

SVX4 Design Report
Henrik von der Lippe, Ray Yarema
November 2, 2001

The second SVX4 design review was conducted at LBNL on October 22. We agree with the review committee that a great deal of progress has been made since the first review that took place in April of this year. In particular, a 64 channel preamp/pipeline test chip was designed, fabricated, irradiated and tested. The test results from that chip permitted selection of a preamplifier design and verification of the pipeline design as well as calling attention to a couple of unknown fabrication and procedural issues. The success of the test chip significantly enhances the chances for a successful SVX4 engineering run. Also since the last review, the back end design was completed except for the I/O drivers. Currently, simulations are being run on increasingly larger portions of the backend design. The present plan calls for 2 more weeks of simulations to complete the backend, followed by 2 weeks of top-level simulations, and 2 weeks of top-level layout, which will result in two full chip designs ready for fabrication on December 7. It is our recommendation that the submission be an engineering run.

Next submission cost analysis

The preamp and pipeline sections of the chip were prototyped at a cost of \$18K. The next chip submission must test the backend portion of the chip, preferably with the front end. Without the front end, testing of the backend would be more difficult and incomplete. A backend chip made on a multi-project run would normally cost \$57K but since it should be built on an epi substrate, the cost would be \$87K (100 chips). A complete SVX4 chip built on epi substrate would cost \$110K (100 chips). The current plan is to build two different versions of the chip to evaluate on chip capacitor bypassing and power routing options. If only one design was submitted, a choice would have to be made between the designs. The potential liability of choosing the wrong design would be a loss of probably 4 months in getting ready for a production submission. If both designs were submitted to a multi-project run, the cost would be \$189K for 100 chips of each design. If both designs were submitted to a dedicated engineering run, the cost would be \$200K. We would get 10 eight inch wafers, which should provide about 3300 chips (1650 of each design). This would provide all the mechanical samples needed as well as all the chips needed for both CDF and DO testing. To expedite the project, the submission of both chips to an engineering run is the best approach. The value of the engineering run needs to be weighed against the risks associated with design and the level of verification, which can be done as described in the next section.

Risk analysis

The simulation tools available today significantly reduce the risk of failure in chip design. At some point however there is a limit to what the tools can do. In the ideal situation, an entire circuit would be modeled with all resistance, capacitance, and inductance, as well as active devices, and tested within the specified operating parameters at the interface to

the system (bond pads) before submission. This ideal, however, cannot be achieved due to several limitations.

General Limitations:

- ❑ Under actual operation, states may persist for milliseconds or even seconds, which are periods of time impractical for simulation
- ❑ The entire design cannot be extracted and simulated with all electrical elements, to the accuracy of a full analog simulation (ELDO/HSPICE)
- ❑ Fabrication guidelines are not completely specified or checked by design tools

Because of the above limitations, it is desirable to strike a balance between prototyping, full verification, and schedule. In SVX4, we have implemented the following plans to balance the multitude of conflicting factors:

- ❑ While not all states are testable, we have developed 23 test vectors that verify the operation of the digital readout circuit. Some vectors include unexpected signal timing for analysis of the most likely problems. These vectors will be run on the final circuit extracted from the layout.
- ❑ We cannot extract resistance or inductance, but all capacitance and active device parameters will be extracted.
- ❑ Simulators with limited accuracy are now available (TMILL) to simulate the whole chip. Our plan is to run TMILL on the entire design to functionally check its operation.
- ❑ Layout verification (DRC and LVS) is done by several tools, some better than others; we intend to use the foundry specified program to ensure the fabrication parameters are met.

Full-chip simulation and foundry specified DRC will be used to mitigate the risks of failure in the backend design, and in the full-IC integration. The successful fabrication and testing of the preamp-pipeline prototype IC significantly reduces the risk contribution of this module to the final design. We believe the steps detailed here result in a reasonable level of assurance and risk for an engineering fabrication run.

Manpower

- ❑ Brad Krieger will be working on this project full-time until it is finished.
- ❑ Henrik von der Lippe will be available approx 25% time through the end of the project.
- ❑ Emanuele Mandelli, an experienced designer who has been working on Atlas Pixels, will be working on this project for at least 3 weeks full-time in November. Emanuele's responsibilities on the Atlas Pixel chip have been mixed mode design and top-level simulation using TimeMill.
- ❑ From Nov 3 through Dec 7, the total available engineering manpower dedicated to the project at LBNL will be at least 10 man-weeks.
- ❑ At the present time little additional effort is expected from the front-end designers at Fermilab.

- We expect to be able to submit the chip by Dec 7, leaving 2 weeks before the Christmas holidays if additional work is needed. The fabrication time is estimated at 8 weeks meaning chips will be available in Feb '02.

Response to Review Committee report

The review committee did an excellent job of providing input to the SVX4 design group. Several issues were raised, some of which were due to the lack of materials available to the review committee. In this section we respond to those issues with additional information or plans to address the issues raised.

Response to Power Issue item 2:

The Front End analog supply (AVDD) current is measured to be 23 mA for 48 channels, with the input transistor bias programmed to be 250 uA per channel. This translates to approximately 0.5 mA/channel, 64 mA/chip, or 160 mW/chip. The Front End digital supply (DVDD) average current draw is approximately 2 mA for the complete chip with a peak transient current around 20 mA (the test chip shows this transient not a problem). The total front end dissipation is thus measured to be 165 mW. The analog supply current for the ADC is designed to be 5 mA for 128 channels including the ramp generator (12.5 mW). The power consumption of the digital backend is 23 mW during digitize and 94 mW to 4mW during readout (read all mode).

Response to Power Issue item 3:

The IBM design manual specifies a special cell for thinox decoupling, which consists of an NMOS in an N WELL, with a series fuse and a series control FET that can be gated off. The aim is to make a decoupling cell with low series resistance and thinox capacitance from VDD to GND that has built in protection against shorts through the thinox. The narrow metal fuse is intended to blow if a short develops, removing the cell from across the power supply. If the fuse does not blow, the control gate of the series FET can be disabled to remove the cell (this requires extra detection logic). Why is this cell architecture recommended? Presumably all forms of thinox are equally vulnerable to shorting. This configuration is probably the simplest, most straightforward and area efficient way of implementing the above mentioned functionality, including short protection. The capacitor is in an N WELL so it can be isolated from the substrate (GND). In the N WELL, an NMOS cap is the most area efficient and puts the device in accumulation instead of inversion. The series FET can be easily integrated as part of the N WELL contact. To summarize, this cell provides a decoupling cap with low series R and built in ways of removing the cell if it is shorted. Since it is a general purpose decoupling cell, used in a variety of circuits, it must have low series R.

There is no analogous documentation from TSMC that addresses how to design on chip decoupling capacitors. We have, however, corresponded with several designers in industry who are doing on chip decoupling effectively using "standard devices" and are

not using the “IBM” approach. In addition, we wanted to make a cell that was easily constructed using available libraries and parametric cells. Therefore, we have taken a somewhat different approach to decoupling cap design. This approach uses thinox capacitance in an NWELL (using a standard PMOS device) and has a different method of short protection. Since we know exactly how many decoupling cells we will place on our chip and the total achieved decoupling capacitance, we do not need or want a “general purpose” cell with low series resistance. In fact, we desire a relatively high series resistance in each cell! This is because we are using the cells to decouple an analog supply that has extremely low resistance bussing across the chip (much less than 1 ohm). A very low ESR capacitor would then give a high Q circuit, which is undesirable. Since we know our total cap value, we can put an appropriate value of series R (several Kohms) in each cell to insure low Q and still have low enough equivalent R to provide adequate decoupling. This has the added benefit of also supplying protection from shorts. A shorted cell does not remove itself, it simply draws DC current from the supply (approx. 1 mA per shorted cell -- not significant unless many tens of cells are shorted). Adding a fusible link would not be effective since it would not open under the low fault current condition. In summary, we have developed our own decoupling cell that uses thinox decoupling (as does the IBM cell) and has protection against drawing large supply currents in case of shorts. The only serious failure mode would be if more than many tens of cells were shorted. This should only occur if there are serious problems with the run, in which case the decoupling cap problems are the least of our worries. It should be pointed out that no failed bypass caps were observed on the preamp/pipeline test chip even after the power supply rails were zapped with an ESD tester.

The bypass capacitors for the preamplifier are located in the pipeline area. Also, the power and ground routing is done with very low resistance grids. A short in one of the capacitors in the bypass array will not significantly affect the voltage applied the preamplifier cells.

Response to Preamp issue item 1:

The preamp/pipeline test chip has been operated at different AVDD supply voltages (2.5, 2.25, and 2.75). The main effect on the preamp of lowering the supply voltage is a change in the rise-time and reduction in the dynamic range.

<u>AVDD</u>	<u>Rise-time</u>	<u>Total range</u>
2.25V	55 ns	1.00V (220 fC)
2.5V	46 ns	1.16V (260 fC)
2.75V	40 ns	1.33V (295 fC)

The change in rise-time is not a problem since it is adjusted anyway by using the programmable BW bits. The dynamic range is sufficient in all cases.

Response to Preamp issue item 2:

The reviewers expressed some concern about the biasing of the preamplifier due to the number of stacked devices. However, the situation is not as precarious as it may appear. The only critical drain-source voltage is really that of the input transistor, M0. If the V_{ds} of M0 goes below V_{sat} , its g_{ds} goes up dramatically and the open loop gain is much reduced. (M0's V_{sat} is approx. 50 mV). V_{ds} on M0 is set to around 200 mV and is determined by the gate voltage on M2 (this bias voltage is called NCAS) and the V_{gs} of M2. M2 is the cascode device for M0. NCAS is generated on chip in the preamp bias circuit so that over process and supply variations, NCAS is always high enough that M0 is correctly biased. Measurements on the test chip show that NCAS can be forced lower by up to 150 mV before causing any significant effect.

M2 is sure to have enough V_{ds} since its source is at about 200 mV, and its drain voltage is set by the V_{gs} of M43, which will be about 500 mV.

The V_{ds} of the PMOS current source, M3, is set by the V_{gs} of M43 (about -500 mV). Thus M3 is well into saturation.

The V_{ds} of the current source cascode, M18, depends completely on the magnitude of the charge injection at preamp reset release which is used to set the preamp output to the "top" of its range. The magnitude of this charge injection is set by a capacitor (which is ratioed to the preamp feedback capacitor) and a voltage above ground, V_{rset} , which is generated in the preamp bias circuit on chip. V_{rset} is generated in such a way that a reduction in AV_{DD} will cause a corresponding appropriate reduction in V_{rset} . V_{rset} is also affected by process variations, so that for example, if M43 (PMOS) V_{gs} has larger than nominal magnitude, V_{rset} will be reduced so that M18 will not be in the linear region. Of course, the V_{rset} level is designed so that after reset release, M18 is not right on the edge of saturation, but has a comfortable margin of at least 100 mV. As a safety precaution, V_{rset} is brought out as an external pad so that it can be overridden if necessary.

The final device of concern, M4, goes linear when the preamp output goes low enough, that is, when the preamp has accumulated enough total charge. This is exactly what sets the dynamic range of the preamp.

Simulations show that over all process corners and supply voltages, the preamp resets to an appropriate voltage near the top of its range, which gives 200 fC or more of range before M4 goes linear.

Response to Pipeline issue item 1

The following changes were made to the SVX3 pipeline control logic:

1. Changed to enclosed 0.25 u design, SEU considered but no special care deemed necessary.
2. Put previously split multiplexer entirely outside cell

3. FIFO made to move in one direction to correct timing problem (SVX3 address FIFO too fast in 0.25um process) and correct logic failure if more than 4 level 1 accepts occur.
4. Small state-machine added to control readout.

The preamp/pipeline test chip verified the functionality of the changes. Four corner simulations including radiation shifts and lower power supply voltage were run on the pipeline logic along with pipeline cells using ELDO to verify operation over all process and operating conditions. These results were done as part of the design but not saved.

Response to ADC issue item 1:

In general, the ADC is in good shape. The design review pointed out a few simulations Jean-Pierre should do to determine the sensitivity of the ADC to non-ideal process variations mainly within one device, but also globally. Some of these simulations have already been done, and no surprises are expected. Jean-Pierre will be available to perform these simulations and any other tasks required to finish the ADC design. He will present the results at an upcoming SVX4 design meeting. A list of items that will be investigated is given below. It is probably about 1 week of work to complete and present.

- Simulations will be done to determine the effect of channel-to-channel V_t variations on the comparator delay
- Simulations will be done to determine the effect of channel-to-channel variations in the comparator hysteresis cap on the comparator delay
- Simulations will be done to determine the effect of channel-to-channel variations in the comparator dynamic threshold cap on the comparator delay
- Simulations will be done to determine the effect of absolute value variations in the comparator dynamic threshold cap on the comparator delay

Response to Backend Digital Core item 1

In general, the digital core circuits are finished at a preliminary level in both design and layout. However, one bug has to be fixed (FIFO data output glitch), and it is clear that considerable verification is required vis-à-vis TimeMill to validate the performance of the circuit. Assuming the design was frozen, this would require approximately one hundred 1.5-hour simulations based on review committee recommendations, in order to explore process and power-supply variations for all 23 test vectors. These simulations should be done post-layout, unless the simulation time increases dramatically (experience of others has been that parasitics don't increase TimeMill simulation time that much). Before this can be done, the circuit needs to be debugged and simulated at the schematic level with several vectors and process/power supply/clock speed corners, and all the layouts finished. When the design appears stable at the schematic level, a description of the operations will be given at a Wednesday meeting, with initial simulation results.

The post-layout "digital core" simulations will cover the total function of the readout FIFO/counter and interface control (IO). Further limited simulations will check the full

chip. For the digital core only, the anticipated tasks are given below. Four man-weeks of total work are estimated.

- ❑ Debug and finalize the schematic, simulate some test vectors at FF-SS/2.25-2.75v/53-65MHz
- ❑ Present readout circuit operation and tests
- ❑ Iterate the FIFO layout
- ❑ Set-up and run the IO layout
- ❑ Wire FIFO to counter and to IO
- ❑ Convert layout, add CMP fill all layers, extract, DRC, LVS in TSMC design kit
- ❑ Run large set of digital core simulations over all corners, include some D0-mode test vectors