SVD Readout Electronics

SVD Internal Review

at KEK

Daniel Marlow
Princeton University

February 2, 2001

Topics

- SVD Readout Requirements and Goals
- Readout Architecture: with emphasis on the changes resulting from the inner tracker task force.
- Radiation tolerant deep-submicron CMOS: (SEU and SELU mitigation)
- Design and Status of various subsystems
  - VA1 chips
  - Hybrid
  - Repeater Docks
  - Trigger

Readout Design Goals

- Retain excellent noise performance of VA1
- Radiation Hard to > 2 Mrad
- Shaping Time of ~500 ns
- Current compensation for DC coupling (or allowance for leaky capacitors)
- Compatibility with SVD1 readout architecture (this is becoming less of an issue).

The VA1 Chip

- 128 channels
- Descendent of Viking (O. Toker et al., NIM A340 (1994) 572.)
- AMS 1.2 um CMOS
- Noise:

$$\mathrm{ENC} = 165\,e^- + 6.1\,e^-/\mathrm{pF}$$

$$@\ 2\ \mu\mathrm{s\ shaping\ time}$$

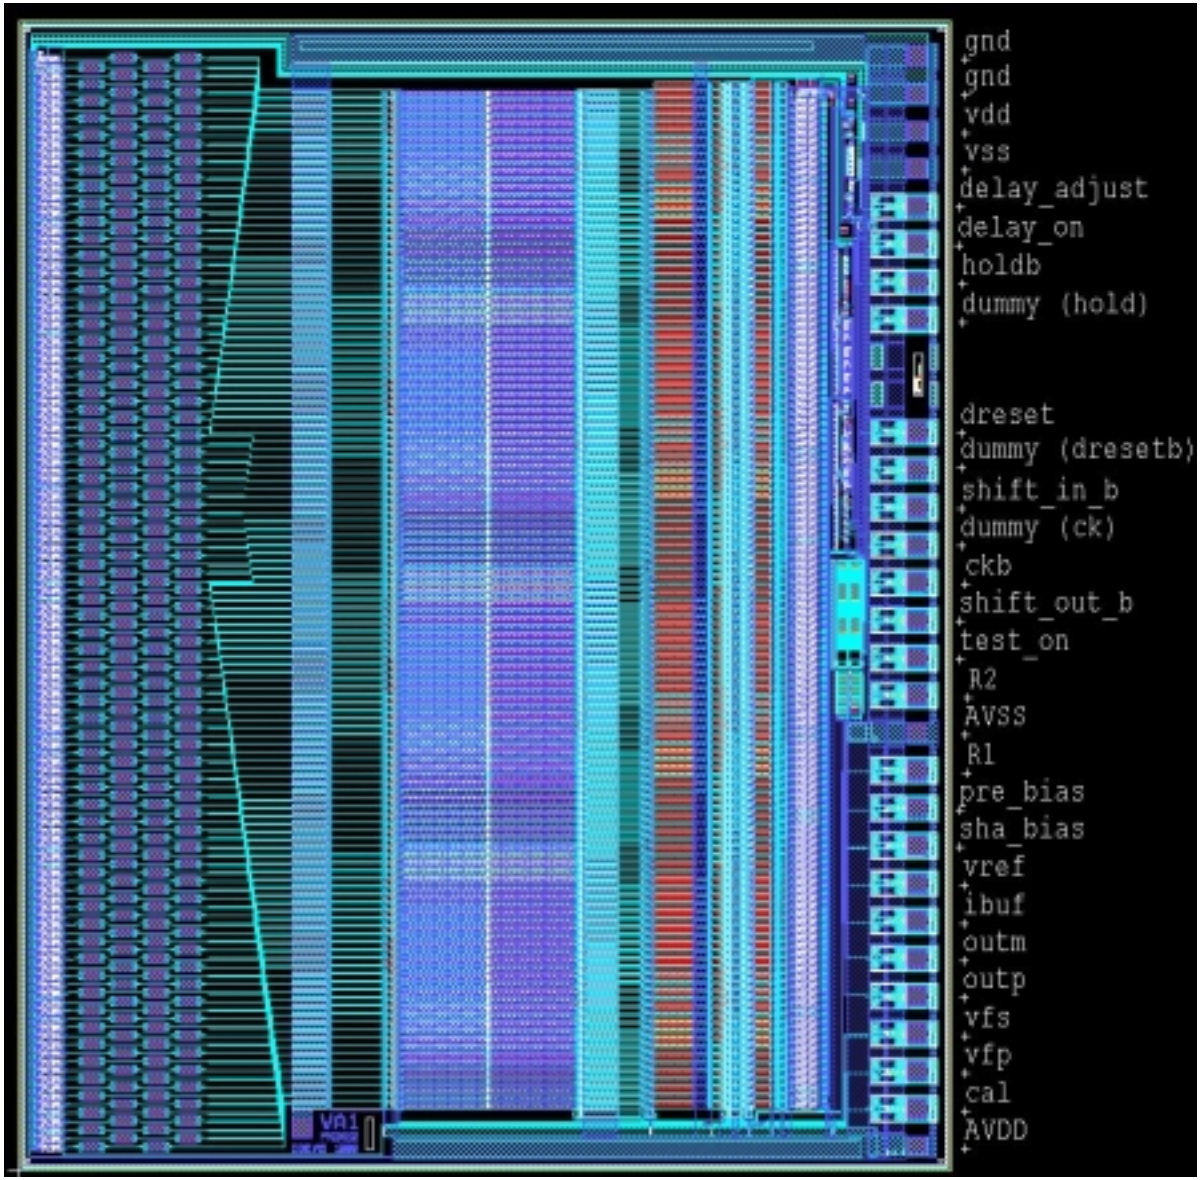

VA1 Principle of Operation

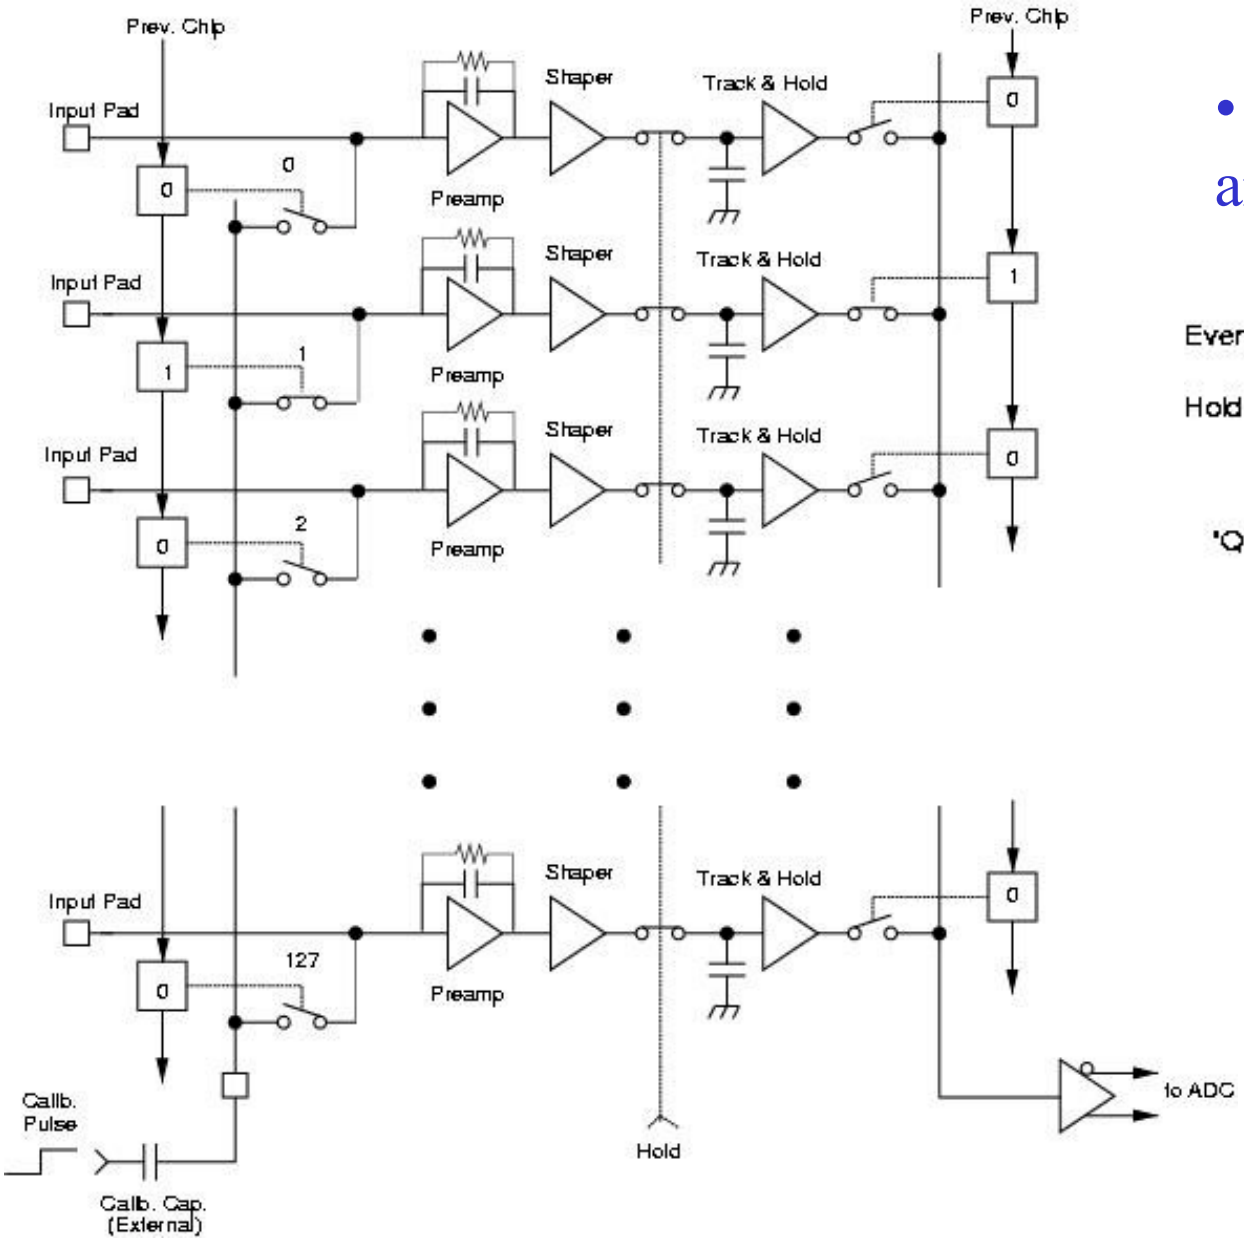

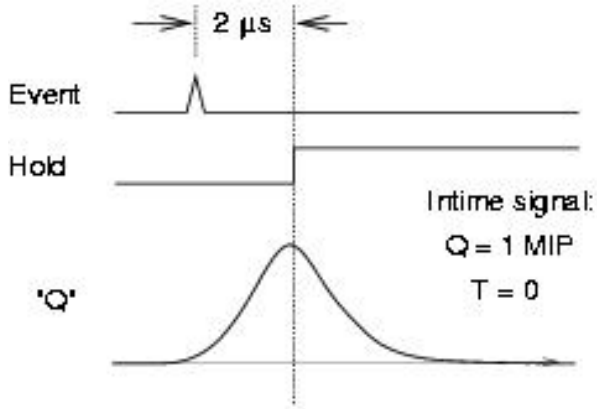

- Classic track and hold architecture
- VA1 shaping time provides latency for Level 1 trigger.
- 5 MHz Serial analog readout.

A Fine Point: the Level 0 Strobe

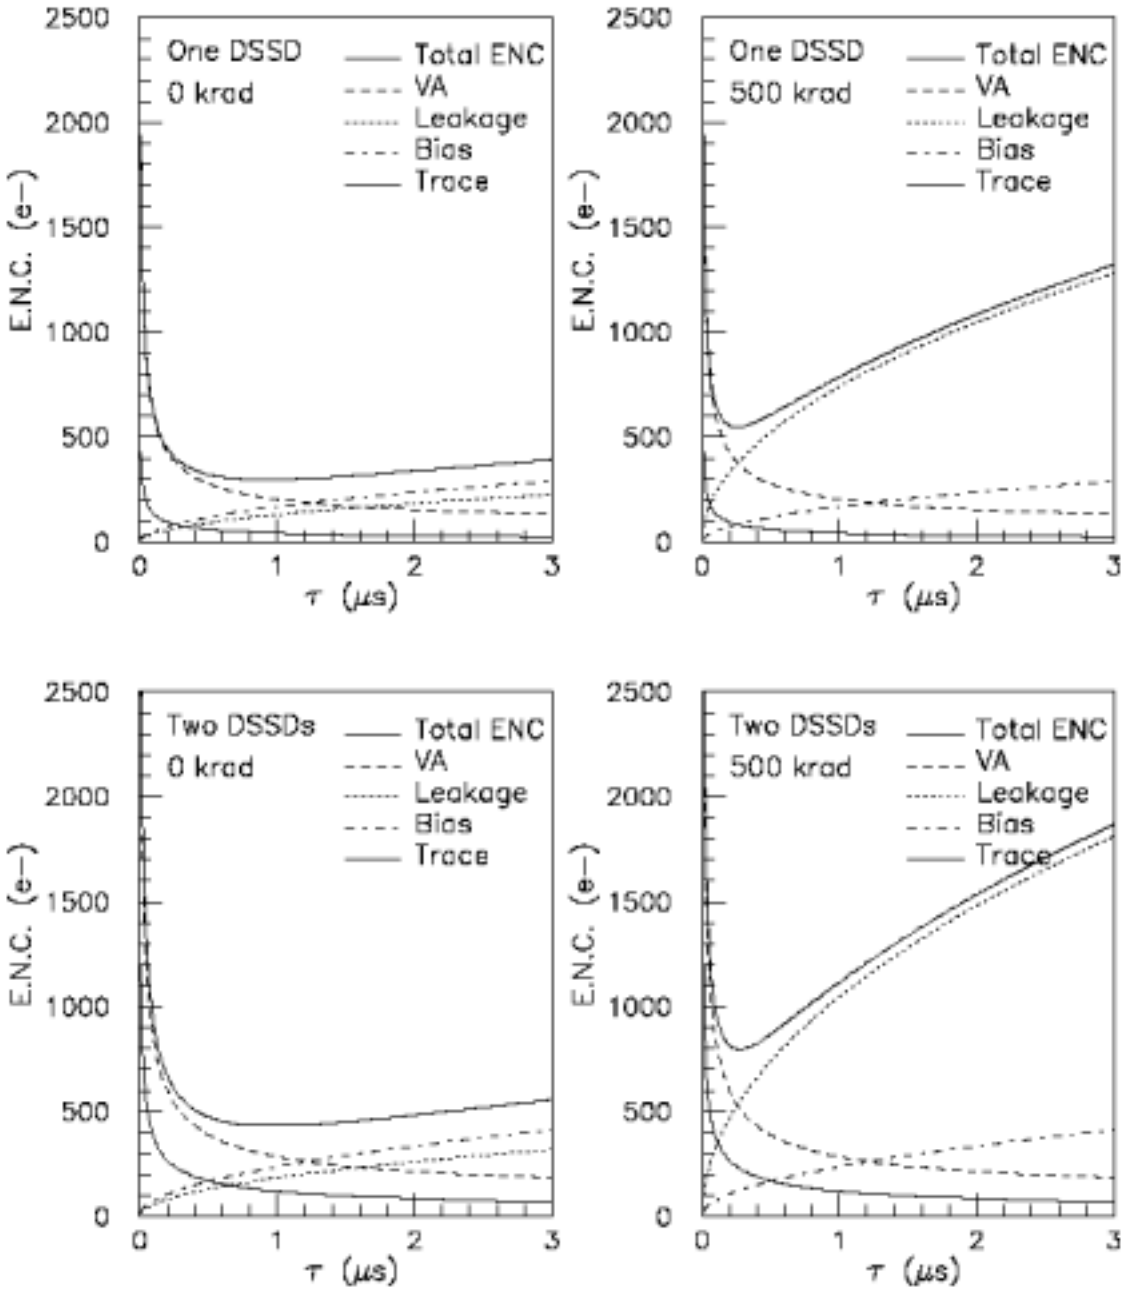

To avoid pileup and leakage effects, we will use a shaping time of ~500 ns, which is much shorter than the 2.2 μs latency of the Belle Level 1 trigger.

A Fine Point: the Level 0 Strobe

- To circumvent the latency difficulty, we have introduced a Level 0 trigger, which is based on simple logical combinations of the ToF counters.

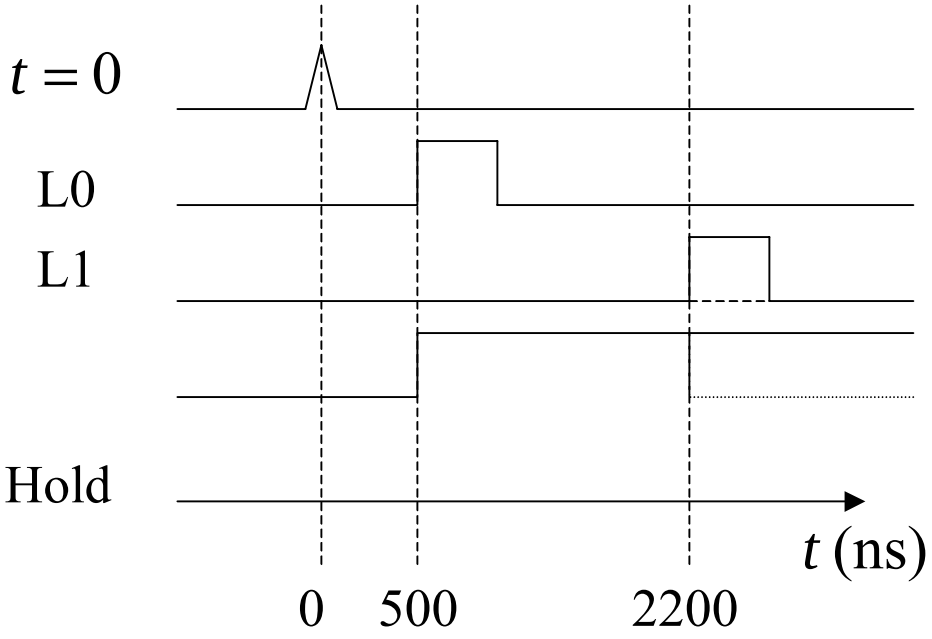

When a Level 0 trigger is issued, the hold is asserted until the Level 1 decision time. If Level 1 fires, the hold is left asserted and a full readout sequence ensues. If the Level 1 does not fire, the hold is cleared and the system is ready for a new event.

Level 0 Strobe: Critical Issues

- **Level 0 deadtime inefficiency**: this is currently

$$\eta = \Delta T_{\mathrm{LO}} R_{\mathrm{LO}} = 2.5\,\mu\mathrm{s} \times 10\,\mathrm{kHz} = 2.5\%$$

We clearly can't tolerate too large an increase.

- **Level 0 latency:** to take full advantage of the Level 0 scheme, we need to go to a 500 ns shaping time. This will require some modification of the ToF trigger, which currently has a latency of ~1 μs.

. . the Level 0 Trigger

One way to reduce the L0 rate is to add extra coincidence requirements. The small-cell CDC layer is a possible candidate for this, since it should have a small drift time.

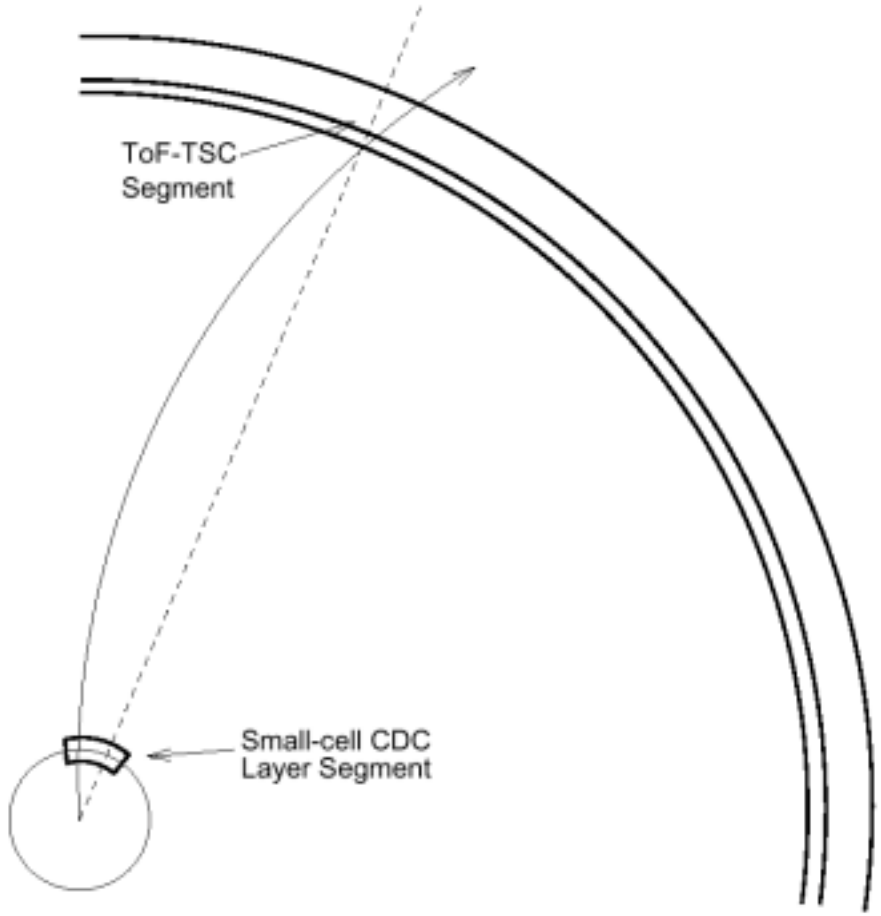

. . the Level 0 Trigger

If we assume 5 mm (diameter) cells and 20 ns/mm drift times, the resolving time will be ~50 ns. We add a factor of two safety margin and take ΔT=100 ns. Thus

$$p \cong R\Delta T < 0.25 \Rightarrow R_{\max} = \frac{0.25}{100\,\text{ns}} = 2.5\,\text{MHz}$$

If the small-cell CDC layer is divided into eight segments, this implies that for total rates of up to 20 MHz (for the entire layer), adding the SC-CDC in coincidence with the ToF will lower the L0 rate.

The VA1TA

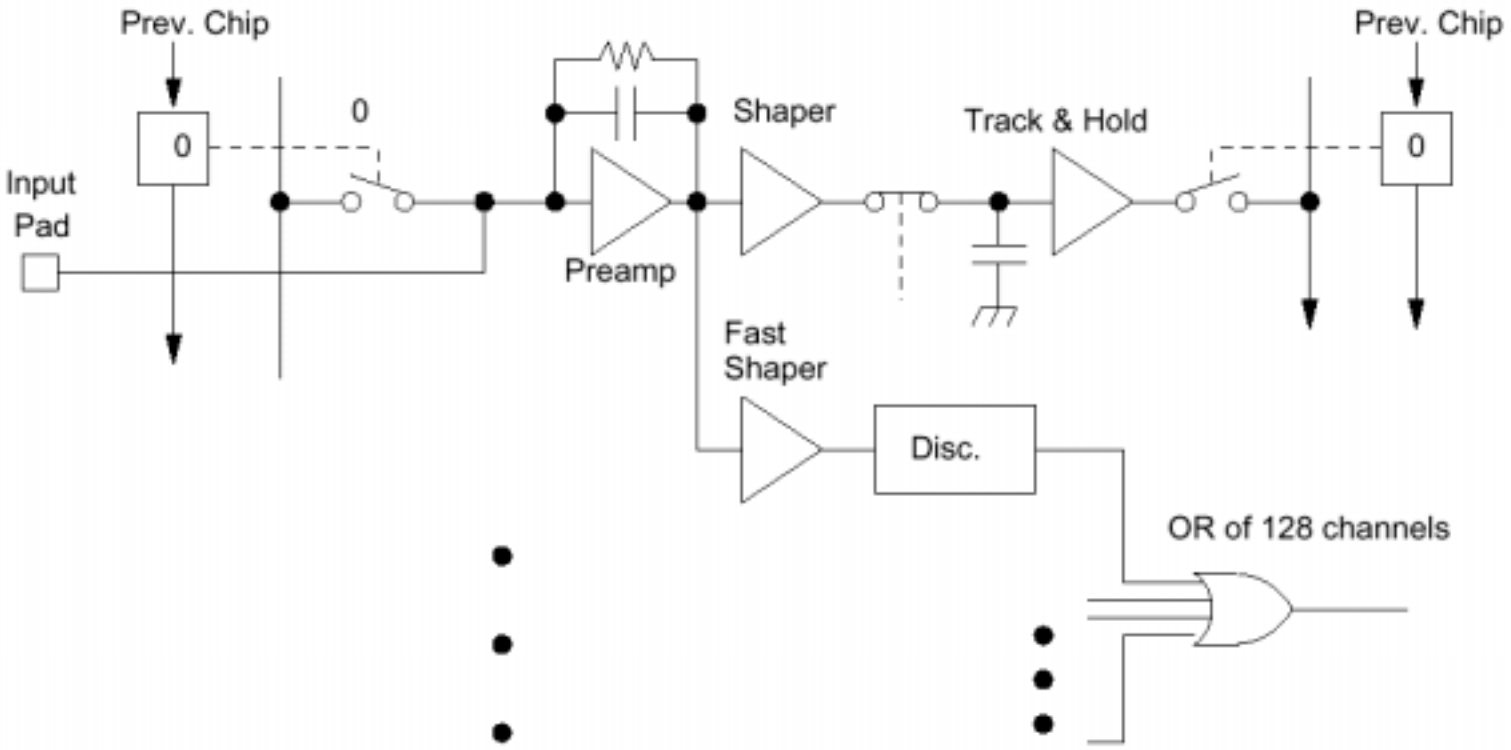

Another way to provide the a Level 0 trigger is to extract prompt trigger information from the VA1 chip itself. This allows for a much more powerful trigger since it is possible to do crude tracking.

The VA1TA

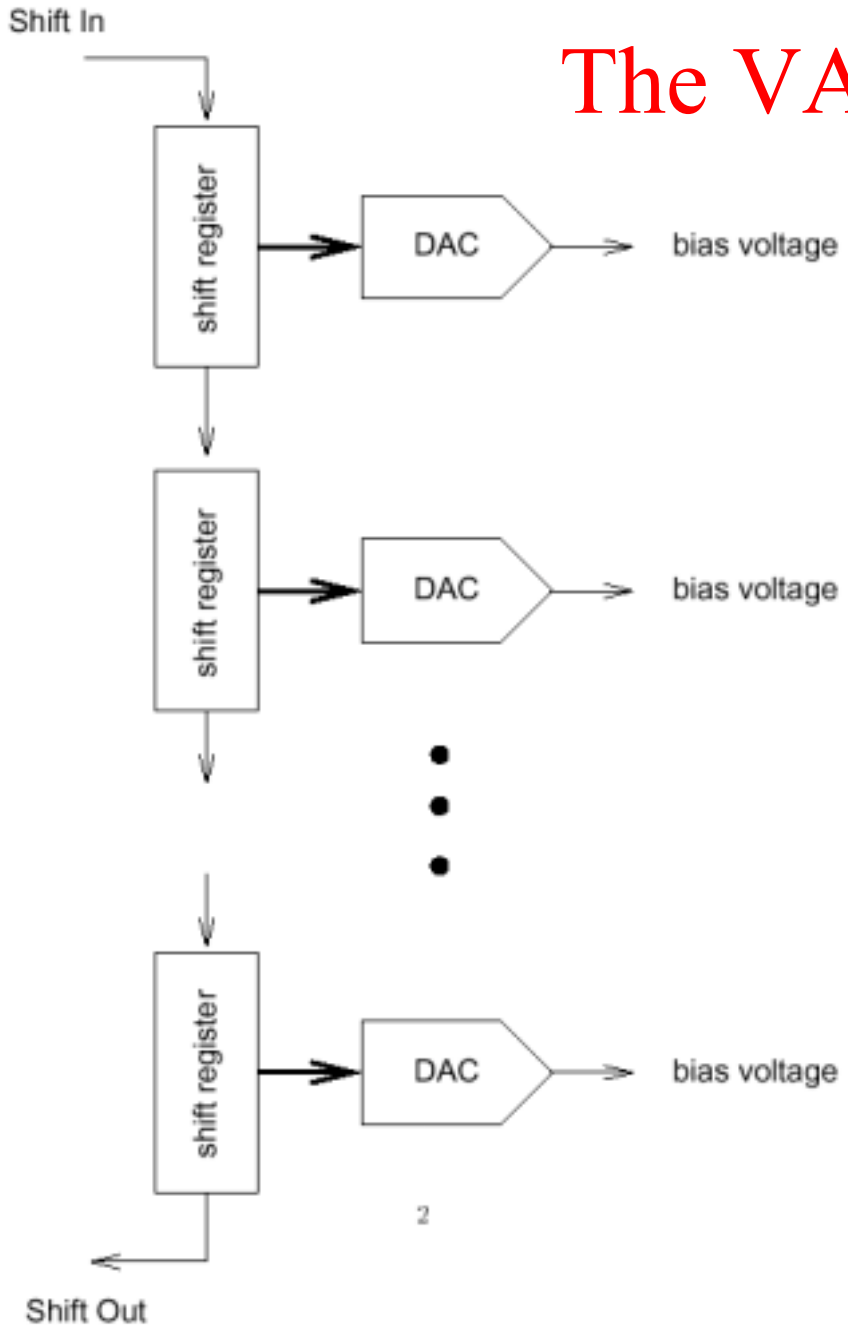

In the VA1, the bias voltages are generated off chip. The added shaper and discriminator circuit blocks of the VA1TA require additional bias voltages. To reduce the number of external connections, the bias voltages for the VA1TA are set using on-chip DACs.

Radiation Hardness

Simple Model for CMOS Radiation Damage

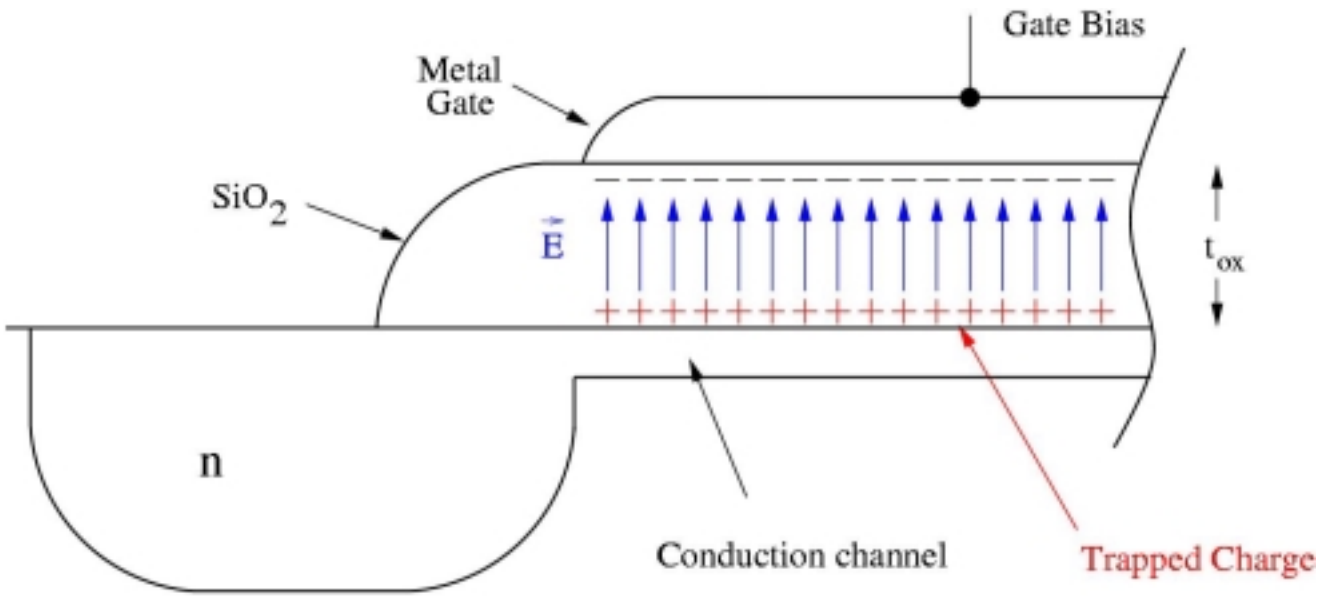

depends on ionization

$$\vec{E} \propto \rho\, \vec{x}$$

$$\Delta V_{th} = \int \vec{E} \cdot d\vec{x}$$

$$\propto \rho \int_0^{t_{ox}} \vec{x} \cdot d\vec{x}$$

$$\propto \rho\, t_{ox}^2$$

Reducing the oxide thickness by half is equivalent to cutting the dose by four.

AMS Process Comparison

Feature Size	t_{oxide}	Projected Tolerance
1.2 μm	250 Å	≡ 200 kRad
0.8 μm	160 Å	500 kRad
0.6 μm	125 Å	800 kRad
0.35 μm	75 Å	2200 kRad

One in fact expects to do considerably better than t_{oxide}^{-2} scaling since tunneling plays an important role in deep submicron processes.

Original VA1 Test Results

AMS 1.2 um
CMOS

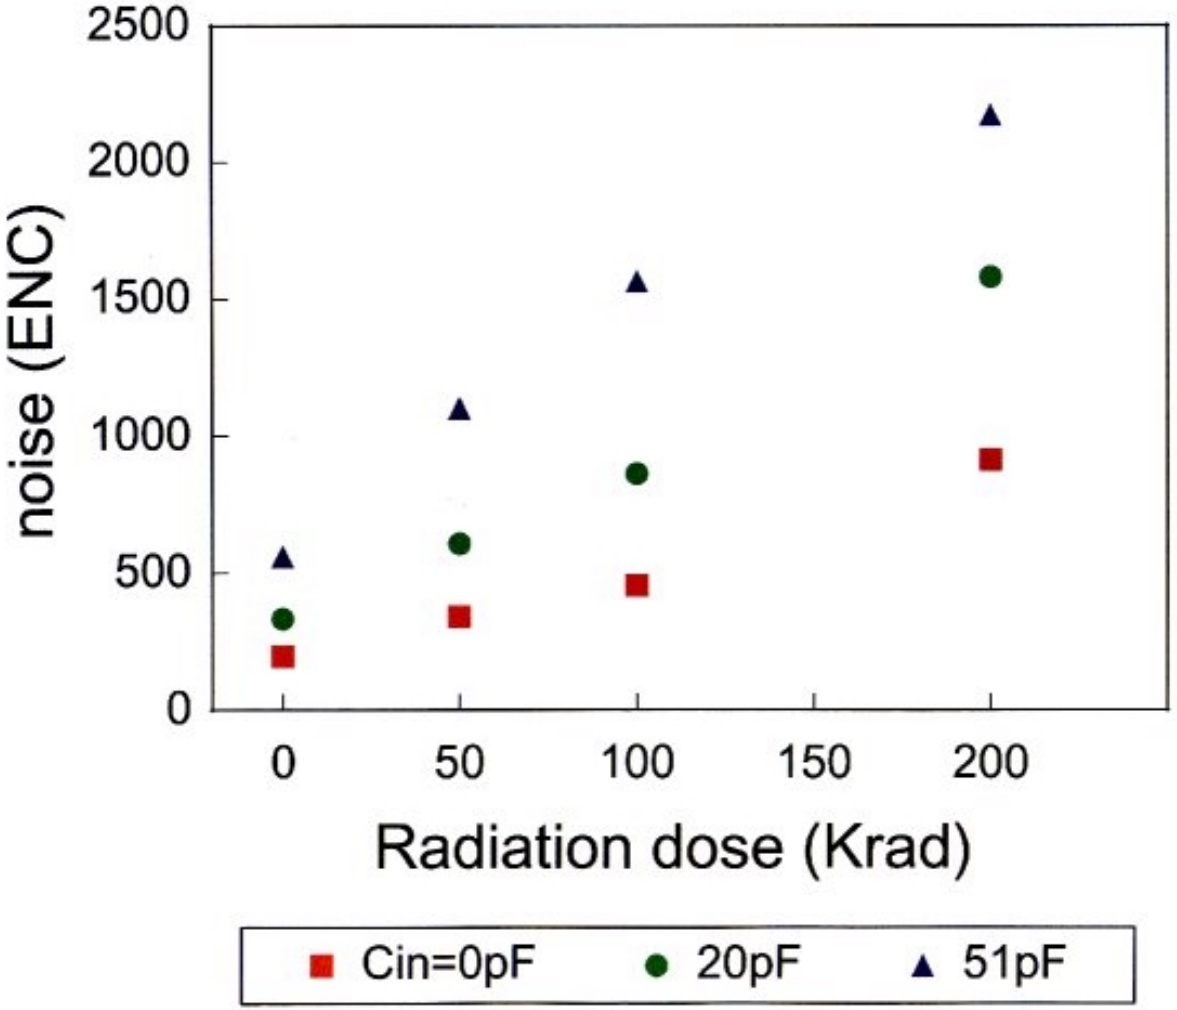

AMS 0.8 um
CMOS

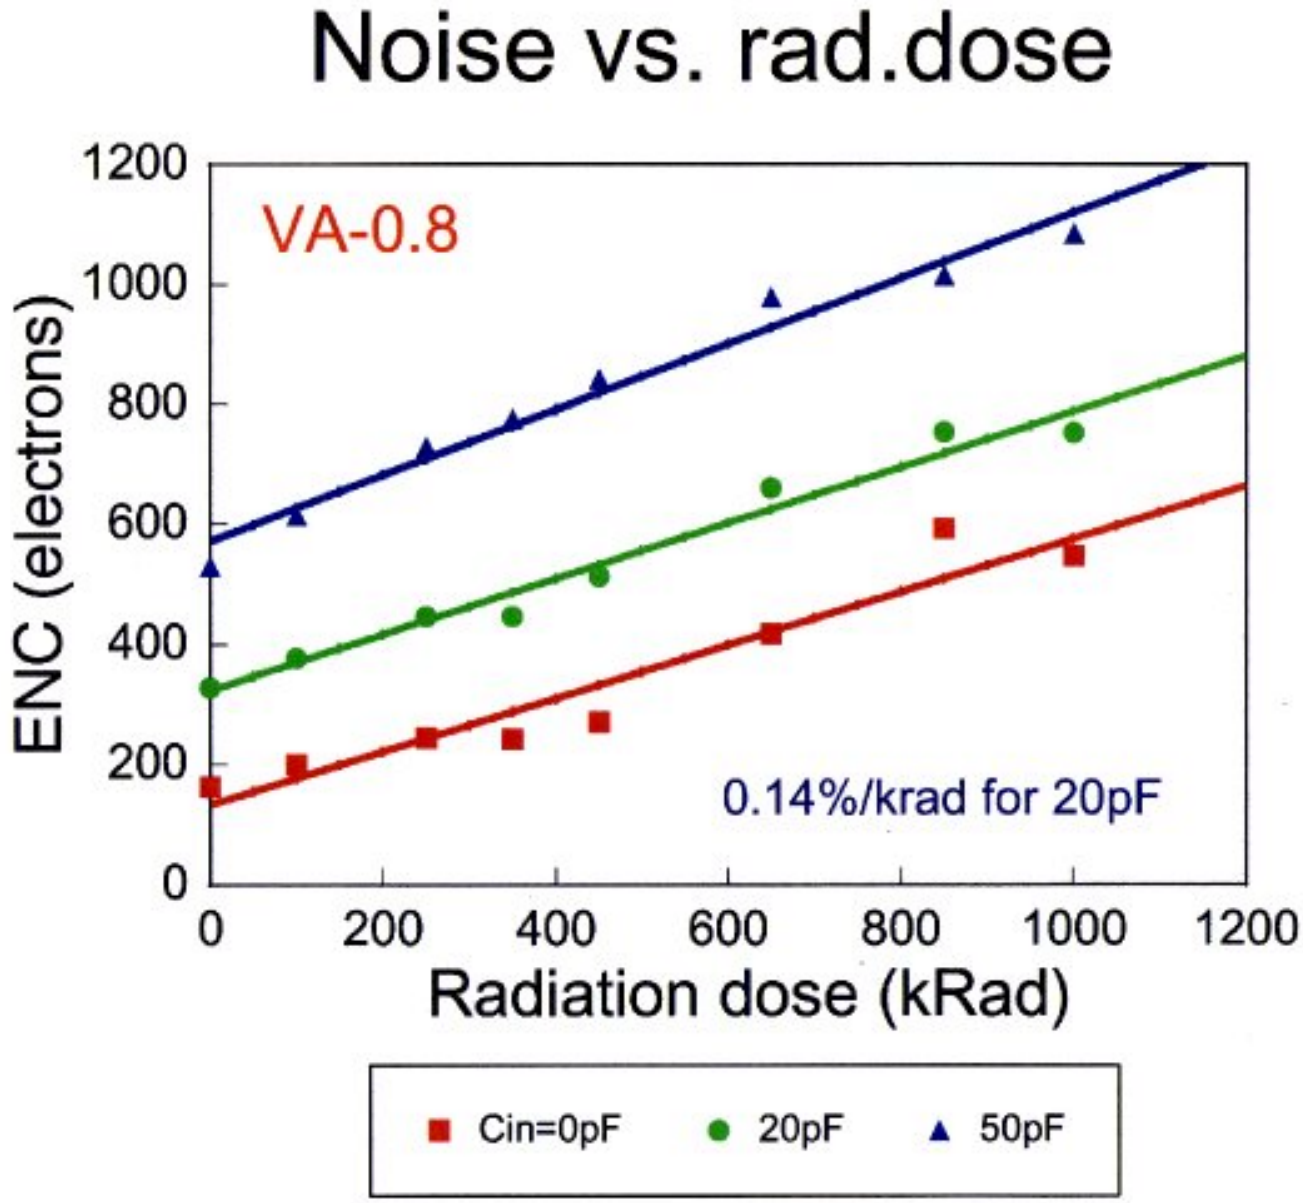

AMS 0.8 um CMOS

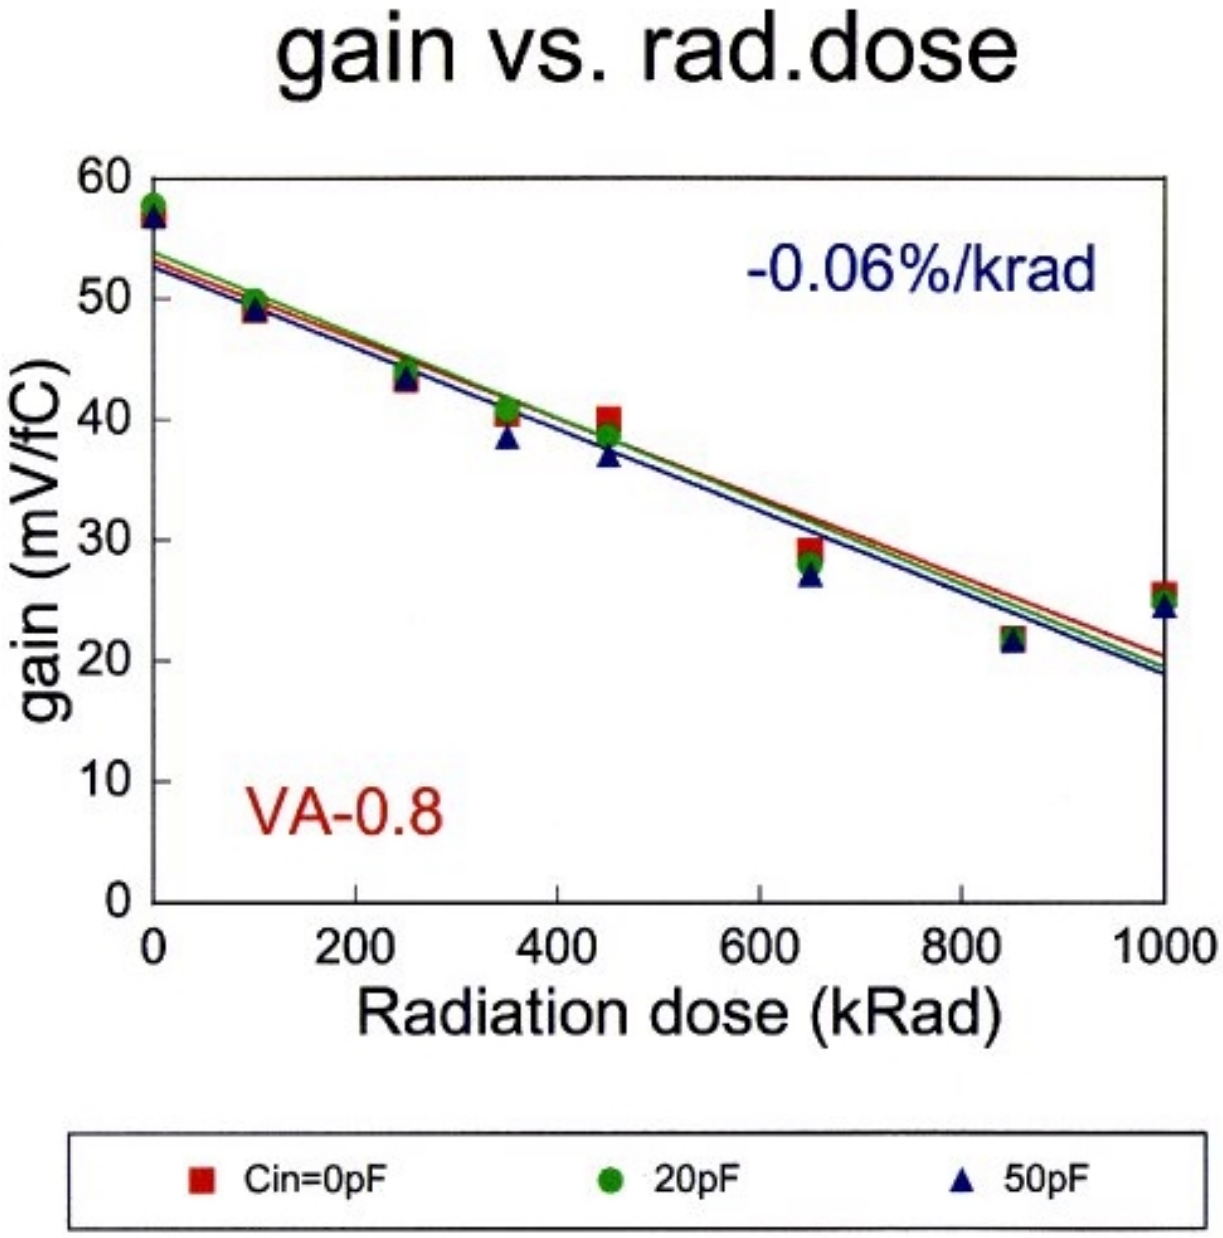

Process Comparison

In the range of interest to BELLE, the noise performance dramatically improves with decreasing feature size, as expected.

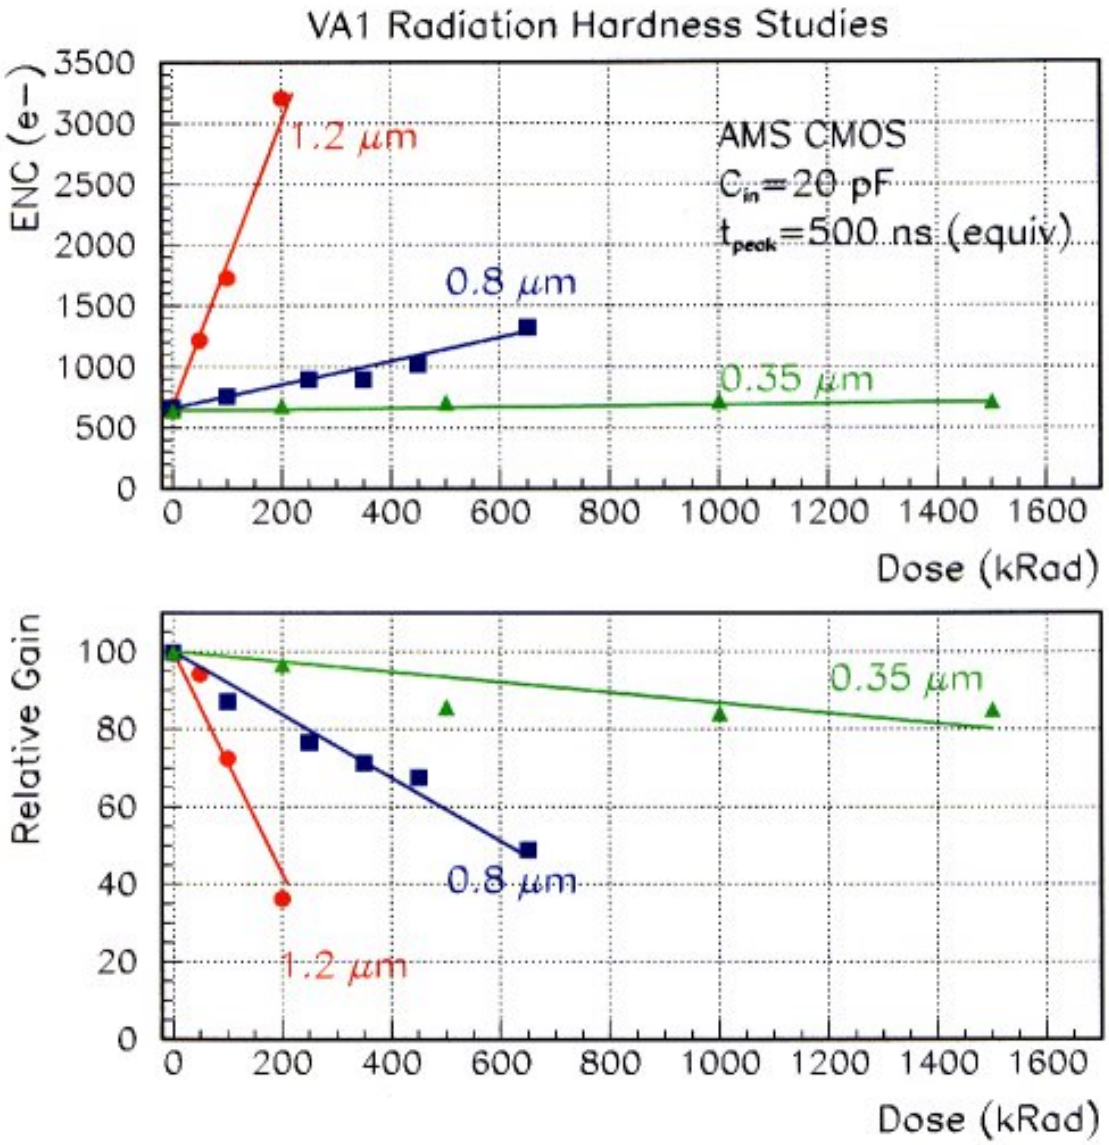

Indeed, the 0.35-um process exhibits phenomenal radiation hardness.

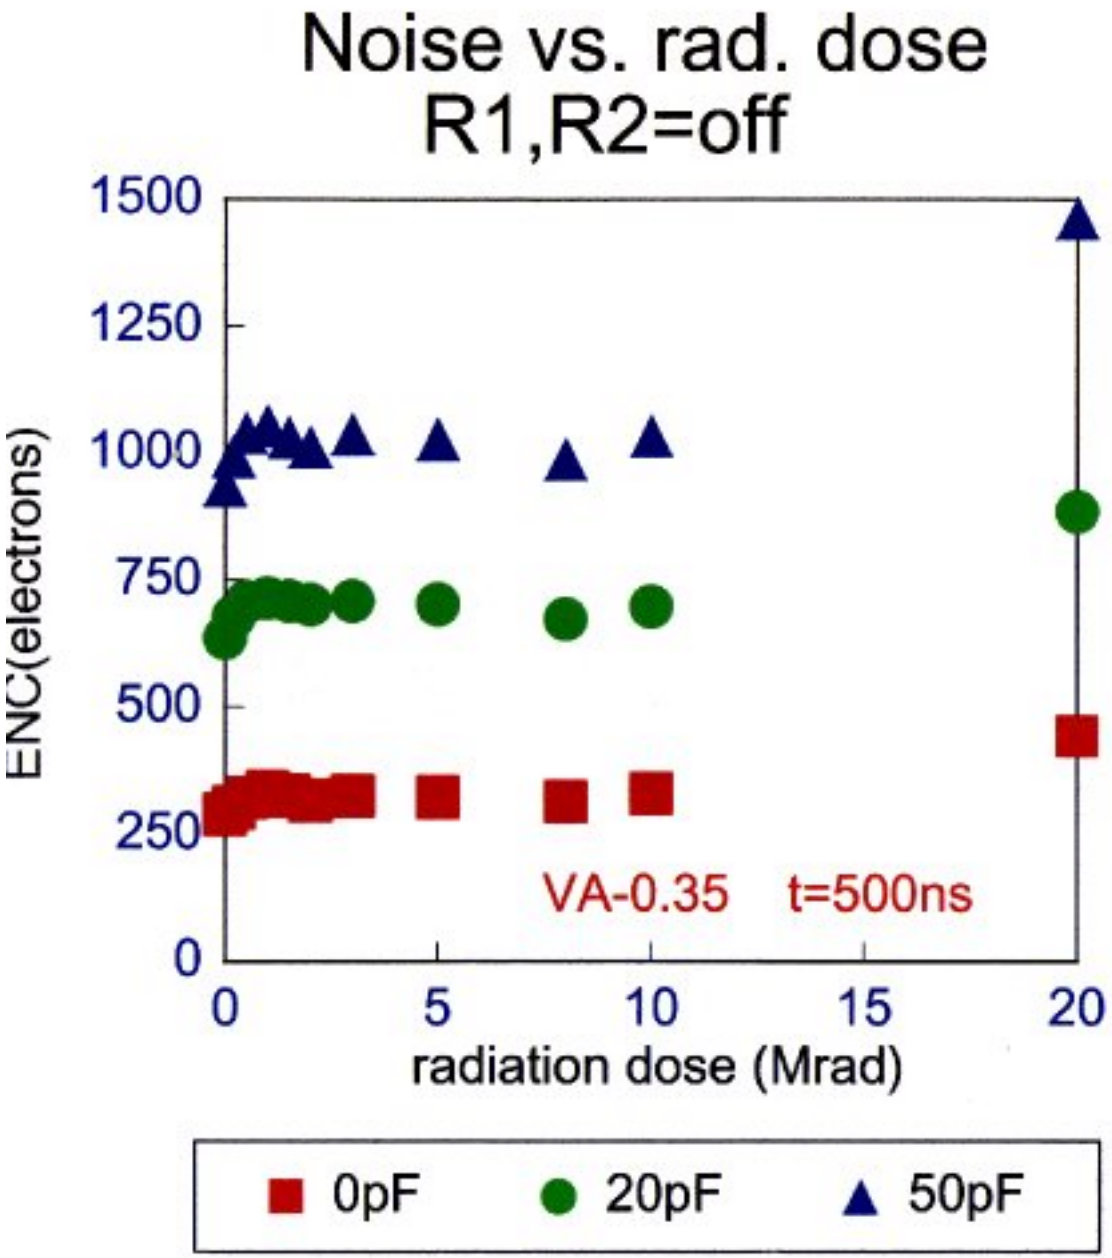

AMS Process Comparison

Feature Size	t_{oxide}	Projected Tolerance	Observed Tolerance
1.2 μm	250 Å	≡200 kRad	200 kRad
0.8 μm	160 Å	500 kRad	1200 kRad
0.6 μm	125 Å	800 kRad	***
0.35 μm	75 Å	2200 kRad	> 20 MRad

Single Event Upset and Latch-up Concerns

- Although it appears that the last run of the VA1 gives us more-than-adequate hardness with respect to integrated dose, the situation with respect to "latch-up" is less clear. In particular, we worry about the effect of large energy deposits localized in space and time.
- Single event upset (SEU)
  - Flip-flops flip when they should flop.
- Single event latch-up (SELU)
  - Parasitic conduction paths form in the IC's substrate.
  - Current draws large enough to destroy the chip through Joule heating can occur.

SEU, SELU, and the VA1

- We have seen latch-up-like symptoms in SVD1.
- It is not clear (to me) how latch-up effects will vary with feature size.
- We have not (yet) carried out systematic latch-up studies with the 0.35 micron VA1.
- Our mode of operation renders us fairly insensitive to latch-up problems in the logic, but we should still be careful, especially for the VA1TA, where there is more on-chip logic.

SELU Mitigation

- Since we had time in the schedule we decided to refabricate the VA1 chip in the AMS 0.35 um CMOS with an epitaxial layer added.
- According to the AMS measurements, this should improve the latchup performance of the VA1 chips.

SEU Mitigation

- The epitaxial layer should also help mitigate SEU effects.

• The logic design of the VA1TA will incorporate 2 of 3 majority logic to detect and automatically correct SEU-induced errors.

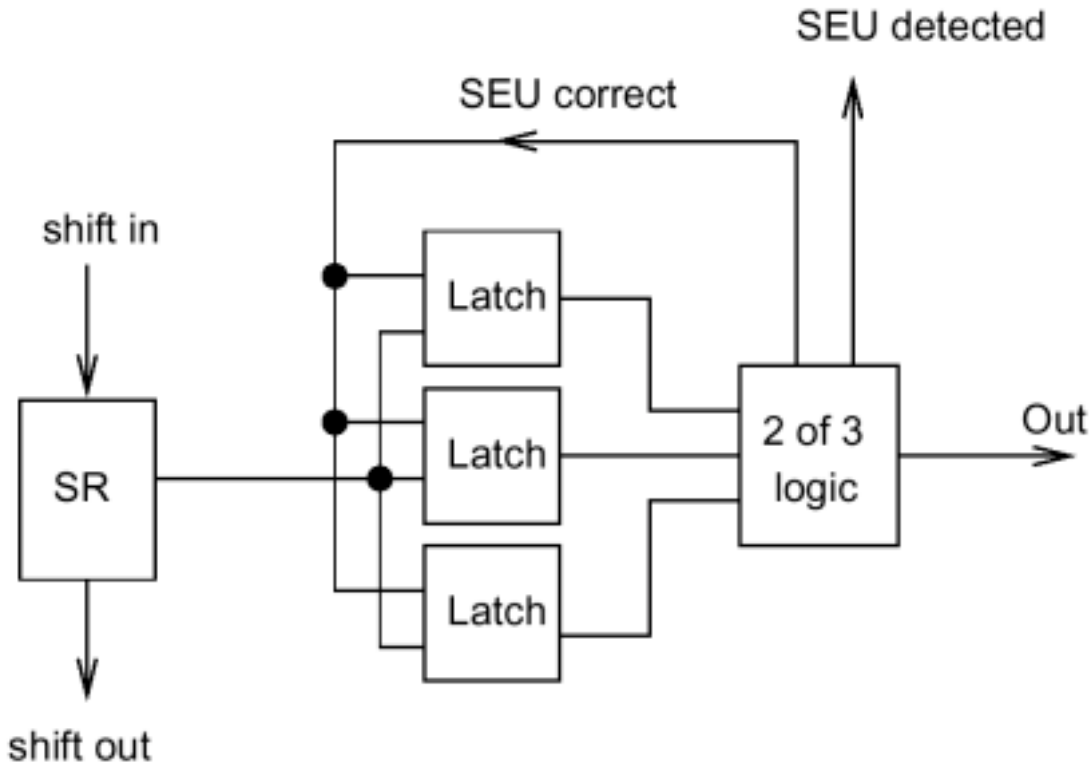

VA1 and VA1TA Status

- Last year we submitted enough wafers of the VA1 design to the AMS 0.35 μm process to ensure that we would have enough chips to complete SVD2.
- That run was successful according to the following criteria:
  - Chip functionality ✓
  - Hardness against integrated dose. ✓
  - Satisfactory "latch-up" performance. ?

. . .VA1 and VA1TA Status

- In the next few weeks we will submit two major engineering runs:
  - VA1 to AMS 0.35 μm *with* epitaxial layer.
  - VA1TA to AMS 0.35 μm *with* epitaxial layer.
- In both cases a successful run will give us enough chips to complete SVD2 (~1600 good die)

System Block Diagram

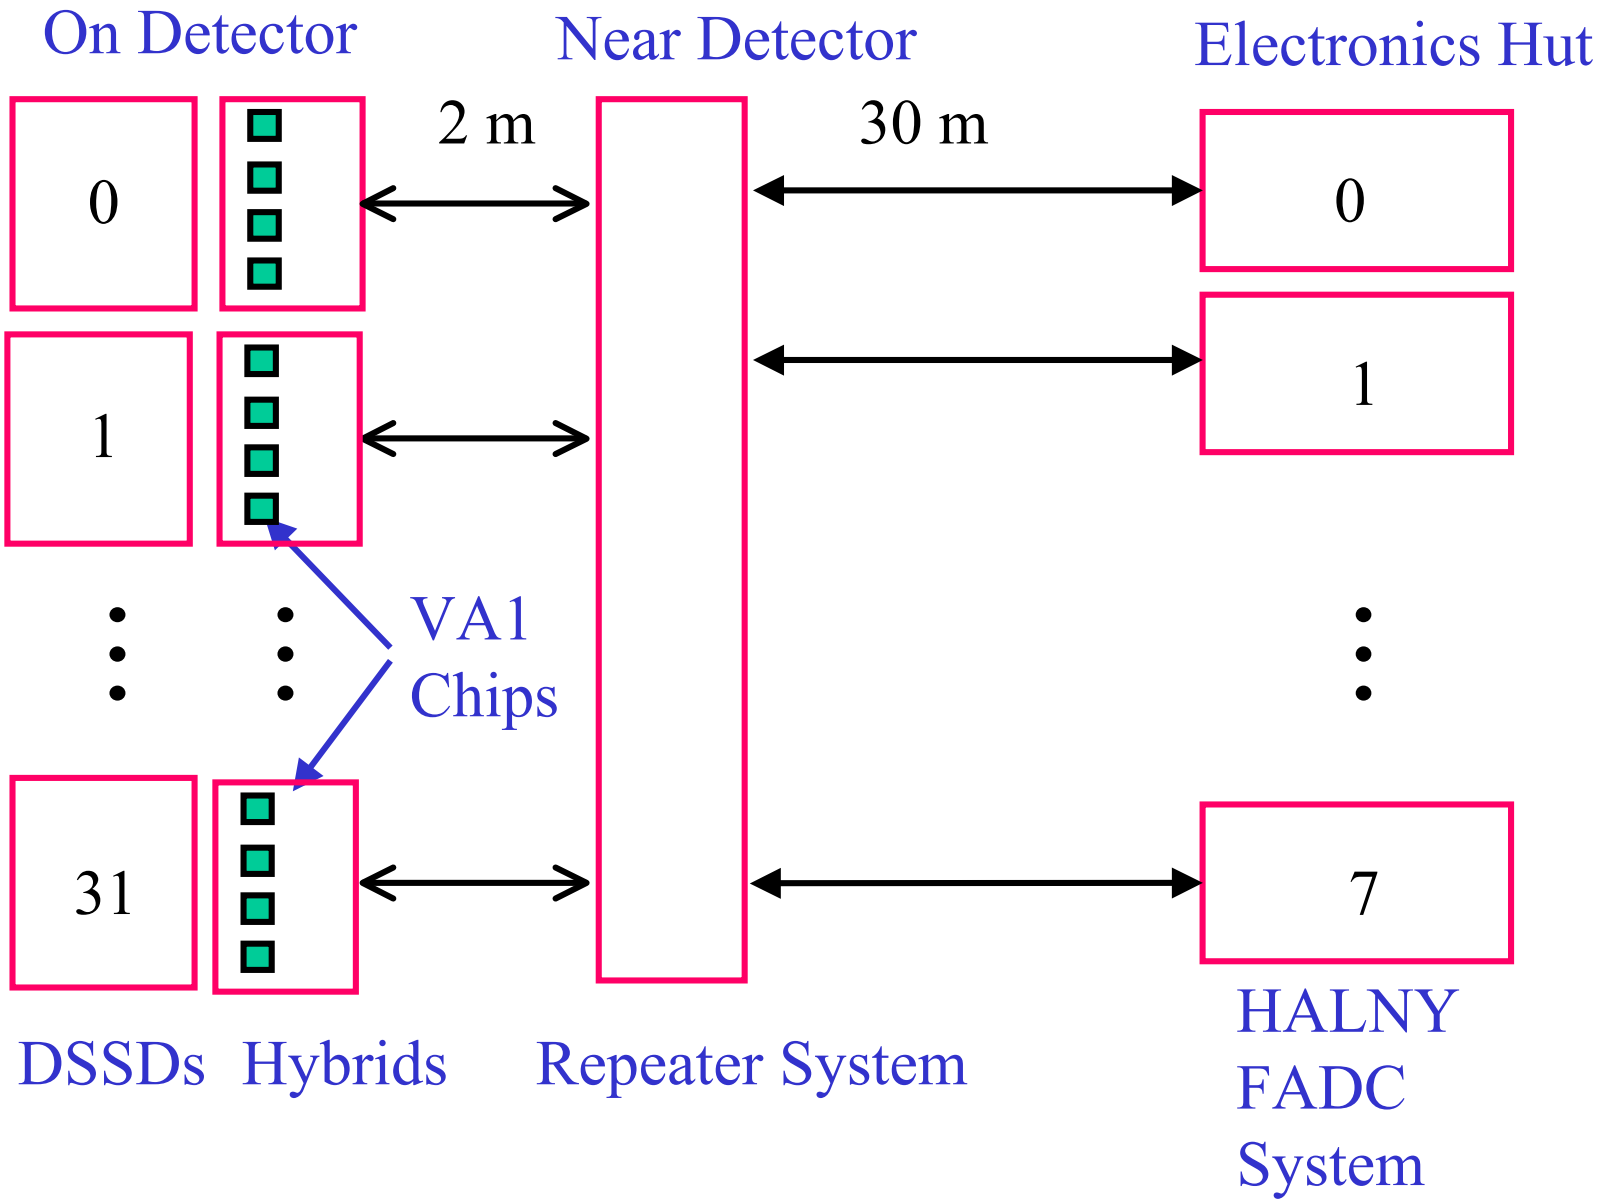

Hybrid

The hybrid contains only a modest number of components apart from the VA1 chips themselves.

The circuit design will change for the VA1TA.

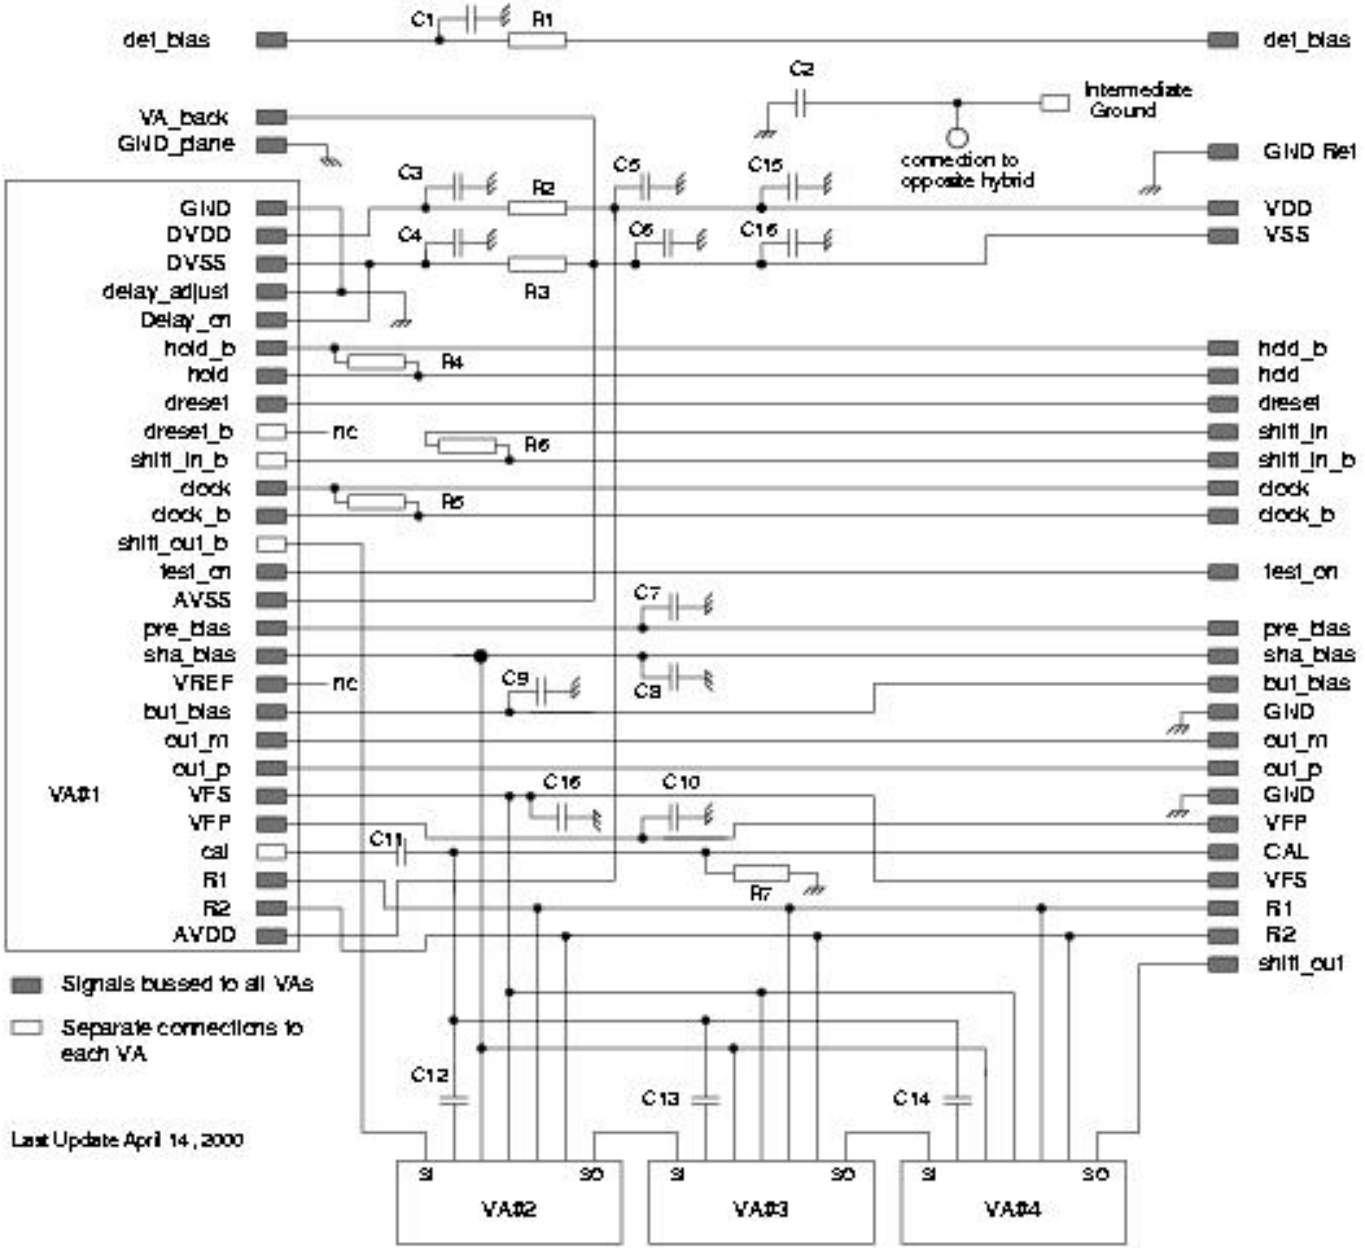

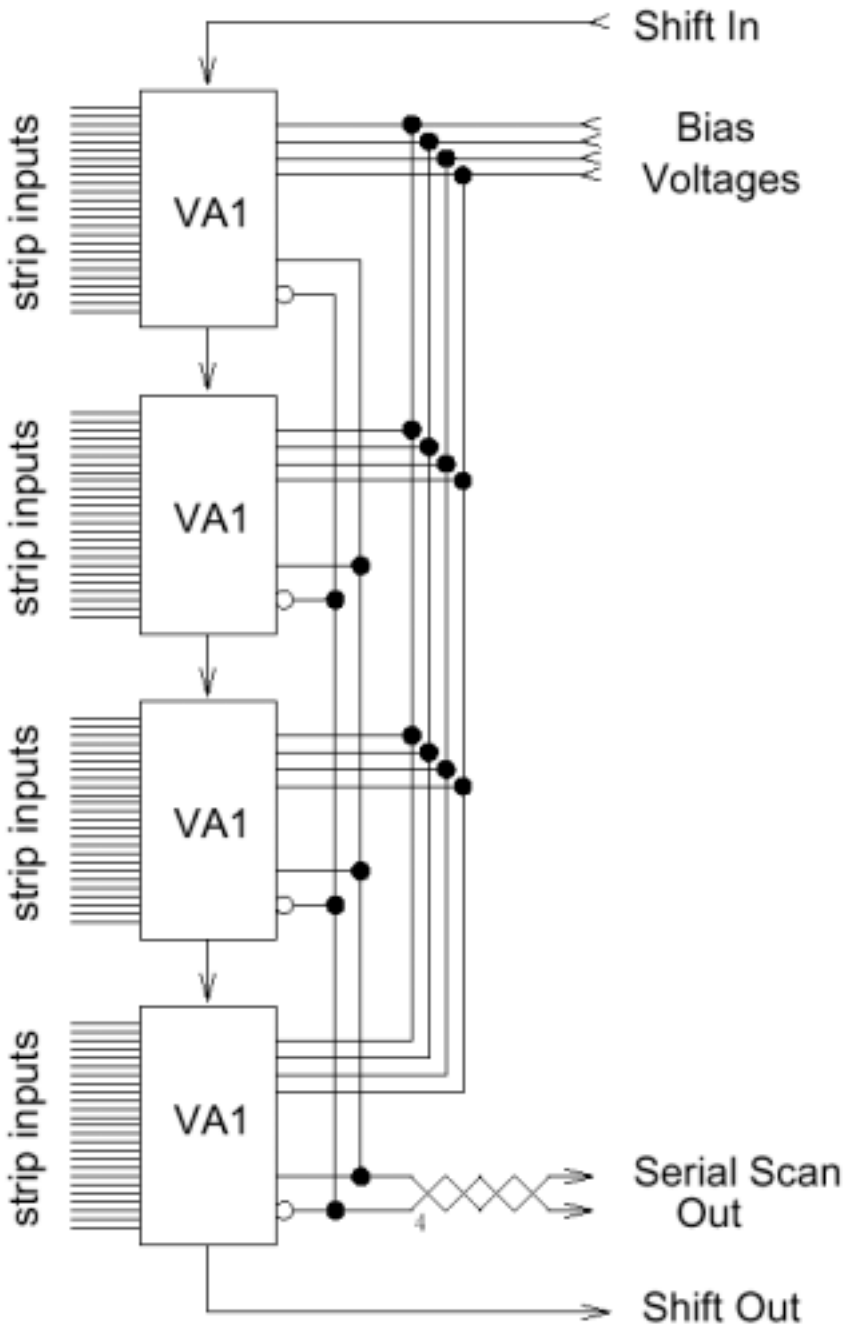

VA1 Hybrid Concept

The bias voltages are supplied externally, but there is just one analog output.

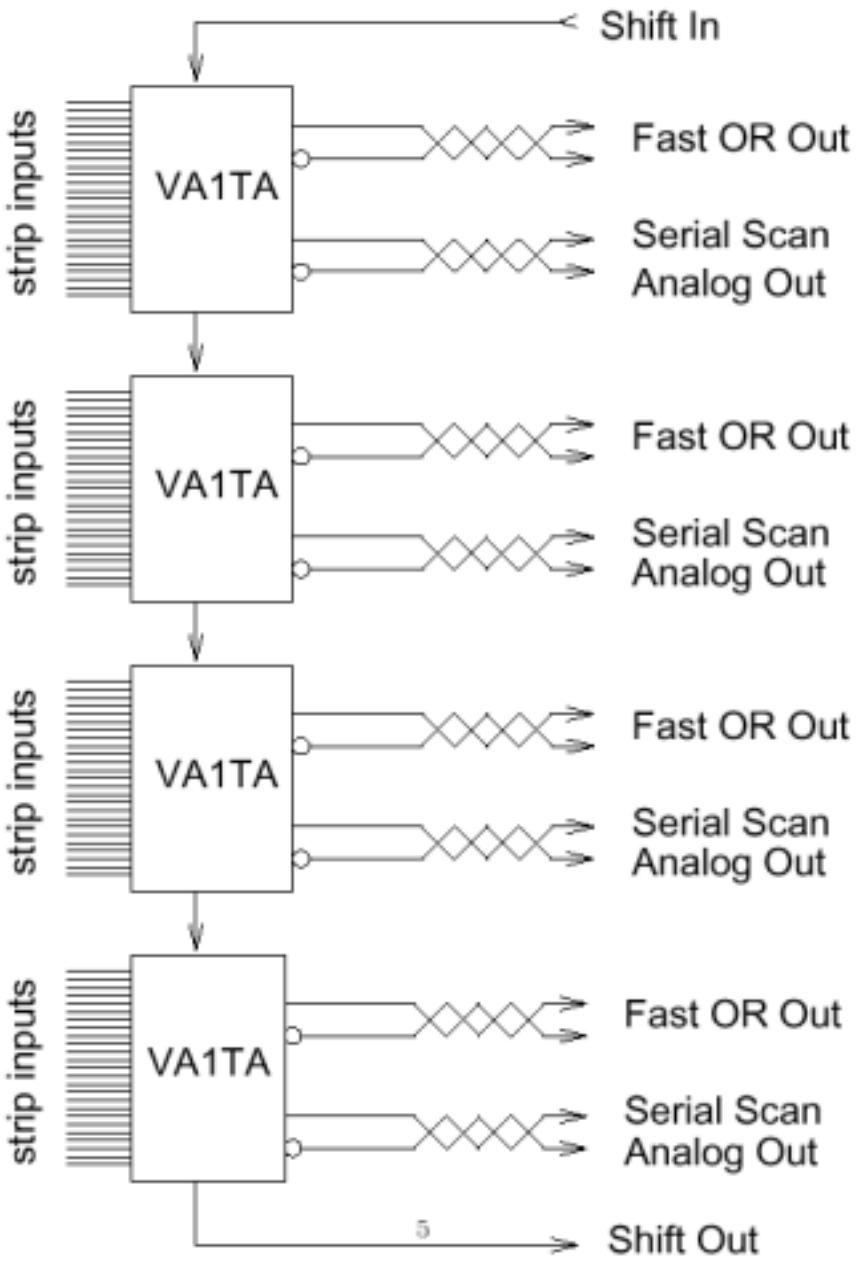

VA1TA Hybrid Concept

The bias voltages are internally generated, but there are many more outputs:

- Four analog outputs
- Four fast-OR outputs for Level 0 trigger.

The outputs are differential, so there is a total of 16 lines.

Hybrid Schedule 2001 Milestones*

Mar 1: Hybrid CDR complete

Mar 24: FR4 hybrid submission to production.

Apr 15: FR4 hybrid production finished

Apr 28: FR4 hybrid tested.

Jun 1: Submission of Kyocera design.

Sep 1: Kyocera delivery to Ideas.

Sep 30: Deliveries to Belle start

> 30 ceramics/week (360 needed)

Dec 31: >260 hybrids delivered

*Estimates assuming 2002 installation.

Repeater System

- Mostly new design
  - Floating ground to accommodate DC coupled detectors.
  - Redesign of many logic elements to improve reliability and maintainability.
- Prototype designed for original SVD 2.0
  - Proof of principle for many items
  - Five-layer SVD presents challenge in terms of a sufficiently compact mechnical package.
  - Progress slowed considerably due to uncertainty regarding Inner Tracking Task Force conclusions.

Halny FADC Board (existing system)

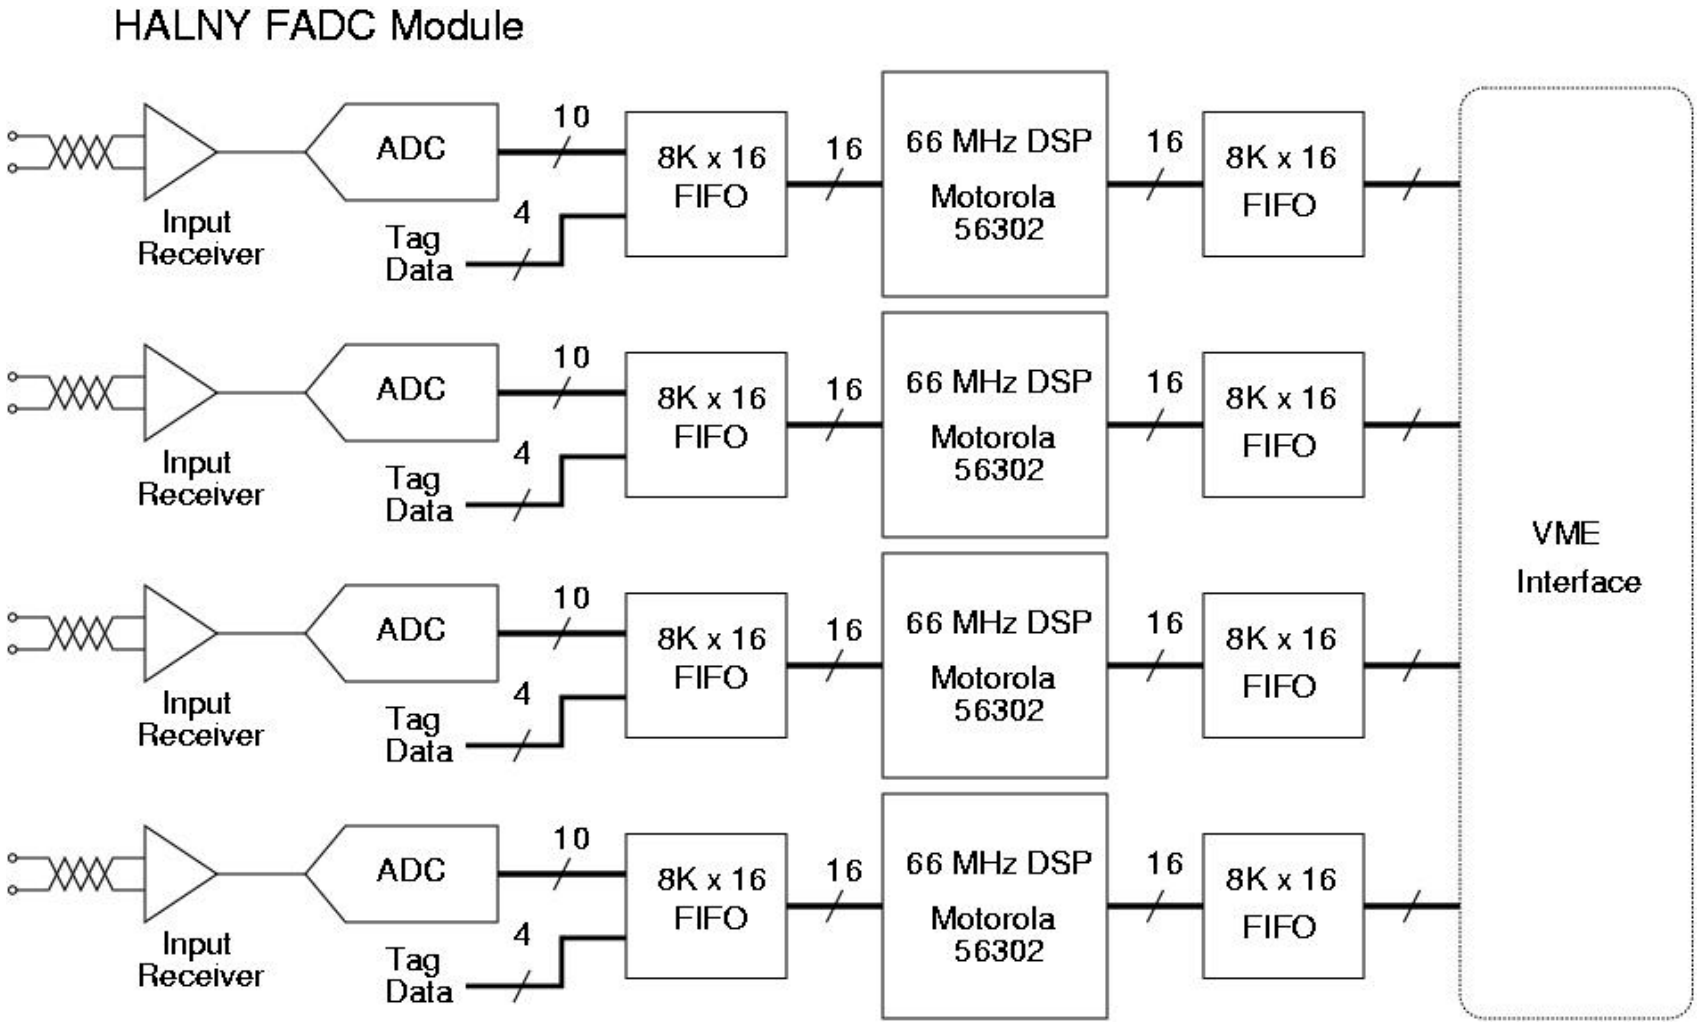

The HALNY is a 6U VME Module Developed by INP Cracow

New System Architecture with Level 1.5

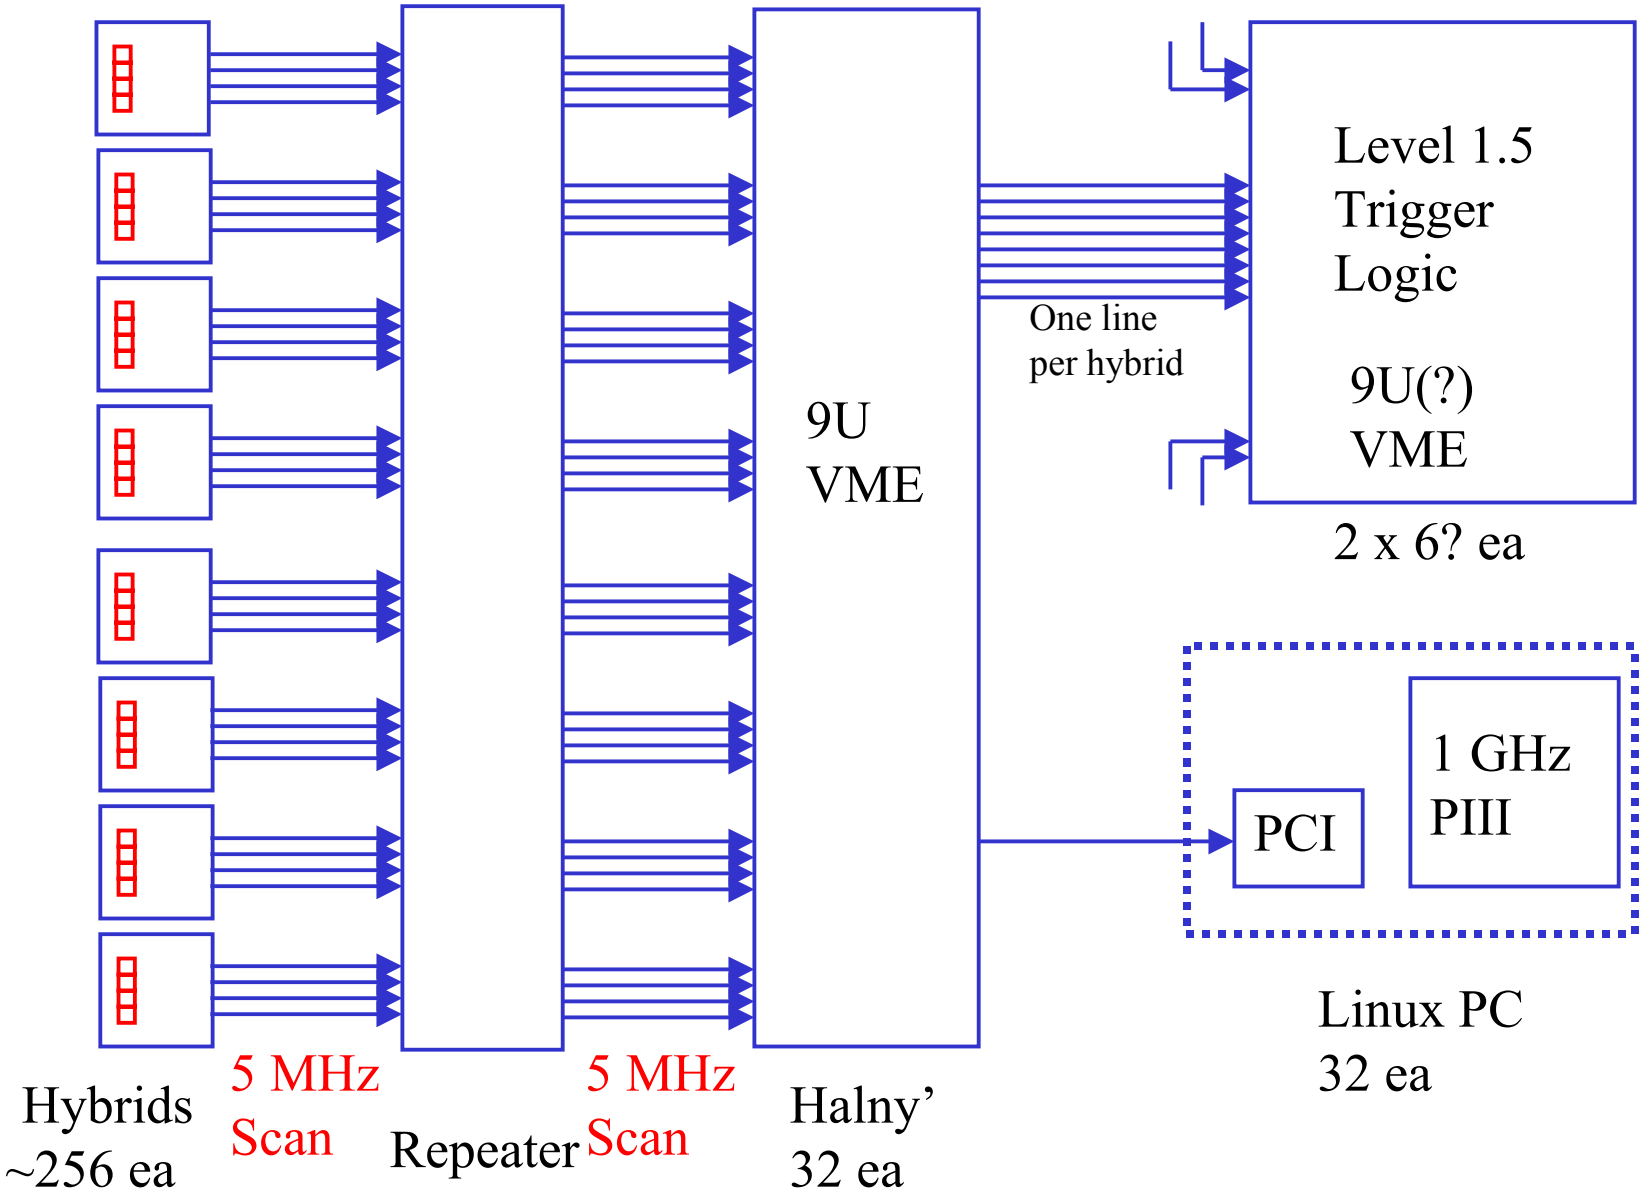

Logic Flow

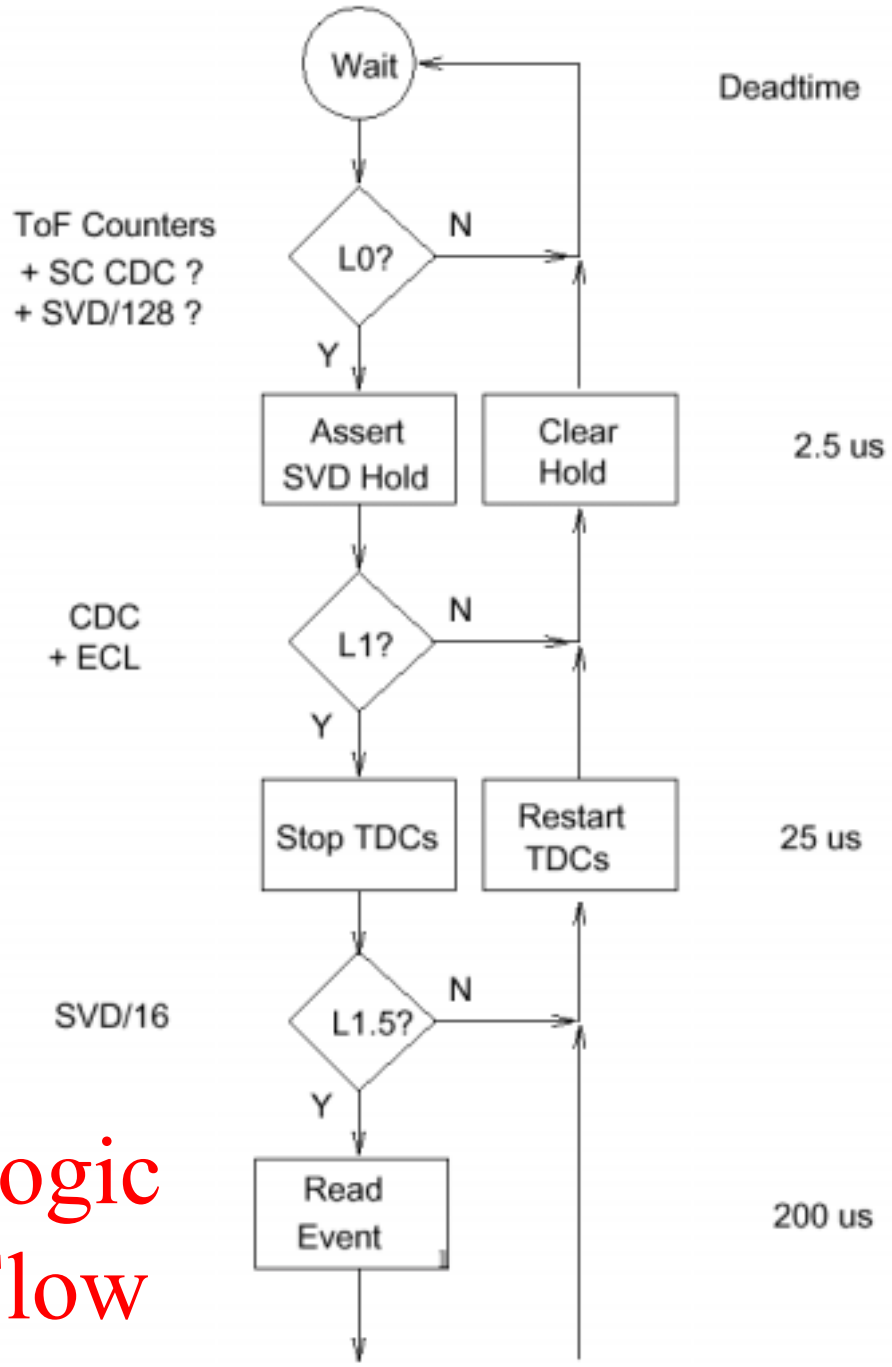

Level 1.5

- Level 1.5 makes use of ADC information picked off at the time of the serial scan.
- Virtual trigger strips with an effective width equal to 16 actual strips can be quickly processed using a massively parallel system of FPGAs.
- Logic capable of measuring an event's *z*-vertex position can be implemented, allowing discrimination against beam-gas events.

Level 1.5 Advantages

- No new chip development needed (although we are progressing with the VA1TA).
- The trigger-bit extraction logic can implement pedestal subtraction and cluster finding, resulting information having “off-line” quality.
- Since more time is available to form the trigger, more sophisticated algorithms are possible, opening up the possibility of higher rejection factors, which would reduce the load on the down stream DAQ.

Trigger Information Extraction

Old

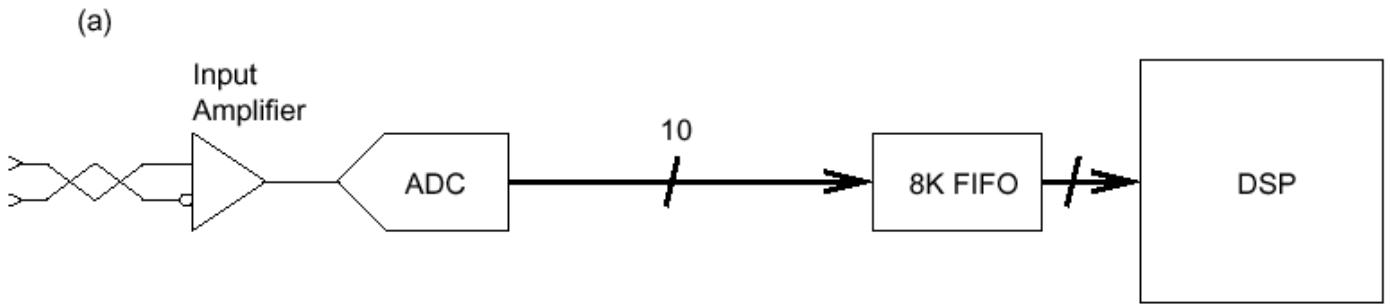

New

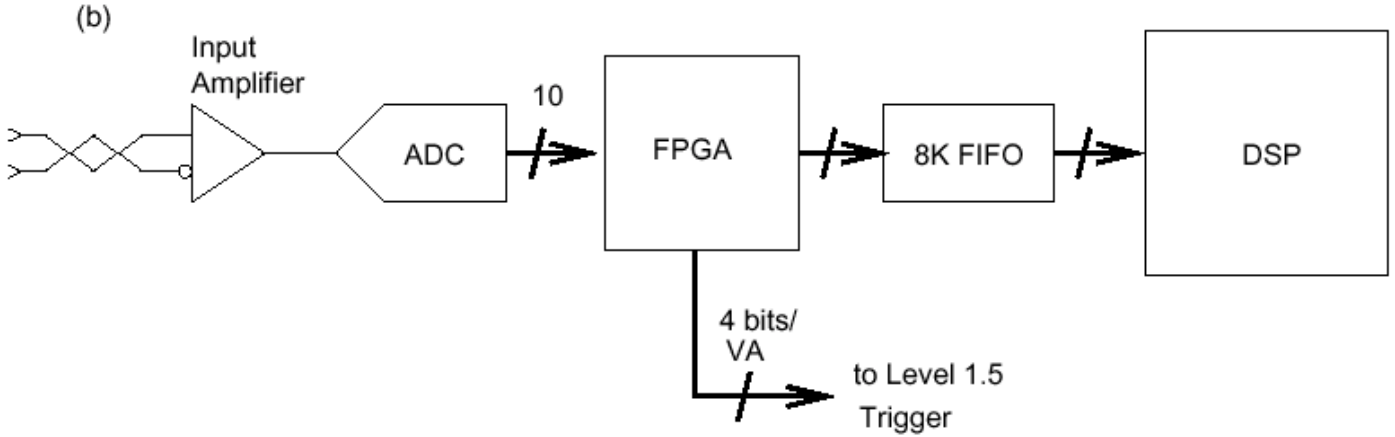

. . .Trigger Information Extraction

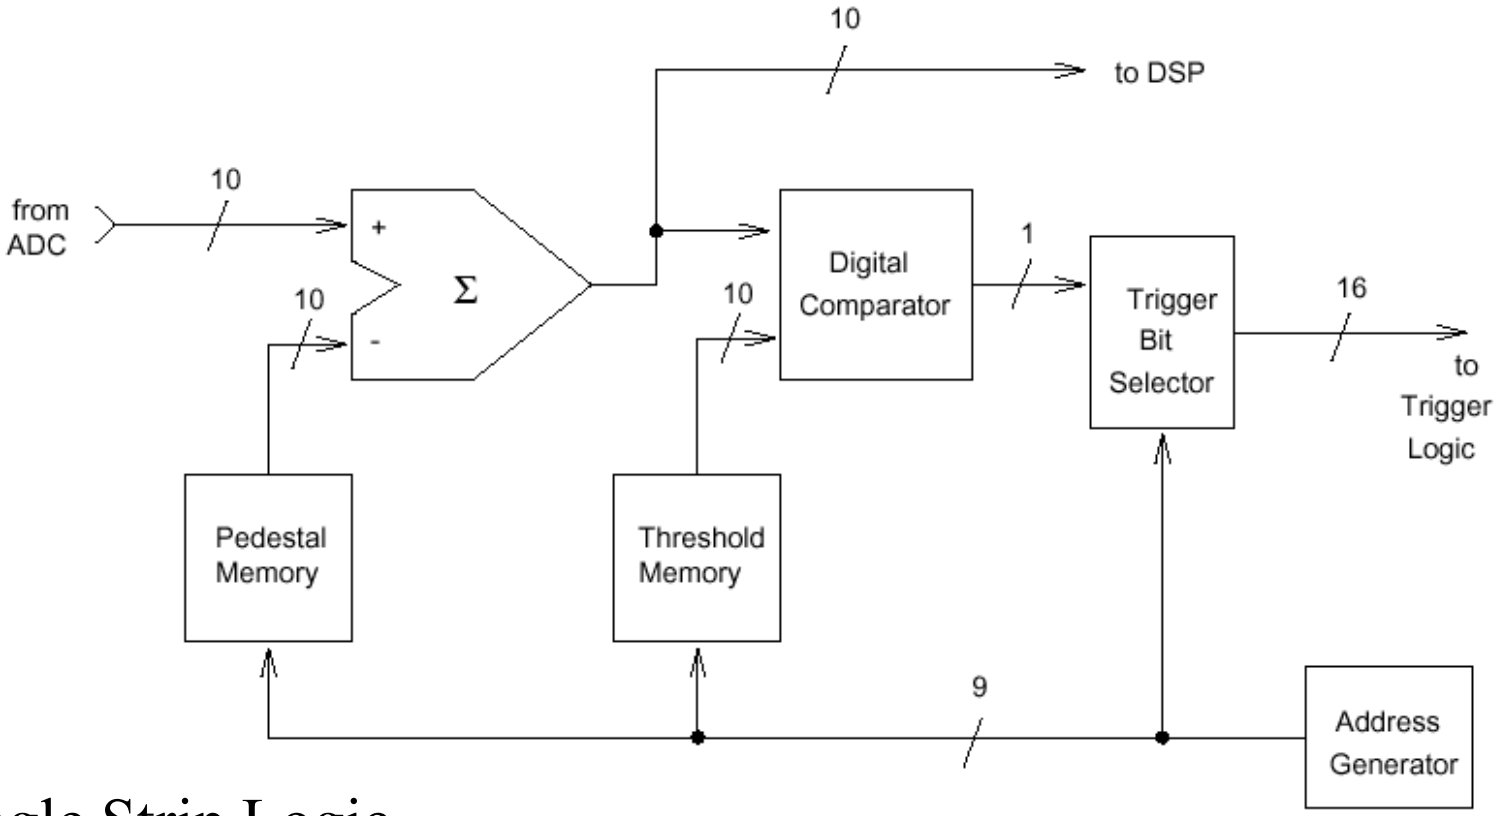

Single Strip Logic

. . .Trigger Information Extraction

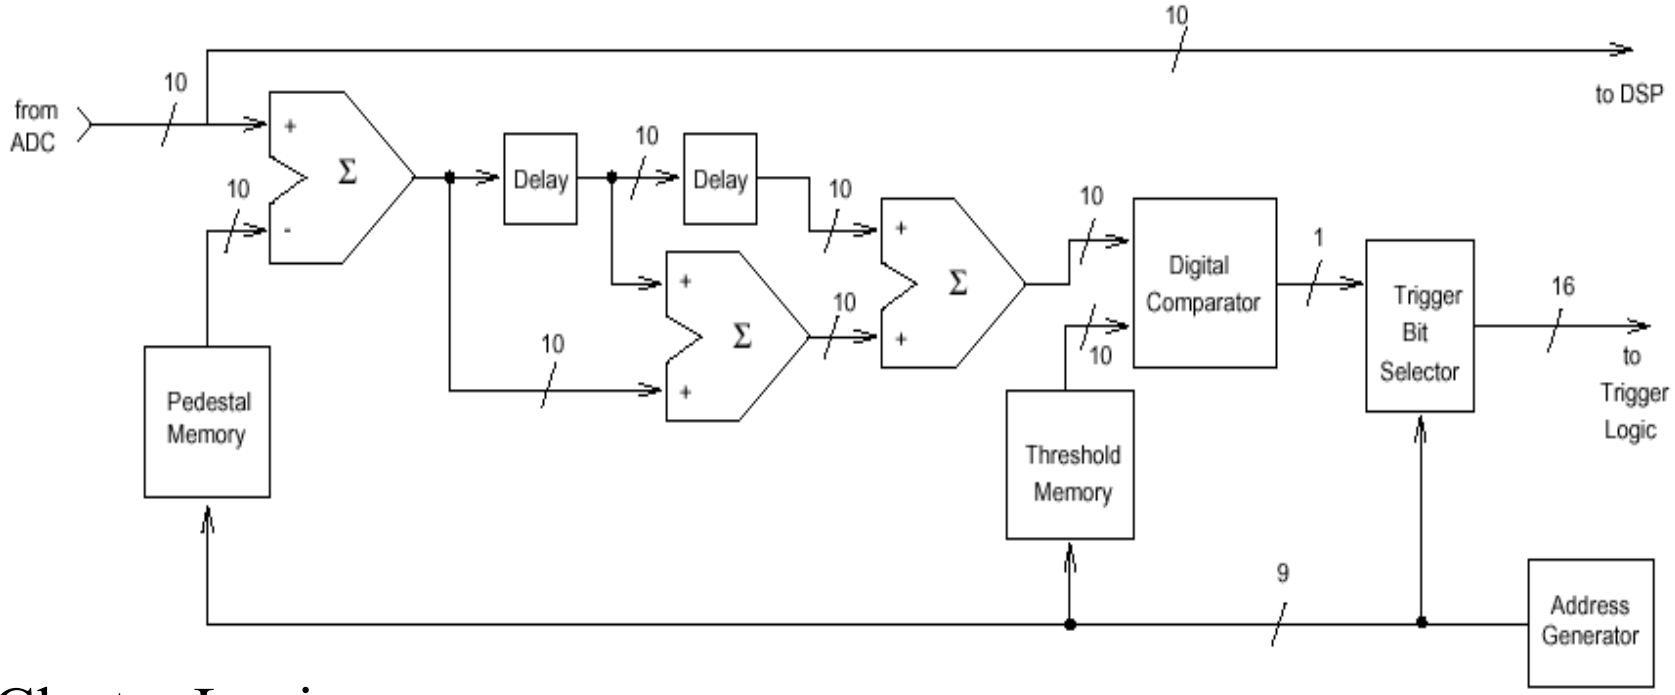

Cluster Logic

A Potential Level 1.5 Drawback

- Level 1 deadtime
- In the current scheme, we read 4 VA1s in series. The scan rate is 200 ns per channel, so the total time is:

$$T_{\text{dead}} = 200\,\text{ns} \times 4\,\text{VA1s} \times 128\,\text{channels} = 102\,\mu\text{s}$$

- The simplest way to mitigate this problem is to add extra ADCs so that fewer VA1s are read in series.

. . . A Potential Level 1.5 Drawback

This requires extra signal lines from the hybrid.

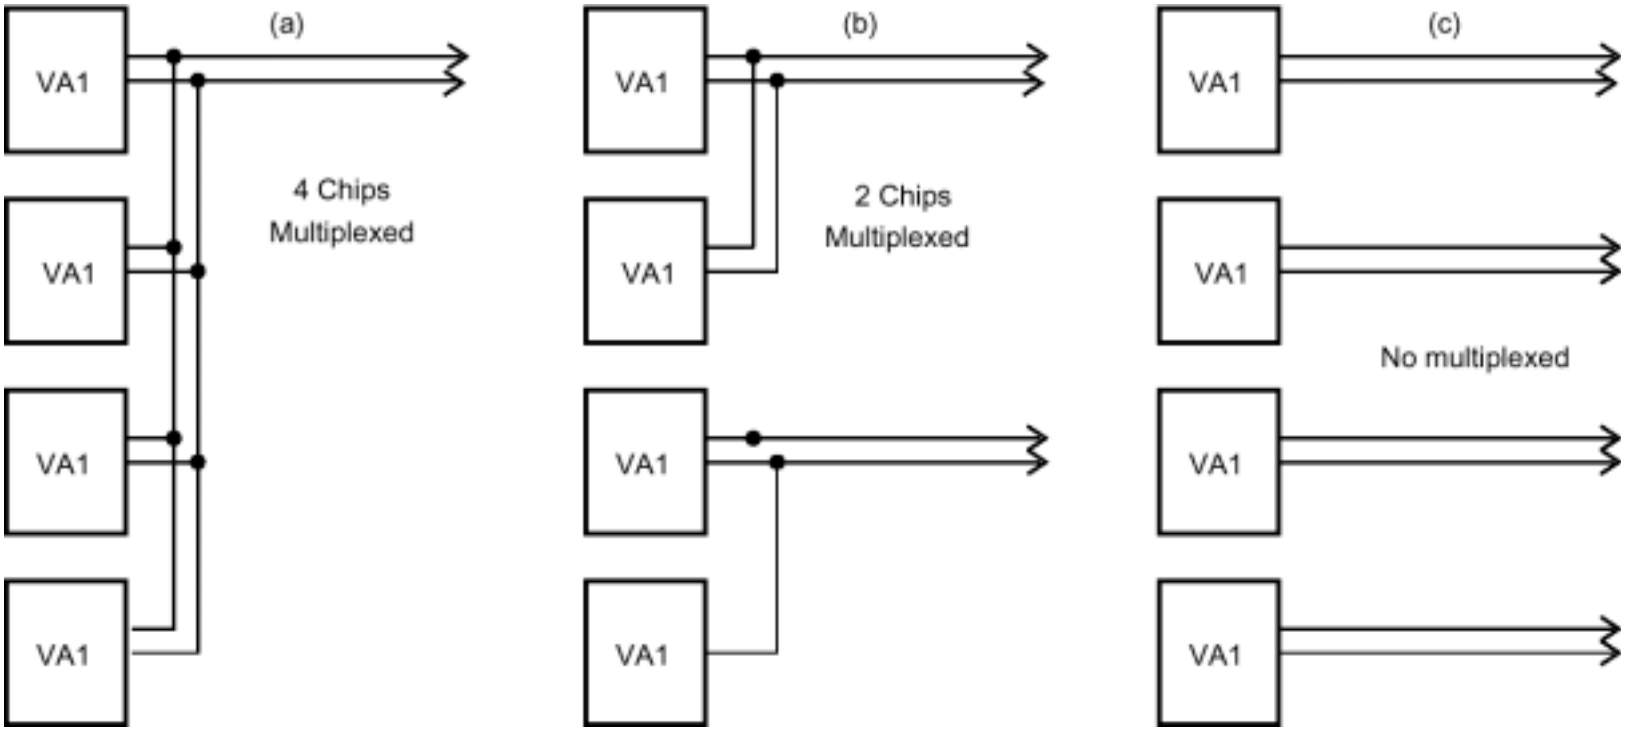

Two lines	Four Lines	Eight Lines
10% @ 1 kHz	5% @ 1 kHz	2.5% @ 1kHz

. . . A Potential Level 1.5 Drawback

- One also needs to worry about the links between the repeater system and the Halny (FADC) boards. Our new Vienna collaborators have suggested the use of fairly standard multi-conductor twisted-pair cable, which appears to meet our requirements. This cable is suitably compact and reasonably inexpensive.
- The main remaining issue is space on the repeater boards. The Cracow group is investigating this.

The Trigger Logic

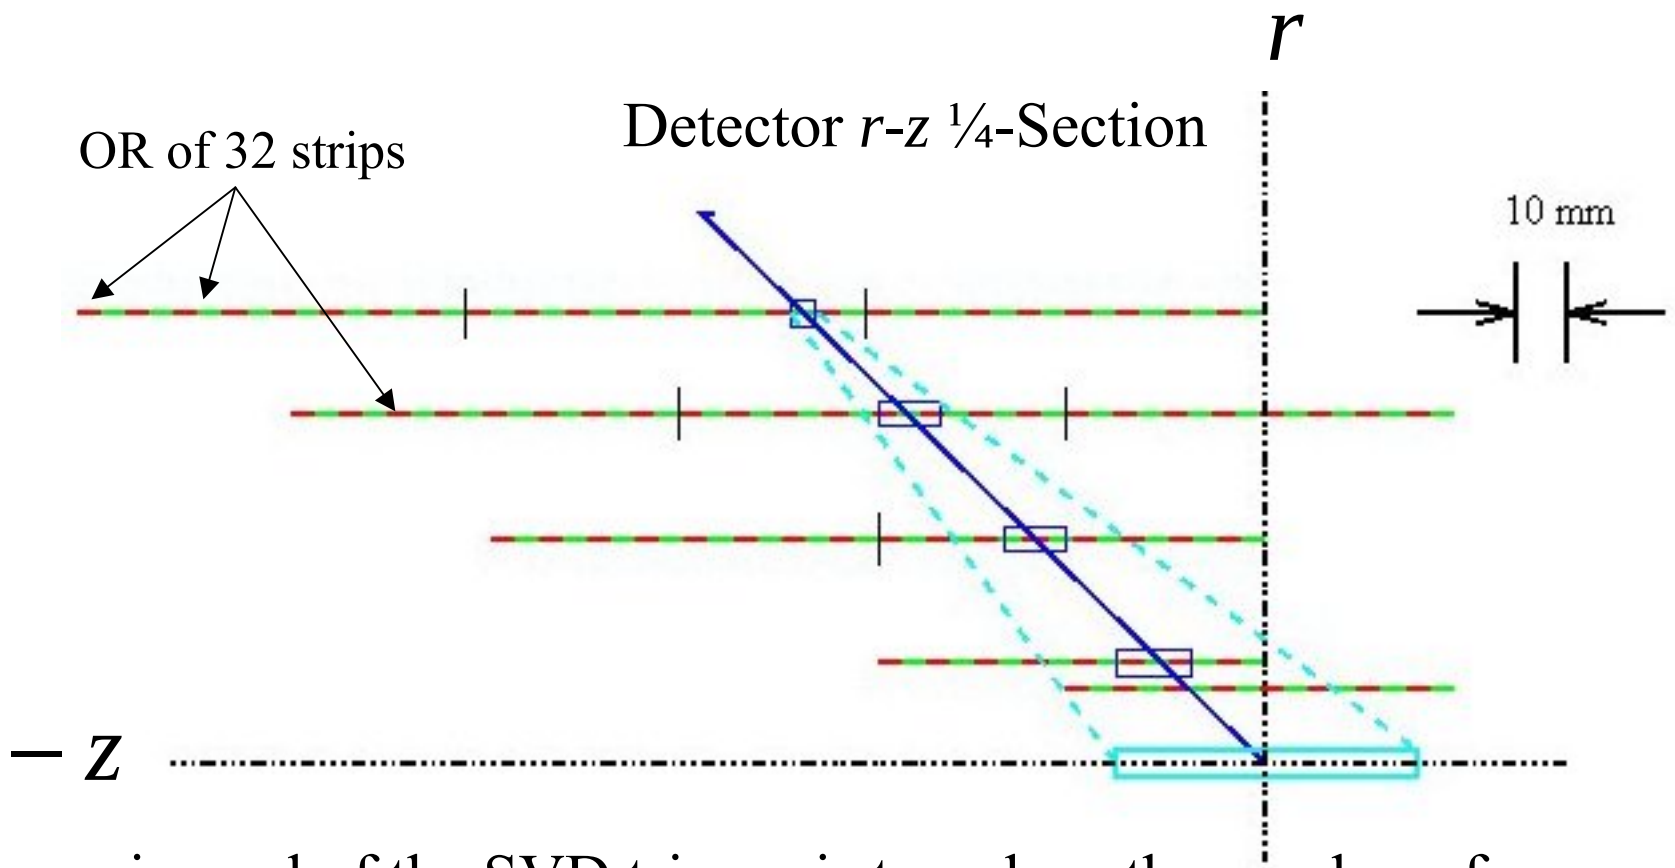

The main goal of the SVD trigger is to reduce the number of spurious triggers resulting from beam-gas interactions by requiring that events have one or more tracks emanating from within a few cm of z=0.

. . . A Possible Implementation

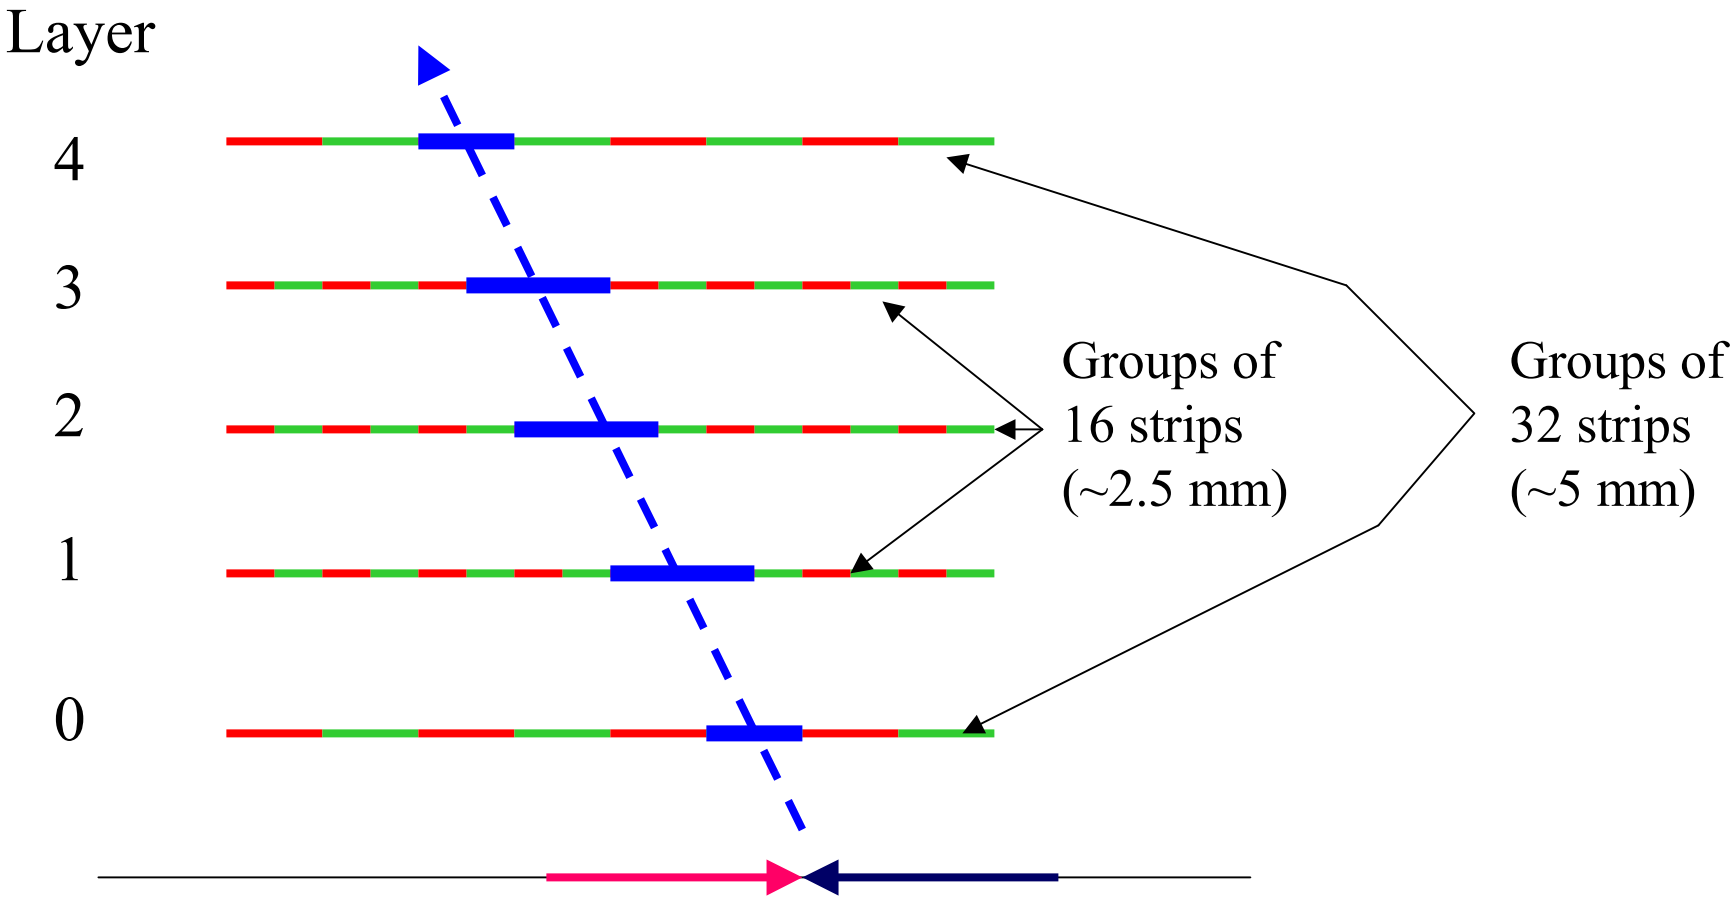

The basic idea is simple: one requires that there be a combination of hits on the five layers that form a straight line, which points in a direction consistent with the IP (z=0).

. . . A Possible Implementation

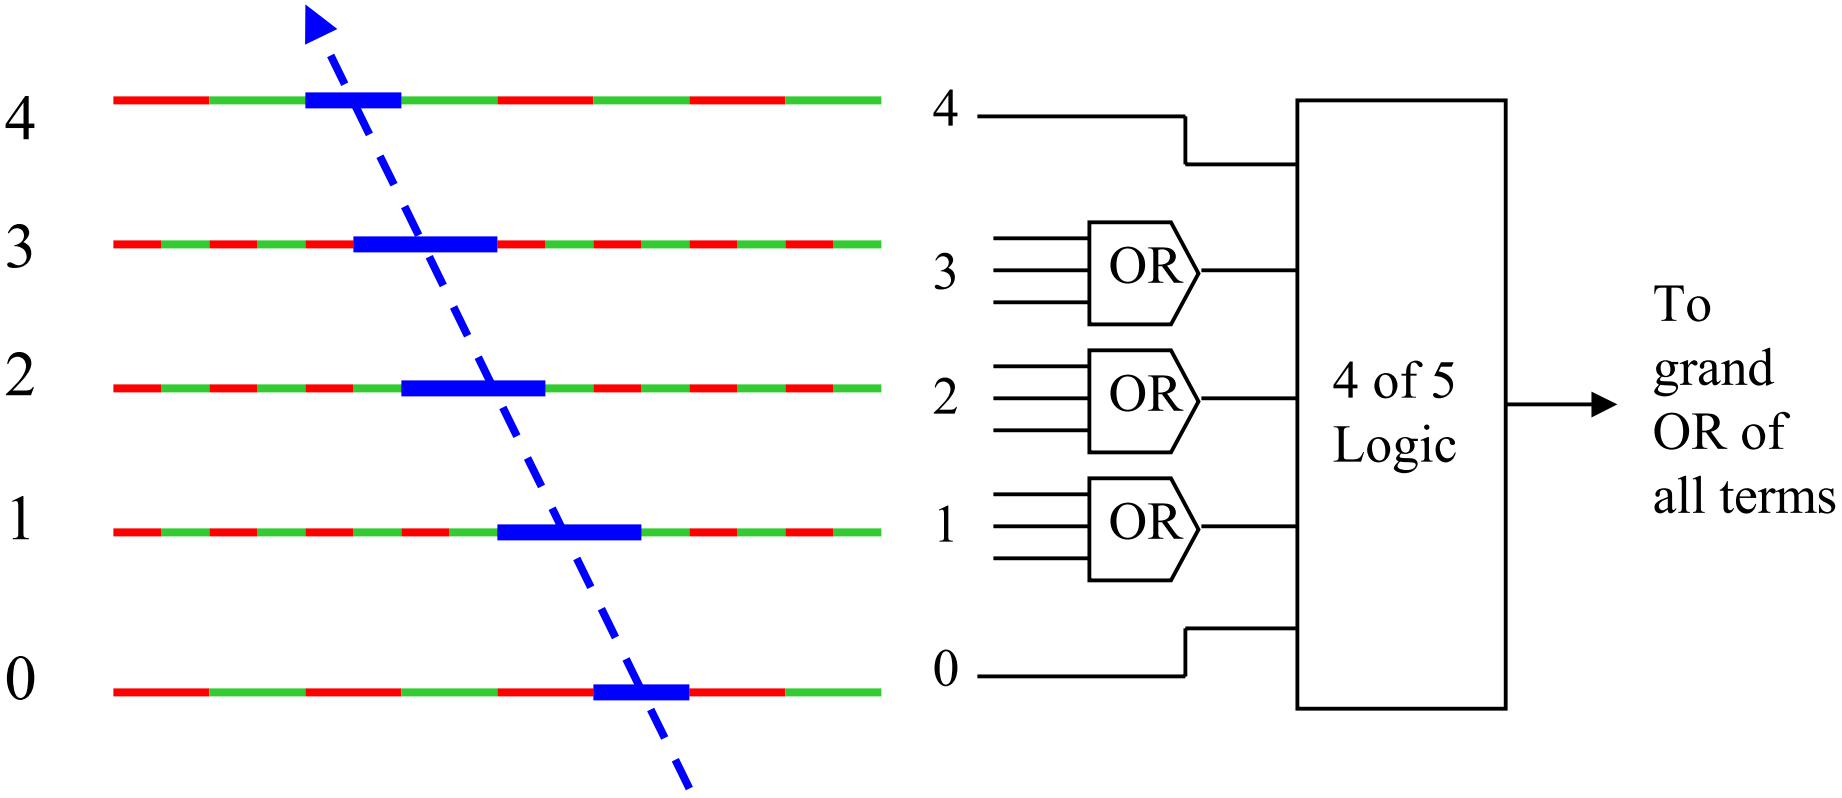

$$\#\,\text{terms} = \#\,\text{inner segments} \times \#\,\text{outer segments}$$

$$N_{\text{terms}} \cong 16 \times 32 = 512$$

Each term can be implemented in two Xilinx CLB's.

Other Considerations

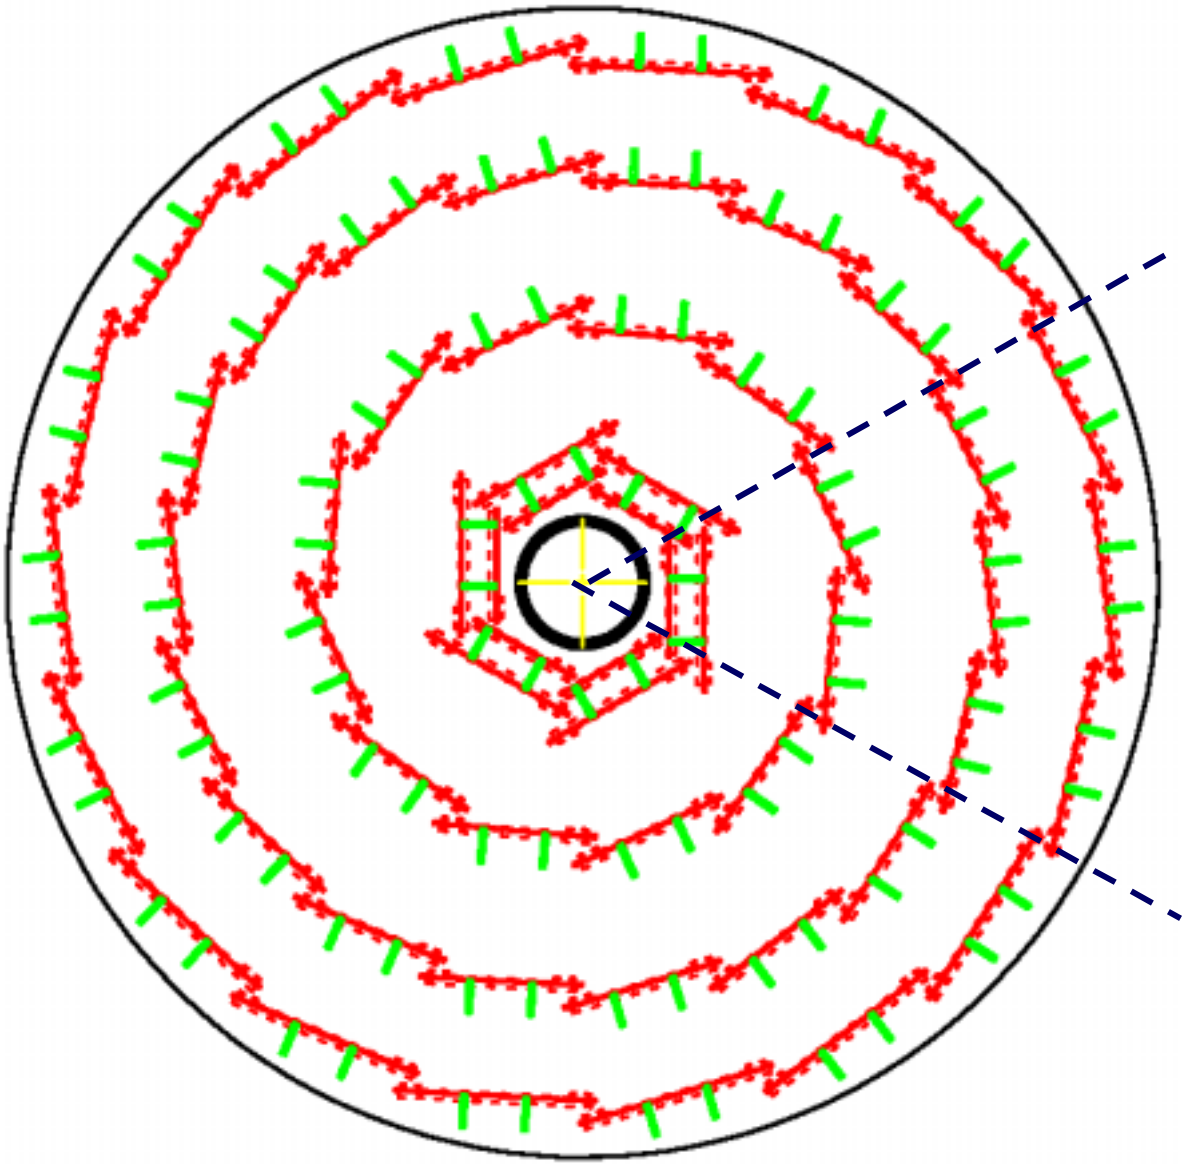

- In azimuth, the layout exhibits a six-fold symmetry.
- To handle one hextant's worth of DSSD's, a board would need to accept about 660 trigger bits.
- In a Level 1.5 scheme this is not a serious obstacle since the trigger bits can be transferred serially.

Serial Data Transfer

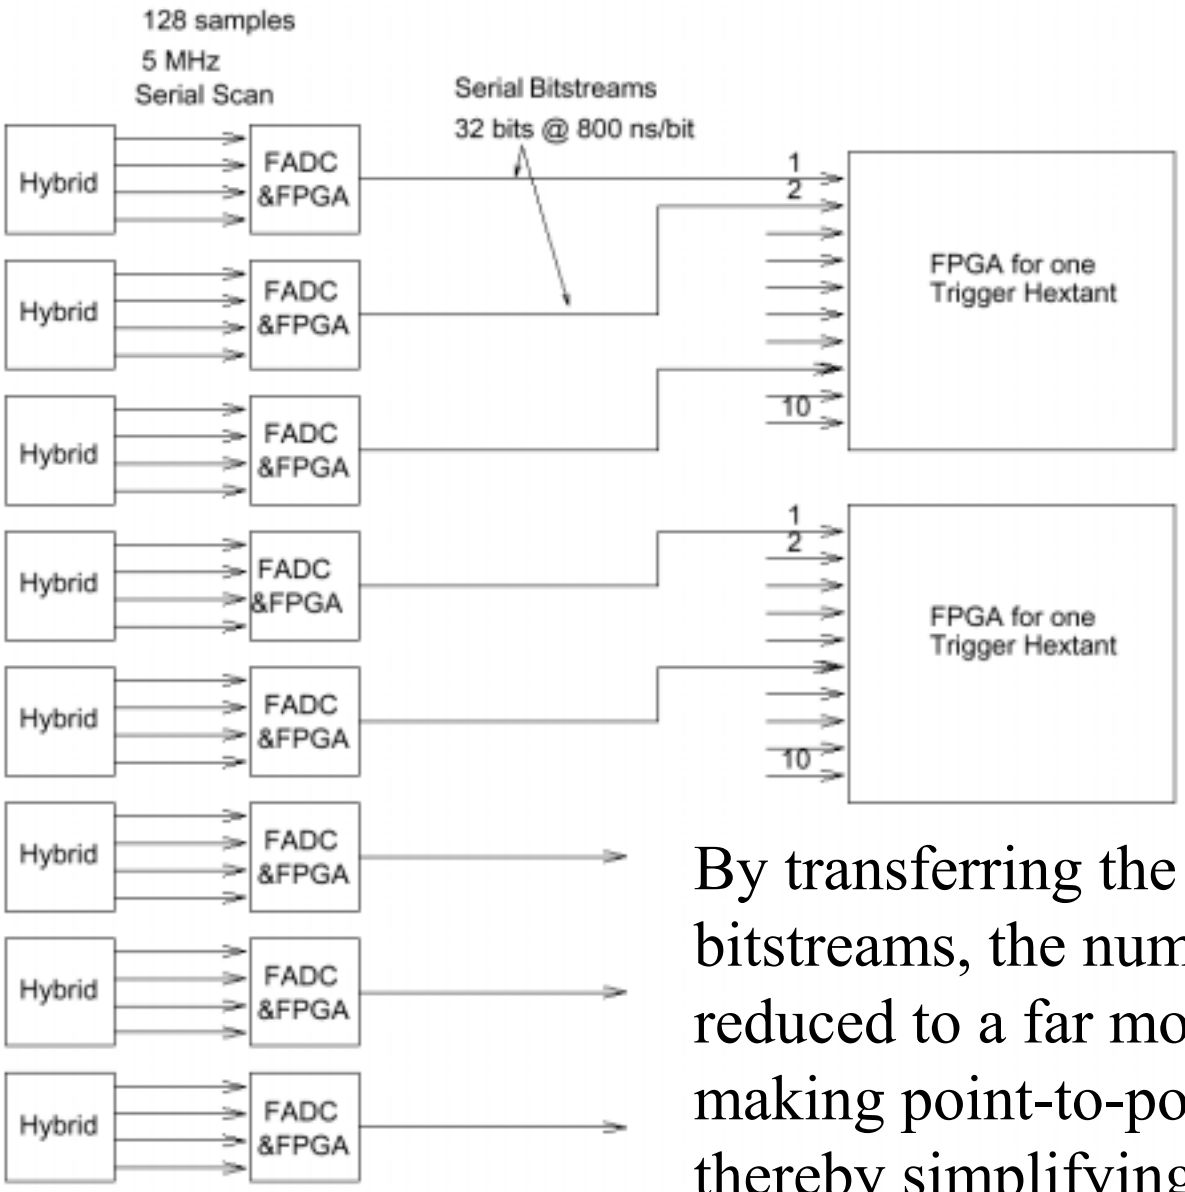

All together there are about 8000 Level 1.5 trigger elements, a potential wiring nightmare.

By transferring the data as serial bitstreams, the number of interconnects is reduced to a far more manageable 256, making point-to-point wiring possible, thereby simplifying system layout.

. . . A Possible Implementation

XILINX®

Virtex™ 2.5 V
Field Programmable Gate Arrays

DS003 (v2.4) October 6, 2000

Final Product Specification

Features

- Fast, high-density Field-Programmable Gate Arrays
  - Densities from 50k to 1M system gates
  - System performance up to 200 MHz
  - 66-MHz PCI Compliant
  - Hot-swappable for Compact PCI
- Multi-standard SelectIO™ interfaces
  - 16 high-performance interface standards
  - Connects directly to ZBTRAM devices
- Built-in clock-management circuitry
- Supported by FPGA Foundation™ and Alliance Development Systems
  - Complete support for Unified Libraries, Relationally Placed Macros, and Design Manager
  - Wide selection of PC and workstation platforms
- SRAM-based in-system configuration
  - Unlimited re-programmability
  - Four programming modes
- 0.22 µm 5-layer metal process
- 100% factory tested

. . . A Possible Implementation

Device	System Gates	CLB Array	Logic Cells	Maximum Available I/O	Block RAM Bits	Maximum SelectRAM+™ Bits
XCV50	57,906	16x24	1,728	180	32,768	24,576
XCV100	108,904	20x30	2,700	180	40,960	38,400
XCV150	164,674	24x36	3,888	260	49,152	55,296
XCV200	236,666	28x42	5,292	284	57,344	75,264
XCV300	322,970	32x48	6,912	316	65,536	98,304
XCV400	468,252	40x60	10,800	404	81,920	153,600
XCV600	661,111	48x72	15,552	512	98,304	221,184
XCV800	888,439	56x84	21,168	512	114,688	301,056
XCV1000	1,124,022	64x96	27,648	512	131,072	393,216

Outstanding Issues

- VA1TA evaluation.
- Hybrid design: no problem expected, but we are starting almost from scratch. One concern is higher power dissipation of VA1TA.
- Repeater board layout. We have added a lot of circuitry, which will make it a tight fit. Fortunately, some circuits (e.g., the bias generators) have been eliminated.
- Detailed trigger simulation and design. A lot of work remains to be done here.