131st Meeting of the Machine Protection Panel

Participants (incomplete): W. Bartman, R. Bruce, A. Butterworth, B. Dehning, M. Christof Poeschl, M. Deile, E.B. Holzer, A. Lechner, N. Magnin, Y. Nie, A. Siemko, M. Valette, D. Valuch, D. Wollmann, M. Zerlauth.

The slides of all presentations can be found on the website of the Machine Protection Panel:
http://lhc-mpwg.web.cern.ch/lhc-mpwg/

1.1 Approval of MPP#130's minutes

- No Actions from 130th MPP.
- No additional comments were received on the minutes; they are therefore considered approved.

1.2 Vertical orbit kicks due to heater firing in main dipoles after beam induced and training quenches (M. Valette)

- On several luminosity fills, during training or beam induced quenches, losses are recorded in IP7 a few milliseconds before the dump. These losses were associated with a small (few tens of µm) vertical orbit oscillation. The quench heaters (QH) firing is thought to be the cause of this.
    - Anton commented he remembered delays of tens of ms for the effectiveness of quench heater firings?
        - The delay between the quench and the QH firing thus between the beam losses and these periodic losses is in this order of magnitude because of the quench discrimination time.
    - Andrzej was surprised the beam is not dumped before the QH are fired.
        - The difference of delays (mentioned by Reiner Denz) between the QPS firing the QH and the QPS opening of the BIS loop can be as long as 5.5 ms. In case of a spurious QH firing the time after current in the heater and the beam dump could be significantly longer.
- The skew dipole magnetic field created by the QH was simulated by L. Bortot, the field from the high field QH is much stronger than the one from the low field QH. In the LHC low field QH are only connected in three MB magnets. The skew dipole (Bx) component reaches 0.7 mT in the beam region with a small sextupole component which was neglected during the studies.
- MADX simulations confirmed that the QH field distortion is indeed the cause of the orbit oscillation and the losses. Since the kick is dependent on beta function and the phase, a systematic study of the beam displacement in the vertical TCP as a function of the MB QH fired was conducted. The maximum beam displacement should be 0.15σ which is not worrying.

- Other magnets will be studied, e.g. for the triplet circuits all QHs on all 4 magnets would be fired at the same time. The superconducting D1, D2 and triplet for HiLumi have a lot of QH circuits with adding magnetic field distortions in the current baseline.
    - Andrzej commented that since we are in an early stage of the HiLumi project the QH powering configuration should be chosen such to minimize the effect of a heater firing on the beam. Furthermore the worst-case connection scenarios should be studied for HL-LHC and the current LHC.

1.3 LHC vs SPS re-phasing (local references) and LHC injection (A. Butterworth)

- There were issues with the synchronisation between the SPS and the LHC on May 5th. Synchronisation requires the RF frequency and the common revolution frequency, which is 7 LHC turns or 27 SPS turns. On 5th of May a single nominal was shot on the TDI because the SPS was on local frequency instead of the LHC's.
- Three levels of mitigations were decided:
    - Improved sequencer checks: the LHC sequencer verifies the two clocks. Originally it was designed to verify only one user, which opened a loophole, when the tage nominal was introduced in the CBCM, which allows 2 cycles with different users in the same SPS super-cycle. Now all users tagged "to_lhc" are checked by the sequencer.
    - SPS clock selection: Checks if the LHC has mastership over the SPS. The operator is not allowed to switch to local, if the LHC is the master. If the LHC is not the master, it is allowed to switch to local clock.
    - Faraday cage HW check of the status of the clock selector: Not yet implemented. A FESA class checks at various points in the cycle the status of the clock. If the FESA class fails, a timeout sends a signal to dump the beam. If the SPS timing is not on LHC reference and destination is the LHC, the beam is dumped.
    - 3 checks during cycle: before injection, at start of ramp, before extraction to avoid dumping at top energy in the SPS.
        - Daniel asked if there is a possibility the selector switches in-between the three checks?
        - Andy answered it is extremely unlikely; this is why it is preferred to check the HW status.
- Conclusion: three levels of mitigation were proposed, two were already implemented and the third is going to be implemented very soon.
    - Daniel asked when the implementation of the third one is envisioned.
        - A. Butterworth replied that it is to be expected in a couple of weeks, as relevant people are on holiday now.
    - Daniel commented it is a nice solution to have three layers, you would need multiple failures to reproduce problem.
    - Andy added it was luck that it was a single nominal, it could have been a full batch.

  - Daniel acknowledged but said it is a foreseen failure case in the design and the passive protection should therefore survive it.

AOB - Issue with creation of PM and mitigation (M. Christof Poeschl)

- There was recently an issue with incomplete Post Mortem events being built, data from some devices were not included in the PM event analysis (although sent to the PM servers). A list of the affected events is available on the slides.
- The current architecture relies on two identical servers, on which data are physically stored twice, in the backend and front end. RDA timeouts lead to missing copies from the FE to the BE server, where the latter are used by the analysis servers.
- Investigation yielded the following: at times when it happened the drives were in high load, with I/O wait times in the order of milliseconds. Since it is based on synchronous calls, long waiting times pose a problem.
- The data from these events were always written to the FE disks, hence no data were ever lost. Incomplete events were (and can always) be reconstructed offline. Back-up processes wrote and read at the same time which are believed to have created the wait.
- Three actions were proposed:
  - Updating the firmware of the Hard Disks Drives
  - The drives were mixed SAS and SATA, they were consolidated to all new SAS drives
  - Reprioritize processes, backend and fronted storage will be firsts and backups last.
- The reprioritization and switch of hardware were done. What is left is to update HDD firmware and modified power settings.
- Last week a drive assigned to the frontend storage failed, the data were reconstructed.
  - Daniel asked if three days later some statistics on the server load and failures were available.
    - It is not the case.
  - Daniel commented it is always reassuring to know the data could be reconstructed.

AOB - ADT firmware update and revalidation before MD1 (D. Valuch)

- New features, presented to 130th MPP, have been implemented in the firmware code; the team is now waiting for an occasion install the new firmware in the LHC and to revalidate with beam. The newest features were motivated by MD requests; these MDs were attributed to MD block 2, thus, the installation and following revalidation should be performed before.
- Before that, the ADT team was frequently solicited for MDs but it was mostly about online-modifications of the operational system. The goal now is to switch to cleaner implementations along with machine critical settings
- There are new excitation modes and features:
  - AC dipole like excitation.
  - The limitation on excitation length within a turn was improved.

  - Decoupling of cleaning and excitation channels, in order to avoid change of cleaning during an MD.
  - New gigabit standards for the BPMs
- Firmware was recompiled and has to be uploaded and revalidated
  - Wolfgang asked if it was possible to gate on a number of bunches.
    - Yes, one can now select individual bunches.
  - Wolfgang commented the derivative of the excitation should be smooth as well.
    - The implementation was done as requested by OP. This will be protected, the people who can access the ac dipole functionality will be stored in a register and need access to a high level application.
- Test and validation of the unrestricted excitation mode should also be conducted. All abort gap cleaning classes underwent major changes and need full revalidation, the perfect time would be the intensity ramp up after MD block 1.
- It is proposed to perform the firmware update just after the MD block. This should take a few hours in the shadow of the high intensity ramp up, with a small train of 12b for an hour.
  - D. Wollmann asked if MDs in block one will not required any of the new functionalities.
    - No, these are necessary updates for MD block 2 only.
    - D. Wollmann commented that the validation tests with beam should be documented. Like this, the responsibles for the MDs know the list of available modes.
  - D. Valuch added that the AC dipole functionality should remain a very expert tool, and never released to the general public.
  - D. Wollmann concluded an appropriate time window will be requested next week to the LHC coordinators in line with the operations for validation of the changes.
    - D. Valuch added some settings' time windows changed and need to be propagate to the corresponding beam processes.
  - Jamie Boyd mentioned there is a 2.5km optics background measurement planned, and asked where does this fit in the program?
    - D. Wollmann answered the 2.5km tests will be done before the intensity ramp up. D. Valuch must be present during the validation cycle, which should probably best be performed after the 2.5km beta* optics tests.
  - Andrzej asked how the decoupling between cleaning and excitation was done.
    - The 4 signals paths are recombined in analog and driving the amplifiers. The cleaning path will not need to be touched anymore, when using the excitation features to avoid a mix up of parameters.
  - Wolfgang asked who is deciding the sequence of tests for the restart after MD block one.

- rMPP will establish list of requested validation fills. In case collimation or the dump team require revalidations these will be added.
- Roderick added they will not request a special test.

AOB - Follow-up of the reduction of the abort gap keeper window (N. Magnin)

- This presentation adds technical details to the LMC talk on Wednesday.
- The Abort Gap Keeper (AGK) inhibits any injection that could lead to beam in the abort gap (previously hard-coded to be 288b long). The longest injection now is 144b so we have an unused 144b space at the end of the circumference, which allows for a shorter AGK window. It is needed soon, any change after EYETS would be without gain (as the beam dump should be replaced by then).
- The known risks:
  - Too short AG: can be detected and cleaned
  - Injection of longer train: MKI max length was lowered
- Timeline: tests of HW and SW is done, deployment on Monday before MD1 then testing for a week during the MDs.
- Tests were done in the lab with signals generator, the outcome confirmed the expected results.
  - The CTRV card used for the implementation is a standard card at CERN and will take the AGK signal and generate one turn later a shorter AGK window.
  - A delay of 40ns is added because of this card, there is a 25ns jitter for the falling edge, the rising edge jitter is still < 15ns, same as before for a 100Mhz system, it just reproduces the one from TSU (input signal).
  - There was no trigger missing after 24h of tests.
- The monitoring is checked with a FESA class, a script locks the drive, it can then no longer change settings.
  - Wolfgang commented in case of a power cut, there is no check of the value saved on the card.
    - This issue is known and possible solutions are under investigations.
  - D. Wollmann asked what happens in case of a reboot.
    - It takes the value of the LTIM database.
- Foreseen future improvement, maybe during TS2 or the EYETS: inhibiting SPS extraction in case of an AGK problem.
- Conclusion: the problem of CTRV driver lock up is still remaining, but it seems to be possible to read the counter configuration through FESA. IQC could verify the AGK length after every injection. In case of AGK instabilities the FIB could interlock and block injection.
  - Wolfgang said it would be good to validate the changes with a first low intensity fill before moving back to.
  - Jamie said one could implement it without using the extra space at first.
  - D. Wollmann concluded the MD block is a good validation time and commented it is a good solution and luminosity will profit from that. After a full validation of functionalities with and without beam, full intensity physics fills can be performed.