

A close look at beam aborts with rise times less than 40 ms from the years
2014-2016. Case studies.

A. Drees

September 2016

Collider Accelerator Department
Brookhaven National Laboratory

U.S. Department of Energy

USDOE Office of Science (SC), Nuclear Physics (NP) (SC-26)

Notice: This technical note has been authored by employees of Brookhaven Science Associates, LLC under Contract No.DE-SC0012704 with the U.S. Department of Energy. The publisher by accepting the technical note for publication acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this technical note, or allow others to do so, for United States Government purposes.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or any third party's use or the results of such use of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof or its contractors or subcontractors. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

C-A/AP/570
September 2016

**A close look at beam aborts with rise times
less than 40 ms from the years 2014-2016.
Case studies.**

A. Drees



**Collider-Accelerator Department
Brookhaven National Laboratory
Upton, NY 11973**

**U.S. Department of Energy
Office of Science, Office of Nuclear Physics**

Notice: This document has been authorized by employees of Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy. The United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this document, or allow others to do so, for United States Government purposes.

A close look at beam aborts with rise times less than 40 ms from the years 2014-2016. Case studies.

A. Drees

September 14, 2016

1 Introduction

In an effort to understand the risks of operating RHIC with an additional delay of 40 ms in the abort system, all beam aborts triggered by loss monitors at store from the years 2014, 2015 and 2016 were analyzed and particularly fast cases selected. The results were presented at the RHIC retreat on Jul 29, 2016 [1]. All beam aborts at injection, during the ramp and at flattop but before the “ev-lumi” event were ignored since the additional delay of 40 ms is proposed for operation at store only.

2 Available Data

For all beam aborts by loss monitors a rise time leading up to the abort trigger was calculated by Yun Luo. The rise time, according to [2] was defined as the time between reaching 10% of the signal maximum and 90% of the signal maximum. Fig. 1 represents a sketch of the




Figure 1: Definition of risetime as the time between 10% and 90% of the loss monitor data. Showcase fill is 19769.

definition. In this example the saturation level of 4530 rad/h for g8-lm5 constitutes the

maximum. The time of the abort trigger is indicated by a black dashed vertical line. In this particular case the rise time appeared to be 19.4 ms. The Post Mortem database provides loss monitor data with a granularity of 720 Hz for each abort. Due to the 720 Hz granularity all rise times constitute an integer multiple of $1/720 = 1.39$ ms.

fill	date	time	ring	species	rise time [ms]
18061	03/17/2014	17:23:13	B	Au	2.78
18070	03/18/2014	21:28:05	Y	Au	1.39
18304	05/11/2014	13:20:13	B	Au	4.17
18361	05/27/2014	23:19:34	B	Au	2.78
18483	06/26/2014	19:06:28	B	p	2.78
18969	04/29/2015	23:17:45	B	p	4.17
19009	05/03/2015	01:49:00	B	p	4.17
19234	06/19/2015	09:28:47	B	p	1.39
19374	01/28/2016	03:19:10	B	Au	0
19519	02/11/2016	17:21:10	B	Au	4.17
19761	04/13/2016	16:29:16	Y	Au	4.17
19769	04/14/2016	16:24:19	B	Au	19.40
19834	04/28/2016	18:26:40	B	Au	2.78
19848	05/01/2016	19:16:49	B	Au	18.05
19894	05/11/2016	09:59:19	Y	Au	4.17

Table 1: List of fills aborted by loss monitors with a signal rise time of ≤ 40 ms

15 cases with rise times below 40 ms were found in 2014 to 2016. The cases are listed in Tab. 1, 4 cases with protons and 11 with Au. 12 happened in the blue ring, the remaining 3 in the yellow ring. Each case is discussed individually below. The figures in this analysis were generated using the PMViewer GUI [3] and LogView [4]. For each case a reason, as quoted in [1], is listed for reference with the more detailed analysis.

2.1 Case 18061

| B | g8-lm5 | 2.78 ms | MCR: FEC reboot while gap cleaning |

No quenches were associated with this abort. Fig. 2 (left) summarizes the event, including the loss data (red line), various thresholds and the beam current data. The data shows a 10 Hz structure (0.1 s steps), representing the fastest of the gap cleaning options of 10Hz, 1Hz or 0.25Hz. The g8-lm5 LM saturates with only 0.3% of the beam lost (and most of that would be lost on the collimators, not at g8-lm5). Due to the ongoing reboot of the FEC (cfe-2a-tune1) the status of the gapCleaner [5] cannot be reconstructed fully. However, as can be seen in Fig. 2 (right), the bunch-by-bunch losses clearly move from the gap to a location in middle of the bunch train. The blue trace is taken 1s before the abort, the red trace shows the data accumulated during the second in which the abort happened. The peak appears 'smeared' due to varying arrival times of particles in the collimator Pin Diodes.

Conclusion: The amount of beam loss is actually small (about 1/3 of one bunch), too small for damage beyond, perhaps, a quench. The gapCleaner, due to its kicker pulse length

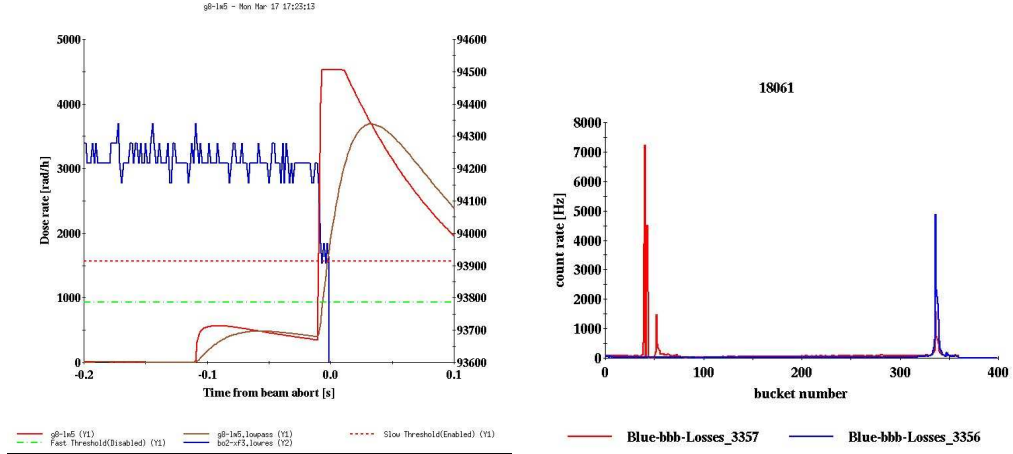


Figure 2: 720 Hz loss monitor and current transformer data from PMViewer (left). Two traces corresponding to two consecutive seconds of bunch by bunch (bbb) loss data from the abort in 18061.

of only about 100 ns [5], affects only one bunch at a time but cannot impact the entire beam. The fastest time for it to switch from one bunch to another is 10 Hz, i.e. 100 ms, plenty of room for an additional 40 ms.

2.2 Case 18070

| Y | y3-lm4 | 1.39 ms | MCR: Lisa steering with a bad activation |

No quenches were associated with this abort. Current transformer and loss data are shown in Fig. 3, left. It is interesting that the losses appear to stall after a fast initial drop over the first 2.8 ms. The losses correspond to 0.7% of the total beam, a bit less than one full bunch. MCR attributed the abort event to a bad lisa [6] steering. Following are two excerpts from the appropriate 2014 lisa logfile:

Tue Mar 18 21:27:21 2014 -1 ERROR: MultiOptimizeScan: error activating - aborting scan

After the error message MCR tries to optimize again:

Tue Mar 18 21:28:08 2014 0 Retrying parallel optimization after failure

This last automatic message and the accompanying retry appears to happen 3 seconds after the time of the beam abort making lisa steering unlikely to be the cause. This timing can be confirmed by checking the logged PS current data for the involved correctors and the timing of the abort- and auxramp2-events as can be seen in Fig. 3 on the right. Indeed, all correctors in question are at nominal values until the ramp-down starts (well after the abort).

In addition, y3-lm4 is an unusual loss monitor to be first responder in a failure caused by corrector magnets around IP6 and IP8. It is, on the other hand, close to a stochastic cooling tank which was indeed moving at the time and losses spike on yi3-ksch3-pd (the pin diodes associated with the yellow horizontal stochastic cooling kickers in sector 3) at the time. The extremely fast loss rise time, however, makes the move of a large mechanical system equally unlikely to be the root cause of this very fast event. Unfortunately, all abort kickers are accounted for and not one PS was alarming or reporting an issue at the time of question.

Conclusion: The abort was not due to a bad lisa activation. The SC kicker as a me-

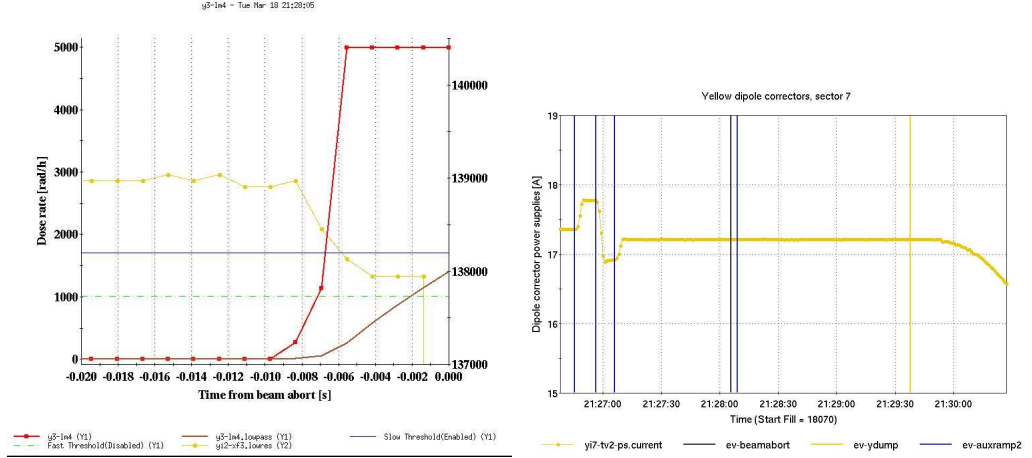


Figure 3: 720 Hz loss monitor and current transformer data from PMViewer. Dipole corrector data during store 18070, before and after the abort.

chanical system is too slow to be the culprit. No other culprit candidate could be found and this case remains a mystery.

2.3 Case 18304

| B | g8-lm5 | 4.17 ms | MCR: FEC reboot while gap cleaning |

No quench associated with this abort. The gapCleaner was set to a 1 Hz frequency and the losses, as shown in Fig. 4, left, clearly show a 1 Hz structure. During the trigger at -1.0 s about 0.1% of the beam was lost and just before the abort another 0.2%. The bunch-by-bunch losses, shown in Fig. 4, right, were accumulated in a second before (trace 516) and during the abort (trace 518). They clearly show how the losses moved from the gap to an occupied bucket. A modest lowering of the trip threshold would have aborted the beam easily 1 s earlier. The FEC (cfe-2a-tune1) apparently rebooted itself spontaneously. A warning

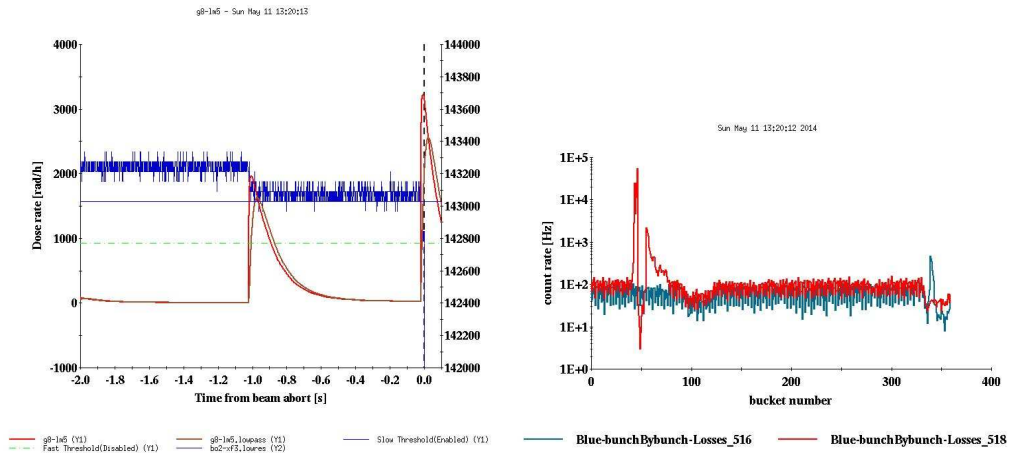


Figure 4: Left: 720 Hz loss monitor and current transformer data from PMViewer. Right: Bunch by bunch losses accumulated two seconds before and the second including the abort.

message (not to reboot with beam in the machine) has been added since.

Conclusion: The gapCleaner, due to its kicker pulse length of only about 100 ns, affects only one bunch at a time and cannot impact the entire beam. Hence the associated beam loss is relatively modest and not close to a damaging amount.

2.4 Case 18361

| B | g10-lm5 | 2.78 ms | MCR: Blue abort kicker prefire |

Module 4 of the blue abort kickers prefired. This was the 15th prefire in run 14. At least one magnet in the blue sector 10 protection bump area quenched (B10DSA4_A3VT). This is the very event we are trying to prevent by adding the 40 ms delay (aka the mechanical switches). Noteworthy is the rise time of 2.78 ms, which isn't even close to the "real" risetime of one turn in a prefire. This demonstrates that +/- 2.78 ms is the limit of accuracy for the rise time determination which is based on 720 Hz data and is thus much slower than 12.8 μ s.

Conclusion: This event can be ignored for the purpose of this study.

2.5 Case 18483

| B | g10-lm5 | 2.78 ms | MCR: APEX, fins of the target clipped through the |
| beam tail |

There were no quenches associated with this event. Fig. 5, left, shows g10-lm5 data. g10-lm5, showing losses caused by an ordinary beam dump, actually did not trigger the beam abort but b12-lm4 did (Fig. 5, right). The actual rise-time for b12-lm4 is well above 100 ms and




Figure 5: Left: 720 Hz loss monitor b12-lm4 data from PMViewer. Right: 720 Hz loss monitor g10-lm5 data from PMViewer

thus of no concern for the extra 40 ms delay. This event got into the list by accident. However, it indicates that a rise time of 2.78 ms can be falsely calculated from the 720 Hz data, here due to the continuous rise of the loss rate after the abort (thus shifting the maximum to after the abort). The actual rise time in this case is less than 1.39 ms (i.e. one 720 Hz bin). In addition, it is suggested to exclude APEX time from operating the aborts with the additional 40 ms. There are, of course, failure modes during APEX that could happen at

any other time as well and therefore need to be checked.

Conclusion: This event can be ignored for the purpose of this study.

2.6 Case 18969

| B | b10-lm4 | 4.17 ms | MCR: Ramp development, abort reason not found |
 | in elog

There were no quenches associated with this event. In principal, this event should not be listed here since it happened during ramp development. However, we need to check if this type of failure could not happen at other times as well. In any case, looking closely at the time of events as shown in Fig. 6, it shows that the permit pull by b10-lm4 comes following an ordinary beam dump. According to the list of RHIC event times, the two are 10 ms apart. b10-lm4, the first loss monitor down stream of the blue abort, is prone to and known for




Figure 6: LogView graph showing beam current data and list of events for end of store 18969.

permit pulls due to “spilled” losses from an ordinary beam dump, particularly during proton operation.

Conclusion: This event can be ignored for the purpose of this study.

2.7 Case 19009

| B | y9-lm4 | 4.17 ms | MCR: Taking a “quick” polarimeter. |

No quenches were associated with this abort. y9-lm4 (yellow!) actually goes above threshold after the abort and the triggering loss monitor is b12-lm4 (blue!). Fig. 7 shows the corresponding data. The rise time for this monitor is approximately 100 ms. Therefore this case is of no interest for this study. y9-lm4, the first loss monitor down stream of the yellow abort, is the yellow LM known for permit pulls due to “spilled” losses from an ordinary beam dump, particularly during proton operation.

Conclusion: This event can be ignored for the purpose of this study.




Figure 7: PM 720 Hz b12-lm4 data

2.8 Case 19234

B	b6-lm3.1	1.39 ms	MCR: Blue Quench Link Interlock, BS2 tripped off early.
---	----------	---------	---

B6Q3 triplet quadrupole quenched in this case. This abort portrays the perfect storm of cavity trips where a faulty damper was not secured when the trip happened. Fig. 8 summarizes the course of events. Shown are the BS2 gap volts (black squares) in kV, the radius (red line) in micron and the beam current transformer data (blue line). The cavity trips 100 ms before the abort ($t=0$). The unsecured tuner starts moving towards the beam, changing the cavities tune in the process. 20 ms before the abort very fast instabilities entail, visible as rapid oscillations of the cavity voltage as well as the radius. Actual beam loss is concentrated in the last $5 \times 1.39 \text{ ms} = 7 \text{ ms}$ before the abort.

Conclusion: Only if the BS2 cavity would have been in the permit (it wasn't) this quench (with or without an additional 40 ms of delay) could have been prevented.

2.9 Case 19374

B	g10-lm20	0 ms	Travis: 0.10 mrad angle bump at IR8
---	----------	------	-------------------------------------

This case happened outside physics, during setup and commissioning and would therefore be done without a modified abort system. However, it was actually a prefire (module 3), at least two magnets quenched in B10DSA4_A3VT and B10QFA3_A2VT. There were only 6




Figure 8: 720 Hz RF BS2 voltage data, radius data and beam current transformer data.

blue bunches in the machine, i.e. a prefire with a very unlucky timing.

Conclusion: This event can be ignored for the purpose of this study.

2.10 Case 19519

| B | g8-lm5 | 4.17 ms | MCR: APEX, squeeze to 50 cm in Blue |

No quenches in this case. The abort happened during an APEX, a time during which no change to the abort system is planned. At the time of the abort were 12 bunches in the machine and a new ramp to squeeze IP8 to 50 cm was tested. As part of the ramp development TbT orbits were triggered every 4 s which include a trigger to the ARTUS tune meter kickers. The losses due to two of those triggers can be seen at -6.5 s and -2.5 s in Fig. 9 (left), precisely 4 s apart. There are two corresponding TbT data sets stored in the orbit archives. Only one bunch (# 31) is affected by the TbT triggers. However, beginning about 15 s before the abort (and 47 s into the ramp), all 12 bunches start to oscillate (± 3 mm peak to peak, approximately 140 Hz frequency) and to show some beam losses according to the bbb beam loss monitor. The amplitude more than doubles about 25 ms before the abort as can be seen in Fig. 9 (right). Even though losses are visible on the loss monitors, there is no noticeable beam loss seen in the current transformers up to 1.39 ms before the abort.

Conclusion: No single cause of abort could be identified for this case although it is likely associated with the ongoing beta-squeeze. One could argue to ignore it since it happened during APEX and during a ramp, if only a squeeze. Alternatively one could establish a maximum allowable oscillation amplitude of, say, ± 2 mm (at store). In this case this ramp could have been aborted 15 s earlier than it actually did.




Figure 9: Left: g8-lm5 PM data from abort of store 18519. Right: The last 0.15 s before the abort depicting 1kHz data from GOFB_10Hz BPM bi8-bh3, current transformer and g8-lm5.

2.11 Case 19761

Y | g8-mlmx.1 | 4.17 ms | MCR: APEX, angle steering at IR6 during lumi leveling

No quenches were associated with this event. This abort happened during an APEX study and would as such have happened without any changes to the abort system. However, this loss scenario is not associated with APEX activities such as the quoted angle steering in IR6 (which is not the root cause for this abort). The beam loss happened very fast (see Fig. 10, left). As can be seen in Fig. 10, right, most of the angle steering in IR6 was already completed when about 15 s prior to the abort angles start drifting again in both IPs and beam is lost at angles much smaller than the ones achieved before. At the time of abort all of the supplies in




Figure 10: left: g8-mlmx.1 PM data from abort 19761. Right: IP8 and IP6 crossing angles during APEX leading up to the abort.

alcove 9A, 9B and 9C tripped on a lead flow fault. Cryogenics replaced a power supply in the 1008B valve box for 8Q3. Yet, a lead flow fault alone does not quite explain the very short

rise time seen in this event. In another similar case (19837, in which the yellow injection kicker got damaged) the rise time was of the order of a few 100 ms.

Conclusion: It is unclear if the likely reason (lead flow fault) is enough to explain the very short rise time. No other reason could be found however. Therefore no recommendation in anticipation of an extra 40 ms delay can be made.

2.12 Case 19769

| B | g8-lm5 | 19.4 ms | MCR: store following 56 MHz study, caused by 10Hz PS |
 | PS

No quenches were associated with this case. This abort is caused by a failure of a GOFB_FB 10Hz PS without any connection to the 56 MHz RF cavities. As can be seen in Fig. 11, the power supply and the orbit start oscillations about 70 ms before the abort (indicated by the dashed black vertical line), significant losses appear about 20 ms before.




Figure 11: Orbit oscillations, power supply current and loss monitor data during 0.15 s before the abort.

Conclusion: In order to trigger an abort 40 ms before the recorded time here, either a BPM based permit or a 10 Hz PS based permit would be needed with a response time of approx. 10 ms.

2.13 Case 19834

| B | b8-lm4 | 2.78 ms | MCR: likely linked to 56 MHz |

There were no quenches associated with this abort. The culprit in this case was not the 56 MHz cavity but operator error. b8-lm4 is the first loss monitor down stream of the blue

collimators with an enabled permit threshold. As can be seen in Fig. 12, left, gap cleaning was in progress at a modest level at 0.25 Hz, one of the three frequency options. This can be confirmed by the configuration setup for the blue gap cleaning kickers (Fig. 12, right). Acquisition mode is set to “3”, i.e. the 4secEvent, minutes before the abort. At the same time the default bucket for gap cleaning, was changed by operators from 347 to 336 and eventually to 334. Note that bucket numbers are offset by +1, starting from 0 and going up to 359. In addition, it can be expected that the kicker pulses are not perfectly centered on the occupied buckets. Timing shifts can occur during the ramp to store. The effect of the

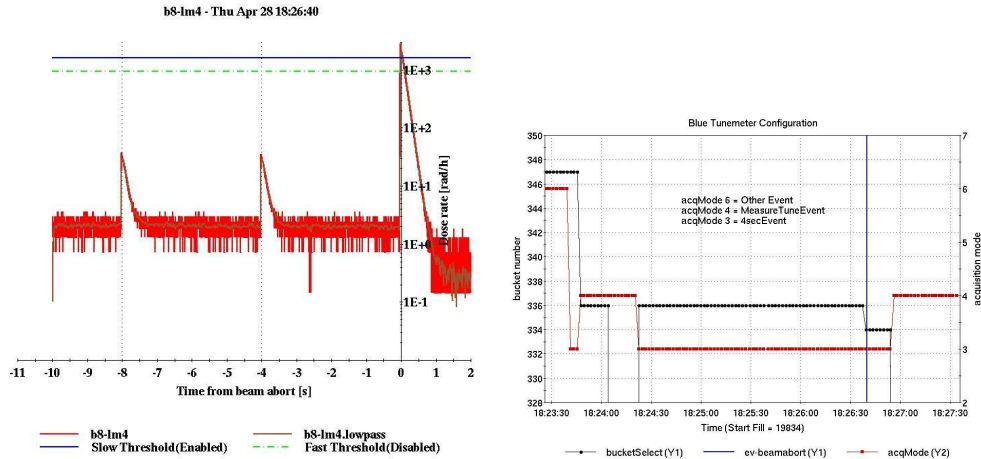


Figure 12: Left: b-lm4 PM data, right: blue tune meter kicker setup before and during the abort.

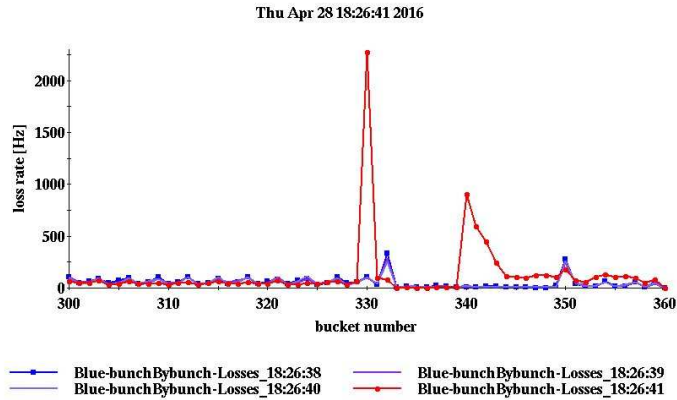


Figure 13: Bunch by bunch loss monitor for the seconds leading to (blue lines) and including the abort (red line).

bucket switch is apparent in Fig. 13. In the seconds leading up to the abort modest losses are evident at buckets 331 and 350. Both buckets are part of the gap. The last data set, including the abort (shown as a red line) shows a peak at bucket 330, the last one of the occupied buckets. The kickers were set to bucket 333, but due to the pulse length of about 110 ns, the last occupied bucket got hit, causing the abort.

Conclusion: There is very little risk for damage since the gap cleaner kickers will only affect one bunch at a time with a repetition rate of 0.1 s or slower. Adding a delay of 40 ms would still have aborted the beam well before the next trigger could have occurred, even if it was set to 10 Hz.

2.14 Case 19848

| B | g2-lm16 | 18.0 ms | MCR: Blue longitudinal pick-up error |

No quenches were associated with this abort. The blue longitudinal pick-up was certainly not involved in this event. Actually, about 700 ms prior to the abort a total of 11 dipole corrector supplies controlled by cfe-5c-ps1 started to fail. It is likely that cfe-5c-ps1 was compromised by the bad vacuum in sector 5, causing the concerted trips of many corrector supplies. One of the power supplies is shown in Fig. 14, together with the PM data for LM g2-lm16, at the maximum of the blue protection bump. Losses in g2-lm16 start showing over 500 ms after the dipole correctors first signs of errors. Only the LM down stream of the blue collimators experience some earlier losses, reaching only about 450 rad/h about 50 ms before the abort.




Figure 14: PM data from g2-lm16, g8-lm5, and bi-7th. LogView current data from bi-5th.

Conclusion: In order to violate the beam permit 40 ms earlier than the current abort time, a PS based permit system appears most promising. A much reduced permit threshold (about 25% of current level) for g8-lm5 could potentially do the same with the added risk of unnecessary and premature permit pulls. In addition, it relies on the collimators being in the proper position.

2.15 Case 19894

| Y | g8-Mlle.1 | 4.17 ms | MCR: y8-q6 QLI |

This event happened during dAu operation. A quench of Y7Q2 was associated with this abort. As can be seen in Fig. 15, the y8-q6 power supply started oscillating 0.5 s before the abort. It is in a steady error state at 0.4 s before the abort. Nevertheless, losses appear only at 0.01 s prior and the abort cannot prevent the quench of Y7Q2. No other loss monitor in the area sees earlier losses either.

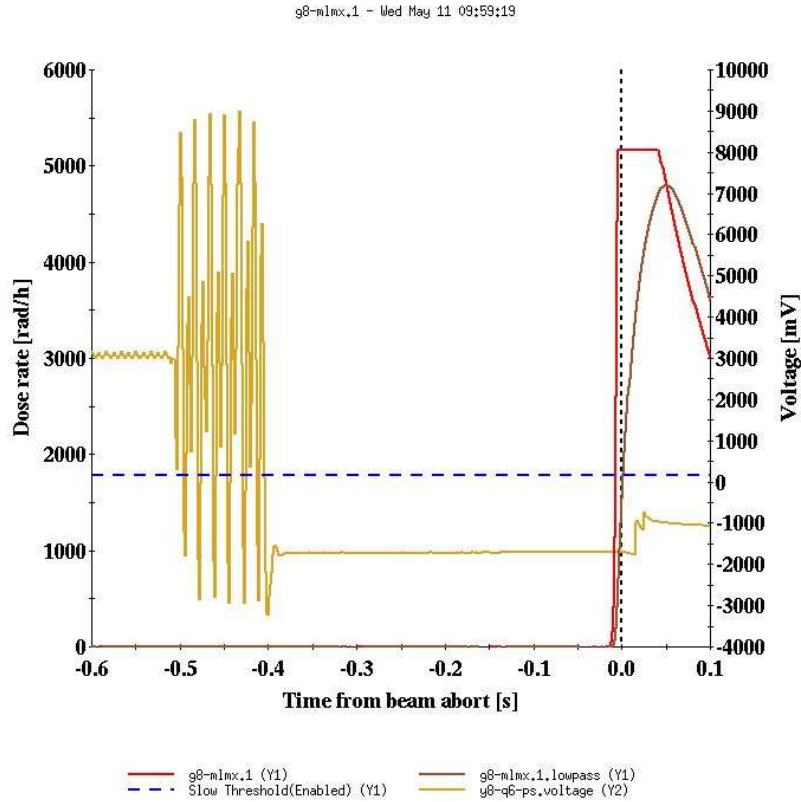


Figure 15: g8-mlmx.1 and y8-q6-ps PM data during the 0.6 s prior to the abort.

Conclusion: Since there are no evident losses early enough only a permit system based on PS data could have prevented the quench. Given the time scale of 0.5 s, an additional 40 ms delay should not pose a problem.

3 Conclusion

Many (but not all) of the 15 studied cases are of no concern. To summarize I put the studied cases into four categories, summarized in Tab. 2.

All aborts in category 3 (not damaging levels) involved the gap cleaner. Permit input other than loss monitors (LM) include BPM, RF (storage cavities) and a permit system based on power supply input. In all cases of category 4 there appears to be enough time to allow an extra delay of 40 ms, provided the new permit systems would be able to act within a few 10 ms. This is a lower limit since in some cases there would have been much more time available than that. Of course, it is unknown if in any of the listed cases damage would have occurred **IF** the beam would have been aborted later.

There is one exception, 19234. In this case the abort even with the existing fast system

#	category	fill	total
(1)	can be ignored	18361, 18483, 18969, 19009, 19374	5
(2)	cannot be explained	18070, 19761	2
(3)	not damaging levels	18061, 18304, 19834	3
(4)	need permit inputs other than LM	19234, 19519, 19769, 19848, 19894	5

Table 2: Cases sorted by category.

happens so late that by the time of the abort most of the beam was either debunched or already lost (there were little to none traces of beam visible on the abort scopes). The resulting damages, however, were modest (compared to a vacuum leak in the injection kickers or a diode failure) and only one triplet quadrupole quenched.

Cases with high damage potential are rare - but not rare enough. In order to make an added 40 ms during physics store conditions as safe as reasonably possible, additional permit inputs such as 10 Hz BF power supplies (19769), RF storage cavities (19234), power supply error states (19894) or BPMs (19519) in addition to significantly reduced LM thresholds for selected LM (19848) should be commissioned. Unnecessary additional down time will be likely, certainly during commissioning. Nevertheless, such times need to be compared to a diode replacement, a vacuum leak in the injection kickers or a scenario in which an experimental detector is damaged in a prefire. Given those additional permit inputs the extra 40 ms delay should be permissible for physics store conditions. For added safety we could consider turning experimental detectors off (and using a fast abort system) during any kind of ramp.

References

- [1] Y. Luo ET AL., “Possible Failure with a 40 ms Delayed Abort”, <https://indico.bnl.gov/getFile.py/access?contribId=16&resId=0&materialId=slides&confId=2127>
- [2] Y. Luo, private communication.
- [3] <http://www.cadops.bnl.gov/Controls/doc/PMViewer/PMViewer.html>
- [4] <http://www.cadops.bnl.gov/Controls/doc/LogView/LogView.html>
- [5] A. Drees et al., “Abort Gap Cleaning in RHIC”, conference proceedings, EPAC 2002.
- [6] A. Drees et al., “Luminosity Optimization Using Automated IR Steering at RHIC”, conference proceedings, EPAC 2004.