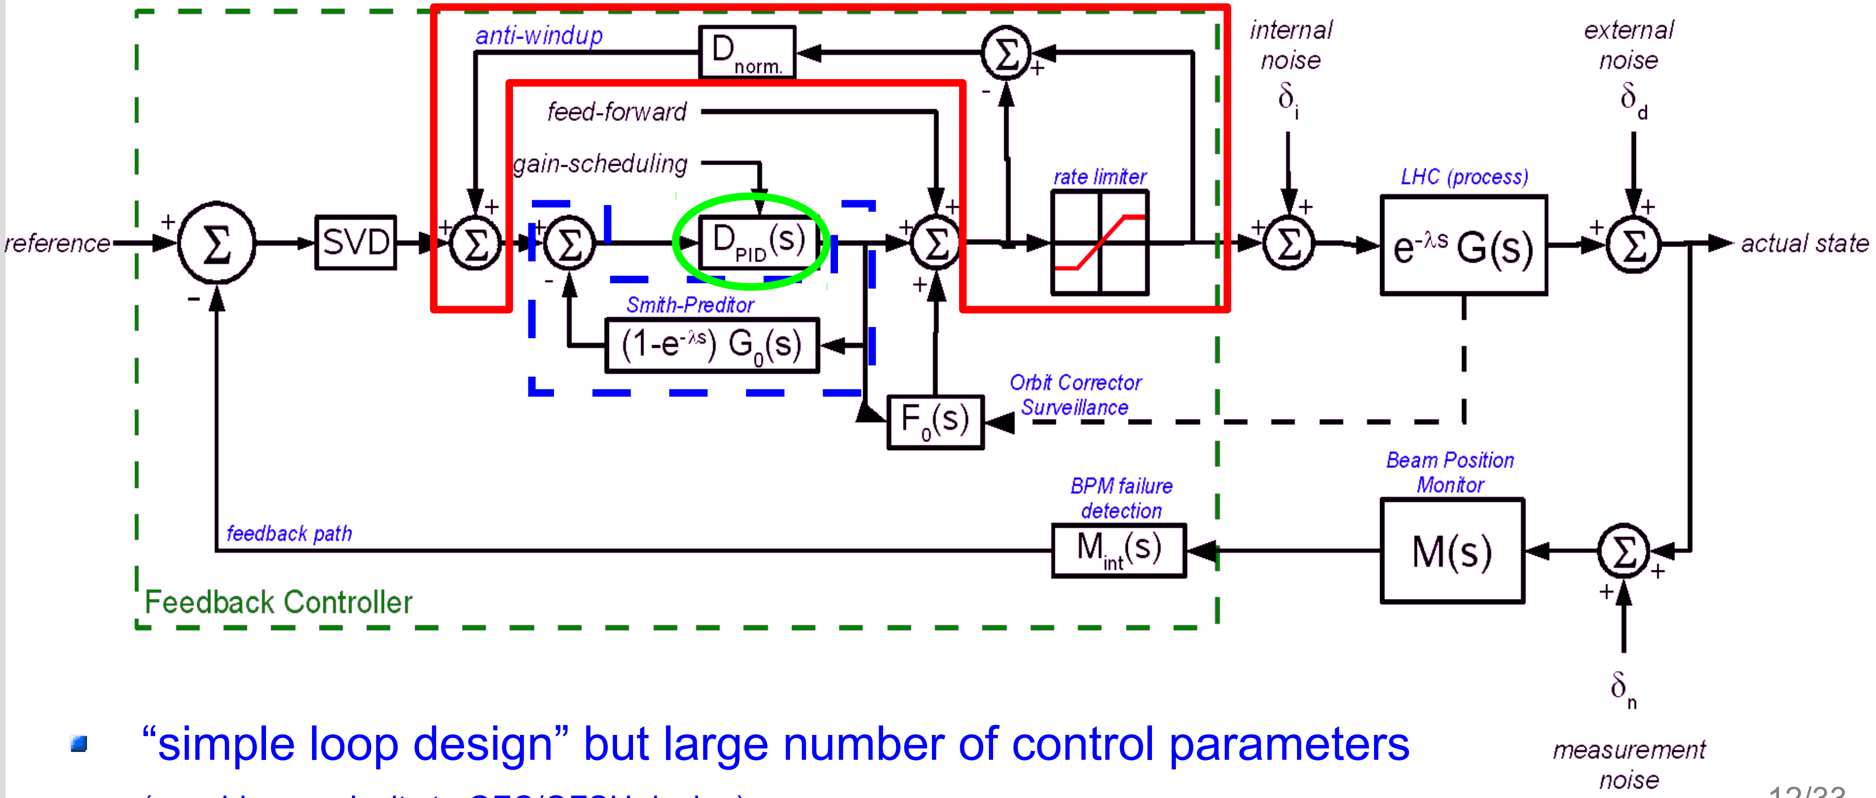

- “simple loop design” but large number of control parameters

  (→ adds complexity to OFC/OFSU design)

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

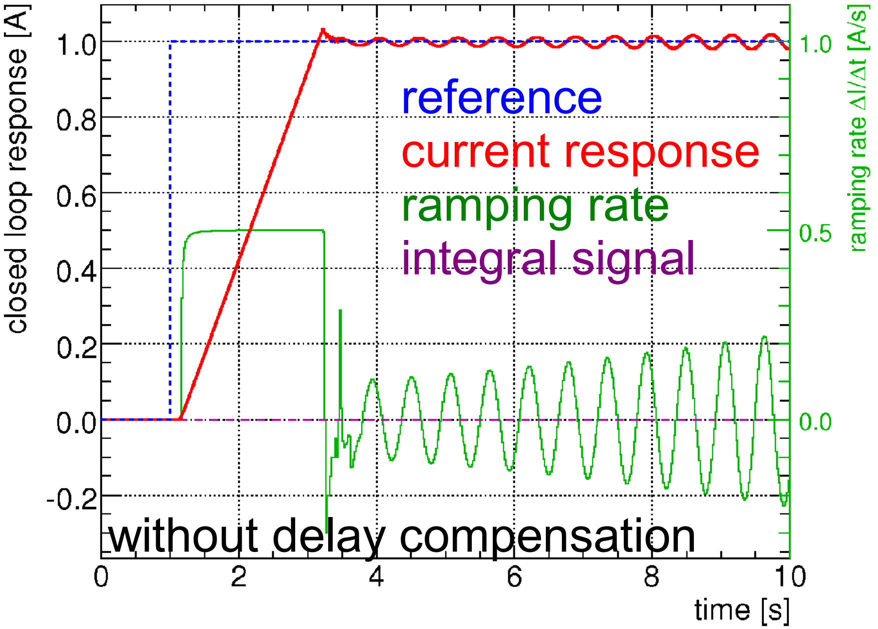

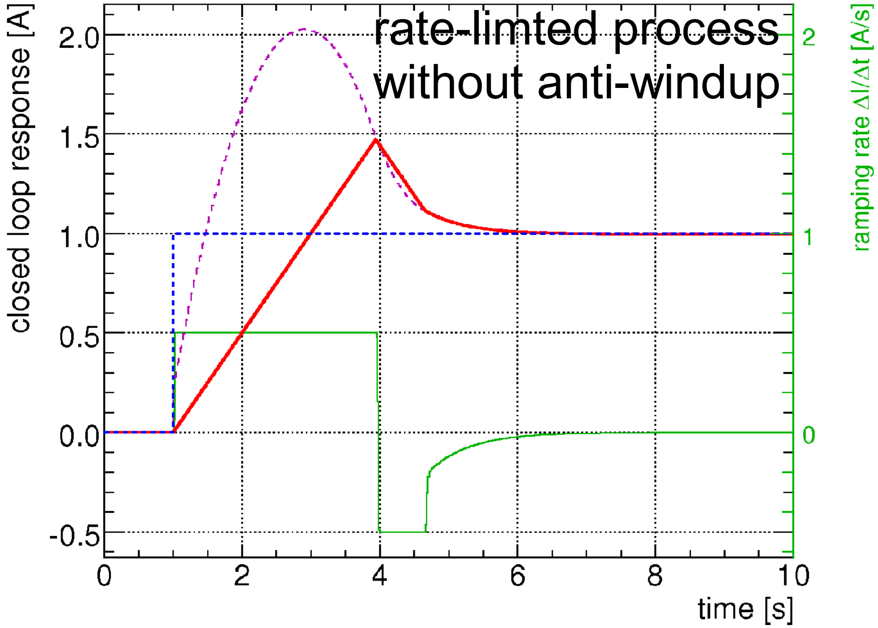

with full delay and windup compensator scheme:

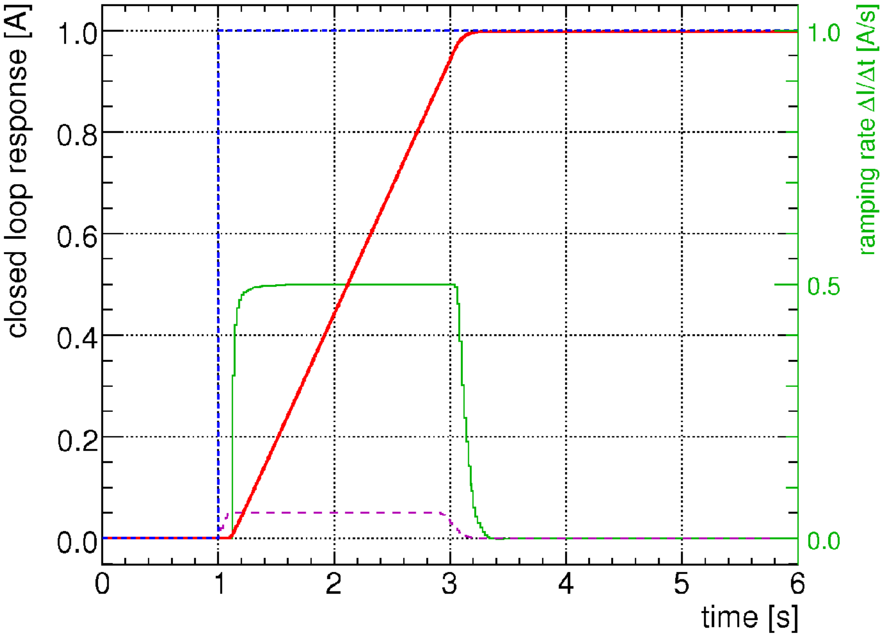

Inter-loop Cross-Dependencies I/II

- Ensemble of feedback loops that simultaneous act on the same beam:
  - beam-based feedbacks on: orbit, energy (radial loop), tune, chromaticity, coupling, fast transverse feedback (damper), synchro-loop, …

  → can/will cause cross talks and even instabilities if not included in design!

- Some choices:
  - Decoupling of the parameter space:
    - Orbit FB (betatron-pertubations) vs. Energy FB (dispersion orbit)
  - Decoupling of operational ranges (either e.g. amplitude or time scales)
    - i.e. Q' tracker being faster than actual Tune-FB loop
  - Introducing a Master-Slave Structure:
    - Energy FB & Q' Tracker sharing the same reference
    - Orbit FB being the slave of the Energy-FB (radial loop)

→ all favour and are more easy to implement in a central controller

Inter-loop Cross-Dependencies II/II

- Multiple FBs and measurements acting on the same RF cavity frequency (N.B. radial position limited by collimator gap)
  - Q' tracker, energy FB (≈'radial loop'), Q'' and other optics measurements
  - strategy: orbit feedback acts as a slave system controlling the RF
    - dispersion orbit is subtracted/not corrected by 'regular Orbit-FB'
    - Energy FB corrects w.r.t. to Q' tracker set reference
- Example: Δp/p reconstruction comparison MADX vs. linear model

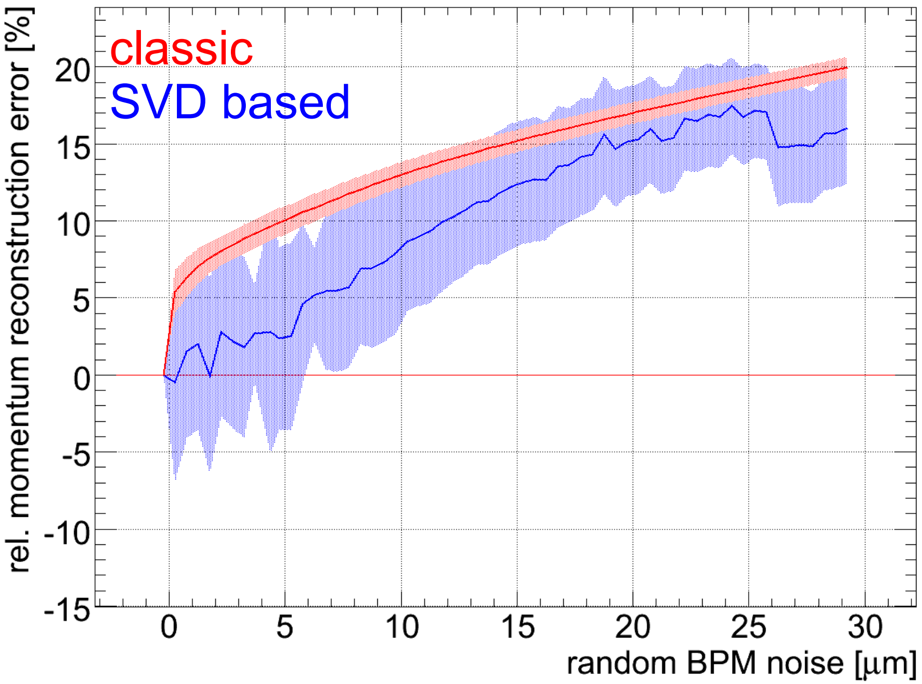

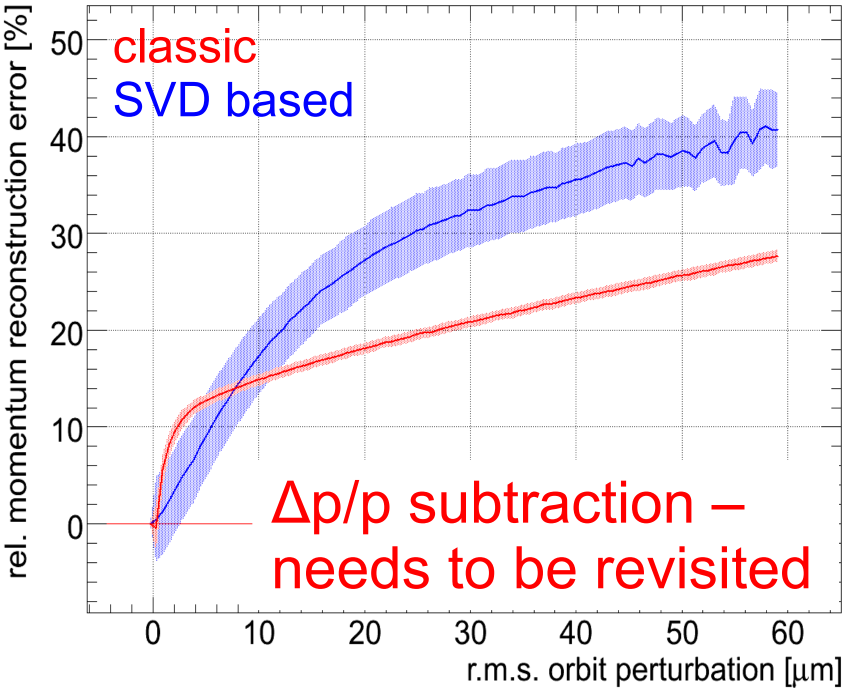

Δp/p subtraction – needs to be revisited

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

- 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 periodic constant load
    - 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

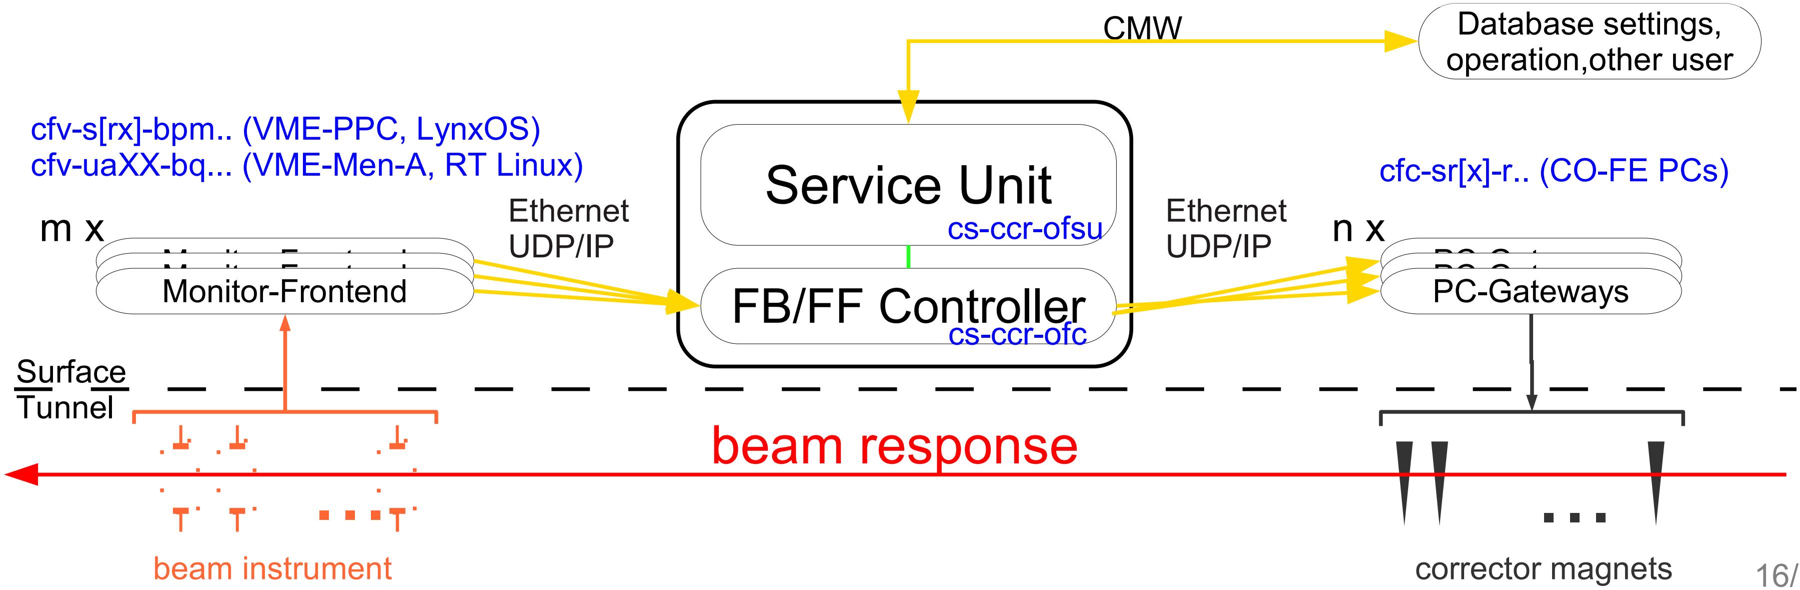

Jitter Compensation
Data Concentration Paradigm I/II

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...

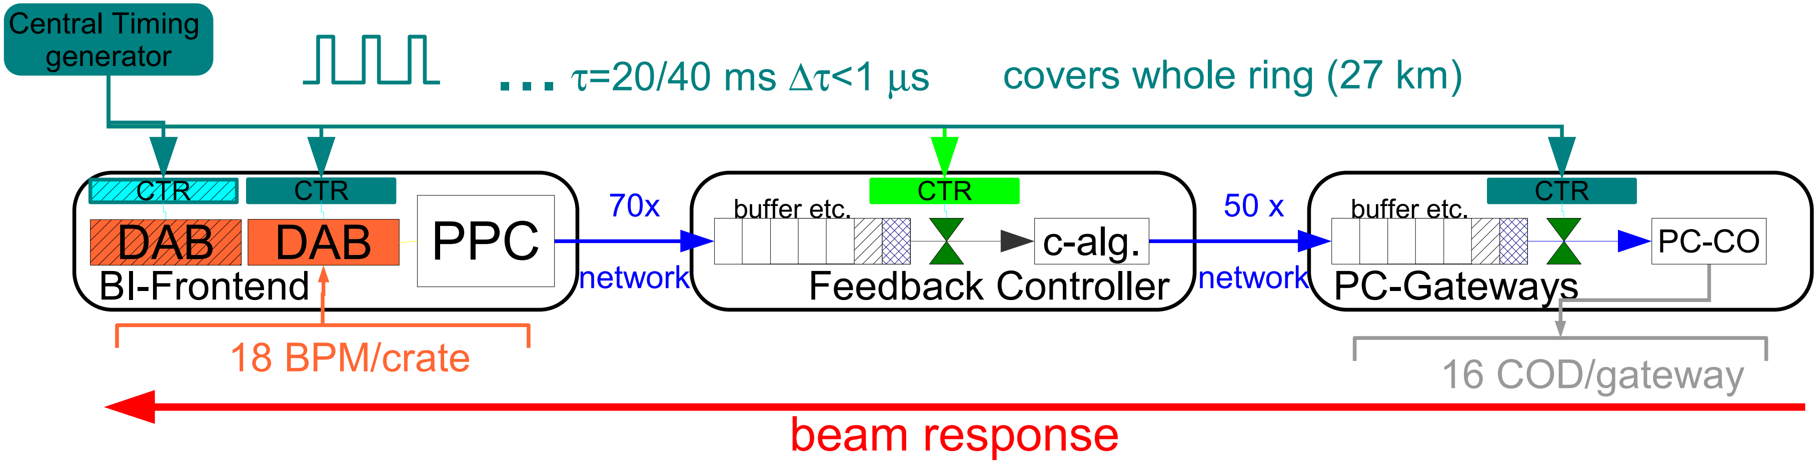

- Single CTR in OFC → single point of failure
  → dropped it in 2011 in favour of retrieving timing from multiple BQBBQLHCs
  → direct UDP software link between BST and OFC for 25 Hz trigger

Jitter Compensation
Data Concentration Paradigm II/II

A) 'Synchronous' trigger derived from BST via custom 'TTL → UDP/IP' translator

- Start acquisition with BST derived/distributed start signal
  - accumulate orbit, Q/Q' and FGC data for 10 ms
- Stop once the accumulation period of 10 ms has passed, drop late data
- PRO: more robust w.r.t. individual front-end crate errors
- CON: unforgiving if packets arrive with extra delays → equal to 'no data'

B) Auto-triggered based on BPM FE synchronicity
(initially foreseen for commissioning but became operational 'default')

- Start with the first BPM packet arriving at OFC
  - accumulate orbit, Q/Q' and FGC data
- Stop once
  a) all expected packet are received, or
  b) a second packet from an already received FE arrived
- PRO: more forgiving if individual front-ends send data late, independent on actual data rate (i.e. can dynamically switch between 1, 10, 25 Hz BPM update rate)
- CON: issues with packets coming in bursts from single front-ends

- Known problem: both schemes work but have issues with either 'RT constraints' or 'constant load conditions' being broken in FESA front-ends

Specific Orbit Feedback Controller (OFC) Structure I/IV

Main streaming taks contained in 'OFBController.cpp', logic flow:

A) <general initialisation>

B) Main Loop

1. Data accumulation loop (free-running or locked at 25 Hz):
   - BPMConcentrator – *nomen est omen*
   - *CODConcentrator* – *nomen est omen*
   - QQPConcentrator, MachineState – *nomen est omen*
2. <validate setting and received data>
3. <update references>
4. EnergyCorrection – radial loop feedback, radial modulation, …
5. OrbitCorrection – *orbit feedback space domain*
   - Wakes up two worker threads performing the two O(n^2) multiplication
6. QQPConcentrator – *tune feedback space and time domain*
7. *<send COD and Q/Q' corrector data>*
8. *<publish/stream OFC state via UDP to OFSU>*
9. *<wait up to 5 ms or for remainder of iteration, service TInterlink requests>*

C) <general de-initialisation/restart>

Specific Orbit Feedback Controller (OFC) Structure II/IV

- Additional independent asynchronuous tasks/threads:
  - Tinterlink – *RPC class executed only once the main task is finished*
    - Light-weight communication layer between OFC and OFSU
    - Blocked except at the end of very main iteration
  - ReferenceOpticsMagic – *OFC-based optics re-computation*
    - Executed (routinely) to account for deselected BPMs, CODs & optics errors
    - High CPU load and risk of stalling the OFC
    - Was added to OFC initially as a “hack”/short-term mitigation on LHC Day-I → should be migrated to OFSU or reliable RT level above
- Typical CPU load on cs-ccr-ofc (note difference peak ↔ avg. load):

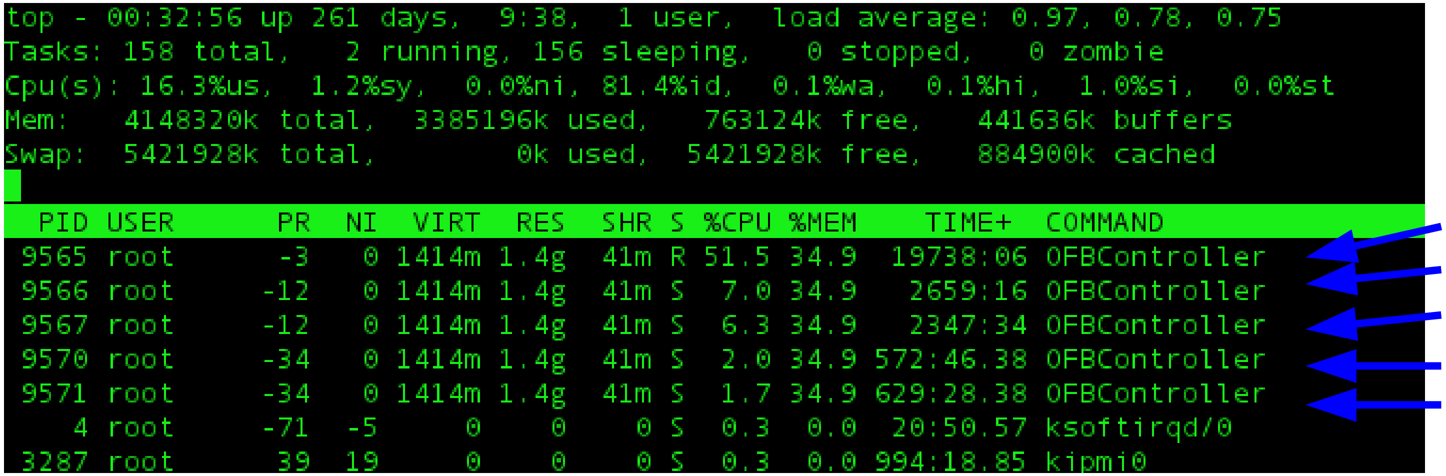

```
top - 00:32:56 up 261 days,  9:38,  1 user,  load average: 0.97, 0.78, 0.75
Tasks: 158 total,   2 running, 156 sleeping,   0 stopped,   0 zombie
Cpu(s): 16.3%us,  1.2%sy,  0.0%ni, 81.4%id,  0.1%wa,  0.1%hi,  1.0%si,  0.0%st
Mem:   4148320k total,  3385196k used,   763124k free,   441636k buffers
Swap:  5421928k total,        0k used,  5421928k free,   884900k cached

  PID USER      PR  NI  VIRT  RES  SHR S %CPU %MEM    TIME+  COMMAND
 9565 root      -3   0 1414m 1.4g  41m R 51.5 34.9  19738:06 OFBController
 9566 root     -12   0 1414m 1.4g  41m S  7.0 34.9   2659:16 OFBController
 9567 root     -12   0 1414m 1.4g  41m S  6.3 34.9   2347:34 OFBController
 9570 root     -34   0 1414m 1.4g  41m S  2.0 34.9 572:46.38 OFBController
 9571 root     -34   0 1414m 1.4g  41m S  1.7 34.9 629:28.38 OFBController
    4 root     -71  -5     0    0    0 S  0.3  0.0  20:50.57 ksoftirqd/0
 3287 root      39  19     0    0    0 S  0.3  0.0 994:18.85 kipmi0
```

Main loop
Orbit-FB-H
Orbit-FB-V
TInterlink
CODConcentrator
(moved to main loop)

Specific Orbit Feedback Controller (OFC) Structure III/IV

- Functional structure, timing diagram & core utilisation (CPU shielding):

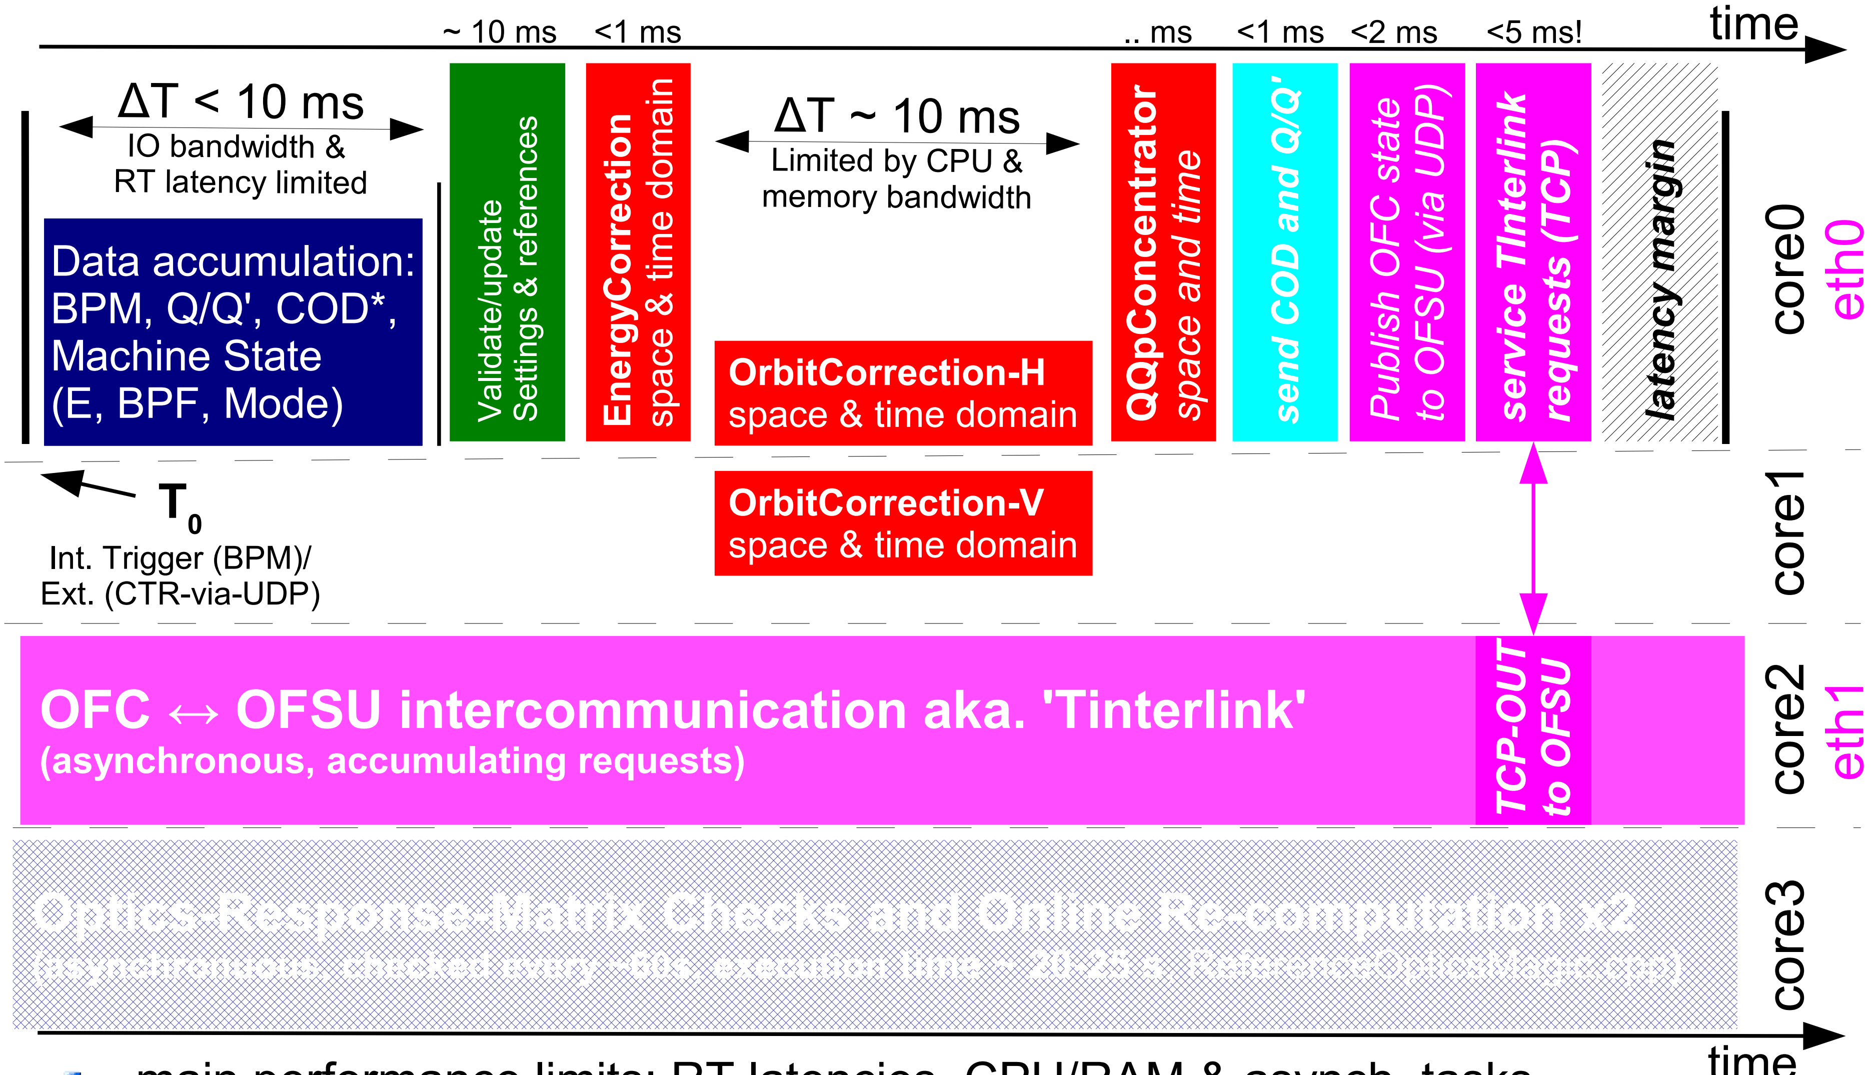

time

~ 10 ms <1 ms .. ms <1 ms <2 ms <5 ms!

ΔT < 10 ms
IO bandwidth &
RT latency limited

Data accumulation:
BPM, Q/Q', COD*,
Machine State
(E, BPF, Mode)

Validate/update
Settings & references

EnergyCorrection
space & time domain

ΔT ~ 10 ms
Limited by CPU &
memory bandwidth

OrbitCorrection-H
space & time domain

QQpConcentrator
space and time

send COD and Q/Q'

Publish OFC state
to OFSU (via UDP)

service TInterlink
requests (TCP)

latency margin

core0 eth0

T_0
Int. Trigger (BPM)/
Ext. (CTR-via-UDP)

OrbitCorrection-V
space & time domain

core1

OFC ↔ OFSU intercommunication aka. 'Tinterlink'
(asynchronous, accumulating requests)

TCP-OUT
to OFSU

core2 eth1

Optics-Response-Matrix Checks and Online Re-computation x2
(asynchronous, checked every ~60s, execution time ~ 20-25 s, ReferenceOpticsMagic.cpp)

core3

time

- main performance limits: RT latencies, CPU/RAM & asynch. tasks

Specific Orbit Feedback Controller (OFC) Structure IV/IV

- Fairly flat C++ Class Hierarchy ↔ reflects io-streaming task:

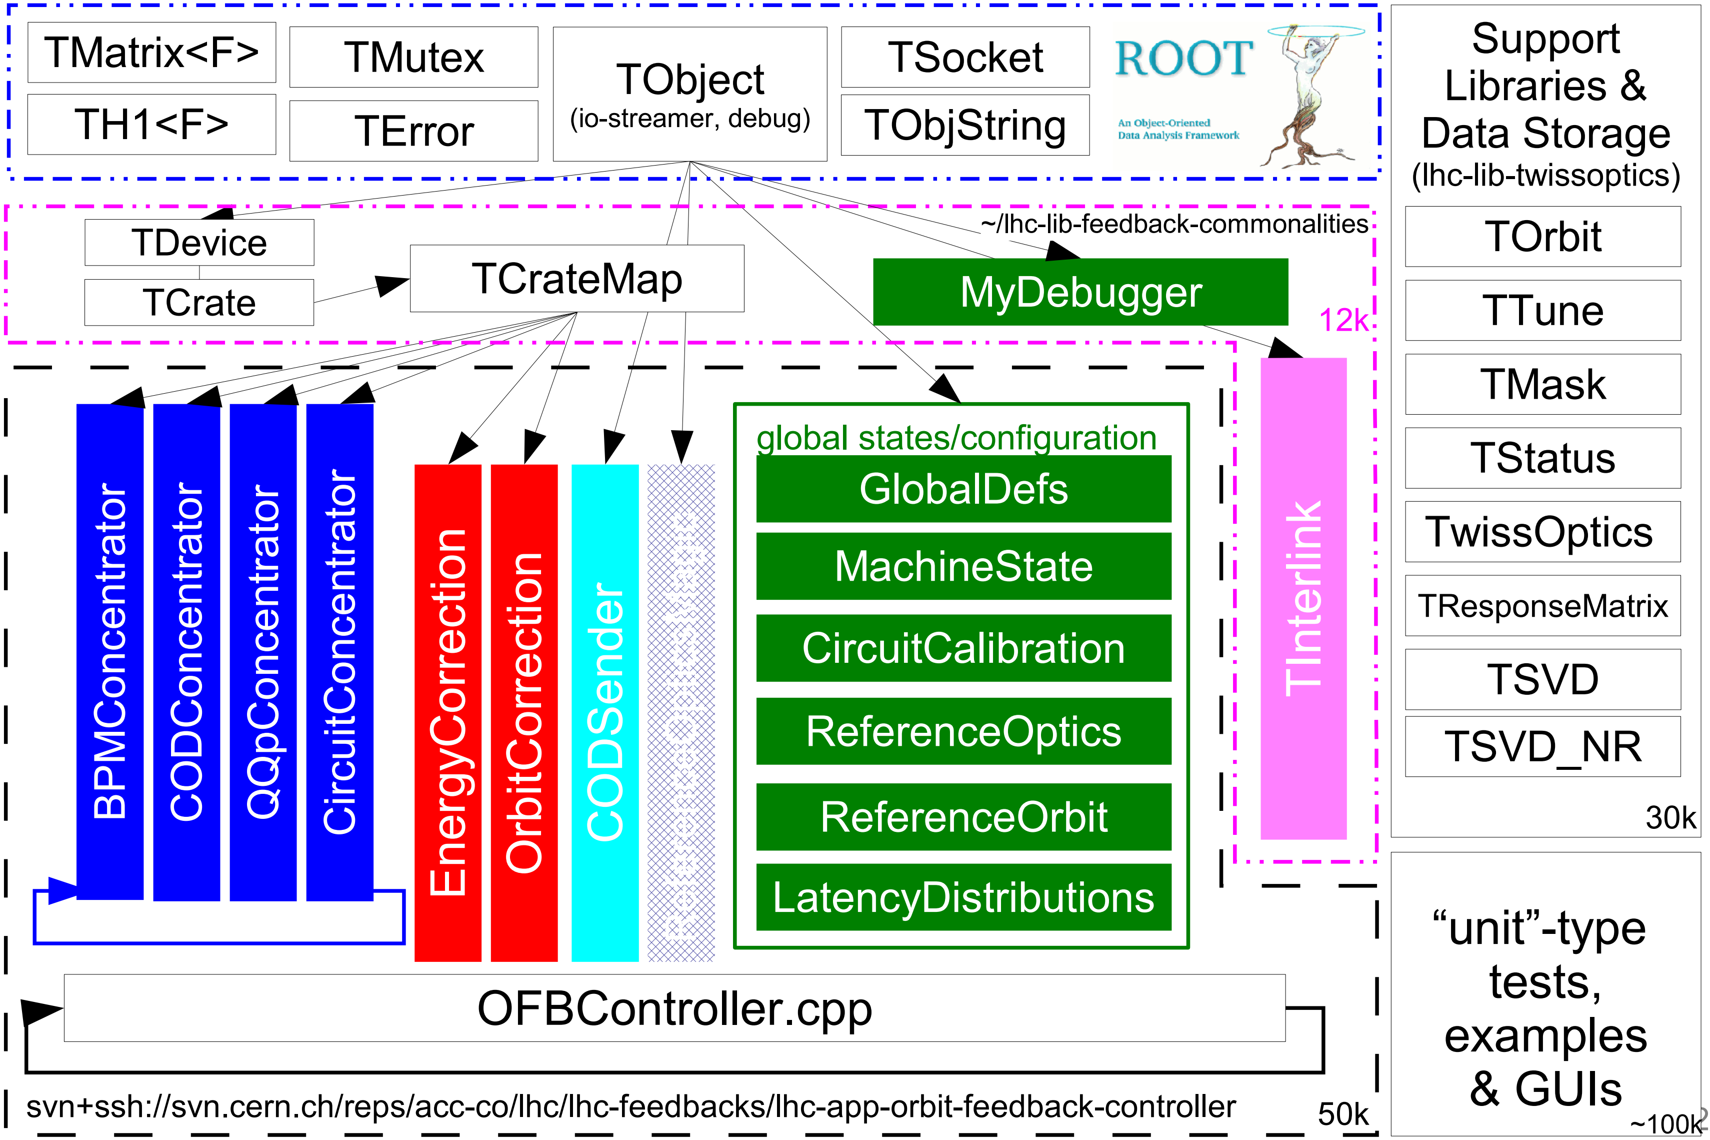

OFC ↔ OFSU Communication to the Outside World I/II
ROOT IO Messages Usage

- Printing to console or (worse) to file is hazardous in an RT environment as it can totally block the process → instead: implemented a circular buffer which re-directs all ROOT, OFC, … Error, Warning, Info and Debugging messages:

```
2011-12-13 07:05:46 - Error in
<BPMConcentrator::CheckDoubleValue(range)>: value +0.000000e+00 at
index 10 in dabTemp is out of range [+1.000000e+01, +1.000000e+02]
```

- Many OFC, data integrity or FB problems can be traced very quickly this way
- Allows quick trace-back to 'class & member' function, e.g. here: 'BPMConcentrator.cpp::CheckDoubleValue(..)' with the message one finds:

```
[..]
unsigned short ttemperature_short = SWAP_USHORT(data.dabTemp[i]);
Double_t ttemperature = CheckDoubleValue(0.1*ttemperature_short,
0.0,tempStatus, i, "dabTemp", 10.0, 100.0); // [10, 100] degC
[..]
```

- short-coming: not often used in CCC and limited internal OFC buffer (~½h)
  → would be nice if this could be logged via OFSU

OFC ↔ OFSU Communication to the Outside World I/II
Tinterlink – ROOT I/O streamers and CINT

- TInterlink implements a light-weight RT RPC for OFC↔OFSU streaming data
  - maintains list of all commands, input, output parameters and link to local objects
- Specific functions are registered as (compile time consistency check):

```
interlink->RegisterFunction(this, (TObjectFunction)&OrbitCorrection::SetOrbitFBStateH,
                  "OrbitFBStateH", kWRITE, TCallBack::kNONE, TCallBack::kBool_t,
                  "sets horizontal OrbitFB state: kTRUE -> on, kFALSE -> off []");

interlink->RegisterFunction(this, (TObjectFunction)&BPMConcentrator::GetLastOrbit,
                  "LastOrbit", kREAD, TCallBack::kTObject, TCallBack::kInt_t,
                  "Get last orbit sample for given plane [Torbit]");

interlink->RegisterFunction(this, (TObjectFunction)&BPMConcentrator::SetGraceTime,
                  "BPMGraceTime", kWRITE, TCallBack::kNONE, TCallBack::kLong_t,
                  "time to wait for incoming packets [us]");
```

- Similar 'get/set + subscribe' paradigm as CMW
- Machine & human readable interface remotely invoked, e.g. via (TCP:8080):
  - list of all available OFC commands: *"get commands"*

- Total ~ 600 commands → propagates to OFSU setting management code complexity: many simple scalar commands like 'switch OFB on/off', gains, …
- TCP/IP known weakness: 'unresponsive client' can block main loop
  → would like to move this to UDP/IP after LS1

Verification of RT Conformity – Latency Histograms

- Latency is measured at every major loop break-point
  - Example – latency of the data concentration loop (for short periods):

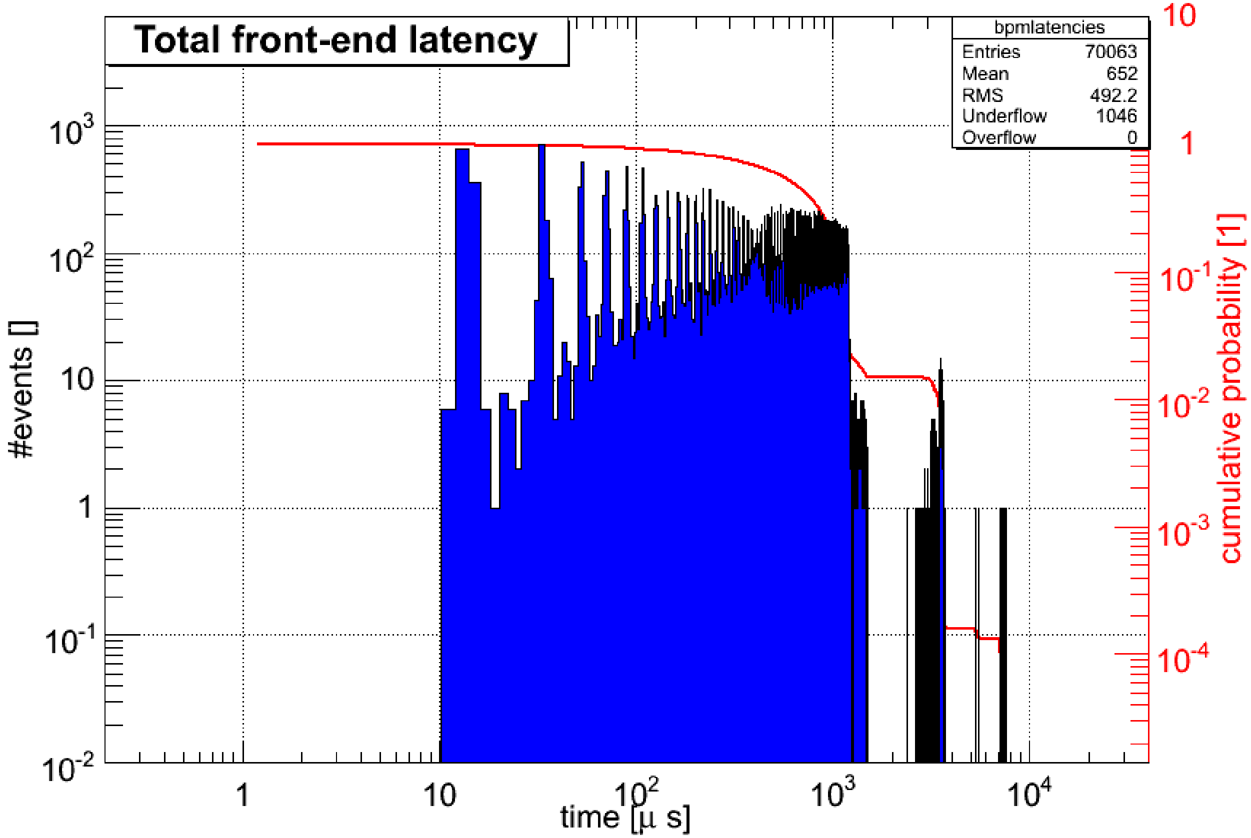

- Helps identifying code sections with weak RT behaviour

→ main performance and quality assurance indicator

Beyond the Minimum-Design Feedback Functionality

"Any beam-based feedback loop is only as strong as its weakest link and beam diagnostics it is based upon."

- Fraction of the OFC code base deals with the primary feedback logic
  (receiving packets → computing corrections → sending new values to FGCs)

- Most OFC code dedicated to non-nominal operation and exception handling
  - Beam measurement quality checks:
    - Detection of erroneous BPMs
    - Q/Q' reconstruction, detection of failures
  - Data and real-time latency integrity checks
    - NaN, Inf, out-of-physical-range (→ checks memory corruption, threading issues, etc.)
    - data packets arriving too early, too frequent, late, or not-at-all
    - settings inconsistency (BPF, ORM-recomputation, ...)

Beam Measurement Quality Checks
Detecting Erroneous BPMs I/II

- Three main lines of defense against BPM errors and faults:
  1. Pre-checks without beam using the in-build calibration unit
     - eliminates open/closed circuits, dead circuits/element candidates
  2. Pre-checks with Pilot and Intermediate beams
     - verifies calibration offset and slope (golden orbit)
  3. Continuous data quality monitoring through Orbit Feedback Controller
     - detects spikes, steps and BPMs that are under verge of failing

→ this is an area where improvements are possible

Details at:
http://lhccwg.web.cern.ch/lhccwg/Meetings/2007/2007.10.23/2007-10-23_LHCCWG-FAULTY_BPM.pdf

Beam Measurement Quality Checks
Detecting Erroneous BPMs II/II

- Orbit feedback procedure in case of a
  - spike: fail-safe choice of assuming that orbit is at reference position
  - step: pause feedback, average orbit before and after detected step and continue from new averaged orbit (N.B. not tested/commissioning at LHC)

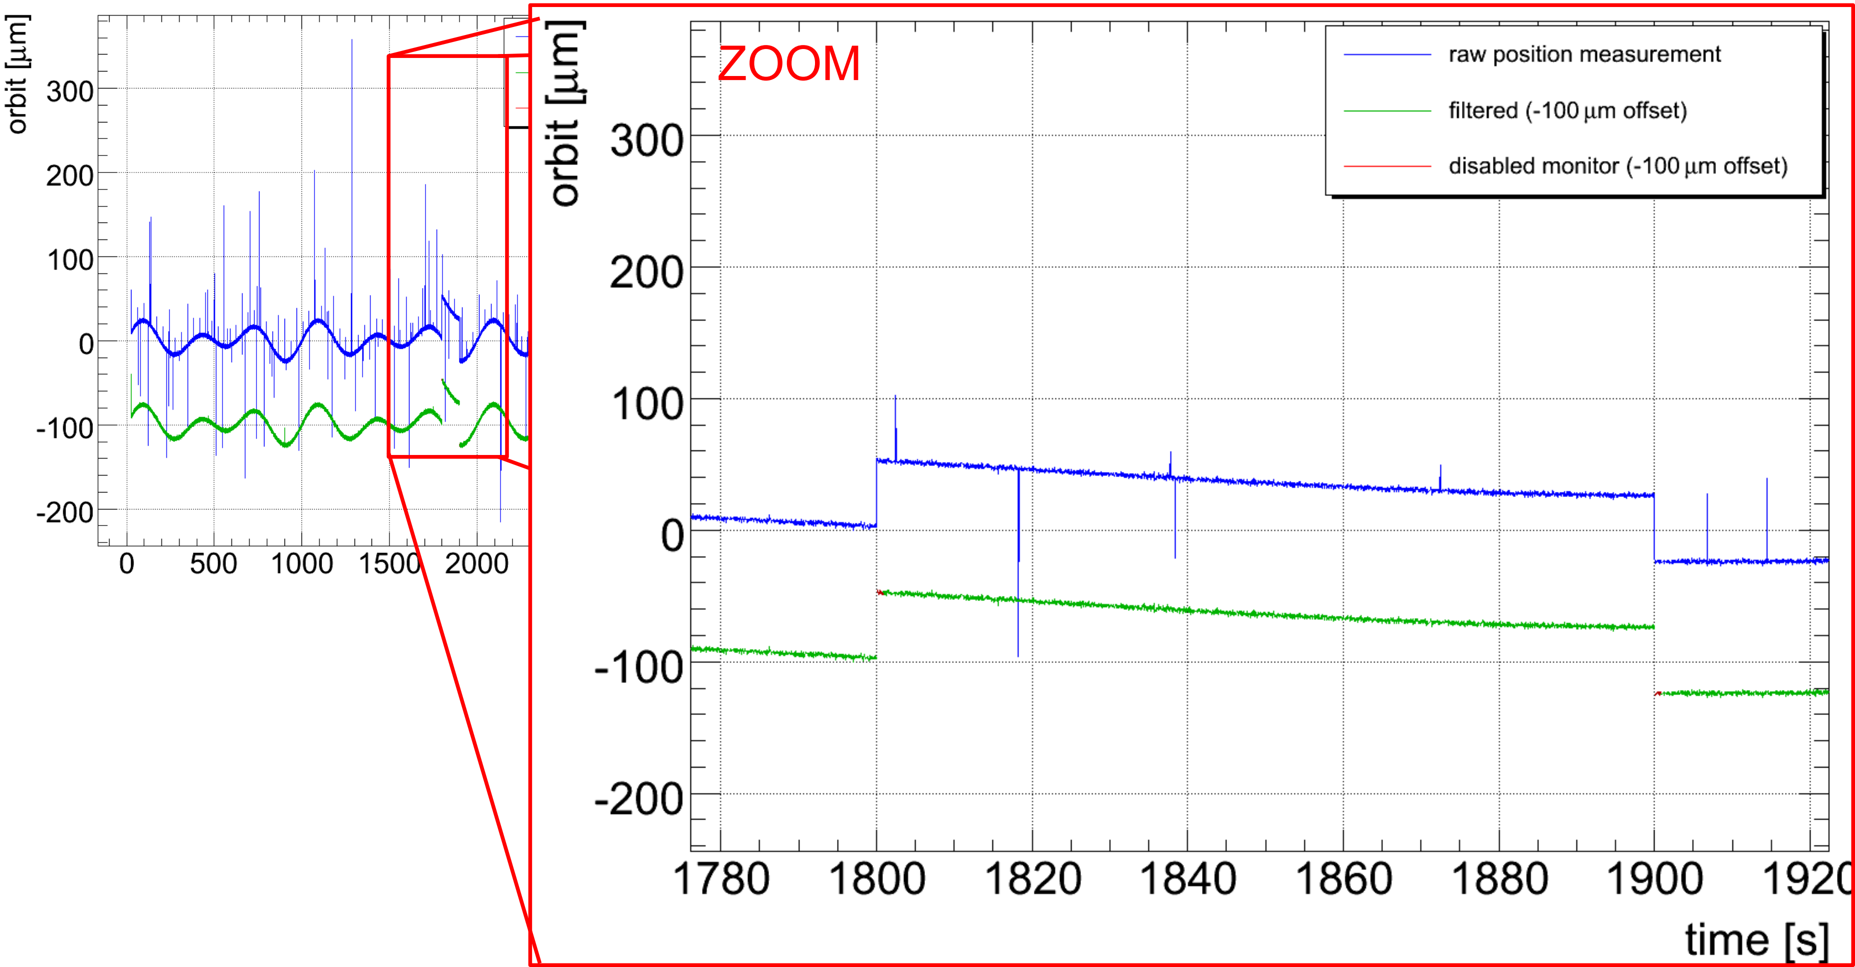

Beam Measurement Quality Checks
Q/Q' Signal Quality Checks I/II

- Q/Q' not a direct beam observable → highly non-trivial detection and tracking
  - strong dependence on beam intensity, filling pattern, particle species, RF settings, ADT, operational mode, …, many cross-links/interferences
  - Subject to many reviews, details in DIPAC'11 & CERN-BE-2011-016

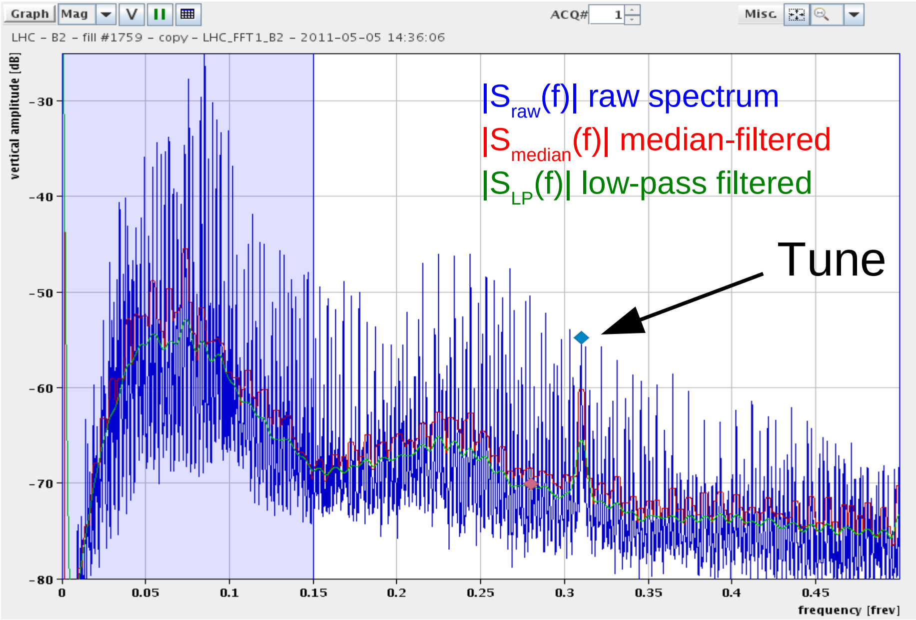

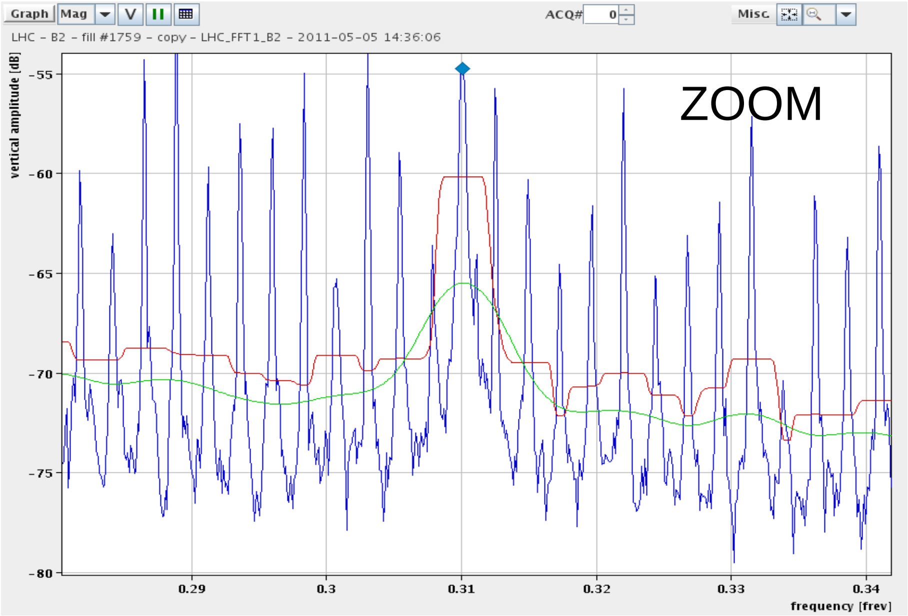

- Of note:
  1. Strong impact on operation, mitigated by detecting S/N deterioration
     → Tune-FB stops to avoid converting noisy/erroneous meas. into real tune drifts
  2. function distributed between BQBBQLHC FESA class and OFC

Beam Measurement Quality Checks
Q/Q' Signal Quality Checks II/II

- Motivation to split Q/Q' logic between BQBBQLHC and OFC:
  - high numerical complexity (tune tracker function limited by FE-CPU performance)
  - High-data throughput → BQBBQLHC performs necessary data reduction
  - More flexible w.r.t. Q/Q' source and tracker modification/deployment (developments still ongoing and plan to upgrade tracker during LS1)

- BBQ FE (8 systems) critical to feedback operation beside being Q/Q' source
  - redistributes timing telegram information (BPF, Energy, Modes, …): enables 'majority voting', thus eliminating single-point-of-failure in OFC
  - Sends data to OFC as well as other “high-performance CCC clients” → interoperability between non- and RT communication most critical

- Main issue: MEN-A CPU running at 80% load, blocking further improvements
  - ie. tracking of multiple peaks to disentangle between mains harmonics, Q_s, interferences etc.

Integrity checks & OFC Changes during LHC Operation I/II

Often “simple” checks to mitigate or fix day-to-day problems, e.g.;

- memory corruption, uninitialised input from systems feeding into the OFC
  → Check every variable against 'NaN', '±Inf', out-of-physically-range (2011)
  (distributed in every module, some performance penalty)
- Energy telegram “noise”, '0000'/'FFFF' → limited dE/dt + majority voting (2011)
- Switching FBs 'OFF→ON' without beam → introduced BPF check
- Tune-FB 'OFF→ON' without Q signal → stability & S/N check
- Noisy Q signal QPS trips → stability and dQ/dt rate-limit checks
- BPM temperature comp. (sensor glitch→orbit glitch) → 'dabTemp' check
- Q_{rev} outside Q acceptance window → stops Tune-FB
- OFSU sending zero-iORM → check/reject/reconstruct iORM
- disabled real-time channels → disables FBs, Q' tracker, …
- FB gains set to zero → check/reject setting
- Switching FBs 'OFF→ON' with fully coupled machine → check $|C^-|$/switch off

Integrity checks & OFC Changes during LHC Operation I/II

- BST/BOBR timing glitch → check/compare inconsistencies for bunch-Count and errorCount
- Lost one beam during ramp → de-coupled/fast re-computation of ORM for single beam operation
- duplicate BBQ/BQS input data → detect and reject
- BPM left in intensity/calibration mode → reject
- TechNet/front-end IT security scans → improved socket and added IP firewall

[..]

- Most of the (in 2006) specified OFC design functionality in place
  - Still, a lot of sanity-checks and initially unspecified code added during four years of operation → should be preserved for post-LS1 operation

- A quite educative point of view on that matter (courtesy Quentin):
  Joel Spolsky, “Things You Should Never Do” (2000)
  http://www.joelonsoftware.com/articles/fog0000000069.html

Specific Architecture 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)
  → Need to improve FB sub-systems to keep up with LHC progress post-LS1

- Feedback architecture, strategies and algorithms are well established
  - same architecture for orbit, energy, Q/Q' and betatron coupling correction
  - OFC split into two domains
    - space-domain (accelerator physics)
    - time-domain (control theory)
      - nested PID, with internal Smith-Predictor and anti-windup loops
  - cross-talk between multiple simultaneously running feedback loops minimised by central controller design and master-slave topology

- Any feedback loop is only as strong as the weakest link
  - Quality and integrity of beam parameter measurement
  - Robustness of front-end and controls infrastructure
  - Loop margin of given settings & correct working point (gains, optics, reference)

Appendix

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

- 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 periodic constant load
    - 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

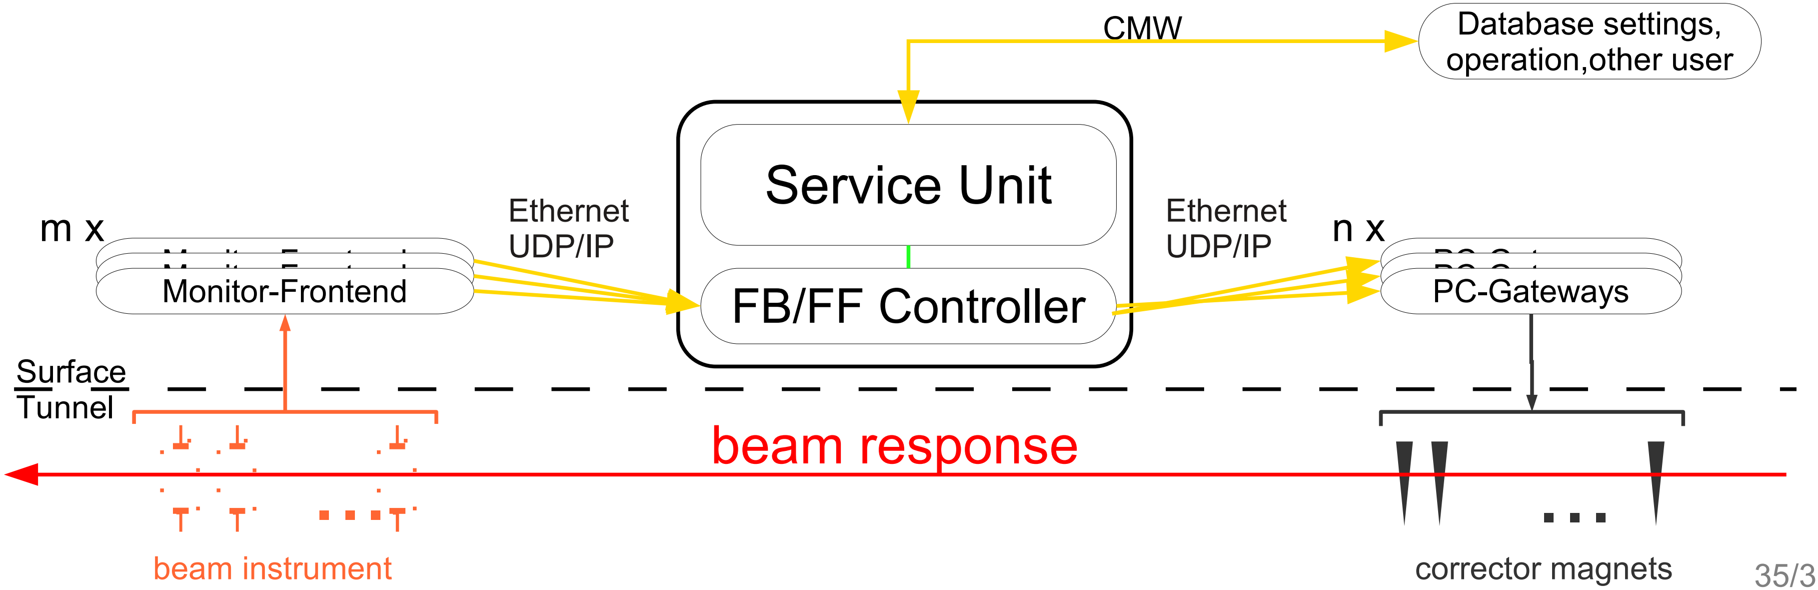

Inter-loop Cross-Dependencies III/III
Cascading between individual Feedbacks

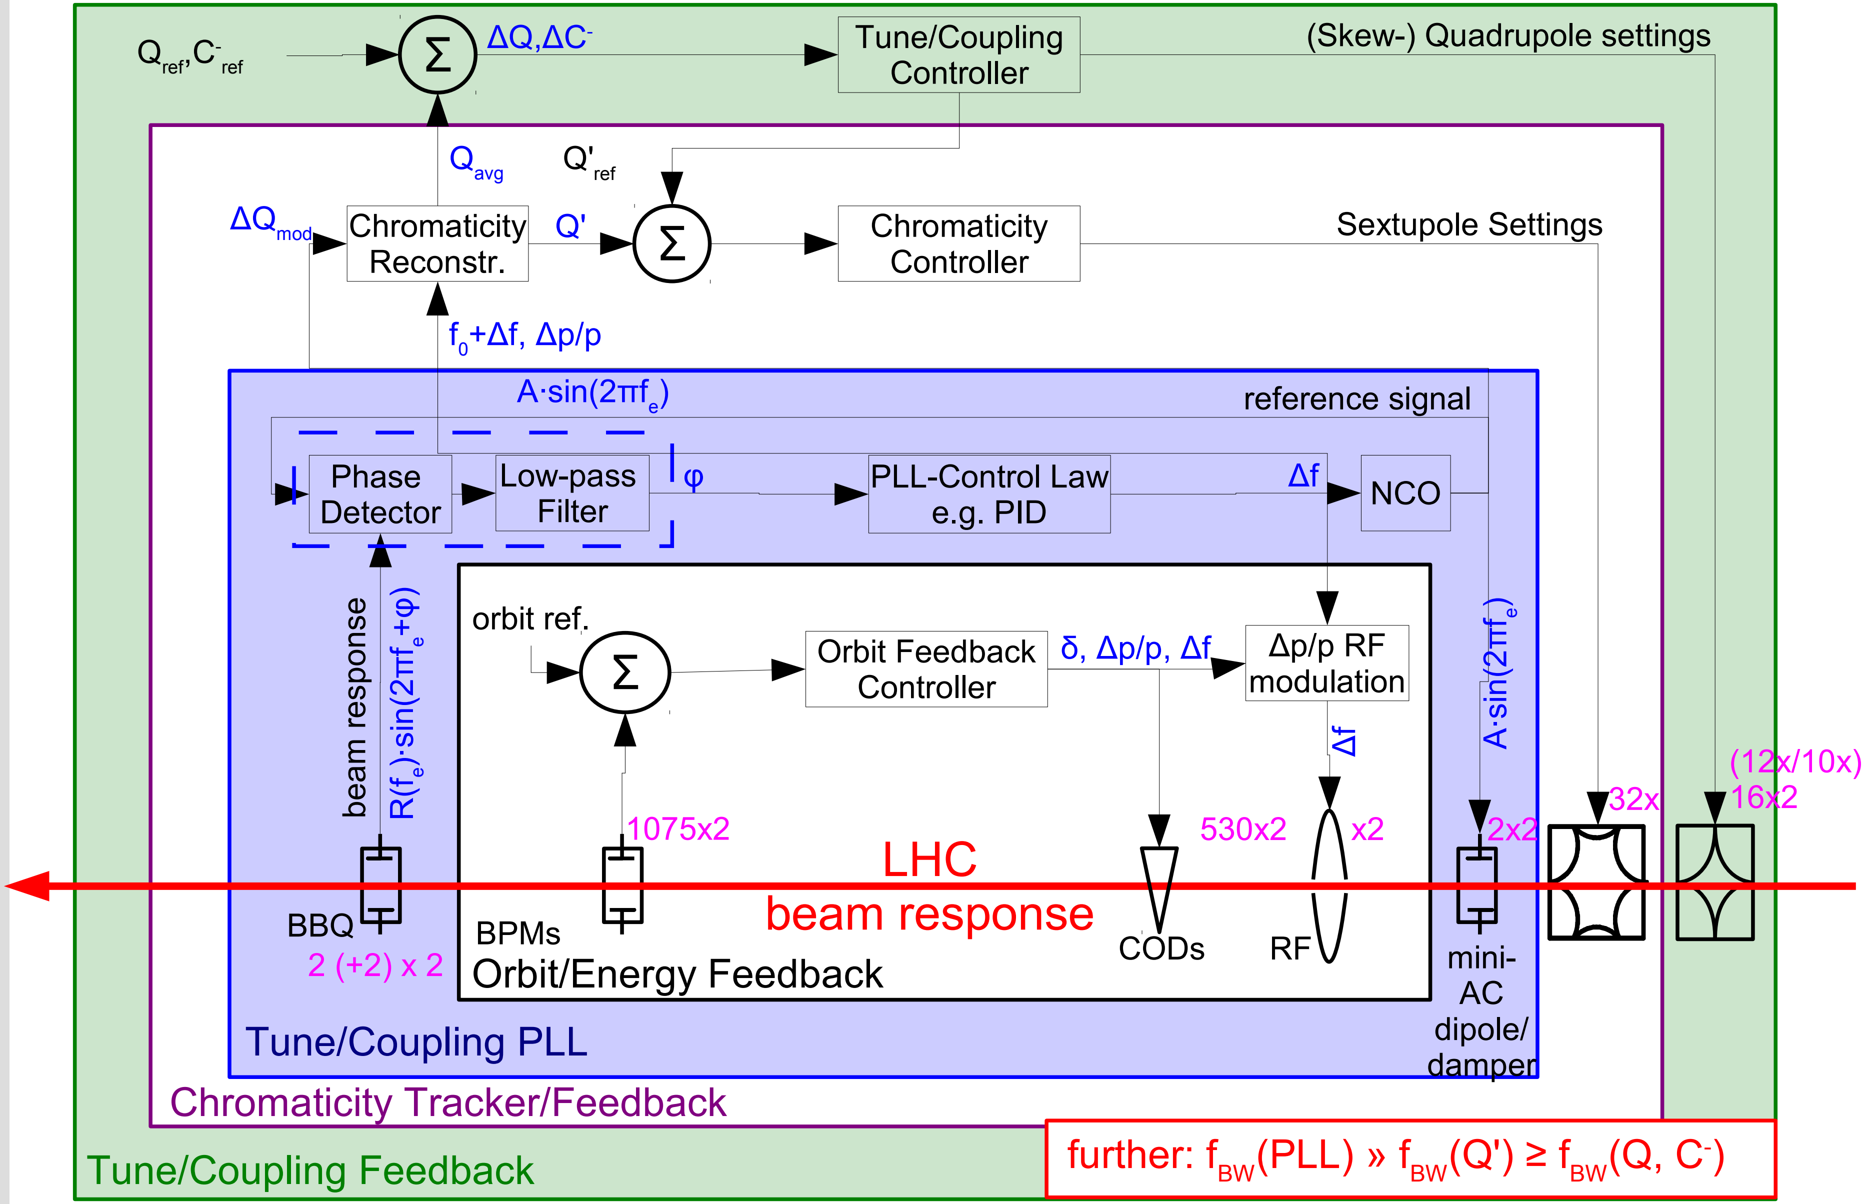

further: $f_{BW}(PLL) \gg f_{BW}(Q') \geq f_{BW}(Q, C^{-})$

Total Number of (FB) Corrector Circuits

- Orbit: 530 correction dipole magnets/plane (71% are of type MCBH/V, ±60A)
  - total 1060 individually powered magnets (60-120 A)
  - ~30 shared between B1 & B2

- Tune:
  - 16x ±600A circuits powered from even IPs (2, 4, 6, 8), 2 families
  - independent for Beam 1&2, but coupling between planes
  - can use them independently, optional use of DS quadrupoles

- Chromaticity:
  - 32x ±600A circuits powered from even IPs,4 families (ΔQ'~1→1A@7TeV)

- Coupling: four skew quadrupoles per arc, 1/2 families
    - Beam 1: 12x ±600A
    - Beam 2: 10x ±600A

- Total: 1130 of 1720 circuits/power converter → more than half the LHC is controlled by beam based feedback systems!

Total Number of (FB) Corrector Circuits
Powering Layout of the SSS Correction Scheme IP4↔IP5

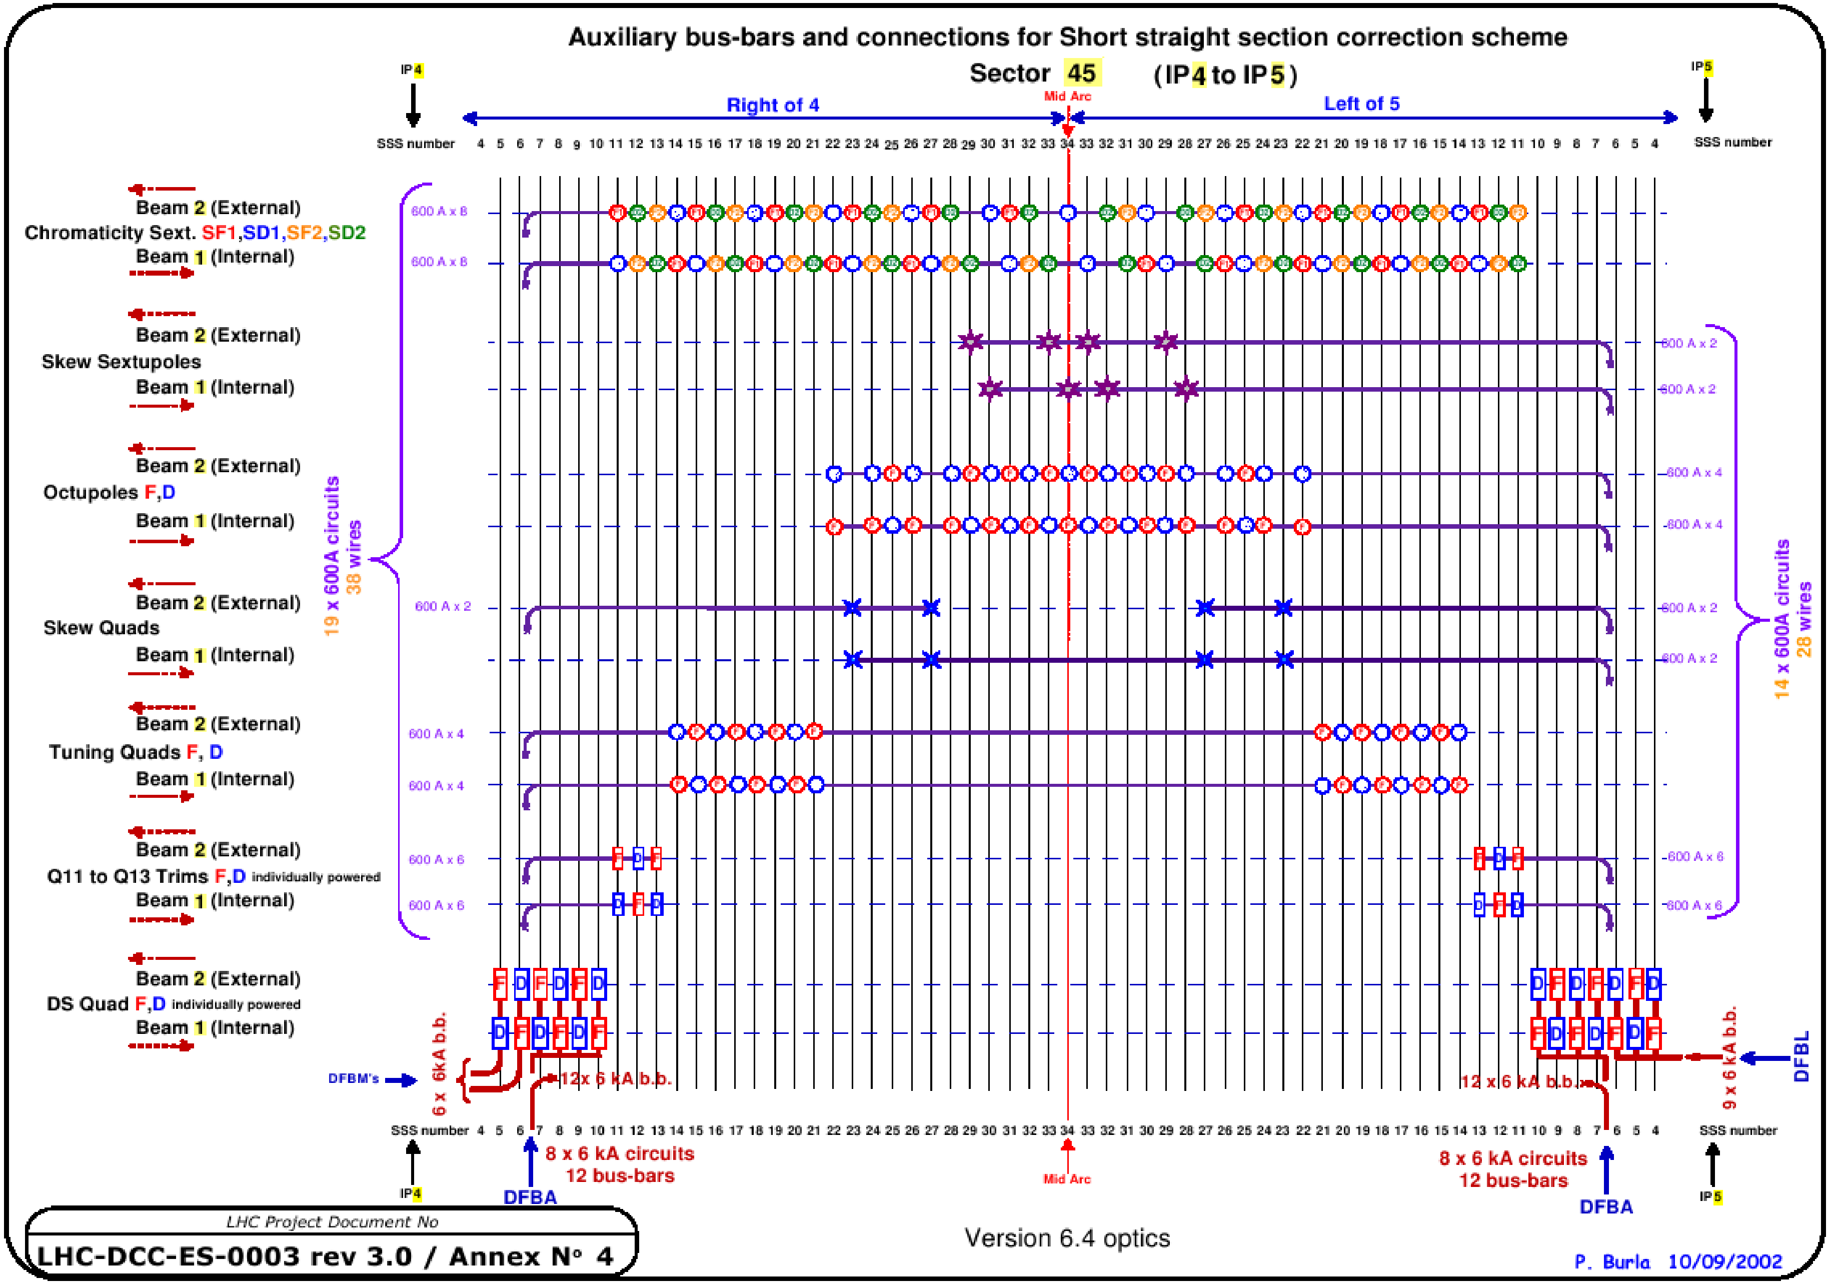

Network Traffic In and Out I/II

- typical '/usr/sbin/iftop' diagnostics on cs-ccr-ofc:

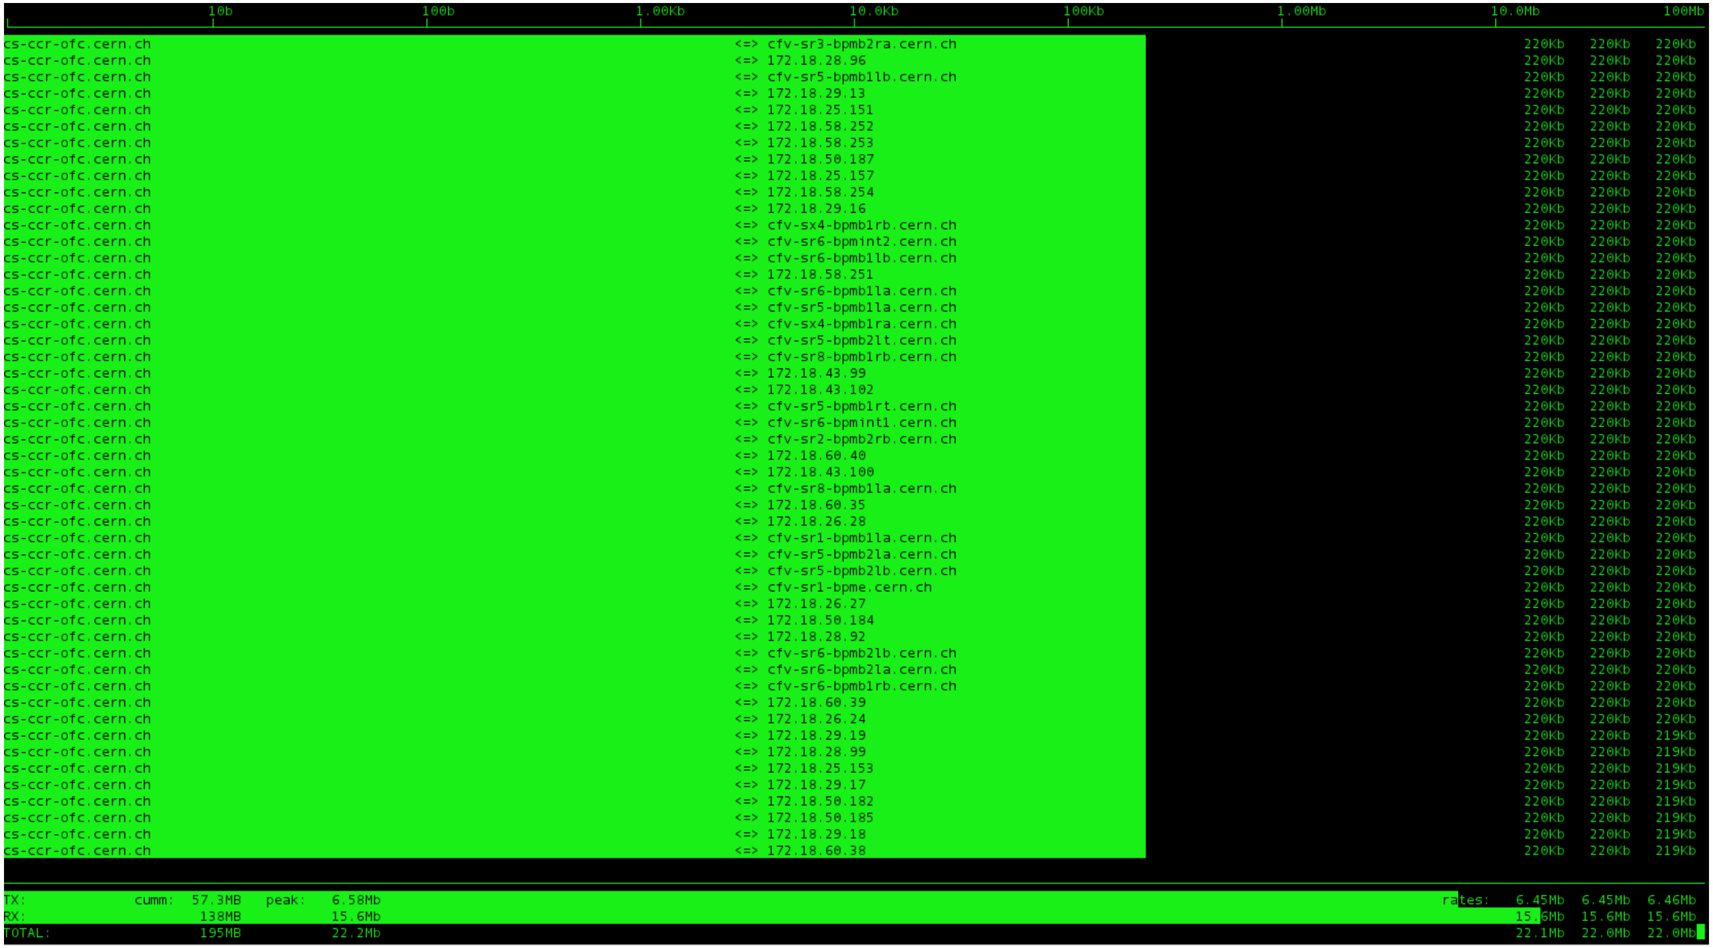

```
                    10b                 100b                1.00Kb              10.0Kb              100Kb               1.00Mb              10.0Mb            100Mb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sr3-bpmb2ra.cern.ch                                  220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.28.96                                             220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sr5-bpmb1lb.cern.ch                                  220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.29.13                                             220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.25.151                                            220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.58.252                                            220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.58.253                                            220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.50.187                                            220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.25.157                                            220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.58.254                                            220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.29.16                                             220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sx4-bpmb1rb.cern.ch                                  220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sr6-bpmint2.cern.ch                                  220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sr6-bpmb1lb.cern.ch                                  220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.58.251                                            220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sr6-bpmb1la.cern.ch                                  220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sr5-bpmb1la.cern.ch                                  220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sx4-bpmb1ra.cern.ch                                  220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sr5-bpmb2lt.cern.ch                                  220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sr8-bpmb1rb.cern.ch                                  220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.43.99                                             220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.43.102                                            220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sr5-bpmb1rt.cern.ch                                  220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sr6-bpmint1.cern.ch                                  220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sr2-bpmb2rb.cern.ch                                  220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.60.40                                             220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.43.100                                            220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sr8-bpmb1la.cern.ch                                  220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.60.35                                             220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.26.28                                             220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sr1-bpmb1la.cern.ch                                  220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sr5-bpmb2la.cern.ch                                  220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sr5-bpmb2lb.cern.ch                                  220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sr1-bpme.cern.ch                                     220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.26.27                                             220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.50.184                                            220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.28.92                                             220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sr6-bpmb2lb.cern.ch                                  220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sr6-bpmb2la.cern.ch                                  220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> cfv-sr6-bpmb1rb.cern.ch                                  220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.60.39                                             220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.26.24                                             220Kb  220Kb  220Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.29.19                                             220Kb  220Kb  219Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.28.99                                             220Kb  220Kb  219Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.25.153                                            220Kb  220Kb  219Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.29.17                                             220Kb  220Kb  219Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.50.182                                            220Kb  220Kb  219Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.50.185                                            220Kb  220Kb  219Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.29.18                                             220Kb  220Kb  219Kb
cs-ccr-ofc.cern.ch                                                    <=> 172.18.60.38                                             220Kb  220Kb  219Kb

TX:             cumm:  57.3MB   peak:   6.58Mb                                                        rates:   6.45Mb  6.45Mb  6.46Mb
RX:                     138MB           15.6Mb                                                                 15.6Mb  15.6Mb  15.6Mb
TOTAL:                  195MB           22.2Mb                                                                 22.1Mb  22.0Mb  22.0Mb
```

- very powerful tool, scroll up/down with 'k' and 'j', 'L' toggles logarithmic display, 't' toggles in/out traffic display, 'h' for help and advanced port/DNS display
- Healthy state: all BPMs, FGC Gateways send with the same data rate

Network Traffic In and Out II/II

- Healthy OFC-OFSU communication:

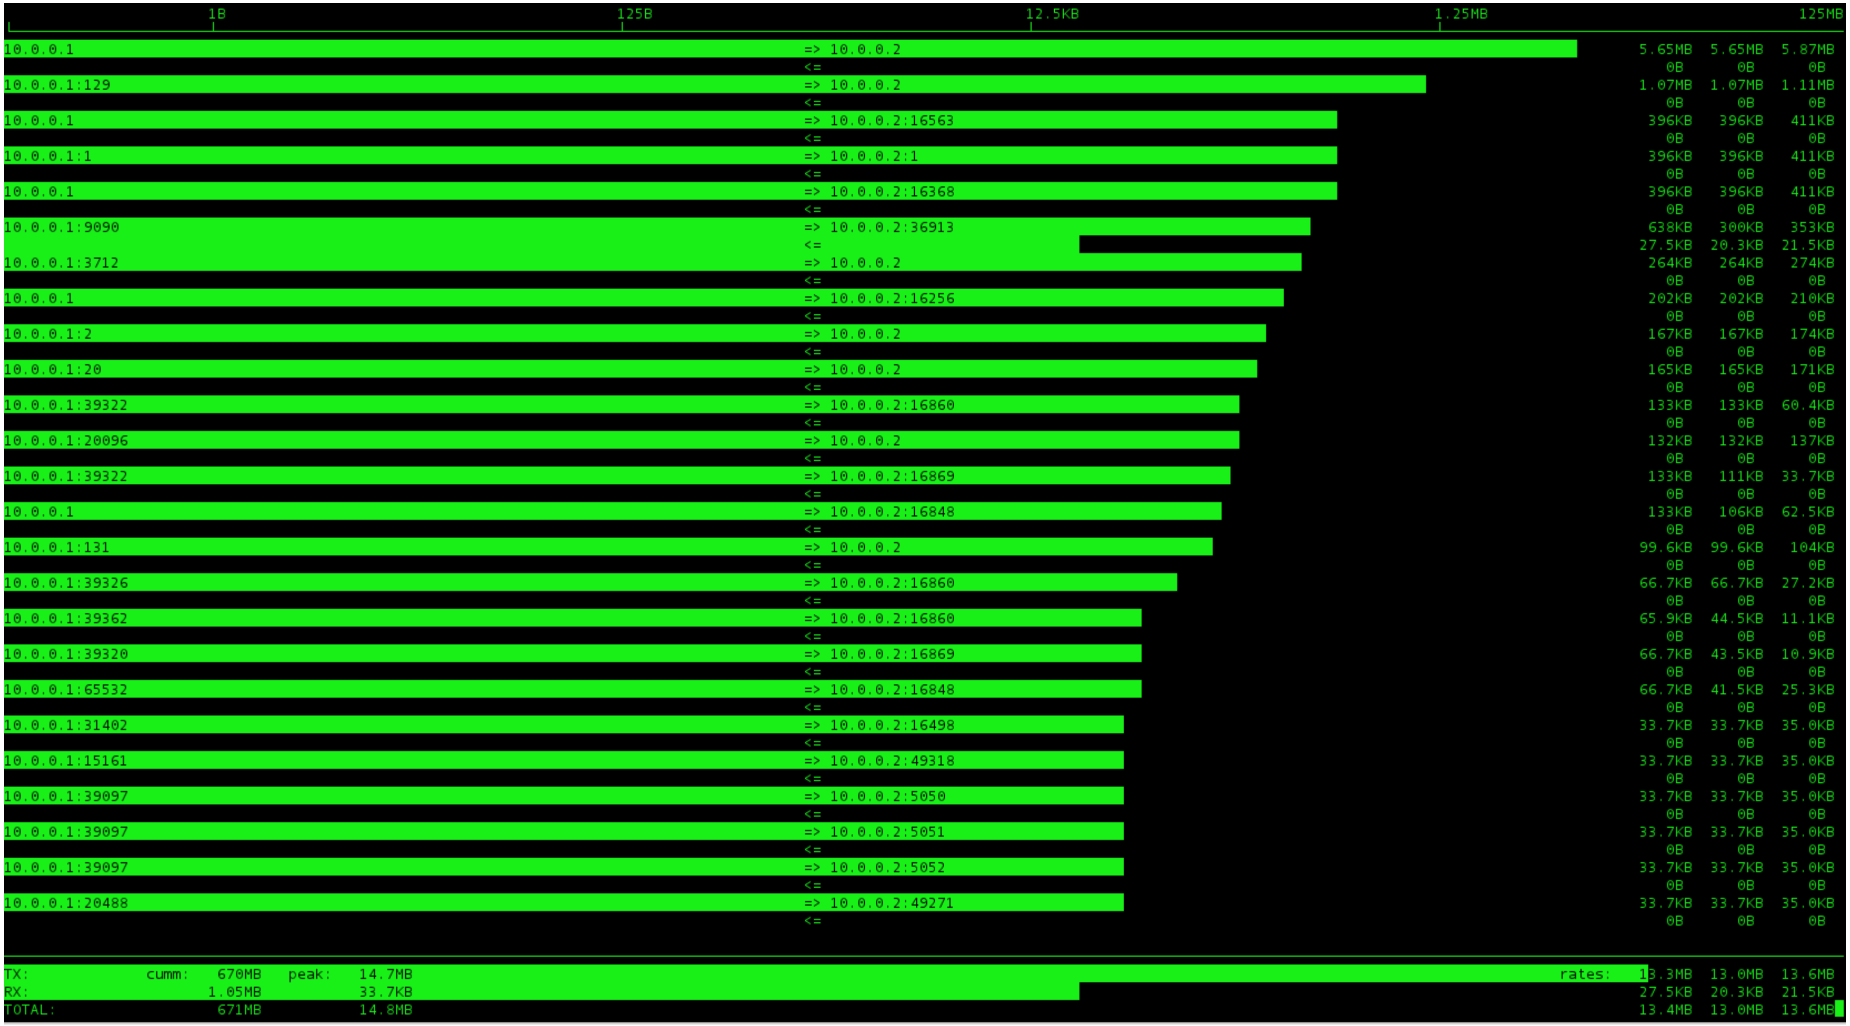

```
                    1B                                      125B                                    12.5KB                                  1.25MB                                  125MB
10.0.0.1                                                                         => 10.0.0.2                                                                           5.65MB  5.65MB  5.87MB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1:129                                                                     => 10.0.0.2                                                                           1.07MB  1.07MB  1.11MB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1                                                                         => 10.0.0.2:16563                                                                      396KB   396KB   411KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1:1                                                                       => 10.0.0.2:1                                                                          396KB   396KB   411KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1                                                                         => 10.0.0.2:16368                                                                      396KB   396KB   411KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1:9090                                                                    => 10.0.0.2:36913                                                                      638KB   300KB   353KB
                                                                                 <=                                                                                    27.5KB  20.3KB  21.5KB
10.0.0.1:3712                                                                    => 10.0.0.2                                                                            264KB   264KB   274KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1                                                                         => 10.0.0.2:16256                                                                      202KB   202KB   210KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1:2                                                                       => 10.0.0.2                                                                            167KB   167KB   174KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1:20                                                                      => 10.0.0.2                                                                            165KB   165KB   171KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1:39322                                                                   => 10.0.0.2:16860                                                                      133KB   133KB  60.4KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1:20096                                                                   => 10.0.0.2                                                                            132KB   132KB   137KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1:39322                                                                   => 10.0.0.2:16869                                                                      133KB   111KB  33.7KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1                                                                         => 10.0.0.2:16848                                                                      133KB   106KB  62.5KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1:131                                                                     => 10.0.0.2                                                                           99.6KB  99.6KB   104KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1:39326                                                                   => 10.0.0.2:16860                                                                     66.7KB  66.7KB  27.2KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1:39362                                                                   => 10.0.0.2:16860                                                                     65.9KB  44.5KB  11.1KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1:39320                                                                   => 10.0.0.2:16869                                                                     66.7KB  43.5KB  10.9KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1:65532                                                                   => 10.0.0.2:16848                                                                     66.7KB  41.5KB  25.3KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1:31402                                                                   => 10.0.0.2:16498                                                                     33.7KB  33.7KB  35.0KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1:15161                                                                   => 10.0.0.2:49318                                                                     33.7KB  33.7KB  35.0KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1:39097                                                                   => 10.0.0.2:5050                                                                      33.7KB  33.7KB  35.0KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1:39097                                                                   => 10.0.0.2:5051                                                                      33.7KB  33.7KB  35.0KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1:39097                                                                   => 10.0.0.2:5052                                                                      33.7KB  33.7KB  35.0KB
                                                                                 <=                                                                                       0B      0B      0B
10.0.0.1:20488                                                                   => 10.0.0.2:49271                                                                     33.7KB  33.7KB  35.0KB
                                                                                 <=                                                                                       0B      0B      0B

TX:             cumm:   670MB   peak:   14.7MB                                                                                                               rates:    13.3MB  13.0MB  13.6MB
RX:                    1.05MB           33.7KB                                                                                                                         27.5KB  20.3KB  21.5KB
TOTAL:                  671MB           14.8MB                                                                                                                         13.4MB  13.0MB  13.6MB
```

- Alternatively: 'netstat -Natn' and 'netstat -Naun' on cs-ccr-of[c/su] indicate if the network sockets are overloaded (via Recv-Q Send-Q)