

Witnessing the Birth of VLSI Design

by Carlo H. Séquin

I have had several lucky breaks in my career. One of them was the opportunity to be immersed in the emergence of VLSI technology and its associated design methodology. I got my PhD in Experimental Physics from the University of Basel, Switzerland, in 1969. My first job was with Bell Telephone Laboratories in Murray Hill, New Jersey. Because of my thesis, in which I studied the behavior of *Interface States in Metal-Oxide-Semiconductor Field-Effect Transistors*, I was placed into Lab 225, which was engaged in building solid state imaging devices based on the brand new CCD (Charge-Coupled Device) technology, which had been conceived there a few months earlier. When I arrived, the group with Mike Tompsett and Gil Amelio had just demonstrated a CCD sensor array with 8 by 8 pixels. Since I have been enamored with geometry ever since high-school, I jumped to the opportunity to design the layouts of much larger imaging arrays, first with 128 by 128 pixels, and eventually (in 1973) with two interlaced fields of 256 scan lines, which was compatible with the American broadcast TV format. The latter device was way larger than any other IC chips of that time: It had $\frac{3}{4}$ of a million MOS electrodes placed in a rectangle measuring $\frac{1}{2}$ inch by $\frac{5}{8}$ of an inch (Figure 1).



Figure 1: *The first solid state image sensor compatible with the American broadcast TV format.*

The success of this chip brought me a job offer from the University of California. In 1976 I went to Berkeley as visiting lecturer, and in 1977 I became a tenured faculty member in the Computer Science Division in the EECS department. Although my background was mostly on the EE side, I was hired into the CS Division because there was an urgent need for expansion. Tom Everhart, then chair of EECS, explained to me that this would put me in a unique and important position. While I already had several acquaintances and friends on the EE side, it would be important for my career to establish close working relations with my CS colleagues. Having good connections in both Divisions should then allow me to help meld together those two Divisions, which at that time were not very congenial competitors. Naïvely I accepted this mission to bring those two factions together. Fortunately I got some unexpected help in this audacious task...

As soon as I was an official faculty member at Berkeley, and no longer a Bell Labs employee, I was contacted by Lynn Conway at Xerox PARC. She invited me to be a consultant in a project to develop a new methodology for the design of large and complex integrated circuits. After a single visit to Xerox PARC, I enthusiastically accepted this additional job offer. The decision was made easy by what I saw at that Lab: Interactive graphics-oriented computer stations were readily used by everybody; prototypes of powerful bit-mapped display panels implied a technological revolution just around the corner; lively discussions were taking place all day long in lounge-like settings furnished

with bean-bags. Overall it felt like I could obtain a glimpse of what the future would soon bring. But most of all, it was the charisma and enthusiasm of Lynn Conway that drew me into this environment. I was excited by the visionary plan of establishing some simple and logical ground rules for the design of integrated circuits, which could readily be taught to a whole class of smart students. Up to this point, my experience with the design of IC chips was more like a magical art – learned by osmosis, slowly transcending from a few old masters to their devoted pupils, who would gradually absorb the mysterious ways in which these devices were brought to life. (I am exaggerating only a little bit.)

The new approach was to extend the system of nested abstractions that was already used in the design of binary logic circuits (e.g., using TTL logic gates) upwards and downwards, so that the abstractions would cover the whole range from the architectural systems level down to the layout of the gates of individual transistors. The technology of choice was the rapidly growing *n*-MOS process that had become stable and well controlled in the early 1970s. The devices were in principle quite simple: a *source* and a *drain* region in the silicon layer, separated by a *channel* that could be turned on and off by the voltage of a metal *gate* placed on top of the thin isolating oxide layer. This geometry could be represented succinctly by a red line (representing the gate electrode) crossing a green line (representing the silicon channel). Suddenly the layout of an integrated circuit was captured by simple and clean geometrical diagrams. This had tremendous appeal for me!

Of course, the method was a little more complicated than outlined above. To make it useful, quite a few details had to be figured out. But these challenges were exactly of the nature that had rendered geometry my favorite subject in high school. So for the next several years I routinely spent one day of every week at Xerox PARC; and this was typically the highlight of the week. It was wonderful to have “one day off,” or at least one day that was quite different from an ordinary “school day.” I worked with Lynn Conway and her team to put together concrete guidelines for the new way of thinking about integrated systems. At several occasions during my consulting days at Xerox PARC, I also met Carver Mead and often engaged in heated discussions of what it really meant “to map the systems architecture onto the 2-dimensional space of a chip” or how to unambiguously specify “a linear array of *n* cells, of size *s*, separated by distance *d*.” At other times I would engage in brain-storming sessions of how to use a computer program to turn the Boolean specifications of a programmable logic array (PLA) into an array of green and red line elements crossing at right angles, thus capturing the basic arrangements of a compact n-MOS realization of circuitry that would perform the specified function.

However, for small PLAs with only a handful of inputs and outputs, the generated layouts were not competitive in compactness with the beautiful, handcrafted layouts done by expert IC layout designers. Therefore many engineers in Silicon Valley, as well as some of my colleagues in academia, dismissed those early results as “toy examples” of no real significance. Also, the sometimes overly enthusiastic statements by Carver Mead, claiming that a proper mapping of the system architecture onto the surface of a silicon wafer would improve layout density by one or two orders of magnitude, gave ammunition to people who were skeptical of those early efforts. But neither Carver nor Lynn let themselves be discouraged by such negative evaluations. They responded by decisively moving along the envisioned path, clarifying one issue after another, and solving problems one at a time as they arose.

Carver Mead saw clearly that Moore’s Law, which predicted a doubling of chip complexity every 18 to 24 months, would soon allow us to place systems on a single chip with tens or hundreds of thousands of individual switching elements; those systems could no longer be designed, drawn, and checked by the traditional “manual” methods. Thus we were rapidly approaching a “complexity barrier” in the design of integrated circuits. And indeed, it was not too long before some computer chips had logic arrays with several dozens of inputs, on the order of a hundred outputs, and more than 200 *min-terms*. Now the computer generated layouts could produce working solutions that could no longer be obtained with manual layout.

I eagerly absorbed those ideas and developments and brought them back to Berkeley. My personal, special graduate course, CS 248, was devoted *Modular MOS LSI Design*. This gave me an opportunity to try out emerging new ideas and carry that feedback back to Xerox PARC. By 1978 Lynn Conway had launched a full-blown effort to capture all the new design concepts in a textbook, and she was making available emerging chapters to whoever was willing to teach such a design course. So this was different of the normal model of developing a text book, where notes accumulated over several offerings of a particular course eventually got distilled into a refined text that documented all the good ideas that had survived this evolutionary process. In this case, ideas developed in brain-storming sessions at Xerox PARC were used to define various lectures and set the overall itinerary of my LSI Design course.

In addition the new methods were also applied in joint research with Dave Patterson. Realizing that the real estate on an IC chip was a limited, precious resource, we carefully evaluated what circuitry and what functions would deliver the most “bang” per square millimeter for making a powerful, general-purpose microprocessor. We then applied the new structured layout methodology for n-MOS circuitry. The result was the RISC (Reduced Instruction Set Computer) principle, and by 1981 our students had realized the first working single-chip RISC (Figure 2).

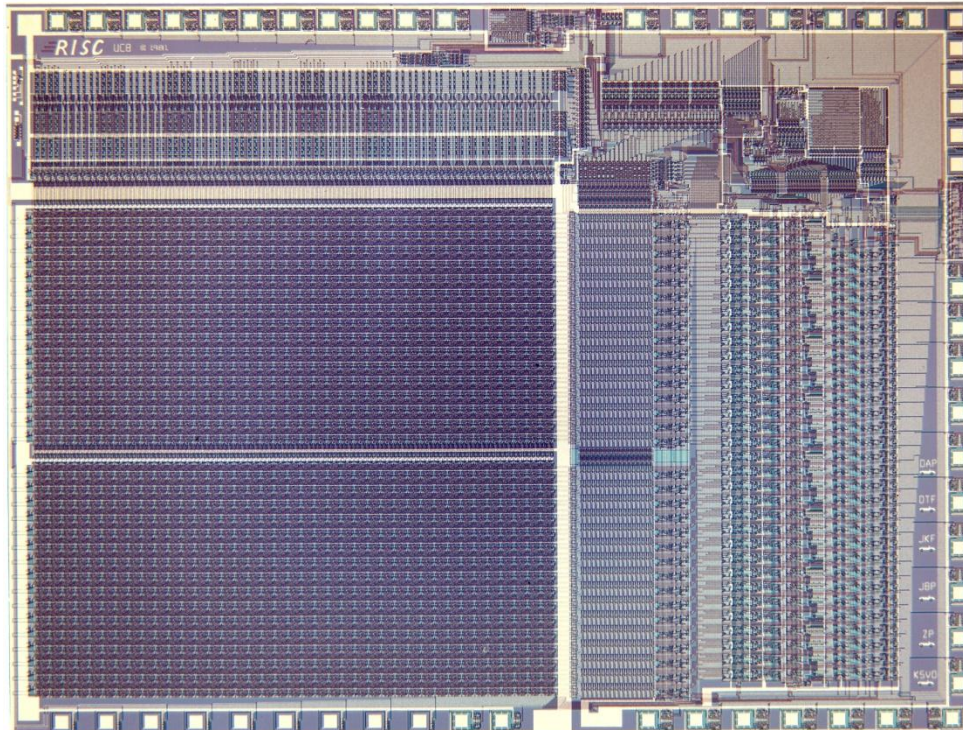


Figure 2: *The first functional RISC chip built by graduate students at U.C. Berkeley.*

Gradually the LSI chip-building activities at Berkeley expanded. For the more ambitious follow-on projects, SOAR (Smalltalk On A RISC) and SPUR (Symbolic Processing Using RISCs), we needed the help and expertise of our colleagues on the EE side, in particular, David Hodges, Alberto Sangiovanni, and Richard Newton. We also needed better and easier-to-use CAD tools to lay out those complicated chips in symbolic sticks format, convert them into compact layouts without design rule violations, and finally verify the proper logic operation and timing behavior of those circuits. John Ousterhout in the CS Division played a major role in this domain; he made a personal commitment to develop a new IC-CAD tool for every new computer chip that this EECS-team designed.

This activity attracted a lot of attention and drew in ever bigger groups of students. Doing IC layout with a user-friendly CAD tool was a lot of fun (almost like today’s video games) and it gave the students a true sense of achievement. By the early 1980s seven faculty and more than 30 students from both EE and CS were working together to develop new powerful IC CAD tools and designing computer chips that gained attention and appreciation in Silicon Valley as well as in academia. Very soon there was no doubt that the IC-CAD effort at UC Berkeley was the best one in academia world-wide, and that this was mostly brought about by a close collaboration of research groups from both EE and CS. Before too long, other schools, like MIT and Stanford, took notice and started to emulate the Berkeley EE+CS collaborative model.

In summary, thanks to the outreach of Lynn Conway, thanks to her enthusiasm and support, and thanks to the exciting ideas emerging in VLSI Design, I was able to start an activity at Berkeley that brought together EE and CS and thus allowed me to make good on the mission originally assigned to me by Tom Everhart. By the early 1980’s the harsh boundaries between the two Divisions had mostly disappeared, many interdisciplinary research groups had formed, and students were freely transitioning from one side to the other. As an example, Manolis Katevenis first did his MS degree under David Sakrison on the EE side and then came to CS Division to do his PhD with me and Dave Patterson; he became the key designer of the successful RISC chip.

The domain of digital integrated circuits will continue to evolve rapidly. I believe that in 1980 few people could truly imagine of how far this technology would develop in the subsequent 30 years. The same is probably true today for the next period of 30 years. Thus we are still very much in need of bold, visionary thinkers of the kind of Lynn Conway and Carver Mead.