Conference on Microelectronics
1982, Adelaide, 12-14 May

Australia's First Multi-Project Chip Implementation System

J.C. MUDGE
Head, VLSI Program, Division of Computing Research, CSIRO, Adelaide
and
R.J. CLARKE
Co-ordinator, VLSI Program, Division of Computing Research, CSIRO, Adelaide

SUMMARY A proven integrated-circuit design methodology, based upon the pioneering work of Mead & Conway, was introduced to Australia in 1981. This has led to a radical increase in the number of integrated-circuit designers.

To complement this, CSIRO has designed a chip implementation system for the rapid, remote fabrication of experimental designs. By combining the pattern-generation files of many independent designs to form one mask set, fabrication costs are reduced by an order of magnitude.

The components of the implementation system are (a) a standard interface to fabrication; and (b) information and logistics management. The first fabrication run will use electron-beam masks, a five-micron n-channel silicon gate process, and a forty-pin package.

1 INTRODUCTION

As the application of microelectronics becomes more pervasive, systems designers, who formerly constructed electronic systems from standard chips (often microprocessors), are now faced with designing complete applications systems on silicon.

A new training need has therefore arisen - the training of new electrical engineering and computer science students and the retraining of of practising designers (in industry and research laboratories). The pioneering work of Mead & Conway (1980) appears to be the most successful approach to solving the teaching need. Courses based on their work are now being taught at over a hundred universities in the United States; several such courses are being introduced in 1982 in Australian universities and institutes of technology.

This increase in design activity has created a demand for chip fabrication that is not met by the conventional manufacturers of chips ((i) firms who design and manufacture integrated circuits; and (ii) teaching and research institutions who own fabrication equipment for student use). The demand comes from both students (in training or retraining) and trained designers who wish to fabricate their prototypes.

The main steps in fabrication are mask making, wafer processing, and packaging. At each step, a designer faces both technical and logistics problems. A sample of these follows:-

1.1 Lack of Technical Specification

In the traditional integrated-circuit firm, chip designers and process designers sit side-by-side. Much technical information passing between the two groups is never documented, and yet mut be known by a designer.

1.2 High Variability Between Vendors

Different firms have quite different conventions (mask polarities, for example), process tolerances, and electrical performance. Even such fundamental electrical parameters as transistor-threshold voltages vary from vendor to vendor.

1.3 High Costs

Although a designer needs only a few chips to test his prototype design, he must pay for the set-up costs and processing of a complete batch run (of around 1,000 chips). Prices range from $30,000 to $50,000 for such a run.

1.4 Vendor Interface Logistics

Because some of the steps, mask making for example, must be done overseas, the problem of communication is exacerbated by geographical distance and international boundaries. Not only must very detailed and sophisticated technical detail be exchanged, but export licences and customs clearances are required. The problems of logistics deter all but the naive designer.

1.5 Unpredictability of Turnaround Time

Because the designer of applications systems offers only a low-volume business to a semiconductor manufacturer, his run is likely to be pre-empted by high-volume designs. This is particularly true of university designs; delays of up to twelve months have been observed, with four months being the common delay through all of the queues.

Clearly, an implementation system is needed which

a. presents, to the designer, a standard technical interface to manufacture;

b. reduces the cost;

c. simplifies the logistics.

Australia's first multi-project chip implementation system, AUSMPC, attempts to do this. Cost savings will occur by having up to forty independent designs share the same fabrication run.

Figure 1 shows the metal layer of a chip containing several independent projects which have been fabricated in the intended way. A particular chip is bonded (input-output connections wired) for just one of the designs, say the one in the centre. Other copies of the chip (wire-bonded differently) are given to the other designers sharing the same fabrication run. This type of implementation is called a Multi-Project Chip (MPC).

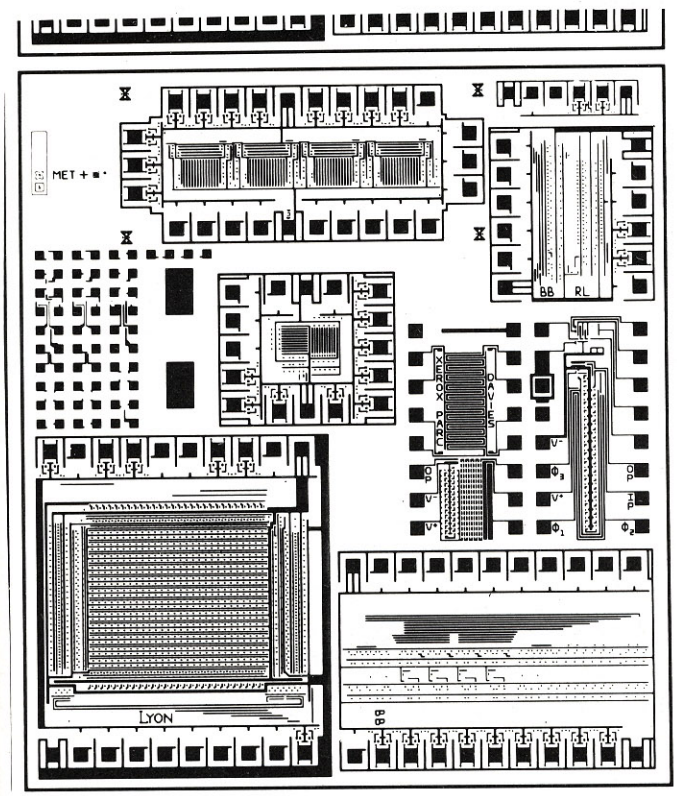

Figure 1

We use a standard interface to fabrication pioneered by Mead & Conway in the United States.

The components of such an interface are:

1. a set of geometric design rules;
2. a standard data format for interchange;
3. a standard set of test structures;
4. an implementation system (responsible for information management, merging of designs, vendor interfacing, etc).

The decision to commit resources in our CSIRO research program to design and implement AUSMPC was not taken lightly. However, we saw no alternative. At this writing, there are just two proven multi-project-chip-implementation systems available to the independent designer. Although a few systems have been developed by US companies to serve the needs of their in-house designers, they are not available to outsiders. The two proven systems are those offered by (a) MOSIS and (b) Synmos Corporation. MOSIS is available only to US ARPA contractors. Synmos, a new Californian firm, would meet most of our needs. However, all steps of fabrication would be carried out in the United States. We wish to use Australian capabilities wherever possible; therefore, we have chosen to design our own implementation system.

2 THE DESIGNER INTERFACE

2.1 Simple Design Rules

A key component of the standard interface to fabrication is the use of a simple scalable set of design rules [Mead and Conway, 1980] and [Lyon, 1981]. The design rules are based upon a single parameter, lambda, a length unit corresponding to the minimum feature size of any geometry within a mask layout. It also corresponds to the worst case misalignment which can be tolerated before circuit performance is significantly degraded.

The value of lambda was determined by averaging design rules from a number of nMOS fabrication processes. The result is a set of rules which, although overdimensioned for any particular fabrication process, will work for a range of processes. There are about 16 design rules for the nMOS process used for the MPC, all based upon the length unit, lambda. As processes improve, one scales the value of lambda down to suit the newer process. Longevity of designs can be attained using scalable rules of this type.

Whereas there are about 16 rules presented to our community of designers with this approach, a traditional integrated-circuit firm would require its designers to be aware of about 120 rules.

2.2 Standard Data Format

We have chosen Caltech Intermediate Form (CIF) as the standard machine-readable form for representing integrated-circuit layouts. We have followed the VLSI research community in the US in this way. Other standards are in use in Europe, but CIF is by far the most widely used. It is detailed in [Mead and Conway, 1980] and [Hon & Sequin, 1980]. Standardised extensions to cover other technologies (such as IIL and CMOS) are being considered.

2.3 Flowchart of the Implementation System

Designers in the Australian MPC community see the implementation system shown in Figure 2. They present their designs in CIF, and, after some delay, packaged chips are returned to them. Inside the box labelled "Implementation" on the flowchart are the tasks we perform, or subcontract. They include merging of designs into a starting frame, converting the data into a patterning format, making masks, processing wafers, dicing the wafers into chips, and mounting and wire bonding the chips into packages.

2.3.1 Selecting and merging designs

Once designs have been submitted to CSIRO for fabrication, several tasks must be performed before the data can be sent for mask making. These are: the checking of designs, the selection of suitable projects, and the merging of designs into a starting frame.

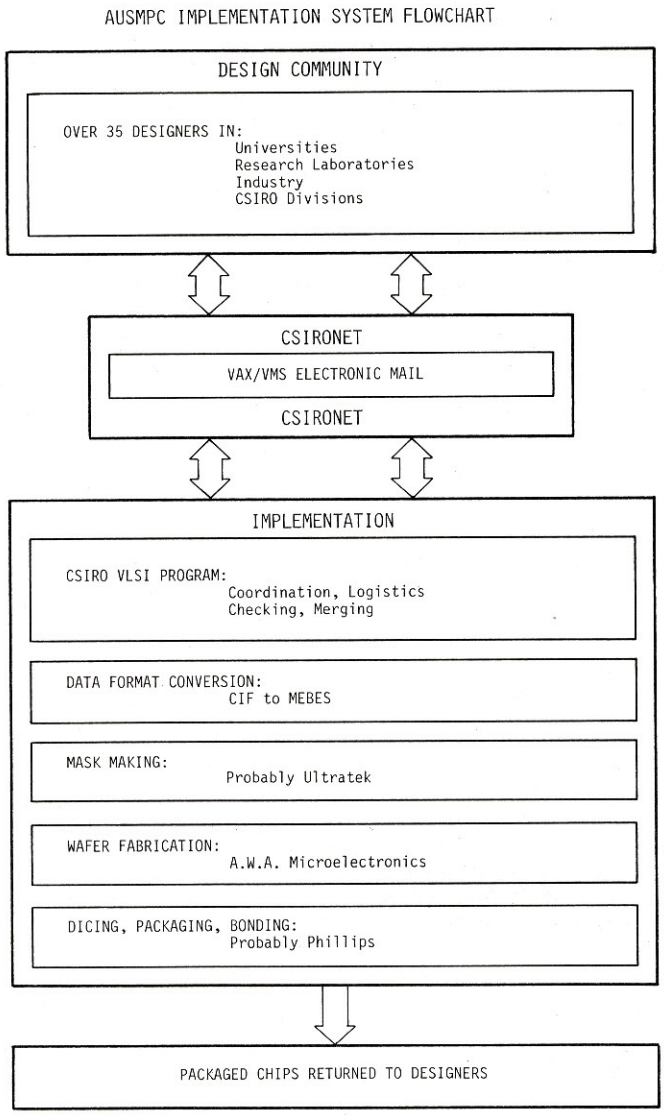

Figure 2 A flowchart of the Multi-Project Chip implementation system

Checking
With current verification programs, it is impractical for the coordinating body to fully check every MPC design for circuit correctness, design rule errors, etc. We can only assume that projects have been checked and deemed correct by the designers. The checking that must be done by the implementation system ensures that designs are submitted in valid CIF and that the size of each project does not exceed that specified by the designer. If the CIF is not valid, it will not be possible to have masks made. If the design boundary is incorrect, a designer's project may accidentally overlap another. Therefore, each designer's CIF file will be processed to determine its bounding box. In doing so, the validity of the CIF is checked by a CIF parser within the bounding box calculation program. The calculated bounding box is used during the merging of projects into the starting frame.

Project Selection
Each designer submitting a design expects to have his project implemented. Naturally, if space permits, all projects will be fabricated. However, if the number or size of projects is too large, some will be deferred. Projects which are likely to enhance our educational goal will be given highest priority.

Merging
Once all designs submitted by the May 31 1982 deadline have been checked, merging of the designs into the starting frame can begin. The merging will be done using a geometric editor to place the projects. Large projects will be placed on chip types with several smaller chips to average about eight projects per chip type. With about five different chip types, we expect to have around 40 projects on the Multi-Project chip set. The starting frame consists of a scribe line profile, alignment marks, layer identifiers, critical dimension testers, etch test patterns, die identification codes, and a series of test structures to be used to determine processing parameters.

2.3.2 Mask making

Following the data-format conversion step (a fifteen-hour computer run), the merged MPC design file is copied to magnetic tape. This pattern-generation tape is then sent to our mask supplier in California.

Because of the pioneering nature of Australia's first MPC effort, we are providing a backup for each step. Hence two sets of masks will be made. The first set will be sent to AWA in Sydney. The second set will be used to run a complete backup of the MPC (fabrication, dicing, packaging and bonding) in the US.

2.3.3 Wafer fabrication

The primary source for wafter fabrication will be the AWA Microelectronics facility in Sydney. AWA are developing an nMOS process suitable for MPC, by extending their existing CMOS process.

Part of our responsibility, as the coordinating body in an implementation system, is testing the wafer-fabrication source for predictability of turnaround and consistency of results. To this end, we supplied AWA with a set of test masks in October, 1981. These masks and other technical data are being used to characterise their new nMOS process.

2.3.4 Packaging

The chips are each 7.2mm by 7.2mm; each project may contain up to 40 input/output connections. Two packaging alternatives are being considered. One is an industry-standard 40-pin package. The other is a 51mm by 25mm ceramic substrate, typically used for hybrid microcircuit applications. Silver paladium conductors are silk-screen printed onto the substrate to provide contacts and connections to the package pins. Each designer will receive five custom-bonded chips containing his project, five unbonded chips in a waffle pack, a colour photograph of the chip type containing his project, and pin-connection documentation.

2.4 Communications Infrastructure

For the operation of AUSMPC, it is necessary for designers to communicate with each other and with our central implementation system. During the progress of MPC, designers need to be reminded of approaching deadlines and supplied with updated technical data. Moreover, designers may wish to share ideas and tools among the community. Communication by post or telephone is intractable for the communication needs of this community.

An Australia-wide electronic mail system is needed. However, such a system does not yet exist. We have rigged an experimental mail system by providing Australia-wide CSIRONET access to the VMS mail system on our laboratory's VAX-11/780 computer. In this configuration, CSIRONET is no more than a terminal switch to a single mail system.

In addition, an experimental link exists between CSIRONET and a network of UNIX computers in the Sydney area.

Submission of final design files will not be done using the electronic mail facility, however, since our file transfer facility is only reliable from one CSIRONET host to another. Most MPC participants use CSIRONET only as a connection to our VAX and so do not have file transfer capability. Therefore, final design files will be submitted on magnetic tape.

2.5 Timetable and Deadlines

Since the MPC is being run over such a short time (3 months), it has been necessary to enforce strict deadlines to monitor the status of AUSMPC 5/82. Designers failing to meet any of the deadlines will not have their projects fabricated. The final submission date must not slide, if fast turnaround from submitted designs to packaged devices is to be achieved, since the MPC is highly dependent upon pre-arranged overseas contractors.

The following deadlines have been set:

MAR	1 1982	FIRST BID FOR SPACE
APR	5 1982	SECOND BID FOR SPACE
MAY	3 1982	FLOOR PLAN AND METAL LAYER (to date) This will provide an accurate estimate of how designs are proceeding. It will tell us circuit complexity, size and progress.
MAY	31 1982	PROJECT REPORT
MAY	31 1982	CUTOFF

We expect to return packaged chips within two and a half months.

Once the chips have been fabricated and returned to the designers, a report on the testing of projects will be required. Information on performance, errors and corrections will be used to evaluate AUSMPC.

2.6 Number of Projects

The chip size for AUSMPC 5/82 will be 7.2mm square (scribe line centre to scribe line centre). Dimensions are summarised in the following:

Overall die dimensions: 7200 x 7200 micrometers

Vertical dimensions:

Top scribe line profile	:	77 micrometers
Setback to starting frame	:	21 micrometers
Starting frame strip	:	370 micrometers
Setback to project area	:	20 micrometers
Maximum project height	:	6627 micrometers
Setback to scribe line	:	8 micrometers
Bottom scribe line profile	:	77 micrometers
Total:		7200 micrometers

Horizontal dimensions:

Left scribe line profile	:	77 micrometers
Setback to project area	:	8 micrometers
Maximum project width	:	7030 micrometers
Setback to scribe line	:	8 micrometers
Bottom scribe line profile	:	77 micrometers
Total:		7200 micrometers

The dice will be arranged 69 per wafer: with about 60 available for projects. The area available on each chip for projects will be 7030 by 6627 micrometers. This area will fit up to 9 projects of size 2.34mm by 2.2mm. A project of this size can have about 7 input/output pads on a side and can contain circuits of considerable complexity (about 500 transistors).

Larger projects will have fewer projects per chip, so there will be fewer instances of that chip type on the MPC wafer. At least 10 instances of each project will be produced. There will be a minimum of 10 wafers in the fabrication run. Therefore, the number of chips of a particular chip type that must be placed on a wafer is equal to the number of projects on that chip type. A chip type containing 5 projects must appear at least 5 times on the wafer and a chip type containing 9 projects must appear 9 times. Therefore, even with 40 projects, only 40 of the chip positions on the wafer need be occupied. The space limitation for the MPC is not the number of chip positions on the wafer, but it determined by the number of different chip types available and by the number of projects that can be accommodated in each chip type. The additional chip positions will contain more instances of the already existing chip types to allow for breakages and damage due to handling.

3 DESIGN TOOLS AND LIBRARY COMPONENTS

3.1 Design Tools

Since we are introducing a new design methodology into Australia, few institutions have the necessary design tools. It was therefore decided to provide a basic set of tools to all participants in AUSMPC.

Three design tools have been distributed:

- a procedural layout language, called BELLE
- a Programmable Logic Array Generator, PLAGEN
- a plotting program, VIEWCIF

BELLE, a procedural-layout language embedded in the Pascal programming language, produces CIF. BELLE provides a very powerful means for parameterising circuit components. Hence, a circuit used in one design can be used in another by simply modifying the parameters for size, number of inputs or outputs, etc. We expect designers to build a library of BELLE procedures for commonly used circuit components such as memory cells, decoders, and comparators.

Since programmable logic arrays simplify the implementation of random logic, PLAGEN was also provided. It accepts logic equations in the form of a truth table and generates the appropriate layout in CIF. In addition to the creation tools described above, we have supplied an analysis program, VIEWCIF, for designers to plot CIF files

for manual analysis. Design rule and structural errors usually can be identified by this technique.

This simple set of design tools provides enough power to complete a modestly complex circuit of about 500 transistors. It is hoped that during this MPC and for subsequent MPCs, many more design tools will be created and become widespread in the design community.

3.2 Library Components

In order to concentrate design effort on systems design, and to minimise duplication of effort, a library containing many important sub-circuits has been distributed.

These have been shown to operate correctly in US MPCs. In addition to the CIF description of each circuit, we generated and distributed further documentation. This included the CIF symbol number; the symbol name; a brief description of the circuit function; the size (in lambda) of the symbol; the position of connections (including all inputs, outputs, power supply connections, and clocks); a block diagram of the symbol, circuit diagrams; length-to-width ratios of all transistors in the circuit (where relevant) and some sample SPICE simulations.

These library components include standard input/output pads, all the subsections of a programmable-logic array, shift register cells on a pitch compatible with the PLA, and superbuffers for driving clock and control lines.

The library and the design tools were distributed on magnetic tape at the end of February.

4 LAUNCHING THE CONCEPT IN AUSTRALIA

Active promotion of the concept began in mid-1981. Initially an invitation was sent to all Australian Universities teaching either Electrical/Electronic Engineering or Computer Science, and to selected Institutes of Technology. Invitations were also sent to various divisions of CSIRO, other research laboratories, and high technology industries which we considered might be involved in suitable research activities for the MPC. About 40 letters were initially sent. The distribution list has since grown to nearly 100 entries.

Three bulletins (in August, October, and December 1981) were published to provide more details and updates on the status of AUSMPC. These have included such information as a description of the implementation system, data formats for receipt of designs and for distribution of library cells and design tools, deadlines, details of the instructors' workshop, suggested reading, access to the electronic mail facility, and some preliminary process information. Considerable interest in the MPC has been stimulated by these bulletins.

Many seminars were given by various members of the VLSI Program to universities, research institutions and industrial firms, as well as at conferences.

4.1 Instructors' Workshop

An intensive three week instructors' workshop to teach the design approach of Mead & Conway to university lecturers and researchers was held in February, 1982. The first week of the course was an intensive lecture period of live lectures, video-taped lectures, and tutorials. [Eshraghian, K.E., and Pucknell, D., 1982]

The second week continued with several lectures, guest lectures and videotapes, while participants also began work on a project of their choice. This format was repeated during the third week.

4.2 Status

The first deadline (March 1, 1982) resulted in 69 bids for space. The list of projects includes the following:

- smart sensor array
- cochlear implant receiver/stimulator
- generator for transcendental functions
- graphics clipper
- error-rate monitor
- PCM signal and line-code generator
- digital filter (several)
- multi-channel analyser
- content-addressable memory
- Kung adder
- testable PLA
- serial data matrix transposer
- radio-astronomy signal processor
- FIFO memory
- direct-memory-access controller
- local-network arbitrator
- ECG pattern recogniser

5 FUTURE DEVELOPMENTS

A second MPC is planned for late 1982 or early 1983. The spacing and number of deadlines will be adjusted as a result of feedback from the first MPC. We expect there to be two or three MPC runs per year in 1983 and 1984.

5.1 Analog and CMOS

To date, there is not a satisfactory standard interface to fabrication for analog circuits, because the electrical behaviour of such circuits depends critically upon process knowledge. There is clearly a demand for an implementation service for analog design; to this end, we have several test structures on AUSMPC to determine if a standard interface is possible.

Several of the application systems in AUSMPC require low-power operation and hence CMOS. Thus we plan to offer a CMOS implementation in 1983. In fact, several of the first AUSMPC designs are nMOS prototypes of CMOS realisations.

5.2 More Advanced Processes

Historically, standard processes (such as those used for MPC designs) have had more relaxed design rules than leading edge production processes. This relaxation can mean as much as a two-generation difference (factor of four) in performance [Mudge, 1981].

Designers of systems on silicon will reach the point where, in order to produce competitive products, they can no longer tolerate such penalties. Silicon foundries will respond by radically improving the characterisation of their advanced processes.

Furthermore, there is likely to be a divergence between processes for standard high-volume chips and processes for custom chips. Although equivalent in minimum feature size, a process for custom chips will be oriented towards the special needs of the system designer. We can already see this happening in the case of interconnect layers in MOS technology. A second layer of metal has been shown to shorten design time. Hence, process

developers with custom chips in mind have diverged from the traditional evolution path followed by the producer of standard chips, who finds a second layer of polysilicon to be adequate.

7 ACKNOWLEDGEMENTS

We gratefully acknowledge the encouragement and technical support of Professor Lynn Conway and Mr Alan Bell of XEROX Palo Alto Research Centre.

CONWAY L., [1981]. The MPC adventures: experiences with the generation of VLSI design and implementation methodologies. Second Caltech Conf on Very Large Scale Integration, January 1981.

LYON R. F. [1981]. Simplified design rules for VLSI layouts. Lambda Magazine First Quarter 1981, pp 54-59.

CONWAY L. [1979]. The MIT '78 VLSI system design course: an outline of the course plan and results. Xerox PARC, August 1979.

CONWAY L., BELL A., NEWELL M., LYON R. and PASCO R. [1980]. Implementation documentation for the MPC79 Multi-University Multi-Project Chip-Set. Xerox PARC, January 1980.

ESHRAGHIAN K. and PUCKNELL D.A. [1982]. An introduction to VLSI system design - course notes. University of Adelaide.

MUDGE J. C. [1981]. VLSI chip design at the crossroads. Proc VLSI 81, Edinburgh, Academic Press, 1981.

STROLLO T., DOUGHTY T., KRASNER G., STONE M., WILNER W. and COHEN D. [1980]. Documentation for participants in the MPC580 multiproject chip-set. Xerox PARC, July 1980.

MEAD C. A. and CONWAY L. [1980]. Introduction to VLSI Systems. Addison-Wesley, 1980.

HON R. W. and SEQUIN C. H. [1980]. A guide to LSI implementation. Xerox PARC, 1980.