

A joint application to the
Science Research Council
for a four year special grant to
support research into
"the computer aided design of very
large scale integrated circuits."



University of Edinburgh

Departments of Computer Science

and

Electrical Engineering

A joint application to the Science Research Council for
a four year special grant to support research into
"the computer aided design of
very large scale integrated circuits".

Principal Applicants:

J.Tansley.
J.Jordan.
J.H.Collins.
P.D.A.Schofield.

Associated Applicants:

R.Kelly
D.J.Rees.

March 1979.

CONTENTS

		Page
1.0	Introduction	1
2.0	The impact of VLSICs on Systems	2
3.0	The Nature of the Problem.	2
4.0	Research Programme.	3
5.0	Resources Requested.	5
6.0	Industrial Collaboration and Support.	9
7.0	Research to be Supported by the Facility.	11
8.0	Academic Staff.	21
9.0	The Wolfson Micro-electronics Institute.	22

1. Introduction

Application is made for staff and computing resources to support the design and research into the methods of design of large scale and very large scale integrated circuits (VLSI). This is seen as covering two distinct, but related, activities: namely:-

- (a) The design of specific circuits to be used in various instrumentation and computing systems.
- (b) Research into, and the development of, design methods and their associated computer aids used in creating VLSI circuits.

This project will play a major part in the development of micro-electronic technology in the UK and therefore a four year special grant is requested.

The report of the SSPDC working party on VLSI makes clear the importance and significance of research and development work in integrated circuit design. In particular, they suggest that Universities should 'single out the area of microcircuit design and realization and its impact on systems'. The Departments of Computer Science, Electrical Engineering and the associated Wolfson Microelectronics Institute (WMI) at Edinburgh are particularly well suited to do this.

The Department of Electrical Engineering is currently establishing an SRC funded Central Micro-electronic Processing Facility (CMPF) to implement the production of silicon circuits, primarily using a silicon gate N-MOS process. Considerable interaction with CMPF staff is expected: there will be a need to generate layouts with minimum effort and time by quickly evaluating experimental circuit design. The Department has considerable experience in designing integrated circuits and research into semi-conductor technology. The WMI has successfully designed and manufactured a number of custom integrated circuits on a commercial basis. All of this activity is related to a broad base of innovative instrumentation and communication systems design.

The Department of Computer Science has designed and implemented a first generation computer aided integrated circuit design system (GAELIC), which was initiated by the WMI. This system is now used by a number of research institutes and industrial organisations. The Department also has considerable experience in the construction of large computer software systems. In particular, the multi-access system implemented on ICL 4/75 and now ICL 2900s demonstrates an ability to develop reliable and usable software. A significant part of the current research work of the Department is concerned with the design and construction of hardware and software for computing systems. The utilization of micro-electronic circuits is a natural continuation of this. At present the VLSI circuit design stage accounts for a large part of the cost of producing prototypes of new chips. The further development of computer aided techniques will reduce this significantly and permit more freedom of experimentation with new devices incorporating VLSI circuit chips.

The competitive position of the UK, in exploiting the LSIC revolution is most likely to be limited by the shortage of people with key skills and a good appreciation of the role and applicability of micro-electronics. It is clear that specialised combined training in

computing, micro-electronics and systems theory and their application is required. The CMPF and the VLSIC design facility can be expected to make a significant contribution to the training problem at both undergraduate and postgraduate levels.

2. Impact of VLSIC's on Systems

Complex large scale integrated circuits are being introduced so rapidly that only a small number of professional engineers and teachers in the UK are fully aware of their applicability and potential. The signal processing power of large scale integrated circuits enables low cost instruments to be developed having great functional complexity, adaptability and maintainability. In particular, micro-processors enable intelligence to be distributed throughout a control system thereby reducing the amount of data that must be transmitted for processing by a central computer. Further, micro-processor based instrumentation systems offer a highly sophisticated user interface. Consequently it can be expected that instruments of considerable complexity will become easier to use. Future instruments will be compact, highly reliable, yet flexible. Flexibility will result from the use of programmable electronic systems. As labour costs rise the trend towards instruments having simpler controls and minimal maintenance will continue. Difficulties commonly experienced with the maintenance of complex instruments will be eased by the inclusion of automatic fault detection and location facilities. Systems using LSI circuits will be designed so that any particular task only uses a subset of available hardware/software elements thereby allowing an alternative route to be taken once a fault condition has been established.

To ensure that the U.K. will be able to be competitive it will be necessary to establish a research base that brings together VLSIC designers, micro-electronic process engineers, computer scientists and research workers studying a wide range of instrumentation systems and transducer problems. Research on novel approaches to instrumentation based on micro-processors and programmable electronics requires access to a VLSIC design facility where circuits can be generated in association with the CMPF with the minimum effort and time. In addition research on design and test automation is required to bring together chip design with the test problems of the end user.

3. The Nature of the VLSI Design Problem

As the components of integrated circuits scale down in size so the number available on a single chip rises, so that a single IC could soon contain 10^{**6} or more active devices. The size of the chips and the physics underlying their manufacture dictate that complete systems are fabricated as a unified activity. The amount of information associated with any one design is therefore large and still increasing. The problem is to control this quantity of information in a sensible way. Clearly computer systems are capable of assisting with this problem and the following are identified as important features of such a system.

- (1) The information associated with a design needs to be safely accessed by a number of designers. They also need to be able to manipulate and create libraries of sub-system structures.
- (2) Where possible sub-system structures need to be placed and connected together automatically.

- (3) Procedures are required to test designs for 'correctness' and verify that they conform to process design rules before they are made.
- (4) Despite certain automatic procedures design will probably proceed iteratively, hence a good interactive design station is required. The user interface should be such that it helps develop a design style that encourages the best utilization of the opportunities offered by design in silicon.

The potential use of VLSI circuits is unknown, as is the impact of an easy to use, comprehensive design facility for these circuits. First generation design facilities such as GAELIC go some way to providing the right kind of design environment. GAELIC, though has certain limitations. It is essentially a geometric editor with post-processors that produce magnetic tapes to drive manufacturing equipment. It falls short in providing good descriptive languages and complete reliable verification of an integrated circuit design.

There are a number of ways to approach these problems and the solution depends on the working environment of the designer and his proposed circuits. Experience in designing with VLSI is essential and needs to be gained concurrently with the development of second generation design systems. The technology-independent concepts of IC design need to be recognised and abstracted. Only in this way will the kernel of a useful design system emerge.

4.0 Research Programme

4.1 Initial Activity

- a) Configure the computing system (see 5.1) in the two Departments and get the operating system running on the design stations and main computer..
- b) Establish local communication between the design stations and VAX, (see 5.1).
- c) Implement standard packages such as GAELIC with "help" documentation so that circuit design may proceed.
- d) Implement data communications with Rutherford Laboratory to use EB facility.
- e) Establish compilers and other utilities.

4.2 Development of a Research Programme

Over the next few years we can expect to see a convergence of semi-conductor technology, computer systems design and data communications with an appropriate impact on the research being conducted in these areas. The provision of this facility along with the CMPF is seen as a significant contribution to developing this very important area.

After an initial stage of commissioning, research work (see section 7) will move on to the system as various facilities become available. Each project will proceed semi-independently, but efforts will be made to co-ordinate this activity. For example, the ULA prototyping system

(see 7.26) once operational would be of great value to others. Similarly as database facilities become available (see 7.28) so it would be appropriate to make user documentation, standard circuit designs etc., accessible through the system. Equally the Department of Electrical Engineering is currently active in several areas of systems research which are expected in the near future to lead to requirements for VLSI components, (see section 7).

J.Tansley & J.Jordan have been identified as the principal co-ordinators with the role of establishing collaboration and feedback to the various groups using the system. As research proceeds, so further proposals will be generated that use the CMPF and the LSIC design facility. We would expect to see other researchers in related disciplines being "boot-strapped" into thinking in terms of modern micro-electronic technology. Collaboration with WMI design staff is expected here, (see section 6.1).

Additionally, activity needs to be initiated to integrate testing facilities into the system. There is a need for both programmable function testing and speed testing. For the first of these a standard test station and interface machine is being developed. Interpreters will need to be developed to provide designers with methods of exercising chips. By adopting standard pulse cards it should not be necessary for designers to be intimately involved in test set-up. The test machine will be close to the CMPF and have 'remote' access for designers through the design station. Speed testing, (outwith research needs), if it is considered important enough, would eventually be supported, at great expense, by equipment from one of the usual vendors.

Research into the methodology of VLSI circuit design and layout is currently being conducted in the Department of Computer Science (see 7.21). The development of this work will be the subject of further research proposals which will use the facility that is the subject of this application.

4.3 Time Scale of the Research Programme

An estimate of the time scale is given in the figure 1.0 shown below:-

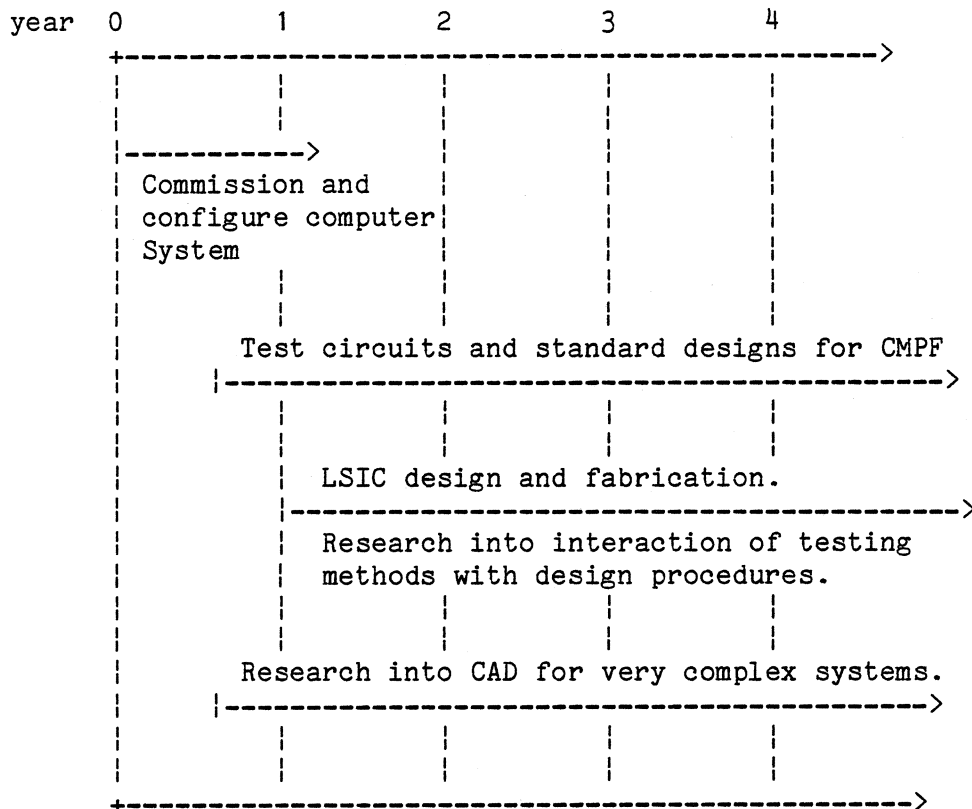


figure 1.0

5.0 Resources Requested

5.1 Computer Systems

A computer system is required that provides, as a minimum, the following:

- a) reasonable computing power
- b) a good interactive user interface
- c) access to the system from local laboratories
- d) good system software and utilities
- e) sharing and structuring of data
- f) a programming language with some degree of general acceptance
- g) some standard software.

Clearly there are a number of possible system configurations that could meet these needs. The following system is therefore proposed, with, we hope, appropriate justification and an indication that it will meet the computational research needs over the next four years. We would expect systems and software developed to be portable and simply incorporated into new and cheaper hardware as it becomes available.

5.11 System Overview

A summary of the system requested can be seen with reference to figure 2.0

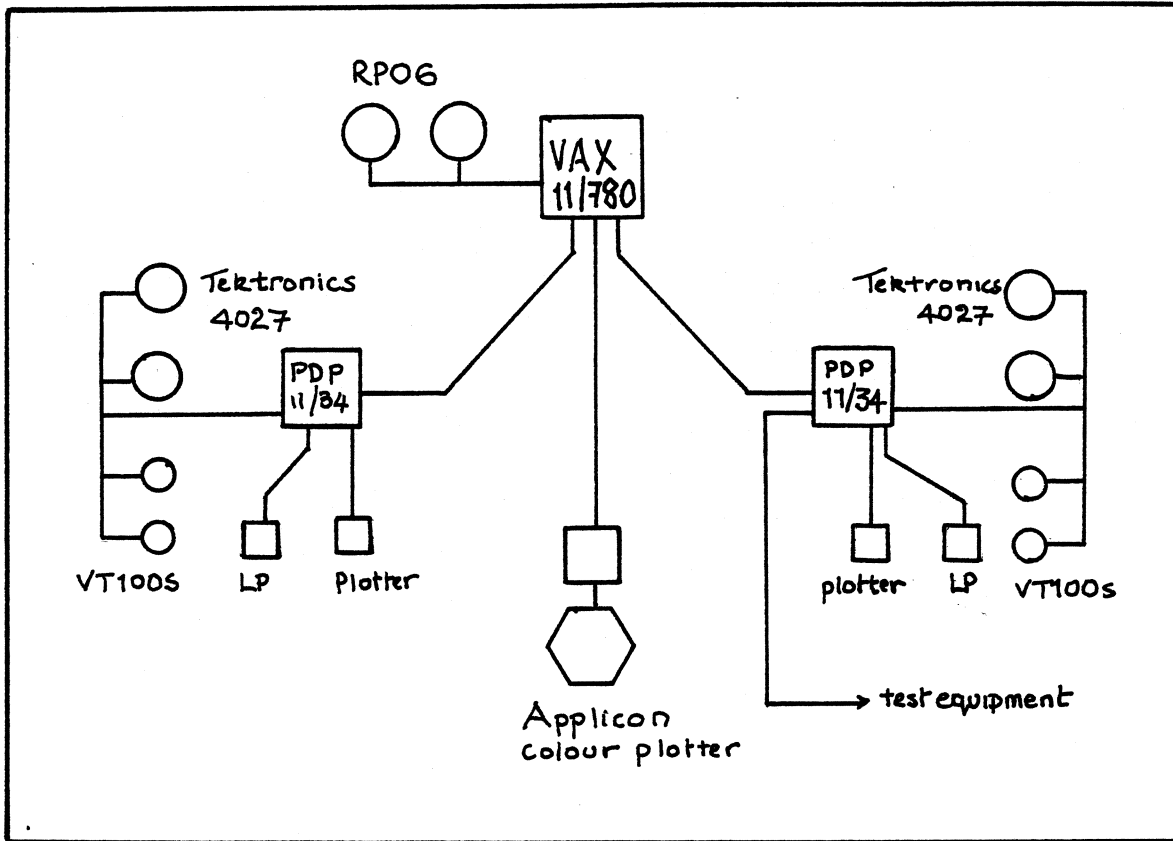


figure 2.0

5.12 Equipment List

1. Computational System

VAX 11/780 computer with
2.0 megabyte of store
2 * RPO6 Disc system
TU16 magnetic tape drive
2 * DMC11 High speed communications link
DUP11 Synchronous link to Rutherford.

Approximate cost £186,330

2. Design Stations (two)

PDP-11/34 with
128kbytes of store
2 * RL01 Disc drives
DMC11 High speed communications link
2 * Tektronix 4027 colour terminals
2 * VT100 terminals
CalComp 1039 plotter
300 lpm line printer

DZ11 communications lines
IB11 IEEE Bus interface

Approximate cost £110,200

3. Fast Colour Plotter

Applicon plotter and associated
Interfaces.

Approximate cost £42,000

5.20 System Characteristics

a) Reasonable Computing Power

A simple quantification of this requirement for VLSI design means that the system, as a whole, should contain at least one dedicated processor that will execute about one million instructions per second. Additionally, it should have a direct user address space of at least 2^{24} bits. These requirements are necessary to deal with the geometric resolution encountered with VLSI chips and the processing associated with semi-automatic procedures, like routing in a reasonable time. The number of computers that meet this requirement is not large, and being limited to sensible prices restricts the choice further. At this stage the alternatives are systems offered by DEC, Prime, Perkin Elmer and the like. Our preference is for a DEC VAX 11/780 for the following reasons:

1. The Department of Computer Science already has a VAX 11/780 that is used for first year teaching. Hence there is good local expertise for this machine.
2. This expertise extends to the provision of compilers and utilities.
3. The VAX 11/780 is now well integrated into the Department's local area communications network. It will be a simple matter to extend this to incorporate the CMPF within the Department of Electrical Engineering.
4. DEC will almost certainly produce 'down range' versions of the VAX. This will eventually mean cheaper and more widely used systems.
5. Local maintenance support for DEC equipment is good and certainly as good as that of the potential competitors.

In order that the VAX system can devote its computing power directly to the VLSI problems it is proposed to connect user terminals via two design stations. Each station contains a processor capable of performing suitable functionally distributed computation, such as generating graphics instructions for plotters and displays. Our preference for this processor is a DEC PDP-11/34.

b) Interactive User Interface

The design station drives the peripherals associated with the VLSI design procedure. The actual design is conducted on one of two colour graphical terminals. The colour display is a powerful addition to this

kind of system, greatly aiding the designer by representing the different semi-conductor substrates in separate colours, (e.g. red for polysilicon green for diffusion etc). This gives a clearer comprehension of these complex circuits and helps visual checking of the circuits as design proceeds. Two additional character terminals are requested to enable users to enter basic circuit descriptions before needing to use the colour displays.

c) Access from Laboratories

The main processor i.e. VAX would be located in the Department of Computer Science air-conditioned machine halls. The design stations would be in separate laboratory space. The Electrical Engineering Department have space allocated near the CMPF. The Computer Science Station would be located in an adjacent hardware laboratory.

d) Systems Software and Utilities

Bell laboratories have now released the UNIX/32v operating system for the VAX machines. This is known to be a comprehensive and usable system, (cf Berkley University). Since UNIX is also available on PDP-11s this makes this system a sensible choice to run. UNIX has many merits, perhaps the most important being its simple malleability for various needs.

e) Application Programming Languages

We would expect the majority of new developments to be coded in ALGOL-68. This language is fast being accepted as a major applications language. Other research work in graphics for example, which may be of use to this project, is being coded in ALGOL-68. (R.Forrest and I.Braid).

f) Sharing and structuring of data

An important part of any design system is the structuring and sharing of design file data. Research work is progressing in this area (see 7.23). In the meantime we would expect to import database software (such as Ingres, which runs under UNIX) to cope with our immediate needs. The two RP06 disc drives (186 mbyte capacity each) are thought to provide sufficient storage space for the proposed level of usage.

g) Standard Software

The system has two roles, one being the support of design methods, the other the design of integrated circuits. To this end we propose to use packages such as GAELIC, SPICE and SUDS etc, to support these immediate needs while new systems are evolving.

5.3 Special Plotter

Although it is proposed to attach CalComp plotters to the design stations these devices will be mainly used to generate simple check plots and plots where considerable accuracy is required. The CalComp style of plotter has two major limitations. The first is that it only draws lines and not shaded areas which are encountered in integrated circuits geometry, and secondly it is slow for any significant density of line drawing. The Applicon plotter overcomes these limitations drawing colour shaded pictures in minutes as opposed to hours. This quick turn round of design plots is thought to be particularly desirable

for advanced teaching courses and research and development work where time is at a premium.

5.4 Access to Rutherford Electron Beam Facility

Part of the manufacturing process that can follow directly from the design stage is the generation of masks used in circuit fabrication. Rather than transport magnetic tapes with a variety of IC descriptions it would be simpler to communicate directly with the Rutherford E-Beam mask making facility via data lines. As part of our software developments we would like to propose an intermediate design file code to which CAD programs could operate and from which the EB facility could produce masks.

5.5 Systems Support Staff

Three people are requested. A VAX manager and two other systems support programmers. One of the support programmers would be assigned to help CMPF and Electrical Engineering staff use the computing facilities. He or she would also be responsible for developing and incorporating manufacture and test equipment into the computer system. The other systems programmer would support the inter-systems communications.

6.0 Industrial Collaboration and Support

6.1 Wolfson Microelectronics Institute

The proposed VLSIC design facility will strengthen the already strong interaction of Academic staff with the contract research and design work of the Institute. All of the LSIC designs innovated by the Department of Electrical Engineering have been implemented by the design staff of the Institute. It is intended that all future research (including SRC funded programmes) involving the design of LSI circuits will, where possible, use the design staff of the Institute on a sub-contract basis. Both parties will benefit from the immediate cross-fertilisation of industrial and research circuit design programmes.

External users of the CMPF will in some cases require design guidance to be able to successfully use the LSIC design facility to produce chip layouts. The work load of Academic staff is such that only a limited time could be devoted to this activity. To overcome this difficulty we propose that an official arrangement be made for the Wolfson Microelectronics Institute to become the body that all external users would use on sub-contract basis to obtain design guidance and complete chip layouts when required. To facilitate these suggestions we propose that the Wolfson Microelectronics Institute be offered the concession of free use of the design facility when servicing these sub-contracts. Use of the facility to service commercial design contracts would be charged at rates to be determined by SRC.

6.2 Other bodies

Discussions are currently underway with the following companies and research organisations:

GEC Computers,
Ferranti,
ICL,
Racal Redac,
RSRE,
PORL,
Pye TMC,
STL

7.0 Research work to be supported by the facility

The following appendix summarises research work currently being conducted by the Departments of Electrical Engineering and Computer Science that is either related to the design of micro-electronic circuits or to their methods of design. In most cases it is expected that extensive use will be made of the requested facility. Progress of some of this work has been limited by their lack of access to good computing power. In some instances as the research develops it may be the subject of further applications for support for either hardware, services or staff specific to that research project.

7.10 Department of Electrical Engineering

7.11 Correlation circuits for measurement systems

J.R. Jordan

Correlation techniques based on flow generated noise have been used for a number of years to measure the flow of abrasive, toxic and corrosive fluids. Recent research has shown that, by careful design of an ultrasonic transducer system, clean fluids can be monitored to an accuracy approaching that of the magnetic flowmeter. Transducers at two points along the flow stream output voltage noise signals resulting from the flow generated noise. The peak of the cross-correlation function of these two signals is time delayed from the origin by a factor inversely proportional to the flow velocity. A requirement therefore existed for a low cost circuit that would detect the peak position of a correlation function and output a signal proportional to the flow velocity. The overloading counter correlator was designed to satisfy this requirement. The NRDC have patented this correlation system and financed the WMI to produce a special purpose large scale integrated circuit for use in correlation flowmeters. These LSI circuits have been successfully produced and are now being used in a commercial version of this correlation flowmeter.

Microprocessors were not available at the start of this project. Subsequent work (financed by the Royal Society) has shown that a competitive correlation flowmeter can result from designing an appropriate interface circuit for microprocessors. This interface circuit could easily be implemented as a special purpose LSI circuit, thereby enabling a very compact implementation of the correlation flowmeter. The additional functional flexibility that results from the use of the microprocessor enables the instrument to satisfy a range of application requirements, for example, flow control can be implemented with flow measurement in one instrument. A discrete I.C. interface system has successfully demonstrated the feasibility of this proposal and a chip circuit suitable for implementation on the CMPF n-MOS process has been investigated. The case of CCD circuit techniques is currently being investigated to assess the possibility of implementing a relay correlator version of the interface circuit. The use of a chip design facility will soon be required.

Multiple correlator systems have been proposed for the remote measurement of, for example, velocity and position. The ready availability of the correlation interface circuits would enable these systems to be implemented at low cost.

7.12 Instrumentation for fault detection and condition mastering in wave energy schemes.

J.R. Jordan and H.W. Whittington

Automatic condition monitoring of wave energy schemes and other complex plant is now a practical possibility. Microelectronic technology has developed to the point where complex data processing functions can be realised by circuits implemented on a single silicon chip. Circuits can be designed to check the condition of a system, predict future states of the system and diagnose possible fault conditions. As a result of this increased data processing capability it can be expected that transducers will be simplified and will therefore become more reliable. These new developments will impact any wave energy scheme likely to be constructed even within the next ten years. It is expected that the Department of Energy will finance the WMI to carry out a study of the instrumentation requirements of fault detection and condition monitoring to assess the possibilities offered by microelectronic technology. This research programme will generate systems suitable for LSIC design studies.

7.13 Speech processing applications for chirp transform processors

P.M. Grant and J.M. Hannah

Speech processing provides one potentially large application area for CCD C-Z-T processors. We are currently investigating the channel vocoder which operates in the frequency domain, transmitting a quantized representation of the short term spectral envelope of speech. In conventional vocoder implementation, this is achieved by a bank of contiguous bandpass filters whose outputs are rectified and low pass filtered. In addition to the spectral envelope a pitch period estimation has to be transmitted before the speech signal can be fully reconstructed by the synthesiser in the receiver. In comparison with conventional time domain transmission systems, e.g., 64 k bit/sec PCM, the spectral envelope is down sampled and quantized to achieve a considerable reduction in transmitted data rate.

We are approaching this work with both practical experimentation and computer simulation. A 64-point CCD sliding chirp-Z-transform processor has been constructed and fully characterised. This processor has demonstrated a transform accuracy of 1.5. Operation of the processor as a spectrum analyser with Hamming weighting has been investigated and sidelobe levels of -36 dB have been measured. Variation in clock frequency permits frequency resolution to be varied in the range 6 Hz - 1 kHz over processing bandwidths 400 Hz - 64 kHz. The simulations we have programmed operate directly on speech waveforms by sampling up to 1.2 sec of signal and inputting it to the computer. The programme then models both the channel vocoder analyser and synthesiser and outputs data into a microprocessor controlled buffer interface, for subsequent playback off line through a loudspeaker. Initial listening tests highlighted deficiencies in the sliding C-Z-T transform due to distortions caused by the non stationary frame to frame sampling. Thus we have now incorporated a direct transform and are investigating details of vocoder design and implementation while assessing relative speech quality at 2.4, 4.8 and 9.6 k bits/sec transmission rates.

These studies have been performed under an SRC CASE studentship with the Advanced Development Division of RACAL Electronics Ltd, by Malcolm

Davie. Davie has fully researched the application of CCD CZT's to channel vocoders and his results are so encouraging that we are now considering a suite of IC's to realise these compact low-power vocoders.

7.14 Novel MOS Device Design and Development

J. Mavor

This is a substantial programme aimed at investigating the potential of MOS integrated circuit technology for novel signal processing applications. Recently, the research has been focussed around MOS/CCD sample - data, analogue structures for the implementation of a key signal processing element: the programmable transversal filter (PTF). Additionally, recursive filters and Fourier transformers have also been realised with MOS/CCD circuits, signal processing with MOS/CCD integrated circuits offer the vital system advantages of low cost, low power in compact form. The research work has been active for over seven years and has been in liaison with the CCD group of RSRE, Malvern. The research is currently supported in the Electrical Engineering Department by DCVD(MOD) and SRC through collaborative programme with J.O. Scanlan of U.C. Dublin

As a result of the promising MOS research in the Electrical Engineering Department, WMI have taken on a number of significant industrial contracts and systems. Multi-tapped delay lines, time delay and integrated devices and advanced PTFs have been pioneered for several significant applications in sonar and general filtering fields.

As a result of the combined research and development in the WMI/EE Department at Edinburgh, we have invented a number of novel, analogue MOS structures e.g. 4-quadrant multipliers and MOS operational amplifiers. Other research has been devoted to advanced systems design based on Z-80 microprocessor to control a number of PTF chips. This enables complex processing to be undertaken in hardware form under microprocessor control.

A significant current development is in adaptive filtering based on the Widrow LMS algorithm. This will permit the concept of 'intelligent' filters which adapt their characteristics to suit the signal form. Such a filter can now be envisaged in single-chips, MOS/CCD I.C. form and should replace a system configured with conventional digital chips for applications including, adaptive equalisation for telephony and medical applications.

7.15 Design Requirements Directly Associated With the CMPF

The main objective of the CMPF is to produce hardware. However, this does imply a significant design requirement for three special cases.

The first is for test patterns. As each major new process stage is commissioned we require a suite of test masks to evaluate the new process and to confirm the design work limits. It is likely that these simple but illustrative circuits will also be of use to newcomers to the field who wish to have a simple introduction to I.C. hardware.

This overlaps with the second category of designs: those for postgraduate training. Ideally, we would like research students (from other Universities as well as Edinburgh) to experience a 'hands on'

design and fabrication training program. The limits of time and cash may restrict this to metallisation of a logic array but there is likely to be significant and recurring design load from this category of students.

The final requirement is for automatic generation of I.C. modules. This is likely to arise during a research program when a critical new design or process is being educated. The chances of success are increased if new circuit segments can be made on their own and then coupled with other circuit sections later. Thus we can approach LSI in a practical way by interconnecting well-evaluated cells or, if only a small number are required, by interconnecting our custom chips with standard commercial units on a hybrid substrate. In either case, the problem is similar and an effective automatic routing system is required.

7.20 Department of Computer Science

7.21 Computer Aided IC Design-Methodology & Models

I. Buchanan & D.J. Rees

The increasing complexity of integrated circuit design necessitates a hierarchical and modular approach to the design task. The design style developed at Caltech, and documented in "Introduction to VLSI Systems" by Mead and Conway, is a first step in this direction.

For many years, computer aided IC design systems have been faced with a poorly structured, essentially monolithic set of design data. This has inevitably led to design systems based on the lowest common denominator, i.e. the physical or geometric description of the design, and implies inherent limitations in the quality of the design tools provided. Properties of the logical structure of the design e.g. connectivity and behaviour, are extremely difficult to quantify or verify.

However, a design philosophy which embraces a top-down, hierarchical, regular, tessellated structure in both its logical and physical representation leads directly to a description language and data structure with those same attributes. Thus a design system which operates within this framework has concurrent access to a consistent set of data covering both the logical and physical description of the circuit, and may therefore provide the designer with a more sophisticated, controllable and efficient set of design aids. In particular, verification procedures e.g. design rule checking and some simulation can be localised to within a block.

A block is regarded as the unit of description. Its external appearance is defined physically as an inviolable rectangular region and logically as a function with a number of connections. Thus, when a block is instanced in a design only its external features need be verified at that level; its interior can be ignored. This gives a substantial performance gain over systems which insist on macro-expanding block descriptions. The advantages of this approach become particularly apparent in view of the regularity and repetition in IC design e.g. memories, busses, registers and parallel data paths.

The traditional physical representation of a design is as a large collection of polygons on each of a number of masks. However, this polygonal form is not suitable for modelling connectivity and its use has a historical rather than a theoretical basis. Variable width wires, or tracks, are a much more attractive model from both a logical and physical standpoint.

The proposed design philosophy and computer aided design system is wire-based and block structured. Analogies with structured programming e.g. modularity, local variables and well defined interfaces, are both valid and instructive.

Current research takes a comprehensive textual description of the logical and physical properties of a circuit and in an initial prototype design aid provides flexible viewing (windowing and zooming) of the design on a colour display and by interactive control verifies the circuit according to a set of geometric design rules, calculates electrical characteristics, produces simulator compatible input files and can trivially convert the original description to a stick diagram. Graphical editing will be strictly limited, if present at all.

7.22 Microprogramming for VLSI

G. Wood & J. Tansley

Microprogramming methods are a strong candidate for the implementation of the control logic of single-chip sequential functions. Its regularity of structure and conduciveness to formal verification are attractive properties, well-suited to the design needs of large scale integration.

At present we have a facility for automatically mapping a canonical behavioural description of a microprogram into microinstruction words of a specified format. With this facility investigations are now possible on how the characteristics of a microprogrammed implementation of the control structure of a processor might be constrained by the fact that the processor (and the control) is implemented on a single chip on silicon. For example, one might investigate whether it would be worth sacrificing more efficient use of control memory for the elimination of additional decoding logic by providing separate fields in the micro-instruction word for each of the mutually exclusive control signals normally grouped together in a single field.

Another topic of investigation is the customising of the sequencing logic for the control of microprograms to the particular microprogram permanently resident on the chip whose function it controls. Currently, sequencing of microprograms is normally performed by standard packages, e.g. AMD 2900 series, which impose constraints on the control structure of the microprogram being implemented - e.g. placement of microinstructions and the "fan-out" of branching. PLA's offer considerable potential for the customised control of microprograms and, due to their extreme regularity of structure, are readily implementable on silicon. Therefore some method for automatically generating a suitable PLA configuration for the control of a specified microprogram can be investigated and developed.

7.23 Databases for computer aided design

M.P. Atkinson

Modern VLSI design techniques require the construction of a large volume of data as a model of the circuit. This model is composed of many aspects of the chip from intended function to geometric details, with complex interrelationships. As the techniques for integrated circuits design are rapidly evolving it is important to establish these models of circuits formally to retain adequate flexibility.

The use of the collected data in the design environment is particularly demanding. Three types of access are required: unpredictable, one off, accesses to investigate any detail or general property which is of interest to the designer; rapid scans of selected parts of the data for supporting such things as interactive graphics, and intricate traversal of the data for rule checking and simulation. The first two forms of access are well supported by modern relational databases with the power of expression given by relational query languages, and performance achieved by links, query optimisation and streaming. Initially such facilities would be provided by importing INGRES, which is implemented for the VAX. Research and development would be necessary to replace components to improve performance.

The final form of access is not well supported by any existing database system. Two problems exist: the programming languages used to write the fairly complex programs involved do not relate well to the manipulations and data structures held in the database, and it is difficult to arrange storage structures which support the operations efficiently. The former problem is critical as it has a significant effect on the ease with which simulation and rule checking programs may be developed. Research is already underway in the department on this aspect of the problem (and funds have already been requested from the SRC to support it) precisely because the importance of this problem was recognised when it was encountered on four earlier CAD projects. Substantial insight has already been gained in the search for a possible solution, which suggests that the locality of a type declaration can be used to imply its persistence and an underlying relation, provided types can be imported into a program module just as predefined constants may be at present. The two research programmes are likely to be mutually beneficial, but each may proceed independently.

Whenever a database is used as a model for designers it has the advantage of providing a single copy of the state of the design facilitating cooperation between designers and reducing the chance of inconsistent or divergent design decisions being taken by the separate designers. But if the designers work on exactly one total version they are overconstrained. It must be possible for a designer to conduct, perhaps extensive, experiments to investigate design proposals and later decide whether to incorporate them or abandon them. During the period of these experiments the other users must be able to proceed undisturbed by them. This requires large transactions scratch pad facilities and arbitrarily deferred updates. Such capabilities can all be developed in conjunction with the proposed relational, but programmer oriented, system possibly using differential files. The development of such a system with appropriate functions, and of good ways of supporting it, is an important area of research with wide applicability.

The investigator who would be responsible for this aspect of the work has already spent eight years working on supporting data for CAD, mainly with the Rainbow Group at Cambridge, and has designed a CAD database system which is in use for aircraft design.

7.24 Semantics and Correctness of Digital Systems

M.J.C. Gordon

Currently work is proceeding on modelling the functional behaviour of digital systems using concepts from denotational semantics. Such concepts are powerful since it is possible to finely tune levels of abstraction. This enables formal descriptions to be clearly structured. The systems considered so far can be described at many levels, for example:

1. In terms of the connections between sequential components such as registers and memories and combinatorial components such as adders, shifters, etc. At this level the primitive operations are the opening and closing of gates.
2. In terms of the transfer of information between registers, memories, arithmetic units, etc. At this level the primitive operations are register transfers and tests - we abstract away from the gate level implementation of these.
3. In terms of higher level data types and their associated operation. For example a digital system emulating a computer could be described by target machine operations on words, bytes, integers etc.

LCF is being examined to see how far it can be used to formalize and reason about levels such as these. In particular, PPLAMBDA theories and ML tactics are being developed to verify the emulation of one level in terms of the operations of a lower one.

As a first case study the verification of the design of a very simple microcoded computer is being examined. This involves formalizing:

1. The architecture of the microprogramming level.
2. The semantics of microcode (i.e. its effect on the host machine state).
3. The semantics of microcode (i.e. its effect on the target machine state).
4. The translation of microcode into a specification for the opening and closing of gates.
5. The translation of machine code into the microcode which emulates it.

Work has begun on proving that microcode is correctly realized by the corresponding gate operations on the host machine, and that machine code is correctly emulated by the corresponding microprograms.

7.25 Microelectronic Circuits for Virtual Address Translation

D. J. Rees

Virtual address translation for computers with a segmented and paged architecture traditionally uses segment and page tables indexed by partitions of the virtual address. The positions of the partitions determine the page size and the size (or potential size) of the tables. Apart from some small machines with segmentation alone, the size of the tables dictates that they should reside in main storage. In order to avoid the inefficiency of perpetually referring to these tables, therefore, a fast associative memory is used to capture the translations of the most recently used few addresses. The form and size of associative memory may vary but they are always regarded as transparent to the systems programmer.

Notwithstanding the widespread acceptance of this scheme, it has a number of drawbacks.

From the architectural point of view, the way segmentation is implemented is unfortunate. The desirable attribute of a segment is that it should define a logical division of the virtual address space which is capable of separate protection. This is indeed achieved by means of placing protection mode information in each segment table entry. However, segmentation is also used to overcome a difficulty which arises purely from the indexing method of address translation. This is the problem of page tables requiring consecutive areas of main store. With paging alone or with large segments the length of page tables may become a considerable embarrassment to an operating system. By reducing the size of the page index partition in the virtual address, the size of individual page tables is reduced and the problem is largely overcome. This poses no difficulty for the many small logical areas which are usually required, but conflicts with the equally common requirement for a few large logical areas. The usual expedient is to allow addresses to overflow between segments without signalling an error. The effect of this is to demean the level of protection possible to a significant extent.

The amount of storage space taken by often sparsely filled segment and page tables can be a noticeable overhead on storage requirements. Even taking the rapidly reducing cost of main store into account, a system which could do without would be an improvement. (The desirability of paging itself as a method of optimising store utilisation is likely to remain for the foreseeable future.)

A remaining short-coming in traditional schemes concerns assistance to the operating system in making its storage scheduling decisions. The technique of placing page usage information (read, written etc.) in the page table entries is difficult to make use of efficiently. An operating system will typically wish to implement a "working set" or a "Least Recently Used" replacement policy. In order to do this the usage information must be inspected and reset at frequent intervals. This is a very time-consuming activity, however. Longer intervals therefore have to be chosen and this militates against the effectiveness of the scheduling.

It is believed that the advent of LSI and VLSI may bring about a solution to these problems. In particular, it should now be possible to

construct very large and sophisticated associative memories. This has never been the case hitherto.

Segment and page tables could be dispensed with altogether where the associative memory is large enough to contain entries for all the pages currently in use by a process. The systems programmer would arrange for each new page entry to be inserted into the associative memory instead of into a page table, and conversely for removal of a page. Access protection can be achieved by storing access bits with the respondent part of each associative memory entry and performing a validation after the appropriate entry has been associatively selected. This gives protection down to the page level instead of just the segment level.

Segmentation is now totally divorced from the address translation mechanism which is as it should be. The design of logical segmentation features can therefore take place with much more freedom, whether they be an integral part of the CPU or whether they are regarded as purely an operating system feature. In either case, the desire for many small segments and some large ones simultaneously can easily be accommodated.

Usage information could be updated in the associative memory entries, just as in page table entries but this does not solve the operating system's difficulties in implementing its desired storage scheduling algorithm. A solution to this would be to regard the associative memory as a stack such that as any entry is matched, it is removed from its present position, the entries above it shifted down one place, and the entry reinserted at the top. By this means the topmost entries will contain the most recently used addresses, and the least recently used addresses will be at the bottom. Any page shifted out of the bottom of the stack would be a candidate for replacement. This mechanism also allows convenient updating of a "written" usage marker each time an entry is reinserted. A "read" marker is no longer required.

It is proposed to investigate the feasibility of constructing an associative memory of the form outlined above in LSI technology. If it proves feasible, devices would be fabricated and incorporated into the design of a prototype test system, probably based around existing microprocessor devices, and an evaluation carried out. A successful trial would naturally lead to the desire to construct an entirely new design of processor to take best advantage of the new device.

7.26 An Uncommitted Logic Array prototyping system.

L.D.Smith & N.H.Shelness.

Uncommitted Logic Arrays now provide 500 or even 1000 logic elements (2-input NAND gates, for example) on a single chip. These can be interconnected to form digital systems by patterning a final layer (Ferranti), or layers (IBM), of metalisation which effectively customises the chip.

A major problem with the use of ULAs to implement prototype systems is the length of time required to produce the artwork for this metalisation layer from whatever primary design description is used. Ferranti, for example, start with a logic diagram and require several weeks, and several #K, in order to produce this artwork (using totally manual methods). IBM, on the other hand, can perform this task in under 3 days. We propose to attack this problem by extending our existing technology independent prototyping system to cope with ULAs.

A basic philosophy that is important with this system is that the prototyping process be as rapid and as correct as possible. This end is best served by a completely automatic system, even though this may lead to "sub-optimal" implementations. This lack of optimality is unimportant for prototypes except where it impacts severely on the performance of the implementation.

Currently we have a trial implementation of a simple technology-independent prototyping system which aids the implementation and documentation of digital hardware using SSI/MSI components, standard (pre-defined) boards, and stitch-welded or wire-wrapped interconnections. The system is based on a simple hardware description language (DL1) (Smith 1976) which is used as primary input, and on the concept that individual processes in the system transform, update, or generate output from a design document which flows through the system.

This system has been successfully used by our SRC-supported distributed computing project (Shelness et al) to aid the construction of interface hardware for their triple-processor computer system.

Our system supports a variety of automatic processes that can be applied to a design document:

- (1) Pure macro processing.
- (2) Syntax checking.
- (3) Flattening of hierarchical descriptions (with automatic resolution of naming clashes).
- (4) Gate level simulation (3-valued time-independent, and 3-valued time-dependent both worst-case and pseudo-random delay).
- (5) Assignment of element-level objects to chips.
- (6) The placement of chips on pre-defined boards (trial implementation only).
- (7) Generation of stitch-weld schedules (for the manually operated Vero-wire "zapgun").
- (8) Automatic generation of logic diagrams (trial implementation only).

It is clear that the early stages of prototyping with ULAs could be identical to those required for prototyping with SSI/MSI TTL. A divergence would occur after the basic design document had been generated. At this stage the element-level description would have to be expanded into ULA-elements and these would have to be placed on the ULA. In order to preserve some vestige of technology-independence it would be necessary for the ULA to be described in DL1. This might require extensions to DL1.

Furthermore it is essential that the placement of elements be done with forethought, or the tracking of interconnections could prove impossible, especially where only one layer of metalisation is provided (Ferranti), and where "tunnelling" (Ferranti) under a connection incurs a performance penalty. IBM use 3 layers of metalisation, no doubt in order to avoid some of these problems.

In order to make the tracking of connections ULA independent it is necessary to use a routing algorithm which, in the jargon of PCB tracking, is able to work with fixed vias, floating vias, or no vias at all. Again the the geometry and topology of the ULA must be expressible in DL1, and this may require further extensions.

Since features such as "tunnels" or vias may impact (severely) on performance, it may be necessary to simulate at least parts of the design at a very low level (device-level). For this purpose a device-level simulator would need to be constructed. It is arguable also that a system-level simulator could prove useful in providing confidence in a design's overall correctness. In both cases extensions are required to DL1 in order to support behavioural descriptions.

8.0 Academic Staff

J.Jordan, Ph.D., B.Sc.

Dr. Jordan has been a lecturer, specialising in Instrumentation and Electronic Circuits in the Department of Electrical Engineering, University of Edinburgh for the past nine years. Prior to this he was a senior engineer with EMI Electronics Limited working on numerical machine control systems and analogue computer design, and has had teaching experience in the control and instrumentation systems area at Teeside Polytechnic. He is currently supervising projects in the areas of microelectronics ion-selective transducers, flow measurement methods using correlation flowmeters based on a novel correlation technique patented through NRDC and optical measurement systems. The production of a custom integrated circuit for a digital correlator has been undertaken by the Wolfson Microelectronic Institute with finance from NRDC, who have also supported the production of four prototype correlation flowmeters using these integrated circuits. The Taylor Instrument Co. Ltd., has recently been awarded a licence to commercially develop and exploit the prototype system. He is an instrumentation systems consultant to Taylor Instrument Co. Ltd., and MACAM Photometrics Limited.

J.Tansley B.Sc.

J. Tansley has been a lecturer in the Department of Computer Science for the past four years specializing in Computer Architecture and hardware design methods. For nine years prior to this he was engaged in computational neurophysiological research. This work was conducted at the Universities of Edinburgh, Manchester and Bristol.

He is currently supervising projects on directly executable languages machines, micro packing and distributed functional systems. He is also supervising the design and construction of a high speed data network.

His own research is currently concerned with the automatic layout and routing of polycells for integrated circuits.

Associated Staff

M.P. Atkinson, Ph.D., B.Sc., Lecturer, Computer Science.

I. Buchanan, B.Sc. Demonstrator

J.H. Collins, Ph.D., B.Sc., Head of Department, Electrical Engineering.

M.J.C. Gordon, Ph.D., B.Sc., SRC advanced research fellow, Computer Science.

P.M. Grant, Ph.D., B.Sc. Lecturer, Electrical Engineering.

J.M.Hannah, Ph.D. B.Sc., Lecturer, Electrical Engineering.

R.Kelly, B.Sc., Lecturer, Electrical Engineering.

J.Mavor, Ph.D., B.Sc., Lecturer, Electrical Engineering.

D.J.Rees, Ph.D., B.Sc., Lecturer, Computer Science.

P.D.A.Schofield, B.Sc., Head of Department, Computer Science.

N.H.Shelness, B.Sc., Lecturer, Computer Science.

L.D.Smith, B.A., Lecturer, Computer Science.

H.W.Whittington, Ph.D., B.Sc., Lecturer, Electrical Engineering.

9.0 Wolfson Microelectronics Institute

The WMI has been active in the development of novel semiconductor devices and integrated circuits since 1970. Much of its work has been collaborative with industry and it has built up a unique position in British Universities as a contract design and development agency in the field of microelectronics. In a field of such obvious industrial significance these industrial contacts have proved invaluable in keeping the University abreast of technological developments. Students and researchers have benefitted directly from the facilities built up by the Institute and gained enormous experience from working with engineers involved in the more commercially oriented contracted work.

The Institute has been instrumental in a number of important developments. The GAELIC suite of CAD programs for integrated circuits was initiated there and developed in conjunction with the Department of Computer Science. Digital integrated circuits have been developed out of research in the Department of Electrical Engineering for correlation flow measurement and at the request of Moore and Wright for their Micro 2000 digital micrometer. Both circuits have been highly successful technically and economically.

More significant perhaps has been the long and close liaison with the CCD research in the Department of Electrical Engineering. This has resulted in the design of a number of novel devices, the invention of several new design techniques and the development of improved systems modules. This work is continuing as described earlier.

The close cooperation of WMI with the Department of Electrical Engineering and Computer Science has proved in the past tremendously beneficial in not only keeping the research of industrial relevance but also providing facilities which were not otherwise readily available. In the present proposal the interaction of integrated circuit engineers, instrumentation specialists and computer scientists will provide a powerful capability to develop not only the new design tools and techniques but also a range of novel devices.