

THE RHIC BEAM ABORT SYSTEM – OPERATION DURING THE RHIC 2001 GOLD RUN*

L. Ahrens[§], J.-L. Mi, W. Zhang, BNL, Upton, New York

Abstract

The RHIC Beam Abort system removes the circulating beams from the two RHIC superconducting rings both under normal conditions at the end of a beam "Store" and in the event of unusual conditions, which cause the beam "permit" to be dropped. The design objective in both cases is to quickly remove the beam without causing any superconducting magnet to quench. In addition it is extremely desirable to avoid spraying any beam into the experiments. The Abort system had been commissioned during the previous (FY2000) RHIC run with gold ions at 70 GeV/nuc and with 5e8 ions in each of 55 bunches. The run reported on here saw an intensity increase to 1e9 ions per bunch and more importantly an energy increase at Store up to the design - 100 GeV/nuc. The second change requires that the Abort kicker magnets operate at higher voltage and also leaves the ring magnets much less margin against beam induced quenches relative to the previous run. Operational experience will be reviewed.

1 INTRODUCTION AND SUMMARY

For the 2001 RHIC physics runs, the RHIC Beam Abort system behaved largely as expected and predicted in earlier descriptive papers [1] and indeed in the RHIC Design Manuel. Deviations from this behaviour are most interesting, but first a short description of some aspects of the system and the run performance will be given.

Each RHIC ring (Blue and Yellow) has an independent beam removal system. A given ring system includes five subsystems each of which contains a pulse forming network (PFN) and a ferrite kicker magnet. The five subsystems are identical, and are fired simultaneously. The time dependence of the resulting beam kick is important (see Figure 1). The 'rise time' (meaning the time for the kick strength to rise enough to move the beam on to the design region of the absorber – a kick of about 1.6 milliradians) is about 1 μ sec. Following this the kick field varies enough to sweep the fifty-five bunches across the face of the beam absorber – to spread the intense heating from the gold bunches. With a 60-bunch pattern in RHIC, the bunch spacing is 214 nsec., so for an unsynchronised kick five bunches fall on the rising edge. The strength of the resulting kick and the size of the absorber face give the system when run at design settings

20% more kick than required. One module could be "off line" and the beam will still get across to the acceptable part of the absorber.

The size of the kick delivered by the system for a given applied voltage has been measured (at injection energy) using the gold beam and the RHIC orbit measuring system [2] each run. A beam position monitor (BPM) located 9 meters into the 24-meter drift region between the kicker and absorber provides an excellent transverse calibration, fitting (in time) the residual turn-by-turn coherent betatron motion (left over from injection) across the occurrence of the kick. From this a voltage for the PFN systems of about 28 kV at 100 GeV/nuc provides the "design" motion at the absorber. The timing to achieve proper synchronization with the abort gap was also initially determined by a short injection study. This timing relies on the RHIC Beam Sync Link and on the local decoding electronics[3].

For most of the run, the system simply did its job and was forgotten. RHIC operated with five missing bunches (which is conservative) creating the abort gap. Beam intensities were essentially at design energy and intensity. Stored beams were dumped typically several times each day. In fact the total number of beam dumps is completely dominated by dumps at injection energy. These are at lower voltage (4 kV) and low intensity (single bunch) but occur typically 10 to 20 times per fill for each ring as test bunches are injected and the injection set-up is prepared. The number of such dumps did not diminish even with the more routine operation near the end of the run. They were a necessary part of that routine.

Radiation Safety issues- respecting the RHIC "Operational Safety Envelope" – added another wrinkle to the requirements from the abort system in conjunction with the RHIC loss monitor system [4] for this run. All dumps were judged as "dirty" or "clean" – which referred to the likelihood not of quenching magnets, but simply of depositing significant beam away from the more heavily shielded sections of the ring. Beam dumped "dirty" in this sense was considered to have been lost in a lightly shielded area and the amount of such, per hour, was required to be held below a fixed limit. Normal running usually was well away from this limit, but during beam studies the limit and this aspect of the system was sometimes relevant.

*Work performed under the auspices of the U.S. Department of Energy

[§] Email ahrens@bnl.gov

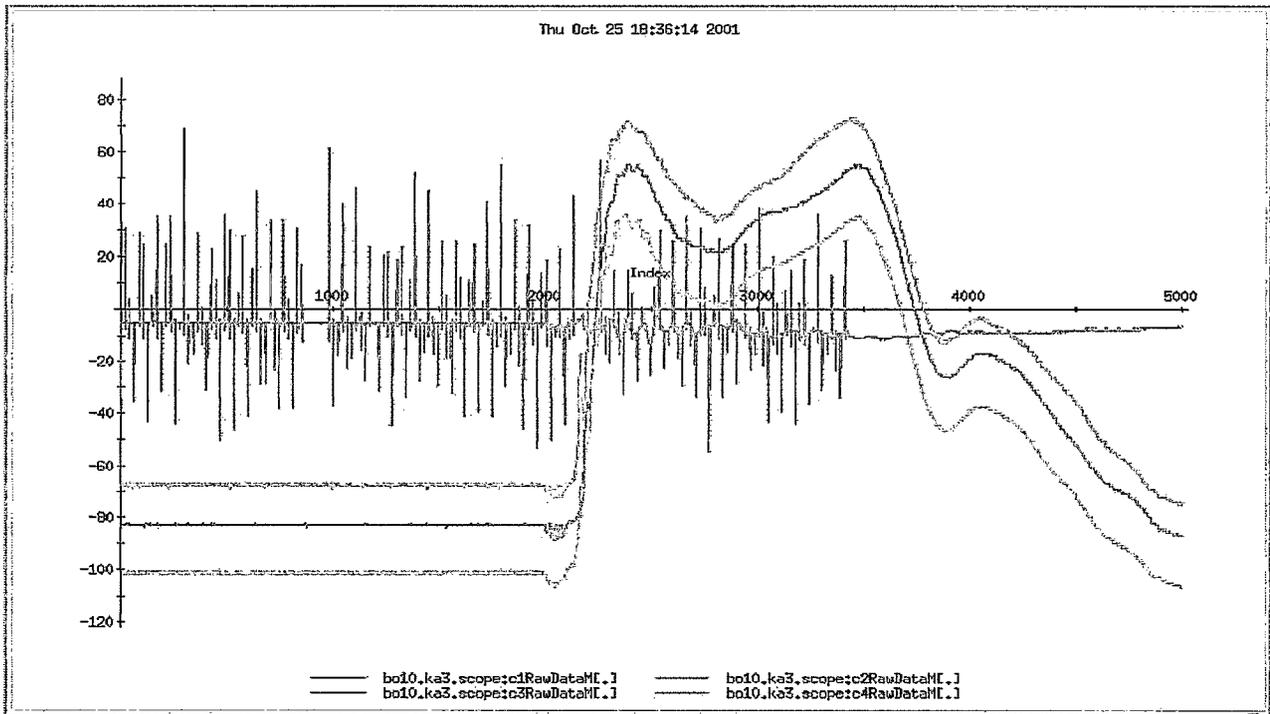


Figure 1 Three PFN Current Transformer Waveforms and the 55 Bunches in the Blue Ring note the abort gaps on the (differentiated) beam signal and the sweeping of the kicker currents

2 PROBLEMS AND SOLUTIONS

First the explanation of some jargon is necessary. A RHIC beam “dump” in e.g. the Blue ring refers to the preplanned removal of beam from that ring. The Abort system is sent an event (a `b_dump` event on the RHIC “time line”) to which the system responds by removing the Blue beam. A RHIC beam “abort” occurs when the beam “permit” is “lost”. The “permit” is a signal, a modulation propagating when beam is “allowed” on a closed loop around the ring. The signal can be interrupted, the permit lost, by permit electronics responding to e.g. a magnet quench, or high losses. In this case the Abort system is not sent a request to remove the beam but rather local Abort electronics senses that the permit is gone, and immediately begins the sequence for removing the beam from both rings. What initiates the action of the Abort system is completely different in the two cases, but beyond a certain point, the Abort system response is identical – get synchronized with the abort gap, and dump the beam. A record of the occurrence of these associated events (the loss of Permit does generate an event just after the fact) with microsecond timing accuracy exists – which helps in reconstructing what happened.

Beam dumps at end of Store occasionally caused magnet quenches. This sort of quench is sensed by the quench link electronics a few hundred milliseconds after the associated dump occurs. The quench appears to be associated with the existence of significant quantities of debunched beam in the ring at the time of the dump.

Some of the beam is surely in the abort gap and hence kicked too weakly. A “gap cleaning” procedure was developed during the run and may explain the decrease in this type of problem as the run progressed [5].

A few-bunch shift in the relative blue abort gap timing was discovered and corrected during a September study period. The shift would cause some beam to fall on the rising kicker edge. Diagnostics were developed to allow checking this timing without a dedicated study. A local BPM signal was added to the oscilloscope display of the current pulse signals from the Abort system PFNs. The resulting scope trace (Figure 1 again) gives the relative timing unambiguously. In fact no further timing problems occurred. The scope traces are available as a “comfort display” in the Main Control Room.

Quenches associated with dumps occurred on average once every few days, and were less frequent during the last week of the run. There seems no evidence for a slow deterioration of the Abort system, as one might fear were the absorber deteriorating. The Blue ring quenched much more frequently than Yellow in both modes, an asymmetry that is not understood, especially given that the Yellow Abort system was the one running at lower voltage (see below).

The most destructive thing the Abort system can do is to kick too weakly, and so to put some – even a tiny fraction – of the circulating beam into a vulnerable spot. A single PFN module discharging “spontaneously” (a “prefire”) can do this since first it starts with only one fifth of the usual kick, and then because it is not

synchronized with the abort gap, it kicks a few bunches with only a fraction of even that kick.

Prefires were a concern early on. The many-hours voltage-holding capability of the PFN thyratrons was recognized as marginal at full design voltage long before first operation [6]. The damage done by prefires was primarily expected to be lost physics running time while the rings recover and are refilled following the expected quench. The rate of occurrence of prefires has never been high enough for this cost to be an issue. Unfortunately a greater cost became apparent once one of these prefires occurred with RHIC at Store and at reasonable beam intensity. The potential cost is to the physics experiments, and in particular to the close-in silicon detectors which make up the Phobos experiment, an experiment located in the 10 o'clock IP between the two Abort systems, (and hence as far away from harm as possible). A few prefires occurred in late September. Each ring system suffered at least one. Our response was to reduce the triggering sensitivity of the Thyatron switches, to replace one marginal tube, to commission circuitry to trigger the remaining PFN's as quickly as possible, and finally, somewhat in desperation, to reduce the running voltage (in the Yellow ring). Following these actions, the gold run continued with ever increasing luminosity for more than six weeks with no further incidents.

In fact it is not known exactly why the modules prefired. The thyratrons are suspect, but the thyratrons installed in Yellow had already been upgraded and are not expected to show a prefire problem at design voltage. Perhaps there is sparking for some other reason in the PFN enclosure. Experience with the final PFN system geometry and for relevant lengths of time has occurred primarily during physics running time. The fine-tuning of the PFN trigger systems (e.g. the Thyatron trigger sensitivity) requires such time, which we will attempt to get before the next run. A search for other weaknesses in the present geometry is presently underway. That a 20% voltage reduction seems to have solved the problem was not expected in advance. This drop removes all the design margin, and is not a comfortable situation.

The attempt to add a fast retrigger was fundamentally constrained by the intentional early on decision to locate the electronics generating the triggers outside of the ring, away from any beam radiation damage. The kicker magnets are, again by design, a competitive vertical limiting aperture at injection. The transit time from kicker to trigger electronics is about 400 ns. So the trigger electronics don't even know that a prefire has occurred until, from the point of view of the bunches seeing the leading edge of the rogue module, it is too late. The circuitry that was added generates triggers for all modules as soon as a current pulse is sensed on any one of them. This second trigger then gets to the other modules only after twice the travel delay plus the time

necessary to create this (high power) trigger. The second trigger occurs about 2.5 μ sec after the unsolicited one. This will help, but it does nothing during the microsecond rising interval of the first PFN. In fact the last "pretrigger event" during the 2001 run occurred with the backup trigger in place, and still caused some damage at the Phobos experiment. Plans for next run are to install the backup trigger banks with only essential electronics in the ring close to the PFNs. There will still be delay associated with the creation of the trigger pulse.

In order to reconstruct what happened when the Abort system faults, usually hours after the fact, we rely on "logged" data. The fast current traces shown above are saved, as are 1 Hz samples of the voltages on the ten PFNs. To save the critical data - right time and time scale - is the challenge. To some extent we were lucky to have clear indications of what had happened from prefire events. These necessarily come without a precursor trigger so the oscilloscope must trigger "internally". A second similar trigger will come shortly crasing the first since the prefire will cause a permit pull. Indeed in this situation the beam permit will be pulled by the Abort system itself once that system realizes that the voltage on one PFNs is not correct. An upgrade in the sophistication of the logging machinery, introducing the ability to very quickly trigger logging on an external event, and more scope storage capacity are planned.

3 ACKNOWLEDGEMENTS

We appreciate the work of John Addessi, Pat Benjamin, Ken Hartmann, Steve Perlstein, Dave Warburton, and Ron Zapasek of the C-A Pulsed Power group in constructing and keeping the Abort system functional, and the efforts of Roger Lee of the C-A Controls group in evolving the "logging" of the many diagnostics for this system.

4 REFERENCES

- [1] H. Hahn et al., "The RHIC Beam Abort Kicker System", Proceedings of the 1999 Particle Accelerator Conference, New York, NY (IEEE, Piscataway, NJ, 2000), pp1100-1103.
- [2] T. Satogata, "RHIC BPM System Performance, Upgrade, and Tools", These Proceedings.
- [3] T. Kerner et al., "V123 Beam Synchronous Encoder Module", Proceedings of the 1999 Particle Accelerator Conference, New York, New York (IEEE, Piscataway, NJ, 2000), pp699-701.
- [4] M. Bei, "RHIC Loss Monitor Threshold System Commissioning in Year 2002 Run", These Proceedings.
- [5] A. Drees et al., "Abort Gap Cleaning in RHIC", These Proceedings.
- [6] H. Hahn et al., "Operation of the RHIC Beam Dump Thyatron CX1575C", AD/RHIC/RD-127, Aug. 1999