71st Meeting of the Machine Protection Panel

Participants: A. Apollonio, T. Baer, D. Belohrad, C. Bracco, F. Burkart, V. Chetvertkova, B. Dehning, M. Deile, E.B. Holzer, R. Jacobsson, E. Nebot, C. Martin, B. Mikulec, G. Papotti, B. Puccio, R. Schmidt, D. Wollmann, J. Wozniak, C. Zamantzas, M. Zerlauth

1 Presentations

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

1.1 MPS Performance and Issues 2012 – (D. Wollmann)

- Daniel shows the statistics of the protection dumps 2012 (in total 536 up to the 8th of November) and compares the distribution of the dump causes to the previous years (2011, 2010):
  - 2012: 35% Not beam induced dumps, 29% Operator dumps, 14% dumps due to beam monitoring, 14% False dumps from MPS, 8% MP tests.
  - 2011: 41% Not beam induced dumps, 27% Operator dumps, 8% dumps due to beam monitoring, 14% False dumps from MPS, 10% MP tests.
  - 2010: 36% Not beam induced dumps, 23% Operator dumps, 14% dumps due to beam monitoring, 12% False dumps from MPS, 15% MP tests
- It can be seen that the number of operator dumps slowly increased over the past years. The increased number of dumps due to beam monitoring compared to 2011 can be explained by beam losses at the end of the ramp and transvers instabilities during the squeeze. The share of UFO-dumps has significantly decreased, although the actual numbers stayed about the same.
- MPS false dumps: As in the previous years the QPS has the biggest share in the false dumps due to MP systems (70%). The dumps due to failures in

the BLM system have significantly increased compared to 2011 (6% -> 20%). Dumps due to failures in the LBDS have been reduced since 2010 from 22% to 5% in 2012.

- All protection dumps are stored in the post mortem database. The dumps above 450GeV are (in addition to the operator entries) classified and commented by MPS experts.
- In addition MPS check lists are regularly distributed by rMPP, to document issues in the different MP systems in the reference period. During the intensity ramp up a check list needs to be filled by the MPS owners before every step in intensity.
- The final check lists are stored in EDMS.
- Daniel shows the top three issues, which appeared in the machine protections and related systems:
    - Reference problems of the Orbit Feedback System Unit (OFSU). In April it was discovered that the reference of the orbit feedback was set to '0' during the squeeze, which caused significant orbit changes in the IRs. The beams were dumped due to losses on the tertiary collimators right of IR2. Due to this issue the next step in intensity increase was postponed until the mitigation via a SIS interlock was introduced and it's correct functioning sufficiently tested. In addition sequencer checks and checks by the operators were introduced. Therefore the OFSU problems, which sometimes still appear were reduced to availability issues not machine safety issues.
    - In April a weakness in the power supply of the LBDS was discovered, which would have caused an asynchronous beam dump, if beam had been present at its appearance. During the following analysis a common mode failure point in a 12V DC power system in the TSU was discovered. In the case of this failure no beam dump would be possible. Following this discovery the beam operation was stopped until a watch dog was deployed, which would force an asynchronous beam dump in case of the common mode failure in the 12V DC powering.

  - One machine protection relevant application for the BSRTs is the monitoring of the abort gap population. Due to RF heating the BSRTs in both beams deteriorated slowly over time. End of August the speed of deterioration of the B2-BSRT increased drastically. The mirror threatened to drop from the support and damage the viewport. Therefore beam operation was stopped and the B2-BSRT removed. As consequence the abort gap could not be monitored anymore. To compensate for that a procedure was introduced to frequently use the abort gap cleaning and observe beam losses due to this. In addition alternative solutions for abort gap monitoring are under investigation. The B2-BSRT was re-designed and re-installed in TS3.
- Daniel presents a whole list with other MPS issues retrieved from the checklists. For each issue it is marked if this issue is understood and if mitigations were performed or are planned for LS1.

Discussion:

- Giulia commented that in the MAC it was mentioned that 'Programmed dumps' should be maybe subdivided into real EOFs and those that are dumped to anticipate another problem.
- Check number of interlocked BPMs dumps (orbit classification in dumps from beam monitoring).
- Send list of MPS false dumps to experts for review, e.g. BLM dumps to Bernd…?
- Richard: LHCb measurement of abort gap population showed to be non-obvious – measurement is fluctuating too much and good absolute calibration is difficult – in the longer term MPP will follow up alternative measures for abort gap population
- Ruediger: Is it possible that an SEU is increasing the gap size above the limit without dumping the beam -> Daniel answered no…

  Chiara to check the TSU fault on the 19th of Aug (most likely a case of wrong analysis only…).

- Bernd: What is the reason for the threshold increase on Q4.R/L6? -> Bottleneck is TCSG, which induces losses of secondary particles in Q4. No primary particles are lost on the Q4.R/L6.
- **Action:** Andrea proposes to present a similar presentation in the availability WG.

1.2 SIS for injector renovation and new machines (J. Wozniak)

- The Software Interlock System (SIS) is a Java based software, complementary to the beam interlock system (hardware). It protects against repetitive faults in the current and following cycles.
- SIS is not a real time system. Millisecond time dependence, not Nano-second, No SIL 4 level , not for human safety. As any software it may fail, although this was rarely experienced.
- Jakub shows the basic architecture of the SIS and the signal paths.
- Key features:
    - Easy to configure
    - Reliable core
    - Flexible to cover various needs.
    - Comes with generic tools and based on common JAPC-MON framework
    - Experience from various installations (SPS, PS, LHC, Linac2, …)
- There are new features for the injectors:
    - Custom triggers (trigger on any conditions).
    - Groovy scripts for easy configuration (conditions, triggers).
    - Counters to filter out noise.
    - ValueReplicator for On-change settings.
        - Bernd comments that this means that the system itself creates signals. Ruediger comments that for an interlock system this is not desirable. Jorg points out that in the SPS and LHC it was tried to repeat all the signals per cycle, to avoid the need of a ValueReplicator.
    - Thresholds as LSA Virtual Devices.
    - Improved diagnostics.

- There exist 8 SIS instances:
  - LHC-SIS with over 2600 subscriptions.
  - SPS-SIS with less than 1000 subscriptions.
    - Bruno asks if the SPS-LHC transfer lines are connected to SPS-SIS. Jorg confirms this.
  - LHC-BIGSISTER with less than 700 subscriptions, which contacts the announcer.
  - SPS-BIGSISTER with 10 subscriptions, which is used for beam quality checks and writes into the logbook.
  - LINAC2-WATCHDOG-SIS with 56 subscriptions, which checks the beam conditions.
  - ISO- SIS (35 subscriptions), alarms, status for Isolde vistar.
  - PSB-SIS (24 subscriptions) – Isolde comparator, talking to the timing.
  - CPS-SIS (27 subscriptions) – Dirac, TFID, nTof comparator.
    - Ruediger asks what the overlap between SIS and the alarm system is. Jakub and Jorg respond that the main difference is that the alarm system does not act.
    - Jorg adds that the SIS was originally designed for the SPS, and is therefore ideal for a cycling machine.
- SIS performance depends on:
  - Data acquisition: based on subscriptions (mostly no gets), data need to be there before evaluation.
  - Logic evaluation: in-memory tree of conditions, primitive types (int, long, boolean), calculation usually very fast (10-100ms). Slowed down by logging & statistics for diagnostics.
  - Export actions: Blocking SET calls over the network via middleware (via FESA class). Inherently unpredictable. Normally between 10-500ms, but can take (in rare cases) up to several seconds.
    - Ruediger asks if it would be possible to directly communicate with the FESA class instead of going via the middleware.

  - Java garbage collection: stops all, therefore fine-tuning was needed, which gave good results. With Java 7 this should be improved. But clearly the SIS is not a real time system.
  - Server machine CPU load.
  - Network hiccups.
  - Frontend load.
- Some LHC-SIS statistics from the past days:
  - Export: 4.8M exports, out of which 967 took 300ms-1s, and 139, which took 1-2s. The majority of sets is done below 100ms.
  - Evaluation: ~700000 evaluations, most done between 10 and 100ms. Some 5000 took 100-200ms.
    - Bernd asks if there are no data, which took longer to export than 2s. Jakub responds that this is true for the time period concerned (about 3 days).
- SPS-SIS statistics is comparable to LHC. Note that the hardware used for the SIS is older and therefore the evaluation is slower. There were 4 cases taking about 10s (time out), when the input was not received due to a non-responding client process.
  - Joerg comments that in the SPS anyhow all the analysis is done at the end of the cycle. Therefore one doesn't need to worry about ms response.
- LINAC watchdog statistics: evaluation speed is comparable to the SIS in the SPS, respectively the LHC-SIS.
- SPS Economy mode triggers on the start cycle with a short delay to wait for NO_BEAM, BEAM data, which then decides if the magnets are ramped or not. This saves about 500kCHF per year.
- Performance conclusions:
  - SIS is good for non-critical calculations (not human safety, not hardware safety with ns response time).
  - SIS can be used, where a very low probability of not acting in time can be accepted.
  - SIS is as good as the environment it is used in (hardware, network, ...).

- Possible improvements:
    - Better server machines.
    - Less loaded frontends.
    - Network improvements (better topology, same network segment, etc.)
        - Bernd asks what time a proxy adds to the export. Jakub responds that he hasn't a number.
    - More than one SIS instance running.
        - Ruediger comments that this probably mainly adds more complexity. It is not clear if this increases the availability. Markus comments that this depends a lot on the system. One may loose availability but especially for the injectors could improve dependability.
        - Bruno asks if one could imagine having more than one SIS instance for LINAC4, comparable to the many BICs. Could this improve the performance? Jakub responds that this would be feasible, but it is not obvious that it makes a big difference to the performance if one uses three trees on one system or one on tree on three different systems. Markus comments that one could probably improve more if this is installed closer to the BIC etc. .
        - Bruno asks if one could run an instance of the SIS directly on the BIC frontend. Jakub answers that, as it is java, one could run it. Bernd asks, what the disadvantages are. Jakub responds there are none. Markus comments that there could be memory leaks etc., which could influence the performance of the front-end. Ruediger mentions that his preference would be to install an independent box in the same crate.
        - Bettina asks if we would be affected by the TGM time out (timing distribution). Jakub answers that this could be the case. Bettina emphasizes that the TGM time out is daily observed in the CCC, which would be worrying. Jakub

mentions that the timing cards, which are installed in the LHC and the SPS, avoid this. The problem is that not all the information is distributed by the timing cards, therefore, one may depend on the TGM. The ultimate solution is installing a dedicated timing card. In the LHC and the SPS only signals from the timing card are used.
        - Markus points out that for new machines these properties can be defined and requested. This is clearly different for existing machines.
- Acknowledged plans:
    - CTF – replace compare classes
    - CPS-SPS TFID check, nTOF intensity.
    - LINAC4 test stand will be implemented.
        - Bruno and Bettina comment that the implementation of the SIS in the LINAC4 test stand is new information.
        - Bernd asks how the implementation of the new JAVA version is done. Jakub answers that this is currently under test and investigation. It will be implemented in LS1. From current experience there are no issues with JAVA7.

1.3 Performance proposal for LHC-BCM – (M. Zerlauth)

- After LS1 the LHC will have new challenges for machine protection (higher beam energies, tighter collimator settings, lower quench levels, etc.)
- The Beam current change monitor (BCM) was a vital part of the HERA machine protection system. It was proposed for the use in LHC already in 2005.
- BI started development in 2010.
- The dynamic range and noise determine the achievable detection level.
- The first performance goal was to detect the loss of ~0.1% within 1ms (~10 turns).
- Recent operational experience: 0.1% loss per turn warning is clearly visible in a reference period (16.-18.10.) with the firmware version

provided by DESY. With these threshold the system wouldn't have caused spurious dumps.

- Over the past 3 weeks only dumps would have been triggered at the end of the fill.
- What threshold level would be possible: During stable beams the resolution could be better than 0.04%. The higher noise during injection (0.1%) is suspected to be due to a dependency of the measurement signal in the toroid to the beam position. The FBCT is sensitive to changes of the orbit by 10um, which is the performance limitation for the current system.
- The current BCM system will generate a dump signal for beams of ~2.2e14p a dump signal after a single loss (one turn) above 0.1% of the beam and 0.02% continuous loss over 10 turns.
- Derived design goals of BCM: it has been shown in the past 3 weeks that a threshold of 0.1% losses per turn has not created any fake dumps. A threshold of 0.15% beam loss per turn is set as initial design goal.
- Future improvements:
    - With energy and beam mode dependent thresholds, a sensitivity of 0.06% loss within one turn at stable beams (~ with 2.2E14p per beam) seems feasible. Although a dependence on beam mode is not recommended for reasons of simplicity.
    - Intensity dependent threshold: recommended, as perturbations seem proportional to total beam intensity, thus, threshold could be set considerably lower for lower intensities.
    - By reducing the position dependence of the toroid a threshold of 0.03% beam loss seems to be realistic.
        - Ruediger comments that the dump threshold should be defined to be an absolute number of lost protons.
- Could such a system limit operation? If this is the case one needs to look at a dependable way to remotely change thresholds.