Follow-up to the silicon damage at CDF of 3/30/02

Original DRAFT 4/17/02
Updated 5/02/02

spalding@fnal.gov
(with Chris Hill and Gino Bolla)

Description of the incident on 3/30/02

On Saturday 4/30/02 at about 15:30 store 1144 was aborted when the TEV quenched. The sequence of events in this abort is as follows: All eight Tevatron RF stations glitched, five of them stayed off and three stayed on (one proton and two antiproton). About five seconds later F1 quenched and aborted the beam. Three other houses also quenched. The loss pattern round the ring just prior to the abort showed high losses in both sets of low-beta quads at B0. While the losses at CDF increased during the five seconds, the main event was the abort, with SVRAD reporting 1.5 Rads on both proton and pbar sides. (For comparison a typical scrape at the start of a store registers about 5 Rads on the proton side and 1 Rad on pbar.) In five seconds the beam would be completely de-bunched so there was effectively no abort gap. The abort kicker rise time is about 2.5 µsec, but Vladimir indicated that the Tevatron will transmit beam around to the B0 quads for only about the initial 150 nsec of the kicker ramp. So we guess that the silicon received a few Rads in about 150 nsec (a rate > 10 MRad/sec).

The state of the silicon system after the incident did not appear to be particularly unusual. Some CAEN supplies tripped on digital overcurrent, and one feature of note is that many of the L00 bias circuits had either blown the fuses on their crowbar circuits or tripped in the power supplies. This feature of the crowbar circuit is under investigation.

Scope of the damage

In this incident six SVX ladders (of 360 total) were damaged, developing a particular set of symptoms know in the CDF silicon group as ìAVDD2 failureî. (AVDD2 refers to an analog voltage supplied to the SVX3 chip.) These six ladders are apparently randomly distributed in the detector. Interestingly, three of the ladders have Hamamatsu double-metal sensors, and three have Micron small-angle sensors. In all three of the Hamamatsu ladders operation of the entire z-sides was lost. For the Micron ladders, the damage is to individual chips.

An AVDD2 failure is characterized by the loss of analog power to the front-end of the damaged chip and the inability to pass the initialization stream from the back-end of this chip to the front-end and subsequently on to all other chips further along the initialization chain. Whereas the drop in analog power may be seen by the power supply system at the time of damage, the initialization failure, and subsequent power drop due to the loss of control of the other chips in the chain is found later, when the initialization sequence is resent.

No AVDD2 failures were found during fabrication or testing, but six such failures occurred during the first year of operation. The failure was intermittent in some of these parts, although their normal state was ìfailî with occasional periods ìworkingî. In a detailed study of these failures (CDF #5817 Affolder) it was determined that a likely cause of the damage was a failure in a silver epoxy connection for AVDD2 to the ìfingerî which provides connection to the front-end of the chip. To reduce the risk of further damage operating procedures were changed to minimize chip power cycles and other thermal changes. No batch-dependence has been found for this failure mode.

The ladders with AVDD2 symptoms have not been operated since their failure was detected. In principle part of the ladder can still be operated (upstream in the initialization chain from the failed chip), although their reliability is unknown.

Circumstantial Evidence

The only silicon power supply trips where on bias voltages (the fuses in the bias crowbar circuits for L00 also blew). No other electrical problems were observed, so an electrical cause for the incident is unlikely. This coincident loss of six ladders therefore points to being beam-loss related. The low dose is WELL below any normal radiation dose concerns, so it appears that parameter most likely responsible for this damage is the dose RATE. This is the primary line of investigation, although further consideration of an electrical link is still ongoing.

Initial Follow Up

After the incident the silicon system (other than the damaged ladders) was operated for one more store. The silicon was turned off on Monday April 1 and has not been operated since, other than for specific testing. The initial follow up consisted of:

Checkout of the radiation protection system
As far as can be determined the SVRAD system functioned OK. The dose was below our abort threshold of 12 Rads delivered within one second. The best estimate of the dose seen by SVRAD is 1.5 Rads.

Discussions with Experts
We had several meetings with the hybrid and chip designers, detector experts and Tevatron experts.

Bench Tests
Attempts to produce damage by shorting the analog voltage at redundant pads on the chips, and by using a 30 μW laser failed. In fact the tests demonstrated that the power supply trips prevent damage to wire-bonds as expected.

Analysis of the earlier AVDD2 failures in light of this new information
Due to the nature of the failures and record keeping (the logging of voltages and currents, and DAQ state changes was improved several months into the run) exact correlation with beam conditions and incidents is difficult. Nevertheless, sufficient evidence exists for four of the earlier failures to indicate that they MAY have been initiated by a beam accident. The four most likely correlations include two cases with an end of store quench, one case where a feeder failure aborted a store and one where the silicon detector was mistakenly under full bias during proton injection at the start of a store. However, one failure was identified when it was first connected at B0. Beam had been delivered, but this ladder had not been powered. Few conclusions can be drawn from this analysis. It is possible that we have a fatal failure mode which can be triggered from different conditions.

Evaluation of Silicon Operating Interlocks and Procedures
Under the assumption that this is a dose RATE dependent failure, the silicon operating procedures were reviewed.

- The main issue for the operating procedures is a determination of whether ì standbyî (bias voltage off, chip voltages on) is a safe condition for this failure mode. Standby is preferred to ì offî (all voltages off) because the transition to off causes a significant thermal cycle for the detector.
- In terms of interlocks, the only device in the accelerator that can cause such a rate is a kicker magnet (either the abort kickers or injection kickers). The interlocks in place so far this run, the SVRAD system, are intended to limit steps in the total integrated dose. They do not protect against a low dose at high rate.

New Considerations for Dangerous Beam Conditions

Note that all of the following conditions have occurred in Run II.

1. DC Beam ñ procedural mitigation
   There is some amount of ì DCî beam in the abort gap during many stores, and is usually reduced using the electron lens. The CDF silicon is in standby for the end of store abort (still assumed to be a safe state at this point). The risk of damage from an unexpected abort is probably low, and the extent of damage would probably be low. CDF should learn to measure the amount of DC beam. This risk can be handled procedurally.

2. p injection at 150 and pbar injection before cogging - procedural
   There is typically 2-3% DC beam after injection, before ramping to 980 Gev. Pbars are injected into the abort gaps and then cogged into the right bunch times. Aborts at this time present a higher risk. Again, CDF is in standby. (Note one AVDD2 failure listed above.) Again, procedural mitigation.

3. RF failure - interlock

The loss of one RF station is not uncommon. Some beam is unbunched, luminosity falls, losses rise, and then the DC beam is cleared out over several minutes (10-15 minutes without the electron lens). A more severe RF problem can cause a quench and abort during this period of high DC beam ñ the scenario that started all this. Beams Division has implemented an interlock to abort the beam within about one µsec of an RF failure. This is quick enough to avoid the creation a significant amount of beam in the abort gap. The signal used is a vector sum of the phase signals from all eight cavities, with a threshold set at 1.5 of the level change when one cavity is lost [ref 1].

4. Loss of the $AA marker ñ interlock
   This occurred near the start of Run II, when the abort was pulled but the kickers failed to fire due to a missing $AA timing marker which indicates the time of the abort gap. It is not known how reliable the marker is. As a result CDF asked that the kickers be set up to wait for up to two turns, and if the marker is still missing to fire without the timing signal. This was to limit the accumulation of a total dose. If rate is the issue, this is the wrong thing to do. There is something like a 1/6 chance of sending one bunch into the B0 quads, leading to probably about 3 or 4 times the dose of the incident of 3/30 in just a few nsec. Beams Division is designing a phase-lock loop system to use a synthetic timing marker so that the kicker will still fire at close to the right time if $AA is missing. There is a discussion within Beams Division whether the system should pull the abort itself if $AA is found to be missing.

5. Injection kicker fail to fire - procedural
   This is known to have happened once this run, on 4/9/01, when there was a synchronization problem between the Main Injector and the Tevatron. SVRAD recorded a dose of 194 Rads in a very short period of time and pulled the abort. The silicon detector was off at the time. This failure is probably rare, and CDF should be in standby during injection.

6. Abort kicker prefire ñ partial mitigation with collimator
   There are 10 abort kicker magnets at A0 (five proton and five antiproton) in the Tevatron. (The abort at C0 is used during injection.) In a prefire one of the kickers self-fires, triggered by internal breakdown. The system reacts by firing all the other kickers 2.5 µsec later. For this 2.5 µsec the ramp up of the field seen by a proton passing through the kicker system (which is about 60m long) is five times slower than the normal kicker ramp up. So Vladimirís 150 nsec ìdanger windowî becomes more like one µsec. Something like 2/3 of the time, beam will be passing through the kickers during the prefire, typically two maybe three bunches. This could result in about 10 times the dose the the 3/30 incident, and over significantly less time. Abort kicker prefires occurred at a rate of about four times a year in Run I. In Run II there have been at least four prefires ñ two without beam and two at C0 with beam but during studies when the silicon detector was off. One of these events produced 5 Rad in SVRAD during the abort, the other apparently zero. The only known mitigation is to add collimation in A sector to

shadow the B0 quads. Beams Division will study whether this looks feasible (including a Nikolai Mokhov simulation). If so, the installation could be in June. In any case, this would reduce, but not eliminate the dose.

Conclusion on Abort Mitigation

Under the assumption that rate is the key, lesser risks can be handled procedurally but new abort interlocks are required to mitigate all but one of the most serious risks. This last risk, kicker prefires, may be the most serious and is not easily mitigated. While the accident rate is not high, there is a probability that the damage could be significantly worse than the 3/30 incident.

Booster Testing

Two series of tests were performed at the Booster. In the first test, on 4/19/02, a single Hamamatsu ladder was exposed in the 8 Gev proton beam in a series of 74 tests, beginning at roughly the expected dose of the 3/30 incident and increasing to about 50 times the dose by the end. Exposure to the sensor and chip were separated, as were the off, standby and on states. In the second series, four ladders (two Micron and two Hamamatsu) were exposed to 24 tests in which a steel target was used to generate spray (producing about a factor 4.5 increase in dose). The dose in each of these tests was in the range 100 to 1000 times the 3/30 incident, delivered in about the same time interval. The DAQ was times to put the chip into various states (acquire. Digitize, readout, preamp reset) at the time of the exposure.

The leakage current increase was measured for each exposure. In the first test, without a target, the relationship between current increase and Booster intensity was determined. In the testing with a target, this served as a calibration for a MIP flux equivalent for the leakage currents.

New non-fatal failure modes were observed in both standby and on conditions ñ where non-fatal means that the ladders could be recovered with subsequent re-initialization. These modes included three features:

- Skipping of the pipeline cell logic as a result of dose in the sensor
- New high-current states of the chips, not seen before, when the chips were dosed
- Chips disappearing from the DAQ completely when the chips were dosed

The behavior was very repeatable. Subsequent analysis of the data revealed that very specific behavior of individual chips repeated. Further analysis will likely provide significant clues to the failure mode. A note will be written on the results of these tests.

Following the Booster tests, a re-analysis of the 3/30 incident found this same current behavior in about 100 ladders.

There were no AVDD2 victims for autopsy. All chip failures were recovered with an initialization sequence.

It is difficult to estimate the fluence in the incident itself. The SVRAD reading is normally considered to be within a factor two or so from the dose at SVXII layer 0. This would suggest a dose of 2-3 Rads, equivalent to about $1x10^8$ MIPS cm^{-2}. This is probably good to within a factor 10. If the loss pattern was the same, a kicker prefire could result in as much as ten times this fluence. In the second series of Booster exposures, the fluence per exposure was determined to be $1\text{-}2x10^{11}$ MIPS cm^{-2}. So we can conclude that four ladders survived a series of 25 exposures each of which was 10-100 times the dose rate in a ìworst-caseî accident condition, and within this series of tests all DAQ conditions of the chip were experienced.

Current Position on Operating the Silicon Detector

The nature of the failure mode is not yet determined, nor is the nature of the link between the accident conditions in the Tevatron and the failure of the ladders been confirmed. Nevertheless, that there is some link seems clear and it is prudent to reduce the likelihood of beam conditions similar to the 3/30 incident.

To reduce the risk for further silicon operation the probability of similar accidents in the Tevatron must be lessened, and the extent of damage from any one incident must at least be shown to be limited at some level. Taking these concerns in the reverse order; since a dose rate effect is the leading candidate for the link, the Booster tests can be taken to provide a limit on the extent of damage expected for a similar or even the worst-case accident. The test results would suggest that while non-fatal failures are very common under these accident conditions, fatalities are more rare ñ even at significantly higher doses. And indeed from the 3/30 currents it appears that 100 ladders experienced the non-fatal failures. Maybe the six were ìweak partsî or simply statistical.

The likelihood of similar accidents is reduced by:

- the RF interlock, which is now in place [ref 1]
- the missing $AA marker interlock ← expected to be implemented in about one week [ref 2]
- the addition of a collimator at A17 ← this will be installed during the June shutdown [ref 3]
- lowering the trip threshold for the analog current in the silicon power supplies so they will trip on the high-current failure mode (with the previous threshold the chips would remain indefinitely in the high-current state)
- the following changes to the CDF operating procedures to avoid running the silicon in abnormal stores:
    - require Tevatron electron lens to be on
    - require LOSTP , B0PHSM and B0PBSM (halo counters gated on beam crossings and gated in the abort gaps) to be below thresholds
    - limits on the spikyness and the rate of rise of these three scalers

(These thresholds will be adjusted as the luminosity increases. Their role is not really to place limits on the losses, but to identify **abnormal** loss behavior.)

With these changes CDF is prepared to operate the silicon detector for physics data-taking while continuing to investigate the AVDD2 failure mode.

References

1. email from Mike Church on the RFSUM threshold
2. email from Greg Vogel on the timescale to implement the missing $AA protection
3. note by Mike Church on the effectiveness of a collimator at A17

Ref 1

Subject: Tevatron RR aborts
Date: Wed, 24 Apr 2002 21:25:24 -0500
From: Mike Church <Church@fnal.gov>
Organization: Fermilab
To: jeff spalding <spalding@fnal.gov>
CC: vladimir shiltsev <shiltsev@fnal.gov>, dmitri denisov
<denisovd@fnal.gov>,
todd johnson <tjohnson@fnal.gov>

Jeff, I have used the datalogger to investigate past Tevatron RF cavity
trips. I have looked for changes in either the total proton cavity
voltage
(T:RFSUM) or the pbar cavity voltage (T:RFSUMA) since 1/1/2002. I have
found 6 trips during Collider stores. These have occurred at the
following
times:

Jan 5, approximately 22:00
Jan 13, approximately 23:00
Jan 19, approximately 20:00
Jan 20, approximately 14:00
Mar 22, approximately 22:00
Mar 24, approximately 23:00.

When these trips occur there is a clear drop in proton bunched beam
lifetime, but no clear drop in luminosity lifetime. There is no clear
drop
in (bunched beam intensity)/(total beam intensity) at the time of the
trip,
which is measurable to <.5%.

Assuming a maximum immediate generation of .5% of DC beam this amounts
to
36E9 DC beam (at 200E9/bunch). The fraction of this beam which would
be
directed to the B0 LB quads in an abort is 150nsec/21microseconds =
.007,
which amounts to .25E9. We previously estimated that the accident of
3/30/2002 which is coincident with the loss of 6 Si ladders was ~60E9
into
the B0 LB quads.

In light of this, I think that setting the Tevatron cavity RF trip
level at
1 cavity is unreasonable. I am going to change the Tevatron RF trip
level
from 1 cavity to 2 cavities, unless you can make a reasonable argument
otherwise. We did lose a store Monday on a single cavity trip.

I have plots of the data if you wish to see them.

Mike

Ref 2

Subject: Re: $AA marker protection
Date: Fri, 03 May 2002 10:50:47 -0500
From: Greg Vogel <vogel@fnal.gov>
To: Jeff Spalding <spalding@fnal.gov>
CC: Mike Church <Church@fnal.gov>, todd johnson <tjohnson@fnal.gov>

Jeff,
We plan to have a modified C479 module ready for bench testing late next
week. Please note that's also when my group gets packed up and moved from our
present location to the booster tower (along with all the rest of the
Controls hardware groups). That will tie us up for a week or two. Still, I
expect that we should be ready to plug one into the Tev before the end of the
month.
Greg

Mike Church wrote:

> Jeff, Greg Vogel has been counting lost $AA markers for 10 days now.
> There have been no unaccounted for lost $AA's during this time (one was
> missed when TLLRF was rebooted). See T:TVAAPD in the datalogger.) Greg,
> could you please give Jeff an estimate when you will be ready to install
> the new 489 module with abort on missing $AA. Thanks. Mike
>
> Jeff Spalding wrote:
>
> > Mike,
> >
> > Any update on implementing the $AA marker protection? We're trying to
> > finish things off so we can hand silicon operation back to the shift
> > crew.
> >
> > thanks
> >
>

Ref 3

A0 Abort Prefire Notes

M Church 04/24/02

There are 5 A0 proton abort kickers and 5 A0 pbar abort kickers. A clean abort requires that all 10 kickers fire at once. A "prefire" refers to the case where a single kicker fires prematurely, presumably due to a thyratron breakdown. In this case, the kicker field will rise to ~5.8kG in ~2.5μsec[1] before the prefire is detected and the remainder of the kickers fire. A prefire is not necessarily synchronous to the abort gap. The kicker length is 1.9m, therefore 5.8kG will generate a kick of 337 μrad (downward) at 980 GeV.

Will a collimator located at the A11 straight section or A17 straight section provide some protection to the CDF detector from a prefire at low beta (LB)? Dean Still and Jerry Annala have made inspections of these straight sections and are confident that a collimator would fit in either location.

Previous aperture scans done at B0[2] indicate that the aperture there is close to the design aperture. The design limiting vertical aperture is at the center of A4Q3 with a beampipe radius of 34mm. The following table shows the vertical beta functions and phase advance from each kicker to the u.s. end of the A0 proton abort, the u.s. end of the A11 warm straight, the u.s. end of the A17 warm straight, and the center of A4Q3 at LB.

Kicker	**β_y at kicker (m)**	**$\Delta\phi_y$ to A0 abort (deg)**	**$\Delta\phi_y$ to A11 (deg)**	**$\Delta\phi_y$ to A17 (deg)**	**$\Delta\phi_y$ to A4Q3 (deg) (mod 360°)**
P1	36.9	20	33	222	152
P2	40.7	17	30	219	149
P3	45.3	14	27	216	146
P4	50.3	11	24	213	143
P5	55.5	9	22	211	141
A5	175.5	x	4	193	123
A4	187.6	x	3	192	122
A3	200.8	x	2	191	121
A2	214.4	x	1	190	120
A1	227.6	x	1	190	120

The vertical beta functions at the A0 u.s. abort, A11, A17, and A4Q3 are 87.0m, 211.0m, 32.2m, 1071.0m, respectively.

In order for the 3σ (@20π-mm-mrad) edge of the beam to just touch the physical aperture at A4Q3 requires single kicker strengths listed in the following table.

kicker	**kicker strength (μrad)**	**deflection @ A0 u.s. abort - 3σ (σ)**	**deflection @ A11 - 3σ (σ)**	**deflection @ A17 + 3σ (σ)**
P1	305	-21.4	-20.9	24.9
P2	265	-17.4	-20.2	21.8
P3	231	-13.9	-15.5	19.2
P4	204	-11.0	-14.8	16.9
P5	185	-9.3	-12.2	15.6
A5	78		-6.0	7.1
A4	75		-4.0	6.8
A3	72		-3.7	6.4
A2	69		-3.3	6.1
A1	67		-3.3	6.1

This table also shows the required location of a collimator at A0, A11 or A17 to also just intercept the 3σ beam envelope. For example, in order for a collimator at A11 to start intercepting beam kicked by a P1 prefire before it hits A4Q3, it must be placed no more than 20.9 beam σ's from the beam center.

In order for the opposite 3σ edge of the beam to just touch the physical aperture at A4Q3 (ie, extinguish the beam at A4Q3) requires single kicker strengths listed in the following table.

kicker	**kicker strength (μrad)**	**deflection @ A0 u.s. abort - 3σ (σ)**	**deflection @ A11 - 3σ (σ)**	**deflection @ A17 + 3σ (σ)**
P1	424	-31.6	-30.8	36.5
P2	368	-26.0	-29.9	32.1
P3	321	-21.2	-23.4	28.5
P4	283	-17.2	-22.4	25.4
P5	258	-14.7	-18.7	23.5
A5	109		-10.2	11.7
A4	104		-7.3	11.2
A3	99		-6.9	10.8
A2	95		-6.4	10.3
A1	92		-6.1	10.3

This table also shows the required location of a collimator at A0, A11 or A17 to also just intercept the opposite 3σ beam edge. For example, in order for a collimator at A17 to intercept any beam kicked by an A1 prefire before it is completely extinguished at A4Q3, it must be placed no more than 10.3 beam σ's from the beam center.

Conclusions:
1) Collimators at A11 or A17 would probably be effective in protecting CDF from proton abort kicker prefires. A17 appears to be more favorable than A11.
2) Collimators at A11 or A17 would not be very effective in protecting against antiproton abort kicker prefires
3) Moving the vertical orbit at A0 downward close to the A0 abort would probably also provide protection against proton abort kicker prefires.

This simple analysis does not consider the effect of an A11 or A17 collimator on "normal" losses.

[1] C. Jensen
[2] 2001 Tev Elog entry #317, 11/26/01; it might be worthwhile to repeat these scans at this time