IMPLEMENTING VLSI SYSTEMS IN A RESEARCH ENVIRONMENT

ROBERT W. HON

DEPARTMENT OF COMPUTER SCIENCE
CARNEGIE-MELLON UNIVERSITY
PITTSBURGH, PA. 15213

This research was sponsored by the Defense Advanced Research Projects Agency (DOD), ARPA Order No. 3597, monitored by the Air Force Avionics Laboratory Under Contract F33615-78-C-1551.

Traditionally, integrated circuits have been designed and implemented only for applications likely to require large numbers of chips. The IC's were realized by closely cooperating experts, each well-versed in some aspect of the art of integrated circuit design and manufacturing. Economics and the detailed knowledge needed prohibited any group unwilling or unable to make a considerable resource commitment from producing integrated circuits. Prototypes and one- or few-of-a-kind systems were built from off-the-shelf components (or not at all). Recently, advances in IC technology and the emergence of new design methodologies have made it possible for people lacking IC manufacturing expertise to design VLSI circuits. While the population able to cast their ideas in silicon is increasing, a previously unseen set of problems has emerged. The remainder of this paper will examine the background of these people and some of the problems that are now critical obstructions in the path to implementing VLSI circuit designs.

The new group of integrated circuit designers comes primarily from research organizations at universities or industrial firms. It is important to stress that they are not in the business of producing IC's for sale, rather, they are building experimental or prototype systems. Members of the group may have physics or engineering or computer science as a main field of research. Their reasons for designing a VLSI system range from the need to implement a particular function for a real-time application to testing the feasibility of a new machine architecture. Designing at the level of individual devices allows these researchers to adjust circuit complexity / execution speed tradeoffs to suit their needs. They also gain the freedom to employ novel or clever structures which might be particularly well-suited to solving the problem at hand. For these designers VLSI is a particularly effective means to an end. At the same time, designs which may conceivably be built in a commercially available

technology lend credibility to such prototyping efforts.

In the semiconductor industry, an integrated circuit ultimately destined for mass manufacturing and sale is created in a multi-step process which involves a number of people with diverse backgrounds. Engineers or computer scientists produce a block diagram of a new function (or processor, or whatever) that they want. After a logic diagram is made, other engineers do the detailed electrical circuit design, keeping in mind the devices available to them in the particular process (e.g. CMOS). Layout specialists then produce specifications for a set of masks to be used in the selective patterning of the layers on the silicon wafer. Once the mask specifications are completed the actual masks are generated in a process which is essentially photographic. From there, the masks are used by the fabrication line to pattern the surface of silicon wafers each containing several hundred copies of the IC. The wafers are separated into individual chips (dicing), mounted in IC packages, and wires are connected from the pads on the circuit to the pins of the package (bonding). A cover is affixed to the package and the circuits are ready for testing.

In a high-volume production setting this entire sequence of steps is usually carried out in-house. Each of the various phases of the process may be affected by the preceding or following phase. For instance the layout people may request modifications to the circuit in order to get around a particularly difficult routing problem. Such problems may require multiple passes over the offending portions of the design. Close interaction between the specialists is possible and is instrumental in producing a working IC. Unlike those working in the specialized world of the semiconductor industry, a scientist in a research environment must play the roles of computer scientist, circuit engineer, and layout specialist (at the very least). Unfortunately he or she does not have years of

experience or even experienced colleagues to draw upon when problems are encountered. How then can he hope to effectively design VLSI systems?

Two recent developments have been instrumental in bringing IC design into the reach of researchers. Design methodologies, such as described in Mead and Conway [1979], have distilled the complexities of semiconductor devices into a straightforward set of design rules and principles. By encouraging the use of geometrically regular structures and hierarchical design, these structured techniques allow complex IC's to be designed in a relatively short time. The abstraction provided by the rules and principles allow IC's to be designed in a "cookbook" manner, largely ignoring the microscopic behavior of semiconductor devices. The other advance is the availability of standard semiconductor processes, n-channel silicon gate MOS for example, which allow the same circuit to be processed by any of a number of manufacturers with comparable results. Aided by the new design methodologies, a researcher can cope with the complexity of a VLSI circuit; the widely available standard processes assure that his or her circuit can be fabricated.

The priorities of these prototype designers reflect a much different emphasis than those of their counterparts in a mass production environment. A designer working on a system which is intended for marketing is almost surely optimizing the design for some combination of high device density and performance. Conversely, a prototype designer is more interested in trying out a research concept (perhaps an unusual interconnection scheme) and thus will use conservative layout rules and liberal timing margins to insure that the concept and <u>not</u> the technology is the limiting factor. Rapid turnaround (mask generation - wafer fabrication - packaging) is a prime requisite to the research designer since tests on the chips provide important feedback. He or she may try a number of different approaches to

a problem, and count on redesigning the chip several times before reaching a conclusion. In this respect the researcher is not unlike a computer programmer who is implementing a software system: the components of the system may undergo radical changes before the result is blessed, and very little optimization is done until the last iteration, if then. Since these designers are interested in their research and not the technology per se, most prefer to remain ignorant of the details of mask generation and wafer fabrication. Indeed, few understand the detailed physics of semiconductor devices.

Once the designs are completed, the question arises as to who will make the masks, process the wafers, and package the chips. One solution is a complete, in-house facility. In view of the large capital investment required, the skilled personnel support needed, and the prospect that much of the equipment will be obsolete in a few years, most research groups cannot afford this approach. In addition such facilities would not benefit from process improvements common in a larger volume industrial installation. Thus research groups must turn to an outside source, at least for masks and wafer processing. However, it may be practical for a research group to own and operate its own dicing and bonding facility.

The cost of subcontracting mask generation and wafer processing leads to the use of multi-project chips -- IC's which are formed from the juxtaposition of several independent designs. Multi-project chips reduce the cost per idea (or design) tested since it usually does not cost n times the cost for one design to have n designs on a chip. Even so, mask generation for large chips costs about $7,000 and the fabrication of 10 to 15 three inch wafers costs about $2,000. Currently a research group might produce one multi-project chip per quarter, composed of approximately 15 projects of several hundred transistors each. Such a chip might be 8 mm on

a side and be further divided by interior scribe lines to facilitate packaging. While one would expect a small percentage yield on a chip of this size, notice that on the average each project takes up only 1/15th of the total area. Therefore it is quite possible for a substantial fraction of the several hundred chips fabricated to contain working versions of most projects.

The realization of rapid mask generation and wafer fabrication turnaround depends on smooth, reliable interaction between the research group, mask house, and fabrication line. An important factor working against smooth interaction is the informal nature of the interfaces that now exist between participants. Mask houses and fabrication lines typically deal on a person-to-person basis with others in the business of manufacturing IC's for sale. Precise documentation of the requirements of each party has been unnecessary, as customers were assumed to have a certain amount of knowledge about the subcontractor's process. The research designer has no such knowledge, worse still, it is often not obvious to either the designer or the subcontractor that an important detail has been overlooked. Misunderstandings about the coordinate system used on a particular pattern generator or the line widths required by a fabrication line can stall the processing of masks or wafers, and ultimately double or triple the time to complete a process which normally requires 3 weeks.

At another level of detail, certain information about a particular phase of the subcontractor's process may speed turnaround or produce a better result. For instance, the manner in which the mask specification data is sorted can significantly reduce mask generation time and cost. Such information is valuable, yet not necessarily available. In the quality conscious semiconductor industry, unusual structures or a feature in an

unexpected place (even if intentional) can cause a delay in processing while the operator verifies the presence or absence of an error. The previous types of problems can be minimized by complete written specification of the requirements of each participant. It is desirable for researchers to be able to treat mask generation and wafer fabrication as "black box" processes, with well defined input requirements and output products. Presently most of this information resides in the heads of various experts at their respective companies; tapping such expertise is a difficult task, but one that will have great benefits.

Researchers' hopes for quick and painless implementation of their IC designs are based on the somewhat naive assumption that a tightly connected sequence of processes can be separated into a series of independent black boxes with immediate success. The simplicity of the goal belies the wealth of detailed knowledge required to effect a solution. Is there really any chance of routinely and quickly implementing prototype VLSI designs? A multi-project chip (9.3mm x 6.3mm, 10 projects) designed at Xerox PARC in the summer of 1978 took 20 weeks for mask generation and wafer fabrication. Largely as a result of that experience, a group at the Massachusetts Institute of Technology were able to have a similar multi-project chip processed in four weeks (thanks to an outstanding effort by all involved parties). The challenge remains to mount a semiconductor industry / research group cooperative effort to implement one-of-a-kind VLSI systems consistently, quickly and inexpensively.

REFERENCE

Mead, C., and Conway, L., Introduction to VLSI Systems, Addison-Wesley, Reading, Ma., at press.