ISBN: 978-3-95450-223-3 ISSN: 2673-7000 doi:10.18429/JACoW-NAPAC2019-WEPLH11

RHIC QUENCH PROTECTION DIODE RADIATION DAMAGE*

A. Drees†, O. Biletskyi, D. Bruno, A. DiLieto, J. Escallier, G. Heppner,
C. Mi, T. Samms, J. Sandberg, Brookhaven National Laboratory, Upton, NY, USA

Abstract

Each of RHIC's superconducting magnets is protected by a silicon quench protection diode (QPD). In total, RHIC has over 800 diodes installed inside the cryostat close to the vacuum pipe. After years of operation with high energy heavy ion beams we experienced a first permanently damaged QPD in the middle of our FY2016 Au Au run and a second damaged diode in the following year. In 2016 the run had to be interrupted by 19 days to replace the diode, in 2017 RHIC could still operate with a reduced ramping speed of the superconducting magnets. Both diodes were replaced and examined "cold" as well as "warm". This paper reports on what we have learned so far about the conditions leading up to the damage as well as the damage itself.

INTRODUCTION

At RHIC an average of about 30 beam induced quenches are reported each year, varying between 1 and over 70 depending on the running mode. In most cases, the involved magnets are the triplet quadrupoles of the low beta insertions. Figure 1 shows a QPD and its assembly inside a superconducting magnet. The actual diode, a few millimeter thick

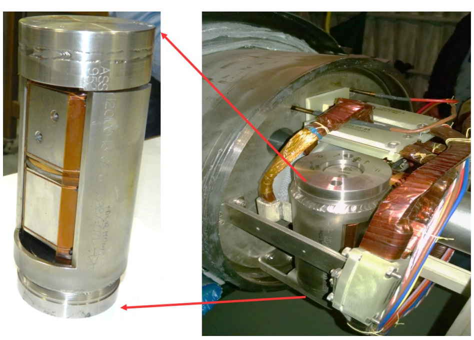

Figure 1: Left: diode assembly, right: installation within a RHIC magnet.

silicon wafer at the center of the assembly, is installed at the same height as and next to the beam pipe. This design has the potential to expose the diode to the highest levels of radiation from beam losses if they originate in the area of its host magnet. Y7-D6 is downstream of the yellow ring's collimators, where higher beam loss rates are expected and routine. However, beam losses in the center of an arc, where B10-D19 is located, are unusual unless the beam trajectory is changed significantly.

The two diodes were damaged in two consecutive years: one in store 19702, in one of the blue arc dipoles (B10-D19) during a dedicated machine development with 206 10^9 Au ions in the blue ring. Two voltage taps reported a beam induced quench: dipoles D15-20 and quadrupoles Q10-20 in arc 10. In principle, any one or several of the 6 magnets sharing a voltage tap could have quenched. It cannot be known which one. The second event occurred at the end of store 20604, damaging the diode in one of the yellow straight section dipoles (Y7-D6) downstream of interaction region 8 (IP8). In this second instance, the beam abort was caused by a failure of the abort kicker trigger circuit during a normal end-of-store beam dump. An unprecedented total of 31 voltage taps reported beam induced quenches in this event. RHIC had 191 10^{11} protons circulating in the yellow ring at the time of the abort.

Both events were indirectly caused by an ongoing effort of preventing damage to experimental detectors after a so-called prefire and the effort of preventing prefires altogether.

ABORT KICKER PREFIRES AND PROTECTION BUMPS

RHIC abort kicker prefires [1, 2], where one of the five abort kicker modules per ring fires spontaneously and asynchronously, happen with a varying frequency of 0-15 incidents per year and per ring, provided the required abort kicker voltage is above the prefire threshold of about 12 kV. For Au-Au operation at 100 GeV and protons at 255 GeV they are operated at about 26 kV. Figure 2 depicts a scope snapshot of the timing of a typical prefire during run16. De-

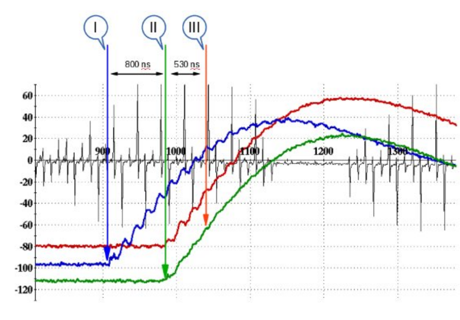

Figure 2: Three abort kicker modules during a prefire of module PFN4 (blue trace). (I) indicates the start time of the prefire, (II) indicates the time the other modules follow suit and (III) indicates the time the sum of all modules reaches default voltage for a clean beam dump. The black trace is the bunched beam.

pending on the exact moment, the fill pattern and the total amount of beam circulating at the time of the prefire, up to about 10 bunches, i.e. 20 10^9 Au ions or 20 10^{11} protons, do not receive enough of a kick for a clean dump. Each of these events typically causes a few quenches of triplet quadrupoles and can lead to serious damage in the experimental detectors.

* Work supported by Brookhaven Science Associates, LLC under Contract No. DE-AC02-98CH10886 with the U.S. Department of Energy
† drees@bnl.gov

ISBN: 978-3-95450-223-3 ISSN: 2673-7000 doi:10.18429/JACoW-NAPAC2019-WEPLH11

Quenches of arc magnets and dipole quenches in general are rare.

To prevent the damage to the experimental detector equipment, RHIC was in run16 operated with a 20 mm orbit bump in an arc downstream of the beam dump and upstream of the experimental detector. These arcs were arc 10 in the blue ring and arc 9 in the yellow ring, each protecting one of the then two experiments. Figure 3 shows the closed orbit with the two prefire bumps and their marginally smaller amplitude compensation bumps. In case of a prefire, due to the large

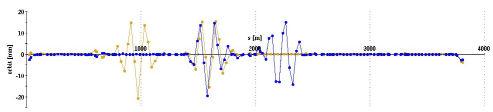

Figure 3: RHIC prefire protection bumps in the Blue and Yellow ring during run16.

orbit amplitude the offending beam that missed the beam dump would end up in the middle of that arc, quenching arc magnets instead of damaging experimental equipment. However, due to the superposition of closed orbit and dispersion function, each and every store caused continuous small losses of off-momentum beam in the arc area, as detailed in Fig. 4. The superposition of dispersion and large orbit

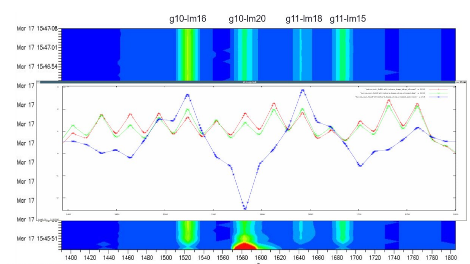

Figure 4: Several minutes of loss map data from the blue arc 10 area with the orbit bump and dispersion function inlaid.

amplitude accounts for a few 10 rad/h continuous loss rate at certain pronounced locations (green color). The names at the top indicate local loss monitors at these locations. At the time of the failure, these losses together with a series of single large loss events such as prefires, accumulated to a dose of about 11 kGy at the edge of the beam pipe, i.e. the location of the QPD. This number is an extrapolation from integrated loss monitor measurements at the outside of the cryostat and activation measurements of the damaged diode after its removal from the cryostat of B10-D19. Due to the elapsed time between failure and activation measurement ($\approx$ 3 weeks) as well as the different locations of loss monitors vs. QPDs, this value is only an approximation. There is no reliable activation measurement available from the second event since RHIC continued to operate and it took more than 4 months between failure and removal of the diode. However, judging from integrated loss monitor data from the area of Y7-D6, it appears that the accumulated dose in run17 at the time of failure was significantly less than for B10-D19 in the year before. However, as stated earlier, Y7-D6 is subject to some regular losses coming from the upstream collimators, albeit not in the month of run17 operation leading up to the damage. So far it was assumed that due to the yearly warm-up and cool-down cycles the up to 80% annealing of semiconductors would lead to an almost complete recovery of irradiated diodes.

QPD MEASUREMENTS

In run16, once we knew there was a problem within one voltage tap (VT), the individual magnets were measured in-situ in the tunnel since it cannot be determined from the outside which of the 6 magnets of one VT has an issue. Figure 5 shows the results of the voltage measurements of all diodes in the VT that hosts B10-D19 (red trace). It clearly identifies B10-D19 as the QPD with a problem. Once it was

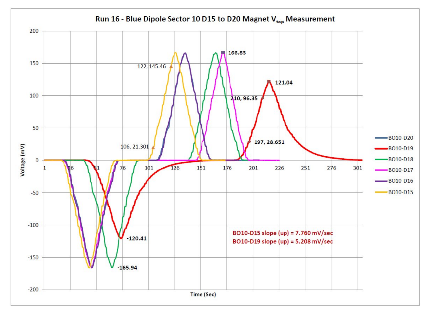

Figure 5: Voltage measurements of the 6 diodes sharing the same VT as B10-D19. All diodes show the same characteristics and amplitude, except for B10-D19 with a reduced amplitude and slope.

determined which diode was damaged, this RHIC sector was warmed up, the cryostat was opened and the diode replaced followed by a cool-down. Since RHIC could still be operated even with the damaged diode in place in run17, the offending assembly was removed once we had warmed up both rings after the run. Figure 6 shows the two damaged diodes from the two runs side by side. The top of each wafer in the picture

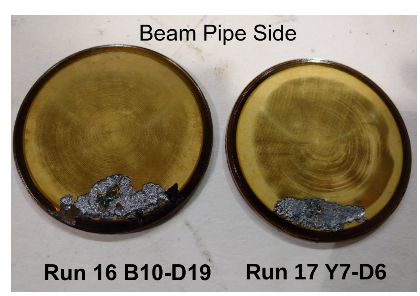

Figure 6: Damaged silicon wafers after removal from the tunnel without their assemblies.

would be the closest point to the beam pipe, the damaged areas are farthest away from it. This lead to the hypothesis that radiation damage to the silicon wafer prior to the magnet quench had altered the forward voltage the most in the area close to the beam pipe and the least far away, thus creating a gradient of the forward voltage across the diameter of the wafer. When the magnet quenched and the diode started

ISBN: 978-3-95450-223-3 ISSN: 2673-7000 doi:10.18429/JACoW-NAPAC2019-WEPLH11

to conduct current, the gradient caused current crowding in the area that became conductive first, i.e. far away from the source of radiation thus causing the burnt silicon. Alternatively, we discussed the possibility that a single event upset could be sufficient to cause current crowding without a history of prior irradiation. At this point we do not have enough data to disprove one or the other explanation. In an attempt to find other potentially radiation altered or damaged diodes, a total of 148 blue and 153 QPD forward voltages were measured in the tunnel while the magnets were still cold. Figure 7 shows the distribution of these measurements. Both rings give rather similar results, with forward voltages

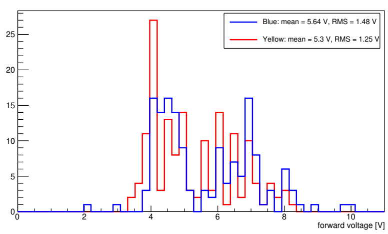

Figure 7: Measurement of the forward bias voltages, when the diodes become conductive, in Blue and Yellow while the magnets were still cold. All magnets except B10-D9 are within a range of 3V to 10 V.

between 3 V and 10 V, a mean of about 5.5 V and an RMS of about 1.5 V in Blue and 1.25 V in Yellow. B10-D19 is not part of this histogram. There is one diode, B10-D9, that is with 2 V lower than all others. However, in that area there is no history of heightened radiation, not in run16 nor in runs prior to that. The various forward voltages do not correlate with the location or the radiation history of the diodes.

Once the damaged diodes were removed from the tunnel they were available for test bench measurements at room temperature, at 77 K and 4.5 K. In order to measure the gradient across the wafer surface along the virtual line between pipe edge and burnt area, the two diodes and one spare diode were lacerated into 5 mm x 5 mm dice by using a laser[1]. When the individual dice were sufficiently electrically isolated from each other, we could measure the individual die's forward voltage one by one. Figure 8, left, summarizes the results of these measurements. Each data point corresponds to one die and is separated by 5 mm from its neighbor. "0" refers to the closest point to the beam pipe edge, however there are still a few additional millimeters between the edge of the wafer and the actual beam pipe. The gradient is clearly present at room temperature and vanishes at 77 K. The 4.5 K data has a significantly larger error and scatter, thus making it difficult to determine a gradient. B10-D19 demonstrates a larger effect, consistent with its history of integrated dosage being larger than that of Y7-D6. In addition, the room temperature forward voltage appears shifted to larger values for the two damaged diodes if compared to the spare. Since the in-situ measurements at 4.5 K of the diodes showed no correlation

[1] US Laser Corp., 20 W, λ 1064 nm, solid state laser (D. Elliott)

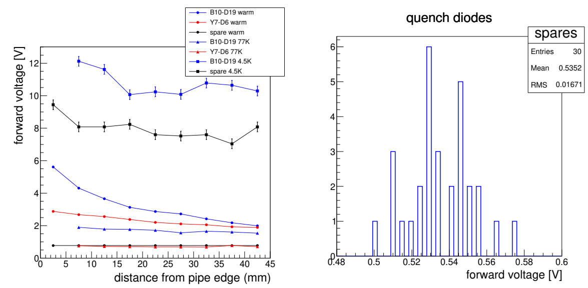

Figure 8: Left: die-by-die forward voltages for two damaged and one spare diode at various temperatures. Right: overall forward voltages of all available spare diodes at room temperature.

to location and history, we are in the process of repeating that measurement but at room temperature. The warm-up is still in progress and the measurements will be performed once the temperature settled. In the meantime, all available spare diodes were measured to confirm the typical mean value of about 0.53 V at a current of 100 mA with small deviations from this value. We expect to be able to pinpoint candidates of pre-damaged diodes with this measurement. We plan to remove one of those candidates to measure their gradient as was done for the two damaged diodes for verification.

CONCLUSION

After two QPDs were damaged in quench events and subsequently removed from the tunnel, we detected a burnt area as far as possible away from the beam pipe edge that is consistent with current crowding. We hypothesized a gradient in the diode's forward voltage across its diameter caused by a history of radiation damage retained in the diode through several cool-down and warm-up cycles. Dicing the diodes by means of lacerating the silicon surface with a laser enabled us to measure the forward voltage die by die in a laboratory setup. We could confirm a strong gradient present in the forward voltage, with its peak close to the beam pipe, i.e. the source of the radiation, and its minimum on the opposite side across the diameter of the wafer. In addition, the room temperature forward voltage was increased significantly in the damaged diodes when compared to our spares. By verifying this shift in warm tunnel diodes we expect to corroborate one or the other of our hypotheses while enabling us to pinpoint other pre-damaged diodes still inside the tunnel and installed in magnets that did not quench to date.

ACKNOWLEDGMENTS

We thank D. Elliott, G. Ganetis, A. Marusic and S. Seaberg for their support, input or discussions.

REFERENCES

[1] A. Drees *et al.*, "RHIC Prefire Protection Masks", BNL-107380-2015-TECH, 2015.

[2] Y. Tan *et al.*, "RHIC Abort Kicker Prefire Report", BNL-105539-2014-TECH, 2014.