The Silicon Foundry: Concepts and Reality

The "Silicon Foundry." You have heard about it; you have been told that it ought to exist; the professionals of the silicon industry certainly need one, and all the systems people who are doing designs in silicon definitely need one too. But where is it? Don't bother looking in any Silicon Valley phone directory yellow pages.

The question of whether any foundries exist today is basically a matter of semantics. The concept of a "silicon foundry" is quite simple: a facility which fabricates an integrated circuit from a design supplied by an independent party. This definition satisfies some people, and many companies today meet this basic criterion. But some other people, especially Professor Carver Mead of Caltech and others in the independent LSI/VLSI design community, attach much greater meaning to the term. To them, a silicon foundry is a facility which offers a very *clean* and *standardized* interface, fast-turnaround service, and a willingness to return either bare wafers or untested, packaged chips. To meet the "clean interface" standard, the design rules and design-data-interchange formats must be standardized, and must be interpretable in only one way. The foundry should not require the designer to know anything about mask or fab considerations such as scribe lines, fiducial marks, mask polarities, bloats or shrinks, etc. (If you don't know what these items are, the point is made. These details are irrelevant to the design process, but must be understood in great detail if you are the one who must draw up the mask specifications.) The process control monitor (PCM) should also be standardized, and should be usable to verify the quality of the processing. In turn, according to this designer/fabricator interaction model, the manufacturer need know nothing about the actual circuit being processed. The manufacturer's responsibility is only to deliver an agreed-upon number of wafers that meet the electrical specifications measured on the PCM which is stepped onto each wafer.

Another requirement for the "ideal" silicon foundry is that it accept orders from small customers as well as large, without demanding a commitment of a minimum number of dollars/year in business (beyond the charge for running one minimum lot of 10 to 20 wafers). Minimum-lot charges currently range from about $2000 for simple processes run by small vendors to $10,000 or more for fast turnaround on a silicon-gate nMOS process.

The cover of this issue of LAMBDA highlights the "ideal" foundry interface. According to this model, the IC fabrication interface should be as straightforward as that for getting your film processed. As the creative designer (or "photographer"), you don't really care what processes Kodak uses in creating your prints, or whether you have positive or negative film in your camera. The important things are that the quality be consistent, the turnaround fast, and the price low. Designers with special needs can ask for special handling and/or variations in the basic process. One could reasonably expect to pay more for this service. Because silicon processing has no equivalent of Kodak to set the processing standards, it probably is not practical to use masks (the most obvious equivalent of film) as the standard interface medium. As we will see later, a carefully defined electronic data format is probably more appropriate.

This article explores the foundry concept in detail, and addresses the needs and desires of both the user and the vendor from technical *and* economic points of view. We differentiate between the more traditional customer-owned-tooling business (COT) (which merely implies some sort of independent fabricator), and the silicon-foundry business that implies the kinds of interfaces and services described earlier.

Who Does What?

One of the first problems a designer faces when the design is finished and it's time to go to fabrication is what route to choose through the maze of companies and agents. You can find an agent to handle the whole job for you; you can handle every step yourself, or you can pick a route in between.

For example, Mel Eklund, Vice President of Integrated Circuit Engineering (ICE) in Arizona, points out that some independent design houses (including ICE itself) also act as agents for outside designers, and use their personal contacts in the mask and fab industries to speed fabrication of a design. This approach allays the fab houses' major fear: that a new designer doesn't know what he is doing, and that therefore, the job will eat up lots of time and engineering resources. Independent design houses are credible, know where problems can develop, and are familiar with the terminology of the trade.

On the other hand, a mask house only furnishes masks to your specs (if you know how to specify masks), and will not help shepherd the wafers through production at the fab facility. Getting your own masks made, and then supervising the fabricator, is by far the highest-overhead mode of operation, and requires the most knowledge on your part. The job is somewhat easier if you work with a mask house that is already supplying masks to your chosen fab facility. In such a case, the mask house already knows many of the detailed mask specifications, and errors are less likely.

The capabilities of silicon fabrication facilities vary a great deal. Some fab facilities have in-house maskmaking; most have a packaging operation for standard package types, and all have some kind of in-house test capability. Those without

in-house maskmaking capabilities will often buy masks for you from their usual vendors.

Independent assembly and testing companies are available, if you need them.

The type of group through which you choose to interface will depend on your background and understanding of the steps, the amount of money you have to spend, your personal contacts, and the turnaround time required.

Those who advocate the establishment of a "true" silicon foundry argue that the lack of standards in the industry, and the requirement that each job be handled one-on-one, engineer-to-engineer, effectively shut out the designers who don't have the expertise and contacts to see a project through fabrication. "True" foundry advocates point out that the lack of standards requires each run to be handled as a special case, thereby driving the costs much higher than otherwise necessary.

As we search for companies who are currently in the custom-fabrication (or customer-owned-tooling (COT)) business, we find that there is no clear line between companies who do custom fab and those who don't. In fact, almost any semiconductor company will take a customer-tooling job, if the price or other aspects of the agreement are appealing. In some cases, the possibility of turning the part into a standard part at a later time is an attraction. Among firms who advertise customer-tooling services, the deal dollar-size needed to attract their interest ranges from about $2000 (for an engineering wafer lot) to over a half-million dollars per year in production business. The minimums these firms quote are by no means the same for every customer. The possibility of large future orders may persuade a fab house to take a small near-term job.

The business climate is another variable. Most companies who take customer-tooling jobs also have a standard product line. If business is slow and standard products aren't moving, they might take on smaller-than-normal customer-tooling jobs to fill the line.

What follows is a sampling (and only a sampling) of the semiconductor companies who are actively seeking COT business today.

American Microsystems, Inc.

Although American Microsystems, Inc. (AMI) currently does not meet the criteria for "a silicon foundry" set forth in this article, it certainly is a major force in the custom-fabrication business. AMI has been fabricating custom-designed circuits since 1974, and COT has been a major business, with individual accountability under the AMI corporate umbrella, since 1978. AMI claims to have done 700 customer-owned circuits over the last six years.

AMI's customer-tooling marketing manager Robert Pecotich summarizes his firm's business as "doing neat things to silicon, putting it in a package, testing it to your specs, and shipping it out."

How does one do business with AMI? Pecotich says, "The key operative statement is 'whatever makes business sense to both of us,' " and adds that it is important to recognize what he calls the "NTA" principle: "No two alike."

AMI is a big company. It handles all aspects of fabrication, including in-house maskmaking. It has tried to minimize the bad aspects of bigness by organizing the company so that a customer always does business through one engineer and one marketing person. In the simplest case, a complete transaction can be handled with only five pieces of paper.

From AMI's side of the table, "business sense" means that somehow, one way or another, AMI would like a minimum of $75,000 to $150,000 worth of annual business from a single customer, depending greatly on how much work AMI has to do in support of the job. For example, if a customer arrives with a pattern generator (PG) tape containing his circuit and a fully-debugged test tape, AMI's initial engineering cost is still $12,000 to $14,000.

But AMI is willing to do any deal "that makes business sense." For example, you can submit working plates with process control monitors (PCMs) stepped in. AMI will guarantee that the wafers will meet PCM specs and mutually agreed upon optical criteria, but it is up to you, the customer, to probe the good die and do the packaging and testing. This sort of arrangement can be handled with five pieces of paper. Typical costs have been $200 to $500 per wafer, in ten-wafer

Dick Anderson, Anderson/Bogert, Los Altos, California (L), Gunnar Wetlesen, VLSI Technology, Inc., Los Gatos, California (C), Mel Eklund, Integrated Circuit Engineering, Scottsdale, Arizona (R)

lots, with turnaround times of two to six weeks (four weeks typical).

AMI has one COT project with an airborne telecommunications customer whose product volumes are small, but who willingly pays a huge premium for light weight and low power consumption. AMI delivers only one or two thousand pieces annually, at $75 each.

For small runs, the preferred "data format" for the foundry is a working plate. For runs above 50,000 pieces, AMI prefers PG tapes, so that it can supervise the maskmaking. (This procedure simply speeds recovery from broken masks.)

According to Pecotich, "The probability of success for a circuit is directly proportional to the quality of the engineer-to-engineer interface"; he urges prospective customers, "Work with us."

Comdial

Now, let's look at the other end of the spectrum of semiconductor companies: a firm that specifically does *not* need production or annual-dollar-volume commitments.

Comdial Semiconductor, a new company and a wholly-owned subsidiary of Comdial Corporation, is making a first-quarter 1981 offering of a line of proprietary telecommunications products aimed at the telephony market. We spoke with Gary Kennedy, Director of Manufacturing, who was the first of a new staff of 35 employees, all hired since July 1979.

Kennedy explained the service side of Comdial's business, and emphasized that Comdial specifically does *not* want to do production-quantity runs. The firm is aiming at very fast turnaround of prototype quantities to service IC designers—especially designers from large production houses. If a design is behind schedule and expedited fabrication will help, Comdial would like to be of service. For example, its price list (Figure 1) shows that an engineering prototype run of ten wafers can be completed in merely five working days. More specifically, only five working days elapse from receipt of working plates to shipment of guaranteed wafers. Comdial also claims to have moved from PG tape to packaged prototype parts in eight working days.

The company does its own processing and packaging, but contracts maskmaking and ion-implantation with outside firms. Comdial can handle up to 1,000 packages (ceramic or cerdip) per month, and can process up to 125 wafers per month, per customer.

Comdial's willingness to do small jobs, the fact that it publishes a process data sheet and price list, and its stated commitment to fast turn-around, all move it closer to the silicon-foundry concept. The company is quite new, and the development of its operations will be interesting to watch.

Synertek

Customer Tooling Marketing Manager Anna Appleby explained that Synertek is putting renewed emphasis on COT business, and that the company is aggressively seeking new customers. But, like AMI and other large production-oriented houses, Synertek is looking for a certain threshold of business in the form of follow-on production. For Synertek, this threshold is about $100,000. As Appleby explains, "it is a matter of matching Synertek's capabilities to identifiable markets, and of setting priorities for the expenditure of limited marketing and engineering resources." The programs must pay off for Synertek, who, Appleby adds, is willing to work with any potential customer.

Schedule of Prices for Fabrication
(n-Channel Silicon-Gate Process)

Run	Number of Wafers	Working Days to Completion	Price per Completed Wafer
Engineering Prototype	10	5	$1,000.00
		10	500.00
		15	400.00
Hot Lot	25	5	750.00
		10	400.00
Short Run	50	10	400.00
		15	300.00
		20	225.00
		25	150.00
Production Limited to 125/month		25	130.00

Wafer price based on completion time

Figure 1. Comdial offers two unique services: very fast turnaround and a published price list. (Prices are of course subject to change without notice.)

Semi Processes, Inc.

Semi Processes, Inc. (SPI) is a small (45-person) San Jose manufacturer of high-voltage, high-powered FETs and gate arrays (see LAMBDA, Fourth Quarter, 1980), and also offers CMOS custom-fab capabilities.

Sales Manager Randy Henniger says that for workhorse metal-gate CMOS, SPI will start a lot of 20 wafers (yielding about 15) for $2500 with about four-week turnaround, or will do the same thing in silicon-gate for about $10,000. But, as always, the assumption is that the customer furnishes working masks with their PCMs stepped in. SPI does not make masks, but will package and test die to a customer's spec.

Precision Monolithics, Inc.

Precision Monolithics, Inc. (PMI) is in the bipolar custom wafer fabrication business. Custom Wafer Fabrication (CWF) Marketing Manager Richard Corbin points out that the CWF operation is completely self-sustaining with a separate operating staff. They offer bipolar processing for linear circuits (including FET input operational amplifiers), Schottky, I^2L, and ECL for digital circuits.

The input to PMI's facility is a sub-master mask. PMI does not do custom design. PMI's customers must know how to do bipolar design and obtain masks. The company will adjust process parameters to suit a customer's needs, and will furnish design-rule guidelines. Processed wafers are the preferred output of its facility. Customer acceptance is based on parametric data from test patterns (which do not have to be PMI's PCMs).

Dick Corbin said that PMI looks for customers with annual production commitments, and is not in the short-run business. PMI's minimum threshold of business is a 100-wafer run and a $500,000 annual business volume.

By the second quarter of 1981, PMI will offer an oxide-isolated silicon-gate CMOS process featuring two-layer metal, five-micron geometries and gate delays of less than 15 ns.

Table 1 lists other companies active in the customer-owned tooling business. Detailed descriptions of each of these companies are given in *LSI Opportunities: Using the IC Service Industry* (Anderson/Bogert, 1980).

Acrain Cupertino, CA	Nitron Cupertino, CA
American Microsystems Santa Clara, CA	Plessy Semiconductor Irvine, CA
Cherry Semiconductor Cranston, RI	Polycore Newbury Park, CA
Comdial Semiconductor Sunnyvale, CA	Precision Monolithics Santa Clara, CA
Dionics Westbury, NY	RCA Solid State Somerville, NY
Exar Integrated Systems Sunnyvale, CA	Semi Processes Santa Clara, CA
General Instruments Microelectronics Hicksville, NY	Solid State Scientific Montgomeryville, PA
GTE Microcircuits Tempe, AZ	Solitron Devices San Diego, CA
Hughes Aircraft Newport Beach, CA	Standard Microsystems Haupauge, NY
Microcircuit Engineering West Palm Beach, FL	Supertex Santa Clara, CA
Monosil Santa Clara, CA	Universal Semiconductor San Jose, CA
Mosfet-Micro Labs Quakertown, PA	VLSI Technology, Inc. Los Gatos, CA
National Semiconductor Santa Clara, CA	ZyMos Sunnyvale, CA

Source: *LSI Opportunities: Using the IC Service Industry* (Anderson/Bogert 1981).

TABLE 1. Customer owned tooling (COT) wafer fabrication firms.

The Silicon Foundry

What are the differences between the customer-owned-tooling business served by these companies and the "silicon foundry" concept proposed by Mead and others? According to Mead, the critical factors are those of standardization and assured access for the "little guy." Many suppliers admit that usually, they are interested only if $100,000 per year in business seems likely. This one criterion shuts many "small" designers out of the market. However, standardization and access to fabrication for the "little guy" may very well be linked to each other. The lack of standardization forces each customer to be treated as a special case. According to AMI, individual attention is required at the engineering level to make the program go smoothly. This procedure undoubtedly increases the probability of success, but it is also self-limiting. AMI, Synertek, and others of their ilk can only hire a limited number of product engineers to deal with customers. The individual attention, and any special handling required during order-entry and wafer-processing, sets the "minimum level of pain" and forces these suppliers to seek minimum commitments in the neighborhood of $100,000 per year.

Mel Eklund of ICE emphasizes the importance of standardized interfaces and of "an extremely efficient inventory-production control scheme that could keep track of small things [and which could] go all the way through, eliminating a lot of expense and hassle." Eklund points out that, long ago, the auto industry learned to deal with the problem of making every car slightly different (color, accessories), yet keeping them all running down the same assembly line. The secret is automation, for which precise standards are essential.

The multi-project chip (MPC) system developed by Xerox (see insert) is an excellent example of the effect of standards and automation on the accessibility of fabrication, and, in turn, of the effect of accessibility on the number of people involved in the process.

Standards

According to Carver Mead, a silicon foundry must be standardized in three critical areas: geometrical design rules, data-interchange format, and the process control monitor (PCM).

Geometrical design rules specify how close the conductors on various layers can be to each other, how far the gate of a transistor must overlap the channel, etc. Currently popular examples of data-interchange formats are Calma data-base or pattern generator (PG) tapes. Over the years, these formats have both developed as *de facto* standards. Not only must the basic data formats be formalized, but additional items such as mask-layer names and numbers must also be specified. The PCM is a set of test structures included on each chip, or (more often) stepped into selected locations on the mask, which can be tested to provide process characterization data. The goal is a test structure that monitors every critical electrical parameter affecting the performance of the actual circuit. If the PCM is adequate, and if the PCM tests reveal no problems, then the foundry can be confident that all stages of the processing took place correctly.

Design Rules

The most obvious way to approach design-rule standards is to establish a set of fixed-value rules with values that most suppliers can meet. This approach has the disadvantage that the rules tend to be least-common-denominator rules from a variety of companies. Such a set of rules translates into circuits with somewhat lesser density and lower performance. A second problem is that processes improve and design rules shrink over time; the "standard" would have to change constantly to keep up with these changes.

Mead and Conway (1980) promote an alternative, in the form of scalable design rules that are described in terms of a parameter, λ, which is a measure of the resolution of the process. For example, the minimum width of polysilicon is stated as 2λ. The minimum width of metal is 3λ. As the process improves (which it does regularly) and the value of λ changes, the rules stay the same. The usefulness of this approach in helping the designer track a moving process is fairly clear. The rules are defined for a generic process such as depletion-load, silicon-gate nMOS, and do not change with time: only the value of λ changes. A designer working in lambda rules designs in terms of lambda, not mils or microns. When the design is complete, the designer can select the value of λ appropriate to the foundry with which he has decided to do business. Software then scales the drawing and outputs it in absolute dimensions. (See "Simplified Design Rules for VLSI Layouts" in this issue.)

This scheme does have disadvantages as well. The rules are generally not optimal for a given process at a given time. In addition, bonding pads and associated circuitry do not scale, and must be handled in some slightly more complex way.

Data Interchange Format

The data interface between the designer and the fabricator

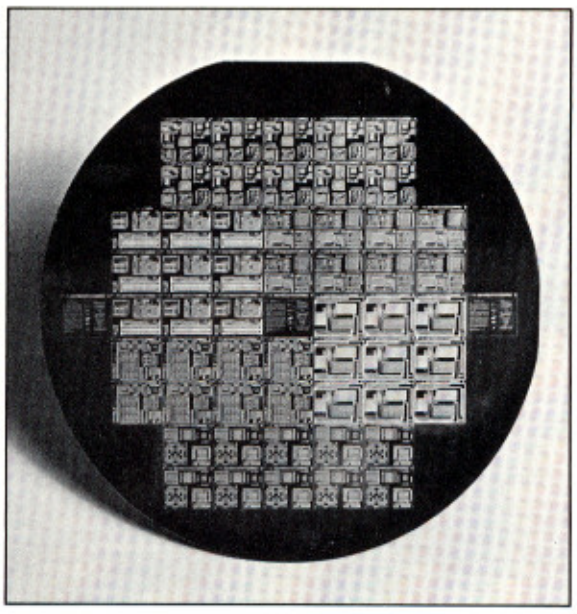

MPC

The multi-project chip (MPC) concept is just what it sounds like: a means of putting more than one design on the same chip. But it doesn't stop there: multiple *chips* can also be placed on a single wafer. The latter embodiment of the concept is especially useful for the large-chip designs that account for so much of industrial IC production.

The MPC approach is ideal for prototyping (which almost always entails surplus chips) because it is extremely effective for reducing costs. Maskmaking and wafer fabrication for a prototype IC typically cost from $5K to $15K. Under the MPC system, 5 to 40 designers share the burden, with individual bills on the order of $500 to $3K.

The first multi-project chips were implemented at Caltech by Professor Carver Mead in 1972. Mead used the technique to allow fabrication of many different student designs with one mask set and one fab run. In 1979 Lynn Conway, Alan Bell, and Martin Newell at the Xerox Palo Alto Research Center developed a semi-automated system to interact with remote designers over the ARPANet, receive and check IC design projects, merge them into multi-project chips, and create the data tape from which E-beam generated masks could be made.

In the fall of 1979, the Xerox PARC LSI Systems Area organized the MPC79 project as a major test of the new VLSI design and implementation concepts. Eighty-two design projects from 124 designers (mostly university students thoughout the U.S.) were implemented in an MPC set. Implementation began on December 4, 1979, and working, packaged chips, custom-wire bonded for each project, were delivered to designers on January 2, 1980—at an average effective cost of less than $500 per project. In the spring of 1980, the Xerox VLSI implementation system was used again, this time in collaboration with DARPA. The MPC580 project implemented 171 projects from 220 designers, in five weeks.

Managing such large amounts of complex data is a challenge to any data-processing system. In this case, it was met handsomely by the ARPANet and by the functional modules of the MPC system. The MPC system built a data-base of design files, and then merged and converted them to form the multiple MPC mask specification. Although the MPC system has been used primarily with nMOS technology, it is technology independent, and CMOS or I^2L applications are well within reason.

These projects demonstrated the feasibility and practicality of remote-entry, fast-turnaround VLSI implementation. Their success was the basis for ISI's development and operation of a regular, scheduled VLSI implementation service for the research community (see main body of article). In addition, the MPC system provides an example for industrial firms to follow for rapid and economical development of new commercial applications of VLSI systems.

involves much more than just picking a well-defined format. The first question is, "How much must the user know in order to work with the fabricator?" One attitude is that the user should only worry about generating a correct design, and should not be concerned with scribe lines, stretches and shrinks, labeling of masks, etc. Regardless of whether one embraces this attitude, the interchange format must be easy to manipulate on a variety of computing systems; it must be formally defined, so as to eliminate misinterpretation, and it must not depend on special characteristics of a particular manufacturer's design system.

Although alternate interchange formats (such as pattern generator (PG) tapes or data-base tapes for Calma or Applicon design systems) have been used quite widely, both have drawbacks. They are either ambiguous, difficult to generate on general-purpose computing systems, and too low-level, or else they have representational limits that reflect the particular system that generated them.

A mask-specification format known as the Caltech Intermediate Form (CIF) was created to overcome the problems of poorly specified or manufacturer-specific data formats. (CIF is specified in Mead and Conway (1980). A detailed description of the format is also available in Hon and Sequin (1980).) CIF has been criticized for having deficiencies of its own, but it still seems to be a major improvement over the alternatives. NCA, Inc. in Sunnyvale (a major supplier of LSI analysis software) has agreed to accept CIF as an input format for its programs. CIF is also being evaluated by CAD manufacturers as a possible universal standard. It is already being used throughout university and Defense Advanced Research Projects Agency (DARPA) communities, especially by people doing MPC work.

The Process Control Monitor (PCM)

Industry opinion seems to be that it will be tough to get companies to adopt a process control monitor (PCM) standard, both because of process variations among fab lines and because of the difficulty of defining a standard that accurately monitors all the relevant parameters. If parameters exist that are not monitored by the PCM, and if these parameters can affect circuit performance, then the PCM is not an adequate means by which the customer can measure the quality of the processing.

PCMs have been criticized for inadequacy in conveying yield information. A PCM can only indicate the quality of the processing at its own site; it cannot indicate defect levels or other yield-determining factors. Gunnar Wetlesen, Vice-President of Engineering at VLSI Technology, Inc. (VTI), points out that independent yield information is most important when a circuit is in the prototyping phase, during which undetected design-rule violations or circuit problems could unknowingly reduce the yield of the customer's circuit. Wetlesen suggests that in this situation, it would be reasonable to devote from 20% to 50% of the wafer to a "canary" circuit that would measure yield independently of the customer's circuit. During prototyping, 99% of all chips are often unused anyway; therefore, the relatively large number of chips dedicated to the yield monitor would not cause a problem. This would be especially true in the case of multi-project-chip (MCP) wafers (see insert). This reasoning flies in the face of "normal" industry logic, but makes a lot of sense

in the prototyping context. This "canary" circuit is a well-characterized, easily tested circuit which provides immediate feedback on the overall quality of the processing by directly testing the ability of the process to deliver a functioning LSI-scale circuit.

Other organizations that recognize the value of a standard PCM include Silicon Systems, Inc. (SSI) of Tustin, California, and the Defense Advanced Research Projects Agency (DARPA) of the U.S. Department of Defense. Gene Potter, Chairman of SSI has said that he believes a universal test structure is very important to the industry—especially to companies like his that have multiple wafer sources. Lt. Col. Duane Adams of the DARPA reports that the National Bureau of Standards and the Jet Propulsion Laboratory associated with Caltech have already developed a standard, and have even tested it (once) on a recent multi-project-chip run. However, the results have not yet been made public.

Promoting Standards

At least two groups with a vested interest in establishing standards are funding activities to promote such standards. Yngvar Lundh, a senior scientist at the Norwegian Defense Research Establishment and a professor of computer science at the University of Oslo, has not only called for the establishment of a silicon foundry, but has also put money where his mouth is. Professor Lundh makes a strong case for system designers' need for access to IC-fabrication facilities. He, too, calls for a clean interface and short turnaround time, and would also stipulate a simple pricing policy and a well-defined process. Three Norwegian agencies (the Royal Norwegian Council for Scientific and Industrial Research, the Norwegian Telecommunications Administration Research Establishment, and the Norwegian Defense Research Establishment) have retained Anderson/Bogert of Los Altos, California, to write the first design book in a series entitled *LSI Circuit Design and Procurement Manual*. The first book will concentrate on the silicon-gate nMOS process. Among other things, the book covers a "test cell" which will be included on *each chip*. Thus, wafer buy-offs or even chip buy-offs could be based on whether test patterns met standard electrical specifications.

Blair Proctor

Larry Matheny and Robert Smith, SynMos, Santa Clara, California.

Richard C. (Dick) Anderson of Anderson/Bogert told LAMBDA that the *LSI Circuit Design and Procurement Manual* will be ready in mid-1981. In the meantime, Anderson/Bogert is offering an "LSI Circuit Design and Procurement Program" in the form of sponsorships of the manual. Each $9,500 sponsorship entitles participants to monthly status reports, preliminary drafts of the manual, an inquiry service, and a two-year post-publication updating service.

In the first report, Anderson/Bogert will develop a set of design rules and corresponding device parameters for silicon-gate nMOS. They expect at least three vendors (presently AMI, Synertek, and ZyMos) to be the fabrication sources in accordance with standard design rules. Anderson/Bogert will also define a standard test cell, with performance parameters and suggested acceptance-test limits. Finally, they will produce the procurement manual, which will feature a design tutorial using standard rules, and which will also include suggestions and standards for procuring tooling, fabrication and testing. At this writing, it is not clear if the design rules will be lambda-based, and the CIF format may not be included.

Alternate Standards

DARPA has funded a continuing "silicon brokerage" activity at the University of Southern California's Information Sciences Institute (ISI), for DARPA-sponsored projects and for some NSF-sponsored projects. ISI has acquired the MPC (see insert) know-how from Xerox, and is now making regular MPC runs. All fab, mask, and packaging services are furnished by outside companies. Design files are submitted to ISI's MPC system (called MOSIS) via the ARPAnet. The design rules and the data format are indeed those of Mead and Conway, but what about the test chip? Lt. Col. Duane Adams of DARPA reports that JPL and NBS are both working on a standard, and that they made a pass at it on the MPC880 run (results not yet published). However, he says, the standard isn't ready yet.

In an effort to locate alternate industrial sources for the services provided by Xerox, Micro Mask, and Hewlett-Packard for the 1979/1980 MPC runs, Lt. Col. Adams also placed an ad in the October 7, 1980, issue of *Commerce Business Daily*, asking for expressions of interest by parties willing to provide fast-turnaround wafer fabrication. He was also interested in maskmaking, packaging, and testing services. The design rules and methodology would be those of Mead and Conway; each wafer would carry multiple die-types, and each die-type would carry multiple projects.

According to Lt. Col. Adams, DARPA intends to qualify several vendors by asking them to fabricate wafers from masks generated in accordance with Mead-and-Conway rules. DARPA would prefer to use CIF as the interface but, for now, will translate CIF to MEBES (Manufacturing Electron Beam Exposure System) format at ISI, have E-beam masks made, and then send the masks to fabrication houses. DARPA intends to locate, qualify, and fund vendors, so as to move the design methodology into the "real world." It will release the list of respondents after funding.

If DARPA wants to "move the methodology into the real world," what, you might ask, is available in the open market right now? The answer depends on who you are. If you have DARPA-sponsored or certain NSF-sponsored projects, then

you may call Danny Cohen at ISI. The rest of you will be interested in the following technical reports and in two new companies that are currently setting up to furnish foundry or broker services for the rest of the design community.

Xerox Palo Alto Research Center has published an excellent technical report, *Guide to LSI Implementation* (Hon and Sequin 1980), which covers a wide range of topics. These topics include design tools, data interchange formats (CIF), mask specifications, procurement of fabrication services, testing, and other important considerations for the designer who needs to understand the details concerning fabrication of LSI designs.

New Companies—New Opportunities

SynMos

Many industry observers lament the lack of new-company start-ups because of the skyrocketing cost of initial capitalization. However, not all deals require millions of dollars worth of processing equipment. In fact, Larry Matheny and Bob Smith founded SynMos on the premise that a well-connected chip broker needs no equipment. The result? SynMos opened for business. "There isn't a company around that couldn't handle a few more wafers, even when they think their capacity is limited," claims Matheny.

Before forming SynMos, both Matheny and Smith were associated with AMI. They felt the Mead-and-Conway revolution coming, and started SynMos to "service all those thousands and thousands of guys who cannot go the custom route, and cannot even go the gate-array route, because they can't make production volume commitments."

Matheny and Smith believe that a growing number of equipment manufacturers' engineers will want to do their own IC designs. When the designs are completed, these engineers will want to obtain a few prototypes (piece parts, not wafers) quickly, with minimal cost and hassle. A broker can minimize the cost just by being well-enough connected to know how and where to shop, and how to contract for services. Further cost reduction is possible by having several customers share mask-making and fabrication costs (sound familiar?).

Matheny believes that by mid-1981, SynMos will offer what he calls shared-silicon technology (SST) fabrication runs (known elsewhere as MPC runs). SynMos plans to accept CIF tapes, and will return small numbers (as few as fifteen pieces) of packaged parts—for a price. The firm is currently prepared to accept PG tapes and to act as silicon brokers (general contractors), obtaining fabrication and packaging wherever they can get the best price and turn around time.

SynMos is open for business today, basically selling expertise—a service. Remarks Matheny, "What would a poor equipment guy say if a fab house called with a message like 'your run bombed'?" SynMos wants to handle all such problems for engineers who don't want be entangled in the communications, technology, and commercial aspects of getting a systems design into silicon. "We're here to streamline the process from the light-bulb to the (working) part," Matheny concludes.

VLSI Technology, Inc.

In contrast to ISI and SynMos, which do not own mask-making, fabrication, or packaging equipment, a new company, VLSI Technology, Inc. (VTI) of Los Gatos, California, (see News Department, this issue) will provide all the facilities and expertise required to manipulate data, make E-beam masks, and fabricate ICs. VTI will accept design files in CIF (or other common data formats) via magtape or via a network facility known as VTINet. VTI will return wafers or packaged chips to customers, and will furnish clean interfaces, including a standard PCM which provides both system-performance parameters and transistor parameters.

VTI's fabrication facility will be set up specifically to service markets that put a premium on service and that require fast turnaround. Gunnar Wetlesen, VP of VTI, observes, "at present, no IC company exists with both high technology for VLSI *and* a service orientation for quick turnaround." In a standard large semiconductor company, with emphasis on productivity in a continuous-flow, standard-product environment, the goal is to maintain inventories. VTI's operating philosophy will be completely different. In the pilot-run environment, says Wetlesen, "it is more important to be effective, rather than efficient." Of course, the rest of production (order entry, production control, and quality control) must also be oriented toward quick turnaround—a mode of operation that, incidentally, requires more highly skilled employees than does that of normal fab facilities.

VTI will provide a full spectrum of services for VLSI designers, including VLSI design courses (available now), design tools, and silicon-foundry service (including MPC-style capabilities). When the fab line goes into operation (Spring 1982), VTI will accept design files in various formats and will deliver products in wafer form or as packaged chips, with an expected three-week turn-around time from design-file in to wafers out.

Conclusion

As you see, where you go and what you pay depends on what you know and on what you can do for yourself. If you have masks, you can easily get wafers processed (if your design rules are compatible with normally available processes). You can achieve mutually agreeable acceptance if you use your fab house's PCM, or perhaps, in the future, some other universally accepted test structure. If you have Calma, Applicon, or PG data, you can readily get masks. At this date, if you have CIF data (and are not DARPA- or NSF-funded), you have a problem—but several companies may soon be able to help you solve it.

Of course, adequate access to fabrication is vital to the custom-IC community. LAMBDA intends to cover this topic regularly. We are especially interested in your comments, criticisms, and experiences; let us hear from you!

—WDJ and DGF

References

Anderson, Bogert. February 1981. *LSI Opportunity: Using the IC Service Industry*. Electronic Trend Publications, Cupertino, CA (408) 996-7401.

Anderson, Bogert. 1981. *LSI Circuit Design & Procurement Manual.*. Electric Trend Publications, Cupertino, CA.

Hon, R.W. and Sequin, C.H. 1980. *A Guide to LSI Implementation*, Xerox PARC Technical Report SSL79-7, Palo Alto, CA.

Mead, C.A. and Conway, L.A. 1980. *Introduction to VLSI Systems*, Reading: Addison-Wesley.