

IT ALL STARTED WITH THE MICROPLEX

As Sir Alexander Fleming once said,

"It is the lone worker who makes the first advance in a subject: the details may be worked out by a team, but the prime idea is due to the enterprise, thought and perception of an individual"

I feel that Sherwood is one such individual. He has had the flashes of insight to recognize what is possible and what to do to make the possible happen in a logical way. One good example of that is the Microplex chip for use with semiconductor microstrip detectors.

The start of this adventure began with Sherwood and Bernard Hyams visiting me at Stanford University in July 1982. They brought with them an example of the silicon microstrip detector which Bernard had created. (slide 2) This device had great potential for use as a nuclear particle detector giving very nice physical dimensional readout, but had the problem that it had an untenable number of signal output connections which had to be dealt with. (slide 3). Bernard brought the detector to a presentation which Sherwood attended, and during the presentation Sherwood figured out how to handle the critical aspects of connecting and reading out the large number of detector elements.

After the talk, Sherwood talked to Bernard, who agreed to collaborate with Sherwood on development of a detector system with an integrated readout chip. But they had a big problem, the readout chip had to be created for the purpose, and it had very exacting requirements. To make the system feasible, a single chip had to have 128 signal inputs, with each input going to a separate channel incorporating a high gain charge integration amplifier, an output sample and hold circuit, a differential analog multiplexer, and a shift register. (slide 4) The charge amplifier needed to have high gain for accuracy, wide bandwidth, negligible input current, and low noise. Additional channel circuits were needed for automatically zeroing the charge integrator and providing the ability for self-test and gain calibration.

Consider how big these requirements looked when compared to the discrete circuitry used by Bernard for the first feasibility tests. The discrete circuits could never be sufficiently compact to permit full instrumentation of a vertex detector and they required a lot of power. Sherwood recognized that if some kind of compact charge amplifier could be made in integrated form, these obstacles could be overcome. But a tremendous amount of effort would be required to achieve this goal.

Sherwood was blessed with being close to the excellent facilities of Stanford University, where there were a number of faculty, research associates, and graduate students working on advancing the state of the art for semiconductor device design and fabrication. I was one of those research associates, being a newly minted doctorate from 1978. One of my main efforts by 1982 had evolved to running the Applicon AGS/860 VLSI design system and using it to design ICs for process development. (slide 5) I was skilled in operating this machine and looking for a challenge. We had at that time one black and white 19 inch raster display terminal, resolution 672 by 504. It was understandably hard to tell exactly what layers you were editing, since they were discriminated only by the dot pattern. That was to cause a problem later during chip development.

So Sherwood and Bernard met with me at Stanford, and we sat around in the Applicon room to discuss the needs of the silicon microstrip detector and the capabilities of Stanford University and the Center for Integrated Systems, which was just in formation then. I was interested in a new challenge, and this certainly would bring one. After I spent some time thinking about possible solutions for the IC design, Sherwood came back and we met again.

Stanford had at that time just done a good job of replicating the current 'hot' integrated circuit process, the all N channel metal oxide silicon or NMOS process which made the HP35 pocket calculator and the Z80 microprocessor possible. This process was a silicon gate transistor process with only two device types: (slide 6) an enhancement transistor with a conduction threshold of +0.8 volts, and a depletion transistor with a conduction threshold of -2.5 volts. The distinction was made by an extra dopant implant in the channel region of the depletion transistors. All

transistors were designed for operation over a 0 to 5 volt range for input and output. Since this was a silicon gate process, the diffusions were aligned to the gates, reducing parasitic capacitances and size. These features were critical to making the chip which Sherwood desired.

When these two types of transistors were combined to make a rudimentary amplifier stage (slide 7), the silicon properties dictated that a stage gain of about 7 would result. Therefore one stage did not have enough gain to make a charge integrating amplifier. Since inverting gain was needed for charge feedback, this dictated a three stage amplifier. This was a sure way to make an oscillator.

Three stages would get the open loop gain up to 340 or more, which was enough to overcome the loading due to the source capacitance of the silicon microstrip detector. But stability of this amplifier would be difficult to guarantee. I was skilled at doing the layout, and could see how to get the amplifier to fit on a chip, but circuit simulation capabilities at that time were rudimentary at best. I did all my designs using only hand calculations (slide 8).

One day Sherwood came in to discuss the project status and see what advances I had made in geometric arrangements, when he decided that the project was not going to get enough progress without more direct involvement by him. So he made the 'mistake' of asking if he could help to move the project along. By then I realized that being able to verify that we had any chance of the chip working correctly was going to take a lot more time than I had. Modern automatic verification software did not exist in those days. The layout and other responsibilities were consuming all of my time.

Sherwood didn't know that he had stepped into a tar pit, and it would be some time before he escaped. He was able to lay hands on a rudimentary SPICE type program which took text files as input, ran in batch mode, and gave line printer graph outputs. He later got some real plotting capability. Using that with some simplified device models, he was able to simulate the detailed operation of our amplifier channel. He was surprised to discover one day almost by accident that we were making a 128 channel

oscillator. That might have had some uses, but not for nuclear physics!

As an engineer, I was happy to make approximations to many things to expedite calculations, but Sherwood was not. His diligence probably saved the project from a long and painful birth. Being a physicist, Sherwood knew that everything was calculable, if you simply had sufficient data and knew the formulas.

Do you remember Teledeltos paper? It is black conductive stuff originally used in FAX machines created in 1939. Sherwood got the modern equivalent of it and made little cutouts of several critical diffusions in the layout. (show piece of paper and a cutout) In this way he was able to get exact numbers for the parasitic resistances of several parts of the layout. These numbers turned out to be essential to getting correct answers for the circuit performance.

Sherwood also measured all the various layer dimensions and calculated the parasitic capacitances for the SPICE program input. In this way were able to get a very detailed schematic for analysis (slide 9). I hate to say it, but now software which I use every day can just from a few keystrokes extract all the transistors, passives, and parasitics components from a layout in seconds. It even outputs a netlist for the SPICE simulator, which is integrated so that a simulation can be run right away.

Some of the later results for the circuit simulation are shown here (slides 10 & 11). We ended up with a fast amplifier that was stable and suited to charge readout.

Along about this time, Sherwood told me that he had finally found a solution to the problem of being able to reliably stop his bicycle when going down the Berkeley hills in rain. (show slide12, brake). This was characteristic of his dogged pursuit of a problem until it was solved.

I looked around my office this week, and found a microplex revision 2 chip used as a paperweight (slide 13). Still existing after 30 years! So I made some photos to show how it came out. Ignore all the dirt on the chip, I was unable to clean it all off. You can see the overall design

arrangement, with the 128 signal input pads in 4 rows at the top, spaced 40 microns center to center. This arrangement was one of the critical ideas which Sherwood and Bernard created. It made possible connection from the large number of detector outputs to the individual amplifier channels. The process basically uses 5 microns as the minimum width of a layer, and the channels are laid out for a total width of just 7 times that, at 35 microns. So the circuitry of one channel is very long and thin. This causes many problems with making connections and parasitics.

The signal flows down vertically from the top to the bottom, where a 128 stage shift register operates an analog multiplexer to output the signals from the channels as a serial analog string. For the following pictures the chip is rotated counterclockwise so that the inputs are at the left. A closeup of the central portion of the chip is shown here (slide 14), with the inputs at the left and the outputs at the right. You can see that one of the design concerns was how to provide power to all the channels, so there are heavy metal lines running vertically across the channels for +5 volts and ground, and smaller ones for signals. Note that in this process, there are only three layers available for interconnection: diffusion, polysilicon, and one layer of aluminum metallization. The modern luxury of 5 (or many more) metal layers did not exist.

Here is the input pad region of the chip (slide 15). Note that at the left side of the chip is another Sherwood and Bernard innovation, the capability for chip self-test and calibration. He was very prescient to recognize that when we tested these chips on a probe station, it would be essentially impossible to probe all 128 inputs. Also, the capacitance of the probes would have made useful results impossible. So by including a set of 4 interleaved calibration lines which drive small capacitors, a set of only 4 input pads could test all 128 channels for functionality. The pads are offset to facilitate both the on-chip wiring and the bonding to the silicon microstrip detector.

The first stage of the amplifier is made with a large amplifier to reduce the thermal noise (slide 16). Sherwood insisted on the importance of low input noise, and I calculated that we would have to use 1.5 milliamps per channel in the first stage to get a reasonable result. We ended up with

a chip which consumed a total of 0.3 amps at 5 volts when running, but in its application, the voltage could be pulsed to reduce average power. This is not low power circuitry! Later CMOS designs used less power, but their noise level was higher.

The feedback capacitor for the charge integrator is between the first and second stages (slide 17). It is a small value made with metal overlying a diffusion. To the right of it is the second stage, followed by another capacitor used for frequency compensation, and the final third stage. Optimization of the frequency compensation components was a critical job which Sherwood did with his SPICE simulations. Otherwise the chip would almost surely oscillated (times 128!)

Signals then go under a big power bus to the output sampler and hold capacitor (slide 18). After the hold capacitor, there is a differential analog multiplexer, which is in turn driven by the digital shift register to its right. The digital logic was done as a simple two phase clocked register so that there was no hazard of its not working. We were still very new to IC design.

A digital output driver (slide 19) permits sending the pulse which reaches the end of the shift register on to the input of the next chip in a daisy chain. This facilitates using arrays of chips for data collection.

All of this circuitry was developed using what was at that time a very new type of circuit operation for us, and we took pains to analyze it carefully and minimize hazards. We were aiming to get fully functional chips on the first process run, and actually did that, sort of (slide 20). Well, the results were very puzzling. Sherwood was testing the chips, and reported that some chips worked perfectly, whereas others would have a large section which did not work at one side or the other. This puzzled us for a few weeks, and eventually he realized that there might be a design fault. Looking at this slide (slide 21, showing closeup of buried contact), you can see that there are some regions where a direct connection is made between a diffusion and a polysilicon element. This required specification of an oxide cut and an extra diffusion implant at that point. Since the layout was done on the black and white terminal, the error in layer dot

pattern was not seen, and one of the necessary mask layers was missing here. After he figured this out, we checked the layout database and verified the problem. Making a corrected mask to add the diffusion fixed the problem.

Sometime later, Sherwood realized that the performance of the chip could be substantially improved for signal to noise ratio if we could cancel the $1/f$ noise inherent in MOS circuits. So he figured out a modification as seen on some schematics previously to add a second sampler per channel output. (slide 22, dual sampler) This allowed us to use true double correlated sampling with differential readout, and virtually eliminate $1/f$ noise. This change was made when the mask data was transferred to American Microcircuits Incorporated (AMI) in Sunnyvale for volume production of parts to be used at SLAC. The final revision 4 version of the chip mask definitions were fabricated in their radiation hardened NMOS process.

(slide 23, full chip again) Despite the visual complexity, you can appreciate that since the Microplex contained 128 identical channels, it is conceptually simple. I only had to do the layout of one channel, and design it in an exemplary manner with Sherwood doing the exacting circuit extraction and analysis, and provide the surrounding external connections. There is basically only one kind of circuit on the chip, but lots of them. The time required to do all of the surrounding layout is similar to the time used in optimizing the core channel layout. The overall project was first discussed on July 29, 1982, with layout started on September 29, 1982. Simulations were completed in September 1983, and fabrication at Stanford started in May 1984. The design was revised to include double correlated sampling and sent to AMI for final production as revision 4 in May, 1986.

That is the Microplex story as I saw it with Sherwood. He brought me quite an interesting adventure.