

Teaching Integrated Circuit and Semiconductor Device Design in New Zealand: The University of Canterbury Approach

Dr Richard J Blaikie^{*}, Dr Maan M Alkaisi, Dr Steven M Durbin, and Dr David R S Cumming[†]
*Department of Electrical and Electronic Engineering,
University of Canterbury, Christchurch, New Zealand*
[†] *Department of Electronics and Electrical Engineering, University of Glasgow, Glasgow, U.K.*
** E-mail: blaikie@elec.canterbury.ac.nz*

Abstract

Teaching the practical aspects of device and chip design in New Zealand presents many problems, including high manufacturing costs, long lead times, and the lack of local industry strength. Nonetheless, it is possible to overcome these issues. This paper describes the courses in these areas at the University of Canterbury, including a practical IC design project that has been running successfully for the past four years.

The IC design project takes final year students through a full custom design using modern design tools and fabrication processes. The design is quite straightforward — a 4-bit arithmetic logic unit — but it emphasises the importance of design, simulation and testing. The final circuits contain a few hundred transistors, so good practice is essential. Twelve designs are integrated on to a single chip to keep costs down, and individual designs are addressed via multiplexers. The designs are fabricated using a 0.5 micron process, accessed through a multi-project vendor (MOSIS). Getting chips back from a manufacturer is significantly more motivating for the students than just performing a paper design.

1. Introduction

The electronics sector now represents around 10% of world trade [1], which can largely be attributed to the engineering and technological advances in the semiconductor and Integrated Circuit (IC) industries. Unfortunately, New Zealand has not participated strongly in the custom IC design sector, with few exceptions [2]. It is perceived that the large capital investments required to perform custom IC design are too great for a country the size of New Zealand. To overcome this perception engineers and technologists must be trained in the practical aspects of IC design, with an emphasis that large up-front costs are not always necessary.

A practical IC design project has been running at the University of Canterbury since 1998. A ‘bottom up’ approach is used to teach final-year students the important aspects of custom silicon IC design. The students’ circuit designs are integrated into a single, multi-project chip, which is fabricated through the services of a multi-project vendor. High-level design methodologies are taught separately through lectures from an outside expert. This approach keeps costs to a reasonable level for a University programme.

This paper will describe this project in the context of the whole Integrated Circuits programme at the University of Canterbury. The next section will summarise the undergraduate and postgraduate courses that are offered. In Section 3 the circuit specifications and design tools used for the IC design project are described and, following this, the methods for assembling the students’ circuits into a multi-project chip are outlined, and details of the chip fabrication are described. In Section 5 broader aspects of the project are discussed, including assessment methods, student feedback and project costs. Finally, some general conclusions are drawn.

2. IC Courses at Canterbury

The Department of Electrical and Electronic Engineering offers a four-year undergraduate degree, together with a mix of both taught and research postgraduate degrees. The first two years of the undergraduate degree concentrate on teaching general background material, and specialization occurs in the final two years.

The IC teaching programme consists of two specific courses on Semiconductor Device Technology and IC design, one each in the third and fourth years, but these are supplemented by other related courses. For example, advanced courses on digital hardware and software design teach many aspects of high-level design and hardware implementation using field-programmable devices, so the IC design courses can concentrate on topics that are

critical to understanding the design of custom silicon ICs, and semiconductor device engineering in general.

The emphasis in the third year course is put on the basic principles of semiconductor device physics, the physics and operation of the most important semiconductor devices, and some basic integrated circuit layout. The layout editor tool used for the advanced circuit design is introduced in this course. Approximately half of the students in the Department take this course (50-60 per annum). Due to the difficulty level of some of the material – particularly the semiconductor physics – and the lack of direct job opportunities in this field in New Zealand, only around 20-30 students choose to take the advanced course in their final year.

The advanced course has as one of its major components a practical IC design project, which is described below. This project is supported by lectures on chip design and layout issues, including guest lectures from practitioners if possible. The course also covers a wide range of semiconductor device and materials issues, with topics including: quantum mechanics and its effects on modern semiconductor devices; analogue integrated circuit design issues; silicon device fabrication and materials processing technologies; semiconductor material properties and growth; emerging technologies; and, high speed transistor structures. This part of the course is supplemented with some practical laboratory sessions in the Department's Micro/Nanofabrication laboratories.

At the postgraduate level a course on Advanced Semiconductor Devices is offered, which looks at issues of fabrication, materials growth and device design. As there is no local industry driving the direction of this course, there is freedom to explore a wide range of topics; in most years a visitor to the Department will offer material in this or a related course. The students taking this course are usually research students from the University's Nanostructure Engineering, Science and Technology (NEST) group, although students from other groups or even other institutes often attend. Use is made of the NEST group's experimental facilities in the course, including electron beam lithography, atomic force microscopy, reactive ion etching and molecular beam epitaxy. For example, in a recent exercise in electron beam lithography one of the course students produced a map of New Zealand on a silicon sample, as shown in Fig. 1. This is the smallest known map of New Zealand, and produced a great deal of media interest. Such attention will help to boost the interest in this important area.




Figure 1. Atomic Force Microscope (AFM) image of a New Zealand map drawn using electron beam lithography.

3. Canterbury IC Design Project

The practical aspects of IC design are taught in the fourth year course, with the emphasis being on getting a design implemented in silicon so that the whole design process can be understood. A major design project is carried out by the students, which involves going from the layout of basic logic gates on silicon to testing and debugging the designs once chips are returned from fabrication.

3.1. Design Specification

The circuit that has been used for this project is a simple 4-bit Arithmetic Logic Unit (ALU), shown in its top level of abstraction in Fig. 2. The circuit takes two 4-bit inputs, A and B , together with a 1-bit carry input, C_{IN} and produces a 4-bit output F together with a 1-bit carry output C_{OUT} . The design is to be implemented in a 0.5 micron CMOS process, and each team of two or three students is given $0.2 \times 0.2 \text{ mm}^2$ of silicon real estate in which their final layout must fit.

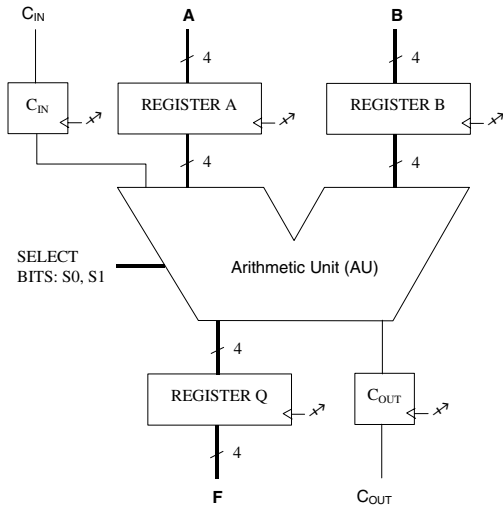


Figure 2. Top level abstraction of the 4-bit Arithmetic Logic Unit (ALU)

Two select bits, S_0 and S_1 , control the circuit's functionality according to the function table shown in Table 1. These functions represent a set of simple arithmetic operations that form the instruction set of the ALU. A simple circuit design such as this has been chosen for two reasons. Firstly it reduces the complexity and size of the final circuit, and secondly it ensures that all students can easily understand the logic operations of the ALU they will design. More rigorous aspects of complex digital hardware and software design are taught in other courses of the students' final year programme [3].

TABLE 1. Function table of the ALU.

Select		Output (F)	
S_1	S_0	$C_{in} = 0$	$C_{in} = 1$
0	0	A (TSF)	A + 1 (INC)
0	1	A + B (ADD)	A + B + 1
1	0	A + B'	A + B' + 1 (SUB)
1	1	A - 1 (DEC)	A (TSF)

The first task for the students to perform is to derive a gate-level circuit from the high-level specifications given in Fig. 2 and Table 1. Immediately there are design decisions to be made, and no single design is correct. The chosen design is simple enough that the weaker students can follow through the gate-level design from a standard textbook example [4]. In many cases the stronger students will explore other possible implementations for the circuit, and some very interesting and efficient designs have been presented to date.

One important aspect of the chosen design is that it is a synchronous circuit, with all input data being stored in registers prior to the ALU operation, and all output data similarly latched at the output. This gives the students the chance to explore the possibility of using a *pipelined* architecture to raise the maximum clock frequency of the circuit. Additional registers can be used to break the ALU operation into smaller, faster combinational blocks, which can be performed sequentially. Many students see this as a challenge to see who can get the fastest possible circuit operation. The constraint of limited available silicon area determines how far this approach can be taken.

3.2. Leaf Cells

Once the gate level circuit design has been finalised the next task is to define a set of basic functional gates (leaf cells) that will be designed and optimised using the layout tools, as described in the following section. At this stage the important outcome is the identification of a small number of logic blocks that, once designed, can be re-used throughout the circuit to reduce the overall design time. The trade-offs between circuit complexity and design simplicity become clear.

A typical set of leaf cells the students have to design is shown in Fig. 3, namely an inverter, a 2-input XOR gate and a 4-input compound logic functional block. In addition, a static latch is provided as a 'standard' library part for constructing the various registers. The presence of a compound function links well with the transistor-level design techniques that are taught as part of the course [5], in which efficient techniques for designing such blocks are presented. For the XOR design the students are directed to recent literature [6] showing them that there are a number of possible circuit implementations for even the most simple logic gates.




Figure 3. Typical set of leaf cells: (a) inverter, (b) XOR and (c) compound gate.

3.3 Design Tools

With the leaf cells defined, the students begin the major tasks of the assignment, namely performing the transistor-level design and mask-level layout of the leaf cells. This is followed by the assembly of these parts into a final circuit layout, with the addition of metal interconnects. They have been introduced to the layout tools and techniques in a previous course, and many groups use the inverter they designed in this course as a starting point.

A number of different design tools are used for performing the layout and simulating the resultant circuit(s). The L-EDIT suite from Tanner Research [7] is used for layout, design rule checking and circuit extraction. Lambda-based design rules are used, but for modern sub-micron IC design these have become somewhat process-specific. Since the target fabrication technology is a 0.5 micron, 3-metal CMOS process, sub-micron design rules are used. The multi-project vendor provides the design rules, transistor models and simulation parameters.

Resource sharing is an issue that arises during this phase of the project, as only one licence for the circuit extraction tool is available. Whilst this causes some anxiety at times, it can provide valuable lessons about project planning and team cooperation. In addition, licences have not been purchased for some add-ons, such as place-and-route tools or standard cell libraries. This makes the layout process rather laborious in parts (another reason for choosing a very simple circuit), but the result is a custom design in every respect.

3.4 Simulations

The circuit extraction tool produces SPICE-like net-list output, and device models are available for the target process. The students are therefore able to simulate their designs at every stage of the process. For simulating leaf cells and small circuit sub-sections the P-SPICE evaluation version that all students have free access to is sufficient, however this cannot handle the full design. At this stage resource sharing is again an issue, as there are only a limited number of licences available for the full version of P-SPICE or Tanner's T-SPICE software.

Simulations of the leaf cells are required to check functionality and obtain gate delay estimates. These are used in static timing analyses order to obtain an estimate for the maximum clock frequency of the circuit. At this stage critical paths can be identified and transistor sizing along these paths can be optimised in order to obtain improved performance.

Hierarchical design is emphasised, so once leaf cell have been designed and tested the students must assemble

this into larger functional blocks prior to laying out the whole circuit. It is suggested that a 1-bit data path be designed, so that four of these units can be assembled to make the final circuit. Simulation of the 1-bit path is also required in order to check functionality.

Simulation of the whole circuit is not requested, primarily due to the resource-sharing constraints. However most students tend to do this anyway as they are keen to ensure that they get a working chip. Note that complete functional testing of even this simple 4-bit ALU is very difficult as it has 11 independent inputs. As the net-list for the extracted circuit will contain a few hundred transistors, the simulation time for an exhaustive test is prohibitive; this is another important lesson. Most students will test a few simple and important cases, but the good students use this experience to learn the art of selecting appropriate test vectors and estimating test coverage. This fits well with a series of lectures that are presented at this time by an outside expert. In these lectures aspects of testing, validation and verification of large designs are discussed from an industry perspective.

4. The Multi-Project Chip Approach

A key aspect of this project is the incorporation of the individual student designs onto a single chip. The chip is then fabricated using the services of a multi-project vendor, with 12 packaged parts and 13 bare silicon die being returned. The multi-project chip approach means that the whole class is acting as a single design team, with each group working on their own (identical) part of the chip. The course coordinator acts as the design team leader, setting deadlines, performing final chip-level integration, and ensuring that all teams submit their designs on time. In addition, the combination of a multi-project chip fabricated using a multi-project vendor keeps costs manageable for a university teaching programme.

4.1 Multiplexed Design Spaces

To minimise costs the final chip is designed to fit within the minimum chargeable area for a single design, which is usually around 5 mm². The chip area used for the project to date is somewhat smaller than this at 2.8 mm², which leaves open the possibility for adding additional design spaces in future.

The floor-plan layout for the chip is shown in Fig. 4, in which there are 12 0.2×0.2 mm² spaces allocated for the students' designs. Surrounding these are 32 bonding pads and bi-directional pad driver circuits, which is insufficient to give each design space a separate set of input and output pins. A set of 11 inputs is therefore applied in common to all the designs.




Figure 4. Floor-plan of the chip showing 12 student design spaces, bond pads and 8-1 output multiplexers.

For output selection purposes the chip is partitioned into two blocks of six design spaces (D1–D6 and D7–D12). The outputs from each block are selected from the six design spaces using 8-1 multiplexers, which requires three multiplexer select inputs for the chip (these are common for the two multiplexers). Two test modes are also provided using the multiplexers', to test the integrity of the input signals and clock signals on the chip.

Power supply and clock inputs are common to all design spaces, although each has its own two-phase clock generation circuit so that only a single-phase clock needs to be supplied and routed around the whole chip. In total, 29 of the 32 bond pads are used.

4.2 Design Integration and Checking

The major workload for the course coordinator begins once all the student designs have been submitted. It is his/her responsibility to integrate all the individual designs into the overall chip layout, and to ensure that there are no fatal errors in any of the designs.

Using the L-EDIT software the design integration is straightforward (the beauty of hierarchical design), so most effort goes into design checking. As the final chip contains around 7,000 transistors this can be a little laborious. Firstly a design rule check is performed on the whole chip to ensure that all design teams have adhered to the design rules throughout. Secondly, and most importantly, a net-list is extracted for the whole circuit and some simple, fundamental connectivity checks are performed. The purpose of these checks is to ensure that there are no fatal flaws in the chip, such as supply to

ground short circuits. A more thorough check of the final circuit is not possible, as a licence for the layout versus schematic tool has not been purchased.

4.3 Use of a Multi-project Vendor (MOSIS)

A major motivation for the students in this project is the prospect of receiving real chips back from a vendor (Fig. 5). The MOSIS multi-project vendor [8] has been used for this project to date, as they provide access to a number of technologies at reasonable cost, with an appropriate minimum number of parts. There has also been previous local experience in the use of this vendor [2].

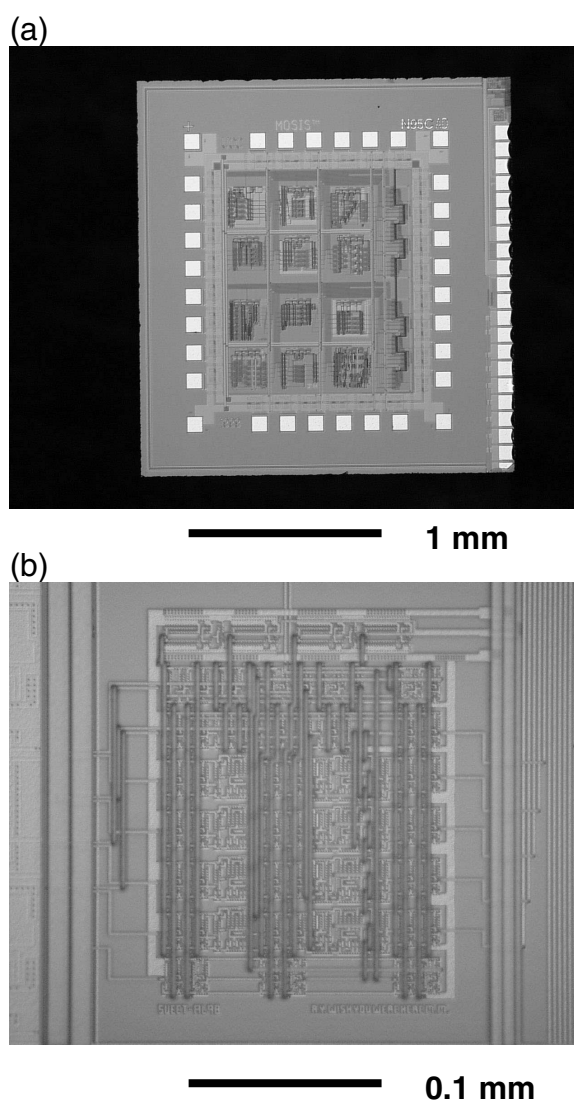


Figure 4. Optical micrograph of (a) the completed chip and (b) one of the student design fields.

Of key importance for this work is the planning of the project to fit the vendor's schedule for design submission. The manufacturing turn-around time is typically around 12 weeks, so designs must be submitted relatively early within the academic year to enable working chips to be returned in time for the students to be able to test them. It would not be possible to run this project in a single-semester course, and the full-year course structure of the Canterbury degree is ideal.

Of the many processes offered by the vendor the HP/Agilent 0.5 micron process was initially chosen for a number of reasons. Firstly it is a sub-micron, multi-metal CMOS process that results in maximum operating frequencies around 200 MHz for the simple design being implemented. This is sufficiently close to 'state-of-the-art' performance for the students to remain highly motivated during the project. Testing is not performed up to these frequencies because of limitations in the chosen packaging and test boards. Secondly, the process is offered on a monthly basis by the vendor, which allows for the possibility of getting working parts fabricated in time even if the proposed deadline is missed. Finally, the cost of using this process is significantly less than if one of the newer 0.35, 0.25 or 0.18 micron process were used. This process has therefore been used for the first three years of the project.

In the fourth year of the project an opportunity arose to join the MOSIS Educational Program (MEP), which is sponsored by US semiconductor industries, including the Semiconductor Industry Association (SIA). This gives free access to a limited number of processes, with a limited number of runs, but was previously restricted to US educational institutes. In 2001 the University of Canterbury was invited to join the MEP, with significant project cost savings resulting. As the HP/Agilent 0.5 micron process is not available under the MEP, the design was transferred to an AMI 0.5 micron process during 2001. This process is similar in many respects, but some modifications to the pad driver protection circuitry were required. From the students' perspective the translation was seamless, as the same design rule set is used in both cases. Unfortunately, of the four designs fabricated through this project, only the latest one (made using the AMI process) has suffered catastrophic failure. An improper translation of pad drivers to this technology is suspected, and a re-design of the pad ring and basic chip layout will be necessary in future years.

5. Discussion and Conclusions

The implementation and administration of this project has been a successful and rewarding experience for all those involved. Approximately 100 students have participated in the first four years, and the project can

manage up to 36 students each year within its current framework. Now that the project gets support from the MEP it will be possible to expand the scope to include more students, both from other courses at the University of Canterbury (including postgraduate projects) and from other institutions.

5.1 Assessment

The students submit two reports during the course of the project, and these are used to assess their performance. The first is due after the design has been sent for fabrication. This outlines the design processes and simulations they performed to produce their final layout, and includes an estimate of the maximum operating frequency for their design. The second is due at the end of the year, once the chips have been returned from the vendor and electrical testing has been carried out. This consists of a functional test report on the circuit, including the reporting of any faults and the identification of their cause. The multi-project nature of the chip is a nice feature of the project at this stage, as it allows students to trouble-shoot non-functional circuits even if their own works perfectly — usually only one or two of the designs are faulty. The characteristics of pipelined and non-pipelined designs can also be demonstrated simultaneously. For internal assessment purposes the first report is worth 20% of the total mark in the course, and the second report carries another 5%.

As part of every course at the University of Canterbury an independent survey of the students is carried out to assess its effectiveness and provide feedback. In 1999 the overall rating of the course was 4.1 out of 5, compared to a University average of around 3.5. Student feedback about the assignment was extremely positive, with many indicating that it was the most rewarding assignment that they had carried out in their four years at university. The students are generally very proud of their achievement of designing a full custom IC (see Fig. 5), and many take one of the die-form chips away with them as a trophy of this achievement.



Figure 5. Students studying their circuit.

5.2 Project Costs

Success in this project is necessary, as it has been the most expensive undergraduate project run within the Department by far. Even with the cost savings made by fabricating a multi-project chip through a multi-project vendor the fabrication costs are approximately \$US3,200 per annum. It costs about half as much as this again to purchase and maintain the design tools. Now that the chips can be fabricated as part of the MEP these costs have reduced dramatically, however the need to retain up-to-date design tools is imperative.

The Department sees the project as beneficial and it is set to continue in future years. As well as providing a practical demonstration to our future electrical engineers that custom IC design is possible in the New Zealand environment, it also provides a practical grounding in the basics of modern microprocessor design and operation for associated digital hardware and software courses.

5.3 Conclusions

A programme has been developed at the University of Canterbury to teach a wide range of aspects of IC and semiconductor device design. Courses are taught in the final two years of the undergraduate degree as well as at the postgraduate level. Students are exposed to a wide range of issues in semiconductor materials and device engineering, areas in which the Department has an active research presence.

An IC design project has been implemented successfully within the final year undergraduate course. A relatively simple circuit has been chosen for the design,

and the students use full custom layout techniques to optimise its performance. Twelve designs are integrated into a single, multi-project chip, and this chip is fabricated using the services of a multi-project vendor. This approach keeps costs down to a reasonable level, as well as teaching important team design skills. The project is set to continue, and it would be possible to expand it to incorporate designs from other university programmes in the future.

This project is an important cornerstone of the IC teaching programme at the University of Canterbury, and has received good feedback from the students involved. It complements the suite of computer software and hardware courses that the Department offers, which are more directly relevant to New Zealand's electronics industry; the aim is to provide a pool of well trained engineers who feel confident to tackle custom silicon designs when the need arises.

6. Acknowledgements

The Department of Electrical and Electronic Engineering is gratefully acknowledged for providing continuing financial support for the project. The enthusiasm, commitment and understanding of the students are also acknowledged as being essential for the success that has been achieved to date.

7. References

- [1] M. Kelly, "2005: a Pause for Silicon", *Physics World*, March 1998, pp 21-22.
- [2] B. Marlow, "Custom IC Design – How and Where", *Proc. ENZCON'98*, Dunedin, N.Z., 31 August – 1 September 1998, pp 94-97
- [3] S.D. Round and W.K. Kennedy, "Teaching of Digital System Design Through Practical Assignments", *Proc. ENZCON'98*, Dunedin, N.Z., 31 August – 1 September 1998, pp 73-77
- [4] M.M. Mano, *Computer Engineering: Hardware Design*, Prentice-Hall, New Jersey, 1998, pp 240-244.
- [5] N.H.E. Weste and K. Eshraghian, *Principles of CMOS VLSI Design: a Systems Perspective*, Addison-Wesley, Reading, Massachusetts, 1993.
- [6] J.-M. Wang, S.-C. Fang and W.-S. Feng, "New and Efficient Designs for XOR and XNOR Functions on the Transistor Level", *IEEE J. Solid-State Circuits*, vol. 27, no. 7, July 1994, pp 780-786.
- [7] <http://www.tanner.com>
- [8] <http://www.mosis.org>