LONGER TERM DIRECTIONS FOR SEMI-CUSTOM VLSI

Gordon B. Hoffman
United Technologies Microelectronics Center,
Colorado Springs, Colorado

Through the convenience of jet travel I have been able to talk both in Detroit and here in Pasadena today--although I admit the time zone change helped, too. As I was cruising along I realized that next year we will begin limited production of a set of semicustom CMOS gate array chips for the jet engine fuel control on the Pratt & Whitney engine that has recently received so much attention due to large orders from Delta and American Airlines. A year or two after this production begins, my life, as well as yours, might well depend on the performance of that controller. It's a sobering thought that tends to bring home the reality of the technologies we develop.

Our newly formed company, United Technologies Microelectronics Center, or UTMC for short, is dedicated to the development and design automation of semicustom circuits. Before UTMC was formed last year, all of the divisions of United Technologies had to go outside for their custom IC needs. It was getting increasingly difficult to get support from the merchant semiconductor industry, particularly where low volumes of devices were required.

In 1979 United Technologies acquired Mostek, and there was hope that Mostek would help alleviate this difficulty in obtaining custom integrated circuit support. Now those who know Mostek realize that "custom" is not a happy word at Mostek. "Custom" represents to Mostek the wrong use of scarce design resources. Mostek does not disagree that there is a strong market need, but for Mostek it has been a conscious business decision not to participate. Mostek did not originally hold this view; in fact, the first one-chip calculator circuit was made by Mostek on a custom contract. The original customer is now bankrupt, however, which gives you an inkling of why Mostek is leery of the custom business.

The custom needs of the UTC divisions range from five- to ten-chip sets a month for esoteric military applications to high-volume chips for automotive applications. Because of the mismatch with Mostek, a high-volume MOS commodity IC supplier, a joint study was launched by UTC and Mostek to

find a solution. As a result of that study UTMC was created last year in Colorado Springs and funded with $22M for Phase 1. Our Phase 1 goal is the development of a CMOS gate array design system by the end of 1982. Mostek will benefit as well as the other UTC divisions, both as a user in its system divisions and in its ability to offer semicustom services to the merchant market using the tools we develop at UTMC.

Gate arrays were chosen as our initial thrust because automation of that design style is possible in the time scale we had to work with, that is, ASAP. Gate array design turnaround time is rapid, and production costs are reasonable, even for the automotive volumes some of our divisions require.

Since our goal from the outset was an automated design system, actual design of the arrays and processing considerations are secondary to CAD requirements. For example, if we can simplify our software or make it easier to use by slightly increasing the wafer processing cost, then we'll live with the extra processing cost. In fact, one of our first decisions was to use two-level metal with our CMOS arrays to make the routing problem easier and the corresponding software task smaller.

Preoccupation with chip area is another concern that needs examination with semicustom circuits. Yields for gate arrays and other semicustom approaches tend to be much higher than custom circuits of corresponding chip size. Several factors are involved including a lower active density, fewer design rule violations particularly with automated design systems, and the fact that a cumulative learning curve generally applies to semicustom circuit cost independent of the actual customization. It's really the number of good die per wafer that determines chip costs, which is in turn determined by a combination of chip size and yield. A gate array twice as large as an equivalent custom chip, but with twice the yield, has the same chip cost, and as it turns out, chip cost becomes less important as system functions require more expensive, high-pinout packages, and as other system integration cost savings are taken into account. I'll talk more about this subject a little later.

We will have the capability to merge several different array metallizations onto the same mask set and, therefore, on the same wafer. This will

allow us to economically process very small quantities of devices, even breadboard quantities.

The complete cycle from the start of logic design to delivery of custom chips should require 3 to 6 months vs 18 months or more for traditional custom circuits. Logic design for the jet engine fuel control I mentioned earlier began last summer and the first complete chip set will be delivered this February.

Despite manual placement and routing of these four chips, each with about 6000 CMOS pairs, only seven months were required for the entire design cycle, so I feel very confident about the 3 to 6 month projection using design automation.

This overall system really represents a type of foundry operation, connecting silicon to order. The concept can, of course, be extended to other forms of semicustom IC's and I'll talk about that later.

Now many of you realize that semicustom concepts, and gate arrays in particular, have been talked about for more than 10 years. Systems similar, at least in concept, to the one I've just described have been tried but without commercial success, that is, until recently. IBM has used gate arrays quite successfully, as has DEC, AMDAHL, and Storage Technology, to name a few. I feel it's instructive to see why semicustom approaches are experiencing a "renaissance" after a long period of "dark ages," which should give an insight into future directions. Let me share with you some of the thoughts along this line that lead to our entry into the semicustom arena.

The first observation was that of transparency. If we can make silicon design transparent to logic designers then we have solved the shortage of silicon designers that so limits the interest of merchant semiconductor suppliers in custom design. Our design automation system is an attempt to do just that. Not only is the initial design time reduced, but so is the time for the inevitable design changes that always seem to occur in custom designs.

When simulation, testability verification, and array routing are complete, the divisions will send us, via DECNET, data that contains test patterns, routing, and identification. We will derive the actual test

equipment programs from these data and forward mask-making and identification information on to Mostek. It is interesting to note that our customers need not devulge the chip's function, and it would take quite an effort to derive its function from the data they supply us.

At Mostek three interconnection masks will be generated: first metal, second metal, and the vias between the two. E-beam mask-making equipment will be used, although we envision going to E-beam direct write on wafer by 1983. Preprocessed CMOS gate array wafers will be inventoried at Mostek with the first level of aluminum already applied to the wafer. We expect the turnaround time for mask-making and the completion of metalization to take two weeks or less, rather than the 13 to 18 weeks normally required for a complete set of CMOS masks and wafer processing from bare silicon.

When Mostek is finished with the metalization, they will deliver the customized wafer to us and we will probe, assemble, final test, and environmentally process the finished devices prior to shipment. Our initial packaging efforts will utilize leadless chip carriers, or a leaded version of these.

The next reason behind the semicustom renaissance is performance. The metal-gate PMOS arrays of the early 70's had a narrow application, since they were not fast enough to replace TTL in most digital systems. Today, with bipolar or CMOS arrays, gate delay times can be achieved such that TTL replacement, including Schottky TTL, is quite practical, greatly expanding the size of the potential market. With CMOS we anticipate average on-chip gate delays under 3 nsec using double-level metal and 3.5 micron gates, and there is still a lot of room for further improvement.

There is also ability to insert technology improvements without modifying the basic functional design. We expect, for example, to be able to process the same array metalization on differently processed arrays for radiation hardened applications or very high-temperature needs.

The last and perhaps most important attraction to semicustom circuits is economics. While traditional logic design costs and associated breadboarding, documentation, and preparation for manufacturing costs continue to increase, computer time costs have dropped dramatically. Some of you remember the discretionary-wired LSI program that was much touted in the

late 1960's. At that time I recall the computer run time costs alone to map and route interconnections to the good die on a two-inch wafer amounted to $2,000. Remember those are 1968 dollars, or about $4,700 in today's rhubarb currency. Based on a recent study, using an IBM benchmark program, computer costs have declined so much that $2,000 spent in 1968 would only cost $40 today, or $17 in 1968 dollars, and computer run costs are continuing to decline. In fact, computer costs may be dropping as fast as MOS RAM prices are, and you can't say that for many items in today's world.

The economic benefits of system integration onto silicon are well known and have been the driving force toward higher levels of integration. But access to higher levels of integration has been limited to high-volume system manufacturers that could justify traditional custom design costs, or to smaller users through microprocessors and related standard products. The immense volume of TTL and other forms of small and medium scale integrated circuits demonstrates that the transition to higher levels of integration is far from complete, and is an attractive area for penetration by semicustom circuitry.

Inflation has exacerbated the problem not only by increasing the cost of the small and medium scale integrated devices themselves, but the cost of putting them into systems and maintaining operation of those systems has risen dramatically. A $100 124-pin semicustom gate array, for example, may be a real bargain if it replaces 60 TTL packages, eliminates a PC board and edge connector, and saves on service and repair costs. Other system advantages include the potential elimination of cooling fans, smaller and less costly power supplies, lower system manufacturing cost, smaller cabinet volume requirements and therefore more cabinet styling latitude, and I could go on to list at least 10 more advantages of system integration but I'm sure you could come up with interesting lists of your own.

Finally, under the topic of economics are "market window costs" to a systems manufacturer. There is a measurable dollar value to getting a product into the market early, or at least before competiton gets there. As experience shows again and again, the highest return on investment belongs to the manufacturer who dominates market share, and it's hard to get this domination if you are a year or two late with your product introduction. With semicustom circuit design, times are short, production ramps up quickly

and, I believe, without the number of unexpected design fixes often required with traditional custom circuits, or conventional designs using TTL for that matter, and necessary changes are quickly implemented.

Market window costs may be a big factor in keeping even high volume manufactuers from initially going the traditional custom route, and I wonder if the switch to custom would occur even if economics were favorable. At the design automation conference last year I asked a European manufacturer of gate arrays, who has fabricated over a thousand different customizations, how many were eventually designed out and replaced by smaller custom circuits. His answer was none, to his best recollection. It seems the engineering and financial resources to make the conversion were always put into new product designs, where the return on investment was consistently higher. I expect this is not an isolated phenomenon.

Now I've talked about what we are doing with gate arrays, why the semicustom approach is experiencing a revival, and the critical role design automation will play. My perspective, admittedly, has been from the MOS point of view, although there is quite an activity in bipolar technology as you know. I now think it's time to make some predictions beyond where my earlier comments leave off.

I don't see gate arrays being a temporary phenomenon. The quick customization will be critical in my applications, despite chip size implications, at least until automated fab operations can produce chips from bare silicon in a really short time frame. There's a good deal of architectural innovations on the horizon for gate arrays as well. We have a big problem to solve with testability, but that problem is endemic with all efforts toward higher levels of integration, and must and will be solved.

The next logical step is to use standard cells, where the design system is identical to gate arrays, except that the wafers are processed from bare silicon using compact, although dimensionally constrained, layouts for each cell type, and resulting in a smaller chip. This approach would be a cost-effective transition from a gate array design for high volume applications once the system design has been stabilized. We intend to make standard cells a key part of our second-phase effort, which begins in 1982.

Following the standard cells I can see a "macro cell" approach where macro cells could be subsystems themselves, not constrained to particular shapes or positions. In the cell library there might be a 32-bit processor, speech synthesizer elements, memory, standard cell logic blocks, and so forth. The design automation requirements are hardly trivial, necessitating, for example, some sort of high-level machine description language. After all, we can't keep dealing with simple logic design inputs as the level of integration continues to rise.

And I couldn't finish this talk without a comment on foundries. It's clear to me that with design automation, connecting silicon to order will become a way of life. The order initiators will, in time, be predominately the custom community rather than the semiconductor manufacturers themselves. To make this practical, design automation is the key. Semiconductor wafer fab and backend operations must evolve from an orientation toward making huge volumes of relatively few circuit types to small volumes of many circuit types. Our plans are to develop the systems I have described today and work with Mostek to define and develop such a foundry operation in the mid 1980's.

Let me leave you with one last thought by proposing a definition of semicustom circuits, and then leaving you with a curious observation.

Semicustom integrated circuits: those ICs which appear custom to the user but are standard to the manufacturer. If this definition is a fair one, and I believe it is, then the distinction between semicustom and custom is bound to disappear. Design automation at the device level is progressing rapidly and transparent custom circuit design can't be far off. When it does become a reality then the only distinction might be entry level, device versus logic, or some other level, but I suspect that will merge as well. I will leave the conclusions to you.