RHIC Run 16 Radiation Upsets

Kevin Brown

2016 RHIC Retreat

July 28-29, 2016

Acknowledgements

- Nick Franco, Bill Eisele
  - Monitored Network problems
  - Coordinated with controls HW
- Controls HW (Charles Theisen, Ralph Schoenfeld)
  - Monitor, maintain, & replace FEC's
  - Check in-ring fiber quality
- John Morris
  - Analysis scripts and web data pages for network and FEC reset statistics
- Al Marusic, John Morris, Jim Jamilkowski
  - ram pattern measurements now logged
  - Watch FEC's very closely
- Roger Smith (JTM's summer intern)
  - Analysis of reset and memory pattern data
- Angelica Drees
  - Just a good resource – keeps close watch on losses
- Kin Yip
  - Simulations of beam loss and radiation environment
- Peter Ingrassia – operations statistics

Outline

- History/Background
- Alcove Network switch resets/failures
- Memory pattern data
- Mitigation strategies

History

- Radiation Upsets in Alcove equipment have been a problem since RHIC began operating
- Problems have been tactically mitigated
  - Going to Radiation resistant VME Power Supplies
  - Using ECC memory in (all ps, not qd or inst) FEC processor boards
  - Using MRAM for WFG boards
  - ado cache data for alcove FECs is stored on file server - not on local RAM that is vulnerable to rad upsets

Todays Vulnerabilities

- Network switches require frequent resets and even fail (2 replaced in Run 14, one replaced in Run 16)
- MADCs become corrupted and stop reporting
- Alcove FECs & WFGs still fail and require resets at a higher rate than in other locations
- QD & BPM FECs can still hang and require resets
- UPS units in alcoves have failed (11C & 5C)
- Fiber runs in RHIC tunnel darken
- Digi TS need AC resets at times
- VME PSs still fail at times (rare)

Alcoves

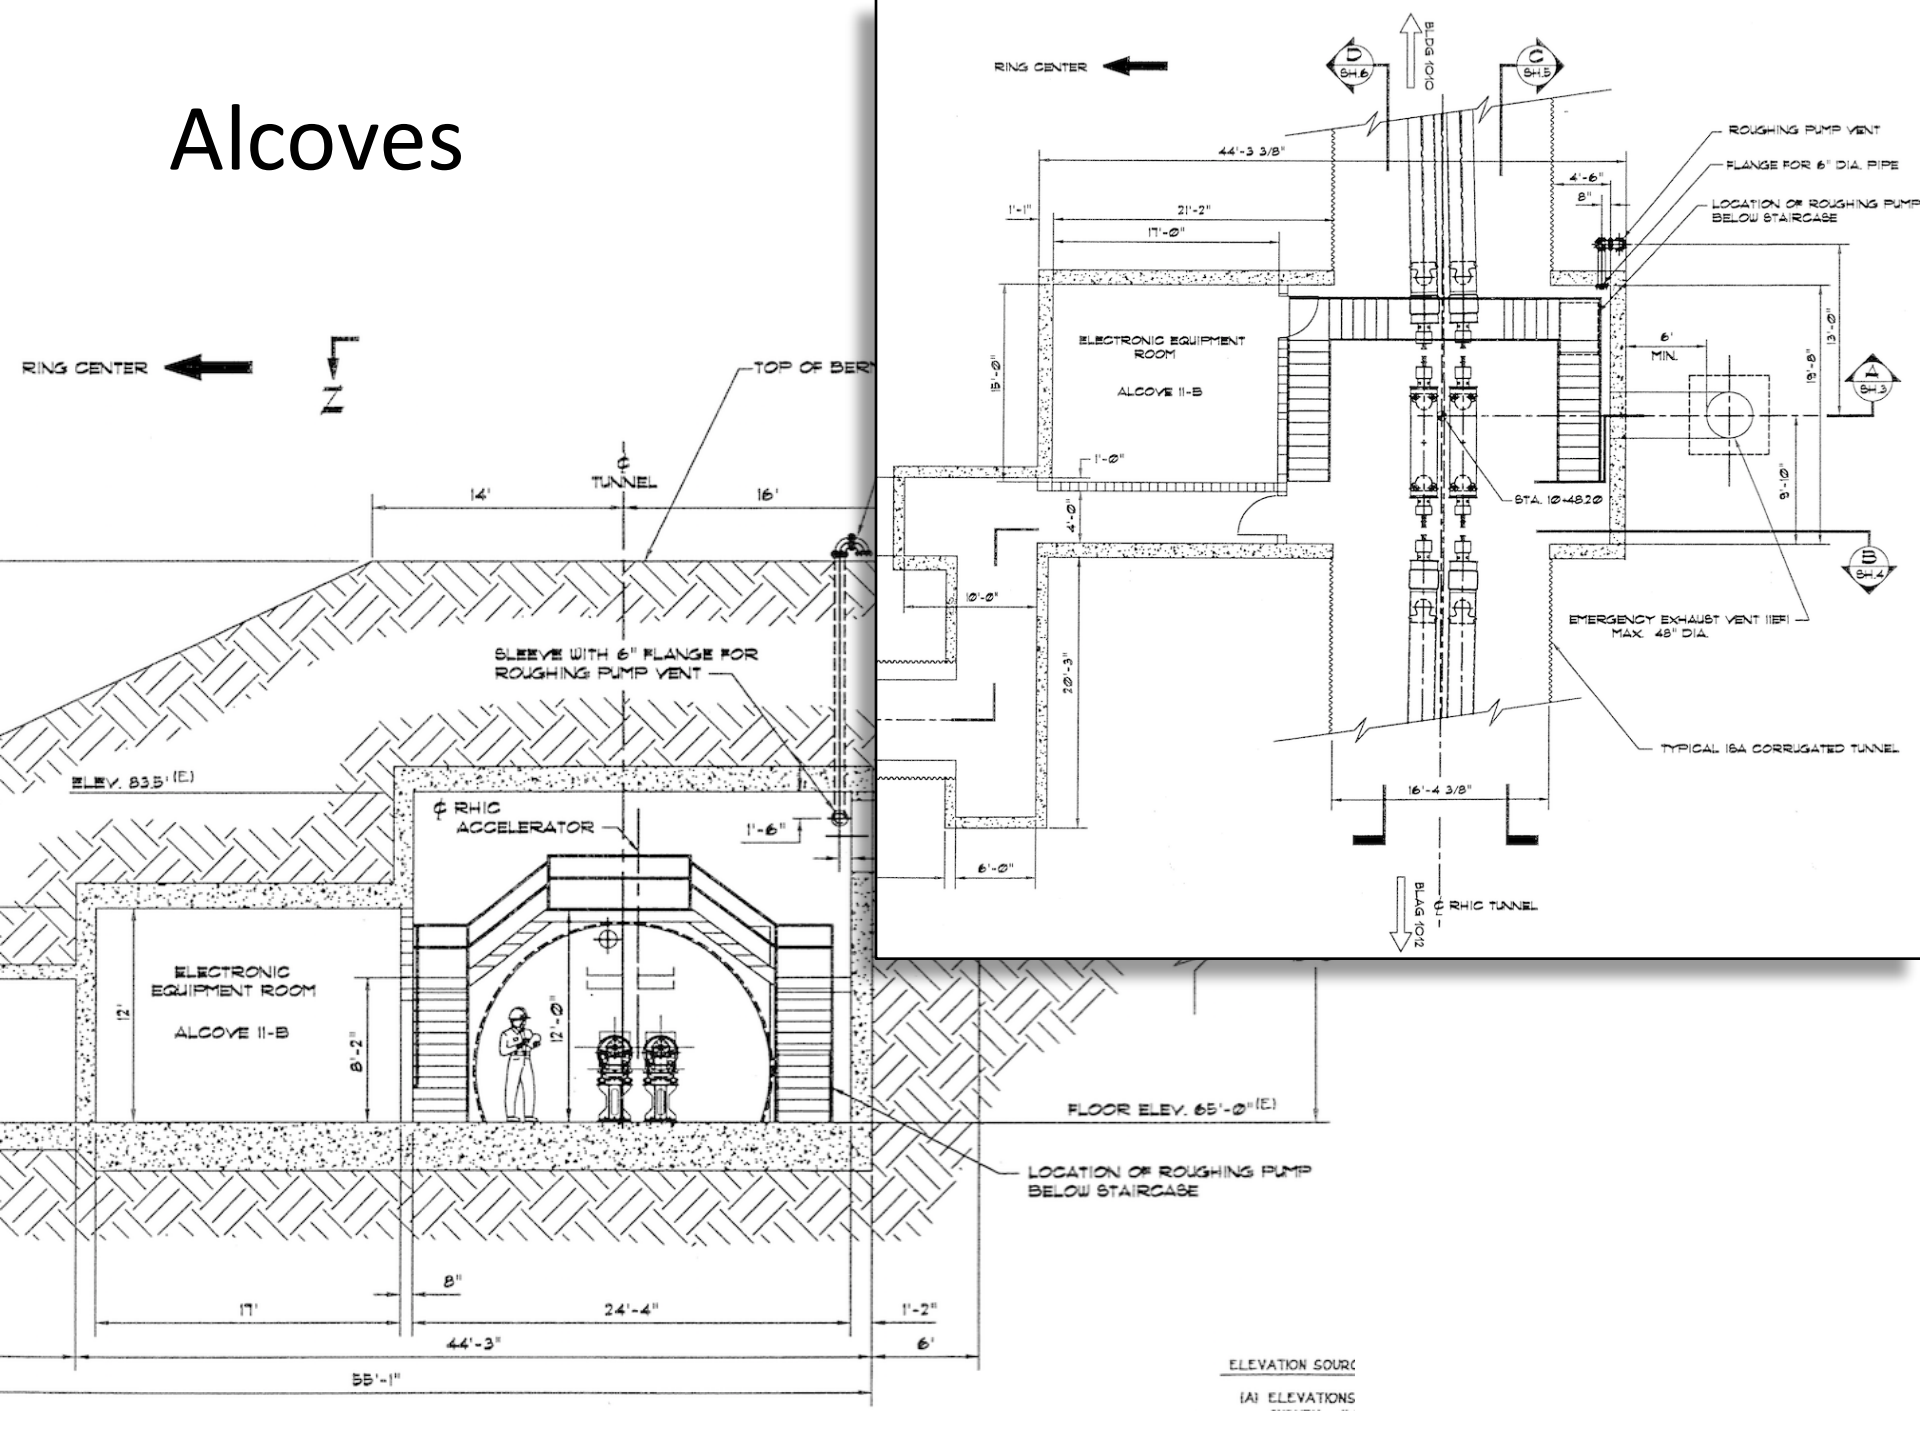
RING CENTER
TOP OF BERM
℄ TUNNEL
14'
16'
SLEEVE WITH 6" FLANGE FOR ROUGHING PUMP VENT
ELEV. 83.5' (E)
℄ RHIC ACCELERATOR
1'-6"
ELECTRONIC EQUIPMENT ROOM
ALCOVE 11-B
12'
8'-2"
12'-Ø"
FLOOR ELEV. 65'-Ø" (E)
LOCATION OF ROUGHING PUMP BELOW STAIRCASE
8"
17'
24'-4"
1'-2"
44'-3"
6'
55'-1"
RING CENTER
BLDG 1010
D SH.6
C SH.5
44'-3 3/8"
1'-1"
21'-2"
17'-Ø"
ROUGHING PUMP VENT
FLANGE FOR 6" DIA. PIPE
4'-6"
8"
LOCATION OF ROUGHING PUMP BELOW STAIRCASE
ELECTRONIC EQUIPMENT ROOM
ALCOVE 11-B
15'-Ø"
1'-Ø"
4'-Ø"
6' MIN.
13'-Ø"
19'-8"
A SH.3
9'-Ø"
B SH.4
STA. 10+48.20
Ø'-Ø"
20'-3"
EMERGENCY EXHAUST VENT (IIEF) MAX. 48" DIA.
TYPICAL 16A CORRUGATED TUNNEL
16'-4 3/8"
6'-Ø"
BLDG 1012
℄ RHIC TUNNEL
ELEVATION SOURC
(A) ELEVATIONS

Equipment Layout in Alcoves

- 18 Alcoves in RHIC, 3 per ARC.
- Two basic layouts – rack use is always the same

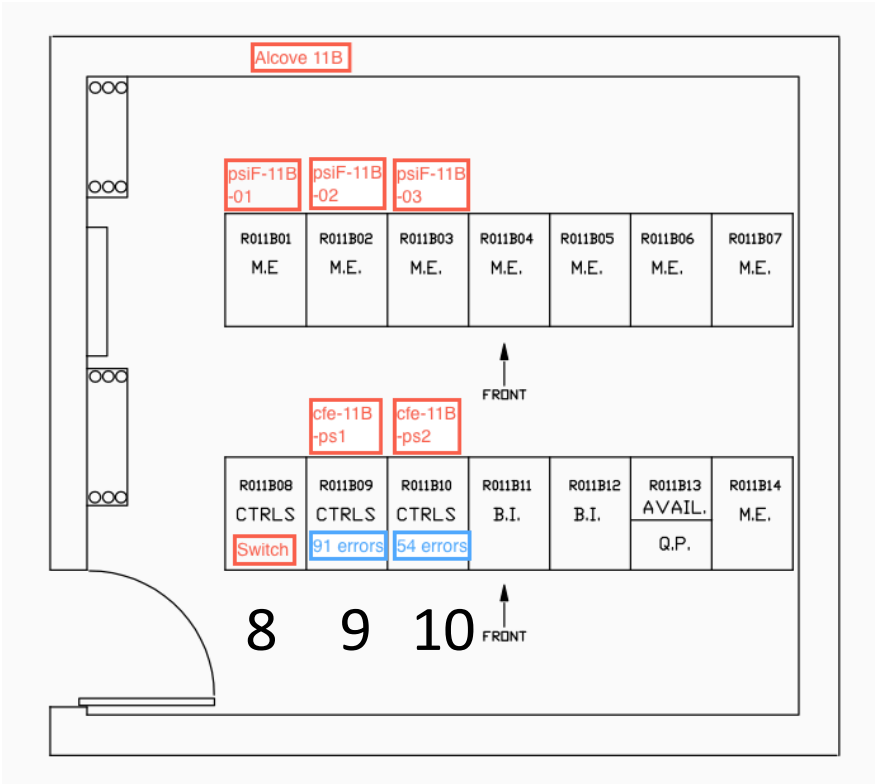

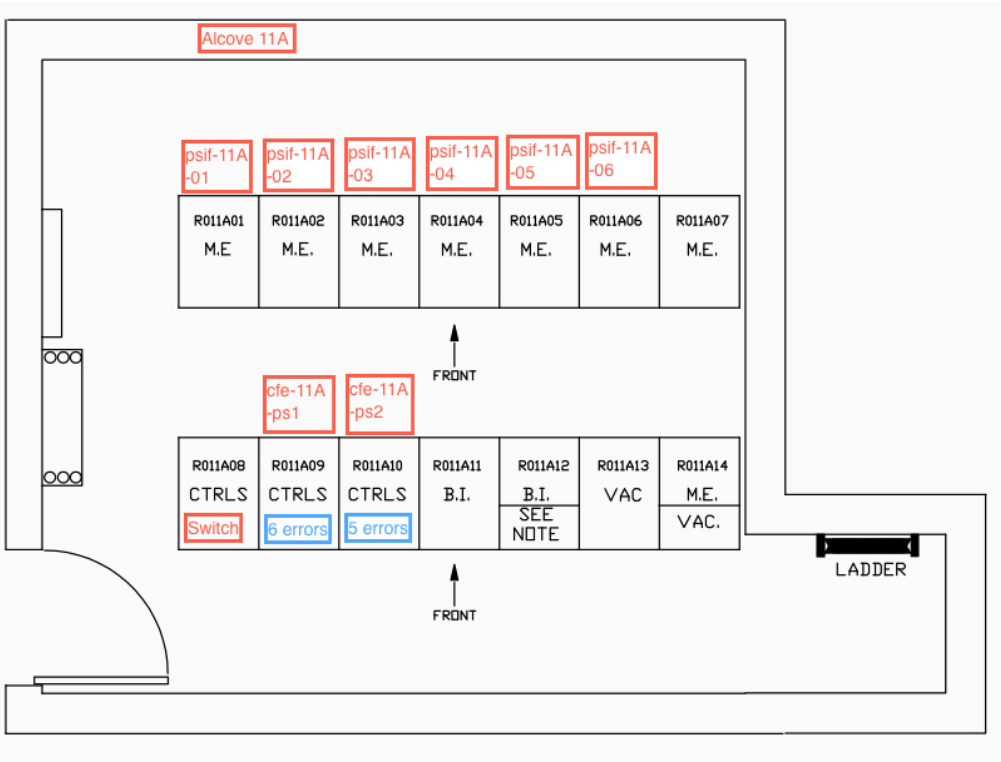

Rack 8

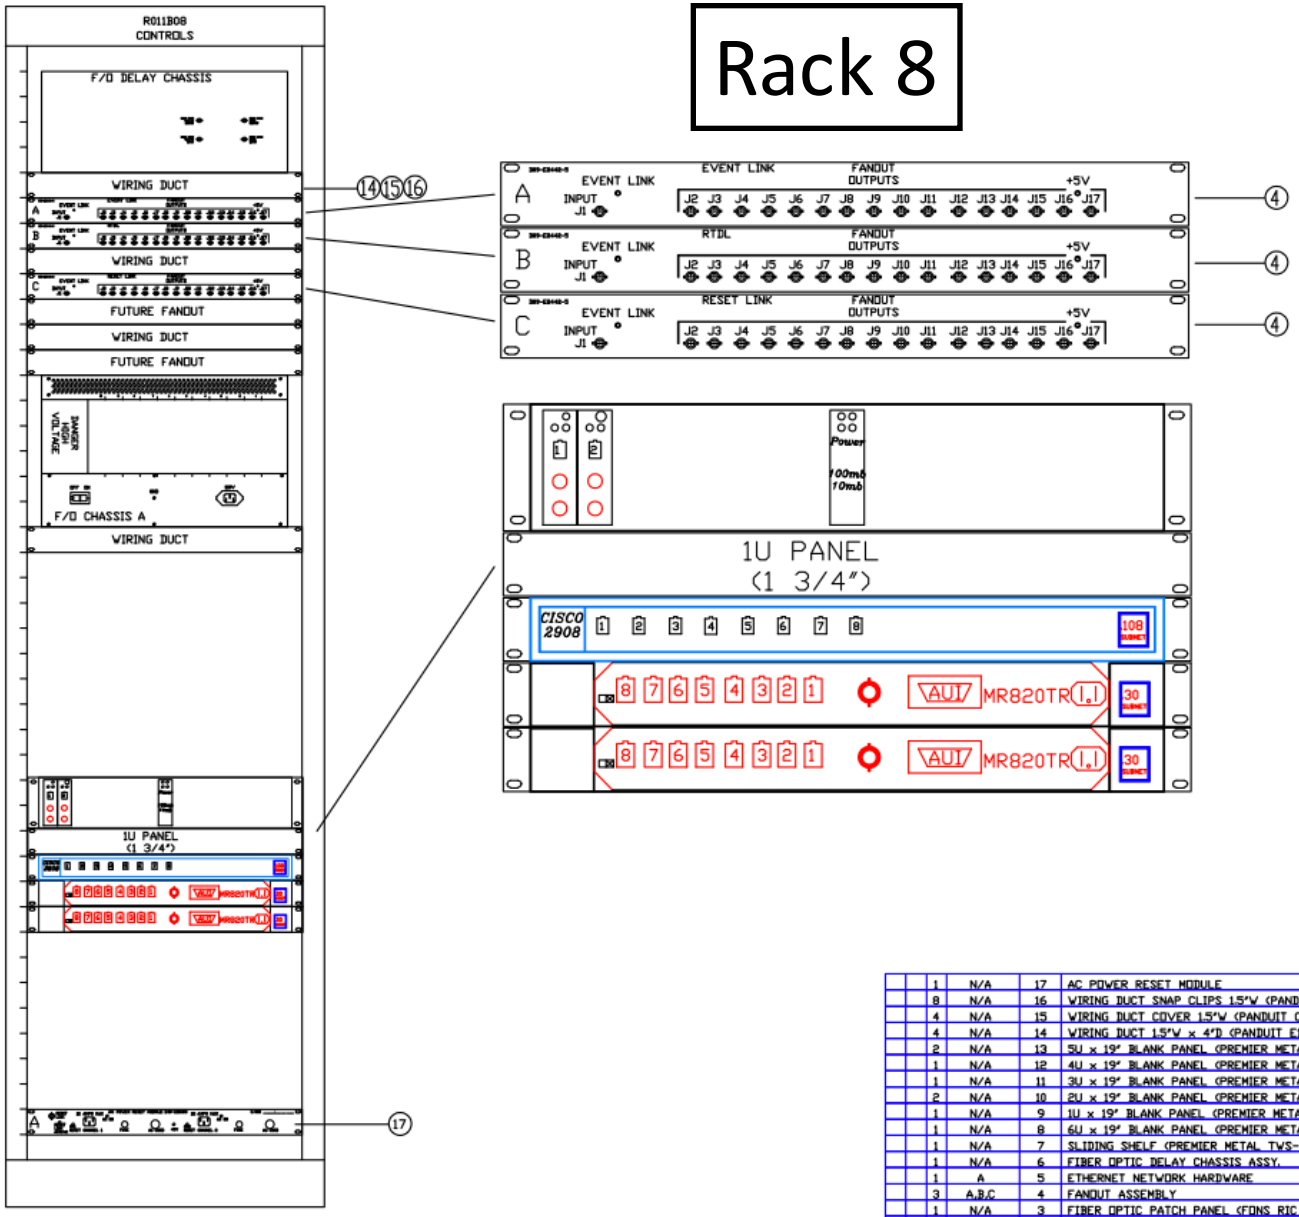

REAR VIEW

		1	N/A	17	AC POWER RESET MODULE
		8	N/A	16	WIRING DUCT SNAP CLIPS 1.5"W (PANDUIT
		4	N/A	15	WIRING DUCT COVER 1.5"W (PANDUIT C1.5
		4	N/A	14	WIRING DUCT 1.5"W x 4"D (PANDUIT E1.5X
		2	N/A	13	5U x 19" BLANK PANEL (PREMIER METAL
		1	N/A	12	4U x 19" BLANK PANEL (PREMIER METAL
		1	N/A	11	3U x 19" BLANK PANEL (PREMIER METAL
		2	N/A	10	2U x 19" BLANK PANEL (PREMIER METAL
		1	N/A	9	1U x 19" BLANK PANEL (PREMIER METAL
		1	N/A	8	6U x 19" BLANK PANEL (PREMIER METAL
		1	N/A	7	SLIDING SHELF (PREMIER METAL TWS-319
		1	N/A	6	FIBER OPTIC DELAY CHASSIS ASSY.
		1	A	5	ETHERNET NETWORK HARDWARE
		3	A,B,C	4	FANOUT ASSEMBLY
		1	N/A	3	FIBER OPTIC PATCH PANEL (FONS RIC 72
		1	N/A	2	F/O CHASSIS A
		1	N/A	1	19" RACK, 77"H RACK SPACE
-3	-2	-1	REF DESIGNATOR	ITEM NO.	
QTY REQD					

Network Switch & FEC Reset Data

FECs may be rebooted or reset for a variety of reasons - some not indicative of a problem with the FEC.

All FEC reboots/resets are counted but the reasons are only recorded when resets are performed with the "fit" application. The data collected only counts resets that appear to be associated with some kind of failure.

At times multiple resets are attempted (and counted in statistics) for a single problem.

Run 14, 100 GeV/n Au-Au Alcove network switch radiation upset locations

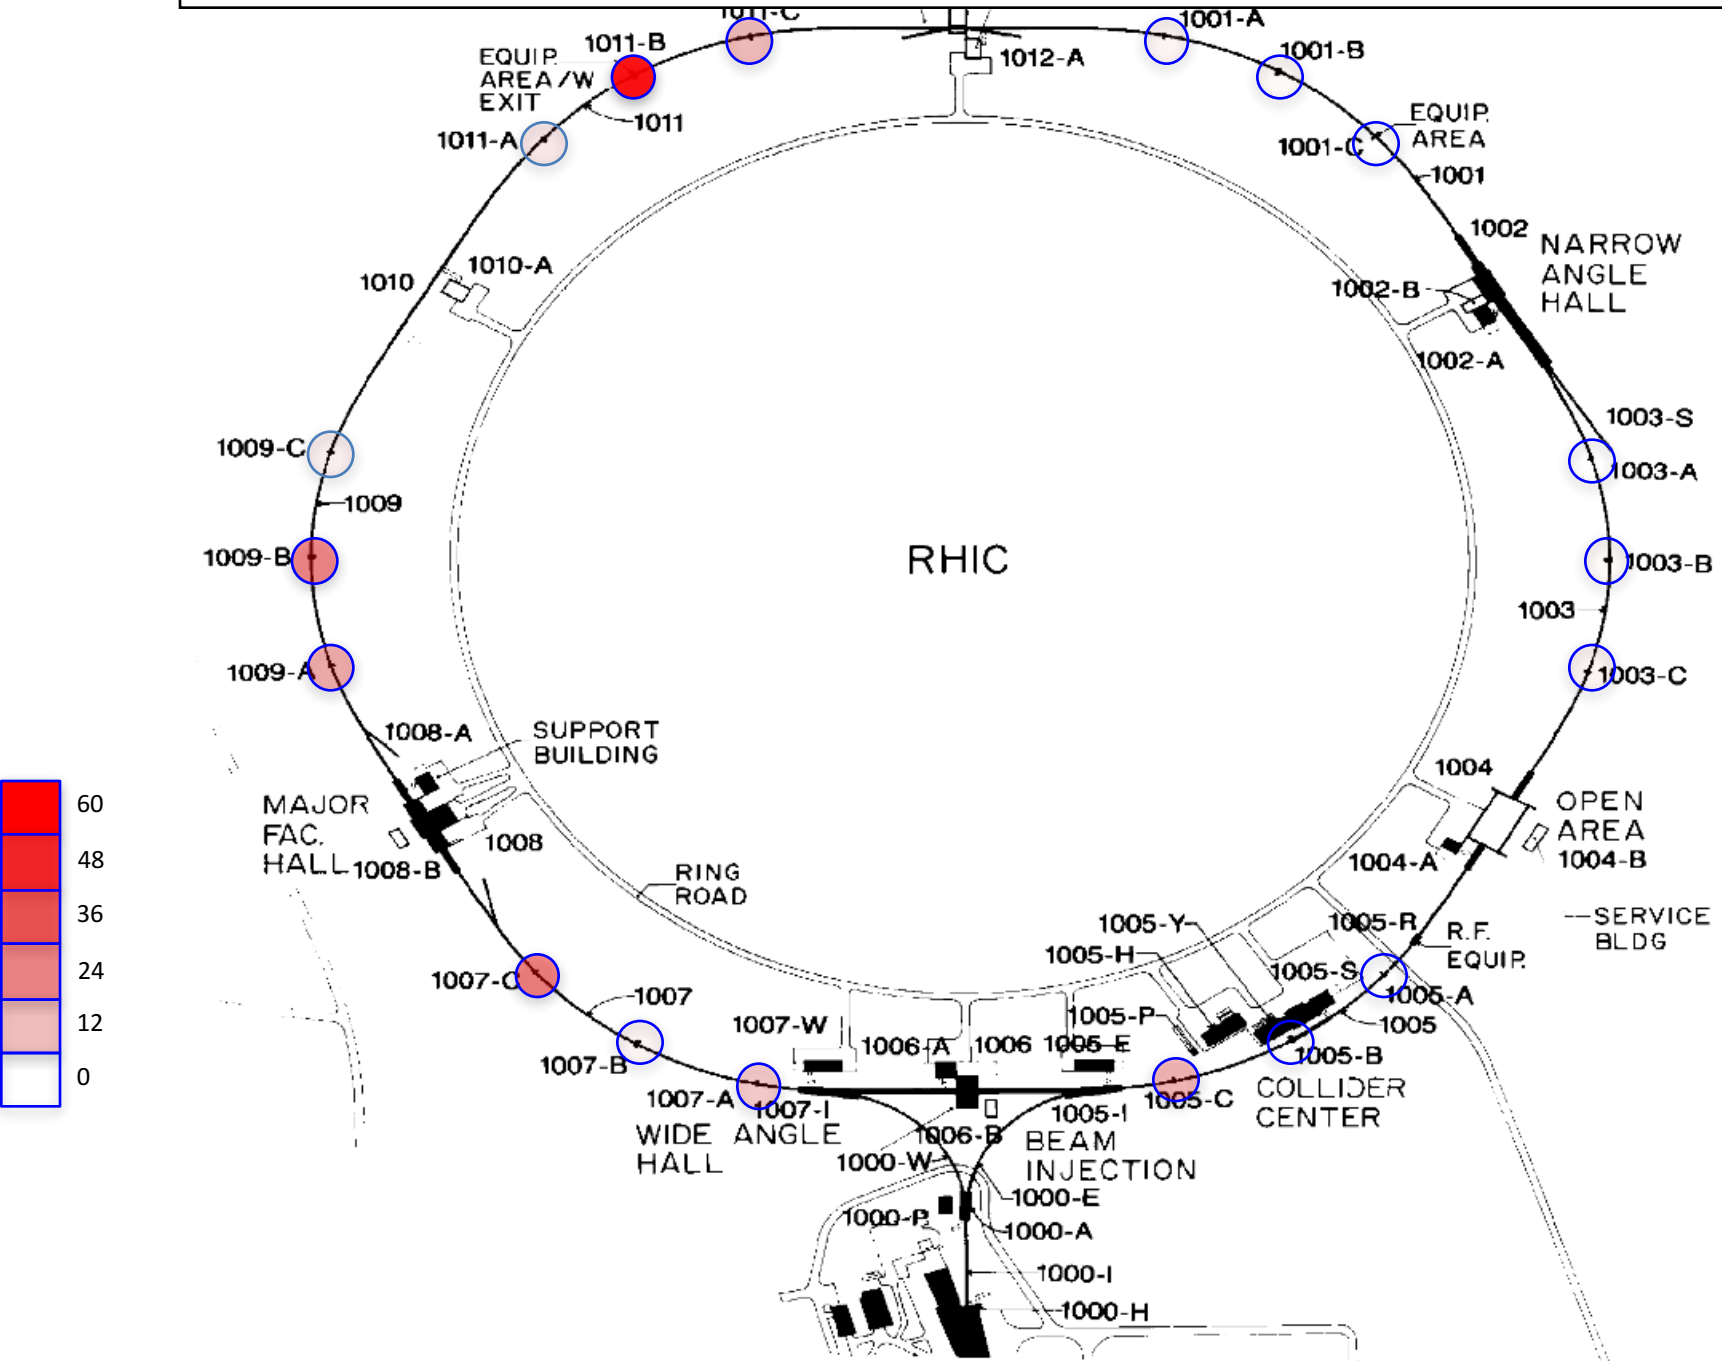

Run 16, 100 GeV/n Au-Au Alcove network switch radiation upset locations

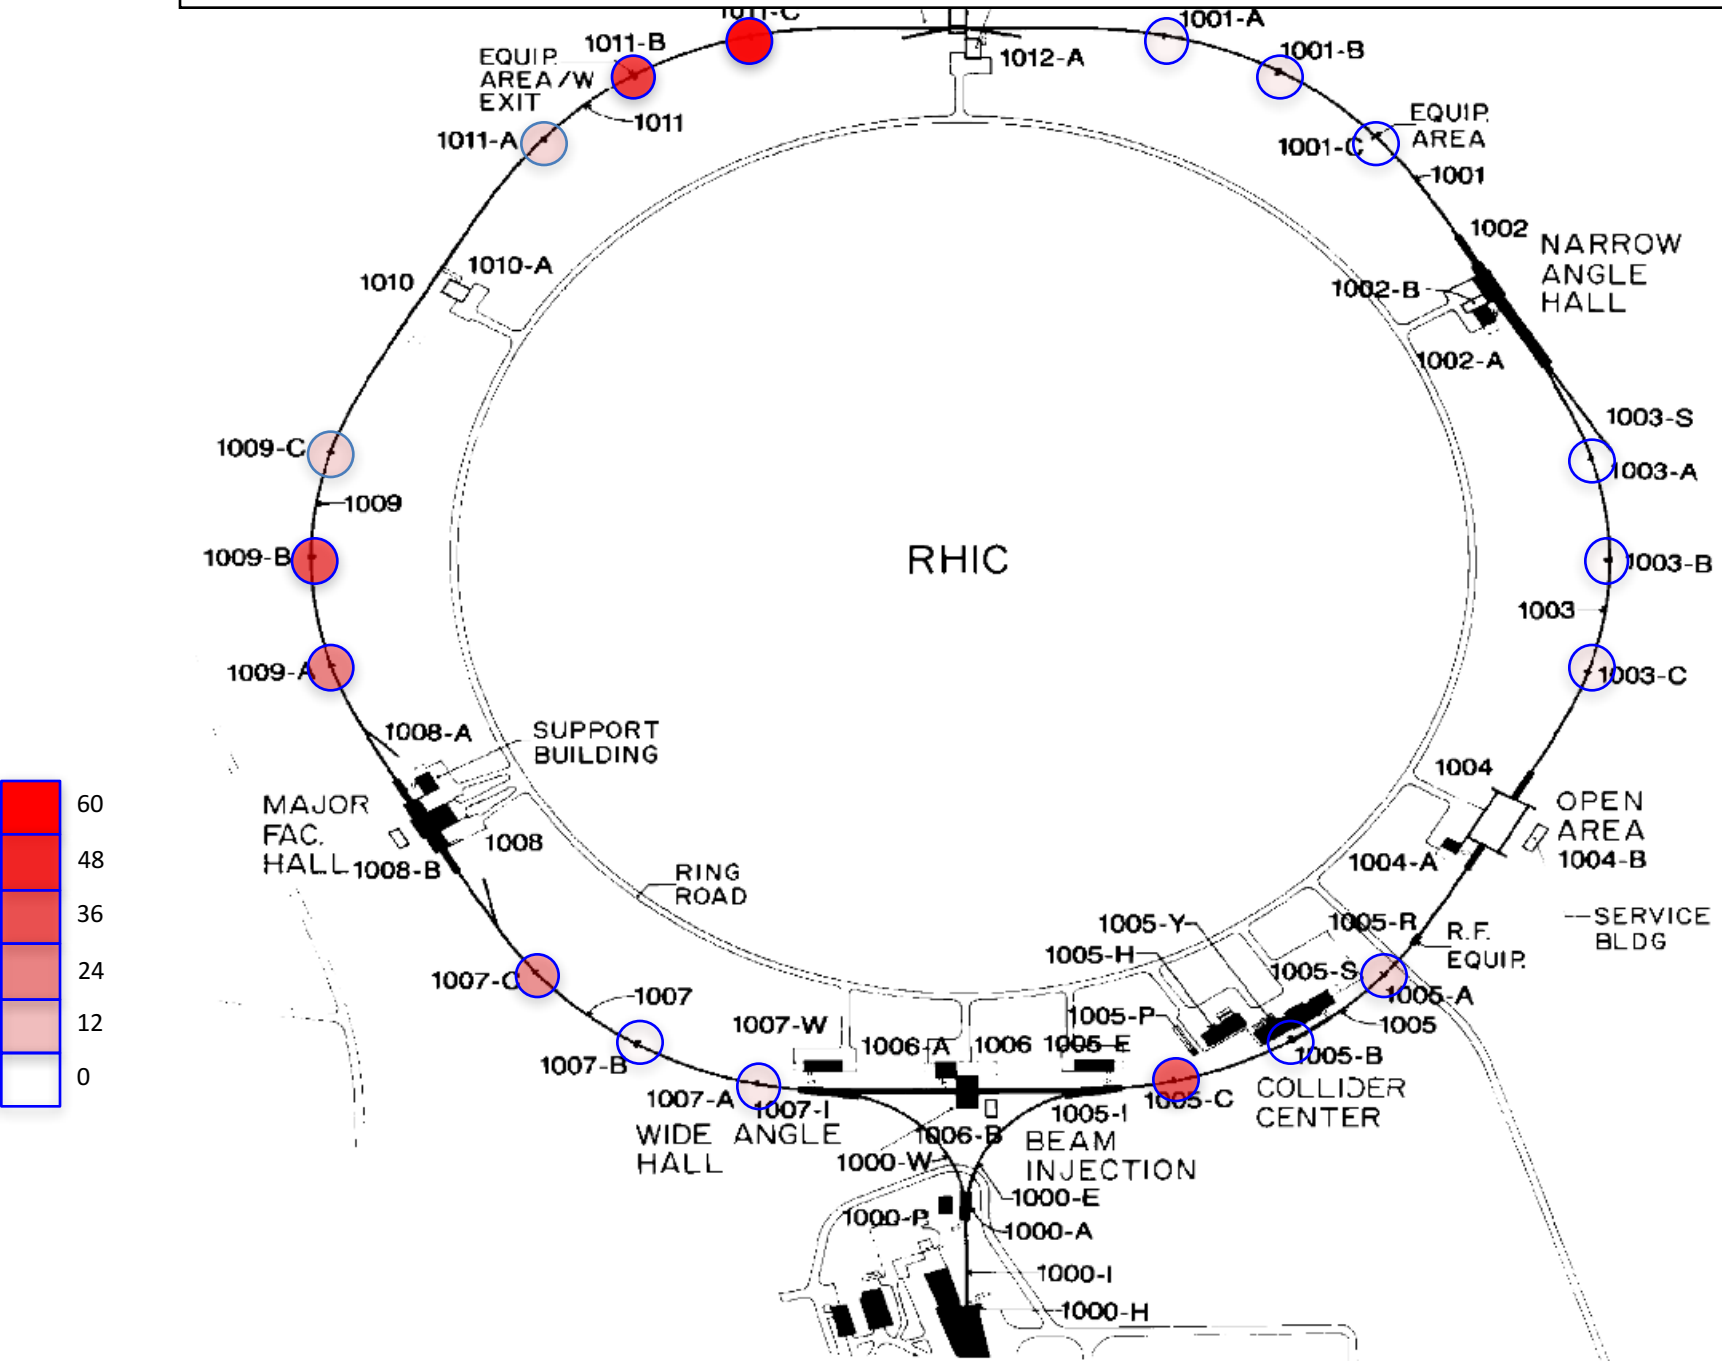

255 GeV p-p Alcove network switch radiation upset locations

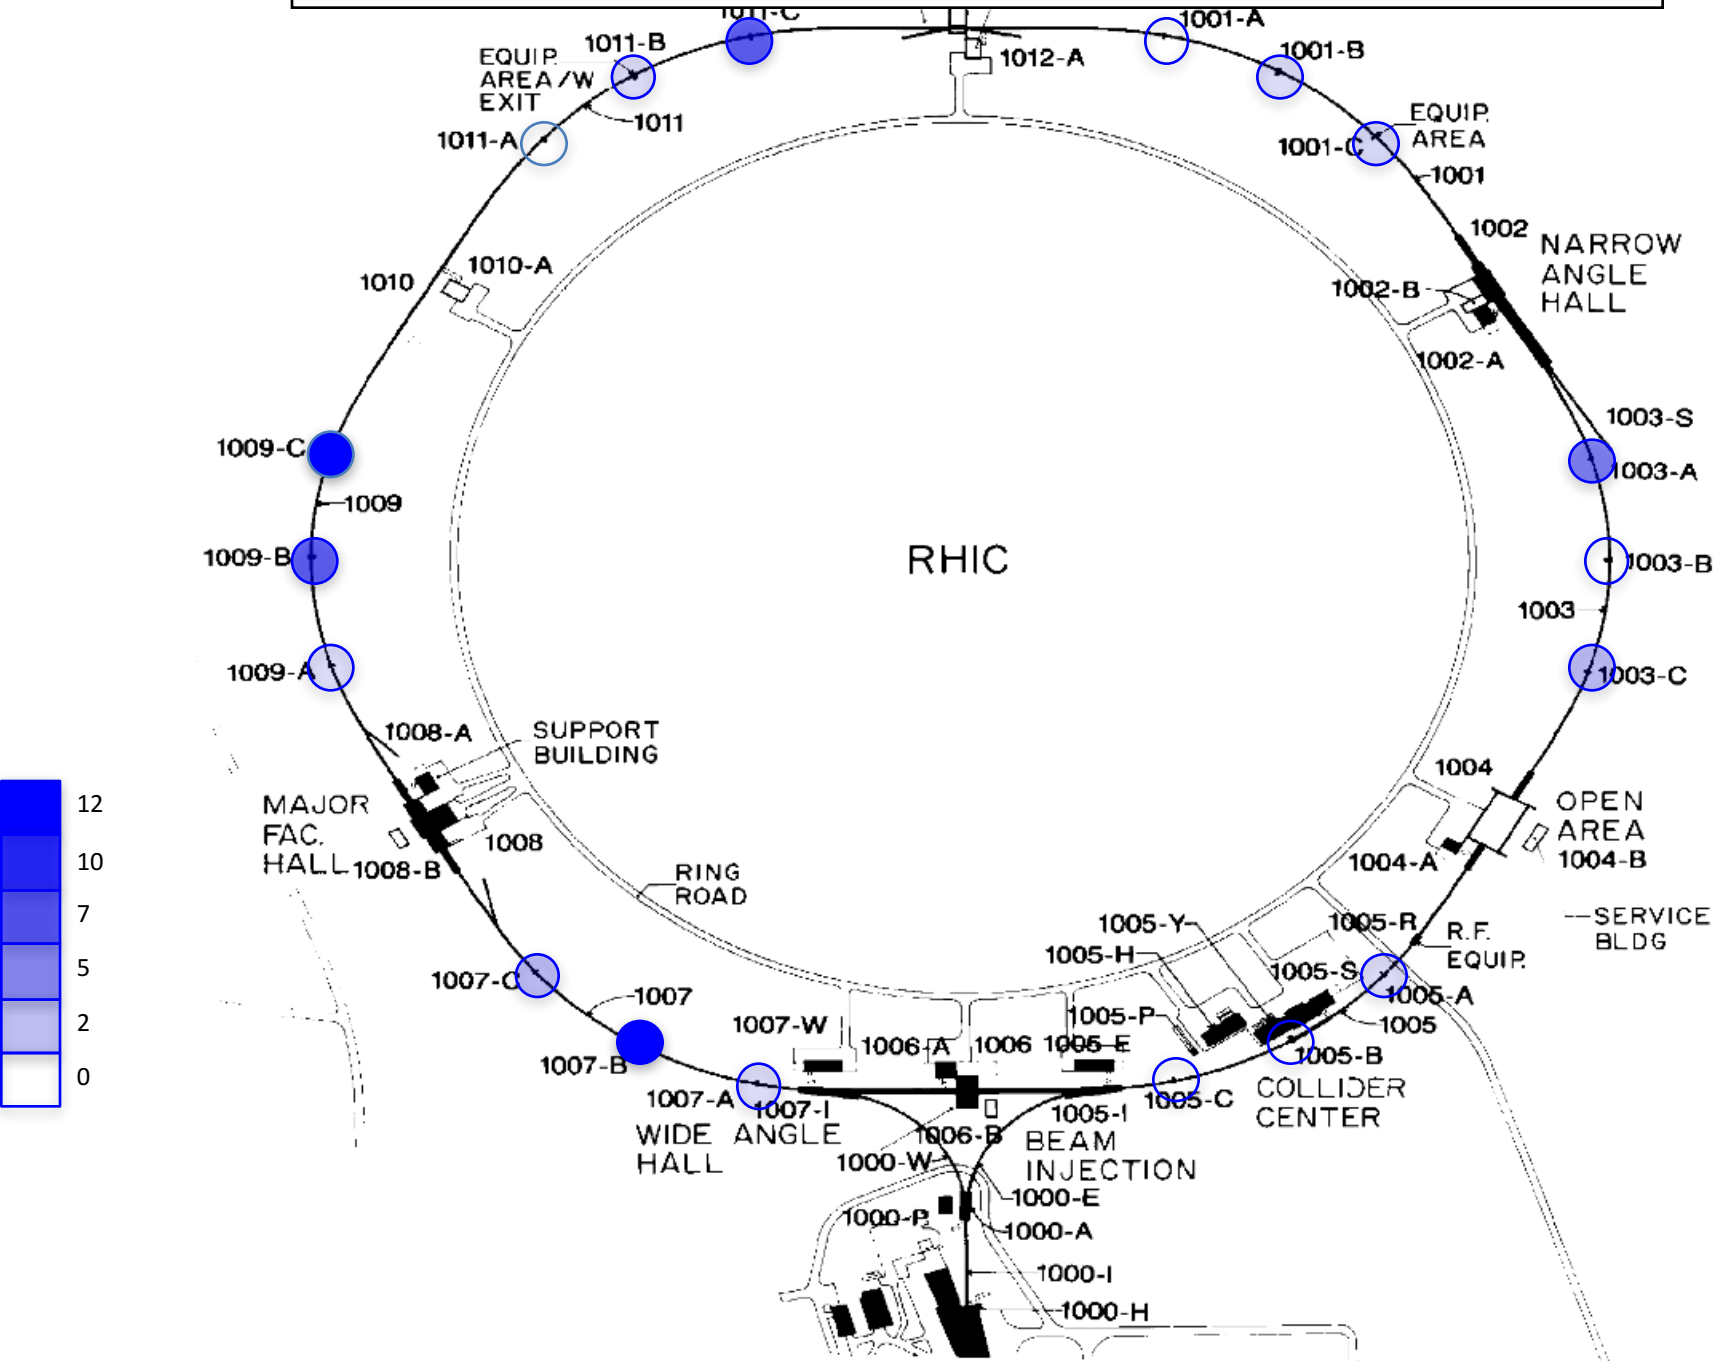

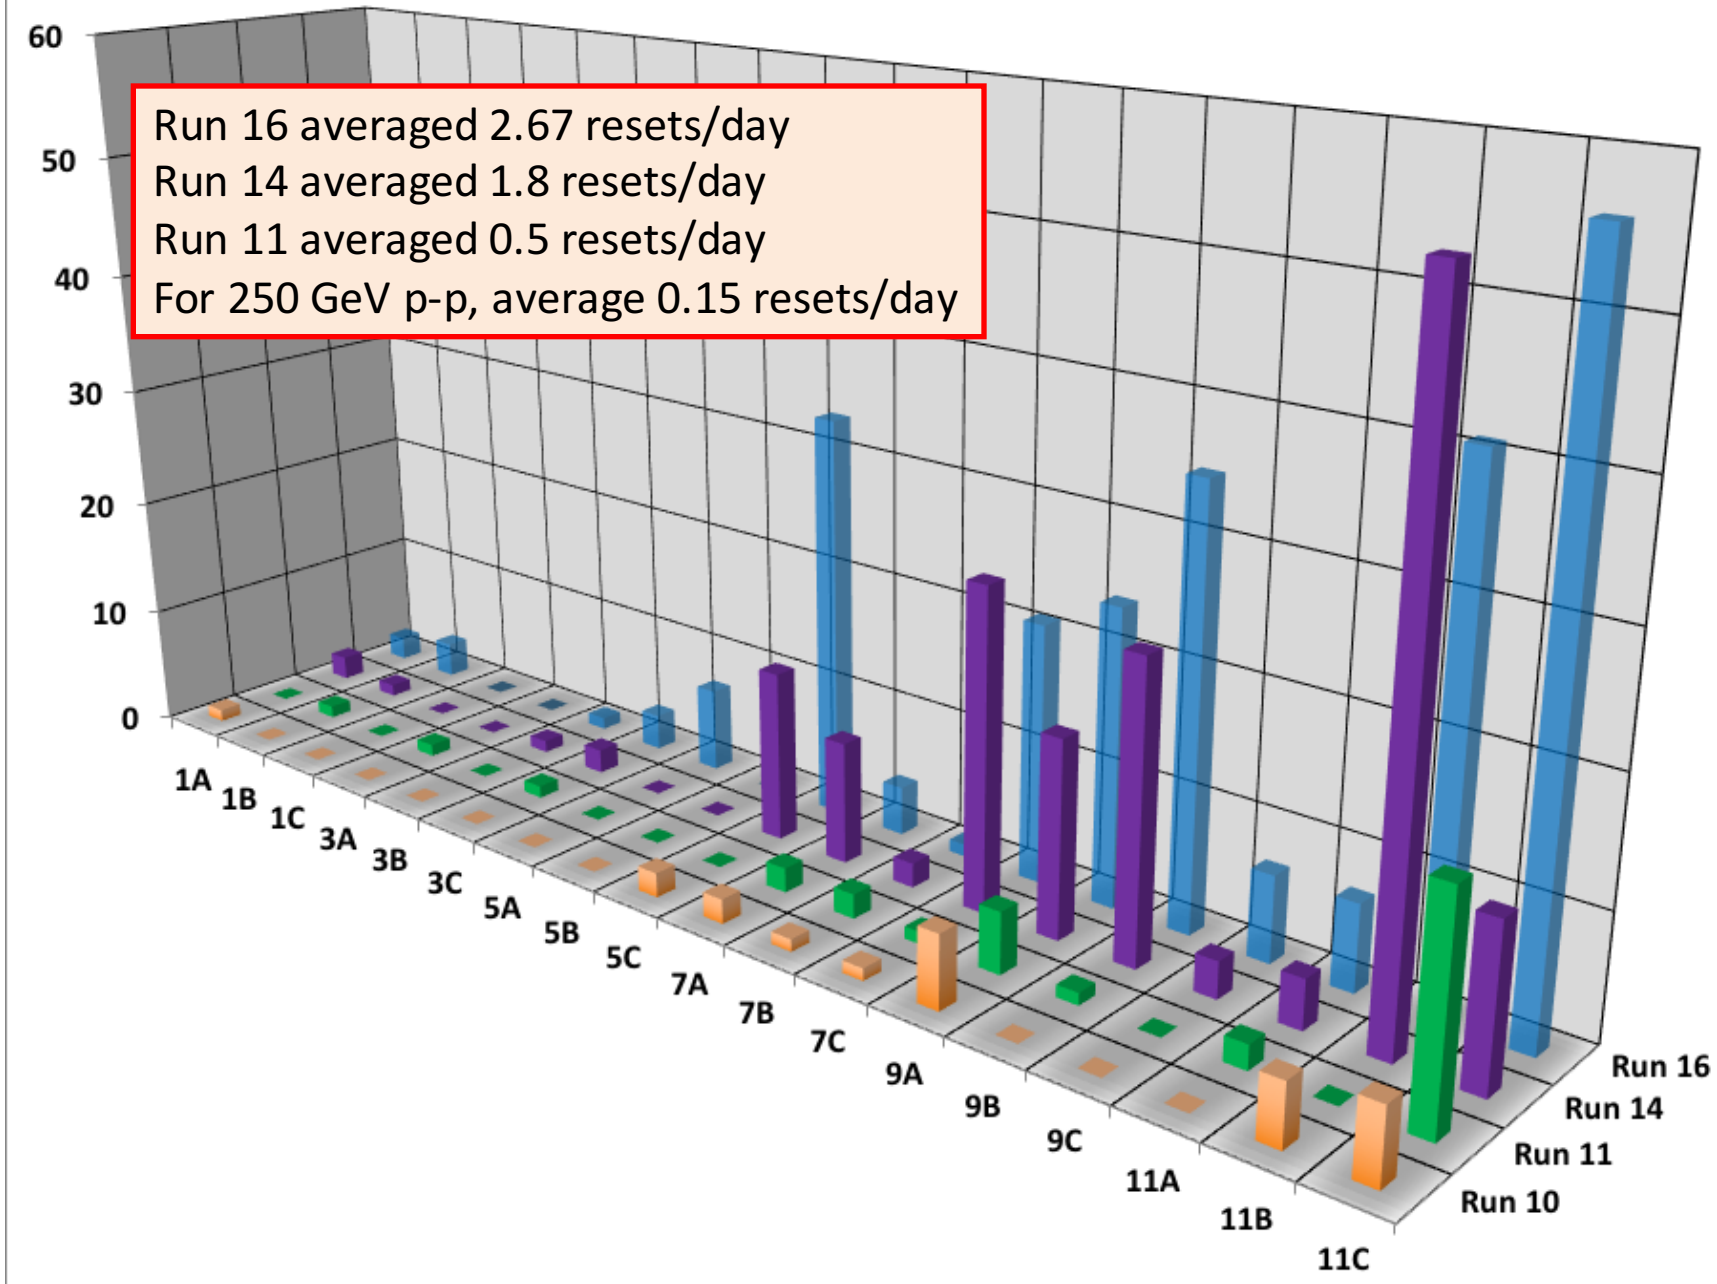
Network Switch Resets per Alcove for 100 GeV/n Au-Au runs
Run 16 averaged 2.67 resets/day
Run 14 averaged 1.8 resets/day
Run 11 averaged 0.5 resets/day
For 250 GeV p-p, average 0.15 resets/day
60
50
40
30
20
10
0
1A
1B
1C
3A
3B
3C
5A
5B
5C
7A
7B
7C
9A
9B
9C
11A
11B
11C
Run 16
Run 14
Run 11
Run 10

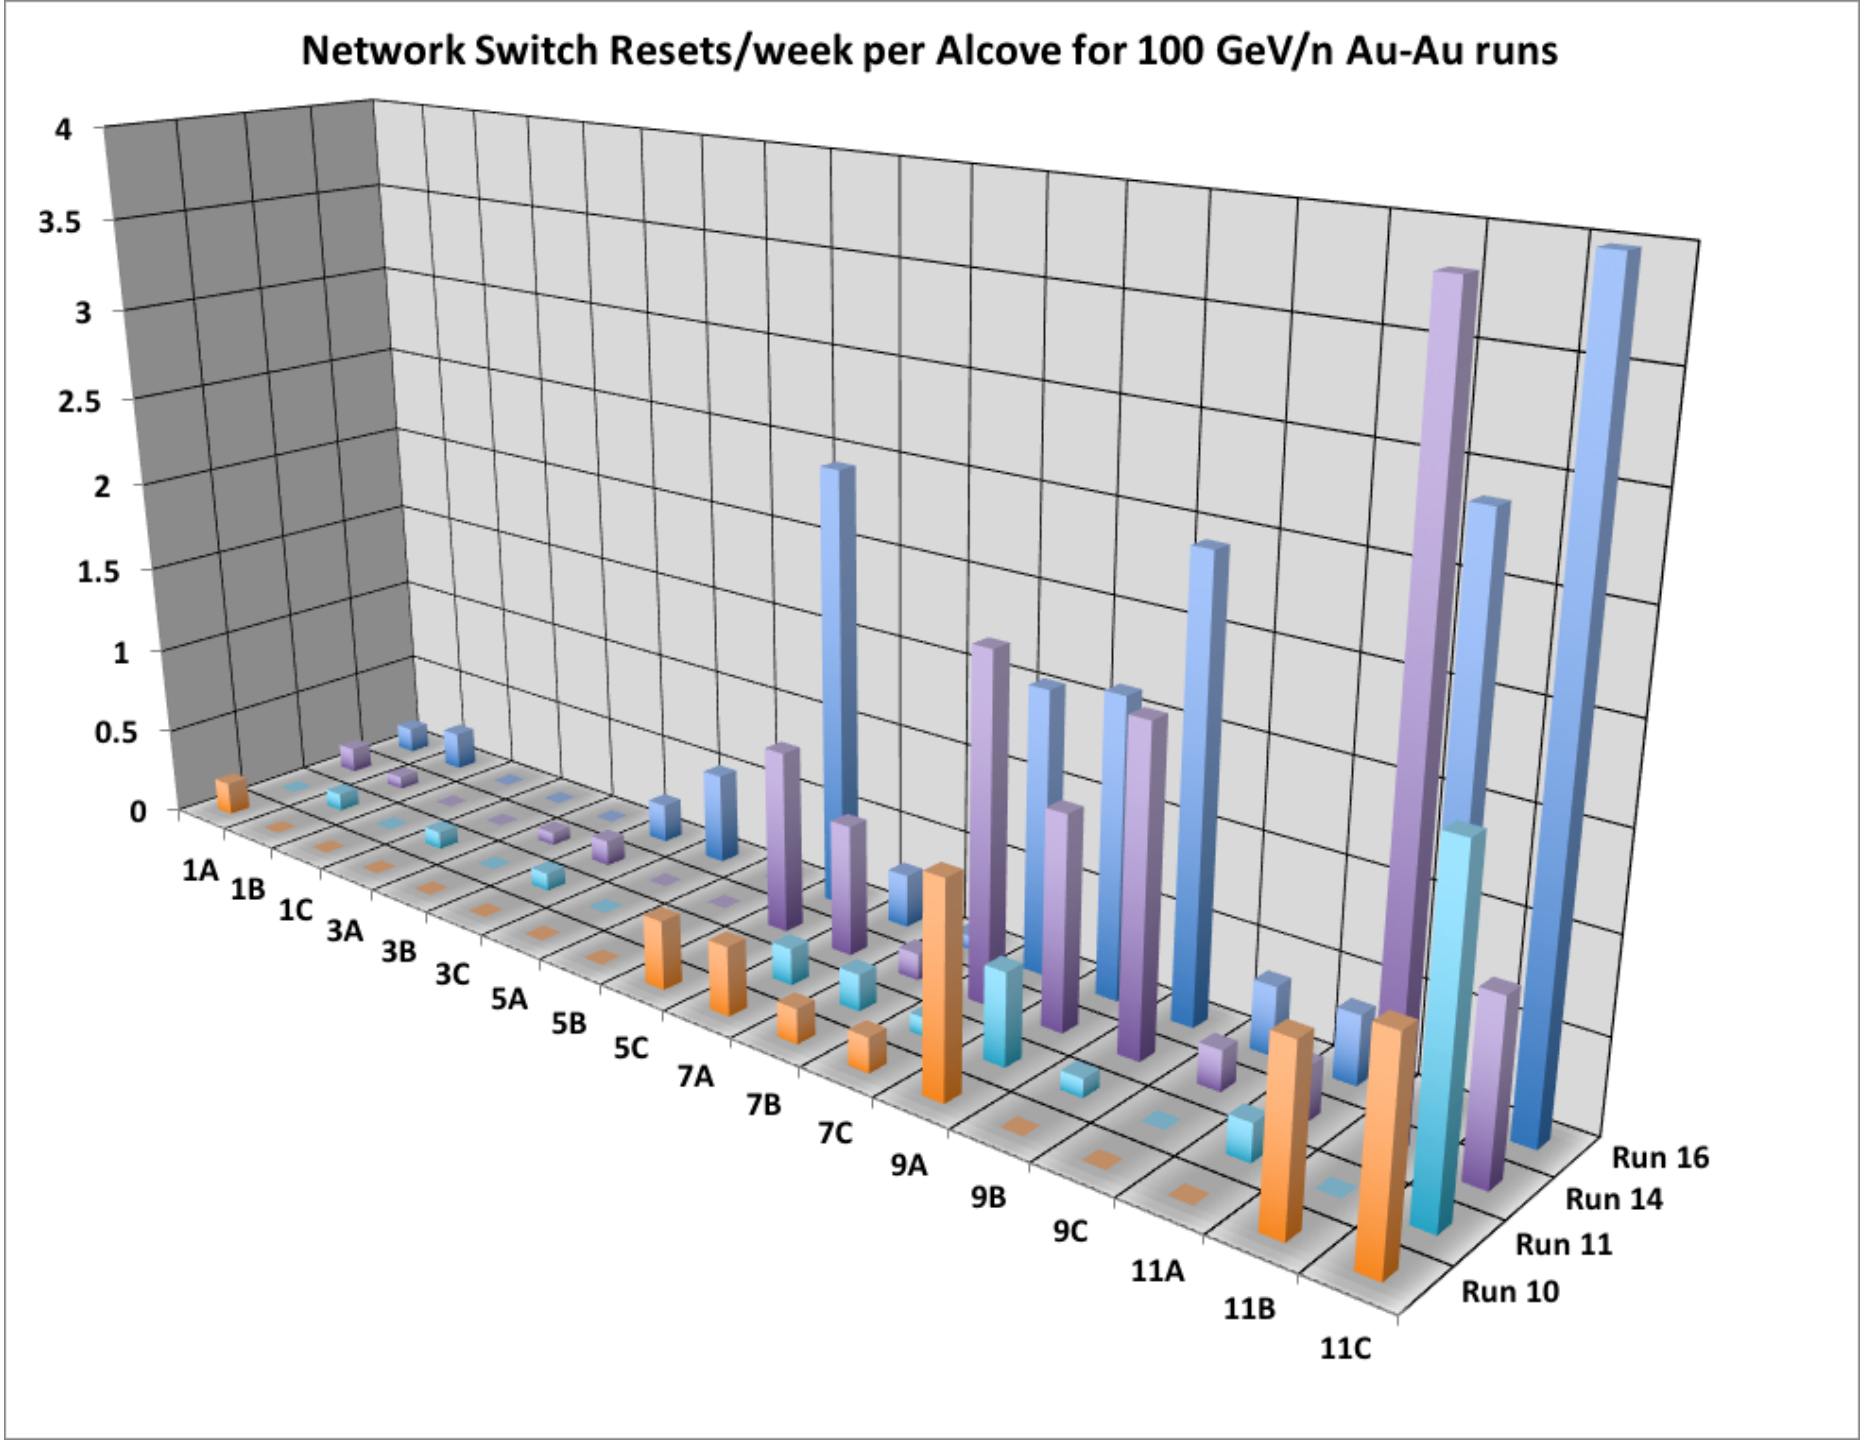
Network Switch Resets/week per Alcove for 100 GeV/n Au-Au runs
4
3.5
3
2.5
2
1.5
1
0.5
0
1A
1B
1C
3A
3B
3C
5A
5B
5C
7A
7B
7C
9A
9B
9C
11A
11B
11C
Run 16
Run 14
Run 11
Run 10

Network Switch Resets/week per Alcove for 100 GeV/n Au-Au Run 16

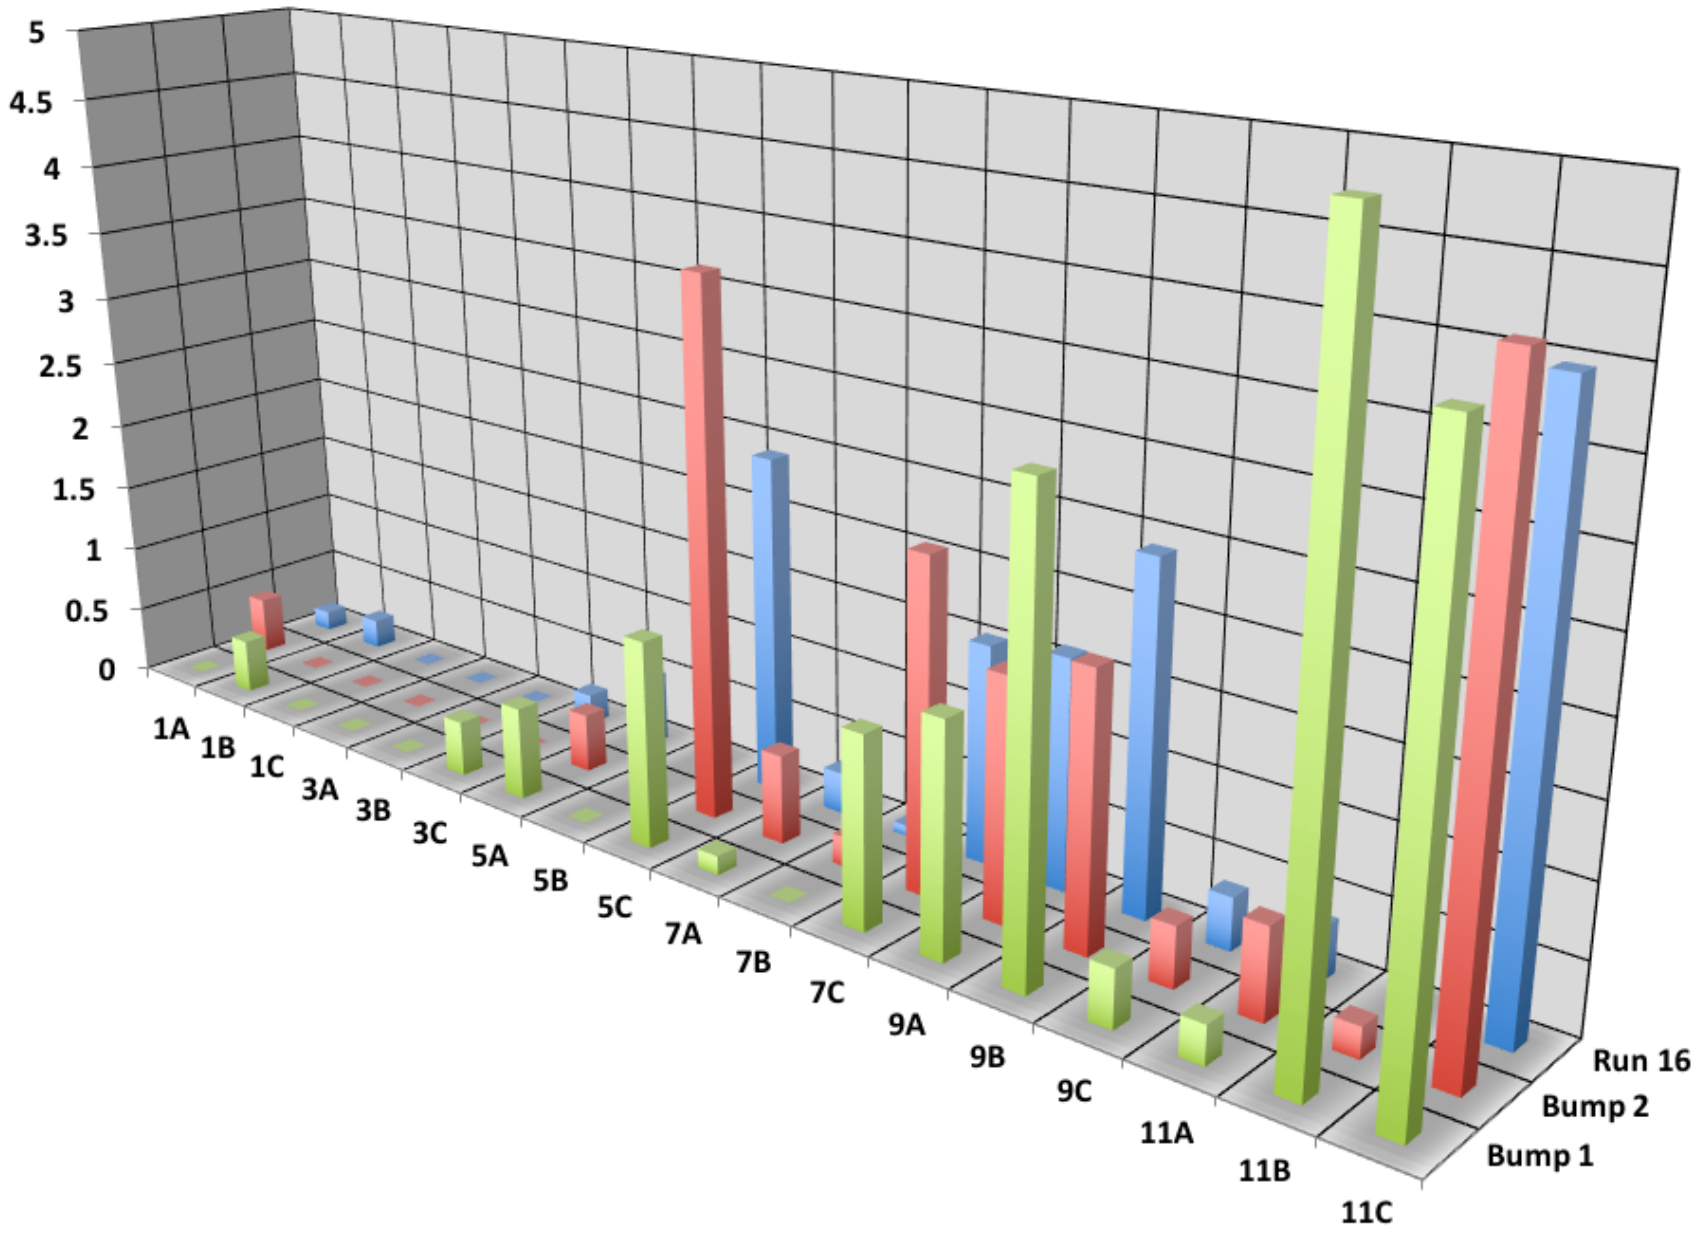

Average Number of Resets/week for different Runs

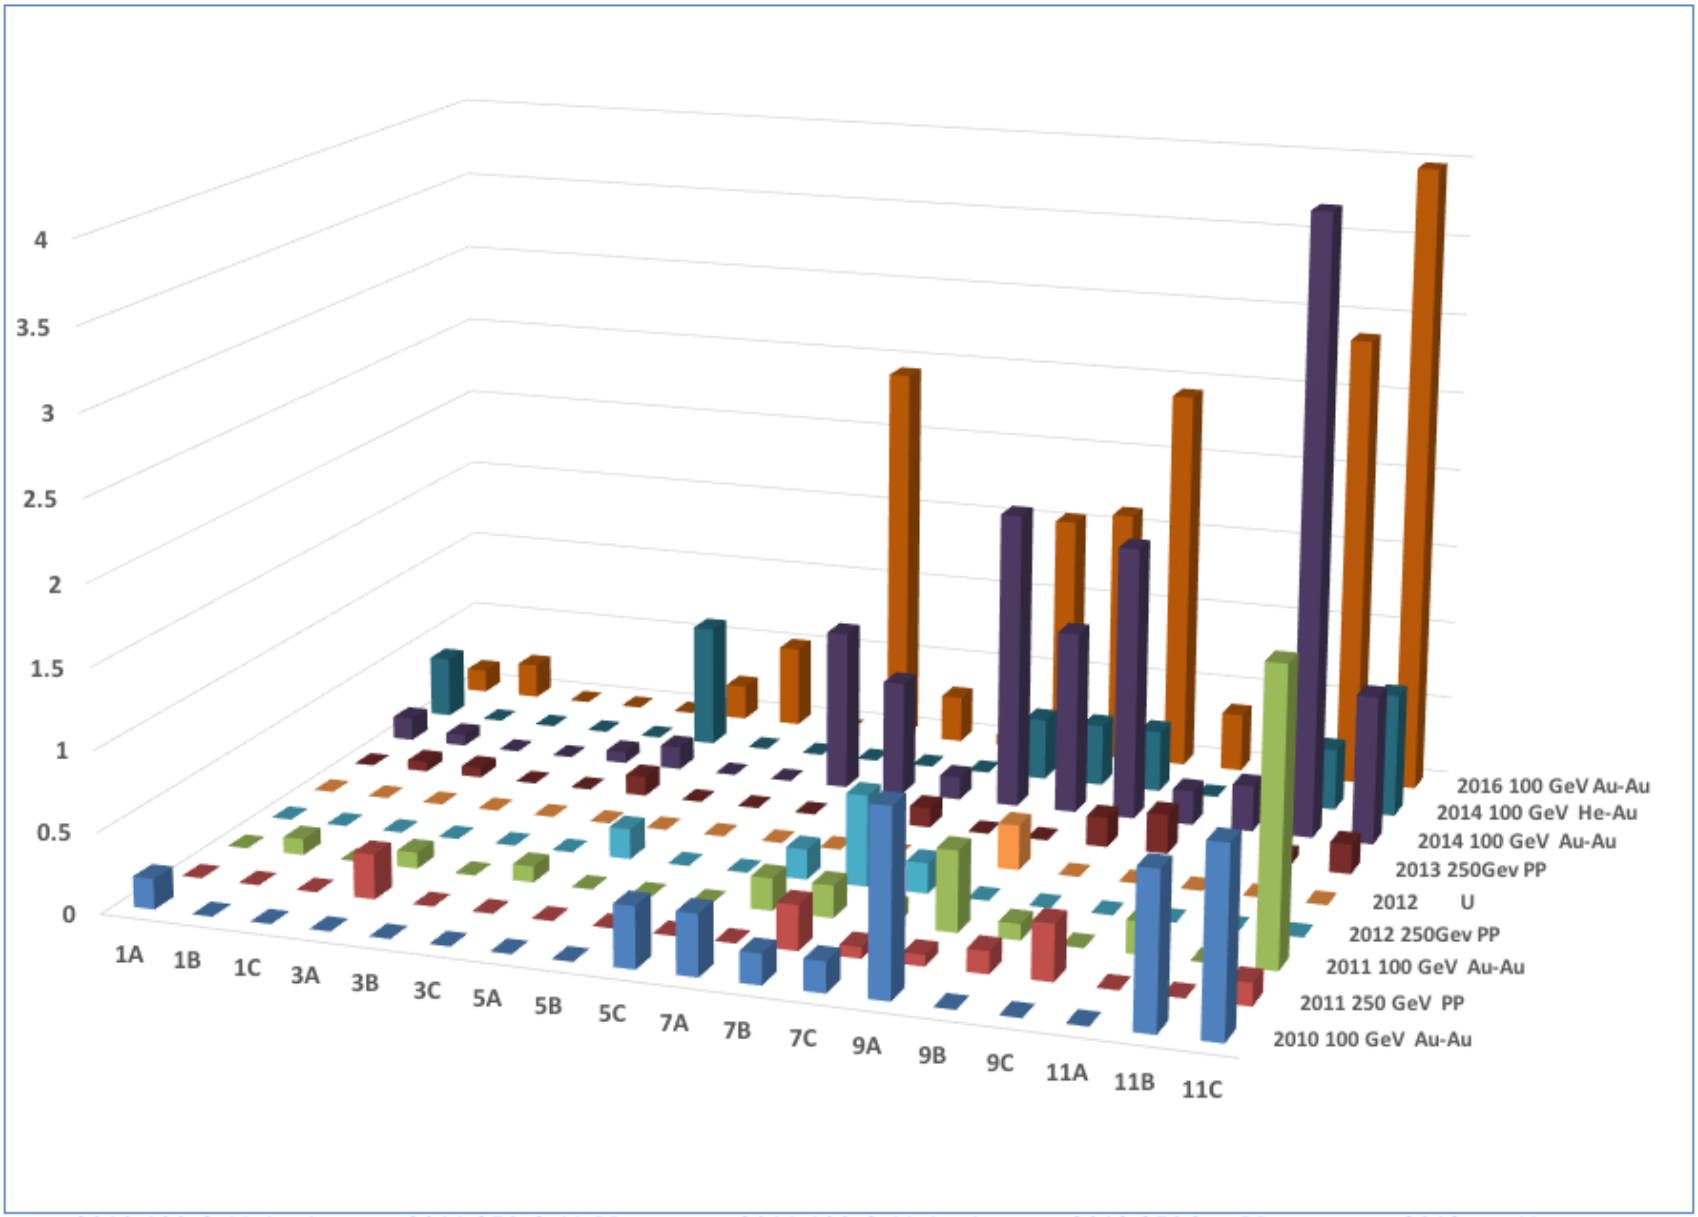

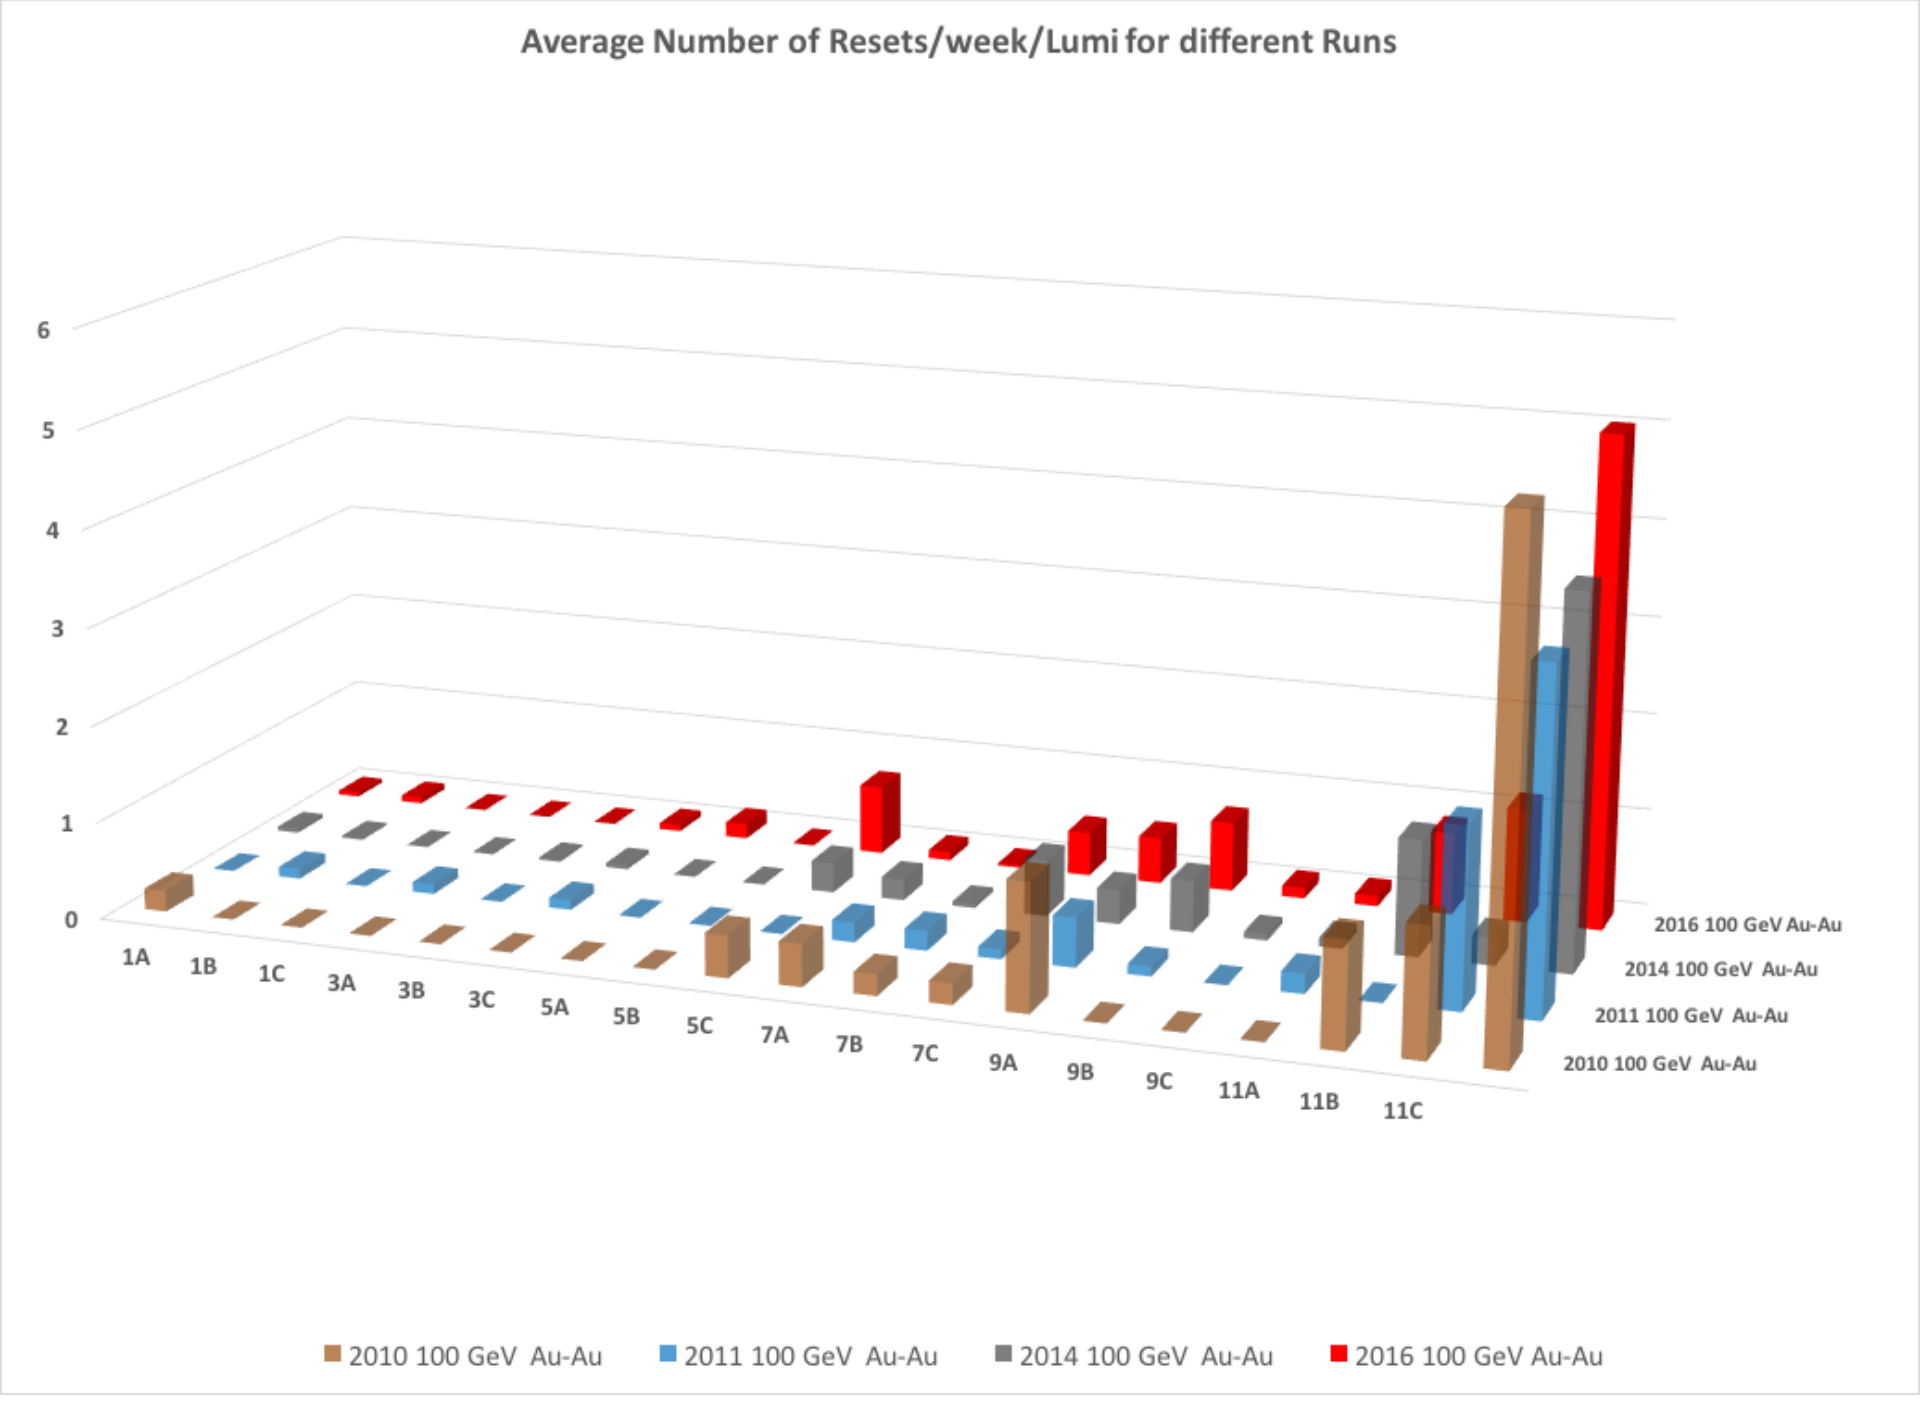
Average Number of Resets/week/Lumi for different Runs
6
5
4
3
2
1
0
1A
1B
1C
3A
3B
3C
5A
5B
5C
7A
7B
7C
9A
9B
9C
11A
11B
11C
2016 100 GeV Au-Au
2014 100 GeV Au-Au
2011 100 GeV Au-Au
2010 100 GeV Au-Au
2010 100 GeV Au-Au
2011 100 GeV Au-Au
2014 100 GeV Au-Au
2016 100 GeV Au-Au

AuAu2016_protBump1 Switch Resets per Week by RHIC Alcove

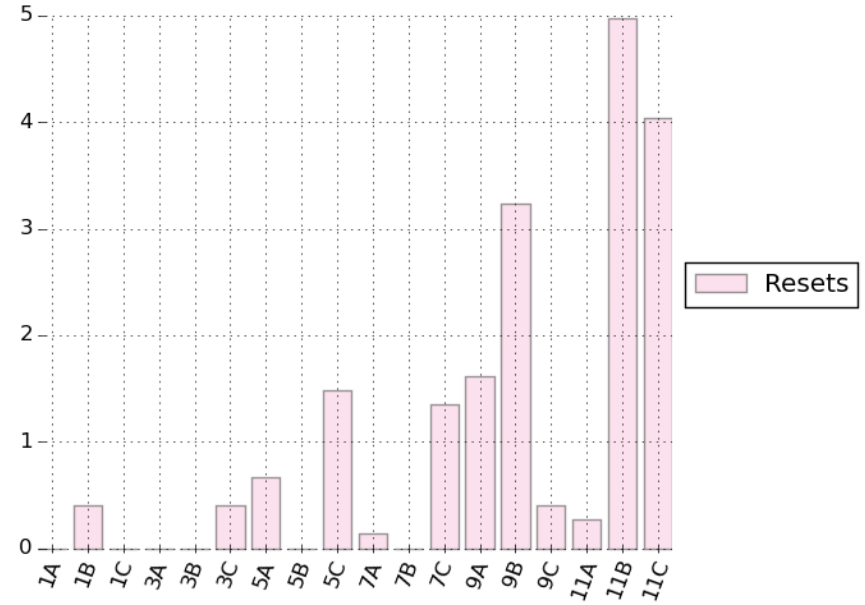

AuAu2016_protBump2 Switch Resets per Week by RHIC Alcove

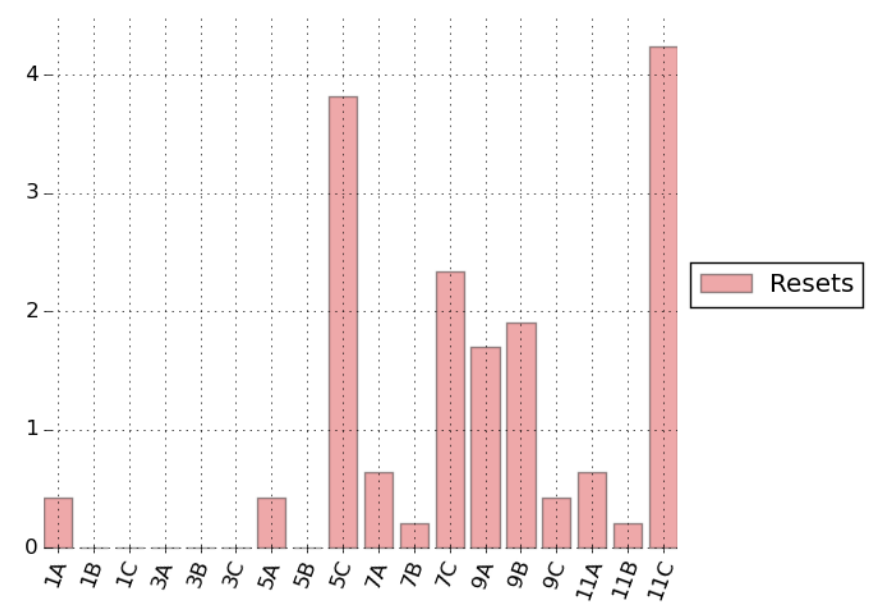

AuAu2016_protBump1 RHIC PS FEC Resets per Week by Location

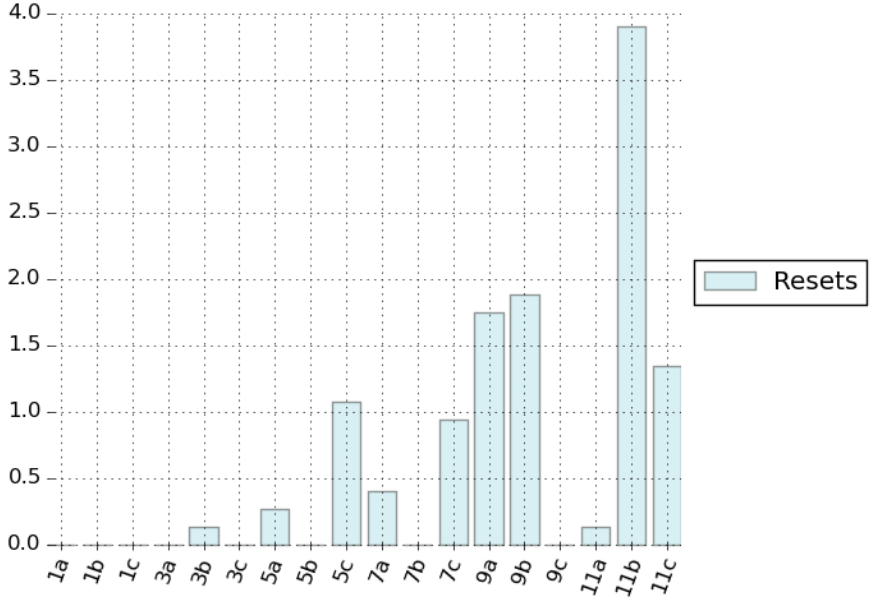

AuAu2016_protBump2 RHIC PS FEC Resets per Week by Location

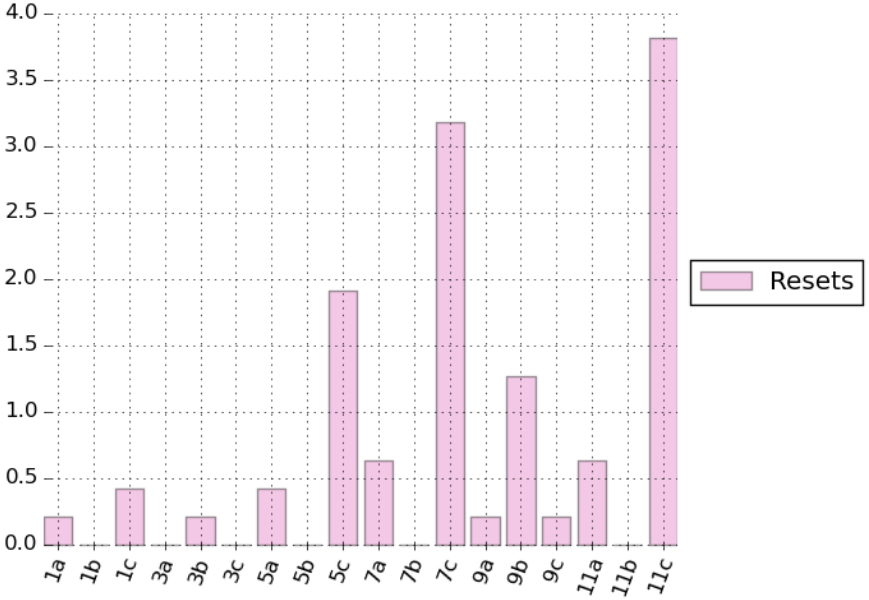

Memory Upset Measurements

- Logging data this year, starting in early Feb.
- Using Ram Disks in Alcove PS FECs
- Place 64 bit pattern repeatedly into 2 Mbytes of memory
- Creates 262,144 samples per FEC
- Patterns are read on down-ramp of every store
- Almost always, changes are at the single bit level [that is: the number of bit changes is small compared to the sample set, so the probability of a single bit changing twice in a sample period is insignificantly small.]

Correlation between Resets and Memory Upset Statistics

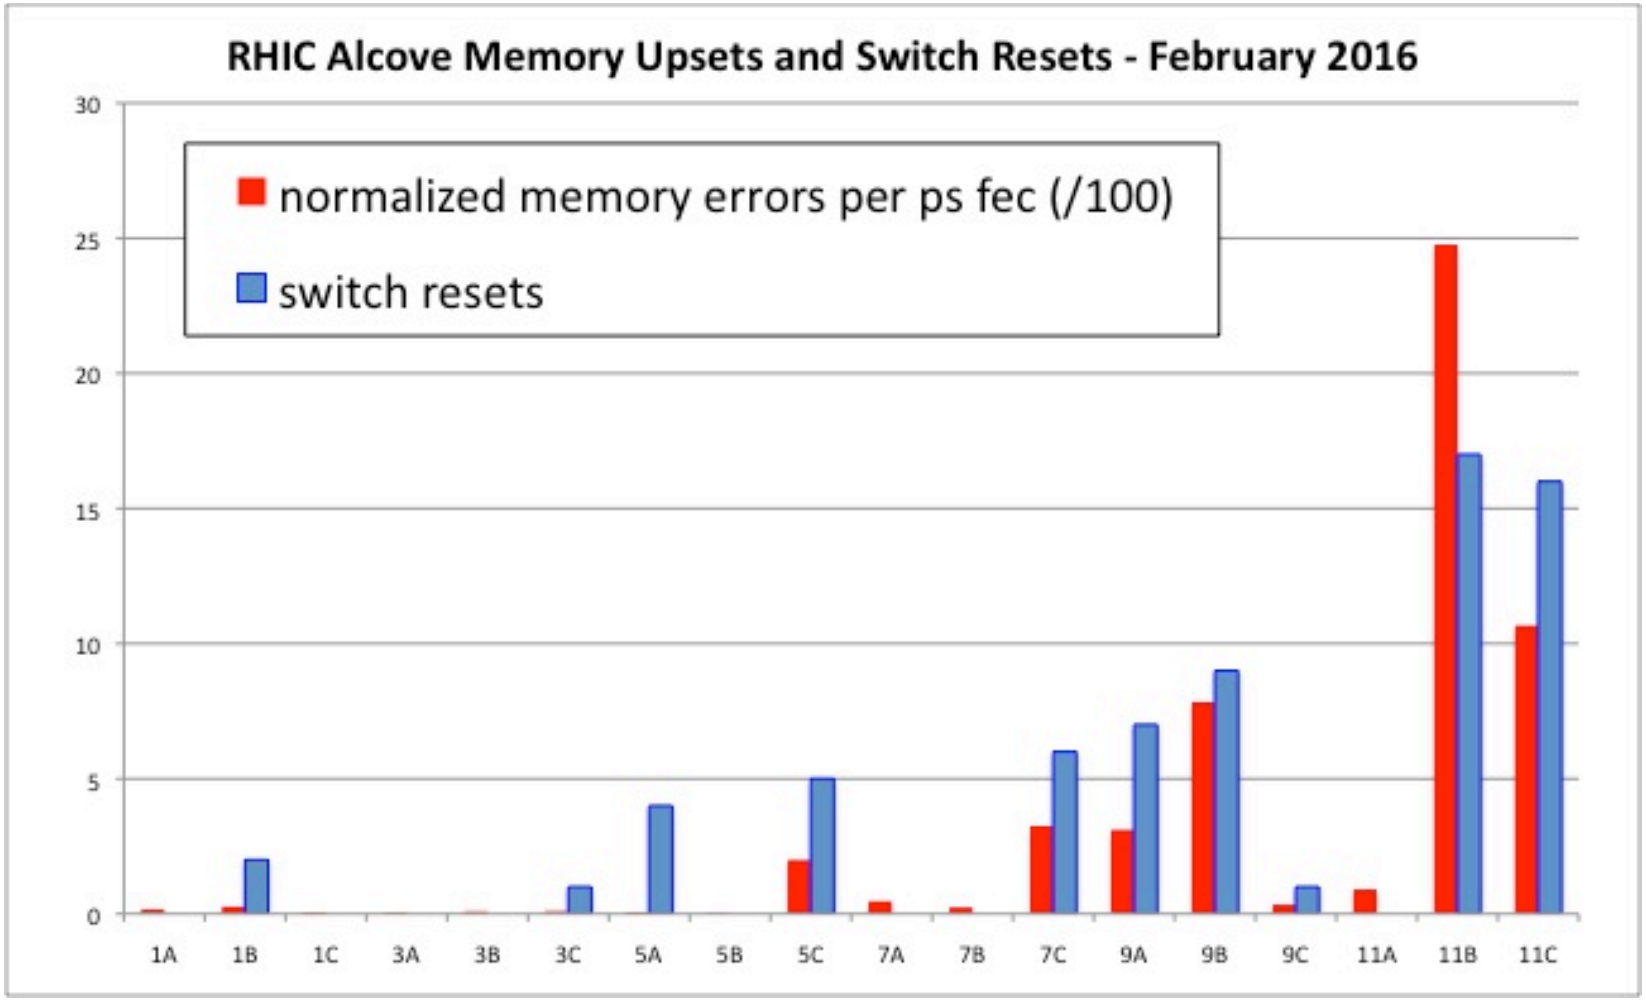

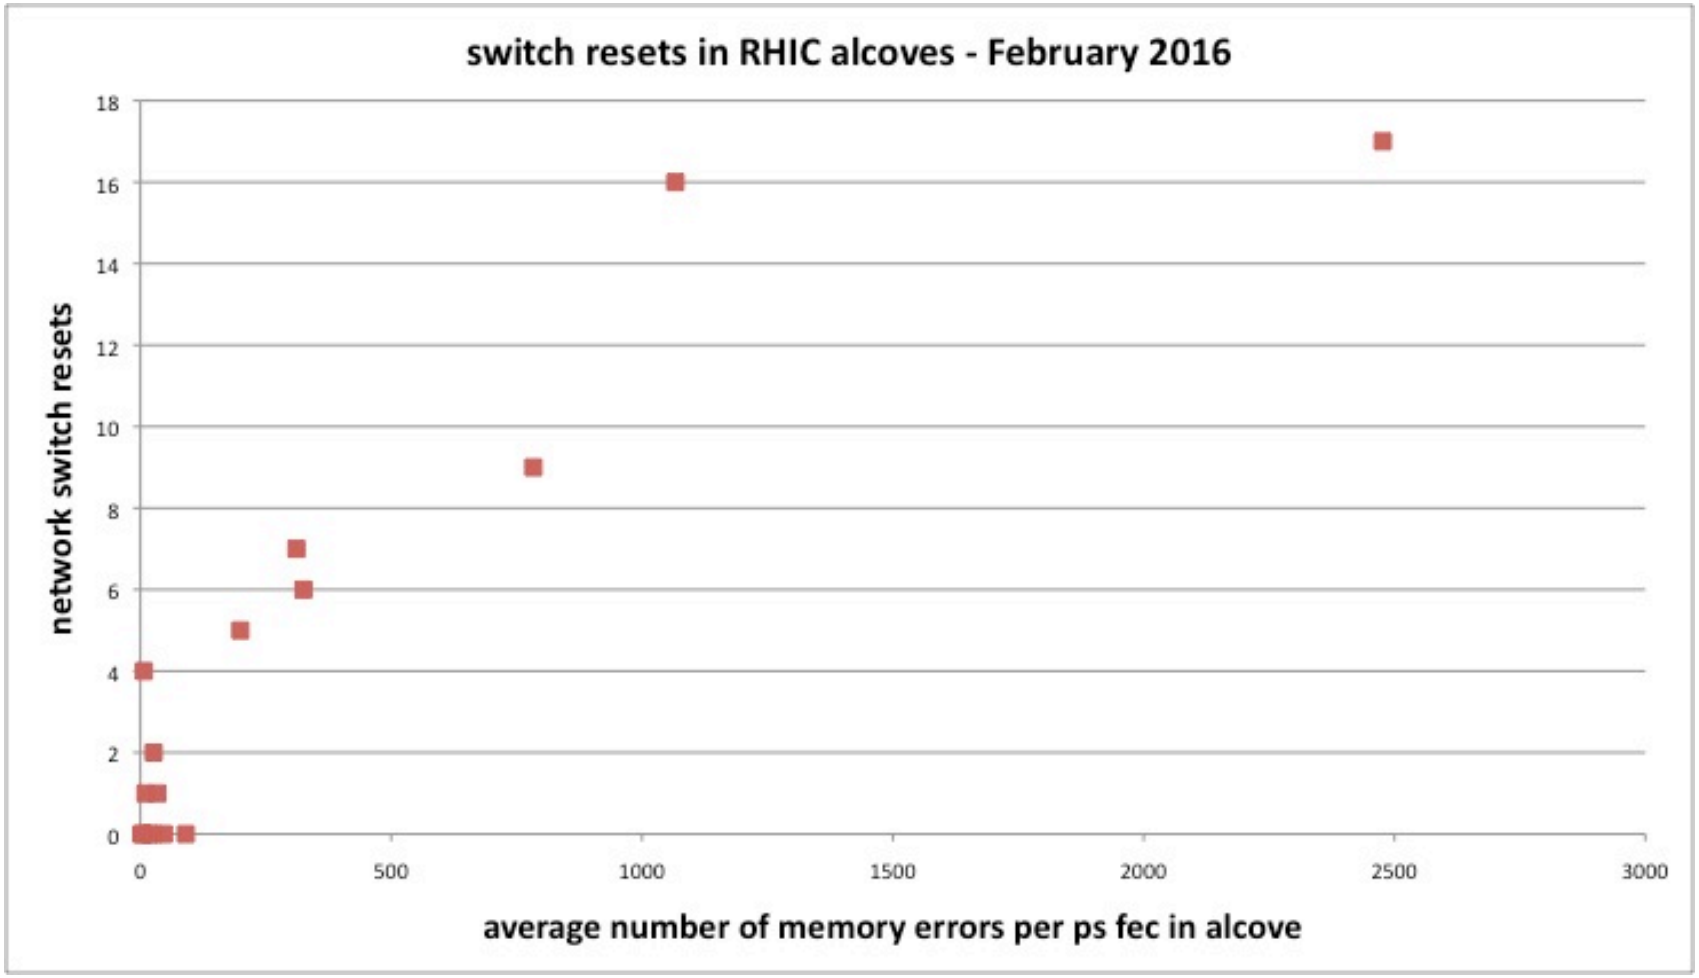
switch resets in RHIC alcoves - February 2016
network switch resets
18
16
14
12
10
8
6
4
2
0
0
500
1000
1500
2000
2500
3000
average number of memory errors per ps fec in alcove

Mitigation Ideas of the Past

- Network switches that use ECC memory
  - Costs are high, >$30k/switch
  - Memory corruption is not the main problem
  - Address corruption hangs the switch – addresses not in ECC
- Shielding Alcoves?
  - Kin's simulations (2014) suggest reduction factors are small (reduction at most ½ with lot's of shielding – not practical)
- Could move sensitive equipment to "quiet" alcoves
  - Some new equipment required & labor intensive
  - Space is tight & distances long
- Could move some network switches out of alcoves to service buildings
  - Distances are large (~1400 ft)
  - Technology issues need investigation
    - Use copper cable instead of fiber?
    - Need to use converters, which may be susceptible and may pose operational challenges.

Improvements since 2014

- Improved on the data collection to better monitor the "health" of the equipment in the alcoves.
  - Automate & Log data from memory pattern scans
  - Monitoring network statistics
- Addition of the protection masks and moving the Experiment protection bumps led to noticeable changes, although Run 16 was still worse than Run 14 ☹
- Removed UPS units from some alcoves (during run).
- Automated (Tape) mass rebooting FECs on down-ramp of every store.
  - helps reduce impact of madc problems
  - No clear evidence it reduces impact of rad upsets in FECs

Future Possibilities

- Small form factor network switches placed in shielded boxes or externally (commonly used in military)
  - Can come NEMA rated
  - Designed for harsh environments
  - May be difficult to find good compatibility with existing hardware.
- External NEMA rated enclosures for network switches & possibly other equipment [cost is roughly $2k per enclosure + installation]
  - Cable runs ~200 ft, could be done with copper (avoids fiber transceivers in the tunnel)
  - No tunnel access to service
- Shield just the network switch in the alcove
- Improve on reset strategies

NEMA† Ratings for External Enclosures

[From NEMA 250-2003]

Comparison of Specific Applications of Enclosures for Outdoor Nonhazardous Locations

Provides a Degree of Protection Against the Following Conditions	Type of Enclosure									
	3	3X	3R*	3RX*	3S	3SX	4	4X	6	6P
Access to hazardous parts	X	X	X	X	X	X	X	X	X	X
Ingress of water (Rain, snow, and sleet **)	X	X	X	X	X	X	X	X	X	X
Sleet ***	...	...	...	...	X	X	...	...	...	...
Ingress of solid foreign objects (Windblown dust, lint, fibers, and flyings)	X	X	...	...	X	X	X	X	X	X
Ingress of water (Hosedown)	...	...	...	...	...	...	X	X	X	X
Corrosive agents	...	X	...	X	...	X	...	X	...	X
Ingress of water (Occasional temporary submersion)	...	...	...	...	...	...	...	...	X	X
Ingress of water (Occasional prolonged submersion)	...	...	...	...	...	...	...	...	...	X

* These enclosures may be ventilated.
** External operating mechanisms are not required to be operable when the enclosure is ice covered.
*** External operating mechanisms are operable when the enclosure is ice covered.

† National Electrical Manufacturer Association (NEMA)

Shielding Small devices?

- Kin Yip is simulating how effectively a small device can be shielded inside an Alcove
- If the results are encouraging, a test will be done in Run 17. [use Alcoves 7b or 9c for proton operations]
- Where in alcove?

Assume Alcove 7B

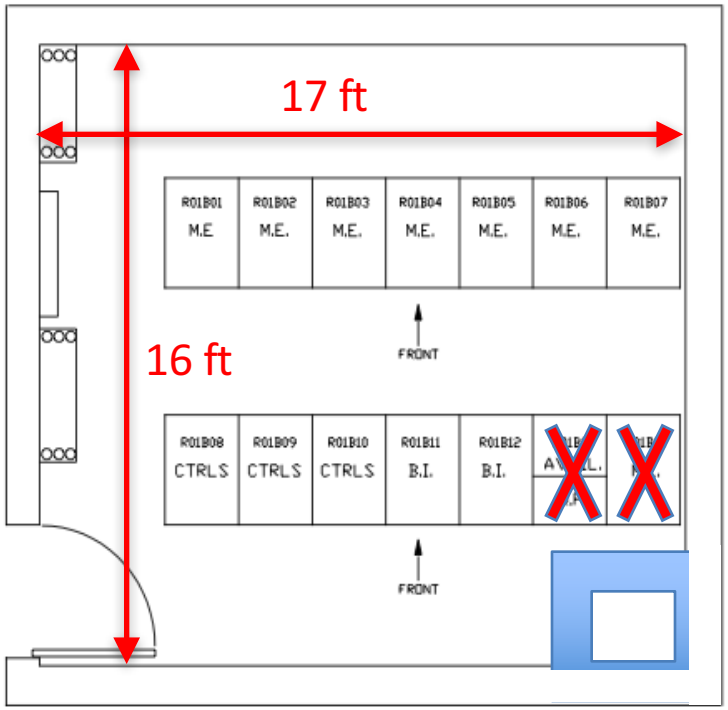

Summary

- Single Event Upsets cause equipment in Alcoves to hang and fail – Run 16 had the highest number of upsets ever.
- Mitigations taken to reduce upsets over the years have kept the impact to a manageable level, although vulnerabilities remain a concern
- Most significant vulnerability is in network switch upsets
- No easy solutions, but there are still things that can be done
  - External enclosures
  - E.g., Improve reaction using memory pattern data

Backup/Auxiliary Slides

Some Terminology

- ECC memory: Error-correcting code memory
  - Detects and corrects most common kinds of internal data corruption.
  - Is immune to single-bit errors
  - Makes use of additional circuitry that checks accuracy of data during I/O operations.
  - E.g., data read from a word of memory will remain uncorrupted even if a single bit in that memory was flipped.
- MRAM: Magneto-resistive random-access memory
  - Non-volatile memory
  - Data is not stored as electrical charge, but by magnetic storage elements (via two ferromagnetic plates), so works on spin, not charge.
- Single Event Upsets, Single Event Latchups, Single Event Gate Rupture, & Single Event Burnouts
  - radiation striking a sensitive node in a micro-electronic device
  - Most common are upsets, which are not permanent damage
  - Result is at least a single bit change

History – 2004

- Front End Processors
  - Power 3E processors lacked ECC memory
  - Some locations had CPUs with ECC memory
    - 1007C, 1009A, 1009C
  - Began process of replacing select CPUs with MVME2112 processors, that have ECC memory
- VME Chassis (250 in field, 40 in Alcoves)
  - 13 VME PS Failures in Alcoves (38%)
  - 3 VME PS Failures in Service buildings (<1%)
- Collaboration: VME Chassis Rad Tested for LHC at CERN (Wiener Chassis now in use)
- Began looking at Rad Resistant PSs = found one that performed well (Vicor PS now in use)

V115 (WFGs)

- Tried ECC memory in WFGs (in 2004)
  - To replace 128K x 32 static ram
  - Plugged into existing memory footprint
  - Initially performed well in tests
  - Didn't perform as well as expected in alcoves
- After 2004 switched to MRAM
- Finished MRAM installation by 2011
- In addition
  - Alcove FECs (for PSs) were upgraded to use ECC memory
  - Stopped using RAM disks to save files
  - Code was modified to save 3 copies of data that had to remain in RAM disks (to avoid corruption)
  - RAM disks for FECs that did not need them were removed (i.e., BPMs and QD FECs)